INTERACTIVE CONTROL ENGINEERING

Delays, Slow Measurements, and State Estimation

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 5

Delays, Slow Measurements, and State Estimation

Controlling what arrives late or cannot be measured

John D. Hedengren
Editor

Advanced Process Solutions, LLC

Interactive Control Engineering: Principles and Applications
Volume 5: Delays, Slow Measurements, and State Estimation
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/v05-delays-estimation/

How to cite this volume. John D. Hedengren, Ed., *Delays, Slow Measurements, and State Estimation*, vol. 5 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 measurements that arrive late and quantities that are not measured at all. Chapters 2 to 9 treat delay: fuel in transit to a fire, a slow analyzer on a column, a digester whose product is measured hours after it is cooked, residence time in continuous manufacturing, a delayed communication channel and a valuation that is a year old. They show what the delay costs and how prediction recovers part of it. Chapters 10 to 13 infer what no sensor reports: product moisture, a temperature inside a kiln, stored ammonia on a catalyst and the position of a vehicle under water.

Volume 1 supplies the models. Volume 9 uses the estimators of this volume inside predictive controllers.

John D. Hedengren
2026

Contents

Preface v

Notation x

1 Introduction: Delays, Late Measurements, and State Estimation 1

1.1 The Question This Volume Answers . . . 1
1.2 A Short History . . . 2
1.3 Notation and Conventions . . . 3
1.4 Delays, Late Measurements, and Estimators . . . 3
1.5 Instruments and Final Elements in This Volume . . . 13
1.6 The Simulator Framework . . . 14
1.7 Working With AI Agents While Studying . . . 18
1.8 Chapter Roadmap . . . 19
1.9 Exercises . . . 21

2 Wood Pellet Grill Simulator for Temperature Control Through a Fuel Delay 24

2.1 Introduction . . . 25
2.2 Smokehouses, Pellet Fuel and the Pellet Grill . . . 26
2.3 The Grill as Modeled . . . 27
2.4 Dynamic Character . . . 34
2.5 Implemented Control System . . . 35
2.6 Eighteen Cases, Measured . . . 37
2.7 Technology in Practice . . . 42
2.8 Limitations . . . 44
2.9 Conclusion . . . 45
2.10 Exercises . . . 47

3 Teaching Composition Control With a Slow Analyzer on a Thirty-Two Stage Column 50

3.1 Introduction . . . 51
3.2 Background . . . 52
3.3 Process Model . . . 53
3.4 Dynamic Analysis . . . 57
3.5 Controller Design . . . 60
3.6 Results . . . 64
3.7 Predictive Control Through a Lagging Analyzer . . . 68
3.8 Estimator Option and Relay Test . . . 71
3.9 Technology in Practice . . . 72
3.10 Limitations . . . 74
3.11 Discussion . . . 75
3.12 Conclusion . . . 76
3.13 Exercises . . . 78

4 Kraft Digester Simulator That Cooks Four Hours Ahead of Its Kappa Analyzer 82

4.1 Introduction . . . 83
4.2 Cooking to a Number: Kraft Pulping and Its Control . . . 84
4.3 The Digester as Modeled . . . 85
4.4 Hours Between the Heater and the Blowline . . . 88
4.5 The Control Layers as Coded . . . 90

4.6 Nine Scenarios, Measured . . . 92
4.7 Technology in Practice . . . 97
4.8 Limitations . . . 99
4.9 Conclusion . . . 99
4.10 Exercises . . . 101

5 Forecasting Frozen-In Stress in a Float Glass Line Simulator 105
5.1 Introduction . . . 106
5.2 Float Glass, Its Annealing and Its Control . . . 107
5.3 The Line From Charger to Cut . . . 108
5.4 Where the Ribbon Keeps Its Stress . . . 114
5.5 Implemented Controller . . . 115
5.6 Cases and Their Measured Outcomes . . . 119
5.7 Technology in Practice . . . 127
5.8 Limitations . . . 129
5.9 Conclusion . . . 130
5.10 Exercises . . . 132

6 Continuous Tableting Line Simulator for Tablet Weight, Dose and Material Tracking 135
6.1 Introduction . . . 136
6.2 Tablets, Granules and Continuous Lines . . . 137
6.3 Model of the Line . . . 138
6.4 Where the Dose Is Decided and Where It Is Seen . . . 144
6.5 Control System as Implemented . . . 146
6.6 Scenarios and Measured Results . . . 148
6.7 Technology in Practice . . . 153
6.8 Limitations . . . 154
6.9 Conclusion . . . 155
6.10 Exercises . . . 157

7 Material Diversion by Residence-Time Tracking in a Simulated Continuous Flow Synthesis 160
7.1 Introduction . . . 161
7.2 Residence Time and Continuous Drug Manufacturing . . . 163
7.3 A Line of Well-Mixed Cells . . . 163
7.4 Dispersion and the Operating Window . . . 168
7.5 Instruments, Loops, Soft Sensor and Supervisor . . . 169
7.6 Nine Cases and a Campaign . . . 173
7.7 Technology in Practice . . . 182
7.8 Limitations . . . 184
7.9 Conclusion . . . 185
7.10 Exercises . . . 187

8 Energy Created by a Delayed Channel in a Bilateral Teleoperation Simulator 190
8.1 Introduction . . . 191
8.2 From Hot Cells to Wave Variables . . . 192
8.3 The Teleoperator, Its Channel and Its Site . . . 193
8.4 The Controllers, Period by Period . . . 198
8.5 The Gain Limit of Force Reflection . . . 200
8.6 Nine Cases and a Challenge . . . 202
8.7 Technology in Practice . . . 206
8.8 Limitations . . . 208

8.9 Conclusion . . . 209
8.10 Exercises . . . 211

9 Pension Funding Ratio Control Through a Year-Late Valuation and a Contribution Cap 214
9.1 Introduction . . . 215
9.2 Funding Rules From Actuarial Spreads to Recovery Plans . . . 216
9.3 A Fund of Two Stocks . . . 217
9.4 The Ratio, Its Gain and Its Lag . . . 221
9.5 Deciding Once a Year . . . 222
9.6 Fifteen and Thirty Years at a Time . . . 225
9.7 Technology in Practice . . . 231
9.8 Limitations . . . 233
9.9 Conclusion . . . 233
9.10 Exercises . . . 235

10 Inferential Moisture Control Near the Sticky Point of a Milk Powder Spray Dryer 238
10.1 Introduction . . . 239
10.2 Energy, Stickiness, Deposits and the Outlet Loop . . . 240
10.3 The Dryer as Modeled . . . 241
10.4 What the Outlet Temperature Stands For . . . 247
10.5 Implemented Controller . . . 248
10.6 Measured Behavior of the Nine Cases . . . 254
10.7 Technology in Practice . . . 261
10.8 Limitations . . . 262
10.9 Conclusion . . . 264
10.10 Exercises . . . 266

11 Controlling an Unmeasured Product Temperature in a Multi-Zone Cathode Calcination Kiln 269
11.1 Introduction . . . 270
11.2 Background . . . 271
11.3 Process Model . . . 272
11.4 Controller Design . . . 278
11.5 Results . . . 281
11.6 Technology in Practice . . . 285
11.7 Limitations . . . 287
11.8 Discussion . . . 288
11.9 Conclusion . . . 289
11.10 Exercises . . . 290

12 Stored Ammonia and a Sensor That Reads Both Ways in Diesel Urea Dosing 294
12.1 Introduction . . . 295
12.2 From Power Plant Stacks to Truck Tailpipes . . . 296
12.3 The Exhaust Line as Modeled . . . 297
12.4 The Doser as Implemented . . . 301
12.5 Ten Cases, Measured . . . 304
12.6 Evidence of Three Kinds . . . 310
12.7 Technology in Practice . . . 311
12.8 Limitations . . . 313
12.9 Conclusion . . . 314
12.10 Exercises . . . 315

13 Navigation Drift as Unimaged Seabed in an AUV Survey Simulator **319**
13.1 Introduction . . . 320
13.2 From Transponder Nets to Aided Dead Reckoning . . . 321
13.3 The Vehicle, the Seabed and the Instruments . . . 322
13.4 Navigation, Guidance and the Autopilots . . . 327
13.5 Measured Behavior of the Cases . . . 330
13.6 Technology in Practice . . . 336
13.7 Limitations . . . 337
13.8 Conclusion . . . 338
13.9 Exercises . . . 340

14 Summary and Outlook **344**
14.1 What the Chapters Showed . . . 344
14.2 Choosing a Method . . . 347
14.3 Where Practice Is Heading . . . 348
14.4 Engineering With AI Agents . . . 350
14.5 Gaps in Delay Compensation and State Estimation . . . 352
14.6 Exercises . . . 356

Answers to Selected Exercises **358**

Glossary **363**

Index of Control Topics **366**

Index of Applications **369**

About the Editor **372**

Volumes in This Series **373**

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: Delays, Late Measurements, and State Estimation

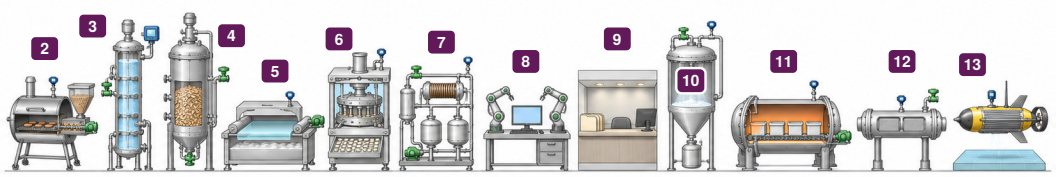

The chapters 2 pellet grill 3 distillation column 4 kraft digester 5 float glass line 6 tableting line 7 flow synthesis 8 teleoperator 9 pension fund 10 spray dryer 11 cathode kiln 12 diesel urea doser 13 survey vehicle

Notice that the units from the grill to the flow synthesis skid carry their product along a line before it is measured, while the kiln, the exhaust canister and the survey vehicle each hide a variable that no instrument reads. Illustration, not to scale.

IN PLAIN TERMS **Acting on News That Arrives Late**

The main idea. A controller can only fix what it has seen, and it often sees late. Material takes time to reach the meter, a test takes time to finish, and some things inside a plant are never measured at all.

The challenge. Acting hard on old news makes a loop swing, and waiting makes it slow. A copy of the plant kept in a computer can fill the gap, but only as well as the copy matches the plant.

Why it matters. The same few ideas steer a grill, a pension fund, a robot hand and a vehicle under the sea.

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 1.4.5, 1.4.6 and 1.4.8, Figure 1.10, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, with each number of the dosing example worked, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.9 with Tables 1.3 and 1.4, Section 1.5 with Figures 1.6 and 1.8, and Section 1.8 with Figure 1.10.

1.1 The Question This Volume Answers

A feedback loop can correct only what it has seen, and it sees the plant late. Material must travel from the actuator to the point of measurement, an analyzer must draw a sample and work on it, and a laboratory reports an hour or a day after the bottle was filled. Some variables are never measured at all, such as the temperature inside a brick, the ammonia stored on a catalyst or the position of a vehicle under water. This volume asks what each delay and each missing measurement costs, and how much of that cost a controller can win back with a model.

Two answers recur. The first is prediction. A model of the plant tells the controller what its own past moves will do before the measurement shows it, which is the idea behind the Smith predictor, internal model control and every predictive controller. The second is estimation. A model run beside the plant and corrected by whatever measurements exist supplies the variables that are late or missing. Kalman filters, moving horizon estimators, soft sensors and laboratory bias updates are forms of this idea. Both answers fail in the same way, when the model is wrong in a direction that the measurements cannot reveal.

Twelve chapters test these ideas, each on a plant with an app that runs in a browser. Chapter 2 and Chapter 3 control one variable through a fuel delay and through a slow analyzer. Chapter 4 and Chapter 5 decide a product quality hours or minutes before an instrument can read it. Chapter 6 and Chapter 7 follow material through pharmaceutical lines whose delays change with throughput.

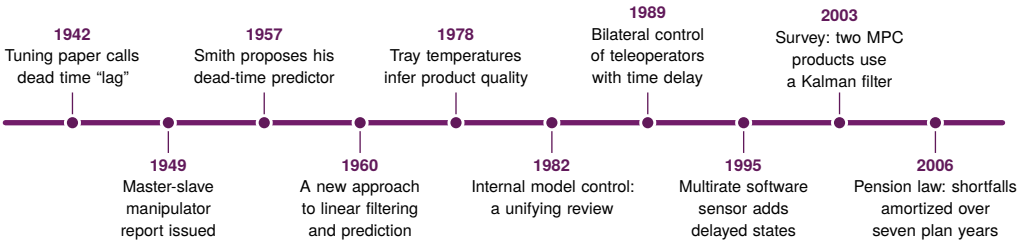

Figure 1.1 Milestones of prediction, estimation and late measurement in the systems of this volume. The dates are from [4] (1942), [5] (1949), [6] (1957), [7] (1960), [8] (1978), [9] (1982), [10] (1989), [11] (1995), [12] (2003) and [13] (2006). Notice that a remote hand and a pension law sit among the plants and the theory: each meets the same late news. Events are evenly spaced, not to a time scale.

Chapter 8 closes a force loop over a delayed channel, and Chapter 9 decides a contribution once a year on a valuation a year old. Chapter 10 and Chapter 11 control a moisture and a temperature that only a model can see, and Chapter 12 and Chapter 13 estimate a stored mass and a position from sensors that cannot report them directly.

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

- State the least error that any controller can leave after a load through a known dead time, and compare a tuned loop against it.
- Build a Smith predictor, say when it pays, and keep its delay right when the delay changes with flow.
- Account for a sampled or late measurement as added dead time, and choose the gain of a controller that acts once per result.
- Test whether a set of measurements can distinguish the disturbances and biases that matter.
- Design a Kalman filter or a laboratory bias update, align a late result with the instant it describes, and choose its weight.

This chapter develops each skill on one small example 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 dead-time compensation and state estimation in fuller detail [1, 2].

1.2 A Short History

The milestones of Figure 1.1 follow prediction and estimation through the kinds of system in this volume. Dead time was read from plant tests early. The second tuning method of a landmark paper of 1942 stepped the valve by hand and read the dead time from the response, though the paper called it lag, a word now kept for a first-order response [3, 4]. Prediction and estimation, the two answers of Section 1.1, then appeared within three years of each other, the predictor in 1957 and the filter in 1960. The two History boxes of this section follow each from its start.

Remote handling met the delay from the other end. A report of 1949 described a master-slave manipulator [5], and by 1989 the bilateral control of teleoperators with time delay was a subject of its own [10]. In the process industries the missing measurement was met with inference. A design of 1978 inferred unmeasured product qualities from secondary measurements, and on a simulated debutanizer four to six tray temperatures improved the steady-state control by as much as 400 % [8]. Internal model control was reviewed as a unifying structure in 1982 [9]. A multirate software sensor of 1995 took process outputs sampled at different rates into an iterated extended Kalman filter and carried the delayed ones as added states [11], the first of the three treatments of Section 1.4.7.

> ⧖ HISTORY **1960 Onward: The Filter That Plants Passed By**
>
> A paper of 1960 recast filtering and prediction in the state-transition form and derived an equation for the covariance of the optimal estimation error [7]. The filter, joined to a state feedback, soon became a standard approach in a wide range of fields, with an estimated 400 patents a year based on it. It had little impact in the process industries, whose units are constrained, nonlinear and hard to model, and whose engineers often had no exposure to its concepts or regarded them as impractical. A survey of 2003 found only two predictive control products that used a Kalman filter. The rest fed back a bias, the measured output minus the predicted one, which assumes a step disturbance at the output that stays constant [12]. Section 1.4.7 shows that bias as the simplest filter of all.

Late measurement is not confined to plants. A law approved in 2006 has a pension plan amortize each funding shortfall in level annual installments over seven plan years [13], and Chapter 9 spreads its deficit over a period that the reader chooses.

> ⧖ HISTORY **1957 Onward: From One Delay to Many**
>
> The predictor that Smith proposed in 1957 splits the model into a fast part without delay and a delay, and so removes the delay from the closed-loop characteristic equation when the model is exact. It was proposed for plants that are stable in open loop, and later work extended it to integrating and unstable ones. Plants with several inputs and outputs were harder, because each path from an input to an output may have its own delay. A design of 1973 assumed one delay on every channel. It kept the predictor's properties but is in general not very useful in practice. Later designs for many delays had to give up either the prediction of the output or a characteristic equation free of delay [6]. Chapter 5 uses the single-loop form on its thickness.

1.3 Notation and Conventions

The chapters share the symbols of the notation page. The setpoint is SP or r, the process variable is PV and the manipulated variable is MV, with controller settings K_c and τ_I. A process model has a gain K_p, a time constant τ_p and a dead time θ_p. In the running example the tank's time constant is τ and the pipe's delay is θ, and both change with the flow Q. A hat marks an estimate, so $\hat{x}$ is the estimate of a state x. A measurement drawn at the instant t_s is reported at $t_s + t_d$, where t_d is the reporting delay, and samples are drawn every T_s. The weight of a bias update is λ.

Two symbols need care across the series. The shared framework uses Q and R as the weights of the predictive controller, so this chapter writes the noise covariances of a Kalman filter as W for the process and V for the measurement, and a scalar pair of variances as q and r_v. The time constant τ_c of an IMC filter is the closed-loop time constant of the SIMC rule, and λ is reserved for the weight of a bias update. It also uses a for the controller action and h for the integration step, so neither letter appears here for anything else. The chapters keep their own notation. Chapter 2 writes θ_t for the auger's transport and θ_m for the delay in its controller's model, Chapter 3 writes τ_a for the analyzer lag and δ for its disturbance state, and Chapter 9 writes θ_v for the age of the valuation. Chapter 8 writes T for the round trip of its channel, and Chapter 13 writes b for the compass bias.

1.4 Delays, Late Measurements, and Estimators

1.4.1 The Running Example

The running example is a dosing station that makes a solution of fixed strength. Water flows at $Q = 1\,\mathrm{m}^3/\mathrm{min}$ through a well-mixed tank of $V_t = 5\,\mathrm{m}^3$, where a pump adds a concentrate in ratio to the measured flow. The dose ratio u, in percent, is the MV. Because the pump follows the flow, each percent adds 1 g/L at any flow, so $K_p = 1\,\mathrm{g/L}$ per percent. The solution then travels through a pipe of $V_p = 5\,\mathrm{m}^3$ to the point where the product is sampled and shipped. The tank outlet concentration x follows Equation (1.1), and the product y at the sample point is the tank outlet that left the tank one

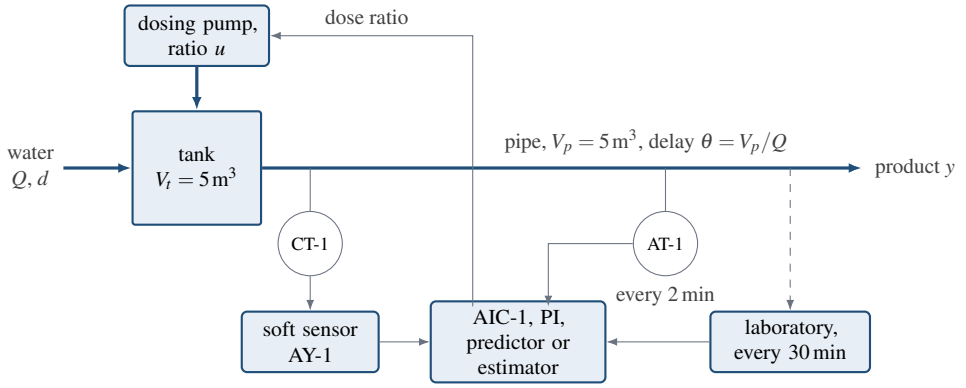

Figure 1.2 The Running Example. A Tank Doses a Concentrate Into a Water Stream, and a Pipe Carries the Product to the Sample Point. The Conductivity Cell CT-1 Reads the Tank Outlet at Once, the Analyzer AT-1 Reads the Sample Point Every 2 min With a Result One Cycle Later, and the Laboratory Reports a Sample From the Same Point 30 min After It Is Drawn.

pipe volume earlier.

$$\tau \frac{\mathrm{d}x}{\mathrm{d}t} = -x + K_p u + d, \qquad \tau = \frac{V_t}{Q}, \qquad y(t) = x(t-\theta), \qquad \theta = \frac{V_p}{Q} \tag{1.1}$$

where d, in g/L, is solute that arrives already dissolved in the water, a load disturbance. At the design flow $\tau = \theta = 5\,\mathrm{min}$, and the design dose of 50 % gives 50 g/L. The plant is first order plus dead time, with a lag and a delay of equal size.

Three instruments see the product, as shown in Figure 1.2. A conductivity cell CT-1 at the tank outlet reads at once, but it reads the dissolved salts of the water as well, so the soft sensor AY-1 that turns it into a concentration is right only while the water stays the same. An analyzer AT-1 at the sample point draws a sample every 2 min and reports it one cycle later. A laboratory receives a bottle from the sample point every 30 min and reports 30 min later, with no bias. The controller AIC-1 is a PI controller in the velocity form of Section 1.6.2, executed every 0.1 min, and the plant is integrated on a step of 0.01 min. Unless a subsection says otherwise, AIC-1 reads an ideal instrument at the sample point, so that the pipe is the only delay. Tuned by the SIMC rule of Volume 2 with $\tau_c = \theta$, it has $K_c = 0.5\,\%$ per g/L and $\tau_I = 5\,\mathrm{min}$.

1.4.2 Dead Time and Achievable Performance

A dead time sets a floor under the error that any controller can achieve. Suppose the water's solute rises by d at $t = 0$. The tank outlet starts to rise at once, but the product shows nothing until θ. Even a perfect controller that recognizes the upset at that instant and cancels it in the tank at once cannot touch the product before 2θ. The product therefore repeats the tank's open-loop rise for one dead time, and the integral of its absolute error (IAE) is at least Equation (1.2).

$$\mathrm{IAE}_{\min} = K_p d \left[\theta - \tau\left(1 - \mathrm{e}^{-\theta/\tau}\right)\right] \approx K_p d \frac{\theta^2}{2\tau} \quad (\theta \ll \tau) \tag{1.2}$$

For the example this is 1.84 g min/L per g/L of upset, and the product rises by at least 0.63 g/L. A setpoint step of size Δr has the floor $\theta\,\Delta r$, since nothing can reach the product sooner.

A PI controller with integral action pays more. Its integral must move the dose by d/K_p in the end, and in the velocity form that move equals K_c/τ_I times the integrated error, so the integrated error (IE) after a load follows Equation (1.3).

$$\mathrm{IE} = \frac{\tau_I}{K_c K_p} d = 2 K_p d\,\theta \quad \text{for SIMC with } \tau_c = \theta,\ \tau_I = \tau \tag{1.3}$$

The IAE equals the IE when the error keeps one sign and exceeds it when the loop overshoots. On the example the SIMC loop leaves an IAE of 10.25 g min/L, 5.6 times the floor, against an IE of exactly 10. The cost of a PI controller grows in proportion to the dead time, while the floor grows with its square while θ is short against τ. Repeating the run at dead times of 1, 2.5, 5 and 10 min with the tank unchanged gives an IAE of 2.00, 5.00, 10.25 and 21.4 g min/L against floors of 0.094, 0.53, 1.84 and 5.68, ratios of 21, 9.4, 5.6 and 3.8. A short delay is cheap in absolute terms, but it is where a PI controller leaves the most on the table relative to what is possible. A long delay costs every controller dearly.

A dead time also caps the bandwidth of the loop, near $1/\theta_p$, as Volume 2 shows from the phase margin. The same cap appears in forms that do not look like a PI loop. In Chapter 8 the loop is a force reflected through a channel with round trip T. Force reflected through a delay can make a contact chatter although the hand, the devices and the wall are each passive [10]. Through a long delay the loop's phase passes -180° at frequencies only $2\pi/T$ apart, so it stays stable at every long delay only if its gain is below one at all frequencies. A wall can then be rendered no stiffer than the hand that holds it, and the largest reflection gain falls from 1.00 at a round trip of 10 ms to 0.211 at 2.6 s. In Chapter 9 the decision interval and the valuation lag play the part of the dead time. Chapter 4 shows the limiting case, in which chips already past the heater when a production change arrives can no longer be reached, and no setting of the heater holds its kappa band through the step.

A fitted dead time is not always a transport delay. Chapter 2 fits a dead time of 4.7 min by the two-point method and 2.8 min by the tangent to a grill whose auger delay is 20 s, because the ignition, burning and the steel add lags in series. Chapter 3 fits 0.21 min to a column that contains no delay at all, the effect of thirty fast modes rounding the foot of the response. Chapter 11 finds no transport delay on its zone loops, so the SIMC rule cannot take $\tau_c = \theta_p$ there, and the scan and the interaction between zones limit the loop instead.

1.4.3 The Smith Predictor and Internal Model Control

A controller that knows the delay can stop waiting for it. The Smith predictor runs a model of the plant without its delay, $G_0(s) = K_p/(\tau s+1)$, beside a copy with the delay, and feeds back the measurement plus the difference between the two, as in Equation (1.4) [1].

$$y_p(t) = y(t) + \hat{x}(t) - \hat{x}(t-\theta_m), \qquad \tau\,\frac{\mathrm{d}\hat{x}}{\mathrm{d}t} = -\hat{x} + K_p\,u \tag{1.4}$$

where θ_m is the model's delay. When the model is exact, y and $\hat{x}(t-\theta_m)$ cancel except for the effect of disturbances, and the controller acts on the undelayed model output. Its characteristic equation then contains no delay. A PI controller with $\tau_I = \tau$ and $K_c = \tau/(K_p\tau_c)$ inside the predictor is exactly internal model control (IMC) with a first-order filter of time constant τ_c [9]. The setpoint response is the filter delayed by θ, and the load response follows Equation (1.5).

$$\frac{y(s)}{d(s)} = \frac{\mathrm{e}^{-\theta s}}{\tau s+1}\left(1 - \frac{\mathrm{e}^{-\theta s}}{\tau_c s+1}\right), \qquad \mathrm{IE} = K_p d\,(\theta + \tau_c) \tag{1.5}$$

The SIMC rule of Volume 2 is this design with the delay approximated, so a Smith predictor tuned with $\tau_c = \theta$ removes no more integrated error after a load than the SIMC PI. Its gain lies in the setpoint response and in the freedom to choose τ_c shorter than θ when the model can be trusted.

Table 1.1 and Figure 1.3 measure that freedom on the example. With $\tau_c = 2.5$ min the setpoint IAE falls from 10.93 to 7.46 and the load IAE from 10.25 to 7.50, and the dose moves four times further in the first instant. The load floor stays far off. The predictor cancels the tank's lag in its model, so an upset in the water decays with the open-loop time constant, and the load IE cannot fall below $K_p d\,\theta$ however short τ_c is. Studies of dead-time compensators find the same pattern, a clear gain over PI control when the delay dominates the lag and little gain on loads when it does not [14]. Reaching nearer the floor needs an estimate of the disturbance itself, the subject of Section 1.4.7, or a feedforward from Volume 3.

Table 1.1 Controllers on the Example at $\theta = \tau = 5$ min. IAE in g min/L After Unit Steps, the Product's Peak Rise After the Load, and the Longest True Delay at Which the Loop Stays Stable With the Model's Delay Held at 5 min

Controller	Setpoint IAE	Load IAE	Load peak (g/L)	Stable to (min)
PI, SIMC ($K_c = 0.5$, $\tau_I = 5$)	10.93	10.25	0.73	15.6
Smith predictor, $\tau_c = 5$ min	9.96	10.00	0.70	16.7
Smith predictor, $\tau_c = 2.5$ min	7.46	7.50	0.67	13.1
Smith predictor, $\tau_c = 1$ min	5.96	6.00	0.65	8.9
Least possible, Equation (1.2)	5.00	1.84	0.63	—

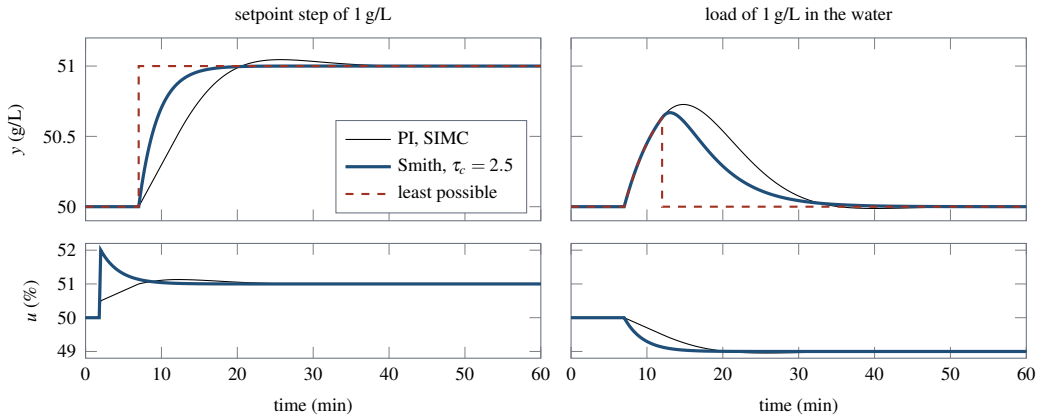

Figure 1.3 A Setpoint Step (Left) and a Load in the Water (Right) at 2 min, Design Flow. Top: Product y. Bottom: Dose Ratio u. Thin Solid: SIMC PI. Thick Solid: Smith Predictor With $\tau_c = 2.5$ min. Dashed: the Least Possible Response, Equation (1.2).

The price of the predictor is its dependence on the model's delay. The last column of Table 1.1 holds the model at 5 min and lengthens the true delay until the loop no longer settles. The SIMC PI tolerates 15.6 min, three times its design, and the predictor with $\tau_c = 1$ min only 8.9 min. A shorter τ_c buys speed with robustness to exactly the error a delay that varies with flow produces, as Section 1.4.4 shows.

Prediction through a delay appears in many forms in this volume. Chapter 5 turns its thickness loop into a Smith predictor with a forming model, and with the model off the same tuning hunts on the delayed gauge with a period of about 54 min. Chapter 2 tunes its loop by the SIMC rule, which it calls an IMC derivation, and its predictive controller carries the pellets in the auger as a shift register of the model's delay, which may be set away from the true one. Chapter 3 derives its PI settings from the IMC formula, and its predictive controller's model leaves out the analyzer lag, so its loop stops settling at a lag of 1.15 min. Chapter 9 projects the fund forward from the last valuation with the recorded contributions, a predictor across a delay of one year, and Chapter 4 predicts the kappa of every slug in the vessel from the heater temperature it received.

1.4.4 Delay That Varies With Flow

The pipe holds a fixed volume, not a fixed time. Material that enters it leaves once a pipe volume of flow has passed, as in Equation (1.6).

$$\int_{t-\theta(t)}^{t} Q(t')\,\mathrm{d}t' = V_p \tag{1.6}$$

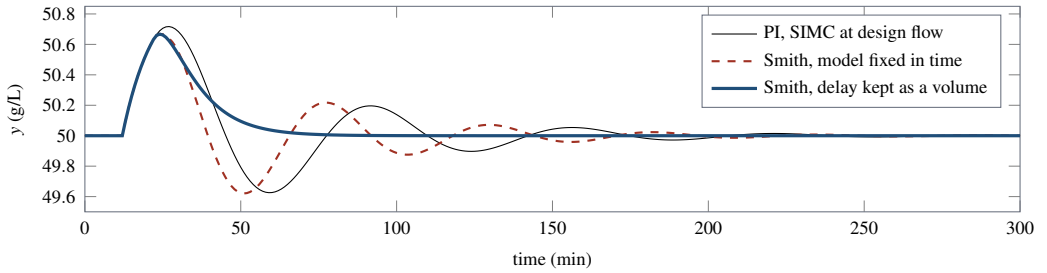

Figure 1.4 A Load of 1 g/L at 2 min With the Flow Halved, So That $\tau = \theta = 10$ min. Thin Solid: SIMC PI Tuned at Design Flow. Dashed: Smith Predictor With Its Model Fixed at 5 min. Thick Solid: Smith Predictor Whose Delay and Time Constant Follow the Measured Flow.

At constant flow this gives $\theta = V_p/Q$. When the flow halves, the delay and the tank's time constant both double to 10 min. A Smith predictor whose model keeps 5 min for both now carries a model error of the kind measured in Table 1.1.

Figure 1.4 repeats the load step at half flow. The SIMC PI stays stable, but its IAE triples from 10.25 to 30.9 g min/L, and the product swings below its setpoint before it settles. The predictor with its model fixed in time rings with a period of about 52 min and an IAE of 26.1. A predictor that measures the flow and keeps its delay as a volume, with its model's time constant and its integral time scheduled on V_t/Q, stays exact. Its response is the design response stretched to twice the time, with an IAE of 15.0, exactly twice the design value. Written in the volume of product that has passed, as Chapter 7 writes its diversion windows, the loop does not change at all.

Real lines mix as they carry. A chain of N equal well-mixed cells of total volume V passes a pulse as the Erlang distribution of Equation (1.7), with a mean residence time $\bar{t} = V/Q$ and a standard deviation of $\bar{t}/\sqrt{N}$ [15].

$$E(t) = \frac{(N/\bar{t})^N t^{N-1}}{(N-1)!}\, \mathrm{e}^{-Nt/\bar{t}} \tag{1.7}$$

One cell spreads a disturbance over a spread as large as its mean, and a hundred cells spread it over a tenth of the mean, close to plug flow. The distribution decides how long material affected by an upset keeps arriving, and so how much product a supervisor must hold back.

Several chapters follow material in this way. Chapter 7 writes its line as 225 cells, follows material by the volume that has passed the sample point, and diverts while the declared 0.1 and 99.9 % points of the coil's distribution, 0.7956 and 1.2328 coil volumes, overlap a window opened by a failing result. Chapter 6 runs a tracker with the plant's own tanks-in-series structure fed by measured rates, and its dryer turns a short feeder error into one or two packages of altered composition. Chapter 4 carries its chip column as slugs advanced on the measured flow, because a grid of 40 cells smeared the design kappa by 2.64 %, and a production rise of 15 % shortens the residence from 5.00 to 4.348 h. Chapter 11 times its feedforward to the window in which the affected material crosses each zone, since the material needs 1710 min to reach the last one. Chapter 5 tags each parcel of ribbon with the speed and thickness of its birth, so that the gauge, 15.2 min downstream, compares like with like.

1.4.5 Sampled and Multirate Measurements

An analyzer adds its own delay to the pipe's. A result drawn at t_s and reported at $t_s + t_d$ is held until the next one, so its age runs from t_d to $t_d + T_s$ and averages Equation (1.8).

$$\bar{t}_{\text{age}} = t_d + \frac{T_s}{2}, \qquad \theta_e = \theta + t_d + \frac{T_s}{2} \tag{1.8}$$

The effective dead time θ_e is what a tuning rule should be given. For AT-1, with $T_s = t_d = 2$ min, the mean age is 3 min and the loop sees $\theta_e = 8$ min. The SIMC loop tuned for the pipe alone then leaves

a load IAE of 16.1 instead of 10.25, against an IE that is still 10, so the error swings to the other side. The rule applied to 8 min gives $K_c = 0.3125$ % per g/L, an IE of exactly 16 and an IAE of 17.0, slower but with far less swing. Each minute of a result's mean age costs about as much as a minute of pipe.

The worst-case age also sizes the equipment downstream. A change that begins just after a sample is drawn is caught by the next sample and reported one cycle after that, so it is detected after $2T_s$ for an analyzer that reports one cycle late, plus any transport to the instrument. Chapter 7 writes this as $t_{det} = 2t_{run} + t_{tr}$ and requires its holding coil to cover it, so that no material passes the diverter before the analyzer has judged it.

The laboratory is too slow to read between results, so a loop on it acts once per result. Let the controller move the dose by g/K_p times each result's error. A result published now describes material made 35 min ago, before the previous move. The move before that is fully in it, since the tank settles to 99.3 % within the 25 min it has had. The error therefore obeys Equation (1.9), with the result two moves behind.

$$e_k = e_{k-1} - g\,e_{k-2} \quad \Longrightarrow \quad z^2 - z + g = 0 \tag{1.9}$$

The loop is stable for $g < 1$ and free of oscillation for $g \le 1/4$, where the two roots meet at $z = 1/2$. The simulated limit is $g = 1.001$, which the 99.3 % settling explains. With $g = 1/4$ the results after a load of 1 g/L stay within 0.05 g/L only from the one published 5 h after the load. With $g = 0.5$ the product overshoots by a quarter of the load and the results settle after 6 h, but the IAE is lower, 99.7 against 119 g min/L. Either way the cost is about ten times that of the pipe alone.

> **DID YOU KNOW? A Controller That Dozes on Purpose**
>
> One answer to a long dead time is the sample-and-hold PID. It acts on the error for a short moment, then holds its output while the correction travels through the process, and only then looks again. The holds slow it, but they allow bolder tuning [4].

Equation (1.9) recurs wherever a correction acts on information two decisions old. Chapter 9 decides once a year on a valuation a year old, and its deficit obeys the same equation with the fund's growth factor beside the first term, so its roots turn complex for spread periods under 3.85 years. Chapter 10 offers moisture control on the laboratory alone, with samples every 30 min reported 40 min later, and with the shipped tuning it oscillates with a growing amplitude. Chapter 7 runs its quality trim once per analyzer result, and Chapter 6 closes its moisture loop once per dryer cell, every 90 s, on a reading that summarizes 480 s of filling and drying. Chapter 4 has three clocks at once, a titrator every 10 min reported 10 min later, an analyzer every 15 min reported after 20, and a laboratory every 4 h reported an hour later. Combining measurements that arrive at different rates and ages is multirate estimation [11], the subject of the next three subsections.

1.4.6 Observability

An estimator can recover a state only if the measurements carry information about it. A linear model $x_{k+1} = Ax_k$, $y_k = Cx_k$ is observable when no initial state other than zero produces an output that stays at zero, which holds when the matrix of Equation (1.10) has full rank n, the number of states.

$$\mathcal{O} = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{n-1} \end{bmatrix}, \qquad \operatorname{rank} \mathcal{O} = n \tag{1.10}$$

A delay never makes a state unobservable. The pipe is a chain of states whose end is measured, and everything in it reaches the instrument in time. A delay makes information late, while a missing measurement can make it absent.

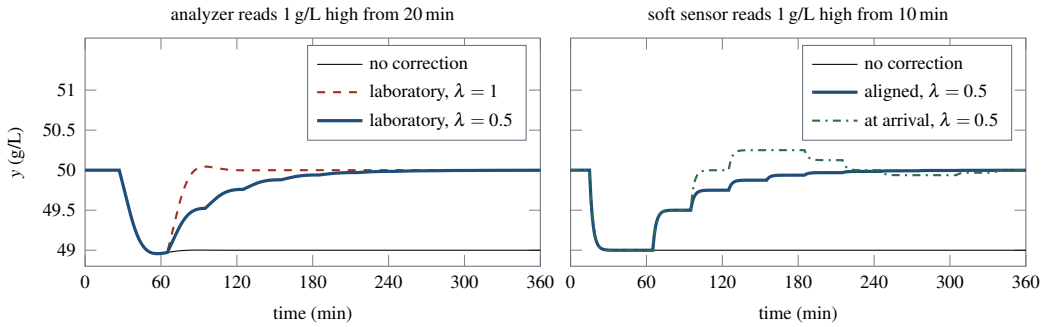

Figure 1.5 Product y When a Measurement Turns Wrong by 1 g/L. Left: Loop on the Analyzer, Whose Bias Is Learned From the Laboratory (Thin Solid: Not Corrected. Dashed: $\lambda = 1$. Thick Solid: $\lambda = 0.5$). Right: Loop on the Soft Sensor, Whose Bias Is Learned From the Laboratory Against the Soft Value of the Sampled Material (Thick Solid) or the Value When the Result Arrives (Dash-Dotted).

The example shows the difference. Sampled once a minute, the tank and five cells of pipe form six states. Add the water's solute d and a bias b of the analyzer, both constant unless they jump. With the analyzer alone, which reads $y + b$, the eight states give a rank of seven. The direction it cannot see is a tank settled at $x = K_p d$ together with $b = -K_p d$, a feed upset that has fully arrived hidden by an equal and opposite drift of the analyzer. Two constant states seen through one measurement are never both observable [16], so two disturbances that persist cannot both be estimated from one measurement. Adding the laboratory, which reads the same product without a bias, restores rank eight. The soft sensor has the same flaw, since its error from the water's salts is a second constant beside d, and the laboratory restores it too.

Figure 1.5 (left) shows what is at stake. At 20 min the analyzer begins to read 1 g/L high, and the loop on the analyzer, retuned for 8 min, drives the product down to hold the analyzer on 50 g/L. The product settles at 49 and stays there, and nothing on the analyzer trend reveals it. A correction from the laboratory pairs each result with the analyzer's reading of the same sample, so that the difference is the bias alone, $\hat{b} \leftarrow \hat{b} + \lambda (y_A - y_L - \hat{b})$. With $\lambda = 1$ the first result that contains the drift, published at 60 min, removes it, and the product is back within 0.1 g/L at 82 min. With $\lambda = 0.5$ the bias is learned by halves and the product returns at 161 min.

> ✦ DID YOU KNOW? **A Plugged Sample Line Reads Calm**
>
> Many analyzers stand away from the process and read a sample carried to them through tubes, filters and regulators. If the sample line plugs, the analyzer measures stale fluid and reports it as the process, with no way of knowing [4].

Observability decides several results in this volume. Chapter 13 finds the compass bias observable from position fixes taken on lines flown in opposite directions, and its estimate of a 1° bias reaches 1.005° by 300 s. Across a canyon, without a fix and without a measurement of the current, the drift is unobservable, and no tuning of the guidance recovers the unimaged strip. Chapter 4 estimates its analyzer's drift against the laboratory exactly as the example does, pairing reports by the instant their samples left the blowline. Chapter 11 corrects only the surface node of its brick model from the exit pyrometer, so the core error is almost untouched. Chapter 12 corrects its stored ammonia along the sensor's sensitivity to each cell's store, and near the bottom of the sensor's V the sensitivities cancel and the observer runs on its model alone. Chapter 10 adds its laboratory bias to the moisture only, so the sticky margin that its overrides read stays 0.9 K optimistic after a change of concentrate, and the laboratory never sees it.

Table 1.2 The Weight of a Bias Update as a Steady Kalman Gain: the Ratio of Drift to Noise It Assumes, the Results Needed to Learn 90 % of a Constant Error, and the Standard Deviation of the Estimate Relative to That of One Result

λ	q/r_v	Results to 90 %	$\sigma_{\hat{b}}/\sqrt{r_v}$
0.1	0.011	22	0.23
0.2	0.050	11	0.33
0.5	0.50	4	0.58
0.8	3.2	2	0.82
1.0	∞	1	1.00

1.4.7 Kalman Filtering and Moving Horizon Estimation

A Kalman filter is the estimator that weighs a model against noisy measurements in proportion to their uncertainties [7]. For a model $x_{k+1} = Ax_k + Bu_k + w_k$ and measurements $y_k = Cx_k + v_k$, with process noise of covariance W and measurement noise of covariance V, each period predicts and then corrects as in Equation (1.11).

$$\begin{aligned} \hat{x}_k^- &= A\hat{x}_{k-1} + Bu_{k-1}, & P_k^- &= AP_{k-1}A^\top + W \\ L_k &= P_k^- C^\top \left(CP_k^- C^\top + V\right)^{-1}, & \hat{x}_k &= \hat{x}_k^- + L_k\left(y_k - C\hat{x}_k^-\right) \end{aligned} \tag{1.11}$$

with $P_k = (I - L_k C)P_k^-$. When the noise statistics are constant, L_k converges to a steady gain, and most apps of this volume use that steady form. An unmeasured disturbance enters as an extra state that the model lets drift as a random walk, and the filter's estimate of it is what makes a predictive controller free of offset [17].

The simplest such filter is the bias update of the previous subsection. A bias that drifts as a random walk with variance q per result, read through noise of variance r_v, has the steady Kalman gain λ of Equation (1.12).

$$\hat{b}_k = \hat{b}_{k-1} + \lambda\left(y_k - \hat{y}_k - \hat{b}_{k-1}\right), \qquad \frac{q}{r_v} = \frac{\lambda^2}{1-\lambda}, \qquad \frac{\sigma_{\hat{b}}^2}{r_v} = \frac{\lambda}{2-\lambda} \tag{1.12}$$

The weight that chapters choose by hand therefore states a belief about how fast the bias moves against how noisy each result is. A constant error ε is learned as $\varepsilon[1-(1-\lambda)^n]$ after n results, and the noise passed into the estimate falls as λ falls. Table 1.2 gives both sides of the trade, and a run of the update on the example's laboratory with noise of 0.1 g/L gives standard deviations of 0.100, 0.055 and 0.030 g/L at λ of 1, 0.5 and 0.2, against 0.100, 0.058 and 0.033 from the equation.

A late measurement must update the state it describes. Three treatments are in use [18]. The model can carry the delay as extra states, a shift register of past values, so that the late measurement is an ordinary output of the augmented model. The filter can store its recent states and inputs, update at the instant the measurement describes and predict forward again, which is exact and costs a store of history. Or it can apply the measurement at its arrival as if it were new, which is cheap and wrong by whatever happened in between.

A moving horizon estimator (MHE) poses the same task as a least-squares fit over a window of the last N_e samples, as in Equation (1.13) [19].

$$\begin{aligned} \min_{x_{k-N_e},\, w} \; & \|x_{k-N_e} - \bar{x}\|^2_{\bar{P}^{-1}} + \sum_j w_j^\top W^{-1} w_j + \sum_j (y_j - Cx_j)^\top V^{-1}(y_j - Cx_j) \\ \text{s.t. } & x_{j+1} = Ax_j + Bu_j + w_j, \qquad x_{\min} \le x_j \le x_{\max} \end{aligned} \tag{1.13}$$

The first term, the arrival cost, summarizes the data before the window through a prior $\bar{x}$ and its covariance $\bar{P}$. With the filter's own prediction and covariance as the arrival cost and no bound active,

the window's solution is the Kalman filter's estimate. The MHE differs only when a bound binds, such as a concentration that cannot be negative, and it pays for that with an optimization at every sample [20]. A clamp applied after a recursive update, as several chapters apply to their biases, keeps the estimate in range but does not let the bound inform the rest of the state.

Chapter 2 runs an extended Kalman filter, Equation (1.11) linearized at every step, with a heat bias as a random walk, and trusts its probe to $r_v = 1\,\mathrm{K}^2$, above its noise, because a tighter filter passes the noise into the feed. Chapter 3 runs a steady filter with a disturbance state and offers an MHE whose reflux records differ from the filter's by at most 2.6×10^{-8}, because no bound ever became active. Chapter 5 runs a steady filter on its annealing zones. Chapter 13 stores 24 s of history, applies each acoustic fix at the instant it describes behind a chi-square gate, and finds that applying the fixes at their arrival more than triples the error.

1.4.8 Soft Sensors and the Laboratory Bias Update

A soft sensor infers a variable that is slow to measure from variables that are fast [8, 21]. Placed upstream of a delay, it takes the delay out of the loop. AY-1 reads the tank outlet, so the loop on it controls a first-order plant with no dead time, and the PI controller of the predictor's inner loop, $K_c = 2$ and $\tau_I = 5\,\mathrm{min}$, applies directly. A load in the water then costs an IAE of 2.50 g min/L, a quarter of the SIMC loop on the product, because the soft sensor sees the upset five minutes before the pipe delivers it. A soft sensor is a Smith predictor whose model output has been replaced by a measurement.

Its weakness is that it is only a model. When the water's salts change at 10 min, AY-1 reads 1 g/L high, the loop drives the product to 49 g/L, and without a correction it stays there, as the thin line of Figure 1.5 (right) shows. The laboratory corrects it with a bias update of the form of Equation (1.12), but only if each result is compared with the soft sensor's value for the same material, as in Equation (1.14).

$$\hat{b} \leftarrow \hat{b} + \lambda \left[y_L(t_s) - \hat{y}_{\mathrm{soft}}(t_s - \theta) - \hat{b} \right] \tag{1.14}$$

The bottle drawn at t_s holds product that left the tank at $t_s - \theta$, and its result arrives at $t_s + t_d$. Stored soft values make the alignment possible. Aligned, the bias error falls as $(1-\lambda)^n$, so with $\lambda = 1$ one result removes it and the product is back within 0.1 g/L at 71 min, and with $\lambda = 0.5$ at 157 min.

Compared instead with the soft value at the result's arrival, the update reads the effect of its own last correction as new error. The result was drawn before that correction reached the product, while the soft value already includes it. The bias error then obeys $\varepsilon_k = \varepsilon_{k-1} - \lambda\,\varepsilon_{k-2}$, the same equation as the laboratory loop of Section 1.4.5. With $\lambda = 0.5$ the product overshoots to 50.25 g/L and returns only at 216 min. With $\lambda = 1$ the roots lie on the unit circle, and the product swings between 49 and 51 g/L with a period of six results, three hours, indefinitely. Aligning a result with the instant it describes turns a loop with a delay into a filter without one.

> ✦ DID YOU KNOW? **Comparing at Arrival Undoes the Update**
>
> From the simulator: with the soft sensor reading 1 g/L high, the bias update with $\lambda = 1$ leaves an IAE of 50.0 when each result is compared with the material it sampled, and 313 when it is compared at its arrival, against 463 with no correction (Table 1.3).

Chapter 10 infers its powder moisture from the air temperatures with the plant's own drying model and corrects it with $\lambda = 0.5$, each result compared with the estimate stored when the sample was drawn, and not with the estimate of the moment. It also reports the trade of Table 1.2, with less powder off specification at $\lambda = 1$ than at 0.2. Chapter 7 runs a copy of its line as a soft sensor with a bias weight of 0.5 against predictions stored at each sample, and notes that a bias corrects a slow mismatch but not a short slug of unmeasured material. Chapter 4 uses weights of 0.8 for its analyzer drift and 0.1 for its predictor, with a gate that holds back an outlying report until later reports agree. Chapter 5 infers thickness at a camera 2.2 min after forming, against 15.2 min for the gauge, and cuts

Table 1.3 What Each Measurement Costs on the Example. Product IAE in g min/L After a Load of 1 g/L in the Water, and After a Measurement Error of 1 g/L Over a Run of 6 or 8 h

Measurement and controller	Delay seen by the loop	IAE
Load of 1 g/L in the water, design flow unless stated		
Least possible through the pipe	pipe, 5 min	1.84
Soft sensor at the tank, exact, PI with $K_c = 2$	none	2.50
Instrument at the sample point, Smith, $\tau_c = 2.5$	pipe, 5 min	7.50
Instrument at the sample point, SIMC PI	pipe, 5 min	10.25
Same at half flow: Smith in volume, Smith in time, PI	pipe, 10 min	15.0, 26.1, 30.9
Analyzer, PI tuned for the pipe, for 8 min	pipe and mean age, 8 min	16.1, 17.0
Laboratory alone, $g = 0.5$ per result	pipe and mean age, 50 min	99.7
Measurement reads 1 g/L high, laboratory bias update		
Analyzer from 20 min: none, $\lambda = 0.5$, $\lambda = 1$	over 6 h	325, 68.0, 39.5
Soft sensor from 10 min: none, $\lambda = 0.5$, $\lambda = 1$	over 8 h	463, 80.0, 50.0
Soft sensor, compared at arrival: $\lambda = 0.5$, $\lambda = 1$	over 8 h	89.7, 313

Table 1.4 The Delay or Missing Measurement of Each Chapter and What Recovers It

Ch.	Delay or missing measurement	What recovers it
2	Pellets ride an auger and burn before they heat, and the lid opens	Lid-aware hold of the feed, and a predictive controller whose filter carries the fuel in transit
3	Analyzer lag of up to 2 min on the distillate	Detuning by phase margin, and a predictive controller with a disturbance state
4	Four hours from heater to blowline, and three instruments on three clocks, one drifting	H-factor feedforward, a predictor of every slug reconciled with the analyzer and the laboratory
5	Thickness read 15.2 min after forming, and stress settled before the polarimeter	Smith predictor, inferred thickness at the camera, thermal-history estimator with a bias
6	Composition seen minutes after the feeders, and a feeder blind while it refills	Refill feedforward, a cascade sampled once per dryer cell, a material tracker
7	Delay and spread that change with throughput, and an analyzer that reports late	Diversion by volume and residence-time quantiles, a soft sensor with a laboratory bias
8	A channel with a round trip of up to 2.6 s, jitter and loss	Reduced reflection, wave variables, passivity control, model mediation
9	A valuation one year late, decided once a year	Spread period, a projected estimate, a predictive plan with a buffer
10	Moisture from the laboratory after 40 min, and the sticky margin never measured	Psychrometric soft sensor with a laboratory bias, feedforward, overrides
11	The brick's core never measured, and a day of transport	Two-node soft sensor and core cascade, feedforward timed per zone
12	Stored ammonia never measured, and a sensor that reads ammonia as nitrogen oxides	Model-based observer, storage setpoint scheduled on temperature
13	Position never measured under water, and acoustic fixes late and sometimes wild	Extended Kalman filter with fixes applied at their instant behind a gate

the rejected glass of a forming upset from 734 to 212 m^2. Chapter 11 infers the unmeasured core of each brick from a two-node model and holds that estimate with a cascade, and Chapter 6 trims its weight and hardness calibrations from each at-line test.

1.4.9 What Each Delay and Missing Measurement Costs

The section reduces to two questions that can be asked of any plant. What does each delay or missing measurement cost, and what recovers it? Table 1.3 answers both for the example, and Table 1.4 names the delay, the missing measurement and the remedy of every chapter. The final chapter returns to the second table with the measured result of each remedy.

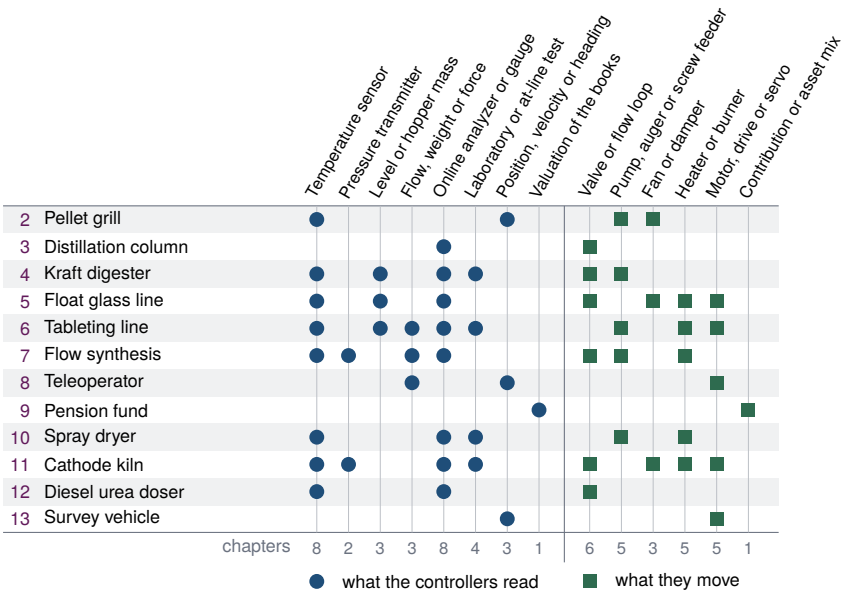

Figure 1.6 What the controllers of each chapter read and move: a round mark where a controller reads the instrument and a square mark where it moves the element, as the chapter's variables table and its technology section name them. A flow set through a flow station counts as a valve or flow loop, and a computed value is marked under the instruments it is computed from. Notice that the analyzers and the laboratory, the slow readings, appear in as many chapters as the thermometers.

Three patterns run through the table. Where the delay is a known transport, in Chapters 2 to 7, the remedy is to predict or to track material, and its quality depends on knowing the delay as the flow changes. Where the measurement is late but exists, a model or a faster instrument takes its place between results, and the slow measurement corrects its bias. Where a variable is never measured, in Chapter 11, Chapter 12 and Chapter 13, the estimate is only as good as the model, and observability decides which errors the remaining sensors can catch. Chapter 8 stands apart, because no model of the channel removes its delay and the remedy is to stay passive. Volume 9 develops the predictive controllers that combine these elements in one optimization.

1.5 Instruments and Final Elements in This Volume

A delay often sits in the instrument rather than in the plant. Figure 1.6 lists what the controllers of each chapter read and what they move. A temperature and an online analyzer or gauge lead the readings, with eight chapters each, and four chapters also wait for a laboratory or at-line test. The three with a laboratory in the loop, Chapter 4, Chapter 10 and Chapter 11, wait from forty minutes to six hours for a result. The final elements are spread more evenly: a valve or flow loop in six chapters, and a pump or feeder, a heater and a motor or drive in five each. In Chapter 9 the instrument is a valuation of the books and the final element a contribution.

An analyzer is slow by its principle. A chromatograph separates the species of a sample in time, so cycle times of minutes are not uncommon even in an online instrument, and some dead time is inevitable [4]. Many analyzers read a sample carried to them by a sample system, and problems with process analyzers often stem from those systems [4]. Such an analyzer stands behind the lag of Chapter 3 and, for liquids, the at-line result of Chapter 7. The top of Figure 1.7 shows a near-infrared analyzer, the kind that reads the blend of Chapter 6 in its feed frame at once, but only through a calibration. The bottom shows a pyrometer, which reads the surface of the bricks leaving

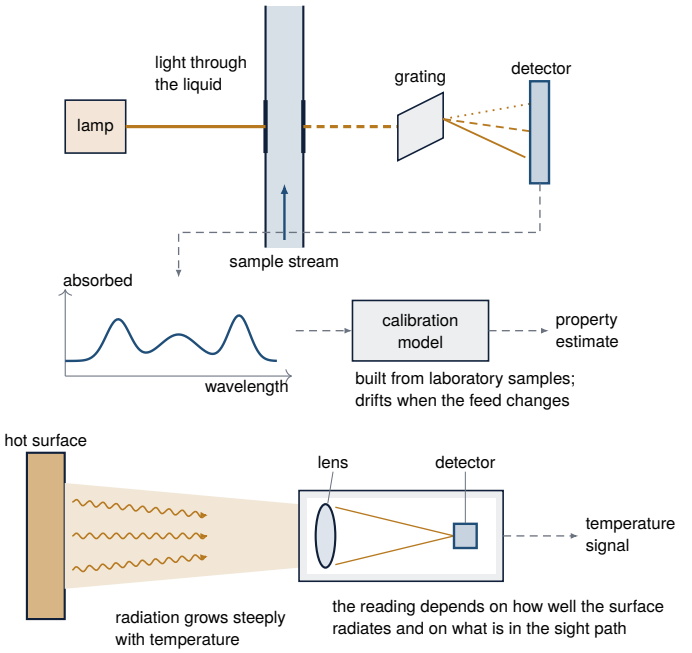

Figure 1.7 Top: principle of a near-infrared analyzer on a sample stream. Bottom: principle of an infrared pyrometer. Notice that the analyzer turns a spectrum into a property through a calibration model, a soft sensor built into the instrument, and that the pyrometer sees only the surface it is aimed at.

Chapter 11 and nothing of their core. Even a transmitter adds dead time when it reports sparingly, and battery-powered wireless transmitters may update only once in several minutes [4].

The estimator usually runs one layer above the loops it serves. A distributed control system places a supervisory level over localized controllers, with cycle times from milliseconds to minutes [22], and a survey of predictive control draws basic control executing about every second, constraint control about every minute and a local optimizer about every hour [12]. Figure 1.8 places the loops of this volume on their platforms. The apps run from the millisecond of Chapter 8 to the yearly decision of Chapter 9.

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 1.5 lists what a reader meets.

1.6.1 The Simulation Engine

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

Three clocks run at separate rates. The plant advances every step h. The controller executes every control period Δt, equal to h unless the app declares a slower controller, and holds its outputs between

Layer	What it does in this volume	Where it runs	How often
Optimizer or decision calendar	an economic optimum for the cook or the kiln, a yearly contribution, a predictive plan	**Local optimizer** over the loops; a **plan year** with its valuation date	30 min to 2 h in the plants; once a year in the fund
Estimator, predictor and supervisor	a Kalman filter, a soft sensor with its bias, a Smith predictor, a material tracker	**Supervisory level** of a DCS; an **application node** beside it; the **vehicle's computer**	each loop scan, up to 1 min; a predictive move every 1 s to 1 min
Base loops	velocity-form PI loops on the estimates, a force loop on each side of a channel	**Localized controllers** of a DCS; a **microcontroller**; a vehicle's **control unit**	1 ms on the teleoperator; 0.1 s to 5 s in most apps; 1.5 min in the kiln
Sensors and final elements	thermometers, analyzers, a laboratory, a valuation; valves, pumps, heaters, drives	wires, a sample line, a bottle to the laboratory, an actuary's report	results from minutes to 6 h late; a valuation a year late

Figure 1.8 The control layers of this volume, the platform that hosts each and how often each acts. Platforms are from [22] and [12] and from the chapters' technology sections, and the periods from the chapters. Notice that the estimator sits between the loops and the optimizer, and that the slowest measurement is often slower than every loop it serves.

Table 1.5 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.

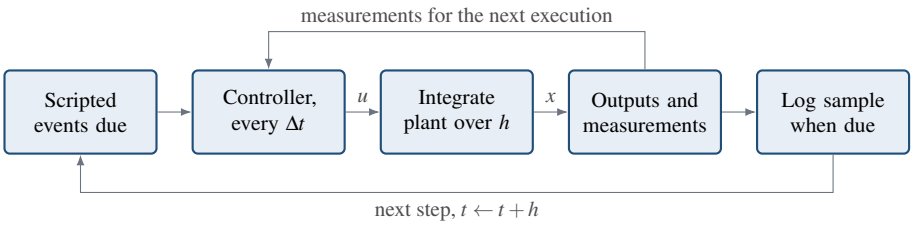

Figure 1.9 Work of the Engine in One Integration Step

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 1.9. 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 (1.15) to (1.18) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.15}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.16}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.17}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.18}$$

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 (1.18) drives I toward the value that reproduces the output applied, which is back-calculation [23]. With no integral action, Equation (1.18) 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 (1.19).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.19}$$

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 (1.20) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q\,(y_k - r_k) + \sum_{j} \Delta u_j^{\top} R\,\Delta u_j + \sum_{i} \rho_i\, s_i^2 \tag{1.20}$$

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 [24]. 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 [17]. 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 [19].

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

1.6.4 Scenarios, Progress and Data

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

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

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

Show Model opens the equations, live values of the variables, the editable parameters and the assumptions. The trends plot SP and PV above the MV for each loop, with actuator limits and marks at setpoint moves, disturbances, mode changes and scripted events. The scorecard reports for each loop the integral of the absolute error (IAE), overshoot, settling time, MV travel and share of time at

a limit, with the app's own measures, and it does not combine them into one score. IAE and travel accumulate from Reset, so two runs are compared at the same clock reading.

Export CSV saves the log as a text file with a header row of column names that carry their units and one row per logged sample. In most apps the time column is in seconds, minutes or hours according to the run length. The log holds the most recent 60 000 samples unless the app sets another size. Apps with the MPC option begin the file with comment lines that start with #, which a script must skip.

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

1.7 Working With AI Agents While Studying

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

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

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

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

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

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

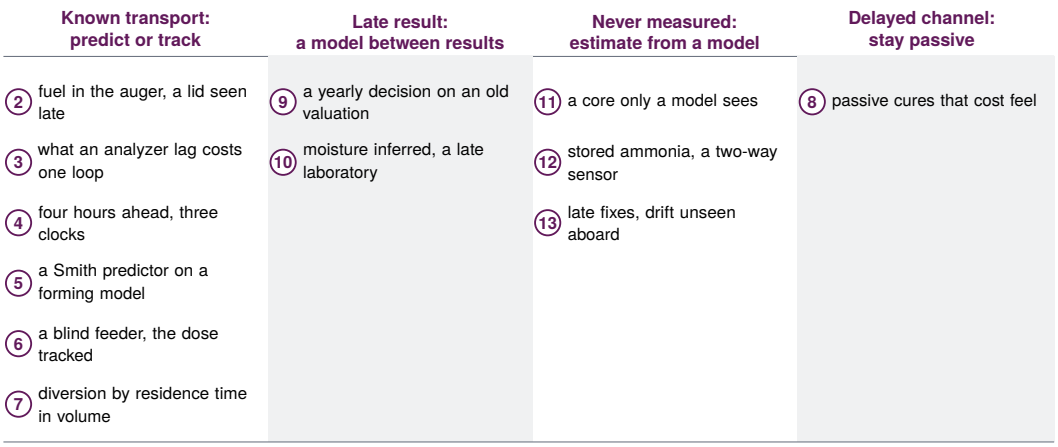

Figure 1.10 Method map of the volume. Each chapter is placed by the three patterns of Section 1.4.9, with the teleoperator apart, and labeled with its lesson, shortened from Table 1.6. The placement is the editor's reading of the two. Notice that most chapters predict through a known transport, and that the three whose variable is never measured rely on an estimator.

Table 1.6 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Wood pellet grill	Identification through a fuel delay, a lid the loop answers too late, and prediction of fuel in transit.	intermediate	3.5 h
3	Thirty-two stage column	What an analyzer lag costs one loop, and a predictive controller whose model omits it.	introductory	2.5 h
4	Continuous kraft digester	Feedforward through four hours, and three clocks reconciled into one estimate.	advanced	2 h
5	Float glass line	A Smith predictor on a forming model, and a stress forecast before it can be measured.	advanced	2.5 h
6	Continuous tableting line	A blind feeder, a cascade sampled per cell, and tracking of the dose.	intermediate	2 h
7	Continuous flow synthesis	Diversion by residence time measured in volume, and a soft sensor with a bias.	intermediate	2.5 h
8	Bilateral teleoperator	The gain limit of force through a delay, and passive cures that cost feel.	advanced	2.5 h
9	Defined-benefit pension fund	A decision once a year on a valuation a year old, with a small valve.	introductory	2.5 h
10	Milk powder spray dryer	A soft sensor of moisture corrected by a late laboratory, and an inferred constraint.	intermediate	2.5 h
11	Cathode calcination kiln	A core temperature only a model sees, held by a cascade on the estimate.	intermediate	3 h
12	Diesel urea dosing	An observer of stored ammonia behind a sensor that reads both ways.	advanced	2.5 h
13	Underwater survey vehicle	A Kalman filter with late fixes, and drift that nothing aboard can see.	advanced	2 h

1.8 Chapter Roadmap

Table 1.6 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 1.10 places the same chapters by the kind of information problem each centers on.

Chapter 3 and Chapter 2 are the first to read. The first isolates one loop and the cost of a measurement lag, and the second adds a true transport delay and a disturbance the loop sees too late. Chapter 5 then shows the Smith predictor on a real line. Chapters 4, 6 and 7 follow material along process lines whose delays move with throughput, and can be read in any order once Sections 1.4.4 and 1.4.5 are familiar. Chapter 10 and Chapter 11 form a path on soft sensors, and Chapter 12 and Chapter 13 a path on observers and Kalman filters, best read after Sections 1.4.6 to 1.4.8. Chapter 8 stands on its own for readers from robotics, and Chapter 9 for readers from finance, where the valuation is the late measurement.

KEY POINTS

- A controller sees the plant late, and a dead time sets a floor under the error that no controller can pass.
- A Smith predictor feeds back what its own moves will do, and it is only as good as its delay, which in a pipe is a volume, not a time.
- A sampled result adds its mean age to the dead time, and a loop that acts once per result two moves old is free of oscillation only for a gain of a quarter or less.
- A bias update is a steady Kalman filter, and each result must be compared with the estimate of the material it describes.

Where else this applies. Prediction across two dead times on a steam heat exchanger in Volume 2, Chapter 4, analyzer biases in a crude tower's predictive controller in Volume 11, Chapter 4, and predictive controllers that join prediction and estimation in Volume 9.

References

[1] J. E. Normey-Rico and E. F. Camacho, "Dead-time compensators: A survey," *Control Engineering Practice*, vol. 16, no. 4, pp. 407–428, 2008.

[2] D. Simon, *Optimal State Estimation: Kalman, H_∞, and Nonlinear Approaches*. John Wiley & Sons, 2006.

[3] J. G. Ziegler and N. B. Nichols, "Optimum settings for automatic controllers," *Transactions of the ASME*, vol. 64, no. 8, pp. 759–765, 1942.

[4] 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.

[5] R. C. Goertz, "Master-slave manipulator," Argonne National Laboratory, Argonne, IL, Report ANL-4311, Mar. 1949.

[6] J. E. Normey-Rico, T. L. M. Santos, R. C. C. Flesch, and B. C. Torrico, "Control of dead-time process: From the Smith predictor to general multi-input multi-output dead-time compensators," *Frontiers in Control Engineering*, vol. 3, p. 953768, 2022.

[7] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[8] B. Joseph and C. B. Brosilow, "Inferential control of processes: Part I. steady state analysis and design," *AIChE Journal*, vol. 24, no. 3, pp. 485–492, 1978.

[9] 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.

[10] R. J. Anderson and M. W. Spong, "Bilateral control of teleoperators with time delay," *IEEE Transactions on Automatic Control*, vol. 34, no. 5, pp. 494–501, 1989.

[11] R. D. Gudi, S. L. Shah, and M. R. Gray, "Adaptive multirate state and parameter estimation strategies with application to a bioreactor," *AIChE Journal*, vol. 41, no. 11, pp. 2451–2464, 1995.

[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] U.S. Congress, "Pension Protection Act of 2006," Pub. L. 109-280, 120 Stat. 780, Aug. 17, 2006, GPO text PLAW-109publ280, 2006.

[14] A. Ingimundarson and T. Hägglund, "Performance comparison between PID and dead-time compensating controllers," *Journal of Process Control*, vol. 12, no. 8, pp. 887–895, 2002.

[15] P. V. Danckwerts, "Continuous flow systems," *Chemical Engineering Science*, vol. 2, no. 1, pp. 1–13, 1953.

[16] 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.

[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] T. D. Larsen, N. A. Andersen, O. Ravn, and N. K. Poulsen, "Incorporation of time delayed measurements in a discrete-time Kalman filter," in *Proceedings of the 37th IEEE Conference on Decision and Control*, vol. 4, 1998, pp. 3972–3977.

[19] 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.

[20] E. L. Haseltine and J. B. Rawlings, "Critical evaluation of extended Kalman filtering and moving-horizon estimation," *Industrial & Engineering Chemistry Research*, vol. 44, no. 8, pp. 2451–2460, 2005.

[21] 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.

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

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

1.9 Exercises

1.1 Theory and calculation · 5 min

A mixing tank with $\tau = 8\,\text{min}$ and $K_p = 1\,\text{g/L}$ per percent feeds a pipe that delays its product by $\theta = 4\,\text{min}$. The water's solute rises by $d = 2\,\text{g/L}$.

(a) Find the least IAE of the product that any controller can achieve after this load. *Checkpoint:* 1.70 g min/L.

(b) Find the SIMC PI settings with $\tau_c = \theta$ and the integrated error they leave after the same load. *Checkpoint:* $K_c = 1\,\%$ per g/L, $\tau_I = 8\,\text{min}$, $\text{IE} = 16\,\text{g min/L}$.

1.2 Concept check · 2 min

A Smith predictor on a first-order plus dead-time plant uses a PI with $\tau_I = \tau$ and the IMC filter time constant $\tau_c = \theta$. Compared with the SIMC PI tuned with $\tau_c = \theta$ on the same plant, what does it do after a step load at the plant's input?

(a) It removes the load in half the integrated error
(b) It leaves the same integrated error
(c) It leaves no error once the delay has passed
(d) It leaves twice the integrated error

1.3 Theory and calculation · 4 min

A tank of $6\,\text{m}^3$ discharges through a pipe of $3\,\text{m}^3$ to a sample point. The flow falls from 1.5 to $0.75\,\text{m}^3/\text{min}$.

(a) Find the tank's time constant and the pipe's delay at both flows. *Checkpoint:* 4 and 2 min at 1.5, 8 and 4 min at 0.75.

(b) A bottle is drawn at the sample point at 10:00 at the low flow. At what time did its material leave the tank, and which soft-sensor value should a bias update compare with its result? *Checkpoint:* 9:56, the value logged at 9:56.

1.4 Theory and calculation · 3 min

An analyzer draws a sample every 6 min and reports it 6 min after the draw. The process behind it has $K_p = 0.5$, $\tau_p = 12$ min and $\theta_p = 3$ min. Find the mean age of the held reading, the effective dead time the loop sees, and the SIMC PI settings with $\tau_c = \theta_e$.

1.5 Theory and calculation · 6 min

A controller acts once per laboratory result and moves its output by g/K_p times the result's error. The plant settles between results.

(a) Each result describes material made before the previous move but after the one before it. Write the error recursion, and find the range of g for stability and the largest g without oscillation. *Checkpoint:* $e_k = e_{k-1} - g\,e_{k-2}$, $0 < g < 1$, $g \leq 1/4$.

(b) A faster laboratory reports each result before the next move, so it describes material made after the previous move. Find the same two limits and the gain that removes a load in one result. *Checkpoint:* $0 < g < 2$, $g \leq 1$, $g = 1$.

1.6 Concept check · 2 min

An analyzer at the end of a pipe reads the product plus a slowly drifting bias, and the water entering the tank carries an unmeasured solute that also drifts. Which addition lets an estimator tell the two apart once both have settled?

(a) An analyzer with a shorter cycle at the same point
(b) A Kalman filter with a longer horizon on the same analyzer
(c) An unbiased laboratory result on the same product
(d) A second identical analyzer beside the first

1.7 Theory and calculation · 4 min

A laboratory bias update uses the weight $\lambda = 0.2$. Treat it as a steady Kalman filter for a bias that drifts as a random walk.

(a) Find the ratio of the drift variance to the result's noise variance that this weight assumes. *Checkpoint:* 0.05.

(b) Find how many results it needs to learn 90 % of a constant error, and the standard deviation of the estimate in units of one result's. *Checkpoint:* 11 results, one third.

1.8 Concept check · 2 min

A soft sensor at a tank outlet is corrected by laboratory results drawn after a pipe, each available one sampling interval after its draw. The update compares each result with the soft sensor's value at the moment the result arrives, with $\lambda = 1$. After a step error of the soft sensor, what does the bias do?

(a) It removes the error with the first result that contains it
(b) It converges as $(1-\lambda)^n$, here in one step
(c) It grows without bound from the first result on
(d) It oscillates without decaying, with a period of six results

1.9 Theory and calculation · 3 min

An at-line analyzer runs a method of 4 min on each sample, and the sample needs 1 min to reach it. Product flows at 1.5 m^3/min past the sample point toward a diverter valve. Find the longest time from the start of a change to the result that reports it, and the volume of hold-up needed between the sample point and the valve.

1.10 ○ Concept check · 2 min

A moving horizon estimator uses the same linear model and noise weights as a steady Kalman filter, with the filter's prediction and covariance as its arrival cost. Its states carry bounds, none of which is active. How does its estimate compare with the filter's?

(a) It equals the filter's estimate

(b) It is smoother, because it fits a window

(c) It lags the filter's by half the window

(d) It is biased toward the arrival cost

CHAPTER 2

Wood Pellet Grill Simulator for Temperature Control Through a Fuel Delay

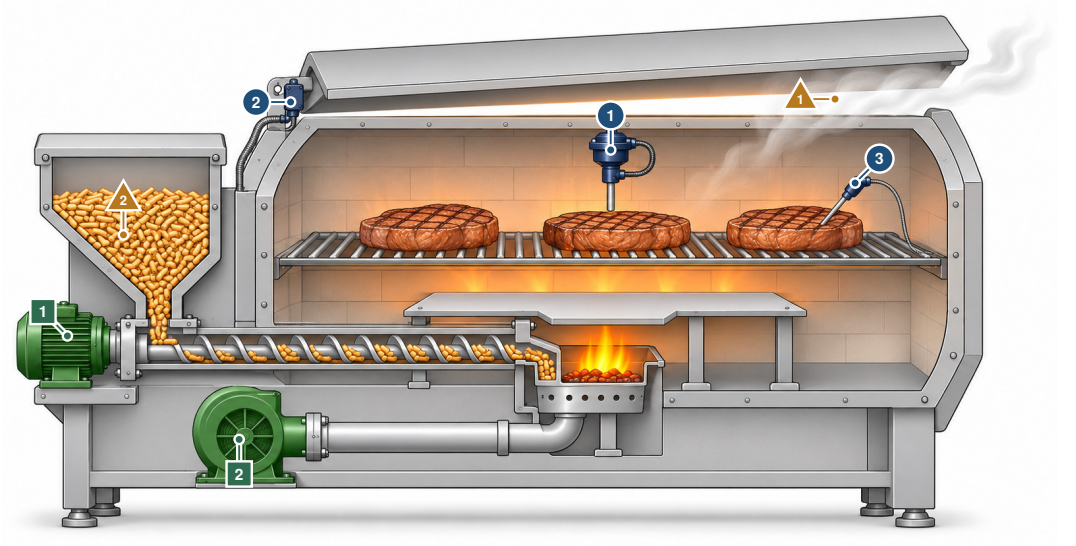

The controller moves 1 motor that turns the auger feeding the pot 2 fan blowing air into the pot

The controller reads 1 grate temperature probe 2 lid switch 3 meat probe

What disturbs it 1 lid opened by the cook 2 damp pellets that burn weaker

Pellets ride the auger from the hopper to the pot and turn to heat only after they catch and burn, while an open lid lets the hot air out within seconds. Illustration, not to scale.

IN PLAIN TERMS Steering a Fire That Is Fed Ahead

The main idea. A pellet grill holds its heat steady by feeding small wood pellets into a fire. A screw carries each pellet to the fire, and the pellet gives off its heat only minutes later, once it has caught and burned.

The challenge. When the lid opens, the hot air escapes within seconds, and a simple controller answers by feeding more pellets. Those pellets burn after the lid is shut, so the grill runs too hot.

Why it matters. A controller that keeps track of the fuel already on its way, and expects the lid to shut soon, holds the grill close to its setting.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the 1957 history segment (Section 2.1), the 1986 history segment and the fact on pellet stoves (Section 2.2), the try-it on the lid (Section 2.4), the try-it on damp pellets and the fact on the pulsed auger (Section 2.6) and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the eighteen measured cases (Section 2.6), Technology in Practice (Section 2.7), Limitations (Section 2.8) and the two open challenges

The control question A pellet grill's controller commands fuel, not heat, and each pellet becomes heat only after it has ridden an auger, caught on the bed and burned. When the lid opens and the chamber air is replaced within seconds, what should the loop do with the feed, and what must it know about the fuel already on its way?

Learning objectives

- Compute the fuel in transit, the heat it carries and the steady heat release from the pellet ledger.
- Fit a first-order plus dead-time model to a feed step, and explain why the dead time depends on the fitting method.
- Derive a SIMC tuning from the fit, and predict the feed a PI loop adds while the lid is up.
- Compare a standard PID, a lid-aware hold and an MPC on the same lid opening, and explain the role of the MPC's assumption about the lid.
- Interpret the heat bias that the extended Kalman filter estimates under damp pellets.

Prerequisites Sections 1.4.2 (dead time and a fitted model), 1.4.3 (prediction through a delay), 1.4.7 (Kalman filter with a disturbance state) and 1.6.2 (velocity-form PID). Section 1.6.3 is needed for the predictive controller.

The app https://apmonitor.com/apps/pellet-grill/

See also Volume 2, Chapter 4 for step tests and tuning on a temperature loop, and Volume 2, Chapter 3 for back-calculation at an actuator limit. Volume 9, Chapter 2 for a burner behind a dead time under predictive control, and Volume 9, Chapter 4 for an offset-free MPC whose estimator carries a fuel bias.

2.1 Introduction

A pellet grill is a thermostat on a wood fire, and its controller cannot add heat to the chamber, only fuel. The fuel it adds becomes heat in three stages. A pellet leaves the hopper on an auger and reaches the burn pot 20 s later, waits about 45 s for the hot bed to light it, and then burns over about 90 s. In the simulator of this chapter, a grill holding 225 °F over a stalled pork shoulder carries 19.6 g of pellets on their way to the fire at any moment, worth 267 kJ of heat, while its steel holds 2137 kJ above the outdoor air. A 5 % step in the feed moves the grate probe by 3.5 °F in five minutes and by 10.5 °F in ten.

The same shape recurs wherever a solid fuel is dosed into a fire. Internal model control of a biomass grate boiler has to be designed around a time-delayed input [1], a small wood pellet boiler emits less carbon monoxide under a nonlinear predictive controller of its airflow [2], and the compensators that handle a known delay, descended from the Smith predictor, work by predicting the delayed effect of the controller's own past moves [3], as Section 1.4.3 shows on the running example of Chapter 1. Model predictive control (MPC) carries that prediction as a model of the plant and is established practice across the process industries [4], and an estimated disturbance added to its model removes the offset that model error would otherwise leave [5].

> ⧖ HISTORY **1957: Predicting Past the Dead Time**
>
> Dead time hurts a feedback loop in two ways: the effects of the controller's moves and of the load take time to reach the measurement, and the controller acts now on an error that began earlier. A loop tuned for speed then oscillates, and one tuned against oscillation turns sluggish. In 1957 Smith proposed a model-based answer, now the best known and most widely used dead-time compensator in academia and industry. It splits the process model into a delay-free part and the delay, and gives the primary controller a virtual signal that anticipates the process output, which removes the delay from the closed-loop characteristic equation when the model is exact [6]. The MPC of this chapter carries the same idea in its own record of the feed in the auger.

The simulator is built against a common belief, that a large disturbance calls for the strongest response a loop can make. The lid of a grill is a disturbance that no controller can counteract while it

lasts, because the chamber air is replaced within seconds, and every pellet fed against it burns after the lid is shut. With the lid up for three minutes, a proportional-integral (PI) loop tuned from a step test raises the feed from 23 to 41.8 % and brings the grate 9.7 °F above its setpoint thirteen minutes after the lid shuts, and brisker tunings overshoot by 28.9 °F and 60.8 °F. A lid-aware variant that holds the feed while the lid switch reads open overshoots by 3.0 °F and settles 8.7 min after the lid shuts, and an MPC that estimates the fuel on the way and assumes the lid will shut at any moment overshoots by 2.7 °F and settles in 0.7 min. Told instead that the lid will stay open, the same MPC feeds 100 % and overshoots by 92.8 °F, so the controller's assumption about the disturbance matters as much as its algorithm.

However, the literatures that touch this plant treat its parts apart. Cooking is modeled as coupled heat and mass transfer inside the meat [7, 8], pellet appliances are characterized by their emissions in steady tests [9], and boiler control is developed for large grates without a lid, a stall or a recipe [1]. This chapter works on a browser simulator of a barrel pellet grill and measures how the fuel in transit, the lid, the meat and the recipe act on one controller, with the true state of the grill logged beside what its instruments report. The lumped model, the velocity-form PID, the offset-free MPC and the recipe logic are established methods. The results are measurements, eight answers to one lid opening and a dozen other cases on a single plant and seed.

Section 2.2 follows the domain from smokehouses and pellet furnaces to the pellet grill. Section 2.3 sets out a 41-state model whose pellet, energy and water ledgers close to integration error, with every equation and constant needed to build the simulator again. Section 2.4 measures the fuel on its way, a feed step with its two dead times, and the lid. Section 2.5 describes the supervisor that lights the fire and declares a flame loss, the standard and lid-aware PID loops, and a predictive controller that carries its own feed history through the auger's delay. Section 2.6 measures eighteen cases, from a feed step and a lid opening through six recipe programs to a set of faults and a graded challenge. Section 2.7 sets the loop on real hardware, and Sections 2.8 and 2.9 state the limits of the model and what the cases show.

2.2 Smokehouses, Pellet Fuel and the Pellet Grill

Cooking meat in smoke under a controlled temperature and humidity began as an industrial process. The smokehouse cooking of meat emulsions was studied in 1987 by raising the smokehouse temperature at 0.42 to 0.98 °C/min and its relative humidity at 1.1 to 3.9 %/min, and by measuring the product's water holding capacity, emulsion stability, shrinkage and texture, with the least shrinkage obtained under a rapid rise of both [10]. The quantities that study controlled, a chamber temperature program and the humidity around a wet product, are the ones a recipe controls in the grill simulated here. Smoke itself comes from incomplete combustion, and the polycyclic aromatic hydrocarbons it deposits on smoked meats are a food-safety concern of their own [11].

Wood pellets made the fire meterable. Densified biomass fuels were characterized with regard to their combustion behavior, with their moisture, ash and heating value among the properties that set how they burn [12]. Pellet stoves for homes were measured across fuel loads, pellet diameters and chimney drafts, and lowering the fuel load raised the emissions of carbon monoxide and hydrocarbons dramatically in a stove that burned continuously, while in a stove that ran intermittently at low load the carbon monoxide rose sharply because of high emissions at the end of each combustion cycle [9]. A grill smoking at 225 °F burns about 0.46 kg of pellets an hour, 23 % of its auger's capacity, so its fire spends most of a cook at low load, and its auger can be driven continuously or in pulses.

The pellet grill came out of the same fuel, through the two patents of the history segment below. The hopper, auger, pot, fan and baffle of the grill simulated here keep that arrangement, and its controller closes a temperature loop on a grate probe, with the recipe programs and lid logic that the problem invites.

> ⧖ HISTORY **1986: A Furnace That Feeds Itself, Then a Barbecue**
>
> A furnace patented in 1986 was meant to run for an extended period without supervision and give a controlled heat output from particulate solid fuel, such as wood waste pellets. A metering wheel took fuel from a hopper, an auger turned by an electric motor moved it along a feed tube to a fire pot, and the fuel was fed intermittently: with no demand for heat, an operating cycle could run, for example, for three minutes every half hour to keep the fire burning [13]. The pellet-fired barbecue patented in 1989 set the pot below the barbecue pan, fed it by an auger under the pan, forced combustion air through it, and covered it with a metallic baffle plate that radiated heat to the food and caught the drippings. With the feed stopped, its burner put itself out relatively rapidly [14].

Larger biomass plants met the same control problem earlier and at a larger scale. Internal model control of a 180 kW biomass grate boiler, designed around a time-delayed input, held high efficiency and low emissions across all load ranges in experiments [1], and a nonlinear predictive controller of the airflow, built on a Hammerstein model of a small wood pellet boiler, cut its carbon monoxide emissions by 35 to 50 % [2]. The grill adds three things those plants do not have, a lid that an operator opens at will, a load that stalls on its own evaporation, and a setpoint that a recipe moves on the clock and on the meat.

> ✦ DID YOU KNOW? **A Pellet Stove's Efficiency**
>
> Certified pellet stoves usually run at 70 to 83 % efficiency, with heating capacities from 8000 to 90 000 Btu per hour [15]. The grill of this chapter burns its pellets at an efficiency of 0.80 and releases 7.6 kW at full feed.

2.3 The Grill as Modeled

The simulated grill is a generic barrel pellet grill rather than a commercial model. A 9 kg hopper feeds an auger rated at 2 kg/h into a burn pot under a baffle and a sloped diffuser and drip tray, a combustion fan delivers up to 36 kg/h of air, a 250 W igniter rod lights the pot, and a single-wall barrel of 0.12 m^3 holds up to two cuts of meat and a water pan, as shown in Figure 2.1. The grate probe TT-101 supplies the controlled variable and the auger feed u_a is the manipulated variable, while the fan follows a schedule on the feed unless it is set by hand or given to the MPC.

Each variable in Table 2.1 carries its role in the loop, controlled (CV), manipulated (MV), disturbance (DV) or state, and the state vector has 41 entries. Ten describe the grill (the hopper, the two pot inventories, the bed, diffuser, shell and gas temperatures, the humidity ratio, the smoke index and the fan speed), seven describe each of two cuts (surface and core temperatures, accessible water, the probe's lag, and bark, smoke and tenderness indices), seven more hold the pan's water and enthalpy, a condensate film on the shell, the steam kit's boiler and reservoir and the lags of the grate and dew-point probes, and ten are ledgers of the pellets fed and burned, the heat released, the energy crossing the boundary and the water entering and leaving. The auger's transport is a delay line of 40 cells outside the integrated state. Units are seconds, kilograms, watts and °C, with kelvin inside the radiation terms, and the interface shows °F by default.

2.3.1 The Pellets' Journey

The hopper empties at the auger's rate, the pellets reach the pot a transport time later, and in the pot they wait as unignited fuel before they burn as the active bed, as shown in Equation (2.1).

$$\begin{aligned} \frac{\mathrm{d}M_h}{\mathrm{d}t} &= -q_{feed}, \qquad q_{feed} = q_{max}\, u_a\, \mathrm{S}\!\left(\frac{M_h}{0.05\,\mathrm{kg}}\right) \\ q_{arr}(t) &= q_{feed}(t-\theta_t) \\ \frac{\mathrm{d}M_u}{\mathrm{d}t} &= q_{arr} - r_{act}, \qquad \frac{\mathrm{d}M_b}{\mathrm{d}t} = r_{act} - r_{burn} \end{aligned} \tag{2.1}$$

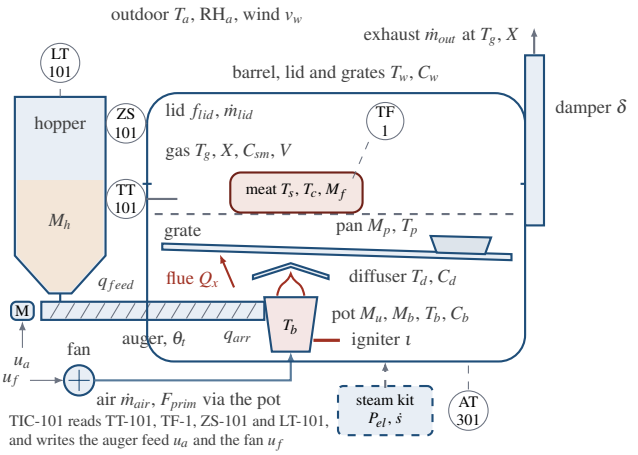

Figure 2.1 Wood pellet grill as modeled. The hopper holds M_h and the auger meters q_{feed}, which reaches the pot as q_{arr} after the transport time θ_t. The pot holds unignited pellets M_u and a burning bed M_b at T_b, lit by the igniter ι and burned in the fan's air $\dot{m}_{air}$, of which the share F_{prim} passes through the pot. The flue gas gives Q_x to the diffuser at T_d, and the chamber gas at T_g with humidity ratio X, smoke C_{sm} and volume V exchanges heat with the barrel, lid and grates at T_w, the meat (surface T_s, core T_c, water M_f) and the pan (M_p, T_p). The lid opens by f_{lid} and exchanges $\dot{m}_{lid}$, the exhaust leaves through the damper δ, and the steam kit and AT-301 are the experimental option. Not to scale.

Table 2.1 Variables of the Grill by Role, With Their Values at the Steady Operating Point of Case 1

Variable	Role	Case 1	Range
TT-101	CV, grate probe	225 °F	SP 150–550 °F
u_a	MV, auger feed	22.9 %	0–100 %, floor 8
u_f	MV, fan (scheduled)	31.5 %	0–100 %
f_{lid}	DV, lid, ZS-101	0	0–1
T_a, RH$_a$	DV, outdoor air	68 °F, 40 %	0–110 °F, 5–100
v_w	DV, mean wind, gusts	2.2 mph	0–30, 0–15 mph
b	DV, air lost at intake	0	0–80 %
d_p	DV, damp pellets	0	0–100 %
b_{TT}	DV, probe bias	0 °F	±20 °F
M_h	state, LT-101	9.0 kg at start	0–9
M_u, M_b	states, pot	5.7 g, 11.4 g	≥ 0
T_b, T_d, T_w	states, steel	643 °C, 376 °F, 147 °F	—
T_g, X, C_{sm}	states, gas	225 °F, dew 105 °F	RH < 100 %
T_s, T_c, M_f	states, per cut	158 °F, 152 °F, 60 %	—
M_p, H_p	states, pan	2 kg	0–4
10 ledgers	states	0	—

where $\mathrm{S}(x) = 3\bar{x}^2 - 2\bar{x}^3$ with $\bar{x}$ the value of x clipped to $[0, 1]$, so the auger starves smoothly as the hopper runs out, and $q_{max} = 2\,\mathrm{kg/h}$. The delay is a line of $\theta_t/\Delta t = 40$ cells at the 0.5 s physics step, filled with the mass fed in each step. In the alternative pulsed drive the motor runs at full speed for $u_a t_{cyc}$ of each cycle of $t_{cyc} = 15\,\mathrm{s}$, and the tube advances one cell only in a step in which the motor turns, so at a duty of 23 % a pellet crosses it in about $\theta_t/u_a = 87\,\mathrm{s}$. Fresh pellets catch only on a hot bed or at the igniter, and the bed burns no faster than its temperature, its inventory and its oxygen

allow, as shown in Equation (2.2).

$$\begin{aligned} r_{act} &= \frac{M_u}{\tau_{act}(1+2d_p)}\,\max\big(f_{ign},\,0.35\,\iota\big) \\ r_{burn} &= \operatorname{smin}\big(r_{kin},\,r_{O_2}\big), \qquad r_{O_2} = \frac{0.85\cdot 0.232\,\dot m_{air}}{1.286} \\ r_{kin} &= \frac{M_b}{\tau_{burn}}\,f_T\Big(\frac{\dot m_{air}}{\dot m_{ref}}\Big)^{0.4} \\ f_{ign} &= \mathrm{S}\Big(\frac{T_b - T_{ign} + 50}{100}\Big) \\ f_T &= \max\Big(\mathrm{S}\Big(\frac{T_b - T_{burn} + 50}{100}\Big),\,0.6\,\iota\Big) \end{aligned} \tag{2.2}$$

where temperatures are in °C, $\iota \in \{0,1\}$ is the igniter, $d_p \in [0,1]$ is the damp-pellet disturbance, 0.232 is the mass fraction of oxygen in air, 1.286 kg of oxygen burn a kilogram of pellets, and 0.85 of the fan's oxygen is usable. The smooth minimum $\operatorname{smin}(a,b) = m\,(1+(m/n)^{16})^{-1/16}$ with $m = \min(a,b)$ and $n = \max(a,b)$ is exactly zero when either rate is, lies within 4.3 % of the true minimum where the two cross, and is exact far from it. With $T_{ign} = 230\,^\circ\mathrm{C}$ and $T_{burn} = 170\,^\circ\mathrm{C}$, no fresh pellet catches on a bed below 180 °C and the bed goes out below 120 °C, so a fire starved of fuel cools, stops lighting new pellets and dies, and a cold pile of pellets stays a pile. The heat released is $Q_{comb} = \eta\,\mathrm{LHV}(1-0.25d_p)\,r_{burn}$ with $\eta = 0.80$ and $\mathrm{LHV} = 17\,\mathrm{MJ/kg}$, 7.6 kW at full feed, and the excess-air ratio of what burns is $\lambda = 0.232\,\dot m_{air}/(1.286\,r_{burn})$.

2.3.2 Air, Flue Gas and the Fire Bed

The fan's air follows its command through a lag, adds to a natural draft through the pot, and loses what a blocked intake or a closed damper takes, as shown in Equation (2.3).

$$\begin{aligned} \tau_{fan}\frac{\mathrm{d}f}{\mathrm{d}t} &= u_f - f \\ \dot m_{air} &= \Big[\dot m_{air,max}\,f\,\frac{\rho_a}{1.204}\,(1-b) + \dot m_{dr}\Big](0.6+0.4\,\delta) \\ \dot m_{in} &= s_k\big[\dot m_{leak}\,\delta\,(1+0.25\,v_w) \\ &\qquad + \dot m_{lid}\,f_{lid}^{1.5}\,(1+0.2\,v_w)\big] \end{aligned} \tag{2.3}$$

where ρ_a is the outdoor air density in kg/m^3, $\dot m_{dr} = 3\times 10^{-4}\sqrt{\max(0,T_b-T_a)/300}$ in kg/s is the draft, $\delta \in [0.2,1]$ is the damper opening, $\dot m_{in}$ is the air entering by leakage and through the lid, and $s_k = \sqrt{\max(0.02,\,(T_g-T_a)/(T_g+273.15))/0.23}$ is a stack factor equal to one at 225 °F in 68 °F air. The flue gas carries the dry air, the dry products and the water of combustion, 0.572 kg per kilogram burned, with 0.007 kg of ash left behind, at a dry flow $\dot m_{fg}$, humidity X_{fg} and heat capacity rate $\dot C_{fg}$. The primary air, $F_{prim} = 0.4$ of the fan's plus the draft, passes up through the bed with the products at the rate $\dot C_{prim}$, and the secondary air mixes in above it at ambient. The bed balance, the mixed flue temperature and the heat the flue gives the diffuser's underside are Equation (2.4).

$$\begin{aligned} C_b\frac{\mathrm{d}T_b}{\mathrm{d}t} &= Q_{comb} + P_{ign}\,\iota - \dot C_{prim}(T_b - T_a) \\ &\qquad - Q_{rb} - K_{pot}(T_b - T_a) \\ T_{mix} &= T_a + \frac{\dot C_{prim}}{\dot C_{fg}}(T_b - T_a) \\ Q_x &= \big(1 - \mathrm{e}^{-UA_x/\dot C_{fg}}\big)\,\dot C_{fg}\,(T_{mix} - T_d) \\ T_{in} &= T_{mix} - Q_x/\dot C_{fg} \end{aligned} \tag{2.4}$$

where $Q_{rb} = 0.9\,\sigma\,(0.012\,\mathrm{m}^2)(T_b^4 - T_d^4)$ is the pot mouth radiating to the diffuser, $K_{pot} = 0.3\,\mathrm{W/K}$ is the pot housing's loss, and T_{in} is the temperature at which the flue gas enters the chamber. In case 1 the bed runs at 643 °C and the fire releases 1733 W at $\lambda = 5.1$.

2.3.3 The Steel and the Chamber Gas

The diffuser, heated from below by radiation and flue gas, loses heat to the gas, the shell, the meat and the pan, and the barrel with its lid and grates loses heat outside, as shown in Equation (2.5).

$$\begin{aligned} C_d \frac{\mathrm{d}T_d}{\mathrm{d}t} &= Q_{rb} + Q_x - hA_d(T_d - T_g) \\ &\quad - Q_{dw} - Q_{df} - K_{dp}(T_d - T_p) \\ C'_w \frac{\mathrm{d}T_w}{\mathrm{d}t} &= Q_{dw} + hA_w(T_g - T_w) + Q_{cw} \\ &\quad - Q_{wf} - Q_{wp} - Q_{wa} - Q_{wo} \end{aligned} \tag{2.5}$$

where $C'_w = C_w + c_w M_{wc}$ includes the condensate film, $h = h_{in}(1 + 1.2 f_{lid})$ because an open lid stirs the chamber, $Q_{dw} = 0.9\sigma A_d(0.55)(T_d^4 - T_w^4)$, Q_{df}, Q_{wf} and Q_{wp} are the heat given to the meat and the pan, Q_{cw} is the latent heat of the film, and the losses outside are $Q_{wa} = (h_{out} + 2.5 v_w^{0.8})A_o(T_w - T_a) + \varepsilon_{out}\sigma A_o(T_w^4 - T_a^4)$ and, with the lid up, $Q_{wo} = f_{lid}\left[(6 + 3v_w)(0.3)(T_w - T_a) + 0.9\sigma A_{open}(T_w^4 - T_a^4)\right]$, with v_w in m/s and the areas $A_w = 1.6$, $A_o = 1.7$, $A_d = 0.30$ and $A_{open} = 0.35\,\mathrm{m}^2$. The chamber gas is well mixed and leaves at T_g, so every stream that enters is written at its own temperature against T_g, as shown in Equation (2.6).

$$\begin{aligned} C_g \frac{\mathrm{d}T_g}{\mathrm{d}t} &= \dot{C}_{fg}(T_{in} - T_g) + \textstyle\sum_i r_i c_{pv}(T_i - T_g) \\ &\quad + \dot{m}_{in,d}(c_{pa} + X_a c_{pv})(T_a - T_g) \\ &\quad + hA_d(T_d - T_g) - hA_w(T_g - T_w) \\ &\quad - \textstyle\sum_f Q_{gf} - hA_p(T_g - T_p) \end{aligned} \tag{2.6}$$

where $C_g = M_d(c_{pa} + X c_{pv})$, $M_d = PV/\big(R_d(T_g + 273.15)(1 + X/0.622)\big)$ is the dry gas the chamber holds, $\dot{m}_{in,d}$ is the dry part of $\dot{m}_{in}$, the sum over i carries the sensible heat of the vapor arriving from each source at its own temperature T_i, and Q_{gf} is the convection to cut f. At 225 °F the gas holds about 110 J/K against 21.2 kJ/K of steel, so it follows the steel and the fire within seconds.

2.3.4 Water in the Chamber

The humidity ratio is fed by the air, the water of combustion, the meat, the pan and the steam kit, and carried out with the exhaust, as shown in Equation (2.7).

$$\begin{aligned} M_d \frac{\mathrm{d}X}{\mathrm{d}t} &= \dot{m}_{fg}(X_{fg} - X) + \dot{m}_{in,d}(X_a - X) \\ &\quad + \textstyle\sum_f r_f + r_p + \dot{s} - r_{wc} \\ p_v &= \frac{PX}{0.622 + X}, \qquad T_{dew} = T_{sat}(p_v) \\ \mathrm{RH} &= \frac{p_v}{p_{sat}(T_g)} \le \frac{P}{p_{sat}(T_g)} \end{aligned} \tag{2.7}$$

where p_{sat} and its inverse T_{sat} are the region-4 equations of IAPWS-IF97 [16], r_f and r_p are the evaporation of the meat and the pan, $\dot{s}$ is the steam, and r_{wc} is condensation on the shell. Above the boiling point the saturation pressure exceeds the chamber pressure, so the relative humidity cannot approach 100 %. In case 1 it is 5.8 % and could not pass 77.6 % however much water went in, while the dew point is 105 °F. A smoke index follows $M_d\dot{C}_{sm} = 1000\,y_s r_{burn} - \dot{m}_{out,d}C_{sm} - 2\times10^{-4}M_d C_{sm}$, with a yield $y_s = 1 + 6\,\mathrm{S}((1.4 - \lambda)/0.6) + 3(1 - f_T)$ that rises when the fire is short of oxygen or smoldering. Its chemistry is not modeled.

Table 2.2 The Cuts. Mass m, Area A, Surface-to-Core Conductance K_{sc}, Surface Share ϕ_s, Water Share, Critical Share and Skin Coefficient k_m

Cut	m (kg)	A (m^2)	K_{sc} (W/K)	ϕ_s	water	crit.	k_m (kg/(s Pa m^2))
pork shoulder	4.0	0.155	1.8	0.30	0.34	0.22	0.9×10^{-8}
beef brisket	5.5	0.25	2.2	0.30	0.32	0.22	0.9×10^{-8}
spare ribs	1.6	0.15	4.0	0.35	0.28	0.32	0.9×10^{-8}
whole chicken[a]	1.8	0.085	1.6	0.30	0.30	0.22	2.0×10^{-8}
thick ribeye	0.6	0.045	1.4	0.35	0.25	0.40	0.9×10^{-8}
second butt	3.0	0.13	1.5	0.30	0.34	0.22	0.9×10^{-8}

[a]Water activity $a_w = 0.98$, all others 0.97. The water share sets $M_{f,0} = \text{water}\cdot m$, and M_{crit} is the critical share of $M_{f,0}$.

2.3.5 The Meat and the Stall

Each cut is an outer layer that exchanges heat and evaporates and a core that the probe reads, with the accessible water held in the core's heat capacity and migrating out to the surface, as shown in Equation (2.8).

$$\begin{aligned} C_s\frac{\mathrm{d}T_s}{\mathrm{d}t} &= Q_g + Q_d + Q_w - K_{sc}(T_s - T_c) - Q_o \\ &\quad - L_v(T_s)\,r + c_w\,r\,(T_c - T_s) \\ C_c'\frac{\mathrm{d}T_c}{\mathrm{d}t} &= K_{sc}(T_s - T_c), \qquad \frac{\mathrm{d}M_f}{\mathrm{d}t} = -r \\ r &= \omega\,\mathrm{smin}\big(k_m A\,[w^2 a_w\,p_{sat}(T_s) - p_v]^+,\; R_{mig}Aw\big) \end{aligned} \tag{2.8}$$

where $C_c' = C_{c,dry} + c_w M_f$, $w = \min(1, M_f/M_{crit})$, $[x]^+ = \max(0,x)$, $R_{mig} = 1.8\,\mathrm{kg}/(\mathrm{h\,m^2})$ caps how fast water reaches a wet surface, $\omega = 0.05$ for a wrapped cut and 1 otherwise, and $L_v(T) = L_0 - (c_w - c_{pv})T$ with $L_0 = 2.501\,\mathrm{MJ/kg}$. The heat arriving is $Q_g = \rho_h h A(T_g - T_s)$, $Q_d = \rho_r(0.85)\sigma A\beta(T_d^4 - T_s^4)$ from the diffuser below and $Q_w = \rho_r(0.85)\sigma A(0.9-\beta)(T_w^4 - T_s^4) + 0.3(A/0.155\,\mathrm{m^2})(T_w - T_s)$ from the shell and the grate, with $\beta = 0.2$ facing the diffuser, and a wrap sets $\rho_h = 0.9$ and $\rho_r = 0.5$ (both 1 otherwise). With the lid up the surface also radiates $Q_o = 0.3 f_{lid}\rho_r(0.9)\sigma A(T_s^4 - T_a^4)$ to the sky. When the gas is above the surface's dew point the surface condenses at $r = 9\times10^{-8}A(p_{sat}(T_s) - p_v) < 0$ and the condensate soaks into the core, and a cut resting on the counter loses $Q_o = 6A(T_s - T_a) + 0.9\sigma A(0.7)(T_s^4 - T_a^4)$ to still air. The cuts are listed in Table 2.2, with $C_s = \phi_s m c_m$, $C_{c,dry} = (1-\phi_s)mc_m - c_w M_{f,0}$ and $c_m = 3300\,\mathrm{J/(kg\,K)}$. No plateau is written anywhere in Equation (2.8). A wet surface cannot heat past the temperature at which its latent loss balances the heat it receives, and the core cannot pass the surface, so the stall emerges from the evaporation and ends as w falls, as in the coupled heat and mass transfer models of convection cooking [7]. The migration cap stands for the water that a shrinking protein network expels toward the surface [8].

Three qualitative indices, from 0 to 100, are labeled as such in the interface. Bark grows as $\dot{B} = \mathrm{S}((T_s - 60)/25)(0.25 + 0.75(1-w))(1 + 200\,\mathrm{S}((T_s - 105)/25))(100 - B)/9000\,\mathrm{s}$ on an unwrapped cut, smoke deposits while the surface is cool at $\dot{S}_m = C_{sm}(0.35 + 0.65\,\mathrm{S}((75 - T_s)/25))(100 - S_m)/(100\cdot 5400\,\mathrm{s})$, and tenderness follows an Arrhenius rate of 110 kJ/mol referenced to 85 °C above a 55 °C onset with a 5760 s time constant. None is a food-safety or doneness prediction.

2.3.6 The Water Pan, the Film and the Steam Kit

The pan is carried as an enthalpy, so that it boils at the pressure's boiling point instead of heating past it, as shown in Equation (2.9).

$$\begin{aligned}
\frac{\mathrm{d}M_p}{\mathrm{d}t} &= -r_p \\
\frac{\mathrm{d}H_p}{\mathrm{d}t} &= Q_{gp} + K_{dp}(T_d - T_p) + Q_{wp} - h_v(T_p)\, r_p \\
T_p &= \frac{H_p}{C'_p} - \left[\frac{H_p}{C'_p} - T_{boil}\right]^+ \frac{M_p}{M_p + 0.005\,\mathrm{kg}} \\
r_p &= 9\times 10^{-8} A_p \left(p_{sat}(T_p) - p_v\right)\psi_p \\
&\quad + \left[\frac{H_p - C'_p T_{boil}}{5\,\mathrm{s}\cdot L_v(T_{boil})}\right]^+ \psi_p
\end{aligned} \tag{2.9}$$

where $C'_p = 400\,\mathrm{J/K} + c_w M_p$, $\psi_p = M_p/(M_p + 0.01\,\mathrm{kg})$ applies to evaporation only, $A_p = 0.06\,\mathrm{m}^2$, $h_v(T) = L_0 + c_{pv}T$ and $T_{boil} = T_{sat}(P)$. An empty pan is steel alone and heats up. The condensate film on the shell grows at $9\times 10^{-8} A_w(p_v - p_{sat}(T_w))$ while the gas is above the shell's dew point and dries back at the same coefficient scaled by $M_{wc}/(M_{wc} + 0.002\,\mathrm{kg})$. The experimental steam kit is a boiler of $C_{sg} = 1.5c_w + 800\,\mathrm{J/K}$ whose 1 kW heater P_{el} tapers off between T_{boil} and $T_{boil} + 2\,^\circ\mathrm{C}$ and whose valve passes $\dot{s} = u_s \dot{s}_{max}\,\mathrm{S}((T_{sg} - T_{boil} + 4)/4)$ up to $\dot{s}_{max} = 1.1\,\mathrm{kg/h}$, both scaled by $M_{res}/(M_{res} + 0.02\,\mathrm{kg})$, with feedwater at 15 °C and a loss of 1.5 W/K.

2.3.7 Instruments and the Lid

Every probe lags its quantity to first order, as shown in Equation (2.10).

$$\begin{aligned}
\tau_{TT}\frac{\mathrm{d}T_{pr}}{\mathrm{d}t} &= T_g - T_{pr}, \qquad \tau_{food}\frac{\mathrm{d}T_{fp}}{\mathrm{d}t} = T_c - T_{fp} \\
\tau'_{dew}\frac{\mathrm{d}T_{dp}}{\mathrm{d}t} &= T_{dew} - T_{dp}
\end{aligned} \tag{2.10}$$

with $\tau_{TT} = 8\,\mathrm{s}$, $\tau_{food} = 15\,\mathrm{s}$ and $\tau'_{dew} = 20\,\mathrm{s}\,(1 + 5\varphi)$ lengthened by the fouling φ. TT-101 reads $T_{pr} + b_{TT}$ plus first-order colored noise of 0.3 °C standard deviation and 2 s correlation, rounded to 0.1 °C. The meat probes add 0.1 °C of noise and read to 0.1 °C, AT-301 reads $T_{dp} - 4\varphi$ with 0.2 °C of noise and is flagged invalid above its 180 °C head rating, the lid switch ZS-101 reads open whenever $f_{lid} > 0.05$, and the level LT-101 reads the hopper to 0.1 kg. The lid moves at $0.5\,\mathrm{s}^{-1}$ toward its target, fully open or shut in 2 s, and a timed opening shuts itself.

2.3.8 Parameters, Initialization and Numerics

The adjustable and hidden parameters are listed in Table 2.3, with the constants $c_{pa} = 1010$, $c_{pv} = 1900$ and $c_w = 4186\,\mathrm{J/(kg\,K)}$ and $R_d = 287.05\,\mathrm{J/(kg\,K)}$. Every constant is an illustrative teaching value, calibrated so that 225 °F with a stalled shoulder and a pan burns about 1 lb of pellets an hour and full feed takes the empty grill to about 590 °F. It is not a measurement of a commercial grill.

A warm case starts from the steady state the grill reaches at its setpoint with the meat and the pan where the case puts them. The plant itself is run for three hours in 2 s fourth-order Runge-Kutta steps (shorter when the fan or probe lag is set faster than 2 s) with the meat, the pan and the ledgers frozen, under a PI loop of $K_c = 1.5\,\%$ per °C and $\tau_I = 600\,\mathrm{s}$ on the true gas temperature with the fan on its schedule, and the auger is then filled at the steady feed. A cold case starts with everything at the outdoor temperature and the grill OFF. The engine of Section 1.6.1 integrates the plant by the fourth-order Runge-Kutta method with a fixed step $\Delta t = 0.5\,\mathrm{s}$, and the historian logs every 5 s. Over an hour of case 4 with the pan and the steam kit running, the pellets balance to $1.2\times 10^{-14}\,\mathrm{kg}$, the energy ledger closes to 1.8 J of 6.64 MJ released and the water ledger to $4.8\times 10^{-7}\,\mathrm{kg}$ of 0.79 kg.

Table 2.3 Model Parameters. Hidden Constants Are Shown but Cannot Be Edited in the Application

Symbol	Meaning	Value	Unit	Visibility
q_{max}	auger feed at 100 %	2.0	kg/h	adjustable
LHV	heating value as fed	17	MJ/kg	adjustable
η	combustion efficiency	0.80	—	adjustable
θ_t	auger transport time	20	s	adjustable
τ_{act}, τ_{burn}	ignition lag, burn time	45, 90	s	adjustable
$M_{h,max}$	hopper capacity	9	kg	adjustable
P_{ign}	igniter power	250	W	adjustable
T_{ign}, T_{burn}	ignition, burning thresholds	230, 170	°C	hidden
$\dot{m}_{air,max}$	fan air at 100 %	36	kg/h	adjustable
$\dot{m}_{ref}$	air at which τ_{burn} applies	12.6	kg/h	hidden
τ_{fan}	fan lag	2	s	adjustable
$\dot{m}_{leak}$, $\dot{m}_{lid}$	leakage, open-lid exchange	2.16, 216	kg/h	adjustable
C_b, C_d, C_w	bed, diffuser, shell	1.2, 4, 16	kJ/K	adjustable
UA_x	flue gas to diffuser	4	W/K	adjustable
h_{in}, h_{out}	convection in, out (calm)	8, 3	W/(m^2 K)	adjustable
ε_{out}	outside emissivity	0.80	—	adjustable
V	chamber volume	0.12	m^3	hidden
$M_{p,0}$, K_{dp}	pan water, tray contact	2, 1.5	kg, W/K	adjustable
P_{el}, $\dot{s}_{max}$	steam heater, valve	1000, 1.1	W, kg/h	adjustable
$M_{res,0}$	steam reservoir	5	kg	adjustable
τ_{TT}, τ_{food}, τ_{dew}	probe lags	8, 15, 20	s	adjustable
$T_{dp,max}$	dew-point probe rating	180	°C	adjustable
P	atmospheric pressure	101.3	kPa	adjustable

Halving the step from 0.5 to 0.25 s reduces the energy residual over the first 25 min of case 4 from 1.30 to 0.028 J, so the residual is integration error and not a bookkeeping term. The noise is seeded, seed 20260928, so every case replays identically.

2.3.9 Assumptions

One well-mixed chamber, one bed, one diffuser, one shell. A real barrel has a hot zone over the pot, a cooler end by the chimney and stratified air. The lumping keeps the fact the lesson needs, that the steel stores far more heat than the air the probe reads, and it fails wherever the position of the meat or the probe in the barrel matters.

Fuel as mass, not particles. The pellets move through a delay and two inventories, and ignition and burning are smooth thresholds of one bed temperature. Pellet size, fines, ash accumulation and the bed's geometry are outside the model, and the heating value, efficiency, time constants and the primary-air share are declared values rather than measurements of a pellet or a pot [12].

Constant transfer coefficients. Convection uses fixed coefficients with a wind term outside, and radiation fixed view factors between gray bodies. The calibration against a pound of pellets an hour makes the totals plausible, not the split between the paths.

A two-node meat. The stall comes from surface evaporation limited by a critical water content and a migration rate, and the constants are chosen so that the default cuts cook in about the times pitmasters quote. Conduction inside the meat, fat rendering, drip loss, collagen chemistry and the smoke ring are not modeled, where a transport model of cooking would resolve them [7, 8].

A single smoke index and a single humidity ratio. Smoke is transported and deposited but has no chemistry, so its yield under a starved fire is qualitative, while real smoke composition sets flavor and the polycyclic aromatic hydrocarbons deposited on the meat [11].

Repeatable numerics. Because the step is fixed and the random sequences seeded, every machine produces the same cook, a repeatability that no grill's microcontroller reading a thermocouple can offer.

2.4 Dynamic Character

2.4.1 Fuel on Its Way and Heat in the Steel

Two inventories set how the grill answers its controller. At 30 minutes into case 1 the tube holds 2.5 g of pellets, the pot 5.7 g that are catching and 11.4 g that are burning, 19.6 g in all and 267 kJ at η LHV, while the steel holds 2137 kJ above the outdoor air. At the 1733 W the fire releases, the pellets already delivered are about 2.6 minutes of fire that the controller cannot take back, and with the auger stopped the fire still releases 80 % of its heat 60 s later, 39 % after 150 s and 11 % after five minutes. The steel's store is eight times larger, which is why the chamber recovers by itself once a lid is shut. The shell passes 841 W to the outdoor air and the flue stream carries 334 W of sensible heat up the chimney, and the rest goes into the meat and the pan as latent heat.

2.4.2 A Feed Step and Two Dead Times

Case 3 holds an empty grill at 225 °F by hand, the auger at 18.2 % and the fan fixed at 27 %, and steps the feed up by 5 % five minutes in. The extra pellets reach the pot exactly $\theta_t = 20$ s after the step, and by 60 s the heat release has moved 12 % of the way to its new value, 91 % by five minutes. Because the manual feed and fan are rounded from the steady values of the warm start, the grill drifts up 0.65 °F over four hours without any step, so the response is taken as the stepped run minus an otherwise identical run with the scripted step removed, which also cancels the probe noise. The grate is then 3.5 °F up five minutes after the step, 10.5 °F at ten and 36.5 °F after four hours. A first-order plus dead-time (FOPDT) model, Equation (2.11), fits the response.

$$\frac{T_{TT}(s)}{U_a(s)} = \frac{K\,\mathrm{e}^{-\theta s}}{\tau s + 1} \tag{2.11}$$

The gain is $K = 4.05$ °C per percent (7.29 °F per percent) and the time constant $\tau = 15.5$ min, while the dead time is 4.7 min by Smith's two-point method and 2.8 min by the tangent at the steepest point. The two fits disagree because the delay is a transport time in series with the ignition lag, the burn time and the lag of the steel, and no single dead time represents that chain. Section 1.4.2 sets this fit beside the other chapters of this volume, among them Chapter 3, which fits a dead time to a column that has no delay at all. The SIMC rule, an internal model control (IMC) derivation, with a closed-loop time constant $\tau_c = \tau/2$ gives $K_c = \tau/(K(\tau_c + \theta)) = 0.31$ % per °C and $\tau_I = \tau = 15.5$ min on Smith's dead time [17], which the application ships as 0.3 % per °C and 15.5 min. The tuning is conservative for a setpoint change, and the lid measures what that costs.

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

2.4.3 The Lid

With the lid up in the default 2.2 mph breeze, the exchange of Equation (2.3) rises to 0.072 kg/s, about 73 W/K of fresh air against a gas that holds 110 J/K, so the chamber air is replaced within seconds. In case 4 the air falls 95.7 °F in three minutes while the shell loses 16.0 °F, and TT-101 bottoms at 128.7 °F. A PI loop's answer to the opening can be written down. The error reaches 53 °C as the lid shuts, so the proportional term raises the feed by $K_c e \approx 16$ %, and the error averages 50 °C over the three minutes, so the integral adds $(K_c/\tau_I)\,\bar{e}\,t_{lid} \approx 2.9$ %, which accounts for the measured rise from 23 to 41.8 %. The proportional part disappears within a minute of the lid shutting because the air recovers on the steel's heat, but the pellets fed meanwhile are already on their way and the integral keeps its 2.9 %. Over the 20 minutes after the lid opens the standard PID feeds 0.180 kg of pellets against 0.161 kg for the lid-aware PID and 0.166 kg for the MPC, and the extra 14 to 19 g, 190 to 250 kJ at η LHV, is what burns into the closed chamber.

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

Table 2.4 Supervisor States and Transitions. TT_0 Is the Grate Reading at Start

State	Actuators	Leaves when
OFF	feed 0, fan 0	Start is pressed
IGNITING	igniter on, feed 50 then 15 %, fan 40 %	TT $\geq$ $TT_0 + 14$ °C, or rising 3 °C in 2 min and TT $\geq$ $TT_0 + 8$ °C (RUNNING), or 900 s pass (FAULT)
RUNNING	the loop in charge, feed floor 8 %	probe failed, TT $>$ 316 °C (600 °F), auger jammed 60 s, or flame loss (FAULT), Shut down (SHUTDOWN)
SHUTDOWN	feed 0, fan 60 %	TT $<$ 60 °C after 300 s (OFF)
FAULT	feed 0, fan 50 % for 10 min	Start is pressed

TRY IT IN THE APP 4 · A Three-Minute Lid Opening

Load the scenario and press Run with the Standard PID as shipped. Watch the feed strip while the lid is up from 20 to 23 minutes: the feed climbs to 41.8 %. Then watch the grate after the lid shuts. It recovers on the steel's heat, keeps rising, and peaks 9.7 °F over the setpoint thirteen minutes later.

2.5 Implemented Control System

The controller scans every $\Delta t_c = 0.5$ s, one physics step. The supervisor, the recipe, the auger drive and the fan act every scan, the PID and the estimator every second, and the MPC every 10 s. Each physics step runs, in order, the controller, the integrator, the auger's transport with the lid, and finally the outputs and the log. Only the MPC's labeled perfect-state teaching mode ever reads a true temperature.

2.5.1 Supervisor

The supervisor of Table 2.4 decides whether there is a fire to control. Start lights the igniter and primes the pot at 50 % feed for 60 s, then trickles at 15 % with the fan at 40 %, and hands over to the loop once TT-101 has risen 14 °C (25 °F) above its value at Start, or has risen 3 °C over two minutes of samples taken every 10 s while standing at least 8 °C above it. From cold, case 2 hands over 5.0 minutes after Start with the grate at 83 °F. A flame loss is declared from a five-minute memory of the grate and the feed, when the lid has been shut for five minutes, the grate has fallen more than 8 °C, the mean feed is above 30 % and the grate is more than 15 °C below the setpoint, because a loop cannot tell a dead fire from a heavy load. A fault stops the feed and runs the fan at 50 % for 10 min, and nothing returns to RUNNING without an explicit Start.

2.5.2 TIC-101 and the Lid-Aware Hold

The standard loop is the velocity-form PID of Section 1.6.2, run every second on TT-101 with the derivative on the filtered measurement, as shown in Equation (2.12).

$$
\begin{aligned}
e_k &= T_{SP,k} - T_{TT,k}, \qquad P_k = K_c e_k \\
D_k &= -K_c \tau_D \dot{\tilde{T}}_{TT,k} \\
s_k &= s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\
\hat{u}_k &= s_k + I_{k-1}, \qquad \bar{u}_k = \min\big(100,\ \max(u_{fl},\ \hat{u}_k)\big) \\
I_k &= I_{k-1} + \frac{\Delta t_{PID}}{\tau_I}\big[K_c e_k + (\bar{u}_k - \hat{u}_k)\big]
\end{aligned}
\tag{2.12}
$$

where $\tilde{T}_{TT}$ is TT-101 through a first-order filter of $\max(1\,\text{s}, \tau_D/8)$, $u_{fl} = 8\,\%$ is the feed floor while RUNNING, and $\Delta t_{PID} = 1$ s. A colder grate raises the feed, so the loop is reverse acting. The symbols map onto Equations (1.15) to (1.18) of Chapter 1 as $s_k = S_k$, $\hat{u}_k = v_k$ and $\bar{u}_k = u_k$, with no output bias, a tracking time $T_t = \tau_I$ and no rate limit, so anti-windup acts on the integral alone [18]. In MANUAL, and whenever another controller or the supervisor has the auger, the integral is reset every second to $I_k = u - s_k$ for the feed u actually applied, so every transfer is bumpless. The defaults are $K_c = 0.3\,\%$

per °C, $\tau_I = 930\,\text{s}$ and $\tau_D = 0$, from sliders of 0.05 to 5 % per °C, 1 to 60 min and 0 to 240 s. The MPC's horizon runs from 5 to 40 min, its move suppression from 0 to 20 and its model delay from 0 to 120 s.

The lid-aware variant runs the same loop until ZS-101 reads open. It then holds the feed at $u_{hold} = \min(\bar{u}_{120}, u_{lid})$, where $\bar{u}_{120}$ is the feed averaged with a 120 s first-order filter that is frozen while the lid is up and $u_{lid} = 30\,\%$ is a cap, keeps holding for 45 s after the lid shuts while the air recovers, and does not call the PID meanwhile, so the integral resumes where it was before the lid opened. The derivative filter is reset to the current reading on release, so the step the air made during the hold produces no derivative kick.

2.5.3 Fan and Auger Drive

By default the fan follows a schedule on the feed that the fire has been getting, $u_f = \min(1,\ (12 + 0.85\,\bar{u}_{90})/100)$ with $\bar{u}_{90}$ the feed through a 90 s first-order lag, so that the air follows the fuel rather than the command and the fan's own effect on the grate does not mask the fuel delay during a step test. Its gain on the grate can take either sign. With the fire short of oxygen more air releases more heat, and with the fire short of fuel more air carries more heat out of the chimney. The continuous auger meters the feed as a mean rate, and the pulsed drive of Section 2.3 turns it into an on-time in every 15 s cycle.

2.5.4 Model Predictive Control

The MPC runs every $\Delta t_m = 10\,\text{s}$ over a horizon of 20 min, $N_p = 120$ steps. Section 1.6.3 describes the option in general. Here the controller's model is the plant's own Equations (2.1) to (2.10) with the declared parameters, the defaults of Table 2.3 fixed when the application loads, the cuts on the grate, the mean weather and no damp pellets or blocked intake, run on the estimate rather than on the truth. The estimated state is $\hat{x} = [M_u, M_b, T_b, T_d, T_w, T_g, T_{s1}, T_{c1}, T_{s2}, T_{c2}, T_{pr}, d]$, where d is a heat bias in watts added to the gas balance of Equation (2.6). The meat's water, the pan, the humidity and the film are carried open loop by the same model, and the hopper comes from LT-101. The pellets in the auger are known without measurement, because the controller keeps its own record of every feed it commanded and passes it through its model's delay, a shift register of $n_d = \theta_m/\Delta t_m$ cells with $\theta_m = 20\,\text{s}$ by default, which may be set away from the plant's θ_t.

At each execution the model is linearized at the estimate by finite differences and discretized exactly for a zero-order hold of Δt_m with the matrix exponential of the augmented system $[A\ B\ c;\ 0\ 0\ 0]\Delta t_m$, with the lid's effect taken as the secant between shut and open. The prediction of TT-101 over the horizon is $\hat{y} = y_{free} + Sz$, where z holds twelve blocked moves of 1, 1, 1, 1, 2, 2, 3, 4, 6, 9, 15 and 75 steps and two slacks, and the quadratic program is Equation (2.13).

$$
\begin{aligned}
\min_{z}\ J = {} & \tfrac{1}{2}\textstyle\sum_{k=1}^{N_p} w_k\,(\hat{y}_k - r_k)^2/(5\,\text{K})^2 \\
& + \tfrac{1}{2} w_{\Delta u} \textstyle\sum_j (u_j - u_{j-1})^2 + w_f\, q_{max} \textstyle\sum_j \tau_j u_j \\
& + \tfrac{1}{2}\left(50\, s_{hi}^2 + 10^4 s_{lo}^2\right) \\
\text{s.t.}\ \ & u_{fl} \le u_j \le 1, \qquad |u_j - u_{j-1}| \le 0.25 \\
& q_{max} \textstyle\sum_j \tau_j u_j \le \hat{M}_h - 0.02\,\text{kg} \\
& \hat{y}_k \le r_k + 15\,\text{K} + s_{hi}, \qquad s_{hi},\ s_{lo} \ge 0
\end{aligned}
\tag{2.13}
$$

where u_j is the feed as a fraction in block j of τ_j seconds with u_0 the feed applied now, r_k is the setpoint, previewed from the recipe when it runs on the clock, $w_{\Delta u} = 10$, $w_f = 0$ by default, and the ceiling is imposed on every other step. The tracking weight is $w_k = 0.3$ for the first three steps and 1 after them, and zero while an announced or programmed lid opening has the lid up and for 30 s after it. An optional stage-hold policy adds the soft bounds $T_{SP} - 3\,°\text{C} - s_{lo} \le \hat{y}_k \le T_{SP} + 3\,°\text{C} + s_{lo}$ until the current stage ends. The fan may be a second decision variable, pulled toward its schedule

and held inside the envelope $u_f \geq 0.554\,u_a$, and the steam valve a third, with the dew point a second output. The dense dual active-set method of Goldfarb and Idnani [19] solves the program, and the first move is applied. If an execution fails, the last plan is followed, and on the third failure in a row the lid-aware PID takes over from the MPC's last output. On case 4 the program has 14 variables and 111 constraint rows, and over 361 executions it needs 0.07 active-set steps on average and 7 at most, with none failing and a median of about 0.5 ms of wall-clock time for a whole execution in the V8 JavaScript engine of the Chromium browsers.

What the MPC assumes about an open lid is a setting. By default it assumes the lid will shut at any moment, so that it predicts the chamber's recovery on the steel's heat and feeds only what the steel and the meat lost. It can instead assume that the lid stays open for the whole horizon, or that an opening lasts as long as a cook announced it would.

The estimator is an extended Kalman filter run every second, the filter of Equation (1.11) linearized afresh at each step, as shown in Equation (2.14).

$$\begin{aligned} \hat{x}^- &= \hat{x} + f(\hat{x}, v)\,\Delta t, \qquad F = I + \frac{\partial f}{\partial x}\Delta t \\ P^- &= FPF^\top + Q\,\Delta t \\ K &= P^- H^\top (HP^- H^\top + R)^{-1} \\ \hat{x} &= \hat{x}^- + K(y - H\hat{x}^-) \end{aligned} \tag{2.14}$$

where v holds the pellets arriving through the model's delay, the fan, the steam valve, the lid switch and the igniter, y holds TT-101 and the meat probes, applied as scalar updates in turn, and the Jacobian is taken by forward differences. The process noise per second is $4\,\mathrm{K}^2$ on T_b, 0.2 on T_g, 0.05 on T_d, 0.01 on T_w, T_s and T_{pr}, 1×10^{-4} on T_c, 1×10^{-9} on the fuel, and $4\,\mathrm{W}^2$ on the bias d, a random walk that makes the controller offset-free [5], the disturbance state of Section 1.4.7. Volume 9, Chapter 4 carries a fuel bias in its estimator for the same reason. The probe is trusted to $R = 1\,\mathrm{K}^2$, above its 0.3 °C noise, because a tighter filter passes the noise into the feed, and the meat probes to $R = 0.0225\,\mathrm{K}^2$. A perfect-state teaching mode replaces the physical estimates with the truth and leaves only the bias to be estimated. The filter starts from the warm or cold start of Section 2.3 computed with the declared parameters, which is the plant's own start when the parameters are at their defaults.

2.5.5 Dew-Point Loop and the Three Models

With the steam kit fitted, AIC-301 is a PI loop from AT-301 to the steam valve, $K_c = 8\,\%$ per °C and $\tau_I = 180\,\mathrm{s}$, which tracks with the valve shut whenever AT-301 is invalid. The application keeps three models apart. The plant is everything in Section 2.3. The identified model is the FOPDT fit of Equation (2.11), with its two dead times, which the PID tuning uses and nothing else does. The controller's internal model is the MPC's copy of the plant's equations, which knows neither the damp pellets, the blocked intake, the gusts, the probe's bias nor a parameter changed from its default, and whose auger delay may differ from the plant's, and whose heat bias absorbs what it gets wrong.

2.6 Eighteen Cases, Measured

The eighteen cases appear in Table 2.5. All of them share one reference configuration (the standard PID at its IMC-based tuning, the scheduled fan, the continuous auger and the MPC's defaults) and depart from it only as their rows say. The figures quoted below all come from the engine run headless on its default seed, stepped exactly as the browser steps it, with each setting made before Run, which gives the same run as a setting carried through Reset (Section 1.6.1). The comparisons quoted were repeated over three to five noise seeds. At the fastest of the app's speeds a sixteen-hour cook plays in 96 s, and the exported file keeps each probe beside the true quantity it measures, the pellets in the tube, in the pot and burning, and the MPC's heat bias. The scorecard counts the share of the scored time with TT-101 within 5 °F of the setpoint, the worst excursion above the setpoint once the loop

Table 2.5 The Eighteen Cases

#	Case and event	Teaching point
1	Steady at 225 °F, stalled shoulder	fuel on its way
2	Cold start, Start at 30 s	ignition, first overshoot
3	Manual, feed +5 % at 5 min	FOPDT, two dead times
4	Lid up 20 to 23 min	feeding a disturbance
5	Overnight shoulder, 3 stages	recipe, the stall
6	Brisket, wrap at 165 °F core	a load change with a lid
7	Ribs, 3-2-1	clock stages, unwrap
8	Smoke, then roast at 350 °F	recipe preview
9	Reverse sear to 450 °F, MPC	carry-over, browning
10	Cold 3 kg butt at 30 min	a lid and a load
11	9 mph gusting wind, 60 to 35 °F	unmeasured losses
12	Damp pellets, MPC	model error, heat bias
13	0.5 kg in the hopper	flame loss, restart
14	0.3 kg pan at 250 °F	a vanishing load
15	Steam kit, dew point 113 °F	dew point, not RH
16	300 °F for 1 h, then 180 °F, MPC	pre-cooling, stage hold
17	Probe +10 °F, intake 30 % blocked	sensor and airway faults
18	Challenge, brisk PID, two lids and a cold load	the whole problem

has reached it, the worst excursion in the 20 minutes after a lid shuts, and the recovery, the time from the lid shutting until TT-101 is back within 5 °F and stays there for five minutes.

2.6.1 The Steady Point and a Cold Start

Case 1 holds 225 °F over the stalled shoulder. Over its scored 80 minutes TT-101 stays within 5 °F the whole time, with an RMS error of 0.55 °F under the standard PID and 0.57 °F under the MPC, at a mean feed of 22.9 %, and the grill burns 0.69 kg (1.52 lb) of pellets in 90 minutes. A steady operating point hides the delay, because nothing asks the fire to change. Case 2 does ask. From cold the supervisor hands over to the loop 5.0 minutes after Start, with the grate at 83 °F, and the standard PID first comes within 5 °F of 225 °F 20.2 minutes after Start and then overshoots by 21.4 °F, because its integral has grown through a climb of more than 100 °F and the pot still holds fuel fed for a colder grill. The MPC is in the band 9.8 minutes after Start and overshoots by 1.9 °F, because it predicts the arrival of what it has already fed.

2.6.2 The Lid and Four Answers

Case 4 opens the lid for three minutes at 20 minutes. Table 2.6 gives eight configurations on the identical opening and seed, and Figure 2.2 traces four of them, and Figure 2.3 draws the overshoot and the time in the band of all eight. The standard PID feeds up to 41.8 % while the lid is up, peaks at 234.7 °F, 9.7 °F over the setpoint, thirteen minutes after the lid shuts, and has not settled back inside the band for good by the end of the hour. Brisker tunings recover sooner by overshooting more, 28.9 °F with a 22.1 min recovery at 0.6 % per °C and 10 min, and 60.8 °F at 1.0 % per °C and 6.5 min. The lid-aware PID holds the feed at 23 % while the lid is up and for 45 s after, feeds no more than 26 % once it resumes, overshoots by 3.0 °F and recovers in 8.7 min. The MPC first cuts the feed to its floor as the probe collapses, then feeds up to 39.1 % to restore the heat that the steel and the meat are losing, and 51 % for a moment as the lid shuts, overshoots by 2.7 °F and is back in the band for good 0.7 min after the lid shuts, 93.9 % of the hour in the band against 53.4 % for the standard PID. Over five noise seeds the overshoot spans 9.0 to 9.7 °F for the standard PID, 2.7 to 3.4 °F for the lid-aware PID and 2.7 to 3.2 °F for the MPC, so those two differ in how fast they settle rather than in how far they overshoot.

The last three rows separate what the MPC knows from what it assumes. With the true state in place of the estimate it overshoots by 1.9 °F instead of 2.7 °F (1.8 to 2.1 °F against 2.7 to 3.2 °F over five seeds), so the estimator costs little. Told that the lid will stay open for the whole horizon, the same controller predicts a chamber it cannot heat, feeds 100 % and overshoots by 92.8 °F, peaking at

Table 2.6 Case 4, the Three-Minute Lid Opening, on One Seed. Feed Is the Largest While the Lid Is Up, and Over Is Measured in the 20 Minutes After It Shuts

Configuration	Feed (%)	Low (°F)	Over (°F)	Recovery (min)	In band (%)
standard PID	41.8	128.7	9.7	—	53.4
PID, 0.6 %/°C, 10 min	61.5	131.0	28.9	22.1	61.4
PID, 1.0 %/°C, 6.5 min	89.1	132.3	60.8	14.7	72.9
PID, pulsed auger	41.4	129.6	12.7	—	45.7
lid-aware PID	23.0	123.6	3.0	8.7	82.5
MPC, lid shuts	39.1	125.1	2.7	0.7	93.9
MPC, perfect state	33.7	124.7	1.9	0.8	93.8
MPC, lid stays open	100	133.5	92.8	7.4	83.5

A dash means the grate had not settled for good by the end of the hour.

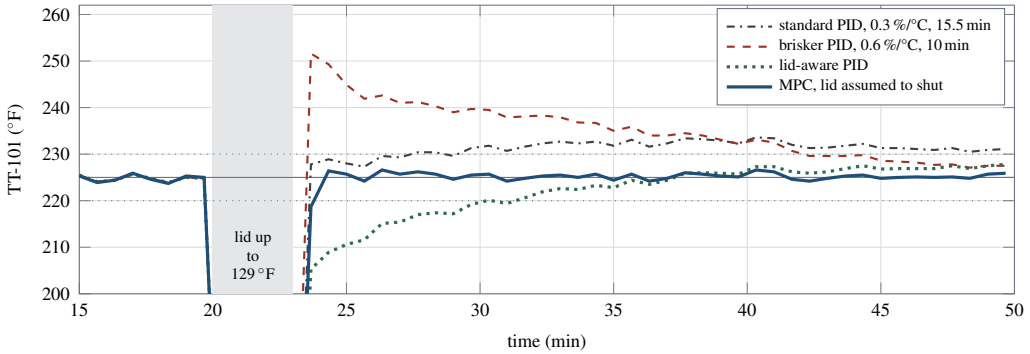

Figure 2.2 Grate probe TT-101 through the three-minute lid opening of case 4, recorded from the running application, with the lid up in the shaded band and the 225 °F setpoint and its ±5 °F band as thin solid and dotted gray lines. The standard PID (dash-dotted) peaks 9.7 °F over thirteen minutes after the lid shuts, the brisker PID (dashed) 28.9 °F within a minute, and the lid-aware PID (thick dotted) and the MPC (thick solid) stay within about 3 °F.

317.8 °F. The opening is the same in both runs and only the assumption differs, which is the point the case is built to make. Pulsing the auger at the same mean feed makes the standard PID worse, 12.7 °F over and 45.7 % in the band, because a pulsed tube advances only while the motor turns and its transport time grows as the feed falls.

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

✦ **DID YOU KNOW? Pulsing the Auger Makes the Same Loop Worse**

From the simulator: on the lid opening of case 4 the standard PID overshoots by 9.7 °F with a continuous auger and by 12.7 °F with a pulsed one at the same mean feed, and its time in the band falls from 53.4 to 45.7 % (Table 2.6). The pulsed delay is longest when the feed is lowest.

2.6.3 Recipes, the Stall and Wrapping

Case 5 puts a 4 kg shoulder on at 9 PM straight from the refrigerator under a program of 225 °F for six hours, 250 °F for three, then 275 °F until the probe reads 200 °F, and a 165 °F hold, while the night cools from 68 °F to 50 °F. With the standard PID the second stage begins at 3:00 AM and the third at 6:00 AM, and the core passes 150 °F at 4.41 h and does not reach 165 °F until 7.59 h. That stall of 3.2 h happens with the grate on target, because the surface is evaporating, and a tighter grate band cannot shorten it. The probe reads 200 °F at 7:41 AM, and the grill burns 6.87 kg (15.1 lb) in

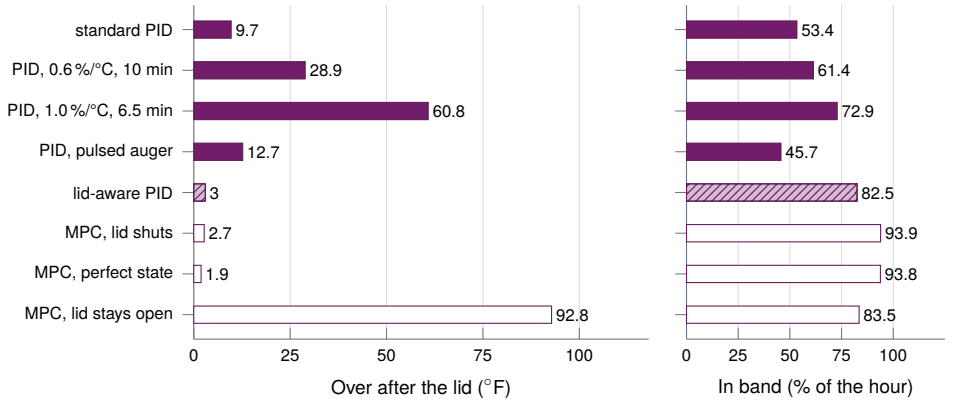

Figure 2.3 The eight answers to the three-minute lid opening of case 4 from Table 2.6: the overshoot in the 20 minutes after the lid shuts (left) and the share of the hour within ±5 °F (right). Dark bars are standard PID loops, the hatched bar the lid-aware PID and white bars the MPC. Notice that brisker PID tunings buy time in the band with larger overshoots, while the MPC told that the lid stays open overshoots most of all.

sixteen hours. The MPC keeps 97.1 % of the night in the band against 87.4 % and finishes at 7:37 AM, although its worst excursion is larger, 14.9 °F against 7.7 °F, because it pre-heats ahead of each programmed step. At 1 PM the core has cooled to 158.6 °F in the 165 °F hold, since a hold holds the chamber and the meat keeps evaporating.

Wrapping acts on the evaporation itself. Case 6 smokes a 5.5 kg brisket from 10 PM at 225 °F until the probe reads 165 °F, prompts a wrap with the lid up three minutes, and runs 250 °F to a 200 °F core. The probe reads 165 °F at 6:38 AM, 8.63 h in, and the wrapped brisket reaches 200 °F at 8:01 AM against 8:16 AM unwrapped. At that hold it keeps 14.3 % of its water against 6.8 % and has a bark index of 28 against 52.5. After the wrap's lid the standard PID recovers in 7.3 min and the MPC in 0.8 min. Case 7 runs spare ribs on the 3-2-1 schedule. The core is 161.6 °F at 3 h, reaches 204.2 °F in the foil by 5 h, falls to 186.3 °F half an hour after the unwrap as the wet surface evaporates, and is back at 195.3 °F at 6 h. After the wrap the standard PID takes 45.4 min to settle, because the wrapped rack suddenly draws less heat and the conservative loop trims the feed slowly, while the MPC settles in 0.7 min, and the grate is in the band for 82.4 % of the six hours against 96.0 %.

2.6.4 Preview and the Stage-Hold Policy

A recipe on the clock is information that feedback cannot use and prediction can, and it can be used too early. Volume 9, Chapter 5 uses published orders the same way, releasing water one travel time ahead of the demand. Case 8 smokes a chicken at 225 °F for 45 minutes and then roasts it at 350 °F to a 165 °F core. The standard PID reaches 345 °F 54.1 min after the step and the chicken is done at 6:07 PM. The MPC sees the step 20 minutes ahead, the grate is already at 256.8 °F when the roast begins, it reaches 345 °F 2.5 min later, and the chicken is done at 5:57 PM. With the stage-hold policy on, the grate is held to 234.9 °F at the step and still reaches 345 °F in 3.1 min. Case 16 holds 300 °F for an hour and then asks for a 180 °F hold. The MPC cuts the feed to its floor well before the transition and the grate is at 218.5 °F when the program changes, 81.5 °F under the stage it was meant to hold, and reaches 185 °F 10.5 min later. With the stage-hold policy it holds 292.8 °F to the end of the stage and takes 23.5 min, while the standard PID holds 300.4 °F, takes 42.4 min, and cuts the feed from 32.5 to no lower than 12.1 %. Using the preview early is optimal for the objective of Equation (2.13) and wrong for a recipe that wanted every minute of each stage, which is what the policy is for. Case 9 shows a limit of the grill itself. A thick ribeye cooked to a 115 °F core and then seared at 450 °F passes its 130 °F pull temperature 8.2 min into the sear, less than a minute after the grate reaches

450 °F, and its core peaks at 135.4 °F on the counter while its surface never passes 169.9 °F, because it is still evaporating, so a convection grill at 450 °F cannot brown a wet steak before its core is done.

2.6.5 Loads the Loop Has to Find

Case 10 puts a cold 3 kg butt on at 30 minutes with the lid up for a minute. The standard PID keeps 96.5 % of the three hours in the band, overshoots by 1.0 °F after the lid and recovers in 7.3 min, while its feed settles from 22.9 to 24.2 % for the new load. Case 11 blows a 9 mph gusting wind across the barrel as the night falls from 60 °F to 35 °F. The standard PID keeps 89.5 % of the four hours in the band and the MPC 93.6 %, the feed rises from 27.6 to 30.4 % in ten-minute means, and the grill burns 2.31 kg against 1.83 kg on a calm 68 °F night. Over five seeds of the gusts the standard PID keeps 82 to 90 % in the band and the MPC 88 to 94 %, so the MPC's advantage holds while its size varies. Case 12 gives the MPC damp pellets that catch more slowly and carry less heat than its model believes. The heat bias reaches −232 W by two hours, a phantom heat loss that stands in for the weaker fuel, the feed climbs from 22.9 to 27.8 %, and the MPC keeps 97.0 % of the two hours in the band, the same as with dry pellets, and overshoots the lid at 70 minutes by 1.9 °F, against 62.8 % for the standard PID. With the model's auger delay set to 0, 60 and 120 s instead of the plant's 20 s, the lid overshoot is 2.1, 1.9 and 3.9 °F. The first two lie within the 1.9 to 2.3 °F that five noise seeds give, and only the 100 s error doubles the overshoot, so the bias absorbs most of a wrong delay, and the lags in the pot rather than the tube dominate the fire's response.

> **TRY IT IN THE APP** **12 · Damp Pellets From a New Bag**
>
> Press Run and watch the heat bias on the MPC card after the damp pellets reach the pot at 30 minutes. It falls to −232 W by two hours, a phantom heat loss that stands in for the weaker fuel, while the feed climbs from 22.9 to 27.8 % and the grate stays in the band for 97.0 % of the two hours.

2.6.6 Faults and the Chamber's Water

Case 13 starts with 0.5 kg in the hopper. The auger runs dry at 1.18 h, the fire burns down what the pot held while the PID raises the feed from 23 to 33 %, and at 75.3 min, with the grate at 181.0 °F, the supervisor declares FLAME OUT and stops the auger. Left alone, the meat's core falls to 114.4 °F by three hours, and refilled and restarted at 90 minutes the grill is RUNNING again 5.7 min after Start. Case 14 cooks at 250 °F over a pan of 0.3 kg. The pan is down to its last 10 g at 47.7 min, the grate, relieved of the pan's latent load, peaks at 259.9 °F, and the dew point falls from 103.7 to 92.1 °F. In case 17 the grate probe begins to read 10 °F high at 20 minutes, and at 50 minutes TT-101 reads 225.9 °F on target while the air around the meat is 216.5 °F. The MPC, fed the same probe, is fooled the same way, since feedback cannot correct its own sensor. As Section 1.4.6 shows for an analyzer, a constant bias seen through the one measurement it corrupts cannot be estimated without a second, unbiased reading. At 80 minutes ash blocks 30 % of the intake, and by 133 minutes the excess-air ratio has fallen from 5.25 to 3.88 while the fire still releases 1617 W. The fire has five times the air it needs, which is also why a small pellet boiler's emissions respond to accurate airflow control [2].

Case 15 fits the experimental steam kit and a dew-point PI that holds 113 °F while a shoulder cooks at 250 °F with no pan. At 1 h the valve runs about 51 %, the relative humidity is only 4.7 % and could not pass 49.1 % at 250 °F however much steam went in, and with the valve shut and wide open the reachable dew point is 90.3 to 126.4 °F. With the steam the core reaches 165 °F at 3.0 h against 3.58 h without, because the surface evaporates less and runs hotter, but 180 °F at 7.55 h against 7.44 h, because the water it kept makes the evaporating period last longer. The steam costs 4.8 kWh and the bark index ends at 76.9 against 95.2.

2.6.7 The Challenge

The challenge holds 225 °F for two hours through a three-minute opening at 20 minutes, a cold 3 kg butt with the lid up for a minute at 45 minutes and another three-minute opening at 75 minutes. The

Table 2.7 The Challenge, Two Lid Openings and a Cold Load Over Two Hours

Configuration	In band (%)	Over (°F)	Recoveries (min)	Pellets (lb)
standard PID as shipped[a]	70.1	73.2	13.1, 8.0, 12.8	2.33
lid-aware PID, shipped gains	85.9	4.6	3.4, 4.2, 3.3	2.19
lid-aware PID[b]	88.4	3.7	2.4, 2.5, 2.4	2.19
lid-aware PID[c]	90.7	2.5	1.5, 1.5, 1.4	2.19
MPC	92.5	3.4	0.7, 0.8, 0.6	2.19

[a] $K_c = 1.2$ % per °C, $\tau_I = 400$ s. [b] $K_c = 2$ % per °C, $\tau_I = 5$ min.
[c] $K_c = 5$ % per °C, the top of its range, $\tau_I = 5$ min.

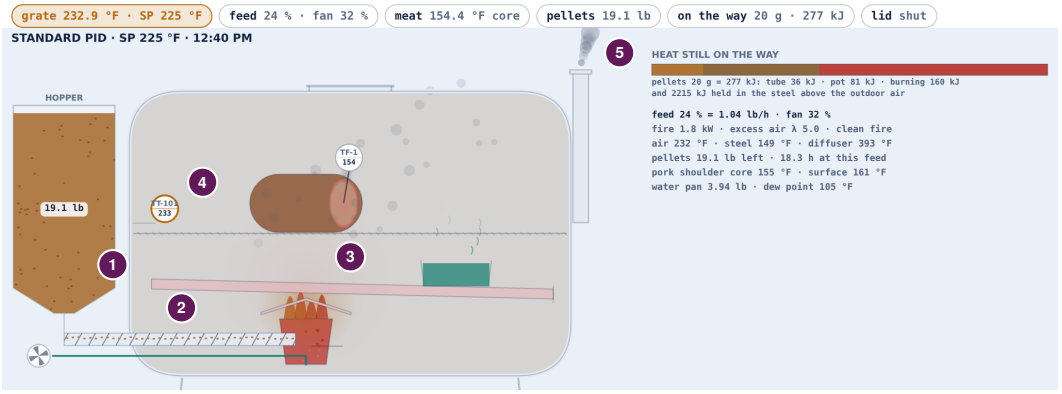

Figure 2.4 The three-minute lid opening of case 4 at $t = 40.0$ min in the Grill view, run with the standard PID and cropped to the plant view (the header, control panel, trends and scorecard are cut). The badge row reads the grate at 232.9 °F, 7.9 °F above the setpoint, and the cutaway shows the shoulder on the grate with its cut end in section, the water pan and the chimney. The callouts mark (1) the hopper at 19.1 lb, (2) the pellets in the auger tube, (3) the fire under its baffle, (4) TT-101 at the grate and (5) the panel that gives 20 g of pellets and 277 kJ still on their way against 2215 kJ in the steel.

target is TT-101 within 5 °F for at least 90 % of the two hours, no more than 10 °F over the setpoint after any of the three openings, and every recovery inside five minutes, judged on the case as given. Table 2.7 gives five configurations. The brisk standard PID it ships with meets no part of the target, and its slowest recovery is 13.1 min. At its shipped gains and at 2 % per °C the lid-aware PID meets the overshoot and recovery parts but not the 90 %, with 85.9 and 88.4 %. With $\tau_I = 5$ min it passes narrowly once K_c reaches 3.3 % per °C, eleven times the step-test gain, and at 5 % per °C it keeps 90.7 % in the band on all five seeds tried, with recoveries of 1.4 to 1.5 min. The MPC meets every part at its defaults with the widest margin, 92.5 % in the band, 3.4 °F over and a slowest recovery of 0.8 min, margins of 2.5 points, 6.6 °F and 4.2 min on the target, and about 6 % of the two hours (7.0 min) is lost to the openings themselves, with the lid up, whatever the controller does. The plant view of the run of case 4 with the standard PID is shown in Figure 2.4, and Table 2.6 gives that run beside the MPC's.

2.7 Technology in Practice

A real grill reads and moves what the loop of Section 2.5 reads and moves. Figure 2.5 draws TIC-101, the lid-aware hold and the MPC with field devices in place of the schematic's symbols.

What Measures It

The controlled variable is the grate temperature that TT-101 reads. Of the contact sensors, thermocouples are by far the most rugged and wide-ranging, while resistance temperature detectors are more fragile but need no reference-junction compensation and are inherently more linear. A thermowell around the element adds a first-order lag through its mass and specific heat, and an air gap between element and well adds a second one [20]. The chapter lumps every such lag into the 8 s of τ_{TT} in Equation (2.10).

> **TECHNOLOGY Sensor: The Grate Thermocouple**
>
> A thermocouple is a junction of two dissimilar metals. In industrial use the wires are protected by a stainless steel or ceramic sheath. Its range depends on its type: about 32 °F to 2300 °F for type K and −300 to 700 °F for type T, and narrower where a tighter accuracy is wanted. Its most common failure is to fail open, or burn out. The open input then picks up noise from nearby motors and drives that reads as a wildly varying temperature, so the instrument is given a burnout mode that drives the reading high or low [20]. TT-101 of this chapter is an 8 s lag with colored noise and an optional constant bias, the fault of case 17, and a failed probe sends the supervisor to FAULT.

What Moves It

The manipulated variable is the auger feed u_a, and the combustion fan follows a schedule on that feed. The furnace patented in 1986 already fed its fuel intermittently and could be controlled by a thermostat or other heat sensitive means, and a blower system supplied the combustion air that flowed into the fire pot [13]. The final element of the chapter is the auger alone. Its delay of $\theta_t = 20\,\text{s}$ is a line of 40 cells, and its two drives differ only in how that line advances.

> **TECHNOLOGY Final Element: An Auger on a Gear Motor**
>
> In both patents an electric motor turns the auger along its feed tube [13, 14], and in process plants powders are metered the same way, by variable-speed screw conveyors. A direct-current motor's power is commonly set by pulse-width modulation, the fraction of each period its switches conduct, and the same time-proportioned output serves an electric oven whose heater can only be fully on or fully off [20]. Pellet stoves need an electrical supply because motorized components feed the fire and blow the air [15]. The application offers both drives, a continuous auger that meters a mean rate and a pulsed one that runs at full speed for a share of every 15 s cycle.

Where the Controller Runs

A grill's controller is the microcontroller reading a thermocouple that the Assumptions of this chapter mention. Embedded digital controls have replaced analog mechanical ones in rotating machines and engines, and supervisory and lower-level loops run with cycle times from milliseconds to minutes [21]. The application's controller scans every 0.5 s, runs the PID and the estimator every second and the MPC every 10 s, all short beside the 20 s transport time and the 15.5 min time constant of the fit. The velocity-form PID of Equation (2.12) is the block that holds the standard loop, the lid-aware hold is a switch around it, and the MPC replaces it with a plan made from its own record of the feed.

What Is Modernizing It

Three developments bear on this loop. The first is the estimator inside predictive control. A survey of industrial MPC found that most products did not carry a process state and fed back an ad hoc bias, while only two used a Kalman filter, one of them an extended filter for a nonlinear model, and it urged practitioners to recognize the rich framework the filter provides [4]. The MPC of this chapter takes that route, with an extended filter and a heat bias. The second is connectivity. Low-cost Ethernet, Internet Protocol and wireless devices are replacing older proprietary technologies in control systems, which brings remote access and raises the likelihood of cybersecurity vulnerabilities and

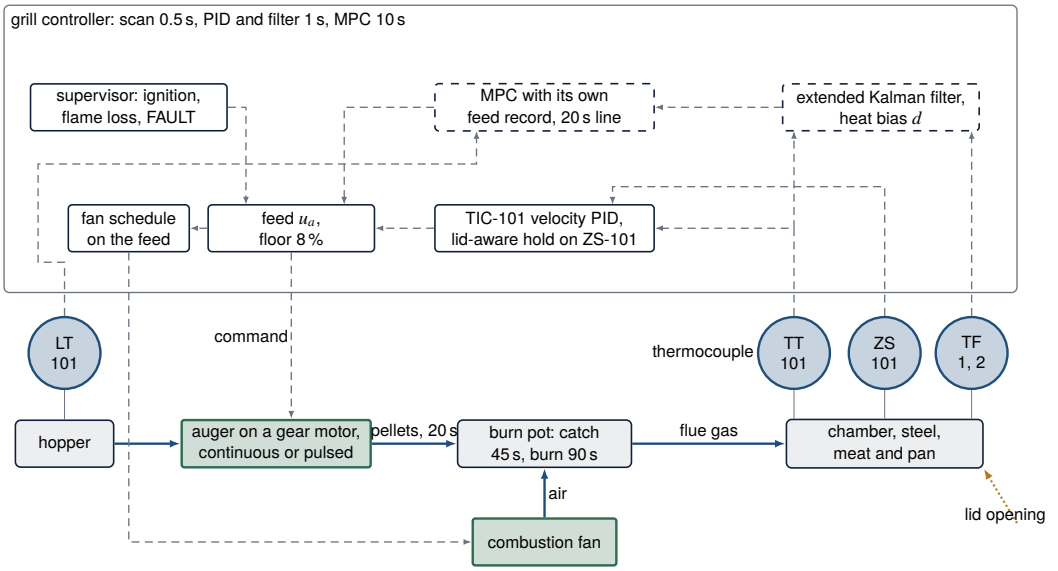

Figure 2.5 TIC-101 and its alternatives on grill hardware. A thermocouple at the grate (TT-101), the lid switch (ZS-101), the meat probes (TF-1, TF-2) and the hopper level (LT-101) feed one embedded controller. The velocity-form PID holds the feed while the lid switch reads open, the dashed MPC and its extended Kalman filter replace the PID when selected, and the supervisor takes the auger for ignition and faults. The fan follows a schedule on the feed. Notice that every path from the controller to the grate passes through the auger's 20 s line and the pot's ignition and burn lags, which only the MPC carries in its model, with its own record of the feed.

incidents [21]. The third is the fuel. Pellets differ in how well they are made and how well they burn, lower quality pellets may give less heat, and third-party certification programs let users identify consistent, reliable pellets [15]. A steadier fuel shrinks the disturbance that case 12 imposes and that the heat bias has to absorb.

> **IN PRACTICE** **Reading a Bump Test**
>
> A process is characterized in the field by an open-loop test. The controller is put in manual, the final element is stepped, and the trend is read for two features: the dead time before the measurement first moves, and the reaction rate, the steepest slope that follows. Dead time tends to stay the same whatever the size of the step, while the reaction rate grows with it, so the step size enters the tuning calculation. Ziegler and Nichols built their second tuning rule on these two readings in their 1942 paper, where they called dead time a lag, a word that now means something else [20]. Case 3 is such a test, run on the auger.

2.8 Limitations

The combustion, heat-transfer and cooking constants are illustrative, chosen so that the totals are plausible, a pound of pellets an hour at 225 °F and about 590 °F at full feed in an empty grill, without being fitted to a grill. The chamber is one well-mixed node and the steel two, so the model cannot say where in the barrel the meat cooks fastest or where a probe should sit. The fire is a mass balance with smooth thresholds, and ash accumulation, pellet quality beyond the damp disturbance, and the chemistry of smoke are outside it. The meat is two nodes, so its stall and its carry-over are qualitatively right while its internal temperature profile is not resolved, and no endpoint in the application is a food-safety assessment.

OPEN CHALLENGE **Delays That Move With the Feed**

Dead-time compensation still has open questions that a pellet fire raises. Extending the properties of compensators descended from Smith's predictor to nonlinear input-output models, such as Hammerstein models, is an open research topic. So is a robust stability test for delays that vary in time that is less conservative than those available [6]. The grill has both: the burn rate depends on the bed's temperature and oxygen, and the pulsed auger's transport time grows as the feed falls. A compensator whose guarantees hold while the delay moves with the feed it sets would be progress.

The controller's model is the plant's own structure with declared parameters, so in the cases the MPC meets model error only through the disturbances they apply, and a parameter changed from its default gives it a plant it was not designed for. A real controller would also face a model of a different structure. The MPC is linearized afresh at every execution and solved as a convex program, and it makes no guarantee beyond its horizon. The supervisor's thresholds are educational and not certified appliance logic, the steam kit and the high-temperature dew-point probe are an experimental teaching configuration, and the lid, the loads and the faults arrive as scheduled events rather than as a cook's habits.

OPEN CHALLENGE **A Biased Probe Looks Like a Good One**

In most industrial predictive controllers a sensor fault is recognized only as a complete failure beyond preset limits. A significant bias or drift that stays within normal limits is generally neither detected nor identified [4]. The probe bias of case 17 is such a fault, and it fools the MPC of this chapter as it fools the PID. Fault detection that finds a bias within the normal range, without slowing the loop, would be progress.

2.9 Conclusion

This chapter studied a pellet grill, a plant whose controller commands fuel that becomes heat through a 20 s transport time, an ignition lag and a burn time, on a simulator that logs every inventory of that fuel beside what its instruments report. At 225 °F the fuel on its way is worth 267 kJ against 2137 kJ in the steel, and a 5 % feed step fits a gain of 7.29 °F per percent and a time constant of 15.5 min with a dead time of 4.7 or 2.8 min depending on the method. A three-minute lid opening leads a PI loop tuned from that step to feed 41.8 % and overshoot by 9.7 °F after the lid shuts, while a lid-aware hold overshoots by 3.0 °F and an MPC that assumes the lid will shut overshoots by 2.7 °F and recovers in 0.7 min, and the same MPC told that the lid stays open overshoots by 92.8 °F. Overnight the meat stalls for 3.2 h with the grate on target, and a recipe previewed by the MPC can pre-heat or pre-cool a stage unless a policy forbids it. The model is lumped and its constants illustrative. A fuller model would add a distributed chamber, a pellet quality that the estimator must identify, a controller whose model differs in structure from the plant, and emissions tied to the excess air. The grill shows that knowing what the actuator has already done decides how a loop behaves when its plant is disturbed.

KEY POINTS

- A pellet grill's controller can add fuel but not heat, and each pellet becomes heat only after it has ridden the auger, caught on the bed and burned.
- A 5 % feed step fits a dead time of 4.7 or 2.8 min depending on the method.
- After a three-minute lid opening a PI loop tuned from that step overshoots by 9.7 °F, a lid-aware hold by 3.0 °F, and an MPC that assumes the lid will shut by 2.7 °F.
- Told that the lid stays open, the same MPC overshoots by 92.8 °F, so its assumption about the disturbance matters as much as its algorithm.

Where else this applies. Volume 9, Chapter 2 puts a burner behind a dead time under predictive control, and Volume 9, Chapter 4 carries a fuel bias in its estimator for the same reason as this grill.

References

[1] C. Schörghuber, M. Gölles, M. Reichhartinger, and M. Horn, “Control of biomass grate boilers using internal model control,” *Control Engineering Practice*, vol. 96, p. 104274, 2020.

[2] R. Stanisławski, R. Junga, and M. Nitsche, “Reduction of the CO emission from wood pellet small-scale boiler using model-based control,” *Energy*, vol. 243, p. 123009, 2022.

[3] J. E. Normey-Rico and E. F. Camacho, “Dead-time compensators: A survey,” *Control Engineering Practice*, vol. 16, no. 4, pp. 407–428, 2008.

[4] 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.

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

[6] J. E. Normey-Rico, T. L. M. Santos, R. C. C. Flesch, and B. C. Torrico, “Control of dead-time process: From the Smith predictor to general multi-input multi-output dead-time compensators,” *Frontiers in Control Engineering*, vol. 3, p. 953768, 2022.

[7] H. Chen, B. P. Marks, and R. Y. Murphy, “Modeling coupled heat and mass transfer for convection cooking of chicken patties,” *Journal of Food Engineering*, vol. 42, no. 3, pp. 139–146, 1999.

[8] R. G. M. van der Sman, “Moisture transport during cooking of meat: An analysis based on Flory–Rehner theory,” *Meat Science*, vol. 76, no. 4, pp. 730–738, 2007.

[9] C. Boman, E. Pettersson, R. Westerholm, D. Boström, and A. Nordin, “Stove performance and emission characteristics in residential wood log and pellet combustion, part 1: Pellet stoves,” *Energy & Fuels*, vol. 25, no. 1, pp. 307–314, 2011.

[10] G. S. Mittal, C. Y. Wang, and W. R. Usborne, “Smokehouse process conditions for meat emulsion cooking,” *Journal of Food Science*, vol. 52, no. 5, pp. 1140–1146, 1987.

[11] P. Šimko, “Polycyclic aromatic hydrocarbons in smoked meats,” in *Safety of Meat and Processed Meat*. New York, NY: Springer, 2009, pp. 343–363.

[12] I. Obernberger and G. Thek, “Physical characterisation and chemical composition of densified biomass fuels with regard to their combustion behaviour,” *Biomass and Bioenergy*, vol. 27, no. 6, pp. 653–669, 2004.

[13] J. P. Traeger, R. J. Traeger, and M. A. Traeger, “Furnace construction,” U.S. Patent 4 619 209, Oct. 28, 1986.

[14] ——, “Pellet-fired barbecue,” U.S. Patent 4 823 684, Apr. 25, 1989.

[15] U.S. Environmental Protection Agency, “Pellet stoves,” Web page, last updated 30 April 2026, 2026, accessed 2026-11-21. [Online]. Available: https://www.epa.gov/burnwise/pellet-stoves

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

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

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

[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] 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.

[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.

2.10 Exercises

2.1 Concept check · 2 min

In case 4 the standard PID raises the feed to 41.8 % while the lid is up, and thirteen minutes after the lid shuts the grate peaks 9.7 °F over the setpoint. What is the main cause of this overshoot?

(a) Pellets fed while the lid was up are still in the tube and the pot when it shuts, and the integral keeps what it gained

(b) The steel lost most of its stored heat with the lid up, and the loop overcorrects as it replaces that heat

(c) The 8 s lag of the grate probe hides the recovery of the air from the loop for several minutes

(d) Anti-windup clamps the running sum while the lid is up, and the sum is released in one step as it shuts

2.2 Concept check · 1 min

The feed step of case 3 is fitted with a first-order plus dead-time model. Smith's two-point method gives a dead time of 4.7 min and the tangent at the steepest point gives 2.8 min, although the auger delay is only 20 s. Why?

(a) The probe noise biases the tangent, which is drawn from two samples near the steepest point

(b) The fuel passes through ignition, burning and the steel's heating in series, and no single dead time represents that chain

(c) The auger's transport time grows as the feed rises, so the delay changes during the step

(d) The scheduled fan adds a second delay that the tangent method cannot see

2.3 Concept check · 2 min

On the lid opening of case 4, the MPC that assumes the lid will shut at any moment overshoots by 2.7 °F. Told that the lid will stay open for the whole horizon, the same MPC feeds 100 % and overshoots by 92.8 °F. What does the comparison show?

(a) The extended Kalman filter loses track of the fuel while the lid is up, so the plan starts from a wrong state

(b) The 20 min horizon is too short to see the lid shut, so the plan ignores the recovery

(c) What the controller assumes about the disturbance's future matters as much as its algorithm

(d) The move suppression of 10 is too weak to stop the feed from reaching its limit

2.4 Theory and calculation · 6 min

The auger delivers $q_{max} = 2\,\mathrm{kg/h}$ at full feed, the pellets have a heating value of 17 MJ/kg as fed and the fire burns them at an efficiency $\eta = 0.80$.

(a) Case 1 holds 225 °F at a mean feed of 22.9 %. Find the pellets burned per hour in steady state and the heat the fire releases. *Checkpoint:* 0.458 kg/h (about 1.01 lb/h) and 1.73 kW.

(b) At 30 minutes into case 1, 19.6 g of pellets are in the tube, in the pot and burning. Find the heat they release in the fire, at η, and how many minutes of fire they represent at the 1733 W the fire releases. Compare the heat with the 2137 kJ the steel holds above the outdoor air. *Checkpoint:* 267 kJ, 2.6 min, and the steel holds 8.0 times more.

(c) With the pulsed auger drive, the tube advances only while the motor turns. Estimate the transport time of a pellet at a duty of 23 % with $\theta_t = 20\,\mathrm{s}$. *Checkpoint:* About 87 s.

2.5 Theory and calculation · 8 min

The feed step of case 3 fits Equation (2.11) with $K = 4.05\,^{\circ}\mathrm{C}$ per percent and $\tau = 15.5\,\mathrm{min}$, with a dead time of 4.7 min by Smith's method and 2.8 min by the tangent. The SIMC rule takes $K_c = \tau/\big(K(\tau_c + \theta)\big)$ and $\tau_I = \tau$ with $\tau_c = \tau/2$.

(a) Find K_c and τ_I on Smith's dead time, and give K_c in percent per °F as the app shows it. *Checkpoint:* 0.31 % per °C (0.17 % per °F) and 15.5 min.

(b) Find K_c on the tangent's dead time, and say which fit gives the more cautious loop. *Checkpoint:* 0.36 % per °C. Smith's dead time gives the more cautious loop.

(c) The shipped loop runs $K_c = 0.3$ % per °C and $\tau_I = 930$ s. In case 4 the error is 53 °C as the lid shuts and averages 50 °C over the 180 s it is up. Estimate the feed the proportional and integral terms add to the steady 23 %. *Checkpoint:* About 15.9 % and 2.9 %, a feed of about 41.8 %.

2.6 Theory and calculation · 4 min

The chamber is at the atmospheric pressure $P = 101.3$ kPa. Steam tables give a saturation pressure of 130.5 kPa at 225 °F (107.2 °C) and of 7.61 kPa at 105 °F (40.6 °C).

(a) Find the highest relative humidity the chamber gas can reach at 225 °F, however much water enters. *Checkpoint:* 77.6 %.

(b) In case 1 the dew point is 105 °F. Find the relative humidity, and say why the humidity loop of the steam kit controls the dew point instead. *Checkpoint:* 5.8 %.

2.7 App challenge · 8 min

Load scenario 4, A three-minute lid opening, and keep the Standard PID as loaded. Run to 60 min at a speed of 60 with the default seed, and read the tiles In ±5 °F, Over SP after a lid and Lid recovery. Press Reset, set Automatic controller to Lid-aware PID, and repeat. Press Reset again, set Automatic controller to MPC, and repeat.

(a) Report the three tiles for each controller. *Checkpoint:* Standard PID 53.4 %, 9.7 °F, not yet. Lid-aware PID 82.5 %, 3.0 °F, 8.7 min. MPC 93.9 %, 2.7 °F, 0.7 min. Within 0.3 %, 0.2 °F and 0.2 min.

(b) The lid-aware PID and the MPC overshoot by nearly the same amount. Explain from the feed strip why the MPC recovers more than ten times faster. *Checkpoint:* The MPC feeds to replace the heat the steel and the meat lost and predicts its arrival. The hold feeds no more than before the lid opened.

2.8 App challenge · 9 min

Load scenario 3, A feed step to identify the fire, and change no setting. The grill stays in Manual with the auger at 18.2 % and the fan at 27 %, and the feed steps up by 5 % at 5 min. Run to four hours at a speed of 600 with the default seed and export the CSV file. Use the noise-free true chamber air, column `T_chamber_C`.

(a) Read the column at 300 s and at 14 400 s, and find the rise and the gain of the whole step in °C per percent. *Checkpoint:* A rise of 20.61 °C (37.10 °F) and 4.12 °C per percent, within 0.05 °C.

(b) The grill drifts up 0.65 °F in four hours with no step at all. Correct the gain for the drift and give it in °F per percent. *Checkpoint:* 4.05 °C per percent, 7.29 °F per percent.

(c) Find the times after the step at which the air has covered 28.3 % and 63.2 % of the rise of part (a), and apply Smith's method, $\tau = 1.5(t_{63} - t_{28})$ and $\theta = t_{63} - \tau$. *Checkpoint:* $t_{28} = 9.75$ min and $t_{63} = 20.2$ min, within 0.1 min, so $\tau = 15.6$ min and $\theta = 4.5$ min.

2.9 App challenge · 8 min

Load scenario 12, Damp pellets from a new bag, with the MPC as loaded. Damp pellets reach the pot at 30 min, when the Damp pellets slider moves to 70 %, and the lid is up for three minutes at 70 min. Run to two hours at a speed of 60 with the default seed and export the CSV file.

(a) Read the MPC's heat bias, column `mpc_bias_W`, and the share in the band, column `metric_in_band_pct`, in the row at 7200 s. Explain the sign of the bias. *Checkpoint:* –232 W, within 3 W, and 97.0 %, within 0.2 %.

(b) Reload the scenario, set Auger delay in the MPC's model to 120 s, and run to 95 min. Read Over SP after a lid, and compare it with the 1.9 °F of the correct 20 s delay. *Checkpoint:* 3.9 °F, within 0.2 °F.

2.10 Application area · 3 min

A 9 kg hopper feeds an overnight cook. In case 5 the grill burns 6.87 kg of pellets in sixteen hours. In case 11 it burns 2.31 kg in four hours of a 9 mph gusting wind against 1.83 kg on a calm 68 °F night.

(a) Find the mean burn rate of case 5 and the pellets left at the end. *Checkpoint:* 0.43 kg/h, and 2.13 kg (4.7 lb) left.

(b) Find what a sixteen-hour cook would need at the windy rate, and what a cook should do before an unattended night in wind. *Checkpoint:* About 9.24 kg, more than the hopper holds.

2.11 Application area · 4 min

In case 17 the grate probe begins to read 10 °F high at 20 min. At 50 min TT-101 reads 225.9 °F on target while the air around the meat is 216.5 °F.

(a) How much colder than displayed does the meat cook, and could the MPC correct it? *Checkpoint:* About 9.4 °F. No: it is fed the same probe.

(b) The chapter's model has one well-mixed chamber. Describe how a practicing engineer would detect a probe bias on a real grill. *Checkpoint:* An independent reference reading at the grate, compared with TT-101 at a steady point.

2.12 Review the agent · 5 min

An AI assistant is asked how to pass the challenge (case 18, Two lid openings and a cold load) with a PID. It replies: "The lid-aware PID cannot meet the 90 % target at any tuning. A gain search reaches only 88.4 % at its best point, 1.11 % per °F (2 % per °C) and an integral time of 5 min, because more gain only adds overshoot after each lid. The MPC is the only way to pass." Find the flaw, and say how you would check it in the app.

CHAPTER 3

Teaching Composition Control With a Slow Analyzer on a Thirty-Two Stage Column

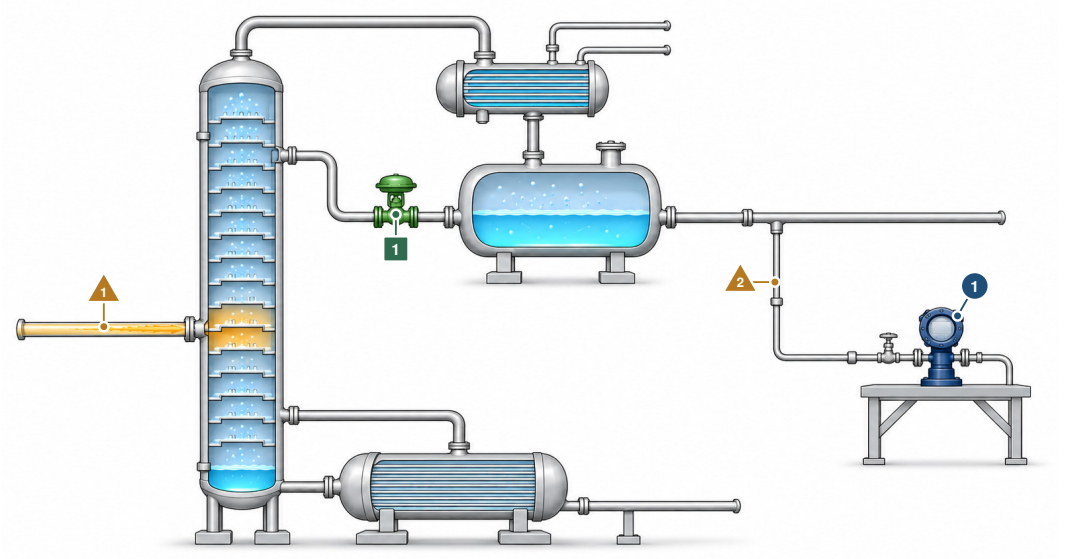

The controller moves 1 reflux valve that sets the reflux ratio

The controller reads 1 analyzer of the distillate's purity

What disturbs it 1 feed richer or leaner in the light component 2 a slower analyzer

The reflux reaches the top trays almost at once, while the analyzer reads the distillate only after the sample has traveled its line and been analyzed. Illustration, not to scale.

IN PLAIN TERMS **Steering a Column by a Late Report**

The main idea. A distillation column splits a liquid mixture by boiling it. Some of the liquid made at the top is poured back down the column, and the more is poured back, the purer the top product becomes. A controller sets that flow from a test of how pure the top product is.

The challenge. The test takes minutes, so the controller always acts on old news. A simple controller then overshoots, and a smarter one that forgets the delay can keep swinging back and forth.

Why it matters. Knowing what the late test costs tells engineers whether to slow the controller, predict ahead, or estimate the purity from faster readings.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on distillation (Section 3.2), the history segment on the relay (Section 3.5.6), the try-it on the reflux limit (Section 3.6.8), the try-it and the fact on the late analyzer (Section 3.7) and the key points
Study	80 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Sections 3.6 and 3.7), Technology in Practice (Section 3.9), Limitations (Section 3.10) and the two open challenges

The control question A thirty-two stage column holds its distillate purity with the reflux ratio, and the only measurement is an analyzer that reports minutes late. How much performance does that lag remove from a fixed tuning, and what can detuning or a predictive controller recover when the controller's model leaves the analyzer out?

Learning objectives

- Derive the ceiling that the overall balance places on the distillate purity, and tell an infeasible setpoint from a poor tuning.
- Fit a first-order plus dead-time model to a reflux step, explain why the column shows a dead time it does not contain, and derive the IMC tuning.
- Predict from the phase margin how an analyzer lag erodes damping, and measure its cost in overshoot, settling and margin to the specification.
- Compare the PID with an MPC whose model omits the analyzer lag, and find where that model error ends the MPC's advantage.
- Judge the PI constants a relay test writes when the relay band, not the column, sets the cycle.

Prerequisites Sections 1.4.2 (fitted dead time), 1.4.3 (IMC tuning), 1.4.5 (age of an analyzer result), 1.6.2 (velocity-form PID) and 1.4.7 (disturbance state). Section 1.6.3 is needed for the predictive controller.

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

See also Volume 6, Chapter 3 for a scanning gauge that adds a whole scan of dead time, and Volume 6, Chapter 2 for the pairing question a single-ended column cannot pose. Volume 13, Chapter 8 for interacting purity loops on one shell, and Volume 13, Chapter 9 for a tray temperature loop on a column with more than one steady state.

3.1 Introduction

A distillation column is among the most heavily instrumented vessels in a chemical plant, and its state is a composition profile rather than a single number. Every tray carries a liquid holdup in contact with a rising vapor, every holdup obeys its own component balance, and the plant presented to a composition controller is therefore of very high order. The controller never sees the profile. What arrives is one composition, measured on one product stream, by an instrument that reports minutes after the sample is drawn. Measurement dynamics of this kind are part of the loop and set a limit on achievable feedback performance that no tuning can recover [1].

A single measurement is nonetheless sufficient for control, because the column response is dominated by one slow mode. A column with thirty-two well-mixed stages has thirty-two eigenvalues, and for the simulator of this chapter they span time constants from 0.025 min to 15.24 min. A reflux move nonetheless produces a top-composition response that a first-order plus dead-time model reproduces almost exactly, because one mode carries essentially all of the steady-state gain and the rest decay before the response is half finished. Order reduction of this kind is a standard result for columns [2, 3], and it is the reason a two-parameter tuning rule works on a thirty-two state plant. The application draws all thirty-two liquid compositions as a live profile beside the trend of the distillate composition, so the order reduction is observed rather than asserted.

The reduction does not extend to the measurement path. A gas chromatograph on a product stream reports on a cycle rather than continuously, and the sample line and the analysis time add to the delay between the tray composition and the value the controller receives. Dead time and lag in the measurement path are not attenuated by feedback and cannot be removed by tuning, and the performance they remove is the subject of a literature of its own [4]. The industrial remedy is to infer the composition from tray temperatures, which respond immediately, and to correct the inference against the slow analyzer [5, 6]. The experiments of this chapter measure the loop degradation that justifies that expense.

This chapter quantifies what an adjustable analyzer lag costs a fixed tuning on a thirty-two stage binary column whose distillate composition is held with the reflux ratio. The topics are process order and model order reduction, apparent dead time, measurement dynamics, nonlinear steady-state gain,

integral windup, and the feasibility limit that an overall balance places on a setpoint, which are core topics of a first course in process dynamics and control [7, 8]. The physical content of the model is one component mass balance per stage, with the equilibrium driving force carried by a constant relative volatility and the flows fixed by constant molar overflow. Section 3.2 places the column in the literature on column dynamics and composition control. Section 3.3 sets out the balances, the instrument, the constants and the numerics in enough detail for a second implementation. Section 3.4 decomposes the column response into a symmetric inventory mode and an antisymmetric profile mode, and shows that fixing the distillate flow at half the feed makes every reflux move a pure internal-flow change, which suppresses the slow mode and sets a hard ceiling of $2x_F$ on the distillate composition. It then fits a first-order plus dead-time model to a step test and derives the PI tuning from it. Section 3.5 gives the discrete control law, a linear model predictive controller offered beside the PID, a moving horizon alternative to that controller's Kalman filter and a relay test that tunes the PID. Section 3.6 carries one tuning across five analyzer lags and two feed upsets, and Section 3.7 compares the predictive controller with the PID on five of the app's cases and across the same lags. Section 3.8 measures the estimator option and the relay test on the same column. Section 3.9 sets the loop on plant hardware, and Sections 3.10 to 3.12 state the limitations, the economic argument and the conclusions.

3.2 Background

Stage-by-stage calculation is the oldest quantitative method in separations. The graphical construction that steps between an equilibrium curve and an operating line was published in 1925 and remains the way a first course explains what a tray does [9]. Its central simplifying device, a constant relative volatility that converts the equilibrium calculation into one algebraic expression in the liquid mole fraction, is also what makes a dynamic stage model small enough to integrate. Every dynamic column model used as a control benchmark since then has been built the same way, one component balance and one equilibrium relation per stage, with constant molar overflow removing the energy balance from the problem [3, 10].

Column dynamics entered the control literature as a question about time scales. Measured columns settle far more slowly than any single tray, and the ratio is large enough that early studies looked for an explanation outside the balances. The explanation follows from the balances themselves. A change in the external flows must change the total holdup of light component in the column, and the products carry that inventory change away at a rate that produces a dominant time constant much longer than any residence time [2]. A change in the internal flows at constant external flows rearranges the profile without changing the inventory and is correspondingly faster. The same distinction determines which upsets are slow and which are fast, and it explains why one lag is usually enough to describe the response a controller acts on.

Composition control itself became a benchmark problem with the pilot column whose two-by-two transfer function matrix was identified from step tests in 1973 [11]. That model has been used ever since to pose the central structural question of column control, which is not how to tune a loop but which flows to pair with which composition. The reflux and boil-up pair, the distillate and boil-up pair and their variants behave very differently under disturbances and at high purity, and the choice interacts with the level loops underneath [12, 13]. Volume 13, Chapter 8 meets the same interaction among three purity loops on one shell. Holding both product compositions at once makes the interaction explicit, and high-purity columns are the standard illustration of an ill-conditioned plant, where the gain in one direction of input space is orders of magnitude larger than in another [14].

The measurement is the weakest element of the loop. An online analyzer is expensive to purchase, expensive to maintain and slow to respond, and the composition it reports is minutes old by the time a controller acts on it. Tray temperatures cost little, respond immediately and are only an indirect indication of composition, so the practical arrangement is to control a temperature and let the analyzer correct the temperature setpoint. Selecting which tray to control is a design problem with

Table 3.1 Variable Classification for the Distillate Composition Loop

Symbol	Role	Range	Unit
x_0	CV, distillate composition	0 to 1	mol frac
x_m	measured CV at the analyzer	0 to 1	mol frac
rr	MV, reflux ratio L/D	1.0 to 10.0	—
x_F	DV, feed composition	0.20 to 0.80	mol frac
F	DV, feed rate	0.50 to 1.50	mol/min
τ_a	DV, analyzer lag	0 to 2	min
$x_1 \ldots x_{30}$	tray liquid compositions	0 to 1	mol frac
x_{31}	bottoms composition, not controlled	0 to 1	mol frac
L, V, B, D	internal and product flows	—	mol/min

several competing criteria [6], and regression of composition on several temperatures at once is the standard way to sharpen the inference [5]. A temperature loop is not free of trouble either. Volume 13, Chapter 9 shows a stripping-tray temperature loop that holds one steady state and drives another the wrong way. None of that structure is needed when the analyzer is fast. All of it is justified by how much the loop deteriorates when the analyzer is slow, which is the quantity this chapter makes measurable.

The industrial motivation is energy. Distillation accounts for a large fraction of the separation energy used by the chemical and refining industries, the reboiler duty is the dominant operating cost of most columns, and reducing the purity giveaway of a poorly controlled column is one of the least expensive available savings [15]. A composition loop that is detuned because its analyzer is slow runs further from the specification limit, and the added distance is paid for in reboiler duty.

> ✦ **DID YOU KNOW? Forty Thousand Columns**
>
> A 2005 report for the United States Department of Energy counted about 40,000 distillation columns in operation in the United States. Distillation made up 90 to 95 % of all separations in chemicals and refining and is usually the method of choice when the relative volatility is 1.2 or more [16]. The column of this chapter has 1.6.

3.3 Process Model

Thirty trays, a total condenser with its reflux drum and a partial reboiler make up the thirty-two well-mixed stages of the simulated column, numbered from the top as shown in Figure 3.1. Stage 0 is the condenser and drum together, stages 1 to 15 are the rectifying trays, stage 16 is the feed tray, stages 17 to 30 are the stripping trays, and stage 31 is the reboiler. The single state of each stage is the mole fraction of the light component in its liquid holdup. Time is carried in minutes throughout the model, and the export writes its time column in the unit that suits the run length, minutes for most runs and hours for a run of an hour or more.

3.3.1 Variable Classification

The loop has one controlled variable, one manipulated variable and three disturbances, and each of them is an operator-set input while the simulation runs. Table 3.1 classifies them together with the states and the two outputs that are reported but not controlled. The bottoms composition is drawn, displayed and logged, and there is no bottoms controller, so the column is single ended.

3.3.2 Flows and Vapor-Liquid Equilibrium

Constant molar overflow is assumed, so the vapor rate is the same on every stage and the liquid rate is constant within each section. The distillate flow is fixed at half the feed, which makes the reflux

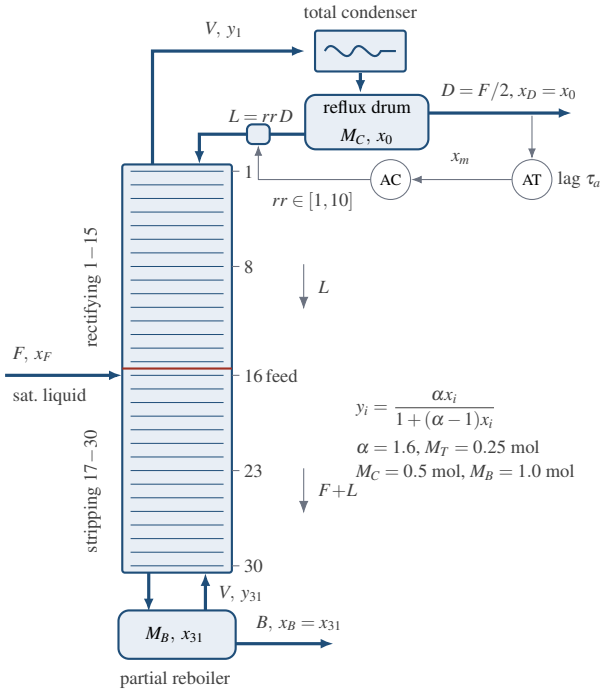

Figure 3.1 Thirty-two stage binary column with the composition loop closed on the distillate. The reflux ratio $rr = L/D$ is the only manipulated variable, the distillate flow is fixed at $D = F/2$, and the analyzer AT reports x_0 through a first-order lag τ_a before the controller AC sees it.

ratio the only degree of freedom, as shown in Equation (3.1).

$$\begin{aligned} D = \tfrac{1}{2}F, \qquad B = F - D, \qquad L = rr\,D \\ V = L + D, \qquad L_S = F + L \end{aligned} \tag{3.1}$$

where D is the distillate flow, B the bottoms flow, L the reflux and the liquid rate through the rectifying section, V the vapor rate through the whole column, and L_S the liquid rate through the stripping section. The feed is a saturated liquid, so it adds to the liquid traffic below the feed tray and leaves the vapor traffic unchanged.

The vapor leaving a stage is in equilibrium with the liquid on it, and the equilibrium is written with a constant relative volatility as shown in Equation (3.2).

$$y_i = \frac{\alpha x_i}{1 + (\alpha - 1)x_i}, \qquad i = 0, 1, \ldots, 31 \tag{3.2}$$

where y_i is the vapor mole fraction of the light component leaving stage i, x_i is the liquid mole fraction on that stage and α is the relative volatility. The default value $\alpha = 1.6$ is a difficult separation. It takes all thirty trays to reach a distillate of 0.9354 at a reflux ratio of 3, and α is adjustable between 1.1 and 3.0 so that the separation difficulty of the same column can be varied over a wide range.

3.3.3 Component Balances

Each stage is a well-mixed liquid holdup in equilibrium with the vapor leaving it, so one component balance per stage closes the model. The condenser receives the overhead vapor and returns liquid of its own composition to the column and to the product, as shown in Equation (3.3).

$$M_C \frac{\mathrm{d}x_0}{\mathrm{d}t} = V\,(y_1 - x_0) \tag{3.3}$$

where M_C is the combined condenser and reflux drum holdup. The condenser is total, so the reflux and the distillate both leave at x_0 and the L and D terms collapse into $V = L + D$.

A rectifying tray receives liquid from the tray above and vapor from the tray below, as shown in Equation (3.4).

$$M_T \frac{\mathrm{d}x_i}{\mathrm{d}t} = L\left(x_{i-1} - x_i\right) - V\left(y_i - y_{i+1}\right), \quad i = 1, \ldots, 15 \tag{3.4}$$

where M_T is the liquid holdup of one tray. The feed tray adds the feed and passes the larger stripping liquid rate downward, as shown in Equation (3.5).

$$M_T \frac{\mathrm{d}x_{16}}{\mathrm{d}t} = F x_F + L x_{15} - L_S x_{16} - V\left(y_{16} - y_{17}\right) \tag{3.5}$$

A stripping tray is the rectifying tray with the larger liquid rate, as shown in Equation (3.6).

$$M_T \frac{\mathrm{d}x_i}{\mathrm{d}t} = L_S\left(x_{i-1} - x_i\right) - V\left(y_i - y_{i+1}\right), \quad i = 17, \ldots, 30 \tag{3.6}$$

The reboiler receives the liquid leaving tray 30, returns vapor of composition y_{31} to the column and withdraws the bottoms product, as shown in Equation (3.7).

$$M_B \frac{\mathrm{d}x_{31}}{\mathrm{d}t} = L_S x_{30} - B x_{31} - V y_{31} \tag{3.7}$$

where M_B is the reboiler holdup. The reboiler is partial, so the vapor it returns is in equilibrium with the sump liquid and the bottoms product leaves at x_{31}. Every term in Equations (3.3) to (3.7) is a molar flow of light component in mol/min, and dividing by a holdup in mol leaves a rate of change of mole fraction in 1/min.

Adding Equations (3.3) to (3.7) cancels every internal term and leaves the overall balance of Equation (3.8).

$$\sum_{i=0}^{31} M_i \frac{\mathrm{d}x_i}{\mathrm{d}t} = F x_F - D x_0 - B x_{31} \tag{3.8}$$

This is the constraint that governs the whole chapter, and it is derived again in Section 3.4.

3.3.4 Instrument

The controller does not act on x_0. It acts on an analyzer reading that carries a first-order lag and additive Gaussian noise, as shown in Equation (3.9).

$$\tau_a \frac{\mathrm{d}x_a}{\mathrm{d}t} = x_0 - x_a, \qquad x_m = \mathrm{clip}_{[0,1]}(x_a + \sigma\varepsilon) \tag{3.9}$$

where x_a is the analyzer state, τ_a is the analyzer lag in min, x_m is the reading passed to the controller, σ is the noise standard deviation in mole fraction and ε is drawn from the unit normal distribution independently at every controller execution, through a Box-Muller transform of the seeded generator of Section 1.6.1, so a run with noise is as reproducible as one without. Setting $\tau_a = 0$ collapses the first expression and the analyzer reports x_0 directly. The lag is implemented as a disturbance rather than as a fixed property of the instrument, so it can be changed while the column is running and the column itself is untouched by the change.

3.3.5 Parameters, Limits and Initial Conditions

Table 3.2 lists every constant of the column, the analyzer and the PID, with the three settings of the predictive controller, whose model and filter are given in Section 3.5.4. The three holdups, the relative volatility and the analyzer noise are adjustable, which allows the tray time constant and the separation difficulty to be varied independently. The reflux ratio limits are hard clamps applied both to the manual input and to the controller output.

Table 3.2 Model Parameters, Limits and Defaults

Sym.	Meaning	Value	Unit
Adjustable			
α	relative volatility (1.1-3.0)	1.6	—
M_T	tray holdup (0.10-1.00)	0.25	mol
M_C	condenser holdup (0.20-2.00)	0.50	mol
M_B	reboiler holdup (0.40-3.00)	1.00	mol
σ	analyzer noise SD (0-0.010)	0.0015	m.f.
Fixed by the model			
N	modeled stages	32	—
n_F	feed stage index	16	—
$rr_{\min}$	lower reflux ratio limit	1.0	—
$rr_{\max}$	upper reflux ratio limit	10.0	—
$\sum M_i$	total liquid holdup	9.00	mol
Operator settings and defaults			
F	feed rate (0.50-1.50)	1.00	mol/min
x_F	feed composition (0.20-0.80)	0.50	m.f.
τ_a	analyzer lag (0-2, step 0.05)	0	min
x_0^{SP}	setpoint (0.700-1.000)	0.970	m.f.
—	band about the setpoint	$\pm$0.010	m.f.
rr	manual reflux ratio, case 1 (step 0.1)	3.0	—
Controller			
K_c	prop. gain (0-300, step 0.6)	60	1/m.f.
τ_I	integral time (0.002-100, 100 is off)	4	min
τ_D	derivative time (0-2)	0	min
u_{bias}	bias, reflux ratio at the start	rr_{man}	—
τ_t	anti-windup tracking time	τ_I	min
Δt	controller sample period	0.02	min
h	integration step	0.002	min
Δt_s	history sample interval	0.2	min
T_s	MPC decision period	0.2	min
T_p	MPC horizon (4-60)	30	min
λ	MPC move suppression (0-3)	0.3	—

The abbreviation m.f. stands for mole fraction.

Every case starts at steady operation. The application solves the thirty-two balances of Equations (3.3) to (3.7) for their steady profile by Newton's method, with the parameters and disturbances in force. A case in manual starts at the steady profile of its manual reflux ratio, and a case in automatic starts at the steady profile of its setpoint, with the manual reflux ratio set to the value that holds it. The reference operating point case starts in manual at $rr = 3$, $F = 1$ and $x_F = 0.5$, where the profile runs from $x_0 = 0.93542$ to $x_{31} = 0.06458$ and stays there exactly, and the other nine cases start at the 0.970 setpoint with $rr = 3.7319$. A setting changed before the run starts, whether the mode, the setpoint, the manual reflux ratio, a parameter or a disturbance, re-initializes the case exactly as a reset would, with the Reset semantics of Section 1.6.1. It therefore changes where the run starts rather than producing a step, and a step or an upset is made after the run has started.

3.3.6 Numerical Integration and Timing

Integration is explicit Euler on a step of $h = 0.002\,\mathrm{min}$, in place of the Runge–Kutta step of Section 1.6.1. The simulation clock advances the column one step of h at a time, every state is clipped to the physical interval $[0, 1]$ after each step, and the analyzer state of Equation (3.9) is advanced on the same clock. The integration step is not a tuning choice. Linearizing the thirty-two balances about

the $rr = 3$ profile gives a fastest eigenvalue of −39.50/min, so the explicit Euler stability limit is $2/39.50 = 0.0506$ min and the integration step sits twenty-five times inside it. The slowest eigenvalue is −0.0656/min, a time constant of 15.24 min, so the stiffness ratio across the state vector is 602.

The controller executes every $\Delta t = 0.02$ min, which is ten integration steps, and one sample is written to the history every $\Delta t_s = 0.2$ min, which is every tenth execution. A sample stamped t carries the analyzer reading the controller acts on at t and the reflux ratio held over the execution period that ends at t, so a reflux move made at an execution first appears in the next sample stamped after it. A change the operator makes between executions reaches the column at the next execution. The speed setting runs from 0.5 min/s to 10 min/s and, as Section 1.6.1 explains, changes the pace of a run and never its result.

3.3.7 Assumptions

The model is deliberately the smallest one that reproduces the behavior under study, and each simplification has a limited domain of validity.

Constant molar overflow. The vapor rate is taken as constant up the whole column and the liquid rate as constant within each section, which removes the energy balance from the problem. This requires equal molar latent heats and negligible heat loss and sensible heat effects. It breaks down for pairs with very different latent heats, and it hides the fact that a real reflux change alters the internal traffic only after the condenser and reboiler duties have responded.

Constant holdup. Every M_i is a constant, so there is no hydraulic dynamics, no downcomer lag, no weir relationship and no level loop on the drum or the sump. Real columns show a liquid hydraulic lag of a few seconds per tray that delays a reflux change on its way down, and thirty of those in series are not negligible against a 7 min composition response. The model therefore reaches the lower trays faster than a real column would.

Fixed distillate split. The distillate flow is $D = F/2$ at all times, as though a perfect ratio controller held it there. This is what makes the reflux ratio the only degree of freedom, and it is also what pins the sum of the two product compositions. A real column has an independent draw, and the choice of which flows to pair with which composition is the configuration question the model does not pose.

Saturated liquid feed. The feed adds its whole rate to the liquid traffic below the feed tray and nothing to the vapor traffic, which is the $q = 1$ case. A subcooled or partially vaporized feed would change the split between L_S and V and move the feed tray optimum.

Constant relative volatility and ideal stages. Equation (3.2) assumes α does not vary with composition, temperature or pressure, and each stage is assumed to reach equilibrium exactly. Real trays have Murphree efficiencies well below unity, and a column running under a pressure profile has a volatility that changes from top to bottom.

Binary mixture and no energy balance. There are two components and no temperatures anywhere in the model. Inferential temperature control, which is what most columns actually run on, cannot be demonstrated on this simulator for that reason.

Instrument. The analyzer is a first-order lag with white measurement noise. A real chromatograph is closer to a sampler with a dead time equal to its cycle (Section 1.4.5), and a pure dead time is harder for a loop to tolerate than a lag of the same magnitude. The results in Section 3.6 therefore understate the damage a real analyzer does.

3.4 Dynamic Analysis

3.4.1 Steady-State Coupling of the Product Compositions

Setting the derivatives in Equation (3.8) to zero and substituting $D = B = F/2$ gives the steady-state relation of Equation (3.10).

$$x_0 + x_{31} = 2x_F \tag{3.10}$$

The relation holds exactly and independently of the reflux ratio, the number of trays, the holdups and the relative volatility, because it is nothing more than the overall component balance with a fixed split.

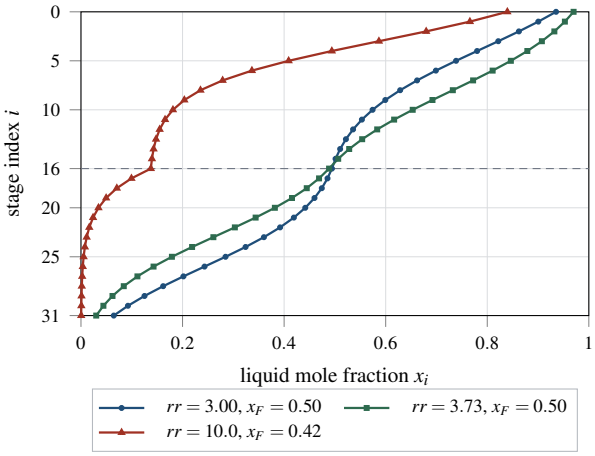

Figure 3.2 Converged liquid composition profiles from Equations (3.3) to (3.7), with circles at $rr = 3.00$, squares at $rr = 3.73$ and triangles at $rr = 10.0$ with $x_F = 0.42$. The dashed line marks the feed stage. Raising the reflux ratio from 3.00 to 3.73 lifts the top and drops the bottom by the same 0.0346, while a feed of $x_F = 0.42$ leaves the distillate pinned at 0.8395 even with the reflux ratio at its 10.0 limit.

The reflux ratio therefore does not change the amount of light component available. It determines only how that amount is divided between the two products.

Two consequences follow. The first is a ceiling. Since $x_{31} \geq 0$, the distillate composition can never exceed $2x_F$ at steady state, so the default setpoint of 0.970 stops being reachable as soon as the feed composition falls below 0.485. The second is a check on every steady state the application produces. At $rr = 3$ the profile settles at $x_0 = 0.9354$ and $x_{31} = 0.0646$, and at $rr = 3.5$ it settles at 0.9622 and 0.0378, both pairs summing to 1.0000 against a feed of 0.5000.

3.4.2 Modal Decomposition of the Linearized Column

The steady-state profiles at $rr = 3$ and $rr = 3.5$ differ by $\Delta x_0 = +0.0267$ and $\Delta x_{31} = -0.0267$, and the holdup-weighted sum of the changes across all thirty-two stages is -5.1×10^{-4} mol, which is 0.012 % of the 4.244 mol of light component the column holds at $rr = 3$. Raising the reflux ratio therefore rearranges the profile without changing the inventory, as shown in Figure 3.2. That is a direct consequence of fixing $D = F/2$. The external flows never move, so the light component that leaves the top faster must leave the bottom slower by the same amount.

Linearizing the balances about the $rr = 3$ profile makes the same point in the frequency domain. All thirty-two eigenvalues are real, so the open-loop plant cannot oscillate, and the two slowest are 15.24 min and 8.17 min. Their eigenvectors are different in kind. The 15.24 min eigenvector has the same sign on every stage, so it is the inventory mode in which the whole column becomes richer or leaner together. The 8.17 min eigenvector changes sign once, near the feed stage, so it is the profile mode in which the top gains what the bottom loses. Decomposing the steady-state gain of x_0 with respect to rr across the thirty-two modes assigns 109 % of it to the 8.17 min profile mode and −2.8 % to the 15.24 min inventory mode, with the remaining thirty modes contributing the small negative balance between them. The slowest dynamics in the column are essentially invisible from the reflux ratio, and a feed composition upset, which does change the inventory, is what brings them out.

3.4.3 Origin of the Apparent First-Order Response

The liquid residence time of a single stage is its holdup divided by the liquid flow through it. At $rr = 3$ and $F = 1$ those are 0.167 min on a rectifying tray, 0.100 min on a stripping tray, 0.250 min in the

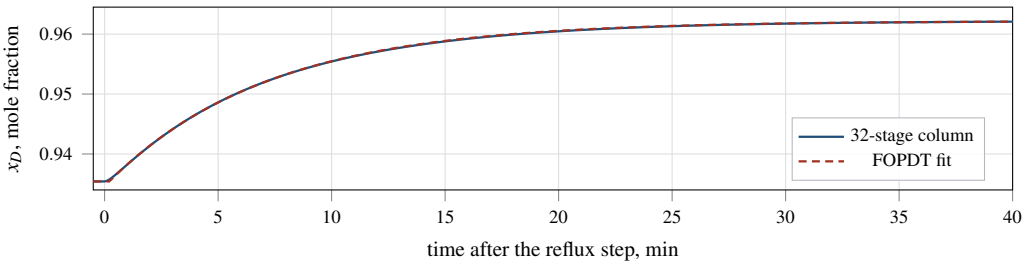

Figure 3.3 Open-loop response of the distillate composition to a reflux ratio step from 3.0 to 3.5 with the analyzer noise set to zero. The solid curve is the column, and the dashed curve is the fitted first-order plus dead-time model, with $K_p = 0.0535$, $\tau_p = 7.07\,\text{min}$, $\theta_p = 0.21\,\text{min}$.

condenser and 0.400 min in the reboiler. No stage in the column holds liquid for longer than 0.40 min, and the distillate responds to a reflux step with an apparent time constant of 7.07 min. The factor of eighteen is not a delay accumulating down a chain. It is the profile mode, whose time constant is set by the holdup that must be redistributed against the separating action of each tray, and it is the reason a column responds far more slowly than any individual stage.

3.4.4 Step-Test Identification

The step test follows the standard identification procedure of a first course [7]. In the reference operating point case, with the analyzer noise set to zero and the loop in manual, the reflux ratio is held at 3.0 on the already converged profile for 5 min and then stepped to 3.5, which gives the response of Figure 3.3. The distillate moves from 0.9354 to 0.9622 in the 60 min after the step for a reflux ratio change of 0.5, so the fitted process gain is $K_p = 0.0535$ mole fraction per unit of reflux ratio. The response reaches 28.3 % of its final value at 2.56 min and 63.2 % at 7.28 min, so the two-point fit gives $\tau_p = 7.07\,\text{min}$ and $\theta_p = 0.21\,\text{min}$. The fitted curve stays within 3.3×10^{-4} mole fraction of the column, 1.2 % of the step.

The apparent dead time of 0.21 min is not a transport delay. Nothing in Equations (3.3) to (3.7) delays anything, and the composition on stage 0 starts moving on the first integration step after the reflux changes. What the two-point fit reads as dead time is the thirty fast eigenvalues rounding the foot of the response, and it is of the same order as the sum of the fastest stage residence times. This is the fitted dead time that is not transport of Section 1.4.2, here on a plant that contains no delay at all. Chapter 2 fits a dead time to a plant that does carry a transport delay and finds the fit far longer than the delay, for the same reason.

Pause. Try Exercise 3.2 before reading on.

3.4.5 Steady-State Gain Nonlinearity

Holding a reflux ratio to convergence at the base feed maps the manipulated variable onto the steady distillate composition, and the map is strongly curved. The model source tabulates it to four decimal places in steps of 0.5 from $rr = 1$ to $rr = 10$, and the table runs 0.7306, 0.8431, 0.9354, 0.9765, 0.9883 and 0.9971 at reflux ratios of 1.0, 2.0, 3.0, 4.0, 5.0 and 10.0. The exact steady states show half a unit of reflux raising the distillate by 0.0571 mole fraction at $rr = 1$ and by 0.00022 at $rr = 9.5$, a factor of 256, which the rounded table would put at 190. The secant gain of 0.0535 measured over the 3.0 to 3.5 step is already 22 % below the local gain of 0.0687 at $rr = 3$, so even a half-unit step test averages over a visible amount of curvature. The same half-unit step test made at $rr = 6$, with the manual reflux ratio set to 6.0 before the run so that the case starts at that steady state, returns a gain of 0.0024, more than an order of magnitude below the 0.0535 at $rr = 3$, on a plant whose equations contain no explicit nonlinearity other than Equation (3.2). A tuning computed at 0.970 is far too slow

at 0.90 and far too aggressive near 0.99, which is the standard argument for controlling a purity loop on the logarithm of the impurity rather than on the composition itself.

3.4.6 Controller Tuning From the Fitted Model

Internal model control for a first-order plus dead-time plant (Section 1.4.3) sets the proportional gain and reset time as shown in Equation (3.11).

$$K_c = \frac{\tau_p}{K_p\,(\tau_c + \theta_p)}, \qquad \tau_I = \min(\tau_p,\, 4\,(\tau_c + \theta_p)) \tag{3.11}$$

with τ_c the specified closed-loop time constant [17]. Specifying $\tau_c = 2\,\text{min}$ with the fitted model gives $K_c = 59.8$ and $\tau_I = 7.07\,\text{min}$. The default tuning is $K_c = 60$ and $\tau_I = 4\,\text{min}$, which is the same proportional gain with the integral time shortened to 4 min, and the exercises of Section 3.13 compare the two.

3.5 Controller Design

3.5.1 Velocity-Form Algorithm

The loop is the velocity-form PID of Section 1.6.2, with the derivative on the measurement. The complete algorithm as coded is given by Equation (3.12), in the loop's own symbols.

$$e_k = x_0^{\text{SP}} - x_{m,k} \tag{3.12a}$$

$$\Delta s_k = K_c\,(e_k - e_{k-1}) - K_c\,\tau_D\,\frac{x_{m,k} - 2x_{m,k-1} + x_{m,k-2}}{\Delta t} \tag{3.12b}$$

$$s_k = s_{k-1} + \Delta s_k \tag{3.12c}$$

$$c_k = u_{\text{bias}} + s_k + I_{k-1} \tag{3.12d}$$

$$rr_k = \text{clip}_{[1,10]}\,(c_k) \tag{3.12e}$$

$$I_k = I_{k-1} + \frac{K_c}{\tau_I}\,e_k\,\Delta t + \frac{\Delta t}{\tau_t}\,(rr_k - c_k) \tag{3.12f}$$

where e_k is the error in mole fraction, s_k is the accumulated proportional and derivative contribution, c_k is the demanded output before the limits are applied, rr_k is the controller output in reflux ratio units at execution k, I_k is the integral state, $\Delta t = 0.02\,\text{min}$ is the sample period, u_{bias} is the manual reflux ratio in force when the case starts or is reset and $\tau_t = \tau_I$ is the tracking time of the back-calculation. The integral used at execution k was accumulated through execution $k-1$, and the derivative acts on the unfiltered reading. Setting τ_I to the top of its range, 100 min, switches integral action off. The integral gain is then zero, the back-calculation term is dropped with it, and the loop is proportional only. In the symbols of Equations (1.15) to (1.18), $s_k = S_k$, $c_k = v_k$, $rr_k = u_k$, $u_{\text{bias}} = u_b$, $\tau_t = T_t$ and $a = +1$. The code forms each increment as the difference of the positional terms for the gains in force, so a gain changed during a run takes effect at the next execution, as in Section 1.6.2.

The controller is reverse acting in the instrument-engineering sense. Raising the reflux ratio raises the distillate composition, so the process gain is positive, and a rise in the measured composition lowers the error and with it the reflux ratio.

3.5.2 Initialization and Bumpless Transfer

A velocity-form recursion has no bias of its own, and its starting output is set by the bias and the integral state. When a case starts or is reset, the bias is set to the manual reflux ratio in force and the integral to zero, so the first execution in automatic returns the positional expression of Equation (3.13).

$$rr_0 = \text{clip}_{[1,10]}\,(rr_{\text{man}} + K_c\,e_0) \tag{3.13}$$

where rr_{man} is the manual reflux ratio. A case that starts in automatic starts at the steady state of its setpoint with rr_{man} at the reflux ratio that holds it, so e_0 is zero apart from the analyzer noise

and the first execution leaves the reflux ratio where it is. Transfers between manual and automatic are bumpless in both directions, as in Section 1.6.2. A switch to automatic adds only $K_c(e_k - e_{k-1})$, which is zero for an unchanged setpoint and a steady reading and is K_c times the setpoint change when the setpoint is raised at the same instant.

A setpoint raised in automatic therefore moves the reflux ratio at the next execution by K_c times the change, the proportional kick that opens the setpoint runs of Section 3.6, and integral action then carries the output to the reflux ratio that holds the new setpoint, which is $rr = 3.7319$ for a distillate of 0.970 at the base feed.

3.5.3 Anti-Windup by Back-Calculation

Anti-windup is the back-calculation term of Equation (3.12). While the output is clipped, $(\Delta t/\tau_t)(rr_k - c_k)$ pulls the integral toward the value that reproduces the applied reflux ratio with a tracking time of $\tau_I = 4\,\text{min}$, so each execution removes 1/200 of the clipped excess and the accumulator s_k is never truncated. With $\tau_D = 0$ and a steady error, the integral settles where the demand exceeds the limit by exactly $K_c e_k$, because $(K_c/\tau_I)\,\tau_t = K_c$. The demand therefore winds no further past the limit than the proportional term, and once the integral has settled the output leaves the limit as the error changes sign. Run XF of Section 3.6 holds the reflux ratio on its upper limit for more than 1200 min, and the demand there settles at 17.83, the limit of 10.0 plus $K_c e = 7.83$. Away from the limits the recursion is algebraically identical to the positional law with the same gains.

The upper limit has a second effect that has nothing to do with windup. At $rr = 10$ the steady distillate is 0.9971 and the local gain is close to zero, so a controller operating near saturation is close to open loop even before the clip engages.

3.5.4 Model Predictive Control

A linear model predictive controller (MPC) is available as a second automatic controller for the distillate. The PID stays the default, and only case 9 selects the MPC. Every $T_s = 0.2\,\text{min}$, or ten controller executions, the MPC optimizes the reflux ratio over a prediction horizon, sends only the first move to the column, keeps that move in force for the period and then optimizes again, the arrangement of Section 1.6.3 [18, 19]. Its model was identified on the simulator at the operating point of the automatic cases, $\bar{x}_0 = 0.970$ with $\overline{rr} = 3.7319$ and the base feed, by averaging the responses to reflux ratio steps of $+0.2$ and -0.2 with the analyzer noise and lag at zero and fitting two first-order lags in series, as shown in Equation (3.14).

$$\begin{aligned} \tau_1 \frac{\mathrm{d}z_1}{\mathrm{d}t} &= -z_1 + K_m\,(rr - \overline{rr} + \delta) \\ \tau_2 \frac{\mathrm{d}z_2}{\mathrm{d}t} &= z_1 - z_2, \qquad \tilde{x}_0 = \bar{x}_0 + z_2 \end{aligned} \tag{3.14}$$

where z_1 and z_2 are model states in mole fraction, $\tilde{x}_0$ is the predicted distillate composition, δ is an unmeasured disturbance in units of reflux ratio, $K_m = 0.02928$ per unit of reflux ratio, $\tau_1 = 7.115\,\text{min}$ and $\tau_2 = 0.566\,\text{min}$. The fit needs no dead time, and its root-mean-square residual of 1.4×10^{-5} is 0.23 % of the response.

The fit restates Section 3.4 at a higher purity. K_m lies 1.4 % above the local gain of 0.02889 that the thirty-two balances give at 0.970, because a ± 0.2 step averages over the curvature of Section 3.4.5, and 45 % below the 0.0535 measured at 0.9354 in Section 3.4.4. The second lag carries the fast modes that the fit of Section 3.4.4 read as 0.21 min of dead time, and a first-order plus dead-time fit of the same response reads 0.48 min instead, with a residual 1.6 times larger. The model contains neither the analyzer lag nor the feed composition, so an analyzer lag is a model error and a feed upset reaches the controller only through the analyzer.

The estimate $\hat{\delta}$ removes offset. A steady-state Kalman filter on Equation (3.14) treats δ as a random walk, the disturbance state of Section 1.4.7 that makes a predictive controller offset-free [20], and it

updates the estimate at every execution whichever controller is acting, as shown in Equation (3.15).

$$\begin{aligned} \bar{z}_k &= \Phi\,\hat{z}_{k-1} + \gamma\left(rr_{k-1} - \overline{rr} + \hat{\delta}_{k-1}\right) \\ \eta_k &= x_{m,k} - \bar{x}_0 - \bar{z}_{2,k} \\ \hat{z}_k &= \bar{z}_k + \ell_z\,\eta_k, \qquad \hat{\delta}_k = \hat{\delta}_{k-1} + \ell_\delta\,\eta_k \end{aligned} \tag{3.15}$$

which is the prediction and correction of Equation (1.11) with steady gains, where Φ and γ are Equation (3.14) sampled with a zero-order hold over $\Delta t = 0.02\,\text{min}$, rr_{k-1} is the reflux ratio applied over the period just ended, $x_{m,k}$ is the analyzer reading that the PID also receives, η_k is the innovation, and $\ell_z = [0.0547 \;\; 0.0360]^\mathsf{T}$ and $\ell_\delta = 10.60$ per unit of mole fraction are the converged Riccati gains for a measurement variance $\sigma^2 = 2.25 \times 10^{-6}$, a state noise variance of $10^{-6}\sigma^2$ and a disturbance noise variance $q_\delta = 0.1\,\sigma^2/K_m^2 = 2.62 \times 10^{-4}$. A richer feed enters $\hat{\delta}$ as the added reflux ratio that would raise the distillate as much.

At each decision a quadratic program on the model of Equation (3.14), sampled at T_s over a horizon $T_p = 30\,\text{min}$ of $N_p = 150$ steps, chooses twelve reflux ratios v_j, held over blocks of $(1,1,1,1,2,2,3,4,6,9,15,105)$ steps, to minimize Equation (3.16) within the hard limits $1.0 \le v_j \le 10.0$.

$$J = \frac{T_s}{\tau_n}\sum_{k=1}^{N_p}\left(\frac{\tilde{x}_{0,k} - x_0^{\text{SP}}}{y_s}\right)^2 + \lambda\sum_{j=1}^{12}\left(\frac{v_j - v_{j-1}}{u_s}\right)^2 \tag{3.16}$$

where $\tilde{x}_{0,k}$ is predicted k steps ahead with $\hat{\delta}$ held, v_0 is the reflux ratio applied now, $\tau_n = 7.7\,\text{min}$, near $\tau_1 + \tau_2$, keeps the first sum independent of the horizon, $y_s = 0.01$ is the half-width of the specification band, $u_s = 9$ is the reflux span and $\lambda = 0.3$ is the move suppression, for weights of 259.7 per squared mole fraction of predicted error and 3.7×10^{-3} per squared move. The symbol λ keeps the app's name for the move suppression here and is not the bias weight of Chapter 1. The program has no specification band, rate limit, terminal weight or feed preview. Its decisions are departures from $\overline{rr}$, a ridge of 10^{-9} relative to the weights keeps it strictly convex, and a dual active-set method [21] solves it. One tuning rule fixed λ and the factor 0.1 in q_δ for each of the seven single-loop studio apps before any comparison was run. Among the settings that move the reflux no more than the PID does in cases 3 and 4, it keeps the largest λ whose true IAE is within 10 % of the lowest.

A plan is rejected on any of three outcomes (a solver status other than optimal, a first move that is not a finite number, or a first move outside the reflux limits). The PID then takes the reflux ratio, and because it follows the applied value at every execution, in the same way that it follows the operator's setting in manual (Section 3.5.2), the transfer is bumpless, and the PID keeps the reflux until a later plan succeeds. When a rejection is imposed 6 min into case 9 with the noise removed, the handover changes the reflux ratio by 0.0006, less than the 0.0008 of each following PID move, and the plan that resumes control changes it by 0.067. The first plan is solved as a case loads, so the controller acts from $t = 0$. It receives only the analyzer reading, its own reflux ratio, the setpoint and the reflux limits, and its constants are fixed, so the thirty-two balances, the fit behind the PID tuning and Equation (3.14) remain three separate models, and a parameter the operator changes reaches only the first.

3.5.5 Moving Horizon Estimation

Each plan may instead start from a moving horizon estimator (MHE), chosen in place of the filter of Equation (3.15) and selected by no case. It solves the window problem of Equation (1.13), which can keep the fitted states inside hard bounds [22]. Writing $\zeta = [z_1 \;\; z_2 \;\; \delta]^\mathsf{T}$ for the augmented state of the

filter, execution k fits the readings $j = s, \ldots, k$, with $s = k - N_e$, by solving Equation (3.17).

$$\begin{aligned} \min_{\zeta_s,\, \omega} \; & \| \zeta_s - \bar{\zeta}_s \|^2_{\bar{P}^{-1}} + \sum_{j=s}^{k-1} \omega_j^{\mathsf{T}} Q^{-1} \omega_j \\ & + \frac{1}{\sigma^2} \sum_{j=s}^{k} \left(x_{m,j} - \bar{x}_0 - z_{2,j} \right)^2 \\ \text{s.t. } & \zeta_{j+1} = \Phi_e \, \zeta_j + \Gamma_e \, (rr_j - \overline{rr}) + \omega_j \\ & 0 \le \bar{x}_0 + z_{i,j} \le 1, \qquad i = 1, 2 \end{aligned} \tag{3.17}$$

In Equation (3.17) $\Phi_e = \left[\begin{smallmatrix} \Phi & \gamma \\ 0 & 1 \end{smallmatrix}\right]$ and $\Gamma_e = [\gamma^{\mathsf{T}} \;\; 0]^{\mathsf{T}}$ repeat the prediction of Equation (3.15), rr_j is the reflux ratio applied after reading j, ω_j perturbs each of the three states, and $Q = \mathrm{diag}(10^{-6}\sigma^2, 10^{-6}\sigma^2, q_\delta)$ and σ^2 are the variances the filter was designed with. The first term ties the window's first state to $\bar{\zeta}_s$, the MHE's own estimate from execution $s-1$ advanced one period, through the steady prior covariance $\bar{P}$ of the filter, held constant. With $N_e = 10$ the window spans 0.2 min, one decision period, and eleven readings, so the dual active-set method of Section 3.5.4 receives 33 decisions and 44 bound rows.

Only declared values set the bounds. A mole fraction lies between 0 and 1, and the same interval is imposed on z_1, the composition that the second lag follows, so in deviation form both states are confined to -0.97 to 0.03. Nothing declared limits δ, which stays free. With no bound active the window problem carries the filter's model, weights and prior and returns the filter's estimate. Should a window stop short of an optimum, return a non-finite value or break a bound, the estimate falls back to the MHE's one-period prediction corrected as the filter would correct it, clipped to the bounds, and a flag in the exported history shows each such case. Whenever the MHE is chosen, at the start or in mid-run, its window opens from the filter's own prediction. The filter keeps running under either choice, and the MHE reads the analyzer, the applied reflux ratio and declared constants only.

3.5.6 Relay Autotuning

A button starts the relay feedback experiment of [23] while the PID holds the distillate in automatic, and none of the ten cases runs it. The relay takes the place of the PID and holds the reflux ratio at $u_0 + d_r$ until the reading rises more than h_r above the setpoint, then at $u_0 - d_r$ until it falls more than h_r below, and so on. The center u_0 is the mean reflux ratio applied over the last 100 executions, 2 min or a tenth of the trend window, so that no single noisy output fixes it. The amplitude $d_r = 0.45$ is five percent of the reflux span, or 0.9 of the distance from u_0 to the nearer limit if that is smaller, and the band $h_r = 3\sigma_0 = 0.0045$ is three times the declared analyzer noise σ_0 whatever noise the run carries. A cycle with a half shorter than four executions is not counted. Once the last two counted cycles match to 5 %, both in period and in the swing $\pm a_r$ of the reading, the first harmonic of the two applied levels, with a correction for the band, yields the ultimate gain K_{cu}, and the Tyreus–Luyben rule converts it and the period P_u into the PI constants of Equation (3.18).

$$K_{cu} = \frac{4 d_r \sin(\pi f_h)}{\pi \sqrt{a_r^2 - h_r^2}}, \qquad K_c = \frac{K_{cu}}{3.2}, \qquad \tau_I = 2.2\, P_u \tag{3.18}$$

Here f_h is the fraction of a cycle spent at the upper level, one half for a symmetric cycle, and K_{cu} has the unit of K_c, the reciprocal of a mole fraction. The written constants carry three significant figures, with $\tau_D = 0$, each inside its input range and τ_I below the 100 min that switches integral action off, and they stay editable. An unconverged test ends after 30 cycles or 40 min, and a second press or a switch to manual or to the MPC cancels it. The PID follows the applied reflux ratio throughout, as in manual (Section 3.5.2), so it resumes without a bump. Because the band delays every switch, the cycle settles where the loop's phase falls $\psi = \arcsin(h_r / a_r)$ short of -180°, and the result note flags any cycle with $\psi \ge 10^\circ$ as one that the band, not the column, has fixed.

Table 3.3 Runs Built on the Reference Operating Point Case

Run	Mode	Change applied	Horizon
STEP	manual	rr 3.0 to 3.5 at 5 min	65 min
L0	auto	SP 0.970 at 2 min, $\tau_a = 0$	62 min
L0.5	auto	L0 with $\tau_a = 0.5$ min	62 min
L1	auto	L0 with $\tau_a = 1$ min	62 min
L1.5	auto	L0 with $\tau_a = 1.5$ min	62 min
L2	auto	L0 with $\tau_a = 2$ min	62 min
D2	auto	L2 with $K_c = 36$, $\tau_I = 5$ min	62 min
N0	auto	L0 with $\sigma = 0.0015$	62 min
XF	auto	x_F 0.50 to 0.42 at 2 min	1296 min
FR	auto	F 1.00 to 1.50 at 2 min	62 min
SF	manual	STEP repeated at $F = 1.50$	65 min

⌛ HISTORY **1984: A Relay Finds the Critical Point**

Most regulators in industry are of the PID type, a large plant may have hundreds of them, and a paper of 1984 observed that many are poorly tuned in practice, in part because simple and robust tuning methods had not been available. The classic Ziegler–Nichols procedure raises the gain of a proportional controller until the loop oscillates and reads the critical gain and period, an experiment that is hard to automate with the swing kept under control. The 1984 answer put a relay in the loop instead. A process with a phase lag of at least 180° at high frequencies can then oscillate at its critical period, and the first harmonic of the relay output gives the critical gain from the relay amplitude and the swing of the output. A relay with hysteresis is less sensitive to measurement noise [23]. The relay of this chapter has such a band.

3.6 Results

3.6.1 Protocol

Every run below starts from the reference operating point case, which loads the column at its steady profile in manual at a reflux ratio of 3.0. The noise, the analyzer lag, the tuning and the mode listed in Table 3.3 are chosen before the run starts, so the case starts at steady operation with those settings, and every step or upset is made after the run has started, as an operator would at the controls. Runs L0 to N0 switch to automatic before the run and so hold the case's own setpoint of 0.9354 with the reflux ratio at 3.0 until the setpoint is raised to 0.970 at 2 min. Runs XF and FR set the 0.970 setpoint before the run, start at the steady state that holds it and step the feed at 2 min. A change made at 2 min reaches the column at the controller execution at 2.00 min, and times below are measured from that execution. The runs were driven through the application's own simulation engine, the one the browser runs, with every change applied through the same calls the controls make. The application's performance panel, the scorecard of Section 1.6.4, reports the integral absolute error, the overshoot, the settling time, the reflux travel and the time at a limit, computed on the analyzer reading from the 0.2 min history against a band of ± 0.01 mole fraction. The measures below are defined on the true distillate composition instead and computed at the 0.002 min integration step over the 60 min after the step. Overshoot is quoted as a percentage of the setpoint step. Settling is the last time the true distillate composition is further than 0.002 mole fraction from setpoint. Travel is $\sum_k |rr_k - rr_{k-1}|$ over every controller execution from the step on. The exported history gives integral absolute errors 2 to 4 % lower and settling times up to 0.2 min earlier, because it is sampled ten times less often. Except where stated, the analyzer noise was set to zero. The settings chosen before the run and the same settings carried through a reset give the same run, as Section 1.6.1 states for every app.

Table 3.4 One Tuning ($K_c = 60$, $\tau_I = 4$ min) Across Five Analyzer Lags

Run	τ_a min	peak m.f.	over %	settle min	IAE m.f. min	travel —
L0	0	0.9726	7.6	13.5	0.0945	3.54
L0.5	0.5	0.9738	10.9	13.7	0.0928	3.87
L1	1	0.9758	16.8	13.9	0.1050	4.24
L1.5	1.5	0.9783	24.0	14.4	0.1312	4.73
L2	2	0.9806	30.7	23.6	0.1687	5.32
D2	2	0.9752	15.0	20.1	0.1487	2.43

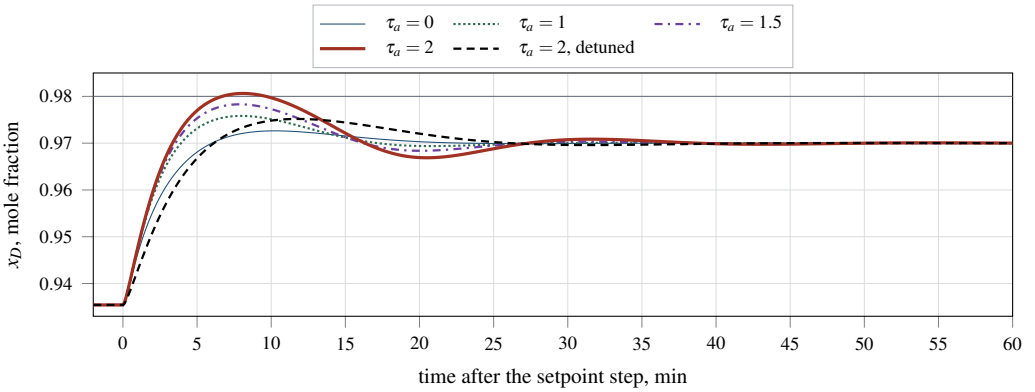

Figure 3.4 Capture of the 0.970 setpoint by the default tuning at four analyzer lags in minutes, after the setpoint is raised from 0.9354 at 2 min: thin solid at no lag, dotted at 1 min, dash-dotted at 1.5 min and thick solid at 2 min. The 0.980 specification limit is the thin gray line. The detuned run, dashed, lowers K_c to 36 and raises τ_I to 5 min at 2 min of lag, which brings the peak back inside the band.

3.6.2 Setpoint Capture With the Default Tuning

Run L0 holds the starting profile in automatic with the default tuning and raises the setpoint from 0.9354 to 0.970 at 2 min, so the loop has to move the distillate by 0.0346 against no disturbance at all. The execution at the step moves the reflux ratio from 3.0 to 5.07, it peaks at 5.10 at the execution 0.10 min later, the composition peaks at 0.9726 10.25 min after the step, and the loop settles inside 0.002 mole fraction of setpoint 13.53 min after the step. The reflux ratio then approaches the 3.7319 that holds 0.970 at steady state and reads 3.7314 60 min after the step. The overshoot is 7.6 % of the step and the composition spends 2.76 min outside the 0.960 to 0.980 specification band, all of it below the band on the way up.

3.6.3 Effect of the Analyzer Lag

Runs L0.5 to L2 repeat L0 with the analyzer lag raised in steps of 0.5 min to the 2 min maximum the application allows, and nothing else changed. The column is identical in all five runs. The only difference is the age of the measurement the controller acts on, and Table 3.4 quantifies the cost. Overshoot rises from 7.6 % to 30.7 %, the settling time rises from 13.5 min to 23.6 min, the integral absolute error rises by a factor of 1.79 and the reflux ratio travels 1.50 times as far. Half a minute of lag lowers the integral absolute error slightly (0.0928 against 0.0945), because the stale reading lets the proportional term push the reflux ratio higher early in the rise. The peak composition crosses the upper specification limit between 1.5 min and 2 min of lag, and at 2 min the loop that was designed to hold 0.970 peaks at 0.9806 and spends 2.95 min above the band, as shown in Figure 3.4. Figure 3.5 draws the same table as bars.

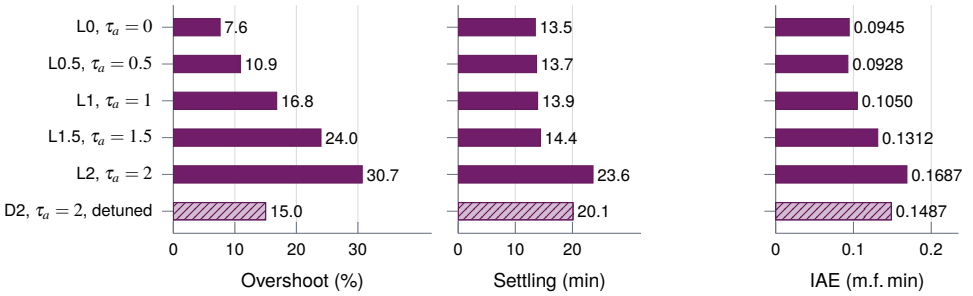

Figure 3.5 One tuning across five analyzer lags, from Table 3.4: the overshoot of the setpoint step (left), the settling time within 0.002 of setpoint (middle) and the integral absolute error (right), with the lag τ_a in minutes. Dark bars are the default tuning, the hatched bar the detuned run D2. Notice that the overshoot grows with every half minute of lag while the settling time barely moves until 2 min, and that detuning about halves the overshoot at 2 min without recovering the settling time of the fast loop.

The identified model explains the whole table. Adding a first-order lag of τ_a to the plant of Figure 3.3 and closing the loop with $K_c = 60$ and $\tau_I = 4$ min gives a phase margin of 73.2° at $\tau_a = 0$, falling to 59.7°, 48.7°, 40.5° and 34.3° at lags of 0.5, 1, 1.5 and 2 min. The gain margin falls from 16.5 to 8.32 over the same range, so the loop remains nominally stable throughout and the lag removes damping rather than stability. That is exactly the shape of Figure 3.4, where the 2 min case swings 0.0106 above setpoint 8.1 min after the step, 0.0031 below at 20.5 min and 0.0009 above at 31.6 min, decaying without diverging.

3.6.4 Detuned Controller at Two Minutes of Lag

Run D2 keeps the 2 min analyzer lag and detunes the loop to $K_c = 36$ and $\tau_I = 5$ min, chosen so that the predicted phase margin at 2 min of lag is 49.9°, close to the 48.7° the default tuning had at 1 min of lag. Every quantity in Table 3.4 improves. The peak falls from 0.9806 to 0.9752 and back inside the specification band, the overshoot falls from 30.7 % to 15.0 %, the settling time falls from 23.6 min to 20.1 min, the integral absolute error falls from 0.1687 to 0.1487 and the reflux travel falls from 5.32 to 2.43. Run D2 is outside the band for 3.42 min, all of it below on the slower rise, against 5.01 min for run L2, of which 2.95 min was above. Detuning further does not help, because $K_c = 24$ with the same integral time settles 0.22 min later than run L2 and raises the integral absolute error to 0.1743. The detuned loop is still 6.5 min slower to settle than the default tuning was with an instantaneous analyzer, and that gap is the part of the performance the lag removes permanently rather than trades against damping.

3.6.5 Effect of the Analyzer Noise on the Manipulated Variable

Run N0 repeats L0 with the default analyzer noise of $\sigma = 0.0015$ mole fraction kept. The trajectory is barely affected. The overshoot moves from 7.6 % to 7.8 % and the settling time from 13.53 min to 13.95 min, while five other noise seeds give 7.3 to 7.8 % and 13.5 to 13.8 min. The reflux ratio is affected far more strongly. Its travel over the 60 min after the step rises from 3.54 to 304.9, a factor of 86, and after the loop has settled the reflux ratio still varies with a standard deviation of 0.091 over a range of 0.62. The plant view at the end of a closed-loop run at 2 min of lag with this noise level is shown in Figure 3.6.

The velocity form makes the arithmetic transparent. The proportional term contributes $K_c(e_k - e_{k-1})$ to every increment, and for white measurement noise the difference of two independent samples has a standard deviation of $\sigma\sqrt{2}$, so the increment has a standard deviation of $K_c\sigma\sqrt{2} = 0.127$ reflux ratio units. The mean absolute value of a zero-mean normal variable is $\sqrt{2/\pi}$ times its standard

deviation, which gives 0.102 per execution and 305 over the 2999 executions that follow the step, against the 303.0 measured over the same executions. The exported history shows a travel of only 31.5 over the 60 min after the step and the performance panel 32.2 over the whole run, because between two samples 0.2 min apart the proportional increments telescope into a single difference of two readings. Raising K_c trades valve movement at this rate for speed, and the derivative time is left at zero on this loop because the second difference of the measurement would amplify the same noise by another factor of $1/\Delta t$.

3.6.6 Feed Composition Disturbance and the Purity Ceiling

Run XF applies the feed composition upset. The case starts in automatic at the steady state that holds 0.970 under the default tuning with no analyzer lag, the reflux ratio at 3.7319, and at 2 min the feed composition is stepped from 0.50 to 0.42. Equation (3.10) sets the new ceiling at $2x_F = 0.840$, well below the lower specification limit, so the setpoint is infeasible at steady state rather than merely difficult. The controller responds as feedback requires. The distillate leaves the band 15.8 min after the step, the reflux ratio reaches its 10.0 limit 28.7 min after the step and never leaves it, and the distillate reads 0.9377 at 30 min after the step, 0.8851 at 40 min and 0.8435 at 60 min. Left to run for 1296 min the column settles at $x_D = 0.8395$ and $x_B = 0.00045$, a sum of 0.8400 against the $2x_F = 0.840$ that Equation (3.10) requires.

Run XF demonstrates that no tuning changes the outcome. The reflux ratio is not saturated because the controller is badly tuned. It is saturated because the setpoint demands more light component in the distillate than the feed contains, and the only remedies are to raise the vapor rate independently of the reflux, to cut the distillate draw below $F/2$, or to lower the target.

3.6.7 Feed Rate Disturbance

Run FR steps the feed rate from 1.00 to 1.50 mol/min, the top of its range, at 2 min with everything else unchanged, and it has no closed-loop effect at all. Over the 60 min after the step the distillate stays at 0.9700 and the reflux ratio at 3.7319, and the run is identical to the same run without the step to the last digit the simulation carries. The reason is in Equation (3.1). Multiplying F by a factor multiplies D, L, V and L_S by the same factor at a fixed reflux ratio, so every right-hand side in Equations (3.3) to (3.7) is multiplied by that factor while the holdups stay where they are. The steady-state profile is therefore independent of the feed rate altogether, and what a feed rate change does is rescale time. Repeating the step test of Figure 3.3 in the application at $F = 1.5\,\text{mol/min}$ returns the same gain of 0.0535 measured at $F = 1$, with the time constant reduced from 7.07 min to 4.71 min and the apparent dead time from 0.208 min to 0.139 min, both within 0.2 % of the factor of 1.5 the scaling requires. The feed rate therefore changes the plant dynamics without displacing the steady state, a behavior easily mistaken for an ordinary well-rejected load change. The tray holdup separates the same two quantities from the other side. Raising it from 0.25 to 1.00 mol divides every tray derivative by four and leaves every steady state exactly where it was, because the holdups enter Equations (3.3) to (3.7) only as divisors. The same step test then returns the same gain with a much longer time constant, 27.1 min against 7.07 min.

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

3.6.8 The Shipped Cases Under the Default Tuning

The app carries ten cases in a fixed order, each with its own completion rule (Section 1.6.4). Every case starts at steady operation, so in cases 2 to 10 the scripted event at 2 min is the first thing that moves the column, apart from the loop's response to analyzer noise in automatic. The upset of cases 4, 5, 9 and 10 is a richer feed, x_F from 0.50 to 0.58, which raises the $2x_F$ ceiling of Equation (3.10) instead of lowering it as run XF does. In automatic the reflux ratio falls toward the 2.90 that holds 0.970 on the richer feed. In manual with the reflux ratio left at 3.73 the distillate climbs to 0.9888 by 30 min. A cut to 2.90 made at once in case 5 acts on the top before the richer feed does, so the distillate dips to 0.9588 at 9.5 min. Case 5 credits a reflux ratio that holds the distillate in the band

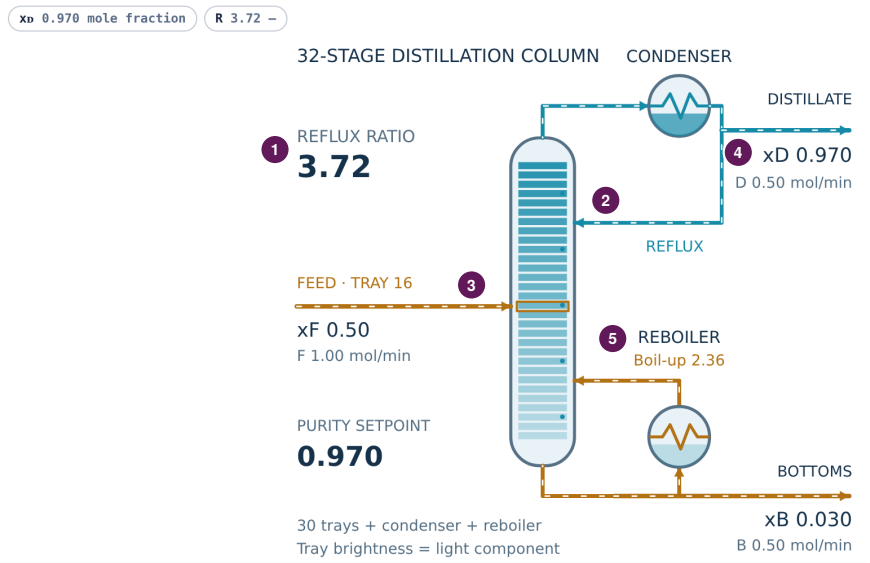

Figure 3.6 The Distillation Column Control Studio holding the 0.970 setpoint with a 2 min analyzer lag and analyzer noise of 0.0015 mole fraction, at 60 min of the reference operating point case, switched to automatic before the run and with the setpoint raised from 0.9354 to 0.970 at 2 min, cropped to the plant view (the header, control panel, trends and performance panel are cut). The default tuning of $K_c = 60$ and $\tau_I = 4$ min has brought the distillate back to 0.970 with the reflux ratio at 3.72. The callouts mark (1) the reflux ratio the controller sets, (2) the reflux returned to the top tray, (3) the feed on tray 16, (4) the distillate whose composition the analyzer reports and (5) the reboiler, whose boil-up follows the reflux because $V = L + D$.

once settled, kept there for 10 min, longer than the 7.07 min time constant, and the cut to 2.90 is credited at 23.2 min.

The ± 0.01 band equals the 0.010 setpoint step of case 3, so settling there is judged within ± 0.002 of the new setpoint, held for 1 min, and the case is credited at 8.6 min. The setpoint of 1.000 in case 8 lies above the 0.9971 the column reaches at the reflux limit, so the reflux ratio works up past 6, 8 and 10 at 4.5, 31.2 and 69.2 min while each unit buys less purity, and the distillate stops at 0.9971. The challenge starts at 0.970 with $K_c = 12$ and $\tau_I = 4$ min and asks for the analyzer reading to stay within ± 0.010 of 0.970 from 19.8 to 30 min through the richer feed at 2 min. With the shipped $K_c = 12$ the worst logged error over that window is 0.0143, so the attempt fails by 0.0043, and the true distillate is still 0.0112 above 0.970 at 20.9 min. The suggested tuning changes only the proportional gain, to $K_c = 60$, and gives 0.0052, 52 % of the tolerance, of which the true deviation is 0.0021 and the rest analyzer noise. Revealing the suggested tuning before the run and carrying it through a Reset give the same run.

> **TRY IT IN THE APP** **8 · Actuator Limit**
>
> Load the case and press Run at 10 min/s. Watch the reflux ratio on the lower strip work up past 6, 8 and 10 at 4.5, 31.2 and 69.2 min, while each unit buys less purity. Once the reflux sits on its limit, the purity stops at 0.9971, short of the 1.000 setpoint, and no tuning closes the gap.

3.7 Predictive Control Through a Lagging Analyzer

Table 3.5 runs cases 3, 4, 8 and 9 and the challenge once under each controller. The controller is chosen before the run starts, so the two runs of a case see identical events and identical noise draws.

Table 3.5 Cases 3, 4, 8, 9 and 10 Run Once Under Each Controller With the Same Analyzer Noise

Case	Controller	IAE (m.f. min)	IAE_x	e_{max} (m.f.)	$t_{0.002}$ (min)	Travel —	Limit (%)
3, setpoint	PID	0.048	0.038	0.0006	5.8	11.7	0
	MPC	0.031	0.016	0.0007	2.4	11.6	0
4, feed	PID	0.061	0.050	0.0045	—	11.4	0
	MPC	0.026	0.008	0.0008	0	6.5	0
8, limit	PID	0.519	0.534	—	—	37.6	36.8
	MPC	0.372	0.386	—	—	12.9	94.2
9, $\tau_a = 1$ min	PID	0.064	0.055	0.0052	17.4	11.4	0
	MPC	0.028	0.019	0.0022	5.8	7.2	0
10, challenge	PID, $K_c = 12$	0.225	0.216	0.0112	—	3.1	0
	PID, $K_c = 60^a$	0.076	0.057	0.0045	18.2	15.3	0
	MPC	0.040	0.010	0.0008	0	9.8	0

IAE uses the analyzer reading and IAE_x the true distillate. e_{max} is the largest true overshoot in case 3 and the largest true departure from setpoint after the feed step otherwise, $t_{0.002}$ the time from the step until the true distillate stays within ±0.002 (0 where it never leaves, a dash where it has not settled), Travel the summed reflux change on the log and Limit the time on a reflux limit. [a]Suggested, $\tau_I = 4$ min.

The PID keeps the shipped $K_c = 60$, $\tau_I = 4$ min and $\tau_D = 0$, the challenge loading $K_c = 12$ and suggesting 60, and the predictive controller keeps the settings of Section 3.5.4 throughout. Unlike Section 3.6, each case starts at 0.970 with the default noise and is read over its whole length on the 0.2 min log, as the performance panel reads it. The noise alone adds about $20\sigma\sqrt{2/\pi} = 0.024$ to a 20 min IAE on the reading, so the table also gives the IAE of the true distillate.

In case 3 the plan made at the setpoint step to 0.980 asks for reflux ratios of 5.74, 6.61 and 6.64 in its first three blocks, against a proportional kick to about 4.3 under the PID, and the true distillate stays within ±0.002 from 2.4 min after the step against 5.8 min, with overshoots of 0.0007 and 0.0006 and a true IAE 59 % lower. The equal logged travel, 11.6 against 11.7, hides that without noise the predictive controller moves the reflux ratio by 6.9 against 0.84, while with noise the PID moves it by 103 over its executions (Section 3.6.5).

In case 4 the estimate $\hat{\delta}$ passes half of its 20 min value, 0.63 or 0.64, at 5.0 min after the feed step under either controller, so the difference lies in the plan. The predictive controller keeps the true distillate within 0.0008 of 0.970, while the PID peaks at 0.9745 at 12.8 min and is still 0.0020 above at 20 min, with a true IAE 6.2 times higher. With the noise removed the true IAEs are 0.0053 and 0.0494 for reflux travels of 0.80 and 0.86, so the gain comes from acting earlier, not from moving more.

Case 8 asks for 1.000, above the 0.9971 that the column holds at the reflux limit (Section 3.4.5). The model places 1.000 at a reflux ratio of 4.76, where the column gives 0.9865, because the secant gain from 3.73 to 10 is 0.0043 and the local gain at 10 is 0.00041, 6.8 and 72 times below K_m. The first plan moves the reflux ratio to 9.17 and the second, 0.2 min later, to 10.0. Each of these plans predicts an overshoot, to 1.0013 in the first, and from 2.8 to 7.6 min the plans hold the reflux below the limit, down to 7.22, while $\hat{\delta}$ books the missing purity as a loss of reflux ratio that reaches −5.3 by 10 min, after which the reflux stays on the limit. The PID passes 6, 8 and 10 at 4.5, 31.2 and 69.2 min. Both end at 0.9971, the predictive controller with an IAE 28 % lower, 0.120 against 0.239 before 30 min and 0.250 against 0.278 after it, because the error in its gain, not any knowledge of the ceiling, drives it to the limit early.

Case 9 is case 4 seen through an analyzer lag of 1 min that Equation (3.14) omits, as shown in Figure 3.7. The predictive controller still scores less than half the PID's IAE, 0.028 against 0.064, with a true peak of 0.9722 at 7.2 min and the ±0.002 band held from 5.8 min after the step, against 0.9752 at 12.0 min and 17.4 min. Against case 4, however, the lag raises the true IAE of the PID by 9 %, from 0.050 to 0.055, and that of the predictive controller by 140 %, from 0.008 to 0.019, whose distillate then cycles about the setpoint with a period near 7 min, down to 0.9681. With the noise removed the cycle shrinks by a factor of only 0.66 to 0.83 per period, and with the noise present the true error between 90 and 120 min still reaches 0.0019 against 0.0002 for the PID.

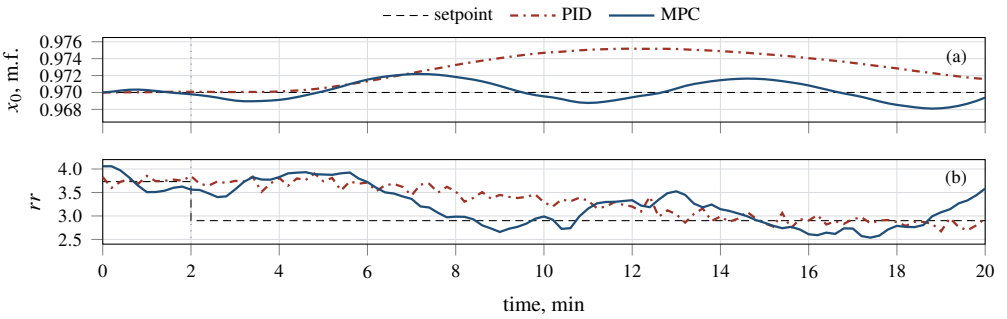

Figure 3.7 Distillate and reflux ratio of case 9 under each controller with the same analyzer noise, sampled on the 0.2 min log. The feed composition steps from 0.50 to 0.58 at 2 min (dotted) and the analyzer lags by 1 min. The PID is dash-dotted and the MPC solid. (a) True distillate composition and setpoint (dashed). (b) Reflux ratio and the steady value that holds 0.970 (dashed), 3.73 on the base feed and 2.90 on the richer one.

> **TRY IT IN THE APP** **9 · MPC With a Late Analyzer**
>
> Press Run and watch the distillate purity after the feed turns richer at 2 min. The MPC answers within minutes, then the purity keeps swinging about the setpoint with a period near seven minutes, because its model leaves out the one-minute analyzer lag. Switch the Controller to PID, press Reset and run again. The PID peaks later and higher, then settles.

Repeating case 9 at the lags of Table 3.4 measures how much of this model error the predictive loop tolerates. Over 20 min with the noise present its true IAE is 0.008, 0.010, 0.019, 0.034 and 0.042 at lags of 0, 0.5, 1, 1.5 and 2 min, against 0.050, 0.052, 0.055, 0.058 and 0.061 for the PID, lower at every lag but growing 5.2 times where that of the PID grows by 21 %. Over 180 min without noise the cycle shrinks per period by a factor of 0.82 at 1 min, 0.88 at 1.05 and 0.94 at 1.10 min, changes by less than 1 % per period at 1.15 min and grows at longer lags, from 1.5 min until the reflux ratio reaches its lower limit of 1.0. At 2 min the distillate cycles between 0.9525 and 0.9851 with a period near 10 min, while the swings of the PID decay by a factor of 0.10 to 0.14 per period, as its 34.3° phase margin implies. The predictive controller answers the upset faster and so has less phase to give up to a lag it does not model, and its largest move suppression, $\lambda = 3$, moves the stability limit only to between 1.25 and 1.5 min. Written into the model as a third first-order lag, outside the application, the analyzer lag restores the margin. With 2 min in its model the controller settles at 2 min of lag with a true IAE of 0.017 over 20 min, and with 1 min it lowers that of case 9 from 0.019 to 0.013, although 2 min in the model against a 1 min analyzer leaves a sustained swing of 0.0059.

Pause. Try Exercise 3.3 before reading on.

> **DID YOU KNOW?** **The Better Controller Loses More to the Lag**
>
> From the simulator: a one-minute analyzer lag added to the richer feed of case 4, which makes case 9, raises the true integral absolute error of the PID from 0.050 to 0.055 and that of the MPC from 0.008 to 0.019 (Table 3.5). The MPC still scores lower, but it gives up far more of its lead.

The predictive controller does not read K_c, so over the first 20 min the challenge reproduces case 4 exactly under the predictive controller and under the suggested tuning. On the judged worst reading error from 19.8 to 30 min it scores 0.0056 against 0.0052 for the suggested $K_c = 60$, both inside the 0.010 tolerance that $K_c = 12$ misses at 0.0143, although its largest true deviation there is 0.0008 against 0.0021, because the judged error is mostly analyzer noise. All 1050 plans were accepted,

and the reflux limit was binding in 566 of the 600 plans of case 8 and in no other, with at most 12 active-set iterations. Averaged over a run, a plan took 0.08 to 0.26 ms under Node.js in three repetitions and 229 and 244 µs in a pair of Chrome runs of case 9, the longest 2.1 ms, which on one desktop computer is 0.12 % of one processor core at one simulated minute per second.

3.8 Estimator Option and Relay Test

3.8.1 The Moving Horizon Estimate

Every case was run twice with the controller set to the MPC, once with each estimator chosen before the run, with the shipped noise seed, for 60 min or, for the challenge, its judged 30 min, and cases 1, 2 and 5 stayed in manual as given. Across the ten cases the two reflux records differ by at most 2.6×10^{-8} and the two readings by at most 7.7×10^{-11} mole fraction, no window was rejected and no plan failed. No bound was ever active, because the filter's purity estimate never left the interval from 0 to 1. Case 8 brought it closest, at 0.9977 under the setpoint of 1.000. Over the 2999 executions of case 9 the two purity estimates differ by at most 1.0×10^{-8} of the noise standard deviation and the two values of $\hat{\delta}$ by 4.8×10^{-9} units of reflux ratio. That residue does not come from rounding. The Riccati iteration that supplies the filter's gain stops after 526 steps with its covariance 2.6×10^{-8} (relative) short of its fixed point. The window performs those Riccati steps afresh at every execution, which separates the two estimates slightly, and with the converged covariance the difference drops to 7.3×10^{-14} of the noise.

Judged against the true distillate, which only the measuring script reads, the two estimators are equally accurate, and no case shows a bound that the filter violates and the MHE holds. Over the 20 min of case 9 the purity estimate has a root-mean-square error of 8.6×10^{-4} and a largest error of 1.7×10^{-3}, while against the analyzer state x_a of Equation (3.9) its error is 3.9×10^{-4}, so the estimate tracks the lagged composition that Equation (3.14) mistakes for the distillate. In case 4, without the lag, the two errors coincide at 3.0×10^{-4}, with a largest error of 1.6×10^{-3}. On a loaded computer under Node.js an estimate of case 9 took a median of 76 µs and a 99th percentile of 244 µs, against 85 and 515 µs for a plan, and each 0.2 min decision period holds ten estimates and one plan. In three Chrome runs on the application's frame loop the estimates and plans of case 9 occupied 1.2 to 1.5 % of a processor core at the 1 min/s speed setting and 5.3 to 8.5 % at 10 min/s, inside the tenth of a core allowed, though within a factor of two of it at the top speed.

3.8.2 The Relay Test and the Constants It Writes

The relay test was pressed 2 min into case 3, its setpoint step removed, with the purity steady at 0.970 in automatic and the default noise. Its center was 3.745, the mean of the last 100 executions, against 3.850 applied at that instant and the 3.732 that holds 0.970, so d_r took its full 0.45 and h_r was 0.0045. The test converged after three cycles and 12.6 min with $a_r = 0.00503$, $f_h = 0.492$, $P_u = 5.36$ min and $K_{cu} = 256$, and it wrote $K_c = 79.9$ and $\tau_I = 11.8$ min, 1.33 times the shipped gain and 2.95 times the shipped integral time. The swing exceeded the band by a factor of only 1.12, so $\psi = 63.5^\circ$ and the note marked the cycle as set by the band. Of the seeds 1001 to 1010, seven converged, with K_{cu} from 116 to 288, P_u from 3.26 to 6.46 min and ψ from 42.5° to 66.2°, and three reached the 40 min limit after 11 to 14 cycles without two agreeing.

The cycle reveals little about the column's ultimate point. The two lags of Equation (3.14) approach -180° without reaching it, and only the zero-order hold over Δt creates an ultimate point for the sampled model, with a period of 0.456 min and a gain of 26,400. The measured period is twelve times longer and the measured gain a hundredth of it. The same relay returns 514 at 11.9 min on the noise-free model, and 767 at 12.3 min on the simulator with the analyzer noise set to zero before the run, where $f_h = 0.518$ and $\psi = 80.6^\circ$. On this loop the swing barely clears the band, and Equation (3.18) then amplifies any change in it, since $\mathrm{d}K_{cu}/K_{cu} = -\left[a_r^2/(a_r^2 - h_r^2)\right] \mathrm{d}a_r/a_r$. The factor is 5.0 at the measured $a_r/h_r = 1.12$ and 37 at the noise-free 1.014. Noise of $\sigma/a_r = 0.30$ widens the swing from

Table 3.6 Shipped, Relay-Written and Two Counterfactual PI Pairs on Cases 3 and 4 Over 60 min With the Shipped Noise Seed

Case	PI pair	IAE (m.f. min)	IAE_x	$e_{\max}$ (m.f.)	$t_{0.002}$ (min)	Travel —
3, setpoint	shipped	0.103	0.042	0.0006	5.8	30.0
	relay	0.109	0.066	0.0001	7.8	39.9
	PI[a]	0.120	0.087	0.0001	11.5	30.0
	PI[b]	0.098	0.036	0.0006	4.6	39.9
4, feed	shipped	0.117	0.059	0.0045	18.2	29.7
	relay	0.154	0.120	0.0051	29.8	39.7
	PI[a]	0.187	0.160	0.0063	35.6	29.8
	PI[b]	0.106	0.044	0.0037	15.6	39.6

Shipped: $K_c = 60$, $\tau_I = 4$ min. Relay: $K_c = 79.9$, $\tau_I = 11.8$ min. [a] $K_c = 60$, $\tau_I = 11.8$ min. [b] $K_c = 79.9$, $\tau_I = 4$ min. Neither counterfactual is shipped or written. Columns as in Table 3.5, over 60 min instead of the case's 20 min.

0.00456 to 0.00503 and so cuts K_{cu} from 767 to 256. The uncorrected $4d_r/(\pi a_r)$ reads 114 and 125, so the band term, not the column, supplies most of either result.

Table 3.6 runs cases 3 and 4 for 60 min under four PI pairs, the shipped and the written ones and two counterfactual pairs that each exchange one constant. The written pair loses both cases. On the setpoint step its scorecard IAE is 0.109 against 0.103 and its true IAE 0.066 against 0.042. On the richer feed they are 0.154 against 0.117 and 0.120 against 0.059, and the true purity returns within ± 0.002 of 0.970 at 29.8 min after the feed step instead of 18.2 min. In both cases the reflux travels 33 % further, the ratio of the two proportional gains, because the analyzer noise sets the travel (Section 3.6.5). The counterfactual pairs locate the loss in the integral time. The shipped gain with the written τ_I does worse still, a true IAE of 0.160 on the richer feed, while the written gain with the shipped τ_I wins both cases, at 0.036 and 0.044. Taken as $2.2P_u$ from a cycle twelve times slower than the ultimate one, the written τ_I lowers the integral gain K_c/τ_I from 15.0 to 6.77 per mole fraction per minute. Across the 14 of 21 tests that converged (the shipped seed and the seeds 1001 to 1020) the written pairs scored 0.101 to 0.148 on the setpoint step, below the shipped 0.103 in 3 of them, and 0.124 to 0.272 on the richer feed, above the shipped 0.117 in all 14. With the noise removed from both case runs the written pair still loses the feed upset, with a true IAE of 0.120 against 0.056.

3.9 Technology in Practice

What Measures It

Chromatographs very commonly monitor and control the composition of distillation products, since one instrument quantifies several species, and an analyzer with its sample system needs much more maintenance than other instruments, with harder fault diagnosis [24]. In a survey reported in a 2013 review, the problems named most often with hardware sensors were maintenance and calibration, and one answer in ten named long dead time [25]. The chapter lumps the sample line and the analysis into the lag τ_a of Equation (3.9).

TECHNOLOGY Sensor: The Process Gas Chromatograph

A process chromatograph injects a precise small volume of sample into a long, narrow column, and a steady flow of carrier gas washes it through. Each species travels at its own rate, so the time it leaves tells it apart, and a detector that only tells carrier gas from anything else is enough. The area under each peak gives the quantity, and the timing is calibrated with a sample of known composition. The separation takes time, so the result arrives late, and ramping the column temperature or switching columns shortens the wait. A leak in the sample line lets outside gas dilute the sample, which a calibration gas sent through the whole line reveals [24].

What Moves It

The manipulated variable is the reflux ratio L/D, with the distillate draw held at half the feed as though by a perfect ratio controller. In a plant control system loops are often nested or cascaded, the setpoint of one taken from the output of another [26], and Figure 3.8, which draws the loop of Section 3.5 on field devices, sets the reflux valve under a flow controller whose setpoint the composition loop writes. The model applies each reflux ratio at once, with no flow loop and no valve lag.

TECHNOLOGY Final Element: The Reflux Control Valve

The reflux runs from the drum back to the top tray through a control valve. Tested at a constant pressure drop, a valve with linear trim passes a flow in proportion to its stem position. In a plant the pressure drop across the valve usually falls as the flow rises, so the installed flow droops at large openings, and equal-percentage trim brings the installed curve closer to a straight line. No trim makes up for pressure lost in restrictive piping, and a large valve can add a lag of tens of seconds to a loop [24]. The application applies each reflux ratio exactly, so all the curvature of Section 3.4.5 comes from the column.

Where the Controller Runs

A column's loops run in a distributed control system, which places a supervisory level over field controllers that hold their process conditions around set points, with cycle times from milliseconds to minutes [26]. A survey of predictive control sketches the usual hierarchy: basic PID loops in that system execute about every second, and a predictive controller above them about every minute [18]. The application sits close to that picture. Its PID executes every 0.02 min, or 1.2 s, its MPC every 0.2 min, and both are short beside an analyzer lag of up to 2 min and the 7.07 min time constant of the column. The velocity-form PID of Equation (3.12) is the function block of the standard loop, and the MPC with its Kalman filter or its moving horizon estimator replaces it when selected.

What Is Modernizing It

Three developments bear on this loop. The first is the analyzer's link to the control system. Older chromatographs give one analog output per species, while most modern ones can send several concentrations over a digital bus and compute derived values from them [24]. The second is the soft sensor, the estimator the Discussion proposes. In one company's survey it served to steady product quality, to save energy by operating close to the specification and to check the online analyzers, and distillation is its largest target in refining and chemicals [25]. The third is tuning that the controller does itself. Adaptive and on-demand tuning of PID controllers has been very successful in the marketplace, while adaptive predictive controllers remain rare, and predictive control has been moving into the distributed control systems themselves [18]. The relay test of Section 3.5.6 is such an on-demand tuner.

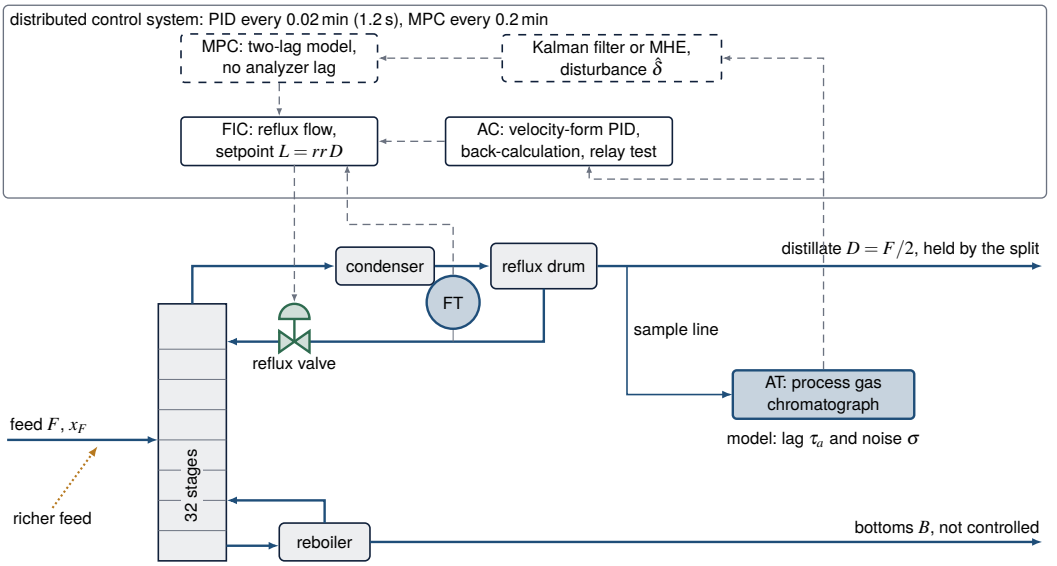

Figure 3.8 The composition loop on plant hardware. A process gas chromatograph (AT) reads a sample drawn from the distillate line, and its result feeds the composition controller (AC) and the estimator of the dashed predictive option. Either controller writes the reflux setpoint $L = rr\,D$ to a reflux flow loop, the plant's way of applying a reflux ratio that the model applies exactly. Notice that in the model the reflux acts on the top tray at once, while the measurement reaches the controller only after the sample line and the analysis, the lag that the MPC's model leaves out.

> **IN PRACTICE When the Analyzer Has Nothing New**
>
> Industrial predictive controllers meet an analyzer whose result is not available at every control execution. A typical solution skips the bias update of that controlled variable for a number of control intervals, and a counter switches control of the variable off if too many executions pass without feedback. A variable whose measurement status is good, and which the operator has enabled, stays under control [18]. The filter of Equation (3.15) corrects its estimate at every execution from a continuous reading, so a chromatograph that reports once per cycle would need such logic first.

3.10 Limitations

The analyzer is a lag and not a sampler. A real chromatograph delivers a new number once per cycle and the number describes the sample taken at the start of that cycle, which is closer to a dead time than to a first-order lag, with the mean age of Equation (1.8). A dead time of 2 min would leave the default tuning a predicted phase margin of 17.1°, half the 34.3° the 2 min lag of run L2 leaves, so Table 3.4 is an optimistic account of what a slow analyzer does.

The column is single ended by construction. Fixing $D = F/2$ removes the second degree of freedom, so dual composition control, the pairing question and the interaction that the relative gain array measures cannot be demonstrated. Volume 6, Chapter 3 measures that interaction on a two-by-two plant read by a scanning gauge. The bottoms composition is reported and logged, and Equation (3.10) means it carries no information the distillate composition does not already carry.

There is no energy balance and there are no temperatures, so inferential control, which is what most industrial columns actually run on, has nothing to be built from. There is no hydraulic lag, no flooding limit, no weeping limit and no pressure, so the reflux ratio can be moved between 1.0 and 10.0 without any of the physical consequences that would stop an operator doing the same thing.

The plant is stiff and the integrator is explicit, so the integration step is fixed at 0.002 min and

cannot be raised. Reducing the tray holdup to its lower limit of 0.10 mol shortens the fastest time constant from 0.0253 min to 0.0101 min and the stability limit from 0.0506 min to 0.0203 min, which is the reason the holdup ranges do not extend further down.

The predictive controller's model is two linear lags fitted at 0.970 on the base feed. It omits the analyzer lag, which brings the loop to the edge of stability at 1.15 min, the flattening of the gain toward high purity, absorbed as -5.3 units of reflux ratio in case 8, and the feed composition, seen only through the analyzer. There is no specification band or reflux rate limit, and the filter takes the reading for the distillate itself, so it is not the soft sensor of the Discussion. The MHE of Section 3.5.5 bounds the states of that same two-lag model, not the column's, so it shares the model's blindness to the analyzer, and with no bound reached on any case it adds nothing to the filter here. The relay test meets a loop with no dead time, where the swing exceeds the band by only 1.01 to 1.48 times, so its K_{cu} depends on the band and on the noise more than on the column, and the Tyreus–Luyben rule is one conservative choice among many.

> **OPEN CHALLENGE** **How Accurate Must the Model Be?**
>
> A survey of industrial predictive control found model uncertainty not addressed adequately. Every product could be detuned for robustness, but the trade-off between performance and robustness was not clear, so it could not be decided when a model was accurate enough for an application. Nearly all products judged model mismatch by brute-force simulation, and robust stability guarantees would greatly cut the time spent tuning and testing [18]. This chapter meets the question in small. The MPC's model leaves out the analyzer lag, its loop stops settling at 1.15 min of lag, and the largest move suppression moves that limit only a little. A test that predicted the limit before the run would be progress.

Finally, the noise is white and additive on the reading. Analyzer drift, a calibration offset and a failed sample line are all more common in practice than Gaussian noise, and none of them is represented.

> **OPEN CHALLENGE** **Trusting an Estimate Between Analyses**
>
> An estimate of composition is only as good as the measurements it is built from. It fails when an input sensor fails, and a statistical estimator should interpolate, never extrapolate. A blocked sample line makes the analyzer itself wrong. The simplest check compares the estimate with the analyzer when a result arrives, so its validity can be judged only then, and a 2013 review calls for another way to validate online. The hard part is to tell normal changes in the process from faults [25]. The analyzer drift, calibration offset and failed sample line that the Limitations name are such faults.

3.11 Discussion

The economics of a column are dominated by the reboiler. Separation work has to be supplied as heat at the bottom and rejected at the top, distillation is responsible for a large share of the process energy used by the chemical and refining industries, and the efficiency of the separation is set as much by how the column is operated as by how it was designed [15]. Composition control decides the operating part. A column held close to its specification limit meets the specification at the lowest energy use, while a column held well inside the limit gives away purity that the customer does not pay for.

How close a loop can be run to a limit is set by its variability, and Table 3.4 quantifies the analyzer directly in those terms. At $\tau_a = 0$ the peak excursion of the default tuning is 0.0026 mole fraction past setpoint. At $\tau_a = 2$ min it is 0.0106, so a setpoint that was safely inside the band at the fast analyzer has to be moved 0.0080 mole fraction away from the limit to give the same protection at the slow one. The detuned run D2 recovers most of that with an excursion of 0.0052, at the cost of settling 6.5 min later than the fast loop, which is time spent off target after every upset. This is the calculation that justifies a temperature inference trimmed by a slow analyzer rather than direct control on the

analyzer itself [5, 6], and it is the same calculation whether the loop is a single tray composition or a full multivariable scheme on both products [13].

The model of this chapter also supports control and estimation methods beyond the implemented PID loop and the linear predictive controller of Section 3.5.4. The thirty-two balances of Equations (3.3) to (3.7) form a complete nonlinear state-space model, so Model Predictive Control can enforce the reflux ratio limits and the specification band as explicit constraints rather than through detuning, and an added independent vapor rate would let one controller coordinate both product compositions. The lagged, noisy analyzer of Equation (3.9) is the standard setting for state estimation, where a moving horizon estimator or a Kalman filter reconstructs the current distillate composition from the delayed reading, a soft sensor (Section 1.4.8) that recovers part of the phase the analyzer removes. APMonitor implements nonlinear MPC and moving horizon estimation for models of exactly this differential algebraic form [27], and the Gekko Optimization Suite provides the same capabilities in Python [28]. The equations and parameters of Section 3.3 are sufficient to reproduce the column in either tool.

3.12 Conclusion

This chapter studied a thirty-two stage binary distillation column as an interactive laboratory for composition control through a slow measurement. The simulator integrates one component balance on each of a total condenser, fifteen rectifying trays, a saturated liquid feed tray, fourteen stripping trays and a partial reboiler, with a constant relative volatility of 1.6, tray holdups of 0.25 mol, explicit Euler on a 0.002 min step and a controller executing every 0.02 min. Fixing the distillate flow at $D = F/2$ makes the reflux ratio the only degree of freedom and forces $x_D + x_B = 2x_F$ at steady state, so a reflux move rearranges the profile without changing the 4.244 mol of light component the column holds, and the 8.17 min profile mode carries 109 % of the steady-state gain while the 15.24 min inventory mode carries −2.8 %. That is why a plant with thirty-two states and no stage slower than 0.40 min responds to a reflux step as a single 7.07 min lag with 0.21 min of apparent dead time and a gain of 0.0535 mole fraction per unit of reflux ratio.

Carrying one tuning across analyzer lags from 0 to 2 min moves the overshoot from 7.6 % to 30.7 % and the settling time from 13.5 min to 23.6 min without touching the column, and the phase margin of the identified model falls from 73.2° to 34.3° over the same range. Detuning to $K_c = 36$ and $\tau_I = 5$ min brings the peak back inside the 0.960 to 0.980 band at 2 min of lag and still settles 6.5 min later than the fast loop. A feed of $x_F = 0.42$ puts the 0.970 target above the $2x_F$ ceiling, and the reflux ratio sits on its 10.0 limit for the rest of the run because no tuning can supply light component the feed does not contain.

The linear predictive controller added beside the PID rejects the richer feed with an integral absolute error of 0.026 against 0.061 and reaches the reflux limit at once when 1.000 is requested, but its two-lag model omits the analyzer, and at 1.15 min of lag its loop stops settling. Selecting the moving horizon estimator for that controller instead of its Kalman filter reproduces the filter's run on every case to 2.6×10^{-8} in the reflux ratio, since no purity bound is reached, and a relay test on the PID writes an integral time from a cycle twelve times slower than the ultimate one, which raises the hour-long integral absolute error of the feed upset from 0.117 to 0.154.

A fuller model would replace the analyzer lag with a sampled dead time, add an independent vapor rate so that both product compositions can be controlled, and add tray temperatures so that an inferential scheme can be built and compared. As it stands, the column is a high-order plant whose limiting element is the measurement rather than the plant itself.

KEY POINTS

- A column of thirty-two stages answers a reflux move like one slow lag, which is why a two-parameter tuning rule works on a thirty-two state plant.
- The fitted dead time of 0.21 min is not a transport delay. It is the thirty fast eigenvalues rounding the foot of the response.
- Carrying one tuning across analyzer lags from 0 to 2 min moves the overshoot from 7.6 to 30.7 % and the settling time from 13.5 to 23.6 min without touching the column.
- The predictive controller rejects the richer feed better than the PID, but its model omits the analyzer, and at 1.15 min of lag its loop stops settling.

Where else this applies. Volume 12, Chapter 4 holds a far taller column through a tray temperature under a slow chromatograph, and Chapter 4 cooks hours ahead of its kappa analyzer.

References

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

[2] S. Skogestad and M. Morari, "Understanding the dynamic behavior of distillation columns," *Industrial & Engineering Chemistry Research*, vol. 27, no. 10, pp. 1848–1862, 1988.

[3] E. F. Wahl and P. Harriott, "Understanding and prediction of the dynamic behavior of distillation columns," *Industrial & Engineering Chemistry Process Design and Development*, vol. 9, no. 3, pp. 396–407, 1970.

[4] A. Ingimundarson and T. Hägglund, "Performance comparison between PID and dead-time compensating controllers," *Journal of Process Control*, vol. 12, no. 8, pp. 887–895, 2002.

[5] T. Mejdell and S. Skogestad, "Estimation of distillation compositions from multiple temperature measurements using partial-least-squares regression," *Industrial & Engineering Chemistry Research*, vol. 30, no. 12, pp. 2543–2555, 1991.

[6] W. L. Luyben, "Evaluation of criteria for selecting temperature control trays in distillation columns," *Journal of Process Control*, vol. 16, no. 2, pp. 115–134, 2006.

[7] 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.

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

[9] W. L. McCabe and E. W. Thiele, "Graphical design of fractionating columns," *Industrial & Engineering Chemistry*, vol. 17, no. 6, pp. 605–611, 1925.

[10] S. Skogestad, "Dynamics and control of distillation columns," *Chemical Engineering Research and Design*, vol. 75, no. 6, pp. 539–562, 1997.

[11] R. K. Wood and M. W. Berry, "Terminal composition control of a binary distillation column," *Chemical Engineering Science*, vol. 28, no. 9, pp. 1707–1717, 1973.

[12] S. Skogestad and M. Morari, "LV-control of a high-purity distillation column," *Chemical Engineering Science*, vol. 43, no. 1, pp. 33–48, 1988.

[13] S. Skogestad, "The dos and don'ts of distillation column control," *Chemical Engineering Research and Design*, vol. 85, no. 1, pp. 13–23, 2007.

[14] 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.

[15] A. A. Kiss and R. Smith, "Rethinking energy use in distillation processes for a more sustainable chemical industry," *Energy*, vol. 203, p. 117788, 2020.

[16] BCS, Incorporated and Oak Ridge National Laboratory, "Materials for separation technologies: Energy and emission reduction opportunities," U.S. Department of Energy, Industrial Technologies Program, Tech. Rep., 2005.

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

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

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

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

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

[22] 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.

[23] K. J. Åström and T. Hägglund, "Automatic tuning of simple regulators with specifications on phase and amplitude margins," *Automatica*, vol. 20, no. 5, pp. 645–651, 1984.

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

[25] M. Kano and K. Fujiwara, "Virtual sensing technology in process industries: Trends and challenges revealed by recent industrial applications," *Journal of Chemical Engineering of Japan*, vol. 46, no. 1, pp. 1–17, 2013.

[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] J. D. Hedengren, R. Asgharzadeh Shishavan, K. M. Powell, and T. F. Edgar, "Nonlinear modeling, estimation and predictive control in APMonitor," *Computers & Chemical Engineering*, vol. 70, pp. 133–148, 2014.

[28] 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.

3.13 Exercises

3.1 ○ Concept check · 2 min

Run FR steps the feed rate from 1.00 to 1.50 mol/min at 2 min with the loop in automatic. Over the next hour the distillate stays at 0.9700 and the reflux ratio at 3.7319. Why does the step move nothing?

(a) At a fixed reflux ratio every flow scales with the feed, so the steady profile does not depend on the feed rate and only the time scale changes

(b) The PID rejects the step within one execution, faster than the analyzer can register the change in the distillate

(c) The feed rate enters only the feed-tray balance, and the reflux ratio overrides that tray before the change reaches the top

(d) The distillate draw takes up all of the extra feed, so the liquid and vapor rates inside the column do not change

3.2 ○ Concept check · 1 min

The reflux step from 3.0 to 3.5 is fitted with a first-order plus dead-time model, which gives $\theta_p = 0.21$ min. The analyzer lag and noise are zero during the test. Where does this dead time come from?

(a) The analyzer, which reports the distillate through a first-order lag

(b) The thirty fast modes of the stage balances, which round the foot of the response

(c) The time the reflux liquid takes to run down the trays to the reboiler

(d) The controller period and the integration step, which hold each reflux move

3.3 ○ Concept check · 2 min

With the analyzer lag at 1.15 min the MPC loop of case 9 stops settling, while the PID with the shipped tuning still settles at a lag of 2 min. Which explanation fits the chapter's measurements?

(a) The MPC has no integral action, so the lag leaves an offset that it cannot remove
(b) The analyzer noise excites the plans, and without the noise the MPC settles at every lag
(c) The MPC's model omits the analyzer lag, and its faster answer leaves it less phase to lose
(d) The reflux ratio reaches its lower limit early, and the clipped plans wind up the estimate

3.4 Theory and calculation · 6 min

The column runs at $rr = 3$ with $F = 1\,\text{mol/min}$ and the holdups of Table 3.2.

(a) Compute D, B, L, V and L_S from Equation (3.1). *Checkpoint:* $D = B = 0.5$, $L = 1.5$, $V = 2.0$ and $L_S = 2.5\,\text{mol/min}$.

(b) Compute the liquid residence time of a rectifying tray, a stripping tray, the condenser and the reboiler, and the ratio of the apparent time constant of 7.07 min to the longest of them. *Checkpoint:* 0.167, 0.100, 0.250 and 0.400 min. The ratio is 17.7, about eighteen.

(c) The feed rate is raised to 1.5 mol/min. Predict the gain, the time constant and the apparent dead time of the same reflux step, from 0.0535, 7.07 min and 0.208 min at $F = 1$. *Checkpoint:* Gain 0.0535, $\tau_p = 4.71\,\text{min}$, $\theta_p = 0.139\,\text{min}$.

3.5 Theory and calculation · 6 min

The step test gives $K_p = 0.0535$ per unit of reflux ratio, $\tau_p = 7.07\,\text{min}$ and $\theta_p = 0.21\,\text{min}$. Use the IMC rule of Equation (3.11).

(a) Compute K_c and τ_I for $\tau_c = 2\,\text{min}$. *Checkpoint:* $K_c = 59.8$ per mole fraction, $\tau_I = 7.07\,\text{min}$.

(b) Compute K_c and τ_I for $\tau_c = 1\,\text{min}$. *Checkpoint:* $K_c = 109$, $\tau_I = 4.84\,\text{min}$.

(c) The shipped tuning is $K_c = 60$ and $\tau_I = 4\,\text{min}$. Compare its integral gain K_c/τ_I with that of part (a), and say which of the two removes a sustained load error faster. *Checkpoint:* 15.0 against 8.46 per mole fraction per minute. The shipped tuning, by a factor of 1.77.

3.6 Theory and calculation · 6 min

Run N0 repeats the setpoint step with the analyzer noise at $\sigma = 0.0015$ mole fraction, $K_c = 60$ and $\tau_D = 0$, and the reflux ratio travels far more than without noise. Treat the reading noise as white.

(a) Find the standard deviation of the proportional increment $K_c(e_k - e_{k-1})$ that the noise alone produces. *Checkpoint:* $K_c\sigma\sqrt{2} = 0.127$ units of reflux ratio.

(b) Find the expected travel $\sum |rr_k - rr_{k-1}|$ over the 2999 executions that follow the step, ignoring the small integral increments. *Checkpoint:* $0.127\sqrt{2/\pi} = 0.102$ per execution, 305 in total.

(c) The exported history samples the reflux ratio every 0.2 min, every tenth execution. Estimate the travel it shows over the same hour, and explain why it is so much smaller. *Checkpoint:* About 30.5 over 300 samples. The ten increments between two samples telescope into one difference of two readings.

3.7 App challenge · 9 min

Load scenario 9, MPC with a late analyzer, with the default seed. The MPC is in charge and the analyzer lags by 1 min. Run to 20 min at a speed of 1 and read the IAE and MV travel tiles. Set Controller to PID, press Reset and repeat.

(a) Report the two tiles for each controller. *Checkpoint:* MPC: IAE 0.028, MV travel 7.22. PID: IAE 0.064, MV travel 11.4. Within 0.002 and 0.1.

(b) Reload scenario 9, set Analyzer lag to 1.5 min before Run, and run to 60 min at a speed of 5. Read the IAE and Saturated tiles. Set Controller to PID, press Reset and repeat. *Checkpoint:* MPC: IAE 0.218, Saturated 12.7 %. PID: IAE 0.122, Saturated 0 %. Within 0.005 and 0.5 %.

(c) Explain why the order of the two controllers reverses between parts (a) and (b). *Checkpoint:* The MPC's model omits the lag. At 1.5 min its swings grow until the reflux reaches its lower limit, while the PID's decay.

3.8 App challenge · 7 min

Load scenario 8, Actuator limit, with the default seed. The setpoint steps from 0.970 to 1.000 at 2 min. Run to 120 min at a speed of 10.

(a) Export the CSV file and find the first row in which `saturated` is 1. Read the Saturated and IAE tiles at 120 min. Why does the LIMIT badge flicker for several minutes before it stays lit? *Checkpoint:* 73.4 min, within one row. Saturated 36.8 % and IAE 0.519, within 0.5 % and 0.005.

(b) The feed composition is 0.50. What ceiling does Equation (3.10) put on the distillate, and is it the limit that stops this run? *Checkpoint:* $2x_F = 1.000$. No. The reflux limit stops the distillate at 0.9971.

(c) Run XF also ends with the reflux on its limit. Name the remedy that each run calls for. *Checkpoint:* Case 8: lower the setpoint, since even a higher reflux limit buys almost nothing. Run XF: cut the draw below $F/2$, raise the vapor independently, or lower the target.

3.9 App challenge · 8 min

Load the challenge, Challenge · return to target, the tenth case in the list, with the default seed. It starts at 0.970 with $K_c = 12$ and $\tau_I = 4$ min, and the feed composition steps to 0.58 at 2 min. Run to 30 min at a speed of 5 each time. The challenge result is the worst $|\texttt{pv} - \texttt{sp}|$ logged from 19.8 min on, which the case's progress list reports and the exported CSV also gives.

(a) Run as loaded and read the challenge result. *Checkpoint:* Worst error 0.0143 against ± 0.01: not met.

(b) Press Reveal suggested tuning, press Reset and run again. Then set Controller to MPC, press Reset and run a third time. Read the challenge result of each. *Checkpoint:* Suggested PID ($K_c = 60$): 0.00518, met. MPC: 0.00555, met. Within 0.0003.

(c) Export the CSV of the last two runs and find the largest $|\texttt{true_pv} - \texttt{sp}|$ from 19.8 min on. Which controller holds the true distillate closer, and why does the judged error not show it? *Checkpoint:* MPC 0.0008, PID 0.0021, within 0.0002. The judged error is read on the noisy analyzer, and the noise dominates it.

3.10 Application area · 4 min

The analyzer lag of Table 3.4 decides how far the setpoint must sit from the specification limit. Under constant molar overflow the reboiler duty is proportional to the vapor rate $V = (rr + 1)D$.

(a) From the peaks of runs L0, L2 and D2, find how far the setpoint must move away from the limit when the lag rises from 0 to 2 min, and what share of that the detuned run D2 recovers. *Checkpoint:* 0.0106 against 0.0026: 0.0080. D2 peaks 0.0052 past setpoint, recovering 0.0054, about 68 %.

(b) The steady map gives 0.9765 at $rr = 4.0$, against $rr = 3.7319$ for 0.970. By what percentage must the vapor rate rise to hold 0.9765 instead of 0.970 at the base feed? *Checkpoint:* V from 2.37 to 2.50 mol/min, 5.7 % more.

3.11 Application area · 4 min

A plant chromatograph on the distillate runs a 2 min cycle. Each result describes the sample drawn at the start of a cycle, is published at its end and is held through the next cycle.

(a) Use Equation (1.8) to find the mean age of the result the controller sees. *Checkpoint:* $t_d = T_s = 2$ min, so $t_d + T_s/2 = 3$ min.

(b) The chapter's Limitations section finds that a pure dead time of 2 min leaves the default tuning a phase margin of 17.1°, against 34.3° for a lag of 2 min. What do you expect for this chromatograph, and what would you do before closing the loop on it? *Checkpoint:* Worse than 17.1°: on the fitted model the loop is unstable, gain margin about 0.85. Detune, or control a tray temperature and trim it with the analyzer.

3.12 Review the agent · 5 min

An AI assistant is asked to improve the composition loop. It replies: "Run the relay autotune in case 3 and keep the constants it writes, $K_c = 79.9$ and $\tau_I = 11.8$ min. A relay test measures the column's ultimate point directly, so Tyreus–Luyben settings from it are better founded than the hand-picked 60 and 4 min, and the longer integral time can only make the loop more robust." Find the flaw, and say how you would check the advice in the app.

CHAPTER 4

Kraft Digester Simulator That Cooks Four Hours Ahead of Its Kappa Analyzer

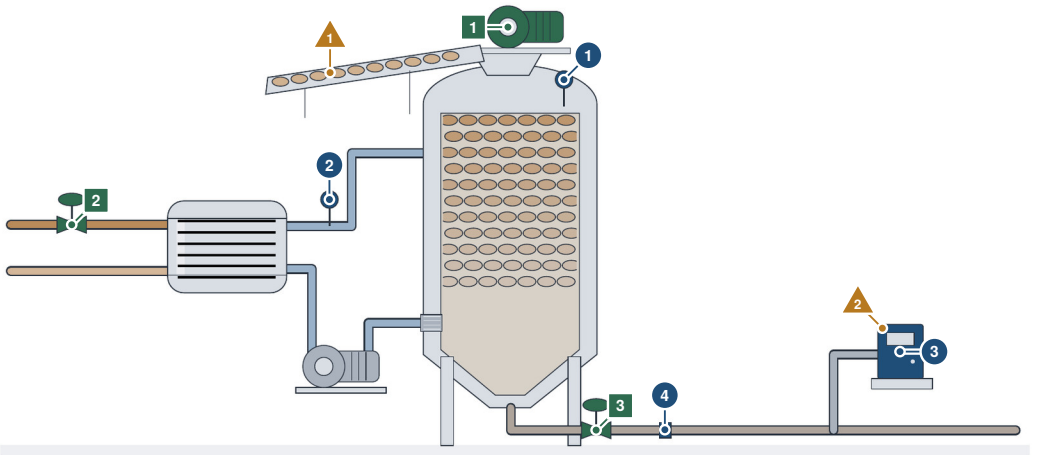

The controller moves 1 chip meter that feeds the chips 2 steam valve of the circulation heater 3 blow that sets the production rate

The controller reads 1 chip level 2 heater outlet temperature 3 online kappa analyzer 4 pulp flow at the blow

What disturbs it 1 a new wood or wetter chips 2 the analyzer drifting

Chips heated at the top reach the blowline about four hours later, and only then can the analyzer far down the line grade them. Illustration, not to scale.

IN PLAIN TERMS **Cooking Wood Chips by a Late Report**

The main idea. A digester cooks wood chips in a hot chemical bath that frees the fibers for paper. The chips sink slowly down a tall vessel, and a heater near the top sets how far they cook.

The challenge. The test that grades the pulp comes about four hours after the heater, and its instrument can slowly drift. The controller must choose what to hold steady between results, and how far to trust each late one.

Why it matters. A poor choice of what to hold can waste more than a hundred dollars of wood, chemicals and steam per tonne of pulp.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history segment and the fact on kraft pulping (Section 4.2), the try-it on the production step (Section 4.6), the fact on the held variable and the try-it on the drifting analyzer (the same section) and the key points
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 4.6), Technology in Practice (Section 4.7), Limitations (Section 4.8) and the two open challenges

The control question A continuous digester cooks its chips for four hours past the heater, and the kappa number that grades the cook arrives late from an analyzer that can drift and later still from a laboratory. What should the heater hold between quality reports, and how should the late and biased reports correct it?

Learning objectives

- Compute the cooking residence, Vroom's rate and the H-factor change that a heater step or a production change brings.
- Explain why a heater step reaches the blowline intact four hours later, and why no heater setting rescues the chips already past the heater.
- Derive the heat-balance and H-factor feedforwards for a production change.
- Compare held variables by their steady loss, and explain why holding the residual alkali makes a new wood worse.
- Reconcile a drifting analyzer with the laboratory, and judge when a better estimate is a worse feedback signal.

Prerequisites Sections 1.4.2, 1.4.4, 1.4.5, 1.4.6, 1.4.8 and 1.6.2. Chapter 2 for a delay in the path of the manipulated input.

The app https://apmonitor.com/apps/kraft-digester/

See also Volume 3, Chapter 6 for a dose computed from a measured upstream variable ahead of a three-hour delay, and Volume 3, Chapter 2 for matching loop structure to where a disturbance enters. Volume 9, Chapter 10 for choosing the setpoint a fast loop holds under a slow quality, and Volume 9, Chapter 8 for an optimum that sits on a constraint.

4.1 Introduction

The digester simulated here makes 1500 t/d of unbleached pulp from 132.98 t/h of dry chips at a design yield of 47.0 %. The chips spend an hour in an impregnation zone at 115 °C and four hours in a cooking zone at about 160 °C, and they leave at kappa 30. When the cooking temperature is raised by 2 °C half an hour into a run, the pulp at the blowline does not change for 4.125 h, then its true kappa falls to 27.1–27.6 and stays there, and the online analyzer reports the new value at 5.08 h on the clock. No measurement of the cook exists during those hours, because nothing inside the vessel reports the lignin left in the chips.

A plant whose product is graded by a property measured long after the product was made is common well beyond pulping. Kappa itself cannot be measured on line inside a batch digester and has been inferred from liquor measurements [1], and the general problem of fusing outputs that arrive at different rates and with different delays has been posed as multirate state estimation for the chemical process industry [2]. A continuous digester adds two features to that class. The vessel is a plug-flow reactor, so a change made at the heater travels down with the chips and arrives intact one residence time later instead of spreading as a first-order lag, and the cook can be steered by two handles, temperature and alkali charge, that reach one quality through different paths. Between quality reports the temperature has to hold some measured variable constant, and the choice of that variable is a question of controlled-variable selection [3, 4].

The simulator corrects two beliefs that a block diagram leaves intact. The first is that a measured variable close to the quality is a safe thing to hold. Holding the residual alkali at the extraction screens with temperature looks reasonable, but on an alkali-hungry wood the residual falls, the heater cools to restore it and kappa rises to 34.8 against 32.7 for the H-factor held on the same wood. The second is that a better estimate of the quality is a better feedback signal. The reconciled kappa removes a drifting analyzer's bias, yet on the same disturbance day without the drift, at a 20 % charge, the H-factor held kappa within ± 2 on the raw analyzer at 11 of 11 seeds and on the reconciled estimate at none, because the estimate is built to wait out short upsets.

The published digester models and controllers are research tools built around fundamental models and a control benchmark [5, 6], and the industrial practice reported for kappa control reduces the

vessel to a first-order plus dead-time response [7]. Neither lets a learner see the H-factor of each slug of chips, the analyzer's and the laboratory's clocks and the truth on one axis, and the cost of each candidate held variable side by side. This chapter works with a kraft digester simulator that runs in a browser, described completely enough for a second implementation, and measures on it how feedforward, reconciliation and the choice of held variable act on a quality that answers hours late. The H-factor, the effective-alkali balance, the bias updates and the loss definition are established methods, and the results are measurements on one teaching model at stated seeds. The model carries the chip column as forty Lagrangian slugs whose lignin and alkali ledgers close to rounding, its estimator and controllers read nothing but declared instruments, and every rate law, constant, tuning and default is stated. Nine graded cases run from a silent blowline to a day with four simultaneous upsets.

Section 4.2 follows kraft pulping and its control. Section 4.3 gives the vessel, its slugs, its instruments and its constants, and Section 4.4 measures how a heater step and a production step travel through it. Section 4.5 sets out the control layers: three PI loops, the H-factor feedforward, the held variable chosen by its loss, the estimator that reconciles three clocks, the kappa trim and the economic optimizer. Section 4.6 measures the nine cases, verifies the code and compares the kinetics with published data, Section 4.7 sets the loops on plant hardware, and Sections 4.8 and 4.9 state the limitations and conclude.

4.2 Cooking to a Number: Kraft Pulping and Its Control

Kraft pulping dissolves the lignin that binds wood fibers in a hot solution of sodium hydroxide and sodium sulfide, and how far the cook has gone is graded by the kappa number of the pulp, a permanganate test whose value is about seven times the lignin left on the pulp in percent [8]. Mills steer a cook by one index of time and temperature, the H-factor, whose origin the history segment below tells. Kinetic models of the cook were developed alongside the index, among them a theoretical model of the kraft process published in 1983 [9], and measured rate constants of the bulk and residual delignification phases for several hardwoods and a pine remain an active subject [10, 11].

> ⧖ **HISTORY** **1957: One Number for Time and Temperature**
>
> The kraft process was patented by Dahl in 1884, and the first kraft mill went into operation in 1890. The rate of delignification rises steeply with temperature. In 1957 Vroom combined cooking time and temperature into one number, the H-factor, the time integral of a reaction rate relative to that at 100 °C, with an activation energy of 134 kJ/mol, so that one hour at 100 °C counts 1 and one hour at 171 °C about 1000. Mills have since cooked to an H-factor target, which predicts the kappa of a cook only while the alkali and the liquor-to-wood ratio stay fixed. Laboratory cooks of one eucalyptus hybrid at eight temperatures from 140 to 175 °C at one H-factor gave kappa numbers with no significant difference, and the yield fell at the hotter cooks [8].

The continuous digester moved the cook from a batch vessel to a column that the chips descend for hours, and the control problem moved with it. A fundamental model of a continuous digester, extended from the Purdue model with a lumped-parameter description of the flow transport, was published for simulation and control in 1997 [5]. Because kappa was not measured inside the vessel, the control literature of the 1990s was inferential. A batch digester was controlled through an extended Kalman filter on liquor measurements [1], a continuous digester through a model identified by subspace methods [12], and the effective alkali of a continuous digester became a controlled variable of its own [13]. A pulp mill benchmark problem for control was posed in 2004 [6], and nonlinear inferential control of kappa at several locations of a continuous digester was formulated with multirate measurements [14]. Prediction and control of the blow-line kappa number were reported in 2005 [15], and an industrial model predictive controller that treats the digester as a first-order plus dead-time process reduced the standard deviation of kappa by 40 % on a dual-vessel continuous digester [7].

Table 4.1 Classified Variables and Their Design Values

Variable	Role	Design	Range
κ	CV, blowline kappa (truth)	30.0	SP 20–45
T_h	CV of the heater loop	160 °C	140–170 °C
h	CV, chip level	5.0 m	SP 1–9 m
v_s	MV, steam valve	58.0 %	0–100 %
n_{cm}	MV, chip meter speed	100 %	0–150 %
T_{SP}	MV of the held variable	160 °C	140–170 °C
AC	MV, alkali charge on wood	20 %	14–26 %
P	production setpoint	1500 t/d	1000–1900
f_B	DV, wood B share	0	0–100 %
m	DV, chip moisture	50 %	40–62 %
s_{wl}	DV, liquor strength	100 %	70–110 %
L_f, L_s	states, lignin pools	0.224, 0.056	kg/kg wood
C, W	states, carbohydrate, acids	0.680, 0.030	kg/kg wood
E	state, effective alkali	50.0 g/L at entry	—
T, H	states, temperature, H-factor	115 °C, 0 at entry	—
E_r	output, residual alkali	10.1 g/L	floor 8 g/L
Y	output, yield	47.0 %	—

✦ DID YOU KNOW? Kraft Makes Most of the Pulp

A 2005 profile for the United States Department of Energy put about 98 % of the country's chemical pulp capacity, or 80 % of all its pulping capacity, on the kraft process. It also found that continuous digesters predominate, being more space-efficient and easier to control than batch digesters [16].

The question of which variable to hold constant has its own history. Optimizing control structures, in which a measured variable is held so that the plant stays near its optimum as disturbances arrive, were formalized for chemical plants in 1980 [3]. The self-optimizing form judges a candidate by its loss, the extra cost of holding it at a constant setpoint against re-optimizing after each disturbance [4], and combinations of measurements extend the candidate set [17]. A digester offers the textbook candidates on one vessel: the cooking temperature, the H-factor, the residual alkali and a model's predicted kappa.

4.3 The Digester as Modeled

The simulated digester, shown in Figure 4.1, is a generic teaching design with no mill as its source. Chips enter the top separator through a chip meter, white liquor is ratioed to them at an alkali charge AC, and the chip column descends through an impregnation zone and a cooking zone. The circulation heater at the top of the cooking zone sets the temperature of the chips that pass it, and they carry that temperature down, cooling slowly to the surroundings. The spent liquor is drawn through the extraction screens, where a titrator reads the residual alkali E_r, and the pulp is blown below them, where an online analyzer and the laboratory read its kappa. The chip level h between the chip meter and the blow is the vessel's one inventory.

The variables are listed in Table 4.1 with their role. Time is in hours, wood flows in tonnes of dry wood per hour, alkali in grams per liter as NaOH and temperature in °C. Five lumped states are integrated together with ten states for each of the 40 slugs. The lumped states are the heater outlet temperature T_h, the chip level h, the column position ϕ in slugs, and two ledgers for the steam used and the heat given to the chips.

4.3.1 One Slug of Chips

Each slug is a closed batch of dry wood and the liquor that travels with it. Its state per kilogram of original dry wood is the bulk lignin L_f, the residual lignin L_s, the carbohydrate C, the alkali-consuming acid groups W (in kilograms of NaOH they consume), the effective alkali E in the liquor,

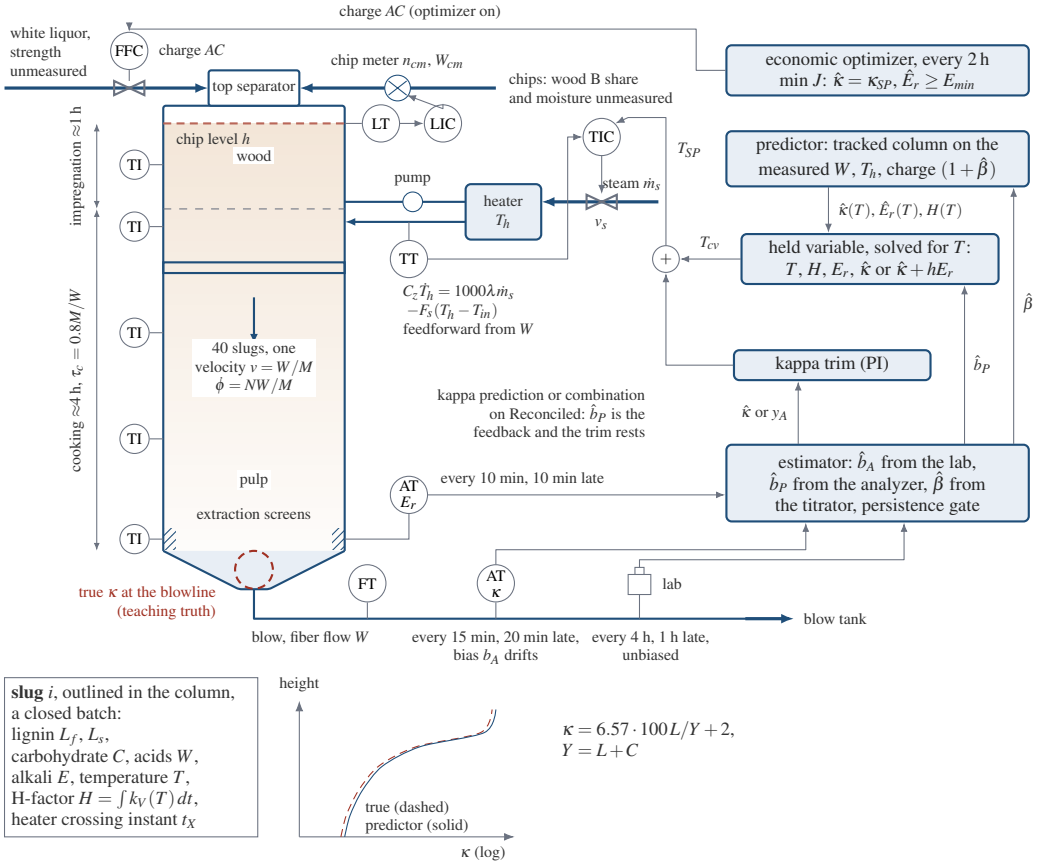

Figure 4.1 Continuous kraft digester as modeled. The column of $N = 40$ slugs is painted from wood to pulp. One slug, outlined, carries its lignin pools L_f and L_s, carbohydrate C, acid groups W, alkali E, temperature T, H-factor H and heater crossing instant t_X. The column moves at one velocity $v = W/M$ set by the blow's fiber flow W. The circulation heater sets T_h with the steam valve v_s under a temperature controller whose setpoint T_{SP} is the held variable's temperature plus the kappa trim. Measurements are open circles, the truth at the blowline and the true profile are dashed (teaching truth), and the inputs of the predictor are tagged rather than wired. Process model only, not to scale.

the temperature T in kelvin and the H-factor H, and its rates in hours are given by Equation (4.1).

$$
\begin{aligned}
k_V(T) &= \exp\bigl(43.2 - 16113/T\bigr), \qquad \frac{\mathrm{d}H}{\mathrm{d}t} = k_V(T) \\
\frac{\mathrm{d}L_f}{\mathrm{d}t} &= -A_L \rho_s k_V E^{0.5} L_f, \qquad \frac{\mathrm{d}L_s}{\mathrm{d}t} = -\frac{A_L}{r_s}\rho_s k_V E^{0.5} L_s \\
\frac{\mathrm{d}C}{\mathrm{d}t} &= -A_C k_C(T)\, E^{0.75} \max(0, C - C_\infty) \\
\frac{\mathrm{d}W}{\mathrm{d}t} &= -A_W k_W(T)\, E\, W \\
\frac{\mathrm{d}E}{\mathrm{d}t} &= -\frac{1000}{LW}\bigl(b_L r_L + b_C r_C + r_W\bigr) \\
\frac{\mathrm{d}T}{\mathrm{d}t} &= -\frac{T - T_a}{\tau_{loss}}
\end{aligned}
\tag{4.1}
$$

where r_L, r_C and r_W are the magnitudes of the lignin, carbohydrate and acid-group rates, ρ_s is the wood's reactivity, LW the liquor-to-wood ratio in liters per kilogram, and k_C and k_W are Arrhenius factors equal to 1 at 100 °C with activation temperatures $E_C/R = 17\,500\,\mathrm{K}$ and $E_W/R = 6000\,\mathrm{K}$. Vroom's rate is 1.019 at 100 °C with the constant 43.2 and 403.6 per hour at 160 °C, and $16113\,R$ is 134.0 kJ/mol. The carbohydrate loss has the higher activation energy (145.5 kJ/mol), so a hotter cook loses more yield at the same H-factor, as the laboratory cooks showed [8]. Kappa and yield follow from the slug's lignin and carbohydrate, as shown in Equation (4.2).

$$\kappa = 6.57 \cdot 100\,\frac{L}{Y} + 2, \qquad L = L_f + L_s, \qquad Y = L + C \tag{4.2}$$

where 6.57 kappa units per percent of lignin on pulp is the commonly quoted proportionality (about seven [8]) and the 2 is a teaching allowance for the hexenuronic acids the test also oxidizes.

A fresh slug enters with $L_f = 0.8L_0$, $L_s = 0.2L_0$, $C = C_0$, $W = W_0$, $T = T_{imp}$, $H = 0$ and $E_0 = 1000\,AC\,s_{wl}/LW$, where $LW = 3.0 + m/(1-m)$, so that 50 % moisture gives 4.0 L/kg and a 20 % charge gives 50 g/L. Wood A, the design wood, has $L_0 = 0.280$, $C_0 = 0.680$, $W_0 = 0.030$ and $\rho_s = 1$. Wood B has $L_0 = 0.290$, $C_0 = 0.670$, $W_0 = 0.050$ and $\rho_s = 0.95$, so it carries more lignin, consumes more alkali and cooks a little slower, and a mix is taken linearly in the share f_B. Each slug also keeps ledgers of the lignin and carbohydrate dissolved and the alkali consumed, which make $L_f + L_s + D_L = L_0$, $C + D_C = C_0$ and $E\,LW/1000 + A = E_0 LW/1000$ linear invariants that the integrator preserves.

4.3.2 A Column of Slugs

The column is carried by the method of characteristics. All 40 slugs ride it at one velocity, chips and free liquor together, and its position advances at the blow's fiber flow, as shown in Equation (4.3).

$$\begin{aligned} \frac{\mathrm{d}\phi}{\mathrm{d}t} &= \frac{N W_{col}}{M}, \qquad \frac{\mathrm{d}h}{\mathrm{d}t} = \frac{W_{cm} - W_{col}}{\rho_A} \\ C_z \frac{\mathrm{d}T_h}{\mathrm{d}t} &= 1000\,\lambda\,\dot{m}_s - F_s\,(T_h - T_{in}) \\ F_s &= \max\!\left(F_c,\ 0.1 F_{c,0}\right) \\ F_c &= 1000\,W_{col}\,(c_w + LW\,c_l) \end{aligned} \tag{4.3}$$

where $M = 664.9\,\mathrm{t}$ is the dry wood held in the two zones, W_{col} the column's descent in tonnes of wood per hour (the blow's flow, zero while the column hangs), W_{cm} the chip meter's flow, $\rho_A = 13.35\,\mathrm{t/m}$ the wood per meter of column, T_{in} the temperature of the slug arriving from impregnation, $\dot{m}_s = v_s\,\dot{m}_{s,max}\,p_s$ the steam in tonnes per hour with $\dot{m}_{s,max} = 90\,\mathrm{t/h}$ and p_s the header pressure as a fraction of design, $c_w = 1.35\,\mathrm{kJ/(kg\,K)}$, $c_l = 4.0\,\mathrm{kJ/(L\,K)}$, $\lambda = 2000\,\mathrm{kJ/kg}$, and $C_z = \tau_h F_{c,0}$ with $\tau_h = 0.1\,\mathrm{h}$ and $F_{c,0} = 2307.2\,\mathrm{MJ/(h\,K)}$ at design. The floor on F_s is a declared stand-in for the heat the circulation trades with a hung column. When ϕ passes 1 the bottom slug is blown with the instant it crossed, every slug moves down one slot, a fresh slug of the present feed enters the top, and the slug that moves into the first cooking slot takes T_h and records its crossing instant t_X. The cooking zone holds $0.8M$, so its residence is $\tau_c = 0.8M/W$, 4.000 h at design and 3.478 h at 1725 t/d, and one slug passes every 7.5 min at design. In this chapter τ_c is this residence and not the closed-loop time constant of Chapter 1. Because the column advances with the fiber flow, the delay of every slug is kept in wood rather than in time, the transport delay defined in volume of Section 1.4.4. Chapter 7 follows its material the same way along a line of well-mixed cells.

A hang-up holds the whole column for 0.5 h while the chip meter keeps feeding the top, and it then drops the distance it missed at once. The slugs in the cooking zone keep cooking while they hang. Every such event, and the weak-liquor hour at a strength of 85 %, the analyzer's drift and its recalibration, is applied once, between integration steps, and marked on the trends, so that no event hides inside the rate equations.

4.3.3 What the Instruments Report

The controllers read only a measurement packet. The heater outlet carries colored noise of 0.15 °C with a 0.02 h correlation time, five thermocouples at fixed heights read 0.3 °C noise, the blow's fiber flow 0.5 %, the level 3 cm and the steam flow 1 %. Three instruments report samples, as shown in Equation (4.4).

$$\begin{aligned} y_E(t_s + \tfrac{1}{6}) &= E_r(t_s) + v_E, & t_s &\in \{0, \tfrac{1}{6}, \tfrac{2}{6}, \dots\} \\ y_A(t_s + \tfrac{1}{3}) &= \kappa(t_s) + b_A(t_s) + v_A, & t_s &\in \{0, \tfrac{1}{4}, \tfrac{1}{2}, \dots\} \\ y_L(t_s + 1) &= \kappa(t_s) + v_L, & t_s &\in \{0, 4, 8, \dots\} \\ \frac{\mathrm{d}b_A}{\mathrm{d}t} &= \delta \text{ after the drift event} \end{aligned} \tag{4.4}$$

where $\kappa(t_s)$ and $E_r(t_s)$ are the true kappa and residual alkali of the slug at the blowline when the sample is drawn, the noise standard deviations are 0.2 g/L, 0.4 and 0.3, and $\delta = 0.3$ kappa per hour. The titrator samples every 10 min and reports 10 min later, the analyzer every 15 min and 20 min later, and the laboratory every 4 h and an hour later, without bias, on bottles drawn at the instants of analyzer samples. A reading holds until the next report, no sample is drawn while the column hangs, and every report is queued with the instant its sample left the blowline. The wood, the moisture, the liquor's strength and the analyzer's bias are not measured. The age of each held result, and the dead time it adds to a loop, follow Section 1.4.5.

4.3.4 Constants, Start-Up and Integration

Table 4.2 gives the constants with their defaults. The kinetic constants are teaching values. A_L and A_C were fitted so that the design slug leaves at kappa 30.000 and yield 47.00 %, with a residual alkali of 10.12 g/L and an H-factor of 1505.5 after cooling 1.72 °C over the cook. Each case starts lined out: one slug of the present feed is marched down the column at the present heater temperature and flow and copied into each slot at its age. The engine of Section 1.6.1 steps 10 s, integrating the lumped states and every slug by the classical fourth-order Runge-Kutta method, and the controller executes and a record is stored once per simulated minute.

4.3.5 Assumptions

The free liquor rides with the chips and there is no axial mixing, so the model has no counter-current wash zone and the cook stops at the extraction screens. Effective alkali stands for the hydroxide and hydrosulfide together, so sulfidity is not a variable. The kinetics are teaching values in the form of a two-pool lignin model and are not any wood's. The heater sets the temperature of each slug that crosses it, so radial gradients, channeling and the steaming of chips before the vessel are outside the model, and the reaction heat is neglected. A wood mix is taken as one wood with averaged composition and reactivity, while a real mix cooks each chip by its own kinetics. The residual alkali floor and the temperature limit are teaching values. A mill sets its floor from trials that keep dissolved lignin from redepositing on the fibers and its limit from the vessel's rating and the pulp's strength, and the model does not simulate redeposition, so the floor is judged as a declared limit. A hang-up holds the whole column, while a real column hangs above some point and leaves a void below it.

4.4 Hours Between the Heater and the Blowline

4.4.1 A Step That Arrives Intact

Case 1 holds the heater at 160 °C and the learner raises the setpoint to 162 °C at 0.5 h. The first slug to cross the heater after the step does so at 0.625 h, the next slug boundary plus the heater's lag, and reaches the blowline at 4.625 h. The slugs that follow leave at a median of 27.32 from 5 to 10 h, within 27.08–27.60 from 6 to 10 h, and the analyzer reports the change at 5.08 h from a sample drawn at 4.75 h. At 160 °C the logarithm of Vroom's rate rises 8.59 % per kelvin, so two degrees multiply the cooking H-factor by 1.186. The plug flow keeps the step's shape. A tanks-in-series grid of 40

Table 4.2 Model Constants. Those Marked Fixed Have No Control in the Simulator

Symbol	Meaning	Value	Unit
N, M	slugs, dry wood in the zones	40, 664.9	—, t
—	impregnation share (fixed)	0.2	—
ρ_A	dry wood per meter	13.35	t/m
T_{imp}	impregnation temperature	115	°C
—	liquor added per kg wood	3.0	L/kg
c_w	heat capacity of wood	1.35	kJ/(kg K)
c_l	heat capacity of liquor	4.0	kJ/(L K)
$\lambda, \dot{m}_{s,max}$	steam heat, valve capacity	2000, 90	kJ/kg, t/h
τ_h	heater time constant	0.1	h
A_L	bulk delignification (fixed)	9.5553×10^{-4}	$\mathrm{h}^{-1}\,(\mathrm{g/L})^{-0.5}$
r_s, f_s	residual pool slowness, share (fixed)	6, 0.20	—
A_C	carbohydrate loss (fixed)	4.0771×10^{-5}	$\mathrm{h}^{-1}\,(\mathrm{g/L})^{-0.75}$
C_∞	surviving carbohydrate (fixed)	0.30	kg/kg
A_W	acid-group consumption (fixed)	0.06	$\mathrm{h}^{-1}\,(\mathrm{g/L})^{-1}$
b_L, b_C	alkali per lignin, carbohydrate (fixed)	0.10, 0.45	kg/kg
T_a, τ_{loss}	surroundings, cooling time (fixed)	30, 300	°C, h
—	delignification multiplier	1.0	—
E/R	plant activation temperature	16113	K
—	analyzer period, delay, noise	0.25, 0.33, 0.4	h, h, —
δ	analyzer drift when on	0.3	h^{-1}
—	laboratory period, delay, noise	4, 1, 0.3	h, h, —
—	hang-up duration	0.5	h
—	weak liquor strength	85	%
p_s, p_a, p_w	steam, alkali, wood price	20, 150, 70	$/t
q	off-target kappa cost	5	$/t per kappa^2
E_{min}	residual alkali floor	8	g/L
T_{max}	cooking temperature limit	170	°C

cells, which an early design of the simulator specified, smeared the design kappa by 2.64 % and failed the refinement test, which is why the column is carried as slugs.

4.4.2 Chips Caught in Transit

A production increase is the disturbance that a fixed cooking temperature cannot absorb. In case 2 the production rises 15 % at 1 h, the residence falls from 5.00 to 4.348 h and the cooking time to 3.478 h, so at a fixed temperature every slug gathers 0.8696 of the design cooking H-factor. The first slugs are 0.3 off by 1.54 h, kappa reaches 32.56 at 5.46 h, and it stays within 32.04–32.61 from 8 to 12 h. The chips already in the cooking zone when the step arrives have crossed the heater, and no later move of the heater reaches them. A sweep of eight heater settings over the worst slug blown from 1 to 8 h measures the consequence. Temperature alone reaches +2.62, temperature with an analyzer trim at gains of 0.15, 0.3, 0.5 and 1.0 °C per kappa reaches +2.75, +3.05, +3.45 and +4.49, the H-factor +2.46, the H-factor with the analyzer trim +2.75, and the kappa prediction on the reconciled estimate +2.59. In seven of the eight the worst slug crossed the heater at the step. No setting of the heater holds ±2 through this step, which is the honestly infeasible variant of the case.

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

4.4.3 The Premise of the Feedforward

The H-factor is worth feeding forward only if one H-factor gives one kappa across production rates. At an H-factor of 1505 and a 20 % charge, an independent bisection on the heater gives 158.45, 160.00 and 161.53 °C at 1300, 1500 and 1725 t/d, and the slugs leave at kappa 29.842, 30.005 and 30.166. The spread of 0.32 is the impregnation's share of the alkali consumption, which shortens with the rate. The engine with the H-factor held and the noise off gives 29.848, 30.010 and 30.170 from 10 to 14 h. The premise holds to a third of a kappa unit on this model, and the laboratory cooks support it within their scatter [8].

4.5 The Control Layers as Coded

The controller runs every minute on the measurement packet. Production is the blow's fiber flow, $W = P/(24 \cdot 0.47)$ in tonnes of wood per hour, so the throughput is set at the bottom of the vessel and the chip level is held at the top.

4.5.1 The Three PI Loops

The heater, level and kappa-trim loops are proportional-integral loops in the velocity form of Section 1.6.2, written for these loops in Equation (4.5).

$$\begin{aligned} e_k &= a\,(r_k - y_k), \qquad s_k = s_{k-1} + K_c(e_k - e_{k-1}) \\ \hat{u}_k &= u_b + s_k + I_{k-1}, \qquad u_k = \min\bigl(u_{max}, \max(u_{min}, \hat{u}_k)\bigr) \\ I_k &= I_{k-1} + \Bigl[\frac{K_c}{\tau_I} e_k + \frac{u_k - \hat{u}_k}{\tau_I}\Bigr]\Delta t \end{aligned} \tag{4.5}$$

where $a = \pm 1$ is the action, u_b a bias and $\Delta t = 1\,\text{min}$. The symbols map onto Equations (1.15) to (1.18) as $s_k = S_k$ and $\hat{u}_k = v_k$, with no derivative term and a tracking time constant equal to τ_I. The heater loop drives the steam valve from 0 to 100 % at $K_c = 2.5\,\%$ per °C and $\tau_I = 6\,\text{min}$. Its bias is a heat-balance feedforward, the valve that lifts the chips from impregnation to T_{SP} at the measured fiber flow, as shown in Equation (4.6).

$$u_b = \frac{100}{\dot{m}_{s,max}} \, \frac{W_f\,(c_w + 4.0\,c_l)(T_{SP} - T_{in,0})}{\lambda} \tag{4.6}$$

where W_f is the measured fiber flow through a 0.02 h filter, limited to 1.5 times the production setpoint and to no less than a tenth of the design flow, and $T_{in,0} = 30 + (T_{imp} - 30)\exp(-0.2M/(300\,W_f))$, in °C, is the declared heat balance's estimate of the arriving chips, 114.72 °C at design. It holds 58.0 % of the valve at design and follows a production move within minutes, before the heater's outlet has fallen. The level loop moves the chip meter from 0 to 150 % of the design wood flow at 10 % per meter with $\tau_I = 2\,\text{h}$, an averaging tuning that lets the level absorb the chip meter's moves. The kappa trim acts on the cooking temperature at 0.15 °C per kappa with $\tau_I = 3\,\text{h}$, reverse acting, within $\pm 12\,^{\circ}\text{C}$. The setpoint of the heater loop is $T_{SP} = \min(170, \max(140, T_{cv} + \text{trim}))$, where T_{cv} is the held variable's temperature.

4.5.2 The H-Factor Feedforward

The held variable H-factor chooses the heater temperature that gives the chips entering the cooking zone now their target H-factor at the measured rate, as shown in Equation (4.7).

$$\begin{aligned} H_{SP} &= H_{imp} + \int_0^{\tau_c} k_V\bigl(T_a + (T_h - T_a)\mathrm{e}^{-t/\tau_{loss}}\bigr)\,dt \\ \tau_c &= 0.8M/W \end{aligned} \tag{4.7}$$

where H_{imp} is the H-factor the predictor's slug gathers in impregnation. The integral is taken by Simpson's rule on 24 intervals. Because $\ln H$ is nearly linear in T_h, Newton's method on $\ln H$ with the slope $(E/R)/(T_h + 272.65)^2$, kept inside a bisection bracket on 140–170 °C, stops when the H-factor is within 10^{-6} of the target or after 30 iterations. A target that the bracket cannot reach returns the bound and is flagged, and the solve never extrapolates. On a fresh instance it reached 162.2698 °C in three iterations against a brute-force scan's 162.2700 °C.

4.5.3 Which Variable the Heater Holds

Five candidates are offered, the temperature itself, the H-factor, the residual alkali E_r, the predicted kappa $\hat{\kappa}$, and the combination $\hat{\kappa} + hE_r$. The last three are computed by the controller's predictor, the declared kinetics run on the design wood and moisture, the commanded charge times $(1 + \hat{\beta})$ and the

Table 4.3 Steady Loss of Each Held Variable ($/t Pulp) at a 20 % Charge, With the Steady Kappa on Wood B

Held variable	Wood B	Moisture	Liquor	Rate	κ on B
temperature	29.45	0.00	12.83	20.12	32.75
H-factor	29.45	0.00	12.83	0.19	32.75
residual alkali	119.68	0.47	190.07	0.00	35.21
kappa prediction	9.90	0.05	0.50	0.63	31.74
combination	7.27	0.08	2.27	0.71	31.54

measured flow, for the slug about to enter the cooking zone. Each is solved for the temperature that holds it at its setpoint by the Illinois variant of regula falsi on 140–170 °C, stopping within 10^{-3} or after 16 iterations, and the solves agreed with 0.001 °C scans to 5.6×10^{-4} °C.

A candidate is scored by its loss, as shown in Equation (4.8).

$$\begin{aligned} L(c,d) &= J\big(T_h(c=c_s,d),d\big) - \min_{T_h} J(T_h,d) \\ J &= \frac{p_w + p_a AC_d + p_s\, s(T_h)}{Y} + q\,(\kappa - \kappa_{SP})^2 \end{aligned} \tag{4.8}$$

where J is the cost in dollars per tonne of pulp of the plant's steady state, AC_d the charge delivered, s the steam per tonne of wood, and c_s the candidate's value at the design point. The candidate is measured as the controller would see it at steady state, with $\hat{\beta}$ fitted so that the predictor's residual equals the plant's. The disturbances are wood B at 60 %, moisture at 56 %, liquor at 90 % strength and production up 15 %. The ideal is found by golden section on 148–172 °C, the held temperature by the candidate's crossing on a 0.25 °C grid with the plant evaluated at that temperature, and the weight h by golden section on $[-2,2]$ minimizing the summed loss over wood B and the weak liquor. At the 20 % charge $h = -0.0869$ kappa per g/L, and the table is given in Table 4.3. Holding the residual alkali with temperature is the worst choice on the wood and the liquor, because both lower the residual and the heater then cools. The kappa prediction and the combination recover most of the loss on the wood, and the H-factor and the residual alkali are the candidates that ignore a production change and a moisture change respectively. Volume 9, Chapter 10 asks the same question of a dissolved oxygen setpoint held under an effluent ammonia that answers in hours.

4.5.4 Reconciling Three Clocks

The estimator keeps its own column of 40 slugs, advanced on the measured fiber flow, given the measured heater temperature as each slug crosses, and cooked with the declared kinetics on the commanded charge, so its predicted kappa $\hat{\kappa}_P$ exists for every slug in the vessel and for the slug leaving now. Each update pairs two sources for the same sample, aligned by the instant the sample left the blowline, as shown in Equation (4.9).

$$\begin{aligned} \hat{b}_A &\leftarrow \hat{b}_A + \lambda_A\big(y_A(t_s) - y_L(t_s) - \hat{b}_A\big) \\ \hat{b}_P &\leftarrow \hat{b}_P + \lambda_P\big(y_A - \hat{b}_A - \hat{\kappa}_P(t_s) - \hat{b}_P\big) \\ \hat{\beta} &\leftarrow \hat{\beta} + g_\beta\left(\beta_s + \frac{y_E - \hat{E}_P(t_s)}{\partial E/\partial \ln AC} - \hat{\beta}\right) \\ \hat{\kappa}(t) &= \hat{\kappa}_P(t) + \hat{b}_P \end{aligned} \tag{4.9}$$

where the first update runs on each laboratory report, the second on each analyzer report and the third on each titrator report, $\lambda_A = 0.8$, $\lambda_P = 0.1$, $g_\beta = 0.2$, β_s is the factor the tracked slug was given when it entered five hours earlier, and $\partial E/\partial \ln AC = 31\,\mathrm{g/L}$ at design. Averaging the factor each report implies for its own slug, rather than integrating the raw residual, keeps the five-hour delay out of the update. Each line is a bias update of the form of Equation (1.12), compared with the value stored for the same sample as in Equation (1.14). The laboratory is the one unbiased instrument, so it alone

makes the analyzer's bias observable (Section 1.4.6). Chapter 10 corrects a soft sensor from a delayed laboratory in the same way. The biases are clamped to ± 15 kappa and $\hat{\beta}$ to ± 0.4, and an update whose partner is missing is skipped. A persistence gate guards the two predictor updates. A titrator report whose implied factor lies more than 0.05 from $\hat{\beta}$, or an analyzer report more than 1.5 kappa from its reconciled prediction, is held back until seven titrator or five analyzer reports on the same side agree, which is more than an hour of reports. On the titrator a run must also stay within 0.05 of the factor that started it. An upset that passes within an hour, such as a weak-liquor hour, therefore never enters the slugs that follow it, while a new wood is believed after its first hour.

4.5.5 The Kappa Trim and the Economic Optimizer

The kappa feedback is Off, a PI on the raw analyzer, or Reconciled. With the temperature, the H-factor or the residual alkali held, Reconciled feeds $\hat{\kappa}$ to the PI. With the kappa prediction or the combination held, Reconciled adds $\hat{b}_P$ to the predicted kappa inside the solve, so the bias is the feedback and the trim rests at zero. The economic optimizer runs every 2 h when switched on and chooses the charge, as shown in Equation (4.10).

$$\begin{aligned} \min_{AC,\,T_h} \; & \frac{p_w + p_a AC(1+\hat{\beta}) + p_s\, s(T_h)}{\hat{Y}} \\ \text{s.t. } & \hat{\kappa}(AC,T_h) + \hat{b}_P = \kappa_{SP}, \quad \hat{E}_r(AC,T_h) \geq E_{min} \\ & 140 \leq T_h \leq T_{max}, \quad 0.15 \leq AC \leq 0.26 \end{aligned} \tag{4.10}$$

The equality removes T_h by the regula falsi solve, a bisection of 22 halvings finds the least charge whose predicted residual stays on the floor, and a golden-section search of 24 steps minimizes the cost along the remaining range, 51 predictor evaluations in all. The optimizer holds the floor on its own predictor with no back-off. On a fresh instance it found 21.237 % at 163.25 °C and \$242.4918/t, against a brute-force search's 21.240 % and \$242.5105/t, which cannot land on the floor.

4.5.6 Three Models Kept Apart

The plant model is Section 4.3 with the true wood, moisture, liquor and analyzer bias. The controller's model is the predictor and the estimator's column, the declared kinetics on the design wood, which know none of the four. The loss table is a designer's offline analysis on the declared model with the rate multiplier at 1 and Vroom's activation temperature. A grep of the controller, estimator, predictor and optimizer sources finds no read of any truth field.

4.6 Nine Scenarios, Measured

The cases appear in Table 4.4. Every case starts lined out at the design point with the defaults above and departs from them as its row says. The numbers below come from the simulator's own engine run without a page at the case's seed (20261030), advanced step by step exactly as the page would advance it, with a learner's setting entered after the run starts where the case asks for it, and for the challenge also at seeds 1–10. The judged quantity is the true kappa of each slug as it is blown, a quantity no instrument reports. The exported file holds the analyzer, laboratory, reconciled and predicted kappa beside the true kappa, whose column name marks it as teaching truth (Section 1.6.4).

4.6.1 Feedforward Through the Rate Step

With the H-factor held in case 3 the heater setpoint jumps from 160.00 to 161.53 °C at the step and the predicted H-factor of the entering chips stays at 1505.0. The chips already past the heater lose their time all the same and peak at 32.46 at 4.48 h, which is the last slug outside ± 1 of 30, and from 7 to 12 h kappa stays within 29.88–30.42. The feedforward ends the upset one cooking residence after the step, while temperature alone never ends it. Volume 3, Chapter 6 makes the same move on a coagulant dose, computed from a measured upstream variable with the slow product measurement demoted to a trim.

Table 4.4 The Nine Scenarios

#	Case and events	Teaching point
1	Temperature held, learner steps T_{SP}, 10 h	the four-hour silence
2	Production +15 % at 1 h, temperature held, 12 h	chips caught in transit, infeasible variant
3	As 2 with the H-factor held	feedforward
4	Wood B 60 % at 1 h, each held variable, 12 h	self-optimizing loss
5	Analyzer drift from 1 h, kappa prediction, analyzer trim, 24 h	multirate reconciliation
6	Wood B at 1 h, H-factor, analyzer trim, 36 h	feedback through the delay
7	Kappa prediction, reconciled, optimizer at 1 h, 12 h	steam against alkali
8	Hang-up at 2 h, H-factor, 10 h	level and over-cooked slugs
9	Challenge: rate, drift, wood B and a weak hour, 24 h	the whole problem

TRY IT IN THE APP **3 · The Same Step With H-Factor Feedforward**

Load the case and press Run at 0.1 h/s. Watch the cooking temperature setpoint jump at the production step at 1 h, then the true kappa (dashed). The chips already past the heater still peak at 32.46 at 4.48 h, and from 7 to 12 h kappa stays within 29.88–30.42.

4.6.2 A New Wood and the Held Variable

Case 4 makes 60 % of the chips wood B at 1 h. Holding the H-factor, kappa reaches 32.59–33.03 from 10 to 14 h with the heater setpoint unmoved at 160.00 °C. Holding the residual alkali, the heater cools to 158.35 °C and kappa reaches 32.50–35.09, because the wood's acid groups consume alkali and the heater tries to put it back. The kappa prediction reads the lower residual as less alkali, heats to 160.73 °C and holds 31.63–32.79, and the combination heats to 160.87 °C and holds 31.44–32.81. The steady loss table ranks them in the same order, $119.68, $29.45, $9.90 and $7.27 per tonne. None reaches 30, because the part of wood B that is slower rather than hungrier reaches no instrument but kappa. Holding the kappa prediction at the design charge, the titrator reads 7.68 g/L at 14 h, under the 8 g/L floor, so the heater can hold kappa but only the charge can hold the floor.

Pause. Try Exercise 4.1 before reading on.

DID YOU KNOW? **The Closest Measurement Is the Worst to Hold**

From the simulator: on wood B, holding the residual alkali measured inside the cook costs $119.68 per tonne of pulp against $29.45 for the H-factor, and its steady kappa is 35.21 against 32.75 (Table 4.3). The heater cools to restore the alkali the new wood consumed.

4.6.3 A Drifting Analyzer and the Laboratory

In case 5 the analyzer's bias drifts upward at 0.3 kappa per hour from 1 h, and the measured traces are shown in Figure 4.2. With the trim on the raw analyzer the loop holds the analyzer a few units above 30 while the true kappa falls to 26.52 at 24 h, the bias then being 6.90. With feedback through the reconciled estimate the true kappa stays within 28.68–29.51 from 12 to 24 h. The laboratory pins $\hat{b}_A$ once every four hours, at 5.56 after the report of 21 h, and the 1.34 it has not yet seen at 24 h leaks into the predictor bias. With no feedback at all the predictor alone holds 29.65–30.49, because nothing in this case changes what the predictor knows.

TRY IT IN THE APP **5 · The Analyzer Drifts and the Laboratory Corrects It**

Press Run at 0.5 h/s on the Clocks view. The trim on the raw analyzer holds the analyzer a few units above 30 while the true kappa (dashed) falls to 26.52 at 24 h. Set Kappa feedback to Reconciled, press Reset and run again. From 12 to 24 h the true kappa now stays within 28.68–29.51.

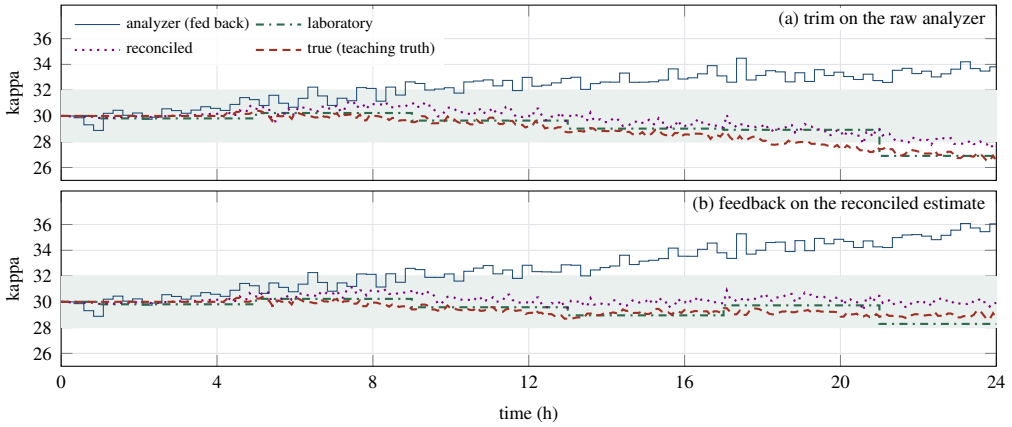

Figure 4.2 Kappa on three clocks in case 5, the analyzer drifting from 1 h. (a) The trim on the raw analyzer holds the analyzer a few units above 30 while the true kappa falls to 26.52. (b) Feedback through the reconciled estimate keeps the true kappa within 28.68–29.51 after 12 h. The analyzer (thin solid) and the laboratory (dash-dotted) hold their last report, the reconciled estimate is dotted, the truth is dashed, and the shaded band is ±2.

4.6.4 Feedback Through Five Hours

Case 6 holds the H-factor on wood B with the analyzer trim and sweeps its gain. At 0.15 and 0.3 °C per kappa the last slug outside ±1 leaves at 24.12 and 12.25 h. At 0.5 kappa undershoots to 27.88 at 14.50 h and is not settled until 33.25 h, and at 0.8 it rings and grows to 20.22–41.10 from 30 to 34 h. The vessel's delay sets how fast the feedback can be, whatever the gain, which is the bandwidth limit of Section 1.4.2.

4.6.5 Steam Against Alkali

Case 7 turns the optimizer on at 1 h. It lowers the charge from 20.00 to 18.83 % and the heater setpoint rises from 159.85 to 160.61 °C. From 7 to 12 h the alkali falls from 425.0 to 396.3 kg per tonne of pulp, the steam rises from 0.814 to 0.817 t/t and the yield from 46.97 to 47.60 %, with 41 of 41 slugs within ±2. The optimum sits on the residual floor of its predictor, and the true residual from 7 to 24 h is 7.72–8.39 g/L, under the floor in 49 % of the steps. A constraint held at its limit on a model is broken about half the time by the model's error, which is the reason a mill backs off from it. As in the data hall of Volume 9, Chapter 8, the cheapest feasible point sits on a constraint.

4.6.6 A Hang-Up

In case 8 the column hangs at 2 h. The level climbs from 5.0 to 8.86 m while the chip meter feeds a column that does not move, and it is 3.87 m at the drop at 2.50 h. Four slugs leave at the drop at kappa 28.56, 29.02, 29.49 and 29.99, the lowest with an H-factor of 1636.0 against 1505.5 for an untouched slug, each equal to the declared kinetics cooked for its own 4.375, 4.250, 4.125 and 4.000 h from crossing to exit. Four hours later the chips that fell past the heater in the drop arrive under-cooked at 31.36 at 6.12 h, and the slug that crossed while the heater settled at 161.00 °C leaves at 28.63 at 6.62 h. The heater, whose load returns at once when the column drops, swings between 154.14 and 161.56 °C after the drop.

4.6.7 The Challenge

The challenge runs 24 h. At 1 h production rises 15 % and the analyzer starts to drift, at 2 h the chips become 60 % wood B, and from 3 h the liquor runs at 85 % strength for an hour. The target is that every slug blown from 12 to 24 h leaves within ±2 of kappa 30 and with a true residual alkali at

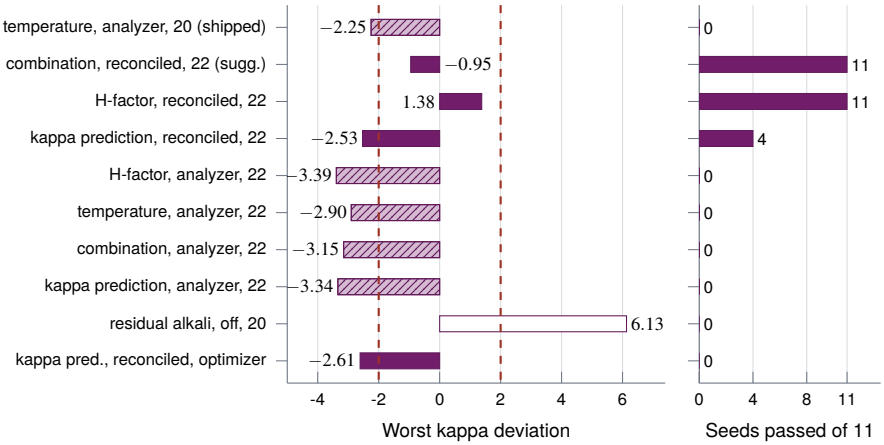

Figure 4.3 The challenge day by structure, from Table 4.5: the worst deviation of the true kappa from 30 over the eleven seeds (left, dashed lines at ± 2) and the seeds that held kappa within ± 2 and the residual alkali on its floor (right). Rows name the held variable, the feedback and the charge in percent. Dark bars are feedback through the reconciled estimate, hatched bars the raw analyzer, the white bar no feedback. Notice that every structure on the raw analyzer falls below -2 because of the drift, and that only the combination and the H-factor on the reconciled estimate pass every seed.

Table 4.5 The Challenge at the Case's Seed and Seeds 1–10. Worst Kappa Deviation and Lowest Residual Over the Eleven, Steam and Alkali per Tonne of Pulp at the Case's Seed

Held, feedback, charge	Pass	Worst	E_r low (g/L)	Steam (t/t)	Alkali (kg/t)
temperature, analyzer, 20 (shipped)	0	−2.25	5.39	0.942	434
combination, reconciled, 22 (sugg.)	11	−0.95	10.07	0.915	487
H-factor, reconciled, 22	11	1.38	10.08	0.914	485
kappa prediction, reconciled, 22	4	−2.53	9.08	0.949	496
H-factor, analyzer, 22	0	−3.39	8.55	0.970	498
temperature, analyzer, 22	0	−2.90	8.85	0.951	492
combination, analyzer, 22	0	−3.15	8.70	0.961	495
kappa prediction, analyzer, 22	0	−3.34	8.58	0.971	498
residual alkali, off, 20	0	6.13	8.72	0.792	404
kappa pred., reconciled, optimizer	0	−2.61	4.83	0.952	465

or above 8 g/L, with every parameter as shipped and the held variable, the feedback and the charge left to the learner. The results over eleven seeds are listed in Table 4.5 and drawn in Figure 4.3. The shipped temperature loop with the analyzer trim keeps all 110 slugs of the case's seed within ± 2 only because the drift's over-cooking cancels wood B's under-cooking, and every one of them is under the floor. The suggested setting, the combination on the reconciled estimate at a 22 % charge, passes at every seed, as does the H-factor on the reconciled estimate. Every structure on the raw analyzer fails because of the drift, and at 20 % 12 of 121 structure-seed runs hold kappa while none also holds the floor. The kappa prediction on the reconciled estimate passes 4 of 11, because its analyzer gate, which counts reports on one side only, admits the weak hour's under-cooked slugs into $\hat{b}_P$ when they reach the analyzer an hour after wood B's, both reading high. Turning the optimizer on with that structure takes the true residual to 5.00 g/L, because the weak hour and wood B reach the titrator five hours after the charge was chosen. The steam and alkali per tonne are reported rather than judged, because holding kappa and the floor on wood B costs more of both than the shipped settings' off-floor pulp. The suggested setting at 16 h is shown in Figure 4.4.

The estimator's gate is also why a better estimate is not always a better feedback signal. On the

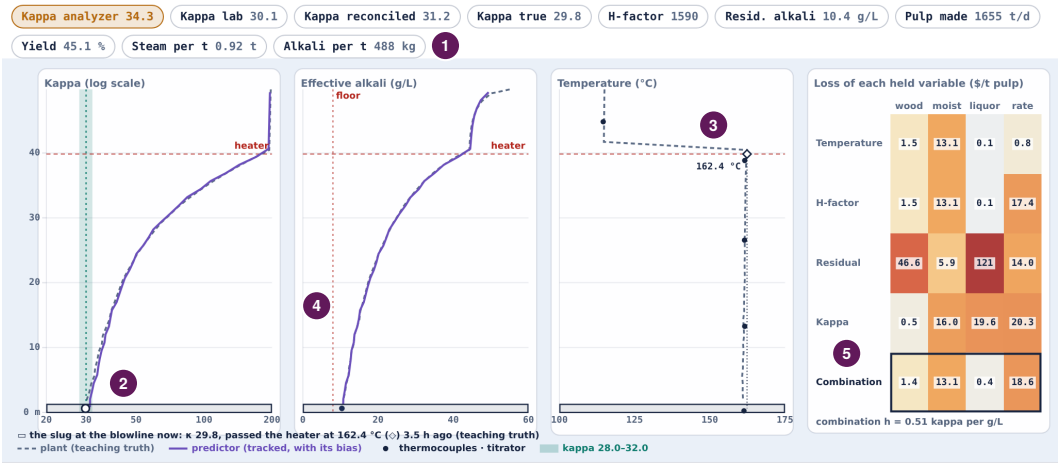

Figure 4.4 The challenge at 16 h with the suggested setting (combination on the reconciled estimate, 22 % charge), Profiles view, cropped to the readings and the profiles. (1) The readings: the analyzer reads 34.3 with its drift, the laboratory 30.1, the reconciled estimate 31.2 and the true kappa 29.8, the chips entering are predicted to gather an H-factor of 1590, the titrator reads 10.4 g/L, and 1655 t/d of pulp leaves at 45.1 % yield, 0.92 t of steam and 488 kg of alkali per tonne. (2) The slug at the blowline, which passed the heater at 162.4 °C 3.5 h earlier. (3) The heater crossing on the temperature profile. (4) The residual alkali floor on the alkali profile. (5) The loss table at the 22 % charge, which boxes the combination's row ($h = 0.51$ kappa per g/L).

same day without the drift, at 20 %, the H-factor held kappa within ±2 on the raw analyzer at 11 of 11 seeds and on the reconciled estimate at none. The reconciled estimate holds back a large new offset for 75 minutes and then moves a tenth of the way per report, which is right for a predictor's bias and slow as the input of a PI loop that already waits five hours for the vessel.

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

4.6.8 Code Verification

The ledgers close to rounding. Over the suggested challenge, 180 slugs that entered and left close their lignin to 1.4×10^{-15}, their alkali to 1.1×10^{-15} and their carbohydrate to 2.0×10^{-15} kg per kg of wood, the delivered charge equals 22 % times the liquor's strength for every slug including the nine of the weak hour, and the vessel's lignin closes to -2.3×10^{-13} of 1230.3 t. Through the hang-up's drop the vessel's lignin closes to -3.4×10^{-13} of 563.2 t. Each slug's H-factor agrees with an independent five-point Gauss-Legendre quadrature of Vroom's rate along its recorded temperature to 1.3×10^{-15}. A slug that crossed the heater after a step in case 1 reached the blowline 4.000 000 h later, and each of 32 slugs at 1725 t/d took 3.4783 h from heater to blowline. Doubling the slugs to 80 moves the suggested challenge's blowline kappa by at most 0.088 % with the noise off, halving the step to 5 s by 0.052 %, and neither changes a verdict at the seeds checked. The loss table, recomputed by an independent bisection and re-optimization, agrees to $0.01/t. A measuring script reproduces on the engine all 249 numbers quoted in the simulator's texts, catalog entries and change record.

4.6.9 Agreement With Published Data

Vroom's rate as coded is 1.01 % above the form with the constant 43.19 at 100, 150, 160, 170 and 171 °C, because of the rounded constant, and one hour at 171 °C gives 1014 against the 1000 quoted [8]. The bulk delignification rate at 35 g/L is 0.0158, 0.0380 and 0.0880 per minute at 150, 160 and 170 °C, inside the ranges 0.0117–0.0371, 0.0160–0.0631 and 0.0244–0.1125 measured for

six woods [11]. Its activation energy is Vroom's 134 kJ/mol, while the measured bulk values are 57–91 kJ/mol, so the model's rate rises 5.58 times from 150 to 170 °C against 2.09–3.22 measured, and its residual pool shares that activation energy. The kinetics are therefore teaching values, close to the published rates at the design temperature and steeper than them away from it. The trends of Section 4.4, constant kappa at constant H-factor and lower yield in a hotter cook, agree in direction with the laboratory cooks [8].

4.6.10 Runtime

At one simulated hour per second in a browser window of 1366 by 768 pixels, three runs of the challenge each took 24.0 s, the longest case's 36 h took 36.0–36.1 s, the largest of 12951 control calls took 10.2 ms, and the optimizer took 4.3–9.6 ms over 15 calls. These are bounds on one machine.

4.7 Technology in Practice

What Measures It

The cook is graded at the blowline, far in time from the moves that decide it. A 2003 report on continuous digesters put the dead time between the quality variables and manipulated inputs such as temperature at 4 to 8 h, and named the variability of the natural feedstock as the other difficulty [18]. Kappa is measured either by an online analyzer or in the laboratory [19]. The chapter keeps both, the analyzer every 15 min with a 20 min delay and a bias that can drift, and the laboratory every 4 h an hour late, in Equation (4.4).

> **TECHNOLOGY Sensor: The Online Kappa Analyzer**
>
> The kappa number is defined by a laboratory test: the volume of 0.1 N permanganate solution consumed by one gram of dry pulp under standard conditions, about seven times the lignin left on the pulp in percent [8]. An online analyzer automates the measurement. It collects its pulp samples itself and judges them by their optical properties through a calibration curve. Mills also keep models that infer kappa from process measurements, to warn the operators when the estimate and the analyzer disagree widely and to watch the analyzer's calibration [19]. A drifting calibration shows on a trend as the analyzer walking slowly away from the laboratory, the fault of case 5.

What Moves It

The cook has two handles. The alkali charge is set by the white liquor ratioed to the chip meter, and the temperature by the circulation heater, whose outlet the heater loop holds with the steam valve under the heat-balance feedforward of Equation (4.6). In a plant control system such loops are often nested or cascaded, the setpoint of one taken from the output of another [20]. Figure 4.5 draws the loops of Section 4.5 on field devices.

> **TECHNOLOGY Final Element: The Steam Valve of the Heater**
>
> A continuous digester is heated by steam injected directly or by circulating the cooking liquor through a heat exchanger [16], the arrangement of this chapter. Excess friction in the stem of a pneumatic control valve makes it move in jerky steps and ignore small reversals of its signal, a fault called stiction that is severe in poorly maintained valves without positioners, and an up-and-down test of the output in manual mode reveals it. An air-to-open valve fails closed on loss of air, and the safer failure for the process decides which action is chosen [21]. The model's valve has no friction, and the heater answers it through one lag of 0.1 h.

Where the Controller Runs

A digester's loops run in a distributed control system, which places a supervisory level over field controllers, with cycle times from milliseconds to minutes [20]. In the 2003 project a fundamental digester model ran beside a mill's system on an application node that exchanged data with it over an open communication protocol [18]. The application's controller executes once a simulated minute

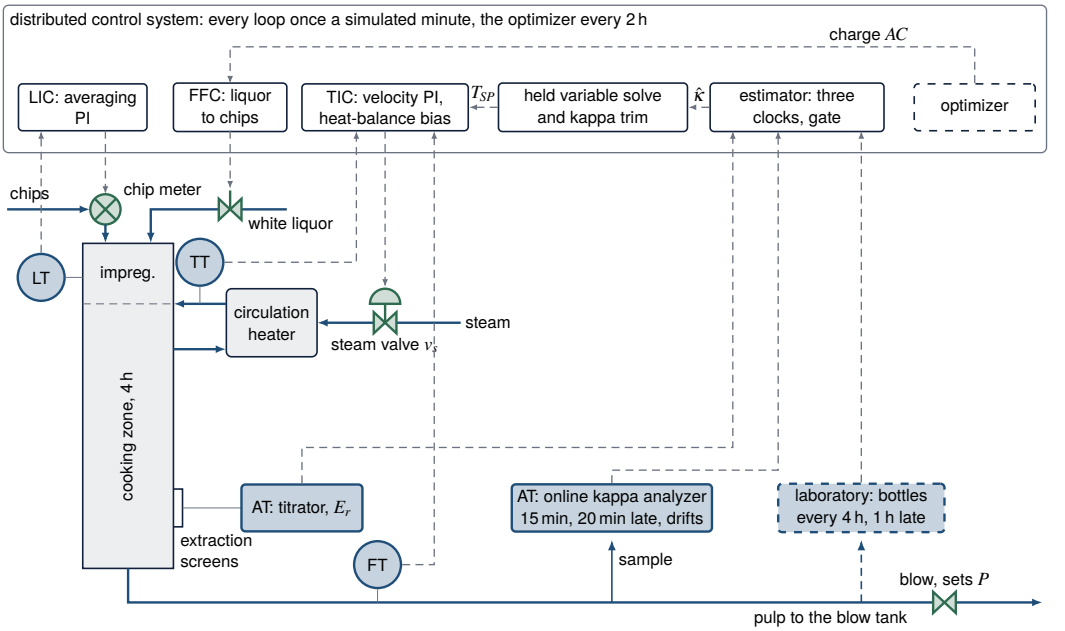

Figure 4.5 The loops of the digester on plant hardware. The heater loop (TIC) holds the circulation heater's outlet (TT) with the steam valve, the level loop (LIC) moves the chip meter, and the white liquor follows the chips at the charge AC. The held variable's solve writes the heater setpoint T_{SP} from the estimator, which reconciles the titrator at the extraction screens, the online kappa analyzer and the laboratory. Notice that every quality instrument sits at or below the screens, so each reads chips that passed the heater hours earlier.

on a plant integrated every 10 s, short beside a slug that passes every 7.5 min and a cook of 4 h. The velocity-form PI loops are the base layer. The held variable's solve, the estimator and the kappa trim sit above them as the supervisory layer that writes T_{SP}, and the economic optimizer runs above that every 2 h.

What Is Modernizing It

Three developments bear on this vessel. The first is predictive control. The 2003 project tested a predictive controller on step-response models against five PI loops, both sampled every 10 min, with the fundamental model standing in for the digester, and the PI loops needed a large tuning filter to avoid sustained oscillations [18]. The second is the adaptive soft sensor. A 2013 review reports a recursive regression model adapted to frequent process changes in a pulp digester, whose estimates fed back to a PID controller improved control [22]. The third is performance monitoring. The same project built a toolbox that judges loops from routine operating data, detecting oscillations and estimating dead time, with methods kept simple enough to supervise hundreds of loops [18]. Case 6, where a larger trim gain rings, is the kind of loop such a tool would flag.

> **IN PRACTICE** **Timing a Change to the Chip Front**
>
> A swing between softwood and hardwood is one of the transitions that digester operators and engineers must manage. One simple policy ramps each operating variable to its new value after a delay that roughly synchronizes the change with the moving front of the chips [18]. Mills align their records the same way. One study compared a chip sample taken at the chip bin with the kappa of pulp blown three and a half hours later [19]. The H-factor feedforward here times its moves by the slug entering the cooking zone.

4.8 Limitations

The vessel is a teaching design whose kinetics are fitted to one design point. Their bulk rate falls within published ranges at the design temperatures, but their temperature sensitivity is Vroom's index rather than a measured one, so a temperature move of several degrees is answered more strongly than the published rates suggest. Plug flow without a wash zone, effective alkali without sulfidity, a heater that sets the chips' temperature at once and an averaged wood mix leave out the radial, chemical and chip-scale structure that fundamental digester models carry [5]. The floor and the temperature limit are declared values, and the model does not simulate the lignin redeposition the floor stands for.

> ⚠ **OPEN CHALLENGE** **A Raw Material No Instrument Reads**
>
> An industry plan of 1999 gave priority to new sensors because the main raw material is naturally variable and because process operations often prevent timely and accurate measurement. Its priorities included cost-effective online sensors for pulp and fiber properties, and control systems that keep performing without deteriorating over time [23]. A 2003 report on continuous digesters added that real-time kappa measurement along the vessel is not generally feasible [18]. In this chapter the part of wood B that cooks slower reaches no instrument but kappa. A sensor that read the reactivity of the chips before the heater would be progress.

The estimator's analyzer gate counts reports on one side of the estimate, so an hour-long upset that follows a lasting change on the same side enters the predictor bias, as the kappa prediction's 4 of 11 shows. The optimizer holds the floor on its predictor without a back-off, so the true residual sits on the floor on average. The loss table is a steady-state analysis at declared prices, and the combination's weight changes with the charge, from -0.0869 at 20 % to 0.514 at 22 %. Every run is one seed of instrument noise at a time, and the challenge's statements rest on eleven.

> ⚠ **OPEN CHALLENGE** **When Should an Estimate Believe a Change?**
>
> Soft sensors must be updated as a process changes, and a 2013 review sets out the trade-off. Recursive updates adapt gradually, can adapt excessively when a process runs in a narrow range for a while, and cannot follow abrupt changes. In one plant a model updated daily from the laboratory kept a bias after each change of operating condition [22]. The estimator of this chapter meets the same choice in its persistence gate. The gate ignores an hour-long upset but admits the weak hour that follows a new wood, so the kappa prediction passes 4 of 11 seeds. A rule that told a lasting change from a passing one would be progress.

4.9 Conclusion

This chapter studied a continuous kraft digester, a plug-flow reactor whose quality is measured hours after it is decided, on a simulator that carries the chip column as 40 Lagrangian slugs whose lignin, alkali and H-factor close to rounding. A two-degree heater step reaches the blowline 4.125 h later and intact, and a 15 % production step leaves chips past the heater that no heater setting can rescue (worst slugs $+2.46$ to $+4.49$). The H-factor feedforward restores kappa one cooking residence later, holding 29.88–30.42. On a new wood the residual alkali is the worst held variable at \$119.68/t and the combination of predicted kappa and residual alkali the best at \$7.27/t, and the reconciled estimate holds kappa at 28.68–29.51 under a drifting analyzer that drives the raw loop to 26.52. On the graded day the combination on the reconciled estimate at a 22 % charge passes 11 of 11 seeds. The kinetics are teaching values whose temperature sensitivity exceeds the measured one. A fuller model would add a wash zone, sulfidity, a back-off for the optimizer's floor and a level test on both gate channels. The simulator shows that on a long plug-flow vessel the decision is what to hold between measurements, not how hard to chase the last one.

KEY POINTS

- In a continuous digester every heater move rides down with the chips and reaches the blowline hours later with its shape intact.
- No heater setting rescues chips already past the heater when production rises, and the H-factor feedforward restores kappa one cooking residence later.
- On a new wood the residual alkali is the worst variable to hold, at $119.68/t, and the combination of predicted kappa and residual alkali the best, at $7.27/t.
- The reconciled estimate removes a drifting analyzer's bias, yet without the drift the raw analyzer is the better feedback signal, because the estimate waits out short upsets.

Where else this applies. Volume 9, Chapter 10 chooses the setpoint a fast loop holds under a quality that answers in hours, and Chapter 10 corrects a soft sensor from a delayed laboratory in the same way.

References

[1] J. H. Lee and A. K. Datta, "Nonlinear inferential control of pulp digesters," *AIChE Journal*, vol. 40, no. 1, pp. 50–64, 1994.

[2] R. D. Gudi, S. L. Shah, and M. R. Gray, "Adaptive multirate state and parameter estimation strategies with application to a bioreactor," *AIChE Journal*, vol. 41, no. 11, pp. 2451–2464, 1995.

[3] M. Morari, Y. Arkun, and G. Stephanopoulos, "Studies in the synthesis of control structures for chemical processes: Part I: Formulation of the problem. process decomposition and the classification of the control tasks. analysis of the optimizing control structures," *AIChE Journal*, vol. 26, no. 2, pp. 220–232, 1980.

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

[5] P. A. Wisnewski, F. J. Doyle, and F. Kayihan, "Fundamental continuous-pulp-digester model for simulation and control," *AIChE Journal*, vol. 43, no. 12, pp. 3175–3192, 1997.

[6] J. J. Castro and F. J. Doyle, "A pulp mill benchmark problem for control: problem description," *Journal of Process Control*, vol. 14, no. 1, pp. 17–29, 2004.

[7] W. Gilbert and B. Allison, "Control of continuous digester kappa number using generalized model predictive control," *TAPPI Journal*, vol. 23, no. 9, pp. 467–476, 2024.

[8] T. E. S. Segura, M. A. Vivian, E. A. Bonfatti Júnior, C. Sarto, F. Schmidt, and F. G. d. Silva Júnior, "Effectiveness of the H-factor for controlling eucalyptus kraft pulping," *Scientia Forestalis*, vol. 47, no. 124, pp. 791–798, 2019.

[9] R. R. Gustafson, C. A. Sleicher, W. T. McKean, and B. A. Finlayson, "Theoretical model of the kraft pulping process," *Industrial & Engineering Chemistry Process Design and Development*, vol. 22, no. 1, pp. 87–96, 1983.

[10] D. Almeida, R. B. Santos, P. W. Hart, and H. Jameel, "Hardwood pulping kinetics of bulk and residual phases," *TAPPI Journal*, vol. 14, no. 10, pp. 652–662, 2015.

[11] D. Almeida, H. Jameel, P. Hart, and R. Santos, "Hardwood kraft pulping kinetics," in *7th International Colloquium on Eucalyptus Pulp*, Vitória, Espírito Santo, Brazil, May 2015. [Online]. Available: https://www.eucalyptus.com.br/artigos/18_Diego+Almeida1.pdf

[12] R. Amirthalingam and J. H. Lee, "Subspace identification based inferential control applied to a continuous pulp digester," *Journal of Process Control*, vol. 9, no. 5, pp. 397–406, 1999.

[13] B. Gough and J. T. Kay, "Kraft continuous digester effective alkali control," *IEEE Industry Applications Magazine*, vol. 3, no. 4, pp. 48–53, 1997.

[14] N. Padhiyar, A. Gupta, A. Gautam, S. Bhartiya, F. J. Doyle, S. Dash, and S. Gaikwad, "Nonlinear inferential multi-rate control of Kappa number at multiple locations in a continuous pulp digester," *Journal of Process Control*, vol. 16, no. 10, pp. 1037–1053, 2006.

[15] R. Rantanen, T. Ahvenlampi, M. Tervaskanto, U. Kortela, and A. Korhonen, "Prediction and control of blow-line kappa number," *IFAC Proceedings Volumes*, vol. 38, no. 1, pp. 495–500, 2005.

[16] 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.

[17] V. Alstad and S. Skogestad, "Null space method for selecting optimal measurement combinations as controlled variables," *Industrial & Engineering Chemistry Research*, vol. 46, no. 3, pp. 846–853, 2007.

[18] F. J. Doyle III, "Final report: Model-based approach to soft-sensing and diagnosis for control of a continuous digester," University of Delaware, Department of Chemical Engineering, Newark, DE, Technical Report DOE-ID13882-F, Jun. 2003.

[19] F. M. Correia, J. V. H. d'Angelo, G. M. Almeida, and S. A. Mingoti, "Predicting kappa number in a kraft pulp continuous digester: A comparison of forecasting methods," *Brazilian Journal of Chemical Engineering*, vol. 35, no. 3, pp. 1081–1094, 2018.

[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] 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] M. Kano and K. Fujiwara, "Virtual sensing technology in process industries: Trends and challenges revealed by recent industrial applications," *Journal of Chemical Engineering of Japan*, vol. 46, no. 1, pp. 1–17, 2013.

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

4.10 Exercises

4.1 ⌕ Concept check · 2 min

In case 4, 60 % of the chips become wood B at 1 h. Holding the H-factor, the true kappa settles near 32.7. Holding the residual alkali at the extraction screens, it climbs to 34.8 by 14 h. Why does the held variable measured in the cook itself do worse?

(a) The titrator reports 10 min late, so the heater oscillates around the residual alkali target

(b) Wood B cooks more slowly, which lowers the H-factor, and only the residual alkali reveals it

(c) The residual alkali target of 10.1 g/L drives the heater to its lower limit of 140 °C

(d) Wood B's acid groups consume alkali, the residual falls, and the heater cools to put it back

4.2 ⌕ Concept check · 1 min

In case 2 the production rises 15 % at 1 h. The chapter finds that no heater setting, with feedforward, analyzer feedback or both, keeps every slug within ±2 of kappa 30 through the step. What limits every structure?

(a) The chips already in the cooking zone have crossed the heater, and no later heater move reaches them

(b) The steam valve reaches full opening when the heat load of the chips rises by 15 %

(c) The H-factor stops predicting kappa once the production rate has changed

(d) The analyzer reports 20 min late, which is too slow for the kappa trim

4.3 ⌕ Concept check · 2 min

On the challenge day without the analyzer drift, at a 20 % charge, the H-factor with the trim on the raw analyzer holds every slug within ±2 at 11 of 11 seeds, and the same structure on the reconciled estimate at none. Why is the better estimate the worse feedback signal here?

(a) The reconciled estimate carries the laboratory's noise, and the trim amplifies it

(b) The estimator holds back a large new offset for more than an hour, then moves a tenth of the way per report

(c) The laboratory's delay of an hour is added to the loop's dead time and makes the trim unstable

(d) With the H-factor held, Reconciled moves the bias into the solve and the trim rests at zero

4.4 Theory and calculation · 6 min

Vroom's rate is $k_V(T) = \exp(43.2 - 16113/T)$ per hour with T in kelvin. The cooking zone holds $0.8M$ with $M = 664.9\,\mathrm{t}$ of dry wood, and the column descends at $W = P/(24 \times 0.47)$ tonnes of wood per hour for a production P in tonnes of pulp per day.

(a) Find k_V at 100 °C and 160 °C, the fractional rise of k_V per kelvin at 160 °C, and the factor by which a step from 160 to 162 °C multiplies the rate. *Checkpoint:* 1.019 and 403.6 per hour, 8.59 % per kelvin, and 1.186.

(b) Find the cooking residence at 1500 and 1725 t/d, and the share of the design cooking H-factor that a slug gathers at the higher rate and a fixed temperature. *Checkpoint:* 4.000 h and 3.478 h, a share of 0.8696.

(c) From the slope of part (a), estimate the heater rise that restores the cooking H-factor at 1725 t/d. The H-factor feedforward of case 3 computes 161.53 °C. Say why it asks for less. *Checkpoint:* About 1.63 °C. The shorter cook cools the chips less.

4.5 Theory and calculation · 6 min

The heater's bias is the heat-balance feedforward of Equation (4.6), with $c_w = 1.35\,\mathrm{kJ/(kg\,K)}$, $c_l = 4.0\,\mathrm{kJ/(L\,K)}$, $\lambda = 2000\,\mathrm{kJ/kg}$, $\dot{m}_{s,max} = 90\,\mathrm{t/h}$, $T_{imp} = 115\,^\circ\mathrm{C}$ and $M = 664.9\,\mathrm{t}$.

(a) Find $T_{in,0}$ and the valve u_b at the design flow, 132.98 t/h, and $T_{SP} = 160\,^\circ\mathrm{C}$. *Checkpoint:* 114.72 °C and 58.0 %.

(b) Find u_b after the production step of case 3, at 152.93 t/h and $T_{SP} = 161.53\,^\circ\mathrm{C}$. *Checkpoint:* About 68.9 %.

(c) Without the feedforward, how far must the heater outlet fall below its setpoint before the proportional term alone, at $K_c = 2.5\,\%$ per °C, opens the valve by the extra amount that part (b) needs? *Checkpoint:* About 4.4 °C.

4.6 Theory and calculation · 7 min

In case 5 the analyzer's bias drifts upward at 0.3 kappa per hour from 1 h. Laboratory bottles are drawn at 0, 4, 8, … hours and reported an hour later, each paired with the analyzer sample drawn at the same instant, and the analyzer bias estimate is updated with $\lambda_A = 0.8$ as in Equation (4.9).

(a) Using Equation (1.8), find the mean age of a held analyzer result and of a held laboratory result. *Checkpoint:* 0.458 h (27.5 min) and 3 h.

(b) Show that, with the noise left out, the estimate lags a ramp that rises by r between bottles by $r(1-\lambda_A)/\lambda_A$ once the update has settled. Find the lag and the estimate after the report at 21 h. *Checkpoint:* A lag of 0.30 and an estimate of 5.40, against a bias of 5.70 in the bottle drawn at 20 h.

(c) Find the true bias at 24 h and the part of it the estimator has not yet seen. Compare with the chapter's 5.56 and 1.34. *Checkpoint:* 6.90, and about 1.50 unseen.

4.7 App challenge · 7 min

Load scenario 1, Open loop: a temperature step and the four-hour silence, with the default seed. Press Run at a speed of 0.1 h/s, press Pause when the clock is between 0.5 and 0.6 h, and raise Cooking temperature setpoint from 160 to 162 °C. Set the speed to 1 h/s and press Run. The case pauses at 10 h. Export the CSV file.

(a) Find the first row after your step at which `truth_kappa_blowline` is below 29, and the time from your step to that row. *Checkpoint:* Between 4.0 and 4.2 h, for example the row at 4.633 h for a step at 0.50 h.

(b) Find the first row at which `kappa_analyzer` is below 28, and the instant its sample was drawn, column `kappa_analyzer_sample_h`. *Checkpoint:* Reported at 5.083 h from a sample drawn at 4.75 h.

(c) Read `truth_kappa_blowline` at 10 h, and find the change in kappa per kelvin of cooking temperature. *Checkpoint:* 27.28, within 0.05, so about -1.36 kappa per kelvin.

4.8 App challenge · 9 min

Load scenario 4, The wood changes: which variable to hold, with the default seed and Held variable as loaded, H-factor. Run at 1 h/s. The case pauses at 12 h. Press Run again and pause at 14 h. Export the CSV file and read `cook_T_sp_C` and `truth_kappa_blowline` in the row at 14 h. Set Held variable to Residual alkali, press Reset and repeat. Set Held variable to Kappa prediction, press Reset and repeat, and also read the titrator, column `EA_residual_gL`.

(a) Report the setpoint and the true kappa for each held variable. *Checkpoint:* H-factor 160.00 °C and 32.69. Residual alkali 158.35 °C and 34.76. Kappa prediction 160.73 °C and 31.80. Within 0.05 °C and 0.05.

(b) Report the titrator reading under the kappa prediction, and say what it means for the alkali charge. *Checkpoint:* 7.68 g/L, within 0.05 g/L, under the 8 g/L floor.

(c) Explain the ranking with the loss table, and why none of the three reaches 30. *Checkpoint:* \$119.68, \$29.45 and \$9.90 per tonne.

4.9 App challenge · 8 min

Load scenario 5, The analyzer drifts and the laboratory corrects it, with the default seed and the settings as loaded: Held variable Kappa prediction and Kappa feedback Analyzer. Run to 24 h at 1 h/s and export the CSV file.

(a) In the row at 24 h, read four columns: the true kappa `truth_kappa_blowline`, the analyzer reading `kappa_analyzer`, the true bias `truth_analyzer_bias` and the estimated bias `bias_analyzer`. *Checkpoint:* 26.52, 33.81, 6.90 and 5.56, within 0.02.

(b) The kappa trim has integral action, yet the analyzer stays more than three units above 30. Estimate the steady error of the trim on a ramp, with $K_c = 0.15\,^{\circ}\mathrm{C}$ per kappa, $\tau_I = 3\,\mathrm{h}$ and the gain of Exercise 4.7. *Checkpoint:* About 4.4 kappa.

(c) Set Kappa feedback to Reconciled, press Reset and run to 24 h again. Find the lowest and highest `truth_kappa_blowline` from 12 to 24 h. *Checkpoint:* 28.68 and 29.51, within 0.02.

4.10 Application area · 4 min

In case 7 the economic optimizer lowers the charge from 20.00 to 18.83 %. From 7 to 12 h the alkali falls from 425.0 to 396.3 kg per tonne of pulp and the steam rises from 0.814 to 0.817 t/t, and the yield of the slug blown at 12 h rises from 46.97 to 47.60 %. The declared prices are \$70 per tonne of dry wood, \$150 per tonne of alkali as NaOH and \$20 per tonne of steam.

(a) Find the saving per tonne of pulp and per day at 1500 t/d. *Checkpoint:* About \$6.22 per tonne, about \$9300 per day.

(b) The true residual alkali then runs at 7.72 to 8.39 g/L, under the 8 g/L floor in 49 % of the steps. Estimate the back-off on the predicted floor that keeps this range above 8 g/L, and its cost, taking $\partial E/\partial \ln AC = 31\,\mathrm{g/L}$. *Checkpoint:* About 0.3 g/L, about 3.6 kg of alkali and \$0.54 per tonne.

4.11 Application area · 3 min

With the challenge's shipped settings, temperature held, the trim on the raw analyzer and a 20 % charge, all 110 slugs blown from 12 to 24 h at the default seed are within ± 2 of kappa 30, and all 110 are under the 8 g/L floor, the lowest at 5.6 g/L.

(a) Why does kappa look good? *Checkpoint:* The drift's over-cooking cancels wood B's under-cooking.

(b) Which instrument shows the problem during the day, and why does a mill keep a residual alkali floor at all? *Checkpoint:* The titrator at the extraction screens, below 8 g/L from about 7.7 h. A floor keeps dissolved lignin from redepositing on the fibers.

4.12 Review the agent · 5 min

An AI assistant is asked for a setting that passes the challenge. It replies: "Hold the Kappa prediction with Kappa feedback on Reconciled and a 22 % charge. It uses the best estimate in the app, and it passes: 110 of 110 slugs within ± 2, worst -1.9 at 15.1 h, lowest residual 9.5 g/L. The loss table's combination only adds a heuristic weight on the residual alkali and cannot do better." Find the flaw, and say how you would check it.

CHAPTER 5

Forecasting Frozen-In Stress in a Float Glass Line Simulator

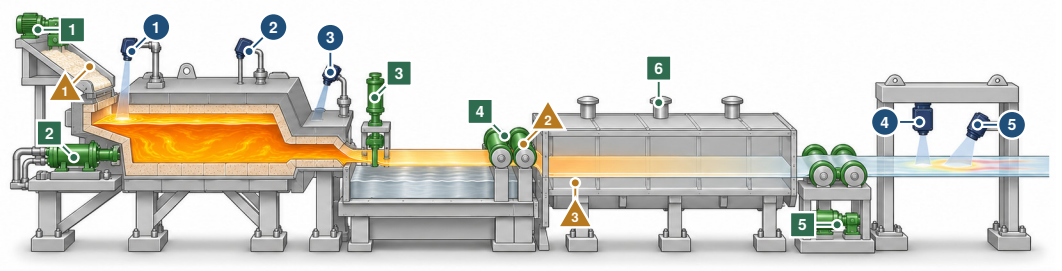

The controller moves 1 batch charger 2 furnace firing 3 tweel, the gate that meters the glass 4 top rolls that set the stretch 5 lehr drive that sets the ribbon speed 6 heating and cooling of the six lehr zones

The controller reads 1 melt level 2 crown temperature 3 canal temperature 4 thickness gauge 5 polarimeter that reads the stress kept

What disturbs it 1 more cullet in the charge 2 top rolls losing their grip 3 a lehr zone failing

Notice how far the two cold-end gauges stand from the top rolls and the first lehr zones, where the thickness and the stress are fixed. Illustration, not to scale.

IN PLAIN TERMS Glass Whose Quality Is Set Before It Is Measured

The main idea. A float line pours melted glass onto a bath of liquid tin and pulls it into a long, flat ribbon. Rollers stretch the ribbon to the right thickness, and a long oven then cools it slowly so that little strain stays locked inside.

The challenge. The gauges that check thickness and strain stand at the cold end, reading glass shaped about a quarter of an hour before. The controller must act on what it predicts, and say plainly when no oven setting can save the glass.

Why it matters. Without a forecast, the line learns of a fault only after many meters of glass are already spoiled.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history segment (Section 5.2), the fact on furnace size (Section 5.3), the try-its on the thickness gauge and the failed lehr zone and the fact on the failed zone (Section 5.6) and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 5.6), Technology in Practice (Section 5.7), Limitations (Section 5.8) and the open challenge

The control question A float line settles the thickness of its glass on the tin and the stress the glass keeps in the first meters of the lehr, minutes before the gauge and the polarimeter at the cold end can read either. How can the line control qualities that are fixed before they can be measured, and when must its controller report that no setting of the lehr can meet the limit?

Learning objectives

- Compute the delays of the line as distances over the lehr speed, and show how they move with the product.
- Derive the stretch an order needs from the forming relation, and explain why top rolls that lose their grip thicken the glass.
- Locate where along the lehr the stress is kept, and explain why the supervisor searches a grid of offsets instead of closing a loop on its forecast.
- Compare a loop on the delayed gauge, the same loop as a Smith predictor and a loop on the mass-balance thickness by the glass each rejects.
- Assess a predictive controller that plans the annealing zones from the job card against the forecast supervisor.

Prerequisites Sections 1.4.4 (delay that moves with flow), 1.4.3 (Smith predictor), 1.4.8 (soft sensor and bias update) and 1.6.2 (velocity-form PID), and Chapter 3 for a quality read late. Sections 1.4.7 and 1.6.3 are needed only for the predictive controller.

The app https://apmonitor.com/apps/float-glass/

See also Volume 9, Chapter 5 for releases planned from a published schedule through a transport delay, and Volume 9, Chapter 4 for a predictive controller made offset free by a disturbance estimate.

5.1 Introduction

A product that carries its own history can be judged only once that history is complete, and float glass shows this plainly. Its thickness is controlled by stretching the ribbon while it is still viscous on molten tin [1]. Its residual stress comes from the temperature differences through its thickness while it cools [2, 3], and it affects how well the glass cuts [3]. The stress is read photoelastically on glass that has cooled [4], and in-line systems have been built to read the thickness in the hot zone of the lehr rather than at its end [5].

The float process pours molten glass from a furnace onto a bath of molten tin, where it spreads and is drawn off as a continuous ribbon [6]. In the line of this chapter a furnace holding 1.27 kt of glass, 2.1 days of its 600 t/day pull, meters the glass under a refractory gate called the tweel. The lehr draws the ribbon 1.611 times as fast as the top rolls release it, which thins it from the 6.0 mm it would settle to on the tin to 4.00 mm, and the ribbon crosses a 130 m lehr at 11.26 m/min. The thickness gauge, 171 m past the top rolls, reads glass formed 15.2 min earlier. The ribbon passes its annealing point 18 m into the lehr and its strain point 21 m further on, and the polarimeter stands another 100 m downstream, 8.9 min of travel.

The plant differs from a loop with a fixed dead time in three ways. Its delays are distances divided by the lehr speed, so they move with the product, from 15.2 min for 4 mm glass to 22.8 min for 6 mm glass at the same pull, a transport delay defined in distance as Section 1.4.4 defines it in volume. Its forming is viscous, which makes width and thickness one mass balance set by two speeds, so top rolls that slip and narrow the ribbon also thicken it. Its stress is a record of the thermal path each piece of glass has followed, fixed over the 21 m in which the glass stops flowing, so a lehr trimmed on the polarimeter acts on glass that was annealed nine minutes before. Upstream, the furnace answers its firing and its pull over hours, its melt with a time constant of 96 min.

A reading at the cold end is easily taken for a report on the glass being made. A gauge at 4.00 mm and a polarimeter at 1.0 MPa describe glass formed a quarter of an hour ago and annealed nine minutes ago, and neither says anything about the glass in the lehr now. The class of problem is broad.

Wherever a quality is measured downstream of the place where it is made, a dead-time compensator of the Smith predictor family can take a known delay out of the loop [7], and an inferential controller can act on an estimate built from faster secondary measurements [8], the two remedies of Sections 1.4.3 and 1.4.8. Steel strip in a continuous annealing furnace likewise takes its temperature along a path whose speed the line sets (Volume 16, Chapter 9). The furnace upstream is the slow half of the same line, and model-based control has been applied to glass melting as well [9]. However, a transfer function with a constant dead time represents neither a delay that changes with the product nor a quality set by a thermal path. It cannot show a stress forecast running ahead of the polarimeter, or a supervisor reporting that no setting of the lehr can anneal the ribbon it has.

This chapter charts where the thickness and the stress of float glass are made and where they are measured. The simulator follows every 3 m of ribbon from the tin to the cut with its own temperature and stress through its thickness, infers the thickness at the width camera from the mass relation, and anneals the ribbon from a forecast made on an estimator's copy of every parcel. Section 5.2 traces the float process, its annealing and its control. Section 5.3 sets out the model of the whole line, from a lumped furnace through viscous forming to a ribbon whose stress follows a Maxwell element at every depth, specified fully enough to rebuild, with the truth, the measurements and the estimates logged apart. Section 5.4 finds where along the lehr the stress is kept and how it depends on the annealing offset. Section 5.5 describes the controller: the PI loops and their feedforwards, the Smith predictor and the inferred thickness, the thermal-history estimator, the supervisor and a model predictive controller built on the estimator. Section 5.6 measures nine cases and a graded product change, which separate what the cold end can report from what the glass in the lehr will keep, and in Section 5.6.7 it sets the predictive controller against the supervisor on a scheduled job change. Section 5.7 sets the loops on plant hardware, and Sections 5.8 and 5.9 state the limitations and the conclusions.

5.2 Float Glass, Its Annealing and Its Control

The float process was described by its inventor in a 1969 review lecture, which set it against the methods used to make flat glass before and at the time of its invention, traced its development and treated three of the main problems solved along the way [6]. Before it, flat glass was drawn from the melt as sheet or cast and then ground and polished as plate [6]. Glass poured on tin spreads toward an equilibrium thickness, and thinner glass has to be stretched [1, 6]. The control of thickness by stretching the ribbon was analyzed in 1977 as a one-dimensional problem for a free ribbon without edge constraints. Hot ribbons were simulated to show their width and thickness attenuating at different rates, and the calculated width was compared with in-plant measurements [1]. Top rolls that hold the edges, as on the simulated line, decide how each stretch is shared between the two.

> ⧖ HISTORY **1950s: Glass Floated on Tin**
>
> Nearly all flat glass in the United States is made by the float process, developed in England in the 1950s and licensed in the United States in 1962. It replaced the plate glass process, whose grinding and polishing could take as much as 3 million Btu per ton of glass. Float glass leaves the bath fire-polished, close to the quality of ground and polished plate, and larger, more efficient furnaces lowered its energy further. The molten glass flows under a tweel and over a lip onto a pool of molten tin about 160 to 190 feet long, kept under nitrogen so that the tin does not oxidize, and takes on the flatness of the tin. The process is so dominant that the terms float glass and flat glass are now used interchangeably [10].

Annealing is older than the float process. How the stress in glass relaxes at the temperatures of its annealing range, and how a cooling schedule follows from the rate of that relaxation, was set out in 1920 [11]. The three-constant viscosity equation that the simulator uses was published in 1925 [12]. Viscoelastic analysis came later. A theory built on measured relaxation data computed the stress history of a glass plate cooled symmetrically from both surfaces in 1965 [2]. Compared with

measured stress distributions in 1969, it agreed for glass cooled from well above the strain point once its numerical formulation was changed, but not for glass quenched from lower temperatures [13]. A model of structural relaxation that carries the memory of the thermal history followed in 1971 [14].

The lehr became a control problem of its own on float lines. An analytical model of a float glass annealing lehr was published in 1977, with radiation between the moving ribbon and its enclosure and gas convection in a flue system. It illustrated how modeling could improve the design and operation of the lehr, and it proposed a linear algorithm for on-line temperature control, checked by computer experiments [15]. The stress a lehr leaves can be measured photoelastically, because glass is birefringent under stress [3, 4], and the residual stress of commercial float glass of several thicknesses has been quantified this way because it affects cutting quality [3]. At the hot end, simplified first-principles models of the combustion chambers of glass furnaces have been developed [16], and model-based control of a glass melting furnace has been reported [9]. Energy is part of the motivation, since the options for decarbonizing this energy-intensive industry have been reviewed [17].

5.3 The Line From Charger to Cut

At its design point the line makes 600 t/day of 4.00 mm glass 3.70 m wide, which leaves 3.40 m of net ribbon once 0.15 m is trimmed from each edge. Positions are measured along the ribbon from the spout. The furnace, the working end and canal, the tweel, the float bath with its top rolls, the six lehr zones and the cold end are shown in Figure 5.1 with every symbol of the equations.

> **DID YOU KNOW? Furnaces Sized in Hundreds of Tons a Day**
>
> A 2002 profile of the United States glass industry found that every float line melts its glass in a side-port regenerative furnace, typically of 500 to 1000 tons per day, large because both the volume and the quality of flat glass demand it [10]. The simulated line pulls 600 t/day.

Table 5.1 classifies the variables. The 33 lumped states are the batch blanket and the glass inventory, the crown, melt and working-end temperatures, four actuator lags, the lehr and top-roll speeds, the flow and thickness leaving the top-roll zone, six zone temperatures and fourteen energy and mass ledgers. About 59 parcels of ribbon, held in 72 slots, carry the rest of the state, each with its position, mass, thickness, width, the order in force when it was formed and ten node values.

5.3.1 Furnace and Tweel

The charger lays a blanket of batch and cullet on the melt, and the blanket melts with a time constant that shortens with temperature and with cullet, as shown in Equation (5.1).

$$\begin{aligned} \frac{\mathrm{d}M_b}{\mathrm{d}t} &= F_{ch} - F_m, \qquad F_m = \frac{M_b}{\tau_b}, \qquad \frac{\mathrm{d}M_g}{\mathrm{d}t} = Y F_m - G \\ \tau_b &= \tau_B \, \mathrm{e}^{-(T_m - T_B)/\theta_b} \, \max(0.2,\ 1 - k_{cul} f_c) \\ Y &= f_c + (1 - f_c) Y_b, \qquad L = \frac{M_g - \rho_m A_f h_0}{\rho_m A_f} \end{aligned} \tag{5.1}$$

Here $F_{ch} = x_{ch} F_{ch,max}$ is the charge, $T_B = 1450\,^{\circ}\mathrm{C}$, and Y is the glass per kilogram of mix, less than one because raw batch loses 17 % of its mass as carbonate gas while cullet loses nothing. The level L is the height of the glass above the depth h_0 in a melter of area A_f. At the design point the charger lays 694 t/day of mix for 600 t/day of glass, and the 17.4 t blanket melts with $\tau_b = 36$ min at the 1446 °C melt. A millimeter of level is 1.06 t of glass. The level enters the tweel flow through the head alone, and the pull loop removes even that, so under automatic control the level integrates the difference between melting and pull behind a melting lag.

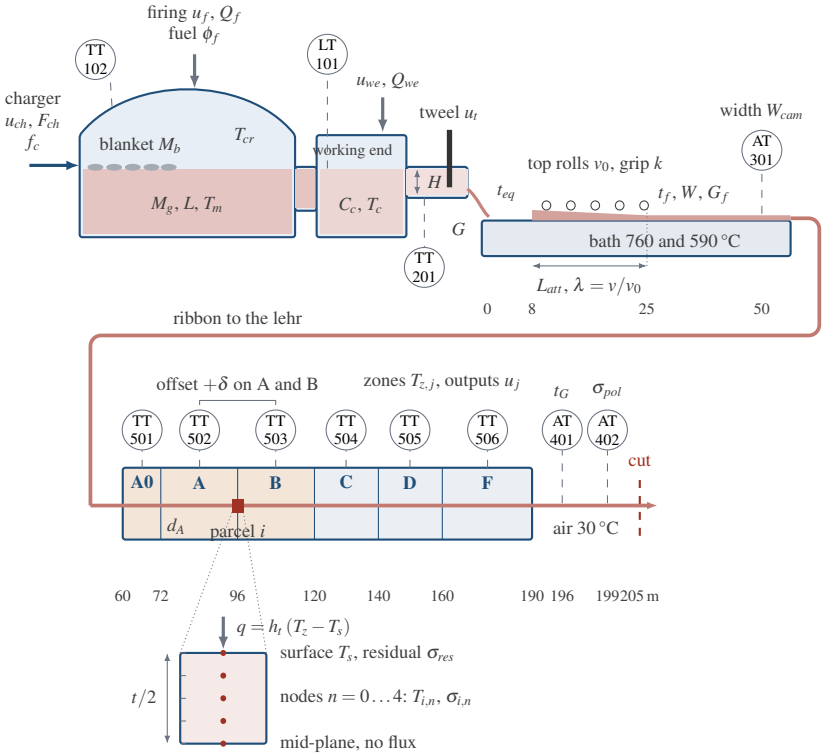

Figure 5.1 Float line as modeled, drawn out of scale with positions in meters from the spout, the furnace and the float bath above and the lehr and the cold end below. The charger lays batch and cullet (fraction f_c) on the melt as a blanket M_b that melts into the glass inventory M_g, whose level L LT-101 reads. The firing u_f heats the crown space (T_{cr}, TT-102), the working end (C_c, T_c) conditions the glass on its way to the canal (TT-201), and the tweel at opening u_t passes the pull G under the head H. On the tin the glass spreads toward t_{eq} and is stretched over L_{att} between the top rolls (v_0, grip k) and the lehr (v), leaving the forming zone at 25 m with thickness t_f, width W and flow G_f. The width camera AT-301 stands at 50 m, the six lehr zones A0 to F hold $T_{z,j}$ with outputs u_j, zones A and B carry the supervisor's offset δ, and zone A can fail with its damper stuck at d_A. The gauge AT-401, the polarimeter AT-402 and the cut follow. Each 3 m parcel i carries five nodes of temperature and stress over its half thickness $t/2$, heated or cooled at its surface by the flux q.

Three heat stores, the combustion space under the crown, the melt and the working end, are coupled as shown in Equation (5.2).

$$\begin{aligned} C_{cr}\frac{\mathrm{d}T_{cr}}{\mathrm{d}t} &= \eta Q_f - UA_{cm}(T_{cr}-T_m) - UA_{cr}(T_{cr}-T_a) \\ C_m\frac{\mathrm{d}T_m}{\mathrm{d}t} &= UA_{cm}(T_{cr}-T_m) - UA_m(T_m-T_a) \\ &\quad - F_m\left[Yc_p(T_m-T_a) + (1-f_c)\,q_r\right] \\ C_c\frac{\mathrm{d}T_c}{\mathrm{d}t} &= Gc_p(T_m-T_c) - UA_c(T_c-T_a) + Q_{we} \end{aligned} \tag{5.2}$$

Here $Q_f = Q_{f,max}\min(x_f, \phi_f)$ is the firing capped by the fuel available ϕ_f and η the share of the firing that stays in the furnace. The term q_r is the reaction and gas heat of raw batch, and $T_a = 25\,°\mathrm{C}$. The working end gives $Q_{we} = x_{we}Q_{wh}$ with its burners for $x_{we} \geq 0$ and $x_{we}Q_{wc}$ with its cooling otherwise. The four actuator positions x follow their commands through first-order lags of 1 min for the charger,

Table 5.1 Roles, Design Values and Ranges of the Line's Variables

Symbol	Role	Design value	Range
L	CV, melt level (LIC-101)	0 mm SP	SP ±3
T_{cr}	CV, crown (TIC-102)	1580 °C	SP 1500–1640
T_c	CV, canal (TIC-201)	1100 °C	SP 1060–1140
$\hat{G}$	CV, pull estimate (FIC-301)	600 t/day	SP 400–850
t	CV, thickness (AIC-401)	4.00 mm	SP 3–8
W_{cam}	CV, width (AIC-301)	3.70 m	SP 3.4–4.0
$T_{z,j}$	CVs, six zones (TIC-501 to 506)	Table 5.2	SP 20–620
$\hat{\sigma}$	CV, stress forecast (supervisor)	1.00 MPa	SP 0.4–2
u_{ch}, u_f	MVs, charger, firing	57.4, 74.4 %	0–100
u_{we}	MV, working end	−46.5 %	±100
u_t	MV, tweel opening	50.0 %	0–100
λ	MV, attenuation v/v_0	1.611	0.8–3.0
v	MV, lehr speed	11.26 m/min	3–25
u_j	MVs, zone heat and cooling	Table 5.2	±100
δ	MV, annealing offset on A, B	6.7 K	0–40
δ_j	MVs, MPC offsets on A0, A, B	—	0–δ_{max}
f_c	DV, cullet in the charge	20 %	0–60
ϕ_f	DV, fuel available	100 %	40–100
k	DV, top-roll grip (nominal 0.70)	100 %	30–100
w, s	DVs, lip wear, t_{eq} shift	0 %	−5–8, ±5
d_A	DV, zone A damper stuck	0 %	0–100
b_A, e_h	DVs, TC bias, estimator error	0	±20, ±30
M_b, M_g	states, blanket, glass	17.4 t, 1.27 kt	—
T_m	state, melt temperature	1446 °C	—
v_0, G_f, t_f	states, forming	6.99 m/min, —	—
parcels	x, m, t, W, t_{SP}, T_n, σ_n	about 59	72 slots

the fuel valves and the working end and 0.5 min for the tweel. At 600 t/day with 20 % cullet the firing is 44.6 MW, 6.43 GJ per metric ton of glass. Of the 26.8 MW that stays in the furnace, 20.8 MW crosses from the crown to the melt and 6.0 MW leaves through the crown and walls. The melt spends 11.8 MW heating the new glass, 4.0 MW on reactions and gas and 5.0 MW on floor and side losses. The glass gives up 2.9 MW between the 1446 °C melt and the 1100 °C canal, 1.7 MW as losses and 1.2 MW to the working-end cooling. With their neighbors held, the melt, the crown space and the working end settle with time constants of 96, 20 and 40 min.

Glass flows under the tweel in proportion to its opening and to the head over the lip, and inversely to its viscosity at the canal, as shown in Equation (5.3).

$$\begin{aligned} G &= G_{ref}\,\frac{x_t}{u_{ref}}\,\frac{H}{H_{ref}}\,10^{\ell(T_{c,ref})-\ell(T_c)}\,(1+w) \\ \ell(T) &= A+\frac{B}{T-T_0} \end{aligned} \tag{5.3}$$

with $H = H_{ref} + L$, $\ell = \log_{10}\eta$ for η in Pa s and w the lip wear the model does not know. The Vogel-Fulcher-Tammann (VFT) form [12] takes $A = -3.13$, $B = 5030$ K and $T_0 = 211$ °C, fitted to typical soda-lime float fixed points. They are 10 Pa s at 1450 °C, 10^3 at 1020 °C, $10^{6.6}$ at the 725 °C softening point, 10^{12} at the 548 °C annealing point and $10^{13.5}$ at the 511 °C strain point. The fit meets them within 0.14 in $\log_{10}\eta$ apart from the annealing point, which it misses by 0.20. At the 1100 °C canal a degree of cooling cuts the flow by 1.46 %, so the conditioning temperature is a flow control, and one percent of opening at the 50 % reference passes 12 t/day.

5.3.2 Viscous Forming on Tin

On the tin the glass would settle to its equilibrium thickness t_{eq}. The lehr draws the ribbon away at v, faster than the top rolls, which grip its edges and turn at v_0, and the grip k decides how much of the stretch $\lambda = v/v_0$ goes into thickness rather than width, as shown in Equation (5.4).

$$\begin{aligned} t(\lambda) &= t_{eq}(1+s)\,\lambda^{-(1+k)/2}, \qquad W = \frac{G_f}{\rho\, t_f\, v} \\ \tau_F \frac{\mathrm{d}t_f}{\mathrm{d}t} &= t(\lambda) - t_f, \qquad \tau_F \frac{\mathrm{d}G_f}{\mathrm{d}t} = G - G_f \\ \tau_F &= L_{att}/v \end{aligned} \tag{5.4}$$

Here s is the fractional shift of the equilibrium thickness with the bath atmosphere, and k runs from 0 for a free sheet that narrows and thins alike to 1 for edges held fast. The thickness and the flow reaching the end of the 17 m attenuation zone lag their inputs by the zone's length over the speed, 1.51 min at 11.26 m/min. With $k = 0.70$ the thickness goes as $\lambda^{-0.85}$, and at a given flow and top-roll speed the width goes as $\lambda^{-0.15}$, so the glass spreads to 3.97 m before it is stretched to 3.70 m and 4.00 mm at $\lambda = 1.611$ with the top rolls at 6.99 m/min. The lehr and top-roll drives follow their commands with a 6 s lag. The glass leaves the top rolls at 860 °C with a viscosity of 4.2×10^4 Pa s, so this is the forming of a liquid, in which the stretch sets the thickness and the mass relation $G_f = \rho W t v$ sets the width.

🕮 **Pause.** Try Exercise 5.3 before reading on.

5.3.3 The Ribbon as Parcels

Past the forming zone the ribbon is carried as parcels $\Delta S = 3$ m long that move at the lehr speed, the tracking of material of Section 1.4.4. The mass crossing the end of the forming zone accumulates until the ribbon has advanced 3 m, and a parcel is then born with that mass m_i, the thickness t_f being formed and the width $W_i = m_i/(\rho\, t_f\, \Delta S)$, so the mass relation holds parcel by parcel. It starts at a uniform 860 °C with no stress. Each parcel also keeps the order t_{SP} in force when it was formed, against which it is later judged. Along its path it sees the bath at 760 °C from 20 to 40 m and at 590 °C from 40 to 58 m, both with a convective coefficient of 15 W/(m^2 K), and a lift-out at 570 °C and 10 W/(m^2 K) up to 60 m. The six lehr zones follow to 190 m, then cold-end air at 30 °C and 12 W/(m^2 K), and a parcel that straddles two of these surroundings sees each in proportion to its overlap.

5.3.4 Heat and Stress Through the Thickness

Both faces of the ribbon are cooled alike, so half of the thickness is modeled, on five nodes from the mid-plane to the surface, as shown in Equation (5.5).

$$\rho c_p \frac{\partial T}{\partial t} = k_g \frac{\partial^2 T}{\partial z^2}, \qquad -k_g \left.\frac{\partial T}{\partial z}\right|_s = h_t\,(T_s - T_z) \tag{5.5}$$

There is no flux at the mid-plane, and $h_t = h_c + \varepsilon\sigma_{SB}(T_z^2 + T_s^2)(T_z + T_s)$, in kelvin, linearizes the radiation about the surface temperature of the previous step. The grid is node-centered with half cells at both ends, and each 1 s step is backward Euler with one tridiagonal solve, which is stable at any step. The heat a parcel gives up is handed to the zones it overlaps on the next step, so the energy ledger of the lehr closes.

The glass is a plate free to expand in its own plane, with a Maxwell element at every depth, as shown in Equation (5.6).

$$\begin{aligned} \frac{\mathrm{d}\sigma_n}{\mathrm{d}t} &= E'\alpha\left(\left\langle\frac{\mathrm{d}T}{\mathrm{d}t}\right\rangle - \frac{\mathrm{d}T_n}{\mathrm{d}t}\right) + \left\langle\frac{\sigma}{\tau}\right\rangle - \frac{\sigma_n}{\tau(T_n)} \\ \tau(T) &= \frac{\eta(T)}{G_r} \end{aligned} \tag{5.6}$$

Table 5.2 The Six Lehr Zones. Output Is the Case 1 Mean From 60 to 120 min, Negative for Cooling

Zone	Span (m)	SP (°C)	Q_H (kW)	Q_C (kW)	h_c (W/(m^2 K))	C_z (MJ/K)	Output (%)
A0	60–72	550	300	500	10	4	−52
A	72–96	$505+\delta$	300	800	10	4	−45
B	96–120	$455+\delta$	250	1000	10	4	−35
C	120–140	360	150	1300	12	4	−44
D	140–160	230	0	1800	25	5	−46
F	160–190	45	0	2800	60	6	−63

Here $E' = E/(1-\nu)$, α is the expansion of the solid glass, G_r the relaxation modulus and $\langle\cdot\rangle$ the mean over the thickness, and the mean terms keep $\langle\sigma\rangle = 0$, the force balance of a free plate [2]. Each step applies the elastic change, then the exact relaxation $\sigma_n \mathrm{e}^{-\Delta t/\tau}$ at the node's mean temperature over the step, complete above 700 °C and absent below 400 °C, and then removes the mean. The stress that would remain if the temperature equalized now is the residual of Equation (5.7).

$$\sigma_{res,n} = \sigma_n + E'\alpha\,(T_n - \langle T\rangle) \tag{5.7}$$

The rest of σ_n is temporary and vanishes with the gradient. Tension is positive and $E'\alpha = 0.831\,\mathrm{MPa/K}$, so each kelvin between the surface and the mean that is frozen in leaves 0.83 MPa. The application reports $-\sigma_{res}$ at the surface as the compression kept.

5.3.5 Lehr Zones and the Cut

Each lehr zone is one heat store, heated electrically in A0 to C and cooled with air in all six, as shown in Equation (5.8).

$$\begin{aligned} C_{z,j}\frac{\mathrm{d}T_{z,j}}{\mathrm{d}t} &= Q_{H,j}\,u_j^{+} - Q_{C,j}\,c_j \\ &\quad - UA_z(T_{z,j} - T_a) + Q_{g,j} \end{aligned} \tag{5.8}$$

with $u_j^{+} = \max(0, u_j)$ and $c_j = \max(0, -u_j)$, while $Q_{g,j}$ counts the heat the parcels gave zone j over the previous step. When zone A fails its heaters trip and its damper sticks at least d_A open, so its heat is zero and $c = \max(c, d_A)$. The zones are listed in Table 5.2. At 205 m each parcel is cut, and the scorecard tallies its net area $\Delta S(W_i - 2e)$ as accepted or rejected on the inspection record. The inspection record holds each instrument's reading of the parcel itself, its own width at the camera, thickness at the gauge and stress at the polarimeter plus that instrument's noise, and the verdict, written at the polarimeter, rejects a thickness reading more than 0.09 mm from the order, a net width reading under 3.31 m or a stress reading over 1.45 MPa. The scorecard checks the truth separately, against a thickness within 0.10 mm of the parcel's order, a net width of at least 3.30 m and a residual surface compression of at most 1.5 MPa.

5.3.6 Positions, Initial State and Numerics

Table 5.3 lists the parameters, and the guard bands are 0.01 mm, 0.01 m and 0.05 MPa. The attenuation zone runs from 8 to 25 m, the width camera stands at 50 m, the bath ends at 58 m, the lehr runs from 60 to 190 m, and the gauge, the polarimeter and the cut stand at 196, 199 and 205 m. Every case starts from the furnace at rest for its pull, crown and canal setpoints and cullet, and from the forming at rest, with the lehr speed from the mass relation and $v_0 = v/\lambda$. The zones start at their setpoints plus the supervisor's first offset, and the ribbon is filled every 3 m from 28 to 202 m with the steady path of one parcel. The zone outputs start where they balance the heat the ribbon gives them, and every loop is seeded by tracking. The lumped states are integrated by classical fourth-order Runge-Kutta with a 1 s step, and the parcels are stepped after each one. Halving the step moves the residual of the stable ribbon by 0.03 %, nine nodes instead of five move it by 1.19 %, and 1.5 m parcels instead of 3 m by 0.28 %, all within the application's own checks.

Table 5.3 Parameters of the Line. The Class Says Whether the Parameter Panel Offers the Value (vis.), Keeps It Out of Sight (hid.) or Cannot Change It (fixed)

Symbol	Meaning	Value	Class
A_f, h_0	melt area, depth at zero level	450 m^2, 1.2 m	hid.
ρ_m	molten glass density	2350 kg/m^3	hid.
C_m, C_{cr}, C_c	heat stores	960, 190, 24 MJ/K	vis.
UA_{cm}	crown to melt	155 kW/K	vis.
UA_{cr}, UA_m, UA_c	losses	3.85, 3.50, 1.60 kW/K	hid.
η, $Q_{f,max}$	heat kept, firing capacity	0.60, 60 MW	vis.
q_r, Y_b	batch heat, glass per kg	0.62 MJ/kg, 0.83	vis.
τ_B	melting time at T_B, no cullet	40 min	vis.
θ_b, k_{cul}	melting scale, cullet effect	60 K, 0.8	hid.
$F_{ch,max}$	charger capacity	1210 t/day	vis.
τ_{ch}, τ_f	charger, fuel valve lags	1, 1 min	hid.
Q_{wh}, Q_{wc}	working-end heat, cooling	1.5, 2.5 MW	vis.
τ_{we}, τ_t	working-end, tweel lags	1, 0.5 min	hid.
G_{ref}	flow at the reference	600 t/day	vis.
u_{ref}, $T_{c,ref}$	reference opening, canal	50 %, 1100 °C	hid.
H_{ref}	head over the lip	0.20 m	vis.
ρ	glass density	2500 kg/m^3	vis.
t_{eq}, k	equilibrium thickness, grip	6.0 mm, 0.70	vis.
L_{att}, T_{form}	attenuation zone, exit	17 m, 860 °C	hid.
τ_d, e	drive lag, edge trim	6 s, 0.15 m	hid.
c_p	glass heat capacity	1200 J/(kg K)	vis.
k_g, ε	conductivity, emissivity	1.5 W/(m K), 0.70	vis.
E, ν	modulus, Poisson ratio	72 GPa, 0.22	hid.
α, G_r	expansion, relaxation	9 ppm/K, 30 GPa	vis.
A, B, T_0	VFT constants	−3.13, 5030 K, 211 °C	vis.
C_z	zone store, A0 to C	4 MJ/K	vis.
UA_z	zone shell loss	100 W/K	hid.
τ_L, τ_{TC}	level, furnace TC lags	1, 1 min	hid.
τ_Z, τ_{cam}, τ_{ga}	zone TC, camera, gauge	18, 12, 12 s	hid.
Δt, W_{min}	tolerance, net width	0.10 mm, 3.30 m	vis.
σ_{lim}	residual stress limit	1.5 MPa	vis.
—	guard bands t, W, σ	0.01, 0.01, 0.05	hid.
σ_{tmp}	temporary tension alarm	10 MPa	hid.
T_a, ΔS	ambient, parcel length	25 °C, 3 m	fixed
—	nodes, parcel slots	5, 72	fixed

5.3.7 Assumptions

The property values, capacities, limits and correlations are illustrative calibrations chosen for teaching, not those of a particular plant or glass. The furnace is three lumped heat stores with a batch blanket and one glass inventory, and the melt's heat capacity is an active 800 t. A real melting tank has convection currents and a temperature field in three dimensions, and fining, seeds, stones and cord are absent, so a cold melt shows only as a growing blanket of unmelted batch.

Forming is quasi-steady and reduced to an equilibrium thickness and one grip, with a first-order lag over the attenuation zone. There is no thickness profile across the ribbon, no edge bulb and no asymmetry between the tin side and the top side. Glass thicker than t_{eq} would need $\lambda < 1$, which a real line does not use.

Heat moves only through the thickness, with both faces alike, constant properties, gray radiation and a conductivity that includes internal radiation. The ribbon is uniform across its width, so there is no edge cooling and no edge stress. The zones change temperature in steps at their boundaries, smoothed only by the parcel length, which the residual hardly notices while the peak temporary tension rises from 4.79 to 5.05 MPa with 1.5 m parcels.

The stress is one Maxwell element per depth with a single relaxation time. Real glass relaxes over a spectrum of times, and its structure relaxes as well [14], so the residual is a proxy for the permanent through-thickness stress, and the 1.5 MPa limit and 1.0 MPa target are teaching values rather than a

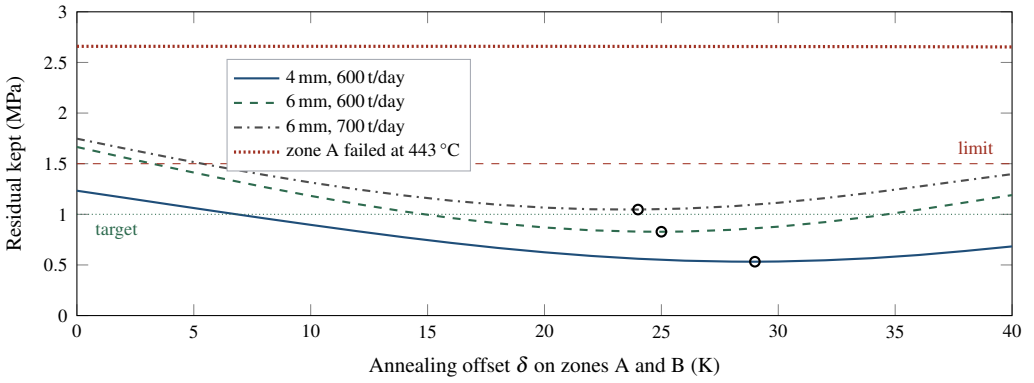

Figure 5.2 Residual surface compression kept by a steady ribbon against the annealing offset added to zones A and B. Each point is one parcel carried from the forming zone to the cut at 1 s steps through fixed zone temperatures, and the circles mark the minima. With zone A failed at 443 °C the offset acts on zone B alone.

standard.

The instruments are lags and filters with colored noise and no failure modes other than the scenario events. The polarimeter is taken to read the permanent stress, and each parcel is judged against its own order with guard bands the size of the instruments' noise. Rejected glass is not returned to the furnace as cullet.

5.4 Where the Ribbon Keeps Its Stress

In case 1 the ribbon enters the lehr at 607 °C with 3.1 K between its mid-plane and its surfaces. Averaged from 60 to 120 min, it passes the annealing point at 78.0 m, in zone A, and the strain point at 98.7 m, just inside zone B. Across those 21 m, 1.8 min of travel, the relaxation time η/G_r grows from 21 s to 24 min. The residual of Equation (5.7), the stress the glass would keep if its temperature equalized at that point, is 1.69 MPa at the lehr entrance and 1.00 MPa at 96 m, where the ribbon is at 519 °C. Over the remaining 94 m of lehr the ribbon cools to 89 °C and its residual changes by less than 0.01 MPa, so it reaches the cut with 1.01 MPa. The stress is settled within the first 36 m of the lehr, and the other 94 m only cool the glass.

🕮 **Pause.** Try Exercise 5.1 before reading on.

The forecast the supervisor makes is this residual carried forward, and its dependence on the annealing offset is the characteristic that shapes the control design, as shown in Figure 5.2 for steady ribbons annealed with each offset from the start. For 4 mm glass at 600 t/day the stress falls from 1.23 MPa on the recipe to the 1.0 MPa target at 6.8 K, close to the 6.7 K the supervisor holds on average from 60 to 120 min in case 1. It reaches a minimum of 0.532 MPa at 29 K and rises again to 0.683 MPa at 40 K. Zone A does nearly all the work, since raising zone A alone by 30 K takes the stress to 0.563 MPa, while raising zone B alone takes it only to 1.213 MPa. Past the minimum the ribbon leaves zone A still at the top of its annealing range, at 547 °C with 0.27 MPa kept at 40 K of offset, and finishes freezing in zone B, which the recipe holds 50 K below zone A, where it gains another 0.41 MPa.

Thicker glass keeps more for the same schedule, because at a given cooling rate the temperature difference through a plate grows with the square of its thickness. For 6 mm glass at 600 t/day the target needs 14.8 K, close to the 14.6 K the supervisor reaches in case 8. At 700 t/day the 6 mm ribbon cannot reach the target at all, with a minimum of 1.047 MPa at 24 K, although 5.4 K takes it under the limit. A 4 mm ribbon at 850 t/day, 42 % faster, needs 15.1 K, so even the largest pull the application allows leaves the lehr feasible. With zone A failed and held at 443 °C the curve is flat, 2.659 MPa

with no offset and 2.653 MPa with 40 K on zone B, because the ribbon passes both fixed points inside zone A.

Two consequences follow for the controller. The sensitivity of the stress to the offset changes sign at the minimum, so an integrating loop closed on the forecast would push the offset the wrong way on the far side of it, and the supervisor searches a grid of offsets instead. And the exit temperature of the lehr says nothing about the stress. On the recipe, cooling zone F from 45 to 20 °C takes the ribbon out of the lehr 19 K colder and to the cut 15 K colder with its residual unchanged, and it only raises the peak temporary tension from 4.56 to 5.05 MPa.

Pause. Try Exercise 5.2 before reading on.

5.5 Implemented Controller

The loops, the estimator and the supervisor run every 5 s on a clock kept in minutes. The controller executes on every fifth 1 s physics step, and each step then integrates the lumped states and moves, heats and stresses the ribbon. The historian logs every 20 s, and the application runs from a quarter of a minute to 20 min of plant per second, 2 min by default. The three clocks are those of Section 1.6.1.

5.5.1 Instruments and the Pull Estimate

No loop reads the truth. Each measurement is the true value passed through a first-order lag or filter, plus colored noise, and each ribbon instrument reads the parcels passing its position. The level has a 1 min lag and 0.02 mm of noise, and the crown and canal thermocouples have 1 min lags with 1 and 0.1 °C. The zone thermocouples have 18 s and 0.3 °C, and the width camera, the gauge and the polarimeter have 12 s filters with 2 mm, 2 µm and 0.03 MPa. The noise is correlated over 12 to 60 s, and the lehr speed is read with 0.02 % of noise. The filters feed the loops and the trends. The inspection record of a parcel takes its own value plus the noise at that scan, because a filtered reading mixes adjacent 3 m parcels, which differ by up to 0.20 mm when the thickness hunts.

No meter measures molten glass, so the pull is estimated from Equation (5.3) with the measured opening, level and canal temperature, multiplied by $1+\beta$ and smoothed over 1 min before the loops use it. The correction β is learned from the cold end. Each parcel reaching the gauge carries its camera width W_{cam}, its gauge thickness t_G and the lehr speed v_B at its birth, and so its own pull $G_{ce} = \rho W_{cam} t_G v_B$. This is compared with the model pull tagged to it when it was formed, and β moves 5 % of the way toward the ratio's excess over one, within ± 8 %. Three percent of lip wear is learned as 2.93 % within 90 min, and the true pull returns to 600.4 t/day.

5.5.2 PI Loops and Their Feedforwards

Every loop is a PI controller written in the framework's velocity form of Section 1.6.2, Equation (5.9).

$$
\begin{aligned}
e_k &= a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c e_k \\
S_k &= S_{k-1} + (P_k - P_{k-1}), \qquad c_k = b_k + S_k + I_{k-1} \\
u_k &= \mathrm{slew}\big(\mathrm{clamp}(c_k, u_{min}, u_{max}),\ u_{k-1}\big) \\
I_k &= I_{k-1} + \Big[\frac{K_c}{\tau_I}\, e_k + \frac{u_k - c_k}{\tau_I}\Big]\Delta t_c
\end{aligned}
\tag{5.9}
$$

Here the sign $a = \pm 1$ fixes the direction, b_k carries the loop's feedforward and $\Delta t_c = 5$ s. These are Equations (1.15) to (1.18) with $c_k = v_k$, the feedforward b_k in place of the bias u_b, no derivative term and the tracking time $T_t = \tau_I$, so the clamp and the slew limit act only on the emitted output and back-calculation on the integral alone. Every loop here has integral action, and the derivative action is switched off in every loop. In manual the integral tracks as in Section 1.6.2, with $I_k = u_{op} - b_k - S_k$, and switching a feedforward moves the integral by the difference between the old and the new feedforward, so that the output does not jump.

Table 5.4 Loops as Shipped. All Are PI in the Velocity Form With Integral Times in Minutes

Tag	CV to MV	K_c	τ_I	Output	a, b_k
LIC-101	level, charger	30 %/mm	200	0–100 %	+1, charge
TIC-102	crown, firing	1.0 %/K	20	0–ϕ_f	+1
TIC-201	canal, working end	5 %/K	15	±100 %	+1
FIC-301	$\hat{G}$, tweel	0.035 %/(t/day)	2	0–100 %	+1
AIC-401	thickness, λ	0.20 per mm	4	0.8–3.0	−1, λ_{ff}
AIC-301	width, v	1.0 min^{-1}	5	3–25 m/min	−1, v_{ff}
TIC-50j	zone j, u_j	3 %/K	5	±100 %	+1

The lehr speed slews at most 3 m/min per minute, zones D and F cannot heat.

The loops are listed in Table 5.4. Their directions follow from the process gains. A low level charges more, a cold crown fires more, a cold canal takes the working end toward heat, a low pull estimate opens the tweel and a cold zone heats, all with $a = +1$. Glass that is too thick must be stretched more and a ribbon that is too wide must run faster, so the two forming loops run with $a = -1$. The charger feedforward is the mix that melts into the estimated pull at the recipe's cullet, $100\,\hat{G}/(YF_{ch,max})$ in percent, and the pull setpoint is ramped at 60 t/day per hour.

5.5.3 Forming Model, Smith Predictor and Inferred Thickness

With the forming model on, the thickness loop takes its bias from the equilibrium thickness, $\lambda_{ff} = (t_{eq}/t_{SP})^{2/(1+k)}$ with the nominal grip, which is 1.611 for 4 mm and 1.000 for 6 mm. The width loop takes the lehr speed from the mass relation, $v_{ff} = \hat{G}/(\rho W_{SP}\hat{t})$, where $\hat{t}$ is the thickness the model says is being formed, so that a change of thickness does not spill into the width. The same model turns the thickness loop into a Smith predictor [7], as shown in Equation (5.10), the predictor of Equation (1.4) with the delay tagged to each parcel instead of fixed in time.

$$\mathrm{PV}_t = t_m + \tilde{t}, \qquad \tau_s \frac{\mathrm{d}\tilde{t}}{\mathrm{d}t} = \left(t_{meas} - t_{pred}\right) - \tilde{t} \tag{5.10}$$

Here t_m is the model's thickness forming now, computed from the measured speed ratio with the nominal grip, no bath drift and the forming lag, t_{pred} is the model thickness tagged at its birth to the glass now under the instrument, and $\tau_s = 15$ s. A transition that the model predicts is then not an error, and a constant mismatch is removed without waiting out the delay on every scan. The instrument is either the gauge or the mass-balance thickness at the camera, an inferential measurement [8] built from the model pull $\hat{G}_B$ and the lehr speed v_B tagged to the glass under the camera, $t_{cam} = \hat{G}_B(1+\beta)/(\rho W_{cam} v_B)$. The camera sees glass 2.2 min after it is formed and the gauge 15.2 min after, and the mass balance is exactly as good as the pull estimate, which is why the pull is learned from the cold end. A soft sensor corrected by a late, exact measurement is also the structure of Chapter 10, where a laboratory corrects a psychrometric estimate of powder moisture. With the forming model off both feedforwards are zero and the loop acts on the raw measurement.

5.5.4 Thermal-History Estimator

The estimator keeps its own copy of every parcel's temperature and stress, stepped at the 5 s control period with the kernel of Equations (5.5) to (5.7), the measured zone temperatures, the tracked positions, each parcel's tagged thickness $\hat{t}$ and an optional heat-transfer error e_h. When a parcel reaches the polarimeter the estimator's residual for it is compared with the reading, and a bias b moves a tenth of the way toward the difference, within ±0.5 MPa. This is the bias update of Section 1.4.8 with a weight of 0.1, which this chapter does not call λ, since λ is the attenuation here. It is aligned by construction, since each reading is compared with the forecast for the parcel it was taken on. In case 1 at 90 min the corrected forecast for the glass entering the lehr and the polarimeter differ by 0.057 MPa. With the supervisor off, the offset held at 6.5 K and a 10 °C bias on zone A's thermocouple, the true

zone runs 10.00 °C below its setpoint, and at 65 min the polarimeter reads 1.35 MPa, 0.35 MPa more than the uncorrected forecast, while b has grown to 0.36 MPa.

5.5.5 Supervisor

Once a minute the supervisor forecasts the residual of the parcel nearest the lehr entrance from the estimator's state, at the measured lehr speed, to the lehr exit with a 4 s step, for offsets on a 2.5 K grid from 0 to $\delta_{max} = 30$ K on zones A and B. A zone whose loop is in automatic and unsaturated is assumed to hold its setpoint, while a zone whose loop is saturated, or any zone in manual, is assumed to stay at its reading and takes no offset. The target offset and the verdict follow Equation (5.11).

$$\begin{aligned} \delta^{*} &= \min\{\delta : \hat{\sigma}(\delta) + b \le \sigma_{SP}\} \\ \text{infeasible} &\iff \min_{\delta} \hat{\sigma}(\delta) + b > \sigma_{lim} \end{aligned} \tag{5.11}$$

The smallest offset is interpolated between grid points and sought only below the grid's minimum, and when no offset meets the target the supervisor takes the offset of least forecast stress. The offset moves toward δ^{*} at 2 K/min. When the lehr is infeasible the supervisor also forecasts the best offset at 95 down to 60 % of the present speed, in steps of 5 %, and names the first pull that would meet the limit. The 4 s step alone makes the forecast from the lehr entrance read about 1.9 % low, which the bias absorbs. A search is used instead of a loop on the forecast because of the sign change in Figure 5.2. Once a minute the estimator also forecasts, for every parcel between 48 and 190 m, the residual it will keep under the present setpoints, and the scene draws these forecasts as ticks on the stress bars.

5.5.6 Annealing Model Predictive Controller

The lehr structure has a third setting beside the manual offset and the supervisor, a model predictive controller (MPC) [18, 19] of the kind Section 1.6.3 describes, that sets zones A0, A and B through three separate offsets δ_j above their recipe setpoints. Zones C, D and F keep their recipe, because a sustained kelvin on any of them moves the residual by less than 0.0001 MPa at the four steady products measured (4 mm at 600 and 680 t/day, 6 mm at 600 and 700 t/day), where zone A moves it by -0.030 to -0.033 MPa/K (and -0.0002 MPa/K for the 6 mm ribbon at 700 t/day, near its minimum in Figure 5.2), zone B by -0.0005 to -0.0031 and zone A0 by $+0.0020$ to $+0.0067$. A warmer A0 sends the glass into zone A hotter, so it then crosses its strain point faster and keeps slightly more. Every zone loop of Table 5.4 acts beneath the plan, the manual offset is unused as under the supervisor, and the supervisor remains the default structure of the application and of every case except case 9.

The controller's model is the estimator's own, which makes it wrong only where the estimator is. At each decision the glass between the lehr entrance at 60 m and the end of zone B at 120 m is divided into one-minute cohorts at the measured lehr speed, $N_a = \lceil 60\,\text{m}/(v \cdot 1\,\text{min}) \rceil$ of them, 6 at 11.26 m/min and at most 20 at the 3 m/min floor. One parcel per cohort is flown by the kernel of Equations (5.5) to (5.7), at the supervisor's 4 s step and at 1 s from 157 to 167 m, from its estimated temperature and stress to the cold end, under the scheduled speed and with the zones at their applied setpoints, or at their readings when saturated or in manual. Glass not yet formed starts at T_{form} with the thickness of its order, or with the largest thickness the top rolls can form at $\lambda = 0.8$ when the order asks for more. Each flight gives the residual x_a the cohort will keep and the largest surface tension z_a it meets in the lehr, which for every product from 3 to 10 mm at 4 to 20 m/min falls at zone F's entrance, 161 to 165 m. A second flight with 1 K more on zone j while the cohort is a minutes past the entrance gives the sensitivities $s_{a,j}$ and $s'_{a,j}$ at the measured speed and the thickness entering

now, and the cohorts are chained as shown in Equation (5.12).

$$\begin{aligned} x_{a+1}[k+1] &= x_a[k] + \textstyle\sum_j s_{a,j}\,\Delta\bar{T}_j[k], \quad & y_k &= x_{N_a}[k] + b \\ z_{a+1}[k+1] &= z_a[k] + \textstyle\sum_j s'_{a,j}\,\Delta\bar{T}_j[k], \quad & y'_k &= z_{N_a}[k] - b \\ \frac{\Delta T_j(s)}{\Delta\delta_j(s)} &= \frac{1+\tau_{n,j}s}{a_{2,j}s^2 + a_{1,j}s + 1}, & & j \in \{\mathrm{A0}, \mathrm{A}, \mathrm{B}\} \end{aligned} \tag{5.12}$$

Here k counts 1 min steps, $x_0[k]$ and $z_0[k]$ are the free forecasts of the cohort entering the lehr during step k, $\Delta\bar{T}_j$ is the mean over the step of zone j's deviation from its reference, y_k is the residual of the cohort leaving zone B and y'_k its peak surface tension. The bias b of the estimator enters the tension with a minus sign, because the surface tension is the temporary stress less the compression the surface keeps. On the steady ribbon the chain's summed sensitivity to zone A agrees with the kernel's steady derivative within 0.7 %.

The last line of Equation (5.12) is each zone's closed loop, identified from ±5 K setpoint steps on the engine at 4 mm and 600 t/day, each averaged about an unstepped run on the same seed so that the noise cancels. The fits give $a_{1,j} = 6.34$, 6.08 and 5.63 min, $a_{2,j} = 20.98$, 13.33 and 10.12 min^2 and $\tau_{n,j} = 4.90$, 4.96 and 4.88 min for A0, A and B, with residuals of 0.03 to 0.07 % of the step. The zero sits at the loops' 5 min integral time, the fitted $a_{2,j}$ lie within 10 % of the 22.2, 13.9 and 11.1 min^2 that the zone store, the loop gain and the cooling capacity give by hand, and at 6 mm no coefficient moves by more than 6.1 %. The fits are sampled exactly at 1 min, and a steady-state Kalman filter (Section 1.4.7) built on them and on the 18 s thermocouple lag estimates each zone's two states from its thermocouple at every 5 s scan, with a measurement variance of 0.09 K^2. The bias b is the output disturbance that makes the plan offset free [20, 21], as the fuel bias in the estimator of Volume 9, Chapter 4 does there. It is the estimator's own polarimeter correction, so each laboratory reading is aligned with the parcel it was taken on, and a reading that arrives between two decisions moves b at once and enters the next plan. In a test, one parcel made to read 0.40 MPa more and reaching the polarimeter 20 s after the 40 min decision raises the A and B offsets of the 41 min decision against the same run without it.

Every 12 control scans, on a clock rounded to the minute and first at $t = 0$, the controller solves the quadratic program of Equation (5.13) over $N = 40$ steps, twice the longest chain, with the offsets held over ten move blocks, two of 1 min, two of 2 min, one each of 3, 4, 5 and 6 min and two of 8 min.

$$\begin{aligned} \min\ & \textstyle\sum_{k=1}^{N} Q\,(y_k - \sigma_{SP})^2 + \sum_{k\in\mathscr{K}} \rho_T\,\varepsilon_k^2 \\ & + \textstyle\sum_{c=L,C} (\rho_c\,\varepsilon_c^2 + \sigma_c\,\varepsilon_c) \\ & + \textstyle\sum_{i=1}^{10} \left[S\,\|\delta_i - \delta_{i-1}\|^2 + n_i R\,\|\delta_i - \delta_0\|^2 \right] \\ \text{s.t. } & y_k \le \sigma_{SP} + \varepsilon_k, \quad y_k \le \sigma_{lim} - g_\sigma + \varepsilon_L \\ & y'_k \le \sigma_{tmp} + \varepsilon_C, \quad \varepsilon \ge 0, \quad k \in \mathscr{K} \\ & 0 \le \delta_{j,i} \le \delta_{max}, \qquad |\delta_{j,i} - \delta_{j,i-1}| \le r_\delta \cdot 1\,\mathrm{min} \end{aligned} \tag{5.13}$$

Here δ_i holds the three offsets of block i, which lasts n_i steps, δ_0 the offsets in force, $\mathscr{K}$ the steps 1, 3, …, 39 at which every band is judged, $g_\sigma = 0.05$ MPa the guard band, $\sigma_{tmp} = 10$ MPa the temporary tension of Table 5.3 that risks a check of the lehr, and δ_{max} and r_δ the supervisor's largest offset and offset rate, 30 K and 2 K/min. The weights are $Q = 200$, $\rho_T = 10^3$, $\rho_L = 10^5$, $\sigma_L = 10^3$, $\rho_C = 10^6$, $\sigma_C = 10^4$, with stresses in megapascals, and $S = 0.1$, $R = 0.002$ per square kelvin. The target term holds the cohort on σ_{SP}, the band above it, priced again with one slack per judged step, makes glass above the target costlier than glass below it, and the two horizon-wide slacks soften the reject level and the check level, so the program is never infeasible. The offsets' box is hard, as is the rate, applied at the first move and at every later block boundary. Thirty offsets and 22 slacks make 52 decision variables, and the QP's regularization is fixed at $10^{-9}(1 + 2\rho_L)$ so that the large check-level price

does not pull the offsets toward zero. Between decisions the applied setpoints slew toward the first move at r_δ, as the supervisor's offset does.

The preview is the job card, the order, pull and width changes scheduled in the case ahead of now. The forecasts convert them into the lehr speed through the mass relation's ratio to the present speed, with the pull ramp and the width loop's 3 m/min per minute slew, and give the glass formed after a change the thickness of its new order. With the preview Off the plan assumes the present order, pull ramp and speed relation for all glass to come, and a change made by the operator is unannounced in both settings. The controller accepts a plan with an optimal status whose moves, predictions, cost and slacks are all finite, whose arrays have their shapes, whose planned offsets lie in their box and whose first entries equal the move. Otherwise the historian records the reason, the supervisor takes over from where the zones stand, zone A's offset becoming its δ while A0 and B ramp back within the rate, and the next valid plan resumes control. Twelve injected solver faults are all rejected, a forced fallback and resumption each move a setpoint by at most the rate over one scan, and no fallback occurs in any case measured below. In manual the MPC idles.

The controller reads only a packet of the estimator's parcels, their tracked positions and tagged thicknesses, the measured speed and zone temperatures, the zones' saturation flags, the applied offsets, the recipe, the target, the limit, the guard band, the check level, the forming range, the offset bounds and the schedule. A source check and a replay test enforce the separation, and they catch 21 of 21 copies of the controller doctored to read a plant or parameter field and 4 of 4 behavioral mutants. In Chrome a decision takes 2.22 to 2.82 ms on average on four cases (4.2 ms at the 95th percentile, 8.0 ms at most), 0.28 % of a core at the default rate of a minute of plant per second and 5.6 % at the top rate of twenty, and a frame guard drops a decision rather than a frame at high speed.

5.5.7 Plant, Control and Estimator Models

Section 5.3 is the simulation model. The control model is the forming model of Equation (5.4) with the nominal grip and equilibrium thickness, the mass relation, and the tweel equation with its learned correction. The estimator's model is the thermal kernel run on measured inputs, the same in structure as the plant and wrong only through the instruments, the tagged thickness and the error e_h. No model is identified inside the application, and the shipped tunings are not derived from one. The MPC holds the estimator's kernel linearized into Equation (5.12) at each decision and the zone loops identified offline at the design point, so a change of the glass properties reaches the estimator and the plan alike, while a change of the zone stores, the grip or the lip wear is model error for both.

5.6 Cases and Their Measured Outcomes

The cases and the challenge appear in Table 5.5 in the order the application lists them. Each has a completion rule, and the challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt (Section 1.6.4). Each case begins from one reference configuration, with the forming model, the mass-balance thickness, the cold-end learning, the charger feedforward and the supervisor on (the MPC in case 9), and every event comes at 20 min unless stated. Case 1 was taken up in Section 5.4. Every value quoted below was measured on the engine, run headless at its default seed. Signals that move with the noise are quoted as extremes of the true value or as averages over the stated window, and the application's 52 checks pass.

5.6.1 Two Clocks: Pull, Fuel and Cullet

Raising the pull target from 600 to 680 t/day at 60 t/day per hour opens the tweel from 50.0 to 56.8 % and speeds the lehr to 12.78 m/min, and the thickness and width never leave their tolerance, with 10 077 m^2 accepted and none rejected in 240 min. With the crown held by its loop, more heat can reach the melt only if the melt cools, and it settles 12.6 °C colder at steady state, reaching 1434.9 °C at 240 min and still falling, while the blanket grows from 17.4 to 23.6 t. The charger feedforward adds the new charge at once, but the batch takes 36 min or more to melt, and the true level falls 0.24 mm, lowest at 69 min, or 0.57 mm at 49 min with a ramp of 400 t/day per hour. Downstream the ribbon

Table 5.5 Case Sequence and What Each Case Teaches

#	Scenario and event	Teaching point
1	Stable ribbon	where the stress is decided
2	Pull 600 to 680 t/day, ramped	two clocks
3	Fuel capped at 70 %	the working end absorbs it
4	Cullet 20 to 40 %	a saving after an upset
5	Bath drift +4 %, gauge	a late measurement
6	Top-roll grip to 60 %	necking thickens
7	Zone A fails, repaired at 100 min	an honest infeasibility
8	Order 4 to 6 mm	a coordinated move
9	Job change, MPC	a plan over the transit
10	300 min product change	graded challenge

Table 5.6 Headline Numbers of the Cases

Case	Measured outcome
1, lehr	548 °C at 78.0 m, 511 °C at 98.7 m
1, stress	1.00 MPa kept by 96 m, 1.01 MPa at the cut
1, furnace	44.6 MW, 6.43 GJ/t, blanket 17.4 t
2, level	true dip 0.24 mm at 69 min
2, furnace	melt 12.6 °C colder at steady state
2, lehr	exit 89 to 101 °C, offset 6.5 to 8.8 K
3, working end	cooling 46.5 to 40.8 %, canal 1100.1 °C
3, loop frozen	canal −11.3 °C, tweel 50.0 to 59.1 %
4, level	+1.02 mm at 37 min, −0.65 mm at 75 min
4, steady	6.43 to 6.18 GJ/t, −3.85 %
5, first reading	gauge 36.1 min, mass balance 22.8 min
5, rejected	734 gauge, 212 mass, 1827 m^2 no model
6, width	camera 3.512 m at 24.5 min, λ to 1.768
6, rejected	579 m^2, 1131 m^2 with the gauge
7, forecast	infeasible at 27.0 min, best 1.96 MPa
7, polarimeter	over 1.5 MPa at 33.5 min, peak 2.72 MPa
7, cut	3121 m^2 rejected, worst accepted 1.46 MPa
8, forming	λ at 0.80 for 0.8 min, t_f peak 6.111 mm
8, rejected	118 m^2, 1037 m^2 without the model

spends 12 % less time in every zone and leaves the lehr at 101 °C instead of 89, and the supervisor raises its offset from 6.5 K at the change to 8.8 K at 240 min to keep 1.0 MPa. The furnace answers in hours and the lehr in minutes.

Capping the firing at 70 % while the furnace needs 74.4 % cools the crown by 23.6 °C and the melt by 15.0 °C by 240 min, and grows the blanket to 22.3 t. The working end absorbs all of it. It cuts its cooling from 46.5 to 40.8 %, holds the canal at 1100.1 °C, and nothing on the ribbon moves. With the conditioning loop's gain set to zero, so that the working end holds its output, the canal falls 11.3 °C, the glass at the tweel is 18 % more viscous, and the pull loop opens the tweel from 50.0 to 59.1 % to keep 599 t/day. The fault still shows where it began, in a firing pinned at its limit and a blanket 4.9 t heavier.

Raising the cullet from 20 to 40 % of the charge makes the charger feedforward cut the mix by 3.8 % at once, since cullet loses no gas. The blanket already on the melt then melts in 29 min instead of 36, and in 26 min once the melt has warmed, so it releases glass that nothing asked for. The true level rises 1.02 mm by 37 min and falls 0.65 mm below its setpoint at 75 min before it settles. At steady state the firing falls from 44.6 to 42.9 MW, 3.85 %, or from 6.43 to 6.18 GJ/t, the blanket shrinks from 17.4 to 12.1 t, and by 240 min the melt runs 6.3 °C warmer. The saving arrives only after the inventories in front of it have moved.

5.6.2 A Thickness Measured Late

In case 5 the bath atmosphere changes so that the glass settles 4 % thicker on the tin, and the loop measures thickness with the gauge. The formed thickness passes 4.05 mm at 20.6 min and the gauge reads it at 36.1 min, 15.5 min later, of which 15.2 min is travel and the rest the gauge's filter. The glass formed from 21.3 to 41.2 min lies more than the 0.09 mm the inspection accepts from its order, and 734 m^2 is rejected. The mass-balance thickness passes 4.05 mm at 22.8 min, the glass out of tolerance is formed only from 21.3 to 26.9 min, and the rejected area falls to 212 m^2. The gauge loop works at all because the forming model makes it a Smith predictor, whose robustness to a delay the model gets right is that of Section 1.4.3. With the model off the same tuning hunts on the delayed gauge. After its first correction the formed thickness swings between 3.85 and 4.13 mm up to 120 min with a period of about 54 min, each swing about three quarters of the one before, and 1827 m^2 is rejected in 120 min.

> **TRY IT IN THE APP** **5 · Thickness-Gauge Delay**
>
> Press Run and watch the thick glass travel toward the gauge, which reads it at 36.1 min, and 734 m^2 is rejected. Press Reset, set Thickness measured by to Mass balance before 20 min and run two hours. The mass-balance thickness passes 4.05 mm at 22.8 min, and the rejected area falls to 212 m^2.

5.6.3 Necking

In case 6 the top rolls lose 40 % of their grip. At the same stretch a grip of 0.42 would form 4.28 mm glass, and the formed thickness reaches 4.235 mm at 23.4 min as the ribbon narrows, because the same pull at the same speed is spread over less width. The camera reads 3.512 m at 24.5 min, and the width loop slows the lehr from 11.26 to 10.53 m/min and has the camera above 3.68 m again by 26.7 min. The thickness loop raises the stretch from 1.611 to 1.768, against the 1.770 that Equation (5.4) gives at $k = 0.42$, since with less grip more of each stretch narrows the ribbon instead of thinning it. The cut rejects 579 m^2, all of it for thickness and 97 m^2 of it narrow as well, and 1131 m^2 when the loop measures with the gauge.

5.6.4 A Lehr Zone That Cannot Be Made Up For

Case 7 fails zone A. Its heaters trip and its damper sticks wide open, and it falls from 511 to 443 °C by the repair at 100 min, with its loop at full heat from 26.3 min. The ribbon then crosses from its annealing point to its strain point in about 6.5 m, from 75 to 81 m at 70 min, instead of 21 m, and both points now lie inside zone A. The forecast does not see the failure at first, because it assumes that a zone whose loop is unsaturated will reach its setpoint, and it stays near 1.0 MPa until zone A's loop saturates. The next forecast, at 27.0 min, finds 1.96 MPa at best against the 1.5 MPa limit. The supervisor declares the lehr infeasible and ramps the offset to its 30 K maximum, reached at 38.4 min. For its first 3 min of infeasibility it names 360 t/day, 60 % of the speed, as a pull that would meet the limit, and from 30.0 min, with zone A at 472 °C and falling, it names none. The polarimeter crosses the limit at 33.5 min, 6.5 min after the forecast and 13.5 min after the failure, and peaks at 2.72 MPa at 105.4 min. The state at 60 min is shown in Figure 5.3.

> **TRY IT IN THE APP** **7 · Lehr-Zone Failure**
>
> Load the case and press Run on the Lehr view. Watch zone A fall below its setpoint, then the supervisor card. It declares the lehr infeasible at 27.0 min, with 1.96 MPa at best, while the polarimeter on the stress trend crosses the 1.5 MPa limit only at 33.5 min.

The cold end rejects 3121 m^2 in all, every square meter for stress, and it accepts no parcel over the limit, the worst accepted keeping 1.46 MPa. The supervisor finds the lehr feasible again at 105.0 min, and good glass is cut from 115.3 min. Within its offsets and its pulls down to 60 % of the speed,

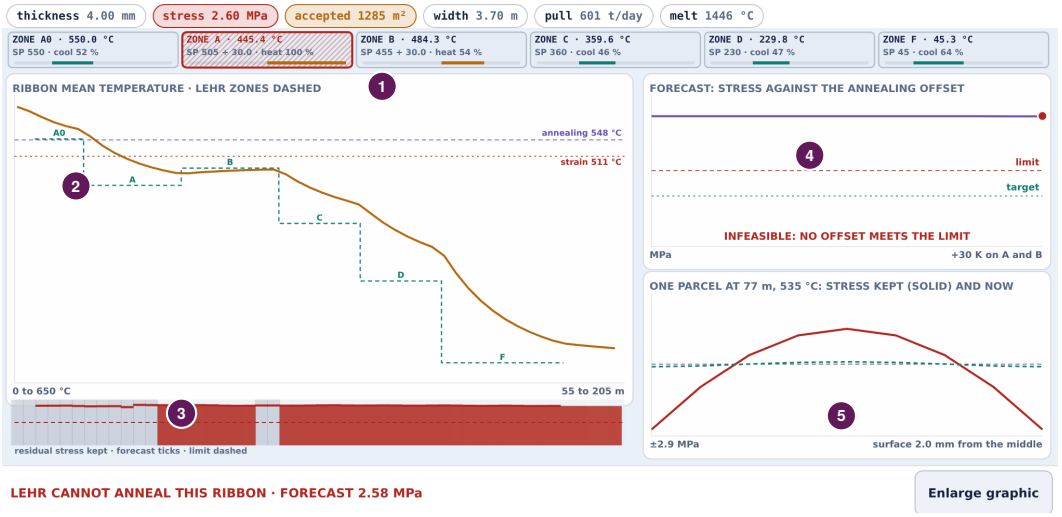

Figure 5.3 The float glass application in case 7, the lehr-zone failure, at $t = 60.00$ min, 40 min after zone A lost its heaters and its damper stuck open, cropped to the plant panel of the Lehr view. The header, the control panel with the supervisor card and the trends are cut. (1) The zone row shows zone A hatched red at 445.4 °C against its setpoint of 505 plus the 30.0 K offset with its loop at 100 % heat, and zone B heating at 54 % at 484.3 °C, while the badges read 4.00 mm, 2.60 MPa, 1285 m^2 accepted, 3.70 m and 601 t/day. (2) The ribbon's mean temperature falls through the annealing and strain points inside zone A and then rides zone B. (3) Beneath it the stress bars are red above the dashed limit wherever the glass has frozen and gray, with red forecast ticks, where it is still above 480 °C. (4) The forecast against the offset is flat near 2.6 MPa, above the limit and marked infeasible, and the status line says the lehr cannot anneal this ribbon, with a forecast of 2.58 MPa. (5) The parcel at 77 m and 535 °C would keep the solid parabola if its temperature equalized now, 2.66 MPa of compression at its surfaces and 1.44 MPa of tension in its middle by the application's own measurement, which the panel draws without printing.

nothing could save the glass annealed while zone A was down, and the forecast said so 6.5 min before the polarimeter could.

5.6.5 Changing the Order

Case 8 changes the order from 4 to 6 mm at 600 t/day. The attenuation feedforward falls to 1.000, since 6 mm is the equilibrium thickness, but the proportional term acts on the same step in the setpoint, and the attenuation sits at its 0.80 limit from 20.0 to 20.8 min. The lehr slows toward 7.50 m/min at its limit of 3 m/min per minute and undershoots to 6.81 m/min at 24.3 min. The formed thickness overshoots to 6.111 mm at 23.8 min, and the camera width swings between 3.34 and 3.99 m, staying outside 3.70 ± 0.05 m until 39.0 min. The feedforwards still move both speeds together, and 118 m^2 is rejected, 89 m^2 of it for thickness, 87 m^2 for width and 20 m^2 for stress, the three overlapping. With the forming model off the formed net width stays below 3.30 m from 20.3 to 58.7 min and 1037 m^2 is rejected. The gauge now reads glass formed 22.8 min earlier and the lehr takes 17.3 min, and the supervisor raises its offset from 6.5 to 14.6 K to keep 1.0 MPa, where an offset held at 6.3 K leaves the 6 mm glass with 1.36 MPa.

5.6.6 The Production Transition

The challenge runs 300 min. At 20 min the order changes from 4 mm at 600 t/day to 6 mm at 700 t/day, and the line starts with the forming model off, the gauge measuring thickness and the supervisor off

Table 5.7 Product-Change Challenge Over 300 min. Worst Is the Highest True Residual Among Accepted Parcels, and Off-Spec Counts Accepted Parcels Outside a Limit

Configuration	Acc. (m^2)	Rej. (m^2)	Worst (MPa)	Off spec	Pass
as shipped	1367	8820	1.245	0	no
suggested[a]	8745	118	1.361	0	yes
model, mass balance, $\delta = 0$	1377	7486	1.501	1	no
same, $\delta = 5$ K	2553	6309	1.521	12	no
same, $\delta = 7.5$ K	8275	587	1.422	0	yes
same, $\delta = 10$ K	8755	108	1.325	0	yes
same, $\delta = 20$ K	8755	108	1.074	0	yes
model, gauge, supervisor	8641	222	1.358	0	yes
no model, mass, supervisor	7921	968	1.061	0	no
no model, gauge, supervisor	1468	8718	1.207	0	no

[a] forming model, mass-balance thickness and supervisor on

with no offset. The target is 8000 m^2 accepted with no accepted parcel outside ± 0.10 mm, 3.30 m or 1.5 MPa. The configurations are compared in Table 5.7 and drawn in Figure 5.4. As shipped the thickness loop hunts on the delayed gauge and 1367 m^2 is accepted against 8820 m^2 rejected, all of it glass formed before the order change. The forming model and the mass balance alone accept 1377 m^2, because 6 mm glass at 700 t/day keeps 1.75 MPa on the recipe. The suggested configuration, the forming model, the mass-balance thickness and the supervisor, accepts 8745 m^2 and rejects 118 m^2, its worst accepted parcels at 1.36 MPa, 0.091 mm from the order and 3.319 m wide, for 837 GJ of fuel. The shipped run falls 6633 m^2 short of the target and fails, while the suggested one clears it by 745 m^2 with no accepted parcel off specification and passes. The supervisor ends at 22.5 K, near the minimum of Figure 5.2, since the target cannot be met at this pull.

A fixed offset also works, and the smallest on the supervisor's grid that passes is 7.5 K, with 8275 m^2. At 10, 12.5, 15 and 20 K it accepts 8755 m^2, one parcel more than the supervisor, while 5 K passes 12 parcels whose true stress reaches 1.52 MPa, and no offset passes one at 1.501 MPa, because the 0.05 MPa guard band is less than twice the polarimeter's 0.03 MPa noise. With the forming model off and the supervisor on, the mass balance accepts 7921 m^2, short of the target, and the gauge 1468 m^2. With the thickness hunting on the gauge, adjacent parcels differ by up to 0.20 mm, yet no parcel outside the tolerance passes, because each parcel's record differs from its own thickness by the gauge noise alone, at most 0.006 mm in this run. The application replays one noise sequence. Run with 20 other sequences the suggested configuration accepts 8695 to 8755 m^2 and passes in every one.

5.6.7 Predictive Control of the Annealing Zones

Case 9 announces on the job card that at 30 min the order changes from 4 mm at 600 t/day to 6 mm at 650 t/day, and it runs the lehr on the MPC of Section 5.5.6 with the preview On. The change reaches the lehr twice. The whole ribbon slows at once, so the thin glass already in the lehr cools more slowly and keeps less, and the thick glass, which needs a much warmer zone A, first enters the lehr at 35.4 min. The structures are compared in Table 5.8 on case 9, case 8, the challenge with its suggested settings, the pull ramp of case 2 and the stable ribbon. Each row has the same seed, duration (360 min, the challenge 300 min), instruments and noise, the MPC keeps the settings of case 9 throughout, and each 3 m parcel is judged by its true residual at the cut. The MPC solved 300 to 360 programs per case and rejected none.

On case 9 the MPC leaves 6 m of ribbon outside the band, all of it above, with a worst cut of 1.114 MPa and 165 K of travel, against the supervisor's 36 m, 1.470 MPa and 241 K. The plan warms zone A0 to 15.2 K at 29 min, before the line slows, because the job card says the thin glass then in

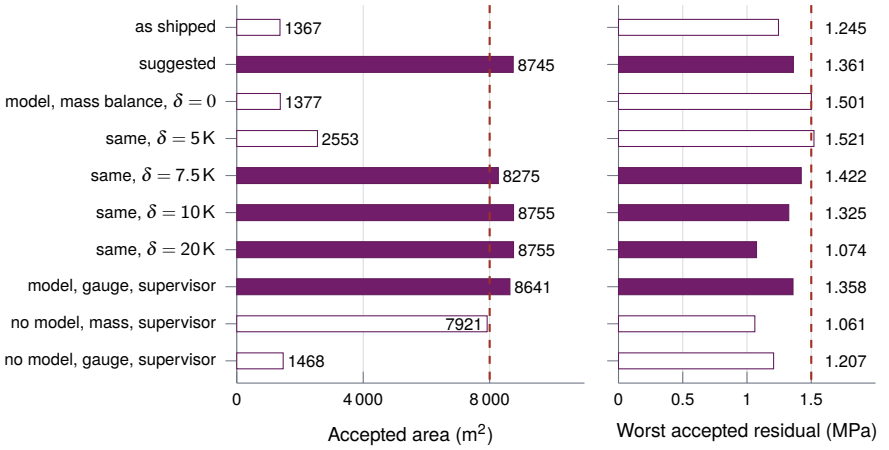

Figure 5.4 The product-change challenge by configuration, from Table 5.7: the accepted area over 300 min (left, dashed line at the 8000 m^2 target) and the highest true residual among accepted parcels (right, dashed line at the 1.5 MPa limit). Dark bars pass the challenge and white bars fail it. Notice that the forming model with the mass-balance thickness is not enough without annealing (third row), and that the supervisor and a fixed offset of 7.5 K or more both pass.

Table 5.8 Annealing Structures Against the Job Changes. Ribbon Outside the Target ±0.1 MPa as Meters Over/Under, the Worst Residual Cut and the Setpoint Travel of Zones A0, A and B Over the Whole Run

Structure	Case 9, 4 to 6 mm, 650 t/day			Case 8, 4 to 6 mm		
	Out (m) over/under	Worst (MPa)	Travel (K)	Out (m) over/under	Worst (MPa)	Travel (K)
offset fixed at 6.5 K	2481/21	1.409	0	2370/21	1.357	0
supervisor, shipped	30/6	1.470	241	33/0	1.456	158
supervisor, retuned[a]	12/6	1.175	61	12/36	1.172	124
supervisor, job-card FF[b]	0/90	1.042	230	0/96	1.047	140
MPC, preview On	6/0	1.114	165	6/9	1.110	173
MPC, preview Off	6/45	1.125	143	9/45	1.118	154
MPC, change 5 min late	9/0	1.119	168	9/15	1.121	179
MPC, original rule[c]	0/324	1.018	22	0/240	1.018	23

Structure	Challenge, suggested			Case 2, pull		Case 1, stable	
	Out (m) over/under	Worst (MPa)	Travel (K)	Worst (MPa)	Travel (K)	Worst (MPa)	Travel (K)
offset fixed at 6.5 K	2199/21	1.464	0	1.075	0	1.024	0
supervisor, shipped	33/0	1.461	100	1.035	155	1.035	130
supervisor, retuned[a]	1686/12	1.175	54	1.035	155	1.035	130
supervisor, job-card FF[b]	0/99	1.058	92	1.037	161	—	—
MPC, preview On	3/0	1.102	198	1.038	143	1.038	154
MPC, preview Off	6/45	1.114	167	1.038	143	1.038	154
MPC, change 5 min late	6/3	1.112	200	1.038	140	—	—
MPC, original rule[c]	0/240	1.043	47	—	—	—	—

No row has ribbon outside the band in cases 2 and 1, a meter over the 1.5 MPa limit or an accepted parcel off specification. [a] Offset rate 10 K/min, largest offset 15 K. [b] From the change on, the larger of the supervisor's offset and the offset it forecasts for the scheduled product. [c] $Q = 3$, $S = 2$, rate capped between block centers.

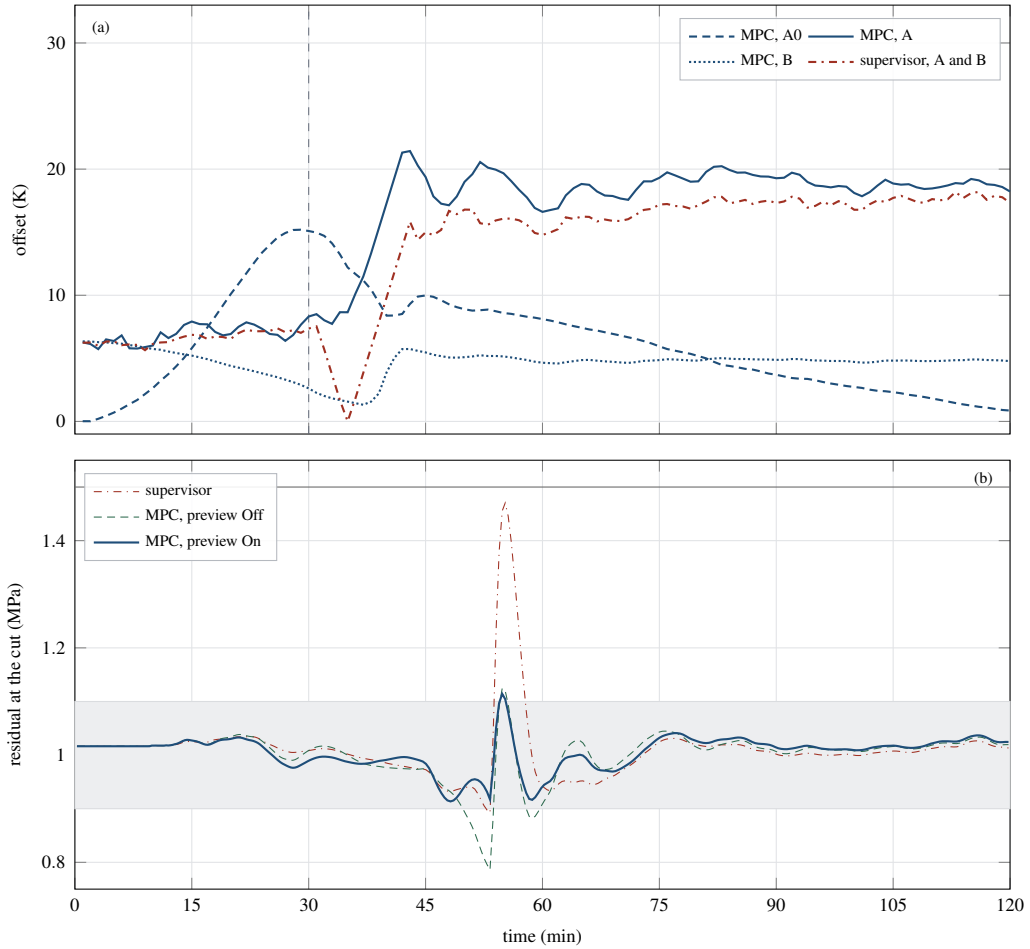

Figure 5.5 Case 9 from 0 to 120 min, the job change at 30 min dashed. (a) Offsets above the recipe applied to zones A0 (dashed), A (solid) and B (dotted) by the MPC with the preview On, all thick, and to zones A and B by the supervisor (dash-dotted). (b) True residual of each 3 m parcel at the cut under the supervisor (thin dash-dotted), the MPC with the preview Off (thin dashed) and On (thick solid), with the 1.0 MPa $\pm$ 0.1 band shaded and the 1.5 MPa limit a thin gray line.

the lehr is about to over-anneal and a warmer A0 makes it keep more. It then raises zone A at the 2 K/min rate to 21.4 K by 43 min, so the first 6 mm glass, cut near 54 min, already finds a warm zone, as shown in Figure 5.5. The supervisor forecasts only the glass entering the lehr. It drops zones A and B to the recipe by 35 min for the slowed thin glass and climbs back after the thick glass has arrived, 15.8 K at 43 min, and its first 6 mm parcels keep up to 1.470 MPa at 55.2 min, 0.03 MPa under the limit. The plan pays for its lead with zone A0, which keeps a mean of 8.9 K from 40 to 60 min while zone A works against it, an offset the supervisor never uses. Case 8, the same order change at 600 t/day and outside the tuning, gives 15 against 33 m, and the challenge, with the MPC in place of the supervisor in the suggested configuration, gives 3 against 33 m while accepting 8755 m^2 against 8745 m^2.

The preview serves the thin glass, not the thick glass. With the preview Off the worst cut is still 1.125 MPa, because the new order is known once it is in force, 5.4 min before the first thick glass reaches the lehr, but 45 m of the slowed 4 mm glass falls under the band, and zone A0 never rises

Table 5.9 Peak Surface Tension in the Lehr on the Two Extreme Job Changes, Before and After the Check-Level Constraint: Peak, Seconds Above the 10 MPa Check Level and Meters Cut Over the Stress Limit

	6 to 3 mm, 600 t/day			4 to 8 mm, 850 t/day		
Structure	MPa	s	m	MPa	s	m
offset fixed at 6.5 K	9.90	0	51	9.98	0	2367
supervisor, shipped	10.32	84	39	10.27	3355	1293
supervisor, retuned	10.32	81	39	10.23	2761	1425
MPC On, without	10.64	160	33	10.43	5045	1104
MPC On, with	9.84	0	75	10.01	16	1383
MPC Off, without	10.32	90	36	10.43	5054	1107
MPC Off, with	10.28	74	39	10.01	16	1374
MPC On, 5 min late, with	10.26	55	66	10.03	49	1356

above 3.7 K before 40 min. A change made 5 min later than the job card announces costs 9 instead of 6 m on case 9 and 24 instead of 15 m on case 8. A schedule known ahead pays on other transport systems too, as Volume 9, Chapter 5 shows for releases made a travel time ahead of published demand. Two classical structures given more than the shipped supervisor were measured beside it. The supervisor retuned once by the rule that chose the MPC's weights, over offset rates from 0.5 to 10 K/min and largest offsets from 15 to 40 K, takes 10 K/min and 15 K and halves the shipped 36 m to 18 m on case 9, but its 15 K cannot anneal 6 mm glass at 700 t/day, and it leaves 1698 m of the challenge outside the band, none of it over the limit. The supervisor given the job card as a feedforward cuts the lowest peaks of all, 1.042 to 1.058 MPa, at the price of 90 to 99 m over-annealed, and of 24 pairs of lead from −4 to 8 min and gain from 0.25 to 1 none leaves less than 45 m on case 9. The pull ramp and the stable ribbon are ties, with no meter outside the band in any row, and on the stable ribbon the MPC moves the zones more, 154 against 130 K.

The weights were chosen on case 9 alone by a rule fixed before the first grid, no meter over the limit, then the fewest meters off target, then the least travel within 3 m of the best. The rule as first written counted only glass above the target, and the winner of its first grid of 108 settings parks the zones hot from the start, with 0 over and 324 m under on case 9 (the last row of Table 5.8). The second criterion was amended after that grid, openly, to count glass off target on both sides, as the target and the supervisor do, and the weights in use come from 150 further settings. Under either rule the MPC leaves less glass over the band on case 9 than the shipped supervisor, 0 or 6 m against 30 m, and both rules retune the supervisor to the same setting.

The surface tension of the ribbon in the lehr was not in the plan as first built, and the measurement exposed it. On the two most extreme changes, 6 to 3 mm at 600 t/day, where the line nearly doubles its speed with 6 mm glass in the lehr, and 4 to 8 mm at 850 t/day, an order the top rolls can form only as 7.25 mm, the plan without the tension chain raised the peak to 10.64 and 10.43 MPa, above the supervisor's 10.32 and 10.27 MPa and the 10 MPa check level, as shown in Table 5.9. The cause was none of the plan's particular choices. The rate cap between block centers gives 10.60 and 10.43 MPa, no band above the target 10.64 and 10.42 MPa, and offsets weighted toward the recipe 10.64 and 10.43 MPa. The peak lies at zone F's entrance, where the surface tension is the temporary stress less the compression the surface keeps, so a kelvin that takes compression off the glass puts tension on its surface. On 6 mm glass at 15.9 m/min a kelvin on zone B lowers the compression kept by 0.024 MPa and raises the peak by 0.030 MPa. The plan was annealing glass that would be rejected anyway, its worst cuts at 2.30 and 1.69 MPa against the supervisor's 2.54 and 1.79 MPa, and the surfaces of that glass met zone F in more tension. The second chain of Equation (5.12) with its soft ceiling at σ_{tmp} was the fix, its price set at ten times the reject level's before any run, because a check risks the ribbon while a reject loses only glass. The one-step forecast of each cohort's peak runs a median of 0.030 to 0.045 MPa above the peak it then meets on four cases, with fifth percentiles of −0.004 to

Table 5.10 Worst Cases Under Every Structure: Meters Cut Over the 1.5 MPa Limit and the Worst Residual Cut (MPa), the Fixed Offset Being 6.5 K

Case	Fixed	Supervisor	Retuned	MPC On	MPC Off
7, zone A failed	918/2.673	915/2.670	915/2.673	915/2.670	915/2.670
damper 50 %	0/1.197	0/1.193	0/1.193	0/1.173	0/1.173
4 to 8, 600 t/day	1818/1.544	0/1.317	0/1.165	0/1.101	0/1.113
4 to 8, 850 t/day	2367/1.976	1293/1.787	1425/1.798	1383/2.004	1374/2.004
6 to 3 mm	51/2.506	39/2.537	39/2.571	75/2.644	39/2.710
pull to 850 t/day	0/1.196	0/1.036	0/1.036	0/1.038	0/1.038
TC +20 K	3669/1.714	138/1.808	126/1.739	135/1.726	135/1.726
TC −20 K	0/1.022	0/1.035	0/1.035	0/1.038	0/1.038
$e_h = -30\,\%$	0/1.357	0/1.396	0/1.415	0/1.261	0/1.260
$e_h = +30\,\%$	0/1.357	0/1.463	0/1.201	0/1.087	0/1.097

Damper stuck from 20 to 100 min, pull ramped at 400 t/day per hour, thermocouple (TC) bias on zone A at 20 min, e_h on case 8.

–0.051 MPa. With the constraint the peaks fall to 9.84 and 10.01 MPa, under the supervisor's in every preview setting, and the price is paid in glass, the meters cut over the stress limit rising from 33 to 75 and from 1104 to 1383, more than the supervisor's 39 and 1293. Equal pricing gives 10.32 and 10.14 MPa with 42 and 1311 m, which only ties the supervisor's peak, and it was not kept. Case 9, case 8, the challenge and cases 1, 2 and 7 are unchanged in every measured column, since there the ceiling never binds.

The worst cases of Table 5.10 take the extreme job changes, the fastest pull ramp, the zone faults and the estimator's heat-transfer error under every structure. No structure accepts a parcel off specification in any of them, and the failed zone A of case 7 defeats every structure alike, with 915 to 918 m over the limit. Against the supervisor the MPC cuts less over the limit with zone A's thermocouple reading 20 K high (135 against 138 m, worst 1.726 MPa against 1.808 MPa) and a lower worst residual under either heat-transfer error (1.261 and 1.087 MPa against 1.396 and 1.463 MPa), and it loses the two extreme changes as described above, with higher worst cuts there as well. On 4 to 8 mm at 600 t/day no structure but the fixed offset cuts glass over the limit, but the preview costs ribbon off target. The plan as first built forecast an order of 8 mm at 5.6 m/min where the line was making 7.25 mm at 6.2 m/min, and once the forecasts were clipped to the thickness the top rolls can form, its 153 m outside the band fell to 81 m, and its travel from 509 to 235 K. The 81 m still exceed the preview Off's 51 m and the supervisor's 48 m, because the rate cap applied at every block boundary makes the plan climb zone A early, while the slowed thin glass is still in it. A cap between block centers leaves 21 and 42 m with the preview On and Off, but 36 m on case 9 instead of 6, because it lets the plan believe that a coarse later block can jump several kelvin and postpone its climb at every decision. The cap at every boundary was kept for case 9, and this loss is its measured cost.

> ✦ **DID YOU KNOW? No Structure Rescues a Failed Zone**
>
> From the simulator: with zone A failed as in case 7, the fixed offset, the shipped and the retuned supervisor and the predictive controller with or without its preview all cut 915 to 918 m of ribbon over the stress limit (Table 5.10). The glass freezes inside the failed zone, where no other zone can reach it.

5.7 Technology in Practice

What Measures It

A float line measures its two qualities where the glass is cold and hard. The gauge stands 171 m past the top rolls and the polarimeter just beyond it, so both read glass shaped and annealed minutes earlier. A 2002 roadmap for the industry found that current sensors and process controls cannot measure and control melting, refining and forming in real time [22]. The chapter's answer is to infer the thickness

at the width camera from the mass relation, and to forecast the stress from the estimator's copy of every parcel, with the cold-end readings kept to correct the inference and the forecast.

> **TECHNOLOGY Sensor: The Polarimeter That Reads Stress**
>
> Glass under stress is birefringent. Polarized light entering it splits into two waves that travel at different speeds, and by the stress-optic law the retardation between them grows with the stress, the thickness and the stress-optical constant of the glass. A polariscope reads that retardation. Newer instruments image it pixel by pixel through a polarization camera, or at several wavelengths at once, the method the latest inline quality systems use. Readings taken by hand at a few points are complex, prone to error and dependent on the operator, and their results vary with who takes them [23]. The chapter's polarimeter is a 12 s filter with 0.03 MPa of noise that reads only the permanent stress.

What Moves It

Four final elements decide the two qualities. The top rolls and the lehr drive set the stretch and the speed, and so the thickness and the width through one mass balance. The heaters and cooling air of the six lehr zones set the thermal path that freezes the stress. Where rolls are driven by alternating-current motors, their speed is set by variable frequency drives, which vary the frequency of the power fed to the motor and are important final control devices in industrial systems [24]. The chapter's drives follow their commands through a 6 s lag, and each zone output acts on its heat and cooling capacity directly, as in Equation (5.8). Figure 5.6 draws the loops of Section 5.5 on these devices.

> **TECHNOLOGY Final Element: Heaters and Cooling Air of the Lehr Zones**
>
> Modern lehrs are modular or zonal, with heat sources placed where they are needed, and many recirculate hot air within each zone and recover waste heat. In 2002 over 90 % of lehrs were fired with natural gas, while electric lehrs, cleaner and cheaper to build but costlier to run, annealed storefront and safety glass [10]. Cooling air is throttled by dampers, multi-vane devices for large flows of air at low pressure [24]. Zones A0 to C of this chapter heat electrically and all six cool with air. A damper stuck open shows on a trend as a zone falling below its setpoint while its loop sits at full heat, as zone A does in case 7.

Where the Controller Runs

Process plants run such loops in a distributed control system, whose supervisory level oversees field controllers that hold their variables at setpoints, with cycle times from milliseconds to minutes [25]. The application executes its loops and its estimator every 5 s on a 1 s physics step, the supervisor once a minute and the predictive controller every 12 scans. The velocity-form PI loops of Table 5.4, with their forming-model feedforwards and the Smith predictor, are the base layer. The estimator and the supervisor, or the MPC in its place, sit above them and write the offsets of the annealing zones.

What Is Modernizing It

Three developments bear on this line. The first is model-based control. The 2002 roadmap planned intelligent process control in stages, from models of process parameters, through stand-alone advisory systems for operators, to neural network systems that adjust the process automatically, and it asked for real-time dynamic models for process control [22]. The estimator and the predictive controller of this chapter are such models run in closed loop. The second is the smart sensor, which reports when it is drifting or failing and calibrates itself [22]. The bias update from the polarimeter does part of that job here, by model rather than by instrument. The third is inline stress measurement. Multi-wavelength photoelasticity is used in the latest inline systems for quality control, where manual point readings are time-consuming under production conditions [23].

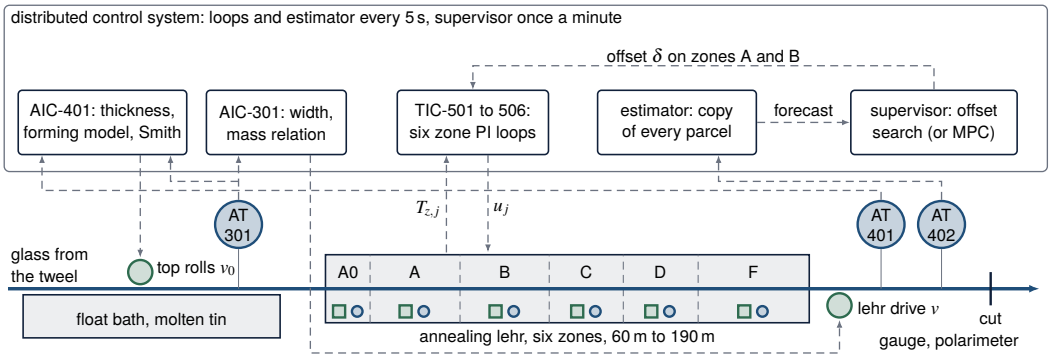

Figure 5.6 The forming and lehr loops of the float line on plant hardware. The thickness loop (AIC-401) moves the stretch through the top rolls, the width loop (AIC-301) moves the lehr drive, and six zone loops (TIC-501 to 506) move the heaters and cooling air of the lehr. The width camera (AT-301) feeds the mass-balance thickness, the gauge (AT-401) the Smith predictor, and the polarimeter (AT-402) the estimator's bias, while the supervisor writes the offset of zones A and B. Notice that both quality instruments stand past the end of the lehr, far downstream of the top rolls and of zones A and B, where thickness and stress are made.

> **IN PRACTICE** **Annealing in Stages, Checked in Polarized Light**
>
> Flat glass is annealed so that it is virtually strain-free and can be cut easily. A 2002 profile describes the lehr in stages. The glass is held just above its annealing temperature, about 1000 to 1050 °F for soda-lime glass, until its stresses relax, then cooled slowly so that no gradient forms between its inside and its surfaces [10]. ASTM C1279 sets out photoelastic test methods for the edge and surface stress of tempered glass, reading edge stress by hand with a polarimeter at two points [23]. The target and the limit of this chapter are teaching values, not those of a standard.

5.8 Limitations

The furnace is lumped, and a cold melt appears only as a growing blanket because glass quality is not modeled. Forming has no cross-ribbon profile, and glass thicker than t_{eq} can be made only with $\lambda < 1$.

The stress model has one relaxation time per depth and no structural relaxation [14], and the ribbon has no edges. The polarimeter is taken to read the permanent stress, but at 199 m the ribbon in case 1 still carries 0.59 K from its middle to its surfaces, a temporary 0.32 MPa of surface tension that a real instrument would see together with the 1.01 MPa kept.

The loops are PI controllers with feedforwards, and the supervisor moves one offset on two zones, while a plant would also schedule the later zones for the product and the pull. Manual and automatic switch every loop together, so a single loop is frozen by setting its gain to zero, as in case 3. The inspection reads each parcel's own value with only the instrument's noise, while a real gauge averages over its spot and its readings must be aligned with the moving ribbon. The forecast treats a zone as healthy until its loop saturates, so the first 6.3 min of the zone A failure pass unseen. Rejected glass is not recycled as cullet.

The MPC's model is the estimator linearized at each decision, at the measured speed and the thickness entering now, so a change of speed or thickness inside the horizon reaches its sensitivities only at the next decision, and its zone loops were identified at 4 mm and 600 t/day and are not refit when an operator retunes them. Its rate cap at every block boundary is pessimistic about the coarse later blocks, which is right on case 9 and costs 81 against 21 m on the 4 to 8 mm change. It leaves zones C, D and F at their recipe because they keep no stress, although on the formable products a kelvin more on zone F lowers the peak surface tension by 0.020 to 0.036 MPa, a lever on the check

level that costs no kept stress and that the plan does not use. Using it would need zone F identified and a limit on the temperature of the glass leaving the lehr, which the model does not constrain and which a kelvin on zone F raises by 0.4 to 0.6 K at the cut.

OPEN CHALLENGE Measuring a Ribbon Before Its Faults Freeze

A 2002 roadmap for the glass industry found that current sensors and process controls cannot measure and control melting, refining and forming in real time, and that an upset can affect product quality for several days. It listed poor control of temperature gradients during forming among the barriers [22]. Here the forecast doubts a failing zone only once its loop saturates, so the first minutes of the zone A failure pass unseen. An estimate that caught a falling zone sooner would be progress.

5.9 Conclusion

This chapter studied a float line, which fixes its two qualities far upstream of the instruments that measure them. The simulator follows every 3 m of ribbon so that the fixing and the measuring can be compared. In the reference case the stress is settled within the first 36 m of a 130 m lehr and changes by less than 0.01 MPa afterward, while the gauge reads glass formed 15.2 min earlier. The stress against the annealing offset has a minimum near 29 K, which argues for a search rather than a loop on the forecast.

The mass-balance thickness at the camera sees a 4 % bath drift 2.3 min after forming instead of 15.5 min and cuts the rejected glass from 734 to 212 m^2. With zone A failed the forecast declares the lehr infeasible at 27.0 min, 6.5 min before the polarimeter, and no glass over the limit is accepted. In the product change the forming model, the inferred thickness and the supervisor raise the accepted area from 1367 to 8745 m^2.

An MPC that plans zones A0, A and B from the same estimator and the job card cuts the ribbon off target on a scheduled job change from 36 to 6 m. On changes too extreme to anneal it briefly exceeds the lehr check level, with 10.01 MPa and 16 s above the check, while sending more glass over the limit.

The model is reduced order, with a lumped furnace, one Maxwell element per depth and a polarimeter that reads only the permanent stress. A fuller model would add structural relaxation, a cross-ribbon thickness profile, recycling of rejected glass and a forecast that doubts a falling zone before its loop saturates. Within these limits the line shows what a cold-end reading can and cannot say about the glass being made.

KEY POINTS

- A float line makes its thickness on the tin and its stress in the first meters of the lehr, minutes before the cold-end instruments read either.
- The delays are distances over the lehr speed, from 15.2 min for 4 mm glass to 22.8 min for 6 mm glass at the same pull.
- The mass-balance thickness sees a 4 % bath drift 2.3 min after forming instead of 15.5 min and cuts the rejected glass from 734 to 212 m^2.
- The stress has a minimum against the annealing offset, so the supervisor searches a grid, and with zone A failed it declares the lehr infeasible 6.5 min before the polarimeter.

Where else this applies. Volume 16, Chapter 9 carries steel strip along a temperature path whose speed the line sets, and Chapter 10 corrects a soft sensor by a late, exact measurement in the same way.

References

[1] O. S. Narayanaswamy, “A one-dimensional model of stretching float glass,” *Journal of the American Ceramic Society*, vol. 60, no. 1-2, pp. 1–5, 1977.

[2] E. H. Lee, T. G. Rogers, and T. C. Woo, "Residual stresses in a glass plate cooled symmetrically from both surfaces," *Journal of the American Ceramic Society*, vol. 48, no. 9, pp. 480–487, 1965.

[3] R. Vivek and K. Ramesh, "Residual stress analysis of commercial float glass using digital photoelasticity," *International Journal of Applied Glass Science*, vol. 6, no. 4, pp. 419–427, 2015.

[4] H. Aben, J. Anton, and A. Errapart, "Modern photoelasticity for residual stress measurement in glass," *Strain*, vol. 44, no. 1, pp. 40–48, 2008.

[5] P. V. Volkov, A. V. Goryunov, A. Y. Luk'yanov, and A. D. Tertyshnik, "Automated multichannel system for in-line monitoring of float-glass ribbon thickness in the hot zone of an annealing furnace," *Glass and Ceramics*, vol. 65, no. 5-6, pp. 144–147, 2008.

[6] L. A. B. Pilkington, "Review lecture: The float glass process," *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, vol. 314, no. 1516, pp. 1–25, 1969.

[7] J. E. Normey-Rico and E. F. Camacho, "Dead-time compensators: A survey," *Control Engineering Practice*, vol. 16, no. 4, pp. 407–428, 2008.

[8] B. Joseph and C. B. Brosilow, "Inferential control of processes: Part I. steady state analysis and design," *AIChE Journal*, vol. 24, no. 3, pp. 485–492, 1978.

[9] R. Bacci di Capaci, L. Musi, A. Landi, L. Sani, and S. Barmada, "Model-based control of a glass melting furnace," *IFAC-PapersOnLine*, vol. 58, no. 14, pp. 253–258, 2024.

[10] J. L. Pellegrino, "Energy and environmental profile of the U.S. glass industry," Energetics, Incorporated, for the U.S. Department of Energy, Office of Industrial Technologies, Columbia, MD, Tech. Rep., Apr. 2002.

[11] L. H. Adams and E. D. Williamson, "The annealing of glass," *Journal of the Franklin Institute*, vol. 190, no. 5, pp. 597–631, 1920.

[12] G. S. Fulcher, "Analysis of recent measurements of the viscosity of glasses," *Journal of the American Ceramic Society*, vol. 8, no. 6, pp. 339–355, 1925.

[13] O. S. Narayanaswamy and R. Gardon, "Calculation of residual stresses in glass," *Journal of the American Ceramic Society*, vol. 52, no. 10, pp. 554–558, 1969.

[14] O. S. Narayanaswamy, "A model of structural relaxation in glass," *Journal of the American Ceramic Society*, vol. 54, no. 10, pp. 491–498, 1971.

[15] G. K. Chui, "Heat transfer and temperature control in an annealing lehr for float glass," *Journal of the American Ceramic Society*, vol. 60, no. 11-12, pp. 477–484, 1977.

[16] O. Auchet, P. Riedinger, O. Malasse, and C. Iung, "First-principles simplified modelling of glass furnaces combustion chambers," *Control Engineering Practice*, vol. 16, no. 12, pp. 1443–1456, 2008.

[17] M. Zier, P. Stenzel, L. Kotzur, and D. Stolten, "A review of decarbonization options for the glass industry," *Energy Conversion and Management: X*, vol. 10, p. 100083, 2021.

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

[19] 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] 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.

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

[22] K. Jamison, J. Eisenhauer, and J. Rash, "Glass industry technology roadmap," Energetics, Incorporated, for the U.S. Department of Energy, Office of Industrial Technologies, Columbia, MD, Tech. Rep., Apr. 2002.

[23] L. Efferz, C. Schuler, and G. Siebert, "Photoelastic measurements and evaluation methods for edge stress in architectural tempered glass," *Glass Structures & Engineering*, vol. 9, no. 3–4, pp. 439–454, 2024.

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

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

5.10 Exercises

5.1 Concept check · 2 min

In case 1 the ribbon enters the lehr at about 607 °C and leaves it at about 89 °C. Where along the lehr is the residual stress it keeps at the cut decided?

(a) At the lehr exit, where the cold-end air takes over the cooling of the glass
(b) In zone F, where the cooling is strongest and the surface tension reaches its peak
(c) Between the annealing and strain points, within the first 36 m of the lehr
(d) On the tin, where the glass is formed and its thickness is fixed for good

5.2 Concept check · 2 min

The supervisor forecasts the stress for offsets on a 2.5 K grid once a minute and picks the smallest offset that meets the target. Why does it not close a PI loop from the forecast to the offset instead?

(a) The forecast is too noisy for integral action, and a search averages the noise out
(b) A PI loop cannot respect the 2 K/min rate at which the offset may move
(c) A forecast made once a minute is too slow for the integral action of a PI loop
(d) The stress has a minimum against the offset, so the loop's gain changes sign past it

5.3 Concept check · 1 min

In case 6 the top rolls lose 40 % of their grip, while the pull and the lehr speed have not yet moved. What happens to the ribbon leaving the forming zone?

(a) It narrows and thickens, because the same mass flow spreads over less width
(b) It narrows and thins, because a sheet held more loosely gives way to the stretch
(c) It thickens at the same width, because the grip acts only on the thickness
(d) It widens and thins, because the edges no longer hold the glass against the stretch

5.4 Theory and calculation · 5 min

The line pulls 600 t/day of glass of density 2500 kg/m^3 as a ribbon 3.70 m wide. The thickness gauge stands 171 m past the end of the forming zone and the width camera 25 m past it.

(a) Find the lehr speed $v = G/(\rho W t)$ for 4 mm and for 6 mm glass. *Checkpoint:* 11.26 m/min and 7.51 m/min.

(b) Find the travel time from the forming zone to the gauge and to the camera for each product. *Checkpoint:* Gauge 15.2 min and 22.8 min. Camera 2.22 min and 3.33 min.

(c) A predictor that holds the gauge delay fixed in time at its 4 mm value is used on the 6 mm product. By how much is its delay wrong? *Checkpoint:* By 7.6 min, a third of the true delay.

5.5 Theory and calculation · 5 min

The forming relation of Equation (5.4) is $t = t_{eq}\lambda^{-(1+k)/2}$ with $t_{eq} = 6.0$ mm and the nominal grip $k = 0.70$, with no bath drift.

(a) Find the attenuation λ that forms 4 mm and 5 mm glass. *Checkpoint:* 1.611 and 1.239.

(b) The grip falls to $k = 0.42$ while the stretch stays at 1.611. Find the thickness formed. *Checkpoint:* 4.28 mm.

(c) Find the stretch that restores 4.00 mm at $k = 0.42$. *Checkpoint:* 1.770.

5.6 Theory and calculation · 7 min

The viscosity of the glass follows the VFT form $\log_{10}\eta = A + B/(T - T_0)$ with $A = -3.13$, $B = 5030$ K and $T_0 = 211$ °C, η in Pa s, and the relaxation time is η/G_r with $G_r = 30$ GPa.

(a) Find the relaxation time at the annealing point, 548 °C, and at the strain point, 511 °C. *Checkpoint:* 20.8 s and 24.1 min.

(b) The canal runs at 1100 °C. By what percentage does a degree of cooling cut the flow under the tweel, Equation (5.3)? *Checkpoint:* 1.46 %.

(c) With $E = 72$ GPa, $\nu = 0.22$ and $\alpha = 9$ ppm/K, find $E'\alpha = E\alpha/(1-\nu)$ and the frozen difference between the surface and the mean temperature that leaves the 1.0 MPa target. *Checkpoint:* 0.831 MPa/K and 1.20 K.

5.7 App challenge · 8 min

Load scenario 5 · Thickness-gauge delay with the default seed. The bath atmosphere changes at 20 min and the glass settles 4 % thicker on the tin. Use a speed of 2 (minutes of plant per second), so each 120 min run takes one minute.

(a) Run to 120 min as loaded. In the exported CSV find the first row at which `thick_gauge_mm` reaches 4.05 mm, and read the Rejected ribbon tile. *Checkpoint:* 36.3 min on the log (the chapter's 36.1 min), within 0.4 min. 734 m^2, within 5 m^2.

(b) Press Reset, set Thickness measured by to Mass balance, run to 120 min and repeat with `thick_massbal_mm`. *Checkpoint:* 23.0 min on the log (the chapter's 22.8 min), within 0.4 min. 212 m^2, within 5 m^2.

(c) The formed thickness passed 4.05 mm at 20.6 min. Account for the gauge's lag with the speed of 11.26 m/min and the gauge's 12 s filter, and say why the mass balance is quicker. *Checkpoint:* 15.2 min of travel plus about 0.2 min of filter. The camera sees the glass 2.2 min after forming.

5.8 App challenge · 8 min

Load scenario 7 · Lehr-zone failure with the default seed. Zone A loses its heaters at 20 min and is repaired at 100 min. Run to 120 min at a speed of 5, which takes 24 s.

(a) In the exported CSV find the first row with `supervisor_infeasible` equal to 1 and the first row with `stress_pol_MPa` above 1.5. How much warning does the forecast give? *Checkpoint:* Rows at 27.3 and 34.0 min (the chapter's 27.0 and 33.5 min), within 0.6 min. About 6.5 min.

(b) Read the Rejected ribbon tile, the share of it rejected for stress in the tile's tooltip, and the Worst accepted stress tile. *Checkpoint:* 3121 m^2, all for stress, within 5 m^2. 1.46 MPa.

(c) The failure starts at 20 min. Explain why the forecast needs about 7 min to see it, and propose a change to the forecast rule. *Checkpoint:* A zone whose loop is unsaturated is forecast at its setpoint. Zone A's loop saturates at 26.3 min.

5.9 App challenge · 9 min

Load scenario 9 · Annealing MPC: a job change planned ahead, with the default seed. The job card changes the order from 4 mm at 600 t/day to 6 mm at 650 t/day at 30 min, and the lehr runs on the MPC with the preview On. Run to 60 min at a speed of 2, which takes 30 s.

(a) In the exported CSV read `mpc_off_A0_K` at 29 min and `mpc_off_A_K` at 43 min, and read the Worst accepted stress tile. *Checkpoint:* 15.2 K and 21.4 K, within 0.3 K. 1.05 MPa.

(b) Press Reset, set Lehr supervisor to On, and run to 60 min. Read `anneal_offset_K` at 35 min, the largest value of `stress_cut_true_MPa` in the file and the Worst accepted stress tile. *Checkpoint:* 0 K. 1.47 MPa near 55 min, within 0.01 MPa. 1.40 MPa.

(c) Explain why the MPC warms zone A0 before the job change, and why the supervisor's Worst accepted stress tile reads less than the worst parcel it cut. *Checkpoint:* The plan knows the line will slow and over-anneal the thin glass. The worst parcels were transition glass the inspection rejected.

5.10 Application area · 4 min

The model assumes the polarimeter reads the permanent stress. At the polarimeter, 199 m along the line, the case 1 ribbon still carries 0.59 K from its mid-plane to its surfaces, and it keeps 1.01 MPa of surface compression.

(a) Take the profile through the thickness as a parabola, so the surface lies two thirds of the mid-to-surface difference below the mean. Estimate the temporary surface tension with $E'\alpha = 0.831\,\mathrm{MPa/K}$. *Checkpoint:* About 0.33 MPa. The engine gives 0.32 MPa.

(b) What would a real instrument at that position read, and how would an engineer make its reading usable for the bias update? *Checkpoint:* About 0.69 MPa of compression. Move it to cooler, more uniform glass or correct the reading with the estimated gradient.

5.11 Application area · 4 min

The batch plant raises the cullet from 20 to 40 % of the charge at 600 t/day. Raw batch yields $Y_b = 0.83$ of its mass as glass and cullet yields all of it.

(a) The firing settles at 42.9 MW instead of 44.6 MW. Find the fuel saved per day and per metric ton of glass. *Checkpoint:* About 147 GJ per day, 0.25 GJ/t.

(b) Find the cut in the mix charged for the same pull, and explain why the melt level still rises before it settles. *Checkpoint:* 3.8 %. The blanket already on the melt melts faster and releases glass nobody asked for.

5.12 Review the agent · 5 min

An AI assistant reviewing the challenge writes: *"The supervisor's search is needless. Close a PI loop from the stress forecast to the annealing offset with a setpoint of 1.0 MPa. Integral action removes any steady error, so the 6 mm glass at 700 t/day will be held at exactly 1.0 MPa once the loop settles, with no grid and no infeasibility logic."* Find the flaw, say what the loop would do, and say how you would check it in the app.

CHAPTER 6

Continuous Tableting Line Simulator for Tablet Weight, Dose and Material Tracking

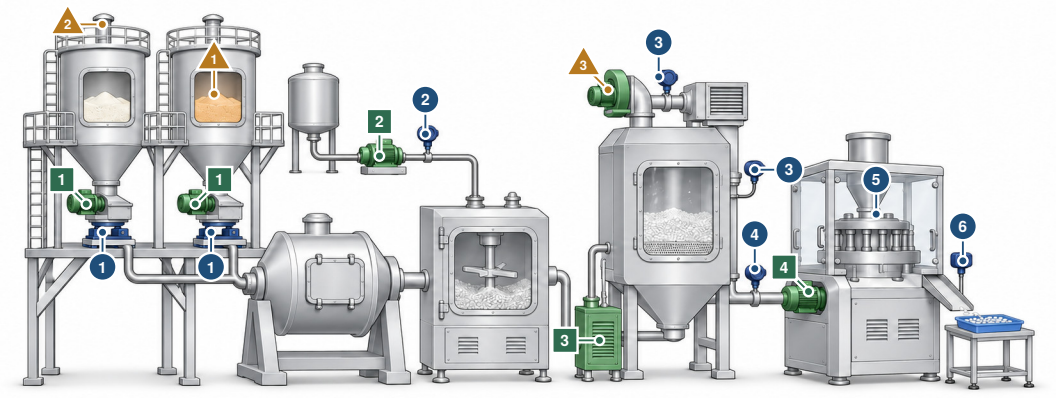

The controller moves 1 feeder screws, one per ingredient 2 binder water pump 3 heater of the drying air 4 press: fill depth, punch gap and speed

The controller reads 1 load cells under each feeder 2 binder water flow 3 drying-air temperature and humidity 4 moisture of each dryer cell as it empties 5 compression force, read as tablet weight 6 at-line weight and hardness tests

What disturbs it 1 a denser excipient lot 2 a feeder refill 3 humid intake air

Notice that the press checks the weight of every tablet, while the medicine content was set at the two feeders on the far left and reaches the press only after the dryer. Illustration, not to scale.

IN PLAIN TERMS **The Right Weight, the Wrong Dose**

The main idea. A tablet line mixes a medicine powder with a filler, makes the mix into small grains, dries them and presses them into tablets. The press judges each tablet's weight by how hard it squeezes, and corrects the weight within a second.

The challenge. How much medicine a tablet holds is set at the feeders, twelve minutes upstream, where the press cannot see it. A feeder being topped up briefly feeds too much, so the line follows each portion of powder to know which tablets to throw away.

Why it matters. A tablet can weigh exactly right and still hold the wrong amount of medicine.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history segment (Section 6.2), the fact on batch size (Section 6.3), the try-its on the feeder refill and the slow analyzer and the fact on the graded shift (Section 6.6) and the key points
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 6.6), Technology in Practice (Section 6.7), Limitations (Section 6.8) and the open challenge

The control question A tablet press holds every tablet's weight from its compression force in about a second, while the feeders fix the composition twelve minutes upstream and run blind while they refill. Which instrument sees each part of a dose error, and what can feedforward, material tracking and a reject gate still guarantee when the dose is wrong?

Learning objectives

- Compute the recipe feeds, the dryer package and the tablets made during an analyzer delay from the line's design data.
- Explain why a tablet at its target weight can carry the wrong dose, and which instrument sees a weight error and which a composition error.
- Derive the feeder's rate gain over a refill and the part of it that the refill feedforward leaves.
- Compare a moisture cascade sampled once per dryer cell with a humidity feedforward on a dryer whose cells see an upset in progress.
- Judge when a reject gate driven by a delayed analyzer, or by a tracker on the measured inputs, protects every accepted tablet.

Prerequisites Sections 1.4.4 (residence-time distributions and material tracking), 1.4.5 (sampled results), 1.4.8 (soft sensors and late references) and 1.6.2 (the PI algorithm and controller action).

The app https://apmonitor.com/apps/continuous-tableting/

See also Volume 3, Chapter 2 for matching a loop structure to where a disturbance enters, Volume 3, Chapter 6 for a dose computed from a measured upstream variable ahead of a long delay, and Volume 3, Chapter 8 for a pharmaceutical formulation line with a ratio block, a release rule and a divert valve.

6.1 Introduction

The simulated tablet press weighs every tablet it makes. It measures the compression force of each one, as the instrumented tableting machine reported in 1954 was built to do [1], converts that force into a weight and holds the weight by moving the depth to which each die is filled. The dose is a different quantity. The content of active pharmaceutical ingredient (API) in a tablet is its weight times the API share of the powder that filled its die, and on a continuous line that share was set by loss-in-weight feeders several kilograms of powder upstream. Batch manufacture could leave the question to the batch, which is accepted or rejected as a whole [2]. A continuous line has to name the tablets a disturbance reached, and the residence-time distribution of each unit [3] is the tool for naming them [4, 5].

The line modeled here is a continuous wet-granulation train. Two loss-in-weight feeders meter the API and an excipient blend into a continuous blender, a twin-screw granulator wets the powder [6, 7], a segmented fluid-bed dryer takes the granules in six cells [8], a mill and a final blender add lubricant, and a rotary press with 24 stations makes 24 tablets a second. A near-infrared (NIR) analyzer reads the blend in the press feed frame, a measurement that has been demonstrated in line on continuous direct compression [9] and used to monitor a commercial continuous run of 28 h [10]. A feeder cannot weigh its own delivery while its hopper is being refilled, and refills are known to disturb the feed rate [11].

What sets this plant apart is where each part of the dose is decided and where it is seen. The weight is controlled at the press in about a second, blind to composition. The composition is fixed at the feeders and carried through 4.38 kg of material in process for a mean of 707 s, spread by three mixers and cut by the dryer into packages of 2160 tablets, so that a feeder error of half a minute arrives as one or two rich or lean packages. The last line of defense is a gate that can reject tablets but cannot correct them, fed by an analyzer whose readings arrive 10 s after their scans, or 240 tablets later. A feeder that delivers nothing cannot be corrected by any loop at all.

The misconception the simulator is designed to correct is that a tablet at its target weight carries its target dose, together with its corollary, that an analyzer at the press protects every tablet it reads. In the second case of the application the true mean tablet weight stays between 249.3 and 250.6 mg

while the package made during a feeder refill reaches the press at 104.77 % LC. The class of problem is general. Any process whose quality is set upstream, carried through a residence-time distribution and released unit by unit needs the same traceability, which a residence-time model of a whole granulation and tableting line has been shown to provide [5], and control of a direct compaction line has combined a dead-time compensator on the API concentration with discarding off-specification powder before the press [12]. However, the control studies of continuous tablet lines cited here are pilot-plant demonstrations built for research [12–14], and a course that meets the press as a transfer function with a dead time has no place for a dose error that the weight loop never measures.

This chapter works with a continuous tableting line simulator that records the true weight, dose and moisture of every tablet apart from what the load cells, the force transducer, the two analyzers and a tracker report, so that a refill, a bridge, a lot change or a slow analyzer can be followed from its cause to the tablets it reaches. Loss-in-weight feeding, tanks-in-series residence-time models, cascade and feedforward control, weight control from compression force and material tracking are all established methods. The simulator joins them in a 133-state line that conserves every component to round-off, turns the stream into dryer packages and individual tablets and runs a measurement-only tracker beside it, and the chapter gives every equation, parameter and default needed to rebuild it. Eight cases and a graded shift separate what each loop, analyzer and gate can and cannot guarantee. The simulator is intended for education in equipment control and runs one locked teaching recipe. Its quality proxies and acceptance limits are illustrative, they have not had quality-assurance review, and neither the application nor this chapter is regulatory or quality-assurance guidance.

Section 6.2 traces the domain from the instrumented press to continuous lines and their residence-time models. Section 6.3 models the line: the recipe, the feeders, the mixers and the granulator, the segmented dryer, the press and the instruments. Section 6.4 separates weight from dose and shows where each is decided and where each is seen, through the tablet spreads, the transport in dryer packages, a refill the load cells cannot weigh and a moisture loop that learns once per cell. Section 6.5 describes the control system as implemented, and Section 6.6 the measured cases and the graded shift. Section 6.7 sets the loops on plant hardware, Section 6.8 states the limitations and Section 6.9 draws the lessons together. The exercises are in Section 6.10.

6.2 Tablets, Granules and Continuous Lines

Instruments on the tablet press date at least to 1954, when the design and construction of an instrumented tableting machine was reported as the third part of a series on the physics of tablet compression [1]. Such a machine measures the forces on its punches, and the force is the only signal the simulated press has of each tablet. It reports how much powder a die received and nothing about what the powder is.

Granules were made in batches. A 2001 paper on trends in granule production noted that a batch offers advantages for quality assurance because it can be accepted or rejected as a whole, while its scale-up may cause problems, and it presented a quasi-continuous line that makes small and large batches on the same equipment [2]. Continuous twin-screw extrusion was then found to wet-granulate lactose robustly and proposed as a possible alternative to high-shear granulation [6], and continuous granulation in the pharmaceutical industry was reviewed in 2005 [7]. By 2016 the continuous manufacture of powder-based pharmaceutical products was the subject of a perspective of its own [15].

A continuous line gives up the batch as the unit of acceptance, and residence-time distributions took its place. The distribution of residence times in continuous flow systems was formalized in 1953 [3], and continuous powder mixing has been characterized by its residence-time distribution [4]. On a continuous dry-granulation tableting line, the distributions of the blender, the roller compactor and the tablet press were measured with NIR, joined into a model of the whole line and checked against the propagation of a step in the API, the traceability that the definition of a batch or lot and the control strategy both need [5]. The feeders can cause such composition steps themselves, since

the hopper refill of a loss-in-weight feeder causes feed-rate deviations [11]. The same traceability decides which material a divert valve sends away on the continuous drug-substance line of Chapter 7.

> ⧖ HISTORY **2002 to 2019: Medicines Begin to Leave the Batch**
>
> Medicines have long been made in batches, while most commodity chemicals, petrochemicals and foods are made continuously. In August 2002 the United States Food and Drug Administration launched an initiative to modernize its regulation of pharmaceutical manufacturing, and a series of guidances followed [16]. A 2015 review by the agency's staff noted that the regulatory definition of a batch rests on an amount of material, not on the mode of manufacture, so a continuous line can still make batches, and it called an understanding of how material flows through the line critical to tracing it [17]. The first end-to-end continuous pharmaceutical process with automated, model-based control ran at bench scale in 2011, and a 2019 summary counted six small-molecule drugs made in part by continuous manufacturing [18].

Measurement and control followed. An in-line NIR method on a continuous direct compression line with two loss-in-weight feeders predicted the API content with root mean square errors of prediction of 0.96 % for the powder at the mixer and 1.37 % for the tablets at the press [9]. At a commercial site, the NIR readings of the blend over a 28 h continuous run averaged 101.17 % LC with a standard deviation of 2.17 % over 12 633 spectra, while 500 of the tablets averaged 100.86 % LC with a standard deviation of 0.4 % [10]. A hybrid of model predictive and PID control with NIR feedback was implemented on the feeders and blender of a continuous direct compaction pilot plant [13], and model predictive control of the main and precompression forces through the compression height and the fill depth held the tablet weight and breaking force together on a pilot-plant press [14]. A control concept for a direct compaction line combined a Smith predictor for the API concentration with model predictive control of the press hopper level [12], and a moving-horizon estimator that reconciles the feeders' mass measurements has been proposed to estimate the composition of the powder they deliver [19]. The dryer between the granulator and the press is a fluidized bed, whose mechanistic models range from single porous granules to the whole bed [8].

6.3 Model of the Line

The simulated line makes a 250 mg tablet at 24 tablets a second, 21.60 kg/h of finished blend, from 20.96 kg/h of powder and 4.611 kg/h of granulation water. Figure 6.1 draws the units, the instruments and the variables of the equations, and Table 6.1 gives each variable its role.

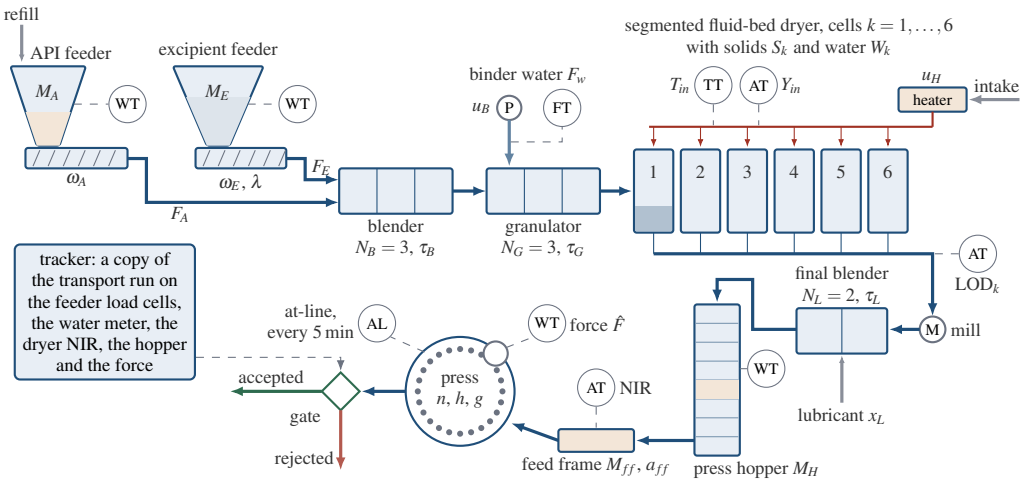

Figure 6.1 The continuous wet-granulation line as modeled. Load cells (WT) weigh the API and excipient feeder hoppers M_A and M_E to give their rates at screw speeds ω_A and ω_E, except during a refill, when the powder at the screw is compacted by D_A or D_E, and a new excipient lot changes its bulk density by λ. The blender and the twin-screw granulator are three tanks in series each, with binder water F_w from the pump u_B on a Coriolis meter (FT). Six dryer cells share one inlet air at T_{in} and humidity Y_{in}, measured by TT and a hygrometer (AT), from a heater u_H, and a dryer NIR (AT) reads each cell's moisture as it discharges. The mill and a two-tank final blender with lubricant feed the press hopper M_H, which passes its slices on first in, first out to the feed frame and its NIR analyzer. The press at turret speed n sets the fill depth h and the punch gap g, measures the force $\hat{F}$ of every tablet and sends ten tablets to the at-line tests (AL) every 5 min. The gate accepts or rejects each tablet on the force, the NIR and the tracker.

> ✦ **DID YOU KNOW? A Batch Set by Running Time**
>
> In the continuous direct compression line that the ICH Q13 guidance works through as an example, the run time at a predefined mass flow rate defines the batch, and market demand calls for batches of 360 to 1080 kilograms [20]. The simulated line makes 21.60 kg/h of finished blend.

Fourteen lumped states hold the two hoppers, screw speeds and refill compactions, the binder pump and the inlet air, together with six running totals, of the API, excipient, water and lubricant fed, of the water evaporated and of the heater's energy. The blender, the granulator, the six dryer cells, the final blender and the feed frame are fifteen vessels that each carry seven components, the API A, the excipient E, water W and lubricant L in kilograms, a granule-size property Z in µm kg, a bulk-density excess R that marks excipient from a new lot, and an inert tracer Q. With the cumulative outflow of the final blender and the pressed totals, the vector has 133 states, integrated together. The dryer schedule, the refills, the press hopper's slices and the tablets are handled as discrete events after every integration step.

6.3.1 Recipe

The recipe is locked. A tablet of $m_t = 250\,\text{mg}$ carries $x_A = 10.0\,\%$ API, which is 25.0 mg, the label claim (LC), at a moisture of $\text{LOD}^* = 2.0\,\%$ loss on drying and with $x_L = 1.0\,\%$ lubricant per kilogram of dried granule. The API share of the powder feed and the powder feed that keeps a press at turret speed n supplied follow, as shown in Equation (6.1).

$$r_A = x_A \frac{1 + x_L}{1 - \text{LOD}^*} = 0.10306, \quad F_s = \frac{N_s n}{60} m_t \frac{1 - \text{LOD}^*}{1 + x_L} \tag{6.1}$$

Table 6.1 Variables of the Line at the Reference Operating Point

Symbol	Role	Reference	Range
$\hat{F}_A$, $\hat{F}_E$	CVs, load-cell rates	2.160, 18.80 kg/h	ratio SP
$\widehat{L/S}$	CV, binder ratio	0.220 kg/kg	SP 0.10–0.35
T_{in}	CV, inlet air	60.1 °C	35–80
LOD_k	CV, cell at discharge	2.0 %	SP 1.2–2.8
$\hat{m}$	CV, weight from force	250 mg	SP 240–260
$\hat{H}$	CV, hardness proxy	100 N	SP 60–140
M_H	CV, press hopper	1.0 kg	SP 0.5–2
ω_A, ω_E	MVs, screw speeds	51.1, 62.0 %	0–100
u_B, u_H	MVs, pump, heater	53.4, 50.2 %	0–100
h	MV, fill depth	7.15 mm	4–12
g	MV, punch gap	3.08 mm	2.4–4.2
n	production rate	60 rpm	30–80
refill	DV, feeder event	at 25 %	or forced
b	DV, bridged share	0	0.85 or 1
λ	DV, excipient lot	0 %	±15
Y_{in}	DV, air humidity	6 g/kg	0–30
block	DV, binder nozzle	open	40 s event
θ_N	DV, NIR delay	10 s	2–120
b_N	DV, NIR bias	0 % LC	±5
M_A, M_E	states, hopper fill	78, 70 %	—
D_A, D_E	states, compaction	0	—
$m_{i,c}$	states, 15 vessels	7 components	—
slices, tablets	discrete states	30 g, 250 mg	—

where $N_s = 24$ is the number of stations. At 60 rpm the powder feed is 20.96 kg/h, of which the API is 2.160 kg/h and the excipient 18.80 kg/h, and the water at the recipe liquid-to-solid ratio $(L/S)^* = 0.22$ is 4.611 kg/h. The acceptance limits of a single tablet are illustrative teaching values, a weight within ±5 % of 250 mg, a dose of 95 to 105 % LC and a moisture of 1.0 to 3.0 % LOD.

6.3.2 Loss-in-Weight Feeders

Each feeder delivers its screw speed times a feed factor that follows the bulk density of the powder at the screw, as shown in Equation (6.2) for the API.

$$\begin{aligned} F_A &= \omega_A \, F_{A,max} \, \tilde{\rho}_A \, (1-b) \, \pi_A \, (1+\varepsilon_A) \\ \tilde{\rho}_A &= 1 + \beta_A \left(f_A - \tfrac{1}{2} \right) + D_A \\ \frac{\mathrm{d}D_A}{\mathrm{d}t} &= -\frac{D_A}{\tau_D} + \frac{\delta_A}{t_R} \, \sigma_R, \qquad \frac{\mathrm{d}M_A}{\mathrm{d}t} = -F_A + F_R \, \sigma_R \end{aligned} \tag{6.2}$$

Here $f_A = M_A / M_{A,cap}$ is the hopper fill, β_A the rise of the feed factor from an empty to a full hopper, D_A the extra compaction a refill gives the powder at the screw, σ_R one while a refill runs and zero otherwise, b the share of the screw blocked by a bridge, π_A a priming factor that rises to one over 8 s once the screw first turns at a start-up, and ε_A a feed-factor variation, first-order colored noise with a 20 s correlation time. A refill starts when the hopper falls below 25 % of its capacity and fills it to 80 % at the constant rate F_R over $t_R = 20$ s for the API and 30 s for the excipient. The screw follows its command through a first-order lag of 1 s. The excipient feeder obeys the same law with its own parameters, priming in 3.2 s and multiplied by $1 + \lambda$, where λ is the bulk-density change of a new excipient lot. A hopper that holds nothing delivers nothing. Loss-in-weight feeders that dose additives into a polymer extruder, with refills and a held weighed flow that is not a held product concentration, are the subject of Volume 12, Chapter 12.

6.3.3 Mixers, Granulation and Granule Size

The blender, the granulator and the final blender are N ideal stirred tanks in series, and every component passes through them by the same law, as shown in Equation (6.3).

$$\frac{\mathrm{d}m_{i,c}}{\mathrm{d}t} = k\,(m_{i-1,c} - m_{i,c}), \qquad k = \frac{N}{\tau} \tag{6.3}$$

where $m_{0,c}$ stands for the inflow divided by k. The blender has $N_B = 3$ and $\tau_B = 45\,\mathrm{s}$, the granulator $N_G = 3$ and $\tau_G = 12\,\mathrm{s}$, and the final blender $N_L = 2$ and $\tau_L = 20\,\mathrm{s}$, so each has the residence-time distribution of N tanks with mean τ and variance τ^2/N [3], the Erlang form of Equation (1.7). The blender receives F_A, F_E and a density excess λF_E, and the binder water $F_w = w_B F_{w,max}(1 - \text{block})$ joins at the granulator inlet, where the pump state w_B follows the command $u_B/100$ through a 2 s lag. The granule size forms at the granulator outlet from the local liquid-to-solid ratio, as shown in Equation (6.4).

$$\begin{aligned} d_{50} &= d_p + (d_{ref} - d_p)\,\min\Big[2.5,\,\Big(\frac{W/S}{(L/S)^*}\Big)^2\Big] \\ \varphi &= \max\big(0,\,1 - d_{50}/d_{ref}\big) \end{aligned} \tag{6.4}$$

where W/S is the water over the solids in the last granulator tank, $d_p = 50\,\mu\mathrm{m}$ the ungranulated powder, $d_{ref} = 300\,\mu\mathrm{m}$ the granule at the recipe ratio and φ the fines index used by the press. The granule size is carried as the flux $Z = (A + E)\,d_{50}$, the mill caps it at 450 µm, and the lubricant enters the final blender at x_L times the discharged granule.

6.3.4 The Segmented Dryer

The six dryer cells run a fixed cycle of 540 s. Cell k fills for $T_f = 90\,\mathrm{s}$ starting at $(k-1)T_f$ modulo the cycle, dries for $T_{dry} = 390\,\mathrm{s}$, discharges for $T_d = 45\,\mathrm{s}$ and stands empty for 15 s, so one cell is always filling and a package leaves every 90 s. A filling cell receives the whole granulator outflow, and each cell is a well-mixed bed of solids S_k and water W_k whose drying follows Equation (6.5).

$$\begin{aligned} \frac{\mathrm{d}W_k}{\mathrm{d}t} &= F_{w,k} - R_c\,\min\Big(1,\,\frac{X_k - X_{eq}}{X_c - X_{eq}}\Big) \\ R_c &= G\,\eta\,\big[Y_s(T_{wb}) - Y_{in}\big], \qquad X_k = W_k/S_k \\ X_{eq} &= X_{e0} + X_{e1}\,\mathrm{RH}(T_{in}, Y_{in}) \end{aligned} \tag{6.5}$$

where $F_{w,k}$ is the water entering while the cell fills, G the dry air per cell, η its contact efficiency, X_c the critical moisture on a dry basis and the rate is zero once $X_k \le X_{eq}$. Drying runs while a cell fills and dries, not while it discharges. The wet-bulb temperature solves $(T_{in} - T_{wb})(1.005 + 1.88\,Y_{in}) = (2501 - 2.36\,T_{wb})[Y_s(T_{wb}) - Y_{in}]$ in kJ/kg by bisection, with the saturation humidity $Y_s = 0.622\,p_s/(P - p_s)$ from a Magnus vapor pressure $p_s = 0.61094\exp[17.625T/(T + 243.04)]$ kPa at $P = 101.325\,\mathrm{kPa}$. A discharging cell empties at the rate $(4/T_d)m_{k,c}$ for every component, and the 1.8 % left after 45 s joins the final blender when the cell turns empty. The inlet air follows the heater through $\tau_h\,\mathrm{d}T_{in}/\mathrm{d}t = T_{amb} + u_H\,\Delta T_{max}/100 - T_{in}$. With the recipe feed and 6 g/kg air, a cell holds 0.524 kg of solids and 0.085 kg of water at the end of its fill, a dry-basis moisture of 0.161, passes below X_c 184.5 s after its fill began, and leaves at 2.0 % LOD with inlet air at 60.1 °C, a heater at 50.2 % and 7.24 kW of heat. The air then has a wet-bulb temperature of 25.4 °C and a constant-rate capacity R_c of 0.342 g/s per cell.

6.3.5 Press Hopper, Feed Frame and Press

The final blender's outflow enters the press hopper as slices of 30 g, which leave from the bottom in the order they arrived, so a dryer package keeps its boundaries until the feed frame mixes it. The feed frame is one well-mixed 90 g holdup refilled from the hopper with a 5 s time constant. The press

starts when the hopper first holds 0.50 kg, stops when it falls below 0.12 kg, restarts above 0.50 kg, and draws the rate $N_s n/60$ times the mean die fill of Equation (6.6).

$$\begin{aligned} \bar{m} &= \rho_f A_d\, h \left[1 - k_\eta (n/n_0 - 1)\right] \\ \rho_f &= \rho_0 \,(1 + k_{\rho w} \Delta\mathrm{LOD})(1 + k_\lambda\, \lambda_{ff}) \left[1 + k_{\rho d} \left(\frac{d_{50}}{d_{ref}} - 1\right)\right] \\ F &= F_{ref} \left(\frac{m}{\rho_t A_d\, g\, SF_{ref}}\right)^{k} (1 - k_w\, \Delta\mathrm{LOD}) \\ H &= H_{ref}\, (F/F_{ref})^{0.8} \left[1 - k_{dw} (n/n_0 - 1)\right] \end{aligned} \tag{6.6}$$

where A_d is the area of the 9 mm die, ΔLOD the feed-frame moisture less 2.0 %, $\lambda_{ff} = R/E$ the lot's density change weighted by the share of new-lot excipient in the feed frame, SF the solid fraction and H a teaching proxy for tablet hardness. With $k = 8$ a tablet 1 % heavier needs 8.29 % more force at the same gap, which is what lets the press weigh every tablet by its force. The reference fill depth is 7.145 mm and the reference gap 3.082 mm, where the force is $F_{ref} = 10$ kN and the hardness 100 N.

Each revolution makes individual tablets, 24 a second at 60 rpm. Tablet i draws two independent standard normal deviates cut off at ± 3, and its weight and dose follow Equation (6.7).

$$m_i = \bar{m}\,(1 + \sigma_m z_{m,i}), \qquad \mathrm{LC}_i = \frac{100\, m_i\, a_{ff}\,(1 + \sigma_a z_{a,i})}{x_A\, m_t} \tag{6.7}$$

where a_{ff} is the API mass fraction of the feed-frame blend and LC_i the dose in percent of label claim. The spreads are $\sigma_m = \sigma_a = 0.6\,\%$ for well-granulated material, widened to $\sigma_m(1 + 3\varphi)$ and $\sigma_a(1 + 8\varphi)$ by fines, which fill dies less evenly and segregate. Each tablet has the moisture of the feed frame. Its true weight, dose and moisture are recorded against the limits, whatever the gate decides.

6.3.6 Instruments

No loop reads a true state. The load cells give each feeder's weight-loss rate through a 4 s filter with 1 % colored noise of 2 s correlation, and give nothing while a refill runs and for 15 s after it. The Coriolis meter reads the water with a 1 s lag and 0.5 % noise, the inlet-air thermometer with a 5 s lag and 0.2 °C, and the hygrometer with a 30 s lag and 0.1 g/kg. The dryer NIR reads each cell's LOD as it starts to discharge, with 0.05 % noise and an optional bias. The feed-frame NIR scans every 2 s, reports the blend's dose with 0.3 % LC of noise and its LOD with 0.05 %, and delivers each reading $\theta_N = 10$ s after its scan with an optional bias b_N. The force transducer reads every tablet with 0.5 % noise, and the press infers the weight and a hardness proxy from it through calibrations made at the recipe moisture and speed, as shown in Equation (6.8). Both are soft sensors in the sense of Section 1.4.8.

$$\begin{aligned} \hat{m}_i &= c_W\, SF_{ref}\, \rho_t A_d\, g\, (\hat{F}_i/F_{ref})^{1/k} \\ \hat{H} &= c_H H_{ref}\, (\hat{F}/F_{ref})^{0.8} \end{aligned} \tag{6.8}$$

The weight calibration ignores the moisture term of Equation (6.6) and the hardness proxy ignores its dwell term. Every 300 s ten tablets are weighed with 0.2 mg of noise and tested for hardness with 3 %, and the results arrive 60 s later. Every deviate is drawn from the seeded generator of Section 1.6.1, so each scenario replays exactly.

6.3.7 Initialization and Numerical Checks

The parameters are collected in Table 6.2. Fixed in the code are the correlation times of the noise, the lags of the Coriolis meter, the thermometer and the hygrometer, the 2 s of rejection after a restart, the heat capacity of 1020 J/(kg K) used for the heater's energy, and the at-line sample of ten tablets. Every case except the start-up begins at the reference operating point, with the mixers at their steady composition, each dryer cell where its schedule puts it at $t = 0$ from one cycle integrated at the

Table 6.2 Line Parameters, Either Adjustable (A) in the Application's Model Dialog or Hidden (H) From It

Symbol	Meaning	Value	Class
$F_{A,max}$, $F_{E,max}$	feeder capacities	4.0, 30 kg/h	A
$M_{A,cap}$, $M_{E,cap}$	hopper capacities	2.0, 12 kg	A
—	refill below, to	25, 80 %	A
t_R	refill time, API, excipient	20, 30 s	A
t_S	weight settling after a refill	15 s	A
β_A, β_E	feed factor rise, full hopper	20, 5 %	A
δ_A, δ_E	compaction by a refill	8, 1 %	A, H
τ_D	compaction relaxation	90 s	A
ε_A, ε_E	feed-factor variation (1σ)	1.0, 0.8 %	A, H
—	screw lag, priming	1, 8 s	H
τ_B, τ_G, τ_L	mixer residence times	45, 12, 20 s	A
$F_{w,max}$	binder pump capacity	8.64 kg/h	A
d_p, d_{ref}	powder, granule size	50, 300 µm	A
—	mill screen, pump lag	450 µm, 2 s	H
T_f, T_{dry}, T_d	cell fill, dry, discharge	90, 390, 45 s	H
G, η	air per cell, contact efficiency	106.2 kg/h, 0.80	A
X_c	critical moisture, dry basis	0.10	A
X_{e0}, X_{e1}	equilibrium moisture	0.004, 0.12	H
T_{amb}	air intake	20 °C	H
ΔT_{max}, τ_h	heater rise, air lag	80 K, 20 s	A
$T_{in,max}$	air the API tolerates	80 °C	A
—	hopper capacity, start, stop	2.5, 0.50, 0.12 kg	A
—	hopper slice	30 g	H
M_{ff}	feed frame holdup	90 g	A
N_s, n_0, d	stations, speed, die	24, 60 rpm, 9 mm	H
ρ_t, ρ_0	true, fill density	1500, 550 kg/m^3	A
$k_{\rho w}$, $k_{\rho d}$, k_λ	fill density factors	0.01, 0.20, 0.75	H
SF_{ref}	solid fraction at F_{ref}	0.85	H
F_{ref}, k	reference force, exponent	10 kN, 8	A
k_w	force loss per % LOD	0.08	A
H_{ref}, k_{dw}, k_η	hardness, dwell, fill loss	100 N, 0.35, 0.10	A
σ_m, σ_a	tablet spreads	0.6, 0.6 %	A
—	growth with fines, σ_m, σ_a	3, 8	H
—	NIR period, noise	2 s, 0.3 % LC	A
—	LOD noise, both NIR	0.05 %	A
—	water, inlet-air noise	0.5 %, 0.2 °C	H
—	hygrometer noise	0.1 g/kg	H
—	load-cell filter, noise	4 s, 1 %	A
—	force noise	0.5 %	H
—	at-line interval, delay	300, 60 s	A

recipe feed, the hopper holding its 1.0 kg setpoint of recipe blend, and the feed factors, screws and pump at the values that deliver the recipe. The start-up case instead begins with every vessel, the hopper and the feed frame empty, the screws, the pump and the heater at rest, the air at T_{amb}, and the press waiting for 0.50 kg in its hopper. At the reference point the line holds 4.376 kg, 0.262 kg in the blender, 0.085 kg in the granulator, 2.819 kg in the dryer, 0.120 kg in the final blender, 1.000 kg in the hopper and 0.090 kg in the feed frame. The states advance by a classical fourth-order Runge-Kutta step of 0.5 s. With the noise off, halving the step moves the peak dose of the refill package in case 2 from 104.676 to 104.677 % LC and the last cell's LOD by less than 0.0001 %, and halving the hopper slice moves the peak by 0.0001 % LC. The component ledger, fed less held less pressed less evaporated, closes to within 5.3×10^{-14} kg in every case, where the application's own check allows

1 µg.

6.3.8 Assumptions

The mixers are ideal tanks in series whose residence time does not change with flow, and nothing segregates in them. Real continuous blenders hold more powder at higher rates, and a cohesive API can separate from its excipient, which would change both the delay and the spread that the tracker assumes.

Granulation is one number, a granule size set by the local liquid ratio, with no nucleation, growth or breakage kinetics. Fines act only through the fill density and the tablet spreads.

Each dryer cell is a lumped, well-mixed bed under one inlet air, with a constant-rate and falling-rate law and illustrative coefficients. The bed temperature, the air leaving each cell, attrition and elutriation are not modeled, where single-granule and bed models are available [8].

The press is lumped, with a power-law compaction curve and a hardness proxy, and the API and the excipient fill and compact alike, so the force carries no information about composition. Capping, lamination, punch wear, relaxation after ejection and differences between stations are not modeled, and the normal spreads cut at three standard deviations mean no tablet is an extreme outlier.

The instruments are lags with noise and additive bias. The feed-frame NIR reads the whole feed-frame blend with 0.3 % LC of noise, far less than the 2.17 % standard deviation reported for NIR on a commercial line [10], and it has no spectral model, sampling volume or calibration maintenance. The refill feedforward reads the hopper mass exactly.

The recipe, the acceptance limits, the hardness proxy and every correlation are illustrative. They were chosen to make the control problem visible and have not had quality-assurance review.

6.4 Where the Dose Is Decided and Where It Is Seen

The dose of Equation (6.7) is a product of two factors that are controlled, measured and disturbed in different places. This section quantifies the three consequences that shape the control problem.

6.4.1 Weight, Composition and the Instruments That See Them

To first order a dose error is the sum of a weight error and a composition error. The force sees the weight, amplified by $k = 8$, and nothing of the composition, because the API and the excipient compact alike in the model. The feed-frame NIR sees the composition of the blend and nothing of the weight or of the spread of individual tablets. The load cells see each feeder's rate, except during a refill. No instrument sees the dose of a tablet.

🕮 **Pause.** Try Exercise 6.1 before reading on.

The tablet spreads set the margin between the blend and the tablet limits. At the default spreads a tablet's dose differs from that of its blend by a standard deviation of 0.837 % and by at most +3.63 and −3.57 %, the truncated extremes, so a blend at a 250 mg mean weight between 98.52 % LC and 101.32 % LC cannot yield an out-of-limit tablet. A Monte Carlo of 4×10^7 tablets puts the share out of limits at 1.25×10^{-4} for a blend at 102 % LC, the default edge of the gate's guard band, 1.0×10^{-2} at 103 % LC and 0.127 at 104 % LC. Fines change the picture. With the 191 µm granules of case 5 the spreads grow to 1.25 and 2.34 %, the dose ratio's standard deviation to 2.62 %, and 5.65 % of tablets fall outside the limits from a blend exactly on target, which no analyzer that reads the blend can detect.

6.4.2 Transport and Packages

A tracer put into the blender at $t = 0$ first reaches a tablet at 637.5 s, is half pressed by 702.5 s and has a mean transit of 707.3 s. Most of that time is spent in the dryer, which holds the granules from the moment they enter a cell to the end of its drying, and in the hopper, whose 1.0 kg is 167 s of production. The dryer also changes the shape of a disturbance. Each cell averages the 90 s it receives into a package of 0.524 kg of solids, 0.540 kg of blend or 2160 tablets, and the hopper keeps the

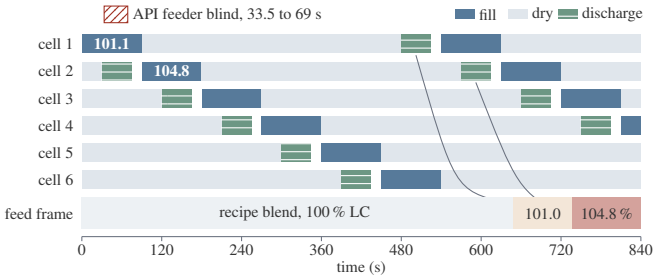

Figure 6.2 Dryer schedule and the refill of case 2. The API feeder is blind from 33.5 s to 69 s and the blender outlet peaks at 93.5 s, so the extra API lands in cell 1, filled from 0 to 90 s, and cell 2, filled from 90 to 180 s. Filling is drawn dark, drying light and discharge hatched. The packages leave at 101.05 % LC and 104.80 % LC when those cells discharge at 480 and 570 s, and reach the feed frame about 168 s later, after the hopper.

packages in order, as shown in Figure 6.2. A feeder error shorter than a fill therefore reaches the press as a step in composition lasting one or two packages rather than as a pulse. The press makes 240 tablets during the default 10 s analyzer delay and 960 during a 40 s one, and every one of them is pressed before the gate learns what the analyzer saw.

🕮 **Pause.** Try Exercise 6.2 before reading on.

6.4.3 A Refill the Load Cells Cannot Weigh

While the API feeder refills, its screw holds the speed the loop last asked for, and the delivered rate follows the density factor of Equation (6.2) alone. From a hopper at 25 %, and ignoring the powder fed during the refill, the ratio of the rate at the end of the refill to the rate before it is given by Equation (6.9).

$$\frac{F_A(t_R)}{F_A(0)} = \frac{1+\beta_A(0.80-\frac{1}{2})+D_A(t_R)}{1+\beta_A(0.25-\frac{1}{2})} \qquad D_A(t_R) = \frac{\delta_A \tau_D}{t_R}\left(1-\mathrm{e}^{-t_R/\tau_D}\right) \tag{6.9}$$

The compaction reaches $D_A = 0.0717$ and the ratio 1.191. With the noise off, the engine's rate runs 18.78 % above its setpoint at the end of the refill, and the feeder delivers 3.31 g more API than asked over the first 400 s, against the 54.0 g a cell receives at the recipe. The load cells see only 0.47 g of the excess, the part delivered after their signal returns at the loop's execution at 69 s. The package that receives most of the powder fed during the blind window leaves at 104.70 % LC and the one before it at 101.03 % LC, while the tracker, which during the refill can only assume that the setpoint was delivered, believes 101.23 % LC and 100.01 % LC. At the default seed the feed-factor variation adds to the error, since a frozen screw cannot follow it either, and the rate peaks 19.9 % high, the blind window from 33.5 to 69 s delivers 2.94 g of extra API, and the packages leave at 101.05 and 104.80 % LC against a belief of 100.10 and 101.09 % LC.

6.4.4 A Moisture Loop That Learns Once per Cell

The dryer NIR reports a cell only when it discharges, one reading every 90 s, and each reading summarizes 480 s of filling and drying, a sampled and held result of the kind treated in Section 1.4.5. At the reference point the discharge LOD falls 0.0935 % per kelvin of inlet air. Air at 16 g/kg needs 69.7 °C for the setpoint and air at 24 g/kg needs 75.2 °C, and with the air held at 60.1 °C the same humidities leave the cells at 3.08 and 3.89 %. A change reaches a cell already in the dryer only in part, as shown in Table 6.3. A humidity step at a cell's age of 270 s carries 0.689 of its full effect on that cell, and a correction of the air at the same age carries only 0.535. The humidity keeps more of

Table 6.3 Share of a Step's Full Effect on the Discharge LOD of One Cell, by the Cell's Age When the Step Occurs, at the Reference Point

Age (s)	0	90	180	270	360	420	480
Y_{in} from 6 to 16 g/kg	1.000	0.896	0.802	0.689	0.500	0.299	0
T_{in} up 1 K	1.000	0.850	0.698	0.535	0.345	0.190	0

Table 6.4 Loops, Defaults and Output Limits. All Are PI in the Velocity Form

Loop	PV to MV	K_c	τ_I	Output
FIC-A	$\hat{F}_A$ to ω_A	12.5 %/(kg/h)	5 s	0–100 %
FIC-E	$\hat{F}_E$ to ω_E	1.6 %/(kg/h)	5 s	0–100 %
FIC-B	$\hat{F}_w$ to u_B	5 %/(kg/h)	4 s	0–100 %
TIC	$\hat{T}_{in}$ to u_H	1.2 %/K	25 s	0–100 %
MIC	$\widehat{\mathrm{LOD}}_k$ to $T_{in,SP}$	8 K/%	200 s	35 °C–80
WIC	$\hat{m}$ to h	0.014 mm/mg	3 s	4–12 mm
HIC	$\hat{H}$ to g	0.0024 mm/N	3 s	2.4–4.2 mm
level	M_H to $F_{s,SP}$	2 (kg/h)/kg	P only	±15 %

its effect because it also raises the equilibrium moisture toward which the falling-rate period relaxes, with a time constant $S(X_c - X_{eq})/R_c$ of about 138 s, while a correction of the air acts mainly through the drying rate over the time that remains. The outer loop is therefore a sampled integrating controller on a plant whose response to each move is spread over the next five to six samples.

6.5 Control System as Implemented

The controllers execute every 1 s, at whole seconds of simulated time, one turret revolution at 60 rpm, on the instruments of Section 6.3. A step follows the order of Figure 1.9, with the discrete events and the tablets after the integration. The historian logs a row every 5 s, and the application runs from 1 to 120 times real time, 30 by default. The loops and their defaults are listed in Table 6.4.

6.5.1 The PI Algorithm

The seven PI loops use the framework's velocity-form PID of Section 1.6.2 as it stands, while the level trim is the proportional law of Equation (6.10). In the symbols of Equations (1.15) to (1.18), each loop has its own action $a = \pm 1$ and output bias u_b, executes every 1 s and uses a tracking time T_t equal to τ_I. No loop here uses derivative action ($\tau_D = 0$), and the back-calculation runs only in loops that have integral action. The output limits of Table 6.4 and any rate limit bound only what reaches the actuator. The weight loop's output slews at most 0.5 mm/s and the hardness loop's at 0.05 mm/s. In manual mode, or whenever a loop does not drive its actuator, it tracks by setting $I_k = u_k - u_b - S_k$, so every return to automatic is bumpless. For a settled scenario, untouched manual stations are initialized from the solved screw speeds, binder pump, heater, fill depth and punch gap. Selecting Manual before Run therefore holds the solved lineup. Station values that the reader sets are restored before priming and survive Reset. The empty start-up keeps its deliberate starting commands.

The moisture loop and the hardness loop have $a = -1$, since a wet cell must raise the air and a soft tablet must close the gap. Their output moves in the same direction as their measurement, so they are direct acting in the convention of Section 1.6.2. The other five PI loops have $a = +1$ and are reverse acting.

6.5.2 Feeders, Ratio Master and Binder

The press speed is the production rate. It slews at 2 rpm/s toward its setpoint, and the ratio master sets the powder demand F_d of Equation (6.1) at the current speed plus a proportional trim on the press

hopper, as shown in Equation (6.10).

$$\begin{aligned} F_{s,SP} &= F_d + \mathrm{clamp}\left[K_H(M_{H,SP} - \bar{M}_H), \pm 0.15 F_d\right] \\ F_{A,SP} &= r_A F_{s,SP} \\ F_{E,SP} &= (1 - r_A) F_{s,SP} \quad \text{or} \quad (1 - r_A)\hat{F}_A / r_A \end{aligned} \tag{6.10}$$

where $\bar{M}_H$ is the hopper load through a 90 s filter and the second excipient setpoint is the measured-API option. The trim is slow on purpose, because the dryer lies between the feeders and the hopper. FIC-A and FIC-E move the screws on the load-cell rates. During a refill and its settling the feeder is volumetric, and its loop tracks the speed it last asked for, and for the first 20 s of a start-up it runs at the speed its calibration gives for the setpoint. With the refill feedforward on, the command is divided by the density factor that the feeder's identified profile predicts, as shown in Equation (6.11).

$$\begin{aligned} \omega_A &= \frac{u_A/100}{1 + 0.95\beta_A(f_A - \tfrac{1}{2}) + \hat{D}_A} \\ \frac{\mathrm{d}\hat{D}_A}{\mathrm{d}t} &= -\frac{\hat{D}_A}{\tau_D} + \frac{0.95\,\delta_A}{t_R}\,\sigma_R \end{aligned} \tag{6.11}$$

The identified profile is 5 % low on both terms, and the loop's upper limit is raised to 100 times the predicted density factor so that the screw can still reach full speed. FIC-B sets the water to $(L/S)_{SP}$ times the measured powder rate, using a feeder's setpoint while it is volumetric.

6.5.3 Moisture Cascade and Humidity Feedforward

TIC holds the measured inlet air on its setpoint every second. It is the inner loop of the cascade, and the humidity enters outside it, so TIC does nothing against a humid intake and MIC answers it only once it has reached a cell. Volume 3, Chapter 2 shows how the place where a disturbance enters bounds what a loop structure can reject. MIC acts only when a cell holding at least 70 % of a nominal fill starts to discharge, on the dryer NIR's reading, and the step it integrates over is the interval since its previous update, nominally 90 s, so the partly filled cells of a start-up do not mislead it. Its output is the inlet-air setpoint, limited to 35 to 80 °C, the highest air the API tolerates. With the humidity feedforward on, the bias of MIC is $b_k = T_{ff}(\hat{Y}_{in}) - T_{ff}(6\,\mathrm{g/kg})$, where $T_{ff}(Y)$ is the inlet temperature the drying model needs for the setpoint at humidity Y, tabulated from 0 to 30 g/kg in steps of 3 at the current production rate and interpolated. Between cells the bias moves the setpoint while the integral does not. With the air fixed, the operator's setpoint is used and MIC tracks. The spray dryer of Chapter 10 also adds a humidity feedforward, because its moisture result arrives late as well.

6.5.4 Weight and Hardness From Force

WIC moves the fill depth to hold the mean force-inferred weight of the tablets made in the last second at its setpoint. HIC moves the punch gap to hold the hardness proxy of Equation (6.8). Each at-line result is a late reference for a soft sensor, as the laboratory is in Section 1.4.8, and trims its calibration by 70 % of the relative difference between the tested and the inferred value, bounded between 0.95 and 1.05 for the weight and between 0.8 and 1.2 for the hardness. Neither loop sees the composition.

6.5.5 Tracker and Reject Gate

The tracker is a copy of the transport run on measurements alone, the tracking of material of Section 1.4.4. It carries the solids, the API, the water and a suspect mass through tanks in series with the plant's structure, fed by the load-cell rates, or by the setpoints while a feeder is volumetric, with the dryer's schedule, the discharge law, each cell's LOD from the dryer NIR, hopper slices scaled to the hopper load cells and a feed frame drawn at the press speed and inferred weight. It never reads a true composition. Powder made before both feeders have held within 3 % of their setpoints for

Table 6.5 Outcomes at the Default Seed. Made and Accepted Count Tablets Outside the Limits, Worst Is the Largest Dose Deviation Accepted

Case, window, configuration	Made	Rejected	Accepted	Worst
1, 1500 s, tracking on	256	3444	0	3.37 %
1, 1500 s, tracking off	256	1644	0	3.37 %
2, 1200 s, feedforward off	609	2712	0	4.90 %
2, 1200 s, feedforward on	0	0	0	3.51 %
3, 1500 s, recipe ratio	5437	7152	0	3.36 %
3, 1500 s, measured API	929	5256	0	3.22 %
3, 1500 s, tracking off	5437	6840	138	9.80 %
4, 1500 s, weight on	0	0	0	3.35 %
4, 1500 s, weight off	15 397	15 101	372	6.45 %
5, 1200 s, tracking on	116	5088	0	3.34 %
5, 1200 s, tracking off	116	0	116	8.75 %
6, 2400 s, cascade	0	0	0	3.86 %
6, 2400 s, air fixed	35 376	37 704	0	3.16 %
6, 2400 s, humidity FF	0	0	0	3.75 %
7, 1800 s, tracking off	4371	11 640	816	10.92 %
7, 1800 s, tracking on	4371	12 672	0	3.36 %
8, 2400 s, trim off or on	0	0	0	3.63 %

20 s, powder granulated while the measured liquid ratio is more than 10 % off, and cells read outside 1.1 to 2.9 % LOD are suspect. The gate rejects a tablet, in this order, during start-up, for 2 s after a restart, when its inferred weight is more than 5 % from the setpoint, when the tracker puts the feed frame outside the guard band $|\mathrm{LC} - 100| > b = 2.0\,\%\ LC$ or its suspect share above 2 %, and when the latest NIR reading is outside the band, or its LOD outside 1.1 to 2.9 %, and for $t_{hold} = 10$ s after. After a start-up the gate opens once the NIR has read inside the band for 20 s, the inferred weight has held within 2 % for 10 s and the tracker holds no suspect material in the feed frame.

6.5.6 Simulation, Control and Internal Models

Section 6.3 describes the simulation model. The controllers carry their own, simpler models, namely the identified feed-factor profile of Equation (6.11), the drying model inverted by the humidity feedforward, which is the plant's own law, the press calibrations of Equation (6.8), and the tracker, which shares the plant's residence-time structure exactly, so that its only errors come from its inputs. No model is identified inside the application.

6.6 Scenarios and Measured Results

The subsections follow the cases in selector order, and the tablets that each configuration made out of limits, rejected and let through are counted in Table 6.5. Each case has a completion rule of the kind described in Section 1.6.4, met when the reader makes the comparison runs it names in Auto and the simulation reaches the state it names. The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt. Unless a change is named, a case runs the reference configuration, with both feedforwards, tracking, the NIR and force gates and both at-line trims on. The values reported here were measured on the engine itself, run without the interface at each scenario's seed, with any named change made at $t = 0$ before the run starts. The application gives the same run when the change is carried through Reset instead. Figure 6.3 draws the worst accepted dose deviation of each configuration against the dose limit.

6.6.1 Start-Up

Case 1 starts every unit empty. The API screw primes over 8 s and the excipient screw in 3.2 s, so at 10 s the blender outlet carries only 55.4 % of the recipe's API share, and the first dryer cell takes all of that lean start. It leaves at 480 s holding only 0.215 kg, dried to 1.09 % LOD by air meant for a full cell. While the press hopper is empty the level trim asks for 2.0 kg/h, 9.5 % above the demand,

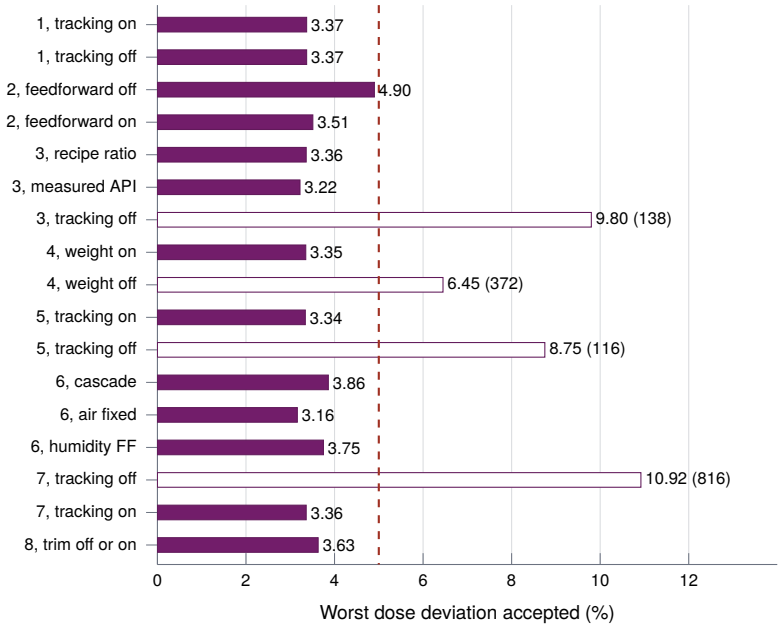

Figure 6.3 The cases by configuration, from Table 6.5: the largest dose deviation among accepted tablets, with the dashed line at the 5 % dose limit. Dark bars accept no tablet out of limits. White bars accept some, and the number accepted out of limits is given in parentheses. Notice that every white bar is a configuration without tracking or without weight control, and that the feeder refill of case 2 with its feedforward off stays inside the limit only because the analyzer's gate reacts in time.

so the next cells leave fuller and wetter, at 2.36, 2.21 and then 2.07 % LOD, as the moisture loop adjusts. The press starts at 605 s. With tracking on the gate opens at 748.5 s, after 3444 rejected tablets carrying 85.0 g of API, of which 256 were truly out of limits. With tracking off it opens at 673.5 s after 1644, and no tablet out of limits is accepted either way. The tracker's extra 75 s is caution about start-up material it cannot vouch for, and it costs 1800 tablets.

6.6.2 A Refill: the Right Weight With the Wrong Dose

Case 2 starts with the API hopper at 26 % and the refill feedforward off. The refill runs from 33.5 to 53.5 s, the feeder loop has its weight signal back at its execution at 69 s, and the blender outlet peaks at 108.4 % of the recipe's API share at 93.5 s. The consequence arrives twelve minutes later, as shown in Figure 6.4. The first package raises the feed frame to about 101.0 % LC from 647.5 s, the second package crosses 102 % LC at 731.5 s and the feed frame peaks at 104.77 % LC at 814.5 s. Throughout, the true mean weight stays between 249.31 and 250.57 mg. In all, 609 tablets are pressed outside the dose limits, and the NIR gate rejects 2712 tablets by 1200 s, 2103 of them within every limit. No over-dosed tablet is accepted, but only just. The first reading above the guard band, from the scan at 730 s, arrives at 740 s, and the first over-dosed tablet is pressed in the next second. The tracker's prediction rises to only 101.06 % LC. With the feedforward on the rate runs at most 2.77 % high while the feeder is volumetric, the feed frame peaks at 100.39 % LC, and no tablet is made out of limits or rejected. Weight checks alone would have passed every tablet of this run.

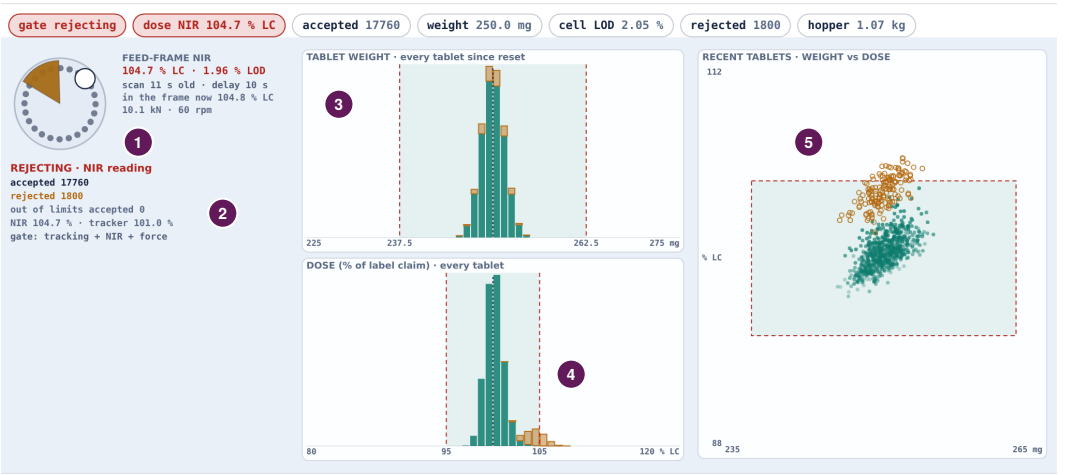

Figure 6.4 The application in case 2, feeder refill, with the refill feedforward off, paused at 815 s, cropped to the plant panel of the Tablets view. The header with the clock, the control panel, the trends and the Learn panel are cut. (1) The feed-frame dial: the package made during the refill fills the feed frame at 104.8 % LC, and the NIR reports 104.7 % LC from a scan 11 s old. (2) The gate's status: rejecting on the NIR while the tracker predicts only 101.0 % LC, with 17 760 accepted, 1800 rejected and none accepted out of limits. (3) The weight histogram of every tablet since reset, centered on 250 mg inside the ±5 % lines, while the status row above it gives the force-inferred weight as 250.0 mg. (4) The dose histogram, with a tail of rejected tablets reaching past 105 % LC. (5) The scatter of recent tablets (every twelfth since 365.5 s, by the application's setting, which the panel does not print), rising in a column of tablets near 250 mg out of its specification box, every one of them a rejected tablet drawn hollow.

> **TRY IT IN THE APP** **2 · Feeder Refill**
>
> Load the case and press Run. When the package made during the refill reaches the press, about 13 minutes in, open the Tablets view. Watch the weight histogram stay centered on 250 mg while the dose histogram grows a tail past 105 % LC. The feed frame peaks at 104.77 % LC at 814.5 s, and 609 tablets are pressed outside the dose limits.

6.6.3 A Feeder That Delivers Nothing

Case 3 stops the API completely for 90 s at 120 s, 54 g of API that never enters the line. FIC-A drives its screw to the limit within 10 s and holds it there for 83 s, and when the arch collapses the rate overshoots to 4.22 kg/h, twice its setpoint, before the loop winds back. No loop can restore the missing API, and the application does not pretend to. With the excipient on the recipe ratio, the package made during the stop falls to 32.4 % LC in the feed frame and 5437 tablets are made out of limits. The tracker, which saw the stop on the load cells, predicts 31.7 % LC, and the gate rejects 7152 tablets and accepts none out of limits. With the excipient following the measured API the composition stays between 95.2 and 102.9 % LC, only 929 tablets are made out of limits, and the cost moves to output, since the line makes less blend and the press hopper falls to 0.439 kg. With tracking off the NIR gate alone lets 138 out-of-limit tablets through before its readings arrive, the worst 9.80 % below label claim.

6.6.4 A Denser Lot: the Right Blend With the Wrong Weight

Case 4 fills the excipient feeder from a lot 8 % denser at 60 s. The excipient rate overshoots by 7.80 % for a few seconds before FIC-E corrects it, and the denser granules reach the feed frame near 745 s and raise the fill density from 550 to 583.0 kg/m^3, the 6.0 % that $k_\lambda \lambda$ predicts. With weight control on, the fill depth falls from a mean of 7.14 to 6.74 mm, and the weight holds at 250 mg apart from a peak of 252.2 mg at 741 s as the front arrives. With it off the tablets reach 265.1 mg, and their dose rises with them while the blend in the feed frame stays between 99.30 and 100.12 % LC. The force gate rejects 15 101 tablets as overweight, and 372 out of limits pass it, 145 on weight, whose force still put them inside ±5 %, and 260 on dose, 33 of them on both counts. The NIR, which reads the blend, sees nothing wrong. This case is the mirror of case 2, and each dose error is invisible to the instrument that watches the other factor.

6.6.5 A Binder Interruption

Case 5 blocks the granulation nozzle for 40 s at 120 s. The pump runs at its limit for 35 s, and when the nozzle clears the liquid ratio overshoots to 0.412 kg/kg at 160 s. The cell that took the dry powder leaves at 570 s with a granule size of 190 µm and 1.40 % LOD, and the next, which took the overshoot, at 318 µm and 2.26 %. The NIR reads between 98.35 and 101.08 % LC throughout, because the blend's mean is right, while 116 tablets are made outside the dose limits, as the fines spread of Section 6.4 predicts. The tracker flagged everything granulated while the ratio was more than 10 % off and rejects 5088 tablets, the affected package and its edges, accepting none out of limits. With tracking off all 116 are accepted, the worst 8.75 % from label claim.

6.6.6 Humid Air and a Cascade Sampled per Cell

Case 6 turns the drying air from 6 to 16 g/kg at 300 s with the humidity feedforward off. The cells already drying finish in humid air whatever the loop does next. With the cascade alone the worst cell leaves at 2.59 % at 480 s, the air setpoint climbs to 70.5 °C, and every cell from 930 s on is within 0.1 % of the setpoint, ten and a half minutes after the change. With the air fixed at 60 °C the cells climb to 3.10 %, 19 of them leave above the 3.0 % limit and 35 376 tablets are made too wet. Told by the dryer NIR which cells were out, the tracker rejects 35 928 tablets and the NIR 1776 more, none out of limits is accepted, and only 34.5 % of the output is kept. With the humidity feedforward the air moves as soon as the hygrometer does, the worst cell leaves at 2.24 %, and every cell from 570 s on is within 0.1 %. The feedforward spends 5.85 kWh of heat over the 40 min against 5.72 kWh for the cascade alone and 4.81 kWh with the air fixed.

🕮 **Pause.** Try Exercise 6.3 before reading on.

6.6.7 A Slow and Drifting Analyzer

Case 7 slows the feed-frame NIR to 40 s at 60 s, bridges the API to 15 % of its rate for 45 s at 120 s, and drifts the analyzer 3 % LC high at 1500 s, with tracking off. The lean package takes the feed frame below 98 % LC at 730.5 s and bottoms at 86.4 % LC at 899.5 s, and 816 tablets out of limits are accepted, the worst 10.92 % below label claim, every one pressed after its powder was scanned and before the scan was reported. With the normal 10 s delay the same bridge lets 103 through, and with tracking on none, because the tracker saw the bridge on the load cells twelve minutes before its package arrived. The drift is the opposite failure. Every reading after it lies outside the guard band, and in the last five minutes the gate rejects 6240 of the 7200 tablets made, all of them within every limit.

Table 6.6 Challenge Configurations Over 3600 s. Worst Is the Largest Dose Deviation Accepted

Configuration	Accepted	Rejected	Out	Worst
as shipped	71 952	14 448	194	9.39 %
all three on	82 992	3408	0	3.72 %
no refill feedforward	75 408	10 992	6	6.09 %
no humidity feedforward	77 280	9120	0	4.01 %
no tracking	83 304	3096	122	8.62 %

TRY IT IN THE APP **7 · NIR Delay and Bias**

Press Run with Reject by material tracking Off, and read the Accepted out of limits tile once the lean package has passed, at about 17 minutes: 816 tablets, the worst 10.92 % below label claim. Reset, set Reject by material tracking to On and run again. None is accepted, because the tracker saw the bridge on the load cells.

6.6.8 A Faster Press

Case 8 raises the press from 60 to 72 rpm at 120 s with the at-line hardness trim off. The feeders are told at once, but the extra powder takes about twelve minutes to arrive, and the press hopper, which swung between 0.90 and 1.10 kg before the change, falls to 0.254 kg at 757 s, while the fuller cells made at the higher rate leave at up to 2.74 % LOD. The dies fill 2 % less at speed, the true mean weight dips to 247.3 mg at 125.5 s, and WIC deepens the fill from 7.15 to 7.30 mm. HIC holds its proxy at 100 N while each tablet is compressed more briefly, and the true hardness settles near 93 N, the factor 0.930 of the dwell term, and reads 92.5 N at 40 min. With the at-line trim on, the ten-tablet tests correct the proxy by bounded steps, and the true hardness is 100.8 N at 40 min with the calibration at 0.918.

6.6.9 The Graded Shift

The challenge runs 3600 s. The API feeder refills at 167.5 s and again at 2014.5 s, the drying air turns to 24 g/kg at 600 s, which needs 75.2 °C inside the 80 °C limit, and the API bridges to 15 % of its rate for 30 s at 1800 s. The shift starts with the refill feedforward, the humidity feedforward and tracking off, and it is met with at least 80 000 tablets accepted and none accepted outside the weight, dose or moisture limits. Table 6.6 sets the configurations side by side. As shipped the gate accepts 71 952 tablets, 8048 short of the target, and 194 of them are out of the dose limits, 4 from the first refill, 189 from the bridge and 1 from the second refill. With all three settings on it accepts 82 992, 2992 above the target, and rejects 3408, and the worst accepted tablet is 2.26 % off in weight, 3.72 % in dose and 2.36 % in LOD at the wettest. Each setting is needed, and the three failures differ in kind. Without the refill feedforward 6 over-dosed tablets pass the NIR, without the humidity feedforward the cells reach 2.94 %, the tracker rejects them as a precaution and only 77 280 are accepted, and without tracking the bridge's package reaches the press before the NIR can react and 122 out-of-limit tablets are accepted.

DID YOU KNOW? Protection That Also Raises Output

From the simulator: in the graded shift as shipped, the gate accepts 71 952 tablets, 194 of them out of limits. With the refill feedforward, the humidity feedforward and tracking all on, it accepts 82 992 and none out of limits (Table 6.6). The settings that protect the dose also keep more tablets.

6.7 Technology in Practice

What Measures It

Three instruments divide a dose between them. The force transducer of the press reports how much powder each die received, the load cells report what each feeder delivers except while it refills, and a near-infrared analyzer in the feed frame reports what the blend contains. The 2023 guidance on continuous manufacturing finds process analytical technology well suited to continuous lines and names in-line near-infrared spectroscopy of blend uniformity or water content among its examples, because it lets disturbances be detected in real time [20]. In the chapter both analyzers are lags with noise and an additive bias, the load cells a filtered rate and the force a reading with 0.5 % of noise, as Section 6.3 sets out.

TECHNOLOGY Sensor: A Near-Infrared Analyzer on the Blend

Molecules absorb light at frequencies at which their atoms vibrate and rotate, in bands unique to each kind of molecule, which lets spectroscopy measure composition [21]. A near-infrared analyzer turns spectra of the flowing powder into an API content through a multivariate model calibrated against a reference method. On one continuous tableting line an analyzer at the blender took spectra from 860 to 1690 nm at 100 a second, averaged to one result a second, and its blend model had a cross-validation error of 0.47 percent by weight [5]. Powders limit how deep the light reaches, so the probe's place and sample size matter [17]. A model that drifts shows on a trend as a step between the analyzer and the true blend, as in case 7. The chapter's analyzer scans every 2 s with 0.3 % LC of noise and reports 10 s late.

What Moves It

The final elements are the two feeder screws, the binder pump, the heater of the drying air, the fill depth and punch gap of the press, and a reject gate after it. The feeders matter most for the dose, because no amount of blending corrects a blend whose ingredients enter at the wrong ratio [16]. In the chapter each screw follows its command through a 1 s lag and the pump through 2 s, and the fill depth slews at most 0.5 mm/s. Figure 6.5 draws the loops of Section 6.5 on these devices.

TECHNOLOGY Final Element: The Loss-in-Weight Feeder and Its Failures

A loss-in-weight feeder meters powder by screw speed under a gravimetric controller that reads its own load cells, and feeders are the most likely source of a content error on a continuous line. Some fluctuation of the feed rate is unavoidable. A hopper refill is the most common cause of a large deviation, because the feeder runs volumetrically while it refills and cannot correct for the change of density. Powder that accumulates in the downspout can break off and raise the content at once. A miscalibrated load cell makes the feeder dispense the wrong rate unseen by its own controller, and shows on a trend as a screw speed that departs from its history while the measured rate stays in range [16]. The chapter's load cells report nothing during a refill and for 15 s after it.

Where the Controller Runs

A distributed control system, used in industries that include pharmaceutical processing, oversees localized controllers that hold their variables at setpoints, with loop cycle times from milliseconds to minutes [22]. A review by the regulator's staff describes a supervisory layer that reconciles measurements of parameters, raw materials, in-process materials and product with a model of the process dynamics to support real-time release [17]. In the application every loop executes once a

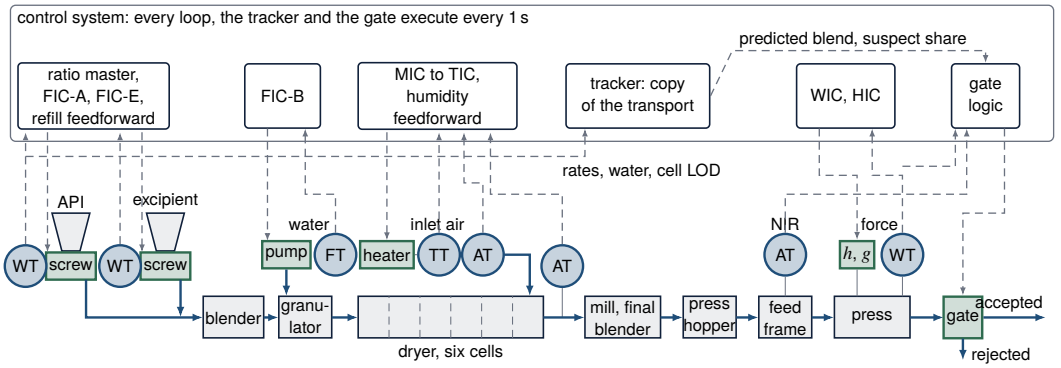

Figure 6.5 The loops of the tablet line on plant hardware. Load cells (WT) of the two feeders drive FIC-A and FIC-E, which move the screws under the ratio master. The water meter (FT) drives FIC-B on the binder pump, the dryer's moisture analyzer (AT at the cell discharge) drives MIC, which sets TIC on the air heater, and the force transducer (WT) drives WIC and HIC on the fill depth h and punch gap g. The tracker copies the transport on the measured rates, water and cell moisture, and the gate logic combines it with the feed-frame analyzer (NIR) and the force. Notice that the analyzer and the gate stand at the press, while the composition they judge was set at the screws on the far left.

second, one turret revolution at 60 rpm. The velocity-form PI loops of Table 6.4 are the base layer, and the tracker and the gate run above them as a copy of the transport driven by measured inputs.

What Is Modernizing It

Three developments bear on this line. The first is the process model inside the control strategy. The 2023 guidance allows process models in commercial control, including the diversion strategy, and defines a soft sensor as a model used in place of a physical measurement to estimate a variable or attribute from measured data [20]. The tracker of this chapter is such a model. The second is real-time release. Continuous manufacturing lends itself to judging the quality of a product from process data, and the volume of data supports multivariate statistical process control, which tracks fewer charts and can detect a changed relation between parameters and quality [17]. The third is advanced control. The same review names model predictive control among the areas in which research should continue [17].

> **IN PRACTICE** **Feeder and Analyzer Limits Tied to the Transport**
>
> The ICH Q13 guidance, published by the United States regulator in March 2023, works through a continuous direct compression line as an example. Its residence-time distributions were measured by replacing the API with a tracer of very similar flow properties. The feeders run gravimetrically except while they refill, and the duration and mass of a refill were studied to limit its effect. The analyzer sits in the press feed frame. The diversion strategy links the feeder and analyzer limits to the distribution from the feeders to the analyzer and to that from the probe to the diversion point, and each predefined segment of a batch carries a unique code [20]. The tracker and the gate of this chapter play those parts.

6.8 Limitations

The mixers are ideal tanks in series with fixed residence times and no segregation, so the tracker's structure is exact and its only errors come from its inputs. A real tracking model is fitted to measured residence-time distributions and carries their uncertainty [5]. The tracker is also an open-loop estimator driven by measured inputs. A moving-horizon estimator that reconciles the feeders' mass measurements [19], of the kind set out in Section 1.4.7, would sharpen exactly the estimate that a refill spoils.

OPEN CHALLENGE Where One Batch Ends

Back-mixing smooths the feeders' fluctuations, and a line in plug flow could not filter them, so a robust continuous line must mix along its length. The same back-mixing blurs the boundary between batches, so a recall of one batch may have to include its neighbors on both sides [16]. A 2015 review calls back-mixing a possible physical limitation to material traceability [17]. The tracker of this chapter knows the plant's mixing exactly, which a real one cannot. Progress would be a diversion rule that carries the uncertainty of measured distributions into the decision.

Granulation is a size proxy, drying is a lumped law with illustrative coefficients, and the press is a power-law compaction with a hardness proxy. The feed-frame NIR is a noisy, delayed reading of the whole feed frame, much quieter than NIR on a commercial line [10].

The recipe and its limits are illustrative. The simulator was built to teach equipment control, its quality proxies and limits have not had quality-assurance review, and nothing in it is regulatory or quality-assurance guidance.

6.9 Conclusion

A continuous tableting line controls the tablet weight at the press and fixes the composition at the feeders twelve minutes upstream, and the simulator described here records what each tablet truly contains apart from what every instrument reports. The force reports the weight magnified eightfold and the composition not at all, and in case 2 a feeder refill keeps the true mean weight within 0.7 mg of 250 mg while its package reaches the press at 104.77 % LC, and the NIR gate catches the first over-dosed tablet with less than a second to spare. The refill feedforward holds the feed frame to 100.39 % LC. The dryer turns the stream into packages of 2160 tablets and its moisture loop learns once per cell, so humid air leaves cells at 2.59 % LOD with the cascade alone and 2.24 % with the humidity feedforward. A 40 s analyzer delay lets 816 under-dosed tablets through a gate that the tracker closes, and a 90 s stop of the API feeder can only be rejected, 7152 tablets with none out of limits accepted. In the graded shift the three settings take the accepted out-of-limit tablets from 194 to none with 82 992 accepted. The model assumes ideal mixers, lumped drying and compaction and a quiet analyzer. A fuller model would add segregation and flow-dependent residence times, a fitted tracker with its uncertainty, a sampling model for the NIR and a state estimator on the feeders. The application puts in a browser a line on which a tablet at the right weight can be seen to carry the wrong dose.

KEY POINTS

- A tablet at its target weight need not carry its target dose: the force sees how much powder fills a die, not what the powder is.
- The composition is fixed at the feeders and carried for a mean of 707 s, cut by the dryer into packages of 2160 tablets.
- A feeder refill keeps the true mean weight within 0.7 mg of 250 mg while its package reaches the press at 104.77 % LC.
- A 40 s analyzer delay lets 816 under-dosed tablets through a gate that the tracker closes, because the tracker sees the bridge on the load cells.

Where else this applies. Chapter 7 diverts material by the same residence-time tracking, Volume 12, Chapter 12 meets the refill of a loss-in-weight feeder on a polymer extruder, and Chapter 10 adds a humidity feedforward to a moisture result that arrives late.

References

[1] T. Higuchi, E. Nelson, and L. W. Busse, "The physics of tablet compression. III. Design and construction of an instrumented tableting machine," *Journal of the American Pharmaceutical Association (Scientific ed.)*, vol. 43, no. 6, pp. 344–348, 1954.

[2] H. Leuenberger, "New trends in the production of pharmaceutical granules: batch versus continuous processing," *European Journal of Pharmaceutics and Biopharmaceutics*, vol. 52, no. 3, pp. 289–296, 2001.

[3] P. V. Danckwerts, "Continuous flow systems," *Chemical Engineering Science*, vol. 2, no. 1, pp. 1–13, 1953.

[4] Y. Gao, A. Vanarase, F. Muzzio, and M. Ierapetritou, "Characterizing continuous powder mixing using residence time distribution," *Chemical Engineering Science*, vol. 66, no. 3, pp. 417–425, 2011.

[5] M. C. Martinetz, A.-P. Karttunen, S. Sacher, P. Wahl, J. Ketolainen, J. G. Khinast, and O. Korhonen, "RTD-based material tracking in a fully-continuous dry granulation tableting line," *International Journal of Pharmaceutics*, vol. 547, no. 1-2, pp. 469–479, 2018.

[6] E. I. Keleb, A. Vermeire, C. Vervaet, and J. P. Remon, "Continuous twin screw extrusion for the wet granulation of lactose," *International Journal of Pharmaceutics*, vol. 239, no. 1-2, pp. 69–80, 2002.

[7] C. Vervaet and J. P. Remon, "Continuous granulation in the pharmaceutical industry," *Chemical Engineering Science*, vol. 60, no. 14, pp. 3949–3957, 2005.

[8] S. T. F. C. Mortier, T. De Beer, K. V. Gernaey, J. P. Remon, C. Vervaet, and I. Nopens, "Mechanistic modelling of fluidized bed drying processes of wet porous granules: A review," *European Journal of Pharmaceutics and Biopharmaceutics*, vol. 79, no. 2, pp. 205–225, 2011.

[9] K. Järvinen, W. Hoehe, M. Järvinen, S. Poutiainen, M. Juuti, and S. Borchert, "In-line monitoring of the drug content of powder mixtures and tablets by near-infrared spectroscopy during the continuous direct compression tableting process," *European Journal of Pharmaceutical Sciences*, vol. 48, no. 4-5, pp. 680–688, 2013.

[10] J. M. Vargas, S. Nielsen, V. Cárdenas, A. Gonzalez, E. Y. Aymat, E. Almodovar, G. Classe, Y. Colón, E. Sanchez, and R. J. Romañach, "Process analytical technology in continuous manufacturing of a commercial pharmaceutical product," *International Journal of Pharmaceutics*, vol. 538, no. 1-2, pp. 167–178, 2018.

[11] W. E. Engisch and F. J. Muzzio, "Feedrate deviations caused by hopper refill of loss-in-weight feeders," *Powder Technology*, vol. 283, pp. 389–400, 2015.

[12] M. Kirchengast, S. Celikovic, J. Rehrl, S. Sacher, J. Kruisz, J. Khinast, and M. Horn, "Ensuring tablet quality via model-based control of a continuous direct compaction process," *International Journal of Pharmaceutics*, vol. 567, p. 118457, 2019.

[13] R. Singh, A. Sahay, K. M. Karry, F. Muzzio, M. Ierapetritou, and R. Ramachandran, "Implementation of an advanced hybrid MPC–PID control system using PAT tools into a direct compaction continuous pharmaceutical tablet manufacturing pilot plant," *International Journal of Pharmaceutics*, vol. 473, no. 1-2, pp. 38–54, 2014.

[14] A. Bhaskar, F. N. Barros, and R. Singh, "Development and implementation of an advanced model predictive control system into continuous pharmaceutical tablet compaction process," *International Journal of Pharmaceutics*, vol. 534, no. 1-2, pp. 159–178, 2017.

[15] M. Ierapetritou, F. Muzzio, and G. Reklaitis, "Perspectives on the continuous manufacturing of powder-based pharmaceutical processes," *AIChE Journal*, vol. 62, no. 6, pp. 1846–1862, 2016.

[16] W. Engisch and F. Muzzio, "Using residence time distributions (RTDs) to address the traceability of raw materials in continuous pharmaceutical manufacturing," *Journal of Pharmaceutical Innovation*, vol. 11, no. 1, pp. 64–81, 2016.

[17] S. L. Lee, T. F. O'Connor, X. Yang, C. N. Cruz, S. Chatterjee, R. D. Madurawe, C. M. V. Moore, L. X. Yu, and J. Woodcock, "Modernizing pharmaceutical manufacturing: from batch to continuous production," *Journal of Pharmaceutical Innovation*, vol. 10, no. 3, pp. 191–199, 2015.

[18] C. Badman, C. L. Cooney, A. Florence, K. Konstantinov, M. Krumme, S. Mascia, M. Nasr, and B. L. Trout, "Why we need continuous pharmaceutical manufacturing and how to make it happen," *Journal of Pharmaceutical Sciences*, vol. 108, no. 11, pp. 3521–3523, 2019.

[19] F. Destro, S. García Muñoz, F. Bezzo, and M. Barolo, "Powder composition monitoring in continuous pharmaceutical solid-dosage form manufacturing using state estimation – proof of concept," *International Journal of Pharmaceutics*, vol. 605, p. 120808, 2021.

[20] U.S. Food and Drug Administration, Center for Drug Evaluation and Research and Center for Biologics Evaluation and Research, "Q13 continuous manufacturing of drug substances and drug products. Guidance for industry," ICH harmonised guideline, https://www.fda.gov/media/165775/download, Mar. 2023.

[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.

6.10 Exercises

6.1 Concept check · 2 min

In case 4 an excipient lot 8 % denser reaches the press with Weight control Off. The tablets reach 265.1 mg while the blend in the feed frame stays between 99.30 and 100.12 % LC. Which instrument on the line can see that these tablets carry too much API?

(a) The feed-frame NIR, because the dose is the API share of the blend it scans
(b) The force transducer, because a heavier tablet needs more force at the same gap
(c) The tracker, because it follows the new lot from the excipient load cells
(d) The excipient load cells, because the denser lot raises the feeder's rate

6.2 Concept check · 1 min

A feeder error that lasts 30 s reaches the feed frame about twelve minutes later as a step in composition lasting one or two dryer packages, not as a short pulse. What gives it that shape?

(a) The blender's three tanks in series spread every error over about twelve minutes
(b) The feed-frame NIR averages its scans over the 10 s of its delay
(c) Each dryer cell mixes its 90 s into one package, which the hopper keeps in order
(d) The press hopper is a well-mixed 1.0 kg tank that blends neighboring packages

6.3 Concept check · 2 min

In case 6 the drying air turns from 6 to 16 g/kg at 300 s. The worst cell leaves at 2.59 % LOD with the moisture cascade alone and at 2.24 % with the humidity feedforward. Why does the feedforward do better?

(a) It raises the gain of the moisture loop, so each cell's reading moves the air further
(b) It replaces the dryer NIR with the hygrometer, which reports each cell's moisture sooner
(c) It holds the inlet air at 80 °C, the highest the API tolerates, until the cells are dry
(d) It moves the air at the hygrometer's reading, so cells already drying are partly corrected

6.4 Theory and calculation · 6 min

Use Equation (6.1) with $N_s = 24$, $m_t = 250\,\text{mg}$, $x_A = 10.0\,\%$, $x_L = 1.0\,\%$ and $\text{LOD}^* = 2.0\,\%$.

(a) Compute the API share r_A of the powder feed, the powder feed F_s at 60 and at 72 rpm, and the API rate at each speed. *Checkpoint:* $r_A = 0.10306$. F_s = 20.96 and 25.15 kg/h, API 2.160 and 2.592 kg/h.

(b) A dryer cell takes the granulator's outflow for 90 s at 60 rpm. Compute its solids, the finished blend they become with lubricant and moisture, and the number of tablets that blend makes. *Checkpoint:* 0.524 kg of solids, 0.540 kg of blend, 2160 tablets.

(c) How many tablets does the press make during a 10 s and a 40 s analyzer delay at 60 rpm, and during the 40 s delay at 72 rpm? *Checkpoint:* 240, 960 and 1152 tablets.

6.5 Theory and calculation · 8 min

A tablet's dose relative to its blend is $(1+\sigma_m z_m)(1+\sigma_a z_a)$ at a 250 mg mean weight, Equation (6.7), with independent standard normal deviates cut off at ± 3 and $\sigma_m = \sigma_a = 0.6\,\%$.

(a) The variance of a standard normal cut off at ± 3 is $1 - 6\varphi(3)/[2\Phi(3) - 1]$, with $\varphi(3) = 0.004432$ and $\Phi(3) = 0.99865$. Compute the standard deviation of the dose ratio. *Checkpoint:* 0.837 %.

(b) Compute the largest and the smallest dose ratio, and the range of blend dose inside which no tablet can fall outside 95 to 105 % LC. *Checkpoint:* +3.63 and −3.57 %. Blend from 98.52 to 101.32 % LC.

(c) The binder interruption of case 5 gives granules of 191 µm. With $\varphi = 1 - d_{50}/d_{ref}$ and the widening by fines of Section 6.3, compute σ_m and σ_a, and say why the feed-frame NIR cannot protect these tablets. *Checkpoint:* $\varphi = 0.363$, $\sigma_m = 1.25\,\%$, $\sigma_a = 2.34\,\%$.

6.6 Theory and calculation · 7 min

Use Equations (6.2), (6.9) and (6.11) for the API feeder, with $\beta_A = 20\,\%$, $\delta_A = 8\,\%$, $\tau_D = 90\,\text{s}$ and $t_R = 20\,\text{s}$. A refill starts at 25 % fill and ends at 80 %. Ignore the powder fed during the refill.

(a) With the screw speed frozen through the refill, compute $D_A(t_R)$ and the ratio of the rate at the end of the refill to the rate before it. *Checkpoint:* $D_A = 0.0717$, ratio 1.191.

(b) With the refill feedforward on, the command is divided by the identified profile of Equation (6.11), which is 5 % low on both terms, and the loop had settled on its setpoint before the refill with the feedforward in place. Find the ratio of the delivered rate at the end of the refill to the setpoint. *Checkpoint:* 1.0085, 0.85 % high.

(c) The chapter reports that, with the feedforward on, the rate runs at most 2.77 % high while the feeder is volumetric. Name two causes of the difference that part (b) leaves out. *Checkpoint:* The feed-factor variation, which a held screw cannot follow, and the 1 s lag of the screw behind a falling command.

6.7 App challenge · 8 min

Load scenario 2, Feeder refill, with the default seed and the settings as loaded (Refill feedforward Off). Run at a speed of 30 until the clock passes 1200 s, about 40 s of wall-clock time, pause and export the CSV file.

(a) In the row at 1200 s, read `metric_rejected`, `metric_false_rejects`, `metric_escapes` and `metric_oos_made`. *Checkpoint:* 2712, 2103, 0 and 609.

(b) Find the highest `api_ff_pct` up to 1200 s and the lowest and highest `m_true_mg`. What do the two ranges say about what the force-based weight loop protects? *Checkpoint:* 104.76 % LC at 810 s, within 0.05. Weight from 249.64 to 250.55 mg, within 0.05 mg.

(c) Set Refill feedforward On, press Reset and run to 1200 s again. Read `metric_rejected` and the highest `api_ff_pct`. *Checkpoint:* 0 rejected, 100.39 % LC within 0.05.

6.8 App challenge · 8 min

Load scenario 6, Dryer humidity change, with the default seed and the settings as loaded: Inlet air setpoint from Cascade and Humidity feedforward Off. Run at a speed of 60 to 2400 s, about 40 s of wall-clock time, and export the CSV file. Press reload, set Inlet air setpoint from to Fixed, run to 2400 s and export again. Press reload, set Humidity feedforward On, run to 2400 s and export a third time.

(a) For each run, find the highest `lod_meas_pct`, the dryer NIR's reading of each cell at discharge, and `metric_heater_kWh` at 2400 s. *Checkpoint:* Cascade 2.56 % and 5.72 kWh. Fixed 3.16 % and 4.81 kWh. Feedforward 2.19 % and 5.85 kWh. Within 0.01.

(b) For the fixed-air run, read `metric_rejected`, `metric_rej_track`, `metric_rej_nir`, `metric_accepted` and `metric_made` at 2400 s, and compute the share of the output kept. *Checkpoint:* 37 704 rejected, 35 928 on tracking and 1776 on the NIR, 19 896 of 57 600 kept, 34.5 %.

(c) The chapter reports that, with the cascade alone, the worst cell leaves at 2.59 % LOD at 480 s. Explain why the highest reading in your first file is lower. *Checkpoint:* The dryer NIR reads each cell with 0.05 % of noise.

6.9 App challenge · 7 min

Load the ninth scenario, Challenge, one shift, three upsets, with the default seed and the settings as loaded. Run at a speed of 120 to 3600 s, about 30 s of wall-clock time, and export the CSV file. In the row at 3600 s, read `metric_accepted`, `metric_rejected`, `metric_escapes` and `metric_worst_api_pct`.

(a) Report the four values and say whether the shift meets its target. *Checkpoint:* 71 952, 14 448, 194 and 9.39 %, within 0.01. Not met.

(b) Set Refill feedforward, Humidity feedforward and Reject by material tracking On, press Reset, run to 3600 s and read the same columns. *Checkpoint:* 82 992, 3408, 0 and 3.72 %.

(c) Set Reject by material tracking Off, press Reset and run again. Read the same columns, and say which upset sends out-of-limit tablets through the gate and why the NIR cannot stop them. *Checkpoint:* 83 304, 3096, 122 and 8.62 %. The bridge at 1800 s.

6.10 Application area · 3 min

The model's feed-frame NIR reads the blend with 0.3 % LC of noise. The chapter cites a commercial line whose blend readings had a standard deviation of 2.17 % over a 28 h run, while 500 of its tablets had 0.4 %.

(a) If the NIR's noise were 2.17 % LC at one standard deviation, what share of readings of a blend exactly on target would fall outside the default guard band of ±2.0 % LC? *Checkpoint:* About 36 %.

(b) What would that do to the gate, and name two ways a plant can use such an analyzer for release. *Checkpoint:* The gate, which holds for 10 s after each out-of-band reading, would reject most of the output.

6.11 Application area · 4 min

In case 1 the line starts empty. With tracking on the gate opens at 748.5 s after 3444 rejected tablets carrying 85.0 g of API. With tracking off it opens at 673.5 s after 1644, and no tablet out of limits is accepted either way.

(a) What do the tracker's extra 75 s cost in tablets and in finished blend, and how does the API in the 3444 rejects compare with 3444 tablets at label claim? *Checkpoint:* 1800 tablets, 0.45 kg of blend. 85.0 g against 86.1 g at label claim.

(b) Give one reason a validated process may still keep the tracker's rule, and one cost of start-up that the simulation leaves out.

6.12 Review the agent · 5 min

An AI assistant is asked how to save the tablets that case 3 rejects, where the API stops for 90 s. It replies: "Add a makeup term to FIC-A. Integrate the API that the load cells show missing, 54 g, and feed it back at twice the setpoint for the 90 s after the arch collapses. The run then averages exactly 100 % LC, the tracker has nothing to reject, and the 7152 rejected tablets are saved." Find the flaw, and say how you would check it.

CHAPTER 7

Material Diversion by Residence-Time Tracking in a Simulated Continuous Flow Synthesis

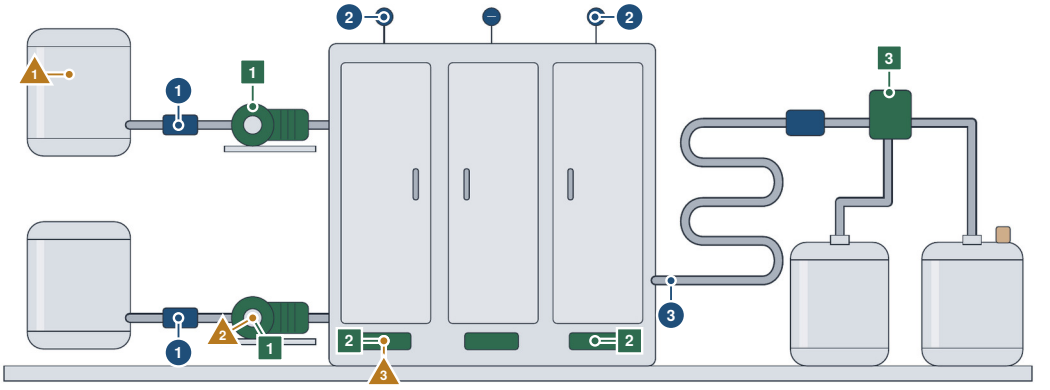

The controller moves 1 feed pumps for substrate A and reagent B 2 circulators that heat or cool the reactor zones 3 divert valve to the accepted or the divert tank

The controller reads 1 feed flow meters 2 zone thermometers 3 analyzer sample point at the coil inlet

What disturbs it 1 a feed A drum stronger than its label 2 an air lock in pump B 3 lost cooling in zone 1

The analyzer samples the product before the long coil, which holds each part until its verdict reaches the valve. Illustration, not to scale.

IN PLAIN TERMS **Following a Medicine Through a Pipe**

The main idea. A continuous line makes a medicine ingredient by pumping two liquids through a warm tube, where they react. A tester checks a sample every few minutes, and a valve then sends each part to a good tank or to waste.

The challenge. The tube spreads a short upset out, so bad product reaches the valve early and stays long. Each test result comes after its sample has moved on, so a long pipe holds the product until the result is in.

Why it matters. Only a line that knows where each liter is can throw out the bad part and keep the rest.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the fact on run time (Section 7.1), the history segment (Section 7.2), the try-its on the air lock and the slow analyzer and the fact on method time (Section 7.6) and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 7.6), Technology in Practice (Section 7.7), Limitations (Section 7.8) and the open challenge

The control question A continuous synthesis has no batch to hold back. When a dispersive reactor spreads an upset and the analyzer reports on material that has already moved on, which part of the stream must be diverted, and does a plug-flow delay or residence-time tracking find it?

Learning objectives

- Compute the mean residence times, spreads and coil quantiles of a line of well-mixed cells, and the plug-flow time from mixer to valve.
- Size a hold-up coil against an analyzer method and a throughput with $t_{det} = 2t_{run} + t_{tr}$, and predict how often a change escapes.
- Explain why a plug-flow delay releases off-specification material at both ends of a spread upset, and measure it on the same material in the app.
- Distinguish the upsets a soft sensor sees through its meters from those only the analyzer sees.
- Separate what a predictive controller gains from its plan and from its target when it is compared with a retuned quality trim.

Prerequisites Sections 1.4.4, 1.4.5, 1.4.8, 1.6.2 and 1.6.3. Chapter 4 for material followed on the measured flow.

The app https://apmonitor.com/apps/continuous-api/

See also Volume 8, Chapter 13 for a flow diversion valve under a supervisor that judges each parcel's history, Volume 8, Chapter 9 for results that arrive after later product has been made, and Volume 8, Chapter 2 for heat removal as the limit on how fast an exothermic reaction may be fed.

7.1 Introduction

A batch of drug substance that goes wrong can be held, tested and rejected as one lot. A continuous line has no such unit. Its product leaves every minute, so when a reagent pump loses its prime or a feed drum is stronger than its label, the question becomes which part of the stream the fault reached, and whether all of that part went to waste. An industry roundtable has named continuous processing a research priority and surveyed its members about the hurdles that slowed its implementation [1], and a review of regulatory readiness has stated that the U.S. Food and Drug Administration supports it on science- and risk-based grounds [2]. Continuous-flow reactors with small channels suit the hazardous and fast exothermic steps of active pharmaceutical ingredient (API) synthesis, because their fast heat transfer prevents the accumulation of heat and the formation of hot spots [3].

The simulator of this chapter is a continuous reaction step of that kind at the scale of a laboratory flow system. Substrate A and reagent B are pumped at 100 mL/min in total, meet in a 15 mL mixer and react through four jacketed zones of 0.40 L, each with its own circulator and temperature loop. A quench stream destroys the unreacted reagent, an at-line high-performance liquid chromatograph (HPLC) samples the quenched stream every 3 min, and a 1.5 L hold-up coil carries the material to a three-way valve that sends it to an accepted tank or to a divert tank. An over-reaction of the product with excess reagent and a thermal decomposition of the substrate compete with the desired reaction, so conversion and impurities pull against each other in temperature, ratio and residence time. As an educational equipment-control simulation, the application uses illustrative kinetics, invented for teaching and standing for no real synthesis, and illustrative quality limits. Neither has had quality-assurance review, and nothing in the application or in this chapter is regulatory or quality-assurance guidance.

✦ DID YOU KNOW? A Run Measured in Months

In the worked example of a continuous drug-substance line with two plug-flow reactors in the ICH Q13 guidance, a risk assessment supports a proposed run time of up to several months [4]. The graded campaign of this chapter runs six hours.

Three properties separate this plant from a reactor with a temperature loop and a dead time. The reactor and the coil are distributed, so a disturbance at the mixer reaches the valve spread out. A four-minute air lock in the reagent pump sends off-specification material to the valve from 24.07 to 35.85 min after it begins, starting 4.58 min before the plug-flow time of 28.65 min and ending 3.20 min after the plug-flow image of the air lock has passed. The analyzer is late by construction, because each result describes a sample drawn one method run and one transport earlier, 3.5 min with the default method, and by then the sampled material has moved on toward the valve. Some disturbances, such as a drum made up 10 % stronger than its label, are seen only in the product.

The simulator addresses the belief that a disturbance reaches the outlet one plug-flow time after it enters and lasts there as long as it lasted at the inlet. That belief is the fixed dead time of a transfer function, and it is also the simplest rule for diverting product. The class of problem is general. Residence-time distributions have been used to define batches and to trace raw materials through a continuous direct compaction line, with the methods stated to extend to wet and dry granulation, hot melt extrusion and capsule filling [5], and to size the waste at the candidate diversion points of a continuous dry granulation line for four feeder upsets [6]. A pasteurizer's flow diversion valve, a tablet press with a delayed composition measurement and a refinery blend with an at-line analyzer have the same shape, and Volume 8, Chapter 13 and Chapter 6 treat the first two. However, the transport delay of an introductory course is a constant, and a constant delay represents neither the spread that a mixed volume gives a disturbance nor a verdict that arrives after the material has moved on. Sections 1.4.4 and 1.4.5 set out both departures on the running example of Chapter 1.

A concern about a gap between process control theory and practice had been voiced by 1950 [7], and a survey of good practice in control teaching centers its discussion on laboratories, software and assessment [8]. A continuous line asks its engineers to answer for every liter it made, and the answer rests on residence time and analyzer timing more than on the gains of any single loop.

This chapter compares, on one simulated flow line and on the same material, a plug-flow delay and residence-time tracking as rules for deciding what a continuous synthesis must divert, and it describes the line in enough detail to be rebuilt. Every method used is established: tanks-in-series residence-time models, ratio control, conditional integration, model-based soft sensors, diversion valves and model predictive control (MPC), whose industrial use has been surveyed [9]. The line has 1205 states. The plant and a measurement-driven copy of it share one set of equations, a supervisor follows material in volume and scores the unselected rule on the same material, and a truth record is kept apart from the instruments. The application pairs the two rules on identical material through eight cases and a six-hour campaign, and a ninth case asks how much of what an optional MPC achieves over the shipped quality trim is the plan and how much is the target it is tuned to.

Section 7.2 reviews the domain. Section 7.3 sets down the model of the line, its chemistry, jackets, coil, pumps and valve, and Section 7.4 works out its residence times, its operating window and the rule that sizes the coil against the analyzer. Section 7.5 describes the instruments, the loops, the soft sensor, the diversion supervisor and the predictive controller. Section 7.6 reports what each case and the campaign measured, including the comparison of the predictive controller with the quality trim and the worst cases under each structure. Section 7.7 sets the line beside the instruments, valves and control platforms of real plants, Section 7.8 states the limitations and Section 7.9 draws the conclusions.

7.2 Residence Time and Continuous Drug Manufacturing

The vocabulary for how a continuous vessel spreads what enters it dates from 1953. The distribution of residence times of a continuous flow system, and its measurement with a tracer, was set out that year [10]. In the same year the dispersion of a solute in slow laminar flow through a tube was traced to the combined action of the velocity profile and radial molecular diffusion [11], so the spread of a real tube depends on its diameter, its flow and the diffusivity of what it carries. A chain of well-mixed cells is the simplest model with an adjustable spread, and it is the one used here.

Continuous processing had been used for decades in other industries while pharmaceutical manufacturing was slow to adopt it, primarily because of regulatory uncertainty [5]. The batch is the unit of release, and the record, the testing and the filing are organized around it. A review of continuous processing in the industry appeared in 2005 [12]. An end-to-end continuous pilot plant then made a finished drug product from chemical intermediates in one integrated system [13]. A plant-wide control strategy was then designed for a system inspired by that pilot plant, as the history segment after this paragraph tells. Two plant-wide MPC designs for an end-to-end pilot plant, one from subspace identification and one from linearized differential-algebraic equations, were later evaluated on a nonlinear simulator of it under uncertain reaction kinetics [14]. A refrigerator-sized reconfigurable system later made four APIs to U.S. Pharmacopeia standards by continuous-flow synthesis and formulation [15]. A review of the transition found continuous drug product already used in approved processes, with continuous API steps under current good manufacturing practice lagging behind it in regulatory filings [16].

> ⧖ HISTORY **2013: A Control Strategy for a Whole Continuous Plant**
>
> The end-to-end pilot plant made a pharmaceutical product in a fully continuous fashion, from the synthesis of its intermediates to the coating of the drug product. A 2013 study designed a plant-wide control strategy for a system inspired by it. Its authors argued that, instead of the fixed recipes of batch work, control of a continuous plant must actively hold it within a safe region of operation, so the key challenge is a strategy for the plant as a whole rather than for isolated units. They split the plant into a hierarchy of levels, chose the loops at each level from sensitivities computed with a plant-wide dynamic model, and in simulation kept the critical quality attributes of the final product close to specification under significant and persistent disturbances [17].

Traceability is what the batch had provided for free. With the residence-time distribution as the tool, methods were proposed for defining a batch, tracing raw materials and choosing the frequency of the sensors on a continuous line [5]. Residence-time models of the units of a dry granulation line were built from tracer tests and used to compute how much material each candidate diversion point would reject in four feeder upsets [6]. The measurements come from process analytical technology, whose sensors and whose monitoring and control methods have been reviewed together with the regulatory view of them [18]. On flow reactors, inline infrared and nuclear magnetic resonance spectroscopy have been combined with an online chromatograph on one modular platform [19], and a perspective on quality by control holds that active process control based on process knowledge and model-based techniques is indispensable in continuous manufacturing [20]. The ICH Q13 guideline, issued in the United States as guidance for industry in March 2023, describes scientific and regulatory considerations specific to the continuous manufacturing of drug substances and drug products [4]. It is cited here as context only.

7.3 A Line of Well-Mixed Cells

At its design point the line feeds 47.62 mL/min of substrate A and 52.38 mL/min of reagent B, both 1.00 mol/L and preheated to 55 °C, which is a molar ratio of 1.10. A quench of 20 % of the feed brings the coil flow to 120 mL/min, and the product reaches the valve at 98.66 % assay with 0.54 % impurities and 0.392 mol/L of P. The feeds, the mixer, the four zones with their jackets and circulators,

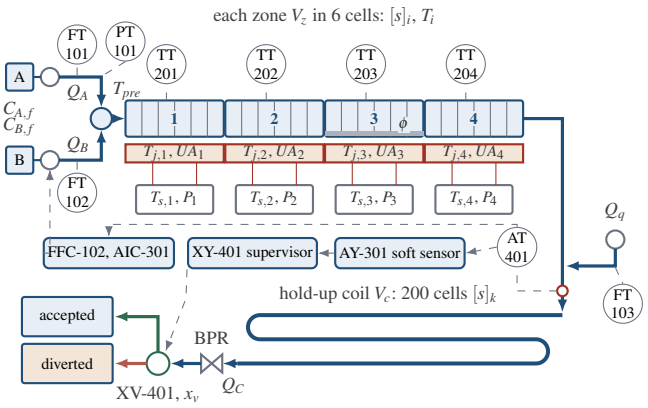

Figure 7.1 Continuous flow synthesis as modeled. Substrate A and reagent B of strengths $C_{A,f}$ and $C_{B,f}$ are pumped at Q_A and Q_B past FT-101 and FT-102 into the mixer at T_{pre}, with PT-101 on the pump discharge. Four zones of volume V_z, six cells each, exchange heat through jackets at $T_{j,z}$ (UA_z) with circulators at supply temperature $T_{s,z}$ and duty P_z, and TT-201 to TT-204 read the third cell of each zone. A deposit ϕ can grow in zone 3. The quench Q_q (FT-103) joins before the sample point of AT-401, and the coil flow Q_C passes the hold-up coil and the BPR to XV-401, which routes it to the accepted or the divert tank. The dashed lines are the signals of the soft sensor, the supervisor, the ratio station and the quality trim.

Table 7.1 What Each Line Variable Is, Its Design Value and Its Range

Variable	Role	Design value	Range
Q_{SP}	SP, total feed flow	100 mL/min	0–200
R_{SP}	SP, molar ratio B/A	1.10	0.9–1.5
$T_{z,SP}$	SP, zones 1–4	55, 60, 60, 60 °C	30–90
A_{SP}	SP, residual A (trim)	0.8 %	0.2–2.5
Q_m, R_m	CV, FT-101 + FT-102	100 mL/min, 1.10	—
$T_{z,m}$	CV, TT-201 to 204	55, 60, 60, 60 °C	—
α_m	CV, AT-401 residual A	0.80 %	—
u_A, u_B	MV, pump speeds	23.8, 26.2 %	0–100
$T^*_{s,z}$	MV, supplies	45.6, 62.3, 59.1, 60.1 °C	0–100
ΔR	MV, ratio trim	0	−0.10–0.30
x_v	MV of XV-401	accept	or divert
δ_A, $\dot{\delta}_A$	DV, strength, drift	0, 0 %/h	±15, ±20
$\delta_{A,as}$	DV, entered assay	0 %	±15
η_B	DV, pump B delivery	100 %	20–110
β_B	DV, FT-102 bias	0 %	±10
c_1	DV, zone 1 cooling	100 % of 200 W	10–100
t_{run}	DV, HPLC method run	3 min	1–15
r_ϕ	DV, deposit growth	0 %/h	0–100
$[s]_i$, T_i	states, 25 reaction cells	7 each	—
$[s]_k$	states, 200 coil cells	5 each	—
$T_{j,z}$, $T_{s,z}$	states	4 each	—
Q_A, Q_B, Q_q	states, pumps	47.6, 52.4, 20 mL/min	—

the quench, the sample point, the coil, the back-pressure regulator (BPR), the divert valve, the two tanks and each instrument appear in Figure 7.1.

Table 7.1 lists the variables by role. The state vector has 1205 elements: seven per reaction cell (A, B, P, I, D, an inert tracer X and the temperature) for the mixer and 24 zone cells, five per coil cell (A, P,

I, D and X) for 200 coil cells, four jacket and four supply temperatures, three pump deliveries, twelve tank contents and seven ledgers of feed, quench and energy. The discrete events are a four-minute air lock in pump B, a new assayed drum of feed A, a stop and a restart of the feeds, a trip reset, a cleaning of zone 3 and a tracer pulse of 10 mmol over 10 s at the mixer.

7.3.1 Chemistry and Rates

Three reactions are carried, as shown in Equation (7.1).

$$\begin{aligned} &\mathrm{A}+\mathrm{B}\rightarrow\mathrm{P}, \qquad \mathrm{P}+\mathrm{B}\rightarrow\mathrm{I}, \qquad \mathrm{A}\rightarrow\mathrm{D} \\ &r_1 = k_1[\mathrm{A}][\mathrm{B}], \qquad r_2 = k_2[\mathrm{P}][\mathrm{B}], \qquad r_3 = k_3[\mathrm{A}] \\ &k_j = k_{j,60}\exp\Big[-\frac{E_j}{R}\Big(\frac{1}{T+273.15}-\frac{1}{333.15}\Big)\Big] \end{aligned} \tag{7.1}$$

where P is the product, I the impurity made when the product meets excess reagent, D the impurity made when the substrate decomposes, T is in °C and $R = 8.314\,\mathrm{J/(mol\,K)}$. Negative concentrations are clipped to zero inside the rates. The activation energies rise from 60 through 95 to 160 kJ/mol, so a hotter reactor converts more of A and makes more of both impurities, the decomposition fastest of all. The heats of reaction are −120, −80 and −60 kJ/mol.

7.3.2 Reaction Cells, Jackets and Circulators

The mixer, cell 0 of volume V_{mix}, and the 24 zone cells of volume $V_z/6$ form one chain carried by the reactor flow $Q_R = Q_A + Q_B$, as shown in Equation (7.2).

$$\begin{aligned} \frac{\mathrm{d}[s]_i}{\mathrm{d}t} &= \frac{Q_R}{V_i}\big([s]_{i-1}-[s]_i\big)+\sum_{j=1}^{3}\nu_{sj}\,r_{j,i} \\ \rho c_p\frac{\mathrm{d}T_i}{\mathrm{d}t} &= \rho c_p\frac{Q_R}{V_i}(T_{i-1}-T_i)+\sum_{j=1}^{3}(-\Delta H_j)\,r_{j,i} \\ &\quad -\frac{UA_z}{6V_i}(T_i-T_{j,z}) \end{aligned} \tag{7.2}$$

where s is one of A, B, P, I, D and X, and ν_{sj} is its stoichiometric coefficient in reaction j. The inlet of the mixer is the mixed feed, $[\mathrm{A}]_{in} = Q_A C_A/Q_R$ and $[\mathrm{B}]_{in} = Q_B C_{B,f}/Q_R$ at T_{pre}, where the true strength $C_A = C_{A,f}(1+\delta_A/100)$ carries the unmeasured strength error δ_A in percent, the slider value plus a drift integrated at $\dot{\delta}_A$ and bounded at $\pm 15\,\%$. The jacket term acts on the six cells of zone z and not on the mixer. Each jacket is fed by a circulator that moves its bath toward the command $T^{*}_{s,z}$ of TIC-20z within power limits, as shown in Equation (7.3).

$$\begin{aligned} C_j\frac{\mathrm{d}T_{j,z}}{\mathrm{d}t} &= \sum_{i\in z}\frac{UA_z}{6}(T_i-T_{j,z})+F_j c_p\,(T_{s,z}-T_{j,z}) \\ C_b\frac{\mathrm{d}T_{s,z}}{\mathrm{d}t} &= F_j c_p\,(T_{j,z}-T_{s,z})+P_z \\ P_z &= \mathrm{clamp}\big(P_z^{req},\,-P_{c,z},\,P_h\big) \\ P_z^{req} &= \frac{C_b}{\tau_c}(T^{*}_{s,z}-T_{s,z})+F_j c_p\,(T_{s,z}-T_{j,z}) \end{aligned} \tag{7.3}$$

Here P_z^{req} is the duty that the circulator's internal controller requests. Away from its limits the supply therefore follows its command as a first-order lag of $\tau_c = 60\,\mathrm{s}$. In this chapter τ_c is the circulator's response time, not the closed-loop time constant of Section 1.3. The cooling capacity $P_{c,z}$ is 200 W for zone 1, scaled by the disturbance c_1, and 100 W for zones 2 to 4, and every circulator can heat by $P_h = 300\,\mathrm{W}$.

7.3.3 Quench, Coil and Quality

At the quench the unreacted reagent is destroyed and the other solutes are diluted into the coil flow $Q_C = Q_R + Q_q$. The coil is a second chain of 200 cells of volume $V_k = V_c/200$, as shown in Equation (7.4).

$$[s]_q = \frac{Q_R}{Q_C}[s]_{24}, \qquad \frac{\mathrm{d}[s]_k}{\mathrm{d}t} = \frac{Q_C}{V_k}\left([s]_{k-1} - [s]_k\right) \tag{7.4}$$

where s is A, P, I, D or X, $[s]_q$ feeds coil cell 0, and nothing reacts after the quench. The quality of a quenched composition follows Equation (7.5).

$$a = \frac{100[\mathrm{P}]}{S}, \quad \iota = \frac{100([\mathrm{I}]+[\mathrm{D}])}{S}, \quad \alpha = \frac{100[\mathrm{A}]}{S} \tag{7.5}$$

where $S = [\mathrm{A}]+[\mathrm{P}]+[\mathrm{I}]+[\mathrm{D}]$, a is the assay, ι the total impurities and α the residual A, all area percentages of the solutes as an in-process HPLC reports them, so dilution does not change them. The strength is [P] itself, which catches solvent at startup. Material is in specification when $a \geq a_{min}$, $\iota \leq i_{max}$ and $[\mathrm{P}] \geq s_{min}$, with illustrative limits of 97.5 %, 1.0 % and 0.33 mol/L. At the design point the sample point carries 98.66 % assay, 0.80 % residual A and 0.54 % impurities, 0.344 % of it I and 0.195 % D, at 99.20 % conversion of A.

7.3.4 Pumps, Pressure and the Zone 3 Deposit

The feed and quench pumps are first-order drives, and the pump discharge pressure is the BPR setting plus laminar drops that scale with the flows, as shown in Equation (7.6).

$$\begin{aligned}
\tau_p \frac{\mathrm{d}Q_A}{\mathrm{d}t} &= \frac{u_A}{100}Q_{A,max} - Q_A \\
\tau_p \frac{\mathrm{d}Q_B}{\mathrm{d}t} &= \eta_B \frac{u_B}{100}Q_{B,max} - Q_B \\
\tau_p \frac{\mathrm{d}Q_q}{\mathrm{d}t} &= \min(q\,Q_m, Q_{q,max}) - Q_q \\
P &= P_{BPR} + \Delta P_c \frac{Q_C}{Q_{C,0}} + \Delta P_z \frac{Q_R}{Q_{R,0}}\,(4 + 40\,\phi)
\end{aligned} \tag{7.6}$$

where u_A and u_B are pump speeds in percent, η_B is pump B's delivery factor (0.1 during the air lock, times the delivery disturbance), $q = 0.20$ is the quench ratio applied to the measured feed flow Q_m, $Q_{C,0} = 120\,\mathrm{mL/min}$, $Q_{R,0} = 100\,\mathrm{mL/min}$, and $\phi \in [0,1]$ is the deposit fraction in zone 3. The deposit grows as $\mathrm{d}\phi/\mathrm{d}t = r_\phi/(100 \times 3600)$ per second with r_ϕ in percent of its full extent per hour, it multiplies the hydraulic resistance of zone 3 by $1 + 40\phi$, 41 times its clean value at $\phi = 1$, and it reduces the jacket transfer of zone 3 to $UA_z(1 - 0.5\phi)$. At design the discharge is 4.30 bar. Above $P_{trip} = 10\,\mathrm{bar}$ a relay stops both feed pumps and the quench until an operator resets the trip and restarts the feeds.

7.3.5 Valve, Tanks and the Truth Record

The divert valve XV-401 splits the coil outflow in proportion to its position x_v, which moves toward 1 (accept) or 0 (divert) at $1/t_{XV}$ per second with $t_{XV} = 4\,\mathrm{s}$. The tanks integrate the volume and the moles of each solute that reach them. After every step the true composition of the last coil cell is judged against the limits, and $x_v Q_C \Delta t$ joins the off-specification volume released whenever that composition fails. A second valve position, driven by the unselected tracking rule, is scored against the same composition, and no controller reads either record. The A backbone ($[\mathrm{A}]+[\mathrm{P}]+[\mathrm{I}]+[\mathrm{D}]$) and the B content ($[\mathrm{B}]+[\mathrm{P}]+2[\mathrm{I}]$) are conserved by all three reactions, so their ledgers (fed, minus the change held in the line, the two tanks and the reagent the quench destroyed) close to within 1×10^{-14} of what was fed in every run made for this chapter, and so does the energy of the cells, jackets and baths.

Table 7.2 Parameters of the Line. Visible Parameters Appear as Sliders in the Model Panel, Hidden Ones Do Not

Symbol	Meaning	Value	Class
V_{mix}, V_z, V_c	mixer, zone, coil	15 mL, 0.40, 1.5 L	hidden
T_{pre}	feed preheat	55 °C	visible
$C_{A,f}$, $C_{B,f}$	feed label strengths	1.00 mol/L each	visible
q	quench per feed flow	20 %	visible
ρc_p	stream heat capacity	1.8 MJ/(m^3 K)	hidden
$Q_{A,max}$, $Q_{B,max}$	pumps at full speed	200 mL/min each	visible
$Q_{q,max}$, τ_p	quench pump, drive lag	100 mL/min, 2 s	hidden
t_{XV}	valve full stroke	4 s	visible
$k_{1,60}$	A + B → P at 60 °C	4 L/(mol min)	visible
$k_{2,60}$	P + B → I at 60 °C	0.0032 L/(mol min)	visible
$k_{3,60}$	A → D at 60 °C	0.0015 min^{-1}	visible
E_1, E_2, E_3	activation energies	60, 95, 160 kJ/mol	visible
ΔH_1	heat of R1	−120 kJ/mol	visible
ΔH_2, ΔH_3	heats of R2, R3	−80, −60 kJ/mol	hidden
UA_z	jacket, per zone	12 W/K	visible
$F_j c_p$	jacket fluid flow	123 W/K	hidden
C_j, C_b	jacket, bath capacities	1.5, 2.0 kJ/K	hidden
τ_c, P_h	circulator response, heating	60 s, 300 W	hidden
$P_{c,1}$, $P_{c,2..4}$	cooling capacities	200, 100 W	visible
P_{BPR}	BPR setting	3.0 bar	visible
ΔP_z, ΔP_c	zone, coil drops	0.20, 0.50 bar	hidden
P_{alarm}, P_{trip}	discharge limits	8, 10 bar	visible
t_{tr}	sample transport	30 s	visible
σ_α, σ_ι	analyzer noise, A and I + D	0.03, 0.015 %	hidden
σ_s	strength noise, of reading	0.5 %	hidden
τ_{TT}, τ_{FT}	thermometer, meter lags	5, 1 s	hidden
a_{min}, i_{max}	assay, impurity limits	97.5, 1.0 %	visible
s_{min}	strength limit	0.33 mol/L	visible

7.3.6 Initial State, Integration and Checks

Table 7.2 lists the parameters, where ΔP_z is the drop per clean zone at 100 mL/min and ΔP_c the coil drop at 120 mL/min. Every scenario except startup and the campaign begins on the steady line, solved cell by cell. Each cell's balances reduce to three reaction extents and an energy balance solved by Newton's method, and a secant search on each supply temperature puts the thermometer cell of its zone on the setpoint. The supplies then sit at 45.58, 62.27, 59.12 and 60.06 °C, and the jackets at 46.39, 62.03, 59.21 and 60.05 °C. Startup fills the line with solvent at the zone temperatures and no flow. The engine (Section 1.6.1) integrates all 1205 states by fourth-order Runge-Kutta with a 1 s step. Halving the step leaves the volume diverted in case 3 at 1.5277 L (within 2 parts in 10^8), because the valve is commanded only at control executions, which fall at the same instants, and moves its worst released assay by 0.009 points (97.778 against 97.769 %). The cell counts are the declared dispersion rather than a discretization, and 3, 6, 12, 24 and 48 cells per zone give design assays of 98.39, 98.66, 98.78, 98.84 and 98.87 %. A 10 s tracer pulse at the mixer reproduces the declared means and variances of the reactor and of the reactor and coil together within 0.05 % at 100 mL/min and 0.07 % at 150 mL/min.

7.3.7 Assumptions

The residence-time distribution is declared as tanks in series, 25 cells in the reactor and 200 in the coil. A real tube disperses by the mechanism of laminar dispersion [11], a coil adds secondary flow

Table 7.3 Steady Assay and Impurities (%) Against Throughput and the Zones 2–4 Setpoint, With Zone 1 at 55 °C and a Ratio of 1.10. An Asterisk Marks an Assay Below Its Limit

Feed	Zone 1 duty	55 °C	60 °C	65 °C
60 mL/min	59 W	99.01 / 0.61	99.09 / 0.76	98.97 / 0.98
100 mL/min	99 W	98.10 / 0.42	98.66 / 0.54	98.90 / 0.72
130 mL/min	128 W	97.14* / 0.35	98.05 / 0.46	98.57 / 0.64
160 mL/min	155 W	96.09* / 0.31	97.32* / 0.41	98.13 / 0.59

and a static mixer adds its own, so a real line has to be measured with a tracer. The conclusions depend on the declared spread, as the cell-count sensitivity above shows.

The chemistry is three reactions with invented constants in one liquid phase of constant ρc_p, with no volume change on reaction and no precipitation. The zone 3 deposit changes only the hydraulic resistance and the jacket transfer, so its own volume and mass are neglected.

The preheater is ideal, the mixer and the quench tee mix instantly, the quench destroys all unreacted B at once, and the coil is not modeled thermally. A real quench takes a finite time, during which the over-reaction continues.

Each jacket is well mixed, each circulator is a bath with a 60 s internal controller and fixed power limits, and there are no heat losses to the room.

The analyzer takes a point sample with Gaussian noise and has no calibration drift, failed injections or sample-line holdup, all of which a real at-line chromatograph adds.

The soft sensor uses the plant's declared kinetics. Its only errors come from what no instrument measures, from noise and lags, and from its corrections, while a real soft sensor also carries model error and would need a wider margin.

7.4 Dispersion and the Operating Window

The two chains of Equations (7.2) and (7.4) declare the dispersion of the line. A chain of N equal well-mixed cells of total volume V passes an impulse as the Erlang distribution of Equation (1.7) in Section 1.4.4, with mean $\tau = V/Q$ and variance $\sigma^2 = \tau^2/N$. The reactor is one 15 mL cell and 24 cells of 66.7 mL, so at 100 mL/min its mean residence is 16.15 min and its standard deviation $[(V_{mix}/Q_R)^2 + 24(V_z/6Q_R)^2]^{1/2} = 3.27$ min. The 200 coil cells of 7.5 mL give 12.5 min at 120 mL/min with a standard deviation of 0.88 min, close to plug flow. The 0.1 and 99.9 percent points of the coil's distribution fall at 0.7956 and 1.2328 coil volumes, 9.95 min and 15.41 min. The reactor spreads a disturbance, and the coil mostly delays it. The plug-flow time from the mixer to the valve is the sum of the two means, 28.65 min, and every coil window is a volume, so a throughput change moves the arrival of material in time without moving it in the coil.

🕮 **Pause.** Try Exercise 7.1 before reading on.

Total flow sets both the residence time and the heat the first zone must remove, as listed in Table 7.3. Complete conversion at 100 mL/min would release 95.2 W, an adiabatic rise of 31.7 K. The reaction starts in the mixer, which leaves at 62.1 °C and 22.3 % converted, so the hot spot is the first cell of zone 1 at 62.3 °C while TT-201, two cells downstream, reads its 55.0 °C setpoint. Zone 1's circulator removes 99 W of its 200 W. The zone leaves at 49.2 °C and 86.1 % converted, having formed 68 % of the decomposition product D, and zone 2's circulator then heats the stream by 29.1 W, while 68 % of the over-reaction product I forms in zones 2 to 4. The recipe is a map against throughput rather than one temperature. At 160 mL/min the design recipe gives 97.32 % assay with 2.26 % residual A, and zones 2 to 4 at 65 °C restore 98.13 % while impurities rise from 0.41 to 0.59 %. At 60 mL/min the longer residence lets the over-reaction grow to 0.76 % impurities. More reagent converts more of A and makes more I, since a ratio of 1.30 gives 99.17 % assay and 0.79 % impurities while 1.00 leaves 3.56 % residual A and 95.97 % assay. Zone 1 at 70 °C instead of 55 gives 1.105 % impurities, half of it D.

The coil has a second job. It holds material between the sample point and the valve while the analyzer decides, so its holding time has to cover the analyzer's worst case. A change that begins just after a sample is drawn is caught by the next sample, one method run t_{run} later, whose result is published one transport t_{tr} and one more run after that. Section 1.4.5 derives the same worst case on the running example. The condition for catching every such change is Equation (7.7).

$$t_{det} = 2\,t_{run} + t_{tr} \le \frac{V_c}{Q_C} \tag{7.7}$$

The comparison is with the coil's mean holding time because a change that has passed through the reactor is already smooth, and its crossing of a limit travels the nearly plug-flow coil at the mean. If the crossing falls at a uniformly distributed phase of the sampling cycle, the fraction of phases at which the change escapes is $(t_{det} - V_c/Q_C)/t_{run}$ whenever that is positive. With the default method $t_{det} = 6.5$ min against 12.5 min of coil. With a 10 min method it is 20.5 min, and Section 7.6.6 measures the cost. Taking the coil's leading edge, 9.95 min at 120 mL/min, in place of its mean gives a conservative form of the same rule.

7.5 Instruments, Loops, Soft Sensor and Supervisor

The controller runs on every second 1 s physics step, a control period $\Delta t_c = 2$ s, before the integration (at $t = 2$, 4, 6 s and so on after a priming execution at $t = 0$, as Section 1.6.1 describes), and the historian logs one row every 10 s, 2160 rows for a six-hour campaign. The application runs at 1 to 600 times real time, 60 by default.

7.5.1 Instruments

No loop reads a cell of the plant except through an instrument. The Coriolis meters FT-101, FT-102 and FT-103 are exact first-order lags of 1 s with multiplicative colored noise of 0.3 % and a 3 s correlation time, and FT-102 carries an optional bias β_B. The zone thermometers TT-201 to TT-204 read the third cell of each zone through a 5 s lag with 0.01 K of noise and a 4 s correlation time, the circulator supply thermometers read through a 3 s lag, and PT-101 reads through a 1 s lag with 0.01 bar of noise and a 2 s correlation time. The analyzer AT-401 takes a physical sample. At the start of each method run, from the first control step on, it draws the true quenched composition at the sample point, and it publishes the result one transport of 30 s plus one run later, with Gaussian noise of 0.03 % on residual A, 0.015 % on impurities and 0.5 % of reading on strength. The assay it reports is 100 minus those two percentages.

7.5.2 The PI Algorithm

Every loop runs the framework's single PI algorithm, written in velocity form, as given by Equation (7.8).

$$\begin{aligned} e_k &= a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k \\ S_k &= S_{k-1} + (P_k - P_{k-1}), \qquad c_k = S_k + I_{k-1} \\ u_k &= \mathrm{clamp}(c_k, u_{min}, u_{max}) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - c_k}{\tau_I}\right]\Delta t_c \end{aligned} \tag{7.8}$$

with $a = \pm 1$ giving the direction of action. These are Equations (1.15) to (1.18) of Section 1.6.2 with no derivative term, no output bias ($c_k = v_k$) and the tracking time $T_t = \tau_I$, so the running sum S_k equals P_k, the limits act on the emitted output only and transfer from manual is bumpless. No output is rate limited. Under the ISA convention every loop but the quality trim is reverse acting.

7.5.3 Flow, Ratio and Zone Loops

FIC-101 and FIC-102 hold the two feed flows in mL/min on the pump speeds with $K_c = 0.3$ % per mL/min and $\tau_I = 10$ s, within 0 to 100 %. Each pump delivers 2 mL/min per percent of speed through

a 2 s drive and a 1 s meter, a loop gain of 0.6. The ratio station FFC-102 sets both setpoints from the total-flow reference Q_{ref} and the measured substrate flow, as shown in Equation (7.9).

$$r_v = \frac{(R_{SP} + \Delta R)\, C_{A,b}}{C_{B,f}}$$
$$Q_{A,SP} = \frac{Q_{ref}}{1 + r_v}, \qquad Q_{B,SP} = r_v\, Q_{A,m} \tag{7.9}$$

where $C_{A,b}$ is the drum label $C_{A,f}$ or the entered batch assay $C_{A,f}(1 + \delta_{A,as}/100)$. The measured ratio is $R_m = Q_{B,m} C_{B,f} / (Q_{A,m} C_{A,b})$, and neither basis sees a drum that drifts after it was assayed. TIC-201 to TIC-204 each circulator's supply command within 0 to 100 °C from the zone thermometer, with $K_c = 2\,^{\circ}\mathrm{C}$ per °C and $\tau_I = 120\,\mathrm{s}$. The circulator's power limit lies downstream of that range, where the back-calculation of Equation (7.8) cannot see it, so while a circulator reports a power limit in the direction its error asks for, the loop runs with $\tau_I = \infty$ for that scan, the integral stops and the proportional term still acts. The quench pump runs in open loop at 20 % of the measured feed flow.

7.5.4 Throughput Scheduler and Quality Trim

FY-100 moves Q_{ref} toward the lower of the setpoint and a ceiling Q_{cap} at the ramp limit $r = 10\,\mathrm{mL/min}$ per minute, and downward at $3r$ while a limit is cutting. Each circulator's cooling duty per unit of measured feed flow is learned through a 600 s filter while it cools with power in hand and the feed exceeds 30 mL/min. With the cooling limiter on, a circulator at full cooling whose zone runs more than 0.5 K above setpoint jumps the ceiling to the flow that 90 % of its capacity supports, and the ceiling then falls 1 mL/min per minute while the condition lasts. With the pressure limiter on, a discharge above 8 bar cuts the ceiling at the same rate. The ceiling climbs back 1 mL/min per minute once every circulator has 15 % of its cooling in hand and the discharge is 0.3 bar under the alarm, and it never falls below 10 mL/min. Both limiters are off in the reference configuration.

AIC-301 trims the ratio from the analyzer's residual A. It executes once per analyzer result, with Δt_c in Equation (7.8) set to the time since the previous result, because between results there is nothing new to act on. The laboratory loop of Section 1.4.5 acts once per result in the same way. A result from a sample with no product in it is skipped. It is direct acting ($a = -1$), since more residual A must raise the ratio, with a target of 0.8 %, $K_c = 0.04$ mol/mol per percent, $\tau_I = 20\,\mathrm{min}$ and the trim ΔR held within -0.10 to 0.30. It is off in the reference configuration, where ΔR is the operator's.

7.5.5 The Soft Sensor

AY-301 starts from its own steady line (from solvent in the two startup cases) and integrates a copy of the 1179 line states of Equations (7.2) to (7.4), in two 1 s Runge-Kutta steps per control period, driven only by the three flow meters, the measured supply temperatures, the batch-assay strength of feed A and the declared T_{pre}, UA_z, kinetics and volumes. Two corrections keep it honest, as shown in Equation (7.10).

$$\frac{\mathrm{d}\hat{T}_i}{\mathrm{d}t} \mathrel{+}= \frac{1}{90\,\mathrm{s}}\left(T_{z,m} - \hat{T}_{z,m}\right), \quad i \in z$$
$$b \leftarrow \mathrm{clamp}\left[b + g\,(y - \hat{y}(t_s) - b),\ \pm b_{max}\right] \tag{7.10}$$

where $\hat{T}_{z,m}$ is the model's thermometer cell passed through the same 5 s lag as TT-20z, y is an analyzer result, $\hat{y}(t_s)$ the model's prediction for that sample stored when it was drawn, $g = 0.5$, and the bias b on assay, impurities and strength is bounded at 3 %, 0.6 % and 0.06 mol/L. Comparing a result with the prediction stored when its sample was drawn is the alignment of Equation (1.14), and the update is the steady scalar filter of Equation (1.12) with $\lambda = g$. The bias is applied to every prediction at once, so it corrects a slow, persistent mismatch well and a short slug of unmeasured material not at all.

7.5.6 The Diversion Supervisor

XY-401 follows material in volume rather than in time, the delay kept in volume of Equation (1.6), as Chapter 4 advances its chip column on the measured flow. It counts the measured volume $W(t) = \int (Q_{A,m} + Q_{B,m} + Q_{q,m})\,\mathrm{d}t$ that has passed the sample point, and material sampled at $W = w$ reaches the valve while W lies in $[w + 0.7956\,V_c,\ w + 1.2328\,V_c]$, the declared quantiles of the coil. In the plug-flow mode an upset seen at the inlet, a measured ratio more than 5 % off its setpoint, a zone more than 3 K off its setpoint or a feed flow more than 10 % off its reference, opens a window from $t_0 + \tau$ to $t_1 + \tau$, with $\tau = V_R/Q_m + V_c/[(1+q)Q_m]$ evaluated at the start of the upset. A failing analyzer result diverts while $W - V_c$ lies between the failing sample and the next good one, and after startup or a restart the valve accepts one plug-flow time after the feeds start. In the residence-time (RTD) mode the valve diverts in three cases, as shown in Equation (7.11).

$$\begin{aligned} &[W - 1.2328\,V_c,\ W - 0.7956\,V_c] \cap [w_a, w_b] \neq \emptyset \\ &W - 1.2328\,V_c < w_{rel} \\ &\hat{q}_k \notin \mathrm{spec} - m, \quad k \geq 200 - n_G \end{aligned} \tag{7.11}$$

where $[w_a, w_b]$ is a window opened by a failing result that the soft sensor did not foresee, reaching back to the last good sample and closed at the first of the confirming run of good results, w_{rel} is the release boundary set at the first confirming sample after startup or a stop, and $\hat{q}_k$ is the predicted quality of coil cell k with its bias. The margin m keeps 25 % of each attribute's design headroom in hand, 0.29 % of assay, 0.115 % of impurities and 0.0154 mol/L of strength, and the guard covers $n_G = \lceil [Q_{C,m}(t_{XV} + \Delta t_c) + 30\,\mathrm{mL}]/V_k \rceil$ cells, 6 at 120 mL/min. The number of confirming results is 1 in the reference configuration. A case that starts on a running line carries the analyzer's last result, computed from its starting state, and if that result fails the valve starts on divert with an analyzer window open from before the start until good results return. The valve cannot change anything upstream of itself, so both modes run on every step and the unselected one is scored on the same material through the second valve position.

7.5.7 Plant, Control and Internal Models

The line of Section 7.3 is the simulation model. The application identifies no model of its own, and the control models that the loops face come from step tests on the line, each the difference between a stepped and an unstepped run in manual from the same state. A 1 °C step of zone 1's supply command moves its thermometer cell by 0.827 °C, 10 % of the way after 40 s and 63 % after 122 s, and zone 2 gives 0.785 °C and 124 s, of which the circulator's own response takes 60 s. The flow loops face 2 mL/min per percent through 3 s of lags. The controllers' internal models are the stoichiometry of the ratio station, the scheduler's learned duty per unit flow, the plug-flow delay τ, the declared coil quantiles and the soft sensor's copy of the line, which has the plant's structure and declared kinetics. The predictive controller of Section 7.5.8 runs that same copy forward and adds a forced response identified by setpoint steps.

7.5.8 A Predictive Quality Controller

Section 1.6.3 describes the MPC option in general. This one plans on the soft sensor's nonlinear model rather than on a small linear model. The Quality control option offers Loops, the default of the application and of cases 1 to 8 and the campaign, and MPC. Under MPC in automatic, a receding-horizon plan replaces AIC-301 and moves three inputs, $u = [\Delta R,\ \Delta T_2,\ \Delta T_3]$, the ratio trim of Equation (7.9) within AIC-301's range of -0.10 to 0.30 and offsets of ± 5 K on the setpoints of TIC-202 and TIC-203. The flow, ratio and zone loops, FY-100 with both limiters, AY-301 and XY-401 act beneath and beside it unchanged. AIC-301 tracks the applied trim, and its switch and tuning leave the panel while the residual A target A_{SP} stays, since the plan holds the same target.

A decision is taken every $T_s = 30$ s, 15 control periods, on a rounded clock and on the priming call, and a plan step is 100 mL of reactor flow rather than a time, because a dead volume through the

cells is the same at every rate. The horizon is $N = 30$ steps, 3 L or 30 min at 100 mL/min and 20 min at 150 mL/min, in ten move blocks, three of one step and two of two, followed by blocks of 3, 4, 5, 5 and 6 steps. The predicted outputs $\hat{y} = [\hat{a},\ \hat{\imath},\ \hat{s},\ \hat{\alpha}]$ at the sample point are given by Equation (7.12).

$$\hat{y}_k = \bar{y}_k + \beta + \sum_{m=1}^{3} g_m(R_{op}, Q_{op}) \left[\xi_{m,k} - u_m(t)\right] \tag{7.12}$$

Here $\bar{y}_k$ is the free response, the soft sensor's 25 reaction cells and four jackets at the declared parameters run forward by fourth-order Runge-Kutta at 4 s from the controller's own estimate, with the four zone loops at their present tuning and commands, the circulators at their declared sizes, the feeds split at the station's present ratio and the total flow following FY-100's own ramp toward its target, sampled every 100 mL. That ramp is the preview, which an advanced switch turns Off to plan on the measured flow held. The term $\xi_{m,k}$ is input m passed through a dead volume θ_m and two equal lags of τ_m steps, from which the input in force, $u_m(t)$, is subtracted, fitted to the averaged up and down setpoint steps of the line in automatic at 100 mL/min, a ratio of 1.10 and the design recipe, each the difference of a stepped and an unstepped run on one seed. The fits give $\theta = 11.50$, 5.75 and 3.25 steps and $\tau = 2.58$, 2.00 and 1.83 steps for ΔR, ΔT_2 and ΔT_3, with root-mean-square residuals of 2.1, 3.6 and 2.0 % of the steady change in assay. The ratio enters at the mixer, 1.615 L or 16.2 steps above the sample point, and the fitted chain passes half its response at 15.9 steps. The gains g_m are interpolated at each decision from a table of the declared steady line over 40 to 200 mL/min and ratios of 0.90 to 1.50, at the mean planned flow Q_{op} and at the true ratio the estimator infers, R_{op}, because the assay is far from linear in the ratio, rising at 9.9 points per unit for a step of 0.02 up at the design point and falling at 14.0 for one down. Stepped again by 0.02 and 1 K, the plant's steady gains agree with the table's within 0.7 % and its half-response points with the chain's within 0.9 step for the ratio and 1.6 steps for the zones. The model leaves out the deposit, a lost circulator until it reports its limit, any change to the kinetics and a cutback before it starts.

The estimator keeps the controller's own copy of AY-301's model, stepped every control period on the same meters, supply thermometers and thermometer correction, with feed A entering at the strength AY-301 believes times a factor σ, an unmeasured input that travels with the material. Each HPLC result is compared with that copy's prediction stored at the sampling time, in the form of Equation (7.13).

$$\begin{aligned} \sigma &\leftarrow \operatorname{clamp}\left[\sigma + g_\sigma (R_{bel}/R^{*} - \sigma),\ 0.80,\ 1.25\right] \\ \beta_j &\leftarrow \operatorname{clamp}\left[\beta_j + 0.3\,(y_j - \hat{y}_j(t_s) - \beta_j),\ \pm\beta_{j,max}\right] \end{aligned} \tag{7.13}$$

where R_{bel} is the ratio the meters and labels gave when that sample's material entered the mixer, one reactor volume before t_s, σ_e the factor in force then, and R^{*} the ratio at which the declared steady line gives its residual A at R_{bel}/σ_e plus the residual $\alpha - \hat{\alpha}(t_s)$. The gain is $g_\sigma = 0.5$, j is $\imath$ or s, the biases are bounded at 0.6 % and 60 mmol/L, and the offset β of Equation (7.12) is $-\beta_\imath$ on assay, $\beta_\imath$ on impurities, β_s on strength and zero on residual A. Repeated results on material already in flight therefore do not wind σ up. With three integrating disturbances for the three measured attributes the estimate is offset-free, the condition stated for disturbance models in MPC [21] and in Section 1.4.7, and $R_{op} = R_{bel}/\sigma$. A result that arrives between decisions moves σ on arrival and the next decision plans with it. The output bias b of AY-301 was tried first and rejected, because a bias learned on material made at one ratio was applied to material made at another.

The plan solves the quadratic program of Equation (7.14) on the shared `LinearMPC` solver.

$$\begin{aligned}
\min\ & \sum_{k=1}^{N}\Big[Q_\alpha(\hat{\alpha}_k - A_{SP})^2 + u_k^{\mathsf{T}} R_u u_k\Big] + \sum_{j=1}^{10} \Delta u_j^{\mathsf{T}} S\,\Delta u_j \\
& + \sum_{k\in\mathscr{K}}\ \sum_{i\in\{a,\iota\}} \big(\ell_i \varepsilon_{i,k} + \ell_i \varepsilon_{i,k}^2\big) + \ell_s(\varepsilon_s + \varepsilon_s^2) \\
\text{s.t. } & \hat{a}_k \ge a_{min} + m_a + c_a - \varepsilon_{a,k} \\
& \hat{\iota}_k \le i_{max} - m_\iota - c_\iota + \varepsilon_{\iota,k} \\
& \hat{s}_k \ge s_{min} + m_s + c_s - \varepsilon_s, \qquad k \in \mathscr{K} \\
& u_{min} \le u_k \le u_{max}, \quad |\Delta u_j| \le \Delta u_{max}, \quad \varepsilon \ge 0
\end{aligned} \tag{7.14}$$

The judged steps $\mathscr{K}$ start at step 3, the dead volume of ΔT_3, take every second step, and include only steps where the free response predicts product, a strength of at least 0.8 s_{min}, so solvent and a startup front are not priced. The margins m are the supervisor's own, so a slack is a predicted diversion and its linear price the predicted off-specification volume. The weights are $Q_\alpha = 10$ per %^{2}, $R_u = \mathrm{diag}(20,2,2)$ and $S = \mathrm{diag}(800,2,2)$ in (mol/mol)$^{-2}$ and K^{-2}, $\ell_a = 4000$ and $\ell_\iota = 10\,000$ per percent, $\ell_s = 100$ per mmol/L, cushions $c = 0.2\,\%$, 0.02 % and 2 mmol/L, and caps $\Delta u_{max} = 0.02$, 0.5 K and 0.5 K per block boundary. The bounds on the inputs and their moves are hard, the assay and impurity bands are soft with one slack per judged step, and the strength is soft with one slack over the horizon. A zone whose circulator reports full cooling is never asked to run colder, nor one at full heating hotter. The program has 59 variables, 30 moves and 29 slacks. The cap applies at every block boundary whatever the block's length, and it binds at a later boundary in 92 of the 239 plans of case 9. Counted instead as a ramp of 0.02 per step between block centers, no case or worst case moves by more than 0.05 L, so the uniform cap was kept.

The controller keeps a solve only when its status is optimal and it is of the right shape, finite in its move, plan, predictions, cost and slacks, starts with the move it reports and keeps every planned input in its box. Otherwise the historian is marked MPC FALLBACK with the reason, AIC-301, which has tracked the trim, takes the ratio without a step, the offsets ramp back at 1 K/min, and the next valid plan resumes. A fallback forced at 40 min in case 9 hands over with a trim step of 0.0000 and offsets moving at most 0.0044 K/s, and the return moves the trim by 0.0009 against its cap of 0.02. Below 30 mL/min or with the feeds stopped the controller holds. With one block and no bound the first move equals the hand least-squares move to 10^{-16}, and with every step its own block it equals the linear-quadratic regulator's $-Kx_0$ to 2×10^{-8}. Under MPC only, XY-401 also remembers each failing result the soft sensor had foreseen, and if a revised bias later predicts that sample in specification the result's own window diverts, marked ANALYZER VERDICT HELD (Section 7.6.11).

The controller reads a packet of AY-301's state and bias, the meters and thermometers, the circulators' reported duties, the loops' commands and tuning, the panel's setpoints and FY-100's ramp, the labels and entered assay and the quality limits with their margins, and the declared parameters once at import. Doctored copies that read the plant, a run parameter, a disturbance or the truth record by 35 routes are all caught by the application's checks, and replaying every decision of case 9 recorded on a plant with changed kinetics, jackets and volumes into a fresh controller gives 140 of 140 identical decisions. In Chrome, measured while other work loaded the machine, a decision takes 3.2 to 3.9 ms on average and 22.0 ms at worst, 0.011 to 0.013 % of a core at real time and 6.5 to 7.7 % at 600 times. A frame guard would drop simulated time rather than frames if a frame overran 25 ms, and it held none, the longest frame taking 23.5 ms.

7.6 Nine Cases and a Campaign

The cases are listed in Table 7.4, numbered as the app's scenario list numbers them (Section 1.6.4). Each starts from the reference configuration (RTD tracking, one confirming result, a 25 % margin,

Table 7.4 Cases and Campaign, With the Event That Defines Each

#	Scenario and event	Teaching point
1	Startup from solvent	the front arrives spread
2	Feed A +30 %/h, 10–30 min	what no meter sees
3	Pump B air lock at 20 min	a slug and its tails
4	160 mL/min, then zones at 65 °C	residence, heat, recipe
5	130 mL/min, zone 1 cooling to 30 %	an infeasible rate
6	10 min method, drum +10 %	analyzer against coil
7	Zone 3 deposit, 60 %/h	pressure as a rate limit
8	Stop at 15, restart at 50 min	held material reacts
9	150 mL/min at 10, drum +8 % at 60 min	quality trim against MPC
10	Six-hour campaign	graded challenge

Table 7.5 Measured Outcomes, Default Seed. RTD Tracking Selected and the Plug-Flow Delay Scored as Its Shadow Unless Stated

Case	Measured outcome
1, strength at valve	0.01 mol/L at 28.2, half 33.9, 95 % at 39.8 min
1, RTD	accepts 42.4 min, 4.45 L diverted, 0 released
1, no confirmation	accepts 38.0 min, worst released 97.85 %
2, trim on	first failing result 42.5 min, trim +0.099
2, RTD and plug	3.59 L diverted, 0 released, plug 139 mL
2, trim off	assay settles 96.12 %, 9.29 L diverted
3, off-spec at valve	44.07–55.85 min, as low as 86.62 %
3, RTD	diverts 43.50–56.23 min, 1.53 L, 0 released
3, plug	diverts 45.50–60.43 min, 177 mL released
4, 160 mL/min	steady 97.32 %, RTD 5.84 L, plug 114 mL
4, ramp 2 mL/min per min	RTD 3.63 L, plug 466 mL
5, no limiter	zone 1 72.73 °C, 1.125 %, 18.0 L diverted
5, cooling limiter	feed down to 54.2 mL/min, 0.790 %, 0 diverted
6, 10 min method	803 mL released, plug 1039 mL
7, pressure limiter	alarm 46.25 min, 8.31 bar peak, 53.4 mL/min at 2 h
7, no limiter	10 bar trip at 71.2 min
8, held hot	impurities to 1.128 %, 3.85 L diverted, 0 released
8, cooled in the stop	at most 0.694 %, 1.86 L diverted

the label basis, the trim and both limiters off) apart from the stated changes, and case 9 starts with the trim on and Quality control on MPC. Each number below is a measurement on the engine, run headless at the default seed with the stated change made before the run starts, which gives the same run as making it and pressing Reset, and times are minutes of simulated time. Where RTD tracking is selected the plug-flow delay is the shadow, and selecting it instead releases exactly what the shadow scored, 177.0 mL in case 3. Table 7.5 gathers the outcomes.

7.6.1 Startup and the First Product

Case 1 starts from a line full of solvent and ramps the feed from zero to 100 mL/min in 10.0 min. The reactor's cells spread the product front, so the strength at the valve passes 0.01 mol/L at 28.2 min, half the design strength at 33.9 min and 95 % of it at 39.8 min, around the 28.65 min plug-flow time of full flow. The first good analyzer result is published at 30.5 min from a sample drawn at 27.0 min, and at that moment the material at the valve has 13.9 % of the design strength and 89.10 % assay. Good material reaches the valve from 37.4 min. The RTD supervisor releases material sampled from the first good sample on, once the coil's spread has carried all earlier material past the valve, and it accepts at 42.4 min after diverting 4.45 L, 0.61 L of it in specification, releasing nothing. With no confirming result the valve opens at 38.0 min, diverts 3.92 L and releases nothing, and the worst unit it accepts is 97.85 % assay, 0.35 points inside the limit.

7.6.2 A Drum That Strengthens Unseen

Case 2 strengthens feed A by 30 % per hour from 10 to 30 min, 10 % in all, with the trim on. The ratio station and the soft sensor both believe the label, so the first warning is a failing result published at 42.5 min, 32.5 min after the drift began. The soft sensor did not foresee it, and the supervisor diverts from 42.5 to 72.4 min, 3.59 L of which 1.39 L was in specification, while off-specification material reaches the valve from 50.4 to 68.7 min. The surprise moves the bias of every prediction at once, including those for material made early in the drift and still in specification, which is the price of a correction that cannot tell when a change began. AIC-301 raises the ratio by 0.099. Without the trim the assay settles at 96.12 % and 9.29 L are diverted in two hours. The plug-flow delay, which opens its window at the failing sample rather than at the last good one, releases 139 mL on the same material.

Pause. Try Exercise 7.2 before reading on.

7.6.3 A Four-Minute Air Lock

Case 3 takes pump B to a tenth of its stroke from 20 to 24 min. FT-102 sees the fall at once, FIC-102 holds pump B at full speed for 182 s, the measured ratio falls as low as 0.19 and the measured feed to 56.9 mL/min, and when the air clears the saturated pump overshoots to a ratio of 3.08 and a feed of 194.0 mL/min. Off-specification material passes the sample point from 31.9 to 43.1 min and reaches the valve from 44.07 to 55.85 min, as shown in Figure 7.2, 11.78 min for a four-minute cause, and the assay at the valve bottoms at 86.62 %. Four results fail, published at 36.5, 39.5, 42.5 and 45.5 min. The soft sensor saw the flows, and its predictions for those four samples, stored when each was drawn, lie within 0.08 points of their true assay, so none is a surprise. Its prediction at the valve stays within 0.11 points of the truth over 30 to 66 min. The RTD supervisor diverts on the prediction from 43.50 to 56.23 min, 1.53 L of which only 0.11 L was in specification, and releases nothing. The screenshot of Figure 7.3 was taken inside that diversion.

> **TRY IT IN THE APP** **3 · Ratio Error**
>
> Load the case and press Run. At 20 min pump B turns red and a hatched slug enters zone 1. Follow it through the four zones, where it widens, and watch the violet prediction bar on the coil. XY-401 diverts from 43.50 to 56.23 min, 1.53 L in all, and releases nothing, while the plug-flow delay scored on the same material releases 177 mL.

The plug-flow delay does worse on the same material. Its analyzer window opens at the failing sample drawn at 33.0 min, which reaches the valve in plug flow at 45.50 min, 1.43 min after off-specification material began to arrive, so it releases 177 mL from the front of the slug, the worst at 94.98 %. The upset was raised at 20.03 min, when the measured feed had already fallen to 84.1 mL/min, so its delay is 34.05 min instead of 28.65 min, and its upset window runs from 54.08 to 60.42 min, opening 10.0 min after the front of the slug reached the valve and closing 4.57 min after its tail had passed. By 1 h it has diverted 1.74 L, 0.50 L of it in specification. A plug-flow delay is wrong at both ends of a spread disturbance, and here it is also timed from a flow that the disturbance itself had changed.

7.6.4 Throughput and the Recipe

Case 4 asks for 160 mL/min at 10 min, ramped at 10 mL/min per minute. The residence falls from 16.15 to 10.09 min in the reactor and from 12.5 to 7.81 min in the coil, and zone 1's circulator has to remove 155 W of its 200 W. The soft sensor sees the shorter residence coming, and the supervisor diverts from 26.7 to 57.1 min, 5.84 L, releasing nothing, while off-specification material reaches the valve from 28.3 to 55.3 min, as low as 96.99 %. The flow followed its reference, so the plug-flow delay had no upset to time and released 114 mL. At 45 min zones 2 to 4 go to 65 °C and the steady assay recovers to 98.13 %, with impurities up from 0.41 to 0.59 % and the hot spot moving to the

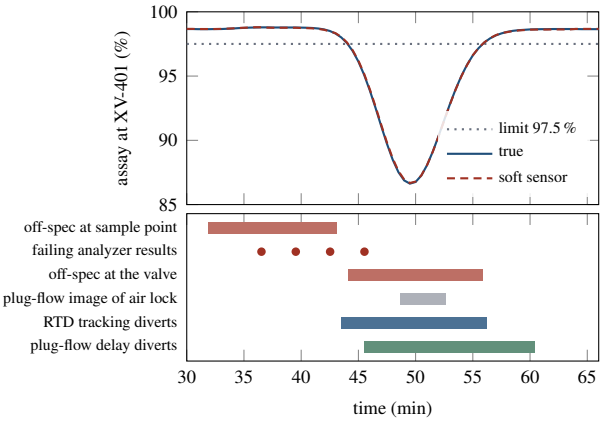

Figure 7.2 Case 3, the four-minute air lock at 20 min. Upper panel: the true assay at XV-401 (solid) and the soft sensor's prediction for the same material (dashed), which never differ by more than 0.11 points, with the 97.5 % limit dotted. Lower panel: when off-specification material passes the sample point and the valve, when the four failing results are published, where plug flow would have placed the air lock, and when each tracking rule diverts.

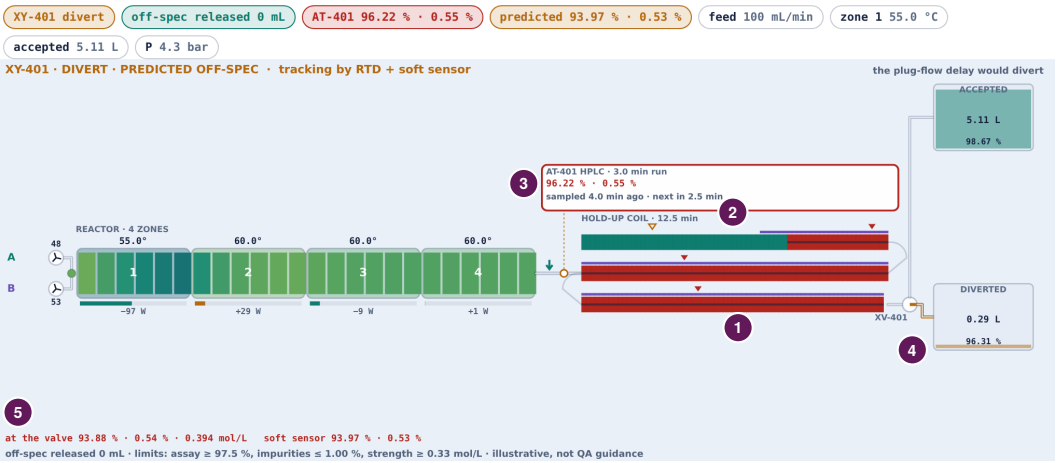

Figure 7.3 The application in case 3, ratio error, at $t = 46.00$ min, 26 min after the air lock began, in the Line view with RTD tracking selected, cropped to the plant panel (the clock is cut). The header reads XY-401 · DIVERT · PREDICTED OFF-SPEC, and the plug-flow delay, scored on the same material, would divert too. (1) The air-locked material fills 155 of the 200 cells of the hold-up coil in red, from the upper pass to the valve, and the soft sensor's violet bar runs over 161 cells, beginning six cells further upstream because of its margin (counts the panel draws but does not print). (2) Filled red triangles mark where the material behind the failing samples of 36.0, 39.0 and 42.0 min now is, and a hollow triangle near the coil inlet marks the sample of 45.0 min still being analyzed. (3) The analyzer box shows the latest failing result, 96.22 % assay and 0.55 % impurities, sampled 4.0 min earlier with the next due in 2.5 min. (4) XV-401 and the tanks, which hold 5.11 L accepted at 98.67 % and 0.29 L diverted. (5) The line under the scene gives 93.88 % assay at the valve against 93.97 % predicted. The zones read their setpoints with duties of −97, +29, −9 and +1 W.

Table 7.6 Drum Arrivals That Escape, at 12 Arrival Times One Minute Apart From 20 to 31 min. t_{det} and V_c/Q_C in Minutes

Method	Feed	t_{det}	V_c/Q_C	Predicted	Measured
5 min	100 mL/min	10.5	12.50	0	0 of 12
6 min	100 mL/min	12.5	12.50	0	0 of 12
6.5 min	100 mL/min	13.5	12.50	0.15	2 of 12, 29–87 mL
7 min	100 mL/min	14.5	12.50	0.29	3 of 12, 85–205 mL
10 min	100 mL/min	20.5	12.50	0.80	10 of 12, 83–923 mL
10 min	70 mL/min	20.5	17.86	0.26	3 of 12, 31–202 mL
10 min	60 mL/min	20.5	20.83	0	0 of 12

end of zone 2 at 70.5 °C. Conversion bought with temperature is paid for in impurity. A ramp of 2 mL/min per minute spreads the change, and the supervisor then diverts 3.63 L while the plug-flow delay releases 466 mL.

7.6.5 Lost Cooling and an Infeasible Rate

Case 5 runs at 130 mL/min, where zone 1 needs 128 W of cooling, and at 10 min its circulator falls to 30 %, 60 W. Zone 1 climbs to 72.73 °C, the decomposition grows fastest, and impurities at the valve reach 1.125 %, off specification from 70.47 min after the fault. Its supply climbs from 42.48 to 64.25 °C by 64.6 min while TIC-201, integral held, lowers the command on its proportional term alone to 7.5 °C by 100 min. The soft sensor sees that supply and tracks the zone thermometer within 0.073 K, so it predicts the rise, and the supervisor diverts everything from 64.6 min, 15.8 min before the first off-specification material arrives, because its margin flags impurities above 0.885 %. In three hours 18.0 L are diverted and 10.1 L accepted. With the cooling limiter on, FY-100 jumps the ceiling at 11.3 min to the 54.8 mL/min that 90 % of the remaining cooling supports, the feed falls as low as 54.2 mL/min, zone 1 never passes 55.94 °C, impurities stay at or below 0.790 %, and nothing is diverted, while 13.5 L are accepted. At 60 mL/min zone 1 needs 59 W. A lower rate is the only way to make material in specification with 60 W, and it accepts 13.5 L against the 10.1 L of the line that ran hot.

7.6.6 The Analyzer's Clock Against the Coil

Case 6 slows the method to 10 min, and at 25 min a drum 10 % stronger than its label goes into service unassayed. The soft sensor cannot help, because nothing measured changed. Off-specification material reaches the valve from 53.9 min, but the first failing result is published at 60.5 min from a sample drawn at 50.0 min, so 803 mL as low as 96.23 % are accepted, and the plug-flow delay would accept 1039 mL. With a 3, 5 or 7 min method the same drum is caught. At 60 mL/min the 10 min method catches it too, with the failing result published at 80.5 min and off-specification material reaching the valve at 83.2 min.

> **TRY IT IN THE APP** **6 · Slower Analyzer**
>
> Press Run and watch the analyzer box, where each result of the 10 min method arrives one run and one transport after its sample. The drum goes in at 25 min, and the first failing result is published at 60.5 min, after 803 mL as low as 96.23 % assay has been accepted. Reset, set Analyzer method run time to 3 min before Run, and the drum is caught.

Whether a drum escapes depends on where its arrival falls in the sampling cycle, and the rule of Equation (7.7) predicts how often, as listed in Table 7.6. With the 10 min method at 100 mL/min the drum escapes at 10 of 12 arrival times against a predicted fraction of 0.80, and the shipped arrival at 25 min is one of them. The 7 min method, safe for the shipped arrival, lets 3 of 12 through, 85 to 205 mL, and 70 mL/min lets 3 of 12 through with the 10 min method. No arrival escapes where t_{det}

is at or below the coil's mean holding time, and the boundary falls where the rule puts it, since a 6 min method ($t_{det} = 12.5$ min) lets none through while 6.5 min lets 2 of 12 through. Hold-up volume, analyzer cycle and throughput are therefore one design decision, and a single run at the shipped arrival time can make an unsafe design look safe.

> **DID YOU KNOW? Half a Minute of Method Time**
>
> From the simulator: at 100 mL/min, drum arrivals one minute apart all get caught with a 6 min analyzer method, while a 6.5 min method lets 2 of 12 through, 29 to 87 mL each (Table 7.6). The detection time has just passed the coil's mean holding time of 12.50 min.

7.6.7 A Deposit and a Pressure Limit

Case 7 grows a deposit in zone 3 at 60 % of its full extent per hour, full at 100 min, with the pressure limiter on. The discharge rises from 4.30 bar to the 8 bar alarm at 46.25 min. PT-101, whose reading carries noise, first reads above the alarm at 46.0 min, and FY-100 then cuts the feed to hold the discharge there, 8.31 bar at most, reaching 53.4 mL/min at two hours. The reactor residence lengthens to 30.2 min, the over-reaction grows, and impurities at the valve rise from 0.54 to 0.707 %, still in specification, with nothing diverted. Without the limiter the discharge reaches 10 bar at 71.2 min and the relay stops the feed pumps and the quench, leaving the line full of hot material. A pressure limit is a capacity limit in disguise.

7.6.8 Material Held Through a Stop

Case 8 stops the feeds at 15 min and restarts them at 50 min. With no reaction heat to remove, zone 1's circulator warms its supply from 45.58 to 55.0 °C during the hold, so it has to swing back when the feeds return, and the hottest cell reaches 63.81 °C at 56.8 min. The reactor contents keep over-reacting with the excess reagent, and they reach the valve from 68.7 to 82.0 min with impurities up to 1.128 %, while the coil's contents, quenched before the stop, stay good. The supervisor diverts from the stop until 87.4 min, waiting from the restart for a confirming result, 3.85 L in all, and releases nothing. With the zones held at 25 °C during the stop, impurities stay at or below 0.694 % and 1.86 L are diverted, all of it in specification.

7.6.9 The Campaign

The challenge runs six hours from a cold start. Feed A batch 2, 6 % strong, arrives at 1.5 h and batch 3, 4 % weak, at 4 h, both assayed, an air lock strikes at 3 h, and batch 3 drifts unmeasured at 12 % per hour from 4.5 to 5 h. The target is no off-specification material released, at least 16.5 mol of P accepted, at most 8 L diverted and both ledgers closed within 1×10^{-6}. The configurations are listed in Table 7.7. As shipped, with the plug-flow delay, no confirmation, the label basis and no trim at 100 mL/min, 4831 mL of off-specification material is released. The suggested configuration releases nothing, accepts 17.23 mol and diverts 6.09 L, of which 0.45 L was in specification, so it passes with 0.73 mol above the target and 1.91 L below the limit, in one run whether it is revealed before Run or carried through Reset. The campaign is judged on the case as given, and a changed model parameter or disturbance voids the attempt. Each element has its own job. The plug-flow delay releases 305 mL even with every other element in place, the trim is what keeps the unmeasured drift from filling the divert tank (14.38 L without it), and 120 mL/min is needed for the product, since 100 mL/min accepts only 14.21 mol. With the trim on, the batch-assay basis adds little, 17.23 against 17.18 mol, since on the label basis the trim itself raises the ratio after batch 2, from 0.042 at 1.75 h to 0.090 at 2 h. Figure 7.4 draws the configurations as a chart.

7.6.10 A Predictive Controller Against the Quality Trim

Case 9 asks for 150 mL/min at 10 min, which shortens the reactor residence from 16.2 to 10.8 min, and puts an unassayed drum of feed A, 8 % strong, into service at 60 min. With Quality control on

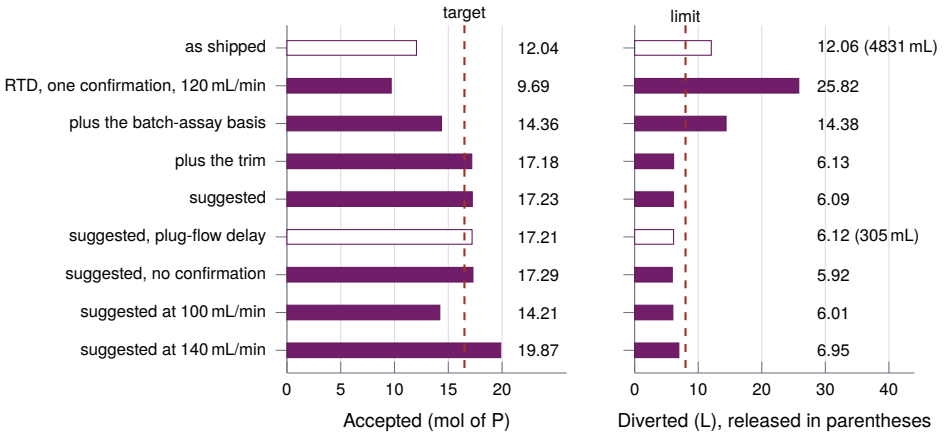

Figure 7.4 The six-hour campaign of Table 7.7 as a chart: product accepted against the 16.5 mol target (left) and volume diverted against the 8 L limit (right) for each configuration. Dark bars release nothing off specification; white bars release the volume in parentheses. Notice that only the plug-flow delay releases anything once the other elements are in place, and that without the trim the unmeasured drift fills the divert tank.

Table 7.7 Campaign Configurations Over Six Hours. Target: Nothing Released, at Least 16.5 mol Accepted, at Most 8 L Diverted

Configuration	Released	Accepted	Diverted
as shipped	4831 mL	12.04 mol	12.06 L
RTD, one confirmation, 120 mL/min	0	9.69 mol	25.82 L
plus the batch-assay basis	0	14.36 mol	14.38 L
plus the trim	0	17.18 mol	6.13 L
suggested[a]	0	17.23 mol	6.09 L
suggested, plug-flow delay	305 mL	17.21 mol	6.12 L
suggested, no confirmation	0	17.29 mol	5.92 L
suggested at 100 mL/min	0	14.21 mol	6.01 L
suggested at 140 mL/min	0	19.87 mol	6.95 L

[a] RTD, one confirmation, batch-assay basis, trim on, 120 mL/min

Loops and AIC-301 on, the trim is still at 0.008 at 20 min, because it learns of the faster material only from analyzer results, and XY-401 diverts from 28.1 to 94.3 min, 4.85 L of which 2.04 L was in specification, as shown in Figure 7.5. Under the MPC the trim reaches 0.112 by 15 min while FY-100 is still ramping, nothing is diverted, and 8.06 mol of P is accepted at 1.178 mol of B per mole against 6.20 mol and 1.513. With the preview Off the trim reaches 0.097 by 15 min, since the meters report the faster flow as it comes and the free response follows it, and again nothing is diverted. The drum is a draw. Both structures first see it in an analyzer result under 98.2 % assay, at 72.5 min under Loops and 75.5 min under the MPC, whose estimator settles at $\sigma = 1.081$, and under the MPC the true assay at the sample point dips to 97.69 %, inside the limit but under the supervisor's margin, until the added reagent has passed the reactor.

Three rows of Table 7.8 carry the comparison: the loops as shipped, AIC-301 retuned once against every knob it has by the rule the MPC was tuned by, and the strongest classical form, which gives the retuned trim every signal the MPC receives. That rule, fixed before any grid and applied on case 9 alone, asks for nothing released, no trip and no more time at a cooling limit than the loops, then the least volume diverted, then within 0.05 L the least reagent and then the least trim travel. Its tie-break was switched from the displayed two-decimal value to the ledger's exact value before any result was

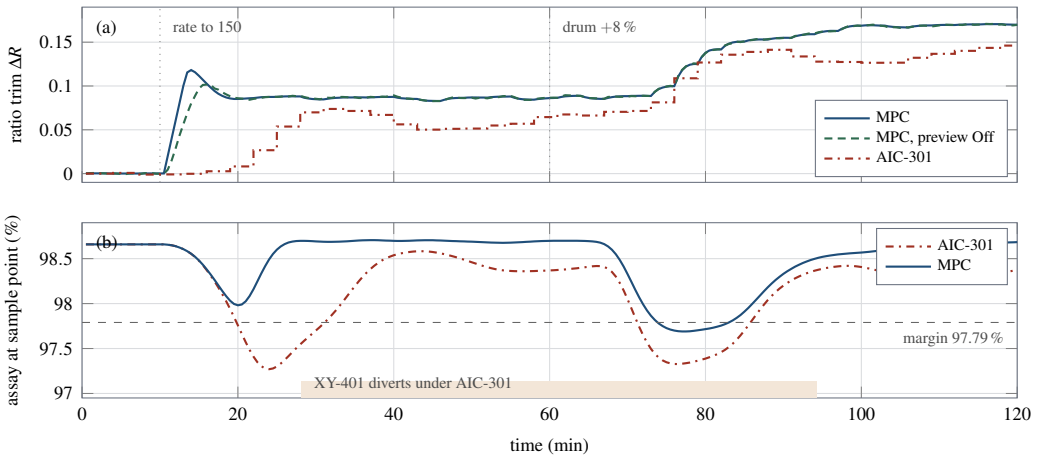

Figure 7.5 Case 9 over two hours, the rate raised to 150 mL/min at 10 min and an unassayed drum 8 % strong at 60 min. (a) The ratio trim as the MPC sets it with its preview (solid) and without it (dashed), and as AIC-301 sets it with Quality control on Loops (dash-dotted). (b) The true assay at the sample point under AIC-301 (dash-dotted) and under the MPC (solid), with the supervisor's margin (dashed) and the shaded interval in which XY-401 diverts under AIC-301, 4.85 L in all. Nothing is diverted under the MPC.

Table 7.8 Volume Diverted (L) and Off-Specification Product Released (mL, in Parentheses Where Any), Two Hours, Six for Case 10 With Its Suggested Settings. Reagent in mol B per mol P Accepted on Case 9

Structure, residual A target	9	B/P	4	2	6	10
Loops as shipped	4.85	1.513	5.84	3.59	7.13 (803)	6.09
AIC-301 on, 0.8 %	4.85	1.513	2.54	3.59	3.60 (809)	6.09
AIC-301 retuned, 0.4 %[a]	0.00	1.181	0.27	1.79	1.79 (33)	5.90
AIC-301 retuned, 0.8 %[b]	1.74	1.280	2.21	3.23	3.60 (815)	6.12
finer retune, 0.4 %[c]	0.00	1.171	1.02	2.51	2.99 (225)	5.91
classical, 0.4 %[d]	0.00	1.224	0.00	0.72	1.79 (25)	6.08
classical, 0.8 %[d]	0.00	1.201	0.58	2.87	2.86 (815)	6.14
MPC, 0.8 %	0.00	1.178	0.00	3.95	4.80 (801)	5.81
MPC, preview Off	0.00	1.177	0.00	3.95	4.80 (801)	5.89
MPC, no zone offsets	0.00	1.178	0.00	3.95	4.80 (801)	5.82
MPC, output bias	0.00	1.192	0.00	7.78	6.11 (807)	5.82
MPC, 0.4 %	0.00	1.222	0.00	0.00	0.00	5.75

[a]Retuned once by the MPC's rule over gain, integral time and target: $K_c = 0.08$, $\tau_I = 40$ min. [b]Target held: 0.08, 20 min. [c]A finer grid of the same rule: 0.06, 40 min. [d]Row *a* or *b* plus a ratio feedforward on FY-100's ramp from the declared steady table and the ratio corrected by a copy of the MPC's σ.

read, and 288 of the 324 MPC settings divert nothing on case 9, so the top settings are equivalent. The product threshold of the bands was moved from the strength limit to 0.8 of it after a worst case, a drift of −20 % per hour, diverted 3.62 L that was all in specification, a design error that leaves case 9 unchanged. The same seed, duration, instruments and noise apply to every row, and the MPC is never retuned per case. Against the shipped loops the gain is mostly the target that tuning chooses. AIC-301 retuned by that rule moves its target from 0.8 to 0.4 % and then diverts nothing on case 9 and 0.27 L on case 4. At an equal target the MPC wins or ties every row. At 0.8 % it diverts nothing on the two ramps, where the retuned trim diverts 1.74 and 2.21 L and the classical form 0.00 and 0.58 L, with the least reagent of those rows, 1.178 against 1.201. At 0.4 % it diverts nothing on all four cases, where the classical form still diverts 0.72 L on the drift and 1.79 L with 25 mL released on the slow analyzer. On the campaign it diverts 5.81 and 5.75 L against 6.14 and 6.08 L. The preview is worth at most 0.08 L, on the campaign, and the zone offsets nothing measurable, every row lying within

0.01 L without them, while they raise the hot spot by 4 K (68.7 against 64.7 °C on the campaign). On reagent alone the MPC does not beat the trim either, since a finer grid of the same rule finds AIC-301 at 0.06 and 40 min diverting nothing on case 9 at 1.171, though that setting loses cases 4, 2 and 6. The output-bias estimator that was rejected diverts 7.78 L on the drift.

🕮 **Pause.** Try Exercise 7.3 before reading on.

The MPC loses two rows at the shipped target. On case 2 it diverts 3.95 L against the loops' 3.59 L and the classical form's 0.72 L, because σ assumes a constant strength and lags a ramping drift. Moving σ fully on each result gives 3.59 L and a quarter of the way 4.67 L, and most of the gap to the classical row is its target, the MPC at 0.4 % diverting nothing. On case 6 every structure at 0.8 % releases about 800 mL, because a 10 min method is too slow for the coil whatever moves the ratio, and the MPC releases 801 mL against the retuned trim's 33 mL. The MPC made 240 decisions in two hours and 713 in the campaign with no fallback.

7.6.11 Worst Cases Under Every Structure

Table 7.9 runs every case and 18 extremes under the loops as shipped, with AIC-301 on, and under the MPC, against every limit the line has, the specification on the true composition accepted, the supervisor's margins, the circulators' power, the 8 bar alarm and the 10 bar trip. Every ledger closes within 1×10^{-6}. The MPC trips the pumps only where every structure does, with the pressure limiter off, releases off-specification product only where the loops also release it, and never spends more than 1.0 min longer than the loops at a cooling limit (6.7 against 5.7 min with three air locks). What it raises is the hot spot, by 4.0 K at startup and by up to 5.3 K with a weak drum, with no limit involved, and on the campaign the zone offsets account for it, the MPC without them peaking at 64.7 °C against 68.7 °C. Its losses beside cases 2 and 6 are the air lock of case 3, 1.68 against 1.53 L, three air locks in a row, 4.10 L against 2.50 L, a drift of 20 % per hour, 5.75 against 4.67 L, and 0.38 L with the cooling limiter where the loops divert nothing.

A review of the application made two behavior fixes, both outside every shipped and suggested run, whose original columns stay bit-identical. First, under the MPC an unassayed drum 15 % strong at 25 min in case 3 released 109 mL and FT-102 reading 10 % high 135 mL, where the shipped loops released none. A failing result that the soft sensor had foreseen opens no window, by the rule of Equation (7.11). The MPC then moves the ratio on the analyzer's evidence, results on the new material shrink AY-301's single bias from -2.95 to -1.14, and material made before the move, still in the coil at 96.2 to 97.3 %, is predicted in specification and released. Neither the band stride, the product threshold, the cushion nor the cap changes either release, and removing the residual A target removes both. The verdict rule of Section 7.5.8 releases nothing in either run, at a price of 3.96 to 4.67 L and 3.37 to 3.79 L diverted, and moves no other case or extreme under any structure. Making every foreseen failing result keep its window was measured too and rejected, because it adds 0.56 L to the campaign. Second, a running case whose last result is off-specification, as on a plant the reader changed before Run, had accepted the coil until the next result, 427 mL with $k_{1,60}$ halved in case 3 under the loops and the MPC alike and 385 mL in case 9 with zone 1 at 10 % and the limiter on. Under the start rule of Section 7.5 both release nothing, diverting 14.29 L under the loops (13.86 before) and 9.66 L under the MPC (8.88) with $k_{1,60}$ halved and 5.36 L (4.98) with the lost cooling, and case 5 at 180 mL/min with the limiter, which starts off specification, falls from 5 mL released to none under every structure.

Two defects of the shipped soft sensor and scheduler were fixed at the same time. AY-301 had been built from the run's parameters, so a reader's change to the plant reached it at once, and with $k_{1,60}$ halved it predicted the true 95.934 % at reset before any result, where the declared model gives 98.659 %. It now runs on the declared values, the model mismatch a real soft sensor faces. FY-100's jump at a cooling limit read the cooling-loss disturbance and now reads the duty the circulator reports, which at full cooling is the same number. No shipped, suggested or quoted number moved with either fix.

Table 7.9 Every Case and Extreme Under Each Structure: Volume Diverted (L) With Off-Specification Product Released (mL) in Parentheses, the Worst Unit Accepted Under the MPC (Assay %, Impurities %, Strength mol/L), and the Hot Spot (°C), Minutes at a Cooling Limit and Peak Discharge (bar) Under the Loops and the MPC. Two Hours, Six for Case 10, Unless Stated, After the Fixes of Section 7.6.11

Case and change	Loops	AIC-301 on	MPC	MPC worst accepted	Hot spot	Cooling limit	Discharge
1 startup	4.45	4.45	4.45	98.48/0.597/0.389	63.9/67.9	0.0/0.1	4.30/4.30
2 drift	3.59	3.59	3.95	97.91/0.546/0.391	63.1/63.1	0.0/0.0	4.30/4.30
3 air lock	1.53	1.54	1.68	97.80/0.627/0.382	64.3/64.0	2.2/3.1	5.52/5.50
4 throughput	5.84	2.54	0.00	97.99/0.624/0.373	70.5/70.5	0.8/0.7	5.09/5.09
5 cooling	8.63	8.79	7.86	98.05/0.885/0.382	73.3/73.2	110.0/110.0	4.69/4.70
6 slow analyzer	7.13 (803)	3.60 (809)	4.80 (801)	96.27/0.539/0.391	63.3/63.3	0.0/0.0	4.30/4.30
7 deposit	0.00	0.00	0.00	98.65/0.630/0.391	62.9/62.9	0.0/0.0	8.31/8.31
8 stop held hot	3.85	3.85	3.49	98.08/0.860/0.390	63.8/63.8	0.6/1.0	4.30/4.30
9 rate and drum	4.85	4.85	0.00	97.69/0.539/0.376	63.8/63.8	0.0/0.0	4.96/4.96
10 as shipped	12.06 (4831)	6.13 (55)	6.21 (43)	97.03/0.639/0.372	65.1/67.9	2.4/3.4	5.52/5.49
10 suggested	6.09	6.09	5.81	97.79/0.629/0.377	65.0/68.7	3.2/4.1	5.67/5.64
5, limiter on	0.00	0.00	0.38	97.77/0.739/0.384	63.1/63.1	10.3/10.9	4.69/4.70
5, 180 mL/min, limiter	5.11	6.38	4.31	97.78/0.863/0.369	64.6/69.1	14.1/14.9	5.34/5.35
5, zone 1 at 10 %, limiter	0.55	0.47	0.55	98.05/0.886/0.384	63.1/63.1	103.8/77.4	4.69/4.70
7, pressure limiter off	0.00	0.00	0.00	98.63/0.540/0.391	62.3/62.3	0.7/0.7	10.00[f]/10.00[f]
7, deposit 100 %/h	0.00	0.00	0.00	98.65/0.664/0.391	62.9/62.9	0.0/0.0	9.34/9.34
3, drum +15 % at 25 min	8.93	3.24 (291)	4.67	98.23/0.553/0.391	63.7/63.7	0.0/0.0	4.30/4.30
3, drum −15 % at 25 min	7.85	5.40	2.52	98.65/0.712/0.342	62.3/67.6	0.0/0.1	4.30/4.30
3, FT-102 +10 %	8.23	3.78	3.79	97.89/0.551/0.388	62.5/62.5	0.0/0.0	4.30/4.30
3, FT-102 −10 %	0.00	0.00	0.00	98.59/0.659/0.375	62.3/62.3	0.0/0.0	4.37/4.37
3, pump B at 60 %	0.00	0.00	0.00	98.46/0.542/0.391	62.5/62.5	0.0/0.0	4.30/4.30
2, +20 %/h for 3 h	4.67	4.67	5.75	97.89/0.552/0.391	63.4/63.4	0.0/0.0	4.30/4.30
2, −20 %/h for 3 h	2.88	2.88	0.00	98.65/0.639/0.347	62.3/62.3	0.0/0.0	4.30/4.30
6, 15 min method	7.13 (805)	5.40 (817)	7.13 (803)	96.25/0.539/0.391	63.3/63.3	0.0/0.0	4.30/4.31
4, 200 mL/min for 3 h	37.25	3.74	0.00	98.01/0.539/0.360	64.9/66.8	4.4/3.7	5.61/5.61
4, 160 then 60 mL/min	0.00	0.00	0.00	98.27/0.704/0.381	63.2/64.4	0.0/0.3	4.91/4.90
8, cooled in the stop	1.86	1.86	1.86	98.27/0.701/0.391	62.6/62.6	22.1/22.1	4.30/4.30
3, air locks at 20, 25, 30 min	2.50	3.78	4.10	97.78/0.624/0.382	65.1/64.7	5.7/6.7	5.52/5.50
9, 10 min stop at 62 min	7.47	7.47	6.83	98.00/0.910/0.376	64.3/64.3	2.3/2.7	4.96/4.96

[f] the 10 bar relay trips the feed pumps and the quench.

7.7 Technology in Practice

What Measures It

The line reads its feeds on three flow meters, its zones on four thermometers, its pump discharge on one pressure transmitter and its product on an at-line chromatograph. In the chapter the meters are 1 s lags with colored noise, the thermometers 5 s lags, and the analyzer a point sample reported one run and one transport late, as Section 7.5 sets out.

> **TECHNOLOGY** **Sensor: An At-Line Liquid Chromatograph**
>
> A chromatograph separates the species of a sample by their time-delayed travel down a packed column, and a detector at the outlet reads each species as it leaves. One instrument can measure the product leaving a reactor, the unreacted reactants and the unwanted products in the same stream. Cycle times of minutes are not uncommon even for on-line chromatographs, and some dead time is inevitable, since the method separates species in time. A leaking sample valve alters the baseline or corrupts the analysis [22]. The ICH Q13 guidance names online HPLC for monitoring the conversion of a chemical reaction and lists periodic system-suitability checks of the analyzer among the controls of a long run [4]. AT-401 draws a perfect sample at the start of each 3 min run, with noise and no drift.

The ICH Q13 guidance finds process analytical technology well suited to continuous manufacturing, because it lets disturbances be detected in real time [4]. In its integrated example the residence-time distribution of the drug-substance segment sets the time frame for monitoring quality,

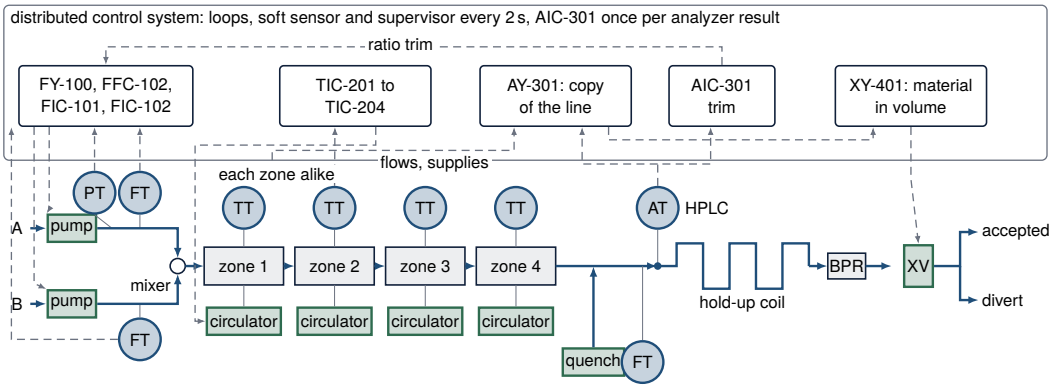

Figure 7.6 The loops of the flow line on plant hardware. Flow meters (FT) on the two feeds drive FIC-101 and FIC-102 on the pumps under the ratio station FFC-102 and the scheduler FY-100, and PT-101 on the pump discharge feeds the pressure limiter. Each zone thermometer (TT) drives its TIC on the circulator's supply. The at-line chromatograph (AT-401, HPLC) at the sample point feeds the quality trim AIC-301, which moves the ratio, and the soft sensor AY-301, which copies the line on the measured flows and supplies. XY-401 follows material in volume through the hold-up coil and the back-pressure regulator (BPR) and sets the divert valve XV. Notice that the analyzer sits upstream of the coil, which must hold the material until its verdict arrives.

which has to allow for the sampling frequency, the test method, the time the measurement takes and the instrument's capability [4]. That is the balance of Equation (7.7).

What Moves It

The final elements are the two feed pumps and the quench pump, the four circulators and the divert valve. In the 2013 plant-wide study the pumps were calibrated volumetric pumps, used both as actuators and as flow indicators [17]. The chapter meters each feed separately, and FT-102 sees the air lock of case 3 at once. The guidance's drug-substance example adds pressure sensors at the discharge of the feed pumps to watch for fouling, and backup pumps at key locations [4], the part PT-101 and the pressure limiter play against the zone 3 deposit of case 7. Figure 7.6 draws the loops of Section 7.5 on these devices.

> **TECHNOLOGY** **Final Element: The Three-Way Divert Valve**
>
> A three-way valve has three ports and sends the flow into one path or the other, like a single-pole double-throw switch. In a solenoid-operated valve a coil pulls an armature against a return spring, so the valve takes its normal position whenever the coil is de-energized, and the solenoid must be much stronger than the spring for the valve to act fast. Process pressure that works against the spring can keep a de-energized valve from reaching its normal position [22]. Diverting takes material out of the stream, so in the guidance's integrated example it changes the mass flow and may need a compensating strategy downstream that considers the residence-time distribution [4]. XV-401 strokes between accept and divert in 4 s, and the supervisor's guard covers that stroke.

Where the Controller Runs

A distributed control system, used in industries that include pharmaceutical processing, holds a supervisory level of control over localized controllers, and its loops run with cycle times from milliseconds to minutes [23]. In the application the flow, ratio and zone loops form the base layer and execute every 2 s. The scheduler FY-100, the soft sensor AY-301 and the supervisor XY-401 run above them on the same period, while the quality trim AIC-301 executes once per analyzer result,

every 3 min with the default method. The supervisor is the layer that holds this chapter's lesson: it follows material in volume and decides at the valve.

What Is Modernizing It

Two developments bear on this line. The first is the process model inside the control strategy. The guidance allows process models in commercial production, including the diversion strategy, and to predict quality attributes in real time. It asks that their developers understand the model's assumptions, naming plug flow against mixed flow, and that the model be maintained and its performance monitored during commercial manufacture [4]. The soft sensor and the two tracking rules of this chapter are such models, and the choice between plug flow and mixing is what the cases measure. The second is active quality control. The 2013 study noted that controlling product properties actively, through feedback and process analytical tools, may be a more efficient way to develop pharmaceutical processes than exhaustively documenting a design space [17]. The predictive controller of Section 7.5.8 is one such controller, and Section 7.6.10 measures what it gains over the trim.

> **IN PRACTICE** **Diversion in a Guidance Example**
>
> The ICH Q13 guidance works through a continuous drug-substance line with two plug-flow reactors. The first reaction's stoichiometry is held by the concentrations and flow rates of the feeds, and the second by feedback of an addition rate on an analyzer's reading of the residual intermediate. The residence-time distribution was modeled for every unit and surge point and confirmed by tracer studies on parts of the commercial equipment. Diversion is triggered by comparing process parameters and analyzer results with predefined criteria, its timing and duration are informed by the distribution, and at start-up material is diverted for a time based on it [4]. AIC-301, the declared cell chains and XY-401 play those parts in this chapter.

7.8 Limitations

The kinetics are invented, the quality limits are illustrative, and neither has had quality-assurance review. A real specification names each impurity and justifies each limit, and a real release decision rests on a quality system, lot genealogy and electronic batch records that this application does not model.

The residence-time distribution is declared rather than computed, so the spread of a real coil, which depends on its flow regime, has to be measured before a tracking rule can be trusted [11].

The soft sensor shares the plant's declared kinetics, so its case 3 predictions are better than a real model's and its 25 % margin is thin.

The analyzer takes a point sample every run, so a change narrower than the sampling interval at the sample point could pass unseen, although every disturbance here enters upstream of a reactor whose 3.27 min standard deviation widens it. During a stop the analyzer keeps sampling the stagnant outlet cell, which a real sample line without flow would not deliver.

> **OPEN CHALLENGE** **A Sample That Speaks for the Stream**
>
> A review by the regulator's staff calls the sampling interface of a continuous line challenging. Industrial experience finds poor measurement performance more often in the sampling system than in the analyzer itself, and on-line and in-line measurements may reduce sampling errors without removing them. The review asks that the process dynamics, such as the residence-time distribution, set the sampling frequency, so that a pulse of variability from a disturbance is resolved [2]. The analyzer of this chapter draws a perfect point sample with no line holdup and no failed injection. Progress would be a sampling system whose delays and failures are measured and carried into the rule that sizes the coil.

The predictive controller's model is the soft sensor's at the declared parameters, with a forced response fitted once at 100 mL/min and the design recipe and scaled by the declared steady line. It knows no deposit, no lost cooling until a circulator reports it, no change to the kinetics and no

cutback before it starts, and its strength factor assumes a constant feed strength, so a ramping drift is learned late (3.95 L against 3.59 L on case 2). The single analyzer bias that AY-301 applies to every batch in the line is kept as the declared design. Under the MPC the supervisor holds a foreseen failing result's verdict, while the shipped structures keep the rule of Equation (7.11), under which AIC-301 releases 291 mL after an unassayed drum 15 % strong. Holding the verdict under every structure would remove that release and change case 6 with AIC-301 on from 3.60 to 4.19 L diverted, a teaching choice left open.

7.9 Conclusion

A continuous synthesis has no batch to hold back, and this chapter has treated that as a problem of following material. Its reactor spreads a four-minute air lock into 11.78 min of off-specification material at the valve, starting 4.58 min before the 28.65 min plug-flow time, and residence-time tracking with a soft sensor diverts all of it in 1.53 L while a plug-flow delay, late at the front and timed from a flow the air lock had changed, releases 177 mL. A drum that no instrument sees is caught only by the analyzer, and the coil must hold longer than twice the method run plus the transport. At 100 mL/min a 10 min method lets the drum through at 10 of 12 arrival times, while a coil holding 20.83 min at 60 mL/min lets none through. Lost cooling at 130 mL/min leaves a choice between 18.0 L diverted and a rate cut to 54.2 mL/min with nothing diverted. The campaign moves from 4831 mL released to none, with 17.23 mol accepted. A predictive controller on the soft sensor's own model diverts nothing through the rate change of case 9, where the shipped trim diverts 4.85 L, but the trim retuned to the target its rule prefers diverts nothing either, and only at an equal target does the plan win every row. The model has invented kinetics, illustrative limits, a declared distribution and a soft sensor that shares the plant's declared kinetics. A fuller study would add a measured coil distribution, individual impurity limits, soft-sensor model error, a finite quench and a quality review of each diversion. The simulated line is a browser plant for practice in following material, rather than signals, to the point where it is routed.

KEY POINTS

- A continuous line has no batch to hold back, so it has to follow its material to the valve and divert the part an upset reached.
- The reactor spreads a four-minute air lock into 11.78 min of off-specification material at the valve, and residence-time tracking diverts all of it while a plug-flow delay releases 177 mL.
- A drum that no instrument sees is caught only by the analyzer, and the coil must hold longer than twice the method run plus the transport.
- Against the shipped loops, the gain of the MPC is mostly the target its tuning chose, and only at an equal target does the plan win every row.

Where else this applies. Volume 8, Chapter 13 judges each parcel's history at a diversion valve, Chapter 6 tracks a tablet line with a late analyzer, and Chapter 4 advances material on the measured flow.

References

[1] P. Poechlauer, J. Manley, R. Broxterman, B. Gregertsen, and M. Ridemark, "Continuous processing in the manufacture of active pharmaceutical ingredients and finished dosage forms: An industry perspective," *Organic Process Research & Development*, vol. 16, no. 10, pp. 1586–1590, 2012.

[2] S. L. Lee, T. F. O'Connor, X. Yang, C. N. Cruz, S. Chatterjee, R. D. Madurawe, C. M. V. Moore, L. X. Yu, and J. Woodcock, "Modernizing pharmaceutical manufacturing: from batch to continuous production," *Journal of Pharmaceutical Innovation*, vol. 10, no. 3, pp. 191–199, 2015.

[3] B. Gutmann, D. Cantillo, and C. O. Kappe, "Continuous-flow technology—a tool for the safe manufacturing of active pharmaceutical ingredients," *Angewandte Chemie International Edition*, vol. 54, no. 23, pp. 6688–6728, 2015.

[4] U.S. Food and Drug Administration, "Q13 continuous manufacturing of drug substances and drug products: Guidance for industry," Center for Drug Evaluation and Research, Docket FDA-2021-D-1047, Mar. 2023. [Online].

Available: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/q13-continuous-manufacturing-drug-substances-and-drug-products

[5] W. Engisch and F. Muzzio, "Using residence time distributions (RTDs) to address the traceability of raw materials in continuous pharmaceutical manufacturing," *Journal of Pharmaceutical Innovation*, vol. 11, no. 1, pp. 64–81, 2016.

[6] J. Kruisz, J. Rehrl, S. Sacher, I. Aigner, M. Horn, and J. G. Khinast, "RTD modeling of a continuous dry granulation process for process control and materials diversion," *International Journal of Pharmaceutics*, vol. 528, no. 1–2, pp. 334–344, 2017.

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

[8] 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.

[9] 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.

[10] P. V. Danckwerts, "Continuous flow systems," *Chemical Engineering Science*, vol. 2, no. 1, pp. 1–13, 1953.

[11] G. I. Taylor, "Dispersion of soluble matter in solvent flowing slowly through a tube," *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 219, no. 1137, pp. 186–203, 1953.

[12] K. Plumb, "Continuous processing in the pharmaceutical industry," *Chemical Engineering Research and Design*, vol. 83, no. 6, pp. 730–738, 2005.

[13] S. Mascia, P. L. Heider, H. Zhang, R. Lakerveld, B. Benyahia, P. I. Barton, R. D. Braatz, C. L. Cooney, J. M. B. Evans, T. F. Jamison, K. F. Jensen, A. S. Myerson, and B. L. Trout, "End-to-end continuous manufacturing of pharmaceuticals: Integrated synthesis, purification, and final dosage formation," *Angewandte Chemie International Edition*, vol. 52, no. 47, pp. 12 359–12 363, 2013.

[14] A. Mesbah, J. A. Paulson, R. Lakerveld, and R. D. Braatz, "Model predictive control of an integrated continuous pharmaceutical manufacturing pilot plant," *Organic Process Research & Development*, vol. 21, no. 6, pp. 844–854, 2017.

[15] A. Adamo, R. L. Beingessner, M. Behnam, J. Chen, T. F. Jamison, K. F. Jensen, J.-C. M. Monbaliu, A. S. Myerson, E. M. Revalor, D. R. Snead, T. Stelzer, N. Weeranoppanant, S. Y. Wong, and P. Zhang, "On-demand continuous-flow production of pharmaceuticals in a compact, reconfigurable system," *Science*, vol. 352, no. 6281, pp. 61–67, 2016.

[16] C. L. Burcham, A. J. Florence, and M. D. Johnson, "Continuous manufacturing in pharmaceutical process development and manufacturing," *Annual Review of Chemical and Biomolecular Engineering*, vol. 9, no. 1, pp. 253–281, 2018.

[17] R. Lakerveld, B. Benyahia, R. D. Braatz, and P. I. Barton, "Model-based design of a plant-wide control strategy for a continuous pharmaceutical plant," *AIChE Journal*, vol. 59, no. 10, pp. 3671–3685, 2013.

[18] L. L. Simon, H. Pataki, G. Marosi, F. Meemken, K. Hungerbühler, A. Baiker, S. Tummala, B. Glennon, M. Kuentz, G. Steele, H. J. M. Kramer, J. W. Rydzak, Z. Chen, J. Morris, F. Kjell, R. Singh, R. Gani, K. V. Gernaey, M. Louhi-Kultanen, J. O'Reilly, N. Sandler, O. Antikainen, J. Yliruusi, P. Frohberg, J. Ulrich, R. D. Braatz, T. Leyssens, M. von Stosch, R. Oliveira, R. B. H. Tan, H. Wu, M. Khan, D. O'Grady, A. Pandey, R. Westra, E. Delle-Case, D. Pape, D. Angelosante, Y. Maret, O. Steiger, M. Lenner, K. Abbou-Oucherif, Z. K. Nagy, J. D. Litster, V. K. Kamaraju, and M.-S. Chiu, "Assessment of recent process analytical technology (PAT) trends: A multiauthor review," *Organic Process Research & Development*, vol. 19, no. 1, pp. 3–62, 2015.

[19] P. Sagmeister, J. D. Williams, C. A. Hone, and C. O. Kappe, "Laboratory of the future: a modular flow platform with multiple integrated PAT tools for multistep reactions," *Reaction Chemistry & Engineering*, vol. 4, no. 9, pp. 1571–1578, 2019.

[20] Q. Su, S. Ganesh, M. Moreno, Y. Bommireddy, M. Gonzalez, G. V. Reklaitis, and Z. K. Nagy, "A perspective on quality-by-control (QbC) in pharmaceutical continuous manufacturing," *Computers & Chemical Engineering*, vol. 125, pp. 216–231, 2019.

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

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

7.10 Exercises

7.1 Concept check · 2 min

In case 3 pump B delivers a tenth of its stroke from 20 to 24 min. Off-specification material reaches the divert valve from 24.07 to 35.85 min after the air lock began, starting 4.58 min before the 28.65 min plug-flow time. Why does it arrive early?

(a) The air lock cut the feed flow, which shortened the residence time of the material behind it

(b) The reactor's 25 well-mixed cells spread the slug, so part of it leaves well before the mean residence time

(c) The analyzer reports 3.5 min after each sample, so the supervisor marks the slug ahead of time

(d) The 200 cells of the hold-up coil give the slug a wider spread than the reactor gives it

7.2 Concept check · 2 min

In case 6 a drum of feed A made up 10 % stronger than its label goes into service unassayed at 25 min. The soft sensor AY-301 predicted the case 3 air lock to within 0.11 points, yet here it cannot help. Why?

(a) It runs on the declared kinetics, so it cannot predict a change in a rate constant

(b) Its bias is bounded at 3 % of assay, which is smaller than the drum's effect

(c) No meter, thermometer or entered assay changes, so its model receives no input that carries the drum

(d) The supervisor stops using its predictions once the method is longer than the coil's holding time

7.3 Concept check · 2 min

In case 9 the MPC diverts nothing, where the shipped loops with AIC-301 on divert 4.85 L. AIC-301 retuned once by the rule the MPC was tuned by moves its residual A target from 0.8 to 0.4 % and also diverts nothing on case 9. What does this comparison show?

(a) Most of the MPC's gain comes from its preview of FY-100's throughput ramp

(b) Most of the MPC's gain comes from its offsets on the zone 2 and zone 3 setpoints

(c) The MPC plans on a more accurate model than the soft sensor that AY-301 runs

(d) Against the shipped loops the gain is mostly the target the tuning chose, and only at an equal target does the plan win

7.4 Theory and calculation · 6 min

The reactor is a 15 mL mixer and four zones of 0.40 L, each six equal cells. The hold-up coil is 1.5 L in 200 cells, and the quench adds 20 % of the feed.

(a) At a feed of 100 mL/min, compute the mean residence time and the standard deviation of the reactor and of the coil, and the plug-flow time from the mixer to the valve. *Checkpoint:* Reactor 16.15 min and 3.27 min, coil 12.5 min and 0.88 min, plug flow 28.65 min.

(b) Repeat the residence times at 160 mL/min, the rate of case 4. *Checkpoint:* Reactor 10.09 min, coil 7.81 min, plug flow 17.91 min.

(c) In case 3 the plug-flow delay times its upset window from the measured feed at the start of the upset, 84.1 mL/min. Compute the delay $\tau = V_R/Q_m + V_c/[(1+q)Q_m]$ and the window it opens for an upset from 20 to 24 min. *Checkpoint:* About 34.1 min, so the window runs from about 54.1 to 58.1 min.

7.5 Theory and calculation · 7 min

Use Equation (7.7) with the 30 s sample transport and the escape fraction $(t_{det} - V_c/Q_C)/t_{run}$ of Section 7.4.

(a) At a feed of 100 mL/min, compute t_{det} and the predicted escape fraction for the default 3 min method and for the 7 and 10 min methods. *Checkpoint:* 6.5, 14.5 and 20.5 min against 12.5 min of coil, so 0, 0.29 and 0.80.

(b) Find the longest method that the coil covers at 100 mL/min by the mean rule, and by the conservative rule that uses the coil's leading edge, 0.7956 coil volumes. *Checkpoint:* 6.0 min and about 4.7 min.

(c) Find the highest feed at which the 10 min method is covered by the mean rule, and the predicted escape fraction at 70 mL/min. *Checkpoint:* About 61 mL/min, and 0.26 at 70 mL/min.

7.6 Theory and calculation · 6 min

Reaction 1 releases 120 kJ/mol. At the design point feed A enters at 47.62 mL/min and 1.00 mol/L in a reactor flow of 100 mL/min, and the stream carries $\rho c_p = 1.8\,\mathrm{MJ/(m^3\,K)}$.

(a) Compute the concentration of A at the mixer inlet, the heat that complete conversion would release, and the adiabatic temperature rise. *Checkpoint:* 0.476 mol/L, 95.2 W and 31.7 K.

(b) In case 5 the line runs at 130 mL/min, where zone 1 needs 128 W of cooling, when zone 1's circulator falls to 60 W. Taking the duty as proportional to the feed, find the feed at which 90 % of the remaining cooling suffices. *Checkpoint:* About 54.8 mL/min, the ceiling FY-100 jumps to at 11.3 min.

(c) Find the feed that would use the full 60 W, and compare it with Table 7.3. *Checkpoint:* About 61 mL/min, and the table gives 59 W at 60 mL/min.

7.7 App challenge · 8 min

Load scenario 3, Ratio error, with the default seed. Run at a speed of 300 and press Pause when the clock is between 62 and 90 min (the tiles do not change in that interval). Read the Off-spec released, Diverted and Shadow tiles of the scorecard.

(a) Report what RTD tracking released and diverted, and what the plug-flow shadow would have released and diverted on the same material. *Checkpoint:* 0 mL released, 1.53 L diverted (0.11 L in specification). Shadow 177 mL and 1.79 L, each within one unit of the last digit.

(b) Press Reset, set Material tracking to Plug-flow delay, and run to the same window. Report the Off-spec released, Worst released and Diverted tiles. *Checkpoint:* 177 mL, worst assay 94.98 %, 1.79 L diverted (0.55 L in specification), each within one unit of the last digit.

(c) The plug-flow image of the air lock reaches the valve from 48.65 to 52.65 min. Use Figure 7.2 to explain where the 177 mL came from, and why the plug-flow delay still diverted more than RTD tracking. *Checkpoint:* The front of the slug, before 45.50 min.

7.8 App challenge · 7 min

Load scenario 6, Slower analyzer, with the default seed. It runs the 10 min HPLC method, and an unassayed drum 10 % strong goes into service at 25 min. Run at a speed of 600 to any time after 1.5 h and read the Off-spec released, Worst released and Shadow tiles.

(a) Report the off-specification volume released, its worst assay, and what the plug-flow shadow would have released. *Checkpoint:* 803 mL at 96.23 %, shadow 1039 mL, each within 5 mL and 0.02 points.

(b) Press Reset, set Analyzer method run time to 7 min before pressing Run, and run again to any time after 1.5 h. Report the off-specification volume released. *Checkpoint:* 0 mL.

(c) Exercise 7.5 predicts that the 7 min method lets a change through at 0.29 of the arrival phases. Explain why the run of part (b) does not contradict it, and what a design check must use instead of one run. *Checkpoint:* Table 7.6: 3 of 12 arrivals escape with the 7 min method.

7.9 App challenge · 9 min

Load scenario 9, Rate and drum: quality trim against MPC, with the default seed. Quality control is on MPC, the residual A target at 0.8 %. Run at a speed of 300 and press Pause when the clock is between 1.9 and 2.1 h. Read the Diverted and Reagent B per mol P tiles. Make each change below after pressing Reset and before pressing Run, and repeat.

(a) Report both tiles under the MPC as loaded, and with Quality control set to Loops (the quality trim stays on). *Checkpoint:* MPC 0.00 L and 1.18 mol. Loops 4.85 L (2.04 L in specification) and 1.51 mol, within 0.03 mol.

(b) Under Loops, set Trim gain to 0.08, Trim integral time to 40 min and Residual A target to 0.4 %, the retuned trim of Table 7.8. Report both tiles. *Checkpoint:* 0.00 L and 1.18 mol.

(c) Set Quality control back to MPC, keeping the 0.4 % target. Report both tiles, and say what the four runs show about the plan and the target. *Checkpoint:* 0.00 L and 1.22 mol.

7.10 Application area · 4 min

A plant engineer proposes to replace the 1.5 L hold-up coil with a 1.0 L coil of the same construction to cut the hold-up, keeping the 3 min HPLC method.

(a) Check the proposal at a feed of 100 mL/min and at 160 mL/min, by the mean rule and by the conservative leading-edge rule. *Checkpoint:* At 100 mL/min the coil holds 8.33 min with its leading edge at 6.63 min, against $t_{det} = 6.5$ min. At 160 mL/min it holds 5.21 min, short of it.

(b) The supervisor's windows use the coil's 0.1 and 99.9 % points. What must be done before those points can be trusted on a real coil, and why? *Checkpoint:* A tracer test on the installed coil.

7.11 Application area · 3 min

In case 7 a deposit grows in zone 3. With the pressure limiter on, FY-100 holds the discharge near the 8 bar alarm by cutting the feed to 53.4 mL/min at two hours. Without it the 10 bar relay trips at 71.2 min.

(a) Compute the reactor residence time at 53.4 mL/min and the share of the design rate left, and say why the impurities rise from 0.54 to 0.707 %. *Checkpoint:* About 30.2 min and 53 %.

(b) What does a trip leave in the line, and what does case 8 suggest an operator should do with it? *Checkpoint:* Hot material that keeps reacting. Cooling the zones during a stop cuts the diversion from 3.85 to 1.86 L in case 8.

7.12 Review the agent · 5 min

An AI assistant is asked how to divert the material of a pump B air lock from 20 to 24 min at 100 mL/min. It replies: "The plug-flow time from the mixer to the valve is 28.65 min, so divert from 48.65 to 52.65 min, plus 1 min on each side for safety, from 47.65 to 53.65 min. A fixed delay is better than the app's plug-flow mode, which times the delay from a flow the air lock had already cut." Find the flaw, and say how you would check it.

CHAPTER 8

Energy Created by a Delayed Channel in a Bilateral Teleoperation Simulator

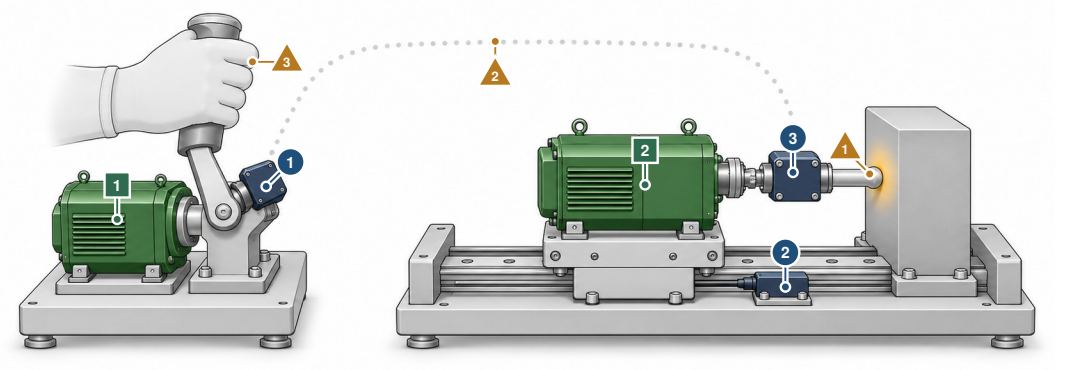

The controller moves 1 motors on the master handle 2 slave drive on the carriage

The controller reads 1 master handle position 2 slave carriage position 3 force sensor at the tool

What disturbs it 1 wall nearer or farther than its drawing 2 longer delay, lost packets 3 operator pushes harder

The handle and the slave share nothing but the delayed link drawn dotted, so the force the tool meets at the block reaches the hand late. Illustration, not to scale.

IN PLAIN TERMS Feeling a Wall Through a Slow Link

The main idea. A person moves a handle, and a robot far away copies the motion. When the robot touches a wall, its push is sent back to the handle, so the hand feels the wall too.

The challenge. The link takes time, so the push arrives after the hand has already moved on. The late push can shove the hand back and make the robot bang the wall, adding energy that no part of the machine put in.

Why it matters. Each fix for this costs something: a softer wall, a robot that lags the hand, or a guess about the wall that can be wrong.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the fact on the Moon's delay (Section 8.1), the two history segments (Section 8.2), the two try-its and the fact on move and wait (Section 8.6) and the key points
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 8.6), Technology in Practice (Section 8.7), Limitations (Section 8.8) and the two open challenges

The control question When a slave's contact force is sent back through a delayed link, the feedback that makes a remote wall feel solid can make it chatter, although the hand, the devices and the wall are each passive. Where does the energy come from, and what does each cure for it cost in feel, position or knowledge?

Learning objectives

- Explain how a delayed channel creates energy at its two ports, and read the ledger that measures it.
- Compute the gain limit K_h/K_{ser} of force reflection through a long delay, and say when the running model's verdict can miss it.
- Estimate the wall that wave variables render from the wave impedance and the delay.
- Compare wave variables, a time-domain passivity controller and model mediation on the same task by stability, felt stiffness, position and task time.
- Judge whether a contact that keeps still proves a passive channel.

Prerequisites Sections 1.4.2, 1.4.3, 1.6 and 1.6.4.

The app https://apmonitor.com/apps/teleoperation-delay/

See also Volume 20, Chapter 13 for a peg pushed into a chamfered hole under force control, Volume 20, Chapter 14 for a teleoperated console that scales motion and filters tremor, and Volume 7, Chapter 14 for the same delayed, lossy channel and a gain that must stay below one at every frequency.

8.1 Introduction

A teleoperator lets a person work where a person cannot go, inside a shielded cell, on the sea floor, inside a patient or on another world. The person moves a master and a slave copies the motion at the remote site. When the slave's contact force is sent back and rendered on the master, the hand feels the remote wall, and the interfaces between the operator and the master and between the slave and its environment become ports through which the teleoperator exchanges energy [1]. A transatlantic robot-assisted operation in 2001 [2] and a ground station that worked a satellite's arm bilaterally through a round trip of about 7 s in 1999 [3] put long links between the hand and the work.

DID YOU KNOW? Six to Eight Seconds to the Moon and Back

In NASA studies of 2023, operators on Earth drove a simulated lunar rover through round trips of 0, 4, 6 and 8 s, driving slower and braking more often under a delay. The mission's communications team now estimates the round trip at 6 to 8 s [4]. Case 9 assumes 2.6 s.

The delay in that link is what sets the problem apart. The hand, the master, the slave and the wall can each only store or dissipate energy, but a channel that returns the slave's force after a delay pairs it with a velocity it never met, and the pair can put out energy that no element put in. Passivity and scattering theory showed why existing bilateral systems are unstable for certain environments [5]. In the simulator of this chapter, force reflection with no delay holds 3.06 N at the wall for 3.06 N on the handle, and its energy ledger books 0.78 mJ over three presses. Given 100 ms of round trip, the same loop bounces off the wall every 0.33 s, and the channel has created 34.1 mJ by the end of the first touch and 15.5 J over 30 s. No element changed, so the energy came from the channel, and a model of the slave with a dead time appended, however exact, could not have shown where it came from.

Three expectations a student brings to this plant fail on it, each by a measured amount. Lowering the reflection gain is expected to trade a little feel for stability, yet the sampled loop, linearized in contact, keeps full reflection stable only below a round trip of about 10 ms, and at 500 ms only below a gain of 0.2245, so that at 0.20, the largest setting of the gain slider below that limit, a press of 3.75 N on the handle drives the peg into the wall at 23.7 N. A contact that keeps still is taken as proof of a passive channel, but at 100 ms and a gain of 0.30 the contact keeps still while the channel creates 110 mJ, and even with no delay a gain of 0.5 creates 12.2 mJ. A cure that makes the channel passive is

expected to cost little, although wave variables, which never create energy at a constant delay, render the wall at 173 N/m against 1243 N/m without delay, and through a dropout a held wave leaves the peg pressing the wall 30.4 mm ahead of the master.

Teleoperation is one member of a wide family. Two dynamic systems that exchange force and motion across a delayed or sampled link can turn the interface between them into a source of energy, and the stability of the pair then rests on the interface's energy books rather than on either system alone. Haptic interfaces that render stiff virtual walls are held stable by the same bookkeeping, demonstrated with 15 ms of delay [6], telemanipulation over the Internet meets delays that vary [7], and networked control systems in general contend with sampling, network delay and packet dropouts [8]. A congestion controller whose dead time is the round trip of the network path it regulates is the subject of Volume 2, Chapter 7.

However, each remedy for delay was demonstrated on its own hardware and judged by its own measure, the completion time of a task, the stability of a contact or the feel of a wall [3, 9–11], and no common bench books the energy of every architecture in one way, whether or not its controller looks at energy. This chapter locates the source of the energy a delayed force-reflecting loop creates, on a teleoperation simulator whose channel carries a ledger at its two ports, and it prices each passivity remedy in feel, position and knowledge. Every architecture used comes from the literature. The teleoperator is described completely enough to be rebuilt. Its two sides learn of each other only through seeded links, its ledger multiplies the same forces and velocities the controllers exchange, and its operator perceives the wall through a filtered feel and a delayed video. The loop is linearized as the application samples it, and the gain limits are compared with the running model. Nine cases and a graded challenge isolate the delay, the gain, the architecture, the network and the model.

Section 8.2 follows remote manipulation from the hot cell to model mediation. Section 8.3 builds the devices, the site, the channel with its energy ledger and the scripted operator, and Section 8.4 the four architectures and the passivity controller, period by period. Section 8.5 derives the gain limit of force reflection through a delay and checks it against the running model. Section 8.6 reports what each case and the challenge measured, Section 8.7 sets the teleoperator beside real sensors, links and practice, Section 8.8 states the limitations and Section 8.9 draws the conclusions.

8.2 From Hot Cells to Wave Variables

Remote manipulation began behind radiation shielding, in the hot cell, with the mechanical master-slave manipulator that the history segment below describes. The experiments of the 1960s separated the two sides by a transmission delay.

> ⧖ HISTORY **1949: Tongs That Followed a Handle**
>
> A report from Argonne National Laboratory dated March 1949 described a device, built and tested, for handling a pair of tongs behind a shielding barrier such as a lead glass window. It was named a master-slave manipulator because the slave tongs moved in exact correspondence with a master handle about three feet away. The test model was entirely mechanical: seven independent linkages of tubes, gears, drums and cables carried three straight motions, three rotations and the closing of the jaws. Most people learned to work it within five to fifteen minutes. Friction and mass in all seven motions had to be kept very low, and backlash impaired operation. The report called it the first completely master-slave manipulator known to exist [12].

Those experiments measured task time. A study of 1963 described several manipulator devices, reported an experiment on the effect of delay on the completion time of a simple manipulative task, and found that under certain circumstances the time with delay could be predicted from measures of performance taken without it [13]. Two years later most operators faced with a delay spontaneously adopted a move-and-wait strategy, and the completion time depended on the number of open-loop moves, which a simple statistical model predicted from the difficulty of the task [14]. Force feedback

was tried next. With long delays, positioning by force feedback alone was possible, but delay coupled with a high loop gain created serious instability, and alternative displays of the feedback force overcame it [9].

Space gave the delay its size. A review of 30 years of research on delay in human teleoperation in space described the experiments on human performance, predictive displays that help the operator overcome the delay, supervisory control ranging from local impedance control on contact to higher-level automation, and wave transformations, and it observed that space teleoperation had tended to avoid the delay by not teleoperating from the ground [15]. In November 1999 a ground station teleoperated the arm of the Engineering Test Satellite 7 bilaterally through a round trip of about 7 s, with a teleoperator built to stay stable through it, and slope-tracing and peg-in-hole tasks showed that kinesthetic force feedback improved performance even at that delay [3].

The account of why delayed force reflection fails came from network theory. A hybrid two-port model treated the interfaces with the operator and with the environment as ports through which the teleoperator exchanges energy, and compared measured or simulated systems with the ideal one [1]. Passivity and scattering theory then explained why existing bilateral systems were unstable for certain environments and gave a law that is stable for any environment and any time delay [5]. Wave variables gave that idea a physically motivated, passivity-based formalism with energy conservation and stability guarantees in the presence of transmission delays [16]. A general architecture that transmits force and velocity in both directions showed that transparency and robust stability (passivity) are conflicting design goals [17], and later work examined the wave reflections and spurious dynamics of passive transmission and laid the basis for applying wave variables to the variable delays of Internet-based telemanipulation [7].

> ⧖ **HISTORY** **1989: The Operator's Arm Sets the Limit**
>
> A paper at a NASA conference on space telerobotics in 1989 modeled bilateral force reflection on one axis for a space robot then in development. In its model the operator's arm and the hand controller play the part that a programmed filter plays in impedance control. The analysis showed that delays as small as 10 ms, from communication or computation, can significantly raise the stiffness and damping the operator must supply to keep the loop stable, and that stiff work pieces raise them further. Tests of five subjects found the operator's stiffness and damping tightly coupled. Scaling the force or position gains can stabilize any such system, the paper concluded, but at the cost of reducing the feel to the operator [18]. The bound of Equation (8.15) makes the same point for this chapter's hand.

Time-domain passivity control replaced the frequency-domain guarantee with bookkeeping. A passivity observer measures the energy flowing in and out of a subsystem in real time, and a passivity controller, an adaptive dissipative element, absorbs at each sample exactly the net energy output the observer measures, which kept a haptic interface stable above 100 N/mm of stiffness or with 15 ms of delay [6]. The method was extended to the two-port network of a teleoperator, stable in contact with a hard wall above 150 kN/m [10], and to communication delays that vary in time [19]. Model-mediated teleoperation went the other way. Instead of sending the slave's sensory data, the slave abstracts it into a simple model of the environment that the master renders without lag, and the slave executes only commands consistent with the model, shown on a one-degree-of-freedom system with 4 s of round-trip delay [11]. The simulator of this chapter books the energy at both ports for every one of these architectures, whether or not its controller uses energy, from the forces and velocities its two controllers exchange.

8.3 The Teleoperator, Its Channel and Its Site

The two devices appear along one axis in Figure 8.1(a) and the site they work on from above in Figure 8.1(b). Master and slave move in one shared frame, in meters, with x running across the wall and y toward it, and their two axes are independent. The integrated state has ten entries, the master's

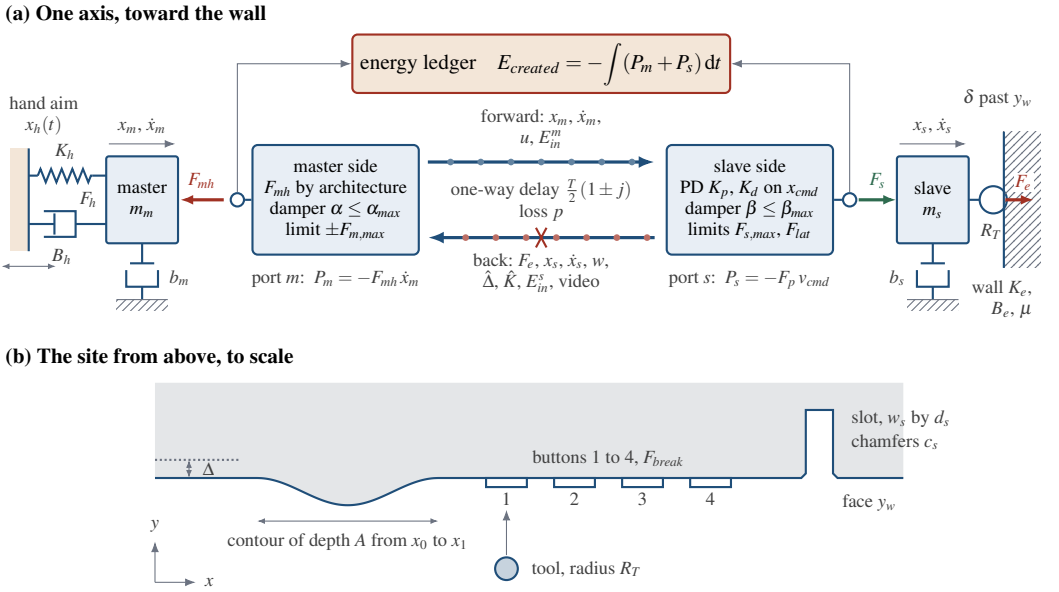

Figure 8.1 The teleoperator modeled by the application. (a) One axis, toward the wall. The hand is a spring K_h and a damper B_h about its aim $x_h(t)$ and pushes the master handle m_m, damped by b_m, while the master's motors push back with F_{mh}. Two links carry one packet each way every millisecond, with a one-way delay that wanders within $\pm j$ of $T/2$ and a loss probability p. The slave drive F_s, a PD on the delayed command limited to $F_{s,max}$ and to F_{lat} across the wall, moves the carriage m_s, damped by b_s, whose tool of radius R_T meets a wall of stiffness K_e, damping B_e and friction μ with the force F_e at a penetration δ past the face y_w. The ledger integrates the power P_m and P_s entering the channel at its two ports. (b) The site in the slave's frame, to scale: the face with its contour bump, four buttons 4 mm proud of it and the slot with its chamfers. The dotted line is the drawing that the operator and the model's prior start from, which in case 7 lies 8 mm beyond the real face, a site offset $\Delta = -8$ mm.

position and velocity $\mathbf{x}_m = (x_m, y_m)$ and $\dot{\mathbf{x}}_m$, the slave's $\mathbf{x}_s$ and $\dot{\mathbf{x}}_s$, and the friction bristles z_L and z_R of the two wall blocks beside the slot, and Table 8.1 classifies the signals. Whatever is discrete, the clicks of the buttons, the seating of the peg, the slip of a bristle, the count of reversals and the oscillation flag, is handled between integration steps and kept out of the derivative function.

8.3.1 The Devices

The master is a mass under the hand and its motors, and the hand is a spring and a damper about the point it aims at, as shown in Equation (8.1).

$$\begin{aligned} m_m \ddot{\mathbf{x}}_m + b_m \dot{\mathbf{x}}_m &= \mathbf{F}_h + \mathbf{F}_{mh} \\ \mathbf{F}_h &= K_h(\mathbf{x}_h - \mathbf{x}_m) + B_h(\dot{\mathbf{x}}_h - \dot{\mathbf{x}}_m) \end{aligned} \tag{8.1}$$

where $\mathbf{x}_h$ is the hand's aim and $\mathbf{F}_{mh}$ the force the motors put on the handle, both held over each 1 ms control period, and b_m is the handle's friction together with the damping its own controller adds. The slave carriage obeys Equation (8.2).

$$m_s \ddot{\mathbf{x}}_s + b_s \dot{\mathbf{x}}_s = \mathbf{F}_s - \mathbf{F}_e \tag{8.2}$$

where $\mathbf{F}_s$ is the drive force and $\mathbf{F}_e$ the force the tool exerts on the site, which the slave's force sensor reads. The two sides share one frame at one scale, so a position on the master is the position it commands at the site.

Table 8.1 Signals of the Teleoperator by Role, With Their Defaults

Symbol	Role	Default	Range
$\mathbf{x}_m$, $\dot{\mathbf{x}}_m$	states, master handle	at the start	site
$\mathbf{x}_s$, $\dot{\mathbf{x}}_s$	states, slave, CV of its PD	at the start	site
z_L, z_R	states, friction bristles	0	$\lvert z\rvert \le \mu F_n/K_t$
$\mathbf{F}_{mh}$	MV, master motors, held 1 ms	0	±20 N an axis
$\mathbf{F}_s$	MV, slave drive, held 1 ms	0	±50 N, F_{lat} in x
$\mathbf{x}_h$	hand's aim, the student's in MANUAL	at the start	site
T, j, p	channel round trip, jitter, loss	0, 0, 0	to 2.6 s, 20, 20 %
—	DV, operator pushes harder	100 %	100–300
Δ	DV, site offset from its drawing	0	±15 mm
w_K	DV, wall stiffness over nominal	100 %	50–400
—	DV, slave drive capacity	100 %	0–100
—	DV, satellite hop, loss burst	events	+500 ms, 1 s
$\mathbf{F}_e$	measured, slave force sensor	—	—
$\bar{F}_h$	perceived, felt force over 50 ms	—	—
P_m, P_s, $E_{created}$	outputs, the energy ledger	0	—

8.3.2 Contact With the Site

The tool is a disc of radius R_T seen end-on. The site is the solid beyond the face $y > y_f(x) + \Delta$, with Equation (8.3).

$$y_f(x) = y_w - A\sin^2\left(\pi\,\frac{x - x_0}{x_1 - x_0}\right), \qquad x_0 < x < x_1 \tag{8.3}$$

and $y_f = y_w$ elsewhere, interrupted by the slot. Seven bodies can touch the tool, the two wall blocks on either side of the slot (each its face, its chamfer and its slot wall, convex near the slot so that each makes one contact), the slot bottom and the four buttons. For each body i the penetration is $\delta_i = R_T - d_i$, with d_i the distance from the tool's center to the body, signed negative inside it, and $\mathbf{n}_i$ the unit normal from the body toward the tool. The rigid bodies push back with Equation (8.4).

$$F_{n,i} = \max\Big(0,\ w_K K_e\,\delta_i + B_e\,\dot{\delta}_i \min\left(1,\ \delta_i/\delta_r\right)\Big) \tag{8.4}$$

where $\dot{\delta}_i$ is the rate of penetration along $\mathbf{n}_i$, w_K the wall stiffness disturbance, and the damping fades in over the first $\delta_r = 0.5$ mm so the force is continuous at the instant of touch. A button is a plate 4 mm proud of the face whose penetration δ is its travel, and its snap spring is Equation (8.5).

$$f_b(\delta) = \begin{cases} F_{break}\,\delta/p_c, & \delta < p_c \\ F_{break} - (F_{break} - F_{low})\,\dfrac{\delta - p_c}{p_s}, & \delta < p_2 \\ F_{low} + (F_{bot} - F_{low})\,\dfrac{\delta - p_2}{h_b - p_2}, & \delta < h_b \\ F_{bot}, & \delta \ge h_b \end{cases} \tag{8.5}$$

with $p_2 = p_c + p_s$, to which a damper $c_b\dot{\delta}\min(1, \delta/\delta_r)$ is added. Once a button bottoms out at $h_b = 4$ mm the tool's front reaches the face behind it and the wall block pushes through Equation (8.4). A button clicks when its travel passes $p_c = 1.5$ mm and re-arms below 0.5 mm, and it lights only in the order 1 to 4, a press out of order starting the sequence again.

Friction acts on the two wall blocks as an elastic bristle along the contact's tangent $\mathbf{t}_i = (-n_{iy}, n_{ix})$, as shown in Equation (8.6).

$$\begin{aligned} &\dot{z}_k = v_{t,k}, \qquad F_{t,k} = -\,\mathrm{clip}\left(K_t z_k + C_t v_{t,k},\ \pm\mu F_{n,k}\right) \\ &z_k \leftarrow \mathrm{clip}\left(z_k,\ \pm\mu F_{n,k}/K_t\right) \quad \text{after each step} \end{aligned} \tag{8.6}$$

where $k \in \{L, R\}$, $v_{t,k}$ is the tool's speed along $\mathbf{t}_k$, and z_k is set to zero whenever its block is out of contact. The buttons and the slot bottom are smooth. The force on the tool is the sum $-\mathbf{F}_e = \sum_i (F_{n,i}\mathbf{n}_i + F_{t,i}\mathbf{t}_i)$. The peg is seated when it is inside the slot within 0.8 mm of the bottom.

8.3.3 The Channel and the Ledger

Each side knows the other only through two links, one each way, built on the framework's shared delay line, `shared/src/framework/core/channel.js`, with its wandering-jitter option on. The same module carries the radio of the vehicle platoon of Volume 7, Chapter 14. Every 1 ms each side sends one packet, which arrives as given by Equation (8.7).

$$t_{arr} = t_k + \tfrac{T}{2}(1 - j) + T\, j\, \omega(t_k) \tag{8.7}$$

where t_k is the time it was sent, T the round trip, j the jitter share and $\omega \in [0, 1]$ a wandering level that ramps linearly, over each interval of $\tau_j = 0.5$ s, to a new value drawn uniformly at the start of the interval, so that the one-way delay wanders within $\pm j$ of $T/2$ as packets behind one queue do. A packet is lost with probability p, independently, and every packet of a burst is lost. Each packet draws two uniform numbers from its link's own generator, whatever the settings, and each link is seeded from the case's seed and its own index, so the losses of one link cannot depend on anything else. A packet that arrives after a newer one is thrown away, so what a receiver holds never goes back in time, and through a gap it keeps the last packet it had. At reset both links are filled as if they had been up with both devices at rest.

The ledger books the power that enters the channel at each port, as shown in Equation (8.8).

$$\begin{aligned} P_m &= -\mathbf{F}_{mh} \cdot \dot{\mathbf{x}}_m, \qquad P_s = -\mathbf{F}_p \cdot \mathbf{v}_{cmd} \\ E_{created}(t) &= -\int_0^t (P_m + P_s)\, \mathrm{d}t \end{aligned} \tag{8.8}$$

where $\mathbf{v}_{cmd}$ is the velocity the channel asks of the slave and $\mathbf{F}_p$ the force the architecture reports at the slave's port, the measured $\mathbf{F}_e$ for position–force and the drive force $\mathbf{F}_s$ otherwise. With no delay and a reflection gain of one, position–force hands the slave the master's velocity and the master the slave's force within one period, so the two terms cancel but for that period. A negative $E_{created}$ is energy in flight, and a positive one is energy the channel has put out that neither port put in.

🕮 **Pause.** Try Exercise 8.1 before reading on.

8.3.4 The Scripted Operator

In AUTO the hand's aim follows a script that reacts only to what an operator can perceive, the force the hand pushes with, filtered over 50 ms as $\bar{F}_h$, and the video of the slave, the lamps and the peg, which arrives with the slave's packets one delay late. A task is a list of steps. The hand moves its aim at 60 mm/s in free space, 30 mm/s over the last 15 mm before an expected surface and 15 mm/s while pressing. It takes a felt force above 0.6 N for contact only once the video shows the slave stopped short of the aim, except under model mediation, where the force is rendered locally and is trusted at once, and it takes one above 4 N for contact whatever the video shows. A step that ends on something perceived ends 0.2 s later, and the hand meanwhile keeps doing what it was doing, except that it stops advancing at once at a felt 8 N. After a rest of 1 s the wall task presses three times at $x = -140$ mm until $\bar{F}_h$ reaches 3 N under the same rule, holds each press 1.5 s, and backs off to 30 mm in front of the face with 0.5 s of rest between presses. A button is pressed until the video shows its lamp, a felt click (a drop of 0.5 N after 1.4 N) or the 8 N cap, and pressed again up to twice if the lamp stays dark a round trip and a second later. The contour is traced from 10 mm before x_0 to 10 mm past x_1 at 20 mm/s, aiming 3 mm past the drawn face. The peg is lined up 3.5 mm off the slot's center and pushed in, and once the video shows it stalled for 1 s while the hand feels more than 1 N, the aim moves across at 3 mm/s, by at most 8 mm, toward the slot as seen. Under the move-and-wait strategy

Table 8.2 Constants of the Model. V Is Adjustable in the Show Model Dialog, H Is Kept Out of It and C Is Fixed in the Source

Symbol	Meaning	Value	Class
Devices			
m_m, b_m	master mass, damping	0.5 kg, 20 N s/m	V
$F_{m,max}$	master force limit, an axis	20 N	V
m_s, b_s	slave mass, damping	2.0 kg, 5 N s/m	V
Site			
K_e, B_e	wall stiffness, damping	5000 N/m, 80 N s/m	V
μ	friction on the wall blocks	0.4	V
δ_r	damping fade-in	0.5 mm	C
K_t, C_t	bristle stiffness, damping	5000 N/m, 60 N s/m	C
R_T	tool radius	5 mm	C
y_w, A	face, contour depth	150, 12 mm	C
x_0, x_1	contour ends	−120, −40 mm	C
—	buttons, centers	−10, 20, 50, 80 mm	C
—	buttons, width by height	18 by 4 mm	C
F_{break}	button break force	2.0 N	H
p_c, p_s, h_b	click, snap, travel	1.5, 0.3, 4 mm	C
F_{low}, F_{bot}, c_b	after the snap, bottom, damping	0.8, 1.2 N, 4 N s/m	C
—	slot center	128 mm	C
w_s, d_s, c_s	slot width, depth, chamfer	12, 30, 2 mm	C
—	clearance a side	1 mm	H
Channel			
τ_j	period of the wandering delay	0.5 s	C
—	burst, satellite hop	1 s, +500 ms	C
—	packets in flight a link	4096	C

the aim moves one increment at 30 mm/s, 10 mm in free space and 2 mm near or in contact, then holds until the video shows the slave within 1.5 mm of it, or a round trip and a half and 0.3 s have passed, and then takes the reaction time. In MANUAL the two aim sliders, moved by the arrow keys, are the student's hand. The operator's stiffness K_h and damping B_h are illustrative values, which the application attributes to the haptics literature.

8.3.5 Counting an Oscillation

The scorecard counts a force reversal each time the contact force turns from rising to falling or back by at least 1.5 N, ignoring a turn within 60 ms of the last and the first turn after the hand starts to back off. A contact episode lasts through bounces while the hand still aims past the slave and ends 50 ms after contact is lost with the hand no longer pushing, and the status reads CONTACT OSCILLATION once an episode holds four reversals. A click starts the episode's count again, and the reversals of the following 0.3 s belong to the snap. An active flag also clears after 1 s without a qualifying reversal, even if the hand keeps aiming past the slave. These are teaching thresholds.

8.3.6 Start, Step and Seed

Every case begins at rest, with the hand's aim, the master and the slave together at the task's starting point, at $y = 100$ mm and so 50 mm short of the face, and both bristles relaxed. The engine is the one of Section 1.6.1. A classical Runge-Kutta step of fourth order and 1 ms integrates the ten states, the two controllers act once a millisecond, and the historian records at 50 Hz, with a 1 kHz scope of the contact force and the ledger beside it. Seeded with 20261001, a case replays bit for bit. With the physics step halved and the control period kept at 1 ms, the press force with no delay and the energy created over the first contact at 100 ms move by a relative 6.4×10^{-6}. All 39 of the application's checks and its 11 checks of the scenario completion rules pass.

8.3.7 Assumptions

Two decoupled axes and a disc. Each device moves in x and y without rotation, so the peg is a disc seen end-on. It can stick by friction but cannot wedge as a real peg does, and the slot task is a

Table 8.3 Controller and Operator Settings, Defaults and Ranges

Symbol	Setting	Default	Range
—	architecture	position–force	4 kinds
k_f	reflection gain	1	0–1.5
K_p	slave stiffness, N/m	2000	100–5000
K_d	slave damping, N·s/m	80	0–200
$F_{s,max}$, F_{lat}	slave force limits, N	50, 50	5–80, 1–50
b	wave impedance, N·s/m	50	5–500
—	lost or late waves	held	or rebuilt
—	passivity controller	off	on
α_{max}	master damper limit, N·s/m	500	0–1000
β_{max}	slave damper limit, mm/s per N	100	0–200
f_u	model update rate, Hz	10	1–50
$\hat{K}_0$	wall stiffness prior, N/m	5000	500–20000
K_h, B_h	hand, N/m and N·s/m	300, 10	to 2000, 60
—	strategy	continuous	move and wait
—	task	three presses	5 tasks

sideways contact rather than a jam.

One frame at one scale. A real master is usually smaller than its slave and its motions are scaled, which multiplies the two ports' powers by different factors and moves every gain limit. A surgical console that scales and filters the surgeon's motion is treated in Volume 20, Chapter 14.

A linear hand on a script. The operator is a spring and a damper of fixed values about a scripted aim, with a fixed reaction time. Real operators stiffen their grip, slow down and adapt to a delay, and the grip alone moves the limits of Section 8.5.

Damped devices. The master carries 20 N s/m of damping and the wall 80 N s/m. With a bare master of 2 N s/m against a wall of 20 N s/m, the loop of Section 8.5 with no delay is stable only below a reflection gain of 0.43, and the running model keeps still at 0.45 and oscillates at 0.5, so these values are what makes full reflection workable at all.

Ideal sensors. Each device measures its own position and velocity exactly, and the slave's force sensor reads the contact force without noise, lag or quantization, so every delay in the loop belongs to the channel and the holds.

A clean channel. Packets are delayed by a wandering amount, lost and dropped in bursts but never corrupted or duplicated, the two ends share a clock, and bandwidth, queueing that depends on the traffic and the compression of the video are not modeled.

Illustrative numbers. The dimensions, stiffnesses and thresholds are values the application chose for teaching, not data from a particular teleoperator.

8.4 The Controllers, Period by Period

Within each millisecond the two sides take turns. The master takes the newest packet from the slave, computes $\mathbf{F}_{mh}$ by the architecture in force, applies the passivity damper if it is on and the motors' limit, and sends its packet. The slave then takes the newest packet from the master, which with no delay is the one just sent, computes its command and its drive force, applies its own damper and limits, and sends its packet, which the master reads on the next period. With no delay the shortest round trip is therefore one period of transport plus the two holds. The ledger books the period's power, and last the operator moves the aim. The settings are listed in Table 8.3.

8.4.1 Position–Force and Position–Position

Position–force sends the master's position and velocity forward and the slave's measured contact force back, as shown in Equation (8.9).

$$\begin{aligned} \mathbf{F}_s &= \mathrm{sat}\left(K_p(\mathbf{x}_{cmd} - \mathbf{x}_s) + K_d(\mathbf{v}_{cmd} - \dot{\mathbf{x}}_s)\right) \\ \mathbf{F}_{mh} &= -k_f\,\tilde{\mathbf{F}}_e \end{aligned} \tag{8.9}$$

where $\mathbf{x}_{cmd}$ and $\mathbf{v}_{cmd}$ are the master's delayed position and velocity, $\tilde{\mathbf{F}}_e$ the slave's delayed force, and the saturation limits the drive to $F_{s,max}$ toward the wall and to $\min(F_{lat}, F_{s,max})$ across it, both scaled by the drive capacity. Position–position runs the same slave and returns the slave's position instead, so the master pulls the handle toward where the slave was, $\mathbf{F}_{mh} = k_f[K_p(\tilde{\mathbf{x}}_s - \mathbf{x}_m) + K_d(\tilde{\dot{\mathbf{x}}}_s - \dot{\mathbf{x}}_m)]$, and the slave's port force is $\mathbf{F}_s$. In both, the force on the handle is clipped to $\pm F_{m,max}$ per axis after any damper.

8.4.2 Wave Variables

The wave architecture sends the forward wave u and returns w, with the wave impedance b, as shown in Equation (8.10).

$$\begin{aligned} \mathbf{F}_{mh} &= \sqrt{2b}\,\tilde{\mathbf{w}} - b\,\dot{\mathbf{x}}_m, \qquad & \mathbf{u} &= \frac{b\,\dot{\mathbf{x}}_m - \mathbf{F}_{mh}}{\sqrt{2b}} \\ b\,\mathbf{v}_{sd} &= \sqrt{2b}\,\tilde{\mathbf{u}} - \mathbf{F}_s, \qquad & \mathbf{w} &= \frac{b\,\mathbf{v}_{sd} - \mathbf{F}_s}{\sqrt{2b}} \end{aligned} \tag{8.10}$$

where the slave's command $\mathbf{x}_{sd}$ integrates $\mathbf{v}_{sd}$. The slave solves its PD, $\mathbf{F}_s = K_p(\mathbf{x}_{sd} + \mathbf{v}_{sd}\Delta t - \mathbf{x}_s) + K_d(\mathbf{v}_{sd} - \dot{\mathbf{x}}_s)$, together with the wave equation in closed form within the period Δt, and a force over its limit is clipped with $\mathbf{v}_{sd}$ recomputed from the clipped force. At each end the power entering the network is then half the square of the wave sent less half the square of the wave received, so a channel that only delays the waves holds the wave energy in flight. A receiver either holds the last wave through a gap or reconstructs it, scaling the wave it uses by Equation (8.11).

$$\kappa = \min\left(1,\ \sqrt{\frac{\max(0,\ \tilde{E}_w - E_{used})}{\frac{1}{2}\,\tilde{w}^2\,\Delta t}}\right) \tag{8.11}$$

where $\tilde{E}_w$ is the sender's cumulative wave energy carried by the newest packet and E_{used} what this receiver has spent, so a lost packet's energy is spent later but never twice.

8.4.3 The Time-Domain Passivity Controller

Over position–force, position–position or wave variables, each port keeps the energy that has entered the network through it, E_{in}, apart from the energy that has left, E_{out}, and sends E_{in} with every packet. A port may put out only what the other side has put in, as far as the channel has told it, and a variable damper takes the excess, as shown in Equation (8.12).

$$\begin{aligned} \varepsilon &= -P_0\,\Delta t - \max\left(0,\ \tilde{E}_{in}^{other} - E_{out}\right), \qquad & P_0 < 0 \\ \alpha &= \min\left(\alpha_{max},\ \frac{\varepsilon}{\Delta t\,|\dot{\mathbf{x}}_m|^2}\right), \qquad & \mathbf{F}_{mh} \leftarrow \mathbf{F}_{mh} - \alpha\,\dot{\mathbf{x}}_m \\ \beta &= \min\left(\beta_{max},\ \frac{\varepsilon}{\Delta t\,|\mathbf{F}_p|^2}\right), \qquad & \mathbf{v}_{cmd} \leftarrow \mathbf{v}_{cmd} - \beta\,\mathbf{F}_p \end{aligned} \tag{8.12}$$

where P_0 is the port's power before its damper and each damper is zero unless $\varepsilon > 0$. The master's damper is a damping on the handle, and the slave's an admittance that backs the command off along the port force. For position–force and position–position that change of velocity also moves the commanded position, and the drift it leaves is walked back at 20 mm/s while the slave's damper is idle and the walk puts no energy out at the port, or puts out only what is in hand. A clipped damper lets the rest through, and the ledger shows it.

8.4.4 Model Mediation

Under model mediation the master renders a model instead of the slave's force. It steps a local copy of the slave, with the same mass, damping, PD and limits, four times per period against the site's drawing shifted by an estimated offset $\hat{\Delta}$ with stiffness $\hat{K}$ and damping B_e, without friction, and renders $\mathbf{F}_{mh} = -k_f\mathbf{F}_{copy}$, the copy's PD force. Like the Smith predictor of Section 1.4.3, the master

closes its fast loop on an undelayed model and leaves the delayed measurements to correct the model. The slave tracks the delayed master as in position–force and fits the site from its own contacts by recursive least squares, as shown in Equation (8.13).

$$\begin{aligned} \delta_{prior} &= F_n/K - \Delta n_y = \varphi^{\mathsf{T}}\theta \\ \varphi &= [F_n,\ -n_y]^{\mathsf{T}}, \qquad \theta = [1/K,\ \Delta]^{\mathsf{T}} \end{aligned} \tag{8.13}$$

where δ_{prior} is the tool's penetration into the drawing, F_n the measured normal force and n_y the drawing's normal. A sample is taken only while the tool presses one rigid body, no button, above 0.5 N and slower than 5 mm/s along the normal. The update forgets with $\lambda = 0.9995$, starts from the prior $\hat{K}_0$ and zero offset with uncertainties of 50 % on $1/K$ and 20 mm on Δ against 0.2 mm of misfit per sample, keeps $1/K$ between $1/50\,000$ and $1/500$, and is published at f_u. The master moves $\hat{\Delta}$ toward each new estimate at no more than 10 mm/s and $\hat{K}$ with a 0.2 s lag, so a correction never kicks the hand.

8.4.5 Limits, Signs and the Models Kept Apart

No loop of this application uses the velocity-form PID of Section 1.6.2. The slave's PD and the master's laws are written out inline with no integral term, which leaves nothing to wind up and no need for anti-windup, and none of the forces is rate-limited. The derivative terms take the velocities as measured, the slave's own and the master's as carried by its packet, with no differencing and no filter. Clips, applied after the dampers, are the only output limits. The application's own sign checks confirm that the slave pushes toward a command ahead of it on each axis, that position–force pushes the handle back from a slave pressing into the wall and that a forward wave moves the slave forward. The four-channel architecture was left out, because it gives the ledger no single pair of force and velocity at each port.

The model of Section 8.3 is not one that any controller holds. The passivity controller keeps only energy books, and under model mediation the master keeps a copy of the slave against its estimated site. The linearization of Section 8.5 is a third model, built for analysis, which the application never runs.

8.5 The Gain Limit of Force Reflection

Pressing a flat part of the face along y, with the contact's damping fully faded in and the aim held, position–force is a linear loop. The delay sits inside a loop closed through the operator's hand, and, as in Section 1.4.2, it caps what the loop may do. In continuous time, with an effective delay T_e, its loop gain is Equation (8.14).

$$\begin{aligned} L(s) &= \frac{k_f\,(K_e + B_e s)(K_p + K_d s)\,\mathrm{e}^{-sT_e}}{D_m(s)\,D_s(s)} \\ D_m(s) &= m_m s^2 + (b_m + B_h)s + K_h \\ D_s(s) &= m_s s^2 + (b_s + K_d + B_e)s + K_p + K_e \end{aligned} \tag{8.14}$$

where D_m is the master under the hand and D_s the slave pressing the wall. At zero frequency $L(0) = k_f K_{ser}/K_h$, where $K_{ser} = K_p K_e/(K_p + K_e) = 1429\,\mathrm{N/m}$ is the slave's PD and the wall in series. Through a long delay the phase of $\mathrm{e}^{-j\omega T_e}$ passes $-\pi$ at frequencies only $2\pi/T_e$ apart, so the loop is stable at every long delay only if $|L(j\omega)| < 1$ at all frequencies. For the 300 N/m hand, whose master and grip are overdamped, the largest $|L|$ is at zero frequency, which gives the bound of Equation (8.15).

$$k_f < \frac{K_h}{K_{ser}} = \frac{300}{1429} = 0.21 \tag{8.15}$$

Through a long delay a wall can be rendered no stiffer than the hand that holds it.

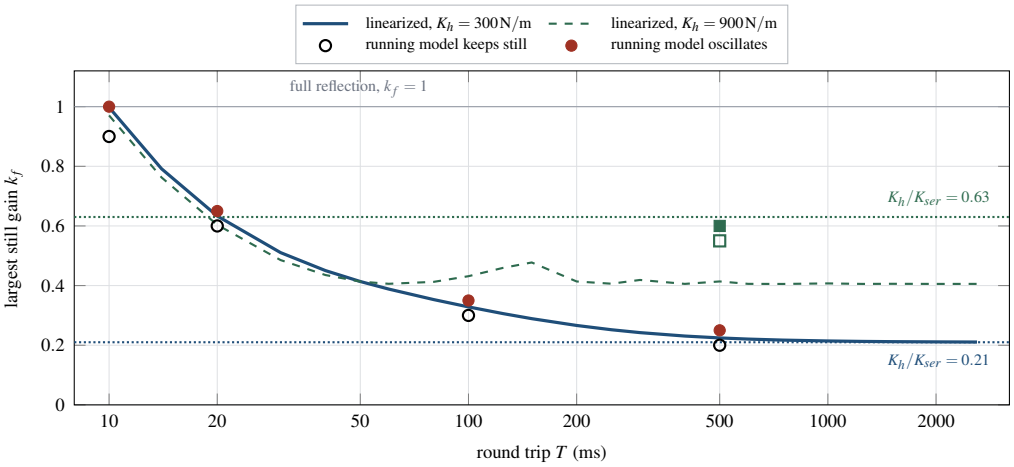

Figure 8.2 The largest reflection gain at which position–force keeps a press still, against the round trip. The curves are the sampled loop linearized in contact for the 300 N/m hand (solid) and for a 900 N/m grip with the same damping (dashed), and the dotted lines their static bounds K_h/K_{ser}. The marks are the running model's verdicts on the three presses of the wall task at each round trip, kept still (open) or flagged as an oscillation (filled), circles for the 300 N/m hand and squares for the firmer grip.

The numbers in this section come from the loop as the application samples it rather than from Equation (8.14). Both plants are discretized exactly over the 1 ms period with the forces held, the slave's PD acts on its own sampled state, the master reads the force N_b periods late and the slave reads the master N_f periods late, with $N_f = N_b = T/(2\Delta t)$, and $N_f = 0$ and $N_b = 1$ with no delay. For each T the largest stable k_f is the gain margin at the frequencies where the delayed loop's phase is $-\pi$. Equation (8.14) agrees with it within 0.5 % at every round trip when T_e is the round trip plus 1 ms for the two half-period holds, and 2 ms with no delay, where the master reads the slave one period late.

The limit falls fast and then levels, as shown in Figure 8.2. It is 1.00 at a round trip of 10 ms, 0.633 at 20 ms, 0.329 at 100 ms, 0.252 at 250 ms, 0.2245 at 500 ms and 0.211 at 2.6 s, approaching the bound of Equation (8.15) from above. Where the instability grows quickly the running model agrees, still at 0.9 and oscillating at 1.0 for 10 ms, still at 0.60 and oscillating at 0.65 for 20 ms, and still at 0.30 and oscillating at 0.35 for 100 ms. At 500 ms it keeps still at 0.20 and oscillates at 0.25, the settings of the slider's 0.05 grid on either side of the limit. Full reflection, the transparent setting, therefore needs a round trip under about 10 ms.

With no delay the limit depends on the press. The 2 ms of sampling leave a limit of 4.51 when the wall's damping is fully engaged, 0.58 with none, and 2.53 with the 61 % of it that a press of 1.53 N engages. A 3.06 N felt press at gain k_f pushes the peg with $3.06/k_f$, so the limit solves itself at 2.20. The running model, given gains beyond the slider's maximum of 1.5 in a counterfactual run, keeps still at 2.0 and oscillates at 2.5. An earlier linearization that left the wall's damping out of the reflected force gave 1.28 here.

Near the boundary the instability is slow, and a verdict from three presses can miss it. At 500 ms the dominant root of the sampled loop, the root of its characteristic equation with the largest modulus, grows by a factor of 14.7 a second at full reflection, with a period of 1.11 s, but at 0.248, past the limit, only by 1.20 a second, while at 0.22 it decays by 3.5 % a second. The firmer grip shows the same effect more strongly. At 900 N/m with 10 N s/m the master follows the hand's force nearly undiminished up to about 40 rad/s, where the slave's reflected force has grown, so $|L|$ peaks near 36 rad/s at 1.55 times its static value and the limit levels at 0.41 rather than at the static 0.63, yet the running model keeps still up to 0.55, where the oscillation grows by only 1.69 a second, and is

Table 8.4 The Nine Cases and the Challenge

#	Change from the reference line	Teaching point
1	none, three presses with no delay	the wall is felt
2	$T = 100\,\text{ms}$	the contact chatters
3	$T = 500\,\text{ms}$	no gain is still and transparent
4	$T = 500\,\text{ms}$, wave variables	passive, and mud
5	$T = 500\,\text{ms}$, passivity controller	a damper that waits
6	as 4, $j = 20\,\%$, $p = 10\,\%$, burst at 4.9 s	a held wave keeps pressing
7	$T = 1\,\text{s}$, model mediation, $\Delta = -8\,\text{mm}$	right until it is not
8	peg, $T = 250\,\text{ms}$	a sideways contact
9	buttons, $T = 2.6\,\text{s}$, $k_f = 0$, move and wait	safe and slow
C	buttons then peg, $T = 400\,\text{ms}$, $p = 5\,\%$	choose an architecture

Table 8.5 Measured Results of Each Case and Variant

Case and configuration	Result
1, as shipped	3.06 N at the wall for 3.06 N felt, 1243 N/m felt, ledger at most 0.78 mJ
1, contour	rms 0.15 mm, in contact throughout
2, as shipped	41 reversals, 25.7 N, 15.5 J, flagged at 2.76 s
2, k_f 0.35 / 0.30	oscillates / still, 10.8 N, 110 mJ
2, position–position	still, 3.8 N, 6.45 N of drag
3, as shipped	master thrown back 94.6 mm, 9.54 J
3, k_f 0.20 / 0.10	still, 23.7 / 27.8 N for 3.75 / 2.50 N felt
3, $K_h = 900\,\text{N/m}$	still to 0.55, flagged at 0.60
4, b = 20, 50, 200	walls of 76, 173 and 1758 N/m, drag 2.28, 3.53 and 6.24 N
5, as shipped	dampers on 14.5 % and 17.2 % of 40 s, 2.95 mJ created
5, $\alpha_{max} = 50\,\text{N s/m}$	274 mJ created, 1.7 s of oscillation
6, held / rebuilt waves	7.16 mJ created, peg 30.4 mm ahead / never positive
7, as shipped	first press 13.78 N for 4.79 N felt, third 3.15 for 3.05 N
7, $\Delta = 0$ / $+4$ mm	every press 3.14 N / the wall never reached
8, as shipped	master across 62 to 187 mm, 93 N on the slot walls, 225 J
8, passivity controller	seated at 3.745 s, 1.1 N on the walls
9, move and wait	272.9 s, 21.8 N
9, model mediation, continuous	17.3 s, 4.36 N
C, as shipped	41 reversals, 108.8 N at the pause
C, model mediation	12 reversals, 4.08 N, done at 18.7 s

flagged at 0.60. At long delays the reverse happens. With full reflection at 1 s the instability is a slow bounce, of period 2.12 s in the linearized loop. On the running model the wall is struck six times, every 2.3 to 2.8 s, with too few reversals in any one press to fill an episode, so the flag never rises while the strikes reach 40.6 N and the channel creates 9.62 J. The ledger, not the reversal count, is the signal that does not miss.

A reduced gain does not make the channel passive either. With no delay, a gain of 0.5 lets the slave do work on the wall of which the hand pays only half, and the ledger books 12.2 mJ over three presses against 0.78 mJ at a gain of one, while a gain of 1.5 returns more on the release than it took on the press and books 7.7 mJ. A passive channel is sufficient for a still contact but not necessary, because the hand, the devices and the wall dissipate what the channel creates.

🕮 **Pause.** Try Exercise 8.2 before reading on.

8.6 Nine Cases and a Challenge

The cases of Table 8.4 are variations on the settings of Tables 8.2 and 8.3, each changing only what its row names. Every figure in this section was measured headlessly, the engine running without its interface at the default seed, in simulated time, on the scenarios of Section 1.6.4, and Table 8.5 collects them.

8.6.1 No Delay, Then a Tenth of a Second

In case 1 the scripted operator presses the wall three times with full reflection and no delay, until the hand feels 3 N. Each press holds 3.06 N at the wall for 3.06 N on the handle, and the master stops 2.06 mm after the felt force first passes 0.5 N, a felt wall of 1243 N/m against the 1429 N/m of the series springs. The presses are done at 13.4 s with a peak of 3.15 N, the ledger never exceeds 0.78 mJ, and the contour, traced at no delay, keeps an rms gap of 0.15 mm in contact the whole way.

Case 2 adds 50 ms each way. The peg first touches at 2.12 s and pushes to 5.99 N, the force reaches the hand 50 ms after it rose, when the hand has already pushed on, and the delayed force throws the master back. The peg leaves the wall at 2.28 s, when the channel has created 34.1 mJ, and strikes again at 2.43 and 2.76 s, with 42.1 mJ and 593 mJ created, the second of which the detector first flags. The ledger turned positive two touches before the status did. Over the run the channel creates 15.5 J, with 41 reversals, a peak of 25.7 N and 6.5 s of oscillation, all of it dissipated by the hand, the devices and the wall, and Figure 8.3 shows the contact during the third press. A gain of 0.35 still oscillates, while 0.30 keeps still with the peg pressing up to 10.8 N, less than a third of it felt, and the channel still creating 110 mJ. Position–position does not chatter (3.8 N at most) but drags the handle toward where the slave was, up to 6.45 N on the way to the wall against 1.19 N under position–force, and the contour traced at 100 ms under position–force bounces along the bump in contact only 40.5 % of the way.

> **TRY IT IN THE APP** **2 · 100 ms of Round Trip**
>
> Load the case and press Run on the Channel view. Watch the energy ledger as the peg first touches the wall at 2.12 s: by the end of that touch the channel has created 34.1 mJ. The status line first reads Contact oscillation during the third touch, at 2.76 s. The ledger turned positive two touches before the status did.

8.6.2 Half a Second With No Cure

Case 3 is the honest infeasible case of position–force. At full reflection the first touch, at 2.62 s, throws the master back 94.6 mm, from 160.1 to 65.5 mm, and the run creates 9.54 J with 2.97 s of oscillation. Every gain tested from 0.25 to 1 oscillates. At 0.20, the slider's next step down, the contact keeps still, but holding 3.75 N on the handle drives the peg into the wall at 23.7 N, a ratio of 6.3, and at 0.10, 2.50 N on the handle is 27.8 N at the wall. A firmer grip moves the limit, to 0.55 at 900 N/m on the running model, although not to its static bound, for the reason of Section 8.5. No choice of gain is both still and transparent at this delay.

8.6.3 Three Passive Cures at Half a Second

Case 4 sends waves with $b = 50\,\mathrm{N\,s/m}$. The ledger never turns positive and holds up to 37.9 mJ in flight, and the contact keeps still, but the master travels 14.4 mm while the felt force rises from 0.5 to 2.99 N, a wall of 173 N/m. The impedance sets the trade. At $b = 20$ the handle drags at most 2.28 N in free space and the wall is 76 N/m. At $b = 200$ free space drags 6.24 N and the wall is 1758 N/m, firmer than with no delay, but the contact bounces for 0.84 s as the waves reflect between the ends. With the task changed to the contour, the trace at 500 ms keeps contact 79.1 % of the way with an rms gap of 1.12 mm.

> **TRY IT IN THE APP** **4 · Wave Variables at 500 ms**
>
> Press Run and watch the ledger, which never turns positive. On the Cell view the master travels 14.4 mm into the wall while the felt force rises from 0.5 to 2.99 N. Press Reset and run with the Wave impedance at 20, then at 200: the rendered wall is 76 N/m, then 1758 N/m, and at 200 the handle drags 6.24 N in free space.

Case 5 returns to position–force at 500 ms with the passivity controller. Over the case's 40 s the master's damper acts during 14.5 % of the time and the slave's during 17.2 %. The slave's acts only

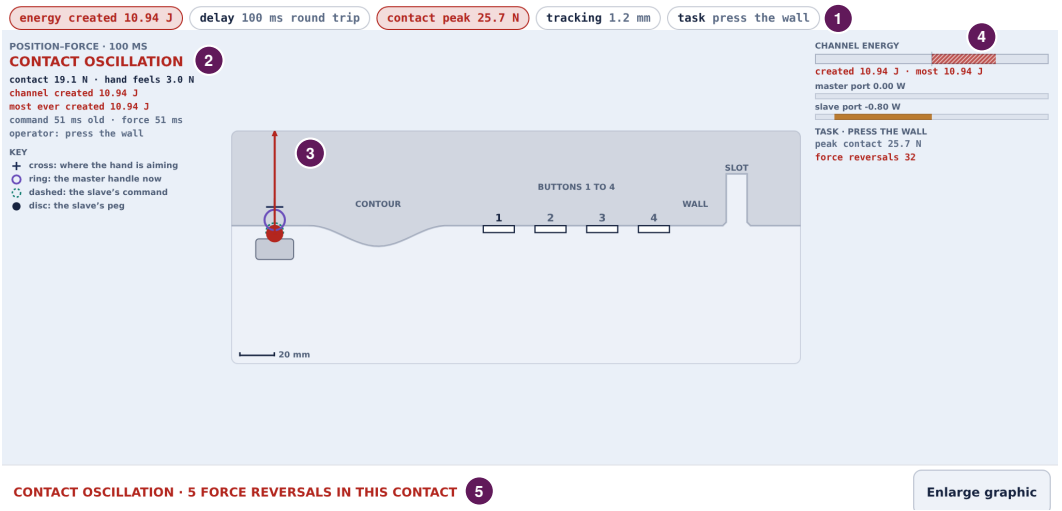

Figure 8.3 Case 2, 100 ms of round trip, at $t = 11.00\,\mathrm{s}$ (the clock is cut) in the Cell view with no control changed, cropped to the plant panel. The peg strikes the wall during the third press with 19.1 N, drawn as the arrow at the peg, while the force reaching the master left the slave 51 ms earlier, before this strike. (1) The badge row: energy created, the round trip and the contact peak. (2) The left column gives the contact force beside the 3.0 N the operator perceives over 50 ms, 10.94 J created by the channel, and newest packets 51 ms old. (3) The peg, the master's ring and the arrow of the contact force at the wall. (4) The channel-energy and task panels count 32 reversals, a peak of 25.7 N and 10.94 J created at most. (5) The status reads Contact oscillation, 5 force reversals in this contact.

while the peg touches the wall, which it does for 29.0 % of the time, and the master's only while a contact force is reaching the handle, up to the 250 ms of the return link after a touch ends, so neither acts in free space. The contact keeps still with a peak of 6.84 N, and the channel creates at most 2.95 mJ against 9.54 J without the controller, while the dampers take 634 mJ. The price is position and feel. The slave's damper lets its command fall up to 18.3 mm behind the master, and the hand travels 10.9 mm into a wall of about 236 N/m, because the slave may not do work the operator has not yet paid for and the payment arrives a round trip late. Limited to $\alpha_{max} = 50\,\mathrm{N\,s/m}$, the master's damper runs out, the channel creates 274 mJ and the contact oscillates for 1.7 s.

Case 6 runs wave variables on a channel whose delay wanders by 20 % and which loses one packet in ten, and both links drop out for 1 s at 4.9 s as the hand backs off from the first press. Holding the last wave, the slave keeps pressing the wall at 2.68 to 2.77 N through the whole second while the master backs off 42.5 mm, and at 13.5 s it still presses at 2.59 N, 30.4 mm ahead of the master, having been off the wall for 0.14 s in all since the dropout. The ledger turns positive, 7.16 mJ, and the second press raises the oscillation flag for 1.39 s. With the dropout removed from the script, a counterfactual run, the jitter and loss alone create 1.15 mJ, and with the jitter and loss set to zero the dropout alone creates 1.71 mJ. With reconstruction by Equation (8.11) the slave lets go at 5.315 s, before the link returns, and the ledger never turns positive. Both lose position, because waves carry velocity and lost velocity is never made up, and with reconstruction the peg is still 11.9 mm ahead of the master at 10 s. These figures belong to the default seed, and the losses move them. On five other seeds the hand is still holding the first press when the link drops on three, whose slaves let go 0.3 to 0.4 s after it returns, and the held waves peak at 0 to 14.1 mJ created, against 0 to 8.1 mJ from the jitter and loss alone, so the split between the two causes lies within the spread. On every seed reconstruction lets go between 5.26 and 5.34 s and never lets the ledger turn positive.

8.6.4 A Model of the Wall at a Full Second

Case 7 has a round trip of 1 s and the real wall 8 mm closer than the drawing the model starts from. On the first press the peg meets the wall early, while the hand still feels nothing, and presses it at 13.78 N while the hand feels at most 4.79 N, and the ledger books 84.4 mJ that the master never put in. The slave's estimate is within 0.5 mm of the truth 28 ms after the touch, and the master's model gets there 1.33 s after it, half a second for the news and the rest for the rate limit. The second press peaks at 4.80 N and the third at 3.15 N for 3.05 N on the handle. With the wall where the drawing says, every press is 3.14 N, with it 4 mm short the presses are 8.86, 4.08 and 3.15 N, and with it 4 mm beyond, the hand presses the model's wall and the peg never reaches the real one. A prior of half the true stiffness costs only the first press, at 4.00 N, and a wall three times stiffer than the prior is estimated at 15 057 N/m.

Pause. Try Exercise 8.3 before reading on.

8.6.5 The Slot, the Moon and the Challenge

Case 8 lines the peg up 3.5 mm off the slot's center and pushes it in with full reflection through 250 ms. The chamfer turns the push sideways, the slot walls push back, and the sideways axis goes unstable as the wall did in case 2. The master swings between 62 and 187 mm across a slot at 128 mm, the peg hammers the slot walls at up to 93 N, the channel creates 225 J, and the peg seats anyway at 3.53 s, by force. A 2 N lateral limit only bounds the swing, to 118 to 138 mm with 28 N on the walls. With the passivity controller the peg seats at 3.745 s with 1.1 N on the walls, as it does with wave variables (1.5 N) and with model mediation at 4.46 s (1.1 N). A disc with no rotation could stick only if pushed sideways more than $1/\mu = 2.5$ times as hard as along the slot, so the jam of a real peg does not occur, and the soft hand fed by the chamfer centers the peg instead. A peg that can stick in a chamfered hole under force control is the subject of Volume 20, Chapter 13.

Case 9 presses the four buttons through the 2.6 s of a lunar link with no force feedback, by move and wait. A 10 mm step in free space takes 3.15 s, a 2 mm step near contact 2.86 to 4.47 s, the longest when the slave presses and the video cannot show it arriving, and the 83 steps take 272.9 s with a peak of 21.8 N. Moving continuously with no force feedback takes 15.0 s but presses blind, at up to 24.0 N. Model mediation, moving continuously, takes 17.3 s at 4.36 N, because the hand feels the local model's buttons click at once, and with move and wait it takes 198.8 s at 2.25 N. Wave variables make five presses out of order and never finish, and the passivity controller never presses a button, because the operator's payment arrives 2.6 s late.

> **DID YOU KNOW? Slow and Careful Still Presses Hard**
>
> From the simulator: through the 2.6 s lunar link with no force feedback, move and wait presses the four buttons in 272.9 s with a peak of 21.8 N, while model mediation, moving continuously, takes 17.3 s at 4.36 N (Table 8.5). Only the local model lets the hand feel each click at once.

The challenge asks for the four buttons in order and the peg seated through 400 ms with one packet in twenty lost, with at most 16 force reversals, a peak under 10 N and the task done within 25 s, while the same task with no delay makes 8 reversals. A run counts only if its round trip stays at or above 400 ms and its loss at or above 5 % at every physics step, so that clean run, which meets the other limits, earns nothing. Table 8.6 compares the configurations as the scorecard shows them when the run pauses, 2 s after the task. As shipped, position–force bottoms the first button and presses on to 13.1 N while the hand feels 2.0 N, then throws the handle 43.2 mm back. When the task is done at 16.3 s it has made 30 reversals at up to 73.3 N, and it keeps hammering, so the pause shows 41 reversals and 108.8 N (41 or 42 and 108.8 to 109.9 N on five other seeds), 2.6 and 10.9 times the limits. Wave variables pass one reversal from the limit, and the passivity controller passes with no margin, because at 450 ms, or with the slave's stiffness 20 % lower, it stalls at a button. Model mediation passes inside the limits by 4 reversals, 5.9 N and 6.3 s, and with margin at 450 and 600 ms,

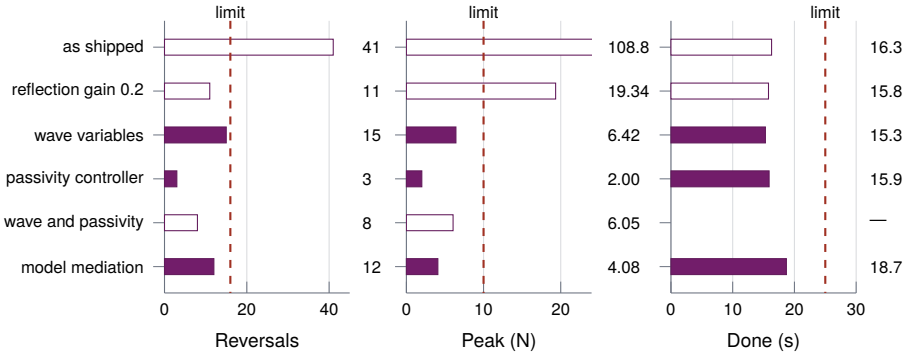

Figure 8.4 The challenge of Table 8.6 as a chart: force reversals against the limit of 16 (left), peak contact force against 10 N (middle) and the time to finish against 25 s (right), as the scorecard shows them when the run pauses. Dark bars belong to configurations that met the target, white bars to those that missed it. The as-shipped peak runs off its axis, and wave variables with the passivity controller never finish. Notice that every configuration that finishes does so within the 25 s, so the reversals and the peak decide the target, and that wave variables pass one reversal from the limit.

Table 8.6 The Challenge as Scored When the Run Pauses

Configuration	Reversals	Peak (N)	Done (s)	Target
as shipped	41	108.8	16.3	missed
reflection gain 0.2	11	19.34	15.8	missed
wave variables	15	6.42	15.3	met
passivity controller	3	2.00	15.9	met
wave and passivity	8	6.05	—	missed
model mediation	12	4.08	18.7	met

at 10 % and 20 % loss, with the slave's stiffness 20 % either way, with 20 % jitter, a 5 Hz update and half the true stiffness as its prior, on three seeds each. It also stays within the limits at 350 ms and with the site 2 mm either way, but the application credits neither run, because the round trip is below its minimum or a disturbance was changed. Figure 8.4 draws the same table against the three limits.

8.7 Technology in Practice

What Measures It

Each side of the teleoperator measures its own position and velocity, and the slave carries a force sensor at its tool, whose reading is what force reflection sends back. The operator also watches the slave on a video that arrives one delay late.

> **TECHNOLOGY** **Sensor: A Strain-Gauge Force Sensor at the Tool**
>
> The force a slave reflects comes from a sensor between its drive and its tool. In a load cell, strain gauges bonded to a metal structure of precisely known elasticity stretch and compress with it, as with a very stiff spring, and a bridge circuit turns their change of resistance into a voltage. Temperature also changes a gauge's resistance, so the bridge cancels it with a dummy gauge that feels the temperature and no strain [20]. Six-axis sensors for force feedback measure three forces and three torques in one compliant body. Strain-gauge designs offer high accuracy but are sensitive to electromagnetic interference and often show cross-axis coupling, in which a load in one direction induces a response in others [21]. The chapter's sensor reads the contact force with no noise, lag or quantization.

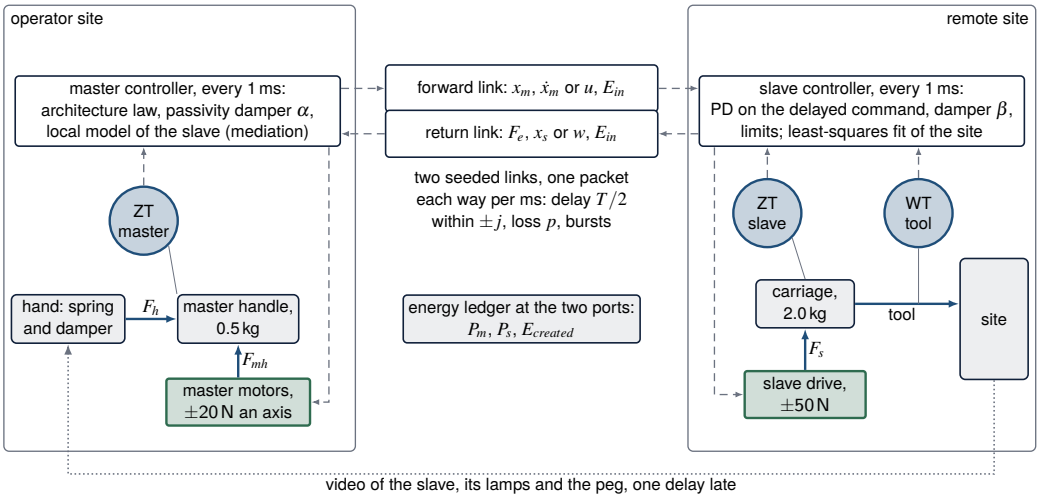

Figure 8.5 The teleoperator of Section 8.3 on its devices. Each side measures its own position and velocity (ZT), the slave reads the contact force at its tool (WT), and the two controllers act once a millisecond on what the seeded links deliver. The ledger books the power at the two ports from the same forces and velocities the controllers exchange, and the dotted path is the video the operator watches, one delay late. Notice that every path from the tool back to the hand passes through the return link or the video, so the hand feels the site only a delay late unless the master's local model renders it.

In the chapter every device measures exactly and the force sensor has no lag, so every delay in the loop belongs to the channel and the holds, as Section 8.3.7 states.

What Moves It

The final elements are the master's motors, limited to 20 N an axis, and the slave's drive, limited to 50 N. The 1989 analysis of force reflection modeled a position-controlled space manipulator as a second-order system, an assumption valid when the arm has no flexible mode in the range of interest, generally below 10 Hz, and took one of 150 lb under a position controller with a bandwidth of 1.5 Hz [18]. The chapter's slave is that kind of model, a 2.0 kg carriage under a PD. On the master side, the same analysis found that the relation between the reflected force and the handle's position is set by the operator and the hand controller, whose characteristics are nonlinear and time varying [18], the part the chapter gives to a linear hand of 300 N/m. Figure 8.5 draws the devices, the two controllers and the links between them.

Where the Controller Runs

Each side runs its own real-time controller, joined to the other only by the network. In the application both controllers act once a millisecond, the physics advances in steps of the same length, and the ledger books each period's power. The layer that holds the chapter's lesson sits between the controllers and the links: the ledger and the passivity dampers act on the forces and velocities the two sides exchange, and under model mediation the master adds a local copy of the slave, stepped four times a period. In a model-mediated teleoperator the slave must also estimate the environment's parameters, which takes time when the environment is complex [22]. The slave of this chapter fits only an offset and a stiffness.

TECHNOLOGY **Platform: The Packet Link Between the Two Sides**

In a conventional bilateral teleoperator, position, velocity and force on both sides are sampled and sent in packets, typically at 1 kHz, for stability and transparency. With the header each packet carries, that rate uses a packet-switched network inefficiently, so perceptual deadband coding sends a sample only when it differs from the last one sent by more than a just noticeable difference, and the receiver fills the gaps, for example by holding the last value [22]. The delay of such a link comes from the signal's time of flight and from queuing and processing at each node, and for a satellite in Earth orbit the node delays dominate [23]. The chapter's links send one packet each way every millisecond and hold the last one through a gap.

What Is Modernizing It

Three developments bear on the delayed channel. The first is model-based, semi-autonomous teleoperation. In studies on satellite servicing, the largest gain came from better situation awareness, through a three-dimensional model of the scene overlaid with the live, delayed video, and from letting operators specify their intended motion precisely, while model-mediated teleoperation reduced the task load without significantly improving task performance [24]. The second is the model of the site itself: for complex geometry, external three-dimensional sensors can capture a precise description of the surface for the master's model [22], where the chapter's model knows only a drawing shifted by one offset. The third is the force sensor. Optical sensors such as fiber Bragg gratings are immune to electromagnetic interference, and capacitive force and torque sensors offer high sensitivity, a simple structure and low power for small tools [21].

IN PRACTICE **Working Through Seconds of Delay**

Ground operators of space robots meet telemetry delays by moving and waiting, and sparse or poorly placed cameras may be an even bigger concern than the delay. In studies on satellite servicing tasks, NASA robot operators kept their keyboard interface, which lets them specify a precise intended motion, over a three-dimensional console, and commanding motion in the frame of the task could improve their performance. The missions put a premium on avoiding failure rather than on time, so the study advises offering several interfaces to choose from at each step of a task [24]. Operators in lunar rover studies estimated that 20 to 100 hours of training would make an operator about 60 % proficient, and full proficiency a yearlong program [4].

8.8 Limitations

The operator is the coarsest part of the model, a linear impedance on a script, where a real operator changes grip and strategy as the delay grows. The devices are point masses in two independent axes, the peg is a disc that cannot wedge, and master and slave move at one scale in one frame. Packets are never corrupted, duplicated or queued behind other traffic, and both ends keep one clock. Model mediation learns a single offset of the whole site along the wall's normal and a single stiffness, so a site that tilts, moves or departs from its drawing in shape is beyond it, and wave variables leave uncorrected the drift that lost waves cause. The four-channel architecture is absent.

OPEN CHALLENGE **Keeping the Local Model True**

A review of model-mediated teleoperation names modeling the environment as its foremost challenge. When the site changes or the estimate is poor, the master's model must be updated quickly, but a sudden change of stiffness or position shows the operator a sudden force, the model-jump effect, which the arm may not stabilize. Fast updates and stability therefore conflict. The approach can be both stable and transparent for relatively simple environments, but existing online estimation cannot make it work efficiently in complex or unknown ones [22]. This chapter's model learns one offset and one stiffness and moves no faster than 10 mm/s.

The oscillation flag counts reversals for teaching. Near the gain limit, and at long delays where the instability is a slow bounce, it can disagree with the linearized loop, and the ledger is then the better witness. All limits, human parameters and challenge thresholds are teaching values, and no figure here describes the performance of a real teleoperator.

OPEN CHALLENGE **Feeling Contact Through a Long Delay**

A predictive display helps an operator position a remote robot through a delay, but it is not feasible once the robot must touch its surroundings, because current models cannot accurately predict what the robot will do next. Bilateral force reflection has been shown impractical and unstable under small to medium delays. Wave variables restore stability, yet delayed force feedback has not proven intuitive to operators, and many researchers have turned to semi-autonomous teleoperation, supervisory or model-based, instead [24]. Progress would be contact feedback, through seconds of delay, that operators find natural and that stays passive on a ledger like this chapter's.

8.9 Conclusion

This chapter has shown a force-reflecting teleoperator, every part of it passive, creating energy in its channel, and has measured that energy at the two ports. With no delay the ledger stays within 0.78 mJ over three presses, while 100 ms of round trip creates 34.1 mJ by the end of the first touch, two touches before the oscillation flag, and 15.5 J over the run. Linearized in contact as it is sampled, the loop tolerates full reflection only below about 10 ms of round trip, and its limit falls toward $K_h/K_{ser} = 0.21$, so that at 500 ms the gain that keeps a press still leaves the peg pressing 6.3 times harder than the hand feels. Wave variables keep the ledger from ever turning positive and render a 173 N/m wall, the passivity controller holds the channel to 2.95 mJ at the cost of a command 18.3 mm behind the master, and model mediation, once its model is right, presses at 3.15 N for 3.05 N felt after a first press of 13.78 N.

Two independent axes, a scripted linear hand and a channel that never corrupts a packet bound what the model can say. A fuller model would add a scaled and rotating master, the four-channel architecture and an operator who adapts to the delay. The simulator is an accessible plant on which the energy books of a delayed interface can be read while it runs.

KEY POINTS

- A teleoperator whose every part is passive can still create energy in its delayed channel, and the ledger shows it two touches before the contact is flagged.
- Through a long delay a wall can be rendered no stiffer than the hand that holds it, and full reflection needs a round trip under about 10 ms.
- A still contact does not prove a passive channel: at 100 ms and a gain of 0.30 the channel creates 110 mJ.
- Each passive cure has a price: a 173 N/m wall with wave variables, a command 18.3 mm behind the master with the passivity controller, and a first press of 13.78 N under model mediation.

Where else this applies. Volume 7, Chapter 14 carries the same delayed, lossy channel, and Volume 2, Chapter 7 regulates a loop whose dead time is a network round trip.

References

[1] B. Hannaford, "A design framework for teleoperators with kinesthetic feedback," *IEEE Transactions on Robotics and Automation*, vol. 5, no. 4, pp. 426–434, 1989.

[2] J. Marescaux, J. Leroy, M. Gagner, F. Rubino, D. Mutter, M. Vix, S. E. Butner, and M. K. Smith, "Transatlantic robot-assisted telesurgery," *Nature*, vol. 413, no. 6854, pp. 379–380, 2001.

[3] T. Imaida, Y. Yokokohji, T. Doi, M. Oda, and T. Yoshikawa, "Ground–space bilateral teleoperation of ETS-VII robot arm by direct bilateral coupling under 7-s time delay condition," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 3, pp. 499–511, 2004.

[4] H. L. Litaker, K. H. Beaton, Z. Q. Li, and J. F. Lewis, "Lunar terrain vehicle (LTV) remote teleoperation studies under four lunar communication latencies," in *2025 IEEE Aerospace Conference*, 2025, pp. 1–21.

[5] R. J. Anderson and M. W. Spong, "Bilateral control of teleoperators with time delay," *IEEE Transactions on Automatic Control*, vol. 34, no. 5, pp. 494–501, 1989.

[6] B. Hannaford and J.-H. Ryu, "Time-domain passivity control of haptic interfaces," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 1, pp. 1–10, 2002.

[7] G. Niemeyer and J.-J. E. Slotine, "Telemanipulation with time delays," *The International Journal of Robotics Research*, vol. 23, no. 9, pp. 873–890, 2004.

[8] J. P. Hespanha, P. Naghshtabrizi, and Y. Xu, "A survey of recent results in networked control systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 138–162, 2007.

[9] W. R. Ferrell, "Delayed force feedback," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 8, no. 5, pp. 449–455, 1966.

[10] J.-H. Ryu, D.-S. Kwon, and B. Hannaford, "Stable teleoperation with time-domain passivity control," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 365–373, 2004.

[11] P. Mitra and G. Niemeyer, "Model-mediated telemanipulation," *The International Journal of Robotics Research*, vol. 27, no. 2, pp. 253–262, 2008.

[12] R. C. Goertz, "Master-slave manipulator," Argonne National Laboratory, Argonne, IL, Report ANL-4311, Mar. 1949.

[13] T. B. Sheridan and W. R. Ferrell, "Remote manipulative control with transmission delay," *IEEE Transactions on Human Factors in Electronics*, vol. HFE-4, no. 1, pp. 25–29, 1963.

[14] W. R. Ferrell, "Remote manipulation with transmission delay," *IEEE Transactions on Human Factors in Electronics*, vol. HFE-6, no. 1, pp. 24–32, 1965.

[15] T. B. Sheridan, "Space teleoperation through time delay: review and prognosis," *IEEE Transactions on Robotics and Automation*, vol. 9, no. 5, pp. 592–606, 1993.

[16] G. Niemeyer and J.-J. E. Slotine, "Stable adaptive teleoperation," *IEEE Journal of Oceanic Engineering*, vol. 16, no. 1, pp. 152–162, 1991.

[17] D. A. Lawrence, "Stability and transparency in bilateral teleoperation," *IEEE Transactions on Robotics and Automation*, vol. 9, no. 5, pp. 624–637, 1993.

[18] J. D. Chapel, "Performance limitations of bilateral force reflection imposed by operator dynamic characteristics," in *Proceedings of the NASA Conference on Space Telerobotics*, vol. 4. Pasadena, CA: Jet Propulsion Laboratory, California Institute of Technology, 1989, pp. 91–100. [Online]. Available: https://ntrs.nasa.gov/citations/19900020524

[19] J.-H. Ryu, J. Artigas, and C. Preusche, "A passive bilateral control scheme for a teleoperator with time-varying communication delay," *Mechatronics*, vol. 20, no. 7, pp. 812–823, 2010.

[20] 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.

[21] S. Y. Lee, J. Y. Sim, D.-Y. Seok, Y. B. Kim, J. Shim, U. Kim, and H. R. Choi, "A redundant-sensing-based six-axis force/torque sensor enabling compactness and high sensitivity," *Sensors*, vol. 26, no. 3, p. 871, 2026.

[22] X. Xu, B. Cizmeci, C. Schuwerk, and E. Steinbach, "Model-mediated teleoperation: Toward stable and transparent teleoperation systems," *IEEE Access*, vol. 4, pp. 425–449, 2016.

[23] J. S. McLaughlin and B. D. Staunton, "Performance constraints and compensation for teleoperation with delay," in *Proceedings of the NASA Conference on Space Telerobotics*, vol. 5. Pasadena, CA: Jet Propulsion Laboratory, California Institute of Technology, 1989, pp. 63–72. [Online]. Available: https://ntrs.nasa.gov/citations/19900020565

[24] P. Kazanzides, B. P. Vagvolgyi, W. Pryor, A. Deguet, S. Leonard, and L. L. Whitcomb, "Teleoperation and visualization interfaces for remote intervention in space," *Frontiers in Robotics and AI*, vol. 8, p. 747917, 2021.

8.10 Exercises

8.1 Concept check · 2 min

In case 2 the hand, the master, the slave and the wall can each only store or dissipate energy. Yet by the end of the first touch, at 2.28 s, the ledger shows 34.1 mJ created. Where does that energy come from?

(a) From the channel, which pairs the slave's delayed force with a master velocity it never met

(b) From the slave's PD, which stores energy in its stiffness K_p whenever the command runs ahead

(c) From the wall's damping, which fades in over the first 0.5 mm and so returns energy on release

(d) From the hand's spring, which stores energy on each press and returns it when the hand backs off

8.2 Concept check · 1 min

In case 2, at a reflection gain of 0.30, the contact keeps still and no oscillation is flagged, but the ledger shows 110 mJ created. What does this show?

(a) The ledger is in error whenever the oscillation flag stays down, since a still contact proves the channel passive

(b) A passive channel is sufficient for a still contact but not necessary, because the hand, the devices and the wall dissipate what it creates

(c) A gain of 0.30 lies above the linearized limit at 100 ms, so the contact sits on the edge of oscillation

(d) The 110 mJ is the one-period lag of the sampled loop, which the ledger books at every gain

8.3 Concept check · 2 min

Case 7 runs model mediation through a round trip of 1 s with the real wall 8 mm closer than the drawing. On the first press the peg pushes the wall with 13.78 N while the hand feels at most 4.79 N, and on the third press 3.15 N for 3.05 N. Why is the first press so hard?

(a) The least-squares estimator needs several presses to converge, so its offset is wrong until the third press

(b) The master renders the slave's measured force a second late, as position–force does, so the hand learns of the touch late

(c) The master renders its local model of the wall, which stands 8 mm too far, so the peg meets the real wall before the hand feels anything

(d) The 10 mm/s rate limit on the master's model drives the hand's aim into the wall during the first press

8.4 Theory and calculation · 7 min

Use Equations (8.14) and (8.15) with the defaults of Tables 8.2 and 8.3.

(a) Compute K_{ser}, the static bound K_h/K_{ser} for the 300 N/m hand, and the stiffness $k_f K_{ser}$ that the hand feels at a gain of 0.20. *Checkpoint:* 1429 N/m, 0.21 and about 286 N/m, just softer than the hand.

(b) For a grip of 900 N/m the static bound is 0.63, but $|L(j\omega)|$ peaks near 36 rad/s at 1.55 times its static value. Estimate the limit through a long delay. *Checkpoint:* About $0.63/1.55 = 0.41$.

(c) At 500 ms and full reflection the dominant root grows by a factor of 14.7 a second with a period of 1.11 s. At a gain of 0.248 it grows by 1.20 a second. Find the growth over one period in the first case and over one 1.5 s hold of a press in the second. *Checkpoint:* About 20 per period, and about 1.31 per hold.

8.5 Theory and calculation · 7 min

For motions slow compared with the round trip, a wave channel whose slave end is held still behaves at the master as a spring of stiffness $2b/T$, the wave impedance over the one-way delay. Behind it the slave's PD and the wall add $K_{ser} = 1429\,\mathrm{N/m}$ in series.

(a) At a round trip of 500 ms, compute the channel's stiffness and the felt wall for $b = 20$ and $b = 50\,\mathrm{N\,s/m}$, and compare them with the 76 and 173 N/m measured in case 4. *Checkpoint:* 80 N/m and 75.8 N/m, 200 N/m and 175 N/m.

(b) With $b = 50\,\mathrm{N\,s/m}$, what wall does the same model predict through the 2.6 s of a lunar link? *Checkpoint:* About 37 N/m.

(c) Which wave impedance would the model need to render a wall as stiff as the 300 N/m hand at 500 ms, and what does a larger b cost in free space? *Checkpoint:* About 95 N s/m. The handle drags more in free space.

8.6 Theory and calculation · 6 min

In case 9 the operator moves and waits through a round trip of 2.6 s. Section 8.3.4 gives the rules: an increment of 10 mm in free space or 2 mm near contact, moved at 30 mm/s, then a hold until the video shows the slave within 1.5 mm of the aim, or until a round trip and a half and 0.3 s have passed, then a reaction time of 0.2 s.

(a) Estimate the time of one 10 mm step in free space, taking the video to show the slave's arrival one round trip after the aim stops. *Checkpoint:* About 3.13 s, against 3.15 s measured.

(b) Compute the longest time a 2 mm step can take, when the slave presses and the video cannot show it arriving. *Checkpoint:* 4.47 s, as measured.

(c) The 83 steps take 272.9 s. Find the mean time per step, and the share of a free-space step spent waiting for the round trip. *Checkpoint:* About 3.29 s a step, and about 83 %.

8.7 App challenge · 8 min

Load scenario 2, 100 ms of round trip, with the default seed, on the Channel view. Run at 1× until the run pauses after the third press, about 16 s on the clock, and read the Force reversals, Peak contact force and Energy created tiles of the scorecard.

(a) Report the three tiles, and when the status line first read Contact oscillation compared with when the ledger first turned positive. *Checkpoint:* 41 reversals, 25.7 N, 15.52 J at most, each within one unit of the last digit. The ledger turns positive at the first touch, two touches before the flag.

(b) Press Reset, set the Force reflection gain to 0.30 before pressing Run, and run until the run pauses. Report the same tiles, and whether the oscillation was flagged. *Checkpoint:* 9 reversals, 10.8 N, 110.3 mJ at most, within 1 mJ. No oscillation flagged.

(c) Is the channel passive at 0.30? Explain with the linearized limit of Section 8.5. *Checkpoint:* No. The loop is stable, since 0.30 is below the limit of 0.329, but the channel still creates energy.

8.8 App challenge · 8 min

Load scenario 5, A passivity controller at 500 ms, with the default seed. Run at 1× until the run pauses after the third press, about 21 s on the clock, and read the Peak contact force, Force reversals, Energy created and Damper dissipated tiles.

(a) Report the four tiles, and compare the energy created with the 9.54 J of case 3, the same delay without the controller. *Checkpoint:* 6.8 N, 8 reversals, about 3 mJ (the tile shows 2.9 mJ), 634.1 mJ, each within one unit of the last digit.

(b) Press Reset, set the Master damper limit to 50 N s/m before pressing Run, and run until the run pauses. Report the same tiles. *Checkpoint:* 7.8 N, 11 reversals, 274.3 mJ, 596.9 mJ, each within one unit of the last digit.

(c) Explain the difference between the two runs from Equation (8.12). *Checkpoint:* The clipped damper cannot take all of the excess, and the rest stays on the ledger.

8.9 App challenge · 9 min

Load the challenge, Four buttons and the peg at 400 ms, with the default seed. Each run ends when it pauses 2 s after the task is done, about 18 s on the clock at 1×. Read the Task, Peak contact force, Force reversals and Challenge target tiles. Reload the scenario before each new configuration and make the change before pressing Run.

(a) Run as loaded, with position–force at full reflection. *Checkpoint:* Done at 16.3 s, 108.8 N, 41 reversals, target not met.

(b) Set the Architecture to Wave variables and run. *Checkpoint:* Done at 15.3 s, 6.4 N, 15 reversals, target met.

(c) Set the Architecture to Model-mediated and run. Give the margins to the limits of 16 reversals, 10 N and 25 s, and say which of the two passing architectures you would put on a link whose delay may grow. *Checkpoint:* Done at 18.7 s, 4.1 N, 12 reversals. Margins 4 reversals, 5.9 N and 6.3 s.

8.10 Application area · 4 min

A surgical master is to move its slave at one fifth of the hand's velocity and to render the contact force five times larger, so that the surgeon feels fine tissue. Use the port powers of Equation (8.8) with no delay.

(a) Find the ratio of the power the channel hands the master to the power the slave takes from the site, and say whether such a channel is a wire. *Checkpoint:* 25. It is not a wire, and the ledger would book energy at every press and release.

(b) The slave moves $0.2\,x_m$ against K_{ser}, and the hand feels five times the force. Find the static loop gain $L(0)$ in terms of K_{ser}/K_h, and say what it means through a long delay. *Checkpoint:* $L(0) = 5 \times 0.2\,K_{ser}/K_h = K_{ser}/K_h$, the same as full reflection, so the bound of Equation (8.15) applies to the product of the two scales.

8.11 Application area · 3 min

A position–force teleoperator for a hot cell is commissioned with one operator whose grip matches the chapter's 300 N/m hand. Its gain is set to 0.20 for a link of about 500 ms.

(a) A second operator holds the handle loosely, at about 150 N/m. Compute the static bound for this grip. *Checkpoint:* About 0.105, so 0.20 is above it.

(b) Give two practices that keep the system safe across operators. *Checkpoint:* Commission against the softest grip expected, and watch the ledger rather than the reversal count.

8.12 Review the agent · 5 min

An AI assistant is asked how to protect a position–force teleoperator that will run on links from 100 to 1000 ms. It replies: "Use the status line as an interlock. When it reads Contact oscillation, set the reflection gain to zero. I checked it on case 2 at 100 ms, where the flag rises at 2.76 s and would have stopped the chatter." Find the flaw, and say how you would check it.

CHAPTER 9

Pension Funding Ratio Control Through a Year-Late Valuation and a Contribution Cap

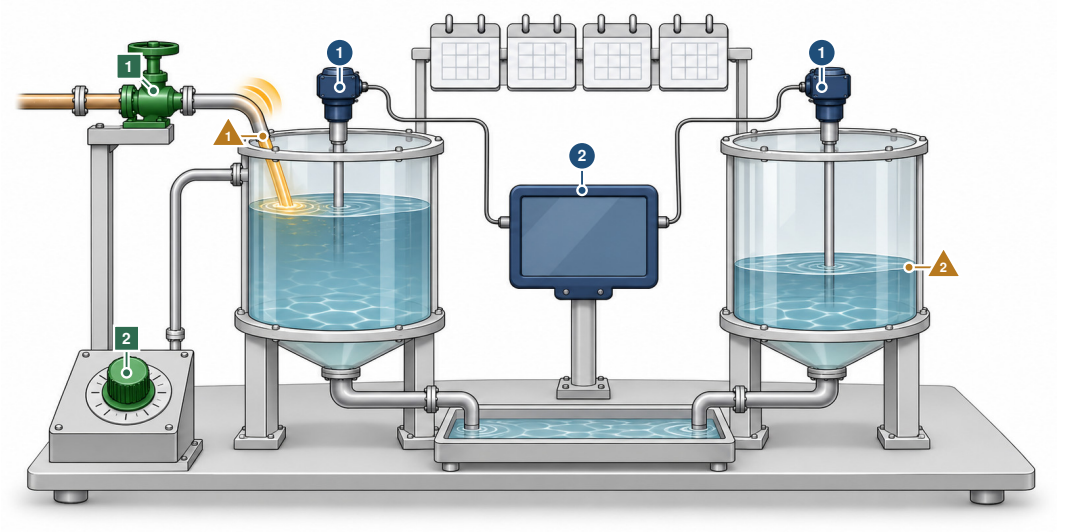

The controller moves 1 the sponsor's contribution 2 the share held in equities

The controller reads 1 assets and liabilities, counted once a year 2 the published valuation, a year late

What disturbs it 1 market returns, a crash among them 2 a fall in the discount rate, which raises the liabilities

The left tank holds the assets and the right tank the liabilities, both counted once a year for a display that shows them a year late, while the contribution valve refills the assets far more slowly than a crash can drain them. Illustration, not to scale.

IN PLAIN TERMS **Refilling a Fund From Last Year's Count**

The main idea. A pension fund holds money today for pensions it must pay for decades. Once a year the employer decides how much to pay in, using a count of the fund's money and promises that is already a year old.

The challenge. A market crash can drain the fund far faster than the largest payment the employer can afford will refill it. And because the count is late, each payment answers last year's fund, not this year's.

Why it matters. The payment rule decides who carries the risk: the employer, through payments that jump, or the fund, through deeper shortfalls.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the two history segments (Section 9.2), the fact on missing filings (Section 9.3), the fact on the late valuation and the two try-its (Section 9.6) and the key points
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 9.6), Technology in Practice (Section 9.7), Limitations (Section 9.8) and the open challenge

The control question A sponsor resets a pension contribution once a year from a valuation a year old, through a cap that refills the fund far more slowly than a market return can drain it. What can a contribution rule achieve under that delay and that cap, and what does a recovery deadline add?

Learning objectives

- Derive the ratio's open-loop growth and the recovery time that a contribution cap sets after a crash.
- Compute the poles of an annual amortization rule on a year-late valuation, and the spread period below which it overshoots.
- Predict the contributions that the calendar's quantum and move limit let a rule apply.
- Explain what a projected estimate sees of a crash, and what a funding buffer buys a predictive plan.
- Judge a funding rule on a declared seed set, not on one path.

Prerequisites Sections 1.4.2, 1.4.5, 1.4.7, 1.6.2, 1.6.3 and 1.6.4. Volume 2, Chapter 10 uses the same calendar on a budget.

The app https://apmonitor.com/apps/pension-funding/

See also Volume 4, Chapter 8 and Volume 4, Chapter 14, stocks with a floor drained faster than they are refilled. Volume 9, Chapter 7, a margin held against an untrusted forecast, and Volume 9, Chapter 2, predictive control behind a dead time.

9.1 Introduction

A defined-benefit pension fund promises benefits that will be paid for decades and holds assets against their present value, so the ratio of assets to liabilities is the level its sponsor, its actuary and its regulator all watch. In the simulator of this chapter, a one-year equity crash of 30 % together with a fall of 1.5 points in the discount rate takes a fully funded fund to a ratio of 0.648 within the year, while a sponsor capped at 12.5 % of payroll can pay at most 2.5 points above the normal cost, which raises the funding ratio by at most 0.00625 a year. The seven-year recovery rule, started by the valuation of year 5, asks for a ratio of 0.8 by year 12, the valuation of year 12 reads 0.698 while the fund stands at 0.707, and the fund regains 0.8 only at 20.23 yr. With the cap at 30 % of payroll the same rule meets the deadline, with contributions that climb to 25.4 %.

Actuaries have treated the setting of contributions as feedback for at least three decades. The classical funding methods add to the normal cost a constant fraction of the unfunded liability [1], and the adjustment was described as a negative feedback component in which the current status is compared with a target and corrective action is taken [2]. Stochastic models then posed the fund level as the controlled variable and the contribution rate as the controlling one, with weighted penalties on the deviations of both [3, 4], and continuous-time models optimized the contribution and the asset allocation together [5, 6]. Treating the fund in a control text therefore needs no translation, because its own literature already speaks of targets, disturbances, delays and quadratic costs.

The structure is general. A stock whose deviation grows on its own, measured once a year and a year late, refilled through a bounded, rate-limited and quantized valve whose authority is small next to the disturbance, describes an endowment drawn down by a spending rule, an insurer's capital, a public debt when interest exceeds growth, a startup's cash, and a storage tank traded against uncertain prices. Volume 2, Chapter 10 treats such a debt with an annual budget, and Volume 4, Chapter 8 such a cash tank with a floor. Each of its parts (dead time, a sampled measurement, an unstable open-loop pole and actuator limits) is standard in process control [7], and Sections 1.4.2 and 1.4.5 treat the first two. The belief the simulator corrects is that a deadline or a faster rule restores a fund. In its crash case every amortization period from 3 to 30 years misses the deadline and ends near a ratio of 0.738, because every one of them sits on the sponsor's cap. A single path also cannot show the trade the rules are chosen by. On seed 11 a thirty-year amortization period ends better funded than a seven-year one, 0.831 against 0.809, on seed 101 it ends at 0.630 against 1.259, and over five declared seeds its lowest published ratio averages 0.598 against 0.695.

However, the actuarial results are stated for independent returns, linear rules without bounds and limiting variances [1, 2, 8], and the recovery rules are written in law with thresholds, deadlines and benefit reductions [9, 10], so a student seldom has one fund on which a rule, its bounds, its calendar and its deadline act together on the same seeded returns. This chapter works with one fund of this kind in full and measures on it what the valuation delay, the sponsor's cap and the return draws do to six ways of setting the contribution. The actuarial models, the amortization rules and the optimizer are established methods from the works cited below. The simulated fund keeps asset and liability ledgers that reconcile at every step, with an annual valuation, a projected estimate, the decision calendar that the economics apps share and a quadratic program, all written out below. Nine cases, a challenge and a reference each isolate one mechanism.

Section 9.2 traces the funding problem from actuarial spreads to recovery plans. Section 9.3 writes the fund as two stocks, with its return, its revaluation, its year-late valuation and its recovery clock. Section 9.4 gives its open-loop and closed-loop character, including the discrete-time poles of the amortization rule with a year of delay and the bound that the cap sets after a crash. Section 9.5 sets out the rules on their calendar, from amortization to the predictive plan. Section 9.6 reports the measured runs, on the five seeds the application declares and on a hundred more. Section 9.7 describes the real valuation, contribution and calendar, Section 9.8 states the limitations and Section 9.9 concludes. The exercises in Section 9.10 use the cases.

9.2 Funding Rules From Actuarial Spreads to Recovery Plans

The arithmetic of a funded pension scheme predates its control theory. The simple model that later stochastic work extends was proposed by Trowbridge in 1952, with an equation of equilibrium that ties the actuarial liability, the normal cost and the benefit outgo together when every assumption is borne out by experience [2]. A funding method fixes the contribution at each valuation as the normal cost plus an adjustment that liquidates the unfunded liability, the difference between the actuarial liability and the fund. In 1988 a simple model with random rates of return was used to study the variability of contributions and fund levels under methods whose adjustment is a constant fraction k of that unfunded liability [1]. Its limiting variances showed a range of spread periods in which a longer period lowers the variance of the contribution and raises that of the fund, and beyond which both rise [2].

Delay entered next, in the work that the history segment below describes [11]. The same group then wrote the problem in the language of optimal control. The 1994 model treated the fund as the controlled variable and the contribution as the controlling variable with weighted penalties on both deviations, the contribution rate risk and the solvency risk [3, 4]. Three ways of adjusting the normal cost as gains and losses arise were compared under random returns [8], and amortizing each gain or loss directly was later found to give more variable funding and contributions than spreading them in proportion, the spread being the contribution control that minimizes mean square deviations [12].

HISTORY 1992 to 2002: The Year-Late Valuation Enters the Theory

An actuarial valuation reports a fund level only after data have been collected and processed and the accounts prepared. A model of 1992 therefore let the adjustment to the contribution depend on the unfunded liability of the year before, and a 1995 analysis compared that one-year delay with no delay at all [2]. With the delay, both limiting variances, of the fund and of the contribution, exceed their values without it, because the information about the fund since the last year has been lost. A trade between the two variances remains, and so does an optimal range of spread periods. The same analysis noted that spread periods of 20 to 25 years were common practice in the United Kingdom [2]. A study of 2002 then controlled the fund ratio rather than the fund level, with a one-year delay in the accounts, and derived its funding policy through the separation principle [4].

Continuous time added the asset allocation as a second handle. A model with several risky assets and random benefit outgo found that a quadratic loss leads to investing more in risky assets when the funding level is low and shifting toward low-risk assets as it rises, which its author called a counterintuitive strategy, and it listed the constraints that complicate the analysis, among them a sponsor's limit on the contribution rate such as twice the normal rate [5]. Minimizing the contribution rate risk together with the solvency risk was shown to lead in several cases to a spread method of funding [6]. For defined-contribution plans, deterministic lifestyling moves gradually from equities to bonds by preset rules [13], which is the family to which the application's glide path belongs, keyed here to the funding ratio rather than to age.

HISTORY 2006: A Seven-Year Rule on Last Year's Ratio

The Pension Protection Act, approved in the United States on August 17, 2006, rewrote the minimum funding rules of single-employer defined-benefit plans. An underfunded plan's minimum required contribution became its target normal cost plus a shortfall amortization charge, which repays each year's shortfall base in level annual installments over seven plan years. A plan was placed in at-risk status when its funding target attainment percentage for the preceding plan year was below 80 percent and, computed with additional actuarial assumptions, below 70 percent. For plan years beginning in 2008 the act let that preceding-year percentage be determined by methods of estimation that the Secretary of the Treasury might provide [9]. The test therefore reads a ratio a year old, as the valuation of this chapter does.

The law fixed the period and the deadline, in the United States through the act of 2006 that the second history segment describes [9]. In the Netherlands the solvency crisis of 2001 to 2004 moved most pension funds to average-wage plans with solvency-contingent indexation [14]. The Dutch Pension Act as it read in the first half of 2015 required a fund below its required own funds to submit a recovery plan reaching them within at most ten years, and measures that may reduce accrued rights when its policy funding ratio stays below the minimum at five consecutive annual determinations [10]. Two facts outside the rules set the stakes. The size of the promise depends on the discount rate, so that state pension liabilities valued at zero-coupon Treasury yields came to $4.43 trillion against $3.20 trillion at rates matched to state debt [15], and the cash a sponsor pays is not free, because capital expenditures fall with mandatory contributions, most visibly at firms facing financing constraints [16].

9.3 A Fund of Two Stocks

The simulated fund is one open defined-benefit plan with a stationary membership, drawn in Figure 9.1 as the tanks and valves it is. Money is measured in units of the initial liability L_0, time in years, and every flow is a rate per year. Table 9.1 classifies the signals. Every coefficient is a teaching value of Table 9.2, chosen for an imaginary plan and not estimated for any fund.

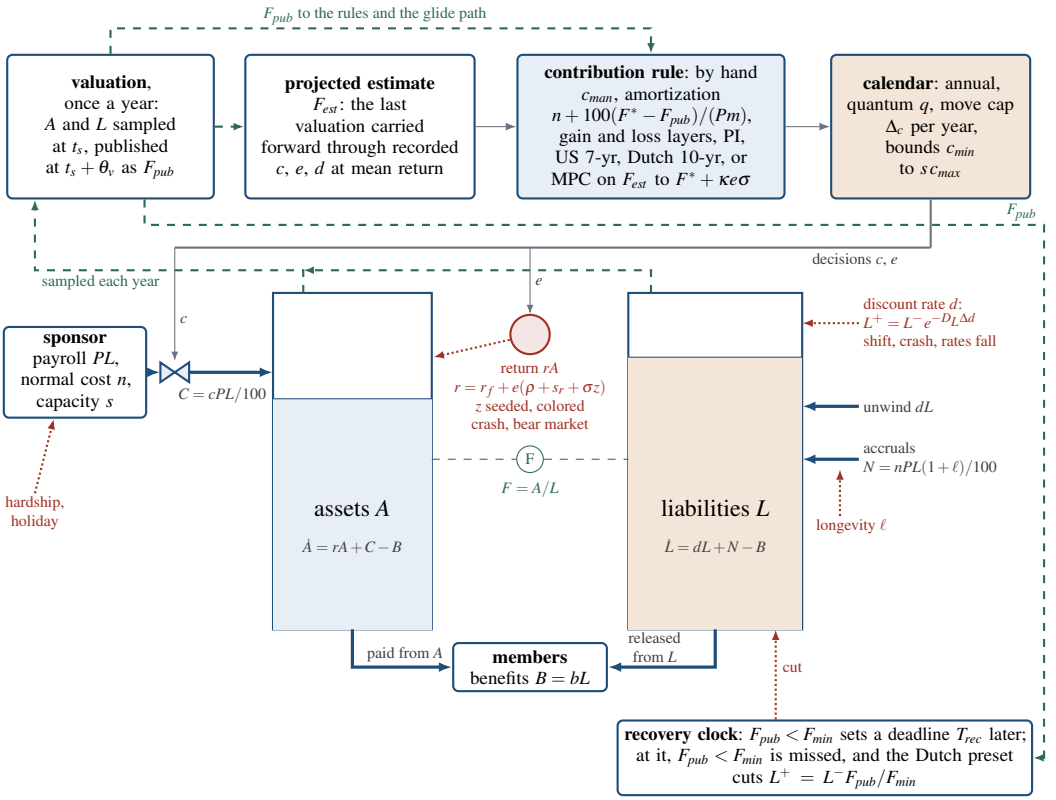

Figure 9.1 The defined-benefit fund as two stocks, not to scale. Contributions C from the sponsor and the return rA fill the assets A, the unwind dL and the accruals N fill the liabilities L, and the benefits B leave both. A move of the discount rate d revalues L through its duration D_L. Dotted red arrows are disturbances, thin solid gray arrows the decisions c and e on the calendar, dashed green lines the measurement path, and thick solid blue arrows the flows of money. The valuation samples A and L once a year and publishes F_{pub} a delay θ_v later, the estimate carries it forward to F_{est}, and the recovery clock may cut L at a missed deadline.

Table 9.1 Signals of the Fund by Role, With Their Defaults

Symbol	Role	Default	Range
F	CV, current ratio A/L, shown, not used by rules	1	—
F_{pub}	measured, last published valuation	1	yearly
F_{est}	estimate projected from F_{pub}	1	—
F^*, F_{min}	SP, funding target and recovery threshold	1, 0.8	0.8–1.4, 0.6–1
c	MV, contribution, % of payroll	10	0–30
e	MV, equity share, by hand or glide	0.6	0–1
s_r	DV, equity return shift and scripted shocks	0	±0.2/yr
Δd	DV, discount-rate shift and scripted moves	0	±0.03/yr
ℓ	DV, additional accrual fraction	0	0–0.2
s	DV, sponsor capacity factor	1	0.2–1
z	DV, seeded colored return noise	seed 11	1–9999
A, L	states, assets and liabilities	1, 1	L_0
—	states, seven ledgers of every flow	0	L_0, %·yr
σ_c	output, time-weighted contribution SD	0	points

9.3.1 Assets, Liabilities and the Ratio

Contributions and returns fill the assets, benefits drain them, and the liabilities grow by the unwinding of the discount and by new accruals while the same benefits leave them, as shown in Equation (9.1).

$$\begin{aligned} \dot{A} &= rA + \frac{c}{100}PL - bL \\ \dot{L} &= dL + \frac{n}{100}PL(1+\ell)\lambda - bL, \qquad F = \frac{A}{L} \end{aligned} \tag{9.1}$$

where $P = 0.25$ is the annual payroll as a fraction of the liabilities, c and $n = 10$ are the contribution and the normal cost in percent of payroll, $b = 0.05$ is the benefit outgo as a fraction of the liabilities a year, d is the discount rate, ℓ the additional accrual fraction and λ the product of 1.08 for each longevity event, one otherwise. Payroll, accruals and benefits are all proportional to the liability, so the fund has no cohorts. Each flow is also integrated into a ledger of its own, and the assets and the liabilities are reconciled against their ledgers at every step. Over all eleven cases the larger residual is $2.6 \times 10^{-14} L_0$, revaluations and benefit cuts included.

9.3.2 Return, Discount Rate and Revaluation

The fund earns the risk-free rate on everything and the equity premium, the noise and any scripted shock on its equity share, as shown in Equation (9.2).

$$\begin{aligned} r &= r_f + e\big(\rho + s_r + \sigma z + j\big) \\ z_{k+1} &= a z_k + \sqrt{1-a^2}\,\nu_k, \qquad a = e^{-\Delta t/\tau_z} \end{aligned} \tag{9.2}$$

where $r_f = 0.016\,\text{yr}$, $\rho = 0.04\,\text{yr}$, $\sigma = 0.16$, ν_k is a standard normal deviate from a mulberry32 generator seeded by the return seed, $\tau_z = 1\,\text{yr}$, and z is held over each weekly step $\Delta t = 1/52$ year. An optional crash year sets $j = \ln 0.75$ for a whole year with probability 0.08, drawn from the return seed at each year boundary, and $j = 0$ otherwise, so a drawn year takes 25 % off equities whatever σ is. A crash event adds $s_r = \ln 0.7$ for one year, so at $e = 0.6$ the assets fall by a factor $0.7^{0.6} = 0.807$ over the year from that term alone. The return is a force of interest, and σ is a scale rather than an annual volatility. Integrated over a year, the noise term's spread is 0.857 times $e\sigma$ in standard deviation, 8.23 % at $e = 0.6$, and consecutive annual integrals are correlated at 0.54, which the application's generator reproduces over 200 000 simulated years against 0.858 and 0.543 for the continuous process. With $r_f + 0.6\rho = d_0 = 0.04\,\text{yr}$ the expected return equals the discount rate, the matched point at which the deterministic fund stays at $F = 1$.

The discount rate is the initial 0.04 yr plus the slider shift and any scripted move, and a change revalues the liabilities through their duration, as shown in Equation (9.3).

$$L^{+} = L^{-} \exp\big(-D_L(d_k - d_{k-1})\big) \tag{9.3}$$

where $D_L = 15\,\text{yr}$ and k counts weekly steps, so a fall of 1.5 points raises L by a factor $e^{0.225} = 1.252$. The revaluation enters its own ledger, separate from accruals. A scripted move may be a step, taken over one week, or a ramp, such as two points over five years. The asset return does not depend on d, so the fund holds no bonds whose price would rise as the rate falls.

9.3.3 The Valuation and the Projected Estimate

The valuation is the framework's sampled analyzer, the instrument of Section 1.4.5 with a sample period of a year. At each integer year $t_k \geq 1$ it publishes the assets and liabilities as they stood θ_v earlier, as shown in Equation (9.4).

$$F_{pub}(t) = \frac{A(t_k - \theta_v)}{L(t_k - \theta_v)}, \qquad t_k \leq t < t_{k+1} \tag{9.4}$$

Table 9.2 Model Constants and Disturbances. All Are Teaching Values, Shown in the Show Model Dialog (M) or Among the Disturbances (D, A for Advanced)

Symbol	Meaning	Value	Range	
F_0	initial funding ratio	1	0.5–1.6	M
r_f, ρ	risk-free return, equity premium	0.016, 0.04/yr	0–0.06, 0–0.1	M
σ, τ_z	noise scale, correlation time	0.16, 1 yr	0–0.3, fixed	M
d_0, D_L	discount rate, liability duration	0.04/yr, 15 yr	0–0.08, 3–25	M
P, n	payroll over L, normal cost	0.25/yr, 10 %	0.1–0.5, 5–20	M
b	benefit outgo over L	0.05/yr	0.02–0.12	M
θ_v	valuation delay	1 yr	0–2	M
—	return seed, crash year j	11, off	1–9999, 0/1	M
s_r, s	return shift, sponsor capacity	0, 1	±0.2, 0.2–1	D
Δd, ℓ	rate shift, extra accrual	0, 0	±0.03, 0–0.2	A
—	crash: $s_r = \ln 0.7$ for 1 yr, $\Delta d = -0.015$	event	—	D
—	rates fall 2 points over 5 yr, longevity $\lambda = 1.08$	event	—	D
—	hardship, s halves for good	event	—	D
Δt	integration step and log interval	1/52 yr	—	—

where $\theta_v = 1$ yr, so the valuation of year 5 describes year 4, and before the first one the published value is the initial state. A held valuation is therefore on average a year and a half old, by Equation (1.8) with $t_d = \theta_v$. A crash that begins at year 3 therefore first appears in the valuation of year 5, because the valuation of year 4 describes the instant before it. The projected estimate $F_{est} = \hat{A}/\hat{L}$ restarts from the published A and L at each valuation and integrates Equation (9.1) forward through every recorded step since the sample, by the midpoint rule, with the recorded contribution, equity share, discount rate, accrual and revaluation, at the declared mean return, $s_r = z = j = 0$. It never reads a realized equity return, random or scripted, so a crash reaches it only through the next valuation. Restarting from the instant the valuation describes and predicting forward through the recorded inputs is the second treatment of a late measurement in Section 1.4.7, with no filter gain, and the projection is a predictor across the delay in the sense of Section 1.4.3.

> **+ DID YOU KNOW? Three Years Without a Reading**
>
> The United States insurer of private defined-benefit plans put the underfunding of single-employer plans at $482 billion in 2021. Its report on fiscal year 2024 showed no estimate for 2022 to 2024, because the plans' annual filings were not yet available, and its model projected about $41 billion by the end of 2024 [17].

9.3.4 Recovery Clock and Benefit Cut

A published ratio below the threshold starts a deadline, and a missed deadline under the Dutch preset reduces the liabilities, as shown in Equation (9.5).

$$\begin{aligned} F_{pub} < F_{min} &\Rightarrow t_{dl} = t_k + T_{rec} \\ F_{pub} \geq F_{min} &\Rightarrow \text{deadline cleared} \\ t_k \geq t_{dl},\ F_{pub} < F_{min} &\Rightarrow \text{missed},\ L^+ = L^- F_{pub}/F_{min} \end{aligned} \tag{9.5}$$

where T_{rec} is seven years by default and ten under the Dutch preset. The deadline is judged by the first valuation at or after it. The Dutch cut, made once, brings that valuation to the threshold, applies to the liabilities as they stand, restates the published ratio and pauses the run. Negative assets would be written off as insolvency, which no shipped case reaches.

9.3.5 Initial State, Step and Checks

Every case starts with $A = F_0$ and $L = 1$, its noise at zero and its valuation equal to the initial state, with the first sample due at year 1. Classical fourth-order Runge-Kutta integrates the nine states on the weekly step, with the contribution, the equity share and the noise held over each step. The

revaluation, the estimate, the valuation, the recovery clock with its cut and the next noise value are applied between steps, in that order, and the historian logs every week. The controller runs every fifth step, $5/52 = 0.0962$ yr, the period the application declares, because a nominal month is not a whole number of weekly steps. Section 1.6.1 describes the engine's clocks, its seed and its Reset. In the matched reference run for thirty years the ratio stays at 1 while assets and liabilities both grow to $e^{0.45} = 1.568\,L_0$. Halving the step in the deterministic crash at a fixed 12.5 % moves the ratio at year 15 from 0.776253 to 0.776240, and a reset with the same seed reproduces a case bit for bit. The application passes its 35 checks, 18 of them its own.

9.3.6 Assumptions

An aggregate fund. Payroll, accruals and benefits are fixed fractions of the liabilities, with no cohorts, mortality tables, salary scales, indexation or taxes. A real valuation projects each member's benefits and would move the normal cost with the membership.

A return independent of the discount rate. The assets carry no duration, so a fall in rates revalues only the liabilities. A fund that hedges with long bonds would see part of the revaluation offset on the asset side.

Colored noise. Annual returns correlated at 0.54 make five-year variance larger than independent years would, a standard deviation of 2.82 times $e\sigma$ against $0.857\sqrt{5} = 1.92$. The actuarial results of Section 9.2 assume independent returns.

Simplified rules. The legal presets carry one threshold, one deadline and one cut, not the statutes' many tests.

Teaching values. Every parameter, threshold, deadline and challenge target is a teaching value. Real ones come from a plan's actuarial valuation, the statute that governs it [9, 10] and its regulator, and none of the numbers here describes any fund.

9.4 The Ratio, Its Gain and Its Lag

9.4.1 A Deficit That Grows by Itself

Dividing Equation (9.1) by L gives the funding ratio as one linear equation in the ratio itself, as shown in Equation (9.6).

$$\dot{F} = (r - g_L)F + \frac{P}{100}c - b, \qquad g_L = d + \frac{n}{100}P(1+\ell)\lambda - b \tag{9.6}$$

where g_L is the growth rate of the liabilities, 0.015 yr at the defaults. At the matched point $r - g_L = 0.025$ yr, so a deficit held at constant contributions grows on its own and doubles in 27.7 yr. The benefits are a fixed fraction of L but come out of A, and a fund below full funding pays a larger share of its own assets. A point of payroll moves $\dot{F}$ by only $P/100 = 0.0025$ a year. After the crash case's fall of 1.5 points in d the open-loop gain $r - g_L$ becomes 0.040 yr, and at the 12.5 % cap Equation (9.6) has an unstable equilibrium at $F = 0.469$. From its low of 0.648 the ratio therefore climbs at only 0.0072 yr, and at the 30 % cap at 0.051 yr. The engine confirms the bound, with the seven-year rule reaching 0.8 at 20.23 yr under the low cap and at 10.60 yr under the high one.

9.4.2 Amortization as a Proportional Gain

The amortization rule sets the contribution to the normal cost plus the deficit spread over m years, $c = n + 100(F^* - F)/(Pm)$, which on Equation (9.6) at the matched point gives a closed-loop pole of $r - g_L - 1/m = 0.025 - 1/m$ per year. The application's own check finds -0.041667 at $m = 15$ against the same value analytically, to 1.6×10^{-13}. The pole is -0.175 at $m = 5$ and -0.0083 at $m = 30$, and it reaches zero at $m = 40$, beyond the slider's range of 3 to 30, so a long enough spread no longer brings a deficit back at all.

The application decides once a year on a valuation one year old. With the contribution held for exactly a year (the application's decisions fall 0.9615 or 1.0577 years apart, a year apart on average)

Table 9.3 Contribution and Allocation Settings. A Marks an Advanced Control

Symbol	Setting	Default	Range
—	contribution rule	amortization	6 rules
—	decision calendar	annual	quarterly, every run
c_{man}	request in MANUAL, % payroll	10	0–30
m	amortization period	15 yr	3–30
c_{min}, c_{max}	contribution bounds (A for c_{min})	5, 30	0–15, 10–30
Δ_c, q	move limit per year, quantum (A)	3, 0.1 points	0.5–10, 0.01–0.5
T_{rec}	recovery deadline (A)	7 yr	3–15
—	holiday above $F_{pub} = 1.2$ (A)	off	on
e	equity share	0.6	0–1
e_{hi}, e_{lo}	glide endpoints (A)	0.6, 0.3	0–1
F_1, F_2	glide band (A)	1.0, 1.3	0.6–1.2, 1–1.6
τ_I, τ_D	PI integral and derivative times (A)	15, 0 yr	0–30, 0–5
H	MPC horizon (A)	15 yr	5–20
κ	funding buffer, in $e\sigma$	2	0–5
w_F, w_c	funding weight (A), smoothing weight	500, 1	10–2000, 0.1–20

and set from the deficit $x = F - 1$ of the year before, the deviation obeys Equation (9.7).

$$x_{k+1} = \phi x_k - \frac{\gamma}{m} x_{k-1}, \qquad \phi = e^{0.025}, \quad \gamma = \frac{e^{0.025} - 1}{0.025} \tag{9.7}$$

whose characteristic roots solve $z^2 - \phi z + \gamma/m = 0$, the equation of a correction that acts on information two decisions old, Equation (1.9), with the fund's growth factor ϕ beside the first term and γ/m in place of g. Without the delay the single root is $\phi - \gamma/m$. At $m = 5$ the delayed loop has roots 0.758 and 0.267 against 0.823 without the delay, at $m = 15$ 0.955 and 0.071 against 0.958, and the roots become complex only for $m < 3.85$ years, $0.513 \pm 0.273\,j$ at $m = 3$. The delay therefore does not slow a deterministic recovery. A rule that acts on last year's larger deficit pays more while the deficit shrinks, which raises the variability of the contribution, as Section 9.6 measures, in the direction the stochastic analysis found for the limiting variances [2]. Equation (9.7) ignores the move cap and the quantum, which bind at the first decision of case 4, where the rule asks for 18.0 points and the cap allows 13.0.

9.5 Deciding Once a Year

Every contribution, whether requested by hand, by a rule or by the optimizer, passes the same calendar and the same shaping before it reaches the valve. Table 9.3 lists the settings.

9.5.1 Annual Decisions, Their Grid and Their Cap

The calendar and shaping functions are those of `shared/src/framework/core/cadence.js`, which the central bank of Volume 2, Chapter 11 also decides through. The annual calendar has an instant at every integer year, the quarterly one at every quarter, and the continuous one makes every controller execution a decision. A decision waits for the first controller execution not earlier than its instant, so on the five-week execution grid the annual decisions fall at years 0, 1.0577, 2.0192, 3.0769, 4.0385 and 5.0000, a pattern that repeats every five years, and the quarterly ones 0.1923 or 0.2885 years apart. At a decision the request c^w is shaped as shown in Equation (9.8).

$$\begin{aligned} \tilde{c} &= q\,\mathrm{round}(c^w/q), \quad \Delta_q = q\lfloor(\Delta_c\tau + \chi_k)/q\rfloor \\ c_k &= \mathrm{clip}_{[c_{lo},c_{hi}]}\big(c_{k-1} + \mathrm{sat}_{\Delta_q}(\tilde{c} - c_{k-1})\big) \\ \chi_{k+1} &= \Delta_c\tau + \chi_k - |c_k - c_{k-1}| \text{ if in } [0,q), \text{ else } 0 \\ c_{hi} &= \max(0, s\,c_{max}), \qquad c_{lo} = \min(c_{min}, c_{hi}) \end{aligned} \tag{9.8}$$

where ties round away from zero, the bounds snap to the grid, τ is the time since the last decision, and the credit χ carries the part of the allowance $\Delta_c\tau$, 2.88 or 3.17 points as τ is 0.9615 yr or 1.0577 yr,

that flooring to whole quanta left unused, so every calendar ramps at Δ_c. At the decision made while a case is set up, τ is a whole calendar step. The contribution then holds until the next decision. A scripted holiday, or the holiday switch while $F_{pub} > 1.2$, sets $c = 0$ outside this shaping, and a cap that falls below the contribution in force cuts it at once through the clip. The equity share, by hand or from the glide path, is set at the same decisions. In MANUAL the learner's request passes Equation (9.8) exactly as a rule's does.

9.5.2 Amortization, the Legal Presets and the Holiday

The amortization rule is the proportional law of Section 9.4 on the published ratio, as shown in Equation (9.9).

$$c^w = n + \frac{100(F^* - F_{pub})}{Pm} \tag{9.9}$$

The gain $100/(Pm)$ is in points of payroll per unit of funding ratio, 26.7 at $m = 15$, and surpluses lower the contribution exactly as deficits raise it, down to c_{lo}. The historical United States preset is Equation (9.9) with $m = 7$ and a seven-year deadline, and the Dutch-style preset uses $m = 10$ with a ten-year deadline and the cut of Equation (9.5). Both are teaching simplifications. The 2006 act amortizes each year's shortfall base in its own seven-year layer and reduces the normal cost by a surplus rather than spreading it [9], which the application's Model dialog and case 5's card also state, and the Dutch text measures its funds against required and minimum own funds rather than a ratio of 0.8 [10].

9.5.3 Gain and Loss Layers

The gains-and-losses rule recognizes each valuation's new deficit apart from the deficit already being repaid, and amortizes it in a layer of its own, as shown in Equation (9.10).

$$\begin{aligned} G_j &= \frac{100(F^* - F_{pub,j})}{P}, \qquad \delta_j = G_j - (G_{j-1} - \Pi_{j-1}) \\ \Pi &= \sum_{t_j > t-m} \frac{\delta_j}{m}, \qquad c^w = n + \Pi \end{aligned} \tag{9.10}$$

where G_j is the deficit at valuation j in points of payroll times years, Π_{j-1} is the payment scheduled before it, and each layer lasts m years. A loss is therefore repaid within m years of the valuation that first reports it, rather than refinanced over a fresh m years at each valuation as in Equation (9.9). This is the structure of the act's shortfall bases [9] and of the direct amortization of gains and losses that the actuarial literature compared with spreading [8, 12].

9.5.4 Proportional-Integral Control

The PI rule is the framework's PID in the velocity form of Section 1.6.2, executed only at decisions on the published ratio, as shown in Equation (9.11).

$$\begin{aligned} e_k &= F^* - F_{pub,k}, \qquad S_k = S_{k-1} + K_c(e_k - e_{k-1}) \\ \tilde{u}_k &= S_k + I_{k-1}, \qquad u_k = \mathrm{clip}_{[c_{lo}, c_{hi}]}(\tilde{u}_k) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{c_k - \tilde{u}_k}{\tau_I} \right] \tau_k \end{aligned} \tag{9.11}$$

where $K_c = 100/(Pm)$, the amortization gain, $\tau_I = 15\,\mathrm{yr}$, τ_k the time since the last decision, and the integral starts at the request in force, ten points, so that the normal cost is carried by I rather than a bias. The symbols map onto Equations (1.15) to (1.18) with S_k the running sum and $\tilde{u}_k$ the output before limiting. As there, only the emitted contribution is clipped and only the integral is back-calculated, here with the tracking time τ_I. A derivative on the measurement, filtered over one year, is available and off as shipped. The output u_k, slewed to the move allowance, is the request c^w

of Equation (9.8), and the integral is back-calculated against the applied c_k, so the move limit, the quantum, a holiday and a falling cap all unwind it. Under every other rule and in MANUAL the PI tracks the contribution in force at each decision, so a switch to it is bumpless.

9.5.5 The Glide Path

With the glide path on, the equity share falls linearly with the published ratio across a band, as shown in Equation (9.12).

$$e = e_{hi} + (e_{lo} - e_{hi})\,\mathrm{clip}_{[0,1]}\Big(\frac{F_{pub} - F_1}{F_2 - F_1}\Big) \tag{9.12}$$

which de-risks as the fund improves, the direction a quadratic loss also chose [5]. At the defaults the share is 0.6 below $F_{pub} = 1.0$ and 0.3 above 1.3, and the expected return falls from 0.040 yr to 0.028 yr, below the discount rate.

9.5.6 A Certainty-Equivalent Recovery Plan

The MPC solves a quadratic program (QP) at each decision on the linear ratio model of Equation (9.6), with the mean return $\bar{r} = r_f + e\rho$ at the declared $r_f = 0.016\,\mathrm{yr}$ and $\rho = 0.04\,\mathrm{yr}$ and the equity share in force, g_L at the discount rate in force, and an exact zero-order hold over the calendar's step T_s (one year, a quarter, or one controller period on the continuous calendar), through the framework's `core/linmpc.js`. From the projected estimate $x_0 = F_{est}$ it chooses the contributions u_j of Equation (9.13).

$$\begin{aligned}
\min_{u,\varepsilon}\ & \sum_{k=1}^{N} w_F (x_k - \bar{F})^2 + 4 w_F (x_N - \bar{F})^2 \\
& \quad + w_c \sum_j (u_j - u_{j-1})^2 + 10^{-3} \sum_k (u_k - n)^2 \\
& \quad + 5000\,\varepsilon^2 + 100\,\varepsilon \\
\text{s.t. } & x_{k+1} = a_d x_k + b_d u_k + c_d \\
& x_k \geq F_{min} + \kappa e \sigma - \varepsilon, \qquad \varepsilon \geq 0 \\
& c_{lo} \leq u_j \leq c_{hi}, \qquad |u_j - u_{j-1}| \leq \Delta_c \tau + \chi
\end{aligned} \tag{9.13}$$

where $\bar{F} = F^* + \kappa e \sigma$ on the declared $\sigma = 0.16$, 1.192 at the defaults, $N = H/T_s = 15$ annual steps of exactly one year, u_0 is the contribution in force, and the input is held over ten move blocks, nine of one year and one of six. The QP has 11 variables and 56 constraint rows. The first move is rounded to the quantum by Equation (9.8), and the plan is drawn in the Recovery plan view. Over the 15 decisions of case 9 the Goldfarb–Idnani dual active-set solver [18] needs a mean of 2.80 iterations and at most 8, with no failure, and the soft floor binds in four of the plans, those of years 8.08 to 11.06. Section 1.6.3 gives the framework's MPC option in general. The buffer $\kappa e \sigma$ is a margin held above a floor against a disturbance the plan does not predict, as the empty flood space of Volume 9, Chapter 7 is held against an inflow its forecast cannot promise.

9.5.7 Views, Status and the Hand Station

The Fund view draws the asset tank beside the liability column with their flows and the return as weather, the Trade-off view the five-seed means of Figure 9.2(b), and the Recovery plan view the realized ratio, the labeled threshold and deadline, and the MPC's last plan from the year it was solved. Badges give the current and published ratios, the contribution, the equity share and the time left to a deadline, and the status line names UNDERFUNDED, RECOVERY PLAN, DEADLINE MISSED, CUT, HOLIDAY, SPONSOR CAP, CONTRIBUTION FLOOR and CRASH as they apply. In MANUAL the learner sets the request, and in either mode the equity share or the glide path. A helper sets the request to the normal cost, which holds the matched reference exactly. A case is

Table 9.4 The Cases, the Challenge and the Reference

#	Change from the defaults, run length	Teaching point
1	MANUAL at 10 %, bear market of −25 % over years 2 to 5, 15 yr	the valve by hand
2	$m = 7$, seed 101, 30 yr, Trade-off view	fast against smooth
3	gain and loss layers, $m = 15$, 30 yr	layers or a spread
4	$\sigma = 0$, $F_0 = 0.9$, $m = 5$, 15 yr	the late valuation
5	$\sigma = 0$, US 7-yr, cap 12.5 %, crash at 3, 15 yr	no cash, no recovery
6	US 7-yr, $e = 0.9$, crash at 3, 15 yr	risk in the recovery
7	glide on, $F_0 = 1.25$, 30 yr	de-risking
8	$F_0 = 1.25$, holiday to 5, crash at 5, 20 yr	the holiday's bill
9	MPC, glide on, crash at 3, 15 yr	a plan on an estimate
ch	US 7-yr, cap 25 %, 20 % at start, $F_0 = 1.25$, crash at 3, rates −1 point over 4 to 8	survive both
ref	$\sigma = 0$, defaults, 15 yr	the matched fund

complete once run for its length with an observation recorded, and the challenge once its verdict is met, under the rules of Section 1.6.4.

9.5.8 Fund, Rules and Plan

Equations (9.1) to (9.5) are the simulated fund. The amortization, layer and PI rules hold no model beyond n and P in their gain. The MPC's model is Equation (9.6) at the declared mean returns and noise scale, with the allocation, the discount rate and the accrual frozen at their present values, started from an estimate that knows the recorded flows and the yields but no realized equity return, random or scripted. A learner's change of r_f, ρ or σ is thus a model mismatch, which the plan meets as it happens. Planning from an estimate with a certainty-equivalent controller is the separation of estimation from control that the 2002 actuarial study derived for a one-year accounting lag [4], with a deterministic projection here in place of its minimum mean-squared-error estimate. The models they stand for are the discrete-time actuarial models with independent annual returns [1, 2, 11], the continuous-time control models [5, 6], and a plan's full actuarial valuation. The application fits nothing to data.

9.6 Fifteen and Thirty Years at a Time

Table 9.4 lists the cases. Each is selected from the dropdown, changed as its row says and run headless on the engine, on its own return seed, 101 for case 2 and 11 for the others, unless another is named. A change made before Run gives the same run as the change followed by Reset, except a disturbance button, whose fired event a Reset clears. Table 9.5 collects the outcomes. Because the plant is stochastic, Table 9.6 repeats the stochastic cases on the five seeds the application declares, 11, 29, 47, 71 and 101, and Table 9.7 the challenge. The scorecard's integrated error grows with the run and is not quoted. The exported CSV file is logged weekly and holds the current, published and projected ratios, the contribution and its request, the equity share, every flow, the contribution SD, the lowest valuation and the times of the last decision and of the last valuation's sample and report.

9.6.1 The Valve by Hand

Case 1 holds the contribution at the normal cost through a bear market that takes 25 % off equities over three years, which at $e = 0.6$ is a factor $0.75^{0.6} = 0.842$ on the assets from that term. On seed 11 the fund peaks at 1.38 at 2.4 yr, falls to 0.945 at 9.19 yr and ends at 1.303, because that seed's draws are kind. On the five declared seeds the same unchanged hand ends between 0.262 and 1.303, three of the five valuations fall below 0.8, and two recovery deadlines are missed. A hand that requests 15 % from the start on seed 11 keeps every valuation at or above 1.000 and ends at 1.554, at a cost of 0.631 rather than 0.421 initial liabilities of contributions.

Table 9.5 Measured Cases on Their Own Seed, 101 for Case 2 and 11 Otherwise, With Final and Lowest Published Ratios, Contribution SD, and Mean and Largest Applied Contribution

Case and configuration	$F(T)$	F_{pub}^{min}	σ_c	$\bar{c}$	c^{max}	note
1, MANUAL at 10 %	1.303	0.954	0.00	10.0	10.0	
1, MANUAL at 15 %	1.554	1.000	0.00	15.0	15.0	
2, $m = 7$	1.259	0.530	8.41	19.3	30.0	missed
2, $m = 30$	0.630	0.386	3.43	13.9	18.2	missed
3, layers, $m = 15$	0.724	0.785	1.92	6.2	11.8	plan
3, amortization	0.797	0.860	2.58	7.2	12.8	
3, PI, $\tau_I = 15$ yr	0.731	0.793	2.04	6.3	11.3	plan
4, $\theta_v = 1$ yr	0.998	0.900	2.35	12.9	17.6	
4, $\theta_v = 0$	0.994	0.900	1.91	12.9	16.6	
4, $\theta_v = 2$ yr	1.002	0.900	2.93	13.0	18.0	
4, $m = 3$, $\theta_v = 1$ yr	1.000	0.900	3.85	12.9	21.0	1.004
5, US 7-yr, cap 12.5 %	0.738	0.648	1.18	11.7	12.5	missed
5, cap 30 %	0.956	0.648	5.63	16.5	25.4	met
5, Dutch 10-yr	0.812	0.648	1.18	11.7	12.5	cut
6, $e = 0.9$	1.499	0.777	3.64	9.2	17.1	
6, $e = 0.3$	0.994	0.710	6.41	14.2	24.2	
7, glide on	0.796	0.858	2.56	6.7	12.8	
7, glide off	1.088	1.186	0.37	5.1	7.0	
8, holiday	1.165	0.879	4.17	6.0	13.2	
9, MPC, $\kappa = 2$	1.331	0.839	2.85	12.9	18.0	
9, MPC, $\kappa = 0$	1.126	0.739	3.52	9.4	14.8	
9, US 7-yr	1.245	0.773	4.40	11.4	19.1	
ch, as shipped	1.379	0.863	3.91	9.3	17.0	not met
ch, suggested	1.158	0.954	1.46	18.9	21.0	met
ref	1.000	1.000	0.00	10.0	10.0	

T, the case's run length. Plan, started at year 30. 1.004, the peak of F.

9.6.2 Spread Period, Layers and the Trade

Case 2 compares spreading the deficit over 7 and over 30 years on the same draws. On seed 11 the longer period halves the contribution standard deviation, 1.86 against 3.74 points, and also ends better funded, 0.831 against 0.809, with a lowest valuation of 0.901 against 0.864. The case therefore ships on seed 101, where the SD falls from 8.41 to 3.43 points while the lowest valuation falls from 0.530 to 0.386 and the final ratio from 1.259 to 0.630, the trade its card describes. Seeds 29 and 71 show the same order, 11 and 47 do not. Over the five seeds the SD averages 2.92 points at $m = 30$ against 6.11 and the lowest valuation 0.598 against 0.695. The application's trade-off view averages the within-path standard deviations over the five seeds for seven periods, as shown in Figure 9.2(b), from 6.57 points and 0.299 at $m = 3$ to 2.92 points and 0.351 at $m = 30$. Each point is a mean over very different paths. Seed 47's funding SD is between 0.699 and 0.726 at every period, because its fund runs into surplus where the contribution sits on its 5 % floor, and seed 11's falls from 0.195 to 0.178 as m rises.

Case 3 replaces the recomputed spread with layers. On seed 11 the layers give the smoothest contribution of the three rules, 1.92 points against 2.58 for the spread and 2.04 for the PI, and the lowest funding, 0.724 at year 30 against 0.797, with a recovery plan started by the valuation of year 30. On the five seeds the order of the contribution SDs reverses, 5.86 points for the layers against 4.33 for the spread and 5.46 for the PI, and the lowest valuation averages 0.638 against 0.653, in line with the finding that spreading is more efficient than amortizing each gain and loss [12]. A PI with $\tau_I = 15$ yr behaves much like the layers, likely because its integral also keeps a deficit on the books after the proportional term has let it go.

9.6.3 The Late Valuation

Case 4 starts a deterministic fund at 0.9 with $m = 5$. With the valuation a year late the contribution climbs from 10 to 13.0, 16.1 and 17.6 points over the first three decisions and the ratio reaches 0.998

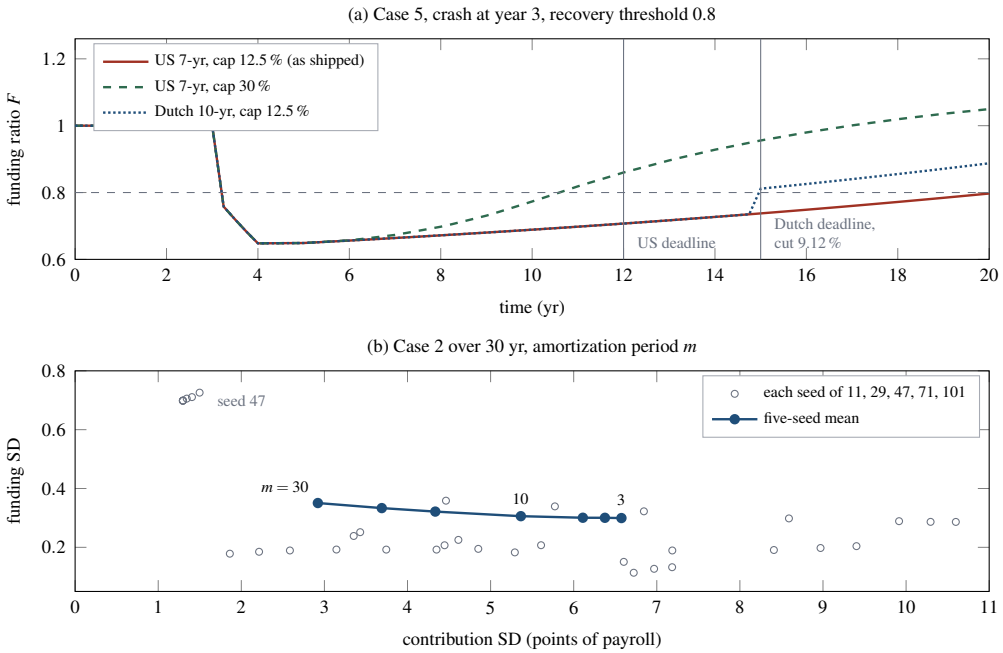

Figure 9.2 Authority and trade-off measured on the application's engine. (a) Case 5, deterministic, under the seven-year rule at the shipped 12.5 % cap (solid), at a 30 % cap (dashed), and under the ten-year Dutch preset (dotted), which cuts the liabilities by 9.12 % at year 15, where the run pauses until Run is pressed again. (b) Contribution SD against funding SD over 30 years for amortization periods of 3, 5, 7, 10, 15, 20 and 30 years, five-seed means (filled circles on a solid line) and each seed (open circles).

at year 15 without overshoot, against 0.994 with no delay. The delay raises the contribution SD from 1.91 to 2.35 points and to 2.93 at two years, where the ratio peaks at 1.003 at 12.02 yr. At $m = 3$ the delayed rule overshoots to 1.004 at 8.08 yr with an SD of 3.85, while without delay it does not, which matches the complex roots of Equation (9.7) below $m = 3.85$. At $m = 20$ the ratio ends at 0.933. The lag is a cost in contribution variability more than in the speed of recovery.

> ✦ **DID YOU KNOW? A Later Valuation, an On-Time Recovery**
>
> From the simulator: in case 4 the deterministic fund ends at 0.994 with no valuation delay, 0.998 with a delay of one year and 1.002 with two, as Table 9.5 prints. The delay does not slow the recovery. It shows in the contribution instead, whose SD rises from 1.91 to 2.35 and 2.93 points.

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

9.6.4 A Recovery the Sponsor Cannot Fund

Case 5 runs the crash on a deterministic fund. The valuation of year 4 still reads 0.999, the one of year 5 reads 0.648 and starts the seven-year clock, and from the decision at year 5 the contribution sits on its 12.5 % cap. The deadline is judged by the valuation of year 12, which reads 0.698 while the fund stands at 0.707, as shown in Figure 9.2(a), and the ratio is 0.738 at year 15. Every amortization period from 3 to 30 years misses the same deadline, ending at 0.7379 for 3 and 5 years and 0.7376 from 7 to 30. At a 30 % cap the rule climbs by its move limit, three points a year, from 12.8 points at year 5 to 25.0 at year 9.04 and 25.4 at year 10, the valuation of year 12 reads 0.819 and clears the

Table 9.6 The Stochastic Cases on the Five Declared Seeds, With Means of the Final Ratio, the Lowest Published Ratio, the Contribution SD and the Mean Contribution, and the Seeds Whose Valuation Fell Below 0.8

Configuration	$F(T)$	F_{pub}^{min}	σ_c	$\bar{c}$	< 0.8
1, MANUAL at 10 %	0.783	0.686	0.00	10.00	3
2, $m = 7$	1.326	0.695	6.11	14.17	3
2, $m = 30$	1.051	0.598	2.92	10.83	3
3, layers	1.277	0.638	5.86	13.44	4
3, amortization	1.168	0.653	4.33	12.28	3
3, PI	1.269	0.653	5.46	13.41	4
6, $e = 0.9$	0.873	0.637	5.47	13.15	4
6, $e = 0.3$	0.826	0.707	6.19	16.35	5
7, glide on	0.882	0.705	4.76	10.78	3
7, glide off	1.503	0.798	4.12	10.22	3
8, holiday	0.871	0.652	6.80	9.46	3
9, MPC, $\kappa = 2$	0.917	0.775	4.12	17.46	3
9, MPC, $\kappa = 0$	0.781	0.673	4.40	13.67	4
9, US 7-yr	0.832	0.686	5.96	14.82	4

Seeds 11, 29, 47, 71 and 101, each run for its case's length.

clock, and the ratio is 0.956 at year 15. Under the Dutch preset the valuation of year 15 reads 0.727, and the cut that brings it to 0.8 reduces the liabilities by 9.12 %, which leaves the fund at 0.812. A deadline sets when the shortfall is counted, and the cap decides whether it can be met.

> **TRY IT IN THE APP** **5 · A Recovery the Sponsor Cannot Fund**
>
> Load the scenario and press Run. Watch the Valuation badge: after the crash of year 3 it still reads 0.999 at year 4 and drops to 0.648 only at year 5. From that decision the contribution sits on its 12.5 % cap, the valuation of year 12 reads 0.698, and the seven-year deadline is missed.

Pause. Try Exercise 9.1 before reading on.

9.6.5 Equity, the Glide Path and the Holiday

Case 6 holds 90 % in equities through the crash. On seed 11 it ends at 1.499 with a lowest valuation of 0.777, while 30 % in equities ends at 0.994 with 0.710 and needs a mean contribution of 14.2 points against 9.2. Once the discount rate has fallen to 0.025 yr, Equation (9.6) at the normal cost lets the ratio rise on its own only above 0.893 with an expected return of 0.028 yr, against 0.481 at 0.052 yr, so the low share needs contributions above the normal cost just to stand still. On seed 29 the order reverses, 0.469 at 90 % against 0.743 at 30 %, with lowest valuations of 0.395 and 0.645. Over the five seeds the high share widens the outcomes, final ratios from 0.351 to 1.499 against 0.658 to 0.994, with four seeds below 0.8 at some valuation against five for the low share. Neither allocation keeps every seed above the threshold, and one favorable seed shows nothing about the others.

Case 7 starts at 1.25 with the glide path on. On seed 11 the share falls to 0.35 at the start, the fund drifts down to 0.796 at year 30 with a mean contribution of 6.7 points, and with the glide path off it ends at 1.088 on 5.1 points, after a surplus that peaks at 2.005 at 14.42 yr while the contribution sits on its floor. The five-seed means are 0.882 against 1.503 at the end, with three seeds below 0.8 either way. The cost of de-risking is expected return, and on these seeds the reduced exposure does not buy a higher lowest valuation, 0.705 against 0.798 on average.

Case 8 imposes a five-year holiday on a fund at 1.25 and a crash at year 5. The contribution is zero until the decision at year 5, the floor then sets it to 5 points at once, and the rule raises it through 5.8, 8.9 and 11.5 to 13.2 points at year 10, the step from 5.8 to 8.9 at the move cap of 3.1 points. Run on the engine without the scripted holiday, the same seed holds 1.905 at year 5 against 1.819, after 0.070 initial liabilities of contributions, and ends at 1.210 against 1.165 with a lowest valuation of

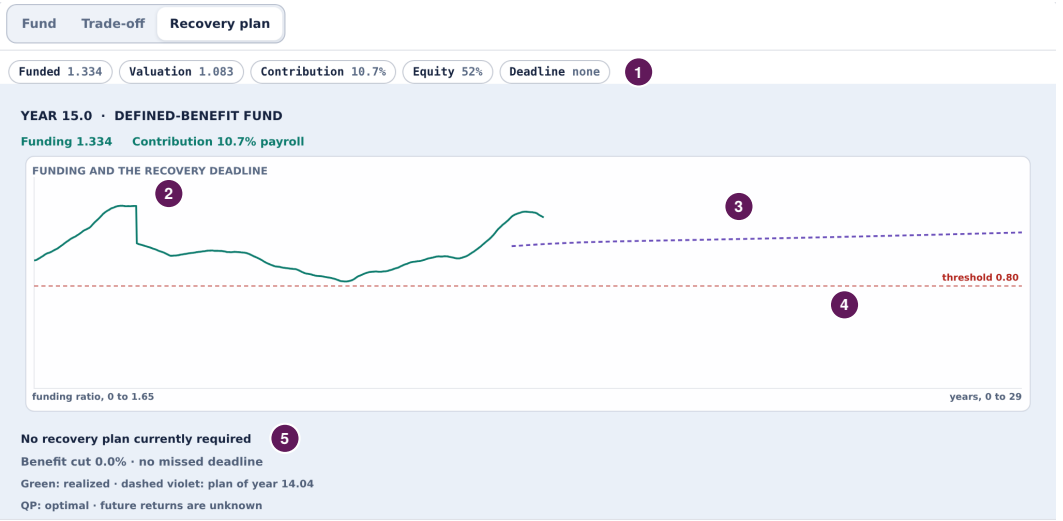

Figure 9.3 Case 9, the MPC with a funding buffer, at $t = 14.98$ yr (the clock is cut, the panel reads year 15.0) in the Recovery plan view, left in AUTO as shipped on return seed 11, cropped to the plant panel. (1) The badges read funded 1.334, valuation 1.083, contribution 10.7 %, equity 52 % and no deadline, so the current and the published ratio differ. (2) The realized ratio, with the crash at year 3, stays above (4) the labeled threshold of 0.80 throughout. (3) The dashed plan is that of the decision at 14.04 yr. It starts from the projected estimate of 1.110 and rises to 1.218 at year 29. (5) The status lines report no recovery plan, no benefit cut and no missed deadline.

0.922 against 0.879. On every one of the five seeds the holiday lowers the final ratio, by 0.016 to 0.093. A surplus spent as a holiday is not there when the crash arrives.

9.6.6 A Plan on an Estimate

Case 9 hands the decision to the MPC with $\kappa = 2$ and the glide path on, as shown in Figure 9.3. Its estimate knows the fall in the discount rate and the revaluation it causes, but not the equity loss. In case 5 with the MPC selected, the estimate reads 0.818 at 3.25 yr against a true 0.774 and a published 1.013. The MPC has held the contribution at the 12.5 % cap since its first decision, because its buffer on the declared $\sigma = 0.16$ sets a target of 1.192 in this deterministic case, while the seven-year rule first moves at year 5. The valuation of year 5 then brings the estimate to the true 0.675. In case 9 on seed 11 the estimate reads 1.153 at 4.5 yr against a true 1.052 and a published 1.423, and over the run it misses the true ratio by a root-mean-square 0.132, and the published valuation by 0.184. With the buffer the MPC ends at 1.331 with a lowest valuation of 0.839 and an SD of 2.85 points, without it at 1.126, 0.739 and 3.52, and the seven-year rule on the same draws at 1.245, 0.773 and 4.40. On the five seeds the buffer raises the lowest valuation from 0.673 to 0.775 on average and the mean contribution from 13.67 to 17.46 points, and no recovery deadline is missed with it, against two without.

> **TRY IT IN THE APP** **9 · A Certainty-Equivalent Recovery Plan**
>
> Run the scenario with the funding buffer of 2, then set Funding buffer to zero and press Reset. With the buffer the MPC ends at 1.331 with a lowest valuation of 0.839. Without it the same draws end at 1.126 with a lowest valuation of 0.739, and the contribution SD rises from 2.85 to 3.52 points.

Pause. Try Exercise 9.3 before reading on.

Table 9.7 The Challenge at Year 15 on the Five Declared Seeds

Seed	Policy	$F(15)$	F_{pub}^{min}	σ_c	total	Target
11	shipped	1.379	0.863	3.91	0.456	SD
29	shipped	0.537	0.478	6.44	0.904	missed
47	shipped	1.286	1.014	3.67	0.359	SD
71	shipped	0.779	0.787	4.55	0.471	floor, end, SD
101	shipped	0.549	0.640	6.75	0.838	missed
11	suggested	1.158	0.954	1.46	0.947	met
29	suggested	1.039	0.940	1.05	0.998	met
47	suggested	1.154	0.980	1.90	0.932	met
71	suggested	1.083	0.976	1.46	0.945	met
101	suggested	1.045	0.941	1.19	0.978	met

Total in L_0. Missed: a deadline, the floor, the end ratio, the SD.

9.6.7 The Challenge on Five Seeds and on a Hundred

The challenge starts a fund at 1.25 with the seven-year rule, a 25 % cap and 20 % of payroll in force, and adds a fall of one point in the discount rate over years 4 to 8 to the crash at year 3. Judged at year 15 on the annual calendar, it asks for every valuation at or above 0.8, a current ratio of at least 1.0, contributions at most 25 % of payroll with an SD of at most 3 points, and no cut. Credit is judged on the case's own seed 11, since a seed, a model parameter or a disturbance changed from the case voids the attempt. There the shipped rule fails on its SD of 3.91 points, 0.91 above the limit, with a lowest valuation of 0.863 and a final ratio of 1.379, and the suggested settings pass with a lowest valuation of 0.954, a final ratio of 1.158, an SD of 1.46 points and a largest contribution of 21.0 %. Table 9.7 gives every declared seed. The shipped rule fails on all five, on its SD alone on seeds 11 and 47, and with missed deadlines and the cap reached on seeds 29 and 101, because a fund at 1.25 lowers its contribution toward the 5 % floor through a crash its valuations do not yet show, reaching 5.0 points at year 4.04 on seed 11, and must then climb back by about three points a year. The suggested settings run the MPC with $F^* = 1.2$, $\kappa = 3$, $w_c = 4$, a glide path from 0.2 to 0.05 in equities over funding of 0.6 to 1.1, and 22 % in force at the start. They pass on all five seeds, ending between 1.039 and 1.158 with lowest valuations between 0.940 and 0.980, by contributing 18.7 to 19.9 points of payroll on average, nearly twice the normal cost, and 0.932 to $0.998\,L_0$ in total against 0.359 to 0.904 for the shipped rule.

Five seeds are a finite sample, so a hundred more were run on the same path. On seeds 1 to 100 the suggested settings pass 98 times, the two failures, seeds 83 and 98, exceeding the SD limit at 3.06 and 3.18 points with every other target met, while the shipped rule passes once, on seed 58. The final ratios of the suggested runs lie between 1.021 and 1.272 and their lowest valuations at or above 0.877.

The comparisons of Table 9.6 keep their direction on seeds 1 to 100. There the thirty-year spread averages a contribution SD of 2.39 points against 5.20 and a lowest valuation of 0.696 against 0.744, and misses a deadline on 39 seeds against 20. Equities at 90 % average a lowest valuation of 0.527 against 0.657 at 30 %, and the MPC's buffer raises it from 0.600 to 0.666 while the seeds with a missed deadline fall from 35 to 17. The layers keep a higher contribution SD than the spread, 4.61 against 3.59 points, at nearly the same average lowest valuation, 0.721 against 0.723. Figure 9.4 draws the five-seed means of Table 9.6 as a chart.

9.6.8 Deciding More Often

The calendar changes little here, because the valuation stays annual. In case 4 the contribution SD is 2.35 points on the annual calendar, 2.42 on the quarterly one with 60 decisions, and 2.46 with a decision at each of 156 executions, and the ratio ends at 0.998 on all three. With the credit χ of Equation (9.8) a finer calendar ramps as fast as the annual one. In case 5 at a 30 % cap the contribution rises over years 6 to 9 at 3.02, 2.99 and 3.02 points a year on the three calendars, and all three meet

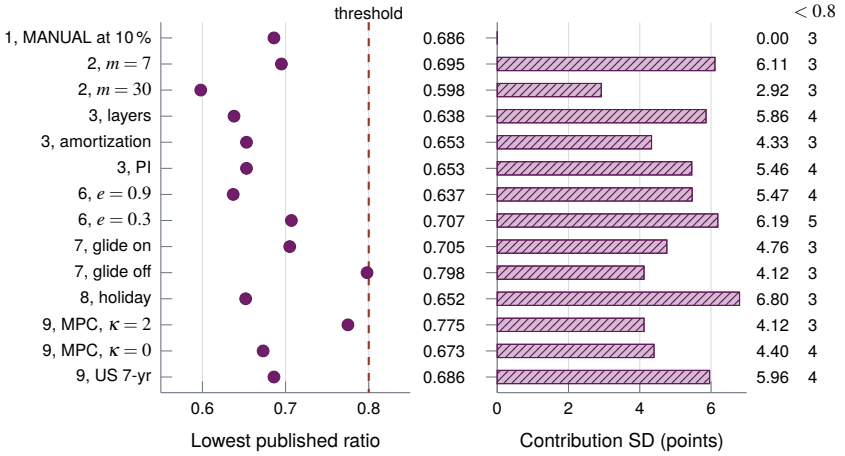

Figure 9.4 The stochastic cases of Table 9.6 as a chart: the five-seed mean of the lowest published ratio (dots, left, against the recovery threshold of 0.8, dashed) and of the contribution SD (hatched bars, right), with the number of seeds whose valuation fell below 0.8 at the far right. Compare rows within one case, because the cases start from different funds and disturbances. Notice that within case 2 the thirty-year spread halves the SD and lowers the lowest valuation, and that within case 9 the buffer raises the lowest valuation from 0.673 to 0.775 at a lower SD.

the deadline of year 12. For the MPC of case 9 a finer calendar lowers the five-seed contribution SD from 4.12 to 3.74 and 3.49 points but also the average lowest valuation, from 0.775 to 0.762 and 0.729, with one deadline missed on each finer calendar, likely because the same ten move blocks then leave the plan free only over its first 2.25 years or 0.87 years. The challenge accepts only the annual calendar.

9.6.9 The Matched Reference

The reference case holds a deterministic fund at the matched point. The ratio stays at 1.000 for fifteen years at the normal cost while both stocks grow at 0.015 yr. Pressed at year 0 before Run, the rates event lowers the ratio to 0.784 at year 5, from which the amortization rule brings it to 1.025 at year 15, and the longevity event leaves it at 0.976.

9.7 Technology in Practice

What Measures It

No instrument reads a funding ratio. The assets have a market value, but the liabilities are a present value that the actuary computes from the promised benefits and a discount rate, and the valuation brings both to one date. In United States law that date, the way assets may be averaged and the yields that discount the liabilities are all written into the minimum funding rules.

> **TECHNOLOGY Sensor: The Actuarial Valuation**
>
> The sensor of a funding loop is the actuarial valuation. Under the 2006 act every determination for a plan year is made as of its valuation date, the first day of the plan year for all but the smallest plans. Assets count at fair market value or at an average of market values reaching back at most to the last day of the 25th month before the valuation month, never below 90 or above 110 percent of market value. Liabilities are discounted at segment rates from a corporate bond yield curve that averages monthly yields over 24 months [9]. Both averages are filters: on a trend the actuarial ratio lags a crash and stays inside its corridor. The valuation of this chapter has no filter and reads market values exactly, a year late.

The valuation of Equation (9.4) stands for this instrument without its filters. It samples both stocks once a year and publishes them a delay θ_v later, and the projected estimate carries the published ratio forward through the recorded flows.

What Moves It

Two handles act on the fund. The contribution is cash from the sponsor, bounded by the sponsor's capacity, and the equity share sets how much of the market return, and of its risk, reaches the assets. In the chapter the contribution passes the quantum, the move limit and the bounds of Equation (9.8) before it reaches the fund, and a cap below the normal cost plus the deficit's spread saturates it, as in case 5.

> **TECHNOLOGY** **Final Element: The Sponsor's Contribution**
>
> The final element is the sponsor's cash contribution. Under the 2006 act the minimum required contribution for a plan year is due $8\frac{1}{2}$ months after the year closes, and a payment made on a date other than the valuation date is adjusted for interest. A plan with a funding shortfall for the preceding plan year pays four required installments, each 25 percent of a required annual payment, due on April 15, July 15, October 15 and January 15 of the following year. An installment paid short is charged interest at the plan's rate plus 5 percentage points [9]. The law thus fixes when cash moves. In the chapter the contribution flows in at the rate in force on every weekly step, and it changes only at the decisions of its calendar.

Where the Controller Runs

The platform of a funding rule is a calendar rather than a computer: the plan year, its valuation date, the actuary's certification and the dates on which contributions fall due. In the application the controller executes every $5/52 = 0.0962\,\mathrm{yr}$, but the shared decision calendar lets the contribution and the equity share move only at its instants, once a year by default. The layer that holds the chapter's structure is that calendar with its shaping, the rules above it and the projected estimate below them. The MPC solves its quadratic program at each decision, from the estimate rather than from the published ratio. Figure 9.5 draws the loop with these devices in place of the schematic's symbols.

What Is Modernizing It

Three developments bear on the late measurement. The first is stochastic projection. The United States insurer of private defined-benefit plans projects its programs with two models that share distributions of investment returns, interest rates and other variables, one running 500 scenarios and the other 5,000, and reports averages and ranges, calling the results estimates, not predictions [17]. The chapter's five declared seeds and its hundred more are the same idea at the scale of one fund. The second is a model corrected by late data: in that report the multiemployer plans' assets reported in the filings for 2022 were higher than the previous year's model had projected, and the new data raised that program's projected net position [17], much as the chapter's estimate restarts from each valuation. The third is a less filtered reading: the 2006 act lets a sponsor determine the minimum contribution from the full corporate bond yield curve without its 24-month averaging, an election revocable only with the consent of the Secretary of the Treasury [9].

> **IN PRACTICE** **Acting on a Ratio Not Yet Certified**
>
> The 2006 act says what a plan assumes while its new ratio is unmeasured. Its benefit limitations key on a funding percentage that the plan's enrolled actuary certifies. If a limitation applied in the preceding plan year, the ratio is presumed unchanged from that year until the actuary certifies the new one. If no certification is made before the first day of the 10th month of the plan year, the ratio is conclusively presumed to be below 60 percent. A plan that stood within 10 points of a limitation is presumed, from the first day of the 4th month, to be 10 points below its preceding-year ratio [9]. The statute holds the last reading, then assumes worse as that reading ages.

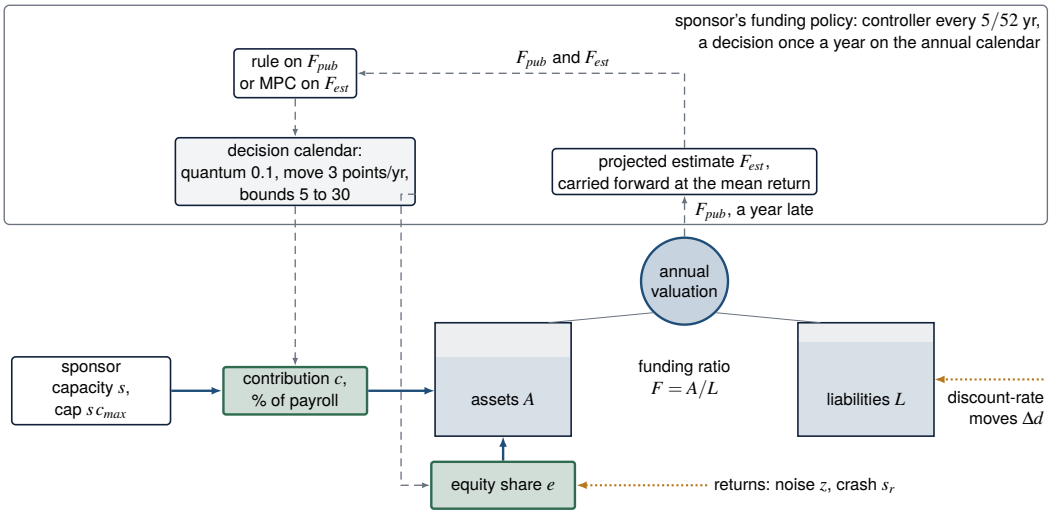

Figure 9.5 The funding loop of Sections 9.3 and 9.5 with a fund's real instruments. Green boxes are the final elements: the sponsor's contribution and the equity share. The blue circle is the measurement, the annual actuarial valuation of both stocks, published a year late. Solid arrows carry money, dashed arrows signals, and dotted arrows the market returns and discount-rate moves that disturb the fund. Notice that the controller runs many times a year, but the contribution and the equity share move only at the decisions of the calendar, and every decision rests on a reading a year old.

9.8 Limitations

The fund is an aggregate with payroll, accruals and benefits proportional to its liabilities, its assets carry no duration, and its returns are colored noise with a one-year correlation, so its five-year risk exceeds that of independent annual returns. The legal presets are simplified. The United States preset spreads the whole deficit over seven years where the 2006 act amortizes each year's shortfall in its own layer, the Dutch preset cuts at the first missed ten-year deadline rather than after five annual determinations below the minimum, and the threshold of 0.8 is a teaching value rather than either text's trigger for a recovery plan. The estimate observes no realized equity return, which keeps the valuation's delay in force for the MPC but discards what a fund knows of its own market value.

> **OPEN CHALLENGE** **A Sponsor Whose Capacity Falls With the Market**
>
> A funding rule has one handle, the contribution, and its size is uncertain. In multiemployer plans, contributions are driven mainly by hours worked and negotiated rates, so a fall in covered work lowers them and can make the minimum requirements hard to meet. Declining work after the 2008 to 2009 financial crisis was a significant factor in the crisis of those plans, and the insurer's projections for fiscal year 2024 still call future contribution levels uncertain and a significant risk [17]. The chapter models the sponsor's capacity only as a slider and a scripted hardship. Progress would be a funding rule judged on paths in which the cap falls with the market that drains the fund.

Every threshold, deadline, cap and challenge target is a teaching value, and five seeds, or a hundred, estimate no probability of success. The application and this chapter give no investment or pension advice.

9.9 Conclusion

This chapter studied a defined-benefit fund written as two stocks whose ratio drifts away at 0.025 yr when left alone and whose valve moves it by 0.0025 a year per point of payroll. Amortization is a

proportional gain with pole $0.025 - 1/m$, and a valuation a year late raises the contribution SD of the deterministic case at $m = 5$ from 1.91 to 2.35 points, while overshoot appears only below 3.85 years. After a crash to 0.648, a 12.5 % cap makes the seven-year deadline impossible for every period, and the Dutch preset closes the gap with a cut of 9.12 %.

On five seeds a thirty-year spread halves the contribution SD of a seven-year one, 2.92 against 6.11 points, and lowers the average lowest valuation from 0.695 to 0.598. The challenge's MPC with a buffer and a glide path passes on all five declared seeds and on 98 of a hundred by contributing nearly twice the normal cost.

The fund is an aggregate with simplified laws. A fuller model would add bonds whose price moves with the discount rate, cohorts with mortality, the statutes' layered bases and multi-year tests, and an estimate that weighs the valuation against the fund's observed market value.

The application turns the trade between the sponsor's volatility and the fund's risk into a measurement on seeded paths.

KEY POINTS

- A pension fund left alone drifts away from full funding, and a point of payroll moves its ratio by only 0.0025 a year.
- A valuation a year late costs contribution variability more than speed of recovery.
- A deadline sets when the shortfall is counted, and the cap decides whether it can be met.
- The projected estimate knows the flows and the discount rate but no realized equity return, so a crash reaches the plan only through the next valuation.

Where else this applies. Volume 2, Chapter 10 steers a public debt through the same annual calendar, and Volume 4, Chapter 8 a cash stock that drains faster than it refills.

References

[1] D. Dufresne, "Movements of pension contributions and fund levels when rates of return are random," *Journal of the Institute of Actuaries*, vol. 115, no. 3, pp. 535–544, 1988.

[2] S. Haberman, "Pension funding with time delays and the optimal spread period," *ASTIN Bulletin*, vol. 25, no. 2, pp. 177–187, 1995.

[3] S. Haberman and J.-H. Sung, "Dynamic approaches to pension funding," *Insurance: Mathematics and Economics*, vol. 15, no. 2–3, pp. 151–162, 1994.

[4] ——, "Dynamic programming approach to pension funding: the case of incomplete state information," *ASTIN Bulletin*, vol. 32, no. 1, pp. 129–142, 2002.

[5] A. Cairns, "Some notes on the dynamics and optimal control of stochastic pension fund models in continuous time," *ASTIN Bulletin*, vol. 30, no. 1, pp. 19–55, 2000.

[6] R. Josa-Fombellida and J. P. Rincón-Zapatero, "Optimal risk management in defined benefit stochastic pension funds," *Insurance: Mathematics and Economics*, vol. 34, no. 3, pp. 489–503, 2004.

[7] 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.

[8] M. I. Owadally and S. Haberman, "Pension fund dynamics and gains/losses due to random rates of investment return," *North American Actuarial Journal*, vol. 3, no. 3, pp. 105–117, 1999.

[9] U.S. Congress, "Pension Protection Act of 2006," Pub. L. 109-280, 120 Stat. 780, Aug. 17, 2006, GPO text PLAW-109publ280, 2006.

[10] "Pensioenwet: Wet van 7 december 2006 houdende regels betreffende pensioenen," wetten.overheid.nl, BWBR0020809, text of 1 Jan.–18 June 2015, Arts. 138, 140, 2006.

[11] S. Haberman, "Pension funding with time delays," *Insurance: Mathematics and Economics*, vol. 11, no. 3, pp. 179–189, 1992.

[12] M. I. Owadally and S. Haberman, "Efficient gain and loss amortization and optimal funding in pension plans," *North American Actuarial Journal*, vol. 8, no. 1, pp. 21–36, 2004.

[13] A. J. G. Cairns, D. Blake, and K. Dowd, "Stochastic lifestyling: Optimal dynamic asset allocation for defined contribution pension plans," *Journal of Economic Dynamics and Control*, vol. 30, no. 5, pp. 843–877, 2006.

[14] E. H. M. Ponds and B. van Riel, "Sharing risk: the Netherlands' new approach to pensions," *Journal of Pension Economics and Finance*, vol. 8, no. 1, pp. 91–105, 2009.

[15] R. Novy-Marx and J. Rauh, "Public pension promises: How big are they and what are they worth?" *The Journal of Finance*, vol. 66, no. 4, pp. 1211–1249, 2011.

[16] J. D. Rauh, "Investment and financing constraints: Evidence from the funding of corporate pension plans," *The Journal of Finance*, vol. 61, no. 1, pp. 33–71, 2006.

[17] Pension Benefit Guaranty Corporation, "FY 2024 projections report," Report, listed on the agency's Projections Report page, last updated January 27, 2026, 2026, accessed 2026-11-21. [Online]. Available: https://www.pbgc.gov/sites/default/files/documents/fy-2024-projections-report.pdf

[18] 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.

9.10 Exercises

9.1 Concept check · 2 min

In case 5 the crash takes a deterministic fund to a funding ratio of 0.648, and the sponsor's cap is 12.5 % of payroll. Every amortization period from 3 to 30 years misses the seven-year deadline and ends near 0.738. Why does the period make almost no difference?

(a) The year-late valuation starts every rule's recovery a year too late

(b) The move limit of three points a year slows a short period as much as a long one

(c) The quantum of 0.1 point rounds the faster rule's larger requests back to the slower one's

(d) Every period asks for more than the cap, so each applies 12.5 % from year 5 on

9.2 Concept check · 1 min

Case 4 starts a deterministic fund at 0.9 with $m = 5$. With the valuation a year late the ratio ends at 0.998 against 0.994 with no delay, and the contribution SD is 2.35 points against 1.91. What does the delay mainly cost here?

(a) Variability of the contribution, with almost no loss in the speed of recovery

(b) Speed of recovery, because every decision acts a year later

(c) An overshoot of the target, because the rule pays on an old deficit

(d) A steady offset, because a proportional rule on old data never reaches the target

9.3 Concept check · 2 min

In case 5 with the MPC selected, the projected estimate reads 0.818 at 3.25 yr while the true ratio is 0.774 and the published valuation 1.013. Why does the estimate sit between the two?

(a) The estimate averages the published valuation with the true ratio it receives each week

(b) The estimate knows the fall in the discount rate and the revaluation, but not the equity loss

(c) The funding buffer of the MPC adds $\kappa e \sigma$ to the estimate it plans from

(d) The estimate integrates the realized return but holds the liabilities at the last valuation

9.4 Theory and calculation · 6 min

Use Equation (9.6) with $P = 0.25$, $n = 10$, $b = 0.05$, an equity share of 0.6, $r_f = 0.016$ yr and $\rho = 0.04$ yr, with no noise.

(a) Find g_L and the open-loop gain $r - g_L$ at the initial discount rate of 0.04 yr, and the time in which a deficit held at the normal cost doubles. *Checkpoint:* 0.015 yr, 0.025 yr and 27.7 yr.

(b) After the crash the discount rate is 0.025 yr. Find the new open-loop gain, the equilibrium ratio at a contribution of 12.5 %, and the rate at which the ratio climbs from 0.648 at that contribution. *Checkpoint:* 0.040 yr, 0.469 and 0.0072 yr.

(c) In case 5 the contribution stays at 10 % until the decision of year 5. Starting from 0.648 at year 4, find the ratio at year 5 and the year at which it reaches 0.8 at the 12.5 % cap. *Checkpoint:* 0.649 and year 20.2.

9.5 Theory and calculation · 6 min

The amortization rule decides once a year on a valuation a year old, and its deviation obeys Equation (9.7) with $\phi = e^{0.025}$ and $\gamma = (e^{0.025} - 1)/0.025$.

(a) Find ϕ and γ, and the roots at $m = 5$ with the delay and the single root without it. *Checkpoint:* 1.0253 and 1.0126, roots 0.758 and 0.267 against 0.823.

(b) Find the spread period below which the roots are complex, and the roots at $m = 3$. *Checkpoint:* 3.85 years, and $0.513 \pm 0.273\,j$.

(c) Relate the result to Equation (1.9). What plays the part of g, and why does the oscillation limit fall at g just above $1/4$ rather than at $1/4$? *Checkpoint:* $g = \gamma/m$. The growth factor $\phi > 1$ moves the limit to $\gamma/m = \phi^2/4$.

9.6 Theory and calculation · 6 min

Case 4 starts at $F_0 = 0.9$ with $m = 5$, a move limit Δ_c of three points a year, a quantum of 0.1 point and the contribution at 10 points.

(a) Find the amortization gain $100/(Pm)$ and the request c^w of Equation (9.9) at the first decision, and the contribution Equation (9.8) allows when a whole year's allowance is available. *Checkpoint:* 80 points per unit of ratio, 18.0 points, and 13.0.

(b) The second decision falls at 1.0577 yr and still sees a published 0.9. Find the allowance, the applied contribution and the credit χ carried forward. *Checkpoint:* 3.17 points, 16.1 points, credit 0.07.

(c) The third decision, at 2.0192 yr, applies 17.6 points without reaching its allowance. What published ratio does that imply? *Checkpoint:* About 0.905.

9.7 App challenge · 8 min

Select scenario 4, "A valuation that arrives late", at the default seed. Speed 1 runs it in about 15 s.

(a) Run 15 yr and read the final funding ratio and the Contribution SD tile of the scorecard. *Checkpoint:* 0.998 and 2.35 points (within 0.01).

(b) In Show Model set Valuation delay to 0, Apply, press Reset and run 15 yr again. Read the same two values. *Checkpoint:* 0.994 and 1.91 points.

(c) Repeat with Valuation delay 2. Read the SD, and explain from Equation (9.7) why the delay raises the SD more than it changes the final ratio. *Checkpoint:* 2.93 points.

9.8 App challenge · 8 min

Select scenario 5, "A recovery the sponsor cannot fund", at the default seed, in the Recovery plan view. Speed 1 runs it in about 15 s.

(a) Run 15 yr. Read the valuation of year 12 from the Published trend, and the final funding ratio. *Checkpoint:* 0.698 and 0.738 (within 0.003).

(b) Set Sponsor contribution cap to 30 %, press Reset and run 15 yr. Read the valuation of year 12, the final ratio and the largest contribution on the Contribution trend. *Checkpoint:* 0.819, 0.956 and 25.4 %.

(c) Set the cap back to 12.5 %, select Dutch-style recovery, press Reset and run until the run pauses at year 15. Read the Benefit cuts tile and the final ratio. *Checkpoint:* 9.1 % and 0.812.

9.9 App challenge · 8 min

Select scenario 9, "A certainty-equivalent recovery plan", at the default seed 11. Speed 1 runs it in about 15 s.

(a) Run 15 yr as shipped, with a Funding buffer of 2. Read the Lowest valuation and Contribution SD tiles and the final funding ratio. *Checkpoint:* 0.839, 2.85 points and 1.331 (within 0.01).

(b) Set Funding buffer to 0, press Reset and run 15 yr. Read the same three values. *Checkpoint:* 0.739, 3.52 points and 1.126.

(c) Set the buffer back to 2, select Historical US 7-year as the rule, press Reset and run 15 yr. Compare the three rules on these draws and say what the five declared seeds add. *Checkpoint:* 0.773, 4.40 points and 1.245.

9.10 Application area · 3 min

The fund holds no bonds, so a fall of 1.5 points in the discount rate revalues only the liabilities, through $D_L = 15\,\text{yr}$.

(a) Find the factor on L and the funding ratio of a fund at 1.0 after the revaluation alone, and after it together with the crash's equity factor of 0.807. *Checkpoint:* 1.252, 0.799 and 0.645.

(b) Outside the model, suppose 40 % of the assets were bonds whose value rose by the same factor and the rest did not move. Find the ratio after the rate fall alone, and name the parameter of Table 9.2 that this hedging would effectively change. *Checkpoint:* 0.879.

9.11 Application area · 4 min

A board asks the sponsor to restore a fund from 0.648 to full funding within seven years by contributions alone. The payroll is $0.25L$ a year and the normal cost 10 points.

(a) Ignoring the fund's growth and returns, how many points of payroll above the normal cost would that take each year, and what total contribution rate? *Checkpoint:* About 20.1 points, a rate of about 30.1 %.

(b) Compare with the case 5 cap of 12.5 %, which raises the ratio by at most 0.00625 a year, and with the application's largest cap. Name one cost to the sponsor that the model leaves out. *Checkpoint:* About 56 years at the cap, and above the 30 % slider maximum.

9.12 Review the agent · 5 min

An AI assistant reviews case 2 for a pension board. It writes, "On return seed 11 a thirty-year amortization period ends better funded than a seven-year one, 0.831 against 0.809, with a lowest valuation of 0.901 against 0.864 and half the contribution SD, 1.86 against 3.74 points. A long period is therefore both smoother and safer, and the board should adopt $m = 30$." Find the flaw, and say how you would check the claim in the application.

CHAPTER 10

Inferential Moisture Control Near the Sticky Point of a Milk Powder Spray Dryer

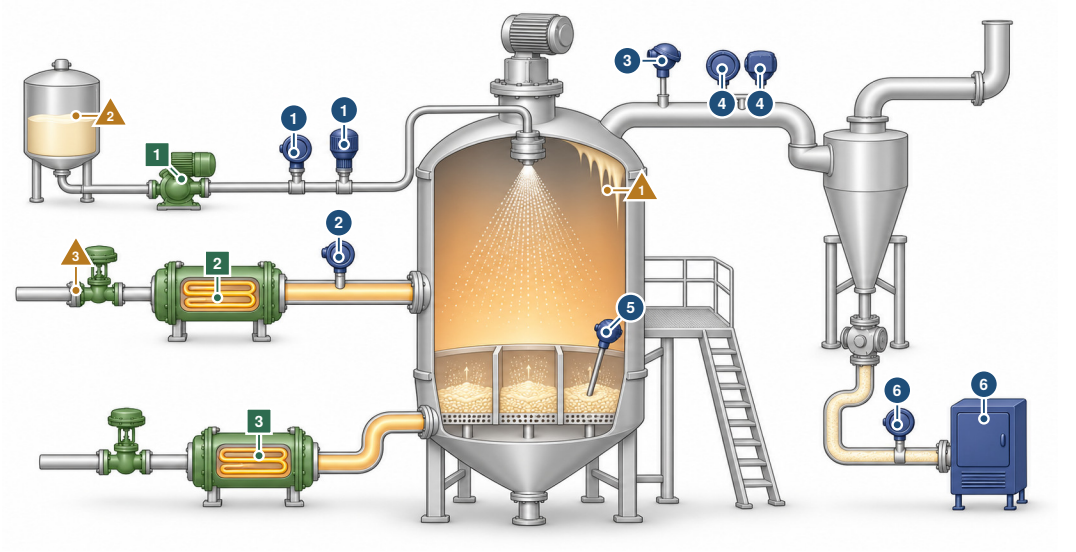

The controller moves 1 feed pump for the milk concentrate 2 heater of the drying air 3 heater of the fluid bed air

The controller reads 1 feed flow and feed solids 2 inlet air temperature
3 exhaust temperature, held by the main loop 4 exhaust humidity and carbon monoxide 5 fluid bed temperature
6 powder sampled for the laboratory

What disturbs it 1 deposit on the hot roof 2 weaker concentrate from the evaporator 3 humid intake air

The exhaust temperature (3) answers within seconds, while the moisture of the powder reaches the laboratory (6) most of an hour later. Illustration, not to scale.

IN PLAIN TERMS Drying Milk by a Stand-In Temperature

The main idea. A spray dryer sprays thick milk into hot air, and the drops dry into powder in seconds. The plant steers by the temperature of the air leaving the dryer, because that reading comes at once.

The challenge. What matters is the water left in the powder, and the laboratory reports it most of an hour late. When the milk or the weather changes, the same air temperature means wetter powder, and powder that is too wet and hot sticks to the walls.

Why it matters. Sticky powder piles up on the hot roof of the dryer, where it can catch fire.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the fact on the moisture rule (Section 10.1), the two history segments (Section 10.2), the try-it on the weaker concentrate and the try-it on the laboratory correction (Section 10.6), the fact on the laboratory alone 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 10.6), Technology in Practice (Section 10.7), Limitations (Section 10.8) and the open challenge

The control question A spray dryer is steered by an exhaust temperature that answers in seconds and judged on a powder moisture that its laboratory reports most of an hour later, while its throughput is bounded by a sticky point that no instrument measures. When does the outlet temperature stand for the powder, and what do a soft sensor, a late laboratory, feedforward and overrides recover when it does not?

Learning objectives

- Explain why an outlet temperature held on its setpoint stops standing for the powder moisture when the concentrate or the intake air changes, and read the feedforward gains from the steady map.
- Compute the glass transition and the sticky margin of the particles leaving the chamber from the Gordon–Taylor rule.
- Trace how a laboratory result, sampled every 30 min and reported 40 min later, corrects a biased soft sensor, and what the bias leaves uncorrected.
- Compare feedforward, moisture feedback, overrides and a throughput optimizer on the campaign's tonnage, sticky margin and energy.
- Judge whether a model predictive controller's gain comes from its model or from its speed.

Prerequisites Sections 1.4.5, 1.4.7, 1.4.8, 1.6.2 and 1.6.4, with 1.6.3 for case 9.

The app https://apmonitor.com/apps/spray-dryer/

See also Volume 4, Chapter 6 for another dryer whose throughput a constraint bounds, Volume 4, Chapter 5 for a food product judged downstream under overrides, Volume 4, Chapter 4 for throughput pushed against limits measured late, and Volume 9, Chapter 4 for offset-free predictive control with an estimated bias.

10.1 Introduction

Spray drying is the preferred way to take the water out of many liquid foods, and it leaves a free-flowing powder [1]. The main difficulty in controlling the process is to meet the residual moisture specification while keeping the powder from sticking to the chamber walls [2]. In the two-stage dryer simulated here, a milk concentrate from an evaporator is atomized into hot air at the top of a chamber, dries in seconds and is finished in a fluid bed below. At the design point 12 000 kg/h of air heated to 190 °C turns 1001 kg/h of skim-milk concentrate at 48 % solids into 498 kg/h of powder at 3.50 % moisture, with the exhaust at 80.0 °C and the particles leaving the chamber 2.93 K inside their sticky limit.

DID YOU KNOW? Five Percent by Federal Rule

United States rules for dairy plants define nonfat dry milk as pasteurized skim milk with only its water removed, holding not more than 5 percent moisture and not more than 1.5 percent milkfat by weight. The powder must be cooled to no more than 110 °F before packaging [3]. The 3.5 % ± 0.3 % specification of this chapter sits inside that ceiling.

The dryer measures its air and sells its powder. The exhaust temperature falls 5.9 K within 2 min of a 10 % step in the feed, while the powder leaving the fluid bed takes 4.0 min to cross its 3.8 % upper limit, and the laboratory, which samples every 30 min and reports each sample 40 min after it was drawn, first reports the change 65 min after the step. The moisture is controlled by holding the outlet temperature with the feed, and the outlet stands in for the moisture for as long as the concentrate and the intake air stay as they were. When the evaporator upstream sends 44 % solids instead of 48 %, the outlet loop keeps the exhaust within 0.39 K of 80 °C while the powder rises to 3.85 %. When the afternoon intake climbs from 8 toward 18 g/kg of water per kilogram of dry air, the powder reaches 4.10 % and the particles leaving the chamber cross their sticky line.

This chapter sets out to dislodge the habit of reading an outlet temperature on its setpoint as powder in specification. Holding a secondary variable in place of an unmeasured quality is the plainest form of inferential control [4] (Section 1.4.8), and the weakness it carries is general. For a simulated industrial debutanizer whose conventional policy held one stage temperature at its setpoint,

a static estimator built from four to six tray temperatures improved the steady-state control of the bottoms quality by as much as 400 % under feed composition disturbances [5]. The structure appears in any plant in which a fast measurement stands for a slow analysis, which is why that study posed the estimator as a general design problem. The dryer adds a second inferred quantity to the first. Its throughput is limited by the stickiness of the particles leaving the chamber, which no instrument measures, and a hotter inlet that buys more powder moves the particles toward that limit while every temperature loop holds its setpoint. A transfer function from feed to exhaust temperature carries neither the moisture nor the margin.

Control studies of spray dryers compare linear and economically optimizing model predictive control with conventional PID and PI strategies on first-principles models [1, 2], and control in drying as a whole has been reviewed across 71 publications [6]. This chapter works with a plant on which the outlet loop can be watched holding its setpoint while the powder leaves its specification, and on which the soft sensor, the laboratory correction and the constraint handling that repair the fault are then added one at a time. The plant is a two-stage milk powder dryer that keeps the true moisture, the true sticky margin and the true roof deposit apart from what its controllers estimate, and closes its water, solids and energy balances to round-off. Psychrometrics, drying curves, the glass transition, soft sensors, bias updating, feedforward and override control are all established practice. The model has 33 states with IAPWS-IF97 saturation pressures, the soft sensor solves the plant's own drying equations at the measured air states, the laboratory is the framework's shared sampled analyzer, and the deposit ignites by a declared rule, all described completely enough to be rebuilt. Nine cases and a graded campaign separate a proxy from the quality it stands for, a model from its bias, and a constraint from the loops that approach it.

Section 10.2 follows the domain from energy surveys to predictive control. Section 10.3 sets out the model, and Section 10.4 reads from its steady states what the outlet temperature stands for. Section 10.5 describes the controller: the PI loops, the soft sensor and its laboratory bias, the feedforward, the overrides, the throughput optimizer and an optional model predictive controller. Section 10.6 reports what the nine cases and the campaign measure, Section 10.7 sets the loops on plant hardware, Section 10.8 states the limitations, Section 10.9 draws the conclusions, and the exercises of Section 10.10 work through the cases.

10.2 Energy, Stickiness, Deposits and the Outlet Loop

Spray dryers burn a great deal of fuel. A survey commissioned in 2000 by the United Kingdom government's energy efficiency program found specific energy consumptions from around 3 to 20 GJ per tonne of water evaporated among spray dryers in the chemicals, foods and ceramics industries, with an average of 4.87 GJ/t, and estimated that about 29 % of the energy supplied to them was wasted [7]. The design of fluid bed and spray dryers for low energy consumption had been treated as one subject in 1984 [8], and the simulated plant pairs the two, a chamber that leaves the powder slightly wet and a bed that finishes it. At its design point it spends 4175 kJ of heat per kilogram of water evaporated, 4.18 GJ/t.

> ⧖ HISTORY **1872: A Patent That Led to Milk Powder**
>
> Spray drying turns a liquid into dry particles by spraying it into a hot gas. Its first observation dates from 1860, and Samuel Percy patented a primitive spray dryer in the United States in 1872. The early devices lacked efficiency and safety. Once those problems were overcome the method suited the food industry, and by the 1920s it was making milk powder, which has remained one of its most important uses. The Second World War, with its need to cut the weight and volume of the food to be carried, pushed the method further, and after the war it reached the pharmaceutical, chemical, ceramic and polymer industries [9].

The drying law in the model descends from the characteristic drying curve, a normalized rate that equals one while the particle surface is wet and falls with the free moisture below a critical value,

whose basis was examined in 1974 [10]. A characteristic drying curve for milk powder was assessed in 2001 for use in computational fluid dynamics modeling [11].

Stickiness has a thermodynamic explanation. Water plasticizes amorphous sugars, and in dehydrated sugar solutions used as models of amorphous foods the glass transition, crystallization and melting temperatures all fell as the moisture rose, with the glass transition of a sucrose and fructose model slightly below the empirical sticky point at every moisture studied [12]. The mixing rule now used to predict the glass transition of a sugar and water mixture from its two components was derived in 1952 for binary copolymers [13]. For skim milk powder, the sticky point measured by the cohesive force between stirred particles in a heated flask followed the predicted glass transition of lactose shifted up by 23.3 K, between 25 and 95 °C [14]. Sticky powder is what stays on the walls. In a pilot-scale co-current dryer the wall deposition flux of skim milk powder was highest, 16 g/(m^2 h), when the outlet particle temperature was furthest above the sticky-point temperature, which happened at the lowest of the three inlet temperatures tested, 170 °C, and the highest product moisture [15].

A deposit is also fuel. An investigation of a series of dryer fires, told in the history segment below, found that such deposits could ignite by themselves [16]. The simulator's rule, a layer of 3 mm on a roof held at 158 °C or more for 20 min, belongs to that line of reasoning but is a teaching value, not a calculated one. The application names the sources of real thresholds, layer-ignition tests to ISO/IEC 80079-20-2 and the dryer guidance of NFPA 61 and VDI 2263.

> ⧖ HISTORY **1984: Fires Traced to the Dryer's Own Deposits**
>
> A company operating a spray dryer for skimmed milk and fat-filled powders asked the Fire Research Station to find the cause of a series of fires. The investigation, reported in 1984, showed that the powders could ignite spontaneously in certain circumstances. It calculated the thickness of powder deposit on the dryer walls needed for a risk of self-ignition and found it close enough to the thicknesses met in practice to suggest that self-ignition was a real problem. It also put forward how fires could start in fluid-bed instantizers and suggested preventive measures [16].

Control came late to drying. A review of 71 publications on control in drying found that control tools began to appear in drying applications only in 1979 and that more advanced concepts followed from about 1998 [6]. For milk powder, a review of modern process control techniques in 2005 described advanced control under development and named models for stickiness and thermodynamic models of spray drying among the developments to come [17]. Predictive control followed. Against a conventional PID strategy, a linear model predictive controller on a two-stage spray dryer raised the production rate by 7.9 %, the profit by 8.2 % and the energy efficiency by 4.1 % over an industrially recorded disturbance scenario [2], and an economically optimizing nonlinear controller improved the cost of drying by 26.7 % against conventional PI control in simulation [1]. The simulator's shipped structure stays between those two ends, with conventional loops, a soft sensor under laboratory correction, feedforward, overrides and a constraint pusher, which is the ground a student covers before a predictive controller can be judged against it, and an optional linear predictive controller, Section 10.5.6, is judged against it in case 9.

10.3 The Dryer as Modeled

The simulated plant dries skim-milk concentrate in a chamber with a rotary atomizer and finishes it in an external fluid bed, as shown in Figure 10.1. The chamber is two well-mixed cells of air, the upper one against the roof around the hot-air disperser and the lower one against the cylinder wall, and each carries one well-mixed population of particles with its own residence. The bed has a drying section blown with heated air and a cooling section blown with room air, and the powder crosses a weir between them. Moisture is written X on a dry basis (kilograms of water per kilogram of solids) and $w = X/(1+X)$ on a wet basis, and the percentages quoted for the powder are wet basis.

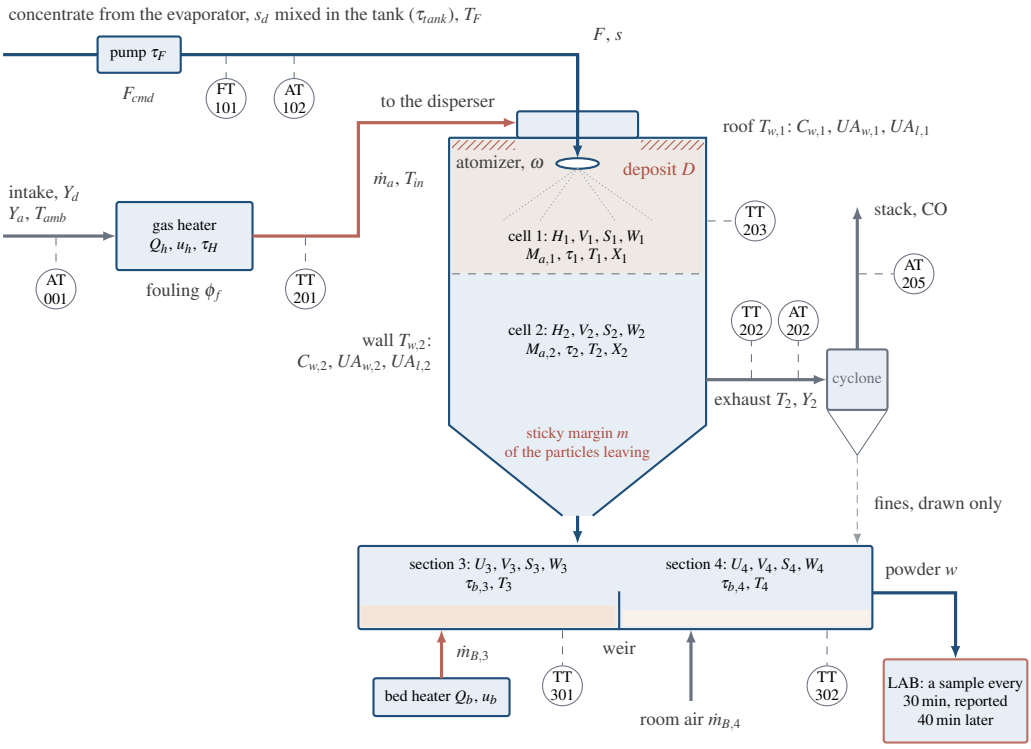

Figure 10.1 Two-stage milk powder dryer as modeled. Intake air at T_{amb} and Y_a passes the gas heater (Q_h, output u_h) and TT-201 and enters the roof through the disperser at $\dot{m}_a$ and T_{in}. Concentrate of solids s, mixed from the evaporator's s_d in a balance tank, is pumped at F past FT-101 and the densitometer AT-102 to the atomizer. Chamber cells 1 and 2 hold air enthalpy H and vapor V and particle solids S and water W, the roof $T_{w,1}$ carries the deposit D, and the exhaust passes TT-202, AT-202, a cyclone that is drawn but not modeled, and the CO analyzer AT-205. Bed sections 3 and 4, joined by a weir, are fed with heated air $\dot{m}_{B,3}$ from the bed heater (Q_b, u_b) and room air $\dot{m}_{B,4}$, and the laboratory samples the powder at the discharge. Dashed lines are measurements.

Table 10.1 gives the controlled, manipulated and disturbance variables beside the 33 states. Each chamber cell carries the enthalpy H_i and water vapor V_i of its air, the solids S_i and water W_i of its particles and one wall temperature, and each bed section carries its energy U_j, the vapor V_j in the air above it and the solids and water of its powder. Six more states hold the two heater duties, the feed flow, the feed solids, the intake humidity and the roof deposit, and the last nine are running totals of the water, solids and energy that enter and leave, the heater energy and the powder made. Because the totals advance with the plant in every integration step, the balances can be checked against the state vector at any instant.

10.3.1 Humid Air

Enthalpies are per kilogram of dry air and referred to dry air and liquid water at 0 °C. The psychrometric relations shared by the plant, the soft sensor and the scene are Equation (10.1).

$$\begin{aligned} h(T,Y) &= c_a T + Y(\lambda_0 + c_v T), \qquad T = \frac{h - Y\lambda_0}{c_a + Y c_v} \\ Y &= \frac{\varepsilon\, p_v}{P - p_v}, \qquad \varphi = \frac{p_v}{p_{sat}(T)} \\ h(T,Y) &+ \big[Y_s(T_{wb}) - Y\big] c_w T_{wb} = h\big(T_{wb}, Y_s(T_{wb})\big) \end{aligned} \tag{10.1}$$

Table 10.1 What Each Variable Does in the Loops, and Its Design-Point Value

Variable	Role	Design value	Range
T_2 (TT-202)	CV, TIC-202	80.0 °C SP	65–100
T_{in} (TT-201)	CV, TIC-201	190 °C SP	150–240
T_3 (TT-301)	CV, TIC-301	55.0 °C SP	40–90
$\hat{w}$ (AY-401)	CV, AIC-400	3.50 % SP	2.5–4.5
$\hat{m}$ (AY-401)	constraint, UIC-205	2.93 K, SP 0.5	0–6
Y_{202} (AT-202)	constraint, AIC-206	49.6 g/kg, 58 max	40–90
F_{cmd}	MV, feed pump	1001 kg/h	0–1500
u_h, u_b	MVs, heaters	74.9, 27.4 %	0–100
$T_{out,SP}$	MV of AIC-400	80.0 °C	65–100
$T_{in,SP}$	MV of FY-500, overrides	190 °C	150–240
s_d	DV, evaporator	48 %	44–52
Y_d	DV, weather	8 g/kg	4–22
T_F, ϕ_f	DVs, feed, fouling	50 °C, 0	35–70, 0–50 %
ω, t_{lab}	DVs, wear, lab	0, 40 min	0–25 %, 20–90
H_i, V_i, S_i, W_i	states, cells 1, 2	see text	—
$T_{w,1}$, $T_{w,2}$	states, roof, wall	147.4, 66.0 °C	—
U_j, V_j, S_j, W_j	states, sections 3, 4	see text	—
Q_h, Q_b	states, heaters	561.5, 22.0 kW	—
F, s, Y_a, D	states	1001, 48, 8, 0	—
w, m	truth, powder, margin	3.50 %, 2.93 K	—

Here Y is the humidity ratio, p_v the vapor pressure, φ the relative humidity, Y_s the saturation humidity and T_{wb} the adiabatic-saturation (wet-bulb) temperature, with $c_a = 1006$, $c_v = 1860$ and $c_w = 4186\,\mathrm{J/(kg\,K)}$, $\lambda_0 = 2.501\,\mathrm{MJ/kg}$, $P = 101\,325\,\mathrm{Pa}$ and $\varepsilon = 0.621945$. The saturation pressure is the IAPWS-IF97 region-4 equation [18], and the wet bulb is found by a Newton iteration kept inside a bracket. The application's checks put p_{sat} within a relative 4.4×10^{-5} of 2.3392 kPa at 20 °C and 101.418 kPa at 100 °C, give air at 30 °C and 50 % relative humidity 13.31 g/kg of water and a wet bulb of 22.0 °C, and agree with a bisection for the wet bulb to 1.6×10^{-12} K from 0 to 250 °C.

10.3.2 The Chamber Cells

Each chamber cell $i = 1, 2$ holds a fixed mass of dry air $M_{a,i}$, so its temperature follows from its two air states through Equation (10.1), and its balances are Equation (10.2).

$$\begin{aligned} \frac{\mathrm{d}H_i}{\mathrm{d}t} &= \dot{m}_a(h_{i-1} - h_i) + P_{i-1} - P_i - Q_{w,i} \\ \frac{\mathrm{d}V_i}{\mathrm{d}t} &= \dot{m}_a(Y_{i-1} - Y_i) + E_i \\ \frac{\mathrm{d}S_i}{\mathrm{d}t} &= \frac{S_{i-1}}{\tau_{i-1}} - \frac{S_i}{\tau_i}, \qquad \frac{\mathrm{d}W_i}{\mathrm{d}t} = \frac{W_{i-1}}{\tau_{i-1}} - \frac{W_i}{\tau_i} - E_i \\ C_{w,i}\frac{\mathrm{d}T_{w,i}}{\mathrm{d}t} &= Q_{w,i} - UA_{l,i}(T_{w,i} - T_{amb}) \end{aligned} \tag{10.2}$$

where $h_i = H_i/M_{a,i}$ and $Y_i = V_i/M_{a,i}$, the particles leave cell i with the enthalpy $P_i = (S_i c_s + W_i c_w)T_{p,i}/\tau_i$ at their own temperature $T_{p,i}$, and E_i is the water they evaporate. The inlet to cell 1 is the heated air, $h_0 = h(T_{amb}, Y_a) + Q_h/\dot{m}_a$ and $Y_0 = Y_a$, and the feed, $S_0/\tau_0 = \dot{s} = Fs$, $W_0/\tau_0 = F(1-s)$ and $P_0 = [\dot{s}c_s + F(1-s)c_w]T_F$, with $c_s = 1500\,\mathrm{J/(kg\,K)}$ for milk solids. The roof exchanges heat with a film at the mean of the inlet and cell temperatures, $Q_{w,1} = UA_{w,1}[(T_{in} + T_1)/2 - T_{w,1}]$, and the cylinder wall with the lower cell, $Q_{w,2} = UA_{w,2}(T_2 - T_{w,2})$. Evaporation adds vapor to the air without adding enthalpy, and that is the whole mechanism by which the spray cools it. The air holdups of 60 kg and 90 kg are 18 and 27 s of flow.

10.3.3 The Drying Curve and the Equilibrium Moisture

The particles in a cell dry at the rate $E_i = S_i k (T_i - T_{wb,i}) f(X_i)$, with $X_i = W_i / S_i$ and $T_{wb,i}$ the wet bulb of the cell's air, through the illustrative characteristic drying curve of Equation (10.3).

$$
\begin{aligned}
f(X) &= \begin{cases} 1, & X \geq X_{cr} \\ \left(\dfrac{X - X_e}{X_{cr} - X_e} \right)^n, & X_e < X < X_{cr} \\ 0, & X \leq X_e \end{cases} \\
T_p &= T_{wb} + (T - T_{wb})(1 - f) \\
X_e &= \frac{X_m C K a_w}{(1 - K a_w)(1 - K a_w + C K a_w)}
\end{aligned}
\tag{10.3}
$$

Here $C = C_0 \mathrm{e}^{H_C / T^*}$ and $K = K_0 \mathrm{e}^{H_K / T^*}$ with T^* the absolute temperature, and a_w is the relative humidity of the cell's air, capped at 0.98. The particle temperature is quasi-steady. It sits at the wet bulb while the surface is wet and rises toward the air as the rate falls, and the heat a chamber particle stores is neglected because the cells hold 3.3 kg of solids against 150 kg of air, although the enthalpy the particles carry is counted. The equilibrium moisture X_e is a GAB isotherm with illustrative constants for skim milk, 5.59 % on a dry basis at 25 °C and a water activity of 0.30 and 2.68 % at 80 °C and 0.17. Atomizer wear ω, in percent of droplet diameter, divides the chamber constant by $(1 + \omega/100)^2$, because the time to evaporate a droplet grows with the square of its size.

10.3.4 The Fluid Bed

Each bed section $j = 3, 4$ is well mixed. Its powder, its steel and the air above it share one temperature, which follows from the section's energy, and its balances are Equation (10.4).

$$
\begin{aligned}
T_j &= \frac{U_j - V_j \lambda_0}{C_b + S_j c_s + W_j c_w + M_{fb} c_a + V_j c_v} \\
\frac{\mathrm{d}U_j}{\mathrm{d}t} &= \dot{m}_{B,j} \left[h_{B,j} - h(T_j, Y_j) \right] + P_{j-1} - P_j \\
&\quad - U A_b (T_j - T_{amb}), \qquad Y_j = V_j / M_{fb} \\
\frac{\mathrm{d}V_j}{\mathrm{d}t} &= \dot{m}_{B,j} (Y_a - Y_j) + E_j \\
\frac{\mathrm{d}S_j}{\mathrm{d}t} &= \sigma_{j-1} - \sigma_j, \qquad \frac{\mathrm{d}W_j}{\mathrm{d}t} = \sigma_{j-1} X_{j-1} - \sigma_j X_j - E_j \\
\sigma_j &= \dot{s}_D \left(\frac{S_j}{\dot{s}_D \, \tau_{b,j}} \right)^{1.5}, \qquad P_j = \sigma_j (c_s + X_j c_w) T_j
\end{aligned}
\tag{10.4}
$$

where the chamber discharge enters section 3, $\sigma_2 = S_2 / \tau_2$ and $\sigma_2 X_2 = W_2 / \tau_2$ with the enthalpy P_2 of Equation (10.2), the drying section is blown with heated air, $h_{B,3} = h(T_{amb}, Y_a) + Q_b / \dot{m}_{B,3}$, and the cooling section with room air, $h_{B,4} = h(T_{amb}, Y_a)$. The powder leaves each section over a weir whose outflow grows as the holdup to the power 1.5, scaled by the design solids flow $\dot{s}_D = 480\,\mathrm{kg/h}$ and the design residences $\tau_{b,3}$ and $\tau_{b,4}$, so the residence rises when the rate falls. It dries on the curve of Equation (10.3) with the much smaller constant k_{bed} of internal diffusion and at the bed temperature, $E_j = S_j k_{bed} (T_j - T_{wb,j}) f(X_j)$.

10.3.5 Stickiness, the Roof Deposit and the Fire

Stickiness is judged on the particles leaving the chamber, as shown in Equation (10.5).

$$\begin{aligned} T_g(w) &= \frac{(1-w)T_{g,s} + k_{GT}\, w\, T_{g,w}}{(1-w) + k_{GT}\, w} \\ m &= \Delta T_{st} - \left[T_{p,2} - T_g(w_2)\right] \\ \frac{\mathrm{d}D}{\mathrm{d}t} &= r_{dep} \min\left[\max(0,-m),5\right] L - r_{clean}\frac{D}{D+0.05} \end{aligned} \tag{10.5}$$

The first line is the Gordon–Taylor rule [13] for amorphous lactose, $T_{g,s} = 101\,^{\circ}\mathrm{C}$, $T_{g,w} = -135\,^{\circ}\mathrm{C}$ and $k_{GT} = 6.7$, which gives 54.9 °C at 3.5 %. The margin m is how far the particles stay below the sticky line $T_g + \Delta T_{st}$, with $\Delta T_{st} = 20\,\mathrm{K}$, so it is positive while they are not sticky, and it is set to ΔT_{st} when the chamber carries less than 1 % of its design load. Sticky fines are taken to recirculate to the roof around the disperser, so the deposit D, in millimeters, grows only while $m < 0$, by up to 5 K of stickiness at $r_{dep} = 2\,\mathrm{mm}$ per kelvin-hour and the load $L = \min(1.5, \sigma_2/\dot{s}_D)$, and air hammers remove it at up to $r_{clean} = 0.3\,\mathrm{mm/h}$. A layer of at least $D_{\mathrm{fire}} = 3\,\mathrm{mm}$ on a roof at $T_{w,1} \geq T_{\mathrm{fire}} = 158\,^{\circ}\mathrm{C}$ held without a break for $t_{\mathrm{fire}} = 20\,\mathrm{min}$ ignites, and an interruption restarts the clock. After ignition the exhaust carbon monoxide rises as $\min[5000, \mathrm{e}^{k_{CO}(t - t_{ign})}]$ ppm with $k_{CO} = 0.5\,\mathrm{min}^{-1}$. The interlock reads the analyzer AT-205, raises an alarm at 15 ppm and shuts the dryer down at 40 ppm, the feed and heaters go to zero, and the run ends a minute later. The sticky criterion, the deposit law and the fire rule are all teaching values.

10.3.6 Actuators, Boundary Conditions and the Product

The heaters, the pump and the two slow boundary conditions all follow their targets to first order, in Equation (10.6).

$$\begin{aligned} \tau_H \dot{Q}_h &= (1-\phi_f)\, u_h Q_{h,max}/100 - Q_h \\ \tau_{Hb} \dot{Q}_b &= u_b Q_{b,max}/100 - Q_b \\ \tau_F \dot{F} &= \mathrm{clamp}(F_{cmd}, 0, F_{max}) - F \\ \tau_{tank}\dot{s} &= s_d - s, \qquad \tau_{amb}\dot{Y}_a = Y_d - Y_a \end{aligned} \tag{10.6}$$

Here ϕ_f is the fouled share of the gas heater, s_d the solids the evaporator sends and Y_d the humidity of the weather. Every kilogram of powder leaving the cooling section, at the rate $\sigma_4(1+X_4)$, is booked in specification while its true moisture is within $3.5\,\% \pm 0.3\,\%$, and off specification otherwise. The specific energy is the heater energy per kilogram of powder, $(Q_h + Q_b)/[\sigma_4(1+X_4)]$ now and $\int (Q_h + Q_b)\,\mathrm{d}t$ over the powder made since reset for the campaign.

10.3.7 Parameters, Initial State and Integration

Table 10.2 gives every parameter. Twenty-eight are shown in the Show Model panel and may be moved while the dryer runs, and 24 are hidden. The declared design point is 1000 kg/h of concentrate at 48 % solids, 190 °C at the inlet, 80 °C at the outlet and 3.5 % in the powder, and the illustrative drying constants were chosen so that the model reproduces it. Every scenario starts from a steady state computed at reset for the configuration in force, its own plus any setting the learner carries through a Reset. The solver finds each chamber cell's temperature and particle moisture by nested Illinois false-position roots, the feed that puts the outlet on its setpoint by a third, and each bed section's temperature and moisture the same way, and with the moisture controller on it also finds the outlet setpoint at which the controller's own reading equals its setpoint. At the design point it returns a feed of 1000.94 kg/h, a gas heater at 74.865 % of its rating and a bed heater at 27.442 %, and the largest relative derivative outside the ledgers at the stored state is 2.5×10^{-15}.

The plant is integrated by the classical fourth-order Runge–Kutta method at a fixed step of 0.5 s, the controllers run every 2 s, and a row goes to the historian every 5 s, so the eight-hour campaign fills 5760 rows. The fastest state is the particle moisture in the lower cell, which relaxes at $1.28\,\mathrm{s}^{-1}$ in

Table 10.2 Parameters of the Dryer. V Is a Value the Show Model Panel Displays and Lets the User Move, H One It Keeps Hidden

Symbol	Meaning	Value		Symbol	Meaning	Value	
$\dot{m}_a$	drying air, dry basis	12 000 kg/h	V	$T_{g,s}$, $T_{g,w}$	glass transition, solids, water	101, −135 °C	V, H
T_{amb}	room and intake air	25 °C	V	k_{GT}, ΔT_{st}	Gordon–Taylor, sticky criterion	6.7, 20 K	V
$M_{a,1}$, $M_{a,2}$	air in the upper, lower cell	60, 90 kg	H	r_{dep}, r_{clean}	deposit growth, cleaning	2.0 mm/(K h), 0.3 mm/h	V
τ_1, τ_2	particle residences	5, 20 s	V	D_{fire}, T_{fire}, t_{fire}	layer, roof, time to ignite	3 mm, 158 °C, 20 min	V
$UA_{w,1}$, $UA_{l,1}$	roof transfer, loss	1700, 150 W/K	H	k_{CO}	smolder growth	0.5 min^{-1}	H
$C_{w,1}$	roof heat capacity	1.5 MJ/K	H	—	CO alarm, shutdown	15, 40 ppm	V
$UA_{w,2}$, $UA_{l,2}$	wall transfer, loss	1000, 340 W/K	H	$\dot{m}_{B,3}$, $\dot{m}_{B,4}$	bed air, drying, cooling	2500, 1500 kg/h	V
$C_{w,2}$	wall heat capacity	2.5 MJ/K	H	$\tau_{b,3}$, $\tau_{b,4}$	powder residence at design	180, 120 s	V
k	drying constant, chamber	1.47×10^{-3} K^{-1} s^{-1}	V	C_b, UA_b	bed steel, loss per section	200 kJ/K, 20 W/K	H
X_{cr}, n	critical moisture, exponent	6.7 % d.b., 1	V	M_{fb}	air above each section	5 kg	H
k_{bed}	drying constant, bed	4.33×10^{-6} K^{-1} s^{-1}	V	$Q_{h,max}$, τ_H	gas heater rating, lag	750 kW, 30 s	V, H
X_m	GAB monolayer moisture	0.05 kg/kg	H	$Q_{b,max}$, τ_{Hb}	bed heater rating, lag	80 kW, 20 s	V, H
C_0, H_C	GAB constant C	0.1402, 1327 K	H	F_{max}, τ_F	feed pump capacity, lag	1500 kg/h, 3 s	V, H
K_0, H_K	GAB constant K	0.2793, 349 K	H	τ_{tank}, τ_{amb}	balance tank, weather lags	300, 1200 s	H
c_s	milk solids heat capacity	1500 J/(kg K)	fixed	w_{spec}, w_{tol}	moisture specification	3.5 % ± 0.3 %	V
$\dot{s}_D$	design solids flow	480 kg/h	fixed	—	lab interval, repeatability	30 min, 0.05 %	V, H

the falling-rate period at the design point, so the product of that rate and the step is 0.64, well inside the stability limit of the method. Halving the step moves the deposit of the push case at 90 min by 5.8×10^{-10} of its value and the humid case's moisture at 120 min by 1.3×10^{-11} points, because the controller scans and the scripted events fall on the same instants at either step. The water, solids and energy ledgers close to 1.5×10^{-13} of their throughput over the moisture-control, push and fire cases. Every noise sequence is seeded at 20261027, so a run replays to the bit, and the speed selector plays the dryer at 1, 5, 20, 60 or 300 minutes per second, five by default, at which the campaign plays in 96 s, without changing its trajectory (Section 1.6.1).

10.3.8 Assumptions

The chamber is two well-mixed air cells, each with one particle population of mean moisture. Real sprays have a size distribution and every particle a history, so some particles leave wetter and hotter than the mean, and those are the ones that stick. Atomizer wear enters only as a slower drying constant.

The particle temperature is quasi-steady, which fails close to the atomizer, where a droplet warms from the feed temperature before it reaches the wet bulb. The powder in each bed section shares one temperature with the steel and the air above it, and powder drier than its equilibrium moisture does not take water back, which matters in the cooling section, whose room air would hold the powder at 4.50 % on a dry basis while it is at 3.63 %.

The drying curve and the isotherm are illustrative constants fitted to the declared design point, not to a milk powder, and the application sets the critical moisture of 6.7 % on a dry basis lower than milk's so that the design particle sits off the sticky line. The sticky line is 20 K above the lactose glass transition, while the measured shift for skim milk powder is 23.3 K [14]. Stickiness depends on the

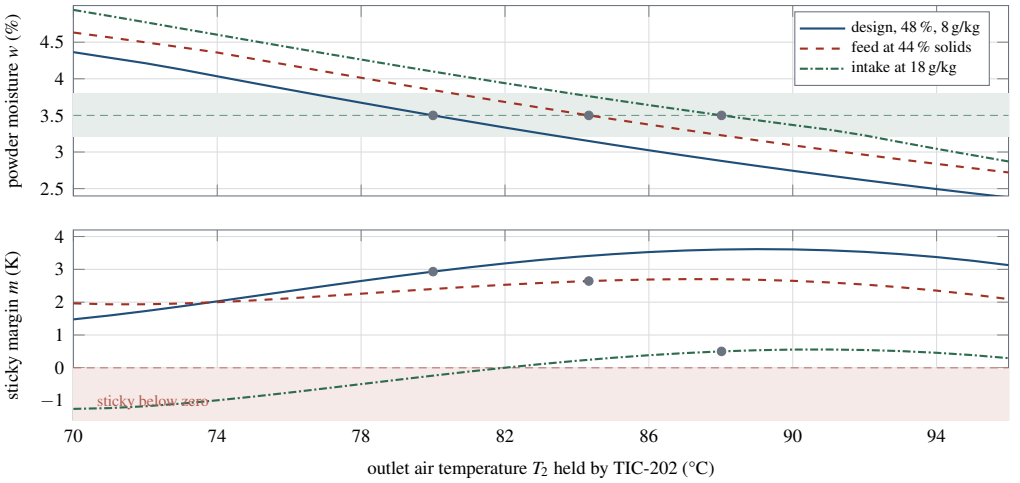

Figure 10.2 Steady-state powder moisture and sticky margin against the outlet temperature held by TIC-202, at a 190 °C inlet and a 55 °C bed. Solid lines are the design point, dashed lines a 44 % concentrate and dash-dotted lines 18 g/kg intake air. The upper band is the 3.2 % to 3.8 % specification, the lower band the sticky region below zero margin, and the dots mark the outlet that gives 3.5 %, 80.0 °C at the design point, 84.3 °C with a 44 % concentrate and 88.0 °C with 18 g/kg intake air.

particle surface rather than on its mean moisture, and the application notes that milk proteins raise the real sticky curve.

The deposit forms only on the roof, and the deposit law, the air-hammer cleaning, the ignition rule and the CO levels are teaching values, with no explosion, venting or suppression model. The air flow is fixed, the walls are two lumped heat capacities, and the cyclone and the ducts are drawn but carry no states. The concentrate arrives through a 5 min balance tank, and in the campaign its solids replay a declared dip of the kind the evaporator makes, one way, without coupling to a model of it. The three-effect milk evaporator of Volume 6, Chapter 4 holds those solids against its own delay. The laboratory has a repeatability of 0.05 % and no sampling error beyond it. All sizes, ratings and limits are teaching values for one product, not a design basis.

10.4 What the Outlet Temperature Stands For

At the design point the gas heater delivers 561.5 kW to the drying air and the bed heater 22.0 kW to the drying section, and the dryer evaporates 503.1 kg/h of water to make 497.9 kg/h of powder, 4219 kJ of heat per kilogram of powder. The upper cell runs at 126.4 °C and 31.8 g/kg, and its particles leave it still in the constant-rate period, at the wet bulb of 45.7 °C and 32.9 % moisture. The lower cell, the exhaust, runs at 80.0 °C, 49.6 g/kg and 15.8 % relative humidity, with a wet bulb of 45.2 °C. Its particles leave the chamber at 4.17 % moisture and a normalized drying rate of 0.44, at 64.8 °C against a glass transition of 47.7 °C, which is the margin of 2.93 K. The drying section at 55.0 °C, blown with air at 56.0 °C, finishes the powder to 3.50 %, and the cooling section brings it to 34.9 °C without drying it further. The walls lose 32.3 kW, 18.4 kW of it from the roof at 147.4 °C and 14.0 kW from the cylinder at 66.0 °C, and the bed loses 0.8 kW.

The outlet temperature is a good proxy only along one curve, as shown in Figure 10.2. Near the design point each kelvin of outlet takes 0.085 % of moisture off the powder and 1.42 % off the throughput, because at a fixed inlet a hotter exhaust carries more of the heat away unused. The curve that gives 3.5 % at 80.0 °C gives it at 84.3 °C when the feed is 44 % solids and at 88.0 °C when the intake carries 18 g/kg, and at a held 80 °C the moisture falls 0.079 % per point of feed solids and rises 0.067 % per g/kg of intake humidity. The horizontal shifts are the feedforward gains, 1.08 K

per point of solids below 48 % and 0.80 K per g/kg of humidity above 8 g/kg, which the application rounds to 1.1 and 0.8. They are not constant. Holding 3.5 % takes 76.8 °C at 52 % solids, 0.80 K per point on the strong side, and 90.3 °C at 22 g/kg, 0.74 K per g/kg.

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

The two disturbances act on the constraint in different ways. A weaker concentrate at a held 80 °C sends the particles out at 4.71 % but cooler, at 59.9 °C, and their glass transition falls to 42.3 °C, so the margin drops only from 2.93 to 2.40 K. Humid intake air raises the wet bulb that the drying rate works against, the particles leave at 4.51 % and 64.5 °C with a glass transition of 44.3 °C, and the margin falls to −0.24 K. A hotter outlet does not always help. It dries the particles, which raises their glass transition, but it also heats them, and the design margin peaks at 3.61 K near 90 °C and then falls. In the humid case the margin is negative below an outlet of 82 °C, peaks at 0.56 K and is 0.50 K at the 88.0 °C that holds 3.5 %. At 22 g/kg, the top of the humidity slider, the outlet that holds 3.5 %, 90.3 °C, leaves the particles 0.55 K past the sticky line, so no outlet setpoint at a 190 °C inlet delivers both.

Throughput is bought with the inlet. At a held 80 °C outlet each kelvin of inlet brings 9.75 kg/h more feed, adds 0.40 g/kg to the exhaust and takes 0.121 K from the margin, and the design particle turns sticky at a steady inlet of 215.7 °C, with 1252 kg/h of feed and 624 kg/h of powder. The roof reaches the 158 °C of the fire rule sooner, at an inlet of 208.5 °C, so between the two a deposit already on the roof can ignite while the powder being made is not sticky.

10.5 Implemented Controller

Six PI loops, the soft sensor, the laboratory and the optimizer execute every $\Delta t_c = 2\,\mathrm{s}$, at the start of every fourth physics step and before its integration, after a priming call at reset. Inside a physics step a scripted event that is due is applied first, the controller runs next when a scan is due, and the integration, the checks for sticky powder, ignition and the CO interlock, the outputs and the historian follow in that order. The structure is drawn in Figure 10.3.

10.5.1 Instruments and the Laboratory

The controllers use the transmitter and laboratory readings, together with the initial deposit specified for the scenario. A transmitter reads the true value where it is mounted, passes it through an exact first-order lag and adds noise that is itself filtered to first order, with the standard deviation of Table 10.3. The laboratory is the framework's shared sampled analyzer, and Section 1.4.5 gives the age its results add to a loop. It draws a sample of the true powder moisture at the bed discharge every 30 min, reports it after the turnaround, 40 min by default, with white noise of 0.05 %, and holds the result until the next one. At reset a sample drawn 30 min earlier is already in analysis and reports at 10 min, so a fresh result arrives at 10, 40, 70 minutes and so on. The first sample drawn after an event at 5 min is reported 65 min after it.

10.5.2 The PI Law

The six loops share the framework's velocity-form PID of Section 1.6.2, written for a loop without derivative action in Equation (10.7).

$$\begin{aligned} e_k &= \gamma(r_k - y_k), \qquad p_k = K_c e_k \\ \Sigma_k &= \Sigma_{k-1} + (p_k - p_{k-1}), \qquad v_k = \Sigma_k + I_{k-1} \\ u_k &= \min\left[u_{max}, \max(u_{min}, v_k)\right] \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - v_k}{\tau_I}\right]\Delta t \end{aligned} \tag{10.7}$$

In it r is the setpoint, y the measurement, $\gamma = \pm 1$ the action and Δt the execution period. It is Equations (1.16) to (1.18) without the derivative term and the slew limit, which no loop here uses.

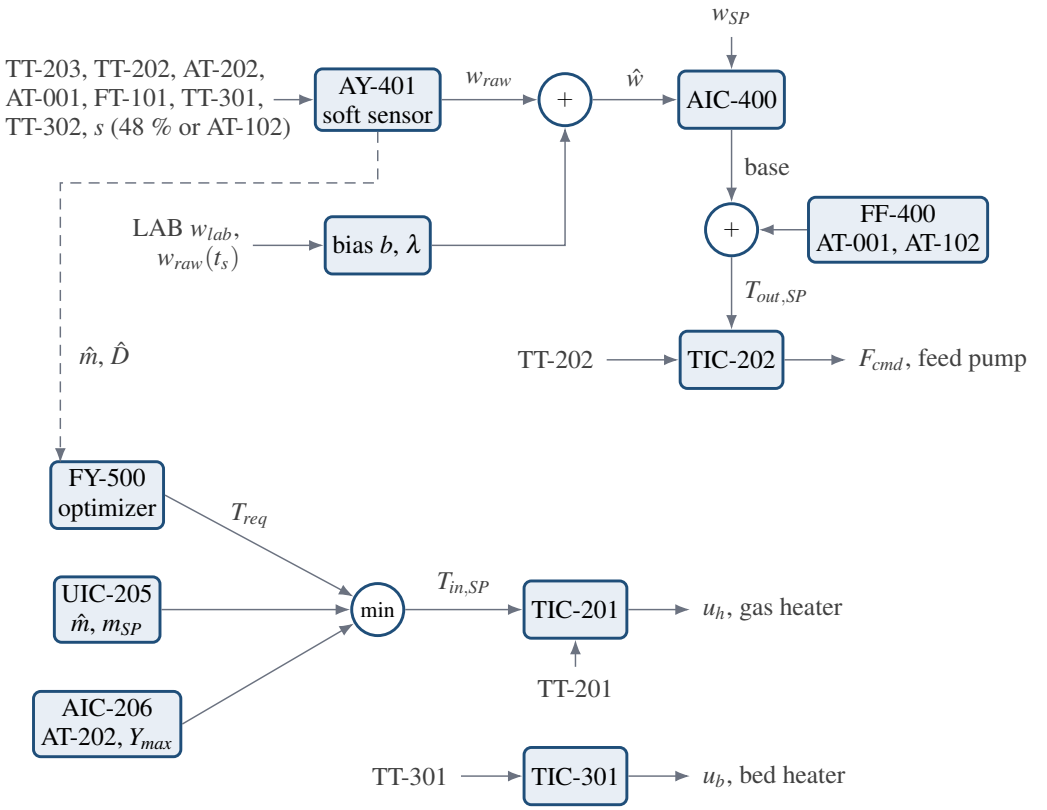

Figure 10.3 Control structure. The soft sensor AY-401, corrected by the laboratory bias, feeds AIC-400, whose outlet setpoint FF-400 biases before TIC-202 moves the feed. The optimizer's request and the two overrides meet in a minimum selector ahead of TIC-201, and the soft sensor's margin and deposit estimates reach the optimizer and UIC-205.

Table 10.3 Transmitters and the Laboratory. Each Noise Has the Correlation Time Shown, and FT-101's Is a Fraction of Its Reading

Tag	Measures	Lag	Noise 1σ, time
TT-201	inlet air	5 s	0.2 K, 3 s
TT-203	upper cell air	5 s	0.3 K, 3 s
TT-202	exhaust air	5 s	0.1 K, 3 s
AT-202	exhaust humidity	30 s	0.1 g/kg, 10 s
AT-001	intake humidity	60 s	0.05 g/kg, 20 s
FT-101	feed flow	1 s	0.3 %, 2 s
AT-102	feed solids	20 s	0.05 points, 10 s
TT-301, TT-302	bed sections	10 s	0.1 K, 5 s
AT-205	exhaust CO	20 s	0.2 ppm, 10 s
LAB	powder moisture	30, 40 min	0.05 points, white

The sum Σ_k always equals the present p_k, and windup is handled on the integral alone, by back-calculation toward the applied output with the time constant τ_I. The exception is AIC-400 on the laboratory alone, which tracks with the longer of τ_I and the time since the previous result. The clamp that follows the feedforward in Equation (10.9) is not one of AIC-400's limits, so a setpoint held there by the clamp does not unwind AIC-400's integral. In manual, or when a loop is not the one

Table 10.4 The Six PI Loops as Shipped. In the ISA Convention the Loops With $\gamma = -1$ Are Direct Acting and Those With $\gamma = +1$ Reverse Acting

Loop	PV	Output	γ	K_c	τ_I (min)	Limits
TIC-202	T_{202}	F_{cmd}	-1	30	1.0	0–1500
TIC-201	T_{201}	u_h	$+1$	1.0	0.5	0–100
TIC-301	T_{301}	u_b	$+1$	6	4	0–100
AIC-400	$\hat{w}$, w_{lab}	T_{base}	-1	5	8	65–100
UIC-205	$\hat{m}$	T_{SO}	-1	5	1	150–240
AIC-206	Y_{202}	T_{YO}	$+1$	1.2	1	150–240

K_c in kg/h per K, % per K, % per K, K per %, K per K and K per g/kg, limits in kg/h, %, %, °C, °C and °C.

selected, it tracks by setting $I_k = u_{op} - \Sigma_k$, so the emitted value is the one applied and the return to automatic is bumpless. At reset a priming call tracks every loop at the steady actuator values of the solved state. Table 10.4 gives each loop's pairing and shipped tuning.

TIC-202 is the classical loop. It raises the feed when the exhaust runs hot, because more water cools the air, and its shipped tuning of 30 kg/h per kelvin and 1.0 min settles an outlet step in under two minutes. TIC-201 holds the inlet with the gas heater and TIC-301 the drying section of the bed with the bed heater.

10.5.3 The Soft Sensor and Its Laboratory Bias

AY-401 infers the moisture now from the measured state of the air, with the plant's own drying curve and a new atomizer, as shown in Equation (10.8).

$$\begin{aligned} \hat{X}_0 - \hat{X}_1 &= \tau_1 k \big[T_{203} - T_{wb}(T_{203}, \hat{Y}_1)\big] f(\hat{X}_1) \\ \hat{Y}_1 &= Y_{001} + F_{101}\, \hat{s}\, (\hat{X}_0 - \hat{X}_1)/\dot{m}_a \\ \hat{X}_1 - \hat{X}_2 &= \tau_2 k \big[T_{202} - T_{wb}(T_{202}, Y_{202})\big] f(\hat{X}_2) \\ b &\leftarrow \mathrm{clamp}\big(b + \lambda\,[w_{lab} - w_{raw}(t_s) - b],\ \pm 2\big) \\ \hat{w} &= w_{raw} + b \end{aligned} \tag{10.8}$$

The feed carries $\hat{X}_0 = (1-\hat{s})/\hat{s}$ of water per kilogram of solids, with $\hat{s}$ either the operator's 48 % or the reading of the densitometer AT-102, subscripts name the transmitters, and each f takes its equilibrium moisture from the relative humidity of the air it is evaluated in. The chamber is solved quasi-steadily by false position at every execution, because a particle spends 25 s there, and the two bed sections are integrated from the chamber estimate by an explicit Euler step of 2 s at TT-301 and TT-302, because the powder spends about 5 min there. The raw estimate w_{raw} is the moisture leaving the modeled cooling section. The same solution gives an estimated margin $\hat{m}$ for the particles leaving the chamber, and the deposit law of Equation (10.5) run on $\hat{m}$ gives an estimated deposit $\hat{D}$. With exact inputs the soft sensor reproduces the plant's steady moisture to 6.9×10^{-18} at five operating points, because it solves the plant's equations. What it cannot know is what it is not told, the true feed solids unless it reads AT-102 and the wear of the atomizer. Chapter 11 infers the unmeasured core of a brick from a two-node model in the same way.

The bias follows the framework's shared update, an exponentially weighted average of the model's residual and the steady scalar filter of Equation (1.12), with $\lambda = 0.5$ and a clamp of ± 2 points, so a constant model error ε is learned as $\varepsilon[1-(1-\lambda)^n]$ after n results. Each result is compared with the raw estimate at the time t_s its sample was drawn, read from a ring buffer of 8192 values at the 2 s cycle, and not with the estimate of the moment, which would chase the plant's own response to the corrections (Equation (1.14)). With the correction switched off the residual is recorded and the bias does not move.

Pause. Try Exercise 10.2 before reading on.

10.5.4 Moisture Control and Feedforward

AIC-400 sets the outlet setpoint from the corrected estimate every 2 s. Feedforward from the intake humidity and the feed solids is added to its output, as shown in Equation (10.9).

$$\begin{aligned} T_{ff} &= K_Y\,(Y_{001} - 8) + K_S(48 - s_{102}) \\ T_{out,SP} &= \mathrm{clamp}(T_{base} + T_{ff},\ 65,\ 100) \end{aligned} \tag{10.9}$$

where T_{base} is AIC-400's output, or the operator's setpoint with the moisture controller off, Y_{001} is in g/kg and s_{102} in percent, each term is present only when its feedforward is switched on, and $K_Y = 0.8\,\mathrm{K}$ per g/kg and $K_S = 1.1\,\mathrm{K}$ per point come from the steady map of Figure 10.2. With the soft sensor set to laboratory only, AIC-400 executes once per laboratory result with Δt equal to the time since the previous result, 30 min, and holds its output in between.

10.5.5 Overrides and the Throughput Optimizer

The inlet setpoint that TIC-201 receives is the lowest of the requested inlet and the enabled overrides. UIC-205 lowers it when the estimated margin falls below its 0.5 K setpoint, and AIC-206 when AT-202 passes its 58 g/kg ceiling. The overrides act on the inlet because, with the outlet held, the inlet sets how much water the air has to carry. Cutting the feed at a fixed inlet heats the particles faster than it dries them. At a 225 °C inlet, a cut from 1343 to 1150 kg/h lifts the steady margin only to 0.03 K, and a cut to 950 kg/h takes it back to −3.46 K, with the particles at 98.9 °C. Each override that is not selected tracks the applied setpoint by external reset, so a change of selection is bumpless. The requested inlet is the operator's, or with FY-500 on it follows Equation (10.10).

$$\begin{aligned} \eta &= \min\Big[\hat{m} - m_{min},\ \frac{D_{max} - \hat{D}}{0.25},\ \frac{95 - u_h}{3}, \\ &\qquad \frac{Y_{max} - 2 - Y_{202}}{2},\ \frac{98 - T_{out,SP}}{2}\Big] \\ \frac{\mathrm{d}T_{req}}{\mathrm{d}t} &= r\,\mathrm{clamp}(\eta, -4, 1) \\ T_{in,SP} &= \min(T_{req}, T_{SO}, T_{YO}) \end{aligned} \tag{10.10}$$

Each headroom is measured in its own unit of caution, the humidity term is present only with that override on, and the last only with the moisture controller on. With $m_{min} = 1.5\,\mathrm{K}$, $D_{max} = 0.5\,\mathrm{mm}$ and $r = 0.5\,\mathrm{K/min}$, the request climbs at r while every headroom exceeds one, more slowly as the smallest approaches zero, and falls at up to 2 K/min when it is negative. It is kept no more than 2 K above the applied setpoint and between 150 °C and the ceiling $T_{in,max} = 225\,^{\circ}\mathrm{C}$, and while an override holds the inlet it may only descend. FY-500 is in effect an integrating controller on the worst normalized headroom, a constraint pusher rather than a model-based optimization, and it comes to rest where the binding headroom is zero. Volume 4, Chapter 6 pushes the line speed of another dryer toward its first binding limit with a model-based push.

Pause. Try Exercise 10.3 before reading on.

10.5.6 Model Predictive Control

A model predictive controller (MPC) is offered as an option for the feed and the inlet. All cases retain the loops as their initial structure apart from case 9, which starts on MPC. A quadratic program is solved at a period $T_s = 10\,\mathrm{s}$ (five control scans) for the inlet request T_{req} and the feed command F_{cmd} over $N = 90$ steps, or 15 min, with the block lengths $(1, 1, 1, 1, 2, 2, 3, 4, 6, 9, 15, 45)$. The first planned values are applied for one decision period, the receding-horizon practice of Section 1.6.3 [19, 20]. The first plan is made on the priming call, so a run under the MPC logs it from $t = 0$. The request

Table 10.5 Response Shapes of the MPC's Model, Equation (10.11). Times in Seconds, Fit Residual RMS in Percent of the Change

Pair	a	θ	τ_1	τ_2	τ_s	Resid.
$F \to \hat{w}$	0.842	36.4	142.5	106.3	2199	0.17
$F \to T_{202}$	0.801	5.9	22.6	11.5	1876	0.12
$F \to Y_{202}$	1.017	10.6	30.5	30.5	1990	0.10
$F \to \hat{m}$	1.081	8.9	33.8	33.8	426	2.52
$T_{req} \to \hat{w}$	0.580	63.6	60.4	205.1	1489	0.20
$T_{req} \to T_{202}$	0.588	13.8	27.8	9.6	1482	0.52
$T_{req} \to Y_{202}$	0.582	81.0	50.3	9.2	1379	0.84
$T_{req} \to \hat{m}$	0.462	49.0	27.6	27.6	1360	3.25
$T_{req} \to u_h$	1	0	10	1	—	—

takes FY-500's place in the minimum selector of Equation (10.10), so UIC-205, AIC-206 and the heater clamp act above the plan as they act above the optimizer, and the feed goes to the pump in TIC-202's place. Meanwhile TIC-202's integral tracks the plan's feed and the outlet setpoint follows TT-202, so that AIC-400 also tracks and the loops stand ready to resume.

The controller's model has the handles $u = (T_{req}, F_{cmd})$ and five outputs, $y = (\hat{w}, T_{202}, Y_{202}, \hat{m}, u_h)$, the corrected moisture estimate, the exhaust temperature and humidity, the estimated margin and the gas heater output. Its gains K_{ij} come from the steady balance that AY-401 solves, evaluated with the declared parameters of Table 10.2, a 50 °C feed and a new atomizer, never with the run's values, and they are recomputed at every decision by forward differences of 1 K and 1 kg/h at the measured operating point (TT-201, the applied feed, TT-301, AT-001 and the solids the soft sensor is told). At the design point, per kg/h of feed and per kelvin of inlet, the moisture gains are 0.00631 and −0.0541 points, the outlet gains −0.0745 and 0.727 K, the margin gains −0.0101 and −0.0243 K, and the heater gains 0 and 0.454 points. Because the margin's feed gain changes sign near the sticky point, that path is omitted from the plan, as it is from the inlet-only overrides. The remaining paths follow Equation (10.11).

$$y_{ij}(t) = K_{ij}\Big[a_{ij}\, g_2(t-\theta_{ij};\tau_{1,ij},\tau_{2,ij}) + (1-a_{ij})\big(1-\mathrm{e}^{-t/\tau_{s,ij}}\big)\Big] \tag{10.11}$$

Here g_2 is the unit step response of two first-order lags in series and the parallel lag τ_s is the chamber walls. The shapes, in Table 10.5, were fitted to steps of 50 kg/h of feed and 5 K of inlet setpoint in both directions from case 2's design point with its event removed, TIC-201 and TIC-301 in automatic and the feed scripted, each against an unstepped run on the same seed over 120 min. At T_s each dead time becomes whole steps on a delay line of its input, the remainder joining τ_1, and the lags are discretized exactly, which gives 38 states (11 for the delays and 27 for the lags). With a constant input the model settles on the linearization $K\Delta u$ to within 1.4×10^{-10}.

Driven by the same steps with the noise off, the plan's model is within 3 % of the plant's change after 120 min for every output but the margin, which falls 0.161 K after the inlet step against a predicted 0.123 K, 31 % more, and whose gain differs by 33 % between steps up and down. The heater path uses nominal lags of 10 and 1 s rather than an independently fitted pair. Its prediction has an RMS error of 6.8 % of the step change, with the correct steady gain.

The optimized variables are increments of each handle, referenced to the currently applied input. A block shares one increment. Each plan minimizes Equation (10.12), in which j counts the blocks,

$\hat{y}_{i,k}$ is a predicted output and ε_i the one slack of output i.

$$\begin{aligned} \mathscr{J} = &-\sum_{k=1}^{N} \frac{F_{cmd,k}}{\mathrm{kg/h}} + \lambda \sum_{j=0}^{11} \left[\left(\frac{\Delta T_{req,j}}{1\,\mathrm{K}} \right)^2 + \left(\frac{\Delta F_{cmd,j}}{10\,\mathrm{kg/h}} \right)^2 \right] \\ &+ \sum_{i=1}^{5} \left(\nu_i \varepsilon_i + \kappa_i \varepsilon_i^2 \right) \\ \text{s.t. } & \underline{y}_{i,k} - \varepsilon_i \le \hat{y}_{i,k} \le \bar{y}_{i,k} + \varepsilon_i, \quad \varepsilon_i \ge 0 \\ & 150\,^{\circ}\mathrm{C} \le T_{req,j} \le T_{in,max}, \quad 0 \le F_{cmd,j} \le F_{max} \\ & -4rT_s \le T_{req,j} - T_{req,j-1} \le rT_s \end{aligned} \tag{10.12}$$

The feed earns one unit per kg/h per step, and $\lambda = 1000$ prices moves. The soft limits are the loops' own. The moisture estimate stays under AIC-400's setpoint, 3.5 %, the margin estimate over the optimizer's floor $m_{min} = 1.5\,\mathrm{K}$, the heater under 95 %, the exhaust humidity under $Y_{max} - 2$ while AIC-206 is on, and the outlet inside the 65 to 100 °C that AIC-400 may ask for. The slack prices (ν_i, κ_i) are $(2 \times 10^5, 2 \times 10^6)$ per point of moisture, $(10^5, 10^6)$ per kelvin of margin, $(2 \times 10^4, 2 \times 10^5)$ per point of heater, $(3 \times 10^4, 3 \times 10^5)$ per g/kg and $(10^4, 10^5)$ per kelvin of outlet, each far above the 90 units a kg/h of feed earns over the horizon. When the initial prediction violates a band, its bound tightens linearly toward the specified limit, taking 60 steps for moisture, 3 for the heater and 6 for each remaining output. The handle limits are hard. At the default ceiling of 225 °C and rate r, the request rises at most 0.083 K and falls at most 0.333 K per block. It cannot rise above the applied setpoint while the deposit estimate exceeds D_{max}. Energy is not priced. The move price was set once on case 9, before any comparison, as the largest λ within 1 % of the best powder in specification among values that reject no plan, leave no sticky minute and move the feed no more than the loops. The model's states are advanced by the inputs actually applied over each period, under either structure, and each output carries an additive bias, the period's mean measurement less the model's present output, held over the horizon. This output-disturbance model makes the plan offset-free on what the transmitters and AY-401 read [21], and no Kalman filter is involved (Section 1.4.7). The offset-free MPC of Volume 9, Chapter 4 carries its disturbance in a Kalman estimator instead. The moisture the plan reads is AY-401's corrected estimate, so the laboratory reaches it only through the bias of Equation (10.8), and it never sees the true moisture, the true solids or a laboratory result before the report. No preview of any event reaches either structure.

The program has 24 input values and 5 slacks and is solved by the shared dual active-set method [22]. Acceptance requires an optimal status, finite values and inputs within their bounds. Otherwise the tracked loops resume control until a subsequent plan passes these checks. Three forced rejections with the noise off hand over with a feed step of 0.011 kg/h and an inlet setpoint step of 0.004 K, and none of the 178740 plans measured below was rejected. The minimum selector still acts at every scan, and a forced ignition in case 7 shuts off the feed and heaters. Manual mode starts at the inlet setpoint after the protective overrides. The MPC reads only what the loops read and the declared values, and nine copies of the controller altered to hand it one truth value or run parameter all fail the application's checks. A warmed plan takes 3.3 ms in Chrome, so the plans need more than a core at the two top speeds, and under the MPC each frame stops stepping after 25 ms, which leaves the logged rows unchanged and plays those speeds at about 40 min/s.

10.5.7 Manual Mode and the Shutdown

In manual the operator sets the feed and both heaters, every loop tracks, and switching to manual copies the current outputs to the manual stations. After a CO shutdown the feed and heaters are held at zero and the loops track zero. A manual helper sets the design point, 1001 kg/h of feed with the heaters at 74.9 and 27.4 %.

Table 10.6 Scenario Events and Teaching Points

#	Scenario and event	Teaching point
1	Manual, feed 1000 to 1100 kg/h at 5 min	fast air, slow powder
2	Outlet 80 to 83 °C at 10 min	the price of a kelvin
3	Feed solids 48 to 44 % at 10 min	a proxy holds the wrong thing
4	Case 3, AIC-400 on the soft sensor	the laboratory corrects
5	Intake 8 toward 18 g/kg at 10 min	feedforward meets the weather
6	Inlet 190 to 225 °C at 10 min	throughput, stickiness
7	3.5 mm deposit, inlet 200 to 225 °C	a fire no loop stops
8	AIC-400, bed 70 °C, inlet 205 °C	a second handle
9	Case 4, FY-500, UIC-205, MPC	one plan, two handles
10	480 min, humid hours, solids dip	graded campaign

10.5.8 The Three Models

Section 10.3 is the simulation model, the plant itself. The shipped loops use no identified dynamic control model. Their response can be characterized by step tests at the design point. At a fixed heater, 100 kg/h more feed lowers the outlet by 5.9 K within 2 min and by 7.3 K at steady state, −0.073 K per kg/h, the slower part coming from the walls, and raises the steady powder moisture by 0.65 points, reaching half of its 30 min change 4.3 min after the step. A point of gas heater raises the inlet by 2.20 K. The shipped controllers carry four internal models. The soft sensor is the plant's own drying model with the design atomizer, the feedforward gains are the slopes of Figure 10.2, the deposit estimate is the plant's deposit law driven by the estimated margin, and the optimizer's headrooms are fixed scalings. None of them knows the feed solids unless AT-102 is read. The optional MPC carries a fifth, the steady model evaluated with the declared parameters of Table 10.2 and joined to lags identified offline by step tests on the application, Section 10.5.6. A learner's edit of a plant parameter therefore reaches the soft sensor but not the MPC, for which it is model mismatch.

10.6 Measured Behavior of the Nine Cases

Table 10.6 lists the cases and the campaign. Each starts from the steady state of the reference configuration, the outlet loop in automatic at 80 °C, the inlet at 190 °C, the bed at 55 °C, the soft sensor assuming 48 % solids with the laboratory correction on, and the moisture controller, both feedforwards, the optimizer and both overrides off, and differs from it only where its row says. All numbers below were computed by the engine outside the browser, stepped exactly as the browser steps it (Section 1.6.4), with the default seed and in minutes of simulated time. Every run exports a comma-separated file of 151 columns with the true moisture, margin and deposit beside every instrument and estimate. A setting that a configuration changes is changed after the case is loaded and before the run starts, as a learner does after a Reset. Carrying the same setting through the Reset instead gives identical runs in cases 1 to 8 and moves the campaign by at most 0.007 t and 8 kJ/kg. Where the noise of the transmitters and the laboratory decides a comparison, the range over eight further noise seeds is given. Table 10.7 collects the outcomes of the three cases that turn on the proxy and of case 9, and Figure 10.4 draws them as bars.

10.6.1 Fast Air, Slow Powder

Case 1 runs in manual with the feed at 1000 kg/h, the gas heater at 74.9 % and the bed heater at 27.4 %, and raises the feed to 1100 kg/h at 5 min. The outlet falls from 80.1 °C to 76.9 °C within 30 s, to 75.0 °C within a minute and 74.3 °C within two, and then creeps to 73.5 °C by 30 min and 73.0 °C by 90 min as the walls cool. The particles leaving the chamber go from 4.15 % to 4.81 % within a minute, but the powder has to cross the bed, and its moisture passes 3.8 % at 9.0 min, reaches half of its 30 min change at 9.3 min, and stands at 4.16 % at 30 min, with the margin down to 1.83 K, and at

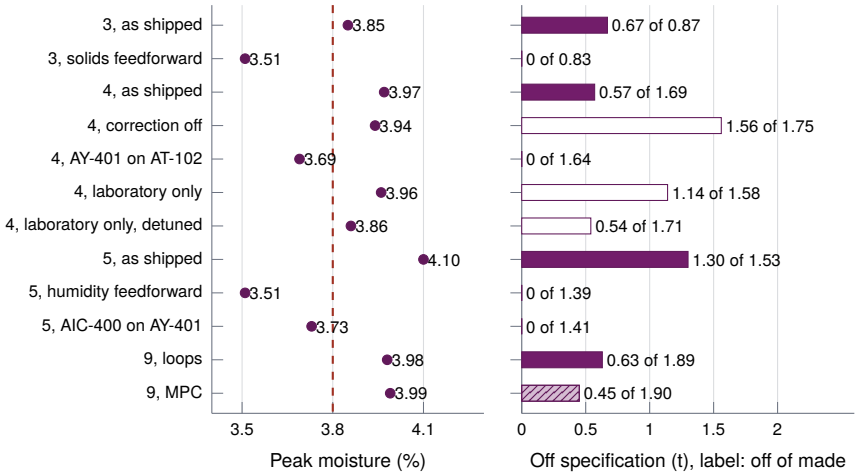

Figure 10.4 The proxy cases and case 9 from Table 10.7: the peak true moisture after the event (left, dashed line at the 3.8 % limit) and the tonnage off specification, labeled with the tonnage made (right). Rows name the case and its configuration, and the detuned row is the table's footnote. Dark bars are the case as shipped, hatched bars add a measurement, a feedforward, the moisture loop or the MPC, and white bars remove the laboratory correction or use the laboratory alone. Notice that in case 4 the rows without the correction and on the laboratory alone at the shipped tuning ship more off-specification powder than the shipped soft sensor, and that every row that reads the true solids or the humidity ships none.

Table 10.7 Truth-Log Outcomes of Cases 3 to 5 and 9. Peak Moisture After the Event, Tonnage Off Specification Out of the Tonnage Made, Worst True Margin and Specific Energy Over the Run

Case and configuration	Peak (%)	Off/made (t)	Worst m (K)	Energy (kJ/kg)
3 as shipped, 120 min	3.85	0.67/0.87	2.38	4820
3 solids feedforward	3.51	0/0.83	2.63	5079
4 as shipped, 240 min	3.97	0.57/1.69	2.27	4990
4 correction off	3.94	1.56/1.75	2.30	4817
4 AY-401 on AT-102	3.69	0/1.64	2.61	5118
4 laboratory only	3.96	1.14/1.58	1.48	5322
4 laboratory only[a]	3.86	0.54/1.71	2.38	4937
5 as shipped, 180 min	4.10	1.30/1.53	−0.26	4177
5 humidity feedforward	3.51	0/1.39	0.47	4589
5 AIC-400 on AY-401	3.73	0/1.41	0.47	4503
9 loops, 240 min	3.98	0.63/1.89	0.61	4817
9 MPC	3.99	0.45/1.90	0.57	4831

[a] detuned, $K_c = 2$ K per % and $\tau_I = 60$ min

4.21 % at 90 min. The laboratory reports 4.11 % at 70 min, from the first sample drawn after the step. The heater works the other way. The feed step is scripted and recurs on every run, so the heater is tested by returning the feed to 1000 kg/h at the scan after the step and raising the heater from 74.9 to 79.9 %. That lifts the inlet to 201.1 °C, the outlet to 84.5 °C within 2 min and 87.0 °C by 30 min, and dries the powder from 3.49 to 2.95 %.

10.6.2 The Price of a Kelvin

Case 2 steps the outlet setpoint from 80 to 83 °C at 10 min. The true outlet is within 0.3 K of its new setpoint 30 s later, stays there after 1.5 min and peaks at 83.17 °C, and the integral of its absolute error from 10 to 60 min is 3.18 K min. The scorecard, which integrates the error of TT-202 and its noise from the start, reads 6.41 K min. A sluggish tuning of 10 kg/h per kelvin and 3 min takes 12.2 min to come within 0.3 K and 13.7 K min. By 60 min the powder is drier, 3.25 % against 3.50 %, and the margin wider, 3.30 against 2.93 K, but the feed has fallen from 1001 to 960 kg/h and the powder from 498 to 476 kg/h, 4.4 % less, while the heaters spend 4402 kJ per kilogram instead of 4219. Dryers are run at the lowest outlet that keeps the powder in specification because every kelvin of it is about 1.4 % of throughput.

10.6.3 Two Ways to Break the Proxy

Case 3 changes the evaporator's product from 48 to 44 % solids at 10 min. Each droplet now carries 1.27 kg of water per kilogram of solids instead of 1.08 kg, and the outlet loop trims the feed to a mean of 931 kg/h over 100 to 120 min while holding the true outlet within 0.39 K of 80 °C, and TT-202 within 0.66 K with its noise. The particles leave the chamber at 4.70 % instead of 4.17 %, the powder passes 3.8 % at 26 min and reaches 3.85 %, and 0.67 of the 0.87 t made in 120 min is off specification. The soft sensor, still told 48 %, gives a raw estimate of 3.40 %, drier than before, because it starts every droplet with the water of a 48 % concentrate, while the laboratory reports 3.78 % at 70 min and 3.87 % at 100 min. With the solids feedforward on AT-102 the outlet setpoint rises to 84.5 °C and the powder stays between 3.48 and 3.51 %, at 5079 kJ/kg over the run against 4820 kJ/kg.

> **TRY IT IN THE APP** **3 · A Weaker Concentrate**
>
> Load the case and press Run. Watch the outlet temperature stay close to 80 °C after the step at 10 min, then the true moisture in the powder moisture trend. It passes 3.8 % at 26 min and reaches 3.85 % while the loop holds its setpoint.

Case 5 turns the weather at 10 min, and the intake humidity follows it through its 20 min lag, 15.8 g/kg at 40 min and 17.5 g/kg at 70 min. At a held 80 °C the exhaust climbs from 49.6 to 60.5 g/kg, 18.9 % relative humidity, and its wet bulb from 45.1 to 47.7 °C. The powder passes 3.8 % at 28 min and reaches 4.10 %, the particles leaving the chamber cross the sticky line at 61 min and reach −0.26 K, 0.40 mm of deposit builds by 180 min and 1.30 t is off specification. The humidity feedforward raises the outlet setpoint to 88.0 °C with the weather and holds the powder between 3.49 and 3.51 % and the margin at 0.47 K or more, at the cost of feed, 922 kg/h at 180 min instead of 1028, and energy, 4660 kJ/kg against 4219 before the weather turned. AIC-400 on the soft sensor also holds the powder, because the soft sensor reads the exhaust humidity, but only after the moisture has reached 3.73 % at 34 min. The feedforward acts on the cause, and the feedback waits for the effect.

10.6.4 A Model Corrected by a Late Laboratory

Case 4 repeats the weak concentrate with AIC-400 on the corrected estimate. At first the loop makes things worse. The soft sensor, told 48 %, reads the weaker concentrate as drier powder, so AIC-400 lowers the outlet setpoint to 78.5 °C while the truth rises, and the powder peaks at 3.97 % at 68 min. The first two results, from powder made before the change, come out 0.06 and 0.05 points low at the default seed and start the bias at −0.03 and −0.04 points, which deepens the dip. The results from 70 to 220 min move the bias to 0.17, 0.32, 0.36, 0.42, 0.47 and 0.46 points, the jump at 70 min coming from the first sample of weak concentrate, and AIC-400 raises the outlet setpoint step by step to 84.4 °C. The powder is back under 3.8 % at 102 min and under 3.6 % at 163 min, and at 240 min it is at 3.49 % while the raw estimate reads 3.05 %, as shown in Figure 10.5. Over three further noise seeds (1, 42 and 20261008), the powder is back under 3.8 % between 80 and 98 min and 0.37 to 0.54 t is off specification, against 0.57 t here, so the default seed lies at the slow end. Without the correction



Figure 10.5 The spray dryer application in case 4, soft sensor and lab correction, at $t = 4.00$ h (240 min) in the Product view, 230 min after the concentrate fell from 48 to 44 % solids, run with no control changed, cropped to the view. The signal path shows (1) the measured inputs, among them TT-203 at 133.3 °C, TT-202 at 84.4 °C, AT-202 at 47.6 g/kg, FT-101 at 882 kg/h and the feed solids assumed at 48.0 %, (2) the drying model's raw estimate of 3.05 % with the chamber exit at 3.56 % and an estimated margin of 3.6 K, (3) the laboratory's 3.56 %, reported 20 min earlier, with a bias of 0.46 % at $\lambda = 0.50$, and (4) the corrected estimate of 3.50 % against a true 3.49 % while the feed is truly 44.0 % solids, with (5) AIC-400 holding the outlet setpoint at 84.4 °C. The sample line shows one sample in analysis and the last two results, 3.70 and 3.56 %, above 1.12 t of powder in specification and 0.57 t off, a powder rate of 400 kg/h and 4990 kJ/kg over the run, and the badges give a true margin of 2.6 K.

the controller holds its estimate on 3.50 % while the powder runs at 3.94 %, and 1.56 t of 1.75 t is off specification. Told the true solids by AT-102, the soft sensor needs no correction. The powder peaks at 3.69 % at 24 min, nothing is off specification, and the bias stays within ±0.043 points. Keeping the laboratory correction off when switching to AT-102, as in the three runs that case 4 asks for, gives a peak of 3.68 % at 24 min, no off-specification powder and 5130 kJ/kg.

> **TRY IT IN THE APP** **4 · Soft Sensor and Lab Correction**
>
> Open the Product view, press Run and run to 240 min. Watch the bias in the laboratory box step at each result, 0.17 points at 70 min and 0.46 by 220 min. The true moisture peaks at 3.97 % at 68 min and is back under 3.8 % at 102 min.

The same run also shows what the correction does not reach. The bias is added to the moisture only, so the chamber exit is still estimated at 3.56 % against a true 4.23 %, and the margin at 3.56 K against a true 2.64 K. With the solids assumed, the constraint estimate that the override and the optimizer read is 0.9 K optimistic, and the laboratory never sees it (Section 1.4.6).

On the laboratory alone, AIC-400 moves once per result, and with the shipped tuning it oscillates with a growing amplitude, as a correction that acts on results older than its previous move does (Equation (1.9)). Within the four hours it swings the outlet setpoint between 78.6 and 98.5 °C and the powder between 2.57 and 3.96 %, and 1.14 t is off specification (0.55 to 1.19 t over the three further seeds), and the setpoint reaches its 65 °C limit at 250 min. Slowed to 2 K per point and 60 min it stops

oscillating but is slow, with the powder at 3.86 % at its worst, still at 3.69 % at 240 min and 0.54 t off. The weight trades speed against noise, as Table 1.2 sets out. With $\lambda = 1$ the bias moves to 0.37, 0.46, 0.40, 0.48, 0.52 and 0.45 points at the results from 70 to 220 min, because each update takes the full laboratory noise and the mismatch of powder sampled 40 min earlier, yet the powder is back under 3.8 % at 77 min and 0.40 t is off specification. With $\lambda = 0.2$ the powder is not back under 3.8 % until 136 min and 0.81 t is off. The ordering holds at every seed tried, with 0.24 to 0.40 t off at $\lambda = 1$ and 0.75 to 0.82 t at $\lambda = 0.2$.

> **+ DID YOU KNOW? The Late Truth Loses to a Corrected Guess**
>
> From the simulator: in case 4 the moisture controller fed by the laboratory alone ships 1.14 t off specification of 1.58 t made, while on the soft sensor under laboratory correction it ships 0.57 t of 1.69 t (Table 10.7). Every result is true but late, and a correction that acts on results older than its previous move oscillates.

10.6.5 Throughput Against the Sticky Line

Case 6 raises the inlet setpoint from 190 to 225 °C at 10 min. The outlet loop sends more feed, 1341 against 1001 kg/h at 120 min, the powder rises from 498 to 669 kg/h, 34 % more, the gas heater settles at 90.7 % after touching 100 % in the transient, and the energy per kilogram falls from 4219 to 3778 kJ/kg. The exhaust carries 63.7 g/kg, and at the steady state of a 225 °C inlet the particles leave the chamber at 66.2 °C and 4.42 %, where the glass transition is 45.2 °C, which is 0.98 K past the sticky line. They cross it at 13.7 min. The deposit grows to 0.29, 1.31, 2.46 and 3.64 mm at 30, 60, 90 and 120 min, the roof passes 158 °C at 20.9 min and reaches 167.2 °C, and the layer reaches the 3 mm that can smolder at 104 min. The sticky-margin override settles the inlet at 211.2 °C, with 1207 kg/h of feed and 601 kg/h of powder over the last 20 min, a lowest true margin of 0.40 K and no deposit, although the roof still passes 158 °C at 32.1 min. The exhaust-humidity override at 58 g/kg does nearly the same at 210.8 °C.

10.6.6 A Fire No Loop Can Stop

Case 7 starts with 3.5 mm of powder baked onto the roof and the inlet at 200 °C, where the roof sits at 153.2 °C, and pushes the inlet to 225 °C at 10 min. The particles are sticky from 12.1 min, the roof passes 158 °C at 16.0 min, and after 20 min of both conditions the layer, now 4.0 mm, ignites at 36.0 min with the roof at 165.0 °C. AT-205 reaches the 15 ppm alarm at 41.8 min and the 40 ppm shutdown at 43.7 min, and the run stops at 44.75 min with 0.46 t made. Lowering the inlet to 200 °C at 15 min keeps the roof at 157.4 °C or less, and there is no fire, although 2.95 mm is still on the roof at 120 min. The sticky-margin override holds the inlet at 212.1 °C, where no new deposit forms, but that inlet puts the roof at 159.4 °C, so the old layer ignites at 45.8 min and the dryer shuts down at 53.5 min. This follows from Section 10.4, where the roof reaches 158 °C at a lower inlet than the one at which the powder turns sticky. The case would not arise as a single push in a plant, but it serves to show a constraint that cannot be recovered once crossed, and a deposit that a controller can only keep from growing.

10.6.7 The Bed as a Second Handle

Case 8 holds the moisture with AIC-400 and raises the bed setpoint from 55 to 70 °C at 10 min, then the inlet from 190 to 205 °C at 60 min. The bed air rises from 56 to 81 °C, the powder dips to 3.25 % at 17.3 min, and AIC-400 answers by lowering the outlet setpoint to 75.5 °C by 59 min, so that, averaged over 50 to 59 min, the feed rises from 1001 to 1066 kg/h and the powder from 498 to 530 kg/h, while the energy falls from 4219 to 4087 kJ/kg. The hotter inlet then brings 1204 kg/h of feed, 597 kg/h of powder and 3935 kJ/kg at 180 min, with the outlet setpoint at 76.2 °C. The price is the margin. It falls from 2.93 to 2.21 K with the lower outlet and to 0.60 K with the hotter inlet, because the chamber now leaves the powder wetter, 4.74 % against 4.17 %, and at 180 min the bed powder at 70 °C and 3.51 % is itself only 4.4 K inside its own sticky line. A second handle on the

Table 10.8 Case 9 Over 240 min, Means of Five Seeds. Over: Minutes of True Moisture Above 3.8 %. Lab: Laboratory Results Above 3.8 %, of 40. Travel: Summed Feed Moves

Structure	In (t)	Off (t)	Made (t)	Over (min)	Lab	Energy (kJ/kg)	Travel (kg/h)
Loops (shipped)	1.322	0.570	1.893	72.9	15	4821	16779
MPC (shipped)	1.452	0.444	1.896	56.4	11	4834	1216
AIC-400 15 K/%, 3 min	1.477	0.413	1.890	52.9	11	4839	16991
Loops in hindsight[a]	1.498	0.427	1.925	53.0	11	4808	16999

[a] best of 560 fixed settings at the 1.5 K floor: AIC-400 at 15 K/% and 3 min, FY-500 at 2 K/min, override off

same product lets the cheaper one do more of the work, and the constraint moves to wherever that handle pushes.

10.6.8 Feed and Inlet in One Plan

Case 9, Feed and inlet in one plan (MPC), repeats the weaker concentrate of case 4 with the soft sensor still told 48 % and with AIC-400, FY-500 at its 1.5 K floor and the sticky-margin override on, while the MPC of Section 10.5.6 sets the feed and the inlet. Each row of Table 10.8 changes only the structure or its tuning, over 240 min using five seeds (the scenario default and integers 1 through 4). The instruments, event schedule and noise realization are matched between structures, and neither receives preview information. In these comparisons off-specification powder is produced only from 10 to 130 min.

Like for like, at the loops' own moisture setpoint, floor, heater headroom and climb rate, the MPC makes the same powder as the shipped loops, 1.896 against 1.893 t, at 0.3 % more energy per kilogram, so it offers no throughput gain. It ships 1.452 t in specification against 1.322 t and 0.444 t off against 0.570 t, better at every seed, and moves the feed far less because TIC-202 answers the noise of TT-202. The whole specification gain is the answer to the laboratory result at 70 min, the first from powder made after the change, as shown in Figure 10.6. Before that result both structures hold the same biased estimate on 3.50 % and ship the same wet powder. On the case's seed the true moisture passes 3.8 % at 17.6 min under the MPC and 21.8 min under the loops, and peaks at 3.99 and 3.98 %. After it, the MPC cuts the feed, and its estimate is back at 3.52 % by 74.7 min and the powder under 3.8 % by 74.5 min. AIC-400 at its shipped 5 K per point and 8 min brings the powder under 3.8 % only at 102.7 min and leaves its estimate above 3.52 % until 220 min.

The loops can buy the same speed by tuning. AIC-400 alone at 15 K per point and 3 min brings the powder under 3.8 % at 74.1 min and ships 1.477 t in specification and 0.413 t off, better than the MPC on both at every seed, and so does the best of 560 fixed settings of the shipped structure chosen in hindsight at the same floor, with 1.498 and 0.427 t. The shipped row keeps AIC-400's shipped tuning, which is slow for this case.

Counterfactual copies of the plan show what drives it, as five-seed means. The feed reward drives the throughput and the inlet delivers it (1.409 t of powder without the reward, 1.368 t with the inlet held), the margin floor bounds it (1.990 t without the floor, sticky for 125 min), and the moisture ceiling alone keeps the powder in specification (2.106 t off without it). The outlet band never binds, and leaving it out changes no digit. On the case's seed the true margin reaches 0.57 K against the loops' 0.61 K. Both margin estimates retain the error from the unmeasured weaker concentrate.

In the feed-solids case, case 3's step with the stack on and the soft sensor reading AT-102, the loops with the solids feedforward and the MPC both ship 1.774 t in specification and none off over 240 min, a tie. Case 8 as loaded is a loss, 4.471 against 4.678 t in specification over 480 min at 4027 against 3975 kJ/kg. The plan sets the inlet itself, so the case's scripted step to 205 °C never happens, and it has no handle on the bed. The same is true wherever a case's lesson is a loop or a scripted setpoint step, in cases 2 to 8, which therefore teach their lessons only on the loops, as the application's guide states.

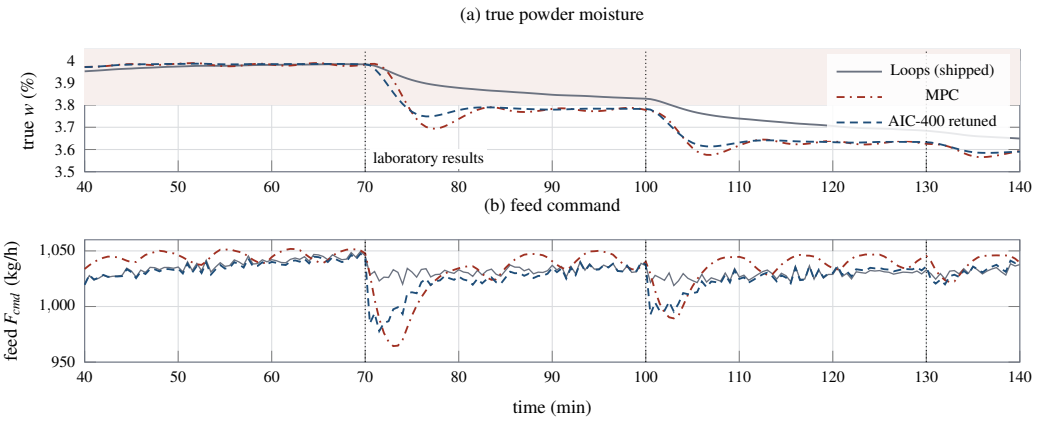

Figure 10.6 Case 9 on its own seed from 40 to 140 min, sampled every 30 s, under the shipped loops (solid), the MPC (dash-dotted) and the loops with AIC-400 retuned to 15 K per point and 3 min (dashed). Dotted vertical lines mark the laboratory results. (a) The true powder moisture, with the region above the 3.8 % limit shaded. After the laboratory result at 70 min the MPC and the retuned loop bring the powder under 3.8 % at 74.5 and 74.1 min, the shipped loop at 102.7 min. (b) The feed command. The MPC cuts the feed by about 70 kg/h within two minutes of the result, while TIC-202's feed carries the noise of TT-202.

Table 10.9 Campaign Configurations Over 480 min, Scored on the Truth Log

Configuration	In (t)	Off (t)	Worst m (K)	Energy (kJ/kg)	Result
as shipped	1.98	2.02	−0.40	4237	fail
AIC-400 on the soft sensor	3.76	0	0.04	4498	fail
and AT-102, both feedforwards	3.76	0	0.06	4500	fail
and the sticky override	3.74	0	0.44	4507	fail
suggested[a]	3.88	0	0.63	4380	pass
suggested, bed at 55 °C	3.78	0	0.61	4496	fail
suggested, optimizer off	3.98	0	0.01	4339	pass[c]
suggested, floor 0.8 K	4.10	0	−0.05	4299	fail
bed 65 °C, sticky override[b]	3.92	0	0.44	4360	pass
as shipped, MPC	3.79	0	0.88	4498	fail
suggested, MPC	3.87	0	0.90	4384	pass

[a] AIC-400, AT-102, both feedforwards, the optimizer at $m_{min} = 1.5$ K, bed 65 °C

[b] as [a] with the sticky-margin override in place of the optimizer

[c] worst margin −0.023 to 0.007 K over three further seeds, one of which passes

10.6.9 The Campaign

The challenge runs 480 min. The intake humidity climbs toward 18 g/kg from 120 min and falls back toward 8 g/kg from 360 min, and from 200 min the concentrate replays an evaporator dip to 45.2 % that recovers to 48 % by 285 min. The target is at least 3.70 t in specification, at most 0.05 t off, no sticky excursion while producing, at most 4450 kJ/kg over the campaign and no CO shutdown. Both the scorecard verdict and challenge credit wait until the full 480 min have elapsed, and the campaign is judged on the case as given, so changing a model parameter or a disturbance voids the attempt. The measured configurations are compared in Table 10.9.

As shipped, on the outlet loop alone at fixed setpoints, the dryer makes 4.00 t, of which 1.98 t is in specification. The powder reaches 4.31 %, the particles are sticky for 191 min from 171 min

with a worst margin of −0.40 K, and 0.90 mm of deposit builds, at 4237 kJ/kg, so the run fails on tonnage, on off-specification powder and on stickiness. The moisture controller alone ships 3.76 t in specification, but the humid hours need a hot outlet and the campaign spends 4498 kJ/kg, over the budget at every seed tried (4486 to 4509 kJ/kg), with the margin down to 0.04 K. The suggested configuration ships 3.88 t, 0.18 t above the target and all of it between 3.44 and 3.56 %, at 4380 kJ/kg, 70 kJ/kg inside the budget, and a worst margin of 0.63 K. Its settings applied before Run and then carried through a Reset give the same trajectory and result. Its optimizer, bound by the estimated margin for more than 479 of the 480 minutes, runs the inlet up to 199.4 °C in the dry hours and down to 175.3 °C in the humid ones. The bed at 65 °C is worth 117 kJ/kg, since the same controls with the bed at 55 °C spend 4496 kJ/kg.

The optimizer's floor is what keeps the true margin positive, and it pays for the optimizer's own speed rather than for the estimator's error. Reading AT-102, the soft sensor's margin stays within 0.15 K of the truth all campaign, yet the worst margin of the suggested run comes at 144 min, with the estimate at 0.62 K, 0.88 K under the floor. The optimizer retreats in proportion to the shortfall, and between 120 and 144 min it lowers the inlet by only 6.9 K while the intake climbs from 8.0 to 15.0 g/kg and takes the true margin from 1.53 to 0.63 K. A floor of 0.8 K ships 4.10 t at 4299 kJ/kg but leaves the true margin at −0.05 K, sticky for 8 min from 138 min. Without the optimizer the same controls ship 3.98 t at 4339 kJ/kg with a worst margin of only 0.007 K at 256 min. That run passes at the default seed, but its worst margin ranges from −0.023 to 0.007 K over three further seeds and two of them fail, so a fixed inlet of 190 °C rides on the sticky line through the humid hours. The suggested run keeps 0.63 to 0.68 K at every seed. The sticky-margin override, a PI loop that retreats as fast as the margin falls, is not the only alternative. With the bed at 65 °C and the override in place of the optimizer, the campaign passes with 3.92 t at 4360 kJ/kg and a worst margin of 0.44 K (0.44 to 0.45 K over the three further seeds), more powder at less energy than the suggested configuration, because the inlet stays near 190 °C in the dry hours and falls only as far as 184.6 °C when the weather turns.

The MPC rows of Table 10.9 are given for information, under the same target and rules. Switched on with nothing else changed, the MPC ships 3.73 to 3.79 t in specification with the margin at 0.88 K or more but spends 4498 kJ/kg, over the budget at every seed (4487 to 4499 kJ/kg), and at seed 2 it also ships 0.058 t off specification. With the suggested settings it passes at all five seeds, with 3.87 to 3.88 t, a worst margin of 0.90 to 0.98 K and 4374 to 4384 kJ/kg, against 0.63 to 0.68 K and 4371 to 4380 kJ/kg for the suggested settings alone.

10.7 Technology in Practice

What Measures It

The loop that steers the dryer reads the exhaust temperature, and the quality that judges it is read by a laboratory. A resistance thermometer in a thermowell is one way to read the exhaust, and it carries two faults that matter here, a lag from its sheath and an offset from its wiring [23].

> **TECHNOLOGY Sensor: A Resistance Thermometer in the Exhaust**
>
> A resistance temperature detector is a winding of fine metal wire, usually platinum, whose resistance rises with temperature almost linearly. A 100 Ω element at 0 °C is very common in industry. It is more fragile than a thermocouple but needs no reference compensation. In a duct it sits in a thermowell, a sheath that lets it be replaced without opening the process but adds to the lag of the reading, and an air gap at the bottom of the sheath adds a second lag. The resistance of two connecting wires adds to the element's and reads as a falsely high temperature [23]. TIC-202 would answer such an offset with more feed and wetter powder while its trend sat on the setpoint.

The chapter gives TT-202 one lag of 5 s and filtered noise of 0.1 K, which stand for the element, its sheath and its transmitter. United States rules for dairy plants ask for as many samples from each

dryer production lot as are needed to assure its composition and quality, tested by official methods of laboratory analysis [3]. The chapter's laboratory samples every 30 min and reports 40 min later.

What Moves It

TIC-202 moves the feed pump, TIC-201 the gas heater of the drying air and TIC-301 the bed heater. The model's heaters follow their commands through lags of 30 and 20 s, and the gas heater loses the fouled share of its output.

> **TECHNOLOGY** **Final Element: The Feed Pump on a Variable-Frequency Drive**
>
> One way to carry out TIC-202's feed command is to set the speed of the concentrate pump. A variable-frequency drive sets the speed of an induction motor by the frequency of the power it applies, and it can limit speed, current, torque, acceleration and deceleration and shut down on overcurrent or overvoltage. Variable-speed machines often cannot change a flow as fast as a control valve, and a stopped pump does not necessarily stop the flow through it [23]. Matching motor speed to the load saved typically 7 to 60 percent of the energy in a range of applications [24]. The model's pump follows its command through one lag of 3 s.

Where the Controller Runs

Loops like these run in a distributed control system. It places a supervisory level over field controllers that hold their process conditions around set points, nests or cascades loops so that one loop sets another's set point, and runs with cycle times from milliseconds to minutes [25]. A 2003 survey of predictive control draws the usual hierarchy. Basic PID loops execute about every second, and above them, about every minute, sits dynamic constraint control, done in the conventional structure by PID algorithms, lead-lag blocks and high and low select logic and in the newer one by a single predictive controller [20]. Here the loops, the soft sensor and FY-500 run every 2 s and the optional MPC every 10 s. The velocity-form PI of Equation (10.7) is the function block of each loop, and the minimum selector ahead of TIC-201, with external reset on the loops it does not select, is the select logic of the conventional structure. Figure 10.7 draws the loops on field devices.

What Is Modernizing It

Three developments bear on this dryer. The first is the soft sensor. A 2013 review of industrial soft sensors prefers a first-principles model wherever one is available with enough accuracy at reasonable modeling effort and computing load, and it reports updating strategies that join a recursive model with an update of the model's output offset [26], the part the laboratory bias plays for AY-401. The second is control for energy. A 2011 guide for dairy plants puts the savings of process control systems at around 5 % or more for many industrial plants without updated systems, and notes that up to 30 % of the heat in a dryer's exhaust can be recovered once the powder it carries is filtered out [24]. The third is predictive control. The 2003 survey found that most industrial predictive controllers carry no process state and feed back through a bias, which it judged probably the best assumption for stable plants without information on the disturbances [20]. The MPC of case 9 carries such an output bias.

> **IN PRACTICE** **Run at the Lowest Outlet the Powder Allows**
>
> A 2011 energy guide for dairy plants states the rule. For the most efficient operation a spray dryer's inlet air should be as hot as possible and its outlet air as cool as possible, within limits set by the product: too hot an inlet damages the product, and too cool an outlet leaves it sticky and too wet. Because the temperature and humidity of the intake air vary while the settings stay fixed, it also recommends measuring them so that operators can adjust the dryer [24]. The humidity feedforward of FF-400 makes that adjustment automatic.

10.8 Limitations

The drying curve, the isotherm, the sticky criterion and the deposit and fire rules are illustrative teaching values rather than correlations fitted to a milk powder, and the thresholds of the fire rule

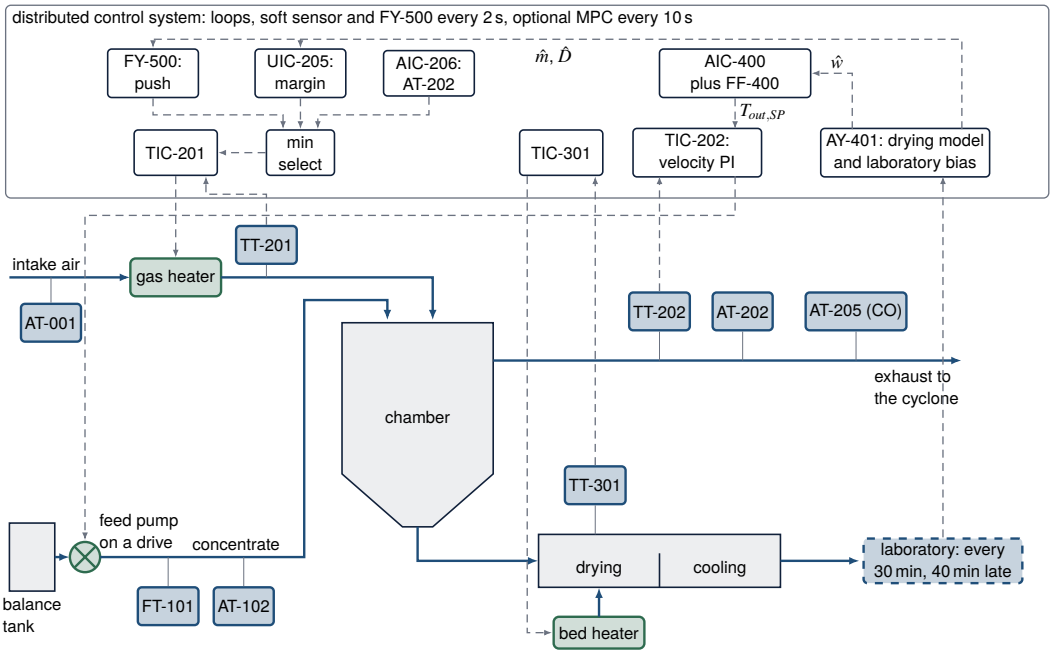

Figure 10.7 The loops of the dryer on plant hardware. TIC-202 holds the exhaust temperature TT-202 by the speed of the feed pump, and AIC-400 with the feedforward FF-400 writes its setpoint from the soft sensor AY-401, which solves the drying model at the measured air and feed (TT-203, AT-202, AT-001, FT-101 and AT-102 are read by it but their lines are not drawn) and takes its bias from the laboratory. The minimum of the push FY-500 and the overrides UIC-205 and AIC-206 sets the inlet that TIC-201 holds with the gas heater, and TIC-301 holds the drying section of the bed. Notice that the laboratory is the only instrument that reads the powder itself.

are not a design basis. The chamber carries one mean particle per cell, with no size distribution, no particle history and no surface stickiness, and the model has no explosion, venting or suppression.

The soft sensor solves the plant's own model, so its only errors come from its inputs, their lags and noise, while a real soft sensor also has structural error for the laboratory bias to absorb. The bias corrects the moisture but not the margin estimate that the overrides and the optimizer read. The optimizer is a constraint pusher with fixed headroom scalings, not a model predictive controller, and the evaporator is a replayed boundary condition rather than a coupled model.

On the laboratory alone, AIC-400 executes once per result, every 30 min. An explicit back-calculation over a step longer than twice its tracking time overcorrects, and with a tracking time of τ_I the integral of a saturated loop grew without bound, to 5.0×10^{12} by 480 min with $\tau_I = 3$ min when the application is rerun with that tracking time. The application now tracks with the longer of τ_I and the time since the last result, which keeps the integral of the same run at or below 425.1 and leaves every shipped case unchanged.

The MPC's model is linear at each decision and identified at one design point, so the sticky margin falls 31 % further after a 5 K inlet step than the plan predicts. The plan reads the same biased estimates as AIC-400 and FY-500, so it cannot ship drier powder before the laboratory reports, it has no handle on the bed and no deposit model beyond a hold on the inlet, and it overrides the scripted inlet steps of cases 6 to 8.

The deposit estimate starts at the scenario's declared initial layer, 3.5 mm in case 7 and zero elsewhere. Two other behaviors limit the shipped structure. UIC-205, which tracks the setpoint applied one scan earlier, is reported as the selected input for 92.9 min of the 240 min of case 9 under

the loops while the estimated margin is 1.44 to 2.99 K against its 0.5 K setpoint, and switching it off adds 0.008 t in specification. The moisture loop's slow response is quantified by the retuned row of Table 10.8.

OPEN CHALLENGE **Measuring What Sticks**

Spray drying still loses product to particles deposited on the walls of the drying chamber [9], and the stickiness behind the deposit reaches no instrument. A 2011 guide for dairy plants names sensors that are inexpensive to install, reliable and able to analyze in real time as an important area for process control, with new sensors still to be developed and demonstrated at commercial scale [24]. In this chapter the laboratory corrects the moisture estimate but never the margin, which stays 0.9 K optimistic in case 4. An online reading of the particles' stickiness, or of the deposit itself, would be progress.

10.9 Conclusion

A milk powder spray dryer is controlled through its exhaust temperature and judged on a moisture its laboratory reports most of an hour later, and this simulator keeps the true powder beside everything its controllers believe. At the design point 12 000 kg/h of air at 190 °C makes 498 kg/h of powder at 3.50 % with the exhaust at 80.0 °C, 4219 kJ per kilogram and the chamber particles 2.93 K inside the sticky line. The same outlet gives 3.85 % when the concentrate weakens to 44 % and 4.10 % with sticky particles when the intake reaches 18 g/kg, and the steady map turns the difference into feedforward gains of 1.1 and 0.8 K. A soft sensor told the wrong solids reads 3.40 % for powder at 3.85 %, and its laboratory bias, learned from results 40 min late, brings the powder back under 3.6 % by 163 min, although the margin it reports stays 0.9 K optimistic. A 225 °C inlet makes 34 % more powder and 2.46 mm of roof deposit in 80 min, and an old deposit ignites even under the override, because the roof reaches 158 °C before the powder turns sticky. The campaign improves from 1.98 t in specification to 3.88 t at 4380 kJ/kg, and a simpler override configuration reaches 3.92 t. An optional MPC of the feed and the inlet makes the same powder as the loops in case 9, 1.896 against 1.893 t, and ships less off specification only by answering the laboratory sooner, 0.444 against 0.570 t, which AIC-400 retuned alone betters at 0.413 t. The model carries mean particles, illustrative drying and stickiness constants and declared fire thresholds. A fuller model would add a particle size distribution, a surface-based sticky criterion, a bias on the margin estimate and a coupled evaporator. The simulator shows that an inferred constraint needs the same care as an inferred quality.

KEY POINTS

- A spray dryer is steered by an exhaust temperature that answers in seconds and judged on a powder moisture that its laboratory reports most of an hour later.
- The same outlet gives 3.85 % when the concentrate weakens to 44 % and 4.10 % with sticky particles when the intake reaches 18 g/kg.
- A soft sensor told the wrong solids reads 3.40 % for powder at 3.85 %, and its laboratory bias, learned from results 40 min late, brings the powder back under 3.6 % by 163 min.
- An inferred constraint needs the same care as an inferred quality.

Where else this applies. Chapter 11 infers an unmeasured core temperature from a model in the same way, and Volume 4, Chapter 6 pushes the line speed of another dryer toward its first binding limit.

References

[1] L. N. Petersen, N. K. Poulsen, H. H. Niemann, C. Utzen, and J. B. Jorgensen, "Economic optimization of spray dryer operation using nonlinear model predictive control," in *53rd IEEE Conference on Decision and Control*, 2014, pp. 6794–6800.

[2] ——, "Application of constrained linear MPC to a spray dryer," in *2014 IEEE Conference on Control Applications (CCA)*, 2014, pp. 2120–2126.

[3] U.S. Department of Agriculture, Agricultural Marketing Service, "Grading and inspection, general specifications for approved plants and standards for grades of dairy products, 7 CFR part 58," Code of Federal Regulations, Title 7, Volume 3, revised as of January 1, 2024, https://www.govinfo.gov/content/pkg/CFR-2024-title7-vol3/xml/CFR-2024-title7-vol3-part58.xml, 2024, sections 58.205(a), 58.240, 58.244 and 58.245.

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

[5] B. Joseph and C. B. Brosilow, "Inferential control of processes: Part I. Steady state analysis and design," *AIChE Journal*, vol. 24, no. 3, pp. 485–492, 1978.

[6] P. Dufour, "Control engineering in drying technology: Review and trends," *Drying Technology*, vol. 24, no. 7, pp. 889–904, 2006.

[7] C. G. J. Baker and K. A. McKenzie, "Energy consumption of industrial spray dryers," *Drying Technology*, vol. 23, no. 1–2, pp. 365–386, 2005.

[8] K. Masters, "Designing fluid bed and spray dryers for low energy consumption," *Drying Technology*, vol. 2, no. 4, pp. 521–549, 1984.

[9] D. Santos, A. C. Maurício, V. Sencadas, J. D. Santos, M. H. Fernandes, and P. S. Gomes, "Spray drying: An overview," in *Biomaterials: Physics and Chemistry, New Edition*. IntechOpen, 2018, pp. 9–35.

[10] R. B. Keey and M. Suzuki, "On the characteristic drying curve," *International Journal of Heat and Mass Transfer*, vol. 17, no. 12, pp. 1455–1464, 1974.

[11] T. A. G. Langrish and T. K. Kockel, "The assessment of a characteristic drying curve for milk powder for use in computational fluid dynamics modelling," *Chemical Engineering Journal*, vol. 84, no. 1, pp. 69–74, 2001.

[12] Y. Roos and M. Karel, "Plasticizing effect of water on thermal behavior and crystallization of amorphous food models," *Journal of Food Science*, vol. 56, no. 1, pp. 38–43, 1991.

[13] M. Gordon and J. S. Taylor, "Ideal copolymers and the second-order transitions of synthetic rubbers. I. Non-crystalline copolymers," *Journal of Applied Chemistry*, vol. 2, no. 9, pp. 493–500, 1952.

[14] C. Hennigs, T. K. Kockel, and T. A. G. Langrish, "New measurements of the sticky behavior of skim milk powder," *Drying Technology*, vol. 19, no. 3–4, pp. 471–484, 2001.

[15] L. Ozmen and T. A. G. Langrish, "An experimental investigation of the wall deposition of milk powder in a pilot-scale spray dryer," *Drying Technology*, vol. 21, no. 7, pp. 1253–1272, 2003.

[16] P. Beever, "Spontaneous ignition of milk powders in a spray-drying plant," *International Journal of Dairy Technology*, vol. 37, no. 2, pp. 68–71, 1984.

[17] D. O'Callaghan and P. Cunningham, "Modern process control techniques in the production of dried milk products – A review," *Le Lait*, vol. 85, no. 4–5, pp. 335–342, 2005.

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

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

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

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

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

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

[24] 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.

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

[26] M. Kano and K. Fujiwara, "Virtual sensing technology in process industries: Trends and challenges revealed by recent industrial applications," *Journal of Chemical Engineering of Japan*, vol. 46, no. 1, pp. 1–17, 2013.

10.10 Exercises

10.1 Concept check · 2 min

In case 3 the evaporator's product falls from 48 to 44 % solids at 10 min. The outlet loop keeps the true outlet within 0.39 K of 80 °C, yet the powder passes 3.8 % at 26 min. Why does a well-held outlet temperature not hold the moisture?

(a) TIC-202 is tuned too slowly for the change, and the outlet error it leaves behind wets the powder

(b) The weaker concentrate cools the inlet air, and TIC-201 cannot restore the heat the air brings

(c) Each droplet brings more water per kilogram of solids, so the same outlet now means a wetter powder

(d) The laboratory reports 40 min late, so the outlet loop works from a moisture that is out of date

10.2 Concept check · 1 min

AY-401 compares each laboratory result with the raw estimate it stored at the instant the sample was drawn, not with its estimate at the moment the result arrives. Why?

(a) The stored value is less noisy, because it is averaged over the 40 min turnaround

(b) The stored values are needed to apply the clamp of ± 2 points to the bias

(c) The result describes the powder at its arrival, and the stored value removes the time spent in the bed

(d) The estimate at arrival already contains the plant's response to earlier corrections, so comparing with it would chase that response

10.3 Concept check · 2 min

UIC-205 protects the sticky margin and AIC-206 the exhaust humidity, and both act by lowering the inlet setpoint rather than by cutting the feed. Why does the inlet carry the overrides?

(a) At a fixed inlet a feed cut heats the particles faster than it dries them, so the margin need not improve

(b) The feed pump saturates at 1500 kg/h, so a feed cut cannot be applied above the design point

(c) The inlet answers within seconds, while the feed pump follows its command with a lag of minutes

(d) A feed cut raises the exhaust humidity, which AIC-206 would then have to undo

10.4 Theory and calculation · 7 min

Use the Gordon–Taylor rule and the margin of Equation (10.5), with $T_{g,s} = 101\,^{\circ}\mathrm{C}$, $T_{g,w} = -135\,^{\circ}\mathrm{C}$, $k_{GT} = 6.7$ and $\Delta T_{st} = 20\,\mathrm{K}$. Moistures are on a wet basis.

(a) Compute the glass transition of the powder at the 3.5 % specification. *Checkpoint:* About 54.9 °C.

(b) At the design point the particles leave the chamber at 4.17 % and 64.8 °C. Compute their glass transition and their sticky margin m. *Checkpoint:* About 47.7 °C and 2.9 K.

(c) At a steady inlet of 225 °C they leave at 4.42 % and 66.2 °C. Compute the margin, and split its change from (b) into the part due to the wetter particle and the part due to the hotter one. *Checkpoint:* About −1.0 K: about −2.5 K from the moisture and −1.4 K from the temperature.

10.5 Theory and calculation · 7 min

The laboratory draws a sample every $T_s = 30\,\mathrm{min}$ and reports it $t_d = 40\,\mathrm{min}$ later, and at reset a sample drawn 30 min earlier reports at 10 min. In case 4 the concentrate weakens at 10 min.

(a) At what time does the first result from a sample drawn after the change arrive, and how long after the change is that? *Checkpoint:* At 70 min, 60 min after the change.

(b) Use Equation (1.8) to find the mean age of a held result, and compare it with the 25 s a particle spends in the chamber and the about 5 min the powder spends in the bed. *Checkpoint:* 55 min, about eleven times the time in the bed.

(c) A constant model error ε is learned as $\varepsilon[1-(1-\lambda)^n]$ after n results. Find the number of results needed to learn 90 % of it at $\lambda = 0.5$ and at $\lambda = 0.2$, and the clock time of that result if the first one to carry the error arrives at 70 min. Use Equation (1.12) to give the standard deviation of the bias in each case for a laboratory noise of 0.05 points. *Checkpoint:* 4 results, 160 min, 0.029 points, and 11 results, 370 min, 0.017 points.

10.6 Theory and calculation · 6 min

FF-400 adds $T_{ff} = K_Y(Y_{001} - 8) + K_S(48 - s_{102})$ to the outlet setpoint, Equation (10.9), with $K_Y = 0.8$ K per g/kg and $K_S = 1.1$ K per point. Take AIC-400 off, so that $T_{base} = 80\,^{\circ}$C, and both feedforwards on.

(a) The concentrate is at 44 % and the intake at 18 g/kg together. Compute the outlet setpoint, and compare the sum of the two terms with the separate shifts of the steady map, Figure 10.2. *Checkpoint:* 92.4 °C, a shift of 12.4 against $4.3 + 8.0 = 12.3$ K.

(b) Near the design point each kelvin of outlet costs 1.42 % of throughput. Estimate the throughput given up in (a). *Checkpoint:* About 18 %.

(c) With the concentrate at 48 % and the intake at the slider's top of 22 g/kg, compute the setpoint FF-400 asks for, compare it with the 90.3 °C that holds 3.5 %, and estimate the moisture with the design slope of 0.085 % per kelvin. What else must change to keep the particles off the sticky line? *Checkpoint:* 91.2 °C, 0.9 K too hot, about 3.42 %. The inlet must come down.

10.7 App challenge · 8 min

Load scenario 1, Manual feed and inlet, with the default seed. Run at 1 min/s to 30 min without touching any control, and export the CSV.

(a) From the columns `Tout_true_C` and `M_true_pct`, read the true outlet temperature at 5, 7 and 30 min, and the time at which the true powder moisture first passes 3.8 %. *Checkpoint:* 80.1, 74.3 and 73.5 °C within 0.1 K. About 9.0 min, 4.0 min after the step.

(b) Read the true powder moisture and `sticky_margin_K` at 30 min, and explain from the holdups of Section 10.3 why the air answers in seconds and the powder in minutes. *Checkpoint:* 4.16 % and 1.83 K, each within one unit of the last digit.

(c) Press Reset and Run. Pause within a minute after the step at 5 min, set Feed rate back to 1000 kg/h and Gas heater to 79.9 %, and run to 30 min. Read the true inlet, the true outlet and the true powder moisture at 30 min. *Checkpoint:* 201.1 °C, about 87.0 °C and 2.95 %, within 0.1 K and 0.01 points.

10.8 App challenge · 9 min

Load scenario 4, Soft sensor and lab correction, with the default seed. Run each part at 60 min/s to 240 min and read the Powder in spec, Powder off spec and Specific energy tiles. Make each change after Reset and before pressing Run.

(a) Run as loaded. From the CSV, also read the peak of `M_true_pct` and its time, and `M_bias_pct` just after the result at 70 min. *Checkpoint:* 1.12 t and 0.57 t, 4990 kJ/kg. Peak 3.97 % at about 68 min. Bias 0.17 points.

(b) Press Reset, set Lab correction to Off on the AY-401 card, and run again. Read the tiles, and the true and estimated moisture at 240 min. *Checkpoint:* 0.19 t and 1.56 t, 4817 kJ/kg. True 3.94 % against an estimate of 3.50 %.

(c) Press Reset, leave Lab correction off, set Feed solids for the soft sensor to AT-102 under Advanced, and run again. Read the tiles and the peak true moisture, and explain why this run needs no laboratory correction. *Checkpoint:* 1.64 t and 0.00 t, 5130 kJ/kg. Peak 3.68 % at about 24 min.

10.9 App challenge · 9 min

Load the challenge, the eight-hour campaign, with the default seed. Run each part at 60 min/s to 480 min and read the Powder in spec, Powder off spec, Worst sticky margin, Specific energy and Campaign target tiles. Reload the scenario before each new configuration and make the change before pressing Run.

(a) Run as loaded. *Checkpoint:* 1.98 t, 2.02 t, −0.4 K, 4237 kJ/kg, not met.

(b) Set the Moisture controller On and change nothing else. *Checkpoint:* 3.76 t, 0.00 t, 0.0 K, 4498 kJ/kg, not met.

(c) Press Reveal suggested tuning, which sets the moisture controller on, the soft sensor on AT-102, both feedforwards and the throughput optimizer on, and the bed at 65 °C. Say which target each of (a) and (b) missed, and where (c) recovers the energy. *Checkpoint:* 3.88 t, 0.00 t, 0.6 K, 4380 kJ/kg, met.

10.10 Application area · 3 min

At the design point the gas heater delivers 561.5 kW and the bed heater 22.0 kW, and the dryer evaporates 503.1 kg/h of water to make 497.9 kg/h of powder.

(a) Compute the heat per kilogram of water evaporated, in GJ/t, and per kilogram of powder, and compare the first with the 4.87 GJ/t average of the energy survey of Section 10.2. *Checkpoint:* 4.18 GJ/t, about 14 % below the average, 4219 kJ/kg of powder.

(b) Case 2 raises the outlet setpoint by 3 K. The powder falls to 476 kg/h and the heat per kilogram rises to 4402 kJ/kg. Explain why plants run at the lowest outlet that keeps the powder in specification, and what one kelvin is worth. *Checkpoint:* About 1.4 % of throughput, some 7 kg/h of powder, per kelvin.

10.11 Application area · 4 min

In case 7 a 3.5 mm layer sits on the roof. Lowering the inlet to 200 °C prevents ignition, while the sticky-margin override, which holds the inlet at 212.1 °C, does not.

(a) Explain from the steady results of Section 10.4 why the override cannot prevent the fire. *Checkpoint:* The roof reaches 158 °C at an inlet of 208.5 °C, below the 215.7 °C at which the particles turn sticky.

(b) With no new deposit, the air hammers remove the layer at $r_{clean} D/(D+0.05)$, Equation (10.5). Estimate the time to take 3.5 mm below the 3 mm of the fire rule, and say what an engineer should do with the dryer meanwhile. *Checkpoint:* About 1.7 h. Keep the inlet below 208.5 °C until the layer is gone or cleaned out.

10.12 Review the agent · 5 min

An AI assistant summarizes case 9 for a plant manager: “On the default seed the MPC ships 0.45 t off specification against 0.63 t for the shipped loops. Its model predicts the moisture ahead of the laboratory, so it cuts the feed before the powder goes wet. Replacing AIC-400 and FY-500 with the MPC is therefore the way to reduce off-specification powder on this dryer.” Find the flaw, and say how you would check it in the app.

CHAPTER 11

Controlling an Unmeasured Product Temperature in a Multi-Zone Cathode Calcination Kiln

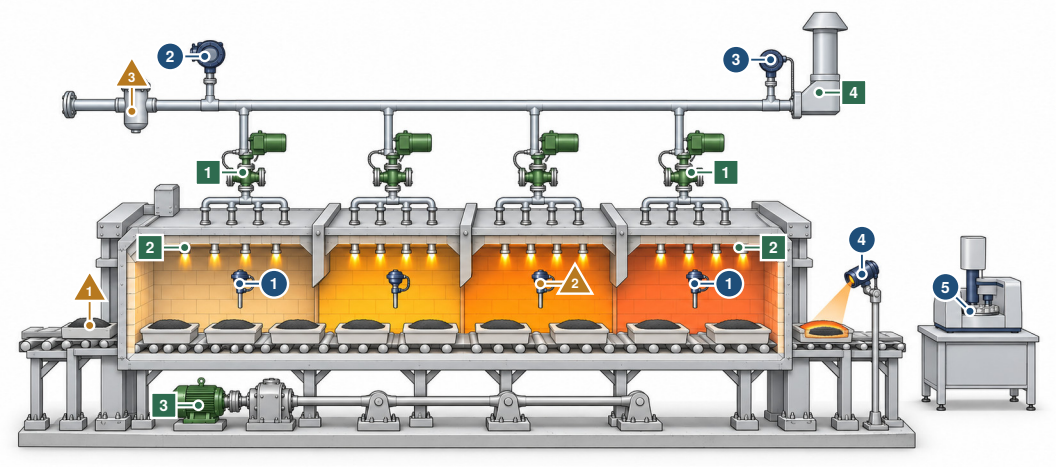

The controller moves 1 gas valves: oxygen share of the blend 2 zone heaters 3 conveyor drive: line speed 4 exhaust damper

The controller reads 1 zone thermocouples, one per zone 2 oxygen in the firing section 3 kiln pressure 4 exit pyrometer on the brick surface 5 laboratory assay, six hours late

What disturbs it 1 denser or wetter bricks 2 a thermocouple that reads high 3 oxygen supply cut back

No instrument reaches the middle of a brick, so the controllers hold the zone readings (1) while the core is estimated from a model and checked only at the exit (4) and in the laboratory (5). Illustration, not to scale.

IN PLAIN TERMS A Kiln Whose Most Important Temperature Is Hidden

The main idea. A long electric oven bakes pressed bricks of powder for more than a day. The baked powder becomes part of a battery. Each section of the oven has its own thermometer and heater. The middle of each brick, where the chemistry happens, has no thermometer at all.

The challenge. The oven can sit right on its settings while the brick centers lag behind, like a thick roast still cold inside a hot oven. The plant must guess the center from a model. A faster belt also leaves less time for the slowest change to finish.

Why it matters. Past a certain belt speed, every brick fails its test.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	11 min	the first page, the history and the fact on the cathode (Section 11.1), the try-it on the biased thermocouple and the try-it on the two cascades, the fact on line speed (Section 11.5) and the key points
🕮	**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 11.5), Technology in Practice (Section 11.6), Limitations (Section 11.7) and the open challenge

The control question A cathode kiln fires compacted bricks for thirty hours, and the core temperature that sets their reactions has no instrument on it. How much of that core do ten well-tuned zone loops control, what does an estimator of the core add, and how fast can the line run before the material leaves the kiln off specification?

Learning objectives

- Compute the parcel mass, residence time and transport delays of the kiln from the conveyor speed, and the surface and core time constants of the brick.
- Identify a zone from a manual step test, and contrast its gain and lag with those of the brick core.
- Explain how a zone loop on setpoint, a biased thermocouple or a surface cascade can leave the product off specification, and name the measurement that reveals each.
- Compare a surface cascade with a core cascade on a soft-sensed temperature, and state which direction of estimator error is dangerous.
- Locate the line speed at which good production peaks, and name the constraint that binds there.

Prerequisites Sections 1.4.4 (delay that moves with flow, material tracking), 1.4.2 (fitted dead time), 1.4.6 (observability), 1.4.8 (soft sensors) and 1.6.2 (velocity-form PID), and Chapter 10 for a quality controlled through an inferred variable.

The app https://apmonitor.com/apps/cathode-kiln/

See also Volume 3, Chapter 2 for choosing the inner measurement of a cascade, Volume 3, Chapter 3 for the speed an inner loop must reach, and Volume 3, Chapter 6 for feedforward from an upstream measurement when feedback sits hours behind the dose.

11.1 Introduction

The thermal processing step that turns a blended precursor into a layered oxide cathode is the most capital-intensive operation in a battery material plant and the one that sets what the plant can make [1]. A nickel-manganese-cobalt hydroxide precursor is blended dry with lithium hydroxide monohydrate, compacted, and fired through a staged thermal cycle under a controlled atmosphere before being cooled, crushed and classified [2, 3]. The properties that a cell eventually reports, the degree of lithiation, the extent of layered ordering, the amount of lithium and nickel site mixing and the quantity of residual lithium left on the particle surface, are all set by the temperature and atmosphere history the material receives inside that kiln [4, 5].

> ⧖ HISTORY **1980: A Layered Oxide Becomes a Cathode**
>
> In 1979 and 1980 John B. Goodenough and his co-workers at Oxford University found that a lithium cobalt oxide could serve as the cathode, the positive electrode, of a lithium battery. Its cobalt dioxide layers are separated by gaps in which lithium ions can be bound without a dramatic expansion of the lattice, and the material showed a very high potential of about 4 to 5 volts against lithium. Paired with a carbon anode, it was the cathode of the commercial lithium battery released in 1991, and new cathode materials have been found since for particular uses [6]. The nickel-rich material this chapter's kiln fires is a layered oxide of the same kind [2].

That history is not measured. A kiln carries a thermocouple in each zone, and a thermocouple reads the zone. Compaction, which is the cheapest capacity a cathode plant can buy because it puts roughly three times as much material in the same machine, is what separates the two: a dense brick has an interior, and that interior is minutes of surface lag and tens of minutes of conduction behind the wall that is heating it. Full-scale modeling of an industrial roller hearth kiln firing ternary precursors reports exactly this separation between the gas and wall temperatures a kiln instrument sees and the temperatures the material containers experience, together with the dependence of that separation on loading and residence time [7].

The class of problem is not confined to batteries. Any continuous furnace that carries a thick charge through a graded profile has an unmeasured product temperature behind a transport delay, and the established answer is a model-based estimate of the material rather than a better measurement of the furnace. Strip temperature in a continuous annealing line was controlled that way from the middle 1990s [8], and the current generation of those controllers is a nonlinear model predictive layer over a distributed thermal model of the strip and the furnace [9, 10]. Kiln control has a parallel history: model predictive control was commissioned on industrial lime kilns in the 1990s against the same combination of long transport delay and slow laboratory feedback [11], and computational fluid dynamics is now routine for the calcination reaction inside them [12]. In both industries the useful measurement of the product is inferred rather than sensed, which places the problem squarely in the soft-sensor literature [13].

Simulation is where the separation can be seen, because on a real kiln the core temperature is exactly as invisible to the engineer as it is to the operator. Assessments of process control education call for practice on the plants that industrial practice actually presents [14, 15], and a cathode kiln is such a plant: a perfectly tuned regulatory layer demonstrably fails to control the variable that matters.

The kiln needs more than a better estimator. It is also the bottleneck of the plant, which makes the conveyor speed the manipulated variable with economic consequence, and every percent of line speed shortens the residence time that the slowest reaction needs. Maximum throughput and maximum saleable throughput are therefore different operating points, and the distance between them is set by a reaction rate inside a brick.

> **✦ DID YOU KNOW? The Heaviest and Costliest Part of the Cell**
>
> The cathode is normally the largest component of a lithium-ion cell by both weight and cost. For commercial nickel-manganese-cobalt cathodes, raw materials make up 75 to 90 percent of the total cost, and further savings can come from the yield and throughput rate of production [2].

This chapter works with a reduced-order kiln model that connects the measurements and manipulated variables of a multi-zone calcination kiln to the inaccessible thermal and reaction history inside a compacted cathode brick, and uses it to quantify the throughput and quality trade-off of a bottleneck unit. Section 11.2 sets out the chemistry of the firing and the shape of the control problem. Section 11.3 develops the model: Lagrangian material parcels, their coupling to the zones, the two-node brick, five reactions, four quality accumulators and the zone, atmosphere and pressure balances. Section 11.4 builds five control layers, from ten zone loops identified by a step test through a throughput feedforward and a cascade on a soft-sensed brick temperature to a supervisory optimizer that sets the line speed. Section 11.5 measures the separation between kiln and product, nine case studies, a sweep of the conveyor speed and the graded challenge. Section 11.6 sets the plant's instruments, heaters and control platform beside the model. Sections 11.7, 11.8 and 11.9 state the limitations, discuss what the results generalize to and summarize them, and Section 11.10 holds the exercises.

11.2 Background

A representative route to nickel-rich cathode active material has five steps. The hydroxide precursor and lithium hydroxide monohydrate are blended dry at a lithium to transition-metal ratio slightly above one, typically 1.03 to 1.05, because lithium is lost during firing. The blend is compacted. The compacts are fired through a long multi-stage thermal cycle. The product is cooled, crushed and classified, and is sometimes washed or surface treated afterward. Published firing cycles for this chemistry span roughly one to one and a half days, so the thirty-hour cycle used here is a design basis for the application rather than a universal recipe.

The chemistry of the firing is not one reaction. Thermal analysis of a nickel-rich hydroxide precursor blended with lithium hydroxide monohydrate resolves distinct events rather than a single

transformation, beginning with the loss of hydrate water near 100 °C, continuing through decomposition of the transition-metal hydroxide near 250 °C and transformation of the lithium hydroxide above 400 °C, and ending with oxidation and lithiation that extend over a much wider temperature range [3]. The layered ordering that follows is slow, strongly activated, and dependent on the oxygen available to it, which is why nickel-rich material is fired under oxygen rather than air and why the oxygen stoichiometry and the defect population respond to the synthesis atmosphere [5].

Two failure modes bracket the firing window from either side. Too little temperature or too little time leaves the layered ordering incomplete and the lithium and nickel sites mixed. Too much of either costs lithium, grows the primary particles and degrades the material in a different way, so the usable window is narrow [2, 4]. A third failure has nothing to do with the firing at all. Residual lithium left on the surface of a nickel-rich particle reacts with atmospheric carbon dioxide and moisture to form lithium carbonate and lithium hydroxide, and it does so preferentially while the product is cooling [16, 17]. The atmosphere control of the cooling section is therefore a quality loop rather than a housekeeping one, and the instrument that finds a fault in it is a dew-point or carbon dioxide analyzer, not a thermocouple.

Against that chemistry, the control problem has the shape of a continuous annealing furnace rather than a stirred tank. The manipulated variables are fast, the measured variables are intermediate, the controlled variable is slow and unmeasured, and the evidence is a laboratory assay taken hours after the material has left. The class was recognized early in steel processing, where model-based strip temperature control replaced furnace temperature control because holding the furnace did not hold the strip [8], and where the current controllers coordinate a distributed thermal model with the line speed in one optimization [10].

11.3 Process Model

The modeled kiln is 40 m long and is divided into ten electrically heated zones of 4 m, each with its own thermocouple, heater and controller, as shown in Figure 11.1. Green compacts travel through it on rollers in eight parallel lanes. The gas space is divided by curtains into three sections that each vent separately, with the gas moving counter-current to the material inside each section.

11.3.1 Variables

Table 11.1 classifies the variables. The ten zone temperatures, the firing-section oxygen fraction and the kiln pressure are measured and controlled. The product temperatures, the reaction extents and the quality index are controlled in the sense that the plant exists to produce them, and none of the three is available to a controller in real time.

🕮 **Pause.** Try Exercise 11.1 before reading on.

11.3.2 Material Transport

Material is carried as $N_C = 40$ Lagrangian parcels, each holding the bricks that occupy one meter of conveyor across all eight lanes, the tracking of material that Section 1.4.4 describes. The parcels share one phase variable rather than forty positions, so their spacing cannot drift, as shown in Equation (11.1).

$$\frac{\mathrm{d}\varphi}{\mathrm{d}t} = v, \qquad x_k = \varphi + k\,p_c + \tfrac{1}{2}p_c, \qquad p_c = \frac{L_K}{N_C} \tag{11.1}$$

where x_k is the center of parcel k and $p_c = 1\,\mathrm{m}$ is the parcel pitch. When φ reaches one pitch the leading parcel has left the kiln, its quality is scored into the campaign ledger, the array is shifted by one index, and a new parcel enters. The new parcel is not given the feed state, because its center is already half a pitch inside the kiln. It is given the feed marched over that half pitch, which makes the traveling wave periodic rather than approximately so.

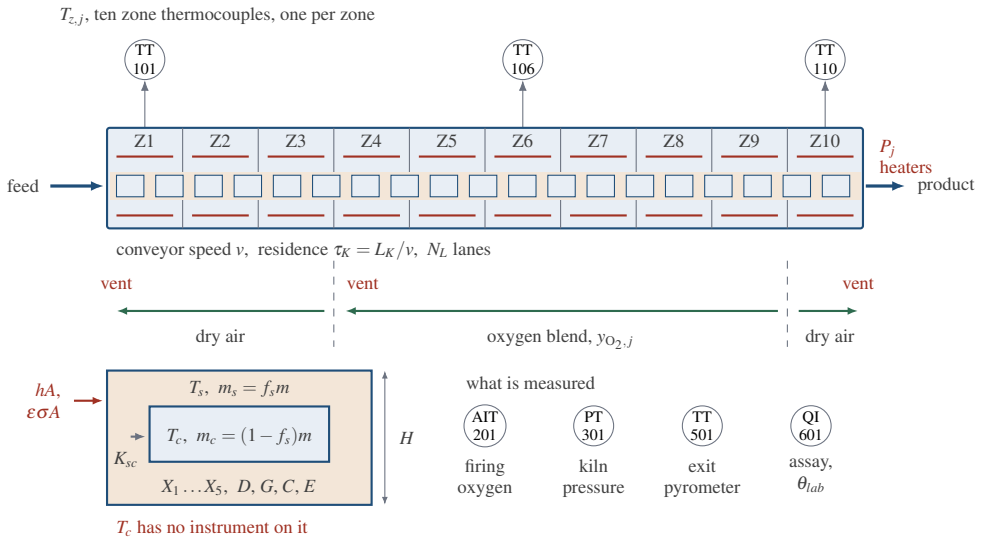

Figure 11.1 Ten-zone cathode calcination kiln, the two-node brick, and the instruments. Each zone carries a thermocouple and a heater, and three gas sections vent separately with the gas counter-current to the material. Ten zone temperatures, one oxygen analyzer, one pressure transmitter, one exit pyrometer and a laboratory assay six hours behind the kiln are the whole instrument list, and the core temperature that sets the reaction rates is not on it.

The mass of a parcel is independent of the line speed, as shown in Equation (11.2), because a faster belt carries the same bricks past faster.

$$m = \frac{N_L \rho_b V_b p_c}{p_b}, \qquad \dot{m} = \frac{m v}{p_c}, \qquad \tau_K = \frac{L_K}{v} \tag{11.2}$$

where $N_L = 8$ lanes, $p_b = 0.34\,\text{m}$ is the brick pitch along the belt and $V_b = 0.0045\,\text{m}^3$ is the brick volume. At $\rho_b = 2500\,\text{kg/m}^3$ each brick is 11.25 kg, each parcel is 264.7 kg, and the design speed gives $\tau_K = 30.0\,\text{h}$ with 353 kg/h of green compact entering. The consequence of the first identity in (11.2) is that raising the line speed is not a larger instantaneous heat load. It is a larger required heating rate, and the zone duty rises only because the material falls further behind the profile.

11.3.3 Coupling Between Parcels and Zones

A ten-zone kiln is not ten isothermal boxes. The muffle, the setters and the gas conduct along the machine, so the environment a brick sees varies smoothly with position while the thermocouple reads it at the middle of its own zone. Each parcel is therefore coupled to the two nearest zone centers in inverse proportion to its distance from them, as shown in Equation (11.3).

$$u = \frac{x_k}{L_z} - \frac{1}{2}, \quad j = \lfloor u \rfloor, \quad w_j = 1 - (u - j), \quad w_{j+1} = u - j \tag{11.3}$$

where $L_z = 4\,\text{m}$ is the zone length. The weighting is a triangular kernel of half-width one zone, sampled on a lattice of one-meter parcels, so it sums to exactly four parcels per zone at every value of φ. Without that property every zone takes a step in loading each time a brick crosses a boundary, and ten zones full of bricks produce a permanent sawtooth on every temperature trend.

Convection sees the linear mean of the two zone temperatures and radiation does not, because the fourth power of an average is not the average of fourth powers. The radiative environment is therefore

Table 11.1 Variable classification

Symbol	Meaning	Units	Range
Controlled, measured			
$T_{z,j}$	zone temperature, $j = 1 \dots 10$	°C	25–880
y_{O_2}	oxygen in the firing section	%	21–98
P	kiln gauge pressure	Pa	−20–80
Controlled, not measured			
T_s	brick surface temperature	°C	25–800
T_c	brick core temperature	°C	25–800
X_4	lithiation extent at the exit	–	0–1
X_5	layered ordering at the exit	–	0–1
Q	CAM quality index	–	0–100
Manipulated			
u_j	zone heater command	%	0–100
f	oxygen share of the gas blend	%	0–100
u_{ex}	exhaust damper	%	5–100
v	conveyor speed	% of design	60–170
$T_{z,j}^{\mathrm{SP}}$	zone setpoint profile	°C	recipe ±
Disturbances			
ρ_b	brick density	×	0.85–1.20
w	free moisture on the feed	×	0.7–1.8
T_{dew}	process-air dew point	°C	−60–20
y_{CO_2}	process-air carbon dioxide	ppm	5–1500
η_6	zone 6 heater rating	×	0.4–1
ϕ_{O_2}	oxygen supply capacity	×	0.4–1
F_{leak}	seal in-leakage	mol/min	0–120
b	thermocouple bias	°C	−25–25

the fourth-power mean of Equation (11.4).

$$T_r = \left[w_j (T_{z,j} + T_0)^4 + w_{j+1} (T_{z,j+1} + T_0)^4 \right]^{1/4} - T_0 \tag{11.4}$$

with $T_0 = 273.15\,\mathrm{K}$. In the annealing ramp, where adjacent zone setpoints differ by 250 °C, using the linear mean in place of (11.4) reports a brick surface tens of degrees too cold and inflates the accumulated gradient penalty more than threefold.

11.3.4 Two-Node Brick

Every parcel carries a surface node holding a fraction $f_s = 0.35$ of the mass and a core node holding the rest, joined by one effective conductance, as shown in Equations (11.5) and (11.6).

$$\begin{aligned} m_s C_p \frac{\mathrm{d}T_s}{\mathrm{d}t} &= hA\,(T_z - T_s) + \varepsilon\sigma A \left[(T_r + T_0)^4 - (T_s + T_0)^4 \right] \\ &\quad - K_{sc}(T_s - T_c) - m_s \textstyle\sum_i \Delta H_i \dot{X}_i \end{aligned} \tag{11.5}$$

$$m_c C_p \frac{\mathrm{d}T_c}{\mathrm{d}t} = K_{sc}(T_s - T_c) - m_c \textstyle\sum_i \Delta H_i \dot{X}_i \tag{11.6}$$

where $m_s = f_s m$ and $m_c = (1 - f_s) m$ are the two node masses, A is the exposed area of the parcel, h the gas convection coefficient, ε the emissivity and σ the Stefan-Boltzmann constant. The conductance

is the lumped Fourier expression of Equation (11.7).

$$K_{sc} = \frac{k_{eff} A_{eff}}{L_{eff}}, \qquad A_{eff} = 2 n_b L W, \qquad L_{eff} = \frac{H}{4} \tag{11.7}$$

where $n_b = 23.5$ is the number of bricks in a parcel, L, W and H are the brick dimensions, and L_{eff} is the distance from a face to the centroid of the half it feeds. With the default geometry $A = 3.18\,\mathrm{m}^2$, $K_{sc} = 6.86\,\mathrm{kJ/(min\,K)}$, the surface time constant is 1.60 min and the core time constant is 22.6 min. At 780 °C the linearized radiative coefficient is 212 W/(m^2 K) against 25 W/(m^2 K) for convection, so above the lithiation temperature the kiln transfers heat by radiation and the gas is present for chemistry rather than for heat.

The consequence of (11.6) is the central behavior of the plant. On a profile that ramps at R degrees per hour, the core settles roughly $0.4R$ degrees below the environment and stays there for the whole ramp. The separation is not a transient and no zone controller can remove it, because the lag is inside the solid and no controller is in that loop.

Pause. Try Exercise 11.2 before reading on.

11.3.5 Reactions

Five reduced-order reactions stand in for the synthesis, as shown in Equation (11.8).

$$\dot{X}_i = k_i(T_c)\,(1 - X_i)\,f_i(y_{\mathrm{O_2}}), \qquad k_i = k_{ref,i}\,\mathrm{e}^{-\frac{E_i}{R}\left(\frac{1}{T} - \frac{1}{T_{ref,i}}\right)} \tag{11.8}$$

where X_1 is dehydration, X_2 the decomposition of the transition-metal hydroxide, X_3 the transformation of lithium hydroxide, X_4 oxidation and lithiation, and X_5 layered ordering. Rates are written against a reference temperature so that each constant in Table 11.2 is a rate at a temperature rather than a bare pre-exponential. Kinetics run on the core temperature, which is what makes the core the variable of interest. The oxygen factors are saturating, as shown in Equation (11.9), because more oxygen stops buying anything once the surface is covered.

$$f_4 = \frac{y_{\mathrm{O_2}}}{K_{\mathrm{O_2}} + y_{\mathrm{O_2}}}, \qquad f_5 = \frac{y_{\mathrm{O_2}}}{K_{\mathrm{O_2},s} + y_{\mathrm{O_2}}}\,X_4 \tag{11.9}$$

Moving from air to 87 % oxygen raises f_5 from 0.46 to 0.78, while moving on from 87 % to 99 % adds three percent more. The dependence of f_5 on X_4 is what makes an under-lithiated brick unable to order however long it is held. The reaction windows are calibrated so that each of the first four is at least fifty times faster at the top of its window than at the bottom, which makes it a switch that completes inside one zone. The fifth is deliberately a slope, eighteenfold over 100 °C, because layered ordering is the reaction that runs out of time when the line speeds up and a switch cannot do that.

11.3.6 Quality Accumulators

Conversion alone does not decide whether material is saleable. Four further states accumulate along the path of each parcel, as shown in Equation (11.10).

$$\begin{aligned} \dot{D} &= k_D\,[\max(0,\, T_s - T_{dam})]^2 \\ \dot{G} &= k_G\,[\max(0,\, |T_s - T_c| - \Delta T_{allow})]^2 \quad \text{for } T_c > T_g \\ \dot{C} &= X_5\, g(T_s)\,(k_{\mathrm{CO_2}} y_{\mathrm{CO_2}} + k_{\mathrm{H_2O}} y_{\mathrm{H_2O}}) \\ \dot{E} &= y_{\mathrm{O_2}}\,\dot{X}_5 \end{aligned} \tag{11.10}$$

where D is overfire damage, G gradient stress accumulated only where the core is hot enough for a gradient to cost product, C surface contamination and E the oxygen exposure weighted by

ordering progress. The contamination sensitivity $g(T_s) = [1 + \exp((T_s - 420)/60)]^{-1}$ switches on as the material cools, which puts the carbonate and hydroxide pickup in the cooling section [16, 17] rather than in the firing zones. The quality index of Equation (11.11) combines them.

$$\begin{aligned} Q = 100 &- w_4(1 - X_4)^2 - w_5(1 - X_5)^2 - w_D D - w_G G \\ &- w_O \left[\max(0, E_{min} - 100E/X_5)\right]^2 - w_C C \end{aligned} \tag{11.11}$$

Material is saleable when $Q \geq 95$, $X_4 \geq 0.99$ and $X_5 \geq 0.97$. The index is a process-quality proxy assembled from quantities this model knows about. It is not a prediction of cell capacity, first-cycle efficiency or cycle life, and the specification numbers on it are teaching defaults.

11.3.7 Zone Energy and Atmosphere Balances

Each zone integrates the energy balance of Equation (11.12).

$$\begin{aligned} C_z \frac{\mathrm{d}T_{z,j}}{\mathrm{d}t} = \eta P_{max} u_j &- UA_j (T_{z,j} - T_a) - Q_{b,j} \\ &- (F_z + F_{lk}) C_{p,g} (T_{z,j} - T_a) \\ &+ f_{rec} F_{in,j} C_{p,g} (T_{z,j+1} - T_{z,j}) - Q_{ax,j} \end{aligned} \tag{11.12}$$

where $Q_{b,j} = \sum_k w_{k,j} q_k$ is the heat drawn by every parcel coupled to the zone, F_z is the purge into it, $F_{in,j}$ is the counter-current gas arriving from the zone downstream, $f_{rec} = 0.25$ is the fraction of that enthalpy that equilibrates with the zone rather than passing through the freeboard, and $Q_{ax,j}$ is a small conductive exchange with the neighbors.

Three gas fractions are integrated per zone, as shown in Equation (11.13).

$$\begin{aligned} n_g \frac{\mathrm{d}y_{s,j}}{\mathrm{d}t} = F_z (y_{s,in} - y_{s,j}) &+ F_{in,j} (y_{s,j+1} - y_{s,j}) \\ &+ G_{s,j} - y_{s,j} \textstyle\sum_s G_{s,j} + F_{lk} (y_{s,a} - y_{s,j}) \end{aligned} \tag{11.13}$$

where $G_{s,j}$ is the generation of species s by the reactions in zone j, being water from reactions 1 to 4, carbon dioxide from residual carbonate, and a negative term for the oxygen taken up by lithiation. The three gas sections each vent separately, so $F_{in,j}$ is zero at a section boundary. With $n_g = 90\,\mathrm{mol}$ per zone and $30\,\mathrm{m^3/h}$ (normal) of purge, the atmosphere settles in under a minute, which makes it the fast subsystem inside a thirty-hour process.

The kiln pressure answers a flow imbalance in well under a second, three orders of magnitude faster than anything else in the model. It is solved algebraically from the flow balance and then lagged by the fan drive, the damper and the transmitter, as shown in Equation (11.14).

$$\begin{aligned} \tau_P \frac{\mathrm{d}P}{\mathrm{d}t} &= P^* - P \\ k_d u_{ex} (P^* + P_{fan}) &= \textstyle\sum_j F_{z,j} + \sum_j G_j + F_{lk} \end{aligned} \tag{11.14}$$

with $F_{lk} = k_L \max(0, -P^*)$. The branch matters: once the algebraic solution goes negative the kiln is pulling room air in through its seals, and the leakage term enters the balance it is the solution of. A kiln below ambient dilutes the oxygen it is paying for and carries ambient moisture to the cooling product, which is why the pressure loop belongs in the quality discussion.

11.3.8 Parameters and Numerics

Table 11.2 lists the parameters that set the behavior discussed here. The state vector has 532 entries: ten zone temperatures, ten heater powers and thirty gas fractions, the conveyor phase and the pressure, and twelve states for each of the forty parcels. Integration is fixed-step fourth-order Runge-Kutta at 0.15 min with a 1.5 min controller scan, and the engine of Section 1.6.1 makes every run repeat exactly.

Table 11.2 Model parameters. Visible parameters are shown in the application and adjustable ones may be changed by the user.

Symbol	Meaning	Value	Status
L_K	heated kiln length	40 m	fixed
N_Z	thermal zones	10	fixed
N_C	material parcels	40	fixed
N_L	parallel lanes	8	adjustable
$L \times W \times H$	brick dimensions	$0.30 \times 0.15 \times 0.10$ m	adjustable
p_b	brick pitch on the belt	0.34 m	adjustable
ρ_b	compacted density	2500 kg/m^3	adjustable
ρ_{loose}	loose blend density	833 kg/m^3	adjustable
C_p	solids heat capacity	0.90 kJ/(kg K)	adjustable
f_s	surface mass fraction	0.35	adjustable
h	gas convection	1.50 kJ/(min m^2 K)	adjustable
ε	brick emissivity	0.80	adjustable
k_{eff}	effective conductivity	0.081 kJ/(min m K)	adjustable
C_z	zone thermal mass	1850 kJ/K	adjustable
UA_z	base zone loss	0.95 kJ/(min K)	adjustable
P_{max}	zone heater rating	6000 kJ/min	adjustable
η	heater coupling	0.92	adjustable
$k_{ref,1}$	dehydration at 100 °C	0.050/min	visible
$k_{ref,4}$	lithiation at 520 °C	0.020/min	visible
$k_{ref,5}$	ordering at 780 °C	0.0145/min	visible
E_5/R	ordering activation	30 000 K	visible
$\Delta H_{1..4}$	reaction enthalpies	267, 134, 78, 47 kJ/kg	visible
T_{dam}	overfire threshold	805 °C	adjustable
ΔT_{allow}	allowable gradient	25 °C	adjustable
T_g	gradient matters above	600 °C	adjustable
n_g	gas hold-up per zone	90 mol	adjustable
τ_P	draft response	3.0 min	adjustable
θ_{lab}	assay delay	360 min	adjustable
Δt	integration step	0.15 min	fixed
Δt_c	controller scan	1.5 min	fixed

The design recipe asks the product to follow 25, 120, 250, 400, 525, 675, 770, 780, 780, 600 and 350 °C at the ten zone boundaries. The zone setpoints that deliver it are not those numbers. They are found by inverting the Lagrangian march so that the brick core meets the recipe at the middle of each zone, and they are 130, 310, 355, 525, 625, 745, 782, 781, 672 and 420 °C. Zone 2 runs 125 °C above its mid-zone product target of 185 °C because the dehydration endotherm is 267 kJ/kg and all of it lands there, and zone 3 then needs little because the endotherm has finished and the counter-current exhaust is doing the work.

11.3.9 Assumptions

Each brick is two thermal nodes joined by one conductance. A real compact has a continuous internal profile, a conductivity that rises as the particles sinter and the porosity closes, and a gas permeability

that decides how quickly oxygen reaches the middle. The single fitted conductance is the largest idealization here, and the model is least trustworthy in the thick, dense, fast-heated case that is also the most interesting one.

Oxygen reaches the core instantly. There is no intra-brick gas diffusion, so a brick too dense to breathe is represented only through its thermal lag. In a real compacted charge the oxygen gradient through the thickness can matter as much as the temperature gradient, and it is the main reason industrial charges are shallow.

Parcel mass is constant through the firing. Water and carbon dioxide leave and oxygen is taken up, which is a 28 % mass loss, but the parcel keeps its green mass in the thermal calculation, which overstates the late-stage heat capacity by roughly a quarter.

Five lumped reactions stand in for the whole synthesis, with constant enthalpies and no dependence on the lithium to transition-metal ratio. The chemistry is nickel-rich NMC and nothing else, with no dopants and no coatings. The gas in each zone is perfectly mixed and axial transport is a directed counter-current purge plus one dispersion coefficient. Cooling below about 420 °C happens outside the modeled kiln, so surface contamination picked up after the discharge is not counted.

11.4 Controller Design

Five control layers are implemented, each adding one element to the one before it. The progression is the argument of this chapter: ten competent zone loops do not control the material, a surface cascade is better and still insufficient, and only an estimator that carries the two-node model reaches the variable that decides the product.

11.4.1 Velocity Form PID Algorithm

All thirteen loops run the velocity-form algorithm of Section 1.6.2, written in this chapter's symbols in Equations (11.15) and (11.16).

$$
\begin{aligned}
e_k &= a\,(\mathrm{SP}_k - \mathrm{PV}_k) \\
\Delta u_k^{P} &= K_c\,(e_k - e_{k-1}) \\
\Delta u_k^{I} &= \frac{K_c}{\tau_I}\,e_k\,\Delta t_c \\
\Delta u_k^{D} &= -\,a\,K_c\,\tau_D\,\frac{\mathrm{PV}_k^{f} - 2\,\mathrm{PV}_{k-1}^{f} + \mathrm{PV}_{k-2}^{f}}{\Delta t_c}
\end{aligned}
\tag{11.15, 11.16}
$$

where $a = \pm 1$ sets the controller action and PV^{f} is the measurement filtered for the derivative path (a 1 min first-order filter on the zone loops, whose derivative time is zero as shipped). The zone and oxygen loops use $a = +1$ and are reverse acting, because both process gains are positive. The pressure loop uses $a = -1$ and is direct acting, because opening the damper lowers the pressure. The demand is clipped to the actuator span, and back-calculation [18] corrects the integral alone, as shown in Equation (11.17).

$$
\begin{aligned}
s_k &= s_{k-1} + \Delta u_k^{P} + \Delta u_k^{D} \\
c_k &= s_k + I_{k-1} \\
u_k &= \min(u_{max},\ \max(u_{min},\ c_k)) \\
I_k &= I_{k-1} + \Delta u_k^{I} + \frac{\Delta t_c}{\tau_t}\,(u_k - c_k)
\end{aligned}
\tag{11.17}
$$

The symbols map onto Equations (1.15) to (1.18) as $s_k = S_k$ and $c_k = v_k$ with no output bias, and the tracking time is $\tau_t = T_t = \tau_I$. At the zone tuning used below each scan therefore removes 6 % of the clipped excess. In the architectures that carry a feedforward the correction of Equation (11.20) is added to u_k after this clip and the sum is clipped to the heater span again, a limit the integral does not see. Manual operation and bumpless transfer are those of Section 1.6.2.

11.4.2 Zone Temperature Loops and Their Identification

Stepping the zone 7 heater from 45.5 to 55.5 % in manual, with the other nine zones in automatic, raises that zone by 9.0 °C within four hours. About 0.4 °C of the rise is drift that the same run shows without the step, because 45.5 % is slightly above the mean duty of the zone, so the step itself is worth 8.5 °C. A least-squares fit of the first-order plus dead-time form of Equation (11.18) to the difference between the two runs over the first 600 min, timed from the controller scan at which the step reaches the heater, returns the constants shown with it.

$$G_z(s) = \frac{K_p\, \mathrm{e}^{-\theta_p s}}{\tau_p s + 1} \qquad K_p = 0.86\,^{\circ}\mathrm{C}/\%, \quad \tau_p = 34.8\,\mathrm{min}, \quad \theta_p < 0.5\,\mathrm{min} \tag{11.18}$$

The gain is small for a heater rated at 100 kW, and the reason is worth reading off the energy balance. The zone does not keep the extra heat. It passes into the bricks, the bricks carry it into zone 8, and the zone 8 controller answers by turning its own heater down. What survives in zone 7 is the extra wall loss, the extra purge enthalpy and a little conduction to the neighbors. The same step moves the brick core by 5.5 °C, reaching 63 % of that change after 63 min, so the product gain is 0.55 °C/% and the product lag is nearly twice the measured one.

The path from heater to zone carries no transport delay, only a 0.6 min heater lag, and the thermocouple adds a 0.9 min lag of its own. The fitted dead time is therefore zero on the zone temperature and about 1 min on the thermocouple reading, a fitted dead time that is not transport in the sense of Section 1.4.2, and the SIMC rule of Equation (11.19) [19] cannot take $\tau_c = \theta_p$ as it would on a plant with real delay.

$$K_c = \frac{\tau_p}{K_p(\tau_c + \theta_p)}, \qquad \tau_I = \min(\tau_p,\ 4(\tau_c + \theta_p)) \tag{11.19}$$

The closed-loop speed is instead limited by the 1.5 min controller scan and by the interaction between neighboring zones, and the application uses $K_c = 4.0\,\%/^{\circ}\mathrm{C}$ with $\tau_I = 25$ min. That gain corresponds to τ_c near 10 min in (11.19), and the reset is shorter than the rule returns because the lag this loop has to remove is the zone rather than the material behind it. Against the cautious settings $K_c = 0.60$ and $\tau_I = 90$ min, an 18 °C setpoint move reaches 95 % of its final value in 22 min rather than 729 min, with 3.6 % overshoot, each response taken against the same run without the step so that the parcel ripple cancels. The core reaches 63 % of its own change in 35 min rather than 121 min, so the retuning makes the core 3.5 times faster while it makes the zone 33 times faster, because the conduction lag inside the brick is untouched by any tuning. In both cases the core moves only about 11.6 °C of the 18 °C the setpoint asked for, because zones 6 and 8 are holding their own temperatures and pulling the material back toward the old profile. Ten zone loops in one kiln are not ten independent loops. The heater clips at 100 % for the first few scans of the fast response, so the tracking time matters here. Absorbing the whole clipped excess in one scan would cut the overshoot to 0.9 % and stretch the 95 % time to 39 min.

11.4.3 Throughput Feedforward

A load cell on the belt and an encoder on the conveyor together measure the mass entering the kiln before any thermocouple moves. The correction that measurement justifies is not a fraction of the zone duty, because most of a zone duty is wall loss and purge enthalpy and the radiant exchange with the bricks is set by a temperature difference rather than by a mass. What scales with the feed rate is the enthalpy the material carries out of each zone, which is the sensible heat of the temperature rise across it plus the reaction heat consumed in it, as shown in Equation (11.20).

$$u_j^{ff} = K_{ff}\, r_j \frac{100\,\dot{m}_0\,[C_p \Delta T_{c,j} + \sum_i \Delta H_i \Delta X_{i,j}]}{\eta P_{max}} \tag{11.20}$$

where $\Delta T_{c,j}$ and $\Delta X_{i,j}$ are read from the design profile across zone j, and r_j is the fractional change of the measured feed rate from its value at the start of the run, the feed the kiln was prefilled with and the profile read at (the design feed in every scenario except the supervisory one, which starts at 90 %), averaged at three points over the window in which the affected material crosses that zone, which runs from jL_z/v to $(j{+}1)L_z/v$ after the measurement. The timing is the part that matters. The transport delay from the load cell to zone 10 is 1710 min at the design speed, so a correction applied to all ten zones at once is compensating nine of them for a disturbance that has not arrived. The windows scale with $1/v$, a delay that varies with flow as in Section 1.4.4. Volume 3, Chapter 6 meets the reverse arrangement, a feedforward from an upstream measurement that replaces feedback sitting hours behind the dose.

11.4.4 Soft Sensor and the Product Cascade

The core temperature estimate is the two-node model of (11.5) and (11.6) without the reaction heat terms and with the linear mean of the two zone temperatures as the radiative environment, marched forward on the *measured* zone temperatures with the estimator's own conductance $\hat{K}_{sc}$ and corrected at the discharge by the exit pyrometer, as shown in Equation (11.21).

$$\hat{T}_{s,N_C} \leftarrow \hat{T}_{s,N_C} + L\left(T_{pyro} - \hat{T}_{s,N_C}\right) \tag{11.21}$$

with $L = 0.15$. Four properties of the estimator follow from this structure. What it adds is the line speed and the brick geometry, which no temperature loop knows. It inherits every bias in the thermocouples that feed it, so it is not an independent check on them. The correction reaches the surface node only, so the core error is almost untouched by it, which is the usual situation with a partially observed state [13] (Section 1.4.6). And the direction of a modeling error decides whether it is dangerous, because an estimator that believes the brick conducts better than it does reports a core that is too hot and the outer loop lowers the firing profile until the reported core is on target. The soft sensors of Section 1.4.8 are corrected by a laboratory bias update, and Chapter 10 holds an inferred moisture that way. This estimator has no such update from the assay.

Pause. Try Exercise 11.3 before reading on.

The outer loop compares the estimate of the material in zone 7 against a product target and trims the setpoints of zones 1 to 8, everything upstream of the cooling section, by a common bias bounded to $\pm 60\,^{\circ}\mathrm{C}$, at $K_c = 1.2\,^{\circ}\mathrm{C}/^{\circ}\mathrm{C}$ and $\tau_I = 180\,\mathrm{min}$. Two variants are provided. The surface cascade holds $\hat{T}_s$ and the core cascade holds $\hat{T}_c$, and the difference between them is the subject of Section 11.5.

11.4.5 Supervisory Optimizer

The supervisory layer owns the conveyor speed and the profile trim. Every 30 min of plant time it marches a brick through each candidate zone profile at each candidate speed using the routine that prefills the kiln, relaxes any zone whose predicted duty exceeds its heater rating toward what that heater can actually hold, and evaluates the constrained problem of Equation (11.22).

$$\begin{aligned} \max_{v,\ \Delta T} \quad & \dot{m}(v) \\ \text{s.t.} \quad & X_{4,f} \geq X_{4,min}, \quad X_{5,f} \geq X_{5,min} \\ & \max |T_s - T_c|_{T_c > T_g} \leq 1.5\,\Delta T_{allow} \\ & E \geq E_{min}, \qquad Q_f \geq Q_{min} + \Delta Q \\ & u_j \leq 100\,\% \ \ \forall j \end{aligned} \tag{11.22}$$

The search is a coarse grid over four profile trims and twelve speeds followed by a two-percent refinement, and the constraint that stops the refinement is reported to the operator as the active limit. When no candidate clears the specification the layer takes the point with the best predicted quality and names the binding constraint, because a constraint the conveyor cannot reach, such as an oxygen

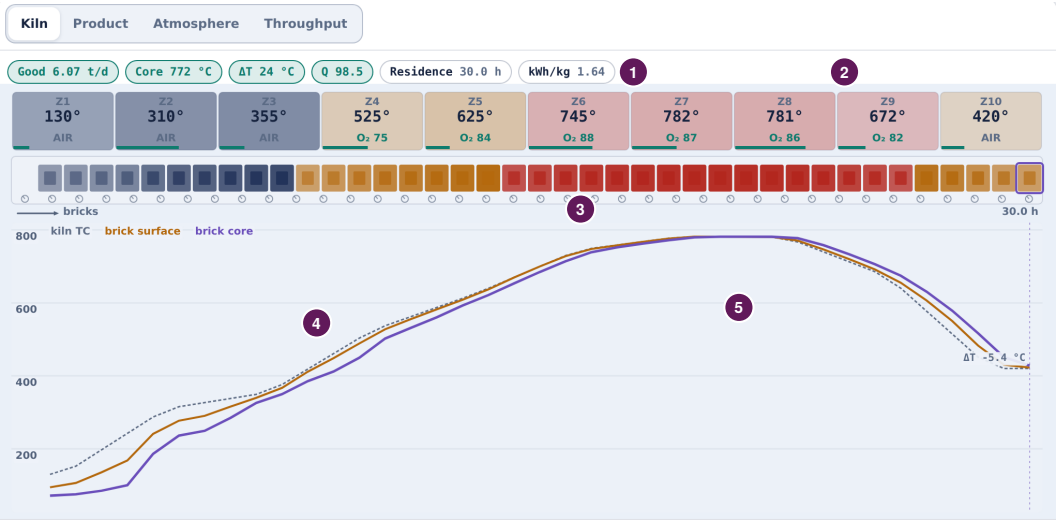

Figure 11.2 The Kiln view of the application at the design point (scenario 5) after 30 h, one residence time, with all ten zone loops in automatic, cropped from the full screen. Callouts: (1) the product tiles, (2) the ten zone tiles with their setpoints and firing-section oxygen, (3) the strip of brick parcels on the belt, (4) the dehydration ramp and (5) the high-temperature hold. The three curves drawn along the kiln are the zone temperature the controllers hold (dotted), the brick surface and the brick core. They coincide through the high-temperature hold and separate wherever the profile ramps, which is the behavior the application exists to show. At this instant the zone 7 trend, cropped from this view, reads 781.9 °C against a 782 °C setpoint while the brick core in that zone is at 772.0 °C.

supply restriction, does not improve by running slower. The core cascade stays in service beneath the layer, which adds its trim to the cascade output and clips the sum to $\pm 60\,^{\circ}\mathrm{C}$, and the conveyor drive ramps toward the chosen speed at 1.2 % of design per minute.

This is a steady-state economic optimization over a dynamic model, evaluated on a slow schedule and handed down to the regulatory layer as setpoints. It is deliberately not a receding-horizon quadratic program of the kind Section 1.6.3 describes, it carries no explicit disturbance model and no state estimator of its own, and it assumes the zone profile it is given will still be there thirty hours later. The supervisor of Chapter 5 has the same structure, a grid search on a forecast from a thermal-history estimator instead of a loop closed on the forecast.

11.5 Results

All results are measured on the application at its shipped settings. Each case starts by selecting its scenario and changing any setting before the run begins, which gives the same run as that setting carried through a Reset (Section 1.6.1), and each run is carried past one full residence time so that the material leaving the kiln reflects the condition being tested. The application ships fifteen scenarios, from open-loop observation through identification and tuning to a graded challenge, and Section 1.6.4 describes how each is credited. The design point (scenario 5) produces 6.07 t/d of cathode active material at a quality index of 98.6, an exit ordering of 0.9862, a worst firing-region gradient of 25.5 °C against an allowance of 25, a busiest heater at 72 % in zone 8, and 1.602 kW h/kg of heater energy per kilogram of product. The screenshot in Figure 11.2 shows that condition.

Table 11.3 Design profile. Zone setpoints, heater loading, firing-section oxygen, and the brick surface and core temperatures at each zone center.

Zone	T_z^{SP} (°C)	u (%)	y_{O_2} (%)	T_s (°C)	T_c (°C)
1	130	17	10	113	78
2	310	63	11	241	194
3	355	25	17	361	336
4	525	45	75	504	474
5	625	24	84	625	607
6	745	64	88	730	719
7	782	45	87	775	772
8	781	72	86	766	770
9	672	28	82	663	681
10	420	22	30	481	503

11.5.1 The Separation Between Kiln and Product

Table 11.3 lists the design profile. Through zones 7 and 8, where the recipe holds flat, the core sits within 11 °C of the thermocouple and within 5 °C of the surface. Through the dehydration and annealing ramps the separation reaches 71 °C, with the dehydration endotherm pulling the core down while the surface stays on the profile. No zone controller is at fault anywhere in the table, and every zone is on setpoint throughout.

The sign of the separation reverses in the cooling section, where the surface leads the core downward. At zone 9 the surface is 18 °C below the core and at zone 10 it is 22 °C below. That reversal is a useful check on the two-node structure, because a model that produced a fixed offset rather than a lag would not show it.

11.5.2 Case Studies

Table 11.4 collects the measured outcomes. Zone 7 integrated absolute error is taken against the *true* zone temperature over the 1800 min following each event, which is one residence time, and from the start of the run where the condition applies from the start (cases A, D and I). Case B is scenario 7 with its cautious tuning, C, E and F are scenarios 9, 10 and 6, G and H are scenario 12 on the two cascades, and D and I are scenario 5 with the oxygen supply or the supervisory layer changed before the run. Figure 11.3 draws the exit ordering, the quality index and the good production rate of each case.

Case B, a heavier load. Ten percent denser bricks are ten percent more tonnes, ten percent more thermal mass and ten percent more reaction endotherm, and the load arrives as a wave that takes a full residence time to reach the exit. With the cautious tuning, adding the feedforward of (11.20) reduces the error summed over all ten zones in the 1800 min after the change from 4641 to 3424 °C min, a reduction of 26 %. With the SIMC tuning it is worth almost nothing (1846 against 1832 °C min), because those loops already hold every zone in the kiln within 0.52 °C, the residual ripple of zone 10 in the undisturbed kiln, and there is nothing left for a feedforward to remove. A feedforward earns its keep against the feedback it is added to rather than against the disturbance.

Case C, a wet feed. Fifty percent more free moisture is a larger endotherm concentrated in the first three zones. The zone loops take it, holding every thermocouple within 0.52 °C, and the material does not: the worst surface-to-core difference anywhere in the kiln rises from 71 to 87 °C and the specific energy rises from 1.602 to 1.632 kW h/kg. A mass feedforward is blind to this disturbance by construction, because moisture does not change the mass on the belt.

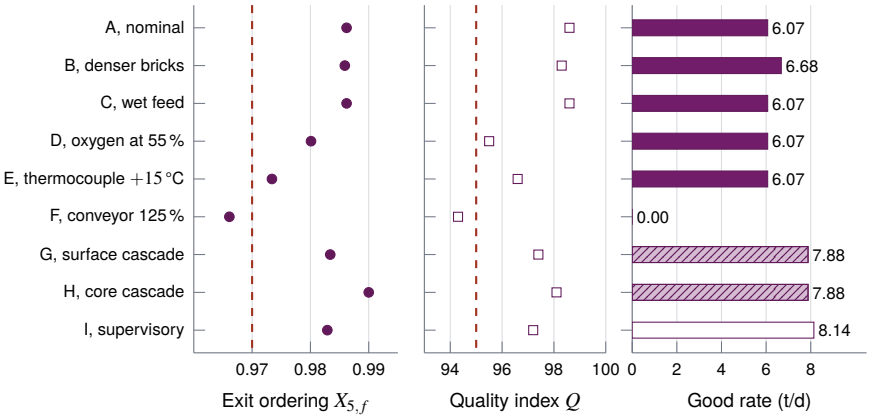

Figure 11.3 The case studies of Table 11.4: the exit ordering (left, dashed line at the 0.970 minimum), the quality index (middle, dashed line at the minimum of 95) and the good production rate (right). Dark bars are the zone loops, hatched bars the two cascades and the white bar the supervisory layer. Notice that in cases D and E every zone reading stays on its setpoint while the ordering falls, and that at the same 7.88 t/d the core cascade orders the material further than the surface cascade.

Table 11.4 Measured case studies. Zone 7 IAE and peak error are against the true zone temperature over the 1800 min after the case's first event (from $t = 0$ for A, D and I), evaluated at the 0.15 min integration step. Exit values are those of the parcel at the discharge at the end of the run. Q uses that same parcel throughout. The app tile instead averages the last sixteen scored exits.

Case	Condition	Zone 7 IAE (°C min)	Peak $\lvert e \rvert$ (°C)	$X_{5,f}$	Worst ΔT (°C)	Q	Good rate (t/d)	SEC (kW h/kg)
A	nominal, 10 zone loops	167	0.27	0.9862	25.5	98.6	6.07	1.602
B	brick density +10 %	262	0.82	0.9859	26.2	98.3	6.68	1.576
C	feed moisture 1.5×	167	0.17	0.9862	24.2	98.6	6.07	1.632
D	oxygen supply at 55 %	167	0.27	0.9801	25.5	95.5	6.07	1.602
E	zone 7 thermocouple +15 °C	26938	15.90	0.9734	24.3	96.6	6.07	1.614
F	conveyor at 125 %	140	0.24	0.9661	31.4	94.3	0.00	1.356
G	surface cascade at 118 %	349	0.57	0.9834	32.8	97.4	7.88	1.409
H	core cascade at 118 %	276	0.45	0.9900	32.9	98.1	7.88	1.417
I	supervisory optimizer	531	19.54	0.9829	31.5	97.2	8.14	1.300

Case D, an atmosphere fault. Restricting the oxygen supply to 55 % of capacity saturates the blend command and lowers the firing-section oxygen. No temperature in the kiln moves. The exit ordering falls from 0.9862 to 0.9801 and the quality index from 98.6 to 95.5, entirely through the oxygen factor f_5 of (11.9).

Case E, a biased thermocouple. A sheath that reads 15 °C high produces the most instructive row in the table. The controller holds the *reading* at setpoint, so the measured error is near zero while

Table 11.5 Conveyor speed against product and production. Good production is the gross rate multiplied by the fraction that meets the specification.

v (%)	τ_K (h)	$X_{5,f}$	Q	Gross (t/d)	Good (t/d)	SEC (kWh/kg)
90	33.33	0.9916	98.86	5.47	5.47	1.754
100	30.00	0.9862	98.62	6.07	6.07	1.602
110	27.27	0.9792	97.55	6.68	6.68	1.494
115	26.09	0.9752	96.80	6.98	6.98	1.440
120	25.00	0.9708	95.58	7.29	7.29	1.404
121	24.79	0.9699	95.34	7.35	0.00	1.394
125	24.00	0.9661	94.27	7.59	0.00	1.358
130	23.08	0.9610	92.62	7.90	0.00	1.316
140	21.43	0.9500	88.03	8.50	0.00	1.252

the error against the true zone temperature integrates to 26 938 °C min over one residence time, 161 times the nominal case over a window of the same length. The exit ordering falls to 0.9734 and the quality index to 96.6. Nothing inside the loop can find this, and the evidence has to come from the exit pyrometer, a second thermocouple, or the assay six hours downstream.

> **TRY IT IN THE APP** **10 · A Thermocouple That Reads 15 °C High**
>
> Load the scenario and press Run. After the thermocouple fault fires, watch the zone 7 strip on the Kiln tab: the dotted TT-107 reading stays on its setpoint while the true zone temperature sits below it. Then open the Product tab as the cold-fired material reaches the exit, where the ordering falls to 0.9734 and the quality index to 96.6.

Cases G and H, two cascades. At 118 % line speed both cascades deliver the same 7.88 t/d, and they do not deliver the same material. The surface cascade reaches an exit ordering of 0.9834 and a quality index of 97.4, while the core cascade reaches 0.9900 and 98.1 for 0.6 % more energy. The surface loop is not mistuned. It is holding the variable it was asked to hold, and the surface-to-core difference is not a fixed offset, so holding the surface does not hold the core. A cascade rejects the disturbances that enter inside its inner measurement, the rule that Volume 3, Chapter 2 demonstrates on a tank, and denser bricks and a faster line both act inside the brick, behind the surface.

> **TRY IT IN THE APP** **12 · A Surface Cascade, and Why It Is Not Enough**
>
> Run the full pass on the surface cascade, then reset, switch to the core cascade and run it again. Both runs end at 7.88 t/d. Compare the ordering at the exit on the Product tab: 0.9834 on the surface cascade and 0.9900 on the core cascade.

11.5.3 Throughput and the Active Constraint

Table 11.5 sweeps the conveyor speed with the recipe unchanged. Gross production rises in direct proportion and the specific energy falls, from 1.754 kW h/kg at 90 % to 1.252 kW h/kg at 140 %, because the fixed losses are spread over more kilograms. Good production, which is the gross rate multiplied by the fraction that meets the specification, has an interior maximum of 7.29 t/d at 120 % and falls to zero at 121, where the exit ordering crosses its 0.970 minimum.

Four constraints could have stopped the line and only one does. The heaters never saturate, the busiest going from 69 % to 76 % across the whole sweep, because the mass in the kiln does not change with line speed and only the required heating rate does. The cooling section is never limiting.

The oxygen exposure is unaffected by speed. What binds is the layered ordering reaction running out of time, with the surface-to-core gradient close behind it. That is the connection the application is built to make, because the capacity of the most expensive machine in the plant turns out to be set by a reaction rate in the middle of a brick.

> ✦ DID YOU KNOW? **One Percent of Line Speed Loses Everything**
>
> From the simulator: at 120 % of design speed the kiln makes 7.29 t/d of good product, its best. At 121 % the exit ordering is 0.9699, just under the 0.970 minimum, and good production is zero (Table 11.5). The specification is a cliff, not a slope.

The supervisory layer of (11.22) does better than any fixed speed, because it buys the lost ordering back with the profile trim before it spends the residence time. Chosen before the run at the design point, it moves the line to 134 % of design speed on its first execution at $t = 0$ and raises the zone 7 setpoint by 19.6 °C at the same instant, 10 °C of trim of its own and 9.6 °C from the proportional action of the core cascade on its first scan, which is the peak error of case I in Table 11.4. Zone 7 is back within 1 °C of its setpoint for good after 55 min. The layer then produces 8.14 t/d of saleable material at a quality index of 97.2 and 1.300 kW h/kg. Against the design point that is 34 % more good production for 19 % less energy per kilogram, and against the best fixed-speed operating point it is 12 % more.

11.5.4 The Graded Challenge

The graded challenge, scenario 15, adds a density change, a wet feed and a derated heater bank to a kiln that starts on the cautious zone loops at design speed. It is judged from 400 to 3600 min: zone 7 must stay within ±3 °C of its setpoint and good production must average at least 7.4 t/d, with the model parameters and disturbances as given. Changing a model parameter or a disturbance voids the attempt. With the shipped settings the worst zone 7 error over the scored window is 2.20 °C, inside the band, but good production averages 6.68 t/d, 0.72 t/d short of the target. The supervisory layer with the SIMC zone tuning raises the average to 8.15 t/d, a gain of 22 %, with a worst error of 0.84 °C, which passes with 0.75 t/d and 2.2 °C to spare, while the specific energy at the end of the 60 h run falls from 1.610 to 1.381 kW h/kg. One change is enough. The conveyor alone at 111 to 116 % on the loops as delivered passes with 7.42 to 7.68 t/d and a worst error near 2.3 °C, while 117 % and faster send the denser, wetter material off specification.

11.6 Technology in Practice

What Measures It

Every zone of the kiln is read by a sheathed thermocouple, and the soft sensor of Section 11.4 is marched on those ten readings. A thermocouple ages in service, and its error stays hidden until the element is taken out and checked.

> TECHNOLOGY **Sensor: A Zone Thermocouple That Drifts**
>
> A thermocouple drifts as its metals change in service, by contamination and oxidation among other causes, and the reading no longer matches the original calibration. Smaller wires and higher temperatures drift faster, and type N has been found about ten times more stable than type K at high temperature. In one set of tests type K drifted by up to about 2 % within 400 h between 1000 and 1100 °C [20]. Checking the transmitter with a simulated millivolt signal does nothing to validate the element itself [21]. On a trend a drifted zone looks perfect, which is case E.

The chapter gives each zone thermocouple a 0.9 min lag and an optional constant bias, the fault of case E. The exit pyrometer reads the brick surface at the discharge, and the laboratory assay arrives 360 min after the material has left. Neither reaches the core.

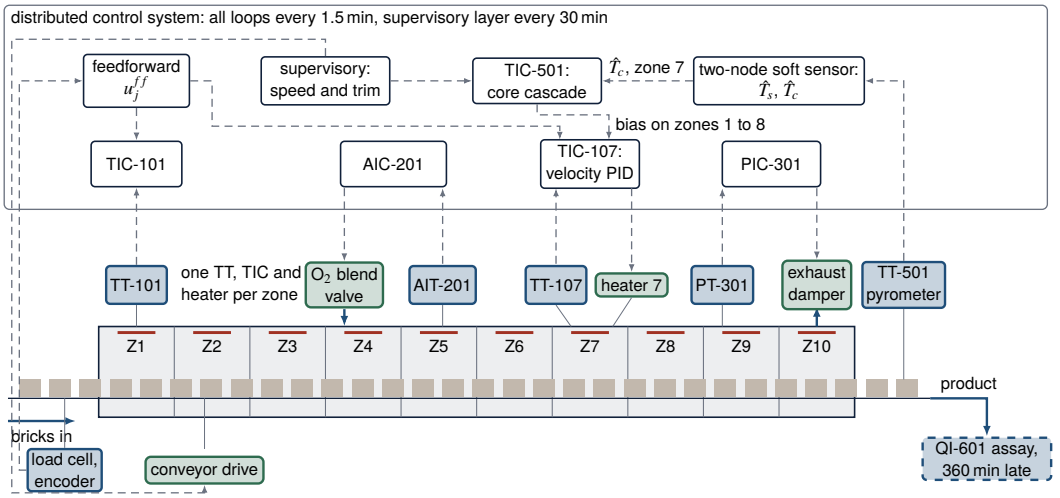

Figure 11.4 The loops of the kiln on plant hardware. Each zone has a thermocouple, a velocity-form PID loop and a heater (zones 1 and 7 drawn), and a feedforward from the load cell and encoder adds to the heater demands as the material reaches each zone. AIC-201 holds the firing-section oxygen with the blend valve and PIC-301 the kiln pressure with the exhaust damper. The soft sensor, marched on the zone readings (lines not drawn) and corrected by the exit pyrometer TT-501, feeds the core cascade TIC-501, which biases the setpoints of zones 1 to 8, and the supervisory layer sets the conveyor speed and a profile trim. Notice that no line reaches the brick core, and that the assay QI-601 corrects nothing.

What Moves It

The zone heaters move the zone temperatures, the oxygen share of the gas blend moves the firing-section atmosphere, the exhaust damper moves the kiln pressure and the conveyor drive moves the residence time. The heater is the element every zone loop acts on, and the model's 100 kW rating is the limit that the supervisory layer relaxes a profile toward.

> **TECHNOLOGY** **Final Element: Resistance Heaters on a Pulsed Output**
>
> In indirect resistance heating an element of a high-resistance material such as silicon carbide or nickel chrome heats the load by radiation, convection or conduction. Heaters that rely on convection are used mainly below about 1250 °F, and those that radiate at higher temperatures. Such heating is precisely controlled and easily automated [22]. An element that can only be switched fully on or off is driven by a time-proportioned output, a pulse train whose duty cycle follows the controller's demand between 0 % and 100 % [21]. The model's heater follows its command through a 0.6 min lag and saturates at its rating.

Where the Controller Runs

Loops like these run in a distributed control system. It places a supervisory level of control over the field controllers, nests or cascades loops so that one loop sets another's set point, and runs with cycle times from milliseconds to minutes [23]. A 2003 survey of predictive control draws the usual hierarchy: basic loops about every second, dynamic constraint control about every minute, and above them a local optimizer, about every hour, that computes an optimal economic steady state and passes it down [24]. Here every loop, the core cascade among them, runs on the 1.5 min controller scan, and the supervisory layer of Equation (11.22) runs every 30 min as such a local optimizer. Each loop is the velocity-form PID of Equation (11.16) as one function block. Figure 11.4 draws the loops on field devices.

What Is Modernizing It

Three developments bear on this kiln. The first is the soft sensor. A 2013 review of industrial soft sensors prefers a first-principles model where one is available with enough accuracy at reasonable modeling effort and computing load, and a survey it reports found the maintenance of models to be the most critical issue in using them [25]. The two-node estimator is such a model, and its conductance is what has to be maintained. The second is measurement. A 2015 sourcebook on process heating notes that emerging technologies can now measure feedstock and melts in real time and that advanced process control lets a heating system be predictive instead of reactive [22]. The third is the sensor itself, where calibration in place, the subject of the first open challenge, would let a drifting zone element be found without taking it out.

> **IN PRACTICE** **Hold the Kiln Slightly Above the Room**
>
> With time most furnaces develop cracks and openings around doors, joints and hearth seals, and a negative pressure draws cold air in through them. A furnace pressure controller reads a pressure gauge in the chamber and regulates the draft fan speed or a damper to hold a slightly positive pressure. Controlled-atmosphere furnaces require positive pressure and special pressure controllers, and steady pressure helps keep the temperature uniform [22]. The chapter's pressure loop does this with the exhaust damper, and below ambient the model's seals leak room air in.

11.7 Limitations

The two-node brick is the limiting idealization. A compact fired from ambient to 780 °C develops an internal profile with more structure than two nodes can carry, its effective conductivity rises as the particles sinter, and the model treats that conductivity as a constant fitted to reproduce a 22.6 min core lag. The application is therefore evidence about the *structure* of the control problem and not about the magnitude of the gradient in any particular brick.

Intra-brick oxygen transport is absent. A dense compact can be oxygen-starved in the middle while the surrounding atmosphere is rich, and that mechanism is the main reason industrial charges are shallow. In this model a brick too dense to breathe shows only a thermal lag, so the density disturbance understates what compaction costs.

The quality index is a process proxy. It is assembled from conversion, overfire exposure, thermal gradient, oxygen exposure and surface contamination, which are the quantities the model knows about, and it says nothing about particle size distribution, crystallite size, cation mixing measured by diffraction, residual lithium by titration, or any quantity an electrochemical cell would report. The specification thresholds are teaching defaults.

The supervisory layer is a steady-state optimization on a slow schedule rather than a receding-horizon controller. It assumes the zone profile it is handed will still be present when the material it is deciding about reaches the discharge thirty hours later, which is exactly the assumption a real implementation would have to relax.

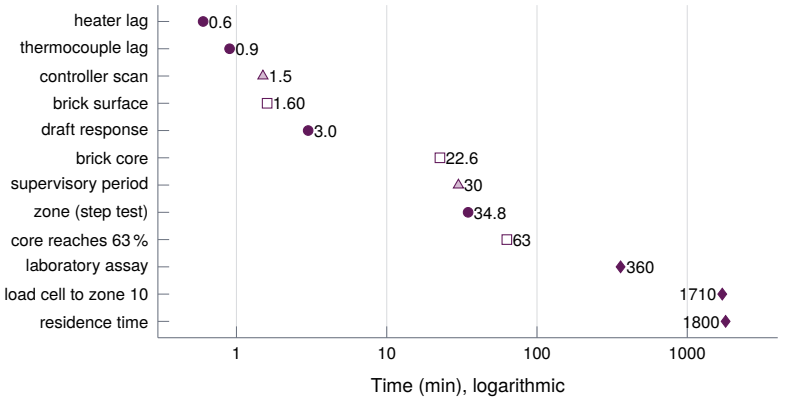

Figure 11.5 The time scales of the kiln, each as the chapter states it. Filled circles are instruments, actuators and the zone, open squares the brick, triangles the control schedule and diamonds the delays of transport and of the laboratory. Notice that the zone loops act on the left of the chart and the product is decided on the right, across more than three decades of time.

> **OPEN CHALLENGE** **A Thermocouple That Reports Its Own Drift**
>
> Unless a thermocouple is removed from time to time for recalibration, nothing during its use shows whether its output is still within specification. Recent work has explored calibrating it in place, against the melting points of reference materials built into the probe or by Johnson noise thermometry [20]. In this chapter a biased zone 7 thermocouple is invisible to its loop and to the soft sensor that runs on it, and only the exit pyrometer, a second element or the assay can find it. A zone element that checks itself in service would close that gap.

11.8 Discussion

The results rest on a separation of time scales between a 1.6 min surface node, a 22.6 min core, a 34.8 min zone and a 30 h transport delay. Figure 11.5 sets these beside the other lags, periods and delays of the chapter. The results in Table 11.4 separate two ideas that are easy to conflate. Cases C, D and E all leave the regulatory layer looking excellent and the product measurably worse, and in each the fault is visible only outside the loop that caused it. That is the generic situation of a plant whose controlled variable is inferred, and it is the reason the soft-sensor literature treats the estimator and the controller as one design rather than two [13].

The second idea is that the estimate is not free. Cases G and H differ by 0.7 points of quality index and 0.0065 of exit ordering, and the better of the two is the one built on a state that is computed rather than sensed. An estimator that believes the brick conducts better than it does reports a core that is too hot, and the outer loop responds by lowering the firing profile until the reported error is zero, which leaves the true core cold with no error visible anywhere in the control system. The application exposes that failure directly by allowing the estimator conductance to be set wrong in either direction.

The throughput result generalizes past cathode kilns. Any unit that is the bottleneck of a plant and whose product quality depends on residence time has an interior maximum in saleable production, and the location of that maximum moves whenever anything about the feed, the utilities or the equipment moves. Locating it requires predicting the quality rather than measuring it, which is the same argument that put model predictive control on lime kilns and annealing furnaces [9, 11]. The model here is deliberately reduced from the computational fluid dynamics that full-scale kiln studies use [7, 12] to something that runs in a browser in real time, which is what makes it usable as a control platform rather than as a design tool.

11.9 Conclusion

This chapter studied a ten-zone cathode calcination kiln as an interactive laboratory for controlling a state that no instrument reaches. The simulator integrates 532 states, ten zone temperatures and heater powers, thirty gas fractions, a conveyor phase and a kiln pressure, and twelve states for each of forty Lagrangian material parcels carrying two temperatures, five reaction extents and four quality accumulators.

A step test identifies the zone loop as a first-order plus dead-time plant with a gain of 0.86 °C/%, a 34.8 min time constant and negligible dead time, while the same step moves the brick core with a gain of 0.55 °C/% and nearly twice the lag. Ten zone loops tuned by SIMC hold every thermocouple within 0.52 °C of a ten percent load change and do not prevent a 15 °C thermocouple bias from costing 2.0 points of quality index, nor an oxygen restriction from costing 3.1, nor a twenty-five percent line speed increase from taking the material off specification entirely.

A core cascade built on a two-node soft sensor reaches an exit ordering of 0.9900 against 0.9834 for the same cascade closed on the brick surface. A supervisory layer that predicts the exit quality raises saleable production from 6.07 to 8.14 t/d while lowering the specific energy from 1.602 to 1.300 kW h/kg. A fuller model would add intra-brick oxygen transport, a sintering-dependent conductivity, a receding-horizon formulation of the supervisory layer, and moving horizon estimation against the delayed assay (Section 1.4.7). The cathode kiln is a plant on which to learn to control what cannot be measured.

KEY POINTS

- The temperature that sets the reactions is in the middle of a brick, and no instrument reaches it.
- On a ramp the core settles below the environment and stays there, and no zone controller can remove the separation.
- A thermocouple 15 °C high leaves its loop looking perfect while the exit ordering falls to 0.9734.
- A core cascade on the soft sensor reaches an exit ordering of 0.9900 against 0.9834 for the same cascade on the surface.
- Good production peaks at 120 % of design speed, where the ordering reaction runs out of time.

Where else this applies. Chapter 5 runs the same kind of supervisor on a thermal-history estimate, and Chapter 10 holds an inferred quality with a laboratory bias update.

References

[1] 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.

[2] W. Li, E. M. Erickson, and A. Manthiram, “High-nickel layered oxide cathodes for lithium-based automotive batteries,” *Nature Energy*, vol. 5, no. 1, pp. 26–34, 2020.

[3] Z. Ye, L. Qiu, W. Yang, Z. Wu, Y. Liu, G. Wang, Y. Song, B. Zhong, and X. Guo, “Nickel-rich layered cathode materials for lithium-ion batteries,” *Chemistry – A European Journal*, vol. 27, no. 13, pp. 4249–4269, 2021.

[4] H. Zhang and J. Zhang, “An overview of modification strategies to improve $LiNi_{0.8}Co_{0.1}Mn_{0.1}O_2$ (NCM811) cathode performance for automotive lithium-ion batteries,” *eTransportation*, vol. 7, p. 100105, 2021.

[5] Q. Xiong, X. Geng, Z. Liu, Y. Yang, H. Wang, P. Long, Y. Yan, H. Li, H. Zeng, S. Luchkin, Y. Lin, J. Wang, and X. He, “Resolving atmosphere-regulated defect cascades in Ni-rich layered cathodes during high-temperature synthesis,” *Energy Storage Materials*, vol. 90, p. 105387, 2026.

[6] The Royal Swedish Academy of Sciences, “Scientific background on the Nobel prize in chemistry 2019: Lithium-ion batteries,” The Royal Swedish Academy of Sciences, Stockholm, 9 October 2019, 2019. [Online]. Available: https://www.nobelprize.org/uploads/2019/10/advanced-chemistryprize2019.pdf

[7] Y. Dai, X. Jin, P. Zhang, P. Li, Z. Shao, M. Yang, C. Cao, and W. Du, “Numerical study on calcination of Ni-Co-Mn ternary precursors in an industrial roller hearth kiln,” *Chemical Engineering Research and Design*, vol. 231, pp. 169–180, 2026.

[8] N. Yoshitani and A. Hasegawa, "Model-based control of strip temperature for the heating furnace in continuous annealing," *IEEE Transactions on Control Systems Technology*, vol. 6, no. 2, pp. 146–156, 1998.

[9] M. Niederer, S. Strommer, A. Steinboeck, and A. Kugi, "Nonlinear model predictive control of the strip temperature in an annealing furnace," *Journal of Process Control*, vol. 48, pp. 1–13, 2016.

[10] S. Strommer, M. Niederer, A. Steinboeck, and A. Kugi, "Hierarchical nonlinear optimization-based controller of a continuous strip annealing furnace," *Control Engineering Practice*, vol. 73, pp. 40–55, 2018.

[11] J. Valiquette, M. Savoie, and M. Leclerc, "Practical aspects of model predictive control implementation on an industrial lime kiln," in *1997 European Control Conference (ECC)*. IEEE, 1997, pp. 2711–2715.

[12] J. Ryan, M. Bussmann, and N. DeMartini, "CFD modelling of calcination in a rotary lime kiln," *Processes*, vol. 10, no. 8, p. 1516, 2022.

[13] 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.

[14] 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.

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

[16] W. M. Seong, Y. Kim, and A. Manthiram, "Impact of residual lithium on the adoption of high-nickel layered oxide cathodes for lithium-ion batteries," *Chemistry of Materials*, vol. 32, no. 22, pp. 9479–9489, 2020.

[17] Q. Tian, R. Song, J. Zhang, Y. Chen, C. Cui, C. Ma, M. Su, and W. Hu, "Formation mechanism and removal strategy of residual lithium compounds on nickel-rich cathode materials," *Progress in Natural Science: Materials International*, vol. 34, no. 6, pp. 1158–1166, 2024.

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

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

[20] R. J. Chesser and N. D. Bull Ezell, "Thermocouple and RTD drift due to neutron irradiation and temperature for space nuclear applications: A review," *IEEE Sensors Journal*, vol. 25, no. 22, pp. 40 965–40 977, 2025.

[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] U.S. Department of Energy, Advanced Manufacturing Office, "Improving process heating system performance: A sourcebook for industry," U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, Washington, DC, Third edition, report DOE/EE-1339, 2015, oSTI ID 1247916. [Online]. Available: https://www.osti.gov/biblio/1247916

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

[24] 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.

[25] M. Kano and K. Fujiwara, "Virtual sensing technology in process industries: Trends and challenges revealed by recent industrial applications," *Journal of Chemical Engineering of Japan*, vol. 46, no. 1, pp. 1–17, 2013.

11.10 Exercises

11.1 ○ Concept check · 1 min

Which of the following quantities is measured on the kiln and available to a controller as the run proceeds?

(a) The brick core temperature T_c in zone 7

(b) The lithiation extent X_4 of the material at the exit

(c) The zone 7 temperature $T_{z,7}$, read by its thermocouple

(d) The quality index Q of the material leaving the kiln

11.2 Concept check · 2 min

Equation (11.6) is the energy balance on the brick core. With the reaction term left out, what sets the time constant with which the core follows its surface, and why can no zone loop shorten it?

(a) The zone thermal mass over its loss, C_z/UA_z, which a higher zone gain shortens

(b) The 1.5 min controller scan, which a faster scan of the zone loops would remove

(c) The 0.9 min thermocouple lag, which derivative action on the zone loops would cancel

(d) The ratio $m_c C_p/K_{sc}$, a property of the brick that no controller acts on

11.3 Concept check · 2 min

The zone 7 thermocouple of the kiln at its design point begins to read 15 °C high, as in scenario 10. The zone loops and the soft sensor are otherwise unchanged. Which statement is right?

(a) The zone runs about 15 °C cold, the ordering rate falls by about a third, and the evidence must come from outside the loop

(b) The zone 7 error against the reading grows at once, so the loop's own scorecard reveals the fault within the hour

(c) The soft sensor sees the core run cold, so the core cascade, if selected, restores the true core temperature

(d) The zone runs about 15 °C cold, but the ordering rate barely changes because the reaction saturates there

11.4 Theory and calculation · 5 min

The kiln of Table 11.2 carries $N_L = 8$ lanes of bricks of 0.0045 m^3 at 2500 kg/m^3, spaced 0.34 m apart along the belt. A parcel is 1 m of belt, the kiln is 40 m long in ten zones of 4 m, and the design residence time is 30.0 h.

(a) Find the brick mass, the number of bricks in a parcel and the parcel mass. *Checkpoint:* 11.25 kg, 23.5 and 264.7 kg.

(b) Find the design conveyor speed and the green feed rate. With the 28 % mass loss of the firing, estimate the product rate in tonnes a day. *Checkpoint:* 1.333 m/h, 353 kg/h and about 6.10 t/d.

(c) Repeat the residence time and the feed rate at 125 % of design speed, and find at both speeds the center of the window in which material weighed at the load cell crosses zone 10. *Checkpoint:* 24.0 h and 441 kg/h. 1710 min at design and 1368 min at 125 %.

11.5 Theory and calculation · 8 min

Use the brick data of Table 11.2: $L \times W \times H = 0.30 \times 0.15 \times 0.10$ m, $k_{eff} = 0.081$ kJ/(min m K), $C_p = 0.90$ kJ/(kg K), $f_s = 0.35$, with a parcel of $m = 264.7$ kg holding $n_b = 23.53$ bricks, an exposed area $A = 3.18$ m^2, a convection coefficient $h = 1.50$ kJ/(min m^2 K) and a linearized radiative coefficient of 212 W/(m^2 K) at 780 °C.

(a) Compute K_{sc} from Equation (11.7). *Checkpoint:* 6.86 kJ/(min K).

(b) Compute the core time constant from Equation (11.6), and the surface time constant from Equation (11.5) linearized at 780 °C with the zone and the core held fixed. Leave out the reaction terms. *Checkpoint:* 22.6 min and 1.60 min.

(c) A first-order node driven by an input that ramps at a steady rate R settles a fixed $R\tau$ behind it. Use this to recover the $0.4R$ of Section 11.3 with R in degrees per hour. The recipe falls from 780 to 600 °C across zone 9. Find R there and the surface-to-core difference $R\tau_c$ at design speed and at 125 %. *Checkpoint:* $0.40R$. 60 °C/h and 22.6 °C. 75 °C/h and 28.2 °C.

11.6 Theory and calculation · 6 min

Use the zone model of Equation (11.18) and the SIMC rule of Equation (11.19).

(a) Take $\theta_p = 0$ and $\tau_c = 10\,\text{min}$. Find K_c and τ_I and compare them with the settings the app uses. *Checkpoint:* 4.05 %/°C and 34.8 min, against 4.0 and 25 min.

(b) An engineer takes the 1 min fitted on the thermocouple reading as θ_p and sets $\tau_c = \theta_p$. Find the settings, and give two reasons from the chapter why they would not work on this kiln. *Checkpoint:* 20.2 %/°C and 8 min.

(c) The same heater step moves the core by 0.55 °C/%. Predict how far the core in zone 7 moves when the zone 7 setpoint is raised 18 °C and the zone loops hold their setpoints. *Checkpoint:* About 11.5 °C.

11.7 App challenge · 10 min

Load scenario 2 · Step one zone and identify it (MANUAL), with the default seed. Before pressing Run, set Zone 7 heater to 55.5 %. Run to 240 min at a speed of 60, which takes 4 s, and export the CSV file.

(a) From columns `z7_C` and `z7_core_C`, find the rise of the zone and of the core between the first row and the last row before 240 min, and the first row at which each has covered 63.2 % of its rise. *Checkpoint:* Zone 8.95 °C at 31.5 min, core 6.09 °C at 60.9 min, within 0.05 °C and one row.

(b) Press the reload button, change nothing, run to 240 min again and read the same two columns at the end. Find the part of each rise that the step caused, and the two gains. *Checkpoint:* Drift 0.43 and 0.58 °C. Step 8.52 and 5.50 °C, gains 0.85 and 0.55 °C/%.

(c) Repeat the 63.2 % times on the difference between the two runs and compare them with Equation (11.18). *Checkpoint:* 34.6 min for the zone and 63.0 min for the core.

11.8 App challenge · 8 min

Load scenario 6 · The kiln is on setpoint and the core is not, with the default seed. The conveyor steps from 100 to 125 % at 30 min. Run to 1800 min at a speed of 120, which takes 15 s, and export the CSV file.

(a) Find the largest deviation of `z7_C` from `z7_sp_C` after 30 min. *Checkpoint:* 0.24 °C, within 0.02 °C.

(b) At 1800 min read the tiles Ordering at the exit, CAM quality index and Good production, and the column `gross_tpd`. *Checkpoint:* 96.6 %, 94.2, 0.00 t/d and 7.59 t/d.

(c) Find the first row with `quality_exit` below 95. Material that entered after the speed change needs 24 h to reach the exit. Explain why the exit fails much sooner. *Checkpoint:* 692 min, within one row. That parcel was 21.6 m from the inlet at the step, ahead of the hold.

11.9 App challenge · 9 min

Load scenario 12 · A surface cascade, and why it is not enough, with the default seed. The bricks become 10 % denser at 60 min and the conveyor steps to 118 % at 900 min. Run each case to 3600 min at a speed of 120, which takes 30 s, and export the CSV file.

(a) Run the scenario as loaded, on + Surface cascade. Read the tiles Ordering at the exit and CAM quality index, and the columns `profile_bias_C`, `z7_surf_C` and `z7_core_C` at 3600 min. *Checkpoint:* 98.3 % and 97.3. Trim 5.2 °C, surface 780.0 °C, core 775.7 °C.

(b) Press Reset, set Control architecture to + Core cascade and run again. *Checkpoint:* 99.0 % and 97.9. Trim 9.6 °C, surface 784.5 °C, core 780.1 °C.

(c) Explain why the surface cascade leaves the core about 4.3 °C short, and name the estimator error that would make the core cascade the worse of the two. *Checkpoint:* The offset grows with $m_c C_p / K_{sc}$ and with the heating rate. A conductance $\hat{K}_{sc}$ set too high.

Accept temperatures within 0.1 °C.

11.10 Application area · 3 min

Equation (11.10) makes the surface contamination of a parcel grow in proportion to $g(T_s) = [1 + \exp((T_s - 420)/60)]^{-1}$.

(a) Evaluate g at 780, 600, 420 and 300 °C. *Checkpoint:* 0.0025, 0.047, 0.50 and 0.88.

(b) Say where in the plant residual lithium picks up carbonate and hydroxide, which instrument watches for it, and why the app understates the pickup. *Checkpoint:* The cooling section and below. A dew-point or carbon dioxide analyzer. The model stops near 420 °C.

11.11 Application area · 4 min

A plant manager weighs line speeds with Table 11.5, taking the specific energy per kilogram of gross product.

(a) Find the heater energy per day at 100, 120 and 125 %, and for the supervisory layer at 8.14 t/d and 1.300 kW h/kg. *Checkpoint:* 9724, 10235, 10307 and 10 582 kW h a day.

(b) Find the energy per extra tonne of good product from 100 to 120 %, compare it with the average, and say what 125 % buys. *Checkpoint:* About 419 kW h/t against 1602.

11.12 Review the agent · 5 min

An AI assistant is asked how to pass the graded challenge of Section 11.5.4 and replies as follows. “At the design point the busiest heater runs at 72 %, so the kiln has 28 % of heater headroom. Heater duty scales with throughput, so the line can run at $100/72 = 139$ % before a heater saturates. Set the conveyor to 130 %: gross production rises to about 7.9 t/d at design density, well above the 7.4 t/d target, and the zone loops will hold zone 7 within ± 3 °C.” Find the flaw and say how you would check it.

CHAPTER 12

Stored Ammonia and a Sensor That Reads Both Ways in Diesel Urea Dosing

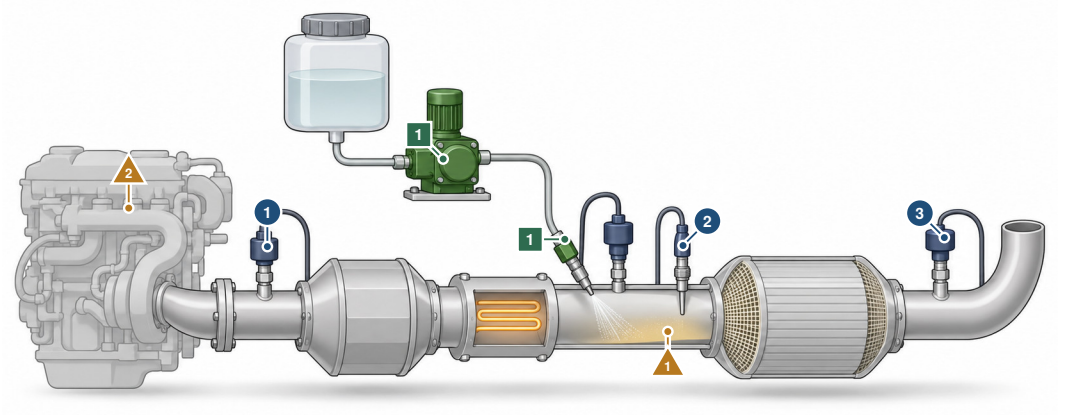

The controller moves 1 urea pump and injector

The controller reads 1 nitrogen oxides leaving the engine 2 gas temperature at the catalyst inlet 3 tailpipe sensor, which also counts ammonia

What disturbs it 1 urea deposits when the gas is too cool 2 engine load: exhaust flow, heat and nitrogen oxides

The ammonia stored in the catalyst brick sets the cleaning but has no instrument, so the controller must infer it from readings (1 to 3) while the tailpipe sensor (3) cannot tell nitrogen oxides from ammonia. Illustration, not to scale.

IN PLAIN TERMS Dosing a Catalyst That Hides What It Holds

The main idea. A diesel truck sprays urea in water into its hot exhaust, where it turns into ammonia. A catalyst then uses the ammonia to turn harmful nitrogen oxides into nitrogen. The catalyst soaks up the ammonia first, like a sponge holding water, and the amount it holds does the cleaning.

The challenge. No gauge reads that amount, and a hot catalyst holds less, so a climb can push the extra out. The one sensor at the tailpipe counts this ammonia as if it were nitrogen oxides. A high reading can mean too little urea or too much.

Why it matters. A controller that trusts that reading alone can keep adding urea while ammonia pours out.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	the first page, the fact on the limits (Section 12.1), the history at the end of Section 12.2, the try-it on the heat-up release and the try-it on the trim loop, the fact on the challenge (Section 12.5) and the key points
🕮	**Study**	50 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 12.5), Technology in Practice (Section 12.7), Limitations (Section 12.8) and the two open challenges

The control question An SCR catalyst converts nitrogen oxides with ammonia it has stored, a store that no instrument measures and whose capacity falls as the brick heats, while the tailpipe sensor reads ammonia as nitrogen oxides. How can a doser hold a cycle's NO_x limit under an ammonia cap when a high reading may mean too little urea or too much?

Learning objectives

- Relate the dose ratio to the stored ammonia and the conversion, and explain why the conversion takes minutes to follow a dose.
- Split a tailpipe sensor reading into its NO_x and ammonia parts, and locate the bottom of the sensor's V.
- Predict a heat-up release from the store and the capacity, and explain why a storage setpoint read at the inlet lowers it.
- Explain why a trim loop on the tailpipe sensor runs away past the bottom of the V, and how an observer's gradient settles the sign.
- Compare feedforward, trim, storage control and MPC on cycle NO_x, slip peak and DEF, and say what each verdict depends on.

Prerequisites Sections 1.4.6 and 1.4.7 (observers and bias updates), 1.6.2 (the trim and fill loops) and 1.6.4 (cases and the challenge), with 1.6.3 for case 9. Chapter 13 is the companion chapter on estimators.

The app https://apmonitor.com/apps/diesel-scr/

See also Volume 7, Chapter 4 for a loop whose process gain reverses at the peak of the plant's curve, Volume 7, Chapter 2 for Arrhenius kinetics that change a plant's character within a few kelvin, and Volume 7, Chapter 3 for a plant whose linearization moves with the operating point.

12.1 Introduction

Ten minutes of city crawl leave the catalyst of the simulated truck near 220 °C with 11.75 g of ammonia stored on it. At 10:00 the truck pulls onto a 2.5 % grade, the inlet thermocouple passes 350 °C 35 s later, and the heating brick can no longer hold what it stored. At 705 s the ammonia leaving the tailpipe peaks at 384.5 ppm while the nitrogen oxides there have fallen to 0.1 ppm. The tailpipe sensor reads 267.5 ppm, which a controller would read as a conversion of only 47.6 % and answer with more urea. The plant has too much ammonia and the measurement says it has too little.

Urea selective catalytic reduction (SCR) was proposed for automotive diesel engines from a technique long established on power plant stacks [1–3], and heavy-duty regulation has since tightened the nitrogen oxide (NO_x) limits again [4, 5]. The control problem it poses belongs to a wider class, in which a reactant held on a surface sets the conversion, the inventory is inferred rather than measured, and the capacity to hold it depends on the operating condition. The three-way catalyst of a gasoline engine, whose conversion depends on the air-to-fuel ratio, the exhaust velocity and the catalyst temperature, is the older automotive member of the class [6]. The SCR case adds a measurement whose sign is ambiguous, because the NO_x sensor at the tailpipe also responds to ammonia [3, 7].

+ DID YOU KNOW? A Limit Cut Again for 2027

From model year 2027 a heavy-duty compression-ignition engine may emit 35 mg of NO_x per horsepower-hour over the transient Federal Test Procedure and the Supplemental Emission Test, and 50 mg over the new Low Load Cycle [8]. The standard set in 2001 was 0.20 g/bhp–hr [9].

This chapter sets out to dislodge two beliefs. The first is that a dosing ratio is a conversion. In the model the conversion follows the stored ammonia, which takes 11.8 minutes to reach 94 % conversion after the dose is raised from 0.8 to 1.0 ammonia per NO_x on a steady highway. The second is that a high tailpipe reading calls for more urea. The sensor's reading against the dose is a V with its bottom near one ammonia per NO_x, and a feedback loop on it pushes the right way only on the left side. A block diagram with a gain and a lag shows neither the stored inventory nor the sign that changes with the state.

The published work treats these elements one at a time. Control-oriented models establish the storage state and its observability [10, 11], kinetic studies model the catalyst [12], and control studies compare open loop, sensor feedback and coverage control [3, 7, 13]. This chapter works with a browser model of a whole SCR line on a transient cycle, in which the true stored ammonia, the true tailpipe NO_x and ammonia and the sensor reading are logged side by side, and measures on it how a feedforward, a sensor trim and a storage controller on an observer meet a cycle limit with an ammonia cap. The sign ambiguity and its resolution by a model are documented in the literature [7, 14, 15], and the observer used here, a model copy corrected along the gradient of its own predicted reading, is a simple member of that family. The model has nineteen states, its ammonia, urea and NO_x ledgers close to round-off, and every parameter is stated so that it can be rebuilt.

Section 12.2 traces the technology from power plant stacks to trucks. Section 12.3 sets out the exhaust line: the engine maps, the urea and its decomposition floor, the three cells of stored ammonia and the tailpipe sensor that reads ammonia as NO_x. Section 12.4 describes the doser: the feedforward, the trim on the tailpipe sensor, the observer of the stored ammonia, the storage loop on its temperature schedule and the added predictive controller. Section 12.5 measures the nine cases and the graded cycle, Section 12.6 sorts the evidence behind the model, Section 12.7 sets the real sensors, dosing hardware and control units beside the model, Section 12.8 states the limitations, Section 12.9 draws the conclusions, and the exercises of Section 12.10 work through the cases.

12.2 From Power Plant Stacks to Truck Tailpipes

Selective reduction of nitric oxide by ammonia over a catalyst was first a technology for power plants [3], and the present status and perspectives of its catalysis were reviewed in 2001 [2]. A vehicle carries the reagent instead as an aqueous urea solution that decomposes in the hot exhaust, a route presented in 2000 as a promising technique to reduce NO_x from automotive diesel engines [1]. In 2006 a control study for a mobile heavy-duty application observed that SCR was well known for power plants and would soon be introduced on European heavy-duty diesels, and that the dynamic operation of a vehicle demanded advanced control [3]. By 2014 the automotive technology had an edited book of its own [16].

The catalysts moved toward copper zeolites. A kinetic model of ammonia SCR over Cu-ZSM-5 was published in 2008 [12], and a copper chabazite, Cu-SSZ-13, was reported in 2010 with excellent activity and selectivity for the reaction [17]. Such catalysts hold ammonia that leaves as they heat. On one commercial Cu-chabazite monolith the total uptake fell from 4.6 to 1.6 g/L as the adsorption temperature rose from 50 to 300 °C, with three desorption peaks centered near 147, 266 and 447 °C [18].

Control followed the stored state. A three-state lumped model of a urea SCR plant was analyzed for controllability and observability in 2006, with an observer for the gas-phase slip and a sliding-mode design [10]. The heavy-duty study of the same year combined a model-based feedforward with feedback on a NO_x sensor cross-sensitive to ammonia, possible only with a procedure that

separated the two signals [3]. A comparison published in 2007 found that open loop can reach high conversion with acceptable slip, that closed loop is needed for robustness to dosing error, aging and engine-out variation, and that an adaptive surface-coverage strategy outperformed sensor feedback with cross-sensitivity compensation over simulated test cycles [13]. Control-oriented models were developed and tested on a two-catalyst system in 2011 [11]. The ambiguity itself became a subject. A closed loop on the cross-sensitive sensor was shown to have spurious equilibria besides the point of interest, to which it nevertheless converges [7], a robust separation of the mixed signal's domain was published [14], and the cross-sensitivity factor itself was estimated by an observer [15].

The limits that frame the problem come from type approval. Euro VI, proposed for about 2013 to harmonize with US and Japanese rules [4], allows 460 mg/kWh of NO_x on the world-harmonized transient cycle (WHTC) of a compression-ignition engine and 10 ppm of ammonia [5]. Its implementing regulation takes the ammonia result as the average concentration over the cycle, weights the WHTC 0.14 cold and 0.86 hot, and sets a NO_x threshold of 1200 mg/kWh for on-board diagnostics (OBD) [19].

> ⧖ HISTORY **2001: A Regulator Doubts the Urea Tank**
>
> In January 2001 the United States Environmental Protection Agency set a heavy-duty standard of 0.20 g/bhp–hr of nitrogen oxides, phased in for diesel engines between 2007 and 2010. The rule noted that selective catalytic reduction was first developed for stationary applications, with a urea solution injected upstream of the catalyst and broken down into ammonia and carbon dioxide, and it did not base the new standard on that technology. Urea would have to be sold at the diesel fuel pump, and because an empty urea tank would not normally affect driveability, operators would have an incentive not to refill it. One commenter proposed a NO_x sensor that would detect the lost control and cut engine power gradually until half of it was gone [9].

12.3 The Exhaust Line as Modeled

The simulated line belongs to a loaded 30 t truck and runs from the turbine outlet to the tailpipe, drawn left to right as the application's Exhaust view draws it, in Figure 12.1(a). An oxidation catalyst heats or cools the gas without chemistry, an injector sprays diesel exhaust fluid (DEF, 32.5 % urea) at its outlet, a mixer leads to a 32 L SCR brick represented by three cells, an optional slip catalyst follows, and a zirconia NO_x sensor sits at the tailpipe. The engine-out NO_x sensor upstream sees no ammonia. The engine integrates in seconds with a step of $\Delta t = 50$ ms, the doser runs every 0.1 s and the log records every 0.5 s, the three clocks of Section 1.6.1.

The variables are classified in Table 12.1. Nineteen states are integrated. Eight are physical, the turbine-out and oxidation-catalyst temperatures and the stored ammonia and temperature of each of the three cells. Eleven are ledgers of ammonia released, reacted, oxidized and slipped, NO_x in and out, tailpipe NO_x and ammonia, engine work, urea and deposit, which make every balance exact.

12.3.1 Engine Maps and the Drive Cycle

The engine is a map, not a combustion model. The road load of a cycle table of speed v and grade β gives the brake power, and three declared functions of power and speed give the exhaust, as shown in Equation (12.1).

$$\begin{aligned} P &= \frac{\left(m_v g(C_{rr}+\beta)+\tfrac{1}{2}\rho C_d A\, v^2 + m_v a\right) v}{\eta_d} + P_{acc} \\ \dot{m} &= 0.045 + 0.0055\, v + 0.00095\, P \\ T_{e,ss} &= 130 + 1.9\,P - 0.0022\,P^2, \qquad x_{in} = 120 + 1.6\,P \end{aligned} \tag{12.1}$$

where $\dot{m}$ is in kg/s, v in m/s, P in kW (capped at 340 kW), the temperature in °C and x_{in} in ppm. At a standstill the engine idles at P_{acc}, and a negative wheel power while rolling cuts the fuel, with $P = 0$, $x_{in} = 0$ and 100 °C at the turbine. The turbine-out temperature follows its map through a

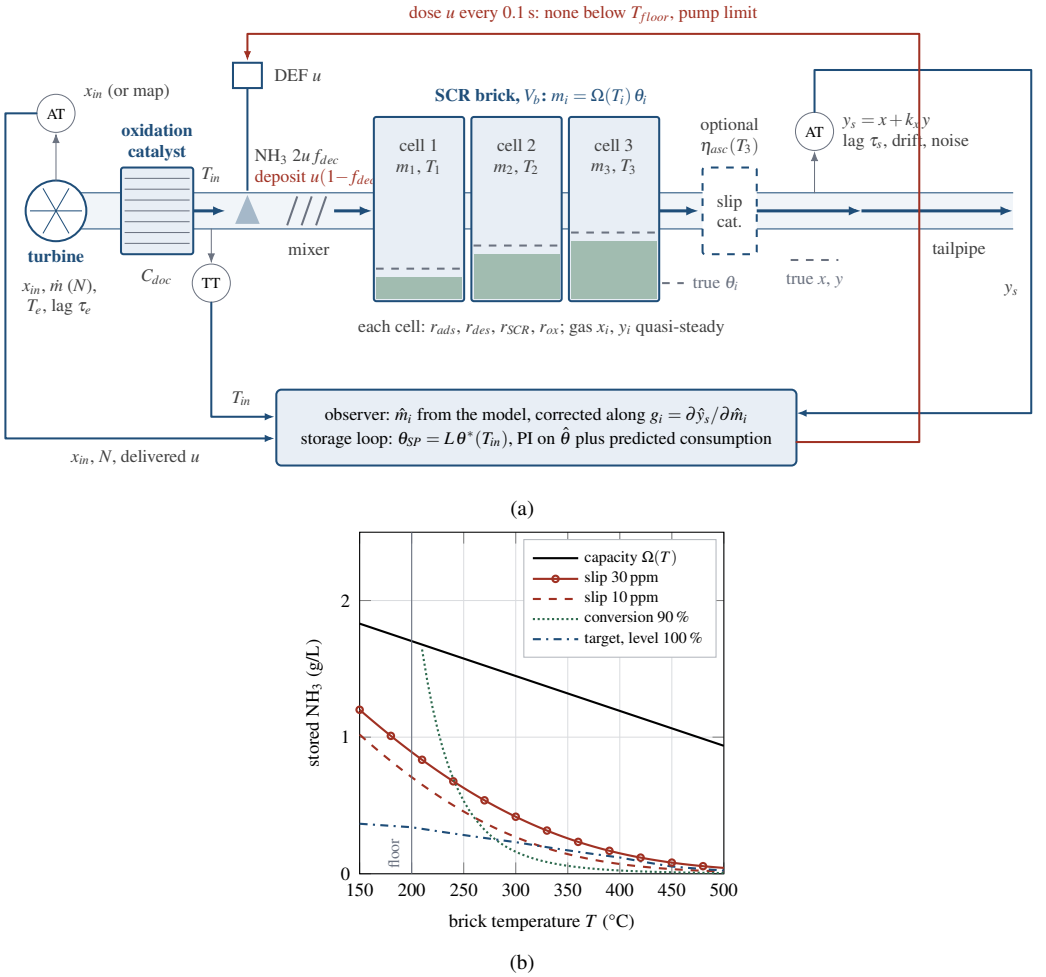

(a)

(b)

Figure 12.1 The SCR line as modeled, and its operating window. (a) The engine maps give the engine-out NO_x x_{in}, the exhaust flow $\dot{m}$ (molar flow N) and the turbine-out temperature T_e. The oxidation catalyst (C_{doc}) sets the inlet temperature T_{in} read by the thermocouple TT. The dose u releases $2u\,f_{dec}(T_{in})$ of ammonia and leaves the rest as deposit. Each cell stores $m_i = \Omega(T_i)\theta_i$ at temperature T_i through four rates, the gas x_i, y_i passing in closed form. The tailpipe sensor AT reads $y_s = x + k_x y$ through a lag τ_s. The observer and storage loop read y_s, T_{in}, x_{in}, N and the delivered urea, and set u. Truth is dashed, the cells filled to the observer's estimate. (b) At the highway point (0.278 kg/s, 294 ppm), from the application's own cell model on a uniform brick: the capacity (solid), the stores at which the gas leaving the brick carries 10 (dashed) and 30 ppm (solid with circles) of ammonia, the store that converts 90 % (dotted), and the storage target at level 100 % (dash-dotted). Not to scale.

lag, $\tau_e \dot{T}_e = T_{e,ss} - T_e$ with $\tau_e = 4$ s. On the 85 km/h highway the engine gives 108.8 kW, 0.278 kg/s, 310.7 °C and 294.1 ppm, a gas hourly space velocity of about 24 000 h^{-1}. On the climb at 72 km/h it gives 244.0 kW, 0.387 kg/s, 462.6 °C and 510.4 ppm. The cycle is a synthetic table, not a regulatory cycle: eight stop-and-go blocks of 75 s, an on-ramp to 85 km/h, three minutes of flat highway, six minutes at 72 km/h on 2.5 %, 90 s of descent at −3 % with the fuel cut, and highway to a stop at 1800 s. Its engine-out NO_x is 4.39 g/kWh, counted as NO_2. A severity factor s scales the truck's mass by $1 + 0.35(s-1)$ and the climb's grade by s.

Table 12.1 Variables by Role, With Values at the Screenshot Instant (Case 4, $t = 705\,\mathrm{s}$)

Symbol	Role	705 s	Unit
x	CV, tailpipe NO_x (truth)	0.087	ppm
y	constraint, tailpipe ammonia slip (truth)	384.5	ppm
y_s	measured, tailpipe sensor	267.4	ppm
$\hat{\theta}$, θ_{SP}	CV of the storage loop, its setpoint	0.270, 0.046	—
u	MV, DEF dosing	2280	g/h
x_{in}	DV, measured engine-out NO_x	510.4	ppm
$\dot{m}$, P	DV, exhaust flow and brake power	0.387, 244.0	kg/s, kW
T_{in}	DV, inlet thermocouple	459.0	°C
m_i	states, stored ammonia (total)	8.91	g
θ_i	coverages $m_i/\Omega(T_i)$	0.12, 0.29, 0.35	—
T_i	states, cell temperatures	460, 447, 422	°C
—	DVs: severity, NO_x map bias, aging, sensor drift, noise, failure	—	—

12.3.2 Oxidation Catalyst, Urea and Deposits

The oxidation catalyst is one thermal node that sets the SCR inlet, and the urea gives two ammonia per molecule once it decomposes, as shown in Equation (12.2).

$$\begin{aligned} C_{doc}\,\dot{T}_{in} &= N c_p\,(T_e - T_{in}) \\ \dot{n}_{NH_3} &= 2\,u\,f_{dec}(T_{in}), \qquad \dot{D} = u\,\big(1 - f_{dec}(T_{in})\big) \\ f_{dec} &= \min\!\Big(1,\ \max\!\Big(0,\ \frac{T_{in} - 170\,^{\circ}\mathrm{C}}{50\,\mathrm{K}}\Big)\Big) \end{aligned} \tag{12.2}$$

where T_{in} is the true gas temperature at the mixer, $N = \dot{m}/M_{ex}$ is the molar exhaust flow, u the urea delivered in mol/s, and D the deposit index. The doser delivers nothing while the inlet thermocouple, a 2 s lag of T_{in}, reads below the floor $T_{floor} = 200\,^{\circ}\mathrm{C}$, and at most 6000 g/h of DEF. One mole per second of urea is 665 280 g/h of DEF, so a dose ratio $\alpha = 1$ (one ammonia per engine-out NO_x) takes 945 g/h on the highway. A ramp from 180 to 250 °C, considered when the model was designed, left the passing controller depositing 15.6 g of urea per cycle at the same floor, against 2.95 g with the ramp used.

12.3.3 Three Cells of Stored Ammonia

The state of each cell is the stored amount per catalyst volume, m_i in mol/m^3, rather than its coverage. A rise in temperature that lowers the capacity then raises the coverage and pushes ammonia out through desorption, instead of making it vanish from the balance. The capacity and the four rates are shown in Equation (12.3).

$$\begin{aligned} \Omega(T) &= A\,\Omega_0\,\max\!\big(0.02,\ 1 - a(T - T_0)\big), \qquad \theta = m/\Omega \\ r_{ads} &= k_a\,C_{NH_3}\,(1-\theta)^{+}\,\Omega \\ r_{des} &= k_d\,\mathrm{e}^{-E_d(1-\alpha_T\theta)/RT}\,\theta\,\Omega \\ r_{SCR} &= A\,k_s\,\mathrm{e}^{-E_s/RT}\,C_{NO}\,\theta \\ r_{ox} &= k_o\,\mathrm{e}^{-E_o/RT}\,\theta\,\Omega \\ \frac{\mathrm{d}m_i}{\mathrm{d}t} &= r_{ads} - r_{des} - r_{SCR} - r_{ox} \end{aligned} \tag{12.3}$$

where C are gas concentrations in mol/m^3, $A = 1$ for a fresh brick and 0.7 for an aged one, and every rate is in mol/(m^3 s) of catalyst. The rates are adsorption on free sites, a Temkin desorption that weakens with coverage, standard SCR ($4\,NH_3 + 4\,NO + O_2$) by NO reacting with stored ammonia, and oxidation of stored ammonia to nitrogen, a lumped coverage description of the kind reduced for control in the cited studies [7, 10].

The gas crosses a 10.7 L cell in 23 ms on the highway and 13 ms on the climb, below the step, so each cell's gas balances are quasi-steady. Both are linear in the outlet at a given coverage and are solved exactly, as shown in Equation (12.4).

$$\begin{aligned} x_i &= \frac{x_{i-1}}{1 + V_c A k_s \mathrm{e}^{-E_s/RT_i}\, \theta_i\, P_{ex}/(RT_i N)} \\ y_i &= \frac{N y_{i-1} + V_c\, r_{des,i}}{N + V_c\, k_a (P_{ex}/RT_i)(1-\theta_i)^{+}\, \Omega_i} \end{aligned} \tag{12.4}$$

where x_i and y_i are the NO and ammonia mole fractions leaving cell i, $x_0 = x_{in}$, $y_0 = \dot{n}_{NH_3}/N$, $V_c = V_b/3$ and $P_{ex} = 101\,325$ Pa. Each cell's ammonia then closes by construction, $N(y_{i-1} - y_i) = V_c(\dot{m}_i + r_{SCR} + r_{ox})$. The cell temperatures follow the gas and the heats of reaction, as shown in Equation (12.5).

$$C_c \frac{\mathrm{d}T_i}{\mathrm{d}t} = N c_p\,(T_{i-1} - T_i) - V_c\left(\Delta H_{SCR}\, r_{SCR} + \Delta H_{ox}\, r_{ox}\right) \tag{12.5}$$

with $C_c = (\rho c)_b V_b/3$, $T_0 = T_{in}$, $\Delta H_{SCR} = -407$ kJ and $\Delta H_{ox} = -317$ kJ per mole of ammonia. No heat is lost to the surroundings. Figure 12.1(b) shows the consequence at the highway flow for a brick of uniform temperature and store. Below about 210 °C no store converts 90 %, and only above 260 °C does a store that converts 90 % leave less than 10 ppm of ammonia in the gas. The window opens as the brick heats, but the capacity that a cool brick filled shrinks at the same time.

12.3.4 Slip Catalyst and Instruments

The optional slip catalyst oxidizes a share η_{asc} of the ammonia leaving the brick, zero below 200 °C and rising linearly to 0.9 at 275 °C of the last cell, and turns 15 % of what it oxidizes into NO. The tailpipe sensor reads NO_x and a share of the ammonia, as shown in Equation (12.6).

$$\tau_s\, \dot{y}_s = x_{tail} + k_x\, y_{tail} + \delta + n(t) - y_s \tag{12.6}$$

where $k_x = 0.7$ is a teaching value (a published control study used 0.8 for a commercial sensor [7]), $\tau_s = 1.5$ s, δ is a drift and n is colored noise of 1 ppm with a 0.5 s correlation time, drawn from the case's seed. The lag is applied in exact zero-order-hold form at each step. The engine-out sensor has the same lag and no noise, and when it fails the controller receives the engine's NO_x map at the measured power instead. The inlet and outlet thermocouples have a 2 s lag. The exhaust flow and the delivered urea are passed to the controller as the engine control unit knows them. The scorecard counts tailpipe NO_x as NO_2 over the engine's positive work, both from the truth and from the sensor's own reading, and the slip peak and time average from the truth.

12.3.5 Parameters, Initialization and Numerics

Table 12.2 lists the parameters. The kinetic constants are teaching values fitted to two declared targets: a storage curve holding 1.09 g/L at 200 °C and 0.28 g/L at 400 °C with 100 ppm of ammonia in the gas, and 97.1 % conversion lined out at $\alpha = 1$ on the highway. The curve is compared with a measured one in Table 12.3, as a comparison rather than a validation. At 500 ppm it holds 1.60, 1.34 and 0.86 g/L at 150, 200 and 300 °C. The measured total uptake on one commercial Cu-chabazite monolith falls to 1.6 g/L at 300 °C [18], and values read from that source's figure place the declared curve at 0.61 to 0.79 of the total uptake and 0.94 to 1.22 of the strongly bound part. From 200 to 300 °C the curve falls to 0.64 of its value, faster than the total uptake (0.82 with water in the feed) and slower than the strongly bound part (0.50), so the heat-up release it predicts is likely generous.

Each case starts at its own declared condition. The warm cases start the line at 220 °C with a fifth of the capacity stored, the cold start at 40 °C with none, and the highway cases lined out at a stated dose ratio, where a bisection on each cell's storage rate gives the steady store and temperature cell by cell from the inlet, with derivatives below 10^{-6}. The observer starts from the true store unless a case says otherwise. The integrator is the framework's classical fourth-order Runge-Kutta method

Table 12.2 Parameters. A: Adjustable in the Model Dialog or the Advanced Controls. F: Fitted Teaching Value

Symbol	Meaning	Value	
m_v, C_dA, C_{rr}	truck mass, drag area, rolling resistance	30 t, 6 m^2, 0.0065	—
η_d, P_{acc}	driveline efficiency, accessories	0.90, 6 kW	—
τ_e, C_{doc}	turbine-out lag, oxidation catalyst	4 s, 10 kJ/K	A
V_b, $(\rho c)_b$	brick volume, heat capacity	32 L, 0.6 MJ/(m^3 K)	A
c_p, M_{ex}	exhaust heat capacity, molar mass	31.5 J/(mol K), 28.8 g/mol	—
Ω_0, T_0, a	sites at T_0, reference, capacity slope	100 mol/m^3, 200 °C, 0.0015 K^{-1}	A, F
k_a	adsorption	5.0 m^3/(mol s)	F
k_d, E_d, α_T	desorption, Temkin	330 s^{-1}, 57.7 kJ/mol, 0.42	F
k_s, E_s	standard SCR	1.3×10^8 s^{-1}, 60 kJ/mol	F
k_o, E_o	ammonia oxidation	2.4×10^7 s^{-1}, 130 kJ/mol	F
ΔH_{SCR}, ΔH_{ox}	reaction heats per mol NH_3	−407, −317 kJ	—
f_{dec}	decomposition ramp	170 to 220 °C	F
T_{floor}, u_{max}	dosing floor, pump	200 °C, 6000 g/h	A, —
k_x, τ_s	cross-sensitivity, sensor lag	0.7, 1.5 s	A
σ_n	sensor noise	1 ppm	A
—	thermocouple lag	2 s	—
η_{asc}	slip catalyst, share to NO	0.9 above 275 °C, 0.15	on/off
A	aging factor (sites, activity)	1 or 0.7	on/off

Table 12.3 The Declared Storage Curve Against a Measured Cu-Chabazite Monolith [18], in g/L

T (°C)	Ω	100 ppm	500 ppm	total, wet[a]	strong[a]
150	1.83	1.39	1.60	2.5	1.7
200	1.70	1.09	1.34	1.7	1.4
300	1.45	0.60	0.86	1.4	0.7
400	1.19	0.28	0.49	—	—

[a] Measured at 500 ppm, read from the source's figure, so approximate. The source's text gives a dry total of 1.6 g/L at 300 °C.

(Section 1.6.1). Halving the step moves the cycle's NO_x by 0.15 % as shipped, 0.07 % under the suggested control and 0.06 % under the feedforward, and moves no verdict. The 1800 s cycle takes 154 ms without a browser, in Node on the build machine, and about 9 s of real time in a browser run 200 times faster than the plant.

12.3.6 Assumptions

The engine and the cycle are declared maps and a synthetic table, so nothing here is a regulatory cycle or an engine's data. Only standard SCR is modeled. The oxidation catalyst makes no NO_2, so the fast SCR reaction, nitrate chemistry and N_2O are absent, which overstates the low-temperature penalty a real system pays. The gas is quasi-steady and the brick is three stirred cells. The three cells are part of the fitted model, not a grid refined to convergence, because the storage curve and k_s were fitted with three, and Section 12.5.9 shows that the challenge verdict depends on that choice. The urea decomposes at the mixer with no droplet evaporation time or wall film, and deposits never decompose again. No heat is lost and no particulate filter sits upstream, so every heat-up reaches the brick faster than it would on a truck with a filter. The cross-sensitivity is a constant. Aging removes 30 % of storage and reaction sites together and nothing else. The 0.46 g/kWh limit, the 10 and 30 ppm ammonia caps and the 0.55 g/kWh OBD threshold are teaching values, discussed with their sources in Section 12.5.

12.4 The Doser as Implemented

Every strategy computes an ammonia request every 0.1 s from one packet of measurements, the engine-out NO_x (or the map), the tailpipe sensor, the inlet thermocouple, the molar flow and the urea delivered at the previous call, and a test in the application fails if the controller or the observer reads any truth field. The request in g/h of DEF then passes the floor (zero while $T_{in} < T_{floor}$) and

the pump's maximum. In Manual the reader sets the DEF rate directly. The feedforward doses $\dot{n}_{NH_3} = \alpha\, x_{in} N$ with $\alpha = 1$ as shipped, where α is the dose ratio.

12.4.1 The Trim on the Tailpipe Sensor

The trim multiplies the feedforward by $1 + u_{trim}$, where u_{trim} is the output of the shared velocity-form PID of Section 1.6.2 on the conversion the sensor implies, as shown in Equation (12.7).

$$\begin{aligned} \eta_s &= 1 - y_s / x_{in}, \qquad & e_k &= \eta_{SP} - \eta_{s,k} \\ P_k &= K_{c,t}\, e_k, & s_k &= s_{k-1} + P_k - P_{k-1} \\ \hat{u}_k &= s_k + I_{k-1}, & u_{trim,k} &= \min\bigl(1,\ \max(-0.5,\ \hat{u}_k)\bigr) \\ I_k &= I_{k-1} + \Bigl[\frac{K_{c,t}}{\tau_{I,t}} e_k + \frac{u_{trim,k} - \hat{u}_k}{\tau_{I,t}}\Bigr] \Delta t_c \end{aligned} \tag{12.7}$$

with $\eta_{SP} = 95\,\%$, $K_{c,t} = 0.5$, $\tau_{I,t} = 30\,\mathrm{s}$, no derivative and $\Delta t_c = 0.1\,\mathrm{s}$. In the symbols of Equations (1.16) to (1.18), s_k is S_k and $\hat{u}_k$ is v_k, and the tracking time is $\tau_{I,t}$. When the pump saturates, the back-calculation drives the integral toward the trim that the applied dose implies. The trim holds by tracking while the floor shuts the doser or the measured engine-out NO_x is below 20 ppm, where no conversion is defined. A low apparent conversion raises the dose, a correct move on an underdosed brick and a wrong one past the bottom of the V. The process gain of the loop changes sign at that bottom, as the gain of a braking wheel does at the friction peak in Volume 7, Chapter 4.

12.4.2 The Stored-Ammonia Observer

The observer runs a copy of the three cells of Equations (12.3) to (12.5) on the measured inputs, with the cell temperatures driven by the inlet thermocouple, and corrects each cell's store along the gradient of its own predicted sensor reading, as shown in Equation (12.8).

$$\begin{aligned} \hat{y}_{s,raw} &= \bigl(\hat{x}_3 + \hat{k}_x \hat{y}_3\bigr) 10^6, \qquad & \tau_s\, \dot{\hat{y}}_s &= \hat{y}_{s,raw} - \hat{y}_s \\ g_i &= \partial \hat{y}_{s,raw} / \partial \hat{m}_i \\ \dot{\hat{m}}_i &= f_i(\hat{m}, \hat{T}) + \frac{1}{\tau_o} \frac{g_i\,(y_s - \hat{y}_s)}{\sum_j g_j^2 + g_0^2} \end{aligned} \tag{12.8}$$

where f_i is the model's storage rate, $\hat{k}_x = 0.7$ is the datasheet value the observer assumes, $\tau_o = 20\,\mathrm{s}$, $g_0 = 0.5\,\mathrm{ppm\,m^3/mol}$, and the gradient is a forward difference of $\max(0.05,\ 0.02\hat{m}_i)$. The observer integrates in two explicit substeps per control period and clips each store at zero. Its capacity assumption κ multiplies Ω_0, with $\kappa = 1$ the right value. The sign of g_i is the model's answer to the ambiguity. Where a larger store would lower the predicted reading (NO_x coming through) g_i is negative, where it would raise it (slip) g_i is positive, and near the bottom of the V both terms cancel and the observer runs on its model alone. In the terms of Section 1.4.6, the stores are observable through the sensor only where its sensitivity to them is not zero. The correction is a fixed-gain relative of the extended Kalman filter of Section 1.4.7, which Chapter 13 runs in full on a vehicle whose position no instrument aboard measures.

12.4.3 Storage Control on a Schedule

The storage loop holds the observer's coverage $\hat{\theta} = \sum \hat{m}_i V_c / \sum \hat{\Omega}_i V_c$ on a setpoint scheduled on temperature, and adds the ammonia the model predicts the reactions consume, as shown in Equation (12.9).

$$\begin{aligned} \theta_{SP} &= \min\bigl(0.95,\ L\,\theta^*(T_s)\bigr) \\ \dot{n}_{NH_3} &= \underbrace{\hat{r}_{SCR} + \hat{r}_{ox}}_{\text{consumption}} + \hat{C}\,\mathrm{PI}_s(\theta_{SP} - \hat{\theta}) \end{aligned} \tag{12.9}$$

where θ^* is linear between 0.20, 0.16, 0.10, 0.05 and 0.025 of the capacity at 200, 300, 400, 450 and 500 °C, T_s is the inlet thermocouple (or, as an option, the observer's brick temperature), L is the storage level, and $\hat{C}$ is the observer's capacity in moles, so the fill loop's output is a share of the capacity per second. The fill loop is the same shared PI with gain $K_{c,s}/60$ for $K_{c,s} = 4\,\mathrm{min}^{-1}$ and $\tau_{I,s} = 60\,\mathrm{s}$, its limits set every call so that the total request stays between zero and the pump's maximum. In the classical strategies, every loop that is not in charge tracks the dose actually applied, so a change of mode or strategy is bumpless (Section 1.6.2).

The schedule is high when the brick is cool, for conversion, and low when it is hot, to leave room for a heat-up release. Reading it at the inlet lowers the setpoint as soon as a climb heats the gas. In case 4 the inlet thermocouple reaches 350 °C at 635.4 s, the brick's mean at 660.8 s and its last cell at 676.6 s.

🕮 **Pause.** Try Exercise 12.2 before reading on.

12.4.4 Model Predictive Control

The option and its general form are described in Section 1.6.3. At a fixed store of 10 mol/m^3 the ammonia in equilibrium with the surface rises from 0.225 to 96.6 ppm between 200 and 450 °C, about 430-fold, so one fixed linear storage model misplaces the slip cap across the cycle. The binding constraint is a release from storage already present when the climb begins, which the scheduled loop meets in the original comparison. The resulting feedback policy depends nonlinearly on the state even though each optimization uses a linear model. The added MPC therefore linearizes Equations (12.3) and (12.4) at each decision rather than using one linear model throughout the cycle [20]. Fresh kinetics, 32 L volume and 0.6 MJ/(m^3 K) heat capacity are frozen at their declared values. A separate instance of Equation (12.8) propagates in every mode from the existing instrument packet and applied dose. It receives the same declared opening prior as Storage, which may match the initial true store. Known service loading updates its nominal prior using its own capacity, without copying the post-event plant store. No storage instrument is added.

The linear representation has eight states, three stores $\hat{m}_i/100$, the previous dose, a constant and three affine output offsets. Central differences of 0.01 mol/m^3 and 0.001 kg/h give the chemical and dose Jacobians. Exact zero-order hold gives a 1 s discrete model. An independent Taylor-exponential calculation differs by at most 5.72×10^{-9} at 200, 300 and 450 °C. The largest tested one-step local output residual is 0.001679 in scaled units. This does not validate a 60 s forecast. The current Jacobians remain fixed over the horizon, with affine offsets for forecast flow, inlet NOx and temperature. The observer includes reaction heat, while the future thermal transport forecast omits it.

The normalized outputs are $z = (\hat{\theta}_f, x/100, y/30)$, with x, y in ppm and $\hat{\theta}_f$ estimated moles divided by fresh nominal capacity. The observer's κ fraction is a separate diagnostic. A residual bias b supplies offset correction [21],

$$b^+ = \mathrm{clip}_{[-100,100]}\left[b + (1 - \mathrm{e}^{-\Delta t_c/30})(y_s - \hat{y}_s - b)\right]. \tag{12.10}$$

Here b is in ppm and adds to predicted NOx. The update has the form of the bias filter of Equation (1.12) with a fixed weight. It cannot separate NOx, ammonia and sensor drift uniquely, as two constant errors seen through one measurement never can (Section 1.4.6). No Kalman gain is used, whereas Volume 9, Chapter 4 carries its unmeasured fuel bias as a state of a Kalman estimator.

The MPC decides every 1 s, holds its first move between decisions and plans 60 samples with the move-block lengths $[1, 1, 2, 3, 4, 6, 8, 10, 12, 13]$ in seconds. For dose v in kg/h, target $r = (\theta_{SP}, 0, 0)$ and nonnegative horizon-wide slacks $s = (s_\theta, s_x, s_y)$, its quadratic objective is

$$\begin{aligned} J = &\sum_{k=1}^{60} \left[400(z_{1,k} - r_{1,k})^2 + 2z_{2,k}^2 + 0.01 v_k^2\right] \\ &+ 0.15 \sum_{j=1}^{10} (\Delta v_j)^2 + 10^4 s_\theta^2 + 20 s_x^2 + 10^3 s_y^2. \end{aligned} \tag{12.11}$$

Table 12.4 The Cases and the Challenge

#	Setting	Teaching point
1	Highway, Manual, lined out at $\alpha = 0.8$	conversion against dose
2	Highway, Manual, lined out at $\alpha = 1.3$	the ambiguous sensor
3	Highway, $\alpha = 1.1$; dosing stops at 120 s	the stored buffer
4	Warm crawl, climb at 600 s; feedforward	the heat-up release
5	Cold start at 40 °C, urban crawl	the floor and deposits
6	Cycle to 15:00, storage control	storage against feedforward
7	Cycle, aged brick, storage control	aging, OBD, infeasibility
8	Highway, lined out at $\alpha = 1.3$, trim	a loop that runs away
9	Full cycle, MPC or Storage, preview Off	constrained dosing and forecast error
ch	Full cycle from 220 °C, trim as shipped	the whole problem

There is zero direct slip tracking weight and no terminal multiplier. Hard constraints impose $0 \le v_k \le 6$, $|\Delta v_j| \le 1.5$ between blocks, nonnegative predicted stores, NOx and slip, and zero dose throughout a block containing any forecast inlet below the selected floor. The first difference uses the applied dose. The measured floor remains a hard actuation backstop. The soft bands are

$$\begin{aligned} \max(0, \theta_{SP,k} - .04) - s_\theta &\le z_{1,k}, \\ z_{1,k} &\le \theta_{SP,k} + .04 + s_\theta, \\ z_{2,k} &\le \frac{.46 \max(0, P_k) 10^6}{3600(46.0055) \max(.1, N_k) 100} + s_x, \\ z_{3,k} &\le 1 + s_y. \end{aligned} \tag{12.12}$$

The NOx ceiling is an instantaneous power-based equivalent of 0.46 g/kWh, not a cycle-integral guarantee. Each output uses one slack for all 60 samples. The shared solver's plans are accepted only when they also satisfy every hard bound over the whole horizon (Section 1.6.3). A rejected first scan holds the previous applied dose, then the tracked Storage loop acts until a valid plan returns. Cold priming solves in the zero-dose box before reporting floor protection. Trim tracking inverts applied dose divided by feedforward only when the latter is nonzero. Its zero-feedforward stand-in cannot represent a positive dose. Storage inversion requires positive estimated capacity, and Manual follows the requested station separately.

The inlet thermocouple is current upstream information, not a future climb announcement. Preview is Off by default and holds current measured inlet, flow and NOx. On uses an explicitly selected public synthetic cycle through nominal thermal lags, with a declared clock error. Storage receives the same forecast and can lower its target using the maximum forecast inlet temperature. Neither law reads the actual future profile or scripted events.

12.4.5 Three Models

Section 12.3 is the plant. From case 1 a reader identifies a conversion against dose ratio and a filling time of minutes. The classical Storage controller's model is the observer's copy, which knows the declared kinetics, V_b and $(\rho c)_b$ exactly, assumes a fresh capacity scaled by κ and $\hat{k}_x = 0.7$, and knows neither the aging, the sensor drift nor the engine's departure from its NO_x map.

12.5 Ten Cases, Measured

The cases appear in Table 12.4. The figures below come from headless runs that advance the engine one step at a time as the browser does, on the default seed, and a setting named in a sentence is applied at the moment that sentence gives. The three ways into a case (opening the page on it, switching to it from another case, pressing Reset) give runs equal to the last bit, as Section 1.6.1 leads a reader to expect.

12.5.1 A Dose Ratio Is Not a Conversion

Lined out on the highway, the brick converts 79.4, 97.1 and 99.6 % of the NO_x at α of 0.8, 1.0 and 1.3, storing 3.0, 7.9 and 16.7 g of ammonia, while the slip rises from 0.5 to 5.3 to 82.8 ppm. Raised from 756 to 945 g/h at 30 s in case 1, the conversion reaches 94 % at 739 s, 11.8 minutes later, because the extra dose first raises the store toward its new level. At 1229 g/h the slip passes 30 ppm at 817 s. Figure 12.2(a) shows the lined-out sensor reading against the dose ratio. It falls with the NO_x to 12.1 ppm at $\alpha = 1.0$ and rises again with the ammonia, reading 60.8 ppm at 0.8 and 59.3 ppm at 1.3. Case 2 walks across it. At $\alpha = 1.3$ the sensor reads 60.0 ppm with 1.3 ppm of NO_x and 82.8 ppm of ammonia. Lowered to $\alpha = 0.8$, the store drains, the slip falls below 2 ppm at 1020 s with the sensor at 36.2 ppm, and by 30:00 the sensor reads 60.4 ppm again, now with 60.4 ppm of NO_x and 0.5 ppm of ammonia. Two doses that differ by 60 % give one reading.

Pause. Try Exercise 12.1 before reading on.

12.5.2 The Buffer and the Heat-Up Release

The store is a buffer of minutes. Lined out at $\alpha = 1.1$ in case 3 the brick holds 12.3 g ($\theta = 0.272$) against the 2.9 g of ammonia a minute that the highway's NO_x consumes at full conversion. When the urea stops at 120 s, the conversion stays above 90, 80 and 50 % for 154, 191 and 247 s. A store lined out at $\alpha = 1.0$ holds 7.9 g and keeps 90 % for 86 s, and one at 1.2 holds 14.9 g for 187 s.

The same store becomes the hazard when the brick heats. In case 4 the feedforward at $\alpha = 1.0$ leaves 11.75 g stored at 600 s, when the climb begins, and the climb produces the two humps of Figure 12.2(b). First the NO_x rises faster than the cool brick converts it, and the sensor reads 296.7 ppm at 627.5 s with 291.5 ppm of NO_x and 3.0 ppm of ammonia. Then the capacity falls under the store, the slip peaks at 384.5 ppm at 705 s, and the sensor reads 267.5 ppm with 0.1 ppm of NO_x. Both readings say dose more, and only the first is right. Storage control on the same seed stores 10.1 g by 600 s, not much less, but its setpoint falls with the inlet temperature and the doser stops, so the NO_x uses the store up instead of fresh urea joining it. Its slip peaks at 32.2 ppm at 692 s. With the schedule read on the brick estimate instead of the inlet, the same loop peaks at 43.1 ppm, and the trim loop, which stores 16.4 g by 600 s, peaks at 1537 ppm.

> **TRY IT IN THE APP** **4 · The Heat-Up Release**
>
> Load the scenario, press Run on the Exhaust view and watch the tailpipe sensor and the ammonia slip as the climb starts at 10:00. Near 705 s the slip peaks at 384.5 ppm while the sensor reads 267.5 ppm, a reading a controller would take for unconverted NO_x.

12.5.3 The Floor and the Cold Start

From 40 °C on the urban crawl of case 5 the doser waits for its floor. At 200 °C the first urea flows at 185 s, and at 10:00 the tailpipe has 3.03 g/kWh with a deposit index of 1.00 g of urea. Lowered to 150 °C, the floor lets urea flow at 104 s into a brick too cold to use it, for 2.93 g/kWh, 3 % less, and 3.04 g of deposit. At 220 °C nothing deposits and the doser waits until 264 s, for 3.26 g/kWh. Every choice leaves more than six times the 0.46 g/kWh limit at 10:00, so the cold start is lost whatever the floor, and the floor trades deposits, not NO_x.

12.5.4 The Trim Loop That Runs Away

Case 8 starts lined out at $\alpha = 1.3$ with the trim on. The sensor's apparent conversion is 80.1 % against a target of 95 %, so the trim adds urea, the slip grows, the reading rises and the trim adds more. The slip passes 100 ppm at 151 s and 300 ppm at 305 s, the trim reaches its +100 % limit at 256 s, and at 10:00 the doser delivers 2458 g/h ($\alpha = 2.60$) with 461 ppm of ammonia and 0.5 ppm of NO_x. Storage control on the same seed drains the overfull brick to its schedule and doses 918 g/h ($\alpha = 0.97$) at 10:00 with 6.3 ppm of ammonia and 11.0 ppm of NO_x. Started lined out at $\alpha = 1.0$ instead, the trim never saturates in 900 s, which is why the case starts on the right side of the V.

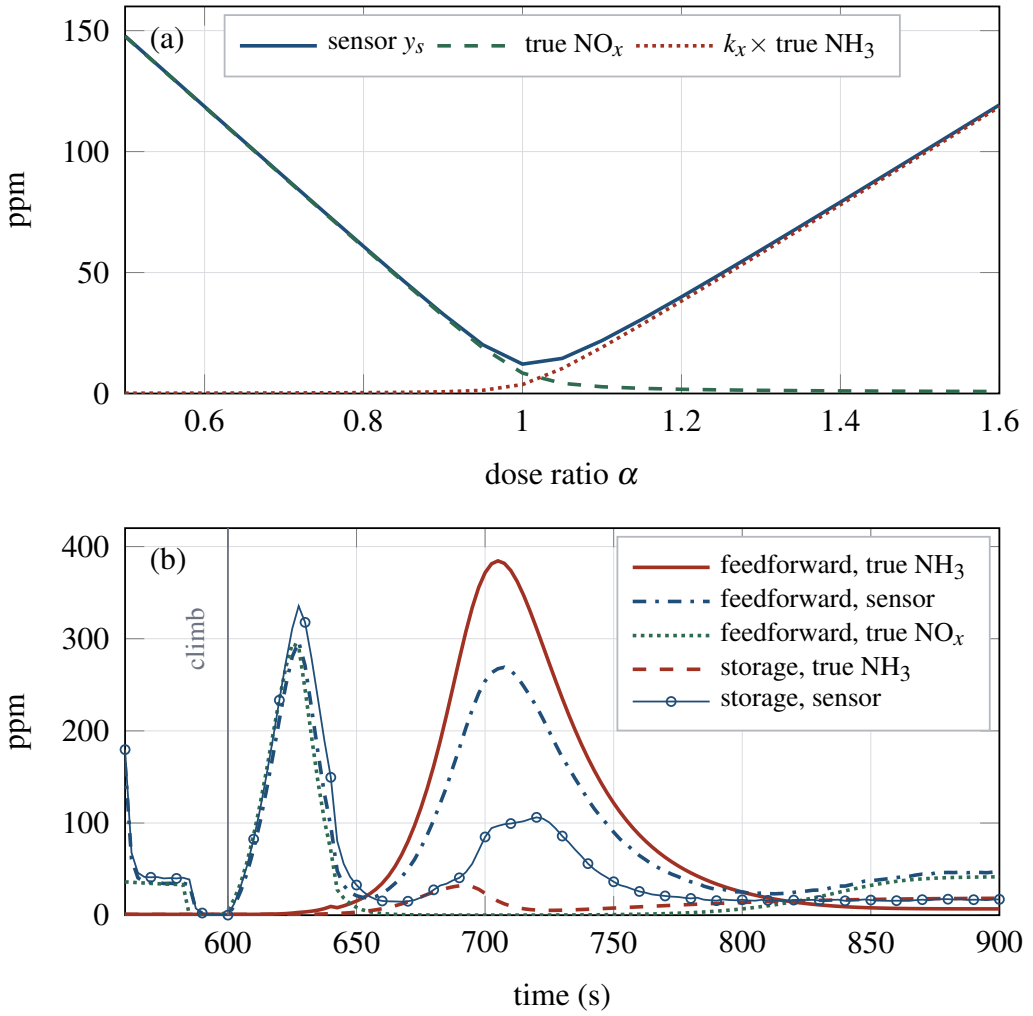

Figure 12.2 Two faces of the ambiguity, from the application's own functions. (a) The lined-out highway point at each dose ratio: the sensor reading (solid), the true NO_x (dashed) and k_x times the true ammonia (dotted). (b) Case 4, the climb from 600 s: under the feedforward (true ammonia solid, sensor dash-dotted, true NO_x dotted) the sensor reads high twice, first on NO_x (296.7 ppm at 627.5 s) and then on released ammonia (267.5 ppm against 384.5 ppm of slip at 705 s). Storage control on the same seed (true ammonia dashed, sensor thin with circles) peaks at 32.2 ppm of slip.

> **TRY IT IN THE APP** **8 · The Trim Loop Runs the Wrong Way**
>
> Press Run and watch the trim and the slip. The trim reaches its +100 % limit at 256 s, and at 10:00 the slip is 461 ppm. Then set Dosing strategy to Storage, press Reset and run ten minutes on the same seed. At 10:00 the slip is 6.3 ppm.

12.5.5 The Observer

From a zero initial estimate with the right capacity on the highway, the coverage error falls from -0.175 to -0.0058 in 120 s and is 0.0019 at 1200 s. Started on the other side of the V at $\hat{\theta}_0 = 0.6$, it falls from 0.527 to 0.126 in 120 s and to 0.021 at 1200 s, more slowly. A wrong capacity assumption biases the stored mass far more than the coverage, +27.0 and −22.3 % at κ of 1.3 and 0.7 against -2.3

Table 12.5 The Challenge Over 1800 s, Default Seed. Limits 0.46 g/kWh, Mean 10 ppm, Peak 30 ppm

Configuration	g/kWh	Peak (ppm)	Mean (ppm)	DEF (g)	Verdict
trim, $\alpha = 1.0$ (shipped)	0.293	1301.9	108.3	893	fails
feedforward, $\alpha = 1.0$	0.564	228.2	17.5	455	fails
feedforward, $\alpha = 0.7$	1.145	40.1	4.8	318	fails
feedforward, $\alpha = 0.6$	1.466	32.3	3.2	273	fails
storage, level 100 % (suggested)	0.423	23.5	5.75	470	passes
schedule on the brick	0.429	32.4	6.1	457	fails
no consumption feedforward	0.448	22.7	5.4	468	passes
both of the above	0.450	23.0	5.2	452	passes
level 90 %	0.470	19.8	5.0	453	fails
level 95 %	0.446	21.6	5.4	462	passes
level 115 %	0.362	29.96	7.2	496	passes
level 120 %	0.344	32.5	7.7	504	fails

and +11.0 %, because the sensor correction holds the coverage that matches the tailpipe. Perturbing every truth field before a control call leaves the dose identical to the last bit.

🕮 **Pause.** Try Exercise 12.3 before reading on.

12.5.6 Storage Control, Aging and the OBD

Over the first 15 minutes of the cycle in case 6, storage control leaves 0.68 g/kWh with a slip peak of 23.5 ppm on the on-ramp, against 0.54 g/kWh and 94.9 ppm for the feedforward. It gives up crawl conversion to keep the on-ramp release under 30 ppm, and over the whole cycle it leads on both, 0.423 g/kWh and 23.5 ppm against 0.564 g/kWh and 228 ppm.

On the aged brick of case 7 the same settings pass the NO_x at 0.428 g/kWh but slip 53.8 ppm at 1231 s, the end of the climb, with a 13.0 ppm mean. Telling the observer the right capacity ($\kappa = 0.7$) gives 0.414 g/kWh and 59.0 ppm, because the schedule is a share of the capacity, and lowering the level to 70 % trades the slip for NO_x, 0.569 g/kWh and 36.0 ppm. The sensor's own cycle number is 0.586 g/kWh aged against 0.490 g/kWh fresh, on either side of the teaching OBD threshold of 0.55 g/kWh, but the fresh brick under the plain feedforward reads 0.792 g/kWh, because its slip reads as NO_x. The threshold is set below the regulation's 1.2 g/kWh [19], but the shipped challenge reads 1.689 g/kWh and case 8 reads 2.876 g/kWh, so some sensor values exceed that threshold. On the severe cycle (severity 1.3, 10.5 % more mass and a 3.25 % grade) the aged brick meets the three limits at none of 19 settings, seven feedforward ratios from 0.6 to 1.2 and twelve storage settings with κ of 1 and 0.7. The lowest peak among them is 38.3 ppm at 1.18 g/kWh, so the case is reported as infeasible.

12.5.7 The Challenge

The challenge runs the whole cycle from a warm start and grades the truth over 1800 s against 0.46 g/kWh of NO_x, a 10 ppm ammonia average and a 30 ppm peak, with the plant as shipped. The NO_x limit is Euro VI's WHTC value and the average its ammonia form [5, 19], here applied to one warm run rather than the weighted cold and hot pair. The peak is the target of a published control study [7]. A 10 ppm peak was considered first. Measured, a peak of 10 ppm or less needs a storage level of 40 % or less, and the least NO_x with it is 0.93 g/kWh (0.92 g/kWh with the true store handed to the controller, a teaching bound), twice the limit, so the regulation's average and the published peak were adopted. Table 12.5 lists the configurations, and Figure 12.3 draws them as a chart.

The shipped trim meets the NO_x by overdosing. Every release reads as NO_x, the trim answers with urea, and the slip peaks at 1302 ppm on the climb. No feedforward ratio passes. At $\alpha = 0.6$ the on-ramp release still peaks at 32.3 ppm and the NO_x is 1.466 g/kWh. Storage control passes with 0.423 g/kWh, a 23.5 ppm peak on the on-ramp and a 5.75 ppm mean, and on five other seeds of the sensor noise it gives 0.422 to 0.424 g/kWh and 23.5 to 23.7 ppm, while the shipped trim fails on

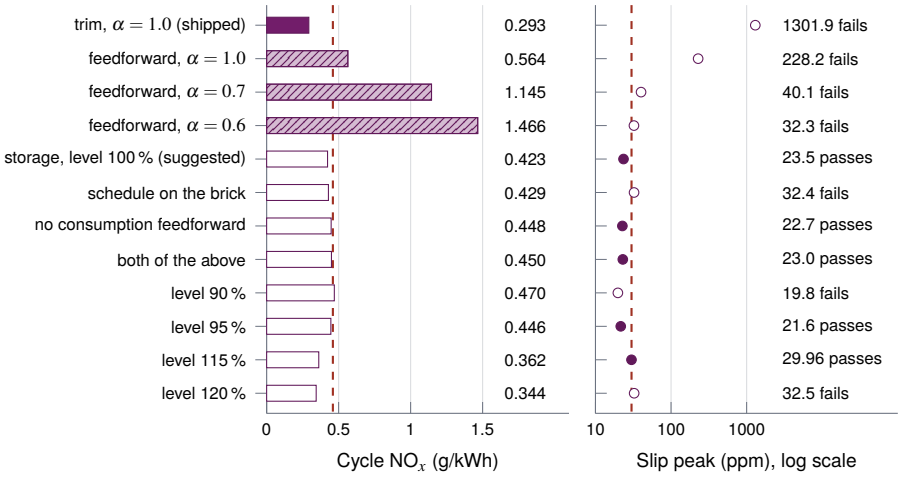

Figure 12.3 The challenge of Table 12.5: the cycle NO_x (left, dashed line at the 0.46 g/kWh limit) and the ammonia slip peak (right, log scale, dashed line at the 30 ppm cap) with the verdict. The dark bar is the shipped trim, hatched bars the feedforward and white bars the storage loop and its variants. Filled dots pass all three limits and open dots fail. Notice that the configuration with the least NO_x has the highest peak, and that the passing storage settings sit in a narrow band just under both lines.

every seed. Levels from 95 to 115 % pass, and every fill gain from 2 to 10 min^{-1} and integral time from 20 to 600 s passes at level 100 %. The inlet lead and the consumption feedforward work as a pair. With the schedule read on the brick estimate and the feedforward on, the loop doses 41.4 g of DEF between 860 and 960 s against 29.6 g as suggested and 35.0 g with both removed, because the feedforward keeps feeding what the hot brick consumes while the late setpoint still sits high, and the peak reaches 32.4 ppm at 958 s. With both removed the cycle passes again. What passes the cycle is the storage loop on the observer, and the lead is what makes the consumption feedforward safe on a climb. With the sensor's cross-sensitivity at 0.8 instead of 0.7 the suggested settings still pass (0.425 g/kWh, 23.5 ppm) and the shipped trim still fails.

> **+ DID YOU KNOW? The Lowest Nitrogen Oxides and the Worst Slip**
>
> From the simulator: the shipped trim leaves 0.293 g/kWh of NO_x over the challenge cycle, the lowest of the twelve configurations of Table 12.5, and still fails, with a slip peak of 1301.9 ppm and a mean of 108.3 ppm. The sensor it trusts reads the released ammonia as NO_x.

12.5.8 MPC and the Information Available to the Doser

Table 12.6 compares full 1800 s runs at the default seed 20261026, 0.05 s physics step and unchanged instruments. One MPC tuning serves every case. A faster Storage rule is frozen before the comparisons: $d\theta/dt \simeq$ fill with repeated closed-loop poles at $-1/\tau$, $\tau = 60/6 = 10$ s, gives $K_{c,s} = 12\,\text{min}^{-1}$ and $\tau_{I,s} = 20$ s. All its other settings equal the original Storage row. No shipped or suggested setting changes. The new comparison case records two full runs separated by Reset, not an emissions pass. Physical criteria here are 0.46 g/kWh, 10 ppm mean and 30 ppm peak. The actual challenge also requires full duration and unchanged parameters and non-event disturbances.

Without preview the MPC reduces NOx but fails the slip peak, while both Storage tunings pass. Its peak exceeds 30 ppm for 35.3 s, with an excess integral of 108.270 ppm s. These are time and cumulative excess, not peak depth. The physical release persists despite a soft predicted cap, as shown in Figure 12.4. With equal public preview the MPC passes while both Storage tunings miss NOx. The

Table 12.6 Full-Cycle Comparison, With One Frozen Tuning per Law. O: Preview Off, P: Public Preview On, E: On With +90 s Clock Error

Preview	Law	NOx (g/kWh)	Peak (ppm)	Mean (ppm)	DEF (g)	Slip integral (ppm s)	Rejected/solves	Fallback (s)	Criteria
O	Storage	.423342	23.505	5.754	470.379	10357.3	0/0	0	pass
O	Fixed Storage	.412347	24.133	6.242	478.419	11235.1	0/0	0	pass
O	MPC	.365405	34.631	7.038	477.686	12668.4	274/1735	273.4	fail
P	Storage	.473256	21.354	5.629	463.382	10132.0	0/0	0	fail
P	Fixed Storage	.467040	22.923	6.027	468.193	10848.8	0/0	0	fail
P	MPC	.395184	26.273	6.459	473.265	11626.2	392/1735	391.8	pass
E	Storage	.529614	21.269	5.353	458.461	9635.8	0/0	0	fail
E	Fixed Storage	.513025	22.124	5.618	464.593	10112.8	0/0	0	fail
E	MPC	.400146	28.644	6.500	476.418	11700.1	510/1735	510.0	pass

+90 s error still passes this fresh case but needs 510 s of fallback. This is an information-dependent comparison.

The aged case without preview gives MPC 0.409 652 g/kWh and 76.937 ppm peak against Storage 0.427 789 g/kWh and 53.757 ppm. The crawl/climb gives 0.136412 and 83.737 ppm against 0.186772 and 32.161 ppm. Sensor drift worsens MPC NOx (0.515580 versus 0.456 280 g/kWh), and failed engine-out sensing plus a map change worsens both NOx and peak (0.548224, 33.836 ppm versus 0.506747, 22.697 ppm). In the high-store ambiguity case both begin above the cap (81.728 versus 82.584 ppm). A 50 L volume mismatch reverses the NOx ranking (MPC 0.346 versus Storage 0.281 g/kWh). No case is retuned to remove these losses.

The matrix contains 134 primary and 66 fixed-rule runs, checking 10.764 million physics steps, plus ten mismatch runs and eight independent integral checks. No finite-value, applied-dose or measured-floor violation occurs. Halving the step retains all physical grades and the 33 original MPC-minus-Storage signs for NOx, peak, DEF and integrated slip. The largest relative changes are 1.09, 3.43 and 0.69 % for NOx, peak and DEF. Fixed Storage has one cold-start/preview-On integrated-slip near-tie reversal, from MPC minus fixed +0.0281 to −0.595 85 ppm s. This is sensitivity evidence, not convergence. The acting observer/reference coverage IAE is 17.933 versus 37.881 s times fractional coverage for Storage and MPC, with 2 Hz dose travel 42071 versus 93 275 g/h and MPC saturation 0.389 %. Fallback changes the observer underlying this metric.

Rendered Chrome at 1× uses 0.276 % of one core for all controller work, with mean solve 642 µs over 12 decisions. At 200× it reaches 201.2 s/s and 60.1 frames/s with 17.92 % controller core share. A separate synchronous full cycle averages 218 µs per solve over 1735 attempts. The 25 ms frame guard has no natural stops in these measurements. Forced stops preserve all 36000 physics steps and 1735 attempts, with identical non-wall-clock CSV data in Chrome. These host timings do not imply a pacing benefit.

12.5.9 The Cell Count Decides the Verdict

The three cells are a declared structure. An independent harness rebuilt the line on the application's own cell functions. It reproduces the application at three cells (feedforward cycle 0.564 g/kWh and 231 ppm against the application's 0.564 and 228), and ran it with the same constants on more cells, with the storage law given the true store as a labeled bound. Table 12.7 gives the result. More cells convert more and release a sharper front, so at level 100 % the peak rises from 24.9 to 42.2 and 66.8 ppm at 3, 6 and 12 cells, and the passing window of levels narrows to 70 % alone at six cells and closes at twelve. The ranking holds at every count, storage control below the feedforward on NO_x and on peak. A six-cell brick with these constants is a different, unfitted catalyst, so the verdict belongs to the three-cell model and the lesson does not.

The heat-up release at its peak is shown in Figure 12.5, the instant case 4 is built around.

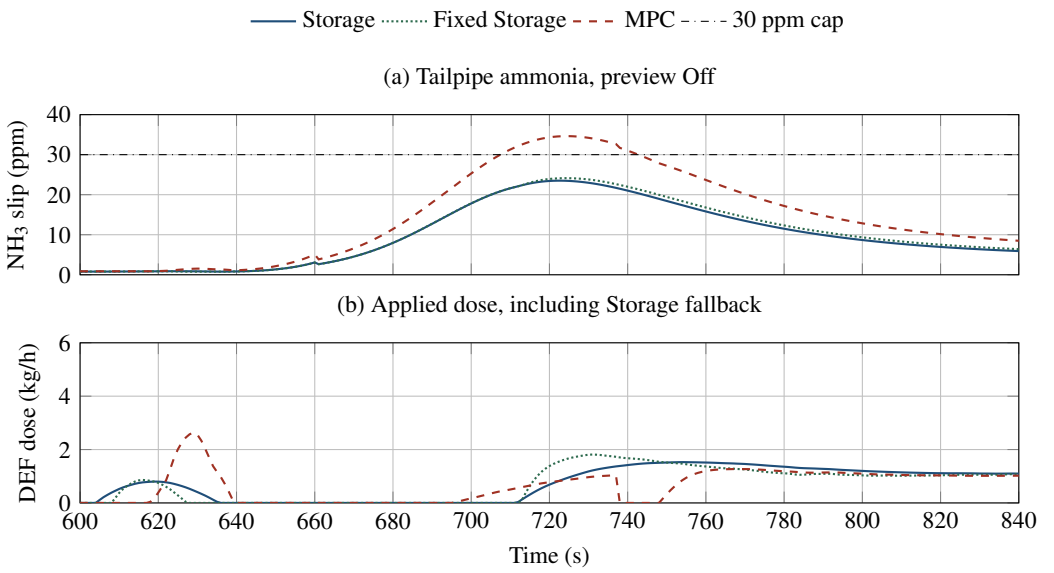

Figure 12.4 The on-ramp release with preview Off on the full-cycle comparison. (a) True ammonia slip and its physical cap (dash-dotted). (b) Applied DEF dose. Common time axes and the same seed show the original Storage settings (solid) beside the fixed rule (dotted) and the MPC (dashed), the MPC run including its Storage fallback. The MPC peak is 34.631 ppm at 724.45 s, during fallback, against original Storage's 23.505 ppm peak at 722.55 s.

Table 12.7 Cell Count With the Same Constants (Independent Harness, Storage Law on the True Store)

Cells	Feedforward cycle	Storage, level 100 %	Case 4 feedforward
3	0.564, 231 ppm	0.406, 24.9 ppm, passes	388 ppm
6	0.487, 225 ppm	0.315, 42.2 ppm, fails	500 ppm
12	0.440, 261 ppm	0.265, 66.8 ppm, fails	598 ppm

12.6 Evidence of Three Kinds

Code verification, agreement with published data and educational evaluation are reported separately. Code verification covers the ledgers, where on a disturbed run of case 4 with the slip catalyst fitted, an aging step and a sensor drift the ammonia closes to a relative 8.7×10^{-16}, the urea to 1.2×10^{-15}, the NO consumed equals the ammonia reacted to 1.7×10^{-14}, and the slip catalyst makes exactly 0.150 of what it oxidizes into NO. The brick's energy balance over 1200 s closes to 0.003 % with a trapezoid on the step values, and the two heats agree with the formation enthalpies. Limiting cases hold, among them a sensor that reads the true NO_x at $k_x = 0$, zero conversion at $\theta = 0$, a store lined out at $\alpha = 1$ that empties to 0 % conversion without urea, and no urea delivered below the floor. The application runs 50 checks, 31 of its own and 19 common to every framework app, among them the step halving above. Agreement with published data is limited to the storage curve of Table 12.3, which lies within one source's band for one catalyst, and to the regulatory numbers taken as teaching values. No engine, catalyst or sensor was fitted to plant data. Educational evaluation consists only of trial runs made while the application was reviewed. The conversion of case 1 reached 94 % after 35 s of wall time and the challenge was passed through the interface in 41.6 s. No learning study has been made.

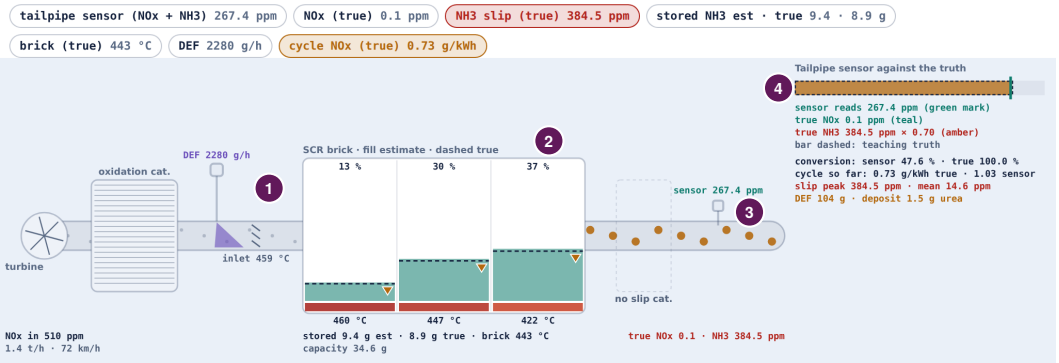

Figure 12.5 The heat-up release of case 4 at its peak, $t = 705\,\text{s}$, on the Exhaust view, as shipped with the feedforward at $\alpha = 1.0$, cropped from the full screen. Callouts: (1) the urea injector with its dose, (2) the three cells of the brick, (3) the tailpipe sensor and (4) the bar of the sensor against the truth. The truck climbs at 72 km/h with 510 ppm of engine-out NO_x, the inlet reads 459 °C and the cells 460, 447 and 422 °C. Each cell is filled to the observer's coverage (13, 30 and 37 %) with the true coverage dashed, and its amber marker shows a store emptying. The tailpipe sensor reads 267.4 ppm against 0.1 ppm of true NO_x and 384.5 ppm of true ammonia, a sensor conversion of 47.6 % against a true 100.0 %, and the status strip, cropped from this view, reads HEAT-UP RELEASE. The observer estimates 9.4 g stored against a true 8.9 g. The trend panels of the full screen, cropped from this view, show over the ten-minute window the crawl's dosing pulses, the storage loop tracking in the background with its setpoint fallen to 4.6 % (0.046 in Table 12.1), and the slip and sensor rising together after the climb.

12.7 Technology in Practice

What Measures It

The doser reads three things in the exhaust: the NO_x leaving the engine, the gas temperature at the inlet of the brick and the tailpipe sensor. The first two are honest. The third is the instrument the whole chapter turns on, because it answers to ammonia as well as to NO_x.

> **TECHNOLOGY** **Sensor: A Tailpipe NO_x Sensor That Also Sees Ammonia**
>
> Exhaust gas sensors are solid-state electrochemical cells, most built on yttria-stabilized zirconia, a ceramic that conducts oxygen ions. Read in the amperometric mode, the cell gives a current of nanoamperes at parts per million of NO_x, and in 2012 only one type of NO_x sensor was sold commercially [22]. In engine tests, when the ammonia in the exhaust rose from about 10 to 30 ppm with no change in nitric oxide, the reading of that commercial sensor rose too [23]. An open circuit in a NO_x sensor is one of the faults on which a heavy-duty engine must derate [8]. The chapter's sensor counts the ammonia with $k_x = 0.7$ behind a 1.5 s lag.

The model keeps that sensor simple, a fixed cross-sensitivity, a lag, a drift and colored noise, so that the ambiguity is the only thing that can mislead the doser. A real sensor's cross-sensitivity varies with the sensor and the gas, which the Limitations section notes and which an observer can estimate.

What Moves It

The single manipulated variable is the dose of diesel exhaust fluid. The pump and the valve that inject it decide what reaches the exhaust, and the temperature of the gas decides how much of it becomes ammonia.

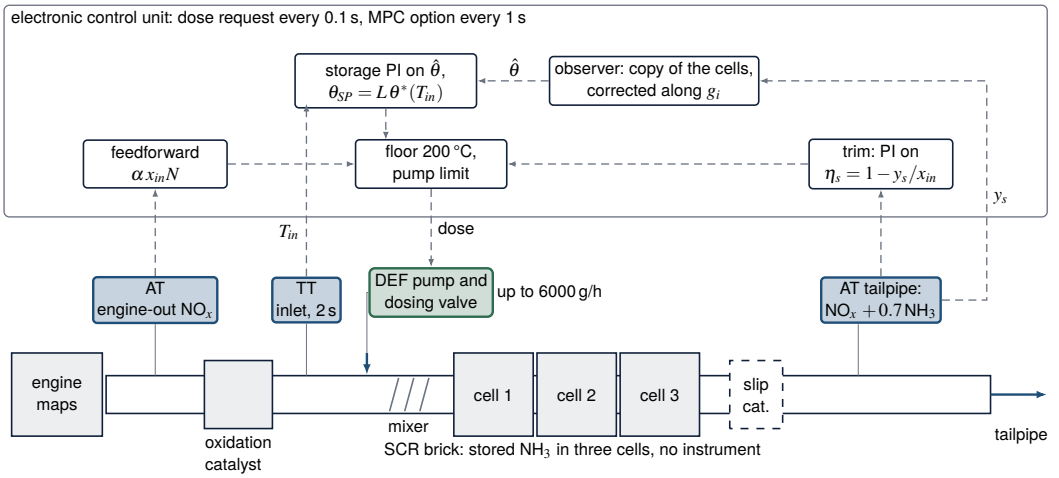

Figure 12.6 The dosing loop on vehicle hardware. The engine-out NO_x sensor AT, which sees no ammonia, feeds the feedforward. The inlet thermocouple TT sets the storage schedule and the dosing floor. The tailpipe sensor AT, which reads $x + k_x y$ with $k_x = 0.7$, feeds the trim and corrects the observer. The observer also reads the engine-out NO_x, the inlet temperature, the flow and the delivered urea (lines not drawn), and the MPC option is not drawn. One electronic control unit computes every strategy every 0.1 s, the request passes the floor and the pump limit, and the pump and dosing valve inject the fluid ahead of the mixer. Notice that no line leaves the three cells of the brick: the store that sets the conversion reaches the controller only through the observer.

> **TECHNOLOGY** **Final Element: The DEF Pump and Dosing Valve**
>
> Diesel exhaust fluid is a water solution of 32.5 % urea [24], generally understood to conform to ISO 22241 [8]. A pump and a dosing valve deliver it, and a tank heater and a tank temperature sensor let the system watch for freezing. A blocked line or dosing valve can be caused by crystallized fluid [24], and dosing may stop through faulty sensors, faults in the delivery system or frozen reductant. Among the data an on-board diagnostic system stores at a fault are the dosing quantity commanded and the quantity actually delivered [24]. The model's doser delivers what it is asked up to 6000 g/h, and nothing while the inlet reads below its floor.

Where the Controller Runs

In a vehicle, control runs in electronic control units, microprocessors with input and output interfaces. A modern car has more than 100 of them, and computer control entered cars in the late 1970s to meet stringent emission requirements [25]. The heavy-duty rules name an aftertreatment control module among the parts whose open circuit must trigger a derate [24]. The sources opened for this section give no scan period for such a module. The chapter's doser computes a request every 0.1 s, and the added MPC decides every 1 s. Every strategy runs in that one controller on one packet of measurements, the loops not in charge tracking the applied dose, and the observer runs its copy of the three cells beside the PI blocks. Figure 12.6 draws the loop on that hardware.

What Is Modernizing It

Three developments bear on this doser. The first is the regulation. The 2027 heavy-duty standards were shown to be achievable with cylinder deactivation, a dual-SCR arrangement and heated dosing of the fluid for low power and cold starts, and with better catalyst formulation, sizing and mixing of the fluid with the exhaust at higher power [24]. Each of these changes the plant the doser sees. The second is the sensor. Work on impedance-based NO_x sensors, read with an alternating current instead of a direct one, aimed at higher sensitivity to NO_x, better long-term stability and the ability to

subtract background interference [22]. In the model a sensor with $k_x = 0$ reads the true NO_x, and the V is gone. The third is a direct reading of the store, the subject of the second open challenge. If it holds up in service, the observer of Section 12.4 would be corrected by a measurement of the very state it now has to infer.

> **IN PRACTICE** **When the Fluid Runs Out, the Truck Slows Down**
>
> Engines that use selective catalytic reduction depend on a constant supply of the fluid. From model year 2027 a heavy-duty engine must derate when the supply falls to 2.5 % of tank capacity or to three hours of operation, when the fluid fails its concentration specification, when a catalyst appears to be missing, or on an open circuit in the tank level sensor, pump, quality sensor, NO_x sensors, dosing valve, tank heater or aftertreatment control module. The derate takes away 1 mile per hour of top speed every five minutes until the scheduled limit is reached [8]. Without the fluid, NO_x could rise to levels comparable to having no NO_x control at all [24].

12.8 Limitations

The engine and the cycle are declared maps, and the kinetics, the storage curve and the decomposition ramp are teaching values fitted to declared targets. The curve sits inside the band of one measured catalyst but is not a fit to it, and its fall with temperature makes the release likely generous. The challenge verdict depends on the declared three cells, since the same constants on six or twelve cells peak at 42 and 67 ppm at level 100 %. Only standard SCR is present, with no NO_2, N_2O, poisoning, heat loss or filter, and urea decomposition is instantaneous. The sensor's cross-sensitivity is a constant, whereas a real one varies with the sensor and the gas and can be estimated [15]. The cycle is one warm run, while type approval weights a cold and a hot one, and the OBD threshold is set below the regulation's so that a short cycle can cross it. A climb announced by a step at 10:00 and an aged brick that loses 30 % of its sites at once are not how a truck behaves, though each isolates one mechanism. The observer knows the plant's structure and kinetics exactly, so it meets model error only through κ, the thermocouple lag, the sensor noise and the ambiguity. A simulated pass is not type approval, and none of these numbers predicts the emissions of a real truck.

> **OPEN CHALLENGE** **Clean Exhaust Before the Catalyst Is Warm**
>
> In stop-and-go traffic or at idle the exhaust can cool until the catalyst system becomes less effective or stops working, and test data indicate that such low-load operation could account for more than half of a truck's NO_x in a typical workday [24]. A 2013 workshop of laboratories and industry noted that most catalyst systems began to work near 200 °C, set a goal of 90 % conversion at 150 °C, and called that lowering extremely difficult [26]. Case 5 meets the same wall: whatever the floor, the cold start is lost. Progress would be a catalyst and a dosing method that convert in the first minutes after a start.

The MPC uses frozen fresh geometry and kinetics, local Jacobians held for 60 s, an approximate thermal forecast without future reaction heat and three horizon-wide soft slacks. It has no future event access or robustness guarantee. Its opening truth-shaped teaching prior, ambiguous residual bias and missing installed slip-catalyst model restrict interpretation. The statement above about the classical observer does not imply that the MPC adapts to changed geometry or chemistry.

> **OPEN CHALLENGE** **A Gauge for the Stored Ammonia**
>
> The NO_x and ammonia sensors around a catalyst give only an inference of how much ammonia it holds, often confounded by the NO_x sensor's cross-sensitivity to ammonia. A radio-frequency method uses the catalyst housing as a microwave resonant cavity and reads the stored ammonia from the change in its dielectric properties. On one engine its readings correlated closely with the inventory estimated from gas measurements, and studies showed its potential as the feedback of closed-loop dosing [27]. A measurement that holds up in service would give this chapter's observer a reading of the very state it infers.

12.9 Conclusion

This chapter studied a diesel SCR line as a plant whose conversion is set by an unmeasured store, whose capacity falls as it heats, and whose tailpipe sensor rises on both sides of the right dose. In the model the conversion takes 11.8 minutes to follow a dose change, the sensor reads 60 ppm at doses 60 % apart, and a climb after a city crawl releases 384.5 ppm of ammonia that the sensor reports as NO_x. A trim on that sensor runs to its limit and slips 461 ppm, while storage control on an observer holds 6.3 ppm on the same seed. Over the graded cycle the storage loop passes at 0.423 g/kWh and a 23.5 ppm peak where no feedforward ratio passes, and the aged brick on a severe cycle passes at none of 19 settings. The verdict belongs to a fitted three-cell model. A fuller model would add NO_2 and the fast reaction, a particulate filter's thermal mass, a sensor cross-sensitivity that varies, a cold and hot weighted cycle, and a storage curve fitted to a measured catalyst. The simulator shows that a doser has to know how much ammonia is stored before it can read its own sensor.

The added MPC lowers cycle NOx from 0.423 to 0.365 g/kWh with preview Off but fails its 34.6 ppm peak, while its public-preview result passes at 0.395 g/kWh and 26.3 ppm.

KEY POINTS

- The catalyst stores the ammonia first, and the stored amount, which no instrument measures, sets how much of the nitrogen oxides is cleaned.
- The tailpipe sensor reads ammonia as NO_x, so two doses 60 % apart give one reading.
- A climb after a city crawl releases 384.5 ppm of ammonia that the sensor reports as NO_x.
- A trim on that sensor runs to its limit, while storage control on an observer passes the graded cycle at 0.423 g/kWh and a 23.5 ppm peak.
- The challenge verdict belongs to the fitted three-cell model.

Where else this applies. Volume 7, Chapter 4 meets a process gain that reverses at the peak of its curve, and Chapter 13 runs a full extended Kalman filter on a state that no instrument aboard measures.

References

[1] M. Koebel, M. Elsener, and M. Kleemann, "Urea-SCR: a promising technique to reduce NOx emissions from automotive diesel engines," *Catalysis Today*, vol. 59, no. 3–4, pp. 335–345, 2000.

[2] P. Forzatti, "Present status and perspectives in de-NOx SCR catalysis," *Applied Catalysis A: General*, vol. 222, no. 1–2, pp. 221–236, 2001.

[3] C. M. Schär, C. H. Onder, and H. P. Geering, "Control of an SCR catalytic converter system for a mobile heavy-duty application," *IEEE Transactions on Control Systems Technology*, vol. 14, no. 4, pp. 641–653, 2006.

[4] T. V. Johnson, "Review of diesel emissions and control," *International Journal of Engine Research*, vol. 10, no. 5, pp. 275–285, 2009.

[5] European Parliament and Council of the European Union, "Regulation (EC) No 595/2009 of 18 June 2009 on type-approval of motor vehicles and engines with respect to emissions from heavy duty vehicles (Euro VI)," Official Journal of the European Union, L 188, 18.7.2009, p. 1, 2009, annex I, emission limits. [Online]. Available: https://www.legislation.gov.uk/eur/2009/595/annex/I

[6] E. P. Brandt, Y. Wang, and J. W. Grizzle, "Dynamic modeling of a three-way catalyst for SI engine exhaust emission control," *IEEE Transactions on Control Systems Technology*, vol. 8, no. 5, pp. 767–776, 2000.

[7] A. Bonfils, Y. Creff, O. Lepreux, and N. Petit, "Closed-loop control of a SCR system using a NOx sensor cross-sensitive to NH3," *Journal of Process Control*, vol. 24, no. 2, pp. 368–378, 2014.

[8] U.S. Environmental Protection Agency, "Control of emissions from new and in-use heavy-duty highway engines," Code of Federal Regulations, Title 40, Part 1036, revised as of 1 July 2024, 2024, sections 1036.104, 1036.111 and 1036.801. [Online]. Available: https://www.govinfo.gov/content/pkg/CFR-2024-title40-vol36/xml/CFR-2024-title40-vol36-part1036.xml

[9] ——, "Control of air pollution from new motor vehicles: Heavy-duty engine and vehicle standards and highway diesel fuel sulfur control requirements; final rule," Federal Register, vol. 66, no. 12, pp. 5002–5193, 18 January 2001, FR Doc. 01-2, 2001. [Online]. Available: https://www.govinfo.gov/content/pkg/FR-2001-01-18/html/01-2.htm

[10] D. Upadhyay and M. Van Nieuwstadt, "Model based analysis and control design of a urea-SCR deNOx aftertreatment system," *Journal of Dynamic Systems, Measurement, and Control*, vol. 128, no. 3, pp. 737–741, 2006.

[11] M.-F. Hsieh and J. Wang, "Development and experimental studies of a control-oriented SCR model for a two-catalyst urea-SCR system," *Control Engineering Practice*, vol. 19, no. 4, pp. 409–422, 2011.

[12] L. Olsson, H. Sjövall, and R. J. Blint, "A kinetic model for ammonia selective catalytic reduction over Cu-ZSM-5," *Applied Catalysis B: Environmental*, vol. 81, no. 3–4, pp. 203–217, 2008.

[13] F. Willems, R. Cloudt, E. van den Eijnden, M. van Genderen, R. Verbeek, B. de Jager, W. Boomsma, and I. van den Heuvel, "Is closed-loop SCR control required to meet future emission targets?" in *SAE Technical Paper Series*, no. 2007-01-1574. SAE International, 2007.

[14] D. Upadhyay and M. Van Nieuwstadt, "Robust separation of signal domain from single channel mixed signal output of automotive urea based selective catalytic reduction systems," *Journal of Dynamic Systems, Measurement, and Control*, vol. 136, no. 1, p. 011012, 2014.

[15] H. Zhang and J. Wang, "NOx sensor ammonia-cross-sensitivity factor estimation in diesel engine selective catalytic reduction systems," *Journal of Dynamic Systems, Measurement, and Control*, vol. 137, no. 6, p. 061015, 2015.

[16] I. Nova and E. Tronconi, Eds., *Urea-SCR Technology for deNOx After Treatment of Diesel Exhausts*, ser. Fundamental and Applied Catalysis. New York: Springer, 2014.

[17] J. H. Kwak, R. G. Tonkyn, D. H. Kim, J. Szanyi, and C. H. F. Peden, "Excellent activity and selectivity of Cu-SSZ-13 in the selective catalytic reduction of NOx with NH3," *Journal of Catalysis*, vol. 275, no. 2, pp. 187–190, 2010.

[18] S. E. Bozbağ, M. Şimşek, O. Demir, D. Şanlı Yıldız, H. B. Özener, G. Hisar, and C. Erkey, "Experimental and theoretical study of NH_3 adsorption and desorption over a Cu-chabazite NH_3-SCR catalyst," *Turkish Journal of Chemistry*, vol. 42, no. 6, pp. 1768–1780, 2018.

[19] European Commission, "Commission Regulation (EU) No 582/2011 of 25 May 2011 implementing and amending Regulation (EC) No 595/2009 with respect to emissions from heavy duty vehicles (Euro VI)," Official Journal of the European Union, 2011, annex III, Appendix 1 (ammonia) and Annex X (on-board diagnostics). [Online]. Available: https://www.legislation.gov.uk/eur/2011/582/contents

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

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

[22] L. Y. Woo and R. S. Glass, "NOx sensor development," Lawrence Livermore National Laboratory, Livermore, CA, Tech. Rep. LLNL-TR-599052, 2012.

[23] ——, "NOx sensor development," Lawrence Livermore National Laboratory, Livermore, CA, Tech. Rep. LLNL-TR-461511, 2010.

[24] U.S. Environmental Protection Agency, "Control of air pollution from new motor vehicles: Heavy-duty engine and vehicle standards; final rule," Federal Register, vol. 88, no. 15, pp. 4296–4718, 24 January 2023, FR Doc. 2022-27957, 2023. [Online]. Available: https://www.govinfo.gov/content/pkg/FR-2023-01-24/html/2022-27957.htm

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

[26] M. Zammit, C. L. DiMaggio, C. H. Kim, C. Lambert, G. G. Muntean, C. H. F. Peden, J. E. Parks, and K. Howden, "Future automotive aftertreatment solutions: The 150 °c challenge workshop report," Pacific Northwest National Laboratory, Richland, WA, Tech. Rep. PNNL-22815, 2013.

[27] P. Ragaller, J. Mandelbaum, L. Lapenta, A. Sappok, J. Pihl, V. Prikhodko, and J. Parks, II, "Direct measurement of ammonia storage on passive SCR systems for lean gasoline NOx reduction using radio frequency sensing," in *Proceedings of the ASME 2019 Internal Combustion Engine Division Fall Technical Conference*, no. ICEF2019-7224, Chicago, IL, 2019.

12.10 Exercises

12.1 ⌕ Concept check · 2 min

Case 2 starts lined out on the highway at a dose ratio of 1.3, where the tailpipe sensor reads about 60 ppm. The dose is lowered to $\alpha = 0.8$ and the brick is left to settle. What does the sensor read then, and why?

(a) About 12 ppm, because the lower dose sits at the bottom of the V where NO_x and slip are both small

(b) About 60 ppm again, because the reading is now made of NO_x that the lower dose leaves unconverted

(c) About 120 ppm, because the NO_x left by the lower dose adds to the ammonia still slipping from the store

(d) About 25 ppm, because the reading falls with the slip while the tailpipe NO_x stays small

12.2 Concept check · 2 min

The storage loop reads its setpoint schedule $\theta^*(T_s)$ at the inlet thermocouple, not at the observer's brick temperature. Why does that choice matter on the climb of case 4?

(a) The inlet thermocouple is quieter than the brick estimate, so the setpoint moves less with noise

(b) The brick temperature is not available to the controller, so the schedule could not be read on it

(c) The inlet passes 350 °C about 25 s before the brick mean, so the setpoint falls and dosing stops first

(d) The inlet is hotter than the brick, so the schedule gives a higher setpoint and keeps the conversion up

12.3 Concept check · 2 min

On the highway the observer is told a capacity 30 % too high ($\kappa = 1.3$) for a fresh brick. Which result does the chapter report, and why?

(a) The coverage and the stored mass are both about 30 % high, since one scales the other

(b) The estimate diverges, because a wrong capacity flips the sign of every gradient g_i of the correction

(c) The coverage is 27 % high and the stored mass only 2 % low, since the sensor fixes the mass

(d) The stored mass is 27 % high and the coverage only 2.3 % low, since the sensor fixes the coverage

12.4 Theory and calculation · 6 min

On the 85 km/h highway ($v = 23.61\,\mathrm{m/s}$) the engine delivers $P = 108.8\,\mathrm{kW}$. Use the maps of Equation (12.1), an exhaust molar mass of 28.8 g/mol, 17.03 g/mol for ammonia and the conversion of 665 280 g/h of DEF per mole per second of urea.

(a) Find the exhaust flow, the engine-out NO_x and the molar exhaust flow N. *Checkpoint:* 0.278 kg/s, 294.1 ppm and 9.66 mol/s.

(b) Find the DEF rate for dose ratios of 1.0, 0.8 and 1.3. *Checkpoint:* 945, 756 and 1229 g/h.

(c) At full conversion, how many grams of stored ammonia does the NO_x use per minute? The store lined out at $\alpha = 1.1$ holds 12.3 g. Compare the time it would last at that rate with the 154 s for which case 3 keeps 90 % conversion after the urea stops. *Checkpoint:* 2.90 g/min, so about 4.2 min, against 2.6 min above 90 %.

12.5 Theory and calculation · 6 min

The highway point lined out at dose ratios of 0.8, 1.0 and 1.3 converts 79.4, 97.1 and 99.6 % of 294.1 ppm of engine-out NO_x and slips 0.5, 5.3 and 82.8 ppm of ammonia. The tailpipe sensor follows Equation (12.6) with $k_x = 0.7$, no drift and no noise.

(a) Find the steady sensor reading at each dose. *Checkpoint:* About 60.9, 12.2 and 59.1 ppm.

(b) Find the conversion η_s that the trim of Equation (12.7) computes at 0.8 and at 1.3, and the sign of its error against $\eta_{SP} = 95\,\%$ in each case. *Checkpoint:* 79.3 and 79.9 %, both below the target, so the trim raises the dose at both.

(c) At 705 s of case 4 the tailpipe carries 0.087 ppm of NO_x and 384.5 ppm of ammonia, the engine-out NO_x is 510.4 ppm and the sensor reads 267.4 ppm. Find the steady value the sensor is heading for, and the conversion a loop computes from the present reading. *Checkpoint:* 269.2 ppm and 47.6 %.

12.6 Theory and calculation · 7 min

The brick holds $\Omega(T)$ of Equation (12.3) with $A = 1$, $\Omega_0 = 100\,\mathrm{mol/m^3}$, $T_0 = 200\,^\circ\mathrm{C}$ and $a = 0.0015\,\mathrm{K^{-1}}$, over $V_b = 32\,\mathrm{L}$. Treat the brick as uniform in temperature and use 17.03 g/mol for ammonia.

(a) Find the capacity in grams at 220 °C, the crawl, and at 459 °C, the inlet on the climb. *Checkpoint:* 52.9 g and 33.3 g.

(b) The crawl of case 4 leaves 11.75 g stored. Find its coverage at both temperatures, and say what the rise means for the desorption rate. *Checkpoint:* 0.222 and 0.353.

(c) Find the storage loop's target in grams at both temperatures for level $L = 100\,\%$, from the schedule θ^* given below Equation (12.9). *Checkpoint:* 10.1 g and 1.52 g.

12.7 App challenge · 9 min

Load scenario 4, The heat-up release, with the default seed, and run at 20× (36 s of wall time per run).

(a) Run the feedforward as loaded past 12:00 and export the CSV. Read the slip peak on the scorecard, and from the columns `truth_nh3_slip_ppm`, `sens_nox_tail_ppm` and `truth_nox_tail_ppm` the time of the peak and the sensor and true NO_x at that row. *Checkpoint:* 384.5 ppm at 705 s (within 1 s), sensor 267.4 ppm, true NO_x 0.1 ppm.

(b) Set Dosing strategy to Storage, press Reset, and run past 12:00 again. Read the slip peak and its time, and the true stored ammonia (`truth_stored_g`) at 600 s in both runs. *Checkpoint:* 32.2 ppm at 692 s. 11.75 g under the feedforward and 10.1 g under storage control.

(c) In Advanced, set Schedule on to Brick, press Reset, and run past 12:00. Read the slip peak and explain the difference from (b). *Checkpoint:* 43.1 ppm, within 0.1 ppm.

12.8 App challenge · 8 min

Load scenario 8, The trim loop runs the wrong way, with the default seed, and run at 20× (30 s of wall time per run). Export the CSV after each run.

(a) Run the trim as loaded past 10:00. From the column `trim_pct`, find when the trim reaches its +100 % limit, and from the row at 600 s read `def_gph`, `alpha_applied` and `truth_nh3_slip_ppm`. *Checkpoint:* About 256 s (within 1 s); 2458 g/h, $\alpha = 2.60$, 461 ppm.

(b) Set Dosing strategy to Storage, press Reset, and run past 10:00. Read the same three columns and `truth_nox_tail_ppm` at 600 s. *Checkpoint:* 918 g/h, $\alpha = 0.97$, 6.3 ppm of ammonia and 11.0 ppm of NO_x.

(c) Explain from the sensor's V why the same tailpipe sensor drives one loop to its limit while the observer, which reads that sensor too, brings the dose back near $\alpha = 1$. *Checkpoint:* The trim's gain has the wrong sign right of the V's bottom; the observer's model knows the side.

12.9 App challenge · 9 min

Load the Challenge, the 30 min cycle, with the default seed. Each run pauses at 30:00 with its verdict; run at 200× (9 s of wall time per run). Read the challenge tile.

(a) Run as shipped (feedforward with the trim). Read the cycle NO_x, the slip mean and peak, the DEF and which limit fails. *Checkpoint:* 0.293 g/kWh, mean 108.3 ppm, peak 1301.9 ppm, 893 g; the ammonia mean and peak fail.

(b) Set Dosing strategy to Storage, leave its settings, press Reset and run. Read the same four numbers and the verdict. *Checkpoint:* 0.423 g/kWh, mean 5.8 ppm, peak 23.5 ppm, 470 g; met.

(c) Set Storage level to 115 %, press Reset and run, then repeat at 120 %. Explain the trade the level makes. *Checkpoint:* 115 %: 0.362 g/kWh, peak 30.0 ppm on the tile, met. 120 %: 0.344 g/kWh, peak 32.5 ppm, not met.

12.10 Application area · 4 min

Euro VI weights the transient cycle 0.14 cold and 0.86 hot and allows 0.46 g/kWh of NO_x on the weighted result.

(a) Storage control leaves 0.423 g/kWh on the warm cycle of the challenge. Taking that as the hot result, find the most the cold run may leave. *Checkpoint:* 0.687 g/kWh.

(b) Case 5 leaves 3.03 g/kWh over the first 10 min of a cold start at a 200 °C floor. If a whole cold run left that much, what hot result would be needed? What does case 5 say about lowering the floor as the remedy? *Checkpoint:* 0.042 g/kWh. Lowering the floor to 150 °C gains 3 % for three times the deposits.

12.11 Application area · 4 min

The on-board diagnostic of the app flags a run when the tailpipe sensor's own cycle number exceeds a teaching threshold of 0.55 g/kWh. Over the full cycle the sensor's number is 0.586 g/kWh for the aged brick under storage control, 0.490 g/kWh for the fresh brick under storage control and 0.792 g/kWh for the fresh brick under the plain feedforward, whose true NO_x is 0.564 g/kWh.

(a) Which runs are flagged, and which of them has a failing catalyst? *Checkpoint:* Two flagged; one failing catalyst.

(b) Why does the fresh brick on the feedforward read higher than its true NO_x, and by how much? What must a diagnostic know to tell aging from overdosing? *Checkpoint:* About 40 % high, from slip read as NO_x.

12.12 Review the agent · 5 min

An AI assistant compares the two strategies of case 9 for a calibration team: "With Drive preview Off, the MPC leaves 0.365 g/kWh of NO_x against 0.423 g/kWh for Storage, for nearly the same DEF, 478 g against 470 g. The MPC is therefore the better doser on this cycle, and switching to it would pass the challenge with a wider NO_x margin." Find the flaw, and say how you would check it in the app.

CHAPTER 13

Navigation Drift as Unimaged Seabed in an AUV Survey Simulator

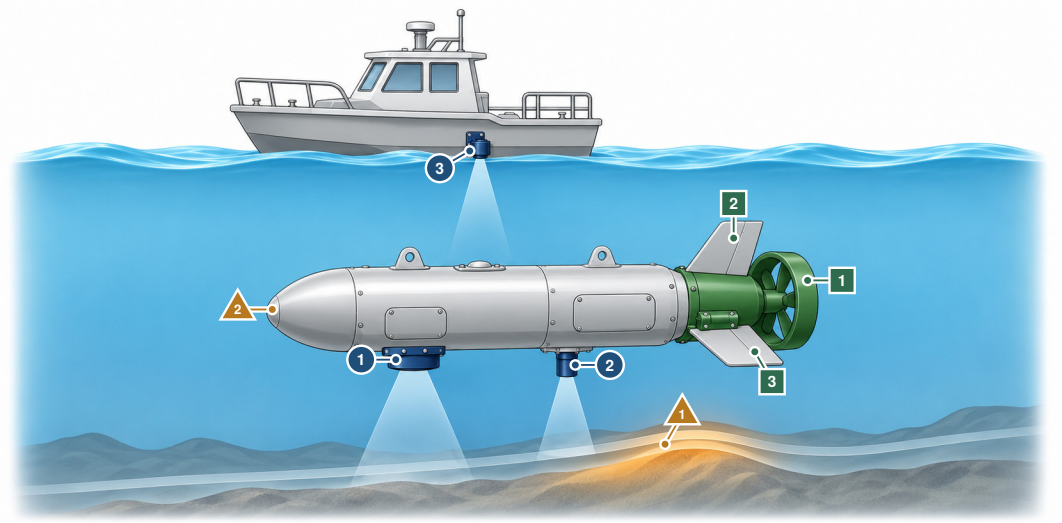

The controller moves 1 thruster that turns the propeller 2 rudder, moved by a servo 3 stern plane, moved by a servo

The controller reads 1 Doppler log: speed over the seabed or through the water 2 altimeter: height above the seabed 3 ship's acoustic transceiver: position fixes that arrive late

What disturbs it 1 ridge the vehicle must climb 2 water current that carries the vehicle off its line

The vehicle steers on what its Doppler log, its compass and the ship's late fixes tell it, so a current or a compass error can carry it off its lines while every reading looks normal. Illustration, not to scale.

IN PLAIN TERMS **A Robot That Maps the Seabed Blind**

The main idea. A robot submarine flies back and forth along planned lines and maps the seabed with sound. Under water no radio or satellite signal reaches it, so it works out where it is from its speed and heading, with a sound signal from a ship now and then.

The challenge. Small errors in that reckoning add up, and the robot steers where it believes it is, not where it really is. A strip of seabed can be missed while every reading on board looks normal.

Why it matters. The map is only as complete as the robot's knowledge of where it went.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on the deep surveys (Section 13.1), the history in Section 13.2, the try-it on the cross current and the try-it on the canyon, the fact on the late fixes (Section 13.5) and the key points
Study	50 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 13.5), Technology in Practice (Section 13.6), Limitations (Section 13.7) and the two open challenges

The control question A survey vehicle under water steers its own estimate of where it is along the planned lines, while the seabed its sonar images is set by where it actually went. Which errors open unimaged strips that nothing aboard can see, and which measurement, guidance or tuning removes each one?

Learning objectives

- Predict the dead-reckoning error of a constant compass bias from the distance to the last fix, and explain why a lawnmower pattern cancels most of it.
- Compute the swath, the overlap that lines spaced 30 m apart leave, and the uncertainty at which the supervisor stops vouching for coverage.
- Compute the steady offset that line-of-sight guidance leaves in a cross current, and show how crab compensation removes it.
- Explain why a drift on water track is unobservable without a fix or a measured current, and why no tuning recovers the strip it leaves.
- Judge when a filter that applies late fixes behind a chi-square gate will refuse honest fixes, at arrival or after a wild first fix.

Prerequisites Sections 1.4.5, 1.4.6, 1.4.7, 1.6.2 and 1.6.4.

The app https://apmonitor.com/apps/auv-navigation/

See also Volume 20, Chapter 6 for a vessel whose observer carries the slowly varying wind and current as a bias state, Volume 20, Chapter 10 for late camera reports moved on by the belt travel since their capture, and Volume 20, Chapter 2 for a path-following loop on a geometric cross-track error.

13.1 Introduction

Autonomous underwater vehicles map the seafloor at a resolution that surface vessels cannot reach, in deep water especially, and they have become routine instruments of marine geoscience, habitat mapping and seabed survey [1]. They do this without being able to measure where they are. Satellite and radio signals do not reach them, so a submerged vehicle integrates its velocity and heading and corrects the result with acoustic positioning, installed beacons or trips to the surface [2]. The control problem that follows has an unusual shape. The autopilots steer the vehicle's estimate of its position along the planned lines, while the product the survey exists for, the seabed its sonar images, is set by where the vehicle actually went.

+ DID YOU KNOW? 191 Surveys at Mid-Ocean Ridges

Over one decade a single deep-sea survey vehicle conducted 191 surveys of the seafloor at mid-ocean ridges, at an average depth of 2000 m and with its navigation precise to a few meters. Its bathymetric and magnetic maps, photographs and maps of hydrothermal plumes all rested on that navigation [3].

The simulator is aimed at one misconception, that a vehicle whose estimate holds its lines has imaged the seabed between them. In the canyon case of Section 13.5.5, 420 s into the run, the estimate is 3.3 m from line 7 while the vehicle itself is 24.6 m from it, and a strip 25 m wide and 260 m long along the edge of the southern block has been left unimaged. Every trend of the vehicle's own signals is unremarkable at that moment. The estimated heading is 3.1° from its line-of-sight setpoint and the speed through the water 0.01 m/s from its own. A block diagram cannot show the difference, because its feedback path carries the measured output, and the output of a survey is a function of a state that no instrument on the vehicle measures. The simulator therefore keeps three things apart, the true vehicle, what its instruments report and what its filter believes, and paints the survey product from the first.

The class of problem is wider than the ocean. Any machine that covers an area it cannot observe, while steering on its own estimate, belongs to it, and coverage path planning is treated as a problem

of robotics in general [4]. Its navigation half is dead reckoning corrected by absolute fixes that are intermittent and late, which a discrete-time Kalman filter can absorb in several ways at different computational costs [5], and occasionally wild enough that a validation gate must refuse them [6]. Its guidance half is path following by an underactuated craft in a current, where a line-of-sight law without integral action or a drift estimate settles off its path [7, 8]. A process loop that holds an inferential estimate on its setpoint while the unmeasured product drifts has the same structure, the soft sensor of Section 1.4.8.

However, the literature takes the two halves apart. Navigation is reviewed as localization [2] and guidance is analyzed as path following [7, 8], and neither measure is the one a survey is judged by, the seabed actually imaged, which depends on both and on the altitude besides. This chapter scores that dependence on a survey simulator that keeps the truth, the instruments and the estimate apart, and grades every run on the seabed imaged. The maneuvering model, the extended Kalman filter, line-of-sight guidance and the chi-square gate are all standard. The simulator combines them into a 14-state vehicle on Prestero's coefficients, instruments that form the only path from the truth to the controller, a filter that applies each delayed fix at the instant it describes by rewinding a 24 s history, a coverage grid painted from the truth under an explicit holiday rule, and a supervisor that reports from its own covariance when coverage can no longer be vouched for. The chapter sets these down in enough detail to be rebuilt. Eight cases and a graded survey sort the errors by their owners, the hull, the guidance, the heading reference, the current, the fixes and the vertical loops.

Section 13.2 recounts how the domain arrived at aided dead reckoning. Section 13.3 builds the vehicle, the seabed, the survey product and the instruments, and Section 13.4 the navigation filter, the guidance, the four autopilots and the mission supervisor. Section 13.5 reports each case and the graded survey. Section 13.6 describes the real instruments and computers, Section 13.7 states the limitations and Section 13.8 draws the conclusions.

13.2 From Transponder Nets to Aided Dead Reckoning

Underwater vehicles first navigated by listening to sound sources placed around them. A pulse-Doppler system tested at sea in 1976 tracked underwater and surface platforms continuously inside a fixed 50 km^2 net of three reference buoys moored 20 m above the bottom, each carrying a transponder for travel times and a continuous-tone beacon for the Doppler shift, and it placed a platform to within 2 to 3 m of the net [9]. The accuracy came at the cost of the net, which had to be laid before any work began. By 1998 commercial acoustic positioning systems offered full-ocean-depth positioning to 5 m absolute and better than 2 m relative accuracy, and an overview of how they were configured still concluded that the vehicle community needed further development [10].

Dead reckoning moved to the center once Doppler sonar became affordable. Low-cost commercial bottom-lock Doppler sonars were reported in 1999 to augment or replace the acoustic time-of-flight systems then common, and the same work analyzed the effect of heading-sensor errors on bottom-track precision, comparing a low-precision magnetic heading sensor with a ring-laser gyroscope on a survey at 2000 m depth [11]. Small survey vehicles were built around this instrument suite. REMUS, described in 1997, was 19 cm in diameter, 134 cm long and 31 kg, had a range of 25 nautical miles at 3 knots on 400 Wh of lead-acid batteries, and navigated by long and ultra-short baseline acoustics, bottom-lock Doppler and GPS [12]. A six-degree-of-freedom model of that vehicle followed in 2001, its coefficients derived from hydrostatics, lift, drag, added mass and the propeller and fin inputs, and its output verified against motion data collected at sea [13]. Its REMUS coefficients are the declared starting set of this simulator.

The answer in current use is an integrated filter, and the history segment after this paragraph describes an early one. Model-aided inertial navigation later brought a kinetic vehicle model and a real-time sea current estimate into a filter that already took USBL positions, pressure and bottom-track velocity, and was evaluated with the USBL and the Doppler log removed or dropped out [14]. A

Table 13.1 Signals of the Survey by Role, With Their Reference Values

Symbol	Role	Reference	Range
$\hat{\psi}$	CV of HDG, estimated heading	ψ_d from LOS	0–360°
u_m	CV of SPD, DVL water track	1.5 m/s SP	0.5–2.3
a_m, z_m	CV of DEP, altimeter or depth	10 m, 38 m SP	4–25, 5–58
θ_m	CV of PIT, measured pitch	θ_{SP}	±20°
$\hat{e}_\perp$	CV of LOS, estimated cross-track	0	—
δ_r, δ_s	MVs, rudder, stern plane	0	±15°
T	MV, thrust	3.645 N trim	0–10 N
θ_{SP}	MV of DEP, pitch setpoint	—	±20°
V_c, set	DV, current and its direction	0, toward 180°	0–0.6 m/s
b	DV, compass bias, unmeasured	0	±3°
ε_D	DV, DVL scale error, unmeasured	0	±2 %
—	DV, USBL lost pings, outliers	10 %, 0	0–80, 0–30 %
k_σ	DV, USBL noise multiple	1	0.2–5
—	DV, USBL outage, ship station	off, (225, 130) m	event, sliders
u, v, r, q, θ, ψ	states, motion	trim	—
n, e, z	states, position and depth	start pose	—
δ_r, δ_s, T	states, actuators	0, 0, trim	as MVs
E, ℓ	states, energy used, distance	0	—
$\hat{n}, \hat{e}, \hat{b}, \hat{V}_{cn}, \hat{V}_{ce}$	filter states	start pose, 0	—
C, holidays	product, from the truth	—	—

review in 2014 still found every alternative fundamentally limited, whether expensive inertial sensors, installed beacons or periodic surfacing [2].

> ⧖ HISTORY **2004: A Lawnmower Pattern Cancels the Drift**
>
> In 2004 a navy's new mine countermeasures vehicle was reported to navigate on an inertial system aided by a Doppler velocity log. An extended Kalman filter, built on a model of the errors, combined the inertial solution with the Doppler log, a pressure sensor and an optional compass, and took GPS fixes at the surface, ranges to transponders, terrain navigation or acoustic fixes from its mother ship when they were available. The growth of its position error depended on the inertial and Doppler errors and on the mission pattern. Low-frequency Doppler errors caused an approximately linear drift on a straight line and tended to cancel on a lawnmower pattern, and the Doppler log's scale error became observable to the filter when position fixes were available or the vehicle turned [15].

Guidance developed beside navigation. An autopilot for the combined speed, steering and diving of a slow vehicle, designed by multivariable sliding modes, was acquiring waypoints by line of sight in 1993 [16]. Line of sight remained the practical law, and because it is susceptible to environmental disturbances, integral action was added to hold a straight path in a constant current [7], as was an adaptive estimate of the sideslip that currents, wind and waves produce [8]. Each of these systems is judged in the end by the chart its vehicle brings back, and the simulator scores its vehicle the same way, from positions the vehicle never sees.

13.3 The Vehicle, the Seabed and the Instruments

The simulated vehicle is a REMUS-class torpedo 0.19 m across and of 30.48 kg, cruising at 1.5 m/s and 10 m above the seabed. It surveys two blocks 260 m wide and 150 m deep, one on each side of a canyon, on ten east–west lines 30 m apart, under a ship that holds station over the canyon, as shown in Figure 13.1. Coordinates are north n, east e and depth z in meters from the south-west corner of the survey, with headings measured clockwise from north.

Table 13.1 sorts the signals by the part each plays. The plant's 14 states are the surge, sway and yaw rate relative to the water, the north and east position, the heading, the pitch rate and pitch, the depth, the rudder and stern-plane angles behind their servos, the thrust behind its motor lag, the energy used and the distance flown.

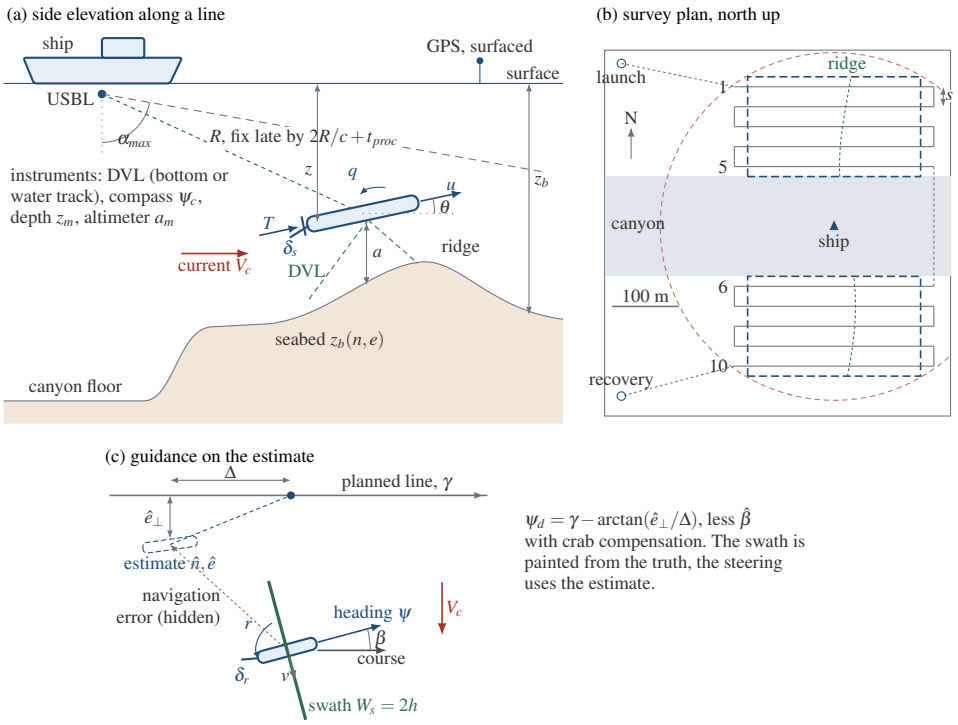

$\psi_d = \gamma - \arctan(\hat{e}_\perp/\Delta)$, less $\hat{\beta}$ with crab compensation. The swath is painted from the truth, the steering uses the estimate.

Figure 13.1 Survey vehicle, seabed and plan as modeled. (a) Side elevation along a line, schematic: the vehicle at depth z and pitch θ, with surge u, pitch rate q, thrust T and stern plane δ_s, at altitude $a = z_b - z$ above the seabed $z_b(n,e)$, its Doppler log (DVL) beams, and the ship's USBL transceiver at slant range R inside a cone of half-angle α_{max} from the vertical. (b) The survey plan to scale: two blocks, ten lines $s = 30\,\text{m}$ apart with 20 m run-ins, the dotted dive, canyon and return legs, the launch and recovery points, the ridge crest, the canyon with the ship over it, and the USBL footprint at 50 m depth. (c) Guidance on the estimate: a line of direction γ, the estimated cross-track error $\hat{e}_\perp$ and lookahead Δ, the hidden navigation error to the true vehicle, whose heading ψ, rudder δ_r, sway v and yaw rate r leave its course an angle β off the heading in the current V_c, and the swath $W_s = 2h$ painted from the truth.

13.3.1 Maneuvering in the Horizontal Plane

Surge, sway and yaw follow Fossen's form [17] with the REMUS coefficients [13], as shown in Equation (13.1).

$$\begin{aligned} m_{11}\dot{u} &= T + X_{u|u|}u|u| + m_{22}vr \\ m_{22}\dot{v} - Y_{\dot{r}}\dot{r} &= Y_{uv}uv + Y_{v|v|}v|v| + (Y_{ur} - m)ur \\ &\quad + Y_{r|r|}r|r| - Y_\delta\, u|u|\,\delta_r \\ -N_{\dot{v}}\dot{v} + m_{33}\dot{r} &= N_{uv}uv + N_{v|v|}v|v| + N_{ur}ur \\ &\quad + N_{r|r|}r|r| + N_\delta\, u|u|\,\delta_r \end{aligned} \tag{13.1}$$

where u, v and r are relative to the water, $m_{11} = m - X_{\dot{u}} = 31.41\,\text{kg}$, $m_{22} = m - Y_{\dot{v}} = 65.98\,\text{kg}$ and $m_{33} = I_z - N_{\dot{r}} = 8.33\,\text{kg}\,\text{m}^2$. The cross terms Y_{ur} and N_{ur} carry the added-mass contributions, and $N_{uv} = -24.0\,\text{kg}$ carries the Munk moment, the pressure moment that turns a slender body broadside to the flow. Y_δ and N_δ are magnitudes, and the signs make $\delta_r > 0$ swing the bow to starboard, the opposite of Prestero's convention. The signed square $u|u|$ keeps the fins' sense in reverse.

Linearized about the straight line at 1.5 m/s, the sway and yaw pair has eigenvalues of $+0.995$ and $-2.27\,\mathrm{s}^{-1}$, because the Munk moment outweighs the fins' restoring moment. The straight line is therefore an unstable equilibrium, and a disturbance doubles every 0.70 s until crossflow drag balances it in a steady turn. This is the distinctive fact of the hull, and it makes the heading autopilot part of the vehicle rather than an option.

13.3.2 Pitch, Depth, Actuators and Energy

The vertical plane keeps pitch and neglects heave, as shown in Equation (13.2).

$$\begin{aligned} m_{55}\dot{q} &= -M_\delta\, u|u|\,\delta_s + M_{uq}uq + M_{q|q|}q|q| - \overline{BG}\,W\sin\theta \\ \dot{\theta} &= q, \qquad \dot{z} = -u\sin\theta \end{aligned} \tag{13.2}$$

where $m_{55} = I_y - M_{\dot{q}} = 8.33\,\mathrm{kg\,m}^2$, $\delta_s > 0$ is trailing edge down (nose down, a dive) as in Prestero, and the center of gravity $\overline{BG} = 19.6\,\mathrm{mm}$ below the center of buoyancy gives a righting moment of 5.86 N m per radian and a natural period of 7.5 s. The vehicle flies along its pitch attitude. At 1.5 m/s a stern plane of $+5^\circ$ holds $\theta = -11.9^\circ$ and a dive of 0.31 m/s, and a pitch of 20° needs 8.3° of plane, while at 1.0 m/s it needs 18.7°, more than the 15° the planes can move.

The vehicle moves over the ground with the water, as shown in Equation (13.3).

$$\begin{aligned} \dot{n} &= u\cos\theta\cos\psi - v\sin\psi + V_{cn} \\ \dot{e} &= u\cos\theta\sin\psi + v\cos\psi + V_{ce}, \qquad \dot{\psi} = r \end{aligned} \tag{13.3}$$

where the uniform current (V_{cn}, V_{ce}) follows its setting through a 30 s first-order lag, and $\dot{\ell} = (\dot{n}^2 + \dot{e}^2)^{1/2}$ accumulates the distance flown. The actuators and the battery are given by Equation (13.4).

$$\begin{aligned} \tau_f\dot{\delta}_j &= \mathrm{sat}_{\delta_{max}}(\delta_{j,cmd}) - \delta_j, \quad j \in \{r,s\} \\ \tau_T\dot{T} &= \mathrm{sat}_{[0,T_{max}]}(T_{cmd}) - T \\ \dot{E} &= P_{hotel} + \max(T,0)\max(u,0)/\eta_p \end{aligned} \tag{13.4}$$

with $\tau_f = 0.1\,\mathrm{s}$, $\delta_{max} = 15^\circ$, $\tau_T = 0.4\,\mathrm{s}$, $T_{max} = 10\,\mathrm{N}$, $P_{hotel} = 30\,\mathrm{W}$ and $\eta_p = 0.35$. At 1.5 m/s the trim thrust $-X_{u|u|}u^2$ is 3.645 N, the propulsion 15.6 W and the total draw 45.6 W, so the 1000 Wh battery would last 21.9 h, with E kept in watt-hours.

13.3.3 The Seabed and the Plan

The seabed is an analytic height field, as shown in Equation (13.5).

$$\begin{aligned} z_b &= 60 + \xi - h_r\, g\,(1-c) + 95\,c, \quad g = \mathrm{e}^{-(e-e_r)^2/2\sigma_r^2} \\ e_r &= 150 + 12\sin(2\pi n/420) \\ c &= S\!\left(\tfrac{n-150}{35}\right) S\!\left(\tfrac{300-n}{35}\right) \\ \xi &= 0.9\sin(\tfrac{2\pi e}{83} + 0.6)\sin(\tfrac{2\pi n}{127} + 1.1) \\ &\quad + 0.6\sin(\tfrac{2\pi(e+0.6n)}{59} + 2.0) \end{aligned} \tag{13.5}$$

in meters, where $S(x) = x^2(3-2x)$ on $[0,1]$ and is 0 below and 1 above, $h_r = 16\,\mathrm{m}$ and $\sigma_r = 24\,\mathrm{m}$. A 60 m shelf carries small undulations ξ, a ridge crosses both blocks from north to south with its crest 43 to 45 m deep, and a canyon 95 m deeper than the shelf runs east to west between the blocks, its walls falling over 35 m, far steeper than a vehicle held to 20° of pitch can follow. On the planned lines, the faces that the vehicle climbs in each line's own direction rise at 17.9 to 24.2°. The plan starts at a launch point north-west of the north block, dives to line 1, flies lines 1 to 5 of the north block east and west in turn with 30 m turn legs, crosses the canyon on its own leg at the depth the leg began at, flies lines 6 to 10 of the south block, and returns to a recovery point south-west of it. Each line runs 20 m past both sides of its block.

13.3.4 The Survey Product

The sonar's swath follows the altitude, as shown in Equation (13.6).

$$h(a) = \min\left(a\tan\theta_s,\ \sqrt{R_s^2 - a^2}\right), \qquad W_s = 2h \tag{13.6}$$

for $a_{min} \le a < R_s$ and $h = 0$ outside that range, with $\tan\theta_s = 2$, $R_s = 50\,\text{m}$ and $a_{min} = 2\,\text{m}$. The swath is 40 m at 10 m of altitude, widest at 89.4 m at 22.4 m, and nothing at 50 m. After every step it is painted from the true position, the true heading and the true altitude onto a grid of 1 m cells covering the chart, along a segment perpendicular to the heading sampled every half cell, while the sonar logs on a survey line. Each cell counts the distinct lines that imaged it. The imaged share C is the fraction of the two blocks' 78 000 cells imaged at least once. A holiday is a run of unimaged cells across the lines, within one column of a block, more than 2 m wide, bounded on both sides by imaged seabed or by the block's edge, inside the part of the block already flown, and at least 3 m long once neighboring columns are merged. A column is examined 8 s after the swath last marked it, so a gap still being crossed is not counted. At 10 m of altitude lines 30 m apart overlap by 10 m, 5 m on each side, and that margin has to absorb the cross-track control error, the current's offset, the sag over the ridge and the navigation error together.

13.3.5 Instruments

The instruments are the only path from the truth into the controller, and every random draw is made in them. The Doppler log reports its velocity in the heading frame, over the ground on bottom track and through the water on water track, as shown in Equation (13.7).

$$\begin{aligned} \begin{bmatrix} f_{bt} \\ s_{bt} \end{bmatrix} &= k \begin{bmatrix} \cos\psi & \sin\psi \\ -\sin\psi & \cos\psi \end{bmatrix} \begin{bmatrix} \dot{n} \\ \dot{e} \end{bmatrix} + w_{bt} \\ \begin{bmatrix} f_{wt} \\ s_{wt} \end{bmatrix} &= k \begin{bmatrix} u\cos\theta \\ v \end{bmatrix} + w_{wt}, \qquad \psi_c = \psi + b + \nu_\psi \end{aligned} \tag{13.7}$$

where $k = 1 + \varepsilon_D$ is a common scale error, w_{bt} and w_{wt} are white noise of 0.01 m/s and 0.03 m/s, and ψ_c is the compass, with a constant bias b and colored noise ν_ψ of 0.1° and 2 s. Bottom track exists while $0.3 < a \le 30\,\text{m}$, and neither exists shallower than 1 m, in the surface bubbles. Pitch is measured to 0.05° (colored, 1 s), depth to 0.02 m, and altitude to 0.05 m (colored, 1 s) within 50 m. GPS reports the position to 1.5 m per axis once a second after 20 s with the antenna out of the water ($z < 0.5\,\text{m}$).

The ship interrogates the vehicle's transponder every 15 s. A fix describes the instant t_f of the reply and reaches the vehicle after the reply has traveled to the ship, been processed there and been sent back to the vehicle, as shown in Equation (13.8).

$$\begin{aligned} z_{fix} &= \begin{bmatrix} n \\ e \end{bmatrix}_{t_f} + \sigma_R\, w + o, \qquad \sigma_R = \max(0.2,\ 0.005\,R)\, k_\sigma \\ t_{arr} &= t_f + 2R/c + t_{proc} \end{aligned} \tag{13.8}$$

where R is the slant range from a transducer 4 m below the surface, $c = 1500\,\text{m/s}$, $t_{proc} = 1.0\,\text{s}$, w is a unit normal pair, k_σ a noise multiple of 1, and o an outlier offset, 15 to 40 m in a random direction at a set rate or 50 m on demand. No fix returns from outside the cone of $\alpha_{max} = 80°$ from the vertical, from above 2 m, during an outage, or on one ping in ten. A fix is therefore a sampled measurement with a reporting delay, in the terms of Section 1.4.5, taken every 15 s and describing an instant 1.2 to 1.4 s before it arrives.

Table 13.2 Parameters. The Class Marks Sliders Under Show Model (Adjustable), Values Kept Out of It (Hidden), Constants Set in the Source (Fixed) and Panel Controls

Symbol	Quantity	Value	Class
m, W	mass, weight	30.48 kg, 299 N	adj., hid.
I_y, I_z	pitch, yaw inertia	3.45 kg m^2	hidden
$X_{\dot{u}}$, $Y_{\dot{v}}$	added mass	−0.93 kg, −35.5 kg	hidden
$N_{\dot{r}}$, $M_{\dot{q}}$	added inertia	−4.88 kg m^2/rad	hidden
$Y_{\dot{r}}$, $N_{\dot{v}}$	cross terms	1.93 kg m/rad, 1.93 kg m	hidden
$X_{u\|u\|}$	axial drag	−1.62 kg/m	adjustable
Y_{uv}, $Y_{v\|v\|}$	lift, crossflow drag	−28.6, −131 kg/m	hidden
Y_{ur}, $Y_{r\|r\|}$	sway from yaw	5.22 kg/rad, 0.632 kg m/rad^2	hidden
Y_δ	rudder lift	9.64 kg/(m rad)	hidden
N_{uv}, N_{ur}	Munk, yaw damping	−24.0 kg, −2.0 kg m/rad	adjustable
$N_{v\|v\|}$, $N_{r\|r\|}$	yaw crossflow	−3.18 kg, −94 kg m^2/rad^2	hidden
N_δ, M_δ	rudder, plane moment	6.15 kg/rad	adjustable
M_{uq}, $M_{q\|q\|}$	pitch damping	−2.0 kg m/rad, −188 kg m^2/rad^2	hidden
$\overline{BG}$	CG below CB	19.6 mm	adjustable
δ_{max}, T_{max}	fin, thrust limits	15°, 10 N	hidden
τ_f, τ_T	fin servo, motor lags	0.1, 0.4 s	adjustable
η_p, P_{hotel}	propulsion, hotel load	0.35, 30 W	adjustable
Q, Q_0	battery, charge at launch	1000 Wh, 100 %	adjustable
—	reserve that aborts	5 %	adjustable
$\tan\theta_s$, R_s	swath, slant range	2, 50 m	adjustable
a_{min}, cell	imaging floor, grid	2, 1 m	hidden
—	holiday width, altitude floor	2, 3 m	adjustable
α_{max}, t_{proc}	USBL cone, processing	80°, 1.0 s	adjustable
—	DVL lock range, noise	30 m, 0.01, 0.03 m/s	fixed
—	compass, pitch noise	0.1° (2 s), 0.05° (1 s)	fixed
—	depth, altimeter noise	0.02, 0.05 m (1 s)	fixed
—	GPS fix, noise	after 20 s, 1.5 m	fixed
—	USBL noise, lost pings	0.005R, 10 %	fixed, DV
q_p, q_b, q_V	filter process noise	5×10^{-4}, 1×10^{-9}, 4×10^{-6}	fixed
σ_b, σ_V	filter priors	2°, 0.1 m/s	control
Δt, Δt_c	step (RK4), controller	0.05, 0.1 s	fixed
—	historian, seed	2 Hz, 20261103	fixed

q_p in m^2/m, q_b in rad^2/s, q_V in m^2/s^3. The vehicle rows are Prestero's REMUS set.

13.3.6 Initial State and Numerics

Table 13.2 collects the parameters. Each case opens trimmed, either on a leg at a stated along-track distance and altitude or depth, or floating at the launch point, with the surge at its setpoint, the thrust equal to the drag, the fins at zero, the current already at its setting and the filter on the true position with a standard deviation of 1.5 m for a start at the surface or 1.0 m for a start underway. The engine of Section 1.6.1 advances the plant by classical fourth-order Runge-Kutta at $\Delta t = 0.05$ s, runs the instruments, filter, supervisor and loops every $\Delta t_c = 0.1$ s, keeps two rows a second in the historian and draws every random number from one generator seeded with 20261103. After each step the vehicle is held at the surface once it reaches it, stopped if its hull touches the seabed and left without its motor once the battery is empty, and the swath is painted. Halving the step moves the lowest altitude of case 8 by 0.013 m and the largest navigation error of case 4 by 0.01 %, and cells of 0.5 m move case 4's imaged share by 0.058 percentage points and its widest holiday from 7 to 7.5 m. Headless, a step takes about 3.5 µs, so the 43 min survey runs in under 0.2 s, and at 60 times real time it plays in 43 s.

13.3.7 Assumptions

Prestero's constant coefficients. They are taken as the declared set for a REMUS-class hull, constant with speed and with linear fin lift up to the 15° limit. The yaw instability they produce is kept as declared, and the heading loop is tuned for this model rather than for a particular vehicle. Fin stall, a

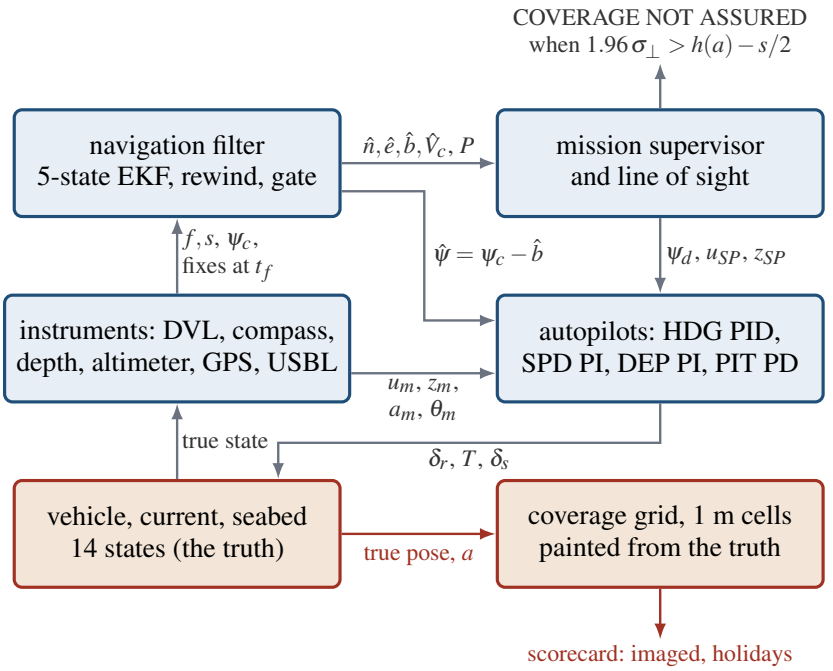

Figure 13.2 Control structure as implemented. The instruments are the only readers of the true state. The filter's estimate feeds the supervisor and, through the bias-corrected heading, the heading loop, and the coverage grid is painted from the truth and read by nothing that steers.

speed-dependent Munk moment and propeller torque are absent, so the model is least trustworthy near the fin limit and far from the cruise speed.

No heave and no roll. The vehicle flies along its pitch attitude with no angle of attack, so depth changes only through pitch, and a turn does not bank the swath. Both fail in steep, fast maneuvers.

A uniform current. The relative-velocity form of Equation (13.3) is exact only for a current constant in space and slow in time. Water track measures exactly the velocity through the water, so the current that bottom track less water track reveals is the current at the vehicle. With shear it would be the current at the tracking layer.

A smooth, charted seabed and a nadir swath. The supervisor knows the seabed of Equation (13.5) exactly. The swath is computed from the altitude straight below, as if the seabed under it were flat, and the sonar's nadir gap, resolution, image quality and sensitivity to pitch and roll are not modeled. Coverage is a yes-or-no grid, and on a steep or rough seabed the nadir altitude misstates the swath.

Generic instruments. The noise and bias models stand for no particular product. The compass bias is constant and independent of heading, and the DVL scale error is not among the filter's states. A heading-dependent compass error, or a scale error, would leave a residual the filter cannot learn.

Teaching thresholds. The line spacing, the 3 m altitude floor, the 2 m holiday, the battery reserve and every supervisor threshold are teaching values. A real survey takes its coverage and gap criteria from a standard such as IHO S-44, and the battery is treated as a store of energy drawn by a constant hotel load and one propulsive efficiency.

13.4 Navigation, Guidance and the Autopilots

Every 0.1 s the controller reads the instruments, advances and corrects the navigation filter, lets the supervisor choose the leg and the setpoints, and runs four loops, as shown in Figure 13.2. Nothing in it reads the true position, heading or altitude.

13.4.1 Dead Reckoning and the Filter

One extended Kalman filter, the filter of Equation (1.11) linearized at every step, carries $X = [\hat{n}, \hat{e}, \hat{b}, \hat{V}_{cn}, \hat{V}_{ce}]^\top$, the position, the compass bias and the current, in every navigation mode. It predicts on the Doppler log and the compass, as shown in Equation (13.9).

$$\begin{aligned} \hat{\psi} &= \psi_c - \hat{b}, \qquad \begin{bmatrix} d_n \\ d_e \end{bmatrix} = \begin{bmatrix} \cos\hat{\psi} & -\sin\hat{\psi} \\ \sin\hat{\psi} & \cos\hat{\psi} \end{bmatrix} \begin{bmatrix} f \\ s \end{bmatrix} \\ \hat{n}^+ &= \hat{n} + (d_n + \omega\hat{V}_{cn})\Delta t_c, \quad \hat{e}^+ = \hat{e} + (d_e + \omega\hat{V}_{ce})\Delta t_c \\ P^+ &= FPF^\top + \operatorname{diag}(q_p\delta\ell, q_p\delta\ell, q_b\Delta t_c, q_V\Delta t_c, q_V\Delta t_c) \end{aligned} \tag{13.9}$$

where (f, s) is the velocity the Doppler log reports, $\omega = 0$ on bottom track, which already contains the current, and $\omega = 1$ on water track, which does not, F is the identity plus $d_e\Delta t_c$ and $-d_n\Delta t_c$ in the position rows of the bias column and $\omega\Delta t_c$ in those of the current columns, and $\delta\ell$ is the distance moved in the step (plus $10^{-8}\,\mathrm{m}^2$ per step). With no Doppler velocity in the surface bubbles, f is the propeller speed $(T_{cmd}/-X_{u|u|})^{1/2}$, $s = 0$ and $\omega = 1$. The bias and the current start at zero with standard deviations $\sigma_b = 2^\circ$ and $\sigma_V = 0.1\,\mathrm{m/s}$. Dead reckoning is this prediction alone. On bottom track with a constant bias, the estimated velocity is the true one rotated by b, so the estimated track is the true one rotated by b about the last fix, and the error follows Equation (13.10).

$$|\hat{x} - x| = 2\sin(b/2)\, d \tag{13.10}$$

Here d is measured from the last fix rather than along the track flown, and $2\sin(0.5^\circ) = 1.745\,\%$. On water track with $\hat{V}_c = 0$ the error grows instead at the current's own speed.

The aided mode applies each USBL fix at the instant it describes. The filter keeps 240 steps (24 s) of its states, covariances and inputs, finds the newest stored step k_f at or before t_f, and updates there by Equation (13.11),

$$\begin{aligned} \nu &= z_{fix} - HX_{k_f}, \qquad S = HP_{k_f}H^\top + \sigma_R^2 I \\ d^2 &= \nu^\top S^{-1}\nu \le \gamma_g, \qquad K = P_{k_f}H^\top S^{-1} \\ X_{k_f} &\leftarrow X_{k_f} + K\nu \\ P_{k_f} &\leftarrow (I - KH)P_{k_f}(I - KH)^\top + \sigma_R^2 KK^\top \end{aligned} \tag{13.11}$$

in which H picks out the position and $\gamma_g = 9.21$ is the 99 % point of the chi-square law for two degrees of freedom. The filter then predicts again through every stored step since k_f with its stored inputs, so a late fix also corrects the bias and the current that those steps used. Of the three ways to fuse a delayed measurement set out in Section 1.4.7, which trade optimality against computation [5], this is the exact and costly one, affordable for five states. Volume 20, Chapter 10 takes a cheaper route for late camera reports and moves each one on by the belt travel since its capture. A fix with $d^2 > \gamma_g$ is refused. In the aided mode the Doppler log's own current, bottom track less water track rotated by $\hat{\psi}$ and averaged over ten samples, is applied once a second as an ungated measurement of $(\hat{V}_{cn}, \hat{V}_{ce})$ with a standard deviation of 0.02 m/s, and through the filter's correlations it also corrects the position that drifted while the current was unknown. The current is thus a slowly varying environmental state carried in the filter, as the bias state of the vessel observer in Volume 20, Chapter 6 carries wind and current. A GPS fix at the surface, in the GPS and aided modes, resets the position to the fix with a variance of $(1.5\,\mathrm{m})^2$, removes its correlations with the bias and the current, which keep what they learned, and clears the history, since no older fix can be applied after it. An advanced option applies each fix as if it described its arrival.

Table 13.3 Loops as Shipped, All at $\Delta t_c = 0.1\,\text{s}$

Loop	PV	MV	a	K_c	τ_I, τ_D, τ_{df}	output
HDG	$\hat{\psi}$	$\delta_{r,cmd}$	+1	1.2	60, 1.2, 0.2 s	±15°
SPD	u_m (lag 1 s)	T_{cmd}	+1	10	5 s, —, —	0–10 N
DEP	z_m	θ_{SP}	−1	8.6	40 s, —, —	±20°
PIT	θ_m	$\delta_{s,cmd}$	−1	3	∞, 1.0, 0.3 s	±15°

K_c in ° per °, N per m/s, ° of pitch per m, ° per °.

13.4.2 Line of Sight and the Current

The supervisor steers the heading toward a point a lookahead Δ ahead on the current leg, and it measures everything on the estimate, by Equation (13.12),

$$\begin{aligned} \psi_d &= \gamma - \arctan(\hat{e}_\perp/\Delta) - c_\beta\,\hat{\beta} \\ \hat{e}_\perp &= -(\hat{n} - n_A)\sin\gamma + (\hat{e} - e_A)\cos\gamma \end{aligned} \tag{13.12}$$

in which (n_A, e_A) is the start of the leg, γ its direction, $\Delta = 20\,\text{m}$ by default, and $\hat{e}_\perp$ is positive to starboard. The cross-track error is a geometric quantity, as in the lane keeping of Volume 20, Chapter 2, but here it is measured on the estimate. With crab compensation $c_\beta = 1$ and $\hat{\beta}$ is the drift angle of the velocity over the ground in the heading frame, $\text{atan2}(s_G, f_G)$ lagged by 2 s, taken from bottom track or from water track plus the current estimate and set to zero below $f_G = 0.2\,\text{m/s}$. Without compensation $c_\beta = 0$, and a steady cross current that needs a crab angle β, with $\sin\beta = V_\perp/U$, can only be met by sitting off the line until $\arctan(e_\perp/\Delta) = \beta$, a steady offset of $\Delta\tan\beta$ downstream. A leg ends when the estimate passes the perpendicular through its end.

13.4.3 The Four Loops

The four loops use the velocity-form PID of Section 1.6.2, Equations (1.15) to (1.18), at Δt_c and with no output bias. In each loop $a = \pm 1$ is the loop's direction and PV^f is the measurement through the exact first-order lag of time constant τ_{df}, so the derivative acts on the measurement. Back-calculation, with the tracking time equal to the integral time, runs only in the heading, speed and depth loops, which have an integral. The clamp acts on the emitted output, and no loop has a rate limit. In MANUAL each loop solves for $I_k = u_{op} - S_k$, which makes it emit exactly the operator's setting, while the pitch setpoint follows the measured pitch, so the return to AUTO is bumpless. The speed loop's integral starts at the trim thrust.

The pairings and tunings are listed in Table 13.3. The heading setpoint is unwrapped next to the estimated heading so that the error is taken the short way round. The speed loop acts on the water-track velocity lagged by 1 s. The depth loop acts in both vertical modes. At constant depth its setpoint is z_{SP}, and in terrain following it is the depth that puts the altitude setpoint over the seabed the altimeter sees, $z_{SP} = \overline{z_m + a_m} - a_{SP}$, the seabed depth lagged by 1 s and held when the altimeter has no reading. Its output, the pitch setpoint, is limited to $\pm 20°$, and because the pitch loop is proportional and derivative only, it holds about 17.7° against the righting moment when that setpoint is pinned at 20°.

13.4.4 The Mission Supervisor

The supervisor walks the plan on the estimated position. It dives from the launch point in depth mode toward the charted seabed at the start of line 1 less the altitude setpoint until the altimeter reads within 3 m of it, flies each line with the sonar logging, holds the depth the canyon leg began at across the canyon, and starts the final ascent where a climb at 17° reaches the surface at the recovery point. In the GPS mode it surfaces at the end of the first line to finish after the interval since the last fix, climbing in a 20 m loiter circle centered 25 m beyond the line's end, waits at the surface for the fix, descends to the altitude setpoint and rejoins the next line 30 m before its start. A battery at its

Table 13.4 Scenarios of the Survey Simulator and What Each Teaches

#	Scenario and settings	Teaching point
1	Manual, free flight at 32 m from line 1	a hull that will not fly straight
2	Lines 1 and 2 in still water	the loops, the turn, the ridge
3	0.4 m/s current toward 180°, lines 1 to 3	crab angle, LOS offset
4	DR, bias 1°, north block from launch	the pattern rotates
5	DR, 0.3 m/s toward 180°, line 5 to line 7	bottom lock lost
6	DR with GPS every 15 min, bias 1°	a reset and its cost
7	USBL, bias 1°, 50 m outlier at 400 s	delay, noise, gate
8	Lines 1 and 2 at 8 m altitude	the pitch limit
9	Full survey, DR, bias 1°, 0.25 m/s, 20 % battery	graded challenge

5 % reserve aborts the mission to the surface. On every line it also judges, from its own covariance, whether the overlap can still be trusted, as shown in Equation (13.13).

$$\begin{aligned} 1.96\,\sigma_{\perp} &> h(a_{plan}) - s/2 \\ \sigma_{\perp}^{2} &= \mathbf{g}^{\top} P_{pos}\,\mathbf{g}, \qquad \mathbf{g} = [-\sin\gamma,\ \cos\gamma]^{\top} \end{aligned} \tag{13.13}$$

where P_{pos} is the position block of the covariance and a_{plan} is the altitude setpoint, or the charted altitude at constant depth. At 10 m the right side is 5 m. When the test holds on a line inside a block, the supervisor records COVERAGE NOT ASSURED for it. The vehicle cannot see a holiday, but it can know when it may be leaving one.

13.4.5 Plant, Filter and Supervisor Models

Equations (13.1) to (13.8) are the simulation model, 14 states with the current's lag and the coverage grid beside them. The controller's internal models are the kinematics of Equation (13.9) with random-walk bias and current, the geometry of Equation (13.12), the drag law behind the propeller speed, the swath law and the charted seabed that the supervisor uses. None of them contains the hydrodynamics of Equation (13.1) or the Doppler scale error. No control model is identified. The tunings of Table 13.3 were set on the simulation model, and every result below is measured on it.

13.5 Measured Behavior of the Cases

Table 13.4 summarizes the eight cases and the graded survey. All of them begin from the configuration of Table 13.3, aided by the USBL with the gate on, in terrain following at 10 m and 1.5 m/s, and each row names what its case changes. The numbers below were measured headlessly at the default seed, stepping the engine as the browser does, and where one differs from the application's instructions guide the engine's value is the one printed. Each variant is set on the case as loaded, before Run, and the application gives the same run when the setting is carried through a Reset (Section 1.6.4). The application's own 41 checks, 24 of them specific to it, and its ten scenario-progress checks pass. Table 13.5 gathers the scorecard of every case.

13.5.1 A Hull That Will Not Fly Straight

Case 1 starts in MANUAL at 32 m depth, 28.8 m above the seabed, heading east at 1.5 m/s with the fins at zero and the thrust at its 3.645 N trim. Undisturbed, it flies straight for the 400 s tested, because nothing breaks the symmetry. Two degrees of rudder for 5 s build 12.1 °/s of yaw rate, and when the rudder is centered the turn does not stop. It settles at 9.85 °/s on a 7.1 m radius, with 5.0° of sideslip and the speed down to 1.22 m/s at the unchanged thrust, which is the unstable pole of Section 13.3 ending where crossflow drag balances the Munk moment. The stern plane behaves differently. A plane of $+5°$ holds the pitch at $-11.9°$ and the depth then grows at 0.31 m/s, and $+10°$ set at the start reaches the seabed 49 s later. Switched to AUTO, with setpoint tracking handing the autopilot the heading and depth the vehicle has, the heading loop stops the turn within 2.8 s, overshooting the captured heading by about 17°.

Table 13.5 Scorecard at the Default Seed. Navigation Error Is the Hidden Truth on the Survey Lines, and a Survey Ends When the Vehicle Surfaces

Case	Imaged (%)	Holidays, widest (m)	Lowest alt. (m)	Nav. error RMS, max (m)	End (s)	Miss (m)
2	22.01	0	6.41	0.56, 1.74	533	—
3	33.11	2, 6	5.95	0.52, 1.73	768	—
4	49.41	2, 7	5.44	6.25, 8.77	1329	—
5	25.69	1, 25	5.55	20.23, 21.56	698	—
6 on DR only	98.80	4, 7	5.44	7.75, 11.35	2574	9.7
6, GPS surfacing	99.61	3, 5	5.34	5.40, 8.75	3075	8.8
6 on USBL	99.96	0	5.58	0.57, 3.65	2574	0.7
7	99.95	0	5.61	0.56, 3.65	2575	0.9
7, gate off	99.94	0	5.61	1.51, 7.05	2577	0.5
7, fixes at arrival	99.90	1, 3	5.65	1.92, 3.65	2584	1.5
8	20.47	2, 5	4.40	0.54, 1.74	536	—
8, stronger depth loop	20.72	0	5.12	0.54, 1.74	536	—
8 at 1.0 m/s[a]	20.41	2, 9	3.77	0.47, 1.19	830	—

[a] Set before Run, so the case starts trimmed at 1.0 m/s.

13.5.2 The Overlap Budget: Autopilots and the Ridge

Case 2 flies lines 1 and 2 in still water on the USBL. The true cross-track error inside the block is 0.49 m RMS on line 1 and 0.46 m on line 2, the navigation error 0.56 m RMS, and 25 fixes are used while one honest fix is refused by the gate. The turn onto line 2 swings 6.0 m off the line inside its 20 m run-in and is 1.8 m off at the block's edge. Over the ridge the depth loop can climb only at its pitch limit, so the vehicle falls to 6.41 m above the seabed on line 2 and its swath narrows to 25.6 m. The neighboring line crosses the same place on the far side of its own climb, higher, so the two still overlap, but by 3 m at the crest instead of 10 m, and no holiday opens. The ridge alone spends 7 m of the overlap before any navigation error. The lookahead barely matters in still water on this agile hull, 0.62 m RMS on line 1 at 5 m and 0.40 m at 40 m.

Case 8 flies the same lines at 8 m, where the swath is 32 m and each side has only 1 m to spare, so the supervisor records COVERAGE NOT ASSURED on both lines before anything has gone wrong. The face that line 2 climbs rises at up to 23.1° along the planned line, while the pitch loop, its setpoint pinned at 20° for 21 s of that climb, holds a median of 17.7° with the planes near −6.7°. The seabed gains on the vehicle, which falls to 4.40 m above it near the crest with a swath of 17.6 m, and two holidays open over the ridge, the widest 5 m. A stronger depth loop, 14.3° of pitch per meter with a 30 s integral time, starts the climb sooner and holds 5.12 m with no holiday, and flying at 10 m holds 6.41 m. Slowing to 1.0 m/s makes it worse, 3.77 m and a 9 m holiday, because the planes run near −13.5° of their 15° and hold a median of only 15.5° of pitch, while the climb is limited by its angle and not by its speed. No gain can climb faster than the pitch limit allows. What the loop decides is how early the climb starts.

13.5.3 A Current Across the Lines

In case 3 a 0.4 m/s current flows south across the east–west lines. The vehicle crabs 15.5° into it on every line, the arcsine of 0.4 over 1.5, and its navigation stays at 0.52 m RMS because bottom track measures the velocity over the ground. Pure line of sight commands the heading, so it leads the line by the crab angle only by sitting downstream of it. The mean offsets inside the block are 5.4, 5.0 and 5.1 m on the three lines at $\Delta = 20\,\text{m}$, against $20\tan 15.5^\circ = 5.5\,\text{m}$, always to the south whichever way the line is flown. The whole pattern slides south together, so the gap opens at the block's upstream edge, two holidays along the north edge, the widest 6 m. A 10 m lookahead halves the offsets to 2.8, 2.5 and 2.6 m with no holiday, and 40 m doubles them to about 10 m and opens a 12 m holiday. Crab compensation, steering the course over the ground that the Doppler log measures, brings them to 1.6, 0.4 and 0.2 m with no holiday, the last two slightly upstream. Line 1 keeps 1.6 m because the setting is in force when the heading loop is primed, so the loop's integral starts by cancelling the crab

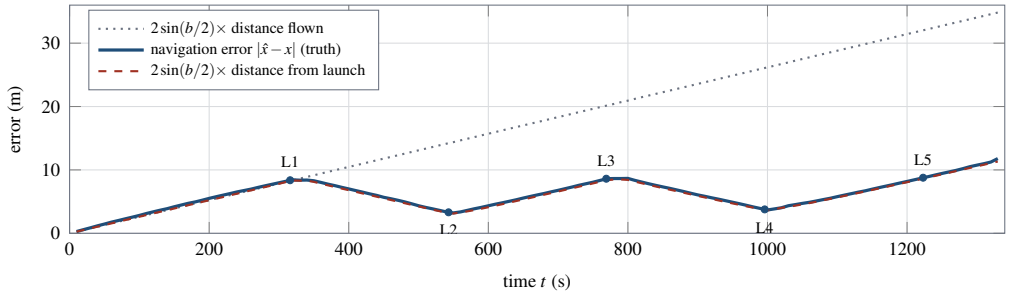

Figure 13.3 Dead-reckoning error in case 4 with a 1° compass bias. The engine's error (thick solid) follows $2\sin(b/2)$ times the distance from the launch point (dashed) and not the distance flown (dotted), and the dots mark the ends of lines 1 to 5.

command and unwinds it over about a minute. Crab compensation is the same remedy that integral and sideslip-compensating line-of-sight laws provide in theory [7, 8].

> **TRY IT IN THE APP** **3 · A Cross Current**
>
> Load the scenario, press Run and watch the trend across the line. The hidden truth sits about 5 m south of every line while the estimate holds it, and two holidays open along the north edge, the widest 6 m. Then set Crab compensation to On, press Reset and run again. The offsets fall to 1.6, 0.4 and 0.2 m with no holiday.

13.5.4 A Compass Bias Rotates the Survey

Case 4 dives from the launch point and flies the north block on dead reckoning with the compass 1° high. Equation (13.10) predicts the error from the distance to the launch point, and the engine follows it to within 0.32 m until the vehicle surfaces at the end, 0.18 m on average, as shown in Figure 13.3. The error is 8.37 m at the east end of line 1, 471 m from the launch point (1.78 %), 3.31 m at the west end of line 2, 182 m away after 814 m flown (1.82 %), 8.61 m at the east end of line 3 and 3.77 m at the west end of line 4 after 1495 m flown. The naive 1.7 % of the distance flown would have given 26 m there, seven times the error measured. A lawnmower pattern is known to cancel low-frequency Doppler errors that drift linearly on a straight line [15], and a constant heading bias cancels the same way here, each line undoing much of the one before.

Pause. Try Exercise 13.1 before reading on.

A rotation keeps the lines parallel and evenly spaced, so no gap opens between them. What it moves is the block's edge. The lines sit about 6 m north of their plan in the middle of the block and 8 m at its east end, and the south edge is left unimaged in two holidays, 7 m wide where the ridge also narrows line 5's swath and 5 m wide at the east end. The supervisor records COVERAGE NOT ASSURED 129 s into the dive, 193 m from the launch point, when its own 95 % uncertainty across the line passes the 5 m of spare overlap, and it records the same flag at the same instant with no bias at all, when dead reckoning stays within 0.24 m RMS, because the filter allows 2° of bias that nothing aboard can rule out. The flag reports what can be vouched for rather than what went wrong. At 2° of bias the error is 12.3 m RMS with a 12 m holiday, and on the USBL it is 0.73 m with none.

13.5.5 The Canyon: A Drift Nothing Aboard Can See

Case 5 finishes line 5 on dead reckoning and crosses the canyon in a 0.3 m/s current flowing south. The supervisor holds 49.8 m across the canyon, the altitude reaches 105.7 m, beyond the altimeter's 50 m as well as the Doppler log's 30 m, and bottom lock is lost for 71.8 s. On water track the current is not in the reckoning, and the error grows from 0.05 to 21.46 m, at 0.298 m/s, and Figure 13.4 shows the run on line 7 at 420 s. Every line of the southern block is then flown about 21 m south of

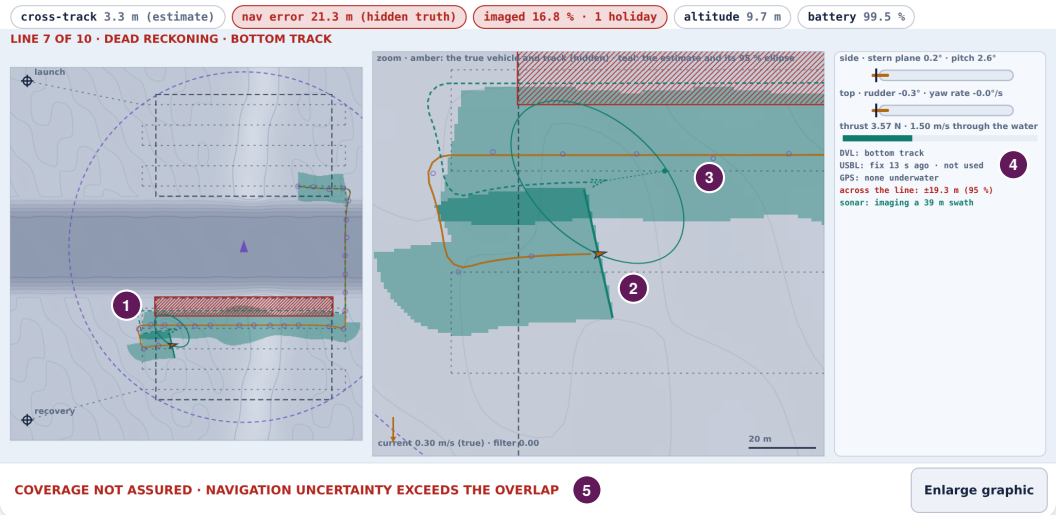

Figure 13.4 Case 5, The canyon, at $t = 7.00$ min (420 s) with no control changed, 34 s after the vehicle turned onto line 7 on dead reckoning, cropped to the Chart view (the application header, the trends, the control panel and the Learn panel are cut). The header line reads line 7 of 10, dead reckoning, bottom track. The survey chart shows the launch and recovery points, the two dashed blocks with their dotted lines, the canyon band with the ship over it inside the dashed USBL circle at the vehicle's depth, the seabed imaged at the east end of line 5, along line 6 and at the start of line 7, and a holiday (1) hatched along the whole north edge of the southern block. The zoom, centered on the true vehicle (2, amber, the hidden truth), shows it south of line 7 while the estimate (3, teal) lies near the line inside its 95 % ellipse, with the line-of-sight point ahead, the 39 m swath across the true vehicle and the unused USBL fixes as open circles. The badges give the cross-track error of the estimate (3.3 m), a navigation error of 21.3 m (hidden truth), 16.8 % imaged with one holiday, an altitude of 9.7 m and 99.5 % of the battery, the inset (4) reports bottom track, a USBL fix 13 s old and not used, and ±19.3 m across the line at 95 %, and the status line (5) reads COVERAGE NOT ASSURED.

where the vehicle believes it is, and the block's north edge is left unimaged in one holiday 25 m wide and 260 m long. The filter's own uncertainty grows honestly across the crossing, and the supervisor records COVERAGE NOT ASSURED on lines 6 and 7.

> **TRY IT IN THE APP** **5 · The Canyon**
>
> Load the scenario, press Run and watch the navigation error trend as the vehicle crosses the canyon. Bottom lock is lost for 71.8 s, and the error climbs to about 21 m while the estimated cross-track error stays near zero. A holiday 25 m wide then opens along the north edge of the southern block.

This is the case that no tuning can recover, and the simulator says so. A lookahead of 10 m leaves a 23 m holiday, 40 m leaves 28 m, and crab compensation leaves 21 m, because without a fix or a current measurement the drift is unobservable (Section 1.4.6), and the application's own checks confirm that every one of these variants is flagged. In still water the same crossing costs nothing, and on the USBL the navigation holds 0.41 m RMS through it, leaving only the current's line-of-sight offset as one 4 m holiday. The remedy is a measurement, a fix from outside or the current measured while the seabed was still in view, the combination that model-aided navigation with sea current estimation was built for [14].

Pause. Try Exercise 13.2 before reading on.

13.5.6 Buying a Fix: Surfacing Against Acoustic Aiding

Case 6 flies the whole survey on dead reckoning with the 1° bias and surfaces for GPS at the end of the first line after each 15 min. The first surfacing follows line 4, at 996 s. Line 7 ends 900.15 s after the last fix of that surfacing, so the second follows line 7, at 2027 s, and a line 7 that ended a fraction of a second sooner would have sent the vehicle on to line 8. Each surfacing climbs from about 50 m in 105 and 107 s in its loiter circle, waits 20 s for the fix, descends and rejoins the next line, which takes 256 s and 3.45 Wh from the end of one line to the start of the next the first time and 257 s and 3.42 Wh the second. The fixes leave the estimate 1.65 and 0.76 m from the truth, against 3.77 and 8.45 m at the ends of lines 4 and 7, but the bias survives them and the pattern goes on rotating, now about each surfacing point. A fix taken between two neighboring lines also changes their spacing. Over the ridge, line 7 runs 5.0 m north of its plan and line 8, flown after the fix, 2.3 m south of its own, and a 5 m holiday opens between them. On the same seed, dead reckoning alone images 98.80 % with four holidays and surfaces 9.7 m from the recovery point at 2574 s on 33.67 Wh, the surfacings image 99.61 % with three holidays, an 8.8 m miss, 3075 s and 40.33 Wh, and USBL aiding images 99.96 % with none, a 0.7 m miss, at 2574 s. A surfacing bounds the error at the fix and corrects nothing that caused it, and in a 0.2 m/s current the two surfacings do not reach the canyon's drift, which leaves a 27 m holiday and a 29.6 m miss.

Case 7 flies the same survey aided by the USBL. The fixes arrive 1.2 to 1.4 s after the instant they describe, with standard deviations of 0.48 to 1.32 m, and 38 of the 172 interrogations bring nothing back. Of those, 14 are lost pings and 24 find the vehicle outside the 80° cone, during the dive, during the final ascent, and on the lines farthest from the ship where the vehicle rises over the ridge. Fixes taken on lines flown east and west see the bias from opposite sides, which makes it observable in the sense of Section 1.4.6, and its estimate reaches 1.005° by 300 s. With the bias removed from the heading the survey images 99.95 % with no holiday, a navigation error of 0.56 m RMS, and surfaces 0.9 m from the recovery point. The multipath fix that arrives at 407 s lies 50 m off at a chi-square distance of 2382 against 9.21 and is refused, and with the gate off it is accepted and the error jumps to 7.05 m before the next fixes pull it back. Applying every fix as if it described its arrival compares it with a position about 1.3 s of travel, some 2 m, further along the line, the gate then refuses 47 honest fixes besides the outlier, and the error RMS more than triples, to 1.92 m. The supervisor records COVERAGE NOT ASSURED once even here, on line 1 at 129 s, because the vehicle is outside the transceiver's cone until 116.5 s and the first fix arrives only at 137 s.

> **+ DID YOU KNOW? About a Second Late, Three Times the Error**
>
> From the simulator: in case 7 every acoustic fix describes an instant 1.2 to 1.4 s before it arrives. Applied at that instant, the fixes hold the navigation error to 0.56 m RMS with no holiday. Applied as if they described their arrival, the same fixes give 1.92 m and a 3 m holiday (Table 13.5).

13.5.7 The Graded Survey

The challenge flies the whole survey launched with 20 % of the battery, in a 0.25 m/s current toward 190°, with the compass 1° high and 3 % of the USBL fixes wild. The target is at least 95 % imaged, no holiday wider than 2 m, never below 3 m above the seabed, surfacing within 20 m of the recovery point and at least 15 % of the battery left, judged when the vehicle surfaces, with the model parameters and disturbances as given. A changed parameter or disturbance voids the attempt. It ships on dead reckoning with a 40 m lookahead, no crab compensation and a soft vertical pair, 3.4° of pitch per meter with a 60 s integral time and a pitch gain of 1.5. Table 13.6 compares the configurations, and Figure 13.5 draws three of its columns against their targets.

As shipped the survey images 83.86 % with nine holidays, the widest 42 m, comes within 1.72 m of the seabed over the ridge and surfaces at 2559 s, 61.3 m from the recovery point. It fails four of the five criteria and meets only the battery, with 16.67 % left. The suggested configuration images 99.95 % with no holiday, never comes closer than 5.50 m, and surfaces 1.9 m from the recovery point

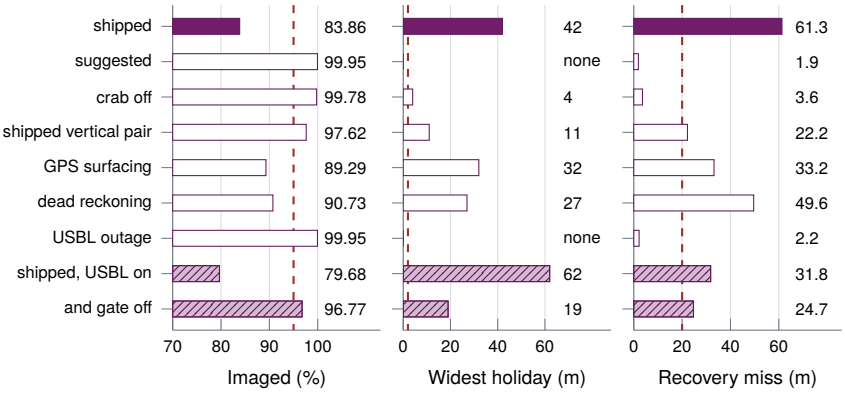

Figure 13.5 The graded survey of Table 13.6: the seabed imaged (left, dashed line at the 95 % target), the widest holiday (middle, dashed line at 2 m) and the miss at the recovery point (right, dashed line at 20 m). The dark bar is the shipped configuration, white bars the suggested one and its single changes, and hatched bars the shipped configuration with the USBL switched on. Notice that switching on the USBL alone images less than dead reckoning did, while the suggested configuration and its USBL outage are the only rows inside all three lines.

Table 13.6 Graded Survey at the Default Seed, With Single Changes From the Shipped and the Suggested Configurations

Configuration	Imaged (%)	Holidays, widest (m)	Lowest (m)	Miss (m)	Battery (%)	Fixes[b]
shipped	83.86	9, 42	1.72	61.3	16.67	—
suggested[a]	99.95	0	5.50	1.9	16.62	125, 7
crab off	99.78	1, 4	5.23	3.6	16.66	123, 3
shipped vertical pair	97.62	12, 11	2.22	22.2	16.61	130, 4
GPS surfacing	89.29	3, 32	5.08	33.2	15.95	—
dead reckoning	90.73	4, 27	4.77	49.6	16.61	—
USBL outage[c]	99.95	0	5.50	2.2	16.62	105, 8
shipped, USBL on	79.68	11, 62	1.11	31.8	16.66	1, 132
and gate off	96.77	13, 19	1.92	24.7	16.68	125, 0

[a] USBL with the gate, $\Delta = 20$ m with crab compensation, 8.6°/m and 40 s, pitch gain 3. [b] Used, refused. [c] From 1150 to 1450 s.

at 2577 s with 16.62 % left. It is inside every limit, the battery by the smallest margin, 1.62 percentage points, and it passes on each of twelve further seeds as well. Each part has an owner. Without crab compensation one 4 m holiday remains, with the shipped vertical pair the vehicle comes to 2.22 m above the ridge and twelve holidays open, GPS surfacing in place of the USBL leaves a 32 m holiday and takes 3078 s, and on dead reckoning the canyon's drift leaves a 27 m holiday whatever else is set, with COVERAGE NOT ASSURED on every line, which is the infeasible variant. The aided filter survives a USBL outage across the whole canyon leg, from 1150 to 1450 s, with no holiday, because it measured the current while it still had bottom lock and had learned the bias from the fixes before.

The most natural first change fails in a way the other cases do not teach. Switching only the navigation of the shipped configuration to the USBL images 79.68 %, less than dead reckoning did, with a 62 m holiday. The first fix to arrive, at 137 s, is one of the wild ones, 24.2 m from the truth, while the filter's own standard deviation across the line has grown to 7.17 m, most of it from the 2° of bias it allows over the dive, so the fix's chi-square distance is 8.04 and the gate accepts it. The filter explains most of it as compass bias, $\hat{b} = -5.41°$ against a true 1°, and its position standard deviation falls to 1.2 m. Every one of the next 132 fixes, 129 of them honest, then lies at a distance of 21 to 479 from a confident and wrong prediction and is refused, and the navigation error reaches 73.4 m.

With the gate off the same configuration uses 125 fixes and images 96.77 %. On twelve other seeds the same change refuses no more than seven fixes and images 96.55 to 96.75 %, still failing for the guidance and the vertical loops. A validation gate is only as tight as the covariance behind it [6], and the first fix of a mission meets it at its widest.

Pause. Try Exercise 13.3 before reading on.

13.6 Technology in Practice

What Measures It

Depth, heading, pitch and roll each have a direct sensor, but no equivalent sensor exists for the two horizontal coordinates [3], so the position is reckoned from a Doppler velocity log and a heading reference and corrected by acoustic fixes, and no instrument aboard measures it. A Doppler log sends four acoustic beams, taken in one analysis as a quarter turn apart around the vertical with the first usually 45° off the forward direction, and turns the Doppler shifts of their echoes into the velocity in the vehicle's frame. Its scale is proportional to the sound speed at the transducer [15]. Bottom-lock logs measure the velocity over the seabed to 0.3 % or better at up to 5 Hz, from altitudes of about 18 to 100 m that vary with the log [3], and a higher frequency buys accuracy at the cost of range [15]. A calibrated magnetic compass gives the heading to the order of 1 to 3° [3], the range of the chapter's compass bias. The chapter's log locks on the bottom within 30 m and falls back to water track above it, which is where the canyon's drift begins.

TECHNOLOGY **Sensor: The Ship's Ultra-Short Baseline Fix**

An ultra-short baseline system needs only a transducer on the hull of the ship. A sonar array in it measures the range and the bearing to the vehicle, and the system needs an alignment calibration with the ship's own positioning system, typically GPS. Every acoustic time-of-flight method needs an accurate sound velocity and is limited by the speed of sound in water, about 1500 m/s, and vehicles deep below the ship are especially challenged, because the sound velocity varies with the temperature and density of the water [3]. The chapter's fix has a noise of 0.5 % of the slant range, reaches the vehicle after the two-way travel and 1.0 s of processing, and fails as a lost ping, outside the 80° cone or as an outlier that the gate must refuse.

What Moves It

The vehicle has three final elements. The rudder and the stern plane each sit behind a servo modeled as a 0.1 s lag and limited to 15°, and the thruster sits behind a 0.4 s motor lag and is limited to 0 to 10 N. The sources opened for this section give no fin or thruster data for a real survey vehicle, so these limits are the model's own.

IN PRACTICE **How a Survey Is Judged on Its Gaps**

A United States hydrographic specification defines a holiday as a gap in gridded data larger than the features the survey claims to detect. On its own surveys it allows no data gap of 3 by 3 grid nodes or more, 3 m by 3 m on a 1 m grid, and the cause of any such gap must be explained in the report of survey. A towed side scan sonar must fly at 8 to 20 % of its range scale above the bottom, and when conditions shorten its effective range the coverage claimed should shrink with it [18]. The international standard asks for 100 % bathymetric coverage in its Special Order and 200 % in its Exclusive Order [19].

Where the Controller Runs

One computer on the vehicle runs the navigation filter, the mission supervisor and the four autopilot loops. The sources opened give no scan period for the vehicle's computer. The chapter's controller reads the instruments, runs the filter and the supervisor and executes the four loops every 0.1 s, and the ship interrogates the transponder every 15 s. Figure 13.6 draws the loops on that hardware.

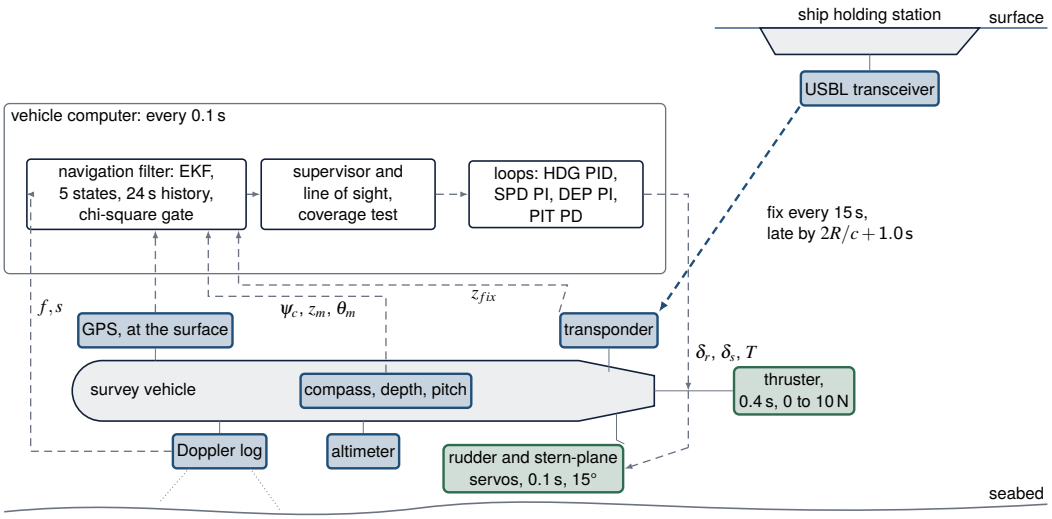

Figure 13.6 The survey loops on vehicle hardware. The Doppler log, the compass and the depth and pitch sensors feed the navigation filter, which also takes GPS at the surface and the ship's USBL fixes. The altimeter feeds the depth loop in terrain following (line not drawn). One computer runs the filter, the supervisor and the four loops every 0.1 s, and the servos and the thruster move the rudder, the stern plane and the propeller. Notice that each USBL fix reaches the vehicle 1.2 to 1.4 s after the instant it describes, and that no line carries the true position: the coverage the survey exists for is never measured aboard.

> **TECHNOLOGY** **Platform: The Navigation Filter and Its Clock**
>
> On a survey vehicle the inertial navigation system computes position, velocity and attitude from high-rate inertial data, and an extended Kalman filter on an error-state model corrects it with the Doppler log, the pressure sensor, the compass and any position fixes. Every measurement must carry the time it describes. When the filter compares a Doppler velocity with its own, a time-stamp error times the vehicle's acceleration becomes a velocity error, and an analysis of the dynamics during surface fixes asked for Doppler time stamps better than 10 ms, preferably with the log reporting its own latency in each message [15]. The chapter's filter keeps 24 s of history so that each fix is applied at the instant it describes.

What Is Modernizing It

The first development is the heading. Heading accuracy is a principal error source of Doppler navigation, and low-cost north-seeking optical gyrocompasses, with a dynamic heading accuracy of the order of 0.1°, effectively ameliorate it [3], and by Equation (13.10) a smaller bias rotates the pattern less. The second is ranging. Acoustic modems with synchronized clocks allow ranging by one-way travel time, in which a ship broadcasts its GPS position and every vehicle in range computes a running fix, while a long-baseline net of seafloor transponders typically reaches up to 10 km with a precision of 0.1 to 10 m [3]. The third is the seabed itself. Terrain-relative navigation and simultaneous localization and mapping correct the position from the vehicle's own bathymetric sonar or cameras against a map, given in advance or built as the vehicle goes, within a sensor range of the order of 10 to 100 m for a sonar [3]. Over the canyon the vehicle flies 105.7 m above the floor.

13.7 Limitations

The vehicle is Prestero's REMUS set with constant coefficients, no heave and no roll, so it teaches the structure of an underactuated, yaw-unstable hull rather than the handling of a particular vehicle,

and its turns do not bank the swath. The current is uniform and unsheared, so water track measures exactly the velocity through the water and the current it reveals is the current at the vehicle.

The seabed is smooth and known exactly to the supervisor, the swath is taken from the altitude at nadir over a locally flat seabed, and the sonar's image quality, nadir gap and resolution are not modeled. The altitude loop sees only the altimeter straight below, with no preview of the seabed ahead. The instruments are generic, the compass bias is constant, and the Doppler scale error is outside the filter.

> **OPEN CHALLENGE** **Navigating Far From Seabed and Surface**
>
> GPS has solved navigation at the surface, and Doppler sonars and landmark navigation have greatly improved it near the seabed. In the mid-depth zone, far from both, only acoustic time-of-flight systems and inertial units measure the horizontal position, which limits how precisely a vehicle can navigate there. A survey of the field expected the advances to come from model-based state estimators, helped by vehicles working near the surface or the seabed at the same time [3]. The canyon of case 5 is a small mid-depth zone, where the chapter's filter loses the current. Progress would be an estimate of the current or the position that holds between the walls.

The filter's covariance is honest only for the errors it models, and nothing in it detects a run of refused fixes or recovers from one, so the lockout of Section 13.5.7 persists to the end of the survey. At the surface the filter reckons on the propeller speed through the steady drag law, which undercounts while the hull coasts through a change of speed. Every threshold the supervisor and the scorecard use is a teaching value.

> **OPEN CHALLENGE** **Navigation and the Survey as One Problem**
>
> Ocean engineers have traditionally worked on determining the position in the water and flown trajectories defined in advance, which suits bathymetry. A survey of the field described studies that analyze vehicle navigation, trajectory planning and the data of the science sensors together, to learn how each design choice affects the accuracy of the mapped process, and noted that observability, uniqueness and lower bounds for such systems were only beginning to be studied [3]. This chapter scores navigation and guidance on the seabed imaged, but on a fixed plan. Progress would be a vehicle that replans its lines from its own covariance to cover the strips it may have missed.

13.8 Conclusion

This chapter has treated an AUV survey as a control problem whose output is never measured. The simulator keeps the true vehicle, its instruments and its filter apart so that the difference can be scored on the seabed. Prestero's coefficients leave the hull unstable in yaw, with a pole at $0.995\ \mathrm{s}^{-1}$, so the heading loop is part of the vehicle. The ridge alone spends the 10 m of overlap down to 3 m, and a 0.4 m/s current holds pure line of sight 5.4 m downstream until crab compensation removes the offset. A 1° compass bias rotates the dead-reckoned survey about the launch point, 1.78 % of the distance from the fix rather than of the distance flown, and over the canyon a 0.3 m/s current carries the vehicle 21.46 m from its estimate and leaves a 25 m holiday that no tuning recovers. GPS surfacing costs about 256 s and 3.4 Wh a time and leaves the bias in place, while USBL aiding learns it and images 99.95 %. The graded survey rises from 83.86 % imaged with nine holidays to 99.95 % with none on thirteen seeds of thirteen, although switching the shipped navigation to the USBL alone locks the gate out after one wild fix. The model omits heave, roll, current shear and the sonar's image. The natural additions are a lockout monitor in the filter, a forward-looking altimeter, a sheared current and a long-baseline net. The survey simulator is a laboratory for the difference between where a vehicle was and where it believed it was.

KEY POINTS

- A survey vehicle steers its own estimate of where it is, while the seabed its sonar images is set by where it actually went.
- A 1° compass bias rotates the dead-reckoned survey about the launch point, 1.78 % of the distance from the fix rather than of the distance flown.
- Over the canyon a 0.3 m/s current carries the vehicle 21.46 m from its estimate and leaves a 25 m holiday that no tuning recovers.
- USBL aiding learns the bias and images 99.95 %, but one wild first fix can lock the gate out.

Where else this applies. Volume 20, Chapter 6 carries wind and current as a slowly varying bias state of its observer, and Volume 20, Chapter 10 moves each late camera report on by the belt travel since its capture.

References

[1] R. B. Wynn, V. A. I. Huvenne, T. P. Le Bas, B. J. Murton, D. P. Connelly, B. J. Bett, H. A. Ruhl, K. J. Morris, J. Peakall, D. R. Parsons, E. J. Sumner, S. E. Darby, R. M. Dorrell, and J. E. Hunt, "Autonomous underwater vehicles (AUVs): Their past, present and future contributions to the advancement of marine geoscience," *Marine Geology*, vol. 352, pp. 451–468, 2014.

[2] L. Paull, S. Saeedi, M. Seto, and H. Li, "AUV navigation and localization: A review," *IEEE Journal of Oceanic Engineering*, vol. 39, no. 1, pp. 131–149, 2014.

[3] J. C. Kinsey, R. M. Eustice, and L. L. Whitcomb, "A survey of underwater vehicle navigation: Recent advances and new challenges," https://robots.engin.umich.edu/publications/jkinsey-2006a.pdf, 2006.

[4] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1258–1276, 2013.

[5] T. D. Larsen, N. A. Andersen, O. Ravn, and N. K. Poulsen, "Incorporation of time delayed measurements in a discrete-time Kalman filter," in *Proceedings of the 37th IEEE Conference on Decision and Control (Cat. No.98CH36171)*, vol. 4. IEEE, 1998, pp. 3972–3977.

[6] Y. Bar-Shalom, X.-R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. Wiley, 2002.

[7] E. Borhaug, A. Pavlov, and K. Y. Pettersen, "Integral LOS control for path following of underactuated marine surface vessels in the presence of constant ocean currents," in *2008 47th IEEE Conference on Decision and Control*. IEEE, 2008, pp. 4984–4991.

[8] T. I. Fossen, K. Y. Pettersen, and R. Galeazzi, "Line-of-sight path following for Dubins paths with adaptive sideslip compensation of drift forces," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 2, pp. 820–827, 2015.

[9] R. Spindel, R. Porter, W. Marquet, and J. Durham, "A high-resolution pulse-Doppler underwater acoustic navigation system," *IEEE Journal of Oceanic Engineering*, vol. 1, no. 1, pp. 6–13, 1976.

[10] K. Vickery, "Acoustic positioning systems. A practical overview of current systems," in *Proceedings of the 1998 Workshop on Autonomous Underwater Vehicles (Cat. No.98CH36290)*. IEEE, 1998, pp. 5–17.

[11] L. Whitcomb, D. Yoerger, and H. Singh, "Advances in Doppler-based navigation of underwater robotic vehicles," in *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, vol. 1. IEEE, 1999, pp. 399–406.

[12] B. Allen, R. Stokey, T. Austin, N. Forrester, R. Goldsborough, M. Purcell, and C. von Alt, "REMUS: a small, low cost AUV; system description, field trials and performance results," in *Oceans '97. MTS/IEEE Conference Proceedings*, vol. 2. IEEE, 1997, pp. 994–1000.

[13] T. Prestero, "Development of a six-degree of freedom simulation model for the REMUS autonomous underwater vehicle," in *MTS/IEEE Oceans 2001. An Ocean Odyssey. Conference Proceedings (IEEE Cat. No.01CH37295)*, vol. 1. IEEE, 2001, pp. 450–455.

[14] Ø. Hegrenaes and O. Hallingstad, "Model-aided INS with sea current estimation for robust underwater navigation," *IEEE Journal of Oceanic Engineering*, vol. 36, no. 2, pp. 316–337, 2011.

[15] B. Jalving, K. Gade, K. Svartveit, A. Willumsen, and R. Sørhagen, "DVL velocity aiding in the HUGIN 1000 integrated inertial navigation system," *Modeling, Identification and Control: A Norwegian Research Bulletin*, vol. 25, no. 4, pp. 223–236, 2004.

[16] A. J. Healey and D. Lienard, "Multivariable sliding mode control for autonomous diving and steering of unmanned underwater vehicles," *IEEE Journal of Oceanic Engineering*, vol. 18, no. 3, pp. 327–339, 1993.

[17] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*. Wiley, 2011.

[18] NOAA Office of Coast Survey, *Hydrographic Survey Specifications and Deliverables, Version 2026.0.03*, National Oceanic and Atmospheric Administration, 2026. [Online]. Available: https://nauticalcharts.noaa.gov/publications/documents/HSSD_2026-0-03.pdf

[19] International Hydrographic Organization, *IHO Standards for Hydrographic Surveys, S-44 Edition 6.1.0*, International Hydrographic Organization, Monaco, 2022. [Online]. Available: https://iho.int/uploads/user/pubs/standards/s-44/S-44_Edition_6.1.0.pdf

13.9 Exercises

13.1 Concept check · 2 min

In case 4 the vehicle dead reckons the north block with its compass 1° high. After 1495 m flown, at the west end of line 4, the navigation error is only 3.77 m, not the 26 m that 1.745 % of the distance flown would give. Why?

(a) The bias rotates the estimated track about the last fix, so the error grows with the distance from the fix, not the distance flown

(b) The filter learns the bias from bottom track on each line, so the error it leaves shrinks from one line to the next

(c) The Doppler log's scale error reverses with the direction of each line, so its share of the error cancels in pairs

(d) The supervisor's COVERAGE NOT ASSURED flag resets the estimate at the end of each line to the charted position

13.2 Concept check · 1 min

In case 5 the vehicle crosses the canyon on dead reckoning in a 0.3 m/s current and then flies lines 6 and 7 about 21 m south of where it believes it is. Why does neither a 10 m lookahead nor crab compensation recover the 25 m holiday?

(a) The lookahead and the crab angle act only once the vehicle reaches the south block, too late to correct line 6

(b) The guidance steers an estimate that is itself 21 m wrong, and nothing aboard can observe that drift without a fix

(c) The depth loop holds the vehicle too high over the canyon walls for the sonar to image the block's north edge

(d) The heading loop saturates its rudder in the turn onto line 6 and loses the line for the rest of the case

13.3 Concept check · 2 min

In the graded survey, switching only the shipped navigation to the USBL images 79.68 %, less than dead reckoning. The Fixes tile ends at 1 used and 132 refused. What went wrong?

(a) Each fix arrives about 1.3 s late, so the gate compares it with a position 2 m behind the vehicle and refuses it

(b) The gate's threshold of 9.21 is too tight for honest fixes at the long ranges of the south block

(c) A wild first fix passes the gate while the covariance is wide, and the confident, wrong filter then refuses the honest fixes

(d) The vehicle stays outside the transceiver's 80° cone for most of the survey, so almost no fix returns

13.4 Theory and calculation · 6 min

Use Equations (13.6) and (13.13) with $\tan\theta_s = 2$, $R_s = 50\,\text{m}$ and lines $s = 30\,\text{m}$ apart.

(a) Compute the swath, the overlap between neighboring lines and the spare overlap $h - s/2$ on each side at altitudes of 10 and 8 m. *Checkpoint:* 40 m, 10 m and 5 m at 10 m. 32 m, 2 m and 1 m at 8 m.

(b) Find the altitude at which the swath is widest, and its width. *Checkpoint:* 22.4 m and 89.4 m.

(c) Find the standard deviation across the line $\sigma_\perp$ above which the supervisor records COVERAGE NOT ASSURED at each altitude of part (a), and say why case 8 is flagged before anything has gone wrong. *Checkpoint:* About 2.55 m at 10 m and 0.51 m at 8 m.

13.5 Theory and calculation · 7 min

Use Equation (13.10) and the case 4 and case 5 results of Section 13.5.

(a) Compute $2\sin(b/2)$ for a bias of 1°, and predict the error at the east end of line 1, 471 m from the launch point, and at the west end of line 2, 182 m from it. Compare with the measured 8.37 and 3.31 m. *Checkpoint:* 1.745 %, about 8.2 m and 3.2 m.

(b) Compute the error that 1.745 % of the 1495 m flown to the end of line 4 would predict, and its ratio to the 3.77 m measured. *Checkpoint:* About 26.1 m, 6.9 times the measured error.

(c) In case 5 bottom lock is lost for 71.8 s in a 0.3 m/s current and the filter's current estimate is zero. Predict the error added over the crossing and compare it with the measured growth. *Checkpoint:* About 21.5 m, against 0.05 to 21.46 m.

13.6 Theory and calculation · 6 min

Line-of-sight guidance without crab compensation meets a steady cross current $V_\perp$ only by sitting off the line (Equation (13.12)). The vehicle moves at $U = 1.5\,\text{m/s}$ through the water.

(a) For the 0.4 m/s current of case 3, compute the crab angle and the steady offset at lookaheads of 10, 20 and 40 m. Compare with the measured mean offsets of 2.8, 5.4 and about 10 m on line 1. *Checkpoint:* 15.5°, and 2.8, 5.5 and 11.1 m.

(b) The graded survey runs in 0.25 m/s toward 190° with a 40 m lookahead as shipped. Find the current across the east–west lines, the crab angle and the steady offset, and compare it with the 5 m of spare overlap at 10 m. *Checkpoint:* About 0.246 m/s, 9.4° and 6.7 m, more than the spare.

(c) Explain why in case 3 the gap opens at the north edge of the block, whichever way a line is flown. *Checkpoint:* The offset is always downstream, south, so the whole pattern slides south together.

13.7 App challenge · 8 min

Load scenario 4, Dead reckoning with a compass bias, with the default seed. Run at 20× until the vehicle surfaces and the run pauses, at about 1331 s on the clock, and read the Seabed imaged, Holidays and Navigation error tiles. Reload the scenario before each new configuration and make the change before pressing Run.

(a) Run as loaded. Report the three tiles. *Checkpoint:* 49.4 %, 2 holidays with the widest 7.0 m, and 6.2 m RMS with 8.8 m at most, each within one unit of the last digit.

(b) Set the Compass bias to 2° and run. Report the same tiles. *Checkpoint:* 48.0 %, 2 holidays with the widest 12.0 m, and 12.3 m RMS with 17.3 m at most, each within one unit of the last digit.

(c) Set the Navigation to USBL, with the bias at 1°, and run. Report the same tiles and the Fixes tile, and explain why the error falls by a factor of about nine. *Checkpoint:* 50.0 %, no holiday, 0.7 m RMS with 3.7 m at most, 64 fixes used and none refused, each within one unit of the last digit.

13.8 App challenge · 7 min

Load scenario 5, The canyon, with the default seed. Run at 20× until the vehicle surfaces and the run pauses, at about 700 s on the clock. Read the Holidays, Navigation error and Fixes tiles, and watch the status line while lines 6 and 7 are flown. Reload the scenario before each new configuration and make the change before pressing Run.

(a) Run as loaded, and watch the navigation error strip while the Doppler log is on water track. Report the tiles and what the status line says on lines 6 and 7. *Checkpoint:* 1 holiday, 25.0 m wide, and 20.2 m RMS with 21.6 m at most, each within one unit of the last digit. Coverage not assured.

(b) Set Crab compensation On on the LOS card and run. *Checkpoint:* 1 holiday, 21.0 m, and 20.5 m RMS with 21.7 m at most, each within one unit of the last digit.

(c) Set the Navigation to USBL and run. Report the tiles, and explain what the aided filter measures that dead reckoning does not. *Checkpoint:* 1 holiday, 4.0 m, 0.4 m RMS with 1.7 m at most, 36 fixes used.

13.9 App challenge · 9 min

Load the challenge, The ten-line survey, with the default seed. Each run ends when the vehicle surfaces and the run pauses, about 2560 s on the clock, or 43 s at 60×. Read the Seabed imaged, Holidays, Lowest altitude, Recovery miss, Battery and Fixes tiles. Reload the scenario before each new configuration and make the change before pressing Run.

(a) Run as loaded. Report the tiles and which of the five criteria fail. *Checkpoint:* 83.9 %, 9 holidays with the widest 42.0 m, 1.7 m, 61.3 m, 16.7 %, each within one unit of the last digit. Four criteria fail, only the battery passes.

(b) Set only the Navigation to USBL and run. Report the tiles. *Checkpoint:* 79.7 %, 11 holidays with the widest 62.0 m, 1.1 m, 31.8 m, 1 USBL fix used and 132 rejected, each within one unit of the last digit.

(c) Press Reveal suggested tuning and run. Report the tiles, and name the setting of the suggestion that deals with each failure of part (a). *Checkpoint:* 100.0 % (99.95), no holiday, 5.5 m, 1.9 m, 16.6 %, 125 fixes used and 7 rejected, each within one unit of the last digit.

13.10 Application area · 4 min

A survey company plans a dead-reckoning survey with long lines and wants to choose the heading reference and the line spacing. Use Equations (13.10), (13.6) and (13.13).

(a) Compute the dead-reckoning error 2 km from the last fix with a compass biased by 1° and with a gyrocompass biased by 0.05°. *Checkpoint:* About 34.9 m and 1.75 m.

(b) At 10 m altitude, the navigation is expected to hold a 95 % uncertainty across the line of 8 m. Find the widest line spacing for which the supervisor's test still vouches for coverage, and name one other error the spacing must leave room for. *Checkpoint:* 24 m. The sag over a ridge, or the current's line-of-sight offset.

13.11 Application area · 3 min

In case 6 each GPS surfacing takes about 256 s and 3.45 Wh from the end of one line to the start of the next.

(a) A vehicle on dead reckoning with GPS surfacing every 15 min is to fly 6 h of survey lines. Estimate the extra time and energy the surfacings cost, and the share of the mission spent on them. *Checkpoint:* 24 surfacings, about 1.7 h and 83 Wh, about 22 % of the mission.

(b) Name what the surfacings do not correct, and what USBL aiding needs instead. *Checkpoint:* The compass bias stays, so the pattern keeps rotating. The USBL needs a ship on station with the vehicle inside its cone.

13.12 Review the agent · 5 min

An AI assistant is asked how to cut the navigation filter's memory and computation. It replies: "Drop the 24 s history and apply each USBL fix when it arrives. A fix is only about 1.3 s late, which is 2 m of travel at 1.5 m/s, well inside the 5 m each swath can spare, so coverage will not suffer." Find the flaw, and say how you would check it.

CHAPTER 14

Summary and Outlook

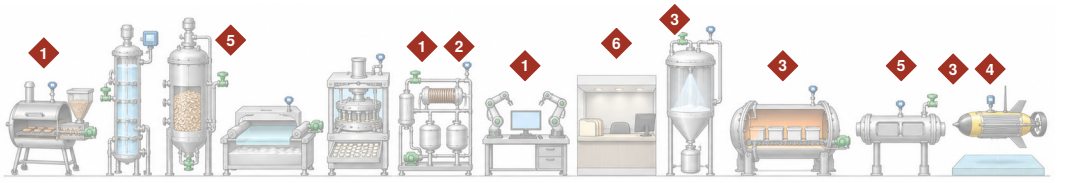

Open gaps 1 moving delays 2 corrections per material 3 unobserved directions 4 gate recovery 5 model resolution 6 calendar decisions

Notice that the gap of moving delays joins the grill, the flow synthesis skid and the teleoperator, while the gap of unobserved directions joins the spray dryer, the kiln and the survey vehicle. Illustration, not to scale.

IN PLAIN TERMS **What Twelve Late Measurements Taught**

The main idea. Every chapter of this book ran something that learns about itself late, or never, from a grill to a pension fund and a vehicle under the sea. Some designs waited for the news. Others predicted it with a model and corrected the model when the news came.

The challenge. A prediction is only as good as its idea of how long the news takes. An estimate cannot fix what no instrument sees, and some loss is already on its way before any controller can act.

Why it matters. Asking what each estimate read, and testing it where the model is wrong, separated the designs that worked from those that failed.

Chapter 1 asked two questions of every plant in this volume. What does each delay or missing measurement cost, and what recovers it? Twelve chapters answered them on a grill, process lines, a teleoperator, a pension fund, an exhaust line and a survey vehicle. This chapter sets the answers side by side, draws a short guide for choosing a method, looks at where practice is heading, reports how the volume was made with AI agents and what they got wrong, and lists the problems that remain open.

14.1 What the Chapters Showed

Table 1.4 named the delay or missing measurement of each chapter and the remedy, without results. Tables 14.1 and 14.2 return to it with the measured result. The first holds the chapters whose measurement is right but late, through a fuel line, an analyzer, a vessel, a ribbon, a reactor or a network. The second holds a valuation a year old and four plants whose key variable no instrument reads. Each row names a structure that recovered the cost, one that did not, and the result with the table or section that reports it. Figure 14.1 sets one like pair of numbers from each row side by side.

The columns of the two tables make five points.

A predictor is only as right as its delay. The MPC of Chapter 2 overshot the grill's lid opening by 2.7 °F when it assumed the lid would shut and by 92.8 °F when told it would stay open, with the same opening in both runs. A model auger delay of 0 or 60 s against the true 20 s cost almost nothing, because the heat bias absorbed it, while 120 s doubled the overshoot (Section 2.6). The predictive controller of Chapter 3 omitted the analyzer lag and stopped settling at 1.15 min of lag, where the PI loop still decayed. On the flow synthesis the plug-flow delay, timed from a flow the air lock had

Table 14.1 What Each Delay Cost and What Recovered It, Chapters 2 to 8

Chapter	Delay	What recovered it	What failed	Measured result
2, pellet grill	A 20 s auger, then ignition and burning. A lid that opens.	A lid-aware hold of the feed. An MPC whose filter carries the fuel in transit.	The PI tuned from the step test. The MPC told that the lid stays open.	Overshoot after the lid 9.7, 3.0, 2.7 and 92.8 °F (Table 2.6). Challenge in band 70.1 shipped, 92.5 % with the MPC (Table 2.7).
3, column	An analyzer lag up to 2 min on a column with 0.21 min of apparent dead time.	Detuning to $K_c = 36$, $\tau_I = 5$ min. A predictive controller with a disturbance state.	One tuning carried across the lags. A predictive model that omits the lag past 1.15 min.	Overshoot 7.6 to 30.7 %, settling 13.5 to 23.6 min (Table 3.4). Feed upset IAE 0.026 against 0.061 (Table 3.5).
4, digester	Four hours from heater to blowline. Three instruments on three clocks, one drifting.	H-factor feedforward. The combination held on an estimate reconciled with the laboratory.	Residual alkali held on a new wood. Feedback on the raw drifting analyzer.	Loss \$119.68 against \$7.27 per tonne (Table 4.3). Challenge passed at 0 of 11 seeds shipped and 11 of 11 reconciled (Table 4.5).
5, float line	Thickness read 15.2 min after forming. Stress settled before the polarimeter.	A forming model that makes the gauge loop a Smith predictor. Thickness by mass balance at the camera.	The same tuning with the model off, which hunts with a period of about 54 min.	Rejected 734 by the gauge, 212 by mass balance, 1827 m^2 with no model (Table 5.6). Challenge accepts 1367 against 8745 m^2 (Table 5.7).
6, tableting line	Dose fixed at the feeders twelve minutes before the press. A refilling feeder. A late analyzer.	Refill and humidity feedforward. A tracker that closes the gate on what the analyzer will judge late.	Tracking off with a 40 s analyzer delay. The line as shipped.	Under-dosed tablets accepted 816 without tracking, none with it (Table 6.5). Challenge out-of-limit accepted 194 shipped, none tuned (Table 6.6).
7, flow synthesis	Delay and spread that change with throughput. A drum only the analyzer sees.	Diversion on residence-time quantiles counted in volume. A coil longer than the detection time.	A plug-flow delay timed from a changed flow. A 10 min method at 100 mL/min.	Air lock: 1.53 L diverted, none released, against 177 mL released (Table 7.5). Drum escapes 10 of 12 times (Table 7.6).
8, teleoperator	A round trip up to 2.6 s, with jitter and loss, inside a force loop.	Wave variables, a passivity controller, or model mediation once its model is right.	Full force reflection. Reflection cut to 0.2. Waves and passivity together.	Challenge reversals and peak force 41 and 108.8 N shipped, 15 and 6.42 N waves, 3 and 2.00 N passivity, 12 and 4.08 N mediation (Table 8.6).

changed, released 177 mL that tracking in volume diverted. A delay kept as a volume of material, as Section 1.4.4 recommends, stays right when the flow moves.

Follow the material, not the clock. Every process chapter that met its target tagged material with what was done to it. The digester carried its chips as slugs advanced on the measured flow. The float line tagged each parcel of ribbon with the speed and thickness of its birth. The tableting line's tracker closed the gate on 816 tablets that the late analyzer would have passed, and the kiln's corrected feedforward waited for the affected material to reach each zone. The rule that sized the synthesis coil, detection time against holding time, predicted the escapes: none where the coil held 20.83 min and 10 of 12 where it held 12.50 min against 20.5 min of detection.

Table 14.2 What Each Late Result or Unmeasured State Cost, Chapters 9 to 13

Chapter	Late or missing	What recovered it	What failed	Measured result
9, pension fund	A valuation a year old, one decision a year, a contribution cap.	A predictive plan on a projected estimate, with a funding buffer and a glide path.	The shipped amortization rule. A seven-year deadline under a 12.5 % cap.	The year-old valuation raises the contribution SD from 1.91 to 2.35 points (Table 9.5). Challenge met on 0 of 5 seeds shipped, 5 of 5 suggested (Table 9.7).
10, spray dryer	Moisture from the laboratory 40 min late. A sticky margin never measured.	A psychrometric soft sensor whose laboratory bias is aligned with each sample, and feedforward.	The outlet loop alone at fixed setpoints. The soft sensor without its correction.	Case 4 off specification 0.57 shipped, 1.56 t uncorrected (Table 10.7). Campaign 1.98 to 3.88 t in specification (Table 10.9).
11, kiln	The brick's core never measured, and 30 h of transport.	A core cascade on a two-node soft sensor. A supervisor that predicts the exit quality.	The same cascade on the surface. Zone loops that hold a biased thermocouple perfectly.	Exit ordering 0.9900 core, 0.9834 surface. Good rate 6.07 to 8.14 t/d. Quality 96.6 with the bias against 98.6 (Table 11.4).
12, urea dosing	Stored ammonia never measured. A sensor that reads ammonia as nitrogen oxides.	Storage control on a model-based observer at a 100 % level.	A trim on the tailpipe sensor. Every feedforward ratio tried.	Trim 0.293 g/kWh with a 1301.9 ppm slip peak. Storage 0.423 and 23.5 ppm, passing. Feedforward at 0.6: 1.466 g/kWh (Table 12.5).
13, survey vehicle	Position never measured under water. Acoustic fixes late and sometimes wild.	An extended Kalman filter that applies each fix at its instant behind a gate.	Dead reckoning. Fixes applied at arrival. The USBL alone switched on.	Error RMS 0.56 at the instant, 1.92 m at arrival (Table 13.5). Survey 83.86 to 99.95 % imaged, 79.68 with the USBL alone (Table 13.6).

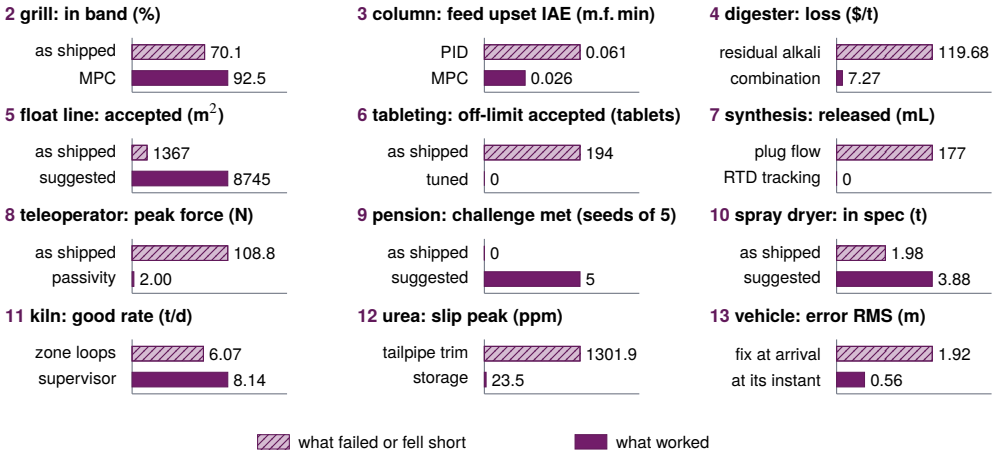

Figure 14.1 One like pair of numbers from each row, as the "Measured result" column of Tables 14.1 and 14.2 prints them. Each panel has its own scale. The hatched bar is the structure or setting that failed or fell short and the solid bar the one that worked; a value of zero has a label and no bar. Lower is better on every panel except the grill, the float line, the pension fund, the spray dryer and the kiln. Notice that no panel compares two gains: each better number came from a change of structure or of what the controller was told, such as a storage observer, a passive channel or a fix applied at its instant.

A late result should correct a model, not drive the loop. The float line's mass balance saw a bath drift 2.3 min after forming instead of 15.5 min and cut the rejected glass from 734 to 212 m^2. The spray dryer's soft sensor, corrected by results 40 min late, took a campaign from 1.98 to 3.88 t in specification, while a loop on the laboratory alone oscillated with growing amplitude. On the digester the reconciled estimate held kappa within 28.68 to 29.51 under a drifting analyzer that drove the raw loop to 26.52 (Section 4.6). The same chapter adds a caution. Its estimate's gate held back a large new offset for 75 minutes, so on a day without the drift the raw analyzer held kappa within ± 2 at 11 of 11 seeds and the estimate at none. An estimate built for a bias is slow as the input of a loop.

What no measurement sees stays wrong. The kiln's zone loops held a thermocouple that read 15 °C high perfectly and cost the product 2.0 points of quality. The spray dryer's laboratory bias corrected the moisture and left the sticky margin 0.9 K optimistic (Section 10.6). Over the survey vehicle's canyon a current carried it 21.46 m from its estimate and left a holiday 25 m wide that no tuning recovered (Section 13.5.5). In each case the estimate was consistent with every measurement it had. Section 1.4.6 names the cause: a direction the measurements cannot see is corrected by nothing.

Some cost is set before any controller acts. Chips already past the digester's heater at a production step could not be rescued by any heater setting, and their worst slugs ran $+2.46$ to $+4.49$ off kappa (Section 4.4). A failed annealing zone sent 915 to 918 m of glass over its limit under every structure (Table 5.10). The teleoperator's stable reflection gain fell toward 0.21, the hand's stiffness over that of the slave and wall in series (Section 8.5). The pension fund's cap made a seven-year deadline impossible for every amortization period (Section 9.9). Each chapter met such a limit with equipment or structure, a diverter, a passive channel or a buffer, and not with gain.

14.2 Choosing a Method

The chapters suggest an order of questions for a loop that sees its plant late. Each names a chapter where the answer decided the result.

1. *What is the delay made of?* Separate transport from lags that a fit reports as dead time and from the age of a held result, $\theta_e = \theta + t_d + T_s/2$ (Equation (1.8)). The grill's 20 s auger fit as 4.7 or 2.8 min, and the column's 0.21 min contains no delay (Chapter 2, Chapter 3).
2. *Does the delay move with flow?* Keep it as a volume or track the material, and never time it from a flow an upset has changed (Chapter 7, Chapter 4).
3. *Can a model or a faster instrument stand in between results?* Put a soft sensor upstream of the delay and correct its bias with each slow result, compared with the estimate stored for the same material (Equation (1.14), Chapter 5, Chapter 10).
4. *Can the measurements separate what matters?* Test observability before building the estimator, and add the measurement or the maneuver that restores it, as the survey vehicle's opposed lines do for its compass bias (Chapter 13, Chapter 11).
5. *Is part of the upset already past the actuator?* Then hold or divert it, and size the holding equipment from the detection time (Chapter 7, Chapter 6).
6. *Does the loop decide on a calendar from a late statement?* Keep the correction per decision near the bound of Equation (1.9) and add a projection across the lag (Chapter 9).
7. *Does the delay sit inside an energy exchange with a person?* Keep the channel passive first and recover feel second (Chapter 8).

Figure 14.2 draws the seven questions in order.

A predictive controller paid where its model carried something the loops could not: fuel in transit, a disturbance state or a preview. It held the grill's challenge at 92.5 % in band where the best lid-aware

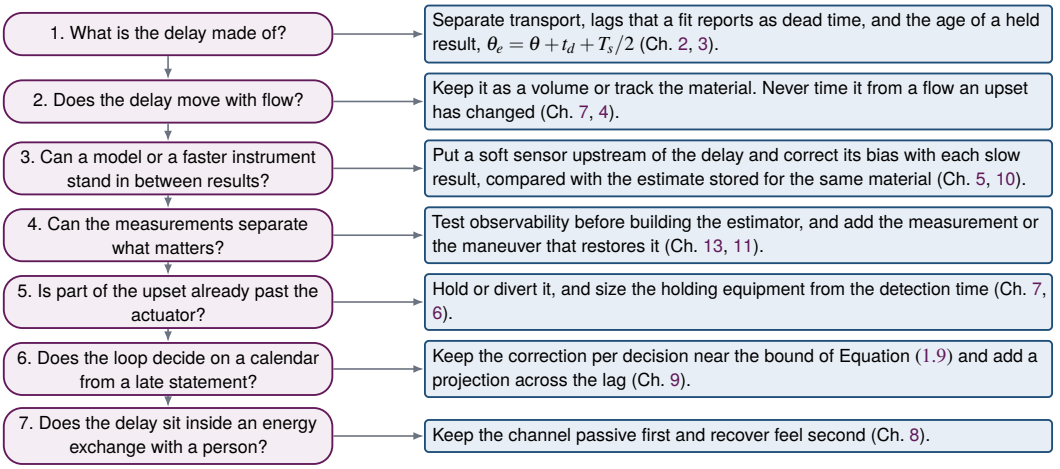

Figure 14.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that what the delay is made of, and whether it moves with flow, come before any estimator or predictor is chosen: a wrong delay cost the column's predictive controller and the flow synthesis's plug-flow timing.

PI reached 90.7 at the top of its gain range (Table 2.7), and it rejected the column's richer feed with an IAE of 0.026 against 0.061. Its advantage often came from the estimate rather than the optimizer. On the spray dryer it shipped 0.444 t off specification against 0.570 t for the loops, and the moisture loop retuned alone shipped 0.413 t (Table 10.8). On the flow synthesis the retuned trim diverted nothing through the rate change, as the MPC did (Table 7.8). Without preview the urea MPC failed its slip peak at 34.6 ppm where storage control passed (Table 12.6). Volume 9 treats that choice in full.

14.3 Where Practice Is Heading

Prediction moves inside larger controllers. Dead-time compensators, from the Smith predictor to its robust extensions, were surveyed in 2008 [1]. Bilateral teleoperation has its own historical survey, in which the delay of the channel is the central problem [2]. A 2016 survey traces model-mediated teleoperation from the late 1980s, a method developed to keep a teleoperator both stable and transparent under arbitrary communication delays, and names the challenges of building a reliable one [3]. In the editor's view, the predictor of this volume is becoming the model inside a predictive controller or a mediated interface, and the open work is keeping its delay right when flow, throughput or a network moves it.

Estimation takes late and out-of-order data as normal. In multisensor tracking, measurements from one target arrive out of sequence, and an exact state update for such arrivals was published in 2002, where earlier solutions were approximate [4]. A real-time algorithm for moving horizon state and parameter estimation was published in 2011 [5]. A 2014 review of underwater vehicle navigation notes that radio signals and satellite positioning do not reach the vehicle and that acoustic communication is low in bandwidth and unreliable. It finds that inertial sensors, installed beacons and periodic surfacing are fundamentally limited, and that advances in acoustic communication and in simultaneous localization and mapping open new possibilities [6]. The survey vehicle of Chapter 13 sits at the start of that path, with fixes stored and applied at their instant.

Soft sensors are learned, and material is tracked by its distribution. A 2021 survey reports that new characteristics of industrial processes degrade traditional soft sensor models and reviews deep learning models for data-driven soft sensors, with the problems met in practical applications [7]. In pharmaceutical manufacturing, regulators describe support for continuous production through science- and risk-based approaches, built on quality by design and process analytical technology [8]. Residence time distributions have been proposed as the basis for defining a batch, tracing raw materials and setting sensor frequency on a continuous line [9], the role they play in Chapter 6 and Chapter 7. In the editor's view, a learned soft sensor still needs a laboratory correction aligned with the material it describes, which is the lesson of Section 1.4.8 and not a property of any model class.

> ⧖ HISTORY **1949 to 2021: From Tongs Behind Glass to Semi-Autonomous Operation**
>
> The timeline of Chapter 1 places a remote hand in 1949. The report of that year described tongs worked behind a shielding barrier, such as a lead glass window, through an entirely mechanical master-slave linkage that an operator learned in five to fifteen minutes [10]. A later study of bilateral force reflection for a space robot found that delays as small as 10 ms, from communication or computation, can significantly increase the stiffness and damping the operator must supply [11]. A 2021 review of remote intervention in space reports that force reflection has proved impractical and unstable under small to medium delays, that wave variables restore stability without feeling intuitive, and that many researchers have turned to semi-autonomous teleoperation [12]. In the editor's reading, the passive channel and the mediated model of Chapter 8 stand between the two ends of that arc.

Hidden states gain instruments of their own. On an exhaust catalyst, ammonia and nitrogen oxide sensors give only an inference of the ammonia stored, confounded by the oxide sensor's cross-sensitivity to ammonia. Radio-frequency sensing uses the catalyst housing as a microwave resonant cavity, its readings have correlated highly with the stored inventory, and earlier studies showed its potential as the feedback of closed-loop urea dosing [13]. A thermocouple gives no indication in use that its output has drifted unless it is removed for recalibration, and a review of drift describes in situ calibration against the melting points of built-in reference materials or by Johnson noise thermometry [14]. The urea doser of Chapter 12 estimates its store with an observer, and the electric kiln of Chapter 11 reveals its biased thermocouple only through the product. In the editor's view, a hidden state that gains an instrument turns part of gap 3 of Section 14.5 into a question of validating that instrument.

Navigation relies on its estimator where no fix reaches. A survey of underwater navigation reports that one deep survey vehicle carried out 191 surveys of the seabed at mid-ocean ridges, at an average depth of 2000 m and with a navigation precision of a few meters. The same survey finds surface navigation effectively solved by satellite positioning, while in the middle of the water column acoustic time-of-flight systems and inertial units are the only effective sensors of horizontal position, and it expects model-based state estimators to be central to the advances there [15]. The survey vehicle of Chapter 13 meets that zone over its canyon. In the editor's view, fixes stored and applied at their instant are the part of its filter that practice will keep.

Soft sensors are maintained against the laboratory. Most industrial predictive control products surveyed in 2003 corrected their predictions with the difference between the measured and the predicted output. That feedback assumes a step disturbance at the output that stays constant, and it rejects disturbances entering at the process input sluggishly [16]. A review of industrial soft sensors holds that a soft sensor should be updated as the process characteristics change and that recursive updates cannot cope with abrupt changes. It reports an industrial application in which a bias between measurement and estimate was unavoidable after the operating condition changed [17]. The bias weight of Chapter 10, which trades speed against noise, and the persistence gate of Chapter 4 are

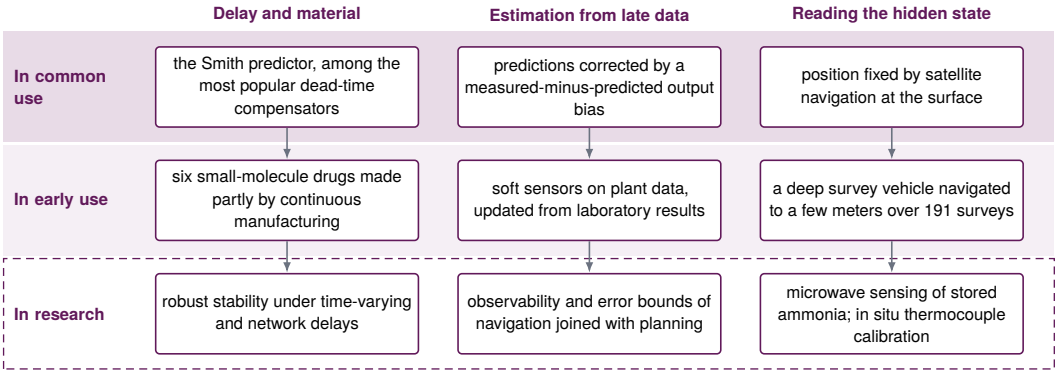

Figure 14.3 Developments in three threads. The top lane holds what the cited source calls popular, solved or used by nearly every product surveyed, the middle lane what it reports in service or on plant data, and the bottom lane what it reports from research, as a proposal or as a need, in the editor's reading. Sources by thread: delay and material [18, 19], estimation from late data [15–17] and reading the hidden state [13–15]. Notice that each thread moves from a fixed correction toward an estimate or a sensor that must itself be checked while the plant runs.

two answers to that tension. In the editor's view, the update rule of a soft sensor is part of its design. Figure 14.3 sets these developments in three threads.

14.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [20]. 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 [21]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

These five are the professional form of the habit introduced in Section 1.7: check before trusting, and know enough to check.

14.4.1 How This Volume Was Made

Each of the twelve apps was built and tested by agents. An agent then wrote the paper on it with verification notes that trace each number to a run. Each paper received a referee report, and at least seven of the twelve are self-reviews by the agent that wrote the paper. Defects went into a log for each app. For this book one agent per chapter reran the numbers its exercises use. The editor set the scope of each round and what each agent could change, and statements that runs contradicted were left for the editor.

The defect logs. The twelve logs give 163 entries a severity. Of these, 37 are wrong behavior of an app, 90 are documentation, 27 are cosmetic or display, and 9 are of other kinds. The kiln alone holds 10 of the 37, and its log opens with nine defects found while its paper's numbers were being measured. Five cases, chosen because each lies in a predictor, a measurement record or an estimator, show what the agents got wrong.

Estimators that knew the plant. Five logs record an estimator, a soft sensor or a controller's start that read the run's hidden settings or the plant's true state. The flow synthesis's soft sensor ran on the run's own kinetics. With a rate constant halved it predicted the sample point's 95.934 % assay exactly, before any analyzer result, where a soft sensor on the declared model predicts 98.659 %. The grill's estimator started from the plant's true state, and 18 of 33 parameters edited before a decision moved the MPC's first move with every measurement unchanged. The pension fund's projection read the run's market returns, which no valuation had reported. No shipped run showed any of them, because a shipped run leaves every hidden value at its declared one. Audits found them by changing a hidden value before Run and asking whether the estimate moved while the measurements did not.

A feedforward that ignored its own delay. The kiln's first throughput feedforward corrected all ten zones the instant the load cell moved, although the material needs up to 1710 min to reach the last zone, and it scaled the whole design duty of each zone. It raised the peak zone error after a density change from 0.17 to 0.98 °C and the kiln-wide error from 4645 to 11 204 °C min, while the scenario's text claimed the opposite. The paper agent found it by measuring the case the text

described. The corrected feedforward waits for the material, and the chapter reports the honest result that a well-tuned zone loop leaves it little to do.

A reading credited to the wrong glass. The float line's inspection credited each parcel with the gauge's filtered output at the first scan after the parcel reached it. That reading mixed the parcel with its neighbors. With the thickness loop hunting, adjacent parcels differed by up to 0.19 mm, the credited reading by up to 0.086 mm against a guard band of 0.01 mm, and four parcels truly out of tolerance were accepted. The defect was first documented as a limitation of one configuration. A sweep of 20 noise seeds then showed the same mechanism failing the suggested challenge at 2 of them, and each parcel is now credited with a reading of its own glass.

A gate that locked itself out. With the survey vehicle's shipped navigation switched to the acoustic fixes alone, the first fix to arrive on the default seed was an outlier 24.2 m off. The filter's uncertainty was still wide, so the gate accepted it and read most of the error as compass bias, -5.41° against a true $+1^\circ$. Its covariance then collapsed and never reopened, and 132 later fixes were refused. The survey imaged 79.68 %, less than dead reckoning. The behavior was kept as a lesson and the texts now say so. The referee then found that the paper called all 132 refused fixes honest, where three were outliers.

Algebra that the simulation overruled. The agent that wrote Chapter 1 drafted two derivations that were wrong until its script disagreed. The laboratory loop's characteristic equation had one delay step too many, and the Smith predictor's delay tolerance looked larger than it was because an output clamp turned a growing oscillation into a bounded cycle. Bisection by simulation with the limits removed exposed both. The spray dryer's referee found a claim about the laboratory weight that was reasoning, not a measurement, and the case was rerun at three weights before the paper kept it.

The book round. The twelve chapter records list 183 rows of numeric verification, and ten report no drift. The digester's text said a trim held its analyzer near 30 where its own figure shows 32.1 to 34.5, and the wording now follows the figure. The urea chapter said no case's sensor reading crosses 1.2 g/kWh, where the shipped challenge reads 1.689. The paper and the chapter now report 1.689 and 2.876 g/kWh. The tableting paper named its controller actions the opposite way to the series convention, and both its chapter agent and the agent writing Chapter 1 flagged it. At the end of that round the editor's list of defects held nine entries for this volume.

Every error above was caught by a check that could fail: a hidden value moved, a case run as its text described it, a sweep over seeds, a scenario changed by one setting, a derivation simulated with its limits removed. For estimation the record adds one rule to those of the shared section. Before trusting an estimate, ask what it read, and test it on a case where the model is wrong.

14.5 Gaps in Delay Compensation and State Estimation

Six gaps belong to the topic of this volume, and the gaps common to agent-era practice follow in Section 14.5.1. In the editor's view none is closed. Each item gives the chapters' evidence, then the state of the literature with its sources. Why each gap is hard and what progress would look like are the editor's view. Figure 14.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the plants in the picture that opens this chapter.

1. *Delays that move with what is not measured.* The flow synthesis's plug-flow delay was wrong once an air lock changed the flow, the grill's pulsed auger lengthened its transport as the feed fell, and the teleoperator's channel jittered. A review of dead-time compensation names the conservativeness of robust stability criteria for time-varying delay as a topic that deserves further investigation [18]. A survey of model-mediated teleoperation calls environment modeling its

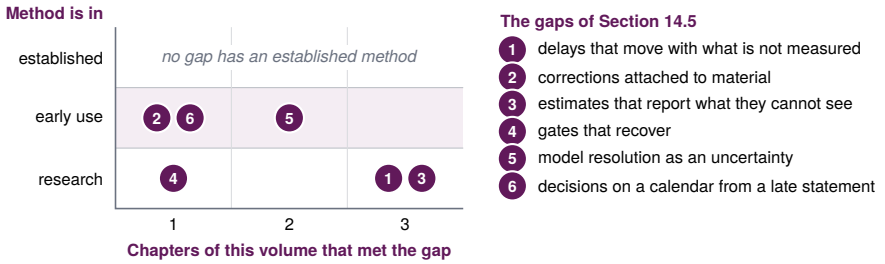

Figure 14.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: three for gap 1 (Chapters 7, 2 and 8), one for gap 2 (7), three for gap 3 (11, 10 and 13), one for gap 4 (13), two for gap 5 (12 and 4) and one for gap 6 (9). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the two gaps met by three chapters each, moving delays and unobserved directions, are both still in research.

foremost challenge, because a sudden change of the local model produces a sudden force [3]. The gap is hard because the delay moves with the very flow or network that the upset disturbs. Online delay estimates that state their own uncertainty are not routine.

2. *Corrections attached to material.* The flow synthesis applies one analyzer bias to every batch in the line, and under the shipped supervisor rule its trim released 291 mL after an unassayed drum 15 % strong. Regulators writing on continuous manufacturing note that poor measurement is often due to the sampling system rather than the analyzer, and suggest using the residence time distribution to set the sampling frequency [8]. A study of traceability finds that the back-mixing a robust line needs confounds the boundaries between batches [9]. The gap is hard because a laboratory result describes material that has since mixed with its neighbors. Bias updates that belong to the material sampled, with weights chosen per stream, are still designed plant by plant.

3. *Estimates that report what they cannot see.* The kiln's surface-only correction, the spray dryer's optimistic margin and the survey vehicle's canyon drift all left the estimate consistent with every measurement. The Open challenge box after this list takes the gap further.

4. *Gates that recover.* The survey vehicle's filter locked out after one wild fix, and the chapter leaves a lockout monitor as further work. A survey of industrial predictive control found that most products detect only sensor failures beyond preset limits, while significant bias and drift within normal limits generally go undetected [16]. The gap is hard because, from inside the filter, a refused honest measurement and a wrong estimate look alike. Rules for reopening a covariance after repeated rejections are heuristic.

5. *Model resolution as an uncertainty.* The urea challenge passes on the declared three cells and fails on six or twelve (Table 12.7), and a grid of 40 cells smeared the digester's kappa by 2.64 %. The same survey found that, until performance and robustness are connected, it will not be possible to say when a model is accurate enough, and that all products but one rely on brute-force simulation to evaluate model mismatch [16]. The gap is hard because the resolution that decides a verdict is rarely the one at which the design was checked. A verdict that depends on the number of cells needs that dependence stated with it.

6. *Decisions on a calendar from a late statement.* The pension fund's plan passed only by contributing nearly twice the normal cost. An actuarial analysis found that a one-year delay in the accounts loses information and raises the resulting variances, that an optimal spread period still exists, and that spreads of 20 to 25 years were common practice [22]. A federal agency's projections treat future employer contributions to multiemployer plans as uncertain and a significant risk [23].

The gap is hard because the cap, the deadline and the delay act together and are set by different parties. Design rules for policies that decide once a year on a valuation a year old, under caps, are thin.

OPEN CHALLENGE Estimates That Report What They Cannot See

Gap 3 runs through the kiln, the spray dryer and the survey vehicle, whose estimates stayed consistent with every measurement while an unseen direction drifted. A survey of underwater navigation reports that researchers are only beginning to investigate properties such as observability and lower error bounds for navigation joined with planning [15]. A review of industrial soft sensors notes that a soft sensor can be checked only when analyzed data are available, and that normal changes are hard to tell from faults [17]. In the editor's view the gap is hard because an unobserved direction leaves no residual to alarm on. Progress would be an estimator that reports which directions no measurement has corrected and how far they may have drifted since.

14.5.1 Gaps Common to Agent-Era Practice

The final chapter of each volume lists gaps specific to its topic. Seven more are common to every topic in the series. In the editor's view none of them is closed.

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [24]. 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 [21]. 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 [25]. 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 [26]. 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 [27]. 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

- A predictor was only as right as its delay, and a delay kept as a volume of material stayed right when the flow moved.
- A late result served best as the correction of a model aligned with the material it measured, not as the input of the loop.
- No estimator recovered a direction its measurements could not see, and cost already past the actuator was met with equipment or structure, not with gain.

References

[1] J. E. Normey-Rico and E. F. Camacho, "Dead-time compensators: A survey," *Control Engineering Practice*, vol. 16, no. 4, pp. 407–428, 2008.

[2] P. F. Hokayem and M. W. Spong, "Bilateral teleoperation: An historical survey," *Automatica*, vol. 42, no. 12, pp. 2035–2057, 2006.

[3] X. Xu, B. Cizmeci, C. Schuwerk, and E. Steinbach, "Model-mediated teleoperation: Toward stable and transparent teleoperation systems," *IEEE Access*, vol. 4, pp. 425–449, 2016.

[4] Y. Bar-Shalom, "Update with out-of-sequence measurements in tracking: Exact solution," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, no. 3, pp. 769–777, 2002.

[5] P. Kühl, M. Diehl, T. Kraus, J. P. Schlöder, and H. G. Bock, "A real-time algorithm for moving horizon state and parameter estimation," *Computers & Chemical Engineering*, vol. 35, no. 1, pp. 71–83, 2011.

[6] L. Paull, S. Saeedi, M. Seto, and H. Li, "AUV navigation and localization: A review," *IEEE Journal of Oceanic Engineering*, vol. 39, no. 1, pp. 131–149, 2014.

[7] Q. Sun and Z. Ge, "A survey on deep learning for data-driven soft sensors," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 9, pp. 5853–5866, 2021.

[8] S. L. Lee, T. F. O'Connor, X. Yang, C. N. Cruz, S. Chatterjee, R. D. Madurawe, C. M. V. Moore, L. X. Yu, and J. Woodcock, "Modernizing pharmaceutical manufacturing: From batch to continuous production," *Journal of Pharmaceutical Innovation*, vol. 10, no. 3, pp. 191–199, 2015.

[9] W. Engisch and F. Muzzio, "Using residence time distributions (RTDs) to address the traceability of raw materials in continuous pharmaceutical manufacturing," *Journal of Pharmaceutical Innovation*, vol. 11, no. 1, pp. 64–81, 2016.

[10] R. C. Goertz, "Master-slave manipulator," Argonne National Laboratory, Argonne, IL, Report ANL-4311, Mar. 1949.

[11] J. D. Chapel, "Performance limitations of bilateral force reflection imposed by operator dynamic characteristics," in *Proceedings of the NASA Conference on Space Telerobotics*, vol. 4. Pasadena, CA: Jet Propulsion Laboratory, California Institute of Technology, 1989, pp. 91–100. [Online]. Available: https://ntrs.nasa.gov/citations/19900020524

[12] P. Kazanzides, B. P. Vagvolgyi, W. Pryor, A. Deguet, S. Leonard, and L. L. Whitcomb, "Teleoperation and visualization interfaces for remote intervention in space," *Frontiers in Robotics and AI*, vol. 8, p. 747917, 2021.

[13] P. Ragaller, J. Mandelbaum, L. Lapenta, A. Sappok, J. Pihl, V. Prikhodko, and J. Parks, II, "Direct measurement of ammonia storage on passive SCR systems for lean gasoline NOx reduction using radio frequency sensing," in *Proceedings of the ASME 2019 Internal Combustion Engine Division Fall Technical Conference*, no. ICEF2019-7224, Chicago, IL, 2019.

[14] R. J. Chesser and N. D. Bull Ezell, "Thermocouple and RTD drift due to neutron irradiation and temperature for space nuclear applications: A review," *IEEE Sensors Journal*, vol. 25, no. 22, pp. 40 965–40 977, 2025.

[15] J. C. Kinsey, R. M. Eustice, and L. L. Whitcomb, "A survey of underwater vehicle navigation: Recent advances and new challenges," https://robots.engin.umich.edu/publications/jkinsey-2006a.pdf, 2006.

[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] M. Kano and K. Fujiwara, "Virtual sensing technology in process industries: Trends and challenges revealed by recent industrial applications," *Journal of Chemical Engineering of Japan*, vol. 46, no. 1, pp. 1–17, 2013.

[18] J. E. Normey-Rico, T. L. M. Santos, R. C. C. Flesch, and B. C. Torrico, "Control of dead-time process: From the Smith predictor to general multi-input multi-output dead-time compensators," *Frontiers in Control Engineering*, vol. 3, p. 953768, 2022.

[19] C. Badman, C. L. Cooney, A. Florence, K. Konstantinov, M. Krumme, S. Mascia, M. Nasr, and B. L. Trout, "Why we need continuous pharmaceutical manufacturing and how to make it happen," *Journal of Pharmaceutical Sciences*, vol. 108, no. 11, pp. 3521–3523, 2019.

[20] 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.

[21] 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.

[22] S. Haberman, "Pension funding with time delays and the optimal spread period," *ASTIN Bulletin*, vol. 25, no. 2, pp. 177–187, 1995.

[23] Pension Benefit Guaranty Corporation, "FY 2024 projections report," Report, listed on the agency's Projections Report page, last updated January 27, 2026, 2026, accessed 2026-11-21. [Online]. Available: https://www.pbgc.gov/sites/default/files/documents/fy-2024-projections-report.pdf

[24] 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.

[25] 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.

[26] 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.

[27] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

14.6 Exercises

14.1 Discussion · 8 min

An agent proposes to shorten the closed-loop time constant of a Smith predictor on a dosing line whose throughput varies by a factor of two, because "the predictor removes the dead time". Using Section 1.4.3, Section 1.4.4 and Chapter 7, explain what the proposal risks and what you would ask for before accepting it.

14.2 Discussion · 8 min

A plant team wants to retire the spray dryer's moisture soft sensor and close the moisture loop on the laboratory alone, "because the laboratory is the measurement the customer pays for". Using Chapter 10 and Sections 1.4.5 and 1.4.8, argue for or against, and say how the laboratory result should be used.

14.3 Discussion · 8 min

In a demonstration an agent-built soft sensor on a reactor "tracks the plant exactly" from the first second of a run in which a kinetic constant was changed. Using Section 14.4.1, explain why that result is a warning, describe the test you would run, and say what a correct soft sensor would have shown.

14.4 Discussion · 8 min

The kiln of Chapter 11 and the survey vehicle of Chapter 13 both report confident estimates while the true state has moved away. Compare the two cases, explain why their estimators could not see the error, and propose one measurement or maneuver for each that would let them.

14.5 Discussion · 7 min

After the gate lockout of Chapter 13, a reviewer proposes to turn the chi-square gate off, since “the gate threw away 132 good fixes”. Using the chapter and Section 14.4.1, evaluate the proposal and state what you would change instead.

14.6 Discussion · 9 min

A manager reads that the spray dryer’s MPC shipped 0.444 t off specification against 0.570 t for the loops and asks for MPC on every dryer. Using Section 14.2, Chapter 10 and Chapter 7, explain where the improvement came from and what comparison you would run before deciding.

Answers to Selected Exercises

Short answers are given for the odd-numbered exercises. Numeric answers from an app are for its default seed and settings unless the exercise says otherwise.

Chapter 1. Introduction

1.1 (a) 1.70 g min/L. (b) $K_c = 1$, $\tau_I = 8$ min, IE $= 16$ g min/L.

1.3 (a) $\tau = 4$, $\theta = 2$ min at 1.5 m^3/min, and $\tau = 8$, $\theta = 4$ min at 0.75 m^3/min. (b) 9:56, the soft value logged then.

1.5 (a) $z^2 - z + g = 0$, stable for $0 < g < 1$, no oscillation for $g \leq 1/4$. (b) $z = 1 - g$, stable for $0 < g < 2$, no oscillation for $g \leq 1$, one result at $g = 1$.

1.7 (a) $q/r_v = 0.05$. (b) 11 results, $\sigma_{\hat{b}} = \sqrt{r_v}/3$.

1.9 9 min and 13.5 m^3.

Chapter 2. Pellet Grill: Control Through a Fuel Delay

2.1 (a). The extra pellets burn into a closed chamber after the lid shuts, and the integral holds its 2.9 %.

2.3 (c). The opening and the controller are the same in both runs, and only the assumption about the lid differs.

2.5 (a) 0.31 % per °C (0.17 % per °F) and 15.5 min. (b) 0.36 % per °C, so Smith's dead time gives the more cautious loop. (c) 15.9 % and 2.9 %, which take the feed from 23 to about 41.8 %.

2.7 (a) Standard PID 53.4 %, 9.7 °F, not yet. Lid-aware PID 82.5 %, 3.0 °F, 8.7 min. MPC 93.9 %, 2.7 °F, 0.7 min. (b) The hold stops feeding a disturbance it cannot fight, and the MPC also feeds what the steel lost.

2.9 (a) About −232 W and 97.0 %. The bias is a phantom heat loss that stands in for the weaker fuel. (b) 3.9 °F, about twice the overshoot with the right delay, and still within the 10 °F the challenge allows.

2.11 (a) About 9.4 °F, and no controller fed TT-101 can correct it. (b) Compare TT-101 with an independent, calibrated reading at the grate at a steady operating point, and log the difference.

Chapter 3. Distillation Control Through a Slow Analyzer

3.1 (a). D, L, V and L_S all scale with F, so the balances change only their speed.

3.3 (c). The lag is unmodeled phase, and the brisk MPC has less phase margin to give up than the PID.

3.5 (a) $K_c = 59.8$, $\tau_I = 7.07$ min. (b) $K_c = 109$, $\tau_I = 4.84$ min. (c) 15.0 against 8.46, so the shipped tuning integrates 1.77 times faster.

3.7 (a) MPC 0.028 and 7.22, PID 0.064 and 11.4. (b) MPC 0.218 with 12.7 % saturated, PID 0.122 with none. (c) The lag the MPC does not model makes its loop unstable at 1.5 min, while the PID keeps a positive margin.

3.9 (a) 0.0143, not met. (b) 0.00518 and 0.00555, both met. (c) MPC 0.0008 against PID 0.0021. The judged error is mostly analyzer noise.

3.11 (a) 3 min. (b) A delay of about 3 min leaves the default tuning unstable on the fitted model. Detune for the delay, or move the loop to a temperature trimmed by the analyzer.

Chapter 4. Kraft Digester Ahead of Its Kappa Analyzer

4.1 (d). The new wood lowers the residual, and the heater answers in the wrong direction for kappa.

4.3 (b). The gate and the small predictor weight slow a loop that already waits five hours for the vessel.

4.5 (a) 114.72 °C and 58.0 %. (b) About 68.9 %. (c) About 4.4 °C, which every slug crossing meanwhile would carry down the vessel.

4.7 (a) Between 4.0 and 4.2 h after the step. (b) Reported at 5.083 h, drawn at 4.75 h. (c) 27.28, about −1.36 per kelvin.

4.9 (a) 26.52, 33.81, 6.90 and 5.56. (b) About 4.4. (c) 28.68 to 29.51.

4.11 (a) Two errors cancel: the drifting analyzer makes the trim cook harder while wood B cooks slower. (b) The titrator, below 8 g/L from about 7.7 h. The floor guards against lignin redepositing on the fibers at the end of the cook.

Chapter 5. Frozen-In Stress on a Float Glass Line

5.1 (c). Between 78.0 m and 98.7 m the relaxation time grows from 21 s to 24 min, and the residual changes by less than 0.01 MPa after 96 m.

5.3 (a). With less grip more of the stretch narrows the ribbon, and the mass relation $G_f = \rho W t v$ then makes it thicker.

5.5 (a) 1.611 and 1.239. (b) 4.28 mm. (c) 1.770.

5.7 (a) About 36.1 min (row 36.3 min) and 734 m^2. (b) About 22.8 min (row 23.0 min) and 212 m^2. (c) Travel of 15.2 min and a filter of 0.2 min. The camera stands 25 m past forming.

5.9 (a) 15.2 K, 21.4 K, 1.05 MPa. (b) 0 K, 1.47 MPa, 1.40 MPa. (c) A warmer A0 makes the slowed thin glass keep more stress. The tile counts accepted parcels only, and the worst parcels were transition glass rejected for thickness, some also for stress.

5.11 (a) About 147 GJ per day, 0.25 GJ/t. (b) 3.8 %. The blanket melts in about 29 min instead of 36 min and the level rises 1.02 mm by 37 min.

Chapter 6. Continuous Tableting Weight and Dose Control

6.1 (b). The dose is the weight times the composition, and only the force sees the weight.

6.3 (d). The cascade learns only when a cell discharges, while the feedforward acts on the measured disturbance.

6.5 (a) 0.837 %. (b) +3.63 and −3.57 %, blend from 98.52 to 101.32 % LC. (c) $\varphi = 0.363$, 1.25 and 2.34 %. The NIR reads the blend's mean, which is on target.

6.7 (a) 2712 rejected, 2103 of them good, none accepted out of limits, 609 made out of limits. (b) 104.76 % LC at 810 s, weight 249.64 to 250.55 mg. (c) 0 and 100.39 % LC.

6.9 (a) 71 952, 14 448, 194, 9.39 %, not met. (b) 82 992, 3408, 0, 3.72 %, met. (c) 83 304, 3096, 122, 8.62 %. The bridge's lean package reaches the press before the NIR reports it.

6.11 (a) 1800 tablets, 0.45 kg. The rejects carry 85.0 g against 86.1 g at label claim, so they are slightly lean on average. (b) The rule accepts only material whose history is known. Omitted: the cost of the API, the disposal and the line time.

Chapter 7. Residence-Time Diversion in Flow Synthesis

7.1 (b). Dispersion in the reactor sends the leading edge of the slug ahead of the plug-flow time and its tail behind it.

7.3 (d). The retuned trim at the lower target matches the MPC on case 9. At an equal target the MPC wins or ties every row.

7.5 (a) $t_{det} = 6.5$, 14.5 and 20.5 min, fractions 0, 0.29 and 0.80. (b) 6.0 min, and about 4.7 min by the leading edge. (c) About 61 mL/min, and 0.26 at 70 mL/min.

7.7 (a) 0 mL and 1.53 L, shadow 177 mL and 1.79 L. (b) 177 mL, 94.98 %, 1.79 L. (c) From the front of the slug, 44.07 to 45.50 min, and its upset window diverted good material to 60.4 min.

7.9 (a) MPC 0.00 L, 1.18 mol. Loops 4.85 L, 1.51 mol. (b) 0.00 L, 1.18 mol. (c) 0.00 L, 1.22 mol. On case 9 the target, not the plan, removes the diversion.

7.11 (a) 30.2 min and 53 %. The longer residence lets the over-reaction grow. (b) A reactor full of hot material with excess reagent. Cool the zones while the feeds are stopped and plan the cleaning.

Chapter 8. Delayed Channel Energy in Teleoperation

8.1 (a). The delayed force multiplies a velocity it never met at the same instant, so the channel puts out energy that neither port put in.

8.3 (c). The hand feels the model's wall, which the peg has not reached when it meets the real one.

8.5 (a) 80 N/m giving 75.8 N/m, and 200 N/m giving 175 N/m, against 76 and 173 N/m measured. (b) About 37 N/m. (c) About 95 N s/m, with more drag in free space.

8.7 (a) 41 reversals, 25.7 N, 15.52 J. The ledger leads the flag by two touches. (b) 9 reversals, 10.8 N, 110.3 mJ, no flag. (c) No. The loop is stable but not passive.

8.9 (a) 16.3 s, 108.8 N, 41 reversals, not met. (b) 15.3 s, 6.4 N, 15 reversals, met. (c) 18.7 s, 4.1 N, 12 reversals, met by 4 reversals, 5.9 N and 6.3 s. Model mediation.

8.11 (a) About 0.105, below the commissioned 0.20. (b) Tune for the softest grip and monitor the ledger, or use a passive architecture.

Chapter 9. Pension Funding Ratio Control

9.1 (d). From the decision of year 5 on, every rule sits on the cap, so the applied contribution, and hence the recovery, is the same whatever the gain.

9.3 (b). The projection carries the recorded discount rate, so it sees the liabilities revalued, but it runs at the declared mean return and never reads the crash.

9.5 (a) $\phi = 1.0253$, $\gamma = 1.0126$, roots 0.758 and 0.267 against 0.823. (b) $m < 3.85$ years, $0.513 \pm 0.273\,j$ at $m = 3$. (c) $g = \gamma/m$, and the limit is $\gamma/m = \phi^2/4 = 0.263$.

9.7 (a) 0.998 and 2.35 points. (b) 0.994 and 1.91 points. (c) 2.93 points. The delayed rule pays on a deficit that has already shrunk, which moves the contribution more without slowing the recovery.

9.9 (a) 0.839, 2.85 points, 1.331. (b) 0.739, 3.52 points, 1.126. (c) 0.773, 4.40 points, 1.245. On five seeds the buffer raises the average lowest valuation from 0.673 to 0.775 at a mean contribution of 17.46 against 13.67 points.

9.11 (a) About 20.1 points above the normal cost, 30.1 % of payroll. (b) At 12.5 % it would take about 56 years. The rate exceeds the application's largest cap of 30 %. Contributions crowd out the sponsor's investment, which the model does not price.

Chapter 10. Inferential Moisture Control of a Spray Dryer

10.1 (c). A weaker concentrate moves the whole map of moisture against outlet temperature, so 80 °C no longer gives 3.5 %.

10.3 (a). With the outlet held, the inlet sets how much water the air must carry, while a feed cut at a hot inlet leaves the particles hotter.

10.5 (a) 70 min, 60 min after the change. (b) 55 min, about eleven times the bed's residence. (c) 4 results by 160 min with a standard deviation of 0.029 points at $\lambda = 0.5$, and 11 results by 370 min with 0.017 points at $\lambda = 0.2$.

10.7 (a) 80.1, 74.3 and 73.5 °C. 9.0 min. (b) 4.16 %, 1.83 K. (c) 201.1 °C, 87.0 °C, 2.95 %.

10.9 (a) 1.98 t in, 2.02 t off, –0.4 K, 4237 kJ/kg, not met. (b) 3.76 t, none off, 0.0 K, 4498 kJ/kg, not met. (c) 3.88 t, none off, 0.6 K, 4380 kJ/kg, met.

10.11 (a) The roof passes 158 °C at a lower inlet than the one that makes the powder sticky, so an override on the margin holds an inlet hot enough to ignite an old layer. (b) About 1.7 h. Hold the inlet below 208.5 °C and clean the roof.

Chapter 11. Cathode Kiln: An Unmeasured Core Temperature

11.1 (c). Table 11.1 lists the zone temperatures among the controlled and measured variables, and the other three among the controlled but unmeasured ones.

11.3 (a). The loop holds the reading, the ordering rate at 780 °C falls to about 0.66 of its value, and only the pyrometer, a second thermocouple or the assay can show it.

11.5 (a) 6.86 kJ/(min K). (b) $\tau_c = 22.6$ min, $\tau_s = 1.60$ min. (c) The core trails its environment by $R(\tau_s + \tau_c) = 0.40R$. Ramps of 60 and 75 °C/h give gaps of 22.6 and 28.2 °C.

11.7 (a) Zone 8.95 °C at 31.5 min, core 6.09 °C at 60.9 min. (b) Step 8.52 and 5.50 °C, 0.85 and 0.55 °C/%. (c) 34.6 and 63.0 min, against $\tau_p = 34.8$ min and the core's 63 min.

11.9 (a) 98.3 %, 97.3, trim 5.2 °C, surface 780.0 °C, core 775.7 °C. (b) 99.0 %, 97.9, trim 9.6 °C, surface 784.5 °C, core 780.1 °C. (c) Denser bricks and a faster line widen the surface-to-core offset, which the surface loop cannot see. An estimator that overstates K_{sc} reports the core too hot.

11.11 (a) 9724, 10235, 10307 and 10 582 kW h a day. (b) About 419 kW h per extra tonne against an average of 1602. At 125 % the kiln uses 10 307 kW h a day and makes no saleable product.

Chapter 12. Diesel Urea Dosing With an Ambiguous Sensor

12.1 (b). At $\alpha = 1.3$ the reading is nearly all ammonia, at 0.8 nearly all NO_x, and the two sides of the V give the same value.

12.3 (d). The tailpipe correction holds the coverage that matches the reading, so a capacity error appears in the grams.

12.5 (a) About 60.9, 12.2 and 59.1 ppm. (b) 79.3 and 79.9 %, both short of 95 %. (c) 269.2 ppm, an apparent conversion of 47.6 %.

12.7 (a) 384.5 ppm at 705 s, sensor 267.4 ppm, true NO_x 0.1 ppm. (b) 32.2 ppm at 692 s; 11.75 g and 10.1 g stored. (c) 43.1 ppm.

12.9 (a) 0.293 g/kWh, 108.3 ppm, 1301.9 ppm, 893 g; fails on ammonia. (b) 0.423 g/kWh, 5.8 ppm, 23.5 ppm, 470 g; met. (c) 0.362 and 30.0 ppm, met; 0.344 and 32.5 ppm, not met.

12.11 (a) The aged brick and the fresh brick on the feedforward are flagged; only the aged brick has lost sites. (b) Its slip reads as NO_x. A diagnostic needs the doser's state, the stored ammonia or the slip, beside the sensor's number.

Chapter 13. AUV Survey Navigation Drift

13.1 (a). With a constant bias every estimated velocity is the true one rotated by b, so the error is $2\sin(b/2)$ times the distance from the launch point.

13.3 (c). A wild first fix passes the wide gate, the filter takes it as compass bias, and every later fix lies far outside the narrowed gate.

13.5 (a) 1.745 %, 8.2 m and 3.2 m. (b) 26.1 m, 6.9 times too large. (c) About 21.5 m.

13.7 (a) 49.4 %, 2 holidays (7.0 m), 6.2 m RMS, 8.8 m max. (b) 48.0 %, 2 holidays (12.0 m), 12.3 m RMS, 17.3 m max. (c) 50.0 %, none, 0.7 m RMS, 3.7 m max, 64 fixes. The filter learns the bias.

13.9 (a) 83.9 %, 9 holidays (42.0 m), 1.7 m, 61.3 m, 16.7 %. Coverage, holidays, clearance and miss fail. (b) 79.7 %, 11 holidays (62.0 m), 1 fix used and 132 refused. (c) All five met: no holiday, 5.5 m, 1.9 m, 16.6 %.

13.11 (a) 24 surfacings, about 1.7 h and 83 Wh, about 22 % of the mission. (b) The bias stays. The USBL needs a ship on station.

Chapter 14. Summary and Outlook

14.1 A good answer notes that the predictor removes the delay only while its model's delay is right. On the example of Chapter 1 the PI loop tolerates a true delay of 15.6 min against a 5 min model and the predictor with $\tau_c = 1$ min only 8.9 min, so a shorter τ_c buys speed with exactly the robustness that a

changing flow uses up. At half flow the predictor fixed in time rings with an IAE of 26.1, while one whose delay is kept as a volume gives 15.0. On the flow synthesis a plug-flow delay timed from a changed flow released 177 mL that tracking by volume diverted. The answer asks for the delay to be kept in volume on the measured flow, a test at the lowest and highest throughput, and the longest true delay at which the loop stays stable.

14.3 A good answer recognizes the flow synthesis case. With a rate constant halved, the soft sensor predicted the sample point's 95.934 % assay exactly before any analyzer result, because it ran on the run's own parameters, while a soft sensor on the declared model predicts 98.659 %. No shipped run showed it, since shipped runs leave the hidden values at their declared ones. The test changes a hidden parameter before Run, holds the measurements, and asks whether the estimate moves. A correct soft sensor starts from the declared model, shows a mismatch until the first analyzer result arrives, and then corrects it through its bias.

14.5 A good answer notes that the gate was not the fault. It accepted one wild first fix, 24.2 m off, while the filter's uncertainty was wide, read it as a -5.40° compass bias, and then refused 132 later fixes because its covariance had collapsed and never reopened. Three of the 132 were outliers, so not all were good. Turning the gate off helped on the challenge but raised the case 7 error from 0.56 to 1.51 m RMS. The better change is a lockout monitor that reopens the covariance after repeated rejections, a tighter prior on the bias, or holding the first fixes until two agree, tested over many seeds.

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.

alignment of a result Pairing a late result with the estimate stored for the same material, so that a correction does not answer its own last move. *Chapter 1.*

analyzer An instrument that measures the composition or another quality of a sample, often slowly and with a delay before it reports. *Chapter 1.*

analyzer lag The delay and smoothing between a change in a product and its appearance on the analyzer reading. *Chapter 3.*

annealing Cooling glass slowly along a set temperature path so that the stress frozen into it stays small. *Chapter 5.*

arrival cost The term of a moving horizon estimator that sums up, as a prior guess, all the data older than its window. *Chapter 1.*

at-line test A test made beside the line on a sample taken from it, with a result some minutes later. *Chapter 6.*

auger A turning screw in a tube that carries pellets or powder forward at a rate set by its speed. *Chapter 2.*

AUV An autonomous underwater vehicle, which runs a planned survey under water with no crew and no cable. *Chapter 13.*

bias update A correction that moves an offset kept by a model or soft sensor by a fraction of each new measured error. *Chapter 1.*

bilateral teleoperation Remote handling in which the operator moves a master device and feels the forces that the remote slave device meets. *Chapter 8.*

calcination Heating a powder to a high temperature so that it reacts and takes its final structure, here to make a battery material. *Chapter 11.*

cascade control Two loops in series in which an outer controller sets the setpoint of an inner, faster loop. *Chapter 11.*

cathode active material The lithium compound in the positive electrode of a battery, which takes up and gives back lithium as it charges. *Chapter 11.*

chip meter A rotary feeder that sets how fast wood chips enter the digester, and so its production rate. *Chapter 4.*

chromatograph An analyzer that separates the parts of a sample as they travel through a column and measures each as it comes out. *Chapter 1.*

contribution The cash a sponsor pays into its pension fund, set in this volume as a percentage of payroll. *Chapter 9.*

covariance A measure of how far an estimate or a noise tends to wander, and of how its parts wander together. *Chapter 1.*

cross-sensitivity The response of a sensor to a gas it is not meant to read, such as ammonia read as nitrogen oxides. *Chapter 12.*

cullet Broken glass returned to the furnace with the raw materials. *Chapter 5.*

dead reckoning Finding position by adding up measured speed and heading over time, which drifts away until a fix corrects it. *Chapter 13.*

dead time The time between a change of the input and the first sign of it in the measurement. *Chapter 1.*

dead-time compensator A controller that uses a model of the delay to act on what its moves will do before the measurement shows it. *Chapter 1.*

DEF Diesel exhaust fluid, a solution of urea in water dosed into the exhaust, where it breaks down into ammonia. *Chapter 12.*

defined-benefit pension A pension that promises a benefit set by a formula, so that the sponsor must cover any shortfall of the assets. *Chapter 9.*

detection time The longest time from a change in the product to the analyzer result that first reports it. *Chapter 1.*

discount rate The interest rate that turns promised future payments into a value today. *Chapter 9.*

disturbance state An extra state of a model that stands for an unmeasured upset and is estimated from the measurements. *Chapter 1.*

diversion Sending product to waste instead of to the product drum while it may be off specification. *Chapter 7.*

Doppler velocity log A sonar that measures a vehicle’s velocity over the seabed or through the water from the frequency shift of its echoes. *Chapter 13.*

effective dead time The transport delay plus the mean age of a held result, which is the delay a tuning rule should be given. *Chapter 1.*

extended Kalman filter A Kalman filter for a

nonlinear model, linearized again at every step, abbreviated EKF. *Chapter 1*.

feedforward A move made from a measured disturbance before the disturbance shows in the controlled variable. *Chapter 4*.

float bath A bath of molten tin on which molten glass spreads into a flat ribbon of even thickness. *Chapter 5*.

funding ratio The value of a pension fund's assets divided by the present value of the benefits it has promised. *Chapter 9*.

H-factor One number that combines cooking time and temperature into a measure of how far a kraft cook has gone. *Chapter 4*.

holiday A strip of seabed that a survey meant to image but missed. *Chapter 13*.

HPLC High-performance liquid chromatography, a chromatograph for liquids that pumps the sample through a packed column. *Chapter 7*.

IAE The integral of the absolute error, the area between the controlled variable and its setpoint counted positive on both sides. *Chapter 1*.

inferential control Control of a quality that is slow or hard to measure through an estimate made from fast secondary measurements. *Chapter 1*.

integrated error The area between the controlled variable and its setpoint with its sign, abbreviated IE. *Chapter 1*.

internal model control A design that places a model of the plant inside the controller and is tuned by one filter time constant, abbreviated IMC. *Chapter 1*.

jitter Variation of a network's delay from one message to the next. *Chapter 8*.

Kalman filter An estimator that weighs a model's prediction against noisy measurements in proportion to how uncertain each is. *Chapter 1*.

kappa number A measure of the lignin left in pulp, read by an analyzer or a laboratory, and the quality a kraft cook aims at. *Chapter 4*.

kraft pulping Cooking wood chips in a hot alkaline liquor that dissolves lignin and frees the fibers for paper. *Chapter 4*.

laboratory bias update A bias update driven by laboratory results, which are slow but free of the instrument's drift. *Chapter 1*.

lehr The long annealing oven through which the glass ribbon is drawn on rollers as it cools. *Chapter 5*.

lignin The natural glue that binds wood fibers together and that pulping removes. *Chapter 4*.

loss on drying The share of a sample's weight that evaporates when it is dried, a measure of its moisture, abbreviated LOD. *Chapter 6*.

loss-in-weight feeder A feeder standing on a scale that sets its delivery rate from how fast its own weight falls. *Chapter 6*.

master-slave manipulator A remote arm whose slave end copies the motions of a master handle that the operator holds. *Chapter 8*.

mean age The average time since a held result was drawn, the reporting delay plus half the sampling interval. *Chapter 1*.

model mediation Giving the operator a local model of the remote site, so that contact feels immediate although the channel is late. *Chapter 8*.

model predictive control A controller that plans its moves over a horizon with a model and applies only the first before planning again, abbreviated MPC. *Chapter 1*.

moving horizon estimation An estimator that fits states to a window of recent measurements by optimization and can respect bounds, abbreviated MHE. *Chapter 1*.

multirate estimation Estimation that combines measurements arriving at different rates and with different delays. *Chapter 1*.

near-infrared analyzer An analyzer that reads a sample's composition from the near-infrared light it absorbs, through a calibration model, abbreviated NIR. *Chapter 6*.

nitrogen oxides Pollutant gases formed in hot combustion that a diesel catalyst must remove, written NOx. *Chapter 12*.

normal cost The value of the benefits that a pension plan's members earn in one year. *Chapter 9*.

observability Whether the measurements carry enough information to recover a state, so that no hidden combination can stay unseen. *Chapter 1*.

observer A model run beside the plant and corrected by measurements to estimate what is not measured. *Chapter 12*.

offset-free control Control that returns the measured variable exactly to its setpoint after a lasting upset, by estimating that upset. *Chapter 1*.

passivity The property of a system that never gives out more energy than it has received, which keeps contact with a wall stable. *Chapter 8*.

PI controller A controller that moves its output in proportion to the error and to the error's running total. *Chapter 1*.

plug flow Flow in which material moves as a slug without mixing, so every part spends the same time in the pipe. *Chapter 1*.

polarimeter An optical instrument that reads the stress in glass from the way the glass changes polarized light. *Chapter 5*.

position fix A position measured from outside the vehicle, such as an acoustic fix, that resets the drift of dead reckoning. *Chapter 13.*

psychrometrics The relations between the temperature, humidity and heat content of moist air. *Chapter 10.*

pyrometer A sensor that reads the temperature of a surface from the heat it radiates, without touching it. *Chapter 11.*

random walk A variable that changes by a small random step at each instant, used to model a slowly drifting bias. *Chapter 1.*

reflux ratio The ratio of the liquid returned to the top of a column to the distillate drawn off. *Chapter 3.*

residence time distribution How long the parts of a slug of material spend in a vessel or line, from first to last. *Chapter 1.*

residual alkali The alkali left in the cooking liquor at the end of a kraft cook. *Chapter 4.*

round trip The time a message takes to reach the remote side plus the time its reply takes to come back. *Chapter 8.*

sample system The tubes, filters and regulators that carry a sample from the process to an analyzer. *Chapter 1.*

sampled measurement A measurement taken at intervals and held between them, so that its value ages until the next one. *Chapter 1.*

scorecard The app panel that totals the error, overshoot, movement and limit time of each loop over a run. *Chapter 1.*

selective catalytic reduction A catalyst that turns nitrogen oxides into nitrogen and water with ammonia from injected urea, abbreviated SCR. *Chapter 12.*

SIMC A tuning rule that sets a PI controller from a first-order plus dead-time model and one chosen closed-loop time constant. *Chapter 1.*

Smith predictor A dead-time compensator that feeds back the measurement corrected by a model run with and without the delay. *Chapter 1.*

soft sensor A model that computes a slow or missing measurement from fast ones that are available. *Chapter 1.*

spray dryer A tall chamber in which a liquid sprayed into hot air dries to powder in seconds. *Chapter 10.*

steady-state Kalman filter A Kalman filter whose gain has settled to a fixed value, as when the noise statistics do not change. *Chapter 1.*

sticky point The combination of temperature and moisture beyond which milk powder turns sticky and builds up on the dryer walls. *Chapter 10.*

stored ammonia Ammonia held on the catalyst's surface, which the catalyst uses up and which no sensor measures directly. *Chapter 12.*

supervisory control A layer above the basic loops that writes their setpoints from a wider view of the plant. *Chapter 1.*

transport delay A dead time caused by material traveling from one point to another, set by a volume and a flow. *Chapter 1.*

tweel The refractory gate that sets how fast molten glass flows from the furnace onto the tin bath. *Chapter 5.*

twin-screw granulator A pair of turning screws that wets a powder blend with a liquid binder and forms it into granules as it passes. *Chapter 6.*

ultra-short baseline An acoustic system on a ship that fixes an underwater vehicle's position from the range and bearing of its replies, abbreviated USBL. *Chapter 13.*

valuation The actuary's statement of a fund's assets and liabilities on one date, published some time later. *Chapter 9.*

velocity form A digital PI or PID that computes only the change of its output at each execution. *Chapter 1.*

wave variables A recoding of force and velocity sent over a delayed channel so that the channel itself stays passive. *Chapter 8.*

wet granulation Binding a powder blend into granules with a liquid, then drying the granules before they are pressed into tablets. *Chapter 6.*

Index of Control Topics

actuator authority, 221
actuator saturation, 67, 331
actuator travel, 66
AI agents
  engineering practice, 350
  open problems, 354
  reproduction, 352
  truth audit, 351
  verification record, 351
  while studying, 18
averaging level control, 90

back-calculation, 61
back-off, 75, 94, 224
bilateral teleoperation, 349
bumpless transfer, 16, 61, 224, 250, 303, 329

cascade control, 147, 280, 329
  choice of inner measurement, 284
certainty equivalence, 224
code verification, 96, 166
composition control, 50
constraint pusher, 251
controlled variable selection, 85
controller action, 60, 146
covariance
  coverage test, 330

data export, 18
dead reckoning
  heading bias, 332
dead time, 1, 45
  achievable performance, 4, 94, 347
  apparent, 59
  fitted, 34
  sampled correction, 222
dead-time compensation, 2, 3, 348
  design procedure, 347
delayed measurements, 319, 337
disturbance rejection, 24
  lid-aware hold, 36

economic optimization, 92
external reset feedback, 251

feedforward, 115, 247, 302
  heat balance, 90
  humidity, 147
  model-based, 90
  refill, 147
  transport delay, 98, 351
feedforward control
  timing, 280
first-order plus dead-time model, 34, 279
friction
  bristle model, 195
fuel in transit, 34, 36

gain margin, 66, 201
gain scheduling, 302

human operator model, 196

infeasible setpoint, 58, 89, 117, 150
inferential control, 11, 51, 116, 238, 269
innovation gate, 328
  lockout, 335
integrator windup, 61
internal model control, 5

Kalman filter, 3, 10, 17, 118, 349
  delayed measurement, 328
  extended, 37, 319, 328
  gating, 352
  input disturbance, 61

laboratory bias update, 11, 91, 116, 250, 347
  persistence gate, 92
line-of-sight guidance, 322
linearization, 58
loop interaction, 279

material tracking, 111, 135, 154, 160, 171, 184, 272, 345
  reject gate, 147
measurement alignment, 352
measurement lag, 50
measurement noise, 55
method of characteristics, 87
modal decomposition, 58
model mismatch, 41, 61, 75, 181, 254
model order reduction, 51

model predictive control, 50, 117, 171, 224, 251
  against classical structures, 347
  disturbance assumption, 37
  in the apps, 16
  move blocking, 118, 172, 251
  nonlinear model, linearized, 36
  setpoint preview, 40, 117, 172
  soft constraint, 224
  successive linearization, 303
model-mediated teleoperation, 193, 209
move-and-wait strategy, 192, 208
moving horizon estimation, 10, 17, 62, 154, 348
multirate estimation, 7, 82

networked control, 192, 208

observability, 8, 332, 338, 347, 354
offset-free control, 37, 61, 118, 172, 253, 303
open-loop instability, 221, 324
out-of-sequence measurement, 348
overall balance, 57
override control, 170, 251

packet loss, 196
pairing, 52
passivity, 190
  energy ledger, 196
  wave variables, 193
phase margin, 66
PID control
  anti-windup, 16, 249, 278, 302
  conditional integration, 170
  velocity form, 16, 35, 60, 90, 115, 146, 169, 223, 248, 278, 302, 329
PID tuning
  internal model control, 60
  SIMC rule, 279
  Tyreus–Luyben rule, 63
plug flow, 88
  Lagrangian slugs, 85
process gain
  sign reversal, 305
  varying, 59
proportional control, 221

quantization, 222

rate limit, 222
ratio control, 146, 170
recursive least squares, 200
relay autotuning, 63, 64
residence time distribution, 6, 135, 160, 349
  tanks in series, 141, 167

sampled analyzer, 74, 88, 142, 219, 248
  detection time, 169, 184
sampled measurement, 7, 145
scenarios, 17, 173, 304
  completion rules, 17
  graded challenge, 41, 94, 119, 152, 178, 205, 225, 260, 285, 307, 334
scorecard, 17, 37, 64
secant gain, 59
selector control, 251
self-optimizing control, 85
  loss, 91
sensor ambiguity, 294, 300
sensor bias, 41, 45, 283
  compass, 325
  drift, 93, 151, 288, 349
setpoint reachability, 58
SIMC tuning rule, 34
simulation engine, 14, 56
simulator framework, 14
small-gain condition, 200
Smith predictor, 5, 25, 116
  delay error, 344
soft sensor, 11, 75, 116, 130, 142, 170, 238, 280
  bias update, 99, 170, 349
  data-driven, 349
  first-principles, 250
state estimation, 1
  open problems, 352
  projected estimate, 220, 232
state observer, 294, 302
  convergence, 306
  output injection, 302
step test, 44, 59, 171, 279
storage control, 302
supervisory control, 14, 35, 117, 280

throughput optimization, 248, 280
time delay, 190
time-domain passivity control, 193
transparency, 193
transport delay, 24, 82, 269
  shift register, 28
  variable, 6, 105, 171
trim loop, 302
two-port network, 193

ultimate gain, 63
underactuated vehicle, 324
valuation delay, 214, 217

Index of Applications

acoustic positioning, 321
active pharmaceutical ingredient, 160
actuarial valuation, 219, 231
ammonia slip catalyst, 300
ammonia storage, 299, 313
amortization, 221
analyzer, 13, 55
annealing lehr, 105
    zones, 112, 128
asset allocation, 224
at-line testing, 142
auger, 27, 43
autonomous underwater vehicle, 319

back-pressure regulator, 164
batch blanket, 108
benefit cut, 220
biomass combustion, 24
bottom track, 325
brick
    two-node model, 274
bulk density, 151
burn pot, 27

calcination kiln, 269
carbon monoxide detection, 245
catalyst aging, 307
catalyst coverage, 299
cathode active material, 269, 270
characteristic drying curve, 240
chip level, 94
chip meter, 85
circulation heater, 85, 97
circulator, 165
cold start, 305, 313
compaction, 270
compression force, 142
condenser, 53
constant molar overflow, 53
continuous annealing furnace, 271
continuous manufacturing, 135, 138, 155, 160, 163
contribution rate, 219, 232
control valve
    installed characteristic, 73
copper zeolite catalyst, 296
crab angle, 329
cullet, 108

dead reckoning, 320
decision calendar, 222
dew point, 30
diesel aftertreatment, 297
digester
    continuous, 82
discount rate, 219
distillation column, 50
divert valve, 166, 183
Doppler velocity log, 321, 322
dose ratio, 302
dosing floor, 299
dosing station, 3
drive cycle, 298
dryer fire, 241

emission limits
    Euro VI, 297
energy consumption
    spray drying, 240, 262
equilibrium thickness, 111
extraction screens, 85

feed frame, 141
flame loss, 35, 41
float bath, 106
float glass, 105, 107
flow chemistry, 160
flow reactor, 161
fluid bed, 241
fluid-bed dryer
    segmented, 141
force reflection, 191, 206
funding ratio, 215

gain and loss layers, 223
gas chromatograph, 51, 73
glass melting furnace, 108
glass transition, 241
glass viscosity, 110
glide path, 224
Gordon–Taylor equation, 245
GPS surfacing, 334

H-factor, 84
hang-up, 87
haptic interface, 192
heat-up release, 305
high-performance liquid chromatography, 161, 182
hold-up coil, 161
holiday (unimaged gap), 325, 336
hopper, 27
hot cell, 192

inertial navigation, 321

kappa number, 84, 97
kiln pressure, 276, 287
kraft pulping, 82

label claim, 139
lactose, 245
lawnmower pattern, 324
layered ordering, 272
liability duration, 219
lid opening, 34
lime kiln, 271
lithium hydroxide monohydrate, 271
loss on drying, 139
loss-in-weight feeder, 140, 153
  refill, 140, 145

master-slave manipulator, 192
Maxwell element, 111
milk powder, 238
  residual moisture, 239
Munk moment, 323

near-infrared analyzer, 142, 153
nickel-manganese-cobalt cathode, 270
normal cost, 219
NO_x conversion, 305
NO_x sensor
  cross-sensitivity, 300, 311

on-board diagnostics, 297, 313
oxidation catalyst, 297

peg-in-hole insertion, 205
pellet grill, 24, 27
pension fund, 214
  defined benefit, 214
polarimeter, 115, 128
product quality
  assay and impurities, 166
psychrometrics, 242
pull estimate, 115
pulp yield, 87
purity giveaway, 75

quality index, 276
quench, 166

reboiler, 53
reboiler duty, 53
recipe program, 39
recovery deadline, 217, 220
reflux ratio, 54
relative volatility, 54
REMUS, 321
residual alkali, 85
residual lithium, 272
residual stress, 105, 129
  temporary and permanent, 111
resistance heating, 286
resistance temperature detector, 261
roller hearth kiln, 270
rotary atomizer, 241

seabed survey, 319
selective catalytic reduction, 294, 296, 297
sonar swath, 325
sorption isotherm, 244
space teleoperation, 193
sponsor, 217, 233
spray dryer, 239
spray drying, 238, 240
stall (meat), 31
steam kit, 41
stickiness, 241, 264
survey coverage, 325, 338

tablet press, 141
tableting, 135
teleoperation, 190
  bilateral, 190, 209
Temkin desorption, 299
terrain following, 329
thermocouple, 43, 285
thickness gauge, 115
top rolls, 111
tweel, 106
twin-screw granulation, 141

ultra-short baseline, 325, 336
urea decomposition, 299
urea deposits, 299
urea dosing, 294, 312

variable-frequency drive, 262
viscous forming, 111
Vogel-Fulcher-Tammann equation, 110

wall deposit, 241
water track, 325
wet-bulb temperature, 243
white liquor, 85
width camera, 115
wood species, 87, 99

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.