INTERACTIVE CONTROL ENGINEERING

Control in Petrochemicals and Polymers

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 12

Control in Petrochemicals and Polymers

Severity, selectivity and grade from cracker to pellet

John D. Hedengren
Editor

Advanced Process Solutions, LLC

Interactive Control Engineering: Principles and Applications
Volume 12: Control in Petrochemicals and Polymers
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/v12-petrochemicals-polymers/

How to cite this volume. John D. Hedengren, Ed., *Control in Petrochemicals and Polymers*, vol. 12 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 covers the chain from olefins to polymers. Chapters 2 to 7 treat cracking, purification and the aromatics and oxide units, where severity, selectivity and runaway set the limits. Chapters 8 to 12 treat polymer reactors and finishing, where heat removal, grade transitions and product properties that are seldom measured are the difficulty.

Each chapter names the technology volume that teaches its main technique. Volume 11 covers the refinery units upstream, and Volume 13 covers other continuous chemical plants.

John D. Hedengren
2026

Contents

Preface v

Notation x

1 Introduction: Quality and Safety Limits From Cracker to Pellet 1
- 1.1 The Question This Volume Answers . . . 1
- 1.2 A Short History . . . 2
- 1.3 Notation and Conventions . . . 4
- 1.4 From Cracker to Pellet on One Chain . . . 4
- 1.5 Instruments and Final Elements in This Volume . . . 12
- 1.6 The Simulator Framework . . . 15
- 1.7 Working With AI Agents While Studying . . . 18
- 1.8 Chapter Roadmap . . . 19
- 1.9 Exercises . . . 21

2 An Ethane Cracking Furnace Simulator for Balancing Schemes and Constrained Severity 24
- 2.1 Introduction . . . 25
- 2.2 Background . . . 26
- 2.3 Process Model . . . 27
- 2.4 Dynamic Analysis . . . 30
- 2.5 Controller Design . . . 30
- 2.6 Results . . . 34
- 2.7 Technology in Practice . . . 40
- 2.8 Limitations . . . 41
- 2.9 Discussion . . . 42
- 2.10 Conclusion . . . 43
- 2.11 Exercises . . . 44

3 A One-Sided Hydrogen Window and a Selectivity Runaway in an Acetylene Converter 48
- 3.1 Introduction . . . 49
- 3.2 Acetylene in Ethylene: Catalyst, Moderator and Runaway . . . 50
- 3.3 The Converter Model . . . 51
- 3.4 Ratio, Trim, Override and a Predictive Layer . . . 56
- 3.5 Measured on the Running Converter . . . 58
- 3.6 Technology in Practice . . . 62
- 3.7 Limitations . . . 65
- 3.8 Conclusion . . . 66
- 3.9 Exercises . . . 68

4 A Propylene Splitter Simulator for Inferential Control and Quality Cascade Audit 71
- 4.1 Introduction . . . 72
- 4.2 Background . . . 73
- 4.3 Process Model . . . 74
- 4.4 Dynamic Analysis . . . 78
- 4.5 Controller Design . . . 78
- 4.6 Results . . . 81
- 4.7 Technology in Practice . . . 86

4.8 Limitations . . . 88
4.9 Discussion . . . 89
4.10 Conclusion . . . 89
4.11 Exercises . . . 91

5 Cyclic Propane Dehydrogenation Beds Whose Regeneration Sets the Next Cycle 95
5.1 Introduction . . . 96
5.2 Regenerative Fixed Beds, From Cracking to Dehydrogenation . . . 97
5.3 The Bank as Modeled . . . 98
5.4 The Control Layers as Implemented . . . 103
5.5 Nine Cases on the Running Bank . . . 104
5.6 Technology in Practice . . . 109
5.7 Limitations . . . 111
5.8 Conclusion . . . 111
5.9 Exercises . . . 113

6 An Ethylene Oxide Reactor Simulator for Selectivity, Runaway and an Optimum Without a Setpoint 117
6.1 Introduction . . . 118
6.2 Background . . . 119
6.3 Process Model . . . 120
6.4 Controller Design . . . 123
6.5 Results . . . 125
6.6 Technology in Practice . . . 130
6.7 Limitations . . . 133
6.8 Discussion . . . 133
6.9 Conclusion . . . 134
6.10 Exercises . . . 136

7 A Paraxylene Simulated Moving Bed Whose Purities Exist Only as Cycle Averages 139
7.1 Introduction . . . 140
7.2 Moving the Ports Instead of the Solid . . . 141
7.3 The Ring as Modeled . . . 142
7.4 How the Ring Responds . . . 147
7.5 Purity Control as Implemented . . . 148
7.6 Nine Cases on the Running Ring . . . 150
7.7 Technology in Practice . . . 154
7.8 Limitations . . . 156
7.9 Conclusion . . . 157
7.10 Exercises . . . 159

8 Fluidized-Bed Polyethylene Simulator for Cascade Control and Grade-Transition Economics 162
8.1 Introduction . . . 163
8.2 Background . . . 164
8.3 Reactor Model . . . 165
8.4 Dynamic Analysis . . . 174
8.5 Controller Design . . . 176
8.6 Results . . . 179
8.7 Grade Transition Case Study . . . 181
8.8 Technology in Practice . . . 185
8.9 Limitations . . . 188

8.10 Conclusion . . . 189
8.11 Exercises . . . 191

9 Polypropylene Loop Simulator to Teach Quality Control Against a Cavitation Limit 194
9.1 Introduction . . . 195
9.2 Polypropylene Process Technology . . . 196
9.3 Process Model . . . 197
9.4 The Constraint and Its Gradients . . . 205
9.5 Control Architecture . . . 206
9.6 Results . . . 210
9.7 Technology in Practice . . . 216
9.8 Limitations . . . 219
9.9 Discussion . . . 220
9.10 Conclusion . . . 220
9.11 Exercises . . . 222

10 Peak Temperature Control Against Ethylene Decomposition in a Tubular LDPE Simulator 226
10.1 Introduction . . . 227
10.2 From a High-Pressure Laboratory to the Pulsed Tube . . . 228
10.3 The Tube as Modeled . . . 229
10.4 How a Peak Forms and What Moves It . . . 235
10.5 Implemented Controller . . . 236
10.6 Nine Cases and the Challenge, as Measured . . . 240
10.7 Technology in Practice . . . 248
10.8 Limitations . . . 250
10.9 Conclusion . . . 251
10.10 Exercises . . . 253

11 A Butyl Rubber Reactor Simulator for Fouling, Actuator Margin and Wash Policy 257
11.1 Introduction . . . 258
11.2 Background . . . 259
11.3 Process Model . . . 260
11.4 Dynamic Analysis . . . 263
11.5 Controller Design . . . 264
11.6 Results . . . 267
11.7 Technology in Practice . . . 274
11.8 Limitations . . . 275
11.9 Discussion . . . 276
11.10 Conclusion . . . 277
11.11 Exercises . . . 278

12 Loss-in-Weight Additive Feeders and the Holdup Their Scales Cannot Weigh 282
12.1 Introduction . . . 283
12.2 Weighing Powder, Dosing Additives and Waiting for the Laboratory . . . 284
12.3 The Compounding Line as Modeled . . . 285
12.4 What the Weighing Boundary Implies . . . 289
12.5 Control and Estimation as Implemented . . . 290
12.6 Ten Cases, Measured . . . 294
12.7 Technology in Practice . . . 297
12.8 Limitations . . . 299
12.9 Conclusion . . . 300

12.10 Exercises 302

13 Summary and Outlook **305**
13.1 What the Chapters Showed 305
13.2 Choosing a Method 309
13.3 Where Practice Is Heading 309
13.4 Engineering With AI Agents 311
13.5 Gaps in Petrochemical and Polymer Control 314
13.6 Exercises 318

Answers to Selected Exercises **319**

Glossary **324**

Index of Control Topics **327**

Index of Applications **329**

About the Editor **331**

Volumes in This Series **332**

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: Quality and Safety Limits From Cracker to Pellet

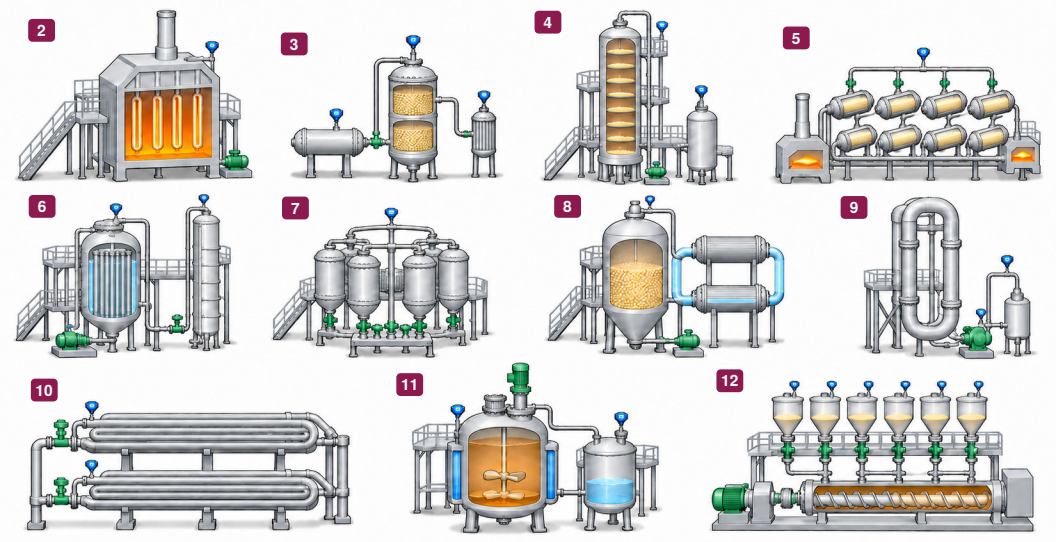

The chapters 2 cracking furnace 3 acetylene converter 4 propylene splitter 5 dehydrogenation beds 6 ethylene oxide reactor 7 paraxylene moving bed 8 gas-phase polyethylene 9 polypropylene loop 10 tubular LDPE reactor 11 butyl rubber reactor 12 additive feeders

Notice that the top row makes and purifies the olefins, the middle row converts or separates them and holds two polymer reactors, and the bottom row holds two more and the feeders that dose additives into the resin. Illustration, not to scale.

IN PLAIN TERMS **From cracked gas to plastic pellets**

The main idea. A petrochemical plant heats gas and light oil until the molecules break into small pieces such as ethylene. Other units clean the pieces and join most of them into long chains that leave as plastic pellets and rubber.

The challenge. Each unit sells a quality, like purity or how easily a plastic melts, that is checked only every few minutes or hours. Each unit also runs near a limit it must never cross, like a tube that gets too hot.

Why it matters. Good control decides how much product can be sold and keeps the plant safe.

THREE WAYS TO READ THIS CHAPTER

Explore	7 min	This page, Section 1.2 with its timeline and two History boxes, Figure 1.2, the Did you know? boxes in Sections 1.4.2, 1.4.4 and 1.4.6, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, working each number of the example P-1, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.8 with Table 1.2, Section 1.5 with Figures 1.7 and 1.10, and Section 1.8.

1.1 The Question This Volume Answers

A petrochemical complex cracks ethane, propane or naphtha into olefins, separates them to polymer grade, converts part of them into intermediates and polymerizes the rest into resins that leave the site as pellets. Each step sells a property that is measured rarely, a purity, a selectivity, a melt index, a Mooney viscosity or an additive content. Each step also runs close to a limit that ends the campaign or damages the plant if it is crossed, a tube metal temperature, a runaway, a decomposition, a cavitating

pump or a refrigeration floor. This volume answers one question: how does each unit hold a quality it seldom sees against a safety or equipment limit it must never cross, and which control techniques recur along the chain from the cracker to the pellet?

Seven control problems define the industry, and most chapters meet more than one of them.

- *Severity is bought with run length.* A hotter cracking furnace makes more product per hour and cokes its tubes faster, and the coke ends the run on a limit no loop controls.
- *Selectivity and runaway are one event.* The wanted and the unwanted reaction share a catalyst, the unwanted one releases more heat and speeds up faster with temperature, so a bed that loses selectivity also loses its margin to runaway.
- *Heat removal sets the capacity of a polymer reactor.* A polymerization releases several megajoules for each kilogram of polymer, and the reactor makes only as much as its recycle gas, its jacket or its injection pattern can carry away.
- *Grades change on a schedule.* A polymer plant makes one grade after another, and the cost of each change is the tonnage made between specifications.
- *Quality is inferred.* Analyzers report every few minutes, laboratories every few hours, and product quality is carried between them by temperatures, gas ratios, balances and models.
- *Some units run in cycles.* Regenerative beds, simulated moving beds, furnace decokes and reactor washes repeat, and they are corrected once a cycle from the outcome of the last.
- *The limit is measured in the wrong place.* The tube metal sits behind coke, a hot spot runs ahead of its thermocouple, and an impeller sees a lower pressure than the transmitter reads.

Eleven chapters answer the question on eleven plants, each with an app that the reader runs in a browser. Chapters 2 to 5 make and purify the olefins: an ethane cracking furnace, an acetylene converter, a propylene splitter and a bank of propane dehydrogenation beds. Chapters 6 and 7 are an oxidation reactor and an aromatics separation. Chapters 8 to 11 are four polymer reactors, and Chapter 12 doses the additives into the resin.

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

- Predict the run length of a fouling or coking surface from its first days and find the severity that maximizes production per elapsed hour.
- Compute how much heat the unwanted reaction makes and how close a cooled bed sits to the end of its cool operating point.
- Say how a polymer reactor removes its heat, what limits it and what spends the margin.
- Estimate the off-grade tonnage of a grade change and the preparations that reduce it.
- Correct an inferred quality from a late laboratory result, and tune a loop that acts once a cycle.
- Name, for each unit, the quality held, how often it is seen, the limit that binds and the structure that guards it.

This chapter develops each of these on one small example and names the chapter that uses it. PID forms and tuning are taken as known from Volume 2, ratio, cascade and feedforward from Volume 3, selectors and overrides from Volume 4, dead time and estimation from Volume 5, run-to-run control from Volume 8 and predictive control from Volume 9. The refinery that supplies naphtha and aromatics is Volume 11.

1.2 A Short History

The milestones of Figure 1.1 follow the chain of Figure 1.2, from the cracker and its first products to the polymer reactors, and then to the control that came to both.

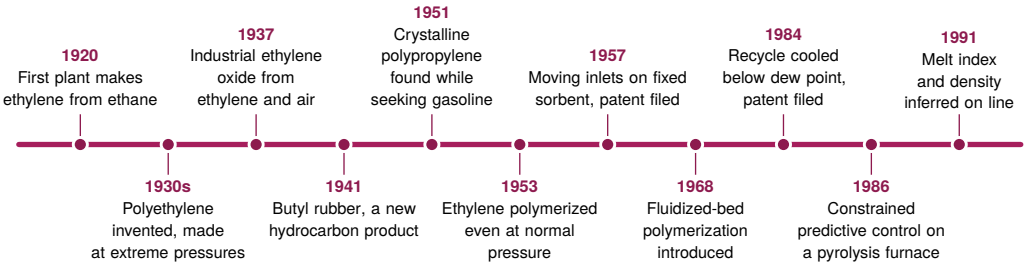

Figure 1.1 Milestones of petrochemicals, polymers and their control. The dates are from [1] (1920), [2] (1930s and 1951), [3] (1937), [4] (1941), [5] (1953), [6] (1957), [7] (1968), [8] (1984), [9] (1986) and [10] (1991). A patent is dated by its filing where the card says so. Notice that the products came first and the control of their quality came last. Events are evenly spaced, not to a time scale.

> ⧖ HISTORY **1920 to 1958: Ethylene and Its First Oxide**
>
> The first plant to make ethylene from ethane started up in January 1920 [1]. The small works-scale plant that followed in 1921 cracked the gas in three-inch tubes inside a gas-fired refractory furnace [11], a furnace of the kind Chapter 2 models. Ethylene oxide was once made by a costly two-stage route through a chlorohydrin, with a large load of effluent. Industrial production by direct oxidation of ethylene with air started in 1937, and processes on oxygen followed in 1958. The direct route replaced the old one because of the cost of chlorine and the effluent, and in 2013 the oxygen process was the only route in use in the European Union [3]. Chapter 6 runs that process in its recycle loop.

A 1941 paper announced butyl rubber as a new hydrocarbon product. Its authors had found that long chains of simple olefins such as the butylenes carry most of the properties of rubber, but that the saturated polymers cannot be vulcanized in the usual way [4]. A little diolefin such as isoprene, copolymerized with isobutylene at low temperature, gave a polymer of very low unsaturation that could still be cured with sulfur [12]. That is the recipe of Chapter 11. In 1957 a patent was filed for a separation in a fixed bed of sorbent whose inlets and outlets are moved in steps, so that the bed behaves as if the solid flowed against the liquid [6]. Chapter 7 moves its ports the same way.

> ⧖ HISTORY **1951 to 1954: Polyolefins at Low Pressure**
>
> The polyethylene of the 1930s needed pressures of 20 to 30 thousand pounds per square inch and gave a branched, low-density polymer. In 1951 J. Paul Hogan and Robert L. Banks, trying to turn propylene into gasoline over a chromium catalyst, made crystalline polypropylene instead. In less than a year they had a process for a stiffer, high-density polyethylene that needed only a few hundred pounds per square inch [2]. Near the end of 1953 Karl Ziegler's group found that ethylene polymerizes even at normal pressure on catalysts that are easy to prepare [5], and early in 1954 Giulio Natta's group polymerized propylene with them [13]. The reactors of Chapter 8 and Chapter 9 run at 21.7 and 38 bar, while the tube of Chapter 10 still runs at 2500 bar.

Polymerization in a fluidized bed of growing powder was introduced in 1968 and now accounts for about 40 percent of the world's polyethylene capacity [7]. A patent filed in 1984 cooled the recycle gas below its dew point and returned the two-phase mixture to the bed, because the polymer a reactor of a given size can make is set by the heat it can withdraw [8]. That is the condensed mode of Section 1.4.4 and of Chapter 8. Control reached these plants through models. A 1986 description of constrained predictive control closed with results from a pyrolysis furnace, where it held three stream temperatures through the fuel gas pressure of three burners [9]. In 1991 the melt index and density in an industrial fluidized-bed reactor were predicted on line from temperatures and gas compositions,

with the model corrected from infrequent laboratory results [10]. Chapter 2 and Chapter 8 build on both ideas.

1.3 Notation and Conventions

Temperatures are in °C, differences of temperature in kelvin, and absolute temperatures in a rate law in kelvin. Mass flows are in t/h and duties in MW. The coil outlet temperature of a furnace is COT and its tube metal temperature TMT. A conversion is X and a selectivity S, the share of the converted feed that becomes the wanted product. An activation energy is written as an activation temperature E/R in kelvin. A resistance to heat transfer is R_f in $\mathrm{m^2\,K/W}$, a bed or polymer inventory is M_b and its residence time τ_s. Melt index is MI, and $q = \ln \mathrm{MI}$. A cycle-to-cycle loop counts time in reports k. The framework's h is its integration step and its a the controller action, so neither is used here for anything else.

The chapters keep their own symbols, and several collide. Chapter 2 writes selectivity as σ, coke resistance as R_c and the analyzer cycle as θ. Chapter 3 uses R for the hydrogen ratio, L for the ethylene lost and a for catalyst activity, as Chapter 5 does. Chapter 8 uses L for bed level and R_H for the hydrogen ratio, while Chapter 9 writes the ratio r_H and the melt flow rate Q. The m_j of Chapter 7 are flow-rate ratios and the m of Chapter 11 is the cooling margin. Mooney viscosity is ML in Chapter 11, and the additive concentrations of Chapter 12 are in ppm by mass.

1.4 From Cracker to Pellet on One Chain

1.4.1 The Olefin, Aromatics and Polymer Chain

Figure 1.2 draws the chain and marks where each chapter's plant sits in it [14]. Ethane, propane or naphtha is cracked with steam in fired furnaces, the cracked gas is quenched and compressed, and a cold train separates it into cuts. The two-carbon cut passes an acetylene converter before ethylene is split from ethane, and the three-carbon cut goes to a propylene splitter, which also takes the propylene of a dehydrogenation plant [15]. Ethylene feeds an oxidation reactor and two polyethylene processes, propylene a polypropylene reactor, and isobutylene with a little isoprene a butyl rubber reactor. Paraxylene comes from the aromatics of a reformer and is separated by adsorption. The resins are compounded with additives into pellets.

The chapters take the plants one at a time, and the techniques they share are developed below on one running example, a small chain called P-1. P-1 cracks ethane in four identical furnaces, one of which is followed here. Part of its ethylene goes to a cooled oxidation tube and 20 t/h to a gas-phase polyethylene reactor, and a bank of regenerative beds makes propylene from propane. Every number of P-1 is declared here or computed from these declarations by the script kept with this chapter's record. The numbers are chosen for clear arithmetic and are not taken from any one chapter.

1.4.2 Furnace Severity, Coking and the Run Length

A cracking coil absorbs heat through its wall, and coke laid down inside the wall adds a resistance between the metal and the gas. With the outlet held by the firing loop, the heat flux q through each square meter is fixed by the cracking duty, and the metal runs above the gas by the flux times the resistance. Coke formation and run length are studied by experiment and model [16], and a simple form lets the coking rate grow by a factor e for every θ_c kelvin of metal temperature. Together these give Equation (1.1).

$$\mathrm{TMT} = \mathrm{COT} + q\,(R_w + R_f), \qquad \frac{\mathrm{d}R_f}{\mathrm{d}t} = \dot{R}_0\, \mathrm{e}^{(\mathrm{TMT}-\mathrm{TMT}_0)/\theta_c} \tag{1.1}$$

where R_w is the clean wall, $\dot{R}_0$ the coking rate at the start of the run and TMT_0 the metal temperature then. Each unit of coke raises the metal by q, and the hotter metal cokes faster. Integrating from a

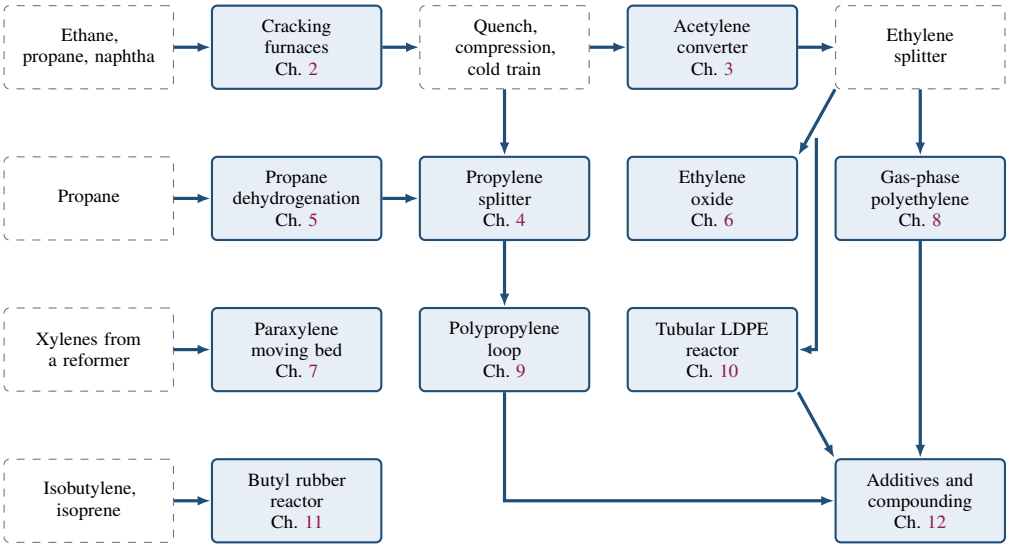

Figure 1.2 The Olefin, Aromatics and Polymer Chain. Filled Blocks Are the Plants of Chapters 2 to 12, Dashed Blocks Lie Outside the Volume. The Two-Carbon Cut Passes the Acetylene Converter, the Three-Carbon Cut Goes to the Propylene Splitter, and Ethylene From the Splitter Feeds the Oxidation Reactor and Both Polyethylene Reactors. Light Ends, Fuel Gas and Most Products Are Omitted.

clean tube gives the metal rise and the run length of Equation (1.2).

$$\Delta\mathrm{TMT}(t) = -\theta_c \ln\!\left(1 - \frac{t}{t^*}\right), \qquad t^* = \frac{\theta_c}{q\dot{R}_0}, \qquad t_{run} = t^*\left(1 - \mathrm{e}^{-\Delta_M/\theta_c}\right) \tag{1.2}$$

where Δ_M is the margin from the clean metal to its limit. The rise is a blow-up in finite time at t^*, not an exponential, and the limit is reached before it.

P-1's furnace cracks 30 t/h of ethane at a COT of 840 °C to 65 % conversion at 80 % selectivity, 15.6 t/h of ethylene. The clean metal reads 1000 °C against a limit of 1090 °C, the flux is 100 kW/m^2, $\dot{R}_0 = 2.5 \times 10^{-7}\,\mathrm{m^2\,K/(W\,h)}$ and $\theta_c = 30\,\mathrm{K}$. The metal starts by rising 0.025 K/h, so $t^* = 1200\,\mathrm{h}$ and the run lasts 1140 h, or 47.5 days, as shown in Figure 1.3a. A straight line through the first day predicts 3600 h, more than three times too long. Half the run spends only 21 % of the margin, and the last 30 K take 103 h. With a constant fouling rate, the limit $\theta_c \to \infty$ of Equation (1.2), the run is simply $\Delta_M/(q\dot{R}_0)$.

A firing loop that holds the COT hides all of this. The outlet stays on setpoint while the firing walks up, and the metal limit is a consequence of the firing, not a variable any loop sets. Chapter 2 derives Equation (1.2) for its own furnace and measures a run of 1073 h. It also shows that four passes balanced to the same outlet temperature leave 16.1 K of metal spread, because a coked pass reads cold at the outlet and hot at the skin. On a furnace whose fourth pass cokes faster, balancing on the outlets gives a run of 841 h against 888 h with no balancing at all.

Severity trades rate against run. A hotter outlet converts more and selects less, and it raises the metal and the coking rate together. The furnace must also stop to decoke, so the quantity to maximize is the production per elapsed hour of Equation (1.3).

$$Y = \frac{P\,t_{run}}{t_{run} + t_{down}} \tag{1.3}$$

where P is the production rate while running and t_{down} the time out of service. On P-1, with selectivity falling 0.8 for each unit of conversion, conversion rising 0.004 per kelvin of COT, the metal following

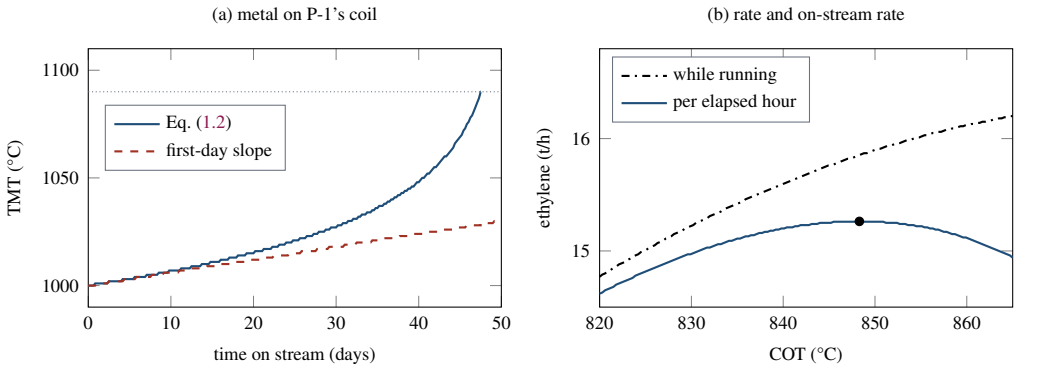

Figure 1.3 P-1's Furnace. (a) Tube Metal Over the Run (Solid), Reaching the 1090 °C Limit (Dotted) at 1140 h, Against the Straight Line Through the First Day (Dashed). (b) Ethylene While Running (Dash-Dotted) and Per Elapsed Hour With a 30 h Decoke (Solid), Whose Flat Maximum (Dot) Lies at 848.3 °C.

the outlet one for one, the coking rate growing as the 1.8 power of conversion and a decoke of 30 h, the design gives 15.20 t/h per elapsed hour. The optimum lies at 848.3 °C with 15.26 t/h and a run of 777 h, as shown in Figure 1.3b. The top is flat, 0.41 % above the design value, while the instantaneous rate keeps rising past it. Neither the conversion nor the rate can find this optimum, because both are still improving where it has passed. Only the run's total, read after the decoke, can find it.

The same arithmetic recurs. Chapter 2 measures an on-stream optimum whose top is flat between 845 and 855 °C. Chapter 11 fouls its jacket at a constant rate, reads its run of 48.9 h from the linear form, and weighs the run against an eight-hour wash with Equation (1.3). Its optimum reactor temperature is warmer than intuition expects, because a colder reactor starts closer to the refrigeration floor, which shortens the run more than the slower fouling lengthens it. Cleaning schedules against fouling are an optimization subject of their own [17]. Chapter 3 loses catalyst activity to green oil over months and Chapter 8 lets its cooler foul. In each case the controlled variable stays on setpoint and the actuator, or a margin computed from it, is the clock.

> **✦ DID YOU KNOW? The Largest-Volume Organic Chemical**
>
> Ethylene is the organic material that industry makes in the largest volume, and most of it becomes polyethylene. In 2013 the olefin plants of Europe had a capacity of about 25 million tonnes of ethylene, about a quarter of the world's [3].

1.4.3 Selectivity and Runaway in Fixed Beds

Many petrochemical reactions run beside an unwanted one on the same catalyst. Ethylene is oxidized to ethylene oxide and burned to carbon dioxide, and acetylene is hydrogenated to ethylene while ethylene is hydrogenated to ethane. When the two rates share their dependence on composition, the ratio of the unwanted to the wanted rate depends on temperature alone, as Equation (1.4) states.

$$\frac{r_2}{r_1} = \frac{r_{2,0}}{r_{1,0}} \exp\Big[-\frac{E_2 - E_1}{R}\Big(\frac{1}{T} - \frac{1}{T_0}\Big)\Big], \qquad S = \frac{1}{1 + r_2/r_1} \tag{1.4}$$

If the unwanted reaction has the larger activation energy, selectivity falls as the bed warms. If it also releases more heat, its share of the heat is larger than its share of the conversion. The heat generated then rises with temperature as one reaction with the effective activation of Equation (1.5).

$$\frac{E_{eff}}{R} = \varphi_1 \frac{E_1}{R} + \varphi_2 \frac{E_2}{R}, \qquad \varphi_2 = \frac{(1-S)\,\Delta H_2}{S\,\Delta H_1 + (1-S)\,\Delta H_2} \tag{1.5}$$

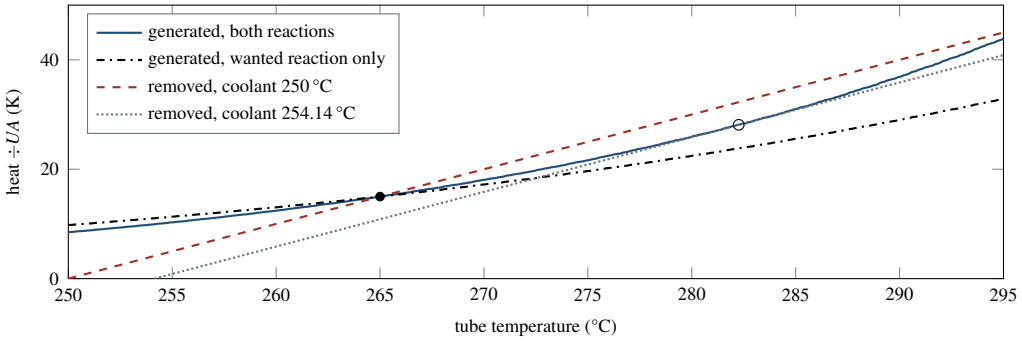

Figure 1.4 P-1's Oxidation Tube. Heat Generated by Both Reactions (Solid) and by the Wanted Reaction Alone With the Same Heat at Design (Dash-Dotted), Against Heat Removed at the Design Coolant (Dashed) and at the Coolant Where the Line Becomes Tangent (Dotted). The Dot Is the Design State and the Circle the Last Cool State.

where $\varphi_2 = 1 - \varphi_1$ is the unwanted reaction's share of the heat.

P-1's oxidation tube has $E_1/R = 8000\,\mathrm{K}$ and 100 kJ/mol for the wanted reaction and $E_2/R = 12\,000\,\mathrm{K}$ and 1300 kJ/mol for the combustion, at a selectivity of 0.85. The 15 % that burns makes 69.6 % of the heat, so $E_{eff}/R = 10\,786\,\mathrm{K}$. Each kelvin raises the ratio of the rates by 1.4 %, and 10 K hotter the selectivity is 0.832.

A cooled tube balances heat generated against heat removed to its coolant. The generation is exponential in temperature and the removal is linear, and a steady state exists only where the two curves cross. The lowest crossing is lost when the removal line becomes tangent to the generation curve, the classical construction of heat generated against heat removed [18]. With the generation rising as one reaction of activation E_{eff}, tangency requires Equation (1.6).

$$T - T_c = \frac{T^2}{E_{eff}/R} \tag{1.6}$$

A tube that runs closer to its coolant than this has a stable cool state, and the ratio of the two sides measures how much of the room is used. Parametric sensitivity studies of tubular reactors and fixed beds with axial profiles give criteria of the same kind [19–21].

P-1's tube runs at 265 °C, 15 K above a coolant at 250 °C, while Equation (1.6) allows about 27 K. A slightly warmer coolant raises the tube more each time, as Figure 1.4 shows. The cool state ends at a coolant of 254.14 °C, only 4.14 K above design, where the tube has reached 282.3 °C and its selectivity has fallen to 0.818. With the same heat at design made by the wanted reaction alone, the margin would be 13.05 K. The side reaction therefore costs two thirds of the coolant margin as well as 15 % of the feed.

Chapter 6 measures the same structure on a twelve-cell tube, where combustion makes 65 % of the heat at 87.2 % selectivity and the cool state ends 4.2 K above the design coolant. Its recycle loop stabilizes the tube, because a hotter tube burns its own oxygen, but the loop turns over in 7.7 min and the peak overshoots on the way. Its chloride moderator has an optimum, where the sign of the process gain changes and a PID started on the wrong side walks the reactor to its trip. Chapter 3 adds a reagent window open on one side [22]. Too little hydrogen puts the product off grade, too much hydrogenates ethylene, and where a bed runs out of acetylene the released ethylene reaction lights the bed off. Its bed rise first falls and reaches the trip only 3.95 min after the hydrogen step, so a temperature limit cannot see a light-off coming. The peak of Chapter 10 is a small runaway lit on purpose and limited by its initiator. Endothermic beds do not run away. The cracking coil of Chapter 2 and the dehydrogenation beds of Chapter 5 cool themselves as they react, and their conversion is limited by the heat supplied or stored. Runaway as a whole is reviewed in [23].

Table 1.1 How Each Polymer Reactor Removes Its Heat, the Limit, and What Spends the Margin

Plant	Heat leaves through	Limit	What spends the margin
P-1 reactor	recycle gas, 15 MW sensible and 5 MW latent	cooler capacity	more production, less condensing agent
Chapter 8	recycle gas, 19.85 and 5.90 MW	cooling at 88.3 % use, sticking 8.0 K away	condensing agent, low-density grades
Chapter 9	jacket on a liquid-full loop	bubble point 3.43 bar below the pressure	hydrogen at 2.44 bar per unit of ratio, inerts, 0.63 bar/K
Chapter 10	four peaks and the jackets	decomposition at 335 °C, peaks at 290	excursions, the pulse valve, a strong initiator lot
Chapter 11	jacket on a slurry	refrigeration floor −106 °C	polymer fouling the wall, reached at hour 49

1.4.4 Heat Removal in Polymer Reactors

A polymerization releases its heat in proportion to the polymer made, so a reactor's capacity is the heat it can remove, as Equation (1.7) states.

$$P_{max} = \frac{Q_{max}}{\Delta H_p}, \qquad \Delta T_{ad} = \frac{\Delta H_p}{c_p} \tag{1.7}$$

where Q_{max} is the heat the reactor can remove, ΔH_p the heat of polymerization per unit mass of polymer and ΔT_{ad} the rise the monomer would see if it converted completely without cooling. P-1 takes $\Delta H_p = 3.6\,\mathrm{MJ/kg}$, so each tonne an hour of polymer is a megawatt. For ethylene with $c_p = 2.8\,\mathrm{kJ/(kg\,K)}$ the adiabatic rise is 1286 K, or 12.9 K for each percent converted. One front that converted 25 % would heat the gas 321 K.

P-1's gas-phase reactor makes 20 t/h and so must remove 20 MW. A fluidized bed loses little heat through its wall, so its heat leaves in the recycle gas. P-1 circulates 150 kg/s at 2.0 kJ/(kg K), entering at 35 °C and leaving the bed at 85 °C, which carries 15 MW. Dry, the reactor could make only 15 t/h. The other 5 MW come from liquid condensed in the recycle cooler that boils again in the bed. At 0.35 MJ/kg that is 14.3 kg/s, 9.5 % of the recycle, and each further percent condensed adds 0.525 t/h of capacity. The gas-phase processes, their cooling and their property correlations are reviewed in [24].

Each polymer reactor of the volume removes its heat a different way, and each way has its own limit and its own consumer of the margin, as Table 1.1 lists. In Chapter 8 the condensing agent that adds cooling also lowers the temperature at which the powder sticks. The loop reactor of Chapter 9 is liquid full, and the hydrogen that sets its melt flow rate raises the bubble point of the liquid toward the pressure the circulation pump needs. The tube of Chapter 10 cannot pass its heat through the wall fast enough, so it converts in four separate peaks, each a few tens of kelvin below the temperature at which ethylene decomposes [25]. The slurry of Chapter 11 fouls its jacket, and the run ends when the jacket reaches its refrigeration floor.

Two lessons follow from the table. The variable that sets quality often spends the cooling margin, as hydrogen does in the loop reactor and the condensing agent and comonomer do in the fluidized bed. The quality loop and the limit then want the same valve in opposite directions, and Chapter 9 lets them meet at a low selector. The second lesson is that the margin is read in a computed quantity. The cooling utilization of Chapter 8, the bubble-point margin of Chapter 9 and the cooling margin of Chapter 11 are each computed from several measurements, and each chapter shows that the computed margin warns earlier than the controlled temperature.

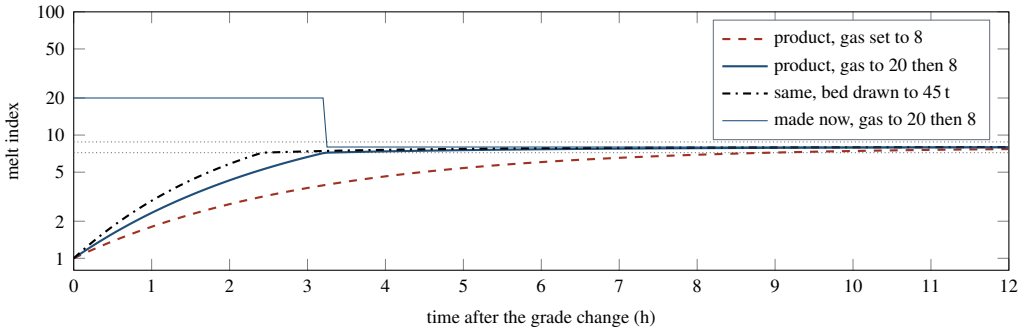

Figure 1.5 A Grade Change on P-1's Reactor From Melt Index 1 to 8. The Product Leaves the Off-Grade Region Below the Band (Dotted, 7.2 to 8.8) at 8.95 h When the Gas Is Set to the Target (Dashed), at 3.23 h When It Is First Driven to 20 (Thick Solid, the Thin Solid Line Being the Melt Index Made at Each Moment), and at 2.42 h With the Bed Also Drawn Down (Dash-Dotted).

> ✦ **DID YOU KNOW? A Reactor More Than a Kilometer Long**
>
> Commercial tubular reactors for low-density polyethylene are typically 1000 to 2500 meters long, built of high-pressure tubes 10 to 15 meters long laid in a serpentine inside a concrete bay. They convert up to 36 percent of their ethylene to polymer, against about 20 percent in an autoclave [26].

1.4.5 Grade Transitions

A polymer property is set by the conditions at the moment a chain grows, but what leaves the reactor is the average of everything the inventory holds [10]. A well-mixed inventory M_b that makes and discharges P obeys the mixing balance of Equation (1.8) for the logarithm of melt index, which averages additively in the usual correlations.

$$\tau_s \frac{\mathrm{d}\bar{q}}{\mathrm{d}t} = q_{inst} - \bar{q}, \qquad \tau_s = \frac{M_b}{P} \tag{1.8}$$

where q_{inst} is the logarithm of the melt index being made now and $\bar{q}$ that of the product discharged. If the gas composition holds q_{inst} constant from the change, the product crosses the edge q_{spec} of the new band after the time and tonnage of Equation (1.9).

$$t_{off} = \tau_s \ln \frac{q_{inst} - q_0}{q_{inst} - q_{spec}}, \qquad W_{off} = P t_{off} = M_b \ln \frac{q_{inst} - q_0}{q_{inst} - q_{spec}} \tag{1.9}$$

where q_0 is the starting value. The production rate cancels. A slower reactor makes less per hour and takes longer by the same factor, so the off-grade tonnage depends only on the inventory and on how far past the target the instantaneous property is driven.

P-1's reactor holds 60 t of powder, so $\tau_s = 3$ h. A change from melt index 1 to 8, with a band of 10 %, must bring the product above 7.2. Setting the gas to make melt index 8 at once leaves 8.95 h and 179 t off grade. Driving the gas to make melt index 20 until the product reaches 7.2, and then to 8, cuts this to 3.23 h and 64.5 t, as shown in Figure 1.5. At half the production rate the same overshoot still costs 64.5 t, over 6.46 h. Drawing the bed down to 45 t before the change cuts it to 48.4 t. Optimal grade transitions have been computed for gas-phase and multistage polyolefin reactors [27, 28].

Chapter 8 derives Equation (1.9), tests it against ten runs of one transition and finds 0.77 to 0.84 t of off-grade resin per tonne of bed against the 0.72 the logarithm gives, because its hydrogen inventory takes 142 min to turn over and the instantaneous property cannot step. Retuning its quality cascade with a drawdown cuts the off-grade tonnage from 179.1 to 77.5 t. The loop reactor of Chapter 9 averages over 52.4 min, and its transition is limited by the bubble-point margin the overshoot in

hydrogen spends. Chapter 10 changes grade through the chain-transfer agent and its peak setpoints, and Chapter 12 times each recipe switch 240 s ahead of the change, the residence of its blender, extruder and transport.

1.4.6 Inferred Product Quality

Few qualities of the volume are measured continuously. A chromatograph reports every few minutes and a laboratory every few hours, so each unit infers its quality between reports from what it does measure, a tray temperature, a gas ratio, a bed temperature or a weighed flow [29, 30]. Three practices recur.

A tray temperature is a bubble point, so it reads pressure as well as composition. The compensation of Equation (1.10) removes the pressure term arithmetically.

$$T_{comp} = T - \frac{\mathrm{d}T}{\mathrm{d}P}\left(P - P_{ref}\right) \tag{1.10}$$

In the propylene splitter of Chapter 4 the composition profile spans 5.80 K over 151 stages, 0.038 K a stage, while pressure moves every tray 3.10 K per bar. A tenth of a bar therefore reads like eight stages of the profile. The same chapter shows that compensating the measurement does not compensate the process, because pressure also changes the relative volatility.

> ✦ DID YOU KNOW? **Purifying Olefins Takes a Share of the World's Energy**
>
> A 2020 review reports that purifying propylene and ethylene accounts for about 0.3 percent of the world's energy use, and it counts the task among seven chemical separations that could change the world [31]. The column of Chapter 4 separates propylene from propane over 151 stages.

A model carries the quality between laboratory results, and each result corrects the model's bias as in Equation (1.11).

$$b_k = b_{k-1} + \lambda\left(y_{lab,k} - \hat{y}(t_{s,k}) - b_{k-1}\right) \tag{1.11}$$

where $\hat{y}(t_{s,k})$ is the model's value at the instant the sample was drawn, not when the result returned, and $0 < \lambda \leq 1$. For a polymer the model is the instantaneous correlation passed through the inventory lag of Equation (1.8). P-1's laboratory draws a sample every 2 h and reports 1 h later, so a result is 1 to 3 h old when it is used. During the direct grade change of Section 1.4.5, a result drawn at 1 h and compared with the model at 2 h would put 0.42 into the bias of ln MI, a factor of 1.53 in melt index, although model and plant agree. Chapter 12 compares each additive assay with its prediction averaged over the sample's own collection interval, and the density estimate of Chapter 10 moves the laboratory value by the change in the mean peak since the sample was drawn.

An inference is only as good as the measurement that corrects it. A frozen chromatograph republishes its last result, and a cascade that trusts it integrates an error nobody is measuring. In Chapter 4 that cascade ends worse than the tray loop alone until a validity hold is armed. When the inlet analyzer of Chapter 3 fails, a hydrogen balance on the outlet analyzer infers the feed acetylene. The weighed flow of Chapter 12 cannot see powder held below the scale, so its loop holds 1.00 kg/h while the product receives 0.600 kg/h, and only the laboratory finds it. Quality control of polymer plants has long been organized around such delayed and inferred measurements [32].

1.4.7 Cyclic Units

Two kinds of cycle appear in the volume. The slow ones are campaigns ended by a decoke, a wash or a catalyst regeneration, and Equation (1.3) prices them. The fast ones are part of the process. The dehydrogenation beds of Chapter 5 store heat during regeneration and spend it on the endothermic reaction, a bed of 36.3 MJ/K cooling 65 K in nine minutes on line. The simulated moving bed of Chapter 7 steps its ports every 90 s, and its purities exist only as twelve-minute averages. In both, a quality is reported once a cycle, and a move of the handle shows only in later reports [33].

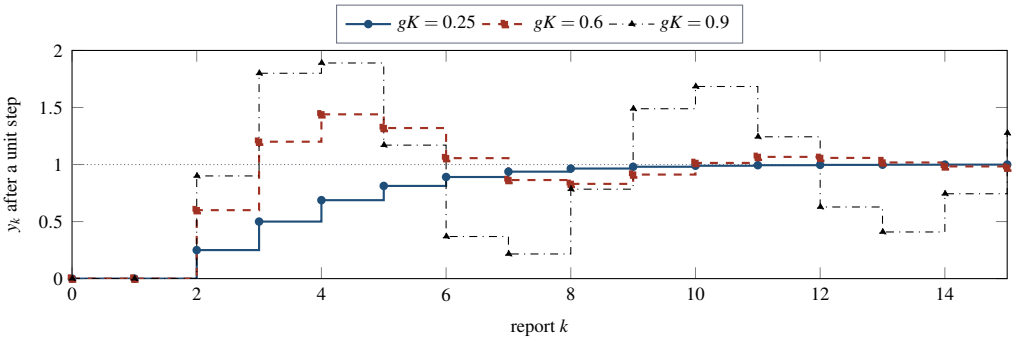

Figure 1.6 A Cycle-to-Cycle Integral Loop on a Plant That Answers Two Reports Late, Equation (1.12) With $d = 2$, After a Unit Setpoint Step: the Double Root (Solid, Dots), a Faster Setting (Dashed, Squares) and One Near the Edge of Stability (Dash-Dotted, Triangles).

A loop on such a unit acts once a report. If the reported quality y_k answers the handle u with gain g after d reports, an integral law that moves the handle by K times each report's error gives the closed loop of Equation (1.12).

$$y_k = g\,u_{k-d}, \qquad u_k = u_{k-1} + K\,(y_{SP} - y_k), \qquad z^d - z^{d-1} + gK = 0 \tag{1.12}$$

For $d = 1$ the loop settles in one report at $gK = 1$. For the delay of two reports that both chapters meet, the characteristic equation is $z^2 - z + gK = 0$. Its roots are a double root at 0.5 for $gK = 0.25$, the fastest setting without overshoot, and lie on the unit circle at $gK = 1$.

P-1's dehydrogenation bank reports its conversion once a cycle and answers the regeneration air by 0.3 points per kelvin two reports later. The integral gain is therefore 0.83 K per point for the double root and 3.33 K per point at the edge of stability. As Figure 1.6 shows, $gK = 0.25$ reaches 5 % of the setpoint at the eighth report, $gK = 0.6$ overshoots 44 % and settles by the thirteenth, and $gK = 0.9$ is still swinging.

The real units add a tail to the pure delay. A regeneration move in Chapter 5 reaches the first bed two slots late, at 0.293 points per kelvin, and a smaller second rise arrives a cycle later as the beds start their next regeneration warmer. Its cycle-to-cycle PI is therefore tuned slower than the pure delay allows, and a predictive controller that knows the delay and the measured feed holds the target sooner. Chapter 7 runs two such loops, on the logarithm of each impurity, with an integral time of two cycles, and draws the region of flow ratios inside which both outlets can be pure, the triangle of equilibrium theory [34]. Its real unit needs a ratio well inside the ideal edge, 0.568 against 0.422, because the beds disperse what the ideal theory sharpens.

1.4.8 Rarely Measured Quality Against Safety Limits

Every chapter of the volume holds a quality that it sees seldom, and every chapter runs near a limit that it cannot afford to cross. The two differ in how they are known. The quality arrives late, every few minutes to every few hours. The limit is measured continuously, but often in the wrong place: behind a coke layer, at a thermocouple that a moving peak has passed, or at a transmitter upstream of an impeller eye. The setpoint that the quality loop may use is what the limit leaves after the largest excursion and the error of its measurement, as in Equation (1.13).

$$\mathrm{SP}_{max} = y_{lim} - e_{max} - e_{meas} - \Delta_{res} \tag{1.13}$$

where y_{lim} is the limit, e_{max} the largest excursion of the controlled variable from its setpoint, e_{meas} how far the true value can run ahead of its measurement and Δ_{res} a declared reserve. Chapter 10 walks its peak setpoints toward exactly this value. Its gas runs up to 10.9 K ahead of every thermocouple while

Table 1.2 For Each Chapter, the Quality Held and How Often It Is Known, the Limit, and the Structure That Guards It

Ch.	Quality, and how it is known	Limit	Guard
2	Conversion on an analyzer every 21 min	Tube metal 1100 °C, behind the coke	One-way override on the COT, decoke trigger
3	Outlet acetylene under 1 ppm on a 5 min chromatograph	Bed rise trip at 70 K	Override at 55 K through a low select on the hydrogen
4	Propylene purity on a 20 min chromatograph, a tray temperature	Reflux and pressure limits	Pressure compensation, analyzer validity hold
5	Conversion on a chromatograph once a 3 min slot	Bed temperature limit	Clamp on the regeneration, a constraint in the MPC
6	Selectivity on a 30 min analyzer	Hot spot, flammable oxygen	One-way override on the production rate
7	Both purities as twelve-minute composites	Edges of the operating triangle	Ratio clamps, purity floors in the MPC
8	Melt index and density from a laboratory, 45 to 105 min old	Sticking temperature, cooling capacity	Six overrides that no switch removes
9	Melt flow rate on a 6 min analyzer behind a 52 min average	Pump cavitation	Bubble-point controller at a low selector, interlocks
10	Melt index every 6 min, density from the laboratory	Decomposition at 335 °C	High-high trip at 325 °C, a supervisor that backs off
11	Mooney viscosity from the laboratory every 2 h	Refrigeration floor	Wash on a schedule or on the cooling margin
12	Additive content from a composite every 3 h, 45 min late	Screw speed, a plugged sock	Level switch, torque diagnostics, a bounded trim

the pressure recovers from a pulse, and its supervisor keeps a reserve of 20 K to the decomposition line.

On P-1's tube the coolant may rise only 4.14 K before the cool state is lost. A coolant loop that swings 1 K leaves 3.14 K, and 1 K of reserve leaves a coolant ceiling of 252.14 °C. A selectivity loop on a half-hourly analyzer that asks for warmer coolant must be stopped well before that, and it learns of its mistake only at the next report. That is why the guards of the volume are fast and act in one direction only. The overrides of Chapter 2 and Chapter 6 can only lower the outlet setpoint and the production rate. Chapter 3 and Chapter 9 take the reagent away from the quality loop through a low selector, and the loop that loses is told what was applied so that it does not wind up.

Table 1.2 lists, for each chapter, the quality held, how and how often it is known, the limit that binds and the structure that guards it. Chapter 13 adds the measured results and asks how each unit holds a quality that is rarely measured against its safety limits.

1.5 Instruments and Final Elements in This Volume

Every loop in this volume begins at a sensor and ends at a final element, the device that changes a flow of material or of energy. Figure 1.7 lists what the controllers of each chapter read and what they move. A temperature is read in ten of the eleven chapters and an online analyzer in nine, and seven chapters act on a quality that is inferred or estimated rather than measured. A heating or cooling duty is moved in nine chapters, a valve or flow loop in seven and a feed of catalyst, initiator or additive in six.

Temperature carries most of the limits of the volume. A thermocouple is a junction of two different metals that behaves like a temperature-sensitive battery [35], as drawn in Figure 1.8. It sits in a thermowell, whose mass adds a first-order lag, and a lag severe enough can compromise the stability of a feedback loop [35]. The gas of Chapter 10 runs up to 10.9 K ahead of every thermocouple while the pressure recovers from a pulse, and Chapter 2 is the only chapter that reads a tube metal.

The analyzers set the pace of quality control. A chromatograph separates the components of a sample in time, as shown in Figure 1.9, so cycles of minutes are not uncommon and some dead time

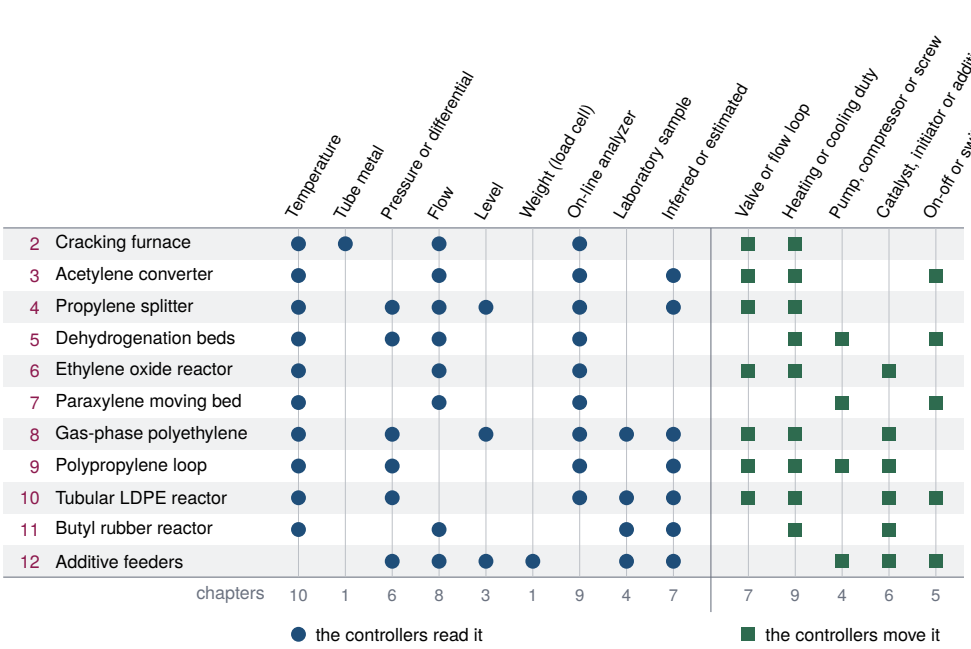

Figure 1.7 Sensors and final elements by chapter. A round mark means that the chapter's controllers or operator read the device, a square mark that they move it. A mark is set only where the chapter's variables table, its instrument text or its section on technology in practice names the device or its reading, so a blank says that the model leaves the device out and not that a real plant lacks it. Notice that only one chapter reads a tube metal and only one weighs its product.

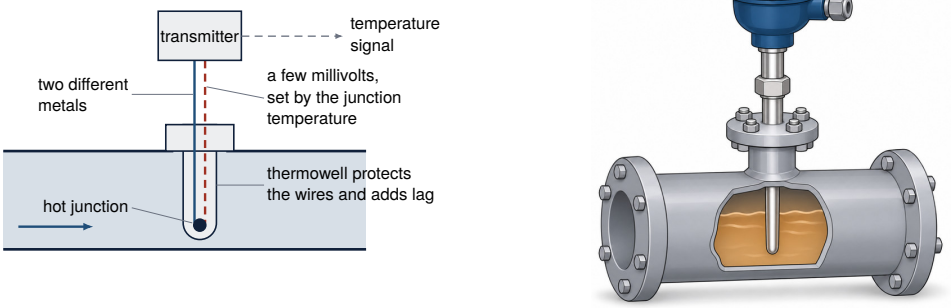

Figure 1.8 Principle of a thermocouple in a thermowell. Notice the metal between the process and the junction. It lets the element be replaced without opening the process, and it delays the reading [35], which is why a moving peak or a hot spot can run ahead of it.

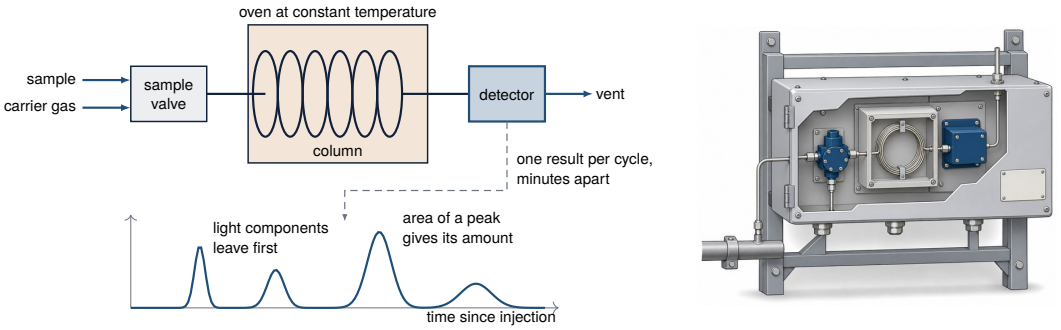

Figure 1.9 Principle of a process gas chromatograph, with a view of the device as installed. Notice that one sample gives one result at the end of the cycle, so a loop on this instrument acts on a value that is always old.

Layer	What it does in this volume	Where it runs	How often
Optimization	picks severity, grade sequence and the time to decoke or wash	**Supervisory computers** send setpoints to the controllers below	every hour for a unit, every day for the plant
Predictive control	moves several handles together against several limits		every minute; 2 s to 0.1 h in the apps
Quality and cycle to cycle control	trims setpoints from analyzers, the laboratory and inferentials	**Distributed control system** field controllers with PID, lead-lag and select blocks; trips in a separate safety instrumented system	at each analyzer, laboratory or cycle report
Overrides, selectors and trips	one-way overrides, low selectors, interlocks and trips		at each scan of the loops below
Regulatory control	holds flows, pressures, levels and temperatures		every second; 1 s to 72 s in the apps
Field devices	thermocouples, analyzers, load cells, valves, pumps and feeders	point-to-point wiring or a fieldbus	continuously

Figure 1.10 The control layers of this volume, the platform that hosts each and how often each acts. The periods of the regulatory, predictive and optimization layers are those of [9, 37], the platforms those of [9, 36], and the ranges in the apps those of the chapters. Notice that each layer sends setpoints to the layer below and acts less often than it.

is inevitable [35]. The chromatographs of Chapter 3 report 3 min after a sample drawn every 5 min, and that of Chapter 4 reports every 20 min. Four chapters wait for a laboratory. The laboratory of Chapter 8 posts each result 45 min after its hourly sample, and the Mooney viscosity of Chapter 11 arrives every 2 h.

The final elements are mostly valves on coolant, fuel and feed, and the polymer reactors add the feeders of catalyst and initiator. Where a limit must never be crossed, a trip acts beside the loops. A safety instrumented system is built of sensors, logic solvers and final elements that bring the plant to a safe state, and it is often independent of the basic process control system [36]. Figure 1.10 places the layers of the volume on their platforms. Regulatory loops act about every second, predictive control every minute and optimization every hour to every day [9, 37]. The apps run their loops every 1 s in Chapter 10 and Chapter 12 and every 72 s in Chapter 2 and Chapter 11, whose predictive controllers plan every 0.1 h.

Table 1.3 What Every App Provides. Standalone Apps Are Described in Their Own Chapters.

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

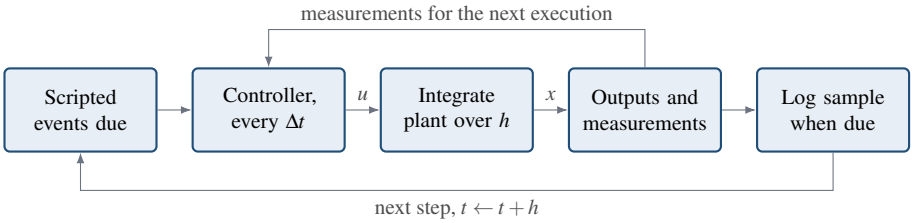

Figure 1.11 Work of the Engine in One Integration Step

1.6 The Simulator Framework

Every chapter of this book has an app that runs in a web browser, collected with the others at https://apmonitor.com/apps/. About 240 apps exist across the series. All but a few are built on one shared framework, so the engine, the PID algorithm, the trends and the data export behave the same way in every chapter. The exceptions are standalone studios, such as the cascade studio and the two Temperature Control Lab apps, which follow the same ideas in their own code. This section describes the framework once and states where the standalone apps differ. Table 1.3 lists what a reader meets.

1.6.1 The Simulation Engine

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

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

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

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

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

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

1.6.2 The Velocity-Form PID Algorithm

One PID implementation serves the PID loops of every framework app. It is the velocity (incremental) form: each execution adds the change in the proportional and derivative terms to a running sum, and the integral is a separate term, as shown in Equations (1.14) to (1.17) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.14}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.15}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.16}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.17}$$

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 (1.17) drives I toward the value that reproduces the output applied, which is back-calculation [38]. With no integral action, Equation (1.17) 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 (1.18).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.18}$$

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 (1.19) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q\,(y_k - r_k) + \sum_{j} \Delta u_j^{\top} R\,\Delta u_j + \sum_{i} \rho_i\, s_i^2 \tag{1.19}$$

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 [39]. 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 [40]. 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 [41].

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 13.

Table 1.4 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Ethane cracking furnace	Balancing the measured variable does not balance the constraint, and the run length is the objective.	intermediate	2.5 h
3	Acetylene converter	A reagent window open on one side, and a light-off a thermocouple cannot see.	advanced	2 h
4	Propylene splitter	A tray temperature is a pressure gauge first, and the analyzer is fast enough.	intermediate	2.5 h
5	Propane dehydrogenation beds	The handle acts on the next cycle's stored heat.	intermediate	2 h
6	Ethylene oxide reactor	The minority reaction owns the heat, and an optimum has no sign.	advanced	2.5 h
7	Paraxylene moving bed	Purities exist only as cycle averages.	advanced	2 h
8	Gas-phase polyethylene	Two time scales two decades apart, and grade changes priced in tonnes.	intermediate	3 h
9	Polypropylene loop	The quality handle spends the pump's margin, so a selector arbitrates.	intermediate	3 h
10	Tubular LDPE reactor	Excursions, not setpoints, spend the margin to decomposition.	advanced	2.5 h
11	Butyl rubber reactor	A flat controlled variable and an actuator counting down.	introductory	2.5 h
12	Additive feeders	A held weighed flow is not a held concentration.	introductory	2.5 h

1.8 Chapter Roadmap

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

Chapters 11 and 12 form a first course: one loop whose actuator is the clock, and a ratio loop whose product is seen only by a laboratory. They rest on Sections 1.4.2 and 1.4.6. Chapters 2, 4, 5, 8 and 9 add a hidden constraint, an inferential, a loop that acts once a cycle, a grade change and a selector. Chapters 3, 6, 7 and 10 combine a selectivity or purity that is measured late with a limit that a fast excursion can reach, and are best left until last. A reader who follows the molecules can read the volume in order, from the cracker to the pellet.

KEY POINTS

- Each unit holds a quality it sees only every few minutes or hours, against a limit it must never cross.
- A firing loop that holds the outlet hides the coking, and only the run's total, read after the decoke, finds the severity optimum.
- A side reaction that releases more heat costs P-1's tube two thirds of its coolant margin.
- The off-grade tonnage of a grade change depends on the inventory and on how far past the target the property being made is driven.
- The guards are fast and act in one direction, because a quality loop learns of its mistake only at the next report.

Where else this applies. Inferred quality returns in Volume 5, overrides and selectors in Volume 4, cycle-to-cycle control in Volume 8 and predictive control in Volume 9.

References

[1] M. Rosenzweig, "A petrochemical centennial approaches," Chemical Processing, August 7, 2019, chemicalprocessing.com/articles/2019/a-petrochemical-centennial-approaches, read November 2026.

[2] American Chemical Society, "Discovery of polypropylene and the development of a new high-density polyethylene," National Historic Chemical Landmark booklet, Office of Communications, Washington,

DC, 1999, designated at Phillips Petroleum Company, Bartlesville, Oklahoma, 12 November 1999; acs.org/education/whatischemistry/landmarks/polypropylene.html, read November 2026.

[3] H. Falcke, S. Holbrook, I. Clenahan, A. López Carretero, T. Sanalan, T. Brinkmann, J. Roth, B. Zerger, S. Roudier, and L. Delgado Sancho, "Best available techniques (BAT) reference document for the production of large volume organic chemicals," Publications Office of the European Union, Luxembourg, EUR 28882 EN, 2017.

[4] R. M. Thomas, I. E. Lightbown, W. J. Sparks, P. K. Frolich, and E. V. Murphree, "Butyl rubber. a new hydrocarbon product," *Rubber Chemistry and Technology*, vol. 14, no. 1, pp. 177–195, 1941.

[5] K. Ziegler, "Consequences and development of an invention," Nobel Lecture, Stockholm, 12 December 1963, 1963, nobelprize.org/uploads/2018/06/ziegler-lecture.pdf, read November 2026.

[6] D. B. Broughton and C. G. Gerhold, "Continuous sorption process employing fixed bed of sorbent and moving inlets and outlets," U.S. Patent 2,985,589, assigned to Universal Oil Products Co., 1961, filed May 22, 1957, issued May 23, 1961.

[7] Y. Liu, "Natural intelligence principle of supervisory-level control for condensed-mode polyethylene fluidized bed reactors: A proportional-integral-derivative perspective," *ACS Omega*, vol. 11, no. 39, pp. 60 064–60 072, 2026.

[8] J. M. Jenkins, III, R. L. Jones, and T. M. Jones, "Fluidized bed reaction systems," U.S. Patent 4,543,399, assigned to Union Carbide Corp., 1985, filed April 3, 1984, issued September 24, 1985.

[9] S. 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] K. B. McAuley and J. F. MacGregor, "On-line inference of polymer properties in an industrial polyethylene reactor," *AIChE Journal*, vol. 37, no. 6, pp. 825–835, 1991.

[11] R. D. Stief, *A History of Union Carbide Corporation: From the 1890s to the 1990s*. Danbury, CT: Carbide Retiree Corps, 1998, chapter 2, read in the copy at carbidecamps.net, November 2026.

[12] R. M. Thomas and W. J. Sparks, "Mixed olefinic polymerization process and product," U.S. Patent 2,356,128, filed 20 October 1939 (continuation in part of an application filed 29 December 1937), granted 22 August 1944, 1944.

[13] G. Natta, "From the stereospecific polymerization to the asymmetric autocatalytic synthesis of macromolecules," Nobel Lecture, Stockholm, 12 December 1963, 1963, nobelprize.org/uploads/2018/06/natta-lecture.pdf, read November 2026.

[14] H. Zimmermann and R. Walzl, "Ethylene," in *Ullmann's Encyclopedia of Industrial Chemistry*. Wiley-VCH, 2009.

[15] J. J. H. B. Sattler, J. Ruiz-Martinez, E. Santillan-Jimenez, and B. M. Weckhuysen, "Catalytic dehydrogenation of light alkanes on metals and metal oxides," *Chemical Reviews*, vol. 114, no. 20, pp. 10 613–10 653, 2014.

[16] M. Geerts, S. H. Symoens, P. A. Reyniers, G. B. Marin, M.-F. Reyniers, and K. M. Van Geem, "Steam cracking coke properties and their influence on furnace run length predictions: Experimental and modeling study," *Industrial & Engineering Chemistry Research*, vol. 59, no. 52, pp. 22 460–22 472, 2020.

[17] J. Tian, Y. Wang, and X. Feng, "Simultaneous optimization of flow velocity and cleaning schedule for mitigating fouling in refinery heat exchanger networks," *Energy*, vol. 109, pp. 1118–1129, 2016.

[18] C. van Heerden, "Autothermic processes," *Industrial & Engineering Chemistry*, vol. 45, no. 6, pp. 1242–1247, 1953.

[19] O. Bilous and N. R. Amundson, "Chemical reactor stability and sensitivity: II. Effect of parameters on sensitivity of empty tubular reactors," *AIChE Journal*, vol. 2, no. 1, pp. 117–126, 1956.

[20] R. J. van Welsenaere and G. F. Froment, "Parametric sensitivity and runaway in fixed bed catalytic reactors," *Chemical Engineering Science*, vol. 25, no. 10, pp. 1503–1516, 1970.

[21] M. Morbidelli and A. Varma, "Parametric sensitivity and runaway in fixed-bed catalytic reactors," *Chemical Engineering Science*, vol. 41, no. 4, pp. 1063–1071, 1986.

[22] A. N. R. Bos and K. R. Westerterp, "Mechanism and kinetics of the selective hydrogenation of ethyne and ethene," *Chemical Engineering and Processing: Process Intensification*, vol. 32, no. 1, pp. 1–7, 1993.

[23] A. Kummer and T. Varga, "What do we know already about reactor runaway? – A review," *Process Safety and Environmental Protection*, vol. 147, pp. 460–476, 2021.

[24] T. Xie, K. B. McAuley, J. C. C. Hsu, and D. W. Bacon, "Gas phase ethylene polymerization: Production processes, polymer properties, and reactor modeling," *Industrial & Engineering Chemistry Research*, vol. 33, no. 3, pp. 449–479, 1994.

[25] C. Kiparissides, G. Verros, and J. F. MacGregor, "Mathematical modeling, optimization, and quality control of high-pressure ethylene polymerization reactors," *Journal of Macromolecular Science, Part C: Polymer Reviews*, vol. 33, no. 4, pp. 437–527, 1993.

[26] European IPPC Bureau, "Reference document on best available techniques in the production of polymers," European Commission, Joint Research Centre, Institute for Prospective Technological Studies, Seville, Tech. Rep., Aug. 2007.

[27] K. B. McAuley and J. F. MacGregor, "Optimal grade transitions in a gas phase polyethylene reactor," *AIChE Journal*, vol. 38, no. 10, pp. 1564–1576, 1992.

[28] M. Takeda and W. H. Ray, "Optimal-grade transition strategies for multistage polyolefin reactors," *AIChE Journal*, vol. 45, no. 8, pp. 1776–1793, 1999.

[29] 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.

[30] 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.

[31] Y. Ren, X. Liang, H. Dou, C. Ye, Z. Guo, J. Wang, Y. Pan, H. Wu, M. D. Guiver, and Z. Jiang, "Membrane-based olefin/paraffin separations," *Advanced Science*, vol. 7, no. 19, p. 2001398, 2020.

[32] M. Ohshima and M. Tanigaki, "Quality control of polymer production processes," *Journal of Process Control*, vol. 10, no. 2–3, pp. 135–148, 2000.

[33] Y. Wang, F. Gao, and F. J. Doyle, "Survey on iterative learning control, repetitive control, and run-to-run control," *Journal of Process Control*, vol. 19, no. 10, pp. 1589–1600, 2009.

[34] M. Mazzotti, G. Storti, and M. Morbidelli, "Optimal operation of simulated moving bed units for nonlinear chromatographic separations," *Journal of Chromatography A*, vol. 769, no. 1, pp. 3–24, 1997.

[35] 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.

[36] 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.

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

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

[39] 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.

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

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

1.9 Exercises

1.1 Theory and calculation · 5 min

A cracking coil starts a run with its tube metal 75 K below the limit, and in the first days the metal rises 0.03 K/h. Its coking rate grows by a factor e for every 30 K of metal temperature, and the outlet and the flux are held, as in Section 1.4.2.

(a) Find t^* and the run length. *Checkpoint:* 1000 h and 917.9 h.

(b) Find the run a straight line through the first days predicts, and the metal rise after 500 h. *Checkpoint:* 2500 h and 20.8 K.

1.2 Concept check · 2 min

A four-pass furnace has one pass more coked than the others. The passes are balanced by trimming their flows until the four coil outlet temperatures agree. What happens to the coked pass?

(a) its metal falls to match the others, because equal outlets mean equal heat flux

(b) its flow is cut and its metal runs hottest, because coke makes it read cold at the outlet and hot at the skin

(c) its flow is raised and its metal cools, because the trim sends more feed to the cold outlet

(d) nothing changes, because balancing moves only the total feed and not its split

1.3 Theory and calculation · 3 min

P-1's furnace can run at 15.6 t/h of ethylene for a run of 1140.3 h, or hotter at 16.0 t/h for a run of 700 h. Each decoke takes 30 h. Which gives more ethylene per elapsed hour, and by how much?

1.4 Theory and calculation · 5 min

P-1's oxidation tube, with the activation temperatures and heats of Section 1.4.3, runs at 265 °C on a less selective catalyst, $S = 0.80$.

(a) Find the combustion's share of the heat and E_{eff}/R. *Checkpoint:* 0.765 and 11 059 K.

(b) Find the room $T^2/(E_{eff}/R)$ at the design temperature and compare it with the 0.85 catalyst. *Checkpoint:* 26.2 K against 26.9 K.

1.5 Concept check · 2 min

An acetylene converter adds hydrogen in ratio to the acetylene it must remove. Too little hydrogen leaves the product off grade, and every mole beyond the need hydrogenates ethylene. Where should the ratio be held?

(a) in the middle of the range between off grade and the trip, to keep equal margins

(b) well above the need, because more hydrogen is always the safe side

(c) at the need exactly, because any error either way costs the same

(d) just on the clean side of the breakthrough edge, because the two failures cost very different amounts

1.6 Theory and calculation · 4 min

A cooled tube carries one exothermic reaction with $E/R = 10\,000$ K and runs at 500 K, 20 K above its coolant.

(a) Find the room $T^2/(E/R)$ and the share of it the tube uses. *Checkpoint:* 25 K and 80 %.

(b) A side reaction raises E_{eff}/R to 12 500 K. Find the room again. *Checkpoint:* 20 K, all of it used.

1.7 Theory and calculation · 4 min

P-1's gas-phase reactor carries 15 MW in its recycle gas as sensible heat, and its condensed liquid takes 0.35 MJ/kg to boil, with $\Delta H_p = 3.6$ MJ/kg and 150 kg/s of recycle. Find the condensed flow and its share of the recycle at 24 t/h.

1.8 Theory and calculation · 6 min

A reactor holds 80 t of polymer and changes grade from melt index 2 to 6 with a band of 10 %, so the product must rise above 5.4.

(a) Find the off-grade tonnage when the gas is set to make melt index 6 at once. *Checkpoint:* 187.6 t.

(b) Find it when the gas first makes melt index 12 until the product reaches 5.4. *Checkpoint:* 64.7 t.

(c) At 16 t/h, how long is the product off grade in (b)? *Checkpoint:* 4.04 h.

1.9 Theory and calculation · 4 min

A cycle-to-cycle loop moves a regeneration temperature once a report, and the reported conversion answers by 0.4 points per kelvin two reports later. With the integral law of Equation (1.12), find the gain K for the double root and the gain at the edge of stability.

1.10 Concept check · 2 min

A propylene splitter runs a tray temperature loop under an analyzer cascade. The chromatograph freezes and keeps republishing its last result, which happens to be slightly high, and the feed then changes. What does the cascade do without a validity check?

(a) it keeps integrating the stale error and walks the tray setpoint away, ending worse than the tray loop alone

(b) it holds the tray setpoint, because a constant reading produces no error

(c) it trips the column, because the analyzer reading stops changing

(d) it tracks the feed change, because the tray loop below it sees the change

CHAPTER 2

An Ethane Cracking Furnace Simulator for Balancing Schemes and Constrained Severity

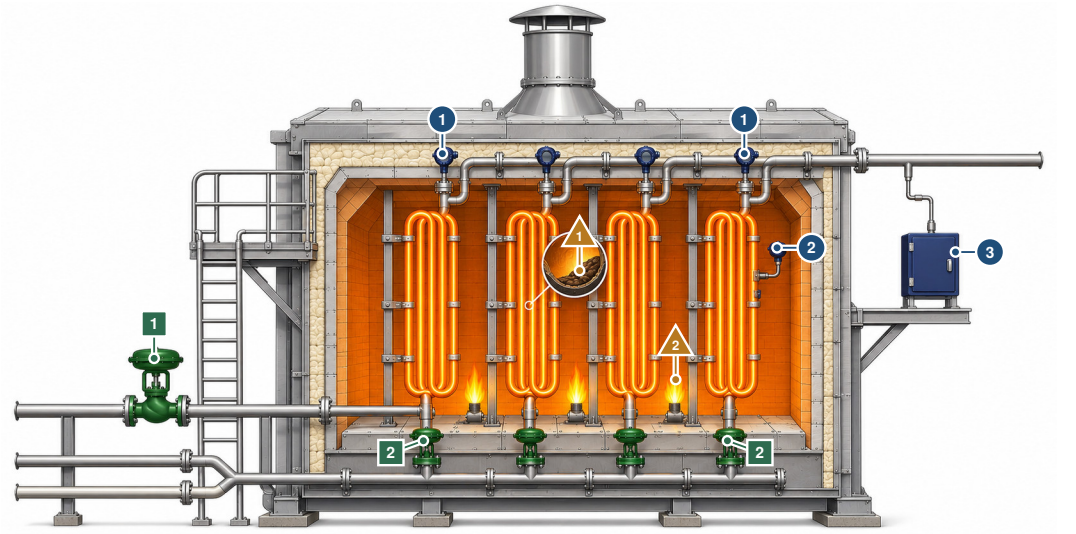

The controller moves 1 fuel valve to the burners, which sets the firebox temperature
2 four pass flow valves, which move feed between the passes

The controller reads 1 four coil outlet temperatures 2 tube metal temperature (one drawn, every pass has one)
3 conversion analyzer on the effluent

What disturbs it ▲1 coke growing inside the tubes ▲2 a degraded burner that makes one pass coke faster
▲ changes in feed rate and dilution steam

Notice that the outlet heads read the gas on one side of the coke, while the limit that ends the run is on the tube metal on the other side. Illustration, not to scale.

IN PLAIN TERMS **Four Hot Tubes and One Hidden Limit**

The main idea. A cracking furnace heats ethane gas in four steel tubes until it breaks apart into ethylene. Each tube slowly grows a crust of carbon, called coke, between the gas and the metal, and the furnace must stop now and then to burn it out.

The challenge. The usual control evens out how hot the gas leaves each tube, yet a crusted tube reads cool there while its metal runs hot. Evening out the wrong reading can end the run sooner than doing nothing.

Why it matters. The whole furnace must stop when its one hottest tube reaches the limit of its metal.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	This page, the History boxes in Sections 2.2 and 2.5.5, the Did you know? boxes in Sections 2.3 and 2.6, the Try it boxes in Sections 2.4 and 2.6, and Key points in Section 2.10. No equations.
Study	50 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 2.6 with Figure 2.2, Section 2.7, Section 2.8 with its Open challenge box, and the second Open challenge in Section 2.9.

The control question A cracking furnace comes off line when one tube reaches its metal limit, yet most pass balancing schemes equalize the coil outlet temperatures on the other side of the coke. What does balancing the measured variable do to the run length, and which measurement can find the severity that makes the most ethylene per elapsed hour?

Learning objectives

- Explain why a coked pass reads cold at its outlet and hot at its skin, and compute the metal spread that equal outlets leave.
- Derive the run length from the self-accelerating coking law and compare it with a straight line through the first day.
- Compare balancing on coil outlet, on tube metal and not at all on a furnace with one degraded burner, from measured run lengths.
- Locate the severity optimum from the on-stream rate, and explain why neither the conversion nor the ethylene rate can find it.
- Judge where a predictive controller's tube metal ceiling belongs against the override and the decoke trigger.

Prerequisites Sections 1.4.2 (coking and run length), 1.4.8 (back-off from a limit), 1.6.2 (velocity form) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/ethylene-furnace/

See also Volume 6, Chapter 4 for fouling that walks one effect toward a hard temperature limit under an override. Volume 9, Chapter 4 for an estimator that carries an unmeasured bias, and Volume 9, Chapter 8 for an optimum that sits on a constraint that moves with the weather. Volume 4, Chapter 2 for a protective override that acts only near its limit.

2.1 Introduction

A balancing scheme equalizes what it measures. On a cracking furnace that is usually the coil outlet temperature, because outlet thermocouples are cheap and reliable and every furnace has them. The constraint that ends the run is on the tube metal, on the other side of a coke layer that is growing at a different rate in every pass. Measured on the furnace modeled here, with four passes starting at coke resistances of 1.0, 2.0, 3.0 and 4.5×10^{-4} m^2 K/W and the flows trimmed until every outlet matches to within half a degree, the tube metal spread is still 16.1 °C. Twenty days later, with the outlets still matched, it is 34.0 °C. That spread is run length. The box comes off when the hottest tube reaches 1100 °C, and three of the four passes arrive there with room they never used.

The misconception the simulator corrects is that balancing a measured variable balances the process. Coke sits between the metal and the gas, so it makes the metal hotter for a given flux and the gas colder. A coked pass therefore reads low on the instrument the scheme uses and high on the instrument the constraint is written on, and the two disagree about which pass needs attention. Measured on a furnace whose fourth pass cokes 1.8 times as fast as the others because of a degraded burner, balancing on coil outlet temperature gives a run of 841 h against 888 h for no balancing at all. The scheme is working perfectly, the instrument is accurate, the tuning is correct, and the result is a shorter run than doing nothing, because the scheme is moving feed away from the pass that needed more of it.

The trap is general. Parallel heat exchangers balanced on outlet temperature rather than on approach, parallel reactors balanced on conversion rather than on bed temperature, and any set of parallel paths where the limit lives on one side of a resistance and the instrument on the other, all have it. The cracking furnace is the clearest example because the resistance grows during the run and the growth is self-accelerating: the coking rate rises by an e-fold every 30 °C of tube metal temperature, so a pass that starts slightly ahead pulls further ahead for the rest of the campaign. Coke formation and its effect on run length prediction is an active experimental and modeling subject [1, 2].

This chapter works with an interactive cracking furnace whose coking rate, radiant coefficient and wall resistance are hidden, so that the divergence must be observed, the balancing variable chosen against a measured run length, and a severity optimum found that neither the production trend nor the conversion trend can locate. The cracking kinetics, the coking correlation and override control are standard. Section 2.2 gives the domain history. Section 2.3 writes out the four pass radiant box, its coking dynamics and its instrument chain in the detail needed to rebuild it, and Section 2.4 derives the metal spread a balancing scheme cannot remove and the positive feedback that sets the run length. Section 2.5 describes the firing and severity loops, the two balancing schemes, the override, the two decoke policies and a model predictive controller that the application offers beside its loops (Section 2.5.5). Section 2.6 measures the scenario sequence, which contains a balancing result that is worse than no control and a severity optimum invisible to every continuous measurement on the unit, and compares the predictive controller with the loops on three scenarios (Section 2.6.6). Section 2.7 names the hardware and the practice behind the model. Sections 2.8, 2.9 and 2.10 give the limitations, the results that carry beyond cracking furnaces and a summary. The application exports its complete time history, as every simulator of this volume does (Section 1.6.4).

2.2 Background

Steam cracking is the largest chemical process in the world by energy consumption, and ethylene is its product. Ethane and dilution steam are heated in a fired tubular coil to about 845 °C in well under a second, and the cracking reactions that make ethylene also make coke, which deposits on the inside of the tube. Every furnace therefore runs a campaign and then comes off line for a decoke, and the campaign length is one of the numbers an ethylene plant is managed by.

> ⧖ HISTORY **1918 to 1921: Ethane Cracked in a Fired Tube**
>
> In the 1910s the chemist George O. Curme Jr. worked at an industrial research institute in Pittsburgh. During the First World War the government, looking for another source of ethylene, asked him about making it from gas oil. Curme suggested ethane instead, because it requires less energy to crack. A small plant was begun at Buffalo, New York, in the summer of 1918. It made its first ethylene in January 1920, and a fire destroyed it later that year. The work moved to a site with plenty of natural gas near Charleston, West Virginia [3]. There a small works-scale plant started in the summer of 1921, and its new tubular cracking process was demonstrated in three-inch tubes inside a gas-fired refractory furnace [4]. Cracking furnaces today still heat the feed in tubes fired from outside by gas or oil burners [5].

Three features of the unit set its control problem and each has its own literature. The first is the coking itself. Coke thickness is not measured, the rate depends on tube metal temperature, on severity and on dilution steam, and run length prediction from coke properties is still being refined against experiment [1]. The second is severity. Higher coil outlet temperature gives higher conversion and lower selectivity, so the ethylene rate and the ethylene yield move in opposite directions, and the effect of coil outlet temperature on furnace performance has been studied directly [6]. The third is the decoke itself, which takes the furnace out of service for about thirty hours and has been optimized as a process in its own right [2]. Scheduling the decokes of a bank of furnaces against a plant demand is a cyclic scheduling problem of the same family as heat exchanger cleaning scheduling [7].

Control developed from the firing. The earliest arrangement is a firebox temperature controlled by fuel gas, with the coil outlet read and corrected by an operator. Coil outlet temperature control on the firing came next and is universal now. Pass balancing arrived with individual pass flow control, and it is almost always configured on coil outlet temperature for the instrument reason given above. Tube skin thermocouples exist on most furnaces and are treated as protection rather than as control measurements, which is the habit this application is built to question. Above the regulatory layer, severity control on an effluent analyzer and constraint handling against the tube metal limit are what

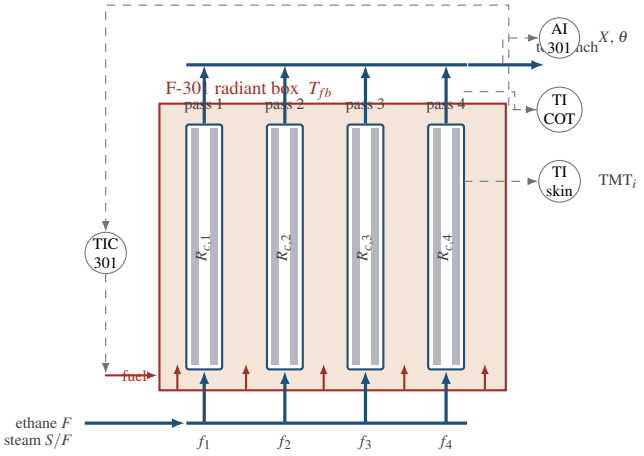

Figure 2.1 Ethane cracking furnace F-301. Ethane F with dilution steam ratio S/F is split into four pass flows f_i entering a radiant box at firebox temperature T_{fb}. Each coil carries a growing internal coke resistance $R_{c,i}$ between the gas and the tube metal, so each pass has its own coil outlet temperature and its own tube metal temperature TMT_i. AI-301 reports ethane conversion X on an analyzer cycle θ. TIC-301 manipulates the firing to hold the average coil outlet temperature, and the four pass flow trims balance the box on whichever variable is selected.

a multivariable controller provides on a modern plant [8], and the same layer now carries emissions and energy objectives as well [9].

The mathematics of what happens when a fired tube runs away is older than any of that. Parametric sensitivity in fixed bed reactors, the property that a small change in an operating parameter produces a disproportionate change in the peak temperature, was given its classical treatment by van Welsenaere and Froment in 1970 [10] and generalized later [11], and Section 1.4.3 applies it to an exothermic tube. The ethylene oxide reactor of Chapter 6 is such a tube. A cracking coil is not a fixed bed and its cracking reactions are endothermic, so it does not run away thermally. What it does instead is the subject of Section 2.6: it degrades on a timescale of weeks through a positive feedback between metal temperature and coking rate.

2.3 Process Model

Furnace F-301 cracks 40 t/h of ethane in four parallel passes of 10 t/h each, with a dilution steam ratio of 0.33 kg/kg, in a radiant box at 1150.9 °C. The coil outlet is 845.0 °C, giving 65.0 % conversion at 80.0 % selectivity for 20.80 t/h of ethylene, with 42.64 MW of absorbed duty, 106.6 kW/m^2 of flux and a tube metal temperature of 1004.9 °C against an 1100 °C limit. Figure 2.1 shows the arrangement.

> ✦ **DID YOU KNOW? The Largest-Volume Organic Chemical**
>
> Ethylene is the organic material that industry makes in the largest volume. In 2013 the plants of Europe could make about 25 million tonnes of it a year, about a quarter of the world's capacity [12]. The furnace of this chapter makes 20.80 t/h at its design point.

2.3.1 Flux, Metal Temperature and the Coke Layer

Each pass absorbs radiant flux from the box through the tube wall and the coke layer. The metal temperature is the gas temperature plus the flux times the resistance on the outside of the gas, Equation (2.1), and the flux is what the firebox can drive through the total resistance, Equation (2.2).

Table 2.1 Classification of Variables at the Start of a Clean Run

Symbol	Role	Nominal	Range
COT	CV, continuous	845.0 °C	820–880 °C SP
X (AI-301)	CV, analyzer	0.650	0.55–0.78 SP
TMT_{max}	constraint	1004.9 °C	$\leq$ 1100 °C
T_{fb}	MV	1150.9 °C	900–1400 °C
pass trims	MV	0 %	±35 %
F	DV, measured	40 t/h	25–52 t/h
S/F	DV, measured	0.33	0.20–0.60
pass skew	DV, unmeasured	1×	0.8–3×
coke factor	DV, unmeasured	1×	0.5–3×
$R_{c,i}$	states	0	growing

The first is the balance of Equation (1.1) in Chapter 1, written for one pass.

$$\mathrm{TMT}_i = \mathrm{COT}_i + q_i\,(R_w + R_{c,i}) \tag{2.1}$$

$$q_i = \frac{h_{fb}\,(T_{fb} - \mathrm{COT}_i)}{1 + h_{fb}\,(R_w + R_{c,i})} \tag{2.2}$$

where $h_{fb} = 0.73\,\mathrm{kW/(m^2\,K)}$ is a linearized radiant coefficient and $R_w = 1.50 \times 10^{-3}\,\mathrm{m^2\,K/W}$ is the clean wall and film resistance. Equation (2.2) is the closed form that results from eliminating the metal temperature. The outlet temperature of each pass then follows from its energy balance, Equation (2.3), in which the absorbed flux pays for heating the hydrocarbon and steam from the 600 °C crossover and for the endothermic cracking.

$$q_i \frac{A}{4} = \dot{m}_i\, c_p\,(\mathrm{COT}_i - T_{in}) + f_i\, X_i\, \Delta H_c \tag{2.3}$$

where $\dot{m}_i$ is the hydrocarbon plus steam flow of pass i, f_i its hydrocarbon flow, $c_p = 2.60\,\mathrm{kJ/(kg\,K)}$ and $\Delta H_c = 4600\,\mathrm{kJ/kg}$ of ethane converted. At the design point each pass absorbs 10.66 MW, 8.31 MW of it for the cracking. Both sides are monotone in COT_i, so the outlet is found by bisection, which converges whatever the coke resistance.

Equations (2.1) and (2.2) carry the lesson of the application. Coke raises $R_{c,i}$, which lowers q_i and therefore lowers the gas temperature the pass reaches, while raising the metal temperature for whatever flux does get through. A coked pass reads cold on the outlet and hot on the skin.

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

2.3.2 Coking, and Why It Accelerates

The coke resistance grows at a rate that depends on the metal temperature, on severity and on dilution steam, Equation (2.4).

$$\frac{\mathrm{d}R_{c,i}}{\mathrm{d}t} = k_c\,\psi_i\, e^{(\mathrm{TMT}_i - T_{c,ref})/\theta_c} \left(\frac{X}{X_0}\right)^n \left(\frac{S/F}{(S/F)_0}\right)^{-m} \tag{2.4}$$

with $\theta_c = 30\,^\circ\mathrm{C}$ per e-fold, $T_{c,ref} = 1004\,^\circ\mathrm{C}$, $n = 1.8$, $m = 0.55$, $(S/F)_0 = 0.33$ and X the conversion of the pass. The factor ψ_i is the furnace coke factor times, for pass 4 only, the pass skew from a degraded burner. The exponential in metal temperature closes a positive feedback: a hotter tube cokes faster, which makes it hotter still. Nothing in a standard control scheme opposes it.

Table 2.2 Model Parameters

Symbol	Meaning	Value	Class
A	radiant surface, all passes	400 m^2	adjustable
h_{fb}	linearized radiant coefficient	0.73 kW/(m^2 K)	adjustable
R_w	clean wall and film	1.50×10^{-3} m^2 K/W	adjustable
k_c	coking rate at reference	2.44×10^{-7} m^2 K/(W h)	adjustable
θ_c	degrees per e-fold in coking	30 °C	adjustable
n	coking against severity	1.8	adjustable
m	steam suppression of coking	0.55	adjustable
X_0	conversion at 845 °C	0.650	adjustable
σ_0	selectivity at X_0	0.800	adjustable
θ	severity analyzer cycle	0.35 h	adjustable
TMT$_{max}$	tube metal limit	1100 °C	visible
decoke time	out of service	30 h	adjustable

2.3.3 Cracking and Severity

Conversion rises with coil outlet temperature and selectivity falls with conversion, Equations (2.5) and (2.6).

$$X = X_0 + \frac{\mathrm{d}X}{\mathrm{d}T}\left(\mathrm{COT} - \mathrm{COT}_0\right) \tag{2.5}$$

$$\sigma = \sigma_0 + \frac{\mathrm{d}\sigma}{\mathrm{d}X}\left(X - X_0\right) \tag{2.6}$$

with $\mathrm{d}X/\mathrm{d}T = 0.00420/^\circ\mathrm{C}$ and $\mathrm{d}\sigma/\mathrm{d}X = -0.52$. The ethylene rate is $FX\sigma$, and because X and σ move in opposite directions and the coking rate goes as $X^{1.8}$, the on-stream rate has an interior optimum that neither X nor the ethylene rate can locate.

2.3.4 Assumptions

The coil is a lumped pass rather than a distributed reactor. There is no axial profile, so the coke thickness is uniform along the tube and the peak metal temperature is at the outlet, which on a real coil it usually is but not always.

Cracking is a conversion and selectivity correlation rather than a free radical mechanism, so the effect of feedstock composition on the product slate cannot be studied and the model is calibrated for ethane only.

The radiant transfer is linearized into a single coefficient. A real firebox has a temperature distribution, and the pass skew parameter is the only representation of it here.

The coke layer grows and is removed completely by the decoke. Real decoking is a controlled burn whose duration depends on how much coke is present, and it does not always return the tube to its clean condition [2].

Tube metal thermocouples are treated as accurate. On a real furnace they are notoriously difficult, they read the tube surface rather than the mid wall, and a substantial part of the argument about balancing on tube metal in industry is about whether the measurement can be trusted.

There is no downstream quench or transfer line exchanger, so the coking that occurs there and often limits a real campaign is absent.

2.4 Dynamic Analysis

2.4.1 Why Equal Outlets Do Not Give Equal Metal

Combining Equations (2.1) and (2.2) and eliminating the flux gives the metal temperature of a pass directly in terms of the firebox and its own resistance, Equation (2.7).

$$\mathrm{TMT}_i = \mathrm{COT}_i + \frac{h_{fb} R_i \left(T_{fb} - \mathrm{COT}_i\right)}{1 + h_{fb} R_i}, \qquad R_i = R_w + R_{c,i} \tag{2.7}$$

Setting every COT_i equal, which is what an outlet balancing scheme achieves, leaves the second term varying with R_i alone. Differentiating at fixed outlet temperature gives the spread a balancing scheme cannot remove, Equation (2.8).

$$\left.\frac{\partial \,\mathrm{TMT}}{\partial R}\right|_{\mathrm{COT}} = \frac{h_{fb}\left(T_{fb} - \mathrm{COT}\right)}{\left(1 + h_{fb} R\right)^2} \tag{2.8}$$

At the moment the outlets of scenario 3 first agree to within half a degree, the firebox is at 1178.9 °C and the mean resistance is 1.77×10^{-3} m^2 K/W, so the numerator is $0.73 \times 334 = 244\,\mathrm{kW/m^2}$, the denominator is $(1 + 1.29)^2 = 5.2$, and the sensitivity is about 4.6×10^4 °C per m^2 K/W. The coke spread of 3.5×10^{-4} m^2 K/W between passes 1 and 4 therefore predicts 16.4 °C of metal spread, and the running model reports 16.1 °C. The spread is not a control failure. It is the resistance between the instrument and the constraint, and no amount of trimming on the near side of it can reach.

> **TRY IT IN THE APP** **3 · Equal outlets, unequal metal**
>
> Load the case and press Run. Touch nothing. Watch the hottest and coldest pass traces on the coil outlet strip come together, within half a degree by 8.2 h. Then watch the Spread across the four tubes trace under the tube metal strip: 16.1 °C at that moment and 34.0 °C by day twenty.

2.4.2 The Positive Feedback That Sets the Run Length

Differentiating Equation (2.4) along a pass with the outlet, and therefore the conversion, held constant gives Equation (2.9).

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\mathrm{d}R_c}{\mathrm{d}t}\right) = \frac{\mathscr{S}}{\theta_c}\left(\frac{\mathrm{d}R_c}{\mathrm{d}t}\right)^2 \tag{2.9}$$

where $\mathscr{S} = \mathrm{d}\,\mathrm{TMT}/\mathrm{d}R_c$ is the metal sensitivity along the path the furnace takes. For one pass trimmed against the others at a fixed firebox it is Equation (2.8). For four passes coking together with the outlet held by the firing, Equation (2.3) fixes the flux, so $\mathrm{TMT} = \mathrm{COT} + qR$ and $\mathscr{S} = q$ exactly. With $\mathscr{S}$ constant the coking rate is $\dot{R}_0/(1 - t/t^*)$ with $t^* = \theta_c/(q\dot{R}_0)$, and the metal rises as $-\theta_c \ln(1 - t/t^*)$, which is Equation (1.2) of Section 1.4.2. At the design point $q = 106.6\,\mathrm{kW/m^2}$ and $\dot{R}_0 = 2.51 \times 10^{-7}\,\mathrm{m^2\,K/(W\,h)}$, so $t^* = 1120\,\mathrm{h}$ and the hottest tube reaches 1100 °C at 1073 h, which is the run measured in scenario 2 in Section 2.6. The coking rate grows as the square of itself, which is a finite-time blow-up rather than an exponential and is why the last week of a campaign is worse than the first month. It is also why a pass that starts a campaign slightly ahead finishes far ahead, and why a degraded burner is visible on a tube metal trend days before it is visible anywhere else. A furnace does not run away thermally in the sense of van Welsenaere and Froment [10], because the cracking reactions are endothermic. It runs away on a timescale of weeks through the coke.

Pause. Try Exercise 2.3 before reading on.

2.5 Controller Design

2.5.1 Velocity Form PID

Both PID loops use the shared velocity-form PID of Section 1.6.2, with the derivative on the measurement, executed every $\Delta t_c = 0.02\,\mathrm{h}$ (72 s). The error and the proportional and derivative terms

e_k, P_k and D_k are those of Equation (1.14), and the rest of the algorithm, in the symbols used below, is Equation (2.10).

$$\begin{aligned} s_k &= s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ c_k &= u_{bias} + s_k + I_{k-1} \\ u_k &= \text{clamp}(c_k, u_{min}, u_{max}) \\ I_k &= I_{k-1} + \left(\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{\tau_t} \right) \Delta t_c \end{aligned} \tag{2.10}$$

Here s_k is the running sum S_k of Equation (1.15), c_k the demanded output v_k of Equation (1.16), u_k the output applied and τ_t the tracking time T_t of Equation (1.17). Neither loop uses derivative action ($\tau_D = 0$), so both are PI. The bias is zero and I is seeded at reset with the initial output. Anti-windup acts on the integral alone, by back-calculation with $\tau_t = \tau_I$ [13], and s_k is never truncated, so the tube metal override can act on the severity loop through I. Manual mode sets I so that c_k equals the operator's value, and in automatic mode the manual firebox value follows the controller output after every execution. Integration is fixed-step fourth-order Runge-Kutta at 0.005 h, and the historian logs a row every 67 steps, 0.335 h.

2.5.2 Firing, Severity and Balancing

TIC-301 holds the flow-weighted average coil outlet temperature on the firebox temperature with $K_c = 3.0$ °C of firebox per °C, $\tau_I = 0.10$ h and output limits of 900 and 1400 °C. Above it, a severity loop holds the analyzer conversion by moving the coil outlet setpoint, with $K_c = 120$ °C per unit of conversion, $\tau_I = 1.5$ h and output limits of 820 and 880 °C. Both loops have $a = +1$, because a cold outlet calls for more firing and a low conversion for a hotter outlet. These are the tuned values most scenarios load. Scenario 1 and the challenge start from cautious values, $K_c = 0.8$ and $\tau_I = 0.6$ h on the firing and $K_c = 30$ and $\tau_I = 6$ h on the severity.

The balancing scheme moves the four pass flow trims to equalize either the four coil outlet temperatures or the four tube metal temperatures. It is implemented as an integral-only trim at rate $K_B = 0.0035$ per °C per hour on the deviation of each pass from the mean of the four, with each trim limited to ±35 % of the pass flow and the trims shifted to sum to zero so that the total feed is unchanged. Which variable it equalizes is the subject of Section 2.6.

2.5.3 The Tube Metal Override

A one-directional override reduces the coil outlet setpoint as the hottest tube approaches its limit, starting 12 °C below it with a gain of 2.0 °C per degree. It can only pull down, and because it is proportional on the hottest tube it switches on and off at the controller scan while the tube sits at the edge of its band. A tube run above its metal limit is replaced at the next turnaround, and a percent of conversion is worth less than that. The severity loop is not badly tuned when the override acts. It is optimizing an objective that has temporarily stopped being the objective, and a single PID has no way to know that. A protective override of the same one-directional kind guards a compressor's surge line in Volume 4, Chapter 2.

The override acts after the severity controller, so the limits of Equation (2.10) never see it. Left alone, the severity integral would see conversion short, wind up toward 880 °C, and leave the proportional pull-down fighting the top of the range. While the override acts, the severity integral therefore receives external reset. Each scan adds $(u_k^o - u_k)\Delta t_c / \tau_I$ to I_k, where u_k^o is the setpoint the override applies, which is the back-calculation of Equation (2.10) with the override's output in place of the clamp. From a decoke until the restart ramp of the next subsection ends, the severity loop tracks the outlet target it held before the decoke, then returns bumplessly.

2.5.4 Decoke Policy

The decoke is started by the operator, after a fixed number of hours (1000 h by default) or when the hottest tube reaches a threshold (1094 °C by default, six degrees under the limit). Each takes the

furnace out for 30 h. The feed is withdrawn no faster than the firebox comes down, so the tube metal does not rise as the coils empty. During the burn each pass carries 0.2 t/h, the firing holds their outlet at the 880 °C coil outlet limit (a firebox near 889 °C, below the firing loop's 900 °C floor), and the coke burns off with a time constant of 0.8 h. The feed and the outlet setpoint then return over 4 h from 880 °C to target, and the outlet peaks at 881.4 °C after a decoke at 700 h in scenario 2. The same choice between a fixed interval and a trigger on a measured margin, against a fouled jacket and a wash, is weighed in Chapter 11.

2.5.5 Model Predictive Control

A selector placed after the existing switches keeps the loops above as the default and offers a model predictive controller (MPC) as the alternative supervisory layer [8, 14]. Section 1.6.3 describes the option in general, and this subsection gives what is specific to the furnace. Under MPC, in automatic, on line and after the restart ramp, one quadratic program takes the place of both the severity loop and the balancing trim. TIC-301 remains in service beneath it, the tube metal override remains armed above it and acts on the planned setpoint exactly as it acts on the severity loop's, and the decoke policy is untouched. Every $T_s = 0.1$ h, which is five controller scans, the plan chooses the coil outlet setpoint COT_{sp} and the trims δ_1, δ_2 and δ_3 of passes 1 to 3, as fractions of the pass flow, and sets $\delta_4 = -(\delta_1 + \delta_2 + \delta_3)$, so feed moves between passes and the total never changes. It predicts $N_h = 10$ periods ahead (1 h, adjustable from 0.5 to 3 h) with move blocks of 1, 2 and 7 periods, decides moves away from the commands in force, applies the first block and holds it until the next plan. Planning begins on the call that primes the run, so a run started under MPC is planned from $t = 0$.

> ⧖ HISTORY **1986: A Pyrolysis Furnace Under Constrained Predictive Control**
>
> The first industrial predictive controllers handled constraints in a somewhat ad hoc way. The next step posed dynamic matrix control as a quadratic program, with the limits on inputs and outputs written into the problem explicitly. The method was first described in a 1983 conference paper, and a fuller description followed in 1986. It closed with results from a pyrolysis furnace, where the controller adjusted the fuel gas pressure of three burners to control the stream temperature at three places in the furnace. The tests showed the input constraints enforced as the furnace moved and the three temperatures decoupled. A later survey counts the algorithm as the second generation of the technology [8]. The plan of Section 2.5.5 solves the same kind of quadratic program every 0.1 h.

The plan's model is the pass model of Equations (2.1) to (2.3) and (2.5), evaluated with the declared parameter values of Table 2.2 and the nominal disturbances of Table 2.1 and linearized by central differences about clean tubes, four passes of 10 t/h and an outlet of 845 °C. Appended to it are the declared firebox lag of 0.035 h, pass flow lag of 0.010 h and outlet transmitter lag of 0.012 h, and TIC-301 itself, executed as in Equation (2.10) at the gain and integral time on its faceplate. Composing five scans with the input held gives Equation (2.11), the furnace's form of Equation (1.18).

$$\tilde{\xi}_{k+1} = \Phi\,\tilde{\xi}_k + \Gamma\,\tilde{v}_k, \qquad \hat{y}_k = \Omega\,\tilde{\xi}_k + \hat{b} \tag{2.11}$$

Tildes mark departures from the design point. The seven states ξ are the firebox temperature, the four pass flows, the filtered outlet and the integral of TIC-301, the inputs are $v = (\mathrm{COT}_{sp}, \delta_1, \delta_2, \delta_3)$, and $\hat{y}$ lists the predicted conversion $\hat{X}$, the flow-weighted outlet, the four tube metal temperatures $\widehat{\mathrm{TMT}}_i$ and the four departures $\hat{z}_i - \bar{\hat{z}}$ of each pass from the mean of the four, where z is the outlet or the metal temperature as the balancing switch selects. The offsets $\hat{b}$ are described below. About the design point a pass outlet gains 0.355 °C and its metal 0.692 °C per degree of firebox, and they lose 10.86 °C and 5.19 °C per t/h of the pass's own flow. Shifting one percent of the feed onto a pass therefore cools its metal by 0.52 °C, while each degree of outlet setpoint warms every tube by 1.95 °C and raises conversion by 0.0042. Stepped by ±1 °C of setpoint and ±1 % of trim on clean tubes, under the tuned and the cautious firing tuning alike, the simulated furnace and the model differ

by no more than 0.24 % of the response from 0.1 to 2 h. Step-test fits were not used, since the plan has to predict the furnace with TIC-301 closed at whatever tuning its faceplate carries and the pass model yields every slope in closed form.

Each pass outlet and tube metal temperature is read on every scan, the same readings the balancing trim uses, so the offset of each from the model is measured rather than estimated. Each entry of $\hat{b}$ is the reading less the model's present value, held constant over the horizon, and it is through these offsets that the coke enters the plan. The firebox and flow states are the model's own response to the commands already sent, and the filtered outlet and the integral are read from TIC-301. Conversion is the only late measurement. The plan computes it from the four outlets with Equation (2.5) and adds a bias $\hat{b}_X$ that each scan moves by $0.04\,(X_{AI} - X_{cp} - \hat{b}_X)$, where X_{AI} is the analyzer reading and X_{cp} a copy of the analyzer, with the same 0.35 h cycle and 0.08 h averaging on the same clock, fed with the computed conversion. The bias is frozen through a decoke and the restart ramp and removes any steady offset in the conversion law [15].

Each plan minimizes Equation (2.12), which adds a linear price on each slack to the terms of Equation (1.19).

$$\begin{aligned} J = & \sum_{k=1}^{N_h} w_X \left(\hat{X}_k - X_{sp}\right)^2 \\ & + \lambda_m \sum_{j=1}^{3} \left[\Delta\mathrm{COT}_{sp,j}^2 + \sum_{i=1}^{3} \left(100\,\Delta\delta_{i,j}\right)^2 \right] \\ & + \sum_{i=1}^{4} \left(\rho_T\,\varepsilon_i^2 + \mu_T\,\varepsilon_i + \rho_z\,\zeta_i^2 + p_z\,\zeta_i\right) \end{aligned} \tag{2.12}$$

Here $w_X = 1/0.006^2$ and X_{sp} is the conversion target. With the severity loop in manual the first term instead holds the flow-weighted outlet on the operator's setpoint at a weight of one per square degree. The move weight is $\lambda_m = 10$, Δ is the change of block j in degrees or fractions, measured from the command in force for $j = 1$, and the factor 100 puts the trims in percent. The slacks ε_i of the tube rows cost $\rho_T = 10^3$ per square degree and $\mu_T = 10^4$ per degree, an exact penalty, so the plan never predicts a crossing that its handles could avoid. The slacks ζ_i of the spread rows cost $p_z = 3$ per degree and $\rho_z = 0.01$ per square degree, and because each row has one slack for the whole horizon, p_z prices the largest excess of each pass over the mean. The rows are Equation (2.13).

$$\begin{aligned} & 820 \le \mathrm{COT}_{sp,k} \le 880, \qquad |\delta_{i,k}| \le 0.35 \\ & \widehat{\mathrm{TMT}}_{i,k} \le \overline{T} + \left[\mathrm{TMT}_{i,0} - \overline{T}\right]^+ e^{-kT_s/\tau_f} + \varepsilon_i \\ & \hat{z}_{i,k} - \bar{\hat{z}}_k \le \left[z_{i,0} - \bar{z}_0\right]^+ e^{-kT_s/\tau_f} + \zeta_i, \qquad \varepsilon_i,\ \zeta_i \ge 0 \end{aligned} \tag{2.13}$$

The first line is hard and repeats the clamps of the severity loop and the trim, with the upper setpoint bound lowered if the coil outlet limit is set below 880 °C and pass 4 entered as two rows per block. Zero trims with any setpoint in range satisfy every hard row and the other rows are soft, so a plan always exists. The ceiling is $\overline{T} = \mathrm{TMT}_{max} - \Delta_o - 1\,^\circ\mathrm{C}$, a back-off of one degree under the override's start with $\Delta_o = 12\,^\circ\mathrm{C}$, read as the override reads it, which is 1087 °C as shipped. With $[\cdot]^+$ the positive part, a row already violated starts at the present reading and closes onto its limit with $\tau_f = 0.1$ h, because a single slack per row would otherwise excuse any later violation up to the present one. With balancing off the trims stay where the operator set them and the spread rows are dropped. Replacing the closing rows by flat ones moves the severity IAE of scenario 10 from 2.176 to 2.178 and the hours over the limit in scenario 8 from 228.51 to 228.45, so none of the results below depends on them.

Five design choices need a reason. The controller is a selector of its own rather than a further value of an existing switch, since it takes the setpoint and the trims together and either switch would misname it. It moves the outlet setpoint rather than the firing, so TIC-301 stays in service, the setpoint keeps the severity loop's range and the override acts on it unchanged. Every tube has its own soft row

Table 2.3 Scenario Sequence

#	Scenario	Teaching point
1	Four tubes that do not agree	the two measurements disagree
2	The coil outlet loop	the constraint is elsewhere
3	Equal outlets, unequal metal	balancing the wrong variable
4	Balancing on the constrained variable	what it costs
5	A burner problem on one corner	worse than no control
6	The optimum neither trend can see	severity economics
7	Dilution steam	a second optimization
8	The tube metal override	objective changes
9	When to decoke	schedules are open loop
10	Balance the passes before the override	one plan, two handles
11	Challenge: sixty days	graded run

instead of one row on the hottest, because a maximum is not linear in the inputs. It decides every 0.1 h rather than every minute, because a minute is not a whole number of 72 s scans and five scans per decision cut the number of plans fivefold. Its estimator is the measured offsets with the analyzer bias rather than a Kalman filter, because every pass is read without delay, the furnace answers in minutes while the coke builds over days, and an input-disturbance Kalman filter, measured once on a separate copy of the code, moved the severity IAE of scenario 10 by at most 0.016 and the hours over the limit in scenario 8 by at most 3.4 h, changing no ranking. Where a disturbance reaches no reading directly, the filter has to carry it as a state, as the fuel bias of Volume 9, Chapter 4 is carried.

The solver's answer is refused when it is suboptimal, holds a nonfinite number or leaves a bound. The severity loop, which has tracked the planned setpoint, and the balancing trim, which has followed the planned trims, then take both handles with one execution each and no bump, the trend is marked and a note is posted, and control returns to the plan at its next acceptable solution. With the severity loop in manual the setpoint stays where the plan left it. The plan reads only what the shipped loops read, which is the pass outlet and metal readings, the outlet transmitter, the analyzer, the integral and tuning of TIC-301, the setpoints and switches, the commands applied, the tube metal limit and margin, and the coil outlet limit. Apart from those two limits, read where the override and the severity loop read them, it never reads a run parameter, the coke or a disturbance, so a plant parameter a learner changes is model error to the plan. The quadratic program is solved by a dual active-set method [16] in a mean of 165 µs per plan in Chrome on the development computer, between 163 and 167 µs over four measurements, which is 0.17 percent of one core at 1 h per second.

2.6 Results

Eleven scenarios run from watching four tubes diverge to a graded sixty-day challenge. Table 2.3 lists them, and Section 1.6.4 describes how a scenario is loaded, credited and exported. Section 2.11 sets runs from several of them.

2.6.1 The Constraint Is Not on the Manipulated Variable

With the coil outlet loop closed in scenario 2, the outlet stays within 0.01 °C of 845 °C for the first 1000 h while the firebox walks up from 1150.9 to 1217.8 °C and the hottest tube from 1004.9 to 1071.8 °C. With the override switched off, the run at this severity is 1073 h, which is 44.7 days, and it ends because of a number nobody is controlling. With the override armed as shipped, the outlet setpoint starts to come down at 1050.2 h, when the hottest tube reaches 1088 °C, but coke grows at any severity, and with nobody decoking a tube crosses 1100 °C at 1134.0 h and spends 16.8 h above it by 1200 h. The firing still has room, because the firebox peaks at 1242 °C against its 1400 °C limit.

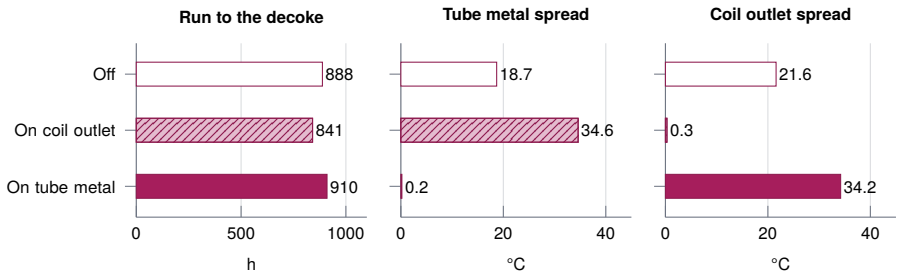

Figure 2.2 The three balancing schemes of Table 2.4 on the furnace with one degraded burner: white bars with balancing off, hatched bars balancing on coil outlet, solid bars balancing on tube metal, with the table's value beside each bar. Notice that the scheme that brings the outlets together drives the tube metal apart and gives the shortest run of the three.

Table 2.4 Balancing Schemes on a Furnace With One Degraded Burner

Balancing	Run (h)	TMT spread (°C)	COT spread (°C)
Off	888	18.7	21.6
On coil outlet	841	34.6	0.3
On tube metal	910	0.2	34.2

It is the tube metal that runs out, and tube metal is a consequence of firing rather than a thing anybody sets.

2.6.2 Balancing on the Wrong Variable

Scenario 3 starts four passes at coke resistances of 1.0, 2.0, 3.0 and 4.5×10^{-4} m^2 K/W with balancing on coil outlet temperature. The outlet spread falls from 12.5 to under 0.5 °C at 8.2 h and stays there, while the tube metal spread rises from 10.8 to 16.1 °C at the same moment, 20.9 °C by day ten and 34.0 °C by day twenty. The override first acts at 336.7 h, and pass 4 first crosses its limit at 455.3 h and spends 4.1 h above it by day twenty.

Balancing on tube metal, in scenario 4 from the same starting coke with the decoke at 1094 °C, converges the metal temperatures and separates the outlet temperatures, which is the same trade in the other direction. Just before the decoke the metal spread is 0.15 °C against 28.8 °C for outlet balancing, and the outlet spread is 16.0 °C against 0.14 °C. The run length barely moves, 396.6 h against 392.4 h, and with balancing off it is 408.8 h, the longest of the three. A one-off difference in coke is partly self-limiting, because with equal flows the coked pass runs a colder outlet, cracks less and therefore cokes more slowly, and its metal spread grows only from 10.8 to 13.2 °C over the run. Balancing on the constraint pays when something keeps pushing one pass ahead, which is the case of the next subsection. The cost is a product cost. Four passes cracking at four different outlet temperatures make four different severities, so the effluent is a blend. On an ethane cracker that is tolerable, and on a naphtha cracker, where the yield structure is far more temperature sensitive, it is the reason this is argued about.

2.6.3 Worse Than No Control

Table 2.4 reports a furnace whose fourth pass cokes 1.8 times as fast as the others because of a degraded burner (scenario 5). Each run starts from clean tubes with the override armed and ends with a decoke when the hottest tube reaches 1094 °C, and the spreads are read on line just before the decoke. Figure 2.2 draws the three rows side by side.

The middle row is the result worth reading twice. Balancing on coil outlet temperature gives a shorter run than not balancing at all, because coke makes pass 4 read cold and the scheme therefore

Table 2.5 Severity Against Instantaneous Rate and On-Stream Rate

COT (°C)	X	σ	C_2H_4 (t/h)	Run (h)	Per elapsed hour (t/h)
835	0.608	0.822	19.99	2516	19.75
845	0.650	0.800	20.80	1073	20.23
850	0.671	0.789	21.18	700	20.31
855	0.692	0.778	21.54	454	20.20
865	0.734	0.756	22.21	186	19.12

takes feed away from exactly the pass that needed more of it. Nothing is broken. The instrument is accurate, the tuning is correct and the loop is doing what it was asked.

> **TRY IT IN THE APP** **5 · A burner problem on one corner**
>
> Run the case to its decoke with Pass balancing on coil outlet and read the Run tile, 841 h. Reset, set Pass balancing to tube metal before pressing Run, and repeat, 910 h. A third run with balancing off decokes at 888 h, later than the run balanced on the outlets.

The first sign of the degraded burner is a single tube metal temperature pulling away from the other three. With outlet balancing, pass 4 is 0.8 °C hotter than pass 1 after three days, 3.0 °C after ten and 19.1 °C after thirty, while the four outlets agree to within 0.2 °C, as shown in Figure 2.3. Neither balancing scheme fixes the burner and neither is supposed to. One of them buys time to schedule the repair and the other spends it.

2.6.4 A Severity Optimum No Continuous Measurement Can Find

Table 2.5 reports five coil outlet setpoints, each held from clean tubes until the hottest tube reaches its 1100 °C limit, at a dilution steam ratio of 0.33. They are measured in scenario 2 with the override off and the decoke triggered on tube metal at 1100 °C, and the last column is the on-stream rate when the furnace comes back on line after its 30 h decoke.

The instantaneous ethylene rate rises monotonically across the whole range and so does the conversion. Neither can find the optimum, because both are still improving where the optimum has already gone past. The only number that finds it is ethylene divided by elapsed time, Equation (1.3) of Chapter 1, and that number cannot be measured until after the decoke.

Pause. Try Exercise 2.2 before reading on.

The top is also very flat, with 20.23, 20.31 and 20.20 t/h at 845, 850 and 855 °C, while the run falls from 1073 h to 454 h across the same ten degrees. Fifteen degrees above the peak the on-stream rate is 5.9 % lower and fifteen degrees below it 2.7 % lower. Optima with flat tops and steep sides should be approached from the flat side. Scenario 6 poses the same question through the severity loop, with the override armed and the decoke at 1094 °C. Conversion targets of 0.62, 0.65, 0.68 and 0.70 decoke at 2026, 1119, 622 and 420 h and give 19.91, 20.23, 20.31 and 20.18 t/h per elapsed hour once the furnace is back on line, so the optimum a learner finds there is the same flat top.

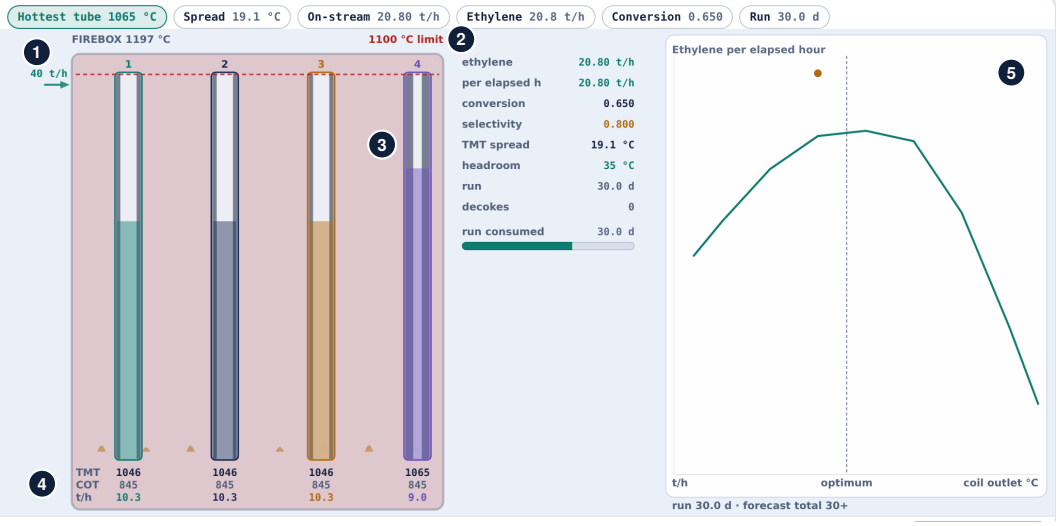

Figure 2.3 The cracking furnace application in scenario 5, the degraded burner case, with balancing on coil outlet temperature at 720 h (30.0 days), cropped to the plant view. The firebox at 1197 °C is drawn with its four coils, each showing its coke layer and its tube metal temperature as a bar against the 1100 °C limit. (1) The status chip reads the hottest tube. (2) The dashed line is the 1100 °C limit. (3) Pass 4, coking 1.8 times as fast, is the hottest tube at 1065.3 °C against 1046.2 °C for the other three. (4) The rows under the coils read each pass's tube metal, outlet and flow: all four outlets read 845 °C (spread 0.16 °C) because the balancing has cut the flow of pass 4 to 9.02 t/h against 10.33 t/h. (5) Beside the furnace readouts is the measured on-stream rate against coil outlet temperature.

> ✦ **DID YOU KNOW? Fastest While Running, Slowest Overall**
>
> From the simulator (Table 2.5). At a coil outlet of 865 °C the furnace makes ethylene faster than at any other setting in the table, 22.21 t/h. Counted per elapsed hour it makes the least, 19.12 t/h, because its run lasts only 186 h before the decoke.

2.6.5 The Graded Challenge

The challenge runs 1400 h, about 58 days, on a furnace whose fourth pass already cokes 1.5 times as fast as the others at a coke factor of 1.2, with measurement noise of 0.6 °C on the outlet and 0.25 % on the analyzer. Dilution steam falls to 0.28 at 200 h, the feed rises to 45 t/h at 600 h and the pass 4 skew rises to 2.2 at 900 h. It starts with the cautious tuning, the outlet setpoint at 852 °C, the severity loop in manual, no balancing, the override armed and no automatic decoke. It is met when, judged at 1400 h, the coil outlet stays within 5 °C of its setpoint on every logged sample from 924 h (the last third, restart ramps included), a tube spends less than 0.1 h above its metal limit, and the furnace finishes above 18.5 t/h of ethylene per hour of elapsed time, with the model parameters and disturbances as given.

As shipped, nobody decokes the furnace. A tube crosses 1100 °C at 538 h, the firebox reaches its 1400 °C limit at 634 h, and the coke then chokes the flux until the outlet falls to the 600 °C crossover temperature, so the run ends with a worst outlet error of 220 °C, 858.2 h over the limit and 10.67 t/h. The suggested configuration (the tuned loops, the severity loop in automatic at 0.650, balancing on tube metal, and a decoke at 1090 °C) decokes at 642.4, 987.0 and 1305.0 h and passes with a worst outlet error of 3.53 °C, an hour into the last restart and 1.47 °C inside the band, no time over the limit

Table 2.6 The Shipped Loops Against the Predictive Controller

#	Structure	Severity IAE (h)	Over limit (h)	First over (h)	Run (h)	C_2H_4 (t)	Per elapsed hour (t/h)	Peak TMT (°C)	Plans, mean solve (µs)
8	loops	129.296	260.46	219.53	–	6163	–	1399.0	–
	MPC	113.709	228.51	251.48	–	6691	–	1398.8	4800, 133
5	loops	0.585	0	–	841.3	17489	20.071	1094.0	–
	MPC	13.772	0	–	1040.4	21290	19.889	1094.0	10405, 108
10	loops	2.986	0	–	–	2439	–	1095.4	–
	MPC	2.176	0	–	–	2427	–	1087.3	1200, 111
	loops, on TMT	2.015	0	–	–	2420	–	1095.3	–

Scenario 8 runs 480 h of coking, scenario 5 (one degraded burner) runs to its decoke and scenario 10 runs 120 h. Severity IAE is the integral of the absolute conversion error on the logged rows. Per elapsed hour divides the ethylene of the run by the run plus its 30 h decoke. Peak TMT is taken over every integration step. No plan was rejected. The mean solve time is wall clock in Node.js on the development computer and varies from run to run.

and 20.67 t/h, 2.17 t/h above the target. The app gives this one run whether the settings are made before Run or carried through a Reset. On four other noise seeds the worst error is 3.52 to 3.93 °C. The decoke trigger with the tuned outlet loop alone also passes (3.95 °C, 19.70 t/h), but neither passes without the other.

2.6.6 The Predictive Controller Against the Shipped Loops

Three scenarios were run under both structures with the same seed, events, instruments and duration, and with measurement noise off. The loops keep each case's own settings and shipped tuning. The plan uses the settings scenario 10 ships, a horizon of 1 h, $\lambda_m = 10$ and $p_z = 3$, with no retuning for any case, and neither structure has a preview. Table 2.6 reports the severity IAE of the scorecard, the time with a tube above 1100 °C, the run, the ethylene made, the peak tube metal and the plans.

The coking case, scenario 8, runs 480 h with nobody decoking. The plan holds the balanced tubes on its 1087 °C ceiling and surrenders only the severity the ceiling requires, where under the loops the override chatters in its band. The first tube crosses 1100 °C at 251.5 h instead of 219.5 h, the time over the limit falls from 260.5 to 228.5 h, the severity IAE from 129.3 to 113.7, and the furnace makes 6691 t of ethylene against 6163 t, 8.6 % more. Once the setpoint reaches its 820 °C floor both structures run away, as the case is meant to show.

Scenario 5 run to its first decoke with balancing on coil outlet temperature is a loss. The plan equalizes the outlets much as the shipped trim does until pass 4 reaches the ceiling late in the run, then moves feed onto pass 4 and cuts severity. Its run lasts 1040.4 h against 841.3 h and makes 21 290 t against 17 489 t, but at a severity IAE of 13.772 against 0.585, so the ethylene per hour of run and decoke is 19.889 t/h against 20.071 t/h, 0.9 % less. The ceiling accounts for all of it. At 1087 °C it sits 7 °C under the 1094 °C decoke trigger, so the plan reaches the trigger only after its trims are at their limit and its setpoint is on its floor, while the override lets the tube into its band at full severity. In a diagnostic run that no scenario ships, with the trigger lowered to 1086 °C under the ceiling, the plan decokes at 791.1 h and the loops at 791.8 h, and both earn 20.040 t/h. The plan carries no price on run length or on the decoke, and the comparison does not retune it for this case.

Scenario 10 was built for this comparison. It starts from the furnace of scenario 5 about 780 h into its run, with coke resistances of 4.0×10^{-4} on passes 1 to 3 and 10.6×10^{-4} m^2 K/W on pass 4, which still cokes 1.8 times as fast, the severity loop in automatic at 0.650, balancing on coil outlet temperature and the override armed. Pass 4's tube starts at 1077 °C, 11 °C under the override's start. Figure 2.4 follows the 120 h under both structures. The plan first equalizes the outlets, to within 0.5 °C at 2.6 h against 10.6 h for the shipped trim. Doing so takes feed off pass 4, down to 8.5 t/h, and pass 4's tube reaches the ceiling at 13.7 h. From then on the plan shifts feed back onto pass 4, reaching the trim limit of 13.5 t/h at 57.2 h, and holds the tube on the ceiling. The override never

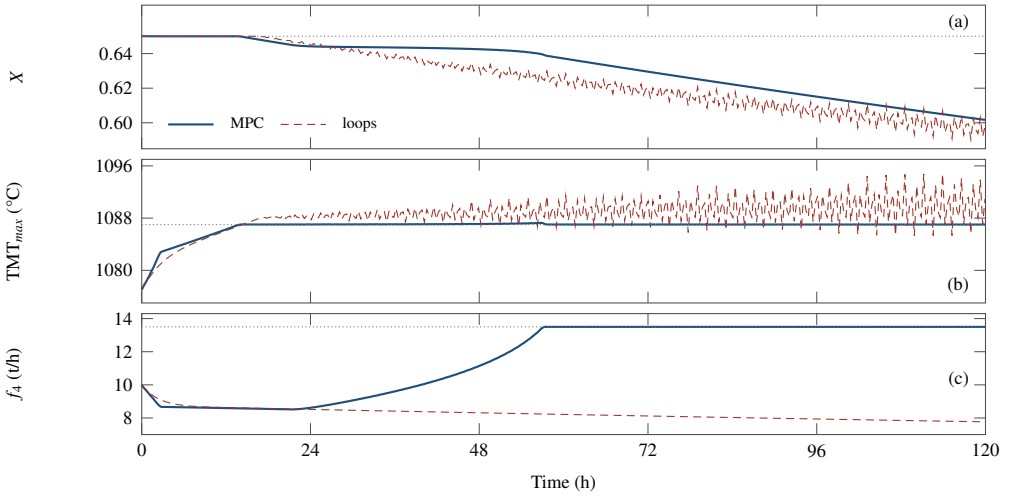

Figure 2.4 Scenario 10 over its 120 h under the predictive controller (thick solid) and the shipped loops (thin dashed), on every logged row. (a) Ethane conversion against its 0.650 target (dotted). (b) Hottest tube metal temperature and the plan's 1087 °C ceiling (dotted), one degree under the override's start. (c) Pass 4 flow and the 35 % trim limit at 13.5 t/h (dotted).

acts and the hottest tube peaks at 1087.3 °C. Under the loops the override first acts at 16.8 h and is acting on 153 of the 359 logged rows, pass 4 falls to 7.8 t/h, and the hottest tube reaches 1095.4 °C (1094.8 °C on the logged rows). Conversion first drops 0.006 below target at 24.7 h under the loops and 25.7 h under the plan. After that the loops' conversion swings with the override, from 0.6275 to 0.6344 over hours 47 to 49 and from 0.6129 to 0.6236 over hours 71 to 73, where the plan's stays within 0.6424 to 0.6427 and 0.6289 to 0.6302. The severity IAE is 2.176 against 2.986 and the peak tube 1087.3 °C against 1095.4 °C, for 2427 t of ethylene against 2439 t, with a mean outlet spread of 44.8 °C against 0.7 °C.

The plan is not the best structure on this case. The loops balancing on tube metal score 2.015, and most of that gap comes from where the plan's ceiling sits rather than from its 1 h view. Balancing on tube metal itself, the plan scores 2.140, so the balancing variable explains 0.036 of the 0.161 difference. The rest comes from the loops letting the hot tube into the override's band, to 1095.3 °C with a mean hottest tube of 1087.85 °C after the first day, while the plan holds 1087.0 °C with a mean of 1086.16 °C. With the override margin cut to 6 °C, which lifts the plan's ceiling to 1093 °C, the plan on tube metal scores 1.858 with no tube above 1093.04 °C, while the loops on tube metal at the same margin score 1.760 but spend 0.87 h above 1100 °C and peak at 1101.4 °C. Over these five days, leaving the flows alone scores best of all, 1.933 under the loops and 2.093 under the plan. Balancing on tube metal leads on days 2 and 3 and trails on days 4 and 5, because feed moved onto pass 4 brings passes 1 to 3 up to the ceiling, 1087.0 °C against 1062.7 °C at 120 h under the plan, and their coke then grows faster, to 8.94×10^{-4} against 6.23×10^{-4} m^2 K/W.

On the challenge, reported for information, choosing the plan and nothing else fails. Nothing decokes in the shipped settings, so every tube runs away once the setpoint is on its floor, with a worst outlet error of 220 °C, 798.5 h over the limit and 11.00 t/h. With the suggested settings and the plan the challenge is met, with a worst outlet error of 3.49 °C, no time over the limit and 21.00 t/h, against 3.53 °C and 20.67 t/h under the suggested loops. The unscored measures show the cost. The severity IAE is 4.978 against 3.015, because the suggested decoke trigger of 1090 °C again sits above the plan's ceiling, and the hottest tube peaks at 1090.1 °C, so no limit the score leaves out is broken.

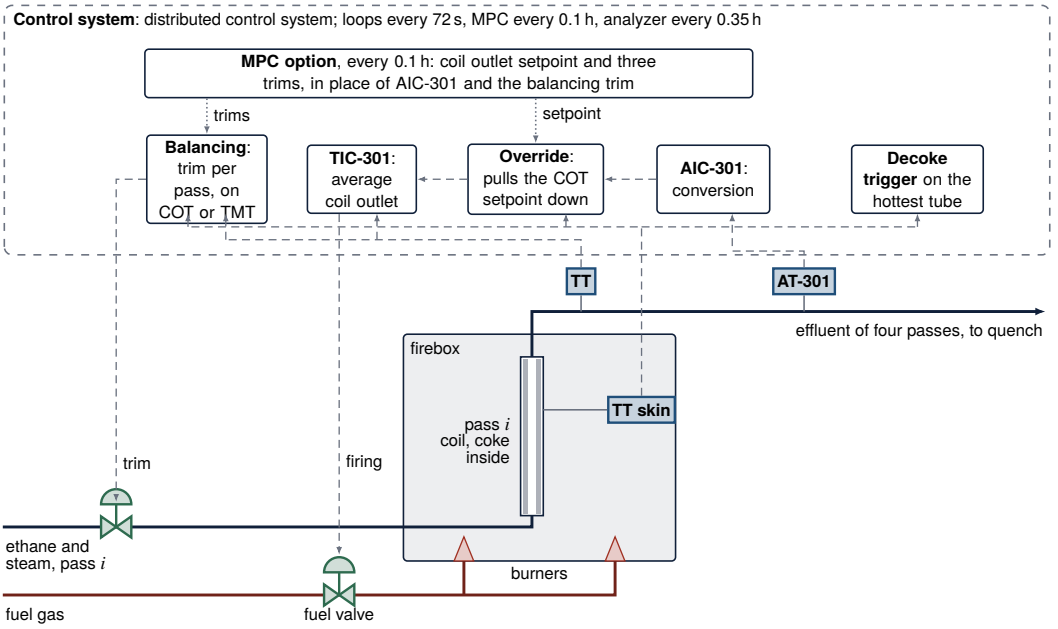

Figure 2.5 The loops of furnace F-301 with their hardware, one pass of four drawn. Blue boxes are sensing devices: TT on the coil outlet of each pass, TT skin on its tube metal and AT-301, the effluent analyzer. Green symbols are final elements: the fuel valve to the burners and the pass flow valve. Dashed arrows are signals, the dotted arrows the outlet setpoint and trims that the MPC option writes in place of the severity loop and the balancing trim. Notice that the balancing block can read either the outlet or the tube metal of each pass, and that only the override and the decoke logic read the tube metal by default. Not to scale, and not a piping diagram.

2.7 Technology in Practice

Figure 2.5 draws the firing loop, the severity loop, the pass balancing and the tube metal override of Section 2.5 with the hardware of a cracking furnace in place of the schematic's symbols.

What Measures It

Each pass carries a thermocouple at its coil outlet, and the firing loop holds the flow-weighted average of the four. Each pass also carries a tube metal reading, which the override, the decoke trigger and the balancing on tube metal use. The severity loop reads an online analyzer on the effluent that reports conversion every 0.35 h. The chapter gives the outlet transmitter a lag of 0.012 h and takes the tube metal readings as exact, which is the assumption the Limitations question.

> **TECHNOLOGY** **Sensor: Tube Metal Temperature**
>
> In furnaces of this kind tube walls commonly run between 700 and 1200 °C. A thermocouple welded to a tube or set into its wall reads continuously, but in the firebox it does not last long, reads the weld rather than the true surface, and is prone to erroneous readings. Operators also survey the tubes with infrared pyrometers and cameras, which receive the energy the wall emits and reflects [17]. Thermocouples are rugged but less accurate than resistance thermometers. An open circuit, their most common failure, can show as a wildly varying reading, so the transmitter is set to go high on it where excess temperature is the danger. A pyrometer sees only the surface, and its main weakness is inaccuracy, since the tube's emittance must be known [18]. The chapter takes its tube readings as exact.

What Moves It

Two kinds of final element carry the handles. The fuel valve to the burners sets the firing, and in a cracking furnace the tubes are heated from outside by gas or oil burners [5]. The chapter models the result as a firebox temperature behind a lag of 0.035 h, within 900 and 1400 °C. The four pass flow valves carry the trims that move feed between passes while the total stays fixed.

> **TECHNOLOGY** **Final Element: The Pass Flow Valve**
>
> Each pass has its own flow control valve, and the balancing scheme trims all four. A pneumatic actuator sets the stem in proportion to its air pressure, against a spring. Friction in the packing makes the stem move in stair steps, called stiction, and leaves it unresponsive to small reversals of the signal. Valves with poor maintenance have shown hysteresis beyond 10 %, and a positioner corrects the gap between signal and stem. The fail position is chosen from the hazard rather than from habit, the way a cooling valve that would starve an engine is made to fail open [18]. The chapter's trims move without friction, behind a lag of 0.010 h, within ± 35 % of the pass flow.

Where the Controller Runs

A distributed control system has a supervisory level that sends set points to the field controllers, its loops are sometimes nested or cascaded, and their cycle times range from milliseconds to minutes [19]. The app's loops execute every 0.02 h (72 s), its MPC every 0.1 h, and its analyzer reports every 0.35 h. The structure maps onto standard blocks: a PI temperature controller on the firing, cascaded under a PI analyzer controller with external reset, an integral trim per pass with its sum held at zero, a proportional override that can only pull the outlet setpoint down, and a decoke trigger in the sequence logic.

> **IN PRACTICE** **Burning the Coke Out**
>
> Almost all plants decoke with steam and air. The air burns the coke off the inside of the coils while steam passes with it to prevent local overheating, at coil outlets of up to 800 °C. A decoke typically lasts 18 to 48 h, and in a European survey decoking took on average about 3 % of the run time. Cyclones or wet scrubbers catch the dust in the gas leaving. Every plant that answered dopes its feed with sulfur compounds against coke, and the burn is pushed hard only near its end, since high temperatures could damage the coils. The European reference document under Directive 2010/75/EU sets out these practices [12]. The chapter's decoke takes 30 h.

What Is Modernizing It

Three developments touch this chapter. The first is the layer above the loops. Crackers use multivariable control with online analyzers, performance and constraint controls, often with online optimization, and for a gas feed the severity they manage is the conversion [12]. The predictive controller of Section 2.5.5 is a small example of that layer, and its history reaches back to a pyrolysis furnace [8]. The second is the coil itself. Coils with less nickel or with an inert inner coating retard coke, mainly in its first layer, and recoating may be needed after a few decokes [12]. Sensors that detect fouling early and high-temperature coatings have been studied as well, and modern cracking furnaces reach fuel efficiencies of typically 92 to 94 % [5]. The third is the heat source. Electric cracking would replace the burners, and with them the handle that TIC-301 moves, by resistance or other electric heating [20].

2.8 Limitations

The coil is lumped rather than distributed, so there is no axial temperature profile and no way to study where along the tube the coke actually forms.

Cracking is a conversion and selectivity correlation calibrated for ethane. A naphtha furnace has a far richer product slate and a much stronger yield dependence on temperature, which is exactly where the balancing trade in scenario 4 becomes expensive, and the application cannot show it.

The tube metal thermocouple is treated as accurate. In industry the argument about balancing on tube metal is largely an argument about whether the measurement can be trusted, and a simulator that assumes it can understates the difficulty of adopting the scheme it recommends.

The transfer line exchanger downstream is absent, along with the coking that occurs in it and often ends a real campaign before the radiant coil does.

The decoke restores the tube exactly and takes a fixed thirty hours. Neither is true, and real decoke duration depends on how much coke is present [2]. The restart is also simplified. The decoke holds the outlet at its 880 °C limit and the setpoint ramps down to target as the feed returns, where a real furnace would bring the outlet up from a cooler steam standby. Dilution steam is tied to the hydrocarbon flow here, so the nearly empty coils cannot be held cool with the firing at its 900 °C floor.

The predictive controller predicts with clean tube slopes, and coke weakens them. With 7.5×10^{-4} m^2 K/W of coke on every pass, the outlet's response to the firebox and the metal's response to the pass flow are both 14.4 % weaker. At the start of scenario 10 the hot pass's metal answers its own flow 26.7 % more weakly than the model does, at −3.80 °C per t/h against the model's −5.19 °C. The plan therefore expects more cooling per tonne it moves than it gets, and only the next measured offset corrects it. The model also leaves out a parameter changed by the learner, which it treats as error, and the decoke and restart ramp, during which it stands aside. The plan has no price on run length or on the decoke, which is why it loses scenario 5. Its back-off of one degree under the override is fixed rather than chosen against the decoke trigger or the thermocouple's error.

The plan's cost in computing is visible at the top speed. At 600 h per second in Chrome the plans take 41.9 % of the wall clock, the page draws 15.3 frames per second against 50.5 under the loops, and the run advances about 456 h per second rather than 600. The logged run is identical at every speed, and at 200 h per second the page keeps 59.6 frames per second.

Finally, the application models one furnace. An ethylene plant runs a bank of them, and the scheduling problem of staggering their decokes against a plant demand is a large part of the real economics [7] and is entirely outside this model.

> ⚠ **OPEN CHALLENGE** **How Fast Will This Furnace Coke?**
>
> Coke formation is still one of the most important problems of steam cracking. Its properties had hardly been studied for a decade, and reported values of some of them differed by orders of magnitude. A 2020 study measured them on industrial samples and showed what the gap costs: across the values in use, the predicted run length of an industrial furnace on ethane ranged from 20 to 227 days. Progress means coke properties measured well enough for run length models to predict [1]. The coking law of Equation (2.4) is exact in the simulator and uncertain on every real furnace.

2.9 Discussion

Three results generalize past cracking furnaces.

The first is that a balancing scheme equalizes what it measures, and if the constraint lives on the other side of a resistance the scheme can make things worse. The measured case here is 841 h against 888 h for no control, and it arises from a sign: the resistance makes the constrained variable rise and the measured variable fall, so the scheme moves the wrong way. Before configuring any balancing loop it is worth writing down what is between the instrument and the limit.

The second is that an optimum can be invisible to every continuous measurement on a unit. Conversion and ethylene rate both rise monotonically through the optimum in Table 2.5, and only production per hour of elapsed time finds it. Where that happens, an optimizer needs a model rather than a measurement, and the model has to include the thing that ends the run [8].

The third is the shape of the optimum rather than its location. Five degrees above the peak costs 0.5 % of the on-stream rate and fifteen degrees costs 5.9 %, while fifteen degrees below it costs 2.7 %. An optimum with a flat top and a steep side should be approached from the flat side, which is a statement about where to put the setpoint rather than about how to find the peak.

OPEN CHALLENGE **A Cracking Furnace Without Flames**

Gas-fired cracking furnaces emit about 366 million tonnes of carbon dioxide a year, and improvements to the present designs cannot reach the cuts now required. Electric heating avoids the heat lost with the flue gas, and resistance heating could be retrofitted to existing crackers. It still has to be tested and demonstrated, and the furnace's heat is woven into the rest of the plant: the steam it raises drives the turbines of the compressors, which would become electric motors. Progress means a demonstrated electric cracker and a plant whose heat integration is redesigned around it [20].

2.10 Conclusion

A cracking furnace ends its campaign when one tube reaches its metal limit, and the instrument that scheme designers reach for is on the wrong side of the coke. In the simulated furnace of this chapter, four passes whose coil outlet temperatures are matched to within half a degree still carry 16.1 °C of tube metal spread, and on a furnace with one degraded burner, balancing on coil outlet temperature gives a run of 841 h against 888 h for no balancing at all and 910 h for balancing on tube metal. The severity optimum is equally hidden. Conversion and ethylene rate both rise monotonically from 835 to 865 °C while ethylene per hour of elapsed time peaks near 850 °C at 20.31 t/h and falls to 19.12 t/h by 865, and the run length falls from 2516 h to 186 h across the same range.

A model predictive controller offered beside the loops shortens the time above the tube limit in the coking case from 260.5 to 228.5 h, but its ceiling 7 °C under the decoke trigger costs 0.9 % of on-stream rate on the degraded burner furnace. A fuller model would add an axial profile so that the location of the coke can be studied, a naphtha feedstock so that the yield cost of unequal pass severities becomes visible, a transfer line exchanger, and a bank of furnaces so that decoke staggering can be scheduled. The furnace as it stands shows what a balancing scheme actually balances.

KEY POINTS

- A balancing scheme evens out what it measures, and on a cracking furnace the limit is on the tube metal, on the other side of the coke from the outlet thermocouple.
- Four passes with outlets matched to within half a degree still carry 16.1 °C of tube metal spread.
- With one degraded burner, balancing on coil outlet temperature gives a run of 841 h against 888 h for no balancing and 910 h for balancing on tube metal.
- Ethylene per elapsed hour peaks near 850 °C while conversion and the ethylene rate keep rising, so no continuous measurement finds the optimum.

Where else this applies. Fouling that walks toward a hard temperature limit under an override returns in Volume 6, Chapter 4, and a protective override that acts only near its limit in Volume 4, Chapter 2.

References

[1] M. Geerts, S. H. Symoens, P. A. Reyniers, G. B. Marin, M.-F. Reyniers, and K. M. Van Geem, "Steam cracking coke properties and their influence on furnace run length predictions: Experimental and modeling study," *Industrial & Engineering Chemistry Research*, vol. 59, no. 52, pp. 22 460–22 472, 2020.

[2] Y. Zhang, W. Cai, and F. Xin, "Simulation and optimization of coil decoking in ethane pyrolysis furnace," *Chemical Engineering Communications*, vol. 196, no. 8, pp. 950–968, 2009.

[3] M. Rosenzweig, "A petrochemical centennial approaches," Chemical Processing, August 7, 2019, chemicalprocessing.com/articles/2019/a-petrochemical-centennial-approaches, read November 2026.

[4] R. D. Stief, *A History of Union Carbide Corporation: From the 1890s to the 1990s*. Danbury, CT: Carbide Retiree Corps, 1998, chapter 2, read in the copy at carbidecamps.net, November 2026.

[5] M. Neelis, E. Worrell, and E. Masanet, "Energy efficiency improvement and cost saving opportunities for the petrochemical industry: An ENERGY STAR guide for energy and plant managers," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-964E, 2008.

[6] K. Barazandeh, O. Dehghani, M. Hamidi, E. Aryafard, and M. R. Rahimpour, "Investigation of coil outlet temperature effect on the performance of naphtha cracking furnace," *Chemical Engineering Research and Design*, vol. 94, pp. 307–316, 2015.

[7] J. Tian, Y. Wang, and X. Feng, "Simultaneous optimization of flow velocity and cleaning schedule for mitigating fouling in refinery heat exchanger networks," *Energy*, vol. 109, pp. 1118–1129, 2016.

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

[9] W. Shen, Z. Tian, L. Zhao, and F. Qian, "Life cycle assessment and multiobjective optimization for steam cracking process in ethylene plant," *ACS Omega*, vol. 7, no. 18, pp. 15 507–15 517, 2022.

[10] R. van Welsenaere and G. Froment, "Parametric sensitivity and runaway in fixed bed catalytic reactors," *Chemical Engineering Science*, vol. 25, no. 10, pp. 1503–1516, 1970.

[11] M. Morbidelli and A. Varma, "Parametric sensitivity and runaway in fixed-bed catalytic reactors," *Chemical Engineering Science*, vol. 41, no. 4, pp. 1063–1071, 1986.

[12] H. Falcke, S. Holbrook, I. Clenahan, A. López Carretero, T. Sanalan, T. Brinkmann, J. Roth, B. Zerger, S. Roudier, and L. Delgado Sancho, "Best available techniques (BAT) reference document for the production of large volume organic chemicals," Publications Office of the European Union, Luxembourg, EUR 28882 EN, 2017.

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

[14] 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.

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

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

[17] S. James, "Reformer and heater tubes temperature accuracy: Comparing practical infrared techniques," Conference paper, 2015, irinfo.org/images/articles/pdf/02-01-2015-james.pdf, read November 2026.

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

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

[20] M. E. H. Tijani, H. Zondag, and Y. Van Delft, "Review of electric cracking of hydrocarbons," *ACS Sustainable Chemistry & Engineering*, vol. 10, no. 49, pp. 16 070–16 089, 2022.

2.11 Exercises

2.1 Concept check · 2 min

Pass 4 of a furnace balanced on coil outlet temperature carries more coke than the other three passes. Compared with the others, what do its outlet and skin thermocouples read, and which way does the balancing scheme move its feed?

(a) Its outlet and its skin both read hot, so the scheme adds feed to pass 4

(b) Its outlet and its skin both read cold, so the scheme takes feed from pass 4

(c) Its outlet reads cold and its skin reads hot, so the scheme takes feed from pass 4

(d) Its outlet reads hot and its skin reads cold, so the scheme adds feed to pass 4

2.2 Concept check · 1 min

Table 2.5 lists five coil outlet setpoints. Which quantity, read on the furnace, locates the best of them?

(a) The analyzer conversion, which levels off at the best setpoint

(b) The instantaneous ethylene rate, which peaks where selectivity losses overtake conversion gains

(c) The selectivity, which is highest at the setpoint with the longest run

(d) The ethylene per hour of elapsed time with the decoke counted, known only after the decoke

2.3 Concept check · 2 min

All four passes coke together while the firing holds the coil outlet. How does the coking rate develop over the run?

(a) It grows as the square of itself, so it would become unbounded at a finite time t^*

(b) It grows exponentially in time, with a fixed e-folding time set by θ_c

(c) It stays at its first-day value, because the held outlet fixes the conversion

(d) It falls slowly, because the growing coke lowers the flux through each tube

2.4 Theory and calculation · 6 min

When the outlets of scenario 3 first agree to within half a degree, the firebox is at 1178.9 °C, the outlets are at 845 °C and the mean resistance $R = R_w + R_c$ is 1.77×10^{-3} m^2 K/W. Take $h_{fb} = 0.73\,\mathrm{kW/(m^2\,K)}$.

(a) Evaluate the sensitivity of Equation (2.8). *Checkpoint:* 4.64×10^4 °C per m^2 K/W.

(b) Pass 4 carries 3.5×10^{-4} m^2 K/W more coke than pass 1. Estimate the metal spread and compare it with the 16.1 °C the model reports. *Checkpoint:* 16.2 °C.

(c) Evaluate the same sensitivity on clean tubes at the design point ($T_{fb} = 1150.9\,^\circ\mathrm{C}$, COT = 845.0 °C, $R = R_w = 1.50 \times 10^{-3}\,\mathrm{m^2\,K/W}$), and say why it is larger. *Checkpoint:* 5.09×10^4 °C per m^2 K/W.

2.5 Theory and calculation · 7 min

The clean furnace runs at its design point: tube metal 1004.9 °C, flux 106.6 kW/m^2, coil outlet held at 845 °C, conversion and steam ratio at their reference values and coke factor one. Take $k_c = 2.44 \times 10^{-7}\,\mathrm{m^2\,K/(W\,h)}$, $T_{c,ref} = 1004\,^\circ\mathrm{C}$, $\theta_c = 30\,^\circ\mathrm{C}$ and a metal limit of 1100 °C.

(a) Find the coking rate $\dot{R}_0$ from Equation (2.4) and the rate at which the metal starts to rise. *Checkpoint:* 2.51×10^{-7} m^2 K/(W h) and 0.0268 °C/h.

(b) Find t^* and the run length to the limit. *Checkpoint:* 1120 h and 1073 h, within 2 h.

(c) Find the run length that a straight line through the first day's metal slope predicts, and its ratio to the run of (b). *Checkpoint:* 3550 h, 3.3 times too long.

2.6 Theory and calculation · 6 min

A planner notes that a coil outlet of 865 °C makes 6.8 % more ethylene than 845 °C and proposes the change. Use Equations (2.5) and (2.6) with $F = 40\,\mathrm{t/h}$, $X_0 = 0.650$, $\sigma_0 = 0.800$, $\mathrm{d}X/\mathrm{d}T = 0.004\,20/^\circ\mathrm{C}$ and $\mathrm{d}\sigma/\mathrm{d}X = -0.52$, and a decoke of 30 h.

(a) Find the conversion, the selectivity and the ethylene rate at 850 °C. *Checkpoint:* 0.671, 0.789 and 21.18 t/h.

(b) The run at 850 °C lasts 700 h. Find the on-stream rate with Equation (1.3). *Checkpoint:* 20.31 t/h.

(c) The rows of Table 2.5 at 835 and 865 °C make 19.99 and 22.21 t/h for 2516 and 186 h. Find their on-stream rates, how far each lies below the result of (b), and answer the planner. *Checkpoint:* 19.75 t/h, 2.7 to 2.8 % below, and 19.12 t/h, 5.8 to 5.9 % below.

2.7 App challenge · 8 min

Load scenario 3, Equal outlets, unequal metal, and change no setting (Pass balancing on coil outlet, default seed). Run at 60 h per second until the clock passes 480 h, pause and export the CSV file.

(a) Find the first row on which `cot_spread_C` is below 0.5 °C. Give its time and the tube metal spread `spread_C` on that row. *Checkpoint:* 8.4 h (the row at 8.375 h) and 16.1 °C, within 0.1 °C.

(b) Read `spread_C` on the last row at or before 240 h and on the last row at or before 480 h, and the largest `cot_spread_C` after the row of (a). *Checkpoint:* 20.9 °C and 34.0 °C, within 0.1 °C. The outlet spread stays under 0.5 °C.

(c) Compare (a) with Exercise 2.4 and explain why the outlet scheme cannot remove the spread, and why it grows.

2.8 App challenge · 8 min

Load scenario 5, A burner problem on one corner, as shipped (Pass balancing on coil outlet, decoke on tube metal at 1094 °C, default seed). Run at 60 h per second until the clock passes 950 h, pause and export the CSV file. Press Reset, set Pass balancing to On tube metal before pressing Run, and repeat. Press Reset again, set Pass balancing to Off, and repeat.

(a) In each file, read `run_hours` on the first row with `decoking` equal to 1. *Checkpoint:* 841.3 h, 910.4 h and 887.9 h, within 0.2 h.

(b) On the row at 719.9 h of each file, read `flow0_tph` and `flow3_tph`. *Checkpoint:* 10.32 and 9.03 t/h, 9.66 and 11.01 t/h, 10.00 and 10.00 t/h, within 0.02 t/h.

(c) Explain from (b) why balancing on coil outlet gives a shorter run than no balancing at all.

2.9 App challenge · 9 min

Load scenario 10, Balance the passes before the override, as shipped (Supervisory control on MPC, default seed). Run at 10 h per second to about 115 h, switch to 1 h per second, and pause when the clock reads between 120 and 121 h. Read the IAE tile of the Conversion SP / PV loop on the scorecard and export the CSV file. Press Reset, set Supervisory control to Loops before pressing Run, and repeat.

(a) Give the IAE under each structure. *Checkpoint:* MPC 2.176 to 2.225, Loops 2.986 to 3.040, as the tile prints them.

(b) In each file, over the 359 rows up to 120 h, count the rows on which `tmt_override` is 1, and find the largest `tmt_max_C`. *Checkpoint:* MPC 0 rows and 1087.3 °C. Loops 153 rows and 1094.8 °C.

(c) Explain why the override never acts under the MPC, and what the plan moves instead.

2.10 Application area · 4 min

A plant installs a decoke system that takes 20 h instead of 30 h. Each run starts from clean tubes, so the run lengths of Table 2.5 do not change.

(a) With Equation (1.3) and the table's ethylene rates and runs, find the on-stream rate at 845, 850, 855 and 865 °C. *Checkpoint:* 20.42, 20.59, 20.63 and 20.05 t/h.

(b) Say where the optimum moves and why, and what setpoint you would choose.

2.11 Application area · 3 min

The predictive controller holds each tube under a ceiling one degree below the override's start, and the override starts Δ_o below the 1100 °C limit. Skin thermocouples read the tube surface, and suppose the one on pass 4 reads 5 °C below the true metal.

(a) Find the ceiling for $\Delta_o = 12\,^{\circ}\mathrm{C}$ and for $\Delta_o = 6\,^{\circ}\mathrm{C}$, and the true metal of pass 4 when its reading sits on the ceiling. *Checkpoint:* 1087 °C and 1093 °C, with true metal 1092 °C and 1098 °C.

(b) Use Equation (1.13) to say what the 13 °C between the limit and the shipped ceiling has to cover, and which margin you would ship.

2.12 Review the agent · 5 min

An AI assistant reads Table 2.6 and writes: “In scenario 5 the predictive controller makes 0.9 % less ethylene per elapsed hour than the loops. The cause is its one-hour prediction horizon, which cannot see the coke that builds over weeks. Set the Prediction horizon to its 3 h maximum and the plan will beat the loops on scenario 5.” Find the flaw, and say how you would check the claim in the app.

CHAPTER 3

A One-Sided Hydrogen Window and a Selectivity Runaway in an Acetylene Converter

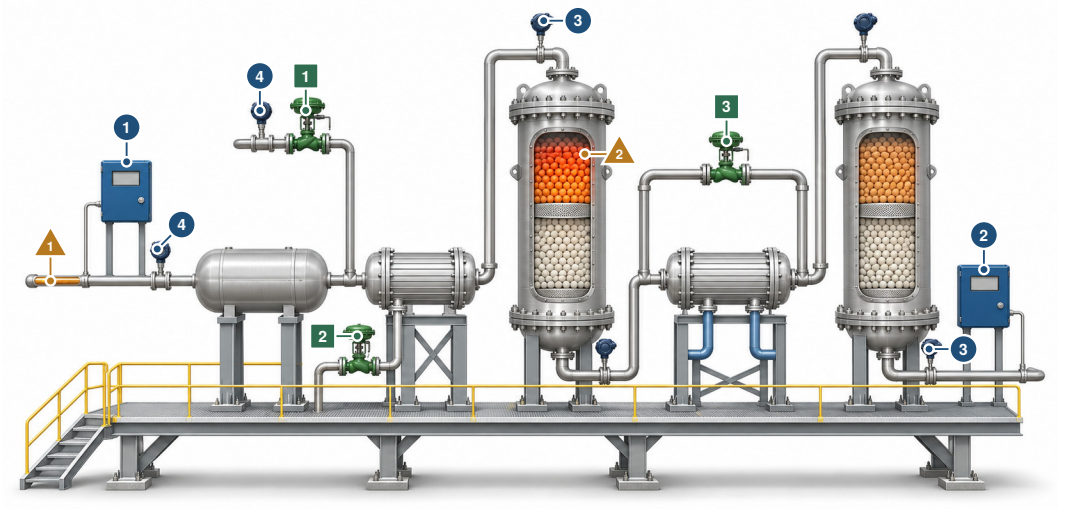

The controller moves 1 hydrogen valve, set in ratio to the acetylene fed
2 steam to the heater, which sets the first bed's inlet temperature
3 hot-gas bypass around the intercooler, which sets the second bed's inlet temperature

The controller reads 1 chromatograph on the feed, ahead of the drum
2 chromatograph on the product, which reports the acetylene left minutes late
3 temperature sensors at each bed's inlet and outlet 4 feed and hydrogen flowmeters

What disturbs it 1 acetylene and carbon monoxide in the feed, which the furnaces change
2 green oil that slowly fouls the catalyst, first at the front of the first bed ▲ a cut in the feed flow

Notice that the feed chromatograph samples ahead of the drum, so the hydrogen can follow a feed change before it reaches the beds, while the acetylene left is known only from the chromatograph at the far end. Illustration, not to scale.

IN PLAIN TERMS **Just Enough Hydrogen, and No More**

The main idea. Gas from a cracking furnace holds a trace of acetylene, which must go before the main gas, ethylene, is made into plastic. A converter adds hydrogen over beds of small pellets, which help it turn the trace into more of the main gas.

The challenge. Too little hydrogen leaves some of the trace and spoils all the product. Too much turns good ethylene into a gas worth less, and the heat it gives off can make a bed run hot before the thermometers know.

Why it matters. The best setting is the least hydrogen that clears the trace, and it moves when the feed changes.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History boxes in Section 3.2, the Did you know? boxes in Sections 3.1 and 3.5, the Try it boxes in Section 3.5, and Key points in Section 3.8. No equations.
Study	45 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 3.5 with Figure 3.4, Section 3.6, and Section 3.7 with its two Open challenge boxes.

The control question Too little hydrogen puts the whole ethylene product off grade, and too much turns ethylene into ethane and can light off the bed. Where in this one-sided window should the hydrogen sit, and what keeps it there?

Learning objectives

- Price both failures of the one-sided hydrogen window in tonnes of ethylene a day.
- Compute the heats, the adiabatic rises and the temperature sensitivity of the two rates.
- Explain why a bed's rise falls while its selectivity collapses, and why an override on it cannot stop a light-off.
- Compare a fixed ratio, feedforwards with a trim and a predictive controller on a graded day.
- Judge what an inlet temperature schedule buys and costs over a catalyst run.

Prerequisites Sections 1.4.3 (selectivity and light-off), 1.4.6 (inferred quality), 1.4.8 (limits and guards), 1.6.2 and 1.6.4. Section 1.6.3 only for the predictive controller.

The app https://apmonitor.com/apps/acetylene-hydrogenation/

See also Volume 7, Chapter 2 for an ignition boundary. Volume 3, Chapter 6 and Volume 3, Chapter 4 for a ratio held by an analyzer trim. Volume 4, Chapter 2 for a protective override.

3.1 Introduction

At 1.3 mol of hydrogen per mol of acetylene the converter simulated here leaves 14.7 ppm of acetylene in its product, and at 1.9 it leaves 0.0002 ppm. The stream is 1400 mol/s (5040 kmol/h) of C_2 gas carrying 1.0 % acetylene, 79 % ethylene and 20 % ethane, and the same 0.6 mol/mol that clears five decades of acetylene raises the ethylene converted to ethane from 0.138 to 0.472 % of the ethylene fed. At the design ratio of 1.5 (75.6 kmol/h of hydrogen) the lined-out outlet holds 0.31 ppm and the loss is 0.200 %, which is 5.4 t of the 2680 t of ethylene the stream carries each day. Below a ratio of 1.43 the product fails the 1 ppm limit of this model and all of it is off grade.

The problem is old and general. Acetylene must be taken from cracked gas to a few ppm before ethylene can be polymerized, palladium catalysts do it selectively, a little carbon monoxide is added to hold the selectivity, and the adiabatic packed beds that carry the reaction are known for thermal runaway in industrial practice [1]. The catalytic literature has two reviews of its own [2, 3]. The runaway belongs to a class wider than this reactor, the parametric sensitivity of exothermic tubular reactors [4], and in particular of reactors carrying consecutive reactions in which the desired product is the intermediate [5]. The control problem has a familiar shape too. A reagent is dosed in ratio to a measured flow, a feedforward answers an upstream disturbance, a slow analyzer trims the dose, and a selector cuts it when a constraint is reached [6]. The shape recurs wherever a reagent is dosed to a parts-per-million specification against a costly side reaction. A water plant that doses coagulant from a raw-water turbidimeter and demotes its slow settled-water measurement to a ratio trim has the same structure (Volume 3, Chapter 6).

> ✦ DID YOU KNOW? **A Trace That Must Go**
>
> In the typical once-through yields of an ethane cracking furnace listed in a 2017 European reference document, ethylene makes up 56 % of the products and acetylene only 0.4 % [7]. The converter of this chapter takes a C_2 stream from 1.0 % acetylene to under 1 ppm.

The simulator corrects two beliefs. The first is that a margin on hydrogen is a margin on safety. The window is one-sided, so the operating point belongs just on the clean side of the breakthrough edge, and every mole beyond it buys ethylene loss and moves the plant toward the runaway. The second is that the runaway is a temperature problem that a temperature limit catches. In case 6 of this model, hydrogen raised by hand from 75.6 to 130 kmol/h empties bed 1 of acetylene in 14 s,

the loss reaches 0.79 % by 30 s after the change while bed 1's temperature rise falls from 30.2 to 26.9 K, and the rise reaches the 70 K trip only 3.95 min after the change. The selectivity is lost before a thermocouple can tell.

However, the published work treats these parts apart. Kinetic studies measure rates on commercial catalysts [8–10], reactor models are built to simulate and optimize an industrial train [10, 11], and catalysis research looks for more selective surfaces [12, 13]. This chapter works with a two-bed tail-end converter simulated in the browser, which records every cell temperature, the true outlet and the true loss beside the two chromatographs. It measures how a hydrogen ratio, its feedforwards, an analyzer trim, a temperature override and a predictive controller keep the plant inside its window through carbon monoxide swings, feed changes, an analyzer failure and catalyst aging. The rate law (the classical one-site form, given teaching constants), the heats of reaction and the controllers are all established methods. What the converter adds is numbers, and two of them matter most: an override on the bed rise cannot see a light-off coming, and the design point is not the point of least loss. The model has 110 states on the shared packed-bed module of the app framework, with carbon and hydrogen closing to round-off and a strict line between the simulator's truth and the instruments, and its nine graded cases give credit only when the plant has answered what the learner did. The general form of the one-sided reagent window and of the light-off is in Section 1.4.3.

Section 3.2 traces the process, its catalyst and its moderator. Section 3.3 builds the model: the rate law, the heats, the beds of cells, the green oil, the feed drum, the heater and intercooler, and the instruments. Section 3.4 describes the controllers, from the hydrogen ratio and its feedforwards to the balance inference and the predictive layer. Section 3.5 runs the cases, from the window and its two costs to the light-off, the end of a catalyst run and a graded day on the furnace cycle. Section 3.6 names the hardware and the practice behind the model, Section 3.7 states what the model cannot do, and Section 3.8 sums up.

3.2 Acetylene in Ethylene: Catalyst, Moderator and Runaway

Cracked gas carries about one percent of ethyne. It reaches the converter either inside a gas that still holds 10 to 20 % hydrogen (the front-end arrangement) or as a stream of ethene and ethane alone, after the light ends have gone (the tail-end arrangement, which needs hydrogen added) [1]. The cracking furnace that makes this stream is the subject of Chapter 2. In either case the ethyne has to fall below about 5 ppm while a larger loss of ethene is economically intolerable [1]. Supported palladium met these demands, on low-area alumina with only 0.01 to 0.1 % of the metal by weight and often as an egg-shell layer [1], and industrial catalysts are commonly modified with silver [13].

> ⧖ HISTORY **A Picture of Palladium Revised, 1993 to 2008**
>
> The first explanation of palladium's selectivity was thermodynamic. Ethyne adsorbs on the metal much more strongly than ethene, so the surface stays covered with ethyne until it is nearly gone. Carbon monoxide, which also adsorbs more strongly than ethene, was thought to help in the same way. A review of 1993 called this classical view an over-simplification. Ethene is hydrogenated even at high ethyne concentrations, a direct route from ethyne to ethane exists, at least two kinds of active site were by then generally assumed, and the effect of carbon monoxide was put down to competition with hydrogen rather than with ethene [1]. In 2008, measurements made during the reaction traced the selective state to carbon from broken-up feed molecules held just under the palladium surface, with unselective total hydrogenation on the hydrogen-saturated β hydride [12].

The classical picture of the selectivity, told in the History box above, is the one-site Langmuir–Hinshelwood law this simulator uses. The changes the catalyst undergoes in use, its steady-state kinetics and the effects of carbon monoxide and promoters have been reviewed at length [2, 3]. Ethyne also oligomerizes on the catalyst to a heavy liquid called green oil, which is a serious problem in practice [1].

Table 3.1 Variables, Roles and Design Values (Case 2, Lined Out)

Variable	Role	Design	Range
A_{out}	CV, outlet acetylene (analyzer)	0.31 ppm	SP 0.1–2
R	MV of the ratio logic, mol/mol	1.50	0.8–2.5
F_{H_2}	MV, hydrogen (kmol/h)	75.6	0–201.6
$T_{in,1}$, Q	CV, MV: bed 1 inlet, steam	41.0 °C, 35.7 %	SP 25–85
$T_{in,2}$, v_2	CV, MV: bed 2 inlet, bypass	58.0 °C, 69.5 %	SP 35–90
ΔT_1, ΔT_2	constraint, bed rises	30.2, 9.8 K	override 55
y_A	DV, feed acetylene	1.00 mol%	0.3–1.8
p_{CO}	DV, carbon monoxide	2.0 ppm	0–10
—	DV, C_2 flow, cooling water (%)	100, 100	50–120, 0–100
—	DV, either analyzer in or out	in service	—
L	output, ethylene lost	0.200 %	budget 0.33
T_i, g_i	states, 48 cells each	Fig. 3.1	—

Plant conditions entered the laboratory as micropilot and kinetic studies, which the History box below follows from 1982 to 1993.

> ⧖ HISTORY **1982 to 1993: Plant Conditions in the Laboratory**
>
> A micropilot-plant study of catalyst performance and aging was reported in 1982 [14]. A review of its results notes that the overall selectivity rose with the carbon monoxide content and that the best carbon monoxide level, hydrogen ratio and temperature depended on the catalyst [1]. In 1993 a kinetic study ran a commercial Pd/Al_2O_3 catalyst in a Berty reactor, which has internal recycle, at tail-end conditions: 0.3 to 1.3 % ethyne, 0.4 to 4 % hydrogen, 0 to 60 ppm carbon monoxide, 0.3 to 2.1 MPa and 299 to 330 K. It tested eight rate expressions and found the selectivity 0.3 to 0.6 without carbon monoxide and 0.7 to 0.9 with it, found the rates in ethyne-ethene mixtures typically two orders of magnitude below that of ethene alone, and showed that the classical model 1 cannot follow the effect of carbon monoxide on the differential selectivity [8].

Front-end kinetics with carbon monoxide followed [9], as did a fixed-bed study with catalyst deactivation and an industrial reactor simulated in plug and dispersed flow [10] and the modeling and optimization of a front-end train [11]. Catalyst design moved to calculation, which proposed nickel–zinc alloys as a cheaper selective catalyst [13]. The sources consulted here do not describe a commercial converter's control system, so the loops below are built from the general methods of ratio, feedforward, selective and inferential control [6, 15] and of industrial model predictive control [16] and describe no licensor's practice. The same contest between a wanted and an unwanted reaction on a hot catalyst returns in the ethylene oxide reactor of Chapter 6, where the unwanted one makes most of the heat.

3.3 The Converter Model

The converter is generic and sized without plant data (Figure 3.1). The C_2 stream from the deethanizer passes a feed drum, hydrogen is added in ratio to its acetylene, a steam heater sets bed 1's inlet, a cooling-water intercooler with a hot-gas bypass sets bed 2's inlet, and a chromatograph samples the gas ahead of the drum and another the product. Time in the engine runs in minutes, with flows in mol/s and heat in watts.

The variables are listed in Table 3.1. The integrated state holds the 48 cell temperatures, 48 green-oil loadings, the heater duty, the intercooler metal, the bypass position, the hydrogen flow and the acetylene and carbon monoxide leaving the feed drum, 102 physical states, and eight ledgers that integrate each of four species into and out of the converter, 110 states in all.

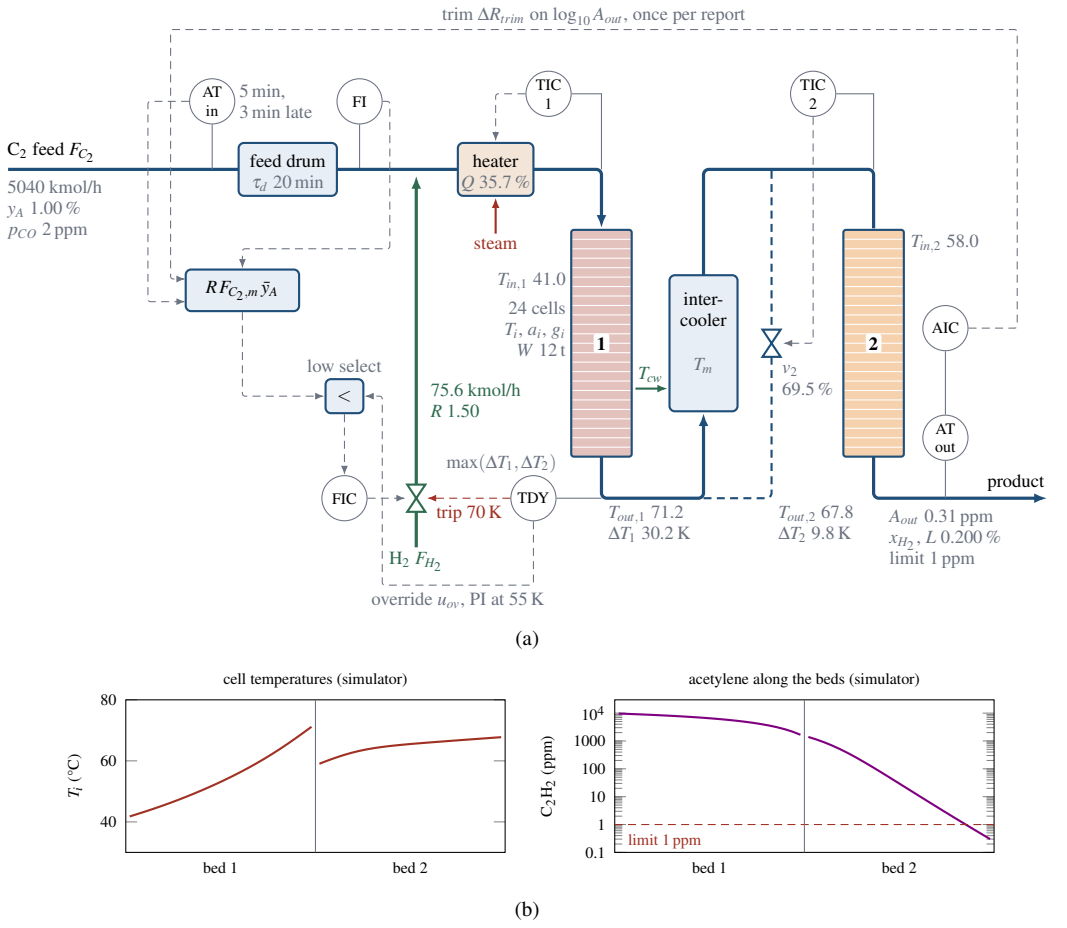

Figure 3.1 The tail-end converter as modeled, case 2 at $t = 1$ min, the design point. (a) After the app's Reactor view. The C_2 feed F_{C_2} with acetylene y_A and carbon monoxide p_{CO} passes the feed drum (lag τ_d), the inlet analyzer (AT in) sampling ahead of it. Hydrogen F_{H_2} follows the ratio logic $R F_{C_2,m} \bar{y}_A$ through a low select against the override u_{ov} from the larger bed rise. The heater duty Q (TIC 1) sets $T_{in,1}$, the bypass v_2 around the intercooler (metal T_m, water T_{cw}) sets $T_{in,2}$, each bed is 24 cells of temperature T_i, activity a_i and green oil g_i, and the outlet analyzer reports A_{out}, x_{H_2} and the loss L to the trim ΔR_{trim}. Printed values are the lined-out design: 75.6 kmol/h and rises of 30.2 and 9.8 K, where the app's panel prints the instruments, 75.5 kmol/h and 30.2 and 10.0 K. A model drawing, not piping. Not to scale. (b) The simulator's cell temperatures and acetylene along the 48 cells at the same instant. The Window view at this instant is Figure 3.3.

3.3.1 Two Hydrogenations on One Kind of Site

The rates per kilogram of catalyst have the classical Langmuir–Hinshelwood form, as shown in Equation (3.1).

$$\begin{aligned} &C_2H_2 + H_2 \rightarrow C_2H_4, \qquad C_2H_4 + H_2 \rightarrow C_2H_6 \\ &r_A = \frac{a k_A p_A p_H}{D^2}, \qquad r_E = \frac{a k_E p_E p_H}{D^2} \\ &D = 1 + K_A p_A + K_E p_E + K_{CO} p_{CO} \end{aligned} \tag{3.1}$$

where $p_i = P F_i / F_t$ is in bar at $P = 20\,\text{bar}$ absolute, a is the cell's activity, and the constants depend on temperature as shown in Equation (3.2).

$$\begin{aligned} k_j &= k_{j,r}\, \mathrm{e}^{-\frac{E_j}{R_g}\left(\frac{1}{T}-\frac{1}{T_r}\right)}, \qquad j = A, E \\ K_{CO} &= K_{CO,r}\, \mathrm{e}^{\frac{Q_{CO}}{R_g}\left(\frac{1}{T}-\frac{1}{T_r}\right)} \end{aligned} \tag{3.2}$$

with $T_r = 60\,^\circ\text{C}$, $k_{A,r} = 3$ and $k_{E,r} = 4 \times 10^{-4}$ mol/(s kg bar^2), $E_A = 25$ and $E_E = 60\,\text{kJ/mol}$, $K_A = 50$, $K_E = 0.01$ and $K_{CO,r} = 2.5 \times 10^4$ bar^{-1}, $Q_{CO} = 30\,\text{kJ/mol}$ and no temperature dependence of K_A or K_E. This is model 1 of Bos et al. [8] without its hydrogen adsorption term, and every constant is a teaching value. At the inlet at 60 °C, $K_A p_A = 9.85$, $K_E p_E = 0.156$ and $K_{CO} p_{CO} = 1.000$ at 2 ppm.

Three behaviors follow from the squared denominator, and one does not. The first is self-inhibition. The ethylene rate at the inlet is 31.0 times lower than with the acetylene gone, as shown in Figure 3.2a, against the two orders of magnitude measured on a commercial catalyst [8]. Without carbon monoxide the factor is 90.7, because carbon monoxide caps the ethylene rate that is released when the acetylene runs out (at 5 ppm the factor is 13.7). The apparent order in acetylene is −0.64 at the inlet and 0.98 at 100 ppm, where the review gives orders around zero or slightly negative and first order at low ethyne [1]. The ethylene rate rises faster with temperature, by 17.3 times from 40 to 80 °C against 3.8 for the acetylene, so a hotter bed is a less selective one, an ordering the published fits share [8]. What the law cannot do is the measured selectivity rise with carbon monoxide. Both rates share D^2, so carbon monoxide divides them by one factor, and the local selectivity $r_A/(r_A + r_E)$ at 1 % acetylene is 0.9896 at 0, 2, 6 and 20 ppm of carbon monoxide. Carbon monoxide is a brake on the catalyst in this model and nothing more.

Pause. Try Exercise 3.2 before reading on.

3.3.2 Heats From Standard Data

The heats of reaction are $\Delta H_1 = -174.26\,\text{kJ/mol}$ and $\Delta H_2 = -136.47\,\text{kJ/mol}$, from the standard heats of formation of acetylene, ethylene and ethane at 298 K in the NIST Chemistry WebBook [17] (226.73, 52.47 and −84.0 kJ/mol). The review gives −172 and −137 MJ/kmol [1]. Heat capacities are held at their values near 60 °C from the same source, 46.66, 29.02, 46.40, 57.40 and 29.13 J/(mol K) for acetylene, hydrogen, ethylene, ethane and carbon monoxide, which makes the feed's heat-capacity flow $\dot{F} C_p = 68.65\,\text{kW/K}$ at design. Hydrogenating the whole 1 % of acetylene therefore heats the gas 35.54 K, and each mol/s of hydrogen that finds ethylene instead adds 1.99 K.

3.3.3 Beds of Cells on the Shared Module

Each bed holds 12 t of catalyst in 24 cells of 0.5 t on the framework's packed-bed module, `core/beds.js`. A cell's composition is algebraic, because the gas crosses it in well under a second, and its temperature is a state on the catalyst's heat capacity, as shown in Equation (3.3).

$$\begin{aligned} \mathbf{F}_{out} &= \mathbf{F}_{in} + \boldsymbol{\nu}^{\mathsf{T}} \boldsymbol{\xi}, \qquad \xi_j = W_c\, r_j(T, \mathbf{F}_{out}) \\ W_c\, c_{cat} \frac{\mathrm{d}T}{\mathrm{d}t} &= \dot{F} C_p\,(T_{prev} - T) - \textstyle\sum_j \xi_j \Delta H_j \end{aligned} \tag{3.3}$$

where $\boldsymbol{\nu}$ holds the two reactions over acetylene, hydrogen, ethylene, ethane and carbon monoxide, $W_c = 500\,\text{kg}$, $c_{cat} = 900\,\text{J/(kg K)}$, T_{prev} is the cell upstream or the bed's inlet, and $\dot{F} C_p$ is evaluated at the converter feed and kept constant through both beds. A cell holds 450 kJ/K, 6.55 s of the gas's heat flow, and a bed 10.8 MJ/K, 2.62 min. The extents are found by Newton's method from zero with a fixed number of iterations and no convergence test, which keeps the derivative a pure function of the state. This app sets the count to seven through the module's option. The app's development record shows that six left a residual of 4.5×10^{-10} of the cell's acetylene and hydrogen at 1.8 % acetylene, 2.5 mol/mol and inlets of 85 and 90 °C, and with seven the residual is at most 1.3×10^{-15} over 729

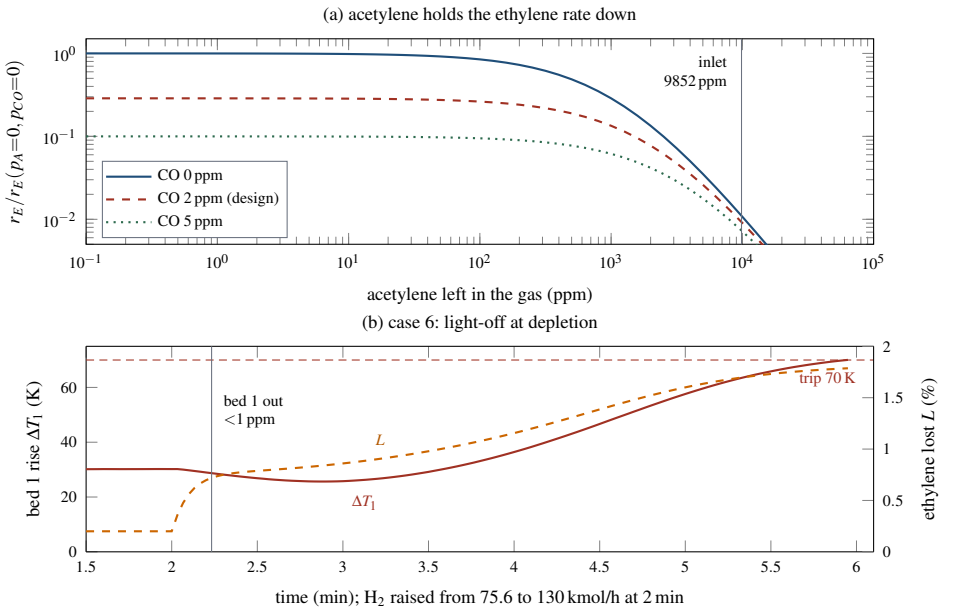

Figure 3.2 Self-inhibition and the light-off it hides. (a) The ethylene rate of Equation (3.1) at 60 °C and 20 bar against the acetylene left, relative to its value with neither acetylene nor carbon monoxide. At the inlet's 9852 ppm it is 31.0 times below its value at zero acetylene for 2 ppm of carbon monoxide. (b) Case 6 with the hydrogen raised by hand from 75.6 to 130 kmol/h at 2 min: bed 1's outlet falls under 1 ppm 14 s later, the loss L (dashed) climbs while the rise ΔT_1 first falls, and the trip comes at 5.95 min.

lined-out points of the Window view and 8.2×10^{-17} at every step of case 6 to its trip. Each of the 48 cells is a one-cell bed of the train, so that it carries its own activity, and the same record shows that a train of two 24-cell beds at uniform activity gives the same flows, heats and derivatives to the last bit.

The cell count is an axial dispersion, and the last ppm depends on it. At design the outlet is 1.01, 0.31 and 0.13 ppm with 12, 24 and 48 cells per bed, while the 1 ppm edge moves only from 1.501 to 1.433 to 1.401 mol/mol and the loss from 0.197 to 0.200 to 0.202 %. The difference reaches a verdict. The challenge's suggested settings, described in Section 3.5, pass at 24 and 48 cells (worst sample 0.53 and 0.49 ppm) and miss at 12, where two of 288 samples exceed 1 ppm (worst 1.39 ppm) and the loss averages 0.394 %. A verdict resting on tenths of a ppm belongs to the 24-cell model.

3.3.4 Green Oil and Activity

Each cell carries a green-oil loading g in weight percent of its catalyst, which grows with the acetylene there and sets the activity, as shown in Equation (3.4).

$$\frac{\mathrm{d}g}{\mathrm{d}t} = c_t\, k_g\, \mathrm{e}^{-\frac{E_g}{R_g}\left(\frac{1}{T}-\frac{1}{T_r}\right)} p_A, \qquad a = \mathrm{e}^{-g/g_0} \tag{3.4}$$

where p_A is the cell's outlet acetylene, $k_g = 0.4\,\mathrm{wt\%/(bar\,d)}$, $E_g = 40\,\mathrm{kJ/mol}$, $g_0 = 6\,\mathrm{wt\%}$, and c_t is a declared time compression, 1 except in case 7, where 1440 makes a minute a day. All three are teaching values.

3.3.5 Feed, Heater and Intercooler

The C_2 flow is $F_{C_2} = 1400\,\mathrm{mol/s}$ times the flow disturbance, and its acetylene and carbon monoxide reach the converter through the feed drum, as shown in Equation (3.5).

$$\begin{aligned} \tau_d \frac{\mathrm{d}y_A}{\mathrm{d}t} &= y_{A,up} - y_A \\ \tau_d \frac{\mathrm{d}p_{CO}}{\mathrm{d}t} &= p_{CO,up} + \Delta_c\, s(t/\Theta) - p_{CO} \end{aligned} \tag{3.5}$$

where $\tau_d = 20\,\mathrm{min}$, $y_{A,up}$ and $p_{CO,up}$ are the slider values, and in the challenge the furnace cycle adds $\Delta_c = 1.5\,\mathrm{ppm}$ times a trapezoid s of period $\Theta = 240\,\mathrm{min}$, low for the first half and high for the second with ramps over an eighth of the period. The bed 1 inlet is $T_{in,1} = T_{feed} + Q/\dot{F}C_p$ with $T_{feed} = 15\,^\circ\mathrm{C}$ and the duty Q following its command through a 0.5 min lag up to 5 MW. The intercooler passes a share $\beta = 0.9\,v_2$ of the gas past the exchanger, as shown in Equation (3.6).

$$\begin{aligned} G &= (1-\beta)\dot{F}C_p, \qquad \varepsilon = 1 - \mathrm{e}^{-UA_g/G} \\ T_x &= T_{out,1} - \varepsilon(T_{out,1} - T_m) \\ T_{in,2} &= \beta T_{out,1} + (1-\beta)T_x \\ C_m \frac{\mathrm{d}T_m}{\mathrm{d}t} &= G\varepsilon(T_{out,1} - T_m) - \phi\, UA_s(T_m - T_{cw}) \end{aligned} \tag{3.6}$$

with $UA_g = 60\,\mathrm{kW/K}$, $UA_s = 400\,\mathrm{kW/K}$, $C_m = 10\,\mathrm{MJ/K}$, $T_{cw} = 30\,^\circ\mathrm{C}$ and ϕ the cooling-water disturbance. The bypass follows its command with a 0.15 min lag and the hydrogen flow loop with 0.1 min, up to 56 mol/s. At design the bypass stands at 69.5 %, which passes 62.5 % of the gas.

3.3.6 Outputs and Instruments

The outlet acetylene is $10^6 F_{A,o}/F_{t,o}$ in ppm, and the ethylene loss is the ethane made as a share of the ethylene fed, $L = 100\,(F_{EA,o} - F_{EA,0})/F_{E,0}$. Thermocouples at each bed's inlet and outlet have a 6 s lag and 0.1 K of noise, and the C_2 and hydrogen meters 0.1 %. Both chromatographs run on the framework's analyzer module, draw a sample every 5 min and report it 3 min later, so a held reading is 3 to 8 min old. The inlet one reads acetylene (0.002 mol% noise) and carbon monoxide (0.03 ppm) ahead of the drum. The outlet one reads acetylene (0.02 ppm), hydrogen (0.005 mol%) and the loss (0.005 %). A chromatograph that has not reported for two periods reads STALE. The trip acts on either bed's true rise above 70 K, latches, closes the hydrogen and stops the run.

The exported file (Section 1.6.4) marks its simulator columns with a `truth_` prefix, out of reach of any controller replayed on it. In that file a column ending `_pct` is a percent of one of five things: the feed acetylene columns of the C_2 feed with hydrogen excluded, the hydrogen columns of the product gas, the loss columns of the ethylene fed, `heater_pct` of the full 5 MW and `bypass_pct` of valve travel, with the flow and cooling-water sliders in percent of design. The Instructions page lists each column's denominator.

3.3.7 Parameters, Initialization and Numerics

Table 3.2 lists the parameters. Every case starts lined out at its controllers' operating point, with the heater and bypass found for the inlet setpoints and the beds settled by the module's steady-state solver. With the hydrogen in Auto, a ratio that would line a bed out above the override's limit starts the plant where the override would hold it, which keeps a setting made before Run from starting an Auto plant past the trip. Classical fourth-order Runge–Kutta advances the state in 1 s steps, control runs every 2 s and a row is logged every 15 s, on the engine's three clocks of Section 1.6.1. The fastest speed runs 60 min of plant per second, and the engine needs 12.7 ms of computation per simulated minute headless in Node on the build host. Lined out, carbon and hydrogen close to 1.5×10^{-15} and 1.2×10^{-15}, and the engine's bed rises match the hand calculation from the conversions, 30.160 and 9.765 K, to the printed digits. Halving the step leaves every challenge verdict and every printed figure unchanged.

Table 3.2 Parameters. A: Adjustable in the Model Dialog. H: Hidden

Symbol	Meaning	Value	
$k_{A,r}$, $k_{E,r}$	rate constants at 60 °C	3, 4×10^{-4} mol/(s kg bar^2)	H, A
E_A, E_E	activation energies	25, 60 kJ/mol	H
K_A, K_E	adsorption constants	50, 0.01 bar^{-1}	H
$K_{CO,r}$, Q_{CO}	CO adsorption, its heat	2.5×10^4 bar^{-1}, 30 kJ/mol	A, H
P	reactor pressure	20 bar	H
W, c_{cat}	catalyst per bed, heat capacity	12 t, 900 J/(kg K)	H
k_g, E_g, g_0	green oil	0.4, 40 kJ/mol, 6 wt%	H
c_t	green-oil time compression	1	A
T_{feed}, Q_{max}	feed temperature, heater duty	15 °C, 5 MW	H
UA_g, UA_s	intercooler gas, water side	60, 400 kW/K	H
C_m, T_{cw}	intercooler metal, water	10 MJ/K, 30 °C	H
b_{max}	largest bypassed share	0.9	H
$F_{H_2,max}$	hydrogen valve capacity	56 mol/s	H
τ_Q, τ_v, τ_H	heater, bypass, hydrogen lags	0.5, 0.15, 0.1 min	H
τ_d	feed drum	20 min	H
Δ_c, Θ	furnace cycle	1.5 ppm, 240 min	H
—	ΔT trip	70 K	A
—	product acetylene limit	1 ppm	A

Table 3.3 Loops, Pairings and Default Tunings

Loop	PV → MV	K_c	τ_I	Output, period
TIC 1	$T_{in,1} \to Q$	1.5 %/K	0.6 min	0–100 %, 2 s
TIC 2	$T_{in,2} \to v_2$	3 %/K	1 min	0–100 %, 2 s
AIC trim	$\log_{10} A_{out} \to \Delta R_{trim}$, rev.	0.05/decade	5 min	±0.6, each report
override	$\max \Delta T \to u_{ov}$	0.3 mol/(s K)	3 min	0–56 mol/s, 2 s

3.3.8 Assumptions

The plant has two reactions on one kind of site with teaching constants, so carbon monoxide brakes the catalyst without raising its selectivity, which a real palladium catalyst does [8]. The direct route from acetylene to ethane, the oligomers other than the green oil and the hydrogenation of the green oil are left out, so carbon and hydrogen close on the two reactions exactly. The gas is ideal at 20 bar with constant heat capacities, there is no pressure drop, heat loss or internal diffusion, and the gas holds neither heat nor composition. Each bed is 24 well-mixed cells, an axial dispersion much larger than a real bed's, and the feed drum is one first-order lag. The hydrogen is pure and its flow loop ideal, and the analyzers carry noise but no drift. The 1 ppm limit, the 55 K override, the 70 K trip and the challenge's 0.33 % loss budget are teaching values. A real limit comes from the ethylene buyer's specification, which the review puts below 5 ppm [1], and a real trip from the catalyst supplier's and the vessel designer's temperature limits, implemented in the plant's safety instrumented system.

3.4 Ratio, Trim, Override and a Predictive Layer

The instrument code hands the controllers one packet of readings, and a test holds them to its keys. The hydrogen loop and the two inlet loops each carry their own Auto and Manual switch, the header's switch throws all three, and each idle controller follows the one in charge, so a switch of mode or of structure causes no bump.

3.4.1 The Shared PI Law

Every PI loop here is the framework's velocity-form law of Section 1.6.2, Equations (1.14) to (1.17), used without derivative action. Only the emitted output is clamped, and the integral is corrected toward the clamped output with a tracking time equal to τ_I. An integral time set to its slider's maximum removes the integral action. The loops are listed in Table 3.3.

3.4.2 Hydrogen in Ratio, With Two Feedforwards

The hydrogen flow is the lower of the ratio logic's request and the override's output, as shown in Equation (3.7).

$$\begin{aligned} F_{H_2} &= \min\bigl(R\,F_{C_2,m}\,\bar{y}_A,\; u_{ov}\bigr) \\ R &= \mathrm{clip}_{[0.8,\,2.5]}\bigl(R_{SP} + \Delta R_{trim} + K_{CO}(\bar{p}_{CO} - 2)\bigr) \end{aligned} \tag{3.7}$$

where $F_{C_2,m}$ is the C_2 meter. On Fixed ratio $\bar{y}_A$ is the design 1 % and there is no carbon monoxide term. On Feedforwards $\bar{y}_A$ and $\bar{p}_{CO}$ are the inlet chromatograph's readings passed through a first-order lag $\tau_{ff} = 20\,\mathrm{min}$, the drum's, and $K_{CO} = 0.15\,\mathrm{mol/mol}$ per ppm. The analyzer samples ahead of the drum, so the feedforward sees a change about 20 min before the converter does, while the outlet analyzer sees its effect only after the drum, the beds and its own delay. The app's development record shows that without the drum, with both analyzers at 5 and 3 min, the feedforward saw an upstream change no sooner than the trim and every structure that passed the challenge with one passed with the other, so the drum is what gives the feedforward its lead.

The trim is reverse acting on $\log_{10}\max(0.05, A_{out})$ against $\log_{10} A_{SP}$ and acts once per outlet report with the analyzer's 5 min period as its step. The logarithm suits a plant whose outlet moves a decade for a tenth of a mol/mol near the edge. The floor of 0.05 ppm is the analyzer's noise floor (its noise is 0.02 ppm), and the development record shows that without it a zero reading drove a limit cycle with a 10 min period between 0 and 8.7 ppm in the challenge. The override is a PI on the larger measured bed rise whose output is a hydrogen flow. While the ratio's request is lower it tracks that request, so it takes over at its limit and hands back without a bump. During development the tuning was settled at a gain of 0.3 mol/(s K) with 3 min after a gain of 0.5 with 1 min formed a growing cycle with the bed's thermal lag that tripped, and the limit at 55 K after a 65 K limit let an Auto run with 2 % acetylene reach the trip. A protective override of the same kind, a PI loop that acts only near its limit, guards the compressor of Volume 4, Chapter 2.

3.4.3 A Balance Inference for a Failed Analyzer

When the inlet chromatograph fails, the feedforward can take the feed acetylene from a hydrogen balance on each outlet report, as shown in Equation (3.8).

$$\begin{aligned} \xi_H &= \frac{F_{H_2,m} - x_{H_2}(F_{C_2,m} + F_{H_2,m})}{1 - x_{H_2}} \\ \xi_E &= \frac{L}{100}\,0.79\,F_{C_2,m} \\ \hat{y}_A &= \frac{\xi_H - \xi_E + 10^{-6} A_{out}(F_{C_2,m} + F_{H_2,m} - \xi_H)}{F_{C_2,m}} \end{aligned} \tag{3.8}$$

where the meters are taken at the instant the report's sample was drawn and x_{H_2} is the outlet hydrogen fraction. Every mole of hydrogen used made ethylene from acetylene or ethane from ethylene, so the hydrogen used less the ethane made, plus the acetylene left, is the acetylene fed. A bias against the inlet chromatograph is updated while that one reports and frozen when it fails, and lined out at 1.0, 1.3 and 1.5 % the inference is within 9.0×10^{-7} percentage points of the truth. A bed-temperature inference would be simpler, but it counts the heat of ethylene burned as acetylene converted. The development record shows its feedforward running away into the trip in case 5 at 33.2 min, which is why the balance replaced it. This is inferential control in the sense of secondary measurements standing in for an unmeasured quality [15].

3.4.4 The MPC on the Ratio and Both Inlets

With the architecture set to MPC, every 2 min a receding-horizon controller [16] of the kind set out in Section 1.6.3 moves R_{SP}, $T_{1,SP}$ and $T_{2,SP}$ for the least ethylene loss. It lines out its own 24-cell copy of the beds, built from the declared parameter defaults with fresh catalyst, for the feed the instruments

Table 3.4 The Nine Cases

#	Case and events	Teaching point
1	Fixed ratio from 1.3, learner steps it to 1.9	the window
2	Design point, learner raises $T_{2,SP}$ to 62 °C	the window map
3	Ratio 1.3, trim on, learner raises its gain	analyzer trim
4	Fixed ratio 2.0, CO 2 → 0.5 ppm at 10 min	override, CO feedforward
5	Acetylene 1.0 → 1.5 % at 10 min	feedforward
6	Hydrogen on Man, learner raises it to 130 kmol/h	runaway, trip
7	Trim on, green oil 1440 times faster	aging, schedule
8	Feedforwards, inlet analyzer fails at 10, acetylene 1.3 % at 20 min	balance inference
9	Challenge, 24 h, fixed ratio, no trim	the whole converter

report, fits the gains of four outputs (log ppm, loss and both rises) by central differences of ±0.05 mol/mol and ±1 K, and plans ten decisions ahead, as shown in Equation (3.9).

$$\begin{aligned} x_{j,k+1} &= a_j x_{j,k} + (1-a_j)u_{j,k}, \qquad a_j = \mathrm{e}^{-T_s/\tau_j} \\ \min\ & \textstyle\sum_k w_L L_k + \sum_k \Delta\mathbf{u}_k^{\mathsf{T}} S\, \Delta\mathbf{u}_k + \text{soft penalties} \\ \text{s.t. } & \log_{10} A_k \le \log_{10} A_{SP}, \quad \Delta T_{b,k} \le \Delta T_{lim} - 5 \end{aligned} \tag{3.9}$$

where both output limits are soft, $T_s = 2$ min, τ_j is the hydrogen loop's lag for the ratio and the bed's $C_{bed}/\dot{F}C_p$ for the inlets, $w_L = 1$ per percent, S is $100\,w_{\Delta u}$ for the ratio and $w_{\Delta u} = 0.5$ for each inlet, the inputs are bounded to 0.8–2.5, 25–85 and 35 °C–90 with moves of at most 0.1, 2 and 2 K, with the moves held over blocks of one, one and eight decisions. The outlet chromatograph corrects the model's log ppm and loss by half the mismatch at each report, the thermocouples its rises at each decision. Each program is solved by the dual algorithm of Goldfarb and Idnani [18] as coded in the framework's quadratic programming module, and a failed plan leaves the setpoints where they were.

3.4.5 Three Models Kept Apart

The simulated plant is the 48-cell model above. Any model a learner identifies from the logged data, the edge against carbon monoxide for one, is a second and separate model. The controllers' models are the ratio, the trim gains, the balance and the MPC's fresh-catalyst copy. Not one of them can read the activity, the true outlet between reports or the true carbon monoxide. A parameter changed in the app reaches the plant and not the MPC's copy, which meets it as mismatch through the chromatograph and thermocouple corrections. The status strip, the window map and the challenge's judge read the truth and act on nothing.

3.5 Measured on the Running Converter

The cases are listed in Table 3.4 and their outcomes in Table 3.5. Every number was measured with the engine run headless at its 1 s step on the default seed, with every case in its shipped form and every learner action made after Run at the instant the text gives. Rerun when these results were written up, the script of the app's technical review reproduced its recorded output exactly, and the challenge was run at four further seeds and at half the step.

3.5.1 The Window and Its Two Costs

Case 1 steps the fixed ratio from 1.3 to 1.9. About 14 min after each step the outlet held 14.7, 1.89, 0.30, 0.010 and 0.0002 ppm while the loss read 0.138, 0.177, 0.201, 0.275 and 0.472 %. The acetylene fell 4.97 decades and the loss rose 3.43 times, and not evenly. The last 0.2 mol/mol cost 0.198 points of ethylene, 3.14 times what the step from 1.3 to 1.5 cost, and lined out the loss climbs 0.40 % per mol/mol between 1.3 and 1.4 and 1.34 between 1.9 and 2.1. The two failures are priced very differently. Breakthrough puts all 2680 t/d of ethylene off grade, while each 0.1 % of loss is 2.68 t/d, so the window is one-sided and the right place is just clean of the edge, as shown in Figure 3.3.

Table 3.5 Case Outcomes Measured on the Engine

Case and condition	Result
1, ratio 1.3/1.4/1.5/1.7/1.9	14.7/1.89/0.30/0.010/0.0002 ppm; loss 0.138/0.177/0.201/0.275/0.472 %
2, $T_{2,SP}$ 58 → 62 °C	edge 1.433 → 1.375; at 1.5: 0.31 → 0.073 ppm, 0.200 → 0.275 %
3, trim gain 0.05	under 1 ppm at 13.3 min (analyzer 18.3); 0.30–0.37 ppm
3, gain 0.3 at 60 min	0.001–6.7 ppm, ratio 1.37–1.76 in the next hour
4, fixed ratio 2.0	loss 0.594 → 1.052 %; override at 40.0 min
4, Feedforwards at 40.8 min	ratio 1.78; override off 1.7 min later; loss 0.76 %
5, fixed ratio	over 1 ppm 2.0 min after; peak 2476 ppm; over 108 of 110 min
5, Feedforwards	peak 1.57 ppm, 2.5 min over; H_2 75.7 → 113.3 kmol/h; loss 0.39 %
6, 130 kmol/h at 2 min	bed 1 out <1 ppm +14 s; loss >1 % +1.58 min; trip +3.95 min at 70.05 K, 111.0 °C, 1.79 %
7, no schedule	over 1 ppm day 137.5; trim at its reach day 193.3; 3830 ppm day 300
7, 5 K per 100 days	over 1 ppm day 177.8; loss 0.446 against 0.334 % on day 150
8, left alone	STALE 18.25 min; over 1 ppm from 23.75; 212 ppm at 2 h
8, inference at 30 min	under 1 ppm 8 s later; loss 0.300 % at 2 h

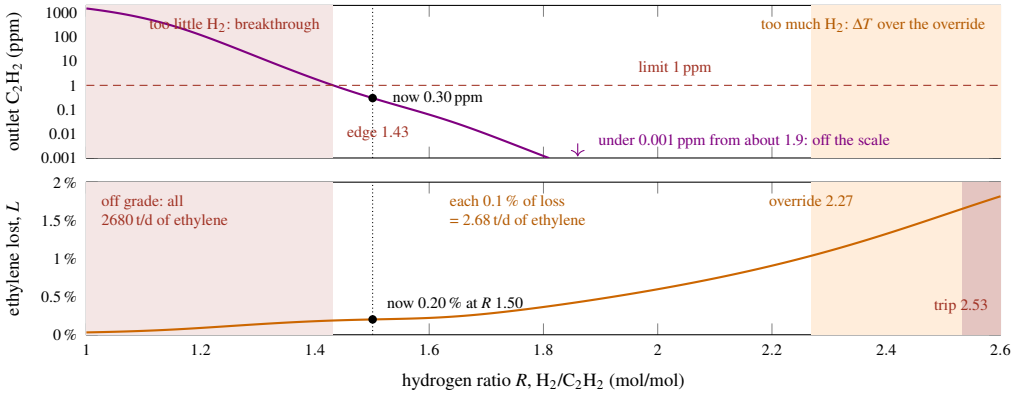

Figure 3.3 The Window view of case 2 at $t = 1$ min, the instant of Figure 3.1: the plant lined out at every ratio with these inlets, carbon monoxide and catalyst. The breakthrough edge is at 1.43, the override is reached at 2.27 and the trip at 2.53, with the two costs.

Further right the lined-out bed 1 reaches the 55 K override at a ratio of 2.27 and the 70 K trip at 2.53. The app's 17-point map places the three crossings at 1.434, 2.268 and 2.534 against 1.432, 2.269 and 2.533 by direct bisection.

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

The window moves. In case 2, warming bed 2's inlet from 58 to 62 °C moves the edge from 1.433 to 1.375, but at the plant's own 1.5 the outlet falls only from 0.31 to 0.073 ppm while the loss rises from 0.200 to 0.275 %. Carbon monoxide moves the edge more, from 1.253 at 0.5 ppm to 1.433 at 2 and 1.981 mol/mol at 5 ppm, while the loss at the edge stays within 0.176 to 0.189 %, which is the one-site law's signature. The design point is also not the cheapest point in the window. Colder inlets favor the acetylene's smaller activation energy, and at 27.5 and 40 °C the edge is at 2.079 with a loss there of 0.112 %, against 0.186 % at the design inlets. A 2.5 K grid over the setpoint ranges finds 0.108 % at 27.5 and 35 °C. The app says the design point is one choice in the window and poses the colder, fatter point as a question on its Instructions page.

> **TRY IT IN THE APP** **2 · The window map and the operating point**
>
> Load the case and press Run. On the Window view, find where the outlet acetylene curve crosses the red 1 ppm line, near a ratio of 1.43. The black dots of the plant sit just to its right, at 1.5. Then raise the Bed 2 inlet setpoint to 62 °C and watch the crossing move left, to 1.375, while the loss at the plant's ratio rises from 0.200 to 0.275 %.

3.5.2 A Trim That Waits on Its Analyzer

Case 3 starts at a ratio of 1.3 with the trim on. The default trim took the outlet from 15.0 to under 1 ppm in 13.3 min, which the chromatograph showed at 18.3 min, and held it at 0.30 to 0.37 ppm with the ratio near 1.50. Raised six times to 0.3 at 60 min, the trim hunted, the outlet swinging between 0.001 and 6.7 ppm and the ratio between 1.37 and 1.76 over the next hour, because each report corrects a sample drawn before the last correction took effect. At a gain of 0.2 the outlet stays at 0.25 to 0.36 ppm.

3.5.3 Carbon Monoxide, Feed and the Override

Case 4 runs fat at 2.0 on a fixed ratio, and at 10 min the carbon monoxide falls from 2 to 0.5 ppm. The outlet acetylene never matters (below 0.000 03 ppm in the run with the switch), but the catalyst burns more of the spare hydrogen on ethylene, the loss rises from 0.594 to 1.052 % and bed 1's rise from 44.7 to 55 K, where the override takes over at 40.0 min and begins to cut the hydrogen from 100.9 kmol/h. Switched to Feedforwards at 40.8 min, the carbon monoxide term takes the ratio to 1.78, the override stands by 1.7 min later, and the loss settles at 0.76 %. The override keeps the beds safe and leaves the loss high. In case 5 the feed acetylene rises half again at 10 min. On a fixed ratio the true ratio falls to 1.00, the outlet crosses 1 ppm 2.0 min after the change, which the chromatograph reports at 8.25 min, and peaks at 2476 ppm. With the feedforward the hydrogen follows from 75.7 to 113.3 kmol/h, the outlet peaks at 1.57 ppm for 2.5 min and the chromatograph never reads above 0.82 ppm, at a loss of 0.39 % because the same ratio on more acetylene is more spare hydrogen.

3.5.4 Light-Off at Depletion and Where the Trip Is Reached

Case 6 puts the hydrogen in the operator's hands, where the override does not act, as shown in Figure 3.2b. Raised from 75.6 to 130 kmol/h (2.58 mol/mol) at 2 min, the hydrogen empties bed 1 of acetylene in 14 s, and the sites the acetylene held go to ethylene. The loss reaches 0.79 % at 2.5 min and passes 1 % 1.58 min after the change. Bed 1's rise first falls, to 25.65 K at 2.87 min, because the reaction has moved to the front cells, whose heat reaches the outlet only after a thermal wave crosses the bed in about 2.6 min. It then climbs to 70.05 K 3.95 min after the change, with the outlet at 111.0 °C, 1.79 % of the ethylene going to ethane and more ethylene burned than made. The jacketed reactor of Volume 7, Chapter 2 shows the same contest between heat that grows exponentially with temperature and removal that grows only linearly, at an ignition boundary. Lined out by hand, 110, 120, 125 and 130 kmol/h give bed 1 rises of 51.1, 60.9, 66.8 and 72.8 K. A plant started at 130 kmol/h, set before Run or kept through Reset, is already past the trip and trips in its first second, which the case's texts say. The lowest flow that completes case 6 is 128 kmol/h, which trips 4.57 min after it is set, while 125, 126 and 127 kmol/h peak at 67.0, 68.2 and 69.4 K in 40 min and never trip.

Pause. Try Exercise 3.3 before reading on.

> **TRY IT IN THE APP** **6 · The runaway: acetylene runs out in bed 1**
>
> Press Run, then raise the Hydrogen flow on the FIC card from 75.6 to 130 kmol/h. Watch the Bed temperature rise trend. Bed 1's rise first falls, to 25.65 K, while the ethylene loss passes 1 % 1.58 min after the change. The rise reaches the 70 K trip only 3.95 min after it, and the trip closes the hydrogen.

At the design C_2 feed no Auto setting reaches the trip. Over sixteen settings and disturbances held 90 min (ratio 2.5, the override at its 60 K top, 1.8 % acetylene, no carbon monoxide, a lost intercooler,

Table 3.6 The Challenge Over 24 h, Default Seed and Seeds 1–4

Configuration	Worst (ppm)	Over /288	Loss (%)	
fixed ratio 1.5, no trim (shipped)	1079	185	0.192	fail
fixed ratio 2.0	0.29	0	0.566	fail
trim alone at 0.07 ppm, gain 0.1	0.70	0	0.374	fail
trim alone at 0.085 ppm, gain 0.1	1.08	1	0.317	fail
Feedforwards at 1.6, no trim	0.28	0	0.455	fail
Feedforwards, trim at 0.3 ppm	1.36	12	0.242	fail
Feedforwards, trim at 0.1 ppm (sugg.)	0.53	0	0.297	pass
seeds 1–4	0.52–0.57	0	0.297–0.306	pass
same, 12 cells per bed	1.39	2	0.394	fail
same, feed acetylene from inference	0.78	0	0.276	pass[a]
MPC, A_{SP} 0.1 ppm	0.64	0	0.354	fail
seeds 1–4	0.48–0.67	0	0.352–0.356	fail

[a] Fails at seed 1 (1.10 ppm, two samples over) and passes at seeds 2 to 4 (0.94 to 0.99 ppm).

inlets at 85 and 90 °C), the largest true rise was 66.2 K. A smaller or richer feed moves the trip line left, to 2.379 at 85 % of the design flow, 2.281 at 70, 2.250 at 50 and 1.581 mol/mol at 1.8 % acetylene, and a ratio past it is held only by the override, which reads the bed outlets and a light-off can outrun. In Auto a ratio setpoint of 2.5 with the C_2 feed cut to 85 % at 20 min trips at 23.57 min, 2.2 with a cut to 70 % at 24.57 min, and Feedforwards at the shipped 1.5 with 1.8 % acetylene and a cut to 50 % at 64.98 min, while 2.2 at 85 % peaks at 64.5 K and 2.0 at 50 % at 67.4 K. A mistuned override can let a fat ratio through too. With a ratio setpoint of 2.5 and the carbon monoxide gone, a gain of 3 with 0.2 min trips at 68.93 min and a proportional gain of 0.05 at the 60 K limit at 42.10 min. The app's technical review tried an override scaled by the C_2 meter, a rate-limited hand-back and a ratio ceiling of 2.0, and none removed the feed-cut trips, so the trip is reachable in Auto by design of the physics, and the trip exists for exactly this.

3.5.5 Green Oil, the Analyzer Failure and an Honest Infeasibility

Case 7 compresses the run, a minute for a day. Bed 1's mean activity falls to 0.54 by day 100 and 0.30 by day 200 (bed 2's 0.86 and 0.64), the trim raises the ratio to 1.65 by day 100 and 1.86 by day 150, the outlet first passes 1 ppm on day 137.5, and the trim reaches its 0.6 mol/mol limit on day 193.3. A schedule of 5 K per 100 days on both inlets holds the outlet under 1 ppm to day 177.8 at a loss of 0.446 % on day 150 against 0.334 without it. On the catalyst left at day 480 (bed activities 0.061 and 0.143) no point meets the limit. A search over the full setpoint ranges with both rises under 55 K, refined to 0.5 K and 0.01 mol/mol, finds at best 9.31 ppm at 85.0 and 88.5 °C and a ratio of 2.49, and allowing the override's 60 K top gives 6.44 ppm. The catalyst must be regenerated. In case 8 the inlet chromatograph fails at 10 min and reads STALE at 18.25 min, the acetylene rises to 1.3 % at 20 min, and the outlet passes 1 ppm at 23.75 min and reaches 212 ppm at 2 h if left. Switched to the inference at 30 min, the outlet is under 1 ppm 8 s later, because the balance had followed the rise all along (1.159 percent at 40 min, 1.295 at 2 h).

3.5.6 A Day on the Furnace Cycle

The challenge runs 24 h. The carbon monoxide swings between 0.5 and 3.5 ppm every 4 h, the acetylene rises to 1.3 % at 6 h, the inlet chromatograph fails at 14 h and the acetylene returns to 1.0 % at 16 h. It is met when every one of the 288 outlet samples carries at most 1 ppm of true acetylene, the true loss averages at most 0.33 % (8.8 t/d) and nothing trips. The budget was set at 0.33 rather than the development plan's 0.5 % because at 0.5 the trim alone passed by aiming deep inside the limit, so the feedforward was not needed. The configurations are compared in Table 3.6, and Figure 3.4 draws the nine run at the default seed against the limit and the budget.

As shipped, the first rise of carbon monoxide breaks through, and the day ends with 185 of 288 samples over and a worst sample of 1079 ppm. A higher fixed ratio fails the other way, and the trim

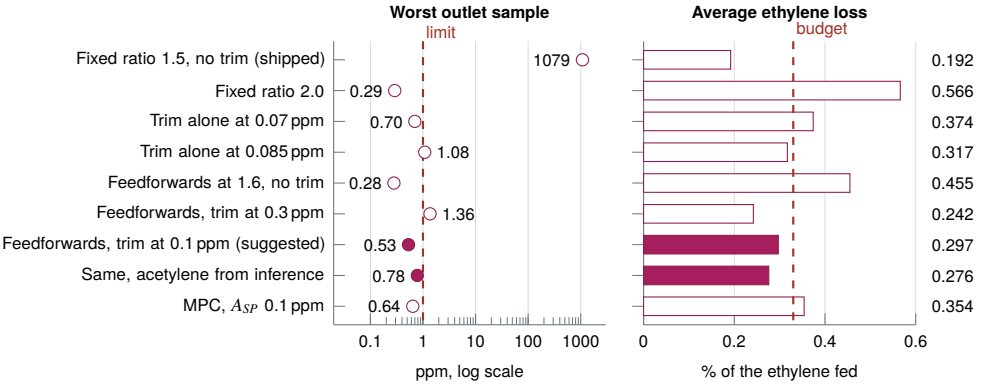

Figure 3.4 The challenge of Table 3.6 at the default seed. Filled marks and bars pass, open ones fail. The dashed lines are the 1 ppm limit on every sample and the 0.33 % budget on the average loss. Notice that only the two configurations with the feedforwards and a trim aimed at 0.1 ppm sit left of both lines, and that the MPC holds every sample but spends more than the budget.

alone holds the samples only aimed at 0.07 ppm, spending 0.374 %. With the feedforwards, the trim on and its setpoint at 0.1 ppm, the worst sample is 0.53 ppm and the loss averages 0.297 %, and the shipped and suggested verdicts hold at seeds 1 to 4 and at half the step. The MPC holds every sample but spends 0.352 to 0.356 %. It made 719 decisions with one failed plan in the day and, after the first hour, kept the inlet setpoints at a mean of 36.4 and 44.7 °C with the ratio setpoint at 1.71, the colder and fatter direction, but its loss enters as a linear cost on a fresh-catalyst copy and the plan does not find the suggested loop's economy.

> ✦ **DID YOU KNOW? Aiming Far Inside the Limit Is Not Enough**
>
> From the simulator (Table 3.6). On the day of the challenge, an outlet trim working alone and aimed at 0.085 ppm, far inside the 1 ppm limit, still let one of the 288 samples through, at 1.08 ppm. With the feedforwards added, a trim aimed at 0.1 ppm kept the worst sample at 0.53 ppm.

The outage leaves a lesson of its own, as shown in Figure 3.5. At 1000 min the suggested configuration still reads the inlet chromatograph's last report, 1.299 % acetylene and 3.0 ppm of carbon monoxide, while the true feed carries 1.041 % and 0.71 ppm. The ratio in force is 1.44, made of the setpoint 1.50, 0.15 mol/mol from the held carbon monoxide and a trim of −0.208, but on the true acetylene the ratio is 1.80, 0.54 mol/mol above the edge of 1.26 that the low carbon monoxide has set and only 0.09 below the override's 1.89. The outlet is at 0.000 007 ppm, the loss at 0.81 %, bed 1's rise at 51.1 K, and the trim is walking the ratio back one report at a time. The day still ends with the average under budget, 0.276 % at that instant.

3.6 Technology in Practice

Figure 3.6 draws the hydrogen ratio, its trim and override, the two inlet loops and the trip of Section 3.4 with the hardware of a converter in place of the schematic's symbols.

What Measures It

The controlled quality is the acetylene left in the product, a few parts per million read by a chromatograph on the outlet. A second chromatograph reads the acetylene and the carbon monoxide of the feed ahead of the drum, and thermocouples read each bed's inlet and outlet. The chapter's two chromatographs draw a sample every 5 min and report it 3 min later, with 0.02 ppm of noise on the



Figure 3.5 The challenge on its suggested settings (Feedforwards, trim on at 0.1 ppm) at $t = 1000$ min, on the Window view, 2.7 h after the inlet chromatograph failed. (1) The badges read the outlet chromatograph at 0.00 ppm, a loss of 0.81 %, a measured ratio of 1.44 and the held 3.0 ppm of carbon monoxide, and (5) the status strip reads ANALYZER STALE. The window, lined out at inlets of 41 and 59 °C, 1.04 % acetylene and 0.7 ppm of carbon monoxide, puts the plant at (3) a true ratio of 1.80, 0.54 above (2) the edge of 1.26, where the outlet curve crosses the 1 ppm line, and 0.09 below (4) the override at 1.89, where the shaded band begins, with the trip at 2.22 and the loss 0.809 % by the chromatograph.

outlet acetylene, and its thermocouples lag by 6 s. Figure 3.7 shows how a chromatograph separates its sample.

> **TECHNOLOGY** **Sensor: Process Gas Chromatograph**
>
> A process gas chromatograph injects a precise quantity of sample at the head of a long, narrow column, and a steady flow of carrier gas washes it through. Each compound leaves at its own time, so a detector that only tells carrier gas from anything else can measure many species. Hydrocarbons ionize easily in a flame, which suits a flame ionization detector, while carbon monoxide does not ionize and needs another kind. The area under each peak, scaled by a factor found with a calibration gas, gives the concentration. The weakness is time: cycles of minutes are not unusual, and that dead time can make a loop oscillate. Many faults start at the sample valve, where a leak can shift the baseline or corrupt the analysis [19]. A failed chromatograph in the chapter holds its last report, a flat trend line.

What Moves It

Three valves carry the handles. The hydrogen valve sets the ratio, the steam valve of the heater sets bed 1's inlet, and the bypass valve around the intercooler sets bed 2's. The chapter gives the hydrogen loop a lag of 0.1 min up to 56 mol/s, the heater a lag of 0.5 min up to 5 MW and the bypass a lag of 0.15 min. A fourth final element belongs to the protection rather than to control.

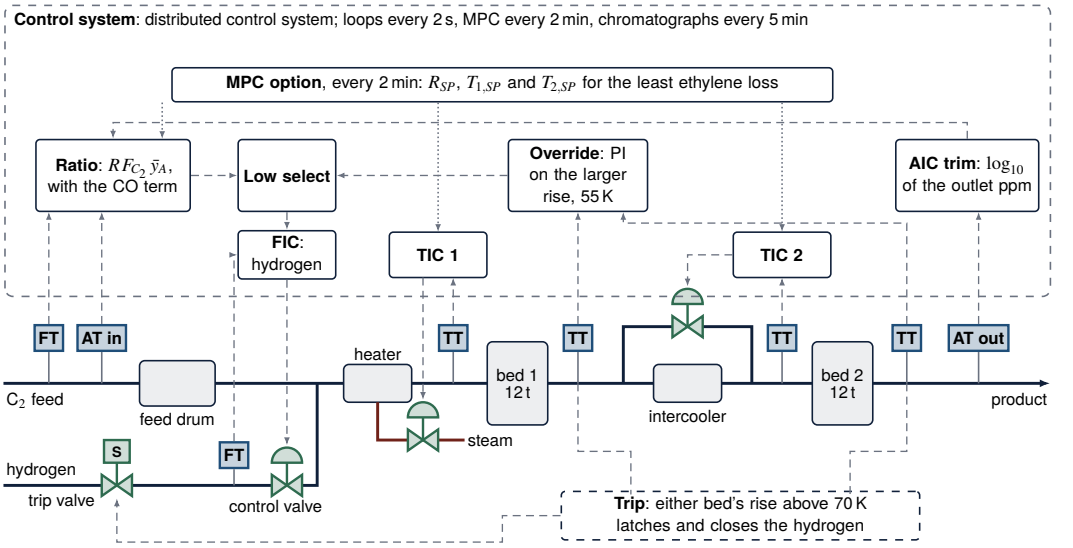

Figure 3.6 The converter's loops with their hardware. Blue boxes are sensing devices: FT on the C_2 feed and on the hydrogen, AT in and AT out, the two chromatographs, and TT at each bed's inlet and outlet. Green symbols are final elements: the hydrogen control valve, the hydrogen trip valve, the steam valve of the heater and the intercooler's bypass valve. Dashed arrows are signals, dotted arrows the setpoints that the MPC option writes. Notice that the override and the trip both read the bed outlets, and that the only reading of the outlet acetylene arrives every 5 min, 3 min late. Not to scale, and not a piping diagram.

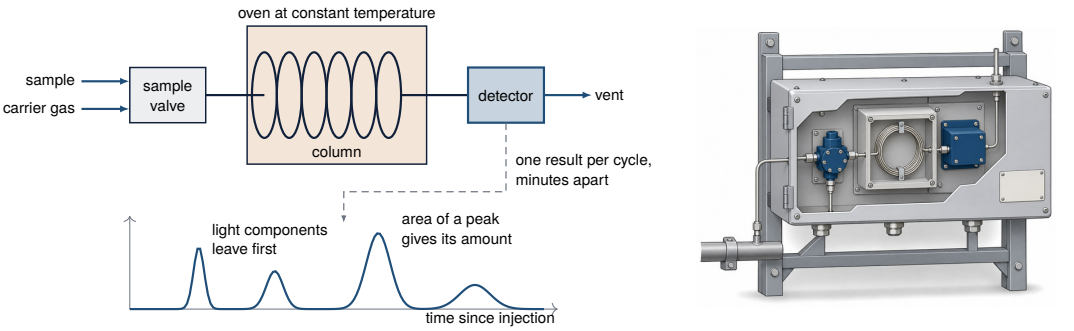

Figure 3.7 Principle of a process gas chromatograph, the instrument behind AT in and AT out. A valve injects a sample into the carrier gas, the column holds back each species for its own time, and the detector sees them leave one by one. Notice that a result exists only at the end of each cycle, which is why the chapter's chromatographs report each sample 3 min after it is drawn and only every 5 min.

TECHNOLOGY Final Element: The Hydrogen Valve and Its Trip

The hydrogen flow controller moves a modulating control valve. A trip that closes the hydrogen belongs to a safety instrumented system, built of its own sensors, a logic solver kept apart from the regulating controls so that one hardware fault cannot disable both, and final elements: on/off valves, or solenoids that force the control valve shut. The valves are generally the least reliable part. A full stroke would shut the unit down, so plants prove them with a partial stroke: a valve that moves well part of the way will very probably move all of it [19]. The chapter's trip reads the true bed rise, latches and closes the hydrogen.

Where the Controller Runs

A distributed control system has a supervisory level that sends set points to the field controllers, its loops are sometimes nested or cascaded, and their cycle times range from milliseconds to minutes [20]. The app's loops execute every 2 s, its MPC every 2 min, and each chromatograph reports every 5 min. The structure maps onto standard blocks: a ratio station with two feedforwards, a PI trim on the logarithm of the outlet acetylene that acts once per report, and a low selector between the ratio's request and the override. An override strategy selects among controller outputs so that only one at a time acts on the process [19]. In a plant the trip runs on hardware of its own.

What Is Modernizing It

Three developments touch this chapter. The first is the analyzer itself. Modern process chromatographs are smart instruments that compute their own results, can calibrate themselves from bottles of calibration gas on a schedule, and send many species over a digital bus [19]. The second is the soft sensor. Analyzers such as gas chromatographs suffer from long measurement delays and high investment and maintenance costs, and one purpose of a soft sensor is to validate an online analyzer against its estimate [21]. The balance inference of Section 3.4 is such an estimate, built from a hydrogen balance. The third is the layer above the loops. Crackers use multivariable control with online analyzers and constraint controls, often with online optimization [7], and a 2003 survey found predictive control well founded in refining and petrochemicals [16]. The predictive controller of Section 3.4 is a small example, and Table 3.6 shows it spending more ethylene than the suggested loops.

> **IN PRACTICE** **Regenerating the Catalyst**
>
> The alkynes in cracked gas are hydrogenated in the majority of European plants, and only two reported not doing it. The catalyst, which may hold palladium and silver among other metals, has a finite life. Most plants regenerate it on site, by thermal oxidation, by purging with steam and air, or by treating it with hydrogen, and send the gas from the regeneration to an abatement system. How often depends on the catalyst, the design and operation and the cracker's feed. A spent catalyst normally goes back to its supplier, who recovers the metal. The European reference document under Directive 2010/75/EU sets out these practices [7]. Case 7 of this chapter ends where only a regeneration helps.

3.7 Limitations

The converter is generic and its sizes, constants and limits serve teaching. The rate law is the classical one-site form, which a commercial palladium catalyst does not follow. Carbon monoxide here slows both reactions by one factor and leaves the loss at the edge between 0.171 and 0.189 % from 0 to 5 ppm, where measurement shows the selectivity rising from 0.3 to 0.6 to 0.7 to 0.9 [8]. A second kind of site would be needed, and it would move every number in the app. The self-inhibition at the inlet is 31 times where measurement gives about a hundred. The 24 cells set the last ppm, so a sub-ppm figure and the challenge's margin belong to this discretization. The green oil, the feed drum and the furnace cycle are teaching closures.

OPEN CHALLENGE A Catalyst That Stays Selective

Supported palladium remains the industrial catalyst of choice. Yet on it acetylene is also hydrogenated past ethylene to ethane, and its molecules couple and grow into green oil, which costs selectivity and activity. A 2023 review calls weakening the adsorption of ethylene and suppressing that coupling, to reach high selectivity with little green oil, the first and foremost challenge for palladium catalysts in industry and academia. Progress means such catalysts made by methods simple and efficient enough for industrial scale [22]. The green oil law of Equation (3.4) is a teaching closure for that chemistry.

On the control side the override reads the bed outlets and acts after the selectivity is gone, the trip is reachable in Auto after a feed cut at a fat ratio, and the MPC plans on a fresh-catalyst copy with a linear loss cost and spends more than the suggested loops. No case models a hydrogen supply upset, a regeneration or a front-end converter.

OPEN CHALLENGE Models That Age With the Catalyst

A 2003 survey of industrial predictive control found no method to verify whether a model had deteriorated after the controller was commissioned, and only two adaptive products on the market, which the authors read as the difficulty of adaptive control in the real world [16]. A 2013 review of soft sensors in the process industries names model maintenance as the most critical issue in their industrial use [21]. Progress means models that follow the plant as it changes and need no maintenance. The predictive controller of this chapter plans on a fresh-catalyst copy and meets the aging catalyst only as mismatch.

3.8 Conclusion

An acetylene converter dresses a parts-per-million specification in a selectivity problem, and this chapter's simulator measures both on two beds of 24 cells whose carbon and hydrogen close to round-off. The window is one-sided. At the design inlets the edge sits at 1.43 mol/mol, breakthrough costs 2680 t/d of product, and each tenth of a percent of loss costs 2.68 t/d. The one-site rate law reproduces self-inhibition, 31 times at the inlet, and a light-off at depletion in which, once the hydrogen is raised by hand, the loss passes 1 % in 1.58 min, well before the bed's rise reaches the trip, while it cannot reproduce carbon monoxide's measured gain in selectivity. Feedforwards from an analyzer ahead of a 20 min drum, with a log-ppm trim, pass a graded day at 0.297 %, where a fixed ratio either breaks through or spends 0.566 % and the predictive controller spends 0.354 %. A balance inference carries the feedforward through an analyzer failure, and at day 480 no setting meets the limit. A fuller model would add a second site so that carbon monoxide raises the selectivity, a predictive override on the outlet chromatograph rather than the bed outlets, and a front-end variant with hydrogen in the feed. The simulator shows that on a one-sided window the cheapest safe hydrogen is the least that clears the edge.

KEY POINTS

- Too little hydrogen puts the whole product off grade, and too much burns ethylene and moves the beds toward a runaway, so the cheapest safe hydrogen is the least that clears the edge.
- At the design inlets the edge sits at 1.43 mol/mol, breakthrough costs 2680 t/d of product, and each tenth of a percent of loss costs 2.68 t/d.
- With the hydrogen raised by hand, the loss passes 1 % in 1.58 min, well before the bed's rise reaches the trip.
- Feedforwards from an analyzer ahead of a 20 min drum, with a log-ppm trim, pass a graded day at 0.297 %.

Where else this applies. A ratio held by an analyzer trim returns in Volume 3, Chapter 6 and Volume 3, Chapter 4, and a protective override in Volume 4, Chapter 2.

References

[1] A. N. R. Bos and K. R. Westerterp, "Mechanism and kinetics of the selective hydrogenation of ethyne and ethene," *Chemical Engineering and Processing: Process Intensification*, vol. 32, no. 1, pp. 1–7, 1993.

[2] A. Borodziński and G. C. Bond, "Selective hydrogenation of ethyne in ethene-rich streams on palladium catalysts. Part 1. Effect of changes to the catalyst during reaction," *Catalysis Reviews*, vol. 48, no. 2, pp. 91–144, 2006.

[3] ——, "Selective hydrogenation of ethyne in ethene-rich streams on palladium catalysts, Part 2: Steady-state kinetics and effects of palladium particle size, carbon monoxide, and promoters," *Catalysis Reviews*, vol. 50, no. 3, pp. 379–469, 2008.

[4] O. Bilous and N. R. Amundson, "Chemical reactor stability and sensitivity: II. Effect of parameters on sensitivity of empty tubular reactors," *AIChE Journal*, vol. 2, no. 1, pp. 117–126, 1956.

[5] M. Morbidelli and A. Varma, "A generalized criterion for parametric sensitivity: application to a pseudohomogeneous tubular reactor with consecutive or parallel reactions," *Chemical Engineering Science*, vol. 44, no. 8, pp. 1675–1696, 1989.

[6] 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.

[7] H. Falcke, S. Holbrook, I. Clenahan, A. López Carretero, T. Sanalan, T. Brinkmann, J. Roth, B. Zerger, S. Roudier, and L. Delgado Sancho, "Best available techniques (BAT) reference document for the production of large volume organic chemicals," Publications Office of the European Union, Luxembourg, EUR 28882 EN, 2017.

[8] A. N. R. Bos, E. S. Bootsma, F. Foeth, H. W. J. Sleyster, and K. R. Westerterp, "A kinetic study of the hydrogenation of ethyne and ethene on a commercial Pd/Al_2O_3 catalyst," *Chemical Engineering and Processing: Process Intensification*, vol. 32, no. 1, pp. 53–63, 1993.

[9] N. S. Schbib, M. A. García, C. E. Gígola, and A. F. Errazu, "Kinetics of front-end acetylene hydrogenation in ethylene production," *Industrial & Engineering Chemistry Research*, vol. 35, no. 5, pp. 1496–1505, 1996.

[10] N. Mostoufi, A. Ghoorchian, and R. Sotudeh-Gharebagh, "Hydrogenation of acetylene: kinetic studies and reactor modeling," *International Journal of Chemical Reactor Engineering*, vol. 3, no. 1, 2005.

[11] R. Gobbo, R. P. Soares, M. A. Lansarin, A. R. Secchi, and J. M. P. Ferreira, "Modeling, simulation, and optimization of a front-end system for acetylene hydrogenation reactors," *Brazilian Journal of Chemical Engineering*, vol. 21, no. 4, pp. 545–556, 2004.

[12] D. Teschner, J. Borsodi, A. Wootsch, Z. Révay, M. Hävecker, A. Knop-Gericke, S. D. Jackson, and R. Schlögl, "The roles of subsurface carbon and hydrogen in palladium-catalyzed alkyne hydrogenation," *Science*, vol. 320, no. 5872, pp. 86–89, 2008.

[13] F. Studt, F. Abild-Pedersen, T. Bligaard, R. Z. Sørensen, C. H. Christensen, and J. K. Nørskov, "Identification of non-precious metal alloy catalysts for selective hydrogenation of acetylene," *Science*, vol. 320, no. 5881, pp. 1320–1322, 2008.

[14] G. C. Battiston, L. Dalloro, and G. R. Tauszik, "Performance and aging of catalysts for the selective hydrogenation of acetylene: a micropilot-plant study," *Applied Catalysis*, vol. 2, no. 1–2, pp. 1–17, 1982.

[15] 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.

[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] P. J. Linstrom and W. G. Mallard, "The NIST Chemistry WebBook: A chemical data resource on the Internet," *Journal of Chemical & Engineering Data*, vol. 46, no. 5, pp. 1059–1063, 2001.

[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.

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

[20] 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] 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.

[22] X. Cao, B. W.-L. Jang, J. Hu, L. Wang, and S. Zhang, "Synthetic strategies of supported Pd-based bimetallic catalysts for selective semi-hydrogenation of acetylene: A review and perspectives," *Molecules*, vol. 28, no. 6, p. 2572, 2023.

3.9 Exercises

3.1 Concept check · 2 min

The converter's lined-out window runs from a breakthrough edge near 1.43 mol/mol to the override near 2.27 mol/mol. Why does the operating point belong just on the clean side of the edge rather than in the middle of the window?

(a) The middle of the window is where the override acts, so a ratio there would fight the override

(b) The outlet chromatograph is most accurate near 1 ppm, so the trim works best close to the edge

(c) Breakthrough puts all the product off grade, while each extra mole of hydrogen costs only a little ethylene

(d) The window is symmetric, so crossing either side costs the same and any point inside it will do

3.2 Concept check · 2 min

At the design inlets the carbon monoxide in the feed rises from 2 to 5 ppm. What does the one-site rate law of this model predict?

(a) The local selectivity rises, so the loss at the 1 ppm edge falls by about half

(b) The edge moves to a lower ratio, because carbon monoxide frees sites for the acetylene

(c) Only the ethylene rate is braked, so the loss falls at every ratio while the edge stays put

(d) Both rates fall by one factor, so the edge moves to a higher ratio and its loss barely changes

3.3 Concept check · 2 min

In case 6 the hydrogen is raised by hand from 75.6 to 130 kmol/h. In the next minute bed 1's temperature rise falls from about 30 to about 26 K while the ethylene loss climbs. Why does the rise fall?

(a) The reaction moves to the front cells, and their heat reaches the outlet only after a thermal wave crosses the bed

(b) The ethylene hydrogenation absorbs heat, so the bed cools while ethylene is turned into ethane

(c) The temperature override cuts the hydrogen as soon as it is raised, which cools the bed for a minute

(d) The bed thermocouples lag by about two minutes, so the reading shown is from before the change

3.4 Theory and calculation · 8 min

Use the standard heats of formation (226.73, 52.47 and −84.0 kJ/mol for acetylene, ethylene and ethane) and the heat capacities of Section 3.3 (46.66, 29.02, 46.40 and 57.40 J/(mol K) for acetylene, hydrogen, ethylene and ethane). The design feed is 1400 mol/s of C_2 gas with 1.0 % acetylene, 79 % ethylene and 20 % ethane, plus 75.6 kmol/h of hydrogen. Neglect the carbon monoxide.

(a) Compute the heats of reaction ΔH_1 and ΔH_2. *Checkpoint:* $\Delta H_1 = -174.26\,\mathrm{kJ/mol}$, $\Delta H_2 = -136.47\,\mathrm{kJ/mol}$

(b) Compute the heat-capacity flow $\dot{F}C_p$ of the feed with its hydrogen. *Checkpoint:* 68.65 kW/K

(c) Compute the adiabatic rise from hydrogenating all the acetylene, and the extra rise from each mol/s of hydrogen that finds ethylene instead. *Checkpoint:* 35.54 K and 1.99 K

(d) At design the loss is 0.200 % of the ethylene fed and the outlet holds almost no acetylene. Estimate the total rise over both beds and compare it with the measured rises of 30.2 and 9.8 K. *Checkpoint:* 39.9 K against 40.0 K

3.5 Theory and calculation · 6 min

Case 5 raises the feed acetylene from 1.0 to 1.5 % at 10 min on a fixed ratio of 1.5 to the C_2 flow at the design 1 %. The change reaches the converter through the feed drum, a first-order lag of 20 min. At design the lined-out 1 ppm edge is at 1.433 mol/mol.

(a) What true ratio of hydrogen to acetylene does the converter see once the change has arrived in full? *Checkpoint:* 1.00

(b) Treating the beds as lined out at every instant, how long after the change does the outlet pass 1 ppm? *Checkpoint:* 1.96 min

(c) What hydrogen flow must the acetylene feedforward deliver at 1.5 % to hold the ratio at 1.5, and why does the loss rise even though the ratio is held? *Checkpoint:* 113.4 kmol/h

3.6 Theory and calculation · 7 min

Use Equations (3.1) and (3.2) with $E_A = 25$ and $E_E = 60\,\mathrm{kJ/mol}$, $Q_{CO} = 30\,\mathrm{kJ/mol}$, no temperature dependence of K_A or K_E, and the inlet values at 60 °C: $K_A p_A = 9.85$, $K_E p_E = 0.156$ and $K_{CO} p_{CO} = 1.000$. Hold the gas composition fixed and compare a catalyst at 40 °C with one at 80 °C.

(a) Compute $k_A(80\,^\circ\mathrm{C})/k_A(40\,^\circ\mathrm{C})$ and the same ratio for k_E. *Checkpoint:* 2.97 and 13.6

(b) Compute the denominator D at both temperatures and the factor $D^2(40\,^\circ\mathrm{C})/D^2(80\,^\circ\mathrm{C})$ that both rates share. *Checkpoint:* 13.00 and 11.55, factor 1.27

(c) Give the rise of each rate from 40 to 80 °C and the change in r_E/r_A. What does it say about a hotter bed? *Checkpoint:* 3.76 and 17.2, ratio up 4.58 times

3.7 App challenge · 8 min

Load case 1, The hydrogen ratio in steps, at the default seed. Before pressing Run, set the Hydrogen ratio setpoint on the FIC card to 1.30 (its shipped value), run to 10 min at speed 5 (5 min of plant per wall second) and export the CSV. Reload the case and repeat with the setpoint at 1.50, then at 1.90. Read `truth_ace_out_ppm` and `truth_loss_pct` in the 10.00 min row of each file.

(a) Give the outlet acetylene and the loss at each ratio. Use a tolerance of 0.5 ppm at 1.30 and 0.003 % on the losses. *Checkpoint:* 14.9, 0.30, 0.00016 ppm and 0.137, 0.200, 0.470 %

(b) By how many decades does the acetylene fall from 1.30 to 1.90, by what factor does the loss rise, and which step costs more ethylene per mol/mol of hydrogen: 1.30 to 1.50, or 1.50 to 1.90? *Checkpoint:* 4.97 decades, 3.43 times, the second step

(c) Price the step from 1.50 to 1.90 in tonnes of ethylene a day, with the 2680 t/d of ethylene the stream carries, and compare it with running at 1.30. *Checkpoint:* 7.2 t/d

3.8 App challenge · 8 min

Load case 6, The runaway: acetylene runs out in bed 1, at the default seed. The Hydrogen loop is on Man. Press Run at speed 1 (1 min of plant per wall second), pause near $t = 2$ min and note the clock, raise the Hydrogen flow on the FIC card from 75.6 to 130.0 kmol/h, and Run until RUNAWAY appears and the run stops. Export the CSV. To repeat the run, reload the case: Reset keeps the 130 kmol/h, and a plant started there trips in its first second.

(a) How long after your change does the trip come? Give it within 0.05 min. *Checkpoint:* 3.95 min

(b) In the CSV, what is `bed1_dT_K` in the last row before the change, and what is its lowest value after it? *Checkpoint:* 30.1 to 30.4 K, then 25.7 to 25.8 K

(c) In the first row after the change where `truth_loss_pct` exceeds 1, what does `bed1_dT_K` read, and what does that say about protecting the plant with the bed rise? *Checkpoint:* 29 to 32 K

3.9 App challenge · 9 min

Load case 9, the challenge (A day on the furnace cycle), at the default seed, and run each configuration below to 24 h at speed 60 (60 min of plant per wall second). The run stops by itself at 24 h. Read the Worst outlet sample and Average ethylene loss tiles, and hover over the So far tile for the number of samples over 1 ppm. Reload the case before each configuration and make every setting before Run.

(a) As shipped: fixed ratio 1.50, Outlet trim off. *Checkpoint:* 1079 ppm, 185 of 288, 0.192 %

(b) Hydrogen from set to Feedforwards and Outlet trim to On, with the Outlet acetylene setpoint left at 0.30 ppm. *Checkpoint:* 1.36 ppm, 12 of 288, 0.242 %

(c) As in (b), with the Outlet acetylene setpoint at 0.10 ppm. Which of the three runs pass, and why does (c) succeed where (b) fails? *Checkpoint:* 0.53 ppm, none over, 0.297 %

3.10 Application area · 4 min

Each 0.1 % of loss is 2.68 t/d of ethylene at the design feed. Over the graded day the suggested loops (feedforwards, trim at 0.1 ppm) average 0.297 % and the predictive controller 0.354 %, and both hold every sample.

(a) Convert both losses to tonnes of ethylene a day and the difference to tonnes a year. *Checkpoint:* 7.96 and 9.49 t/d, about 560 t a year

(b) Give the two reasons the chapter gives for the predictive controller spending more, and one change you would test before trusting it on a plant. *Checkpoint:* model and cost

3.11 Application area · 3 min

Case 7 compresses a catalyst run so that a minute is a day. Without a schedule the outlet first passes 1 ppm on day 137.5. With 5 K per 100 days on both inlets it holds to day 177.8, at a loss of 0.446 % on day 150 against 0.334 without it.

(a) What does the schedule buy, and what does it cost in tonnes of ethylene a day on day 150? *Checkpoint:* 40.3 days for 3.0 t/d

(b) What ends the run even with the schedule, and what does a plant do then? *Checkpoint:* regeneration

3.12 Review the agent · 5 min

An AI assistant asked to make the converter robust against breakthrough writes: “Set the hydrogen ratio setpoint to 2.2 on a fixed ratio. The window puts the override at 2.27 and the trip at 2.53, so 2.2 leaves a wide margin on the breakthrough side, and the temperature override at 55 K guarantees that no Auto setting can reach the 70 K trip.” Find the flaw in the claim and say how you would check it in the app.

CHAPTER 4

A Propylene Splitter Simulator for Inferential Control and Quality Cascade Audit

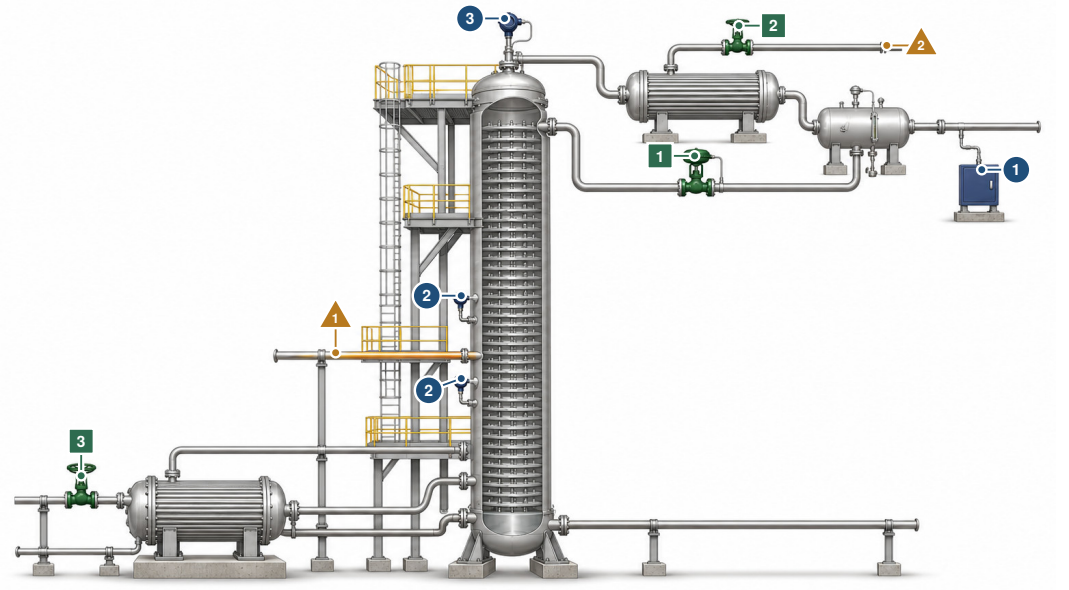

The controller moves 1 reflux valve, set as a ratio to the distillate flow
2 cooling water valve, which holds the overhead pressure
3 steam valve to the reboiler, set by the operator rather than a loop

The controller reads 1 product chromatograph, one result every twenty minutes
2 control and reference tray thermometers 3 overhead pressure transmitter

What disturbs it 1 changes in the feed's propylene content and rate
2 cooling water, cold at night and warm in the afternoon ▲ a chromatograph that freezes on its last result

Notice that the thermometers sit partway down the tower, close to the feed, while the only reading of the product's purity comes from the chromatograph at the far end of its line, every twenty minutes. Illustration, not to scale.

IN PLAIN TERMS A Thermometer Standing In for a Test

The main idea. A tall tower of trays splits propylene, which becomes plastic, from propane, which boils at almost the same point. The tower is so slow that a thermometer part way up stands in for the test of the product, which comes every twenty minutes.

The challenge. The thermometer also feels the pressure, and a tray far from the top sees changes the product never feels. If the test freezes on its last answer, a controller that trusts it does more harm than good.

Why it matters. This tower is often the biggest user of heat on its site, and plastic makers take only very pure propylene.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? box in Section 4.1, the History boxes in Section 4.2, the Try it boxes in Sections 4.4 and 4.6, the second Did you know? box in Section 4.6.7, and Key points in Section 4.10. No equations.
Study	50 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 4.6 with Figure 4.2, Section 4.7, Section 4.8 with its Open challenge box, and the second Open challenge box in Section 4.9.

The control question A propylene splitter is about half settled four hours after a reflux step, so its step test does not finish within a shift and a tray temperature must stand in for a product analyzed every twenty minutes. Which tray, which pressure correction and which layer above it hold polymer grade, and what is the analyzer above the tray loop actually for?

Learning objectives

- Explain why an open-loop step test cannot identify this column within a shift, and read the settling time from the impurity drift.
- Compute the pressure term of a tray bubble point, and what compensation removes and leaves behind.
- Choose a control tray for fidelity to the product rather than for sensitivity, and justify the choice with the measured feed step.
- Compare the inferential, cascade and analyzer-only structures with a healthy and a frozen analyzer, and state what the validity hold buys.
- Evaluate a predictive controller that removes giveaway against the minutes it spends below the specification, and set a back-off.

Prerequisites Sections 1.4.6 (inferred quality), 1.4.8 (back-off from a limit), 1.6.2 (velocity-form PID), 1.6.3 (predictive control option) and 1.6.4 (scenarios and Reset).

The app https://apmonitor.com/apps/c3-splitter/

See also Volume 5, Chapter 3 for a small column whose composition loop does feel a slow analyzer, and Volume 5, Chapter 10 for an inferred quality corrected by a delayed laboratory. Volume 3, Chapter 2 for the structure of a cascade before its tuning. Volume 9, Chapter 4 for offset-free predictive control with a bias estimator, and Volume 9, Chapter 9 for a controller that spends a band.

4.1 Introduction

The first experiment in a process control course is an open-loop step test. A manipulated variable is moved, the controlled variable settles, and a gain, a time constant and a dead time are read off the record. On a propylene splitter that experiment cannot be performed. The column modeled here is 53 percent complete four hours after a three percent step in reflux ratio and 98 percent complete after twenty-four, so a student who watches for a shift sees a ramp and fits an integrator to a plant that is not one. The rate of change of distillate impurity falls below 20 ppm/h only after 715 min. Every practical consequence follows from that single fact. The product analyzer is not the bottleneck it is usually assumed to be, the tray temperature that replaces it measures something other than what the specification is written on, and the architecture that a textbook would recommend is not the one that performs best.

The shape is not unique to propylene. Any separation carried out close to both a minimum stage count and a minimum reflux ratio has a small stage-to-stage driving force, a large liquid inventory relative to the product rate, and a composition response measured in hours rather than minutes [1]. Ethylene splitters, xylene columns and deisobutanizers share it, and so does any unit whose quality measurement is intermittent while its inferential is continuous and biased. What generalizes is not the distillation but the structure: a slow plant, a fast measurement that is a proxy, a slow measurement that is the truth, and a supervisory loop that must decide which to believe. Section 1.4.6 sets out the practices that recur across this volume, and Volume 5, Chapter 10 meets the same structure in a dryer whose moisture a delayed laboratory corrects. Inferential estimation from secondary measurements

has been studied for distillation since the 1980s [2] and remains active as soft sensing and disturbance rejection are combined [3]. The control literature on this particular column is narrower [4, 5] and concentrates on what advanced control adds rather than on what the underlying measurement means.

> ✦ DID YOU KNOW? **Two Separations, a Share of the World's Energy**
>
> Purifying propylene and ethylene accounts for about 0.3 % of the world's energy use, a 2020 review reports, and it is counted among seven chemical separations that could change the world [6]. The reboiler of this chapter's column takes 25.4 MW continuously.

A student handed a first-order transfer function and a set of tuning rules never encounters the difficulty that dominates this service, which is that neither available measurement is both fast and faithful. The misconception the simulator is built to correct is that an analyzer with a twenty-minute cycle is slow. Against a composition response that needs about twelve hours to settle it is fast, and measurement on this column is limited by fidelity and validity rather than by speed. This chapter works with an interactive propylene splitter in which the stage count, the holdup, the volatility and the bubble-point pressure coefficient are hidden, so that the inferential must be built, the control tray chosen, and the analyzer layer audited rather than assumed. Its scenarios reach conclusions on tray selection and cascade architecture that contradict the usual rules of thumb, and each conclusion is measured rather than asserted.

The chapter is organized as follows. Section 4.2 recalls the Fenske and Underwood limits that make the column tall and the history of its control. Section 4.3 states the 159-state column with its instrument chain, in the detail needed to rebuild it, and Section 4.4 shows what a step test can and cannot identify on it. Section 4.5 describes the five quality control settings, the pressure-compensation options, the analyzer validity hold and a model predictive controller. Section 4.6 reports the measured results, from the tray loop and pressure compensation through tray selection, the architecture audit, floating pressure and the graded challenge to the comparison with predictive control. Section 4.7 names the hardware and the practice behind the model. Sections 4.8, 4.9 and 4.10 state the limitations, the three lessons that carry beyond distillation and a summary, and the exercises follow in Section 4.11.

4.2 Background

Propylene and propane differ by one double bond. At 18 bar their bubble points differ by 6.2 °C and the relative volatility is 1.12. Polymer-grade propylene at 99.5 mol% or better is a polypropylene feedstock and refinery-grade propylene at 70 percent is not, so the separation is worth doing and expensive to do.

> ⧖ HISTORY **From Short-Cut Methods to Column Dynamics, 1920s to 1965**
>
> A survey of distillation dynamics and control, first given in 1992, looks back over some seventy years of study. From the 1920s into the 1950s the work mainly treated the steady state, with graphical and short-cut models. When the digital computer arrived in the 1950s, research turned to rigorous models that simulate the dynamic as well as the steady behavior, and to control. Up to about 1965 the leading contributions came from industrial researchers in the United Kingdom and the Netherlands, who did pioneering work on dynamic models and simulation, on controlling both product compositions at once, on choosing the control configuration and on the interactions between loops. The same survey noted that most columns respond sluggishly, almost like a pure integrator [7].

Two results from the 1930s and 1940s fix how expensive. Fenske's total-reflux analysis, developed while fractionating Pennsylvania gasoline [8], gives the minimum theoretical stage count, which for this duty is 72. No reflux ratio is large enough to shorten the column below it. Underwood's treatment of minimum reflux at infinite stages [9] gives 11.4 for the same duty, so no stage count is large enough to make the column cheap. Real splitters therefore carry 150 to 250 trays, frequently

in two shells because a single column that tall will not stand, and run reflux ratios of 14 to 20. The reboiler on the column modeled here is 25.4 MW continuously, which on many olefin plants is the largest single heat consumer on the site.

Control of the service developed in three stages. Before continuous analysis was available the column was run on a tray temperature and a reflux flow, with laboratory samples once a shift to confirm that the temperature still corresponded to the product. Tray temperature is a bubble point and therefore a function of composition and pressure together, so the arrangement was tolerable only while pressure was held constant and the feed did not change. Process gas chromatographs made a direct composition measurement available from the 1960s and made the modern cascade possible, with the analyzer adjusting the setpoint of the temperature loop rather than the valve. The estimation literature formalized what a secondary measurement can and cannot recover [2], and the configuration question of which flows to pair with which compositions was settled in the same period by studies of dual composition control and its interaction [1, 10]. Model predictive control of the splitter followed once the economics of an energy consumer this size justified it [4, 5].

> ⧖ HISTORY **A Temperature Loop Under an Analyzer, 1973 to 1991**
>
> Product composition measurements were often expensive, unreliable and late, so the most popular way to control a column's product became a tray temperature. To be sensitive, that tray sits some distance from the column's ends, where it may not follow the product well, and works of 1973 and 1990 answer with a fast temperature loop cascaded with an analyzer. A 1978 estimator that inferred composition from temperatures and flows proved very sensitive to model errors on high-purity columns. In 1991 a regression on about five temperatures spread along the column worked in simulation and in implementation, could also correct for pressure, and almost removed the measurement delay. By 1989 an industrial predictive controller on benzene and toluene towers was adjusting a tray temperature setpoint [7].

Energy became the second driver. Shinskey's treatment of distillation in an energy conservation context introduced floating pressure, the practice of letting the column run at the lowest pressure the condenser can hold, because relative volatility rises as pressure falls and the separation becomes cheaper [11]. The practice is standard now and the savings are still being quantified for this specific column [12, 13]. Whether a saving created in reflux ratio is collected in megawatts is a separate question, treated in Section 4.6. The alternatives to distillation are under active development, with membranes and metal-organic frameworks offering polymer-grade propylene at a fraction of the energy [14], but the installed base of superfractionators is large and will be operated for decades.

4.3 Process Model

Column C-401 is a 200-tray superfractionator with a total condenser, a reflux drum, a partial reboiler and a sump, operating at 18 bar on a feed of 700 kmol/h containing 72 mol% propylene. Reflux returns at 7000 kmol/h for a distillate of 494 kmol/h, a reflux ratio of 14.17. Figure 4.1 shows the arrangement, the instrument tags and every symbol used in the equations that follow.

Table 4.1 classifies what the loops interact with. The plant is expressed in the units of olefin operation, flows in kmol/h, holdups in kmol, pressure in bar, temperature in degrees Celsius and time in minutes. Every material balance therefore carries a factor of $1/60$. One hundred and fifty-nine states are integrated: 152 compositions covering the total condenser, 150 theoretical stages and the partial reboiler, the two holdups, the overhead pressure, and four actuator and hydraulic lags.

4.3.1 Stage Material Balances

Each theoretical stage carries a balance on the light key given by Equation (4.1), with vapor-liquid equilibrium Equation (4.2).

$$M\frac{\mathrm{d}x_n}{\mathrm{d}t} = L_{n-1}x_{n-1} + Vy_{n+1} - L_n x_n - Vy_n \tag{4.1}$$

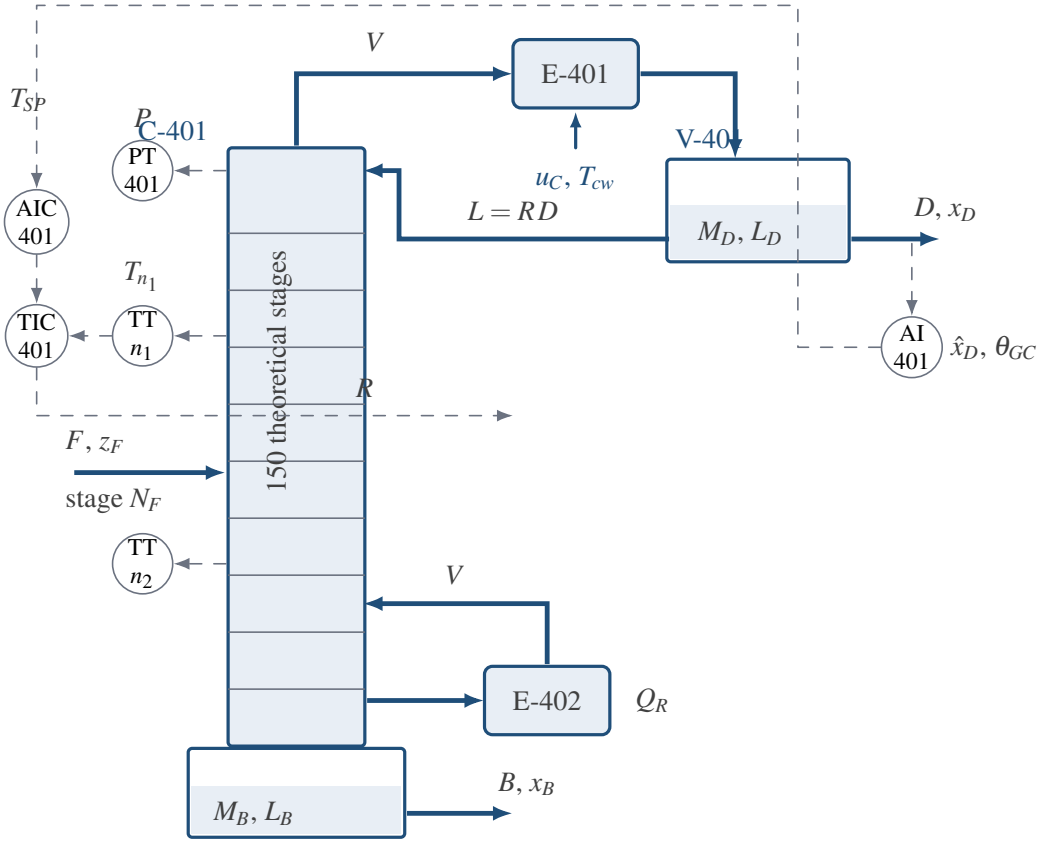

Figure 4.1 Superfractionator C-401 with its condenser E-401, reflux drum V-401, reboiler E-402 and sump. Feed F of composition z_F enters at stage N_F. Boilup V leaves the top, is condensed against cooling water at T_{cw} through valve u_C, and collects in V-401 at holdup M_D and level L_D. Reflux L is set as the ratio R times the distillate D, which leaves at propylene mole fraction x_D. The sump holds M_B at level L_B and the bottoms B leaves at x_B. TT-n_1 and TT-n_2 are the control and reference tray thermometers, PT-401 measures the overhead pressure P, and AI-401 is the product chromatograph of cycle time θ_{GC}. TIC-401 manipulates the reflux ratio and AIC-401 sets its setpoint T_{SP}.

$$y_n = \frac{\alpha x_n}{1 + (\alpha - 1) x_n} \tag{4.2}$$

where M is the liquid holdup per stage in kmol, x_n and y_n are propylene mole fractions in the liquid and vapor leaving stage n, L_n and V are the internal liquid and vapor flows in kmol/h, and α is the relative volatility. Constant molar overflow is assumed, so V is the same on every stage and L_n changes only at the feed.

The holdup is 14.6 kmol per stage, which is about 2200 kmol of liquid against a reflux of 7000 kmol/h. Liquid turns over in eighteen minutes and composition responds in hours, because near a pinch the stage-to-stage driving force is small and re-inventorying the profile is far slower than moving liquid through it.

The stage count is deliberately not reduced. Collapsing 150 stages into 40 with an effective relative volatility of $\alpha^{3.75} = 1.53$ preserves the total separation power and looks like a defensible reduction. It is not. The reduced column makes 99.99 mol% propylene at a reflux ratio of 10, below the true Underwood minimum of 11.4, where the 150-stage model reaches only 95.98 mol%, because the coarse stepping walks over the tangency between the equilibrium curve and the operating line. The

Table 4.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which the Scenarios Start Unless They Set Otherwise.

Symbol	Role	Nominal	Range
x_D (AI-401)	CV, intermittent	99.53 mol%	99.0–99.9 SP
T_{n_1} (TT)	CV, continuous	46.71 °C	44.0–50.5 °C SP
P (PT-401)	CV, continuous	18.0 bar	15.0–21.0 bar
R	MV	14.17	8–24
u_C	MV	65 %	0–100 %
V	MV, operator set	7494 kmol/h	4000–11 000 kmol/h
F	DV, measured	700 kmol/h	450–900 kmol/h
z_F	DV, unmeasured	0.72	0.60–0.88
T_{cw}	DV, measured	28 °C	12–36 °C
analyzer freeze	DV, unmeasured	off	on/off
$x_0 \ldots x_{151}$	states	profile	0–1
M_D, M_B	states	700, 450 kmol	0–1400, 0–900
P, lags	states	18 bar	—

pinch is the reason superfractionators exist, and the application carries an automated test that fails if it is ever lost.

4.3.2 Relative Volatility and the Tray Thermometer

Relative volatility falls with pressure, Equation (4.3), which is the entire economic content of floating pressure.

$$\alpha(P) = \alpha_0 - \frac{\mathrm{d}\alpha}{\mathrm{d}P}\,(P - P_0) \tag{4.3}$$

with $\alpha_0 = 1.12$ at $P_0 = 18\,\mathrm{bar}$ and $\mathrm{d}\alpha/\mathrm{d}P = 0.0075/\mathrm{bar}$.

A tray temperature is a bubble point, so it depends on composition and pressure together, Equation (4.4).

$$T_n = x_n\, T^{sat}_{C_3^=}(P) + (1 - x_n)\, T^{sat}_{C_3}(P) \tag{4.4}$$

$$\begin{aligned} T^{sat}_{C_3^=}(P) &= T^{sat}_{C_3^=,0} + \frac{\mathrm{d}T}{\mathrm{d}P}\,(P - P_0) - \beta\,(P - P_0)^2 \\ T^{sat}_{C_3}(P) &= T^{sat}_{C_3,0} + \left(\frac{\mathrm{d}T}{\mathrm{d}P} + \delta\right)(P - P_0) - \beta\,(P - P_0)^2 \end{aligned} \tag{4.5}$$

where the bubble points at P_0 are $T^{sat}_{C_3^=,0} = 45.6\,^\circ\mathrm{C}$ for propylene and $T^{sat}_{C_3,0} = 51.8\,^\circ\mathrm{C}$ for propane, the pressure slope is $\mathrm{d}T/\mathrm{d}P = 3.10\,^\circ\mathrm{C/bar}$, the propane curve is steeper by $\delta = 0.15\,^\circ\mathrm{C/bar}$, and the curvature is $\beta = 0.05\,^\circ\mathrm{C/bar^2}$. The composition term and the pressure term of Equations (4.4) and (4.5) are what the application is about. The composition term spans 5.80 °C from the top of the column to the bottom, about 0.038 °C per stage. The pressure term is about 3.1 °C per bar (3.12 °C/bar at stage 70, where the liquid is 82 percent propylene). A tenth of a bar therefore reads like almost seven percent of the reflux ratio, and a 0.4 bar excursion is more temperature than the entire composition profile between the control tray and the top of the column.

4.3.3 Why the Manipulated Variable Is a Ratio

With constant molar overflow and the drum level on the distillate valve, $D = V - L$. Reflux is 7000 kmol/h and distillate is 494 kmol/h, so one percent on the reflux flow is fourteen percent on the product rate, and the bottoms composition moves from 6 to 29 mol% propylene for it. Ratioing reflux to distillate instead gives Equation (4.6).

$$D = \frac{V}{1 + R}, \qquad L = R\,D \tag{4.6}$$

Table 4.2 Model Parameters. Adjustable Parameters Can Be Changed in the Model Dialog, and Hidden Parameters Are Fixed.

Symbol	Meaning	Value	Class
α_0	volatility at 18 bar	1.12	adjustable
$\mathrm{d}\alpha/\mathrm{d}P$	volatility fall with P	0.0075/bar	adjustable
$\mathrm{d}T/\mathrm{d}P$	bubble point vs. pressure	3.10 °C/bar	adjustable
M	holdup per stage	14.6 kmol	adjustable
N	theoretical stages	150	fixed
N_F	feed stage	84	hidden
τ_h	reflux to sump transit	8.0 min	adjustable
τ_V	reboiler thermal lag	2.5 min	adjustable
C_V	overhead capacitance	12.0 kmol/bar	adjustable
k_C	condenser duty coefficient	6.55	adjustable
θ_{GC}	analyzer cycle	20 min	adjustable
τ_T	thermowell lag	0.6 min	adjustable
x_D^{spec}	polymer grade	99.5 mol%	adjustable
x_B^{max}	propylene in bottoms	8.0 mol%	adjustable

One percent on R is then one percent on the product. The LV configuration is not wrong in general and the configuration question has a long literature [1, 10]. It is wrong at a reflux ratio of fourteen, where the manipulated flow is fourteen times the flow that closes the material balance.

4.3.4 Condenser, Pressure and Holdups

Overhead pressure is a small capacitance driven by the imbalance between boilup and condensation, Equation (4.7).

$$C_V \frac{\mathrm{d}P}{\mathrm{d}t} = V - k_C\, u_C\, f\left(T_{C_3^=}^{sat}(P) - T_{cw}\right) \tag{4.7}$$

where $C_V = 12.0\,\mathrm{kmol/bar}$, $k_C = 6.55\,\mathrm{kmol/h}$ per percent of valve per kelvin of approach, u_C is the cooling water valve in percent, and f is a condenser condition factor that is 1.0 when clean. At design the valve sits at 65 % with 17.6 K of approach. The resulting pressure time constant is about half a minute (0.55 min), which makes pressure the fast loop on a column whose composition is the slow one. The drum and sump integrate their own flow imbalances. A reflux change reaches every rectifying stage at once, because constant molar overflow carries no tray hydraulics, and a single eight-minute lag τ_h carries it to the stripping section and the sump.

4.3.5 Assumptions

The model is calibrated to produce the right sensitivities rather than to size equipment, and six assumptions bound what it can teach.

The separation is binary. A real feed carries propadiene, methylacetylene, ethane and C4 traces, and polymer-grade specifications include several of them, so the reader would not learn here that a product can fail on a contaminant this separation does not address.

Constant molar overflow is assumed with no energy balance, so subcooled reflux and changes in feed enthalpy do not alter the internal traffic the way they do on a real column.

Relative volatility is a single number linear in pressure. Real propylene/propane volatility varies along the column with temperature and composition, and the variation matters most where the driving force is smallest.

Tray efficiency is fixed and folded into the stage count. Efficiency on a real splitter falls at turndown and near flooding, which is a common way for a column to lose separation while every instrument reads normally.

There is no hydraulic capacity limit. Flooding, weeping, entrainment and tray pressure drop are absent, so the application cannot show the constraint that usually stops a splitter being pushed, which is that the column floods before the reboiler runs out.

Liquid transport is one eight-minute lag rather than 150 weir calculations, so a reflux change arrives at the sump smoothed rather than as a wave, while the rectifying section sees it without delay.

4.4 Dynamic Analysis

4.4.1 What Can Be Identified

The distillate composition response to a three percent step in reflux ratio (14.17 to 14.60, which moves the distillate from 99.528 to 99.638 mol%) is 29 percent complete at two hours, 53 percent at four and 98 percent at twenty-four, and it reaches 63 percent only at 312 min. In any window a student can watch, the record looks like a ramp. Fitting first order plus dead time to the first two hours returns a time constant of 3.1 h against the 5.2 h at which the step is 63 percent complete, and a gain 43 percent low, with a tuning to match. The application therefore reports the rate of change of distillate impurity in parts per million per hour, filtered over about five minutes, as an explicit indicator of whether an experiment has finished. Below about 20 ppm/h the column is as settled as it becomes inside a shift, and the design condition reaches that threshold 715 min after a three percent ratio step.

> **TRY IT IN THE APP 1 · The step test that does not finish (MANUAL)**
>
> Load the scenario, stay in Manual and raise the Reflux ratio from 14.17 to 14.60. Run at 200 min/s and watch the Drift badge rather than the purity trend. The badge climbs as the column starts to move and first reads under 20 ppm/h only at about 715 min, although one pass of liquid through the column takes eighteen minutes.

The tray temperature is not faster in open loop. Stage 70 moves half of its eventual change within ten minutes of the step, because the extra reflux reaches every rectifying stage at once, gives almost all of it back by twenty minutes as the liquid lag delivers the same reflux to the stripping section, and then follows the same composition front as the product. It is 36 percent complete at two hours and 63 percent complete at 312 min, exactly as the distillate is, and a first-order fit to its first two hours captures only the initial transient. What a shift-long window identifies on either signal is a slope. The inner loop is fast because feedback overdrives the ratio, as Section 4.6 shows, rather than because its measurement settles sooner. The outer analyzer loop is tuned from the requirement that it be slower than the inner loop by a wide margin, which is the standard cascade argument [15]. Volume 3, Chapter 2 shows the other half of that argument on a level loop: an inner loop rejects the disturbances that enter inside its own measurement path and few others.

4.4.2 Sensitivity Against Fidelity

The steady-state gain from reflux ratio to control tray temperature, for a one percent increase at constant boilup, is −0.047 °C at stage 70, which is as sensitive as any stage above the feed, and −0.016 °C at stage 25, roughly a third of it. Sensitivity is maximized in the stripping section, where the most sensitive stage is 127 at −0.39 °C, about eight times the sensitivity of stage 70. Sensitivity is not the criterion, however, because a stage far from the product responds to disturbances the product does not see. Section 4.6 reports the measured comparison.

4.5 Controller Design

4.5.1 Velocity Form PID

Every loop in the application uses the shared velocity-form PID of Section 1.6.2, Equations (1.14) to (1.17), executed every $\Delta t_c = 0.1$ min, with the derivative taken on the measurement through a first-order filter of time constant τ_f. The derivative path is configured ($\tau_f = 3.0$ min on TIC-401) but

$\tau_D = 0$ in every scenario, so every loop runs as PI. The bias is zero and the integral is seeded with the starting output at reset.

Anti-windup acts on the integral term alone, with the tracking time $T_t = \tau_I$, and the running sum of proportional and derivative increments is never truncated. The tracking time matters little here. In the uncompensated pressure case of Section 4.6 the demand reaches 27.3 while the ratio sits on its 24.0 stop, and a counterfactual run that absorbs the whole excess in one scan ($T_t = \Delta t_c$) leaves the ratio's later undershoot within 0.02 of the shipped value, although it arrives about 70 min sooner. No loop is slew limited. The reflux flow, R times the distillate flow filtered over 3 min, is clamped again to 3000–12 000 kmol/h after the PID, a limit TIC-401 does not know about, but it cannot bind within the ratio and boilup ranges. Transfer between Manual and automatic is bumpless in both directions, as Section 1.6.2 describes. The integration step is 0.02 min of fixed-step fourth-order Runge-Kutta, and every run, including its seeded noise, reproduces exactly (Section 1.6.1). The fastest run speed is 200 minutes of plant time per wall-clock second, so a twelve-hour settling time passes in a few seconds.

4.5.2 The Quality Control Settings

The quality control selector offers five configurations. The first four are described here and audited in Section 4.6, and the fifth, a model predictive controller, is the subject of Section 4.5.5.

The action a below is that of Equation (1.14). *Off* leaves the reflux ratio in manual. *Inferential* closes TIC-401 from the control tray temperature to the reflux ratio (limits 8 to 24), with $a = -1$ because more reflux ratio is a colder tray. *Cascade* adds AIC-401 from the product analyzer to the tray temperature setpoint (limits 44.0 to 50.5 °C), with $a = -1$ because a colder tray setpoint is more purity. *Analyzer only* removes the temperature loop and puts the analyzer straight on the ratio station, with $a = +1$ because more reflux ratio is more purity. Each scenario loads one of two tunings. The cautious set, used where the reflux is in manual and as the challenge's starting point, is $K_c = 0.6/°\mathrm{C}$ with $\tau_I = 400$ min on TIC-401 and $K_c = 0.15\,°\mathrm{C}$ per mol% with $\tau_I = 1200$ min on AIC-401. The tuned set, used in every closed-loop scenario, is $K_c = 5.0/°\mathrm{C}$ with $\tau_I = 60$ min (70 min on the differential) and $K_c = 0.9\,°\mathrm{C}$ per mol% with $\tau_I = 240$ min. The analyzer-direct case runs $K_c = 8$ per mol% of purity with $\tau_I = 240$ min.

Two further options change what the inferential measures. Pressure compensation subtracts the bubble-point pressure term arithmetically using a reference pressure, Equation (4.8), which is Equation (1.10) of Chapter 1.

$$T_n^{comp} = T_n - \frac{\mathrm{d}T}{\mathrm{d}P}\left(P - P_{ref}\right) \tag{4.8}$$

The differential option replaces the single tray with the difference between two trays, which removes the pressure term structurally because both trays move together. Section 4.6 reports where each option holds and where it fails.

4.5.3 Analyzer Validity

A frozen chromatograph republishes its last good result rather than failing visibly, so trends look calm while the outer loop integrates against an error nobody is measuring. The validity option inhibits the supervisory loop while the analyzer result is flagged bad, holding the tray temperature setpoint at its last value and tracking it so that the loop resumes without a bump. The measured consequence of omitting it is in Table 4.4.

4.5.4 Pressure, Levels and Floating Pressure

PIC-401 manipulates the cooling water valve, with $a = -1$ because more valve is a lower pressure, and $K_c = 12\,\%/\mathrm{bar}$ and $\tau_I = 1.2$ min in the tuned set (3 %/bar and 8 min in the cautious one). The two level loops are deliberately detuned so that the drum and sump absorb upsets rather than pass them downstream, with gains of 40 and 25 kmol/h per percent and integral times of 20 min in the tuned set (22 and 14 kmol/h per percent, 30 min, in the cautious one). Pressure minimization ramps

the pressure setpoint down at 0.006 bar/min while the valve command is below 92 %, toward a floor of 15.5 bar, and back up at three times that rate when the command exceeds 97.5 %, never above the operator's setpoint.

4.5.5 Model Predictive Control

The fifth quality setting replaces the single-number loops by a plan, built on the predictive control option of Section 1.6.3. Every $T_s = 5$ min, fifty controller scans, a quadratic program chooses the reflux ratio R and, under floating pressure, the setpoint P^{SP} of PIC-401 over $N = 48$ decisions, four hours, and the first move is held until the next decision [16, 17]. Three choices depart from the design first proposed, which named the cooling water valve as the second handle and reflux as the energy cost. The plan writes the pressure setpoint rather than the cooling water valve, because with the valve held at 65 % after the cooling water steps from 28 to 34 °C the pressure passes the 19.6 bar relief margin within a minute (19.65 bar). The pressure is a decision only under floating pressure, and otherwise stays at the operator's setpoint. Reflux is priced but is not an energy cost, because the boilup is an operator setpoint and unspent reflux leaves as distillate.

The controller's model is identified at the design point by steps of ± 0.15 in R, ± 0.25 bar in P^{SP} through PIC-401 and ± 0.02 in z_F, whose central differences over 48 h were fitted by Equation (4.9).

$$\begin{aligned} \Delta y(t) &= \Delta u \sum_{i=1}^{7} \left[a_i \, \phi_i(t) + b_i \, \psi_i(t) \right] \\ \phi_i &= 1 - e^{-t/\tau_i}, \quad \psi_i = 1 - \left(1 + t/\tau_i\right) e^{-t/\tau_i} \end{aligned} \tag{4.9}$$

Here Δu is the step, Δy the change in the compensated stage-70 reading of (4.8) or in $\ell = \ln(100 - x_D)$ with x_D in mol%, τ_i runs over 0.5, 2, 6, 20, 60, 200 and 700 min, and a_i and b_i are least-squares coefficients. The impurity is modeled in logarithms because its response is one-sided. With two states per lag and a random-walk disturbance d on the feed path the model has 43 states. Its gains on ℓ are -0.878 per unit ratio, $+0.646$ per bar and -26.6 per unit z_F, its residuals are 0.02 to 0.44 percent of each response, and an independent test with smaller steps agrees on the ratio response of ℓ to 0.6 percent at four hours. A steady-state Kalman filter reads the compensated tray once a minute, with a measurement variance of $(0.02\,^{\circ}\mathrm{C})^2$, and corrects the feed path and d, and at each new analyzer result a bias on ℓ moves as $b \leftarrow b + 0.3\,(\ell_{GC} - \hat{\ell} - b)$, which has the form of Equation (1.11) with $\lambda = 0.3$. Together they make the prediction offset free [18], as the fuel-bias estimator of Volume 9, Chapter 4 does for a drum boiler. The analyzer reaches the plan with its shipped cycle and sample age, noise, freeze and validity hold.

The plan minimizes Equation (4.10) over twelve move blocks whose lengths grow from one decision for the first four to fifteen, with a last block of three,

$$\begin{aligned} J = c_R \sum_{k=0}^{N-1} R_k + \rho_\sigma \sigma^2 + \mu_\sigma \sigma \\ + \sum_{j=1}^{12} \left[S_R \, \Delta R_j^2 + S_P \, (\Delta P_j^{SP})^2 \right] \end{aligned} \tag{4.10}$$

subject to Equation (4.11).

$$\begin{aligned} &\hat{\ell}_k \le \ln\left(100 - x_D^{SP}\right) + \sigma, \;\; k \ge 12, \;\; \sigma \ge 0 \\ &R_{lo} \le R_j \le R_{hi}, \quad |\Delta R_j| \le 0.5 \\ &P_{fl} \le P_j^{SP} \le P_{op}^{SP}, \quad |\Delta P_j^{SP}| \le r_f T_s \end{aligned} \tag{4.11}$$

Here $\hat{\ell}_k$ is the corrected prediction k decisions ahead, x_D^{SP} the purity target, Δ the move from the previous block or from the applied value, $c_R = 1$, $S_R = 300$, $S_P = 10$ per bar^2, and the one slack σ is

priced at $\rho_\sigma = 10^4$ and $\mu_\sigma = 10^3$, so it is zero whenever the band can be met. The purity row is soft and one-sided, so purity above the target is giveaway the reflux price removes. It starts one hour out, because a ratio move reaches only 9.9 percent of its effect on ℓ within 60 min. The ratio row is hard, $R_{lo} = \max(8, 3000/D_f)$ and $R_{hi} = \min(24, 12000/D_f)$ on the 3 min filtered distillate D_f, as the ratio station reads them. The pressure row is hard, with P_{op}^{SP} the operator's setpoint, $r_f = 0.006\,\mathrm{bar/min}$, and P_{fl}, never below the operator's floor, the pressure at which the valve would stand at $u_{lim} = 95\,\%$, from (4.7) at its declared values, Equation (4.12).

$$\begin{aligned} V_{SP} &= k_C\, u_{lim}\left(T_{C_3^{=}}^{sat}(P_{fl}) - \hat{T}_{cw}\right) \\ \hat{T}_{cw} &= T_{C_3^{=}}^{sat}(P_m) - V_{SP}/(k_C\, u_C) \end{aligned} \tag{4.12}$$

Here $\hat{T}_{cw}$ is the cooling water temperature the condenser behaves as if it had, from the boilup setpoint, the valve command and the measured pressure, filtered over 5 min. On scenario 11 this floor lies 0.001 bar above the true one on average. A dual active-set solver [19] handles the 24 block decisions, and S_R and the disturbance intensity were chosen once, on scenario 11, for the least quality IAE among settings that rejected no plan.

Under the plan the three PI loops track the applied ratio. When the solver reports anything but an optimum, or the first move is non-finite or outside the box, the plan is discarded and the inferential loop takes the reflux. Forced at four instants with the noise off, the hand-over steps the ratio by 7.1×10^{-4} to 2.2×10^{-3}, and a switch to manual leaves the ratio exactly where the plan put it. No plan was infeasible or rejected on any case of Section 4.6.7. The plant runs with the run's parameters, the instruments add lags, sampling and noise, and the model is the fit at the design point, so a plant edit is a mismatch, while an edit of $\mathrm{d}T/\mathrm{d}P$ or P_{ref} reaches the plan through the compensated reading. Solving one plan took 1.40 ms on average in Chrome on scenario 11, at a load average of 1.57.

4.6 Results

Twelve scenarios progress from an unfinishable step test to a graded challenge. Scenario 1 attempts the step test and finds that it does not finish, and scenario 2 finds that neither the tray temperature nor the purity settles inside a shift. Scenario 3 walks a thermometer along the temperature profile, scenario 4 closes the inferential and shows that a temperature is not a purity, and scenarios 5 and 6 break it with pressure. Scenario 7 chooses a tray and finds that the sensitivity rule is only half a rule, scenarios 8 and 9 audit the architecture, scenario 10 takes or fails to take the energy saving of floating pressure, and scenario 11 sets the predictive controller against the cascade. Scenario 12, the graded challenge, asks for the column to be restructured. Each case is credited separately when the learner performs the listed actions and the simulation reaches the listed state. Several ask for a rerun after Reset, which restarts the case from time zero and keeps the learner's own settings (Section 1.6.4), so the comparison a scenario teaches is also the condition for completing it. The challenge is judged on the case as given, and changing a model parameter or a disturbance setting voids the attempt. Credit accumulates toward certificates of progress and of completion for the module.

4.6.1 The Inferential Holds a Temperature

Stepping the tray temperature setpoint 0.09 °C down at 60 min moves the reflux ratio from 14.14 to 14.48, 2.4 percent, after overdriving it to 14.87, returns the temperature to within 0.005 °C of setpoint in 208 min, and raises the distillate from 99.527 to 99.614 mol% over the following hours, overshooting to 99.634 near eight hours. The loop has held a temperature, and the temperature corresponds to a purity only at the feed composition and pressure in force when the setpoint was written down.

4.6.2 Pressure Compensation and Its Limit

A 0.4 bar pressure setpoint move raises every bubble point in the column by about 1.25 °C, which is more temperature than the entire composition profile between the sensitive tray and the top. With an

Table 4.3 Feed Composition Change Held on Two Control Trays. Peak Departure From the Pre-Step Purity and Composition at 15 h.

Measure	Stage 25	Stage 70
Relative sensitivity	1	3
Peak departure (mol%)	+0.099	−0.373
Composition at 15 h (mol%)	99.489	99.298

uncompensated measurement the controller cannot win and tries anyway. The reflux ratio runs from 14.14 to its stop at 24.0 within 32 min of the step, and the distillate climbs to 99.89 mol% by 300 min against a specification of 99.5. When the compensation is switched on at 300 min with the pressure still high, the loop unwinds, undershoots to 13.56 near 750 min and settles at 14.37, 1.7 percent above where it started, with the distillate at 99.47 mol%. The residual is a process effect, because the higher pressure has lowered α from 1.1200 to 1.1170. Selected before the run instead, the compensated loop moves the ratio by at most 2.9 percent, after a dip to 12.48 in the first two minutes of the step, where the pressure transmitter (0.1 min filter) runs ahead of the thermowell (0.6 min lag) and the arithmetic briefly overcorrects.

The differential measurement removes the pressure term more completely at the measurement. A tenth of a bar moves a single tray by 0.312 °C and moves the difference between stages 90 and 70 by 0.0012 °C, which is 0.4 percent of it. Through a 0.6 bar pressure setpoint step held from 60 min, however, the compensated single tray lets the distillate fall by at most 0.092 mol% and settle at 99.44, while the differential, held near its 0.50 °C setpoint, lets it fall 0.63 mol% to 98.90 at 14.3 h, and it is still only 99.18 at 25 h. Pressure also changes relative volatility (1.1200 to 1.1155 here) and therefore changes the profile a given differential corresponds to. Compensating the measurement is not compensating the process.

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

4.6.3 Tray Selection

Stage 70 has three times the sensitivity of stage 25 and loses. Table 4.3 reports the 72 to 78 percent feed composition step at 60 min held on each, with the stage-70 case run as the scenario directs, by moving the tray, resetting and taking the setpoint (46.71 °C) from the reading.

Stage 25 is nearer the product and is therefore a more faithful proxy for it. Its purity is still moving at fifteen hours and settles at 99.495 by 50 h, while stage 70 settles at 99.28. The most sensitive stage in the column is 127 at about eight times stage 70, and it is in the stripping section, which makes it an excellent inferential for the bottoms and a poor one for the distillate.

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

4.6.4 Architecture Audit

Table 4.4 reports the architectures against the same feed composition step, with a healthy analyzer and then with the analyzer frozen, each on the scenario that runs it. In the frozen cases the chromatograph stops updating at 60 min with its last result 0.07 mol% high, and the feed steps a minute later. Figure 4.2 draws the five rows side by side.

Three results follow. With a healthy chromatograph the analyzer alone outperforms the cascade, which contradicts the usual recommendation. It departs 0.063 mol% against 0.143 and is back at 99.51 at fifteen hours while the cascade is at 99.47, although both reach 99.50 by thirty. The analyzer cycle time barely matters. Moving it from 5 to 60 min changes the peak departure only from 0.061 to 0.070 mol%. Against a response that is 63 percent complete at 5.2 h and settles in about twelve, a twenty-minute analyzer is a fast measurement. The far smaller column of Volume 5, Chapter 3 is the opposite case, where an analyzer lag of a few minutes already takes the loop out of its band.

Table 4.4 Architecture Comparison on the 72 to 78 Percent Feed Composition Step. Peak Departure From the Pre-Step Purity and Distillate Composition at 15 h and 30 h, All in mol%.

Architecture (scenario)	Peak	15 h	30 h
Inferential alone (7)	−0.373	99.298	99.283
Cascade (9)	−0.143	99.472	99.499
Analyzer only (8)	+0.063	99.512	99.499
Cascade, frozen, hold at 600 min (9)	−0.712	98.906	98.903
Cascade, frozen, hold from start (9)	−0.516	99.190	99.170

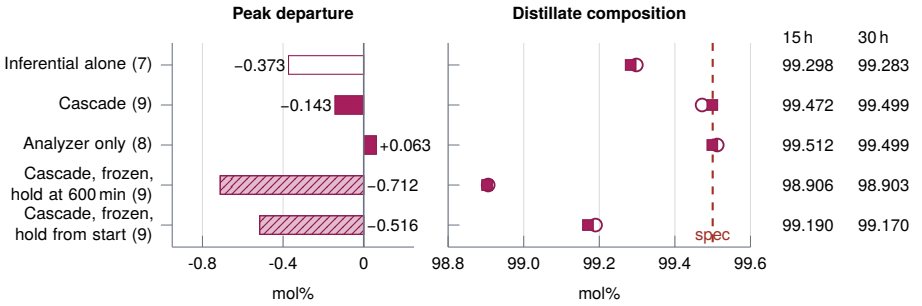

Figure 4.2 The architectures of Table 4.4 on the 72 to 78 percent feed composition step. (a) Peak departure from the pre-step purity: the open bar has no analyzer, solid bars a healthy one, hatched bars a frozen one. (b) Distillate composition at 15 h (open circles) and 30 h (filled squares) against the 99.5 mol% specification (dashed line). Notice that the analyzer alone departs least while the chromatograph is healthy, and that a frozen chromatograph left in the cascade until 600 min takes the column furthest below the specification.

Pause. Try Exercise 4.1 before reading on.

The inferential earns its place elsewhere. With the chromatograph frozen and the validity hold not armed until 600 min, the outer loop keeps integrating a stale error and walks the tray setpoint from 46.80 to 46.99 °C, the purity falls to 98.81 at 473 min and settles at 98.90, which is worse than the inferential alone. Arming the hold from the start freezes the setpoint at 46.80 °C, and the cascade then settles at 99.17, exactly where the inferential alone settles on the same case when its setpoint is 46.80 °C rather than the commissioned 46.71 °C. The remaining gap to the inferential row is the trim the analyzer applied before it failed. The case for the inner temperature loop on this column is therefore availability rather than speed.

TRY IT IN THE APP 9 · What the inferential is actually for

Run the case as loaded. The chromatograph freezes at 60 min and the feed goes richer a minute later. Watch the Tray temperature setpoint walk from 46.80 to 46.99 °C until the hold comes on at 600 min, while the purity falls to 98.81 and settles at 98.90 mol%. Press Reset, set Analyzer validity to Hold on bad quality before Run, and the cascade settles at 99.17.

4.6.5 Floating Pressure

Enabling pressure minimization on a cold night, with cooling water at 16 °C, walks the pressure setpoint to its 15.5 bar floor in 417 min and raises relative volatility from 1.1200 to 1.1388. At fixed boilup the improvement appears first as recovery. By 1000 min the cascade has lowered the reflux ratio to 13.86, against 14.11 on the same case at a fixed 18 bar, the distillate has risen from 494 to 504 kmol/h, and the bottoms carry 0.21 mol% propylene instead of 6.0. The reboiler duty does not



Figure 4.3 The propylene splitter application near the end of scenario 9, the analyzer validity scenario, at 1800 min with the quality cascade in automatic, cropped to the plant view. The chromatograph froze at 60 min reading 99.592 mol%, the feed went from 72 to 78 percent propylene at 61 min, and the validity hold was armed at 600 min. (1) The status chips read the true purity, the drift, the reflux ratio (12.58), the duty, the recovery and the bottoms. (2) The condenser, its cooling water valve and the reflux drum sit above the shell, which is drawn in elevation with its 150 theoretical stages colored by propylene mole fraction, and the reboiler and sump below it. (3) The control tray thermometer and (4) the feed sit on the side of the shell. (5) The inset plots the composition and temperature profiles against stage number with the control tray and the feed stage marked.

move. It stays at 25.40 MW, because the boilup is a number an operator typed and no loop is reducing it. Cutting the boilup to 7341 kmol/h at 1000 min takes the duty to 24.88 MW, half a megawatt, with the distillate still on grade. That cut collects a small part of what the lower pressure created. In the model the design split (494 kmol/h of distillate at 99.5 mol% with 6 mol% propylene in the bottoms) needs 6100 kmol/h of boilup at 15.5 bar against 7433 kmol/h at 18 bar, 20.67 MW against 25.19 MW, and the cascade settles there (ratio 11.35, bottoms 6.01 mol%) when the boilup is cut to 6100 kmol/h and the cooling water stays cold. The saving is rented from the weather. When the cooling water warms to 34 °C at 1200 min the valve saturates, the pressure touches 19.03 bar, and the minimization backs the setpoint up to 17.99 bar. A column cut to 6100 kmol/h settles at 17.4 bar, where that boilup is far too little, and its bottoms reach 27 mol% propylene by 2400 min. The distinction between creating a saving and collecting it is the point of the scenario, and it is why floating pressure is usually commissioned together with a duty minimization rather than alone [11, 12].

4.6.6 The Graded Challenge

The challenge runs 2400 min from the design point on a configuration that is wrong three ways, with an uncompensated stage-25 inferential, the cautious tuning and no supervisory loop. Four disturbances arrive: feed composition 72 to 77 percent at 120 min, cooling water 22 to 33 °C at 500 min, feed rate 700 to 790 kmol/h at 900 min, and feed composition 77 to 70 percent at 1200 min. It is met when the distillate is held within 0.12 mol% of the 99.5 mol% target from 1200 min to the end with propylene in the bottoms below 8 mol% at 2400 min. As loaded, the column never goes off specification. It gives purity away instead, up to 99.76 mol%, with a worst error after 1200 min of 0.259 mol%, while the bottoms reach 39.5 mol% propylene after the feed rate step and end at 23.8. The suggested configuration compensates the stage-70 measurement, closes the cascade with the

Table 4.5 Quality Structures Against Model Predictive Control on the Feed Composition Cases

Case (window)	Structure	IAE (mol% · min)	Giveaway (kmol)	Off-spec (min)	More than 0.01 below (min)	Lowest purity (mol%)	Propylene product (kmol)	Plans, fallbacks
F (900 min)	inferential	175.1	0.2	767.1	751.0	99.156	8193	—
	cascade	47.5	0.2	789.5	774.1	99.386	8037	—
	analyzer only	37.8	3.3	0.0	0.0	99.512	7904	—
	predictive	13.6	1.2	0.0	0.0	99.502	7954	180, 0
D, $\theta_{GC} = 60$ min (900 min)	inferential	175.1	0.2	767.1	751.0	99.156	8193	—
	cascade	48.3	0.2	792.0	777.9	99.375	8040	—
	analyzer only	37.5	3.3	0.0	0.0	99.510	7905	—
	predictive	13.3	1.2	0.0	0.0	99.504	7953	180, 0
Scenario 11 (1800 min)	inferential	220.4	13.6	452.1	407.1	99.269	15672	—
	cascade	123.3	8.1	355.8	327.8	99.364	15775	—
	analyzer only	144.0	6.9	817.0	778.6	99.323	15778	—
	predictive	23.2	1.0	942.3	324.2	99.459	15805	360, 0

IAE, giveaway and off-specification minutes (below 99.5 mol%) are the scorecard's, on the true distillate composition from $t = 0$, and the next column counts the minutes more than 0.01 mol% below over the same steps. The lowest purity follows the first event. The reboiler duty is 1371 GJ at 900 min and 2743 GJ at 1800 min in every row. Cases F and D are noise free, so their inferential and cascade rows differ in the third decimal from Table 4.4. Diagnosis runs on scenario 11, not part of the application: handed the true present purity, the predictive controller spends 689.3 min below specification, 192.2 of them more than 0.01 below, and with its band from the first step 797.0 min. With the target set above the specification it spends 336.1 min below with 1.78 kmol of giveaway at 99.51 mol%, 111.0 min with 4.26 kmol at 99.53, and none with 7.35 kmol at 99.55.

validity hold, applies the tuned set, floats the pressure to a 16.2 bar floor and raises the boilup to 8500 kmol/h. It passes with a worst error of 0.103 mol%, below target at 1208 min, inside the band with 0.017 to spare, and 5.72 mol% in the bottoms, although the condenser valve is saturated from the cooling water step onward. Every case starts its column at steady state on the case's own boilup and reflux ratio, so the higher boilup acts from the first scan through the reboiler lag and the configuration gives one run whether it is chosen before the run or carried through Reset. An automated check confirms that the suggested configuration meets the target, so the challenge cannot be set without a demonstrated solution.

4.6.7 Comparison With Model Predictive Control

Table 4.5 sets the closed-loop structures of Table 4.4 against the predictive controller. Case F is scenario 8's feed step, 72 to 78 percent at 60 min with a 20 min analyzer, and case D is the same with a 60 min cycle. Scenario 11 adds a step to 70 percent at 900 min, floats the pressure toward a 15.5 bar floor and carries thermowell noise of 0.01 °C and analyzer noise of 0.02 mol%. Every row has the same seed and instruments, the analyzer-fed structures work to the same 99.5 mol% target, the shipped structures keep their shipped tuning, the predictive controller keeps the settings of scenario 11 on every case, and none receives a preview.

The predictive controller has the least integrated error on every case, the least giveaway on scenario 11, and the least on cases F and D among the structures that hold the specification. The inferential loop and the cascade give away only 0.2 kmol on cases F and D because they spend more than 760 min below it. On case F the plan integrates 13.6 mol% · min against 37.8 for analyzer direct, because it reads the feed change in the tray. Its disturbance estimate passes 10^{-3} at 69.1 min, nine minutes after the step, and by 180 min it has cut the ratio from 14.02 to 13.29, while analyzer direct has taken 0.04 off its ratio. A 60 min analyzer cycle hardly matters to either (13.3 against 37.5), because the chromatograph samples rather than delays.

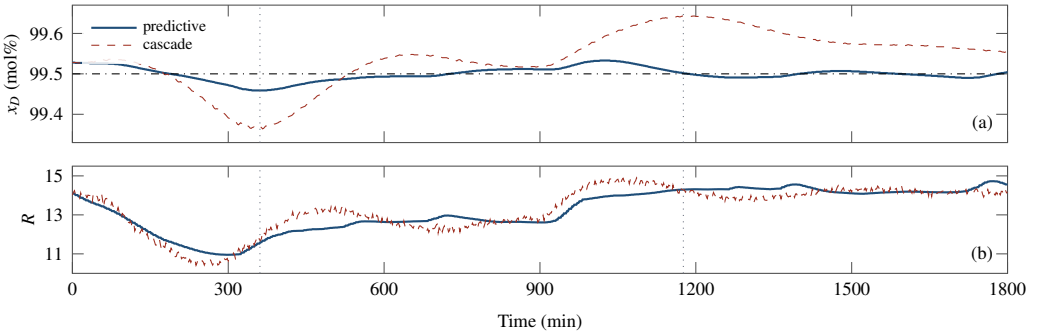

Figure 4.4 Scenario 11 under the predictive controller (thick solid line, blue) and under the cascade (dashed line, red). (a) True distillate composition against the 99.5 mol% specification (thin dash-dotted line). (b) Reflux ratio. The dotted verticals mark 361 min, the predictive controller's lowest purity, and 1176 min, the cascade's highest after the lean feed.

> **DID YOU KNOW? Hours Below Grade With a Fast Inner Loop**
>
> From the simulator (Table 4.5). On case F, the feed step with a 20 min analyzer, the cascade spent 789.5 min of the first 900 min below the 99.5 mol% specification. The analyzer working alone on the reflux ratio spent 0.0 min there.

On scenario 11 the plan's lowest purity is 99.459 mol% at 361 min against the cascade's 99.364 at 366 min, and after the lean feed the cascade carries the purity to 99.643 mol% at 1176 min while the plan stays at or below 99.533, as shown in Figure 4.4. The plan also walks the pressure setpoint onto its floor, about 16.26 bar, by 295 min, where the shipped minimization stops at 16.37 bar. The duty is 2743 GJ in every row, and the reflux the plan does not spend leaves as product, 15805 against 15 775 kmol of propylene.

The plan loses the off-specification minutes of scenario 11, 942.3 against 355.8 min, while the minutes more than 0.01 mol% below are level, 324.2 against 327.8. The root is that the target equals the specification for every structure and the reflux price holds the plan's prediction on that line, so errors of a few hundredths decide the side. Estimation error is only part of it, since with the true present purity the plan still spends 689.3 min below. The rest comes from the model's prediction, whose bias on ℓ reaches 0.304 here against 0.064 on case F, and from the first hour of each plan, which the band leaves free. With an exact estimate the plan still rides the line, 268.5 min below on case F run to 1800 min and 115.0 min on it with a 5 min cycle, all within 0.01 mol%. The free hour costs no start-up minutes, since no case has a predictive off-specification minute in its first hour, and a band from the first step still loses the tile, 797.0 against 355.8 min. The back-off in the note to Table 4.5 removes the loss at a price. Raised 0.01 mol% at a time, the target first keeps every minute within 0.01 of the specification at 99.53 mol%, with 4.26 kmol of giveaway, and at 99.55 mol% the plan spends no minute below while giving away 7.35 kmol, still less than the cascade's 8.07.

The model's reach ends where its inputs end. On the challenge with the suggested configuration and the plan in place of the cascade, the worst error after 1200 min is 0.038 mol% against 0.103, but the run spends 1166.1 min below specification against 811.5 and reaches 99.727 mol% after feed rate and cooling water steps the model does not contain.

4.7 Technology in Practice

Figure 4.5 draws the quality loops of Section 4.5, the pressure loop and the predictive option with the hardware of a splitter in place of the schematic's symbols.

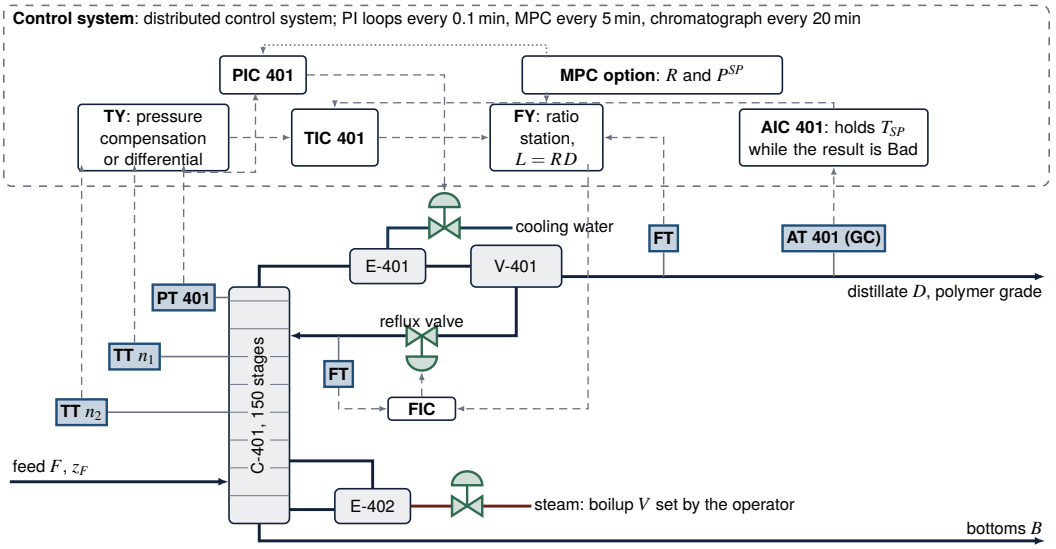

Figure 4.5 The splitter's loops with their hardware. Blue boxes are sensing devices: PT 401 on the overhead, the control and reference tray thermometers TT n_1 and TT n_2 in their thermowells, flow transmitters on the distillate and the reflux, and AT 401, the product chromatograph. Green symbols are final elements: the reflux valve in its flow loop, the cooling water valve and the steam valve, which the operator sets. Dashed arrows are signals, dotted arrows the moves that the MPC option writes. Notice that the only reading of the product arrives every 20 min, while everything the tray loop acts on is continuous. Not to scale, and not a piping diagram.

What Measures It

The controlled quality is the propylene in the distillate, read by a process chromatograph on the product line. Chromatographs are a very common choice for monitoring and controlling distillation columns [20]. Beneath it the tray loop reads a temperature in a thermowell and the overhead pressure. The chapter's chromatograph reports every 20 min, its thermowell lags by 0.6 min and its pressure transmitter is filtered over 0.1 min.

> **TECHNOLOGY** **Sensor: Resistance Thermometer in a Thermowell**
>
> A resistance thermometer is a fine wire of pure metal, usually platinum, whose resistance rises almost linearly with temperature, and 100 Ω at 0 °C is a common reference value. The measuring current warms the wire, so the reading errs high unless the current is small or pulsed [20]. Platinum elements span about −200 to 800 °C and are chosen for high accuracy and precision, best at the lower temperatures [21]. A thermowell lets the element be changed without opening the column, but its metal adds a first-order lag, and an air gap at the bottom of the well adds a second one, which can unsettle a loop. Transmitters close together on a column can disagree where the flow past them differs [20].

What Moves It

Two valves carry the loops. The reflux valve sets the internal traffic of the column through its flow loop and the ratio station, and the cooling water valve sets the overhead pressure. A third, the steam to the reboiler, is left to the operator, which is why the floating-pressure saving of Section 4.6 appears in megawatts only when someone cuts the boilup. The chapter's cooling water valve saturates on a warm afternoon.

TECHNOLOGY Final Element: The Reflux Valve Under a Ratio Station

The reflux is moved by a control valve inside a flow loop, and that loop's setpoint comes from a ratio station. The station measures a wild flow, here the distillate, multiplies it by an adjustable constant and hands the result to the captive flow's controller. Excess friction in the valve makes the stem stick and then jump, and leaves it deaf to small reversals of its signal. Poorly kept valves without positioners can show more than 10 % of this hysteresis, which up-and-down steps of the output in manual reveal [20]. The chapter's reflux is the ratio times the distillate filtered over 3 min, and a 0.09 °C setpoint step moves the ratio by 2.4 %.

Where the Controller Runs

A distributed control system has a supervisory level that sends set points to its field controllers, and its loops are sometimes cascaded, one loop's set point taken from another's process variable, at cycle times from milliseconds to minutes [22]. The app's PI loops execute every 0.1 min, the predictive controller every 5 min, and the chromatograph reports every 20 min. The structure maps onto standard blocks: an arithmetic block for the compensation or the difference, a PID block whose setpoint is written by the analyzer's PID, a ratio station, and a flow loop. The validity hold has a ready home in fieldbus function blocks, which pass each value with a status of Good, Bad or Uncertain. A smart device can flag its own data as Bad, and a block that receives a Bad input can be configured to hold its output at the last Good value [20].

What Is Modernizing It

The first development is the heat pump. In mechanical vapor recompression the overhead vapor is compressed so that it condenses hot enough to drive the reboiler. The practice is limited to mixtures whose relative volatility is close to one, which makes propane/propylene splitters important examples, and heat-integrated columns are projected to save up to 50 % more than recompression [23]. The same guide reports that floating pressure has cut winter energy use by up to 25 % in case studies. The second is the soft sensor. A 2009 survey counted 439 industrial soft-sensor applications, 331 of them on distillation, and one of their purposes is to validate an online analyzer against the estimate [24]. The compensated tray temperature of this chapter is a soft sensor of the plainest kind. The third is the layer above the loops. A 2003 survey found predictive control well founded in refining and petrochemicals [16], and the predictive controller of Section 4.5.5 removes giveaway at the price of minutes below the specification.

IN PRACTICE Giveaway Shows on the Shift Log

An energy guide for petrochemical plants advises a routine. Compare the design reflux with the ratios each shift runs, and compare the steam used with the reflux ratio and the purity every day. Many plants purify more than the specification asks, and the cure costs little: lower the reflux in small steps until the purity just meets it, which lowers the reboiler duty. In one reported case a debutanizer ran at maximum reflux and boilup to stay on grade through feed changes, and an online chromatograph kept it on specification so much more often that the energy saved paid for the analyzer in three months [23].

4.8 Limitations

The model is a teaching representation of a superfractionator and not a design tool. The separation is binary, so contaminant specifications cannot be studied. There is no energy balance and no hydraulic capacity limit, so the constraint that most often stops a real splitter, flooding before the reboiler runs out, is absent. Tray efficiency is constant, which removes a common way for a column to lose separation while every instrument reads normally.

The instrument chain is idealized. The chromatograph is a perfect sampler with a cycle time and two explicit fault modes, freeze and bias, and a real analyzer has more ways to be wrong than either. The thermowell is one lag, and matched high-resolution sensors on a real column drift relative to

one another, which is the practical reason a differential measurement is harder to maintain than its steady-state performance suggests.

Finally, the pressure-compensation coefficient is exact in the model. On a real column it is fitted at commissioning from a steady-state model, and a coefficient carried through two control system migrations without being re-derived is one of the more common quiet faults in this service. The application can show what a wrong coefficient does only by having a learner type one in.

The predictive controller's model was identified at the design point with stage 70 as the control tray, so another tray is a mismatch only the analyzer bias corrects, and beyond the horizon smaller steps fall 3.5 percent short of its ratio response and 20 percent short of its feed response at 48 h. The boilup, the feed rate and the cooling water reach it only through the tray and the analyzer. Its target is the specification, with no back-off by default.

> **OPEN CHALLENGE** **Keeping an Inferential True**
>
> Tray temperature control assumes that a temperature held at its setpoint keeps the product on specification. When the feed composition or rate of a multicomponent column changes, that assumption fails, and pressure moves the temperatures as well. Distillation is the main target of soft sensors, and in a 2009 survey of engineers that a 2013 review reports, accuracy lost to changes in the process was the problem named most often, at 29 %, which makes model maintenance the most critical issue. Progress means estimators that stay accurate without maintenance [24]. The compensation coefficient of this chapter is such a model.

4.9 Discussion

Three conclusions from this column generalize beyond distillation.

The first is that the speed of a measurement must be judged against the plant and not against a clock. A twenty-minute analyzer is slow on a heat exchanger and fast on a superfractionator, and on this column the measured penalty for moving the analyzer cycle from 5 to 60 min is 0.009 mol% of peak departure. Effort spent shortening that cycle would be effort wasted, while effort spent on validity logic is worth 0.27 mol% of settled composition when the chromatograph freezes.

The second is that a proxy measurement should be selected for fidelity rather than for sensitivity. The most sensitive tray in the column is a poor inferential for the distillate because it responds to a great deal the distillate does not see. The rule of thumb that selects the steepest point of the temperature profile answers a different question, which is where the signal-to-noise ratio is best, and the two answers coincide only when the sensor is near the product.

The third is that removing a disturbance from a measurement is not the same as removing it from the process. The differential temperature cancels pressure at the thermometer almost perfectly and still lets the distillate fall 0.63 mol% through a held 0.6 bar pressure change, because pressure enters the column twice, once through the bubble point and once through the relative volatility. A compensation scheme derived from the measurement equation alone cannot know that.

> **OPEN CHALLENGE** **Separating Propylene Without a Tall Column**
>
> Distillation needs more than 100 stages to make polymer-grade olefins, because the relative volatility of propylene and propane is low, and a 2020 review sees membranes as the promising alternative. A membrane with a selectivity above 10 could be paired with distillation to save energy, and one above 35 could work alone. Yet the carriers that make some membranes selective are unstable, through the reduction of silver ions, the disproportionation of copper ions and poisoning by sulfides, which limits long operation, and framework membranes still need ways to be made at large scale. Progress means membranes that stay selective in long operation and can be built at industrial scale [6].

4.10 Conclusion

A propylene splitter is a plant on which the first experiment a control course teaches cannot be performed. The simulated column of this chapter is 53 percent complete four hours after a three

percent reflux ratio step and needs 715 min to fall below 20 ppm/h of impurity drift, so the identification must be done on a secondary measurement and the primary measurement must arrive intermittently. The tray temperature that results is a bubble point, and its 3.10 °C/bar pressure term is large against a composition profile spanning 5.80 °C over 151 stages, which is why an uncompensated measurement runs the reflux ratio into its stop on a 0.4 bar excursion. Measured on the same feed disturbance, the analyzer placed directly on the ratio station departs 0.063 mol% against 0.143 for the cascade and 0.373 for the inferential alone, and the ordering reverses the moment the chromatograph freezes, where a cascade that keeps integrating the stale result until its validity hold is armed settles at 98.90 mol% and is worse than having no cascade at all. Selecting the control tray for sensitivity rather than fidelity costs 0.19 mol% of composition at fifteen hours. A model predictive controller, the fifth setting, cuts the giveaway of scenario 11 from 8.1 to 1.0 kmol and loses its off-specification minutes, 942 against 356 min, by riding the specification. A fuller model would add an energy balance so that reflux subcooling and feed enthalpy reach the internal traffic, a hydraulic capacity limit so that flooding can end an experiment the way it ends a real push, and a multicomponent feed so that a product can fail on something this separation does not address. The superfractionator is an accessible browser platform for learning what a soft sensor measures and what a supervisory layer is for.

KEY POINTS

- The column answers so slowly that a step test does not finish in a shift, so a tray thermometer stands in for a product analyzed every twenty minutes.
- A tray temperature is a bubble point, and its pressure term of 3.10 °C/bar is large against a profile of 5.80 °C.
- On the feed step the analyzer alone departs 0.063 mol% against 0.143 for the cascade, and with the chromatograph frozen the cascade without its hold settles at 98.90 mol%.
- Choosing the tray for sensitivity rather than fidelity costs 0.19 mol% at fifteen hours.

Where else this applies. A fast inferred measurement corrected by a slow assay returns in Volume 5, Chapter 11 and Volume 19, Chapter 10, and an inferential measurement carrying bias and noise in Volume 15, Chapter 3.

References

[1] S. Skogestad, “Dynamics and control of distillation columns: A critical survey,” in *Dynamics and Control of Chemical Reactors, Distillation Columns and Batch Processes*. Elsevier, 1993, pp. 11–35.

[2] M. T. Guilandoust, A. J. Morris, and M. T. Tham, “Estimation and control of distillation product composition using tray temperature measurements,” in *Dynamics and Control of Chemical Reactors and Distillation Columns*. Elsevier, 1988, pp. 203–208.

[3] F. Al Kalbani and J. Zhang, “Inferential composition control of a distillation column using active disturbance rejection control with soft sensors,” *Sensors*, vol. 23, no. 2, p. 1019, 2023.

[4] V. Gokhale, S. Hurowitz, and J. B. Riggs, “A comparison of advanced distillation control techniques for a propylene/propane splitter,” *Industrial & Engineering Chemistry Research*, vol. 34, no. 12, pp. 4413–4419, 1995.

[5] A. I. Hinojosa and D. Odloak, “Using Dynsim to study the implementation of advanced control in a propylene/propane splitter,” *IFAC Proceedings Volumes*, vol. 46, no. 32, pp. 33–38, 2013.

[6] Y. Ren, X. Liang, H. Dou, C. Ye, Z. Guo, J. Wang, Y. Pan, H. Wu, M. D. Guiver, and Z. Jiang, “Membrane-based olefin/paraffin separations,” *Advanced Science*, vol. 7, no. 19, p. 2001398, 2020.

[7] S. Skogestad, “Dynamics and control of distillation columns: A critical survey,” *Modeling, Identification and Control*, vol. 18, no. 3, pp. 177–217, 1997.

[8] M. R. Fenske, “Fractionation of straight-run Pennsylvania gasoline,” *Industrial & Engineering Chemistry*, vol. 24, no. 5, pp. 482–485, 1932.

[9] A. J. V. Underwood, “Fractional distillation of multicomponent mixtures,” *Industrial & Engineering Chemistry*, vol. 41, no. 12, pp. 2844–2847, 1949.

[10] H. W. Andersen, M. Kümmel, A. Nørgaard Hansen, and K. Nielsen, "Tuning of dual-composition distillation column control," *Chemical Engineering Science*, vol. 44, no. 3, pp. 619–630, 1989.

[11] F. G. Shinskey, "Distillation," in *Energy Conservation Through Control*. Elsevier, 1978, pp. 251–283.

[12] M. Porru, R. Baratti, and J. Alvarez, "Energy saving through control in an industrial multicomponent distillation column," *IFAC-PapersOnLine*, vol. 48, no. 8, pp. 1138–1143, 2015.

[13] T. Liu, "Research on multi-objective energy-saving optimization of propylene distillation column," *Measurement and Control*, vol. 58, no. 4, pp. 515–524, 2024.

[14] H. Abedini, M. Asgari, M. Watt Coull, A. Shariati, and M. R. Khosravi-Nikou, "Efficient production of polymer-grade propylene from the propane/propylene binary mixture using Cu-MOF-74 framework," *Separation and Purification Technology*, vol. 276, p. 119172, 2021.

[15] 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.

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

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

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

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

[20] 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] A. Boghossian, J. Brown, and S. Zak, "3.2: Temperature sensors," in *Chemical Process Dynamics and Controls*, P. Woolf, Ed. University of Michigan open textbook, LibreTexts, 2023.

[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] M. Neelis, E. Worrell, and E. Masanet, "Energy efficiency improvement and cost saving opportunities for the petrochemical industry: An ENERGY STAR guide for energy and plant managers," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-964E, 2008.

[24] 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.

4.11 Exercises

4.1 ○ Concept check · 2 min

On the feed composition step of scenario 8, moving the analyzer cycle time from 5 to 60 min changes the peak purity departure of the analyzer-only loop only from 0.061 to 0.070 mol%. Which explanation is right?

(a) The loop is tuned so slowly that it never acts on a single analyzer result

(b) The cycle is short against a response that is 63 percent complete only at 5.2 h

(c) The tray temperature carries the feed change to the loop before the analyzer reports

(d) The analyzer noise is larger than anything a shorter cycle could have detected

4.2 ○ Concept check · 2 min

The most sensitive stage in the column is stage 127, about eight times as sensitive to the reflux ratio as stage 70. Why is it a poor control tray for the distillate?

(a) Its thermowell noise is amplified eight times by the higher sensitivity

(b) It sits below the feed, so a change in the reflux ratio never reaches it

(c) It responds to a great deal that the distillate composition does not see

(d) Its temperature is fixed by the reboiler duty and does not move with reflux

4.3 Concept check · 2 min

The difference between stages 90 and 70 moves only 0.0012 °C when the pressure rises by a tenth of a bar, against 0.312 °C for a single tray. Held near its setpoint through a 0.6 bar pressure step, it still lets the distillate fall 0.63 mol%. Why?

(a) The differential loop drives the reflux ratio onto its limit of 24 after the step

(b) The two thermowells have different lags, so the difference reads the transient as composition

(c) The difference cancels the pressure term only while the pressure is changing, not after

(d) Pressure also lowers the relative volatility, which moves the profile the difference stands for

4.4 Theory and calculation · 6 min

The splitter separates propylene from propane at a relative volatility $\alpha = 1.12$ at 18 bar. For a binary pair, the Fenske minimum stage count and the Underwood minimum reflux ratio for a saturated liquid feed are

$$N_{min} = \frac{\ln\left[\frac{x_D}{1-x_D}\,\frac{1-x_B}{x_B}\right]}{\ln \alpha}, \qquad R_{min} = \frac{1}{\alpha - 1}\left[\frac{x_D}{z_F} - \frac{\alpha\,(1-x_D)}{1-z_F}\right].$$

(a) Find N_{min} for the design products, $x_D = 0.99528$ and $x_B = 0.05986$. *Checkpoint:* 71.5 stages.

(b) Find R_{min} for $x_D = 0.995$ and $z_F = 0.72$. *Checkpoint:* 11.35.

(c) Floating pressure takes the column to 15.5 bar. Use Equation (4.3) to find α there, repeat (a) and (b), and compare the design reflux ratio of 14.17 with R_{min} at both pressures. *Checkpoint:* $\alpha = 1.13875$, 62.4 stages, $R_{min} = 9.81$.

4.5 Theory and calculation · 5 min

The column boils up $V = 7494\,\mathrm{kmol/h}$ at a reflux ratio $R = 14.17$, and the drum level holds $D = V - L$ under constant molar overflow.

(a) Use Equation (4.6) to find the distillate and the reflux flows. *Checkpoint:* 494.0 kmol/h and 7000 kmol/h.

(b) The reflux flow is raised by 1 % at constant boilup. Find the new distillate flow and its change in percent. *Checkpoint:* 424.0 kmol/h, −14.2 %.

(c) The reflux ratio is raised by 1 % instead. Find the new distillate flow and its change in percent, and explain why the manipulated variable is the ratio. *Checkpoint:* 489.4 kmol/h, −0.93 %.

4.6 Theory and calculation · 8 min

Stage 70 holds liquid at $x = 0.82$. Use Equations (4.4) and (4.5) with the parameters of Section 4.3, and take the steady-state gain of stage 70 as −0.047 °C per 1 % increase of the reflux ratio.

(a) Find the rise of the stage-70 temperature when the pressure goes from 18.0 to 18.1 bar at constant composition. *Checkpoint:* 0.312 °C.

(b) Express that rise as the change of reflux ratio, in percent, that an uncompensated tray loop would make to cancel it. *Checkpoint:* 6.6 %.

(c) The pressure goes from 18.0 to 18.4 bar. Find the error left in the compensated reading of Equation (4.8) with $P_{ref} = 18\,\mathrm{bar}$ and $\mathrm{d}T/\mathrm{d}P = 3.10\,^{\circ}\mathrm{C/bar}$, in °C and in percent of reflux ratio. Then name what makes the ratio settle 1.7 % above its starting value in scenario 5. *Checkpoint:* 0.0028 °C, 0.06 %.

4.7 App challenge · 8 min

Load scenario 1, “1 · The step test that does not finish (MANUAL)”, at the default seed. The case starts in Manual. Before Run, set Reflux ratio to 14.60. Run at speed 200 (200 min per wall-clock

second) until the clock passes 900 min, pause and export the CSV. A run longer than an hour exports its time column as Time (h), in hours, so convert each time below from minutes.

(a) In column `settle_ppm_h`, the drift rises above 100 ppm/h soon after the step. Find the first later row below 20 ppm/h. *Checkpoint:* 715 min.

(b) In column `xD_pct`, read the rows at 120.0 and 240.0 min. The step finally takes the distillate from 99.528 to 99.638 mol%. What fraction of the change has arrived at each time? *Checkpoint:* 99.560 and 99.586, 29 and 53 percent.

(c) Explain what a first-order fit to the first two hours would report about this column, and why the drift column is the better guide. *Checkpoint:* A time constant too short and a gain too low.

4.8 App challenge · 8 min

Load scenario 5, "5 · The thermometer is a pressure gauge first", at the default seed with no change. The pressure setpoint steps from 18.0 to 18.4 bar at 60 min, and the measurement switches to Pressure-compensated at 300 min. Run at speed 200 until the clock passes 2400 min (about 12 s), pause and export the CSV. A run longer than an hour exports its time column as Time (h), in hours, so convert each time below from minutes.

(a) Column `reflux_saturated` turns to 1 when the ratio sits on a limit. Find the first such row after 60 min, and read the commanded ratio in column `reflux_ratio_cmd` of that row. *Checkpoint:* 91.8 min, 24.00.

(b) In the row at 300.0 min, read `xD_pct`, `tray_raw_C` and `tray_comp_C`, and explain the difference between the two temperatures. *Checkpoint:* 99.892, 47.410 °C and 46.170 °C.

(c) After 300 min, find the lowest `reflux_ratio_cmd` and its time, and read `reflux_ratio_cmd` and `alpha` in the row at 2400.0 min. *Checkpoint:* 13.56 at 750 min, then 14.37 and 1.1170.

4.9 App challenge · 10 min

Load scenario 11, "11 · The answer arrives tomorrow; the plan is made now", at the default seed. Quality control is on MPC. For each run below, run at speed 200 until the clock passes 1800 min (about 9 s), pause and export the CSV, and read `metric_giveaway_kmol` and `metric_offspec_min` in the row at 1800.0 min and the lowest `xD_pct` with its time. A run longer than an hour exports its time column as Time (h), in hours, so convert each time below from minutes.

(a) Run the case as loaded. *Checkpoint:* 0.98 kmol, 942.3 min, 99.459 at 361 min.

(b) Press Reset, set Quality control to Cascade before Run, and run again. *Checkpoint:* 8.07 kmol, 355.8 min, 99.364 at 366 min.

(c) Press Reset, set Quality control back to MPC and Propylene purity target to 99.53 mol% before Run, and run again. Explain why the minutes below specification fall. *Checkpoint:* 4.26 kmol, 111.0 min, 99.491.

4.10 Application area · 4 min

A site energy engineer commissions pressure minimization on this column. A month later the reboiler duty has not changed. Using the results of scenario 10, explain why, what must be added for the saving to appear as energy, and what that addition must do on a warm afternoon.

4.11 Application area · 3 min

An operations manager wants to use the simulator to decide how far the feed can be pushed above 790 kmol/h by raising the boilup. Name two things the model leaves out that decide that answer on a real splitter, and say what the app would show instead.

4.12 Review the agent · 5 min

An AI assistant reviewing this column writes: “Table 4.4 shows that the analyzer alone departs 0.063 mol% on the feed step against 0.143 mol% for the cascade, and is back at 99.51 at fifteen hours while the cascade is at 99.47. The TIC-401 inner loop therefore adds nothing. Put AIC-401 straight on the ratio station and retire the tray thermowells.” Find the flaw, and say how you would check it in the app.

CHAPTER 5

Cyclic Propane Dehydrogenation Beds Whose Regeneration Sets the Next Cycle

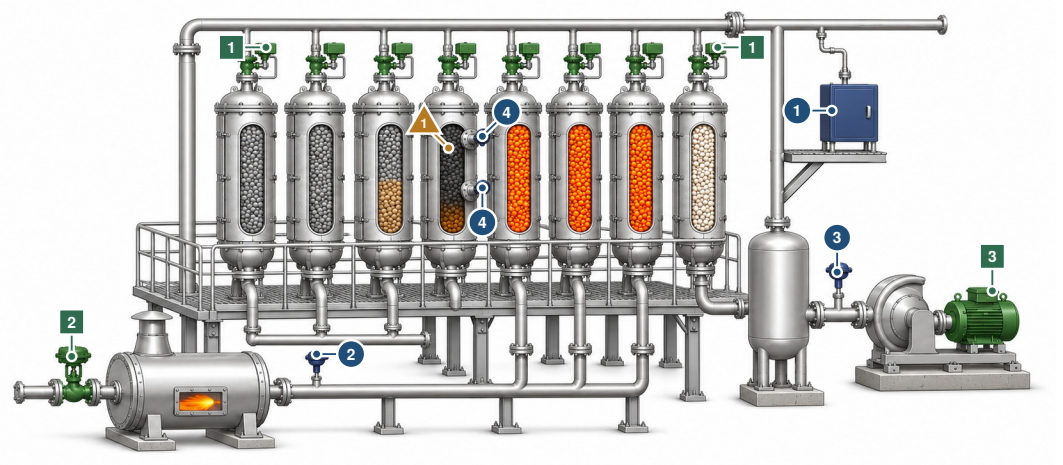

The controller moves 1 valves that switch each bed between its four phases, worked by the sequencer
2 fuel valve of the air heater, which sets the regeneration air temperature
3 compressor drive, whose speed holds the vacuum on the beds

The controller reads 1 product analyzer, one result for each three-minute slot 2 regeneration air thermometer
3 suction pressure transmitter, which reads the vacuum 4 bed thermocouples, which read every part of a bed

What disturbs it 1 coke on the catalyst, which grows as the catalyst ages ▲ a bed whose valves fail
▲ more propane fed to each bed

Notice that three beds make product while three others take the hot air, so a change to the air reaches only beds not yet on line, and that the beds of the model lie horizontal rather than upright. Illustration, not to scale.

IN PLAIN TERMS **Beds That Run on Stored Heat**

The main idea. Eight beds of catalyst take turns turning propane gas into propylene. While a bed makes product, nothing heats it. It spends heat it stored like a battery, then goes off line to be reheated with hot air.

The challenge. The only way to change what a bed makes is to heat it more before its next turn. So every correction arrives one or two turns late. When a bed breaks down, the order of the turns matters more than any controller.

Why it matters. The way the turns are set decides whether the plant keeps its output, its target or a steady flow of product.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	This page, the History boxes in Section 5.2, the Did you know? boxes in Sections 5.1 and 5.5, the Try it boxes in Section 5.5, and Key points in Section 5.8. No equations.
Study	45 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 5.5 with Figure 5.5, Section 5.6, and Section 5.7 with its Open challenge box.

The control question A dehydrogenation bed on line spends heat it stored in its last regeneration, so no move reaches the beds now making propylene. How far can a regeneration temperature set once a slot hold the conversion through a lost bed and a feed step, and what does the schedule decide first?

Learning objectives

- Compute a bed's heat store, its cooling on line and the equilibrium conversion at the bed pressure.
- Explain why a regeneration move reaches the conversion one to two cycles late, and why a deeper vacuum buys selectivity rather than conversion.
- Tune a cycle-to-cycle PI for a delay of reports and compare it with a predictive controller that measures the feed.
- Choose a seven-bed schedule after a lost bed and price it in conversion, propylene and product ripple.

Prerequisites Sections 1.4.7 (cyclic units and cycle-to-cycle control), 1.4.3 (endothermic beds), 1.6.2 and 1.6.4. Section 1.6.3 only for the predictive controller.

The app https://apmonitor.com/apps/propane-dehydrogenation/

See also Volume 8, Chapter 9 for corrections made while results are still in flight, and Volume 8, Chapter 10 for learning from one layer to the next. Volume 9, Chapter 2 for a step-response predictive controller behind a dead time, and Volume 9, Chapter 4 for stored energy drawn down fast and restored slowly.

5.1 Introduction

Nothing heats a dehydrogenation bed while it makes propylene. In the generic unit simulated here, each of eight beds holds 33 t of chromia catalyst and inert material at 1.1 kJ/(kg K), 36.3 MJ for every kelvin, and receives 10 t/h of propane at 600 °C and 0.50 bar a. In its nine minutes on line the reactions absorb 2.44 GJ, the bed supplies 2.36 GJ of that from its own heat by cooling 65.0 K, and its conversion falls from 69.9 to 44.7 %. Steam purge, nine minutes of hot air that also burns the 10.1 kg of coke laid down, and an evacuation return it to the line at 622.6 °C after 24 min. Three beds are on line at every instant, staggered by one 3 min slot, and their product is compressed to a propylene splitter of the kind treated in Chapter 4.

> **DID YOU KNOW? The Third Source of Propylene**
>
> A 2023 review counts propane dehydrogenation as the third-largest source of propylene in the world, with 4.5 percent of the total capacity [1]. The bank simulated here sends 14.57 t/h of propylene to its splitter (Table 5.1).

The simulator addresses a belief that a block diagram encourages, that a controller corrects the conversion by acting on the reactor that makes it. In this plant no move reaches a bed on line. When the regeneration air is raised 10 K, the first bed to feel it comes on line 5 min later only 1.2 K warmer, and the bank's conversion has made 78 percent of its final rise of 2.93 points 26 min after the move. The correction is made on the next batch of stored heat, from a measurement of the last one, which is the structure of run-to-run and batch-to-batch control in semiconductor and batch chemical manufacture [2, 3]. A second belief is that a better equilibrium buys conversion. Lowering the suction from 0.45 to 0.35 bar a raises the equilibrium conversion of propane at 600 °C from 63.4 to 67.6 %, while the bank's conversion moves from 57.80 to only 58.15 % because the heat stored in the beds, not the equilibrium, sets how much propane they can convert. What the vacuum does buy is selectivity, 88.0 to 90.7 %. Fixed beds that store the heat of one period for the next are a class of their own, of which reverse-flow reactors [4] and the regenerative crackers and dehydrogenators of the refinery [5, 6] are members.

The literature on this process divides along disciplinary lines. Reviews treat the catalysts and the competing technologies [1, 7–9], kinetic studies fit transient rates on a single catalyst [10], and a dynamic model of a commercial reactor follows one bed through reaction and regeneration for

operational optimization [6], while the run-to-run literature draws its examples from batch and semiconductor processes [2]. This chapter works with one simulated bank whose beds, sequencer, regeneration air, vacuum and analyzer interact. It measures how far a regeneration temperature chosen once a slot can hold the conversion through a lost bed and a feed step, and what the schedule decides before any controller acts. The kinetics, the shared bed stack, the velocity-form PI law and the linear model predictive controller (MPC) are all established methods. The results are measurements on a single synthetic unit, and the most consequential is that no regeneration within the bed limit lets seven beds with three on line reach the 57 % target (52.6 % at 680 °C air), so that the schedule, not the loop, decides the twelve-hour challenge. The model has 114 states, its energy and carbon ledgers close to 4×10^{-12} GJ and 2×10^{-8} kmol over every case, each run starts on its own cyclic steady state, and the controllers read instruments only. Nine cases each carry a completion rule written against the state of the bank. The general picture of a cyclic unit corrected once a report is in Section 1.4.7.

Section 5.2 follows the regenerative fixed bed from catalytic cracking to dehydrogenation and places the control problem among run-to-run methods. Section 5.3 sets out the bank: the reactions and equilibrium, a bed on the shared cell stack, purge, regeneration and evacuation, the surge and the compressor, the sequencer and the measures of performance, the instruments and the assumptions. Section 5.4 describes the control layers, from the two regulatory loops to the cycle-to-cycle PI and the MPC. Section 5.5 runs the nine cases, from where the heat goes and the regeneration as the handle to a lost bed, aging, the vacuum and the twelve-hour challenge. Section 5.6 names the hardware and the practice behind the model, Section 5.7 states what the model leaves out, and Section 5.8 sums up.

5.2 Regenerative Fixed Beds, From Cracking to Dehydrogenation

The Houdry process cracked oil over fixed catalyst beds, and Eugene Houdry's contributions to catalytic cracking are reviewed in [11], while the History box below follows its first commercial unit and the equipment it relied on. The sequence of a modern cyclic dehydrogenation bed, reaction, purge, regeneration and evacuation, keeps the same elements, and the same reliance on valves and vacuum.

> **HISTORY** **Valves on Timers, 1937**
>
> In the early 1920s the French engineer Eugene Houdry began looking for a catalyst that would make gasoline from lignite, and he later turned to petroleum. His process gained a way to regenerate the catalyst after only ten minutes of use, and in March 1937 the first fully commercial unit went into operation. It processed 15 000 barrels of petroleum a day with a molten-salt heat control technique and motor-operated valves controlled by timers [12]. A paper of 1947 turned to the equipment such fixed beds relied on: steam-jet ejectors that evacuated the air from the reactors, removed oil vapors during the purge periods and held a vacuum on the closed heat-transfer salt systems [5].

The chemistry of propane dehydrogenation suits this arrangement. The reaction absorbs heat and is limited by equilibrium, so it needs a high temperature and benefits from a low pressure, while selectivity is hard to hold at the temperatures the reaction needs [7, 9]. In a cyclic reactor the catalyst loses activity quickly to carbon deposits and is restored by periodic regeneration [6]. Alkane dehydrogenation over supported chromium oxide is reviewed in [13], the mechanisms by which catalysts deactivate, fouling by carbon among them, are reviewed in [14], and transient kinetic models of propane dehydrogenation have been fitted on a platinum–tin catalyst [10]. Reviews compare the commercial dehydrogenation technologies and their operating windows [8, 9]. A recent review tabulates the cyclic fixed-bed process known as Catofin against the continuous moving-bed alternative, with chromia on alumina, fixed beds, 600 °C to 700 °C, a cycle of 15 to 30 min and a propane conversion of 48 to 65 % [1]. A dynamic model of a commercial Catofin reactor for isobutane describes beds of chromia on alumina cycling through dehydrogenation, regeneration, evacuation,

purge and reduction every 18 to 30 min, and names the gradual loss of heat in the bed together with deactivation as the operating concern [6].

> ⧖ HISTORY **Chromia and the Heat Problem, From 1933**
>
> Chromium oxide on alumina became the commercial catalyst of choice for dehydrogenation after the work of Frey and Huppke, published in 1933. Companies that set out to build a process on it met three problems: supplying enough heat to the reactor, controlling the temperature closely enough to hold the conversion up and the degradation products down, and regenerating the catalyst. Every industrial process that followed uses the heat of burning the coke in the dehydrogenation step, and the cyclic version runs a series of adiabatic fixed beds that alternate quickly between reaction and regeneration. In the 1970s and 1980s the main interest was butane, dehydrogenated toward synthetic rubber. By 1999 propane and isobutane had become more important [13].

Control of such a unit is layered. The sequencer owns the valves, ordinary loops hold the regeneration air and the compressor suction, and the conversion is corrected once a cycle from an analyzer, which places it in the family of run-to-run, iterative learning and repetitive control, where each run is adjusted from the measured outcome of the last [2, 3]. Iterative learning on a part built layer by layer is the subject of Volume 8, Chapter 10. Model predictive control, widely applied as supervisory control in the process industries [15], fits the same cadence when its model is written slot by slot. The literature reviewed here does not describe a cycle-to-cycle loop on a commercial unit, and the simulator's loop is offered as a teaching design rather than as industrial practice.

5.3 The Bank as Modeled

The bank is generic and drawn from no plant or licensor, as shown in Figure 5.1. Eight horizontal adiabatic reactors A to H take turns in a fixed rotation. A bed on line receives propane F_{in} from a charge heater at T_{feed} and sends its effluent to a product header, a bed in regeneration receives hot air at T_{air}, a bed in purge receives steam, and an evacuated bed is drawn to the vacuum header with no flow. The product passes a surge volume V_s at the suction pressure P_s to a compressor at speed n_c, and the beds run ΔP above the suction. The regeneration air leaves an air heater at T_{ah} and is raised to T_{air} by fuel gas injected through a valve. The engine's clock is the minute, flows inside a bed are in mol/s, heats in GJ and temperatures in °C. The engine and its clocks are those of Section 1.6.1. Each bed carries twelve states, six cell temperatures T_i and six coke masses $m_{coke,i}$, and the plant adds five more and thirteen ledgers, 114 in all.

Table 5.1 gives each variable's role and its value in case 2 at $t = 0$, which is the design point.

5.3.1 Reactions and Equilibrium

Three reactions run on line, the dehydrogenation $C_3H_8 \rightleftharpoons C_3H_6 + H_2$, thermal cracking $C_3H_8 \rightarrow CH_4 + C_2H_4$, and coking $C_3H_6 \rightarrow 3\,C + 3\,H_2$, at the rates of Equation (5.1).

$$
\begin{aligned}
r_1 &= a\,k_1 \left(\frac{P}{P_0}\right)^{n-1} \left(p_{C_3H_8} - \frac{p_{C_3H_6}\,p_{H_2}}{K_p}\right) \\
r_2 &= f_{cr}\,k_2\,p_{C_3H_8}, \qquad r_3 = a\,f_{ck}\,k_3\,p_{C_3H_6} \\
k_j &= k_{j,0} \exp\left[-\frac{E_j}{R}\left(\frac{1}{T} - \frac{1}{T_r}\right)\right] \\
K_p &= \exp\left(\frac{\Delta S_1}{R} - \frac{\Delta H_1}{RT}\right), \qquad a = a_{cat} \exp\left(-\frac{m_{coke}}{c_D W}\right)
\end{aligned}
\tag{5.1}
$$

where p_k are partial pressures in bar at the bed pressure P, $P_0 = 0.5\,\text{bar}$, $T_r = 873.15\,\text{K}$, $n = 0.3$ is the order of the dehydrogenation in total pressure, a is a cell's activity, W its mass, $c_D = 1.6\,\text{kg}$ of coke per tonne of bed lowers the activity to 37 percent, and a_{cat}, f_{cr} and f_{ck} are the disturbance settings. Cracking is thermal and does not depend on the activity. With $\Delta H_1 = 129.5\,\text{kJ/mol}$ and

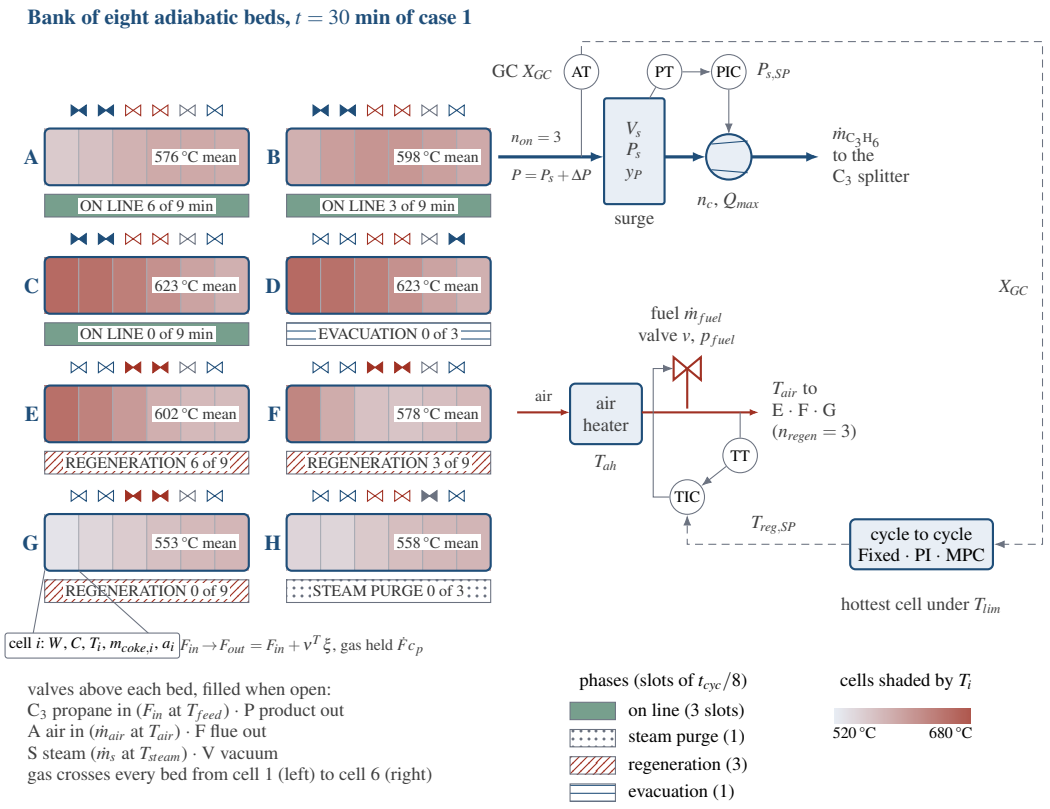

Figure 5.1 The eight-bed bank as modeled, redrawn from the Bank view of case 1 at $t = 30$ min with every station held in Manual. Beds A, B and C are on line (6, 3 and 0 min into their nine), D is evacuated, E, F and G regenerate in 660 °C air and H is purged with steam. Cells are shaded by temperature, from 523 °C at G's inlet to 658 °C at C's, and the valve sets show which gas each bed receives. The product side carries 14.35 t/h of propylene at a suction of 0.446 bar a, and the GC (AT) feeds the cycle-to-cycle layer, which writes the TIC's setpoint. Bed A's own chart over the same run is Figure 5.2. Not to scale.

$\Delta S_1 = 139.2\,\mathrm{J/(mol\,K)}$, $K_p = 0.334$ bar at 600 °C, and the equilibrium conversion of pure propane, $X_{eq} = \sqrt{K_p/(P+K_p)}$, is 63.3 percent at 0.5 bar a and 50.0 percent at 1 bar a. The thermodynamic polynomials of NASA Glenn [16] give $\Delta H_1 = 130.01$ kJ/mol, $\Delta S_1 = 139.28\,\mathrm{J/(mol\,K)}$ and $K_p = 0.315$ bar at 600 °C, or 62.2 and 48.9 %, so the app's equilibrium sits about one point high. The difference comes almost entirely from the app's ΔH_1, 0.5 kJ/mol under the tabulated value, and is kept as a teaching value.

5.3.2 A Bed on the Shared Cell Stack

Each bed is built on the packed-bed cell stack shared by the framework's reactor applications, which solves a cell's composition algebraically and carries its temperature as a state. The stack is assembled here as six one-cell beds per reactor so that each cell carries its own coke and activity. On line, cell i

Table 5.1 Roles and Design Values of the Variables

Variable	Role	Design	Range
X_{GC}	CV, bank conversion (GC)	57.8 %	SP 45–68
T_{air}	CV, air (TIC)	660.0 °C	SP 600–720
P_s	CV, suction (PIC)	0.450 bar a	SP 0.25–0.80
$T_{reg,SP}$	MV, cycle to cycle	660 °C	600–T_{lim}
v	MV, fuel valve	49.9 %	0–100
n_c	MV, compressor speed	74.8 %	0–100
t_{cyc}	set, cycle time	24 min	16–40
F	set, feed per bed on line	10 t/h	6–14
—	set, with a bed out	leave the gap	3 options
p_{fuel}	DV, fuel gas pressure	100 %	40–110
α_c	DV, compressor available	100 %	40–100
a_{cat}	DV, catalyst activity	100 %	50–100
f_{ck}, f_{cr}	DV, coke make, cracking	1, 1	0.5–3
S, $\dot{m}_{C_3H_6}$	outputs	88.0 %, 14.57 t/h	—
T_i, $m_{coke,i}$	states, per cell (96)	—	—
x_v, T_{air}, P_s	states, plant	—	—
n_c, y_P	states, plant	—	—

obeys Equation (5.2).

$$\begin{aligned} F_{out,i} &= F_{out,i-1} + \nu^T \xi_i, \qquad \xi_{j,i} = W\, r_j(T_i, F_{out,i}) \\ C\,\frac{\mathrm{d}T_i}{\mathrm{d}t} &= \dot{F} c_p\,(T_{i-1} - T_i) - \textstyle\sum_j \xi_{j,i}\,\Delta H_j \\ \frac{\mathrm{d}m_{coke,i}}{\mathrm{d}t} &= 3\,M_C\,\xi_{3,i} \end{aligned} \tag{5.2}$$

where $F_{out,0} = F_{in}$ is pure propane, the first cell's inlet temperature T_{i-1} is T_{feed}, $W = 5.5\,\mathrm{t}$, $C = 6.05\,\mathrm{MJ/K}$, ν holds the stoichiometry of the three reactions, $\Delta H_2 = 82\,\mathrm{kJ/mol}$, $\Delta H_3 = -20\,\mathrm{kJ/mol}$, M_C is the molar mass of carbon, and $\dot{F} c_p$ is the propane fed times 160 J/(mol K), held along the bed. Each cell's extents are found by six Newton iterations from zero with no convergence test, so a sweep of the bed depends only on its inputs and runs replay bit for bit. At 10 t/h, $F_{in} = 62.99\,\mathrm{mol/s}$ and $\dot{F} c_p = 10.08\,\mathrm{kW/K}$, so the gas would need 10.0 min to change a cell's temperature by itself, longer than the on-line period. The bed's temperature therefore follows its heat of reaction, and the cell nearest the inlet, where the propane is richest, cools first and furthest. The stack itself is checked apart from any plant, against the tanks-in-series answer, the plug-flow limit and its own energy closure, and its outlet error halves when the number of cells doubles.

5.3.3 Purge, Regeneration and Evacuation

Off line a different gas crosses the bed. In regeneration the coke laid down on line burns and releases its heat in the cell where it lies, so a bed that made more coke is reheated more, whatever the air temperature. Purge, regeneration and evacuation differ only in the gas, as shown in Equation (5.3).

$$\begin{aligned} C\,\frac{\mathrm{d}T_i}{\mathrm{d}t} &= \dot{m}_g c_g\,(T_{i-1} - T_i) + \Delta H_b\, k_b(T_i)\, m_{coke,i} \\ \frac{\mathrm{d}m_{coke,i}}{\mathrm{d}t} &= -k_b(T_i)\, m_{coke,i} \end{aligned} \tag{5.3}$$

where in regeneration $\dot{m}_g c_g$ is the air per bed, $\dot{m}_{air} = 40\,\mathrm{kg/s}$ at 1150 J/(kg K), the first cell's inlet is T_{air}, $\Delta H_b = 32.8\,\mathrm{MJ/kg}$ and $k_b = 1\,\mathrm{min}^{-1}$ at 650 °C with an activation energy of 100 kJ/mol, first order in the coke with the oxygen in excess. In purge the gas is $\dot{m}_s = 3\,\mathrm{kg/s}$ of steam at $T_{steam} = 450\,^{\circ}\mathrm{C}$ and nothing burns, and an evacuated bed, or one out of service, holds its state. The air's time constant on a cell, $C/(\dot{m}_{air} c_{air}) = 2.19\,\mathrm{min}$, is short against the nine-minute regeneration,

so the air front crosses the bed and leaves its inlet cells near T_{air}. One regeneration air header serves every regenerating bed, as shown in Equation (5.4).

$$\begin{aligned} T^{*}_{air} &= T_{ah} + \frac{\eta\,\mathrm{LHV}\,\dot{m}_{fuel}}{n_{regen}\,\dot{m}_{air}\,c_{air}} \\ \dot{m}_{fuel} &= x_v\,\dot{m}_{f,max}\sqrt{p_{fuel}} \\ \tau_{air}\frac{\mathrm{d}T_{air}}{\mathrm{d}t} &= T^{*}_{air} - T_{air}, \qquad \tau_v\frac{\mathrm{d}x_v}{\mathrm{d}t} = v - x_v \end{aligned} \tag{5.4}$$

where $T_{ah} = 600\,^{\circ}\mathrm{C}$, $\eta = 0.98$, $\mathrm{LHV} = 47\,\mathrm{MJ/kg}$, $\dot{m}_{f,max} = 0.36\,\mathrm{kg/s}$, p_{fuel} is the fuel gas pressure as a fraction of design, n_{regen} is the number of beds regenerating, $\tau_{air} = 0.1\,\mathrm{min}$ and $\tau_v = 0.05\,\mathrm{min}$. At design the valve stands at 49.9 % and burns 0.647 t/h, and with fewer beds regenerating the same fuel heats less air further.

5.3.4 Surge, Compressor and Vacuum

The beds' effluent mixes in the surge volume, whose pressure and propylene share follow the balances of Equation (5.5).

$$\begin{aligned} \frac{V_s}{RT_s}\frac{\mathrm{d}P_s}{\mathrm{d}t} &= \sum_{on} F_{out} - n_c\,\alpha_c\,Q_{max}\frac{P_s}{RT_s} \\ N_s\frac{\mathrm{d}y_P}{\mathrm{d}t} &= \sum_{on} F_{out}\,(y_{in} - y_P), \qquad N_s = \frac{P_s V_s}{RT_s} \end{aligned} \tag{5.5}$$

where F_{out} is a bed's total outlet flow, y_{in} the propylene share of the beds' combined effluent, $V_s = 1500\,\mathrm{m^3}$ at $T_s = 40\,^{\circ}\mathrm{C}$, $Q_{max} = 23.3\,\mathrm{m^3/s}$ at full speed, α_c the compressor's available capacity, and the speed follows its command with 0.1 min. The propylene to the splitter is the compressor's molar flow times y_P, and the bed pressure is $P = P_s + \Delta P$ with $\Delta P = 0.05\,\mathrm{bar}$. The surge holds 25.9 kmol at 0.45 bar a, about 1.4 min of the bank's 310.7 mol/s, and at design the compressor runs at 74.8 % speed. Because its capacity is a volume at suction, a deeper vacuum needs more speed for the same moles.

5.3.5 The Sequencer and the Measures of Performance

The cycle is cut into eight slots of $t_{cyc}/8$, and every phase lasts a whole number of them, three on line, one in purge, three in regeneration and one evacuated. Each bed in service holds a position in a ring, its position sets its phase, and at every slot boundary all beds move one position on, so the hottest bed comes on line and the one on line longest goes to purge. A bed taken out of service leaves an empty position. The setting *With a bed out* leaves the gap, so that for three slots of each cycle only two beds are on line, or removes the hole when it reaches a regeneration position (seven beds, three on line, regeneration of two slots) or an on-line position (seven beds, two on line). The sequencer belongs to the plant, moves on at the end of each integration step, and switches every valve at a slot boundary in one step.

The bank's conversion and selectivity are taken over each slot, $X = 1 - \int\sum F_{\mathrm{C_3H_8},out}\,dt \,/\, \int\sum F_{in}\,dt$ and $S = \int\sum F_{\mathrm{C_3H_6},out}\,dt \,/\, \int\sum (F_{in} - F_{\mathrm{C_3H_8},out})\,dt$, which equal one bed's averages over its whole on-line period. The product ripple is the peak-to-peak propylene to the splitter within a slot as a percentage of its mean. Catalyst life is a projection from a teaching wear law, as shown in Equation (5.6).

$$\mathrm{life} = \frac{L_{ref}}{\exp\left[\frac{E_w}{R}\left(\frac{1}{T_w} - \frac{1}{T_i}\right)\right]} \tag{5.6}$$

where the bar is the mean over every cell in service and the last cycle, $E_w = 340\,\mathrm{kJ/mol}$, $T_w = 923.15\,\mathrm{K}$ and $L_{ref} = 0.54\,\mathrm{years}$, so the life halves about every 14 K near 650 °C. The projection does not lower the activity within a run.

Table 5.2 Parameters. A: Adjustable in the Parameters Panel. V: Fixed and Stated in the App. H: Fixed and Not Displayed

Symbol	Meaning	Value	
—	bed mass, heat capacity	33 t, 1.1 kJ/(kg K)	V
ΔP	beds above the suction	0.05 bar	A
T_{feed}	charge heater outlet	600 °C	A
T_{lim}	bed temperature limit	680 °C	A
$k_{1,0}$, E_1	dehydrogenation	0.012 mol/(s kg bar), 120 kJ/mol	A, V
n	order in total pressure	0.3	V
ΔH_1, ΔS_1	dehydrogenation	129.5 kJ/mol, 139.2 J/(mol K)	V
$k_{2,0}$, E_2	cracking	8×10^{-4} mol/(s kg bar), 200 kJ/mol	A, V
$k_{3,0}$, E_3	coking	1.4×10^{-4} mol/(s kg bar), 150 kJ/mol	A, H
c_D	coke for activity $1/e$	1.6 kg/t	A
k_b, E_b	coke burn rate	1 /min at 650 °C, 100 kJ/mol	V, H
ΔH_b	heat of coke burn	32.8 MJ/kg	V
c_p	gas per mol of propane fed	160 J/(mol K)	H
$\dot{m}_{air}$	air per regenerating bed	40 kg/s	A
c_{air}	air heat capacity	1150 J/(kg K)	H
T_{ah}	air heater outlet	600 °C	A
LHV, η	fuel heating value, share	47 MJ/kg, 0.98	V, H
$\dot{m}_{f,max}$, τ_v	fuel valve, lag	0.36 kg/s, 0.05 min	H
τ_{air}	air mixing lag	0.1 min	V
$\dot{m}_s$, T_{steam}	purge steam	3 kg/s, 450 °C	V
—	steam heat capacity	2300 J/(kg K)	H
V_s, T_s	surge volume	1500 m^3, 40 °C	A, V
Q_{max}	compressor at suction	23.3 m^3/s	A
τ_c	compressor speed lag	0.1 min	H
E_w, L_{ref}	wear law	340 kJ/mol, 0.54 y	V

5.3.6 Instruments, Initialization and Numerics

A gas chromatograph (GC) analyzes a composite sample of each slot's product and reports X_{GC} 1 min after the slot ends, with 0.1 point of noise, holding between reports. The regeneration air thermocouple carries a 0.1 min lag and 0.3 K of noise, the suction transmitter 0.001 bar of noise, the bed thermocouples read every cell, and the sequencer's state and the feed are known exactly. Every run starts on the cyclic steady state of its own settings. One bed is integrated through the ring with the engine's own step until it repeats to 10^{-7} K, and each bed is placed at the state of its ring position, so a bank started this way is periodic to the integrator. From a uniform 620 °C a bed repeats to 0.1 K within four cycles. Integration is fourth-order Runge–Kutta with a 1 s step, which puts every slot boundary on the grid, the loops execute every 5 s and the historian logs every 10 s. At the top playback speed of 120 min per second the twelve-hour challenge needs 6 s, provided the browser sustains it. The exported file (Section 1.6.4) sets each GC report beside the true slot conversion, so an explanation can be checked against the simulator's own record. Halving the step moves the cyclic start by less than 10^{-6} K and the design conversion by 0.012 points. Over every case the bank's energy closes on its ledgers to 4.0×10^{-12} GJ and its carbon to 2.2×10^{-8} kmol. Table 5.2 lists the parameters.

5.3.7 Assumptions

The beds are adiabatic and lose no heat, and six cells stand for plug flow, which gives 57.8 percent conversion at design against 59.5 with twelve cells and 60.4 with twenty-four, while the start and end temperatures move less than 1.3 K. The gas's heat capacity flow is that of the propane fed, and the reactions' change of it is left out. Coke is carbon from propylene, burned first order with the oxygen in excess, the air bringing nearly 190 times the oxygen the coke needs. A real cycle also reduces the catalyst with hydrogen before it goes back on line, and its valve sequence takes tens of seconds, while here a bed is purged with steam and evacuated with no flow and its valves switch at once. The GC reports a slot 1 min after it ends, faster than a real analyzer and sample system. The rate order in

total pressure, the cracking and coking rates, the 88 percent selectivity at design and the wear law are teaching values. The design conversion of 57.8 percent and the 24 min cycle lie inside the 48 to 65 percent and 15 to 30 min tabulated for the Catofin process [1], while the selectivity has not been compared with a source. The bed temperature limit of 680 °C and the vacuum alarm are teaching values. Real limits come from the catalyst supplier and the licensor, and the trips that guard them from the plant's safety requirements under IEC 61511.

5.4 The Control Layers as Implemented

The sequencer owns the valves, and no controller moves a bed between phases. One Manual and Auto switch governs the whole bank. In Manual the fuel valve and the compressor speed are the operator's stations and every controller tracks what is applied, so a return to Auto, or a change of architecture, is bumpless. Each execution builds the measurement packet, runs the cycle-to-cycle layer when a GC report has arrived, then the two loops.

5.4.1 Two Regulatory Loops

Both loops use the framework's velocity-form PID law of Section 1.6.2, Equations (1.14) to (1.17), with the derivative set to zero and a direction $a = \pm 1$ on the error. Only the emitted output is clamped, and the integral is corrected toward it with a tracking time of τ_I. The TIC ($a = +1$) holds T_{air} by the fuel valve at 0.5 % per K and 0.3 min within 0 to 100 %. The PIC ($a = -1$) holds P_s by the compressor speed at 300 % per bar and 2 min. Both execute every 5 s. Either loop becomes proportional only when its integral time is set to the slider's maximum, 10 min for the TIC and 60 min for the PIC.

5.4.2 The Cycle-to-Cycle Layer

Above the TIC the regeneration setpoint $T_{reg,SP}$ comes from one of three architectures. *Fixed* holds the operator's setpoint, 660 °C as shipped. The *cycle-to-cycle PI* is the same routine called only when a GC report arrives, once a slot, with Δt the time since the last report, $a = +1$, 3 K per point of conversion and 15 min, and its output clamped to 600–T_{lim}. The integral used at report k was accumulated through report $k-1$, so the integral part of a move reaches the air one report after the error that caused it. The air it sets reaches the first bed two slots later, so the loop meets a delay of a cycle, not a time constant. Equation (1.12) gives the pure-delay form of such a loop. A correction computed from results that arrive while later runs are already under way is also the problem of the vial filler of Volume 8, Chapter 9, and the simulated moving bed of Chapter 7 is this volume's other unit whose quality is known only once a cycle.

The *MPC*, the predictive option of Section 1.6.3 written slot by slot, plans $T_{reg,SP}$ at every report over $N = 16$ slots in four move blocks of 1, 1, 1 and 13 slots, on a finite step response model of the slot conversion, as shown in Equation (5.7).

$$
\begin{aligned}
& y_k = b + \textstyle\sum_{j=1}^{16} h_j (u_{k-j} - u_0) + \sum_{j=1}^{16} g_j (w_{k-j} - w_0) \\
& \min \ \textstyle\sum_k w_X (y_k - X_{SP})^2 + w_{\Delta u} (\Delta u_k / 5\,\mathrm{K})^2 \\
& \qquad + 0.06\, \ell_{life}\, u_k \\
& \text{s.t. } 600 \le u_k \le T_{lim}, \qquad u_k + \delta \le T_{lim} - m + \sigma
\end{aligned}
\tag{5.7}
$$

where u is $T_{reg,SP}$, w is the measured feed per bed, held over the horizon, h_j and g_j are the slot-by-slot increments of step responses identified on the model with +5 K and +0.5 t/h steps, one set for each schedule (eight beds, seven with three on line, seven with two), $h_1 = h_2 = 0$ because the air reaches the bank two slots late, b is the last report less the model's value of it, δ is the last measured hottest regenerating cell less the setpoint then in force, $w_X = w_{\Delta u} = \ell_{life} = 1$, $m = 5\,\mathrm{K}$, and the slack σ is priced $50\sigma^2 + 5\sigma$. At eight beds the identified gains are 0.288 points per K and -2.66 points per t/h. The linear cost on u stands for catalyst wear. Each plan is a quadratic program, solved in the framework's linear MPC module with the Goldfarb–Idnani algorithm, a dual active-set method [17].

When a plan comes back infeasible, not finite or outside its bounds, the tracked PI takes the setpoint without a bump.

5.4.3 Alarms and Three Models

BED TEMPERATURE LIMIT flags any cell above T_{lim}, VACUUM LOST flags a suction more than 0.02 bar above its setpoint for 2 min, and SCHEDULE GAP and BED OUT OF SERVICE report the sequencer. None of them trips the unit. Three models are kept apart. The plant is the model of Section 5.3, a reader of the trends can identify a second one, such as the 0.29 points per K and the two-slot delay of case 3, and the controllers' own model is the MPC's step response, which knows the schedule and the feed but neither the coke on any cell nor the activity.

5.4.4 Design Choices

Eight decisions taken while the app was built either depart from its plan or settle a value the plan left open. First, the dehydrogenation is of order 0.3 in total pressure, a teaching value. From 0.45 to 0.35 bar a the design cycle's conversion moves $+0.35$, $+0.07$ and -0.67 points at orders 0.3, 0.5 and 1, so a law of first order in the partial pressures would make the vacuum lower the conversion, the reverse of why the process runs under vacuum. Second, case 8 therefore teaches that the vacuum buys selectivity, because the stored heat caps the conversion. The selectivity gain holds at every order tried, 2.70, 2.46 and 1.80 points, while the conversion's sign does not. Third, the phases are fixed whole slots, because a rotation with a fixed number of beds on line needs them. Fourth, the sequencer and the GC are written for this app rather than taken from the shared stage sequencer and analyzer, since the shared sequencer starts and stops stages against a demand with run and rest times instead of rotating a ring, and the GC's composite must follow the slot length. Fifth, one air header and one TIC serve all regenerating beds. Sixth, each reactor is six one-cell beds of the shared stack so that each cell has its own coke. Seventh, catalyst life is a projection that does not feed back on the activity, and aging is a disturbance. Eighth, the challenge's conversion target is a band, every hourly average within ± 2 points of 57 % with a fuel limit, rather than a floor. A fixed 660 °C with two beds on line meets "at least 57 %" in every hour from the second, at 57.85 to 64.89 %, with no control and 0.647 t/h of fuel, so a floor would grade burning fuel rather than control.

5.5 Nine Cases on the Running Bank

The cases are listed in Table 5.3 and the principal measurements in Table 5.4. All values were measured headlessly on the app's engine, at its default seed and with the 1 s step the browser uses. Changes a case requests are applied once the run is under way, at 4 min where a case says after the first slot, and the challenge was repeated at seeds 0 to 3 and at half the step. Slot conversions are the simulator's composite over each slot.

> **✦ DID YOU KNOW? Losing a Bed Can Raise the Conversion**
>
> From the simulator: when bed D fails in case 5 and the seven beds close up with two on line, the Fixed regeneration gives 64.8 % conversion, against 57.8 % with the gap left, because each bed now regenerates longer. The price is propylene, 10.66 t/h against 12.73 t/h (Table 5.4).

5.5.1 Where the Heat Goes

Case 1 follows bed A from the instant it comes on line, with both stations held in Manual. Its mean temperature falls from 622.6 to 557.6 °C in nine minutes, the inlet cell from 658.1 to 540.1 °C and the outlet cell only from 581.1 to 574.9 °C, while its conversion falls from 69.9 to 44.7 % and averages 57.8 % at a selectivity of 88.0 %. The equilibrium at its mean temperature falls with it, from 71.5 to 46.0 %, so the bed runs close to equilibrium throughout and its conversion is limited by the heat left in it, as shown in Figure 5.2. The stored-heat arithmetic closes independently. The reactions

Table 5.3 The Nine Cases

#	Case and events	Teaching point
1	Manual, bed A just on line	where the heat goes
2	Auto, Fixed 660 °C, Schedule view	the ripple and its averaging
3	Regeneration air 660 to 670 °C	the handle, a cycle late
4	Cycle 24 to 32 min (16, 40)	the temperature swing
5	Bed D out at 12 min	three ways to run seven beds
6	Feed +15 % at 12 min, then PI or MPC	tuning for a delay
7	Catalyst aged at 12 min	coke as fuel, selectivity lost
8	Suction 0.45 to 0.35, then 0.30 bar a	vacuum buys selectivity
9	Challenge, 12 h: bed D out at 3 h, feed +15 % at 7 h	schedule and loop together

Table 5.4 Case Outcomes Measured on the Engine

Case and condition	Result
1, bed A on line	622.6 to 557.6 °C, X 69.9 to 44.7 %
2, PIC 300 / 40 / 1000 %/bar	ripple 5.1 / 2.6 / 9.0 %
3, air +10 K	slot X 57.80 to 60.73 %, S 86.9 %
4, cycle 16 / 24 / 32 / 40 min	X 56.1 / 57.8 / 58.4 / 56.8 %
4, same	life 5.0 / 3.0 / 2.1 / 1.6 y
5, gap, Fixed	57.8 %, 12.73 t/h, ripple 56.6 %
5, three on line, Fixed / PI	48.0 / 52.6 % (air at 680 °C)
5, two on line, Fixed / PI	64.8 / 57.0 %, 10.66 / 9.68 t/h
6, Fixed / PI / MPC	53.95 % / holds from 76 / 49 min
7, aged, Fixed	start 625.1 °C, X 56.3, S 85.8 %
8, 0.45 / 0.35 bar a	X 57.80 / 58.15, S 88.0 / 90.7 %
8, 0.30 bar a asked	held at 0.337, VACUUM LOST in 2 min

absorb 2.443 GJ over the period, the bed gives 2.359 GJ (36.3 MJ/K × 64.98 K) and the 600 °C feed the remaining 0.084 GJ. Over the whole cycle the air brings 2.182 GJ, the burn of the coke 0.332 GJ, the feed 0.084 GJ, and the purge steam takes away 0.155 GJ, which again sums to the reaction heat. The steam cools the bed to 553.4 °C, the air returns it to 622.6 °C with its hottest cell at 658.1 °C, and 0.4 kg of coke is left unburned for the next period.

> **TRY IT IN THE APP** **1 · One Bed Through Its Cycle (Manual)**
>
> Load the scenario and press Run, then watch bed A on the Bed view. In its nine minutes on line its mean temperature falls from 622.6 to 557.6 °C and its conversion from 69.9 to 44.7 %. Keep running through the purge and the regeneration and see it come back on line at 622.6 °C, ready to spend the same heat again.

Case 2 puts the eight beds together. The three beds on line are 0, 3 and 6 min into their periods, so the bank's instantaneous conversion saws between 54.1 and 62.1 % every slot and the propylene to the splitter between 14.12 and 14.87 t/h, a ripple of 5.1 %. Loosening the PIC to 40 % per bar and 15 min halves the ripple to 2.6 % while the suction stays within 0.445 to 0.455 bar a, because the surge itself averages the flow once the speed stops following it, and tightening it to 1000 % per bar and 1 min passes the sawtooth on at 9.0 %. The price of the loose loop is speed. A setpoint step from 0.45 to 0.40 bar a settles within 0.005 bar in 2.3 min as shipped and 8.3 min at 100 % per bar, and not within 54 min at 40 % per bar. This is averaging control applied to a gas inventory under vacuum.

5.5.2 The Regeneration as the Handle

Case 3 raises the regeneration air from 660 to 670 °C at 4 min and follows the response, as shown in Figure 5.3. The first slot to move is the one ending at 12 min, 57.92 against 57.80 %, made in part by a bed that finished its regeneration in the hotter air. A bed comes on line above 629 °C at 18 min, the slot conversion passes 60 % at 27 min and has made 78 % of its final change by 30 min and 98 % by 54 min. A second, smaller rise arrives one cycle later, when the beds start their next

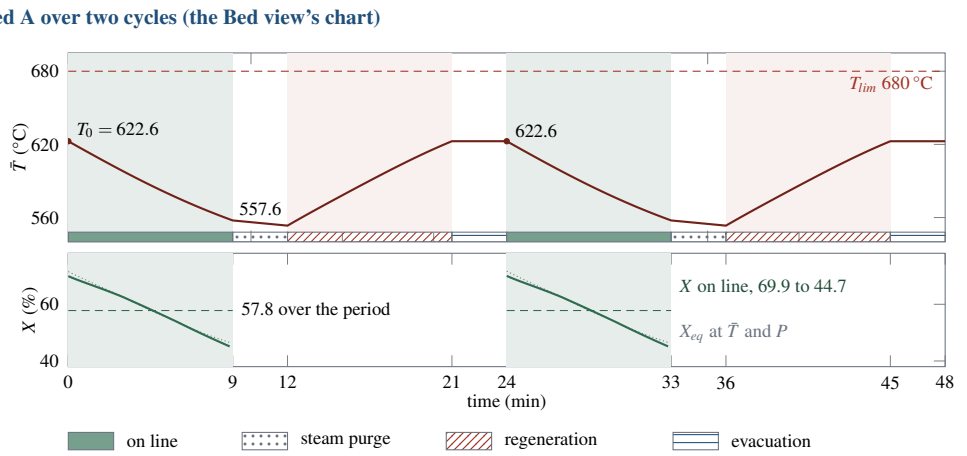

Figure 5.2 Bed A over two cycles, the Bed view's chart in case 1 with every station held in Manual. Top: its mean temperature (solid), falling on line from 622.6 °C and recovering in regeneration, against the 680 °C limit (dashed), over the phase strip. Bottom: its conversion while on line (solid) against the equilibrium at its mean temperature (dotted), with the average over each on-line period (dashed).

regeneration from a warmer end, and the bank settles at 60.73 % with every bed coming on line at 630.8 °C, 0.293 points per kelvin of air. The selectivity falls from 88.0 to 86.9 %, the fuel rises 0.106 t/h from 0.647 t/h, and the projected catalyst life falls from 3.0 to 2.0 years.

🕮 **Pause.** Try Exercise 5.1 before reading on.

> **TRY IT IN THE APP 3 · Regeneration Heat Is the Handle**
>
> Press Run and, after the first slot, raise the Regeneration air setpoint from 660 to 670 °C. Watch the Start of cycle badge: a bed comes on line above 629 °C at 18 min, and the slot conversion passes 60 % only at 27 min. The bank settles at 60.73 %, one to two cycles after the move.

Case 6 makes the delay a tuning problem. With the regeneration fixed, the first slot after the 15 % feed step reads 55.50 % and the bank settles at 53.95 %, the reactions now absorbing 2.633 GJ per bed and period against 2.443 GJ. Handed the setpoint at 18 min, the shipped PI raises the air step by step to 671 °C and holds the GC within 0.5 point of the target from 76 min, peaking at 57.06 %. The MPC holds from 49 min, peaking at 57.03 %. At its first report it moves the air to 668.7 °C against the PI's 661.9 °C, because its model knows both the delay of its own moves and the feed it has measured. Frozen so that it does not see the feed, it holds from 70 min, so the measured feed bought most of its lead. A tighter PI, 5 K per point and 12 min, holds from 52 min with a 0.44 point overshoot. Over a grid of gains from 2 to 8 K per point and integral times from 6 to 30 min, the most aggressive settings overshoot to as much as 58.6 % or never settle within the band, which is why the shipped tuning is slow.

5.5.3 Cycle Length, a Lost Bed, Aging and Vacuum

Lengthening the cycle in case 4 widens the swing at both ends. At 32 min a bed comes on line at 637.2 °C instead of 622.6 °C and its conversion runs from 76.4 down to 39.5 %, yet the average barely moves, 58.4 against 57.8 %, because the air puts in the same heat per minute. The selectivity falls from 88.0 to 87.8 %, the ripple rises from 4.6 to 8.4 % and the projected life falls from 3.0 to 2.1 years. At 16 min the bed starts at 608.0 °C and converts 56.1 % with 5.0 years of life, and at

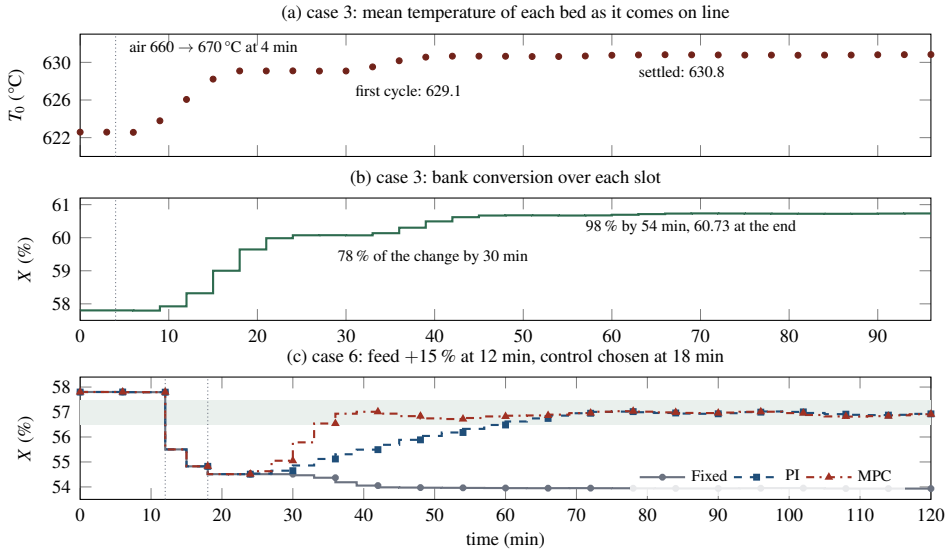

Figure 5.3 The regeneration acts a cycle late. (a) Case 3, the mean temperature of each bed as it comes on line after the air is raised 10 K at 4 min. (b) The bank's slot conversion in the same run, drawn over the slot it belongs to. (c) Case 6, the slot conversion after the feed per bed rises 15 % at 12 min, with the regeneration left Fixed (solid, dots) or handed at 18 min to the cycle-to-cycle PI (dashed, squares) or the MPC (dash-dotted, triangles). The shaded band is ±0.5 point of the 57 % target.

40 min the conversion itself falls to 56.8 % as the end of each period runs too cold. The model counts none of the costs of a short cycle, the valve strokes, purge steam and vacuum.

Case 5 loses bed D at 12 min. Leaving the gap keeps the conversion at 57.8 %, but the propylene falls from 14.57 to 12.73 t/h and ripples 56.6 % each cycle as the empty slot passes. Seven beds with three on line close up at 21 min and give a steady 12.39 t/h, but every bed now regenerates for two slots, starts at 596.9 °C, and the conversion falls to 48.0 %. Seven beds with two on line close up at 15 min, the conversion rises to 64.8 % on the longer regeneration, and the propylene falls to 10.66 t/h. The cycle-to-cycle PI holds 57.0 % with two on line at 636.6 °C air. With three on line it drives the air to the 680 °C limit and stops at 52.6 %. A sweep of the air temperature explains why. With three on line the conversion is 43.5, 45.7, 47.9, 50.3 and 52.6 % at 640, 650, 660, 670 and 680 °C air, so no regeneration under the limit reaches 57 % at 10 t/h per bed, and it would take a feed of 8.5 t/h (57.25 %). This infeasibility is deliberate, and the case states it.

The aged catalyst of case 7, at 75 % activity and 1.5 times the coke, lays down 88.6 instead of 80.9 kg of coke per cycle, whose burn releases 2.91 instead of 2.65 GJ, so the beds come on line hotter, at 625.1 °C. The conversion falls only from 57.8 to 56.3 %, because a slower catalyst draws less heat early, while the selectivity falls from 88.0 to 85.8 % and stays there. The PI restores 57.0 % with 662.5 °C air at a projected life of 2.6 instead of 2.9 years.

🕮 **Pause.** Try Exercise 5.3 before reading on.

In case 8 the suction setpoint is lowered to 0.35 bar a. The conversion rises only from 57.80 to 58.15 % while the selectivity rises from 88.0 to 90.7 % and the propylene from 14.57 to 15.11 t/h, as cracking and coking slow at the lower partial pressures, and the compressor needs 96.2 % of its speed. Asked for 0.30 bar a, it runs at full speed, holds 0.337 bar a, and VACUUM LOST appears 2 min after the request.

🕮 **Pause.** Try Exercise 5.2 before reading on.

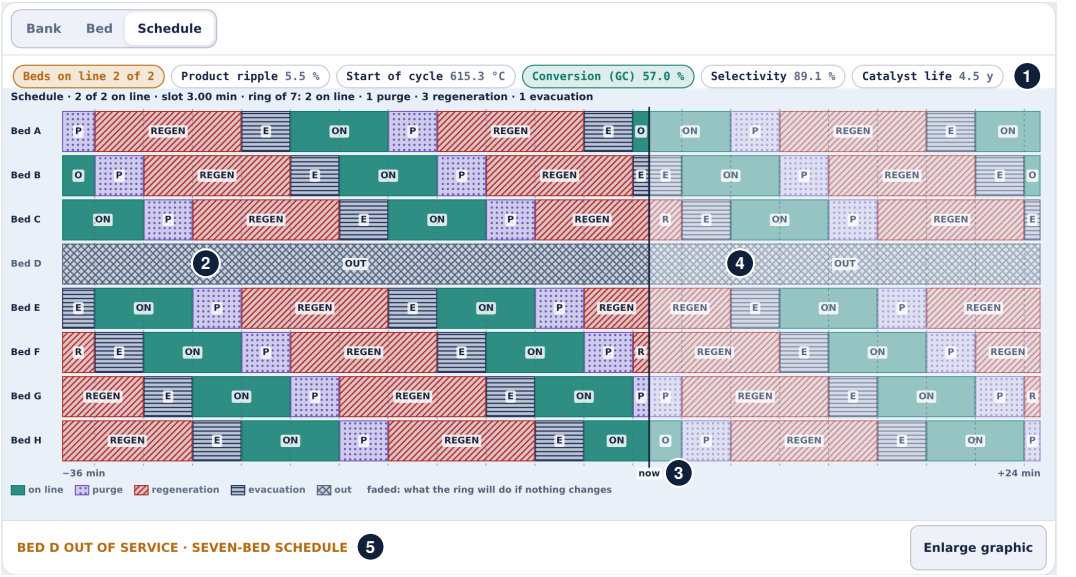

Figure 5.4 The challenge on its suggested settings, the cycle-to-cycle PI with seven beds two on line, at $t = 718$ min on the Schedule view, cropped to the plant. (1) The badges read 2 of 2 beds on line, a ripple of 5.5 %, a start-of-cycle temperature of 615.3 °C, a GC conversion of 57.0 %, a selectivity of 89.1 % and a projected life of 4.5 years. (2) Bed D, out of service since 3 h, while the seven-bed ring runs two on line, one in purge, three regenerating and one evacuated. (3) The present instant. (4) What the ring will do over the next 24 min, faded. (5) The status line that reports the bed out and the seven-bed schedule.

5.5.4 Twelve Hours With a Lost Bed and More Feed

The challenge runs 12 h. Bed D fails on line at 3 h and the feed per bed rises 15 % at 7 h. It is met when every clock hour's conversion from the second hour lies within ±2 points of 57 %, no cell exceeds 680 °C, the ripple stays at or under 10 % in every slot of the last two hours and the fuel averages at most 0.55 t/h, judged on the simulator's truth at every step. The configurations are compared in Table 5.5, and the suggested run is shown in Figure 5.4. As shipped, with Fixed regeneration and the gap left, the product ripples 58.6 % and the conversion drifts to 53.8 % after the feed rises, 3.19 points from the target at hour 11. The suggested settings pass, and the MPC on the same schedule passes too. Their margin is thin but steady. Over the default seed, seeds 0 to 3 and the halved step the worst hour lies 1.70 to 1.75 points from the target for the PI and 1.69 to 1.73 for the MPC, always hour 4, the close-up transient in which the bank's slot conversion briefly reaches 61.4 %. The MPC solves 239 plans with none rejected, and ties the PI at 1.72 against 1.73 points and 0.494 against 0.497 t/h. Its advantage is recovery speed after a measured feed change, not the graded outcome. Figure 5.5 sets the table's graded columns against their limits.

Three on line fails whatever the control, 8.14 points with the PI and 9.14 with the MPC, because the two-slot regeneration cannot reach the target under the limit. Even with the feed cut to 7.5 to 8.5 t/h per bed it misses by 2.89 to 3.08 points at hour 4. The PI with the gap holds the conversion within 1.10 points but fails the ripple and the fuel. The Fixed regeneration with two on line overshoots to 64.89 %. The tighter PI that settled case 6 sooner passes with 0.83 points, yet it overshoots case 6 by 0.44 point, and the shipped tuning was kept. The price of passing is throughput, 11.49 t/h of propylene on average against 13.66 t/h as shipped.

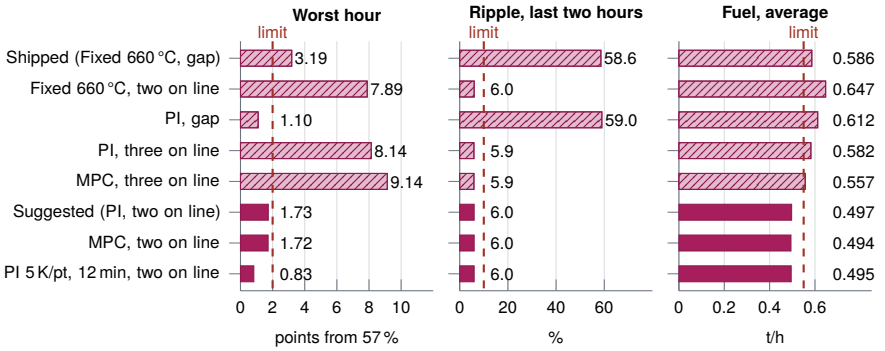

Figure 5.5 The challenge configurations of Table 5.5 against three of the four limits that grade them: the worst hourly conversion, the ripple in the last two hours and the average fuel. Solid bars pass the challenge, hatched bars fail it, and the dashed line in each panel is the limit. Notice that the gap fails on the ripple, three on line and Fixed regeneration fail on the conversion, and every failing configuration also burns more fuel than the limit allows. The fourth limit, no cell above 680 °C, holds in every row.

Table 5.5 Challenge Configurations Scored at 12 h

Configuration	Worst (pt)	Hottest (°C)	Ripple (%)	Fuel (t/h)	Verdict
shipped (Fixed 660 °C, gap)	3.19	659.6	58.6	0.586	fails
Fixed 660 °C, two on line	7.89	658.3	6.0	0.647	fails
PI, gap	1.10	674.0	59.0	0.612	fails
PI, three on line	8.14	671.3	5.9	0.582	fails
MPC, three on line	9.14	666.9	5.9	0.557	fails
suggested (PI, two on line)	1.73	658.1	6.0	0.497	passes
MPC, two on line	1.72	658.1	6.0	0.494	passes
PI 5 K/pt, 12 min, two on line	0.83	658.1	6.0	0.495	passes

5.6 Technology in Practice

Figure 5.6 draws the loops of Section 5.4 with the hardware of a cyclic bank in place of the schematic's symbols.

What Measures It

The conversion that the cycle-to-cycle layer corrects comes from a process gas chromatograph that analyzes each slot's product. Analyzers of this kind suffer from large measurement delays and high investment and maintenance costs [18], and the report a minute after each slot in this model is faster than a plant's. The regulatory loops read a thermocouple on the regeneration air header and an absolute pressure transmitter at the compressor suction. The model gives the thermocouple a 0.1 min lag with 0.3 K of noise and the transmitter 0.001 bar of noise, so neither a plugged impulse line nor a heavily damped signal appears in it.

> **TECHNOLOGY** **Sensor: An Absolute Pressure Transmitter on the Suction**
>
> An absolute pressure transmitter is a differential pressure instrument whose low side is sealed to a vacuum chamber, so any pressure above a perfect vacuum reads as a positive difference. It has no vent on the low side and a single impulse line to the process [19]. The impulse line can plug with debris, and where plugging is chronic the line is filled with a harmless fluid or purged with a steady flow. A transmitter damped too heavily draws a smooth trend that looks like tight control yet understates a real change [19].

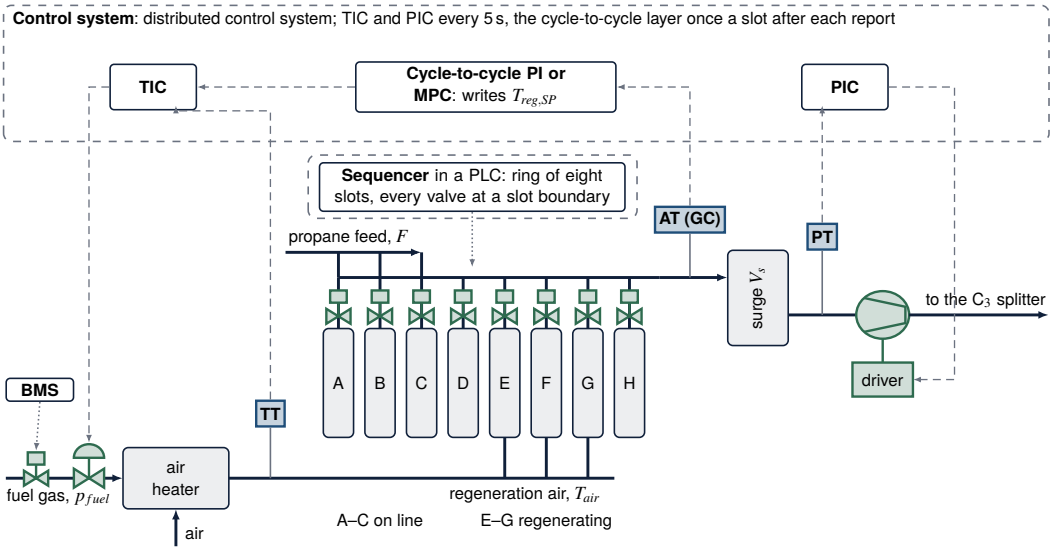

Figure 5.6 The bank's loops with their hardware. Blue boxes are sensing devices: TT on the regeneration air header, PT at the compressor suction and AT, the chromatograph on the product header. Green symbols are final elements: the fuel gas valve of the air heater, the compressor's driver and the on/off valves of the beds, one set drawn per bed. Dashed arrows are signals, dotted ones belong to the separate burner trip system and the sequencer. Notice that the cycle-to-cycle layer acts only through the air for beds that are not on line, and only once a slot. Purge steam, the vacuum header and the feed heater are left out. Not to scale, and not a piping diagram.

What Moves It

Three kinds of final element act on the bank. The fuel gas valve of the air heater is a throttling control valve, which the model follows with a 0.05 min lag and a flow that rises with the square root of the fuel gas pressure. The compressor's speed sets how deep the vacuum is, and its capacity, not the loop, ends case 8 at 0.337 bar a. The on/off valves of every bed belong to the sequencer, not to a controller. A real valve sequence takes tens of seconds, where the model switches every valve in one step.

> **TECHNOLOGY** **Final Element: On/Off Valves That Switch the Beds**
>
> A bed changes service through on/off valves, the fluid equivalent of an electrical switch, used to route a process fluid to different places. Ball, plug, butterfly, gate and globe valves serve this duty, and a large one opens to a nearly unimpeded path. A pair of pneumatic pistons on a rack and pinion can turn a butterfly valve, and a box of limit switches on its stem tells the control room where the valve stands. Tight shut-off is hard to keep in hot service, because bubble-tight seats usually need soft materials that limit the temperature range [19]. In the model every valve switches at once at a slot boundary, and a bed whose valves fail is taken out of service, as in case 5.

Where the Controller Runs

A distributed control system holds loops like these at a supervisory level that sends set points to field controllers, with loops nested or cascading and cycle times from milliseconds to minutes [20]. Programmable logic controllers (PLCs) serve such a system as local controllers within the supervisory scheme, programmed for logic, timing, counting and PID control [20]. The sequencer of this chapter is a state machine of that kind, whose output depends on its ring position as much as on its inputs, and IEC 61131-3 names the Sequential Function Chart among its five languages for writing one [19]. In the app the TIC and the PIC run every 5 s as velocity-form PID blocks. The cycle-to-cycle PI is

the same block called only when a report arrives, once a 3 min slot, and the MPC sits above both as a supervisory layer that writes $T_{reg,SP}$.

> **IN PRACTICE** **The Regeneration Air Heater Has Its Own Guard**
>
> The fuel valve that the TIC moves is a modulating valve. The burner it feeds is also watched by a burner management system, which purges the firebox with air for several minutes before ignition, lights the burner, and then watches for low fuel pressure, high fuel pressure and loss of flame. On any of them it halts the fuel with double block and bleed valves. Staff impatient to restart are often tempted to bypass its purge time. Such safety instrumented systems in the process industry are the subject of ANSI/ISA-84.00.01 [19]. The chapter's alarms trip nothing, and its fuel gas pressure disturbance is one of the inputs such a system watches.

What Is Modernizing It

Model predictive control has a solid foundation in refining and petrochemicals [15], and here it pays in recovery after a measured feed change, not in the graded outcome. Soft sensors estimate product quality from other process measurements where an online analyzer is not available, and they are also used to validate online analyzers [18], which would let a bank correct between chromatograph reports. Microprocessor burner controls time the purge more accurately, make it harder to override and connect to a network for remote monitoring [19]. On the catalyst side, a constant activity over the life of a chromia charge can be held by raising the reaction temperature gradually with time on stream [13], the slow drift that case 7 compresses into one step.

5.7 Limitations

No operating unit supplied the sizes, rates, heats or limits, all of which are teaching values. The kinetics have the standard forms with a rate order in total pressure chosen so that the vacuum helps, the equilibrium sits about one point above the NASA Glenn data, and the selectivity of 88 percent is not checked against a source. Six cells per bed leave the conversion 1.7 points under twelve, the heat capacity flow of the gas is fixed, and the reduction step and the valve sequence of a real cycle are left out. One air header serves every regenerating bed, so the regeneration of individual beds cannot be trimmed, and catalyst life is projected rather than fed back on the activity.

The control layers have their own simplifications. The GC reports a minute after each slot and its noise is 0.1 point, faster and cleaner than a real analyzer, so the cycle-to-cycle loop's information is better than a plant's. The MPC's step responses were identified on the model itself, so its model matches the plant's gain exactly at design, and it knows neither the coke nor the activity. The challenge's margin of 0.25 to 0.31 points is thin, although it held at every seed tried and at half the step. The MPC ties the PI on the graded outcome and wins only on recovery speed after a measured feed change.

> **OPEN CHALLENGE** **A Catalyst That Survives Its Cycles**
>
> Chromia catalysts lose activity over many cycles in ways that coke alone does not explain. Coke formation grows with the number of cycles, part of the chromium is reduced for good, and a spinel of chromium and alumina formed during cycling is linked to the loss. A 1999 review concluded that more characterization work was needed to understand this deactivation fully [13]. A 2023 review still names poor cycling performance and heavy carbon deposits as limits on industrial dehydrogenation, and the environmental risk of chromium as a limit on chromia catalysts [1]. Progress would be a measured law of activity against cycles, the piece this chapter's model replaces with a projected catalyst life.

5.8 Conclusion

A cyclic dehydrogenation bed spends heat it stored in the previous regeneration, and this simulator puts eight such beds, their sequencer, the regeneration air, the vacuum and the analyzer on one

clock. A bed on line cools 65.0 K and gives 2.36 GJ of the 2.44 GJ its reactions absorb, so 10 K more regeneration air raises the bank's conversion 2.93 points, first in the slot that ends 8 min after the move and fully one to two cycles later. A deeper vacuum buys 2.7 points of selectivity and 0.35 of conversion, because the stored heat, not the equilibrium, caps what a bed converts. With a bed lost, seven beds with three on line cannot reach the target under the bed limit (52.6 %), seven with two on line can, and a cycle-to-cycle PI or an MPC then holds every hour of the challenge within 1.73 and 1.72 points. The kinetics, the selectivity and the wear law are teaching values, and the analyzer is faster than a real one. A fuller model would add per-bed regeneration trim, the reduction step, a slower analyzer and catalyst aging that feeds back on the activity. Even without them, the bank shows control that acts on the next cycle rather than the present one.

KEY POINTS

- Nothing heats a dehydrogenation bed while it makes propylene: it spends heat stored in its last regeneration.
- 10 K more regeneration air raises the bank's conversion 2.93 points, but fully only one to two cycles later.
- A deeper vacuum buys 2.7 points of selectivity and 0.35 of conversion, because the stored heat caps what a bed converts.
- With a bed lost, seven beds with three on line cannot reach the target under the bed limit, while seven with two on line can, under a cycle-to-cycle PI or an MPC.

Where else this applies. Correcting the next run from the measured result of the last is also the work of Volume 8, Chapter 9, Volume 8, Chapter 10 and Chapter 7.

References

[1] C. Zuo and Q. Su, "Research progress on propylene preparation by propane dehydrogenation," *Molecules*, vol. 28, no. 8, p. 3594, 2023.

[2] Y. Wang, F. Gao, and F. J. Doyle, "Survey on iterative learning control, repetitive control, and run-to-run control," *Journal of Process Control*, vol. 19, no. 10, pp. 1589–1600, 2009.

[3] J. H. Lee and K. S. Lee, "Iterative learning control applied to batch processes: An overview," *Control Engineering Practice*, vol. 15, no. 10, pp. 1306–1318, 2007.

[4] Y. S. Matros and G. A. Bunimovich, "Reverse-flow operation in fixed bed catalytic reactors," *Catalysis Reviews*, vol. 38, no. 1, pp. 1–68, 1996.

[5] P. Freneau, G. Kelso, and A. W. Hoge, "Application of steam-jet ejectors in the Houdry fixed-bed catalytic cracking process," *Journal of Fluids Engineering*, vol. 69, no. 1, pp. 69–79, 1947.

[6] Z. Nawaz, "Dynamic modeling of CATOFIN® fixed-bed iso-butane dehydrogenation reactor for operational optimization," *International Journal of Chemical Reactor Engineering*, vol. 14, no. 1, pp. 491–515, 2016.

[7] J. J. H. B. Sattler, J. Ruiz-Martinez, E. Santillan-Jimenez, and B. M. Weckhuysen, "Catalytic dehydrogenation of light alkanes on metals and metal oxides," *Chemical Reviews*, vol. 114, no. 20, pp. 10 613–10 653, 2014.

[8] M. M. Bhasin, J. H. McCain, B. V. Vora, T. Imai, and P. R. Pujadó, "Dehydrogenation and oxydehydrogenation of paraffins to olefins," *Applied Catalysis A: General*, vol. 221, no. 1-2, pp. 397–419, 2001.

[9] Z. Nawaz, "Light alkane dehydrogenation to light olefin technologies: a comprehensive review," *Reviews in Chemical Engineering*, vol. 31, no. 5, 2015.

[10] M. P. Lobera, C. Téllez, J. Herguido, and M. Menéndez, "Transient kinetic modelling of propane dehydrogenation over a Pt–Sn–K/Al_2O_3 catalyst," *Applied Catalysis A: General*, vol. 349, no. 1-2, pp. 156–164, 2008.

[11] A. G. Oblad, "The contributions of Eugene J. Houdry to the development of catalytic cracking," in *Heterogeneous Catalysis*, ser. ACS Symposium Series. American Chemical Society, 1983, pp. 61–75.

[12] American Chemical Society, "The Houdry process for the catalytic conversion of crude petroleum to high-octane gasoline," National Historic Chemical Landmark booklet, Office of Public Outreach, Washington, DC, 1996, acs.org/education/whatischemistry/landmarks/houdry.html, read November 2026.

[13] B. M. Weckhuysen and R. A. Schoonheydt, "Alkane dehydrogenation over supported chromium oxide catalysts," *Catalysis Today*, vol. 51, no. 2, pp. 223–232, 1999.

[14] C. H. Bartholomew, “Mechanisms of catalyst deactivation,” *Applied Catalysis A: General*, vol. 212, no. 1-2, pp. 17–60, 2001.

[15] 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.

[16] B. J. McBride, M. J. Zehe, and S. Gordon, “NASA Glenn coefficients for calculating thermodynamic properties of individual species,” NASA Glenn Research Center, Technical Publication NASA/TP-2002-211556, Sep. 2002.

[17] 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.

[18] 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.

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

[20] 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.

5.9 Exercises

5.1 Concept check · 2 min

In case 3 the regeneration air setpoint is raised from 660 to 670 °C at 4 min. The TIC holds the new air temperature within a minute, yet the first slot conversion to move is the one ending at 12 min, and the bank reaches its final conversion only one to two cycles later. Why?

(a) The air heats only beds in regeneration, which come on line slots later, and the beds then start their next regeneration warmer

(b) The TIC is tuned slowly, so the air temperature itself takes most of a cycle to reach the new setpoint

(c) The gas chromatograph reports each slot late, so the conversion has moved long before the report shows it

(d) The surge volume holds the product for a cycle, so the splitter sees the higher conversion a cycle late

5.2 Concept check · 2 min

In case 8 the suction setpoint is lowered from 0.45 to 0.35 bar a. The equilibrium conversion of propane at 600 °C rises by about four points, but the bank’s conversion rises only from 57.80 to 58.15 %. What does the vacuum buy, and why?

(a) Nothing lasting, because the compressor cannot hold 0.35 bar a and the suction drifts back

(b) Selectivity, because the stored heat caps the conversion while cracking and coking slow at the lower partial pressures

(c) Conversion only after a cycle, because the beds must first cool to the new equilibrium temperature

(d) Throughput only, because the lower pressure lets the compressor pass more volume at the same speed

5.3 Concept check · 2 min

In case 7 the catalyst drops to 75 % of its fresh activity and makes 1.5 times the coke, with the regeneration fixed at 660 °C. Why does the bank’s conversion fall only from 57.8 to 56.3 %?

(a) The cycle-to-cycle PI raises the air to make up the lost activity

(b) The coke covers only the outlet cells, so the inlet cells convert as before

(c) A slower catalyst draws less heat early, and the extra coke burned in regeneration reheats the beds

(d) The equilibrium limits the conversion, so the activity matters only far from equilibrium

5.4 Theory and calculation · 7 min

Each bed holds 33 t of catalyst and inert at 1.1 kJ/(kg K) in six equal cells, and receives 10 t/h of propane (44.10 g/mol) at 600 °C. The gas carries 160 J/(mol K) per mole of propane fed, and in regeneration a bed receives 40 kg/s of air at 1150 J/(kg K).

(a) Compute the bed's heat capacity. Over its nine minutes on line in case 1 its mean temperature falls 64.98 K while the reactions absorb 2.443 GJ. How much of that heat does the bed give, and how much does the feed supply? *Checkpoint:* 36.3 MJ/K, 2.359 GJ and 0.084 GJ

(b) Compute one cell's heat capacity C, the molar feed F_{in}, the heat capacity flow $\dot{F}c_p$ and the time the gas alone would need to change a cell's temperature, $C/(\dot{F}c_p)$. Compare it with the time on line. *Checkpoint:* 6.05 MJ/K, 62.99 mol/s, 10.08 kW/K, 10.0 min

(c) Compute the air's time constant on a cell and compare it with the nine-minute regeneration. What does each comparison say about where a bed's temperature comes from on line and off line? *Checkpoint:* 2.19 min

5.5 Theory and calculation · 8 min

The app's dehydrogenation has $\Delta H_1 = 129.5\,\mathrm{kJ/mol}$ and $\Delta S_1 = 139.2\,\mathrm{J/(mol\,K)}$, with $K_p = \exp(\Delta S_1/R - \Delta H_1/(RT))$ in bar and $R = 8.314\,\mathrm{J/(mol\,K)}$. The equilibrium conversion of pure propane is $X_{eq} = \sqrt{K_p/(P+K_p)}$, and the bed pressure is $P = P_s + 0.05\,\mathrm{bar}$.

(a) Compute K_p at 600 °C. *Checkpoint:* 0.334 bar

(b) Compute X_{eq} at 600 °C for suction setpoints of 0.45 and 0.35 bar a. Give one decimal. *Checkpoint:* 63.3 and 67.5 %

(c) The NASA Glenn polynomials give $K_p = 0.315\,\mathrm{bar}$ at 600 °C. Compute X_{eq} at 0.50 bar with it. *Checkpoint:* 62.2 %

(d) In case 8 the same change of suction moves the bank's conversion from 57.80 to 58.15 %. What fraction of the rise in X_{eq} is that, and what caps the rest? *Checkpoint:* about 0.08

5.6 Theory and calculation · 7 min

Model the bank's cycle-to-cycle loop as Equation (1.12): the GC reports once a slot (3 min), the conversion answers the regeneration setpoint with $g = 0.293\,\mathrm{points/K}$ after $d = 2$ reports, and the integral law moves the setpoint by K kelvin per point of error at each report. Ignore the proportional part of the PI and the tail of the real response.

(a) Compute the integral gain K for the double root and at the edge of stability. *Checkpoint:* 0.853 K and 3.41 K per point

(b) The PI's integral part moves the setpoint by $K_c\,\Delta t/\tau_I$ per point of error at each report. Compute it, and gK, for the shipped 3 K per point and 15 min, the tighter 5 K per point and 12 min, and 8 K per point with 6 min. *Checkpoint:* gK = 0.176, 0.366 and 1.17

(c) Say what the model predicts for each tuning and compare with the chapter's case 6. *Checkpoint:* no overshoot, overshoot, unstable

5.7 App challenge · 8 min

Load case 3, Regeneration heat is the handle, at the default seed. Press Run at speed 1 (1 min of plant per wall second) and pause between 4 and 6 min. Set the Regeneration air setpoint on the TIC card to

670.0 °C, then run at speed 5 to 96 min, pause and export the CSV. A run longer than an hour exports its time column as Time (h), in hours, so convert each time below from minutes.

(a) Read the Start of cycle badge and the Selectivity, Fuel and Catalyst life tiles. In the CSV, read the column `truth_x_slot_pct` in the 96.00 min row. Use tolerances of 0.1 K, 0.02 points and 0.01 t/h. *Checkpoint:* 630.8 °C, 86.9 %, 0.75 to 0.76 t/h, 2.0 y, 60.73 %

(b) Against the design values (start 622.6 °C, conversion 57.80 %), compute the conversion gained per kelvin of air and per kelvin of start-of-cycle temperature. *Checkpoint:* 0.293 and 0.357 points per K

(c) Use Equation (5.4) at steady state with three beds regenerating to compute the fuel at 660 and at 670 °C, and compare with the tile. *Checkpoint:* 0.647 and 0.755 t/h

5.8 App challenge · 8 min

Load case 5, A bed out of service, at the default seed. Before pressing Run, choose Cycle-to-cycle PI under Regeneration control and Seven beds, two on line under With a bed out. Run to 120 min at speed 5 and export the CSV. Then choose Seven beds, three on line, press Reset (it keeps your settings), run to 120 min again and export. A run longer than an hour exports its time column as Time (h), in hours, so convert each time below from minutes.

(a) For each schedule read `treg_C`, `treg_saturated` and `truth_x_slot_pct` in the 120.00 min row, and average `prod_c3h6_tph` over the rows from 99 to 120 min (one seven-slot cycle). Use tolerances of 1 K, 0.1 points and 0.05 t/h. *Checkpoint:* two: 636 to 637 °C, 0, 57.0 %, 9.68 t/h; three: 680.0 °C, 1, 52.6 %, 13.34 t/h

(b) Why does the PI hold the target with two on line and not with three? *Checkpoint:* three on line regenerates for two slots

(c) The chapter's sweep shows what feed per bed three on line would need. Which schedule gives more propylene at the target? *Checkpoint:* two on line, unless the feed drops to 8.5 t/h

5.9 App challenge · 8 min

Load case 9, the Challenge (twelve hours, a bed out and more feed), at the default seed. Run it as shipped at speed 120 (120 min of plant per wall second, about 6 s) to 12 h and read the So far tile. Then choose Cycle-to-cycle PI under Regeneration control and Seven beds, two on line under With a bed out, press Reset, and run again. Drop to speed 10 near 3.5 h and 4.5 h to read the tile, then finish the run.

(a) Give the four values on the So far tile at 12 h for each run. *Checkpoint:* shipped: ±3.2, 660, 58.6, 0.59; PI: ±1.7, 658, 6.0, 0.50

(b) What does the tile's worst hour read at 3.5 h and at 4.5 h in the second run, and which event made the worst hour? *Checkpoint:* ±0.1 then ±1.7 points; the ring closing up after bed D fails

(c) Which parts of the target does each run meet, and what does passing cost? *Checkpoint:* 11.49 t/h against 13.66 t/h

5.10 Application area · 5 min

A valve on bed D fails and the bed must stay out for a week. The splitter downstream needs a steady feed and the plan calls for 57 % conversion. Case 5 gives, against 14.57 t/h of propylene with eight beds: leave the gap, 57.8 % and 12.73 t/h rippling 56.6 %; three on line, 48.0 % and 12.39 t/h; two on line, 64.8 % and 10.66 t/h; two on line with the cycle-to-cycle PI, 57.0 % and 9.68 t/h. Compute the propylene each choice loses per day, and say which one an operator should run and what else could change.

5.11 Application area · 4 min

Case 4 runs the cycle at 16, 24, 32 and 40 min. The conversion is 56.1, 57.8, 58.4 and 56.8 %, and the projected catalyst life 5.0, 3.0, 2.1 and 1.6 years. Each bed runs one valve sequence a cycle, from reaction through purge and regeneration to evacuation. How many valve sequences does the bank make per day at 16, 24 and 32 min? Which costs of the cycle time does the model count, and which does it leave out?

5.12 Review the agent · 5 min

An AI assistant asked to make case 6 recover faster writes: “The bank’s conversion answers the regeneration setpoint within one slot, so treat it as a first-order lag of about 3 min. Lambda tuning then gives the cycle-to-cycle PI a gain of 8 K per point and an integral time of 6 min, which will hold the target within two slots of the switch instead of the shipped tuning’s hour.” Find the flaw and say how you would check it in the app.

CHAPTER 6

An Ethylene Oxide Reactor Simulator for Selectivity, Runaway and an Optimum Without a Setpoint

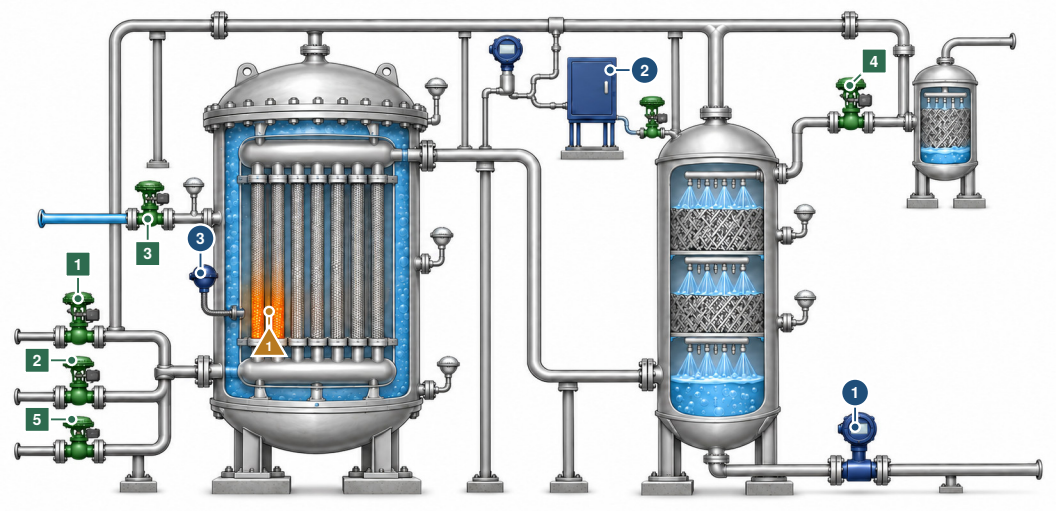

The controller moves 1 oxygen feed valve, which sets the production 2 ethylene feed valve 3 coolant valve, for the coolant temperature the oxygen loop sets 4 valve of the carbon dioxide wash 5 chloride dosing valve

The controller reads 1 production meter on the ethylene oxide made 2 cycle gas analyzer for oxygen, ethylene and carbon dioxide 3 hot spot thermocouple, which drives the override

What disturbs it 1 catalyst activity that changes ▲ a wash column that loses circulation

Notice the hot zone inside the tubes, which the controller reaches only through the oxygen feed and the coolant, and which the model gives to every tube alike. Illustration, not to scale.

IN PLAIN TERMS A Reactor Where the Waste Makes the Heat

The main idea. Ethylene oxide is made from ethylene gas and oxygen on silver, in thousands of thin tubes standing in boiling water. The small share of ethylene that burns gives off most of the heat.

The challenge. A tube a few degrees too warm can overheat on its own, while a cooler one leaves spare oxygen nearer to catching fire. A trace of a chlorine compound helps the silver up to a point, and a simple controller cannot tell which side of that point the plant is on.

Why it matters. The plant must make product while it keeps clear of two dangers that pull opposite ways.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History boxes in Section 6.2, the Did you know? boxes in Sections 6.1 and 6.5, the Try it boxes in Section 6.5, and Key points in Section 6.9. No equations.
Study	40 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 6.5 with Figure 6.2, Section 6.6, and Section 6.7 with its Open challenge box.

The control question On an ethylene oxide reactor the eighth of the ethylene that burns makes two thirds of the heat, the cool steady state of the tube ends a few degrees above the design coolant, and the chloride that sets the selectivity has a maximum. Which handle protects which limit, what keeps the reactor off its cliff, and why can no fixed-sign loop hold the chloride at its optimum?

Learning objectives

- Compute the heat share of the unwanted reaction and explain why the reactor runs hotter as it makes less product.
- Locate the end of the cool steady state on the open tube, and explain how the recycle loop holds the reactor off it and why it still overshoots.
- Justify pairing production with the oxygen feed and an override that cuts the rate rather than the coolant.
- Explain why one chloride PID settles below the selectivity maximum and walks to a trip above it, and compute the resolution of a step search.
- Evaluate an economic predictive controller against a fixed chloride policy on production, margin and oxygen tracking.

Prerequisites Sections 1.4.3 (selectivity and runaway), 1.4.8 (one-way guards and back-off), 1.6.2 (velocity-form PID), 1.6.3 (predictive control option) and 1.6.4 (scenarios and Reset).

The app https://apmonitor.com/apps/ethylene-oxide/

See also Volume 7, Chapter 2 for an ignition boundary a few kelvin from the operating point, and Volume 7, Chapter 4 for a setpoint at the peak of a plant characteristic. Volume 4, Chapter 2 and Volume 4, Chapter 6 for protective overrides on a hard limit. Volume 9, Chapter 8 for a supervisory layer that chooses setpoints under a moving optimum, and Volume 9, Chapter 4 for offset-free predictive control with a disturbance estimate.

6.1 Introduction

A selectivity of 87.2 % sounds like a process in which the unwanted reaction is a detail. On an ethylene oxide reactor it owns the energy balance. Epoxidation releases 105 kJ/mol of ethylene and complete combustion releases 1327, so the 12.8 % of the ethylene that burns produces 65 % of the 40 MW coming off the tube bundle. Everything difficult about the unit follows from that ratio. The reactor is nine thousand narrow tubes standing in boiling water because the heat has to leave within a centimeter of where it is made. The combustion carries the larger activation temperature, 11 800 K against 8200, so it accelerates faster with every degree, which means the reactor gets hotter because it is making less of what is wanted. Section 1.4.3 derives this split for two parallel reactions in general, on the small oxidation tube of the running example.

+ DID YOU KNOW? Carbon Dioxide for Every Tonne of Oxide

Carbon dioxide is the main side product of ethylene oxide, 0.2 to 0.3 tonnes for every tonne of oxide. Industrial reactors convert only 7 to 15 percent of their ethylene on each pass, because the selectivity falls as the conversion rises [1]. In the reactor of this chapter 12.8 % of the ethylene burns.

The misconception the simulator corrects is that a minority reaction is a minor term. A second one follows. Measured on the model at the design gas composition, 4.2 °C above the design coolant temperature the tube has no cool steady state at all: the peak goes from 267.0 °C at 244.8 °C of coolant to 290.1 °C at 248.8, and from 248.97 °C the nearest solution is on the far side of seven hundred. That is parametric sensitivity, the classical property of a fixed bed in which a small parameter change produces a disproportionate change in the peak temperature [2, 3], and the safe region ends at a point rather than tapering. Volume 7, Chapter 2 measures the same contrast of exponential generation against linear cooling on a jacketed stirred tank, which can ignite after a jacket change of a few kelvin.

A third result is the one worth the most. The recycle loop around the reactor is a stabilizer: burning more oxygen leaves less oxygen to burn. Measured with the coolant fixed in manual and the override off, a thirty-five percent catalyst activity step settles the tube back at 267.6 °C on 5.68 mol% oxygen instead of 7.20, on a catalyst that would ignite a once-through tube. The recycle is worth more to this reactor than any controller on it, and a single-tube analysis would badly overstate how fragile the unit is. It also arrives late. The tube answers in about a minute and the loop takes eight minutes to turn over once, so the peak reaches 286.6 °C on the way, nineteen degrees past where it ends up and eight past the alarm.

This chapter works with an interactive ethylene oxide reactor whose kinetics, chloride response and flammability correlation are hidden, so that the energy split must be read from the model, the sensitivity boundary set against the recycle loop that keeps the reactor off it, and a chloride optimum approached by a method other than a controller. Its scenarios include one in which an identical PID controller settles from one side of an optimum and another in which it trips the reactor from the other. The epoxidation kinetics, parametric sensitivity and override control are established subjects, and the chapter applies them to one plant.

The chapter is organized as follows. Section 6.2 recalls how the process came to be a history of selectivity and the three features that set its control problem. Section 6.3 states the twelve-cell multitubular reactor inside its recycle loop, with its chloride dynamics and its moving flammability limit, in the detail needed to rebuild it. Section 6.4 describes the regulatory loops, why production is paired with the oxygen feed, the override that acts on the rate, the three ways to set the chloride and an optional economic model predictive controller. Section 6.5 reports the measured results, from the energy split and the cliff through the recycle, the chloride optimum, the search and the two conflicting margins to the graded challenge and the comparison with predictive control. Section 6.6 names the hardware and the practice behind the model. Sections 6.7, 6.8 and 6.9 state the limitations, the three lessons that carry beyond epoxidation and a summary, and the exercises follow in Section 6.10.

6.2 Background

Ethylene oxide is made at a scale of tens of millions of tonnes a year and is the starting point for ethylene glycol, polyester and most nonionic surfactants. Direct oxidation over silver replaced the chlorohydrin route, as the History box below tells, and the entire subsequent history of the process is a history of selectivity. Early silver catalysts reached the mid seventies of a percent. Modern high selectivity catalysts reach the high eighties, and the improvement came from promoters: cesium, rhenium and the chloride moderator that is the subject of four of the chapter's scenarios [4].

> ⧖ HISTORY **From Chlorine to Oxygen, 1937 and 1958**
>
> The route that direct oxidation replaced was a costly two-stage one. Ethylene reacted in the liquid phase with hypochlorous acid to an ethylene chlorohydrin, which hydrated lime then turned into the oxide, with a selectivity of about 80 percent and with ethylene dichloride and calcium chloride as unwanted by-products. Industrial production by direct oxidation started in 1937 with a process that fed ethylene and air, and in 1958 processes fed with oxygen instead of air were introduced. Air plants had to send part of their recycle gas to a second purge reactor, and they emitted more and yielded less. The price of chlorine and the effluent load retired the old route, and in 2013 the oxygen route was the only one in use in the European Union [5].

Three features set the control problem.

The first is the heat. The reactor is multitubular with boiling coolant on the shell, and heterogeneous models of industrial reactors of this type resolve the axial profile because the peak is what matters [6]. Parametric sensitivity is the property that governs how close such a reactor can be run to its limit, and the criteria for it were developed for fixed beds generally [2, 3] and are surveyed across chemistries in the safety literature [7].

Table 6.1 Classification of Variables at the Design Operating Point

Symbol	Role	Nominal	Range
EO rate	CV, continuous	8.00 kmol/min	2–12 SP
y_O	CV, analyzer	7.20 mol%	4–9 SP
y_E	CV, analyzer	24.00 mol%	15–32 SP
y_C	CV, analyzer	5.00 mol%	2–12 SP
T_{hot}	constraint	267.1 °C	alarm 278, trip 296
y_O^{lim}	constraint	8.60 mol%	moves
F_O	MV	7.606 kmol/min	0–14
T_c	MV	244.8 °C	200–280 °C
F_E	MV	9.461 kmol/min	0–18
w	MV	67.3 %	0–100 %
Cl	MV, optimizing	2.50 ppm	0–8 ppm
activity	DV, unmeasured	1×	0.55–1.5×
T_i, y_E, y_O, y_C, Cl	states	—	—

HISTORY Heat Made Against Heat Removed, 1928 to 1970

The picture behind this chapter's cliff is older than the process. In 1928 Semenov considered an exothermic reaction and drew the heat it releases, which rises exponentially with temperature, against the heat removed by heat transfer, which rises linearly. His diagram has a stable operating point and a critical one, and a coolant warmer than the critical value lets the reaction ignite itself. That work laid the groundwork for a run of runaway criteria. In 1970 van Welsenaere and Froment set critical conditions for fixed beds from the locus of temperature maxima, and Morbidelli and Varma drew the boundary where the reactor becomes sensitive to small changes in its inputs [7].

The second is the flammable envelope. The cycle gas is oxygen and hydrocarbon in one vessel at fifteen bar, kept out of the flammable region by ballast, and the limiting oxygen concentration is a function of the gas rather than a constant. Flammability limits move with temperature, pressure and diluent, and their determination at elevated conditions is its own experimental problem [8], with estimation methods used where measurement is impractical [9]. The practical consequence for control is that an oxygen controller holding a fixed setpoint is holding a moving margin. Volume 4, Chapter 6 guards a solvent vapor limit in a dryer with an override of the same one-directional kind.

The third is the chloride. A few parts per million of an organic chloride on the silver block the sites that burn ethylene, and the benefit saturates. The same chloride blocks the sites that do anything at all, and that penalty does not saturate. The result is a selectivity curve with a maximum, and the industrial practice of finding and holding that maximum is one of the oldest real-time optimization problems in the chemical industry. It is solved by periodic perturbation rather than by feedback, for the reason Section 6.5 measures.

Above all of it, a modern plant carries a multivariable controller that arbitrates between production rate, selectivity and the two constraints [10]. The application implements that regulatory layer and an optional economic MPC.

6.3 Process Model

Reactor R-501 is a bundle of 9200 tubes carrying 3.10 kg of silver catalyst each, in boiling coolant at 244.8 °C. Cycle gas at 205 °C enters at 320 kmol/min carrying 24.00 mol% ethylene, 7.20 mol% oxygen and 5.00 mol% carbon dioxide in a methane ballast, and makes 8.00 kmol/min of ethylene oxide at 87.21 % selectivity. Figure 6.1 shows the arrangement.

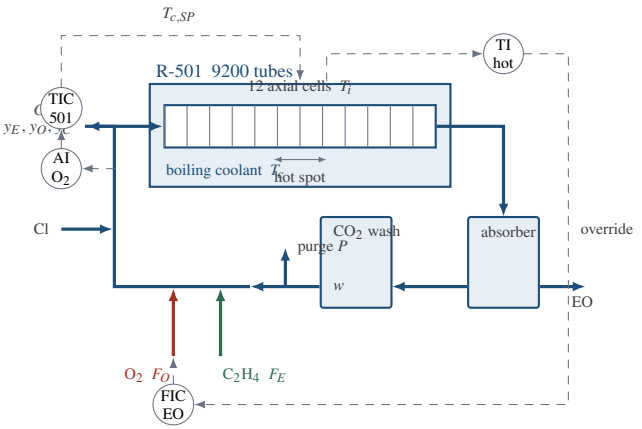

Figure 6.1 Ethylene oxide reactor R-501 inside its recycle loop. Cycle gas G of composition y_E, y_O and y_C enters twelve axial cells at temperatures T_i, cooled by boiling coolant at T_c. Product is washed out in the absorber, a slipstream w removes carbon dioxide, a purge P leaves, and fresh oxygen F_O, ethylene F_E and chloride Cl are added before the gas returns. TIC-501 holds the cycle gas oxygen by moving the coolant setpoint $T_{c,SP}$, FIC-EO holds production on the oxygen feed, and the hot spot thermocouple drives a one-directional override on the production setpoint.

6.3.1 Two Reactions on the Same Sites

Both reactions are first order in oxygen and half order in ethylene, inhibited by carbon dioxide, Equations (6.1) and (6.2).

$$r_1 = k_1 e^{-E_1/RT}\, y_O \sqrt{y_E}\, \phi\, \phi_{act} \tag{6.1}$$

$$r_2 = k_2 e^{-E_2/RT}\, y_O \sqrt{y_E}\, \phi\, \phi_{act}\, (1 - \beta_{Cl}) \tag{6.2}$$

with $E_2/R = 11\,800\,\mathrm{K}$ and $E_1/R = 8200\,\mathrm{K}$, and $\phi = 1/(1+4.6\,y_C)$ the carbon dioxide inhibition. The single difference between the two activation temperatures is why selectivity falls with temperature: the reaction that destroys the product is the one that accelerates faster. Since the two rates share their composition terms, their ratio depends on temperature alone, as Equation (1.4) states.

The heat released follows Equation (6.3), and the ratio in it governs the whole reactor. Weighted by these heats, the heat generated rises with temperature as one reaction with the effective activation of Equation (1.5).

$$q = r_1\,\Delta H_1 + r_2\,\Delta H_2, \qquad \frac{\Delta H_2}{\Delta H_1} = \frac{1327}{105} = 12.6 \tag{6.3}$$

🕮 **Pause.** Try Exercise 6.1 before reading on.

6.3.2 One Axial Cell and the End of the Cool Branch

Each cell balances generation against two removal terms, Equation (6.4).

$$C\,\frac{\mathrm{d}T_i}{\mathrm{d}t} = q_i - \dot{n} c_p (T_i - T_{i-1}) - UA\,(T_i - T_c) \tag{6.4}$$

Generation is exponential in temperature and both removal terms are linear in it, so the steady-state residual has one, two or three roots, the construction of heat generated against heat removed in Section 1.4.3. The working reactor is on the lowest of them, and the steady-state solver that initializes each case solves for it by scanning upward and taking the first sign change rather than by iterating, because a fixed point iteration converges to whichever root it happens to find. Where no crossing

Table 6.2 Model Parameters

Symbol	Meaning	Value	Class
E_1/R	epoxidation activation	8200 K	adjustable
E_2/R	combustion activation	11 800 K	adjustable
ΔH_1	heat of epoxidation	105 kJ/mol	adjustable
ΔH_2	heat of combustion	1327 kJ/mol	adjustable
UA	coolant transfer per cell	0.0200 kW/K	adjustable
β_{max}	burning blocked at saturation	0.80	adjustable
$\mathrm{Cl}_{1/2}$	chloride for half of it	1.20 ppm	adjustable
k_{Cl}	activity lost per ppm	0.340 per ppm	adjustable
τ_{Cl}	chloride uptake	55 min	adjustable
G	cycle gas circulation	320 kmol/min	adjustable
holdup	gas in the loop	2450 kmol	adjustable
θ	selectivity analyzer cycle	30 min	adjustable
T_{hot} alarm	hot spot alarm	278 °C	visible

exists the tube has no cool steady state, and the solver returns an explicit runaway flag rather than a hot root. A setting chosen before a run starts it at that setting’s steady state, unless the protection would trip on that state at once, in which case the run starts at the case’s own state and the setting acts from $t = 0$. The coolant helper and the shipped checks read that flag, while the running simulation integrates Equation (6.4) for all twelve cells with $C = \dot{n} c_p \tau_{tube}$ and $\tau_{tube} = 1.15\,\mathrm{min}$.

6.3.3 The Chloride Curve

Chloride blocks the burning sites with a benefit that saturates and all sites with a penalty that does not, Equations (6.5) and (6.6).

$$\beta_{Cl} = \frac{\beta_{max}\,\mathrm{Cl}}{\mathrm{Cl}_{1/2} + \mathrm{Cl}} \tag{6.5}$$

$$\phi_{act} = e^{-k_{Cl}\,\mathrm{Cl}} \tag{6.6}$$

A saturating good against a compounding bad is a curve with a maximum, and a maximum is a place where the process gain is zero and changes sign.

6.3.4 A Limit of Flammability That Moves

The limiting oxygen concentration is not a constant. Ethylene is the fuel and lowers it, and carbon dioxide is a heavy inert and raises it, Equation (6.7).

$$y_O^{lim} = y_0 - a\,(100 y_E - 24) + b\,(100 y_C - 5) \tag{6.7}$$

The correlation is a straight line fitted near the design composition and is clamped outside the range where it was fitted, because an extrapolated line is a fiction. The consequence for control is that the carbon dioxide the wash column removes for activity reasons is the same carbon dioxide that keeps the loop out of the envelope.

6.3.5 Assumptions

One tube stands for 9200 and the bundle is assumed identical. Real bundles have a distribution of flow and packing, and a runaway starts in the tube at the end of that distribution rather than in the average one.

Twelve lumped axial cells replace a distributed profile, so the hot spot is located to about one twelfth of a tube and the model understates the true peak.

Two lumped reactions represent a network that also makes acetaldehyde and can burn the oxide once formed.

The chloride effect is a fitted shape rather than a site balance, with one first-order lag. Real chloride behavior depends on the hydrocarbon carrying it and has a memory measured in days [4].

The gas loop is one well stirred tank with no transport delay. A real loop is an absorber, a compressor and a wash column in series.

The flammability limit is linear in two compositions and clamped. Real envelopes are measured surfaces applied with a margin nobody negotiates [8, 9].

6.4 Controller Design

6.4.1 Velocity Form PID

Every loop uses the shared velocity-form PID of Section 1.6.2, Equations (1.14) to (1.17), with back-calculation anti-windup on the integral alone and the tracking time $T_t = \tau_I$. Every loop here runs with $\tau_D = 0$, so only the proportional and integral increments appear. The action is $a = -1$ for the oxygen loop, because a hotter coolant leaves less oxygen in the cycle gas, and for the carbon dioxide loop, because more wash leaves less carbon dioxide. It is $a = +1$ for the production, ethylene and chloride loops. A loop in manual tracks the output actually held, which makes the return to automatic bumpless, and during a trip, which forces the feeds to zero and the wash to full, every loop tracks the forced values in the same way.

One limit acts after the PID output and is unknown to it. A carbonate wash that has lost circulation caps the applied slipstream below the 100 % the carbon dioxide loop believes it has. With half the circulation lost (scenario 11) that loop commands 100 % from 68 min onward while 50 % reaches the column, and its integral settles at 100 % with a demanded output, the output before limiting v_k of Equation (1.16), of 157 %, back-calculated against its own limit rather than against what the column delivers. The controller scan is 0.05 min and integration is fixed-step fourth-order Runge-Kutta at 0.01 min (Section 1.6.1).

6.4.2 Why Production Is on the Oxygen Valve

At steady state in a closed recycle loop every oxygen molecule fed is either reacted or purged, Equation (6.8), so the oxygen feed sets the production rate and the coolant temperature cannot.

$$F_O = \dot{n}_{O_2,reacted} + P\,y_O \tag{6.8}$$

Measured in scenario 2 with the oxygen feed fixed and the ethylene and carbon dioxide loops holding, two degrees of extra coolant, stepped 10 min into the run, moves production from 8.00 kmol/min to 7.91 kmol/min within two hours, and the cycle gas oxygen settles at 6.80 mol% instead of 7.20. Production first rises to 8.43 kmol/min within eight minutes and then falls back as the loop composition thins, so the coolant has gain on the rate only for about one loop turnover. At steady state with every feed fixed the same two degrees give 7.91 and 6.22 mol%. What the coolant sets is where the loop composition comes to rest, which is to say how much unreacted oxygen is standing in a fifteen bar loop that must never be flammable, and that is a more important thing for it to decide than the rate.

6.4.3 The Override Acts on the Rate

Cooling the reactor would bring the hot spot down and would leave more unreacted oxygen in the cycle gas, walking the loop toward the envelope. The two constraints want opposite things from the coolant valve. At steady state with every feed fixed, cooling from 244.8 °C to 241.0 °C takes the hot spot from 267.1 °C to 262.5 °C and takes the oxygen margin from 1.40 mol% to −0.50 mol%, while heating to 248.8 °C gives 272.0 and 1.82. The only handle that moves both margins the right way is the oxygen feed, and it is the one that costs money. The override is therefore one-directional on the production setpoint, Equation (6.9), one of the fast one-way guards of Section 1.4.8.

$$\mathrm{EO}_{SP} \leftarrow \mathrm{EO}_{SP} - K_o\,[T_{hot} - (T_{alarm} - m)]^{+} \tag{6.9}$$

The default gain is $K_o = 0.35\,\mathrm{kmol/min}$ per degree and the margin $m = 8\,^\circ\mathrm{C}$. The trimmed value is the setpoint of the production loop, so the override reaches the oxygen valve through that loop rather than through any integral state, and T_{hot} is the model's hottest cell with no thermocouple lag.

Pause. Try Exercise 6.3 before reading on.

6.4.4 Three Ways to Set the Chloride

With the Loops supervisor selected, the application offers manual, a PID on the selectivity analyzer, and a step and search optimizer. Only two of them can work, and Section 6.5 measures which.

6.4.5 Economic Model Predictive Control

Section 1.6.3 describes the predictive option of the framework. This one differs in its objective, which prices production and selectivity instead of weighting the distance from setpoints. The optional MPC coordinates $u = [F_O, T_{c,SP}, \mathrm{Cl}_{feed}]$ while the ethylene and carbon dioxide PI loops remain regulatory. Manual handles remain fixed. A 19-state nominal model contains the 12 cell temperatures, coolant, three recycle fractions, catalyst chloride, and two regulatory commands. Central differences of the declared balances at the design point, rather than current plant parameters or hidden activity, supply the affine matrices A_d, B_d, E_d, c_d, C_d in Equation (6.10).

$$x_{k+1} = A_d x_k + B_d v_k + E_d \hat{d}_k + c_d, \qquad \hat{y}_k = C_d x_k + y_0 \tag{6.10}$$

Here x and v are state and command deviations, $\hat{d}$ contains estimated constant disturbances, and y comprises EO rate, selectivity in percentage points, cell-8 temperature, oxygen, and oxygen minus its flammability limit. The first four disturbances enter through E_d, while the fifth shifts the hot output in y_0. Known regulatory setpoints and purge enter the affine term. Symmetric nonlinear steps of ± 0.1 kmol/min oxygen feed, ± 0.5 °C coolant and ± 0.25 ppm chloride give at most 0.547% dominant-output response error at 20, 60 and 180 min. This is local validation, not an ignition model.

The first solve occurs at $t = 0$. Thereafter $T_s = 2$ min spans 40 unchanged controller scans. The 180 min horizon has $N = 90$ steps with move blocks $[1, 1, 2, 4, 8, 16, 58]$, or 21 command decisions. A constrained quadratic program [11] minimizes

$$\begin{aligned} J = \sum_{k=1}^{90} T_s[-a_E \hat{E}_k - a_S \hat{S}_k + a_S \kappa (\mathrm{Cl}_{feed,k} - \mathrm{Cl}_*)^2] \\ + \sum_{b=1}^{7} \sum_{i=1}^{3} 0.1(\Delta u_{i,b}/s_i)^2 + \sum_j (\alpha_j \varepsilon_j + \beta_j \varepsilon_j^2). \end{aligned} \tag{6.11}$$

The local derivatives of $M = E - 0.62E/(S/100)$ are $a_E = 0.289083$ and $a_S = 0.065215$, where E is EO in kmol/min and S is percent selectivity. The declared chloride curve gives $\mathrm{Cl}_* = 2.492825$ ppm and $\kappa = 0.348829$ percentage points/ppm^2. The block-boundary move Δu uses scales $s = [0.1\ \mathrm{kmol/min}, 1\,^\circ\mathrm{C}, 0.1\ \mathrm{ppm}]$. A 10^{-7} Hessian ridge regularizes the program. There is no terminal cost or disturbance preview.

Command boxes are hard: $0 \le F_O \le 14$ kmol/min, $200 \le T_{c,SP} \le 280$ °C and $0 \le \mathrm{Cl}_{feed} \le 8$ ppm. Production and oxygen requests are soft upper bounds, with $(\alpha, \beta) = (10^3, 10^3)$. The hot bound is 1 °C below the override start, 269 °C in the new case. The oxygen-envelope bound backs off 0.2 mol% from the existing trip guard. Both use $(10^4, 10^3)$. Each bound L_j initially expands to $L_j + [\hat{y}_{j,0} - L_j]^+ \exp(-kT_s/10)$. One nonnegative ε_j per bounded output is shared across the entire horizon. Excess above this funnel can therefore persist up to that same slack without a separate later penalty. These priced bounds do not guarantee nonlinear plant constraint satisfaction. The override and trip retain authority.

A steady-state Kalman gain corrects 17 physical states plus five constant states for activity, the three gas balances and hot-spot offset [12]. Its five readings are EO, oxygen, ethylene, carbon

Table 6.3 Hot Spot Against Coolant Temperature at Design Composition

Coolant (°C)	Hot spot (°C)
244.8	267.0
246.0	271.0
248.0	281.0
248.8	290.1
248.9	293.6
249.0	no cool solution (741)

dioxide and the hottest cell. Selectivity and hidden activity are not measurements. The diagonal process covariance is 10^{-4} for each temperature, 10^{-10} for each gas fraction and 10^{-6} for chloride. Disturbance variances are $[10^{-5}, 10^{-10}, 10^{-10}, 10^{-10}, 10^{-2}]$. Measurement variances are the squares of $[0.02, 0.03, 0.05, 0.03, 0.1]$ in the native measurement units, in the stated order. Known applied commands drive prediction, including the commanded wash before its unmeasured failure cap. The instrument filters remain in the simulation but are omitted from the controller model. Disturbance estimation corrects prediction offsets, but economic ceilings do not enforce zero setpoint error. The losing PIDs track applied commands. A rejected, nonfinite or out-of-range plan returns authority to those loops, marks fallback, and resumes MPC after a valid solve. Chloride search waits a full dwell on handover.

6.5 Results

Thirteen scenarios run from the energy split to a graded challenge, in the order the app lists them: 1, “Which reaction makes the heat” (the energy split), 2, “Five loops and what each one owns” (pairing in a recycle), 3, “The hot spot is not the outlet” (moving constraints), 4, “Where the cool solution stops existing” (parametric sensitivity), 5, “The recycle is the real stabilizer” (slow negative feedback), 6, “Chloride, in both directions” (a flat maximum), 7, “A PID above the optimum” (positive feedback), 8, “The same controller below it” (conditional correctness), 9, “Step, wait, compare, reverse” (search resolution), 10, “Two margins, opposite directions” (constraint conflict), 11, “The wash column is a safety device” (inert double duty), 12, “The reaction you do not want, priced” (economic coordination), and the challenge, “A shift on a fresh charge”. Scenarios 5 and 7 are demonstrations, and the other ten numbered cases are experiments. Scenarios 7 and 8 differ only in a starting condition and are meant to be run back to back. Each case is credited separately when the learner performs the listed actions and the simulation reaches the listed state (Section 1.6.4). The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt, and the credit accumulates toward certificates of progress and of completion.

> ✦ **DID YOU KNOW? Cooling the Reactor Can Spend the Safety Margin**
>
> From the simulator: with every feed fixed, cooling from 244.8 to 241.0 °C lowers the hot spot from 267.1 to 262.5 °C, and it takes the oxygen margin from 1.40 to −0.50 mol%, past the limit of flammability (Table 6.6). The move that looks safe for the silver is not safe for the gas.

6.5.1 The Energy Split and the Cliff

At the design point the combustion produces 65.0 % of the 40 MW duty at 87.21 % selectivity. Table 6.3 reports the coolant temperature swept on the open tube at the design gas composition.

The cool solution ends at 248.97 °C, so 4.2 °C is the entire distance between the design point and a tube that cannot stand up. The steps accelerate because the exponential was always going to outrun the lines and the only question was where.

Scenario 4 runs the same coolant steps, to 249.4 °C, on the full loop with every station in manual, and the cliff does not appear. The hot spot reads 268.5, 270.8 and 272.4 °C at the ends of the three

Table 6.4 Chloride Against Selectivity and Hot Spot, All Loops Holding

Chloride (ppm)	Selectivity (%)	Hot spot (°C)
1.50	86.77	254.9
2.10	87.15	262.1
2.40	87.21	265.8
2.70	87.19	269.6
3.30	87.01	277.5

Table 6.5 One Chloride PID, Two Starting Points

Start (ppm)	Reaches (ppm)	Outcome
1.50	1.97	still rising at 10 h, 2.09 at 50 h
3.30	4.76	trips on hot spot at 381.8 min

steps (at 60, 110 and 240 min) and never exceeds 273.8 °C, because a hotter tube consumes its own oxygen and the cycle gas falls from 7.20 mol% to 6.33 mol%. The cliff belongs to the tube at fixed composition. What keeps the reactor off it is the subject of the next subsection.

6.5.2 The Recycle Stabilizes and Arrives Late

With the coolant fixed in manual, the override off and the catalyst made thirty-five percent more active at 20 min (scenario 5), the cycle gas oxygen falls from 7.20 mol% to 5.68 mol% by 4 h and the tube settles at 267.6 °C, almost exactly where it started. The production loop stays closed and ends with the oxygen feed at 7.59 kmol/min against 7.61, so the stabilizing action is the loop composition rather than the feed. On the way the peak reaches 286.6 °C at 25.9 min, nineteen degrees past where it ends up and eight past the alarm, and it stays above the alarm for 7 min, because the tube answers in about a minute and the loop takes eight minutes ($2450\,\text{kmol}/320\,\text{kmol/min} = 7.7\,\text{min}$) to turn over once.

> **TRY IT IN THE APP** **5 · The Recycle Is the Real Stabilizer**
>
> Load the scenario, which starts with the oxygen loop in Manual and the override Off, and press Run for four hours. Watch the Hot spot and Cycle gas oxygen plots. The catalyst turns more active at 20 min, the hot spot peaks at 286.6 °C at 25.9 min, and it settles at 267.6 °C as the oxygen falls from 7.20 to 5.68 mol%. The coolant stays fixed throughout.

Both halves of that result matter. A steady-state analysis of this loop says the reactor cannot run away. A dynamic one says the silver has been where silver should not be.

6.5.3 The Chloride Optimum and the Controller That Cannot Hold It

Table 6.4 reports the chloride swept with all four composition and rate loops holding their setpoints and the override off, each value stepped in the feed 10 min into a run from the design point and read 20 h after the step. Scenario 6 reproduces the two ends, 86.77 and 87.01 % with the hot spot at 255.0 and 277.5 °C, at the end of each step (500 and 960 min).

The maximum is near 2.5 ppm (2.49 at the design composition and rate) and it is flat: a part per million either side costs less than half a point of selectivity, 0.45 below and 0.29 above. The hot spot over the same range is not flat at all, moving 22.6 °C. The variable that is nearly invisible in the economics is enormous in the temperature, and it is the temperature that ends the campaign. Figure 6.2 draws the table on one chloride axis.

Now put a PID on it. Table 6.5 reports the same controller, the same setpoint of 87.15 % and the same reactor, started on either side of the maximum.

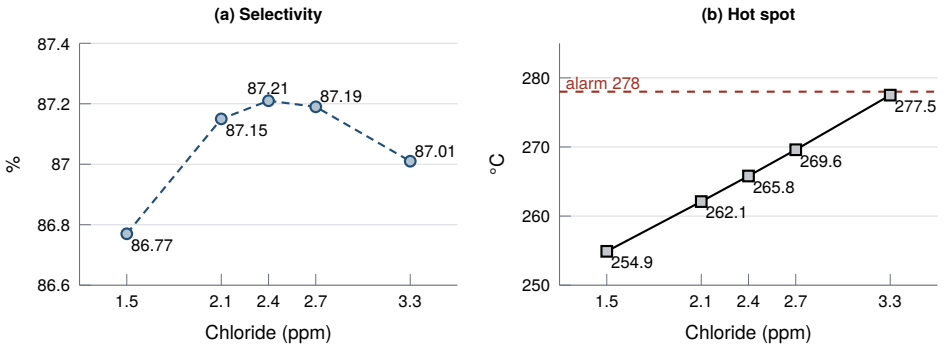

Figure 6.2 The chloride sweep of Table 6.4, with all four composition and rate loops holding. Circles on a dashed line are the selectivity, squares on a solid line the hot spot, and the dashed red line in (b) is the hot spot alarm. Notice the scales: across the whole sweep the selectivity moves less than half a point, while the hot spot climbs 22.6 °C, to 277.5 against an alarm at 278 °C.

Above the maximum the controller pushes chloride up because up is what raises selectivity everywhere it has ever been, and here up lowers it. The error grows, so the controller pushes harder, so the error grows. The loop does not oscillate, it walks, and it walks the coolant up with it until the unit trips. In scenario 7 the hot spot starts at 277.4 °C, passes its alarm at 23.2 min and reaches the 296 °C trip at 381.8 min, with 4.76 ppm on the silver against a command of 5.28 and a coolant request of 272.3 °C. Below the maximum the same controller is stabilizing: in scenario 8 chloride reaches 1.97 ppm at ten hours and 2.09 after fifty hours. Each case starts from its converged feed operating point; the controller and tuning are the same, while the initial compositions follow those different chloride/coolant stations.

> **TRY IT IN THE APP** **7 · A PID Above the Optimum**
>
> Load the scenario, which starts with the chloride on its PID and the override Off, run ten hours and watch the Selectivity and Hot spot trends. The chloride climbs from 3.30 to 4.76 ppm and the unit trips at 381.8 min. Then load scenario 8, “The same controller below it”, and run ten hours more: the same loop brings the chloride to 1.97 ppm and nothing trips.

That is the uncomfortable part. The two runs differ in their starting operating conditions, and nothing in the loop, the tuning, the trend or the alarm list distinguishes the case that works from the case that trips the reactor. A feedback controller assumes the sign of the process gain, and at an optimum the gain is zero and the sign is undefined.

Pause. Try Exercise 6.2 before reading on.

6.5.4 What a Search Can and Cannot Resolve

The step and search optimizer climbs, overshoots, reverses and then hunts across the top forever. Measured over fifty hours from 1.50 ppm it makes twenty-four moves of 0.3 ppm on a two-hour dwell, and after the first ten hours its chloride lives between 1.84 and 2.96 ppm around a true optimum near 2.5, reversing five times. The hunting is not a defect to be tuned out. A search that settles has stopped measuring the gradient.

The width of the hunt follows from arithmetic rather than from tuning. A step of size h at a distance d from a smooth peak changes the objective by roughly $|S''|dh$, so a step can be told from the smallest difference the search trusts, σ, only where $d > \sigma/(|S''|h)$. Here the curvature is about 0.7 % of selectivity per part per million squared, the step is 0.3 ppm and the search ignores differences under 0.04 %, so no step within about 0.2 ppm of the top can be resolved. The chloride uptake (55 min) and

Table 6.6 The Two Constraints Move Opposite Ways With the Coolant

	All feeds fixed		Scenario 10	
Coolant (°C)	Hot spot (°C)	O_2 margin (mol%)	Hot spot (°C)	O_2 margin (mol%)
241.0	262.5	−0.50	262.8	0.52
244.8	267.1	1.40	267.1	1.40
248.8	272.0	1.82	271.6	2.18

the half-hourly analyzer let only part of each step register before the two-hour dwell ends, which is why the turns come at 1.8 and 3.0 ppm, about two steps either side of the top. With 0.5 % of analyzer noise $\sigma/(|S''|h)$ exceeds two parts per million, and the second half of scenario 9 shows it: ten reversals in two days instead of five and, after the first ten hours, a band of 1.54 to 2.66 ppm on the shipped noise sequence, while one of four other sequences wanders from 0.02 to 4.67 ppm. The consolation is in the same arithmetic. The optimum is flat because the curvature is small, so half a part per million off costs less than a tenth of a point, and the quiet hunt averages 87.16 % against 87.21 at the peak.

6.5.5 Two Margins in Opposite Directions

The fixed-feed solver now closes all three material balances to below 10^{-9} kmol/min. The cold equilibrium is outside the oxygen envelope; it is not a safe operating point. Runs with no safe fixed-feed equilibrium use a safe case composition line-up, or the declared design point if no such line-up exists, with the chosen stations acting from $t = 0$. Table 6.6 reports the coolant swept twice, as a steady state with every feed fixed and as scenario 10 runs it, with the production and oxygen loops open and the ethylene and carbon dioxide loops holding their compositions, read at the end of each step.

With the ethylene and carbon dioxide held at their setpoints the limit of flammability stays near 8.6 mol%, so the scenario loses less margin on the cold side, but the two margins still move in opposite directions. The hot step is also not gentle in the run. The hot spot reaches 286.4 °C 11 min after the coolant goes to 248.8 °C and stays above its alarm for 10 min before settling. With the two loops then closed and the override on, the rate is trimmed for only 6 min while the oxygen loop takes the coolant back.

No direction of coolant movement improves both margins at once. Cutting the oxygen feed moves both the right way and costs production, which is why the override acts there. When two constraints conflict on a handle, the move that relieves both is almost always the one that gives up the objective.

6.5.6 The Graded Challenge

The challenge runs fifteen hours on a fresh catalyst charge that becomes more active twice, with a carbonate wash that loses part of its circulation, with cautious tuning and three structural decisions switched off. It is met when the cycle gas oxygen is held within 0.12 mol% of setpoint through the last third, from 594 min to the end, the hot spot stays below its 278 °C alarm throughout, and the unit does not trip. The shipped settings miss the oxygen band, with a worst late error of 0.314 mol%; the suggested settings pass at 0.105 mol%. Their hot-spot peaks are 272.6 and 274.4 °C, with no trips. On seeds 1–4 the suggestion gives 0.137, 0.083, 0.109 and 0.092 mol%; one of the four fails. Retuning only the oxygen loop ($K_c = 7$, $\tau_I = 18$ min) passes at 0.077 mol%, but leaves only 1.56 mol% of oxygen margin at the end, against 1.75 for the full suggestion. Leaving composition loops in manual with the other suggested settings leaves 1.59. Restoring the chloride to 2.50 ppm raises final selectivity by 0.21 points but also raises the final hot spot by 8.4 °C.

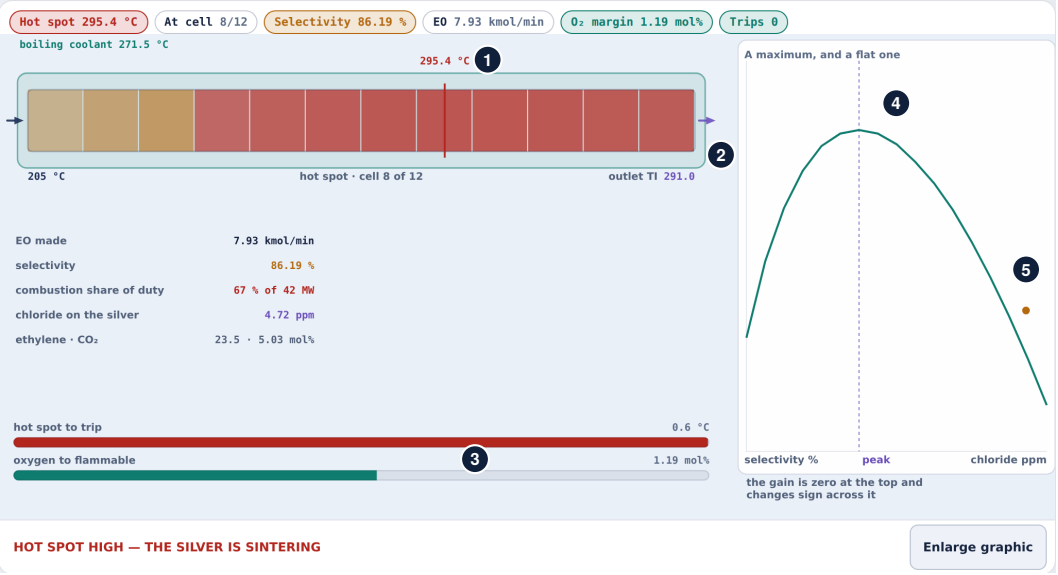

Figure 6.3 The ethylene oxide application running scenario 7, the chloride PID started above the optimum with the override off, at 378 min, about four minutes before the trip. One tube is drawn the long way with its twelve axial cells shaded by temperature and the hot spot marked where it is (1), 295.4 °C in cell 8 of 12 against 291.0 °C at the outlet thermocouple (2). Bars under the tube (3) show the two margins, 0.6 °C from the hot spot to its trip and 1.19 mol% from the cycle gas oxygen to the limit of flammability. The inset draws the model's selectivity against chloride at the design composition and rate, with its peak marked at 2.5 ppm (4), and places the present operating point (5), 4.72 ppm and 86.19 %, far to the right of the peak and 0.2 points above the curve, which is drawn for the design gas.

6.5.7 Measured Economic and Tracking Trade-offs

Case 12 reduces activity from 1.00 to 0.85 at 60 min and raises the rate request from 8.0 to 8.8 kmol/min at 240 min. Table 6.7 uses seed 20260921, identical instruments, noise, events and initial conditions, with no preview or case-specific MPC retuning. The search case runs 2880 min and the other rows run 900 min. EO and oxygen IAE use the shared scorecard definitions. The EO total is the application's own scorecard total, the left-rectangle sum of the rate in kmol/min over the 0.01 min physics steps while the unit is not tripped. The reported $\bar{M}$ is the time average of $E - 0.62C_2$used in product-equivalent kmol/min, not dollars per hour.

MPC yields $\bar{M} = 2.338943$, against 2.093185 with search and 2.340381 with the best tested fixed-chloride policy. Its deficit to that comparator is 0.06145%, a near tie. The hindsight sweep tests 1, 1.25, 1.5, 1.75, 2 and 2.25 ppm, selecting 1.5 ppm after the first 0.01 min step so the initial state is shared. The two lowest settings trip. MPC makes 7706.8 versus 7689.2 kmol and peaks at 269.783 versus 271.529 °C, but oxygen IAE rises from 35.081 to 76.098 mol%·min. All four main policies spend zero minutes above the alarm or oxygen trip guard; this comparison establishes no alarm-avoidance advantage. Oxygen-feed travel is 5.099 versus 7.992 kmol/min, and coolant travel 16.371 versus 30.108 °C.

The wash case with manual chloride remains a loss: MPC EO IAE is 68.884 versus 3.216 kmol; automatic chloride reduces it to 1.060. Both suggested challenge policies pass. Their oxygen IAE is 57.908 for MPC and 54.131 for loops, and mean margins are 2.329 and 2.322. Their worst late errors are 0.053 and 0.105 mol%; shipped MPC misses at 0.198 mol%.

Table 6.7 Measured Control Comparison With Unchanged MPC Settings

Case	Policy	EO IAE kmol	O_2 IAE mol%·min	Peak hot °C	EO total kmol	$\bar{M}$ equivalent	V_{F_O} kmol/min	V_{T_c} °C	Alarm min	Plans/ rejected	Mean ms
12	Search	1171.42	53.95	277.73	6567.5	2.093185	23.31	44.45	0.00	0/0	0.00
12	Fixed 2.5 ppm	711.99	29.05	277.14	7026.9	2.215296	25.31	42.65	0.00	0/0	0.00
12	Hindsight 1.5 ppm	74.44	35.08	271.53	7689.2	2.340381	7.99	30.11	0.00	0/0	0.00
12	MPC	31.48	76.10	269.78	7706.8	2.338943	5.10	16.37	0.00	450/0	1.39
3	Loops, manual Cl	34.75	15.57	274.15	7201.7	2.370899	5.75	15.20	0.00	0/0	0.00
3	Loops, search	52.09	68.29	274.15	7200.9	2.367571	7.17	36.75	0.00	0/0	0.00
3	MPC, manual Cl	21.97	15.06	273.44	7199.2	2.370326	2.46	10.92	0.00	450/0	1.34
3	MPC, auto Cl	20.79	15.05	273.45	7201.9	2.370540	2.53	10.61	0.00	450/0	1.19
11	Loops, manual Cl	3.22	5.65	269.93	7197.0	2.291114	0.17	2.23	0.00	0/0	0.00
11	MPC, manual Cl	68.88	1.16	269.00	7131.2	2.275079	0.39	1.82	0.00	450/0	1.38
11	MPC, auto Cl	1.06	1.48	269.00	7199.2	2.291223	0.26	1.36	0.00	450/0	1.38
9	Search	101.86	220.53	273.28	23039.4	2.308247	9.38	95.86	0.00	0/0	0.00
9	MPC	6.52	55.05	267.30	23045.1	2.310983	1.30	14.43	0.00	1440/0	1.36
2	Search	1165.78	46.10	278.09	6706.5	2.165959	18.04	31.62	9.96	0/0	0.00
2	MPC	13.40	0.94	269.22	7860.1	2.476808	1.37	6.33	0.00	450/0	1.40
13	Loops, shipped	85.91	412.24	272.56	7211.4	2.310520	1.30	31.16	0.00	0/0	0.00
13	Loops, suggested	103.54	54.13	274.41	7188.3	2.322046	15.42	112.31	0.00	0/0	0.00
13	MPC, shipped	48.12	92.99	272.76	7243.0	2.331915	11.20	27.45	0.00	450/0	1.21
13	MPC, suggested	44.06	57.91	271.00	7218.5	2.329163	14.14	28.48	0.00	450/0	1.24

Case 12 is aging plus a rate request, case 3 fresh catalyst, case 11 wash loss, case 9 chloride search, case 2 a rate request at 60 min, and case 13 the unchanged challenge. V is total command variation from the scorecard. $\bar{M}$ is normalized product-equivalent kmol/min. The hindsight row selects the best of six fixed chloride settings after the first 0.01 min step, preserving the initial state but using foreknowledge. Zero values in the final column indicate no optimization. Times depend on the host. Every row has zero oxygen trip-guard duration and zero oxygen-feed and coolant hard-bound saturation.

Removing only the chloride-curvature cost still causes a late trip. A production-only objective makes 7717.7 kmol but lowers mean margin to 2.277, against 2.339 for MPC. Halving the physics step changes MPC product by -0.000493 kmol and mean margin by $-6.33e-07$, preserving the comparison order. The application checks cover the estimator, fallback and converged material balances. The comparison table reports current Node planning times; these machine-dependent times cover planning rather than all initialization work.

6.6 Technology in Practice

Figure 6.5 draws the loops of Section 6.4 with the hardware of a real unit in place of the schematic's symbols.

What Measures It

The oxygen loop reads the cycle gas oxygen from an analyzer, and the ethylene and carbon dioxide loops read theirs the same way. The selectivity reaches the chloride loop only once every 30 min, from an analyzer of the kind that suffers from large measurement delays and high investment and maintenance costs [13]. The hot spot that drives the override is the model's hottest cell, read with no thermocouple lag, while Figure 6.3 shows the outlet thermocouple reading 291.0 against a hot spot of 295.4 °C.

> **TECHNOLOGY** **Sensor: A Paramagnetic Oxygen Analyzer**
>
> No other industrial gas is nearly as paramagnetic as oxygen, so an analyzer that senses the pull of a magnetic field on its sample reads oxygen and ignores the rest. Nitric oxide is one of the few gases that interfere [14]. The classic design, a small dumbbell hung in the field, answers in 3 to 10 seconds, set mainly by how fast the gas in the cell is exchanged, and it is sensitive to vibration [15]. A span check above the oxygen of air needs pure oxygen as its standard, and an analyzer system needs much more maintenance than other instruments [14].

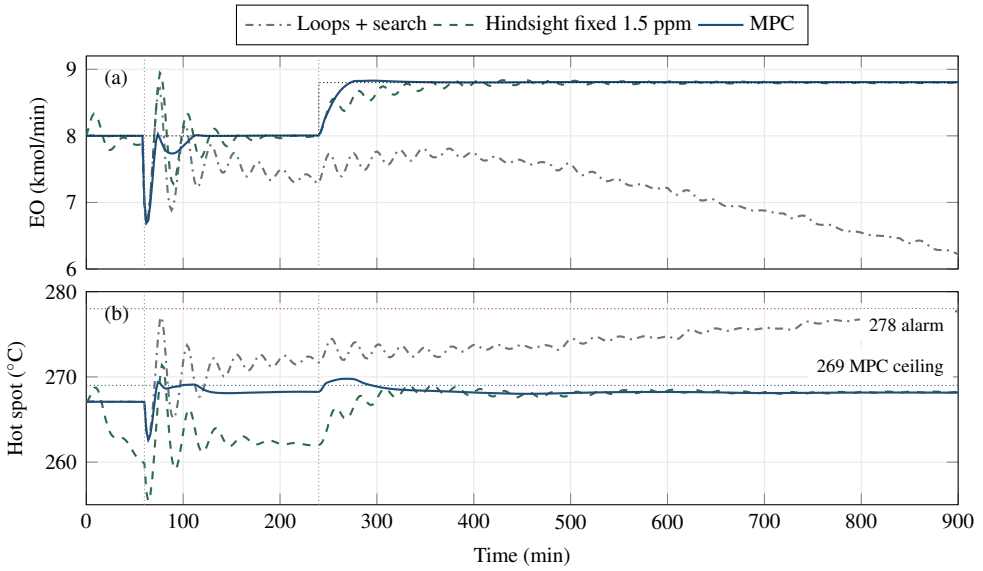

Figure 6.4 Full-duration aging and rate-request comparison: loops with the chloride search (dash-dotted), hindsight fixed 1.5 ppm (dashed) and MPC (solid). Activity falls at 60 min and the request rises at 240 min (vertical dotted lines). The upper dotted line in (a) is the requested rate, and the dotted lines in (b) are the labeled alarm and MPC ceiling. MPC and hindsight fixed 1.5 ppm produce 7706.8 and 7689.2 kmol, with mean margin indices 2.338954 and 2.340378. The fixed comparator uses a disclosed six-setting hindsight sweep and switches after the first 0.01 min physics step. MPC has worse oxygen IAE, 76.053 versus 35.108 mol%·min. No shown run exceeds the alarm.

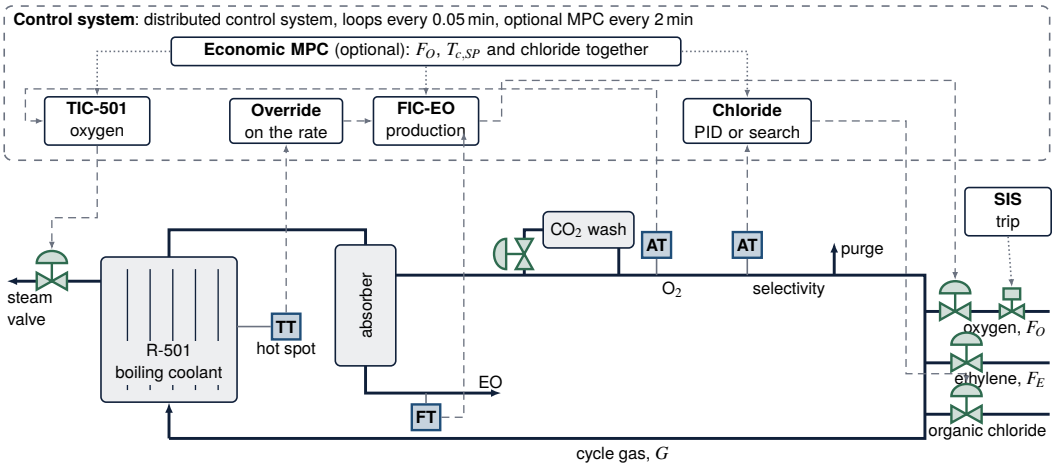

Figure 6.5 The reactor's main loops with their hardware. Tagged boxes are sensors, valve symbols final elements, dashed arrows signals and dotted ones the optional MPC and the trip. The steam valve sets the boiling pressure and so T_c. Notice that the hot spot reaches the oxygen valve only through the production setpoint. The ethylene and carbon dioxide loops, the trip's sensors, the stripper and the compressor are left out. Not to scale.

What Moves It

Five final elements act on the loop. The oxygen feed valve carries the production loop and the override. The coolant is boiling on the shell, where the reaction heat is removed and recovered by

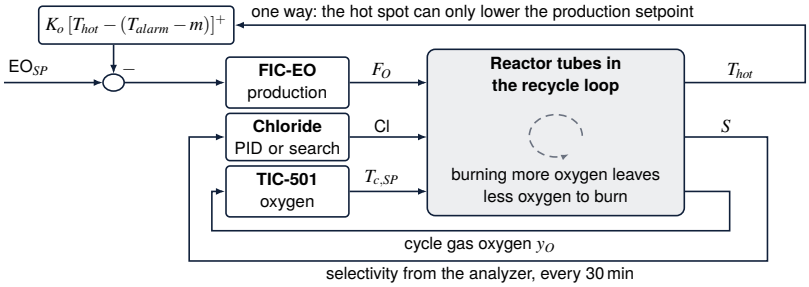

Figure 6.6 The regulatory structure as blocks. The hot spot can only lower the production setpoint, through the one-way trim of Equation (6.9). Notice the dashed loop inside the plant: the recycle's own feedback. The ethylene and carbon dioxide loops and the optional MPC are left out.

raising steam [5], and the boiling temperature of a liquid is a function of the absolute pressure it experiences [14], so the coolant setpoint $T_{c,SP}$ is in practice a pressure on the boiling side. The wash slipstream passes through a hot potassium carbonate absorber and returns to the recycle gas, and the chloride is an organic compound, typically ethyl chloride or dichloroethane, dosed into the recycle gas [5]. The model caps the wash when the column loses circulation, as in scenario 11.

> **TECHNOLOGY** **Final Element: The Oxygen Feed Control Valve**
>
> An industry code for oxygen piping counts process control valves as the most severe class of service in gaseous oxygen. They throttle at a high differential pressure, with fast, turbulent flow striking the trim, the body and at least eight valve diameters of the pipe downstream. Such valves are generally globe, modified ball, eccentric plug or butterfly types. In an oxygen system the oxygen is the oxidizer, the metals, seals, lubricants and any particles or oils are the fuel, and particle impact, compression heating and friction are the ignition sources. The design answers with burn-resistant alloys where ignition can occur and with cleaning for oxygen service [16].

Where the Controller Runs

A distributed control system holds loops like these at a supervisory level that sends set points to field controllers, with loops nested or cascading and cycle times from milliseconds to minutes [17]. The override is a selection in which only one controller output at a time acts on the process [14], here the lower of the operator's production setpoint and the one the hot spot allows. The trip belongs to a safety instrumented system with its own sensors and a logic solver separate from the regular controls, whose final elements are on/off valves or solenoids that force a control valve shut [14]. In the app every loop scans each 0.05 min in velocity form, the chloride search moves once a dwell, and the economic MPC sits above the loops and plans every 2 min. Figure 6.6 draws the same structure as blocks, with the recycle as the slow feedback inside the plant.

IN PRACTICE **Methane Ballast and a Lost Coolant**

Every site that answered the survey behind the European reference document dilutes its recycle gas with methane. Against nitrogen, methane allows more oxygen, limits the temperature rise better, needs less energy to compress and lets the purge be burned as fuel, and nitrogen serves only for a limited time when methane is short. A purge of about 0.1 to 0.2 percent of the recycle gas keeps ethane, argon and nitrogen from building up. The document lists a lost coolant supply among the abnormal events: the reactor overheats, may run away, and its pressure relief system must then act [5]. The trip in this chapter forces the feeds to zero and the wash to full instead.

What Is Modernizing It

Model predictive control has a solid foundation in refining and petrochemicals [10], and the economic version of this chapter prices production and selectivity instead of setpoint errors. Soft sensors estimate quality between analyzer reports and are also used to validate online analyzers [13], which would let a chloride search judge a step sooner than the half-hourly selectivity allows. Paramagnetic sensors built as microelectromechanical devices answer faster because their small dead volumes exchange the gas quickly [15], and a laser analyzer tuned to a single absorption line tells one compound from another even where their absorption bands overlap [14].

6.7 Limitations

One tube stands for the whole bundle, so the tube-to-tube distribution that determines where a real runaway starts is absent.

Twelve axial cells understate the true peak, and a real profile is sharper than twelve cells can render [6].

The chloride model is a fitted shape with a single lag, and real chloride has a memory of days and depends on the carrier [4]. The reader does not learn here how long a chloride change actually takes to settle on a plant.

The flammability correlation is linear in two compositions and clamped outside the range it was fitted in. The clamp is honest and it means the application cannot be used to explore compositions far from design [8].

The gas loop is one well stirred tank, so the transport delays that every composition loop on a real plant contends with are absent.

The optional economic MPC uses nominal local slopes rather than the full nonlinear industrial optimization problem [10]. Its model omits ignition curvature, analyzer lags, hot-spot migration and direct nonlinear activity effects on the EO/selectivity map. Estimated activity is a disturbance coordinate, not a calibrated catalyst measurement. The shared slacks and finite horizon provide no invariance proof. Wash failure and the suggested challenge demonstrate that prediction correction does not remove an economic tracking trade-off.

OPEN CHALLENGE **Why Chloride and Rhenium Work**

The high-selectivity catalysts of industry carry rhenium and alkali promoters on the silver, and an organic chloride dosed in the feed blocks sites on its surface. How the promoters act is still argued. A 2025 study calls the mechanism of rhenium promotion a topic of debate and finds a clear understanding of its effect limited [1]. A 2024 study finds a systematic picture of the oxygen on working silver still lacking, and essential to understanding the promoters [18]. Progress would be a mechanistic model that predicts where the chloride optimum sits and how it moves as a charge ages, in place of the fitted curve of this chapter.

6.8 Discussion

Three results generalize past epoxidation.

The first is that a minority reaction can own the energy balance. Thirteen percent of the ethylene makes sixty-five percent of the heat, and any process with a large disparity in heats of reaction deserves the same arithmetic before its cooling is designed.

The second is that a fixed-sign controller cannot hold an optimum, and the failure is conditional rather than absolute. The same loop settles from one side and walks to a trip from the other, and the condition is invisible from inside the loop. Where an objective has a maximum the first question is not how to optimize it but whether there is a measurement that is zero there [19]. Volume 11, Chapter 13 holds such an optimum with an ordinary PI loop, because its analyzer measures a difference that is zero at the peak and signed on either side. Volume 7, Chapter 4 regulates at the peak of a friction curve, where the process gain also passes through zero and reverses. Where there is no such measurement, the shipped Loops structure uses a search, and the resolution of that search is the noise on the objective divided by the product of the curvature and the step.

The third is that two constraints conflicting on one handle is a design finding rather than a tuning problem. Cooling this reactor helps the catalyst and hurts the flammability margin, and no setting of the coolant valve satisfies both. The move that relieves both constraints is the one that gives up the objective, and recognizing that early is what separates an override structure that works from a selector that oscillates.

6.9 Conclusion

An ethylene oxide reactor at 87.2 % selectivity makes 65 % of its 40 MW of heat with the 12.8 % of the ethylene that burns, which is why it is nine thousand narrow tubes in boiling water. In the simulated reactor of this chapter the peak goes from 267.0 °C at the design coolant temperature to 290.1 °C four degrees higher and has no cool solution above 248.97 °C, a total margin of 4.2 °C. The recycle loop stabilizes a thirty-five percent activity step that would ignite a once-through tube, settling the peak back at 267.6 °C on 5.68 mol% oxygen, and it overshoots to 286.6 °C on the way because the tube answers in a minute and the loop turns over in eight. The chloride optimum near 2.5 ppm is flat in selectivity, worth less than half a point across a part per million either side, and steep in hot spot, worth 22.6 °C across the same range, and an identical PID controller creeps toward its setpoint from 1.50 ppm and trips the reactor from 3.30 ppm at 381.8 min. The economic MPC reaches a 269.783 °C peak versus 271.529 for the best tested fixed-chloride policy, with nearly tied mean margins of 2.338943 and 2.340381 and larger oxygen IAE, 76.098 versus 35.081 mol%·min. A fuller model would add a tube-to-tube distribution so that a runaway can start where a real one starts, a finer axial grid, a chloride model with the memory a real catalyst has, and a transport delay in the recycle loop. The reactor shows where feedback stops being the right tool.

KEY POINTS

- The reaction nobody wants owns the heat: the eighth of the ethylene that burns makes about two thirds of the reactor's heat.
- The cool steady state ends 4.2 °C above the design coolant, and the recycle loop keeps the reactor off that cliff, but late enough to overshoot to 286.6 °C.
- The chloride optimum near 2.5 ppm is flat in selectivity and steep in hot spot, and one PID settles from below it and trips the reactor from above it.
- Cooling helps the hot spot and hurts the flammability margin, so the override cuts the production rate.

Where else this applies. Holding a peak where the process gain changes sign is also the problem of Volume 7, Chapter 4 and Volume 11, Chapter 13, and a one-way override on a hard limit guards Volume 4, Chapter 2 and Volume 4, Chapter 6.

References

[1] C. J. Keijzer, P. T. Weide, K. H. Helfferich, J. Zieciak, M. de Ridder, R. Dalebout, T. L. Lohr, J. R. Lockemeyer, P. van den Brink, and P. E. de Jongh, "Insight into the influence of Re and Cl on Ag catalysts in ethylene epoxidation," *Catalysis Science & Technology*, vol. 15, no. 2, pp. 323–333, 2025.

[2] R. J. van Welsenaere and G. F. Froment, "Parametric sensitivity and runaway in fixed bed catalytic reactors," *Chemical Engineering Science*, vol. 25, no. 10, pp. 1503–1516, 1970.

[3] M. Morbidelli and A. Varma, "Parametric sensitivity and runaway in fixed-bed catalytic reactors," *Chemical Engineering Science*, vol. 41, no. 4, pp. 1063–1071, 1986.

[4] D. Ren, H. Xu, J. Li, J. Li, and D. Cheng, "Origin of enhanced ethylene oxide selectivity by cs-promoted silver catalyst," *Molecular Catalysis*, vol. 441, pp. 92–99, 2017.

[5] H. Falcke, S. Holbrook, I. Clenahan, A. López Carretero, T. Sanalan, T. Brinkmann, J. Roth, B. Zerger, S. Roudier, and L. Delgado Sancho, "Best available techniques (BAT) reference document for the production of large volume organic chemicals," Publications Office of the European Union, Luxembourg, EUR 28882 EN, 2017.

[6] Z. Nawaz, "Heterogeneous reactor modeling of an industrial multitubular packed-bed ethylene oxide reactor," *Chemical Engineering and Technology*, vol. 39, no. 10, pp. 1845–1857, 2016.

[7] A. Kummer and T. Varga, "What do we know already about reactor runaway? a review," *Process Safety and Environmental Protection*, vol. 147, pp. 460–476, 2021.

[8] A. Pekalski and H. Pasman, "Distinction between the upper explosion limit and the lower cool flame limit in determination of flammability limit at elevated conditions," *Process Safety and Environmental Protection*, vol. 87, no. 1, pp. 47–52, 2009.

[9] L. G. Britton and B. K. Harrison, "Using chetah to estimate lower flammable limit, minimum ignition energy, and other flammability parameters," *Process Safety Progress*, vol. 33, no. 4, pp. 314–328, 2014.

[10] 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.

[11] 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.

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

[13] 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.

[14] 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.

[15] K. Jasek, M. Pasternak, and M. Grabka, "Paramagnetic sensors for the determination of oxygen concentration in gas mixtures," *ACS Sensors*, vol. 7, no. 11, pp. 3228–3242, 2022.

[16] European Industrial Gases Association, "Oxygen pipeline and piping systems," European Industrial Gases Association, Brussels, EIGA Doc 13/20, 2020.

[17] 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.

[18] M. Guo, N. Dongfang, M. Iannuzzi, J. A. van Bokhoven, and L. Artiglia, "Structure and reactivity of active oxygen species on silver surfaces for ethylene epoxidation," *ACS Catalysis*, vol. 14, no. 13, pp. 10 234–10 244, 2024.

[19] 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.

6.10 Exercises

6.1 Concept check · 2 min

At 87.2 % selectivity the combustion makes 65 % of the heat. Why does the heat released in a cell of this tube rise faster with temperature than the epoxidation rate alone would suggest?

(a) The epoxidation releases more heat per mole as the catalyst warms

(b) The carbon dioxide made by the combustion inhibits both reactions less when hot

(c) The cycle gas carries less heat out of a cell as its oxygen is consumed

(d) The combustion has the larger activation temperature and 12.6 times the heat

6.2 Concept check · 2 min

Scenarios 7 and 8 run the same chloride PID with the same tuning and the same 87.15 % setpoint. Started at 3.30 ppm it trips the reactor at 381.8 min. Started at 1.50 ppm it creeps toward 2.1 ppm. What distinguishes the two runs?

(a) The sign of the gain from chloride to selectivity, which reverses near 2.5 ppm

(b) The integral time, which is too short for the chloride uptake from the higher start

(c) The hot spot override, which acts in scenario 8 and is switched off in scenario 7

(d) The analyzer cycle, which reports the selectivity too late when the chloride is high

6.3 Concept check · 2 min

The hot spot is above the start of its override. Why does the override cut the production setpoint instead of lowering the coolant temperature?

(a) The coolant drum answers too slowly to reach the hot spot before the alarm

(b) A cooler reactor leaves more unreacted oxygen and walks the loop toward flammability

(c) Lowering the coolant cuts production at once, and the override exists to keep it

(d) The coolant has no lasting effect on the hot spot once the recycle has settled

6.4 Theory and calculation · 6 min

The reactor makes 8.00 kmol/min of ethylene oxide at 87.21 % selectivity, with $\Delta H_1 = 105\,\mathrm{kJ/mol}$ and $\Delta H_2 = 1327\,\mathrm{kJ/mol}$ per mole of ethylene.

(a) Find the fraction of the heat made by the combustion. *Checkpoint:* 65.0 %.

(b) Find the ethylene burned in kmol/min and the total duty in MW. *Checkpoint:* 1.173 kmol/min, 39.95 MW.

(c) Use Equation (1.5) with $E_1/R = 8200\,\mathrm{K}$ and $E_2/R = 11\,800\,\mathrm{K}$ to find the effective activation of the heat generated. Taking each quantity to grow by $(E/R)/T^2$ per kelvin, compare the growth of the heat and of the epoxidation rate at the 267.1 °C hot spot. *Checkpoint:* 10 540 K, then 3.61 % and 2.81 % per kelvin.

6.5 Theory and calculation · 7 min

At the design point the oxygen feed is $F_O = 7.606\,\text{kmol/min}$, the purge $P = 1.20\,\text{kmol/min}$ and the cycle gas oxygen 7.20 mol%, with 8.00 kmol/min of ethylene oxide at 87.21 % selectivity. Each mole of ethylene oxide uses half a mole of oxygen and each mole of ethylene burned uses three.

(a) Use Equation (6.8) to find the oxygen reacted, and check it against the stoichiometry. Find the oxygen used per mole of product. *Checkpoint:* 7.520 kmol/min and 0.940.

(b) In scenario 2, with the oxygen feed fixed, two degrees of extra coolant bring the Selectivity tile to 86.92 % and the cycle gas oxygen to 6.80 mol%. Find the production the oxygen balance allows. *Checkpoint:* 7.91 kmol/min.

(c) Find the production gained per kmol/min of extra oxygen feed at constant selectivity, and the time the loop takes to turn over once (holdup 2450 kmol, circulation 320 kmol/min). Use them to explain the rise of production to 8.43 kmol/min within eight minutes of the coolant step and its fall afterward. *Checkpoint:* 1.064 and 7.66 min.

6.6 Theory and calculation · 6 min

The step and search optimizer of scenario 9 moves the chloride by $h = 0.3$ ppm and ignores selectivity differences under $\sigma = 0.04$ %. Near the maximum the selectivity curve has a curvature $|S''| = 0.7$ % per ppm^2.

(a) Find the distance from the top inside which a step cannot be told apart, $d = \sigma / (|S''|\,h)$. *Checkpoint:* 0.19 ppm.

(b) Repeat with 0.5 % of analyzer noise in place of σ. *Checkpoint:* 2.38 ppm.

(c) Taking $\Delta S = \tfrac{1}{2}|S''|d^2$, find what being 0.5 ppm from the top costs, and compare it with the measured hunt, which averages 87.16 % against 87.21 at the peak. Then find the step that would resolve 0.1 ppm at $\sigma = 0.04$ % and say what it would do to the hunt. *Checkpoint:* 0.09 points against 0.05, and 0.57 ppm.

6.7 App challenge · 8 min

Load scenario 5, "5 · The recycle is the real stabilizer", at the default seed. It starts with the Oxygen concentration loop in Manual and the Hot spot override Off, and the catalyst activity steps to 1.35 at 20 min. Run at speed 20 (20 min per wall-clock second) until the clock passes 240 min, pause and export the CSV. A run longer than an hour exports its time column as Time (h), in hours, so convert each time below from minutes.

(a) Find the largest `hot_C` and its time, and read `metric_alarm_min`, the minutes above the 278 °C alarm, in the row at 240.0 min. *Checkpoint:* 286.6 °C at 25.9 min, then 6.9 min.

(b) In the row at 240.0 min, read `hot_C`, `o2_pct` and `o2_cmd`, and say what stopped the tube. *Checkpoint:* 267.6 °C, 5.68 mol%, 7.592 kmol/min.

(c) Press Reset, set Hot spot override to On before Run, run again to 240 min and repeat (a). Explain why the override helps so little. *Checkpoint:* 283.2 °C at 24.9 min, then 4.6 min.

6.8 App challenge · 8 min

Load scenario 10, "10 · Two margins, opposite directions", at the default seed. The Production and Oxygen concentration loops start in Manual, and the coolant steps to 241.0 °C at 20 min and to 248.8 °C at 220 min. Run at speed 60 (60 min per wall-clock second) until the clock passes 480 min, pause and export the CSV. A run longer than an hour exports its time column as Time (h), in hours, so convert each time below from minutes.

(a) Read `hot_C` and `o2_margin_pct` in the rows at 220.0 and 480.0 min. *Checkpoint:* 262.8 and 0.52, then 271.6 and 2.18.

(b) Find the largest `hot_C` after 220 min and its time, and read `metric_alarm_min` at 480.0 min. *Checkpoint:* 286.4 °C at 231.3 min, then 10.4 min.

(c) Switch both loops to Auto and Hot spot override to On, run another hour, export again and count the minutes with `override_active` equal to 1 after the switch. Explain what moved each margin back. *Checkpoint:* 6 min.

6.9 App challenge · 10 min

Load the challenge, "Challenge · a shift on a fresh charge", at its default seed. For each run below, run at speed 60 (60 min per wall-clock second) until the clock passes 900 min, about 15 s, pause and export the CSV. A run longer than an hour exports its time column as Time (h), in hours, so convert each time below from minutes. The graded number is the worst $|\texttt{o2_pct} - \texttt{o2_sp_pct}|$ over the rows from 594 min, which the progress panel also reports as the worst oxygen error. Read `o2_margin_pct` in the row at 900.0 min.

(a) Run the case as loaded. *Checkpoint:* 0.314 mol%, margin 1.73 mol%.

(b) Press Reset, set Oxygen loop gain to 7.0 and Oxygen integral time to 18.0 min before Run, change nothing else, and run again. *Checkpoint:* 0.077 mol%, margin 1.56 mol%.

(c) Press the reload button beside the scenario list, press Reveal suggested tuning and run again. Say why the suggestion is worth having although its graded number is worse than in (b). *Checkpoint:* 0.105 mol%, margin 1.75 mol%.

6.10 Application area · 3 min

A plant engineer proposes running the carbonate wash harder, taking the cycle gas carbon dioxide from 5.00 to 4.00 mol% to gain catalyst activity. Show Model lists the coefficients of Equation (6.7), the limiting oxygen lost per mol% of ethylene and gained per mol% of carbon dioxide, as $a = 0.115$ and $b = 0.180$. Find the new limit of flammability and the margin at 7.20 mol% oxygen and 24.0 mol% ethylene, and the gain in rate from the inhibition term $\phi = 1/(1 + 4.6\,y_C)$. What does the oxygen loop do with the gain?

6.11 Application area · 4 min

An operations engineer wants to set a coolant temperature ceiling on a real reactor from the 248.97 °C at which the cool state of this tube ends. Apply Equation (1.13) with a coolant loop that swings 0.5 °C, no measurement error and 1 °C of reserve. Then name two simplifications of the model that make the real margin smaller than the simulator's 4.2 °C.

6.12 Review the agent · 5 min

An AI assistant reviewing scenario 7 writes: "The trip at 381.8 min is a tuning problem. A chloride PID gain of 1.20 ppm per percent is aggressive for a loop behind a 55 min uptake and a 30 min analyzer. Cut Chloride PID gain to 0.30 and the loop will settle at its 87.15 % setpoint from 3.30 ppm, as it does from 1.50 ppm in scenario 8." Find the flaw, and say how you would check it in the app.

CHAPTER 7

A Paraxylene Simulated Moving Bed Whose Purities Exist Only as Cycle Averages

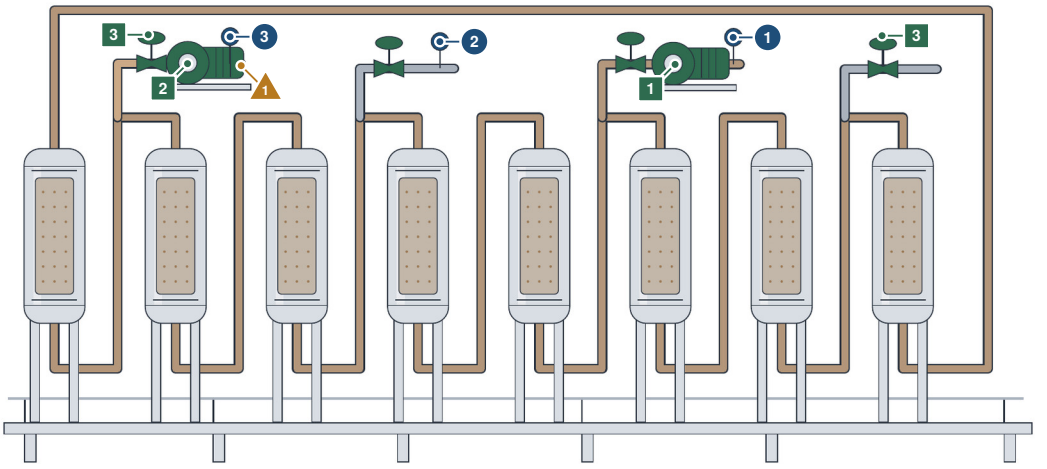

The controller moves 1 extract pump, for the purity of the extract 2 feed pump, for the paraxylene lost to the raffinate 3 port valves, stepped one bed every switch by the sequencer

The controller reads 1 extract purity analyzer, one report a cycle 2 raffinate purity analyzer, one report a cycle 3 feed analyzer, which the predictive controller reads

What disturbs it 1 a feed that turns richer in paraxylene ▲ an adsorbent that slowly loses capacity

Notice that the three analyzers report once a cycle from the streams entering and leaving the ring, while the port valves move the four streams one bed forward every switch, two beds apart in the chapter's model. Illustration, not to scale.

IN PLAIN TERMS A Purity That Exists Only as an Average

The main idea. A solid holds paraxylene, the raw material of polyester, a little more tightly than the look-alike liquids around it. The solid fills a ring of beds, and the inlets and outlets step one bed on at set times, as if the solid moved.

The challenge. So the liquid that leaves keeps changing, and its purity is known only as an average over a full turn. The flows can be set right only once a turn, while the feed changes and the solid wears out.

Why it matters. In the simulated plant a change made once a turn keeps the product pure, and no fixed setting does.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? box in Section 7.1, the History boxes in Section 7.2, the Try it boxes in Sections 7.4 and 7.6, the Did you know? box in Section 7.6, and Key points in Section 7.9. No equations.
Study	45 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 7.6 with Figure 7.5, Section 7.7, and Section 7.8 with its Open challenge box.

The control question A simulated moving bed has no steady state, and its purities exist only as twelve-minute averages reported once a cycle. Which flow-rate ratios keep both outlets on specification, and how much can a correction made once a cycle hold through a richer feed and an aging adsorbent?

Learning objectives

- Compute the zone flows, the port flows and the desorbent ratio of the ring from its flow-rate ratios and switching time.
- Construct the ideal triangle from the isotherm and explain why the simulated unit needs a margin from its edges.
- Explain the cyclic steady state and why a purity exists only as a cycle average.
- Compare cycle-to-cycle PIs on the logarithm of each impurity with a repetitive MPC through a feed step and an aging adsorbent.
- Name the pump that delivers off its meter from the pattern of the two purity analyzers.

Prerequisites Sections 1.4.7 (cyclic units and cycle-to-cycle control), 1.6.2 and 1.6.4. Section 1.6.3 only for the predictive controller. Chapter 5 runs a cycle-to-cycle loop on another cyclic unit.

The app https://apmonitor.com/apps/smb-paraxylene/

See also Volume 8, Chapter 11 for an adsorbent whose capacity is not measured, Volume 8, Chapter 9 for corrections made while results are still in flight, and Volume 8, Chapter 10 for learning from one layer to the next. Volume 9, Chapter 2 for a step-response predictive controller behind a dead time.

7.1 Introduction

No reading of the extract stream of a simulated moving bed is its purity. In the eight-bed ring simulated here, every 90 s the four ports that admit desorbent and feed and withdraw extract and raffinate move one bed forward, and the liquid leaving at the extract port comes from a different bed after each move. Ten cycles after a start-up, when every switch repeats the previous one shifted by one bed to better than one part in a million, the paraxylene content of the extract (desorbent-free) still runs from 99.40 to 99.98 % inside a single switch, while its average over the twelve-minute cycle is 99.89 %. Only that average is steady, only that average is sold, and only that average can be measured as a composite sample once a cycle.

DID YOU KNOW? Most of the Feed Comes Back

A classical paraxylene plant drawn in a 2010 study makes 250 thousand tonnes a year. The other isomers that leave its adsorber go to an isomerization reactor and return, and that recycle loop makes up the major part of the adsorber's feed [1].

Two beliefs that a block diagram leaves in place are the simulator's targets. The first is that a quality variable has a value at every instant that a controller can track. Here a purity exists only over a whole cycle, it arrives 2 min after the cycle ends, and a move of a flow shows fully only several cycles later. The second is that being inside a design region means being on specification. The triangle of equilibrium theory, which turns an isotherm into the two flow-rate ratios that decide both purities [2, 3], is drawn for an ideal unit with no dispersion and no mass-transfer resistance. With the dispersion of thirty mixing cells per bed, the left edge that protects the extract sits at $m_2 = 0.422$ on the triangle at $m_3 = 0.82$, while the simulated unit needs $m_2 \geq 0.568$ to reach 99.7 %. Processes that are corrected once per run, batch, lot or cycle from the measured outcome of the last form a class of their own, treated as run-to-run, iterative learning and repetitive control [4, 5], and a repetitive model predictive controller has been applied to a simulated moving bed [6]. Section 1.4.7 sets out the integral loop that such a unit closes once a report, and Volume 8, Chapter 10 learns in the same way from one layer of a part to the next.

However, a student rarely meets that class on a plant whose cycle can be watched. Models of industrial paraxylene units are written to study design, revamping and aging [7–9], controllers for simulated moving beds are reported on sugar and pharmaceutical separations [10–12], while the standard examples of run-to-run control come from batch reactors and semiconductor lots [4]. This chapter follows one simulated eight-bed paraxylene ring whose sequencer, cycle composites, triangle and purity controllers can be watched switch by switch. It asks what fixed ratios, cycle-to-cycle feedback and a repetitive predictive controller can hold through a richer feed, an aging adsorbent and a sixty-cycle challenge, and how much of the answer belongs to the bed's declared dispersion. The mixing-cell model, the triangle, the velocity-form PI law and the linear predictive controller are standard. What the chapter adds are measurements on one synthetic unit, among them that the challenge's verdicts hold at the shipped thirty cells per bed at five seeds and with the step halved, but change at 20, 45 and 60 cells. The model has 973 states, its paraxylene balance closes to 1.9×10^{-14} of the feed per cycle, its triangle matches an independent derivation to 2×10^{-16}, and its controllers read the analyzers and meters only. Nine graded cases judge both the learner's actions and the condition of the ring.

Section 7.2 follows the process from the patent to equilibrium theory. Section 7.3 sets out the ring as modeled: mixing cells, pumps and flow-rate ratios, cycle averages, the triangle and the instruments. Section 7.4 shows how the ring reaches its cyclic steady state, where the simulated unit meets its specification against the triangle, and which ratio moves which purity. Section 7.5 describes the fixed ratios, the cycle-to-cycle PIs and the repetitive MPC, and the design choices behind them. Section 7.6 reports the nine cases, Section 7.7 names the hardware and the practice behind the model, and Section 7.8 states what the model leaves out. The exercises at the end of the chapter run the cases in the app.

7.2 Moving the Ports Instead of the Solid

A true countercurrent of adsorbent and liquid would separate two components continuously, but it requires moving the solid. The arrangement that avoids it, told in the first History box below, keeps the beds still and advances the ports, so that relative to the ports the solid moves the other way. One of the first industrial applications was the recovery of paraxylene from mixed xylenes, and the units known publicly as UOP's Parex, IFP's Eluxyl and Toray's Aromax run in the liquid phase near 180 °C and 9 bar [8]. Paraxylene is the raw material of terephthalic acid, which is made into polyethylene terephthalate [1].

> ⧖ HISTORY **Fixed Beds With Moving Ports, 1957**
>
> A patent filed by Broughton and Gerhold in 1957 and issued in 1961 kept the sorbent still in a series of fixed beds and moved the inlets and outlets instead, shifting them downstream in equal steps so that the liquid met the sorbent as in a countercurrent. The patent aims the method at sorbents too fragile to be carried from zone to zone without losing their activity to attrition. To shift the ports it offers either a manifold of valves opened and closed by timed electric switches or a circular valve whose plug turns in a housing ringed with ports, two inlet passages alternating with two outlet passages [13]. In 1989 an article analyzed counter-current and simulated counter-current adsorption separation processes together [14].

Design moved from trial to equilibrium theory, as the second History box below tells, and its construction for nonlinear isotherms in simulated moving bed units [3] is the triangle in the plane of the zone II and zone III ratios that the simulator draws. Models of the industrial paraxylene unit followed [7], and a later study built one on measured xylene equilibria and kinetics on ion-exchanged faujasite as a tool to train operators and to choose operating conditions [8], while the loss of adsorbent capacity over a unit's life became a modeling question of its own [9]. A study of an industrial-scale unit states the specification the simulator adopts, paraxylene in the extract above

99.7 % with recoveries around 97 %, and models each bed with forty finite volumes to stand for its axial dispersion [1].

> ⧖ HISTORY **Drawing the Region in Advance, 1993**
>
> In 1993 a design procedure for the four-section unit was built in the frame of equilibrium theory, for a binary mixture with a desorbent of intermediate strength and a constant-selectivity isotherm, with mass-transfer resistance and axial mixing neglected. Requiring complete separation gave explicit bounds on the flow-rate ratios of the four sections, a region that can be drawn in advance from the adsorption constants and the feed composition. The same work analyzed the disturbances a unit meets, set criteria for a robust design and compared the theory with experiments on a six-port simulated moving bed run in the vapor phase [2]. In 1997 the construction was carried to nonlinear chromatographic separations [3].

Control came later than design. Model-based control of a simulated moving bed was reported for the separation of fructose and glucose [10], a repetitive model predictive controller that corrects the unit from one cycle to the next was applied to a simulated moving bed chromatography system [6], and automatic control of simulated moving beds was developed alongside [11]. The optimization and control of chromatography have been treated together [12]. Simulated moving beds serve the sugar, petrochemical and pharmaceutical industries, and their systematic optimization on the fully discretized cyclic model has been demonstrated [15]. Model predictive control in general is widely applied in the process industries [16]. The records read for this chapter do not describe the purity control of a commercial paraxylene unit, and the simulator's loops are offered for teaching, not as a description of how industrial units are run.

7.3 The Ring as Modeled

The unit is generic, taken from no licensor's design or operating plant, as shown in Figure 7.1, and its composition profile is unrolled in Figure 7.2. Eight beds of 20 m^3 and void fraction $\varepsilon = 0.40$ stand in a closed ring, numbered 1 to 8 in the direction the liquid flows. With the ring shift s $(0,\dots,7)$ the desorbent D enters at the inlet of bed s, the extract E leaves at the outlet of bed $s+1$, the feed F enters at the inlet of bed $s+4$ and the raffinate R leaves at the outlet of bed $s+5$ (bed numbers modulo 8), which makes four zones of two beds each. Zone I strips the paraxylene off the solid, zone II pushes the other C_8 back from the extract port, zone III holds the paraxylene back from the raffinate port and zone IV cleans the liquid that recirculates to zone I. Every switching time t_s the shift advances by one, and eight switches make a cycle. The engine's clock is the minute, flows are in m^3/s internally and reported in m^3/h, and compositions are liquid volume fractions. The feed is the C_8 cut with a paraxylene fraction x_F of 22 % at design, and the rest of the liquid is desorbent.

Table 7.1 gives each variable's role, its value at the design point, $m = (1.40, 0.62, 0.82, 0.30)$ at $t_s = 90$ s, and its range.

7.3.1 Mixing Cells, Mass Transfer and the Isotherm

Each bed is $N_c = 30$ mixing cells in series of volume $V_c = V/N_c$. In each cell the liquid fractions of paraxylene (A) and of the other C_8 aromatics lumped as one component (B) and their adsorbed amounts are states, as shown in Equation (7.1).

$$\begin{aligned} \varepsilon \frac{\mathrm{d}c_i}{\mathrm{d}t} &= \frac{Q_z}{V_c}\,(c_{i,\mathrm{in}} - c_i) - (1-\varepsilon)\frac{\mathrm{d}q_i}{\mathrm{d}t} \\ \frac{\mathrm{d}q_i}{\mathrm{d}t} &= k\,(q_i^{*} - q_i), \qquad q_i^{*} = \frac{N K_i c_i}{1 + K_A c_A + K_B c_B} \end{aligned} \tag{7.1}$$

where $i = A, B$, Q_z is the liquid flow of the zone the bed is in, $c_{i,\mathrm{in}}$ is the outlet of the upstream cell (or a port node), q is the adsorbed amount as the liquid volume it would make per volume of solid, k is the linear-driving-force coefficient and q^{*} the competitive Langmuir equilibrium with the desorbent's

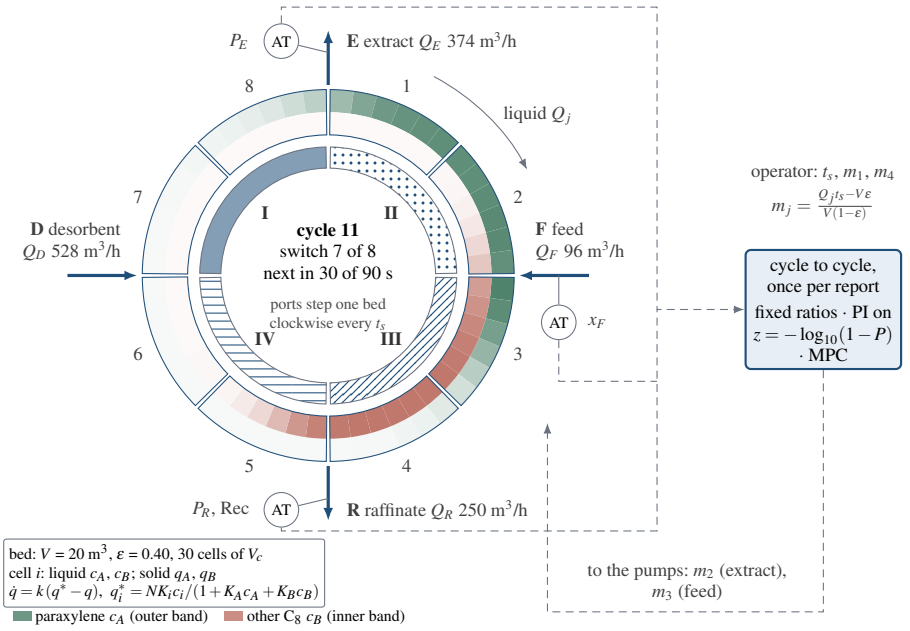

Figure 7.1 The eight-bed ring as modeled, redrawn from the Ring view of case 1 at $t = 130$ min, cycle 11, switch 7 of 8, with the unit at its cyclic steady state since cycle 10. Each bed's wedges are shaded by the paraxylene (outer band) and the other C_8 (inner band) in their liquid, the zones carry their own patterns, and the four ports sit where the sequencer places them at ring shift $s = 6$. The analyzers (AT) report cycle composites of the extract, the raffinate and the feed to the cycle-to-cycle layer, which sets m_2 and m_3. Not to scale.

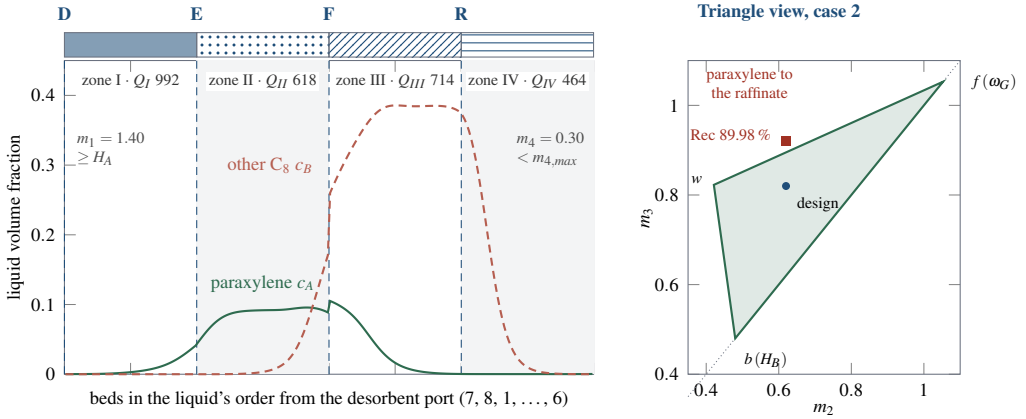

Figure 7.2 Left, the profile of the ring of Figure 7.1 unrolled from the desorbent port, paraxylene solid and other C_8 dashed, with the zone flows in m³/h. Right, the Triangle view of case 2 with m_3 held at 0.92 (square), above the upper edge wf, where the recovery over cycle 11 is 89.98 %. The dot is the design point. Left of the edge wb the other C_8 reaches the extract.

competition folded into the saturation capacity N and the affinities K_i. The Henry constants are $H_i = NK_i$, so $H_A = 1.20$, $H_B = 0.48$ and the selectivity K_A/K_B is 2.5 at design. Temperature and

Table 7.1 Variables, Their Roles and Design Values

Variable	Role	Design	Range
P_E	CV, extract analyzer	99.893 %	SP 99.00–99.95
P_R	CV, raffinate analyzer	99.868 %	SP 98.50–99.95
m_2	MV, extract pump	0.62	0.30–0.90
m_3	MV, feed pump	0.82	0.50–1.20
t_s	set, switching time	90 s	45–150
m_1, m_4	set, zones I and IV	1.40, 0.30	1.0–1.8, 0–0.6
x_F	DV, feed paraxylene	22 %	15–30
T	DV, adsorber temperature	177 °C	167–187
—	DV, aging rate	0	0–2 %/h
$c_{D,B}$	DV, other C_8 in desorbent	0	0–0.5 %
b	DV, one pump's meter bias	0	±5 %
σ	DV, analyzer noise	0.01 point	0–0.05
Rec, $\dot{m}_{PX}$	outputs	99.53 %, 18.10 t/h	—
c_A, c_B, q_A, q_B	states, per cell (960)	—	—
a	state, capacity (fresh = 1)	1	0.5–1

aging enter as shown in Equation (7.2).

$$\begin{aligned} K_i(T) &= K_{i,0}\exp\Big[-\frac{\Delta H_i}{R}\Big(\frac{1}{T}-\frac{1}{T_0}\Big)\Big] \\ k(T) &= k_0\exp\Big[-\frac{E_k}{R}\Big(\frac{1}{T}-\frac{1}{T_0}\Big)\Big] \\ N &= N_0\,a, \qquad \frac{\mathrm{d}a}{\mathrm{d}t} = -r_{age} \end{aligned} \tag{7.2}$$

where $T_0 = 177\,^\circ\mathrm{C}$, $\Delta H_A = -18\,\mathrm{kJ/mol}$, $\Delta H_B = -14\,\mathrm{kJ/mol}$, $E_k = 20\,\mathrm{kJ/mol}$, a is the capacity as a fraction of fresh and r_{age} the aging rate, held above $a = 0.5$. At 187 °C H_A falls to 1.081, the selectivity to 2.44 and k rises to $0.337\,\mathrm{s}^{-1}$, and at 167 °C they are 1.339, 2.56 and $0.266\,\mathrm{s}^{-1}$. The thirty cells stand for the bed's axial dispersion, the same device as the forty volumes per bed of an industrial model [1]. They are a declared physical choice rather than a converged grid, and the purities depend on them strongly. At the design point the extract purity is 97.18, 99.12, 99.89 and 99.98 % with 10, 15, 30 and 60 cells per bed.

7.3.2 Ports, Pumps and Flow-Rate Ratios

Four pumps realize the four zone flows. The desorbent, extract, feed and recycle pumps deliver Q_D, Q_E, Q_F and Q_{IV}, and the raffinate is what is left, as shown in Equation (7.3).

$$\begin{aligned} m_j &= \frac{Q_j\,t_s - V\varepsilon}{V(1-\varepsilon)}, \qquad j = I,\ldots,IV \\ Q_I &= Q_{IV} + Q_D, \qquad Q_{II} = Q_I - Q_E \\ Q_{III} &= Q_{II} + Q_F, \qquad Q_R = Q_{III} - Q_{IV} \\ Q_p &= (1+b_p)\,Q_{p,SP} \end{aligned} \tag{7.3}$$

where m_j is the ratio of zone j's liquid flow to the solid's apparent flow of one bed per switch, $Q_{p,SP}$ is the setpoint a pump's meter reads and b_p is a meter bias that makes one pump deliver more than its meter shows. The zone flows the beds see are rebuilt from what the pumps deliver, so a bias moves the true m while every meter still reads its setpoint. At design the zone flows are 992, 618, 714 and 464 m^3/h, with 528 m^3/h of desorbent and 96 m^3/h of feed in and 374 and 250 m^3/h of extract and raffinate out. At each port node the incoming streams mix, as shown in Equation (7.4).

$$\begin{aligned} c^{I}_{i,\mathrm{in}} &= \big(Q_{IV}\,c_i^{IV,out} + Q_D\,c_{D,i}\big)/Q_I \\ c^{III}_{i,\mathrm{in}} &= \big(Q_{II}\,c_i^{II,out} + Q_F\,c_{F,i}\big)/Q_{III} \end{aligned} \tag{7.4}$$

where $c_F = (x_F, 1-x_F)$ and $c_D = (0, c_{D,B})$. The extract is drawn at the outlet of zone I and the raffinate at the outlet of zone III, each at the composition of the last cell there. A switching time entered during a switch takes effect at the next one, so every switch is whole and the flows, which follow t_s, never change inside it.

7.3.3 Cycle Averages and the Cyclic Steady State

Purity, recovery and production are integrals over each whole cycle of eight switches, as shown in Equation (7.5).

$$\begin{aligned} P_E &= \frac{\int Q_E c_A \, dt}{\int Q_E (c_A + c_B)\, dt}, \qquad P_R = \frac{\int Q_R c_B \, dt}{\int Q_R (c_A + c_B)\, dt} \\ \mathrm{Rec} &= 1 - \frac{\int Q_R c_A \, dt}{\int (Q_F x_F + Q_D c_{D,A})\, dt} \\ \dot{m}_{PX} &= \rho_{PX} \int Q_E c_A \, dt \,/\, \Delta t_{cyc} \end{aligned} \tag{7.5}$$

where the compositions are those at the extract and raffinate ports, $\rho_{PX} = 861\,\mathrm{kg/m^3}$ and $\Delta t_{cyc} = 8 t_s$. The recovery counts the paraxylene the raffinate did not take. At the cyclic steady state it equals the extract's paraxylene over the feed's, while after a change it does not count what the ring takes up or gives back. The extract-based definition read 73.15 % in the cycle after case 3's feed step, purely from the ring filling, while the loss-based one read 98.43 %, which is why the app uses the latter. The unit is declared at its cyclic steady state when, at every switch of a whole cycle, the profile differs from the one a switch earlier shifted by one bed by less than 10^{-6} in any liquid fraction, as shown in Equation (7.6).

$$\max_{b,n,i} \left| c_i(b+1, n, t_k) - c_i(b, n, t_{k-1}) \right| < 10^{-6} \tag{7.6}$$

where b is the bed, n the cell and t_k the instant of switch k.

Pause. Try Exercise 7.2 before reading on.

7.3.4 The Triangle

For the competitive Langmuir isotherm and a feed of fractions c_A^F and c_B^F, an ideal unit with no dispersion and no mass-transfer resistance gives both outlets pure inside a triangle in the (m_2, m_3) plane [3]. Its vertices $b = (H_B, H_B)$ and $f = (\omega_G, \omega_G)$ lie on the diagonal and w above it, where $\omega_G > \omega_F$ are the roots of Equation (7.7), with $\kappa_A = K_A c_A^F$ and $\kappa_B = K_B c_B^F$.

$$(1+\kappa_A+\kappa_B)\,\omega^2 - \left[H_A(1+\kappa_B) + H_B(1+\kappa_A)\right]\omega + H_A H_B = 0 \tag{7.7}$$

The two edges through w, wb and wf, are the lines of Equation (7.8), in that order.

$$\begin{aligned} \left[H_A - H_B(1+\kappa_A)\right] m_2 + \kappa_A H_B m_3 &= H_B(H_A - H_B) \\ \left[H_A - \omega_G(1+\kappa_A)\right] m_2 + \kappa_A \omega_G m_3 &= \omega_G(H_A - \omega_G) \end{aligned} \tag{7.8}$$

Left of wb the other C_8 is not pushed back in zone II and reaches the extract, and above wf the paraxylene is not held in zone III and reaches the raffinate. Zone I needs $m_1 \geq H_A$, and zone IV needs $m_4 < m_{4,max} = \frac{1}{2}\{\Theta - (\Theta^2 - 4H_B m_3)^{1/2}\}$ with $\Theta = H_B + m_3 + \kappa_B(m_3 - m_2)$. At design $\omega_G = 1.05280$, the base runs from 0.480 to 1.053, the vertex is $w = (0.421, 0.823)$, the upper edge at $m_2 = 0.62$ is $m_3 = 0.895$, the left edge at $m_3 = 0.82$ is $m_2 = 0.422$ and $m_{4,max} = 0.416$.

The app's triangle was re-derived by hand, and the vertex found as the intersection of the two lines agrees with the app's closed form to 2.2×10^{-16}. The construction is used here only for the competitive Langmuir case its source names in its title [3] and in the standard form given above. In the linear limit w falls on (H_B, H_A). The simulation itself reproduces the edges when it is made nearly ideal. With 60 cells per bed and $k = 3\,\mathrm{s^{-1}}$ the raffinate impurity at $m_2 = 0.62$ jumps from 0.003 % at $m_3 = 0.890$ to 0.50 % at 0.900, bracketing the line's 0.895. Zone IV fails between $m_4 = 0.41$

Table 7.2 Parameters. A: Adjustable in the Parameters Panel. V: Fixed and Stated in the Model Dialog

Symbol	Meaning	Value	
V, ε	bed volume, void fraction	20 m^3, 0.40	V
N_c	mixing cells per bed	30	V
N_0	saturation capacity, fresh	1.2 (0.9–1.5)	A
$K_{A,0}$, $K_{B,0}$	affinities at T_0	1.0, 0.4	A
ΔH_A, ΔH_B	heats of adsorption	−18, −14 kJ/mol	V
T_0	reference temperature	177 °C	V
k_0	mass transfer at T_0	0.30 s^{-1} (0.2–0.4)	A
E_k	its activation energy	20 kJ/mol	V
ρ_{PX}	paraxylene density	861 kg/m^3	V
—	analyzer delay after a cycle	2 min	V
—	cyclic-steady-state tolerance	10^{-6}	V

(0.096 % extract impurity) and 0.42 (3.1 %) against the limit 0.416. The left edge is a desorption wave and is smeared by dispersion, but it converges as the bed sharpens. At $m_2 = 0.42$ the extract impurity is 0.97 % with 60 cells and $k = 3\,\mathrm{s}^{-1}$ and 0.22 % with 150 cells and $k = 8\,\mathrm{s}^{-1}$, against the line's 0.4216.

7.3.5 Instruments, Initialization and Numerics

Composite samplers on the extract, the raffinate and the feed fill over each whole cycle, and three analyzers report the extract's paraxylene purity, the raffinate's other-C_8 purity (both desorbent-free) and the feed's paraxylene 2 min after the cycle ends, with 0.01 point of noise, holding between reports. An independent trapezoid integral of the outlet streams over 20 cycles of case 3 under PI control agrees with the plant's cycle records to 6.7×10^{-5} points for the extract and 4.8×10^{-5} for the raffinate, and with zero noise each report equals its record exactly, 2.008 min after the cycle (the first step past the delay). The recovery estimate uses the analyzers and the metered flows, so a meter bias biases it. The four flow meters read the pumps' setpoints, and the adsorber temperature and the sequencer's state are known. Every case except the start-up begins on the cyclic steady state that its settings produce, solved at load by integrating the ring switch by switch with the engine's own step until the profile repeats to 10^{-11}. Integration is fourth-order Runge–Kutta with a step of 0.5 s, the controller runs every 5 s and a log row is written every 15 s, on the engine of Section 1.6.1. Halving the step moves the cyclic steady state's purities by less than 10^{-7} points. Twelve ledgers integrate every port stream inside the integrator, and the paraxylene balance closes to 1.9×10^{-14} of the feed per cycle at the cyclic steady state. The state vector holds 960 cell states, the capacity and the twelve ledgers, 973 in all. Table 7.2 lists the parameters.

7.3.6 Assumptions

Three pseudo-components stand for the liquid. Metaxylene, orthoxylene and ethylbenzene are one lump with one selectivity, and the desorbent is folded into the isotherm, so the triangle is the one for a non-adsorbing solvent while a real desorbent adsorbs and competes. The isotherm and mass-transfer constants serve teaching, set so that the eight-bed unit reaches 99.7 % at 90 s, and N is an effective capacity rather than a pore volume. Liquid volumes are additive and the liquid incompressible. The ports switch in one step, so the rotary valve's flushing and the line volumes are left out. The bed temperature is uniform and set, the heat of adsorption is not balanced, and aging lowers the capacity only, compressed from years into the run. Industrial units have more beds and a longer cycle than eight beds of 20 m^3 at 90 s. The specification of 99.7 % and the recovery target of 97 % are the figures of an industrial-scale study [1], and the limits on m_2, m_3 and the 0.05 feed-pump minimum are teaching values, because real ones come from the downstream customer, the licensor's design and the equipment.

The desorbent ratio is the clearest consequence of the teaching isotherm. At the point the PIs settle at the unit uses $(m_1 - m_4)/(m_3 - m_2) = 3.79$ cubic meters of desorbent per cubic meter of

feed. The industrial-scale study uses 264.14 against 220.15 t/h of feed, 1.20 by mass, and calls a desorbent consumption below 1.5 classical, with a paraxylene Langmuir constant of 5.865 m^3/kg against 1.777 m^3/kg for meta- and orthoxylene, a selectivity of 3.3 [1]. At a selectivity of 2.5 even the ideal unit at the vertex w, with $m_1 = H_A$ and m_4 at its limit of 0.375, needs $(1.20 - 0.375)/(0.823 - 0.421) = 2.05$, and the dispersion of eight short beds and the margins on m_1 and m_4 need the rest.

7.4 How the Ring Responds

7.4.1 From a Ring of Desorbent to the Cyclic Steady State

Case 1 starts the ring full of desorbent at the design point. The feed's components spread round it cycle by cycle, and the paraxylene produced rises from 1.8 t/h in cycle 1 to 12.7 in cycle 2 and 18.1 t/h by cycle 6. The extract over cycle 1 carries only 10.0 % of the paraxylene fed, 92.6 % by cycle 3, while the loss-based recovery already reads 99.98 % because almost none reaches the raffinate. The largest profile change at a switch falls from 1.0×10^{-1} in cycle 1 to 4.1×10^{-7} in cycle 10, by a factor of 2.7 to 3.8 per cycle from cycle 5, and the cyclic steady state is declared at the end of cycle 10, 120.00 min. Inside the first switch of cycle 11 the extract still runs from 99.40 to 99.98 % around its cycle average of 99.89 %. The approach is a ring that fills, not a time constant, and any change of a flow, the switching time or the adsorbent starts it again.

> **TRY IT IN THE APP** **1 · Start-up: From a Ring of Desorbent to the Cyclic Steady State (Manual)**
> Load the scenario, press Run and watch the Cycle badge and the status strip. The paraxylene produced climbs from 1.8 t/h in cycle 1 to 18.1 t/h by cycle 6, and CYCLIC STEADY STATE appears at the end of cycle 10, at 120.00 min. Even then the extract inside one switch runs from 99.40 to 99.98 %.

Pause. Try Exercise 7.1 before reading on.

7.4.2 The Triangle Against the Simulated Unit

The triangle locates the edges, but not where the simulated unit meets its specification, as shown in Figure 7.3. Each curve is found by bisection on the cyclic steady state at fixed $m_1 = 1.40$, $m_4 = 0.30$ and $t_s = 90$ s. With the shipped thirty cells the recovery reaches 97 % at $m_3 = 0.8828$ for $m_2 = 0.62$, below the ideal upper edge at 0.8951, and the extract reaches 99.7 % only for $m_2 \geq 0.568$ at $m_3 = 0.82$, far to the right of the ideal left edge at 0.422. The specification therefore holds in a narrow wedge near the triangle's upper edge, and its corner lies close to the point the PIs settle at, $(0.564, 0.854)$, where the unit makes 25.81 t/h at a recovery of 97.87 %. With 60 cells the recovery edge moves up to 0.8945 at $m_2 = 0.62$, nearly on the ideal line, and the extract edge to 0.518, while with 15 cells the extract edge needs $m_2 \geq 0.74$ at $m_3 = 0.82$, and with 10 cells at $m_1 = 1.40$ and $m_4 = 0.30$ no point reaches 99.7 % at all. The case 2 point above the upper edge, $(0.62, 0.92)$, holds the extract at 99.938 % while the recovery settles at 89.98 %, as shown in the right panel of Figure 7.2.

7.4.3 Which Ratio Moves Which Purity

Each ratio sits on the edge that protects one outlet, so the extract purity answers mostly to m_2 and the raffinate purity mostly to m_3. From steps of $+0.01$ in the cyclic steady state, measured in decades of impurity, the relative gain of the pairing of extract with m_2 is 0.872 at the design point and 0.914 at the PIs' point. Central differences give 0.857 and 0.907, and step tests on the running engine read ten cycles later give 0.882 and 0.927 in percent. The crossed pairing would have relative gains near 0.13 and 0.09, and each of its loops would see its gain change seven- to elevenfold as the other closed. The raffinate purity's gain to m_3 changes 8.4-fold in percent between the design point and the PIs' point (-2.5 to -21.0 points per unit of m_3 by central differences) but only 1.9-fold in decades of impurity (-8.1 to -15.2), which is why both loops act on $z = -\log_{10}(1 - P)$.

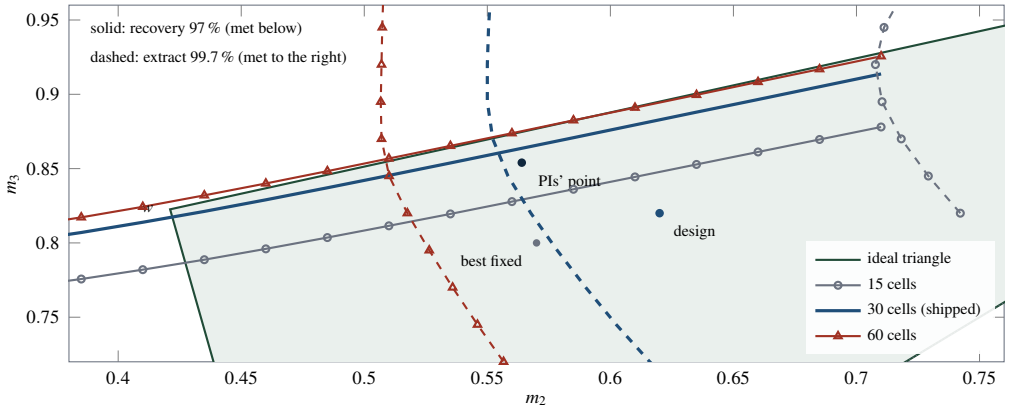

Figure 7.3 Where the simulated unit meets its specification, at three cell counts per bed, against the ideal triangle. Solid curves, recovery of 97 %, met below each curve. Dashed curves, extract purity of 99.7 %, met to the right. Circles mark 15 cells, triangles 60 cells, and the thickest curves the shipped thirty. Cyclic steady states at $m_1 = 1.40$, $m_4 = 0.30$, $t_s = 90$ s, fresh adsorbent and a 22 % feed. The shipped thirty cells leave a narrow wedge whose corner lies near the PIs' point.

7.5 Purity Control as Implemented

One Manual and Auto switch governs the ring. In Manual, m_2 and m_3 are the operator's stations, and m_1, m_4 and t_s are the operator's in every mode. In Auto the *Purity control* setting chooses *Fixed ratios*, which holds the stations where the operator put them and reads no analyzer, the *cycle-to-cycle PI*, or the *repetitive MPC*. Each controller execution builds a measurement packet from the analyzers, the meters, the sequencer and the temperature transmitter, and a test of the app fails if the packet carries anything else. Whatever the architecture, m_2 is clamped to 0.30–0.90, m_3 to 0.50–1.20, and m_3 is held at least 0.05 above m_2, the feed pump's minimum. The PIs track whatever is applied when they are not in charge, so a change of architecture or a return to Auto is bumpless.

7.5.1 The Cycle-to-Cycle PIs

Both loops are the velocity-form PID routine of Section 1.6.2 with the derivative set to zero, called only when a report arrives, as shown in Equation (7.9). The regeneration loop of Chapter 5 has the same structure.

$$
\begin{aligned}
e_k &= \pm\left(z_{SP} - z_k\right), \qquad z = -\log_{10}(1-P) \\
S_k &= S_{k-1} + K_c\left(e_k - e_{k-1}\right), \qquad \hat{u}_k = u_0 + S_k + I_{k-1} \\
u_k &= \min\left(u_{max}, \max\left(u_{min}, \hat{u}_k\right)\right) \\
I_k &= I_{k-1} + \left[\left(K_c/\tau_I\right) e_k + \left(u_k - \hat{u}_k\right)/\tau_I\right]\Delta t_k
\end{aligned}
\tag{7.9}
$$

where u is m_2 for the extract loop (direct acting, +) and m_3 for the raffinate loop (reverse acting, −), Δt_k is the time since the last report, one cycle in normal running, and S accumulates the proportional moves without clamping. The anti-windup is the back-calculation of Equation (1.17) with the tracking time τ_I, and a move cut by the feed-pump minimum is tracked as well. The defaults are $K_c = 0.06$ per decade and $\tau_I = 24$ min, two cycles, for both loops, with setpoints of 99.75 and 99.40 %. The integral used at report k was accumulated through report $k-1$, so its contribution lags the error that causes it by one report. Section 1.4.7 treats such a loop on a plant that answers reports late. An integral time at its slider's maximum of 600 min makes a loop proportional only.

7.5.2 The Repetitive MPC

The MPC plans m_2 and m_3 at every report over $N = 9$ cycles in three move blocks of 1, 1 and 7 cycles, on finite step responses of both purities in decades, as shown in Equation (7.10).

$$\begin{aligned} z_k &= \beta + \textstyle\sum_{j=1}^{9} \left[h_j^{(2)} \tilde{m}_{2,k-j} + h_j^{(3)} \tilde{m}_{3,k-j} \right. \\ &\quad \left. + h_j^{(F)} \tilde{x}_{F,k-j} \right] \\ \min\ & \textstyle\sum_k 20\, w_P \,(m_{2,k} - m_{3,k}) + \sum_j \frac{w_{\Delta u}}{0.01^2} \|\Delta m_j\|^2 \\ &\quad + \textstyle\sum_{o \in \{E,R\}} (400\, s_o^2 + 40\, s_o) \\ \text{s.t. } & z_{o,k} \geq z_{o,SP} - s_o, \quad s_o \geq 0, \quad m \text{ within its limits} \end{aligned} \tag{7.10}$$

where a tilde marks the deviation from the value in force over the last cycle, Δm_j is the move of block j, h_j are the increments of step responses identified on this model at commissioning with steps of $+0.01$ in m_2 and m_3 and $+2$ points of feed paraxylene from the PIs' point, x_F is the feed analyzer's report held over the horizon as a measured disturbance, β is the offset between the latest report and the model's prediction of it, renewed at every report, and $w_P = w_{\Delta u} = 1$. The final gains are 9.35 decades per unit of m_2 for the extract and -15.26 per unit of m_3 for the raffinate. Re-identified on the present model, the step responses reproduce the stored ones to four decimals for m_2 and m_3 and to 10^{-4} for the feed. The linear cost pushes the feed $m_3 - m_2$ up until both purities sit on their setpoints, which the soft constraints treat as floors. The plan is computed as a quadratic program by the shared linear MPC code of Section 1.6.3, with the dual algorithm of Goldfarb and Idnani [17], and a plan that comes back infeasible, not finite or out of range hands m_2 and m_3 to the tracked PIs without a bump. The model is not re-identified as the plant ages or warms, and those changes are left to β, as the step-response controller of Volume 9, Chapter 2 leaves its model error to a bias. The feed history starts from the feed analyzer's first report, which fills it, and not from the simulated feed composition.

7.5.3 Design Choices

Ten decisions taken while the app was built depart from its plan or settle what the plan left open, and each was measured. The beds carry 30 cells instead of 10, because ten reach 99.7 % at a 97 % recovery only with m_1 at its 1.8 limit and a feed of at most $m_3 - m_2 = 0.09$, 45 % of the design feed. The pairing is extract with m_2 and raffinate with m_3 rather than the crossed pairing the plan proposed, on the relative gains above, and the loops act on the logarithm of the impurity. The recovery is loss based, for the accumulation reason given with Equation (7.5). The feed step is 22 to 30 % paraxylene, because 26 moves the PIs' point only to 97.34 % recovery at its cyclic steady state. Case 4 teaches a production peak in the switching time rather than a region that closes, because the region does not close inside the stable range of $t_s \geq 45$ s and $m_1 \leq 1.8$. The honest infeasible case is the dirty desorbent of case 8. The MPC is identified once with an offset updated every cycle. The challenge adds a 30 t/h production floor, because without it a point far inside the triangle passes by making less. The sequencer, the composite samplers and the analyzers are written for this app, since the series' shared stage sequencer runs a bank of on and off stages against a demand and its analyzer samples an instant, while here the whole ring moves on a clock and the analysis is a cycle composite. Fixed ratios exists so that the challenge can ship in Auto with the suggested setting a control value.

The start-up logs nothing it has not measured. Before the first whole cycle and the first analyzer report at 14.01 min, the cycle averages, the recovery estimate, the production, the feed analysis and the profile change are logged blank, so their trends begin with a gap. The two loop measurements are the exception. They repeat their setpoints until the first report, because the series' loop metrics need a finite measured value, and the logged report count of zero marks those rows. The plant (the ring of Section 7.3), the gains and delays a student fits to the cycle records, and the MPC's step responses, which know neither the cell count nor the capacity, are three different models.

Table 7.3 The Nine Cases

#	Case and events	Teaching point
1	Manual, start from a ring of desorbent	the cyclic steady state
2	Manual at design, m_3 to 0.92, then m_2 to 0.44	which edge protects which outlet
3	Manual at the PIs' point, feed 22 to 30 % at 36 min	cycle-to-cycle feedback
4	PIs on, t_s 90 to 60, then 45 s	production against purity margin
5	Manual, capacity −5 % at 24 min	the triangle shrinks
6	Manual, extract pump +3 % off its meter at 24 min	a fault's signature
7	Manual, adsorber +10 K at 24 min, later 167 °C	temperature moves the region
8	PIs on, m_1 to 1.30, m_4 to 0.34, dirty desorbent	desorbent and infeasibility
9	Challenge, 60 cycles, aging, feed step at cycle 20	the whole problem

Table 7.4 Case Outcomes Measured on the Engine

Case and condition	Result
1, start-up	1.8 / 12.7 / 18.1 t/h in cycles 1, 2, 6; steady at 120.00 min
2, m_3 0.92	Rec 96.84 (cycle 3), 90.07 % (cycle 11); 24.5 t/h
2, m_2 0.44	P_E 99.47 to 96.04 %
3, Fixed	Rec 98.43 (cycle 4), below 97 % from cycle 8, 96.21 %
3, PI / MPC at 98 min	cycles 8–11 / 8–9 below 97 %; 36.1 t/h
3, PI / MPC / blind MPC from $t = 0$	lowest Rec after the step 97.63 / 98.41 / 97.90 %
4, t_s 120 / 90 / 75 / 60 / 45 s	21.2 / 25.8 / 28.2 / 29.0 / 18.9 t/h
5, aged, Manual / PI	Rec 93.20 % (cycle 10) / inside from 63.0 min
6, biased, Manual / PI	P_E 99.489, P_R 99.65 % / $P_E \geq 99.7$ at 98 min
7, +10 K / 167 °C	Rec 87.65 % / P_E 98.95 %
8, m_1 1.30 / m_4 0.34	D/F 3.59, 24.96 t/h / D/F 3.73, 22.91 t/h
8, dirty desorbent	P_E 95.97 % (cycle 36), m_2 0.899, 4.35 t/h

7.6 Nine Cases on the Running Ring

The cases are listed in Table 7.3 and the principal measurements in Table 7.4. All values come from the app's own engine, stepped outside the browser at the browser's 0.5 s step and default seed, from the case as loaded. Settings a case asks for after the first cycle are entered at 13 min, and every number in the app's own texts was regenerated twice by script with identical output. Cycle values are the true cycle averages of Equation (7.5), and report times are those of the analyzers.

DID YOU KNOW? A Shorter Switch Can Make Less

From the simulator: with the PIs holding both purities, the ring makes 25.8 t/h of paraxylene at a switching time of 90 s and 29.0 t/h at 60 s, but only 18.9 t/h at 45 s, less than the 21.2 t/h at 120 s (Table 7.4).

7.6.1 Crossing the Edges

Case 2 runs at the design point and raises m_3 from 0.82 to 0.92 after the first cycle, above the upper edge. The extra feed raises the paraxylene produced from 18.1 to 24.5 t/h, but zone III no longer holds the paraxylene, the recovery falls to 96.84 % in cycle 3 and 90.07 % by cycle 11, and the raffinate purity to 97.28 %. Back at $m_3 = 0.82$ with m_2 lowered to 0.44, the extract purity falls from 99.47 to 96.04 %, although the point is still inside the ideal triangle, whose edge there is at 0.422. Case 7 moves the region instead of the point. Ten kelvin hotter, the vertex moves from $(0.421, 0.823)$ to $(0.392, 0.760)$, the point is left above the upper edge and the recovery falls to 87.65 % by cycle 10 while the extract stays at 99.872 %. At 167 °C the triangle grows past the point the other way, $w = (0.454, 0.892)$, the recovery returns to 98.22 % and the extract falls to 98.95 %, because zone II at the same m_2 no longer pushes back the more strongly held other C_8.

7.6.2 Feedback Once a Cycle

Case 3 enriches the feed from 22 to 30 % paraxylene at 36 min with the flows fixed at the PIs' point. The recovery first rises to 98.43 % in cycle 4 as the ring takes up the extra paraxylene, then the report at 98.0 min reads below 97 % for cycle 8 and the recovery falls to 96.21 % by cycle 13 (96.10 % at the new cyclic steady state), while the extract becomes purer, 99.915 %. Handed the ratios at that report, the PIs lower m_3 to 0.821 and m_2 to 0.526, cycles 8 to 11 stay below 97 %, both analyzers hold within 0.05 and 0.1 point of their setpoints from 194 min, and the richer feed yields 36.1 t/h. The MPC switched at the same report leaves only cycles 8 and 9 below 97 % and settles at the same time. Running from the start, the MPC never lets the recovery below 98.41 % after the feed step, against 97.63 % for the PIs. The same MPC with its feed analyzer frozen holds 97.90 %, so about two thirds of its advantage comes from the measured feed and one third from its plan.

TRY IT IN THE APP 3 · A Richer Feed and the Cycle-to-Cycle PI

Press Run and watch the Cycle view through the feed change at 36 min. When the report at 98 min shows the recovery below 97 %, switch to Auto with Cycle-to-cycle PI chosen under Purity control. The PIs lower m_3 to 0.821 and m_2 to 0.526, both analyzers hold close to their setpoints from 194 min, and the richer feed yields 36.1 t/h.

Case 5 removes 5 % of the capacity at 24 min. A resin whose capacity nobody measures is the subject of Volume 8, Chapter 11. Every Henry constant falls with it, H_A from 1.200 to 1.140, and the vertex moves to $(0.400, 0.781)$. The point leaves the region at 24.01 min and the recovery falls to 93.20 % by cycle 10. Switched to the PIs at 48 min, m_3 falls to 0.817 and m_2 to 0.547, the point is inside again from 63.0 min, the recovery is above 97 % from cycle 9, and the aged unit makes 24.0 t/h against 25.8 fresh. Case 6 makes the extract pump deliver 3 % more than its meter reads, 12.0 m^3/h, which lowers the true m_2 and m_3 together by 0.025. The extract purity falls to 99.489 % while the raffinate purity rises to 99.65 %, a pair that a recycle pump delivering 3 % too much (all four m up by 0.029) would reverse. In Auto the PIs command $m_2 = 0.590$ and $m_3 = 0.879$, so the beds see 0.565 and 0.855, and the extract analyzer reads 99.7 % again at 98 min. The fault is now hidden in the setpoints, and the recovery estimate from the meters reads 97.78 % against a true 97.88 %.

Pause. Try Exercise 7.3 before reading on.

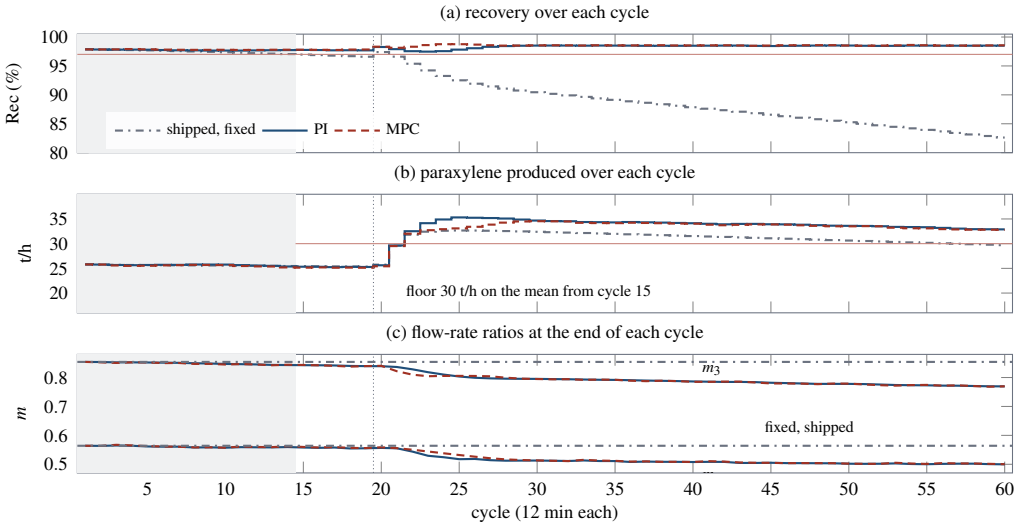

Figure 7.4 The challenge cycle by cycle. True cycle averages for the shipped fixed ratios (dash-dotted), the suggested PIs (solid) and the MPC (dashed), with the adsorbent aging 0.6 %/h throughout and a 30 % feed from cycle 20 (dotted). The shaded cycles 1 to 14 are not judged. The fixed point loses its recovery to aging while the PIs and the MPC walk m_2 and m_3 down.

7.6.3 Switching Time and Desorbent

At a fixed m every flow scales with $1/t_s$, so a shorter switch passes more feed per hour. With the PIs holding both purities, case 4 at 60 s first pushes the paraxylene to 37.4 t/h in cycle 3 while the extract falls to 99.38 %, the extract PI raises m_2 to 0.636, and the unit settles at 29.0 t/h against 25.8 t/h at 90 s. The switching time of 75, 120 and 45 s gives 28.2, 21.2 and 18.9 t/h. At 45 s the extract PI needs $m_2 = 0.744$, which leaves a feed of $m_3 - m_2 = 0.106$ and a desorbent ratio of 10.42. The triangle does not move with t_s, because equilibrium theory has no mass-transfer rate in it. What shrinks is the wedge of Figure 7.3, and the extract PI finds its edge.

Case 8 cuts the desorbent with the PIs on. Lowering m_1 to 1.30 dips the recovery to 97.43 % in cycle 4, the raffinate PI trims m_3 to 0.843, and the desorbent ratio falls from 3.79 to 3.59 for 24.96 t/h. Raising m_4 to 0.34 lets the other C_8 through zone IV, the extract PI raises m_2 to 0.594, and the ratio returns to 3.73 for 22.91 t/h. The desorbent then returns from its recovery column carrying 0.1 % of the other C_8, which enters zone I upstream of the extract port. The extract falls to 98.58 % within two cycles and 95.97 % by cycle 36 while the extract PI drives m_2 to 0.899, against its 0.90 limit, and the production to 4.35 t/h. No setting of the five handles can restore it. A grid of 1500 points followed by a pattern search over m_1 to m_4 and t_s from 45 to 150 s finds at best 99.284 % (at a recovery of 58.6 %) and 98.973 % with the recovery at 97 % or more, while the same best point with a clean desorbent exceeds 99.7 %. The cure is the desorbent column.

7.6.4 Sixty Cycles on an Aging Adsorbent

The challenge runs sixty cycles of 12 min. The adsorbent loses 0.6 % of its capacity every hour from the start and the feed changes to 30 % paraxylene at 228 min, the start of cycle 20. It is met when every cycle from cycle 15 has a true extract purity of at least 99.7 % and a recovery of at least 97 % and the paraxylene produced over those cycles averages at least 30 t/h. The configurations are compared in Table 7.5 and the cycles in Figure 7.4. As shipped, fixed ratios at the PIs' fresh point lose the recovery to aging from cycle 15, down to 82.62 % by cycle 60. The PIs pass with a lowest extract purity of 99.745 % (cycle 51), a lowest recovery of 97.47 % (cycle 23, just after the feed change) and

Table 7.5 Challenge Configurations Scored Over Cycles 15 to 60

Configuration	P_E min (%)	Rec min (%)	Mean (t/h)	Verdict
shipped, fixed 0.564 / 0.854	99.798	82.62	30.55	fails
fixed at design 0.62 / 0.82	99.912	97.57	23.39	fails
best fixed 0.57 / 0.80	99.735	97.95	27.00	fails
PI at half gain (0.03)	99.756	96.84	32.72	fails
suggested, PI	99.745	97.47	32.78	passes
MPC	99.745	97.77	32.44	passes
20 cells, PI / MPC	99.732 / 99.740	97.80 / 97.77	25.03 / 24.49	fail
45 cells, PI / MPC	99.746 / 99.746	97.17 / 97.77	36.36 / 36.11	pass
45 cells, fixed 0.54 / 0.80	—	—	30.48	passes
60 cells, PI	99.746	96.9976	37.87	fails
60 cells, MPC	99.745	97.77	37.65	passes
60 cells, fixed 0.52 / 0.79	—	—	31.83	passes

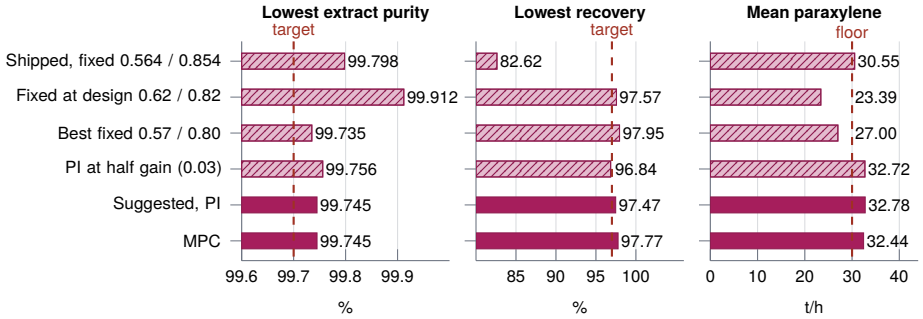

Figure 7.5 The configurations of the upper block of Table 7.5, at the shipped thirty cells per bed, against the three targets of the challenge: the lowest extract purity and the lowest recovery over cycles 15 to 60 and the mean paraxylene made. Solid bars pass, hatched bars fail, and the dashed line in each panel is the target. Notice that every fixed point fails on one target, the shipped point on the recovery and the two others on production, and that the PI at half gain misses the recovery by a fraction of a point. The lower block, with other cell counts, is not drawn.

32.78 t/h, walking m_3 from 0.854 to 0.769 and m_2 from 0.564 to 0.500. The MPC passes with 99.745, 97.77 and 32.44 t/h, holding the recovery at 98.14 to 98.76 % through cycles 21 to 28 against 97.47 to 98.44 % for the PIs, and solves 59 plans with none rejected. No fixed point on a 0.01 grid passes. The best, $(0.57, 0.80)$, holds both purities for all sixty cycles with only 27.00 t/h. At five seeds (default and 0 to 3) the PIs' lowest recovery lies between 97.34 and 97.47 %, always in cycle 23, and the MPC's between 97.75 and 97.77 %, the shipped run fails identically, and halving the step changes no verdict. Figure 7.5 sets the configurations of the table's upper block against the three targets.

The lower block of Table 7.5 repeats the challenge with the cells per bed changed in a copy of the source. At 20 cells nothing reaches the 30 t/h floor. At 45 cells the PIs and the MPC pass and so does the fixed point $(0.54, 0.80)$. At 60 cells two fixed points pass, the best making 31.83 t/h, the MPC passes, and the PIs as tuned miss the recovery by 0.0024 points in cycle 23, identically at half the step. The lesson that no fixed point passes, and the margin by which the PIs pass, therefore belong to the shipped dispersion. The ring at the instant Figure 7.1 redraws is shown in Figure 7.6.

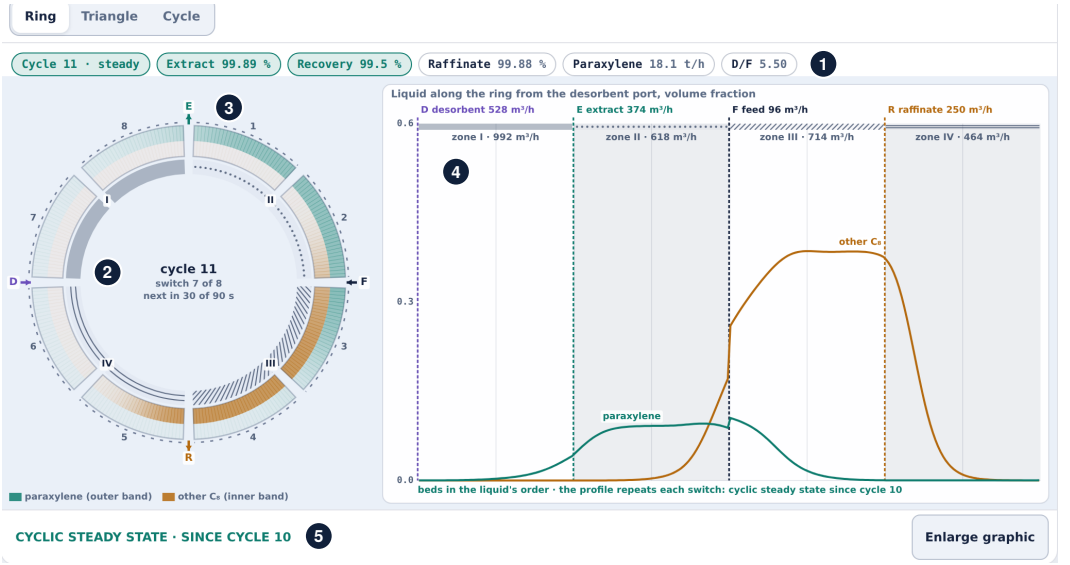

Figure 7.6 Case 1 at $t = 130$ min on the Ring view, in Manual at the design point. (1) The badges give the last cycle's extract purity of 99.89 %, recovery of 99.5 %, raffinate purity of 99.88 %, 18.1 t/h of paraxylene and a desorbent ratio of 5.50. (2) The ring's center reads cycle 11, switch 7 of 8, next in 30 of 90 s. (3) The four ports D, E, F and R sit where the sequencer has put them, with the zones I to IV between them. (4) The profile carries the zone flows of 992, 618, 714 and 464 m^3/h and the port flows of 528, 374, 96 and 250 m^3/h. (5) The status strip reads CYCLIC STEADY STATE SINCE CYCLE 10.

7.7 Technology in Practice

What Measures It

Figure 7.7 draws the ring of Section 7.3 with the hardware of a simulated moving bed in place of the schematic's symbols and the control system above it. The purities that the controllers correct come from analyzers behind composite samplers that fill over each cycle, and the model gives them a report 2 min after the cycle with 0.01 point of noise. Plant analyzers such as gas chromatographs and near-infrared spectrometers suffer from large measurement delays and high investment and maintenance costs [18], so the model's analyzers are faster and cleaner than a plant's, as Section 7.8 says. The four flow meters read the pumps' setpoints, and the meter bias of case 6 is the one fault the model gives them. An analyzer that reads the liquid inside the adsorber, rather than a composite of what has left it, is the subject of the box below.

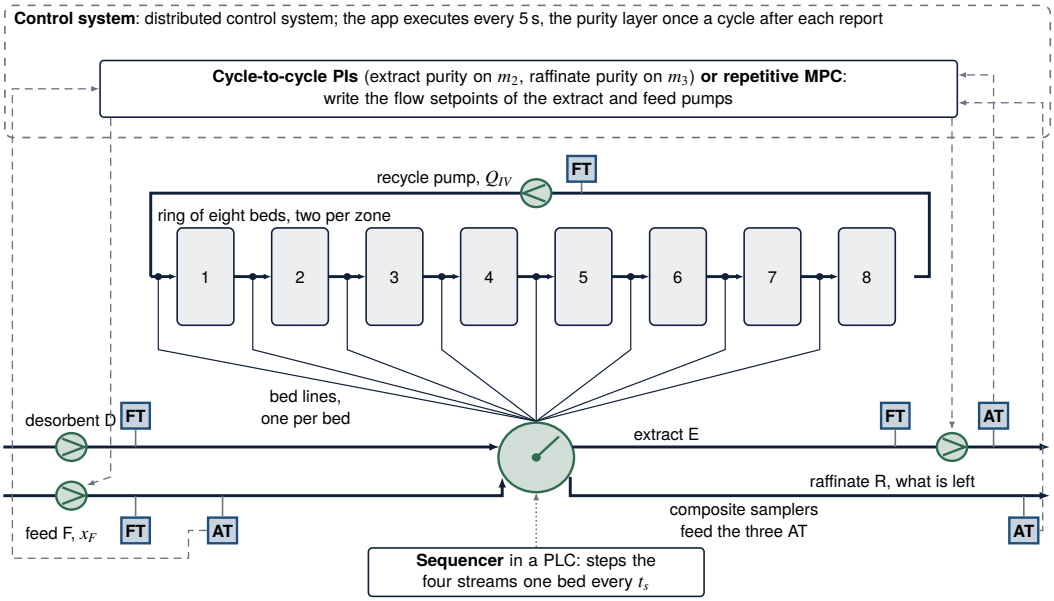

Figure 7.7 The ring's loops with their hardware. Blue boxes are sensing devices: the flow transmitters FT on the pumps and the three analyzers AT behind their composite samplers. Green symbols are final elements: the desorbent, feed, extract and recycle pumps and the rotary valve, which joins every bed line to the four streams. Dashed arrows are signals, the dotted one the sequencer's step command. Notice that the purity layer acts once a cycle, and only through the flow setpoints of the extract and feed pumps. The desorbent and recycle flows, which set m_1 and m_4, and the switching time are the operator's. Not to scale, and not a piping diagram.

TECHNOLOGY **Sensor: A Raman Analyzer on the Adsorber**

Raman scattering shifts the frequency of a small part of a monochromatic beam through coupling with the vibrations of the molecules it meets. The effect is so weak that it came into wide use only after powerful lasers arrived in the 1960s [19]. One paraxylene process design monitors the concentration profiles in its adsorber by on-line Raman spectroscopy for advanced dynamic control [20]. An optical analyzer can fail quietly. Its window must be kept clean of dust and condensation, and an analyzer fed stale liquid from a plugged sample line has no way of knowing it and reports a value that no longer follows the process [21].

What Moves It

Two kinds of final element act on the ring. Four pumps set the zone flows, and in the model each delivers its setpoint times one plus its meter bias, with the raffinate taking what is left. Only the extract and feed pumps belong to the purity loops, through m_2 and m_3. The ports move through the sequencer and its valve, not through a controller, at the switching time the operator sets.

TECHNOLOGY **Final Element: The Rotary Valve That Moves the Ports**

A typical industrial unit has 24 beds in two chambers of 12, each bed joined to one rotary valve by its own bed line, and pumps circulate the liquid between the chambers so that they act as a single unit [20]. At any moment only four bed lines carry the feed, the desorbent, the extract and the raffinate, and their positions are switched as the composition profile moves down the chamber [22]. A line that has just carried one stream holds its residue when it carries the next, which can contaminate the product. Designs answer with extra streams that flush the lines or with separate lines and individual on/off valves for each stream [23].

Where the Controller Runs

A distributed control system holds loops like these at a supervisory level that sends set points to field controllers, with loops nested or cascading and cycle times from milliseconds to minutes [24]. Programmable logic controllers (PLCs) serve such a system as local controllers within the supervisory scheme, programmed for logic, timing, counting and PID control [24]. The sequencer that steps the ports is a state machine, whose output depends on its position in the ring as much as on its inputs, and IEC 61131-3 names the Sequential Function Chart among its five languages for writing one [21]. In the app the controller executes every 5 s, while the cycle-to-cycle PIs are velocity-form PID blocks called only when a report arrives, once a 12 min cycle. The MPC sits above them as a supervisory layer that writes m_2 and m_3 and hands them back to the PIs when a plan fails.

IN PRACTICE **Keep the Desorbent Clean**

The desorbent leaves the adsorber in the extract and the raffinate, is recovered at the bottom of their two distillation columns and is recycled to the chambers. A slipstream goes to a small rerun column that rejects heavy components and desorbent degradation products, and filters may remove adsorbent fines. With these safeguards the desorbent stays good enough for the molecular sieve to last more than ten years. A very small amount of water may be injected into the feed to keep the adsorbent's selectivity high [22]. Case 8 shows the other side: with other C_8 in the desorbent, no setting of the five handles restores the extract.

What Is Modernizing It

Model predictive control has a solid foundation in refining and petrochemicals [16], and here it pays mostly through the feed analyzer it reads as a measured disturbance. Soft sensors estimate product quality from other process measurements where an online analyzer is not available, and they are also used to validate online analyzers [18], which would give the ring a second reading to hold against its analyzers. An analyzer that reads the concentration profile inside the adsorber, as in the box on the Raman analyzer above, gives a controller the profile that the chapter's composite samplers average away. New sorbents that would separate the xylenes with less energy are the subject of the Open challenge box in Section 7.8.

7.8 Limitations

Every size, rate and limit is a teaching value, and no operating unit supplied them. The isotherm, its heats and the mass-transfer coefficient were chosen so that eight short beds reach 99.7 % at 90 s, not fitted to any adsorbent, and the selectivity of 2.5 sits below the 3.3 of a published industrial-scale model. As a result the unit needs a desorbent ratio of 3.79 where that study uses 1.20 by mass and calls 1.5 or less classical [1]. The other C_8 are one lump, the desorbent does not compete in the triangle, and the heat of adsorption, the rotary valve and the line volumes are left out.

OPEN CHALLENGE **A Sorbent Beyond the Zeolite**

Paraxylene is separated industrially on faujasite-type zeolites, which need relatively high temperatures and pressures in a simulated moving bed and give a paraxylene selectivity usually below ten. Sorbents and membranes that would purify paraxylene with less energy are therefore sought [25], while a 2011 review of microporous materials found the number of satisfactory results still limited [26]. The bar is high, because a zeolite kept in clean desorbent lasts more than ten years [22]. Progress would be a sorbent more selective than the zeolite that survives years of cycling. It would lower the desorbent ratio that this chapter traces to its selectivity of 2.5.

The thirty cells per bed are a declared dispersion, and the purities move with them, from 97.18 % at 10 cells to 99.98 % at 60 at the design point. The challenge's verdicts move too. At 45 and 60 cells a fixed point passes, at 60 the PIs as tuned miss the recovery by 0.0024 points, and at 20 cells nothing reaches the production floor. At the shipped 30 cells every verdict holds at five seeds and at half the

step, and the student cannot change the cell count, but the lesson that fixed ratios cannot pass belongs to this bed and not to simulated moving beds in general. The PIs' margin on the recovery is 0.34 to 0.47 point. The MPC's advantage over the PIs comes mostly from the feed analyzer, and it was identified on the simulator's own model. The analyzers are faster and cleaner than real ones, and the protection of the unit is outside the model.

7.9 Conclusion

A simulated moving bed has no steady state to control, and this simulator puts its eight beds, its stepping ports, its cycle composites and its triangle on one clock. From a ring of desorbent the unit repeats itself to one part in a million after ten cycles, while its extract still saws from 99.40 to 99.98 % inside each switch. The triangle derived from the isotherm matches an independent derivation to 2×10^{-16} and the near-ideal simulation reproduces its edges, but with the shipped dispersion the extract needs $m_2 \geq 0.568$ where the triangle's edge is at 0.422. Paired extract with m_2 and raffinate with m_3 (relative gains 0.87 and 0.91) and acting once a cycle on the logarithm of each impurity, two PIs recover a richer feed, an aged adsorbent and a drifting meter, and they pass sixty cycles of aging at 32.78 t/h where the best fixed point makes 27.00. The repetitive MPC holds the recovery higher through the feed change, mostly because it reads the feed analyzer. The constants are teaching values, the desorbent ratio of 3.79 is about three times the 1.20 by mass of a modeled industrial unit, and the challenge's verdicts belong to thirty cells per bed. Further work would fit the isotherm to measured equilibria, add the competing desorbent and the rotary valve's flushing, and re-identify the MPC as the adsorbent ages. The simulator shows that in a cyclic process the controlled variable is a cycle and the controller is a correction from one cycle to the next.

KEY POINTS

- Nothing in a simulated moving bed settles inside a cycle: the extract still swings from 99.40 to 99.98 % within one switch, and only its cycle average is steady.
- The ideal triangle locates the edges, but with the shipped dispersion the extract needs $m_2 \geq 0.568$ where the triangle's edge is at 0.422.
- Paired extract with m_2 and raffinate with m_3 and acting once a cycle on the logarithm of each impurity, two PIs pass sixty cycles of aging at 32.78 t/h where the best fixed point makes 27.00.
- The repetitive MPC holds the recovery higher through the feed change, mostly because it reads the feed analyzer.

Where else this applies. Correcting the next cycle, run or layer from the measured result of the last is also the work of Chapter 5, Volume 8, Chapter 9 and Volume 8, Chapter 10.

References

[1] G. Bergeot, D. Leinekugel-Le-Cocq, L. Wolff, L. Muhr, and M. Bailly, "Intensification of paraxylene production using a simulated moving bed reactor," *Oil & Gas Science and Technology – Revue d'IFP Energies nouvelles*, vol. 65, no. 5, pp. 721–733, 2010.

[2] G. Storti, M. Mazzotti, M. Morbidelli, and S. Carrà, "Robust design of binary countercurrent adsorption separation processes," *AIChE Journal*, vol. 39, no. 3, pp. 471–492, 1993.

[3] M. Mazzotti, G. Storti, and M. Morbidelli, "Optimal operation of simulated moving bed units for nonlinear chromatographic separations," *Journal of Chromatography A*, vol. 769, no. 1, pp. 3–24, 1997.

[4] Y. Wang, F. Gao, and F. J. Doyle, "Survey on iterative learning control, repetitive control, and run-to-run control," *Journal of Process Control*, vol. 19, no. 10, pp. 1589–1600, 2009.

[5] J. H. Lee and K. S. Lee, "Iterative learning control applied to batch processes: An overview," *Control Engineering Practice*, vol. 15, no. 10, pp. 1306–1318, 2007.

[6] S. Natarajan and J. H. Lee, "Repetitive model predictive control applied to a simulated moving bed chromatography system," *Computers & Chemical Engineering*, vol. 24, no. 2–7, pp. 1127–1133, 2000.

[7] M. Minceva and A. E. Rodrigues, "Modeling and simulation of a simulated moving bed for the separation of p-xylene," *Industrial & Engineering Chemistry Research*, vol. 41, no. 14, pp. 3454–3461, 2002.

[8] ——, "Understanding and revamping of industrial scale SMB units for p-xylene separation," *AIChE Journal*, vol. 53, no. 1, pp. 138–149, 2007.

[9] P. S. Gomes, M. Minceva, and A. E. Rodrigues, "Operation of an industrial SMB unit for p-xylene separation accounting for adsorbent ageing problems," *Separation Science and Technology*, vol. 43, no. 8, pp. 1974–2002, 2008.

[10] K.-U. Klatt, F. Hanisch, and G. Dünnebier, "Model-based control of a simulated moving bed chromatographic process for the separation of fructose and glucose," *Journal of Process Control*, vol. 12, no. 2, pp. 203–219, 2002.

[11] G. Erdem, S. Abel, M. Morari, M. Mazzotti, M. Morbidelli, and J. H. Lee, "Automatic control of simulated moving beds," *Industrial & Engineering Chemistry Research*, vol. 43, no. 2, pp. 405–421, 2004.

[12] S. Engell and A. Toumi, "Optimisation and control of chromatography," *Computers & Chemical Engineering*, vol. 29, no. 6, pp. 1243–1252, 2005.

[13] D. B. Broughton and C. G. Gerhold, "Continuous sorption process employing fixed bed of sorbent and moving inlets and outlets," U.S. Patent 2,985,589, assigned to Universal Oil Products Co., 1961, filed May 22, 1957, issued May 23, 1961.

[14] D. M. Ruthven and C. B. Ching, "Counter-current and simulated counter-current adsorption separation processes," *Chemical Engineering Science*, vol. 44, no. 5, pp. 1011–1038, 1989.

[15] Y. Kawajiri and L. T. Biegler, "Optimization strategies for simulated moving bed and PowerFeed processes," *AIChE Journal*, vol. 52, no. 4, pp. 1343–1350, 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] 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.

[18] 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.

[19] Y. Liu and H. Gao, "Raman: Theory," Chemistry LibreTexts, Supplemental Modules (Physical and Theoretical Chemistry), page 1850, CC BY 4.0, 2023, last updated January 30, 2023.

[20] J. C. G. da Silva, "Hybrid separations and adsorption/reaction processes: The case of isomerization/separation of xylenes," Ph.D. dissertation, Faculdade de Engenharia da Universidade do Porto, 2015.

[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] H. Falcke, S. Holbrook, I. Clenahan, A. López Carretero, T. Sanalan, T. Brinkmann, J. Roth, B. Zerger, S. Roudier, and L. Delgado Sancho, "Best available techniques (BAT) reference document for the production of large volume organic chemicals," Publications Office of the European Union, Luxembourg, EUR 28882 EN, 2017.

[23] S. Qian, "New processes and adsorbents for p-xylene production," Ph.D. dissertation, Faculdade de Engenharia da Universidade do Porto, 2019.

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

[25] M. Rahmani, C. R. M. O. Matos, S.-Q. Wang, A. A. Bezrukov, A. C. Eaby, D. Sensharma, Y. Hjiej-Andaloussi, M. Vandichel, and M. J. Zaworotko, "Highly selective p-xylene separation from mixtures of C8 aromatics by a nonporous molecular apohost," *Journal of the American Chemical Society*, vol. 145, no. 50, pp. 27 316–27 324, 2023.

[26] K. A. O. Santos, A. A. Dantas Neto, M. C. P. A. Moura, and T. N. Castro Dantas, "Separation of xylene isomers through adsorption on microporous materials: A review," *Brazilian Journal of Petroleum and Gas*, vol. 5, no. 4, pp. 255–268, 2011.

7.10 Exercises

7.1 Concept check · 2 min

Case 1 starts the ring full of desorbent, and the unit is declared at its cyclic steady state at the end of cycle 10. Inside the first switch of cycle 11 the extract still runs from 99.40 to 99.98 % paraxylene. Which statement is correct?

(a) The unit has not settled, and the profile still changes by more than the tolerance between switches

(b) The swing is the noise of the composite sampler, which the analyzer averages out over a cycle

(c) The swing comes from ports drawn between beds, which the sequencer corrects at the next switch

(d) What repeats is each switch shifted by one bed, so only the cycle average of 99.89 % is steady

7.2 Concept check · 2 min

The app defines the recovery as one minus the paraxylene that the raffinate takes over the paraxylene fed. In the cycle after case 3's feed step, a recovery defined as the extract's paraxylene over the feed's reads 73.15 %, while the loss-based one reads 98.43 %. Why does the app use the loss-based definition?

(a) The extract-based value falls while the ring takes up the extra paraxylene, though none of it is lost

(b) The extract analyzer reports 2 min late, so the extract-based value lags the cycle it describes

(c) The extract carries desorbent, which dilutes the extract-based value until the column recovers it

(d) The raffinate is the larger stream, so a loss-based value has the smaller measurement error

7.3 Concept check · 2 min

In case 6 one pump starts to deliver 3 % more than its meter reads. Over the next cycles the extract analyzer falls and the raffinate analyzer rises, while every flow on the screen still reads its setpoint. Which pump is it?

(a) The recycle pump, which raises every zone flow and so all four m_j together

(b) The extract pump, which takes more out ahead of zone II and lowers m_2 and m_3 together

(c) The feed pump, which raises m_3 alone and sends more paraxylene toward the raffinate

(d) The desorbent pump, which raises the flows of zones I to III and so m_1, m_2 and m_3

7.4 Theory and calculation · 6 min

The ring has eight beds of $V = 20\,\mathrm{m}^3$ with void fraction $\varepsilon = 0.40$, and switches every $t_s = 90\,\mathrm{s}$. Use Equation (7.3).

(a) Compute the four zone flows in m^3/h at the design point $m = (1.40, 0.62, 0.82, 0.30)$. Give one decimal. *Checkpoint:* 992.0, 617.6, 713.6 and 464.0 m^3/h

(b) Compute the desorbent, extract, feed and raffinate flows and the desorbent ratio Q_D/Q_F. *Checkpoint:* 528.0, 374.4, 96.0 and 249.6 m^3/h; 5.50

(c) The PIs settle at $m_2 = 0.564$ and $m_3 = 0.854$ with the same m_1 and m_4. Show that the desorbent ratio equals $(m_1 - m_4)/(m_3 - m_2)$, compute it and the feed, and compute the feed at the same m with $t_s = 60\,\mathrm{s}$. *Checkpoint:* 3.79; 139.2 and 208.8 m^3/h

7.5 Theory and calculation · 7 min

At the design point the Henry constants are $H_A = 1.20$ and $H_B = 0.48$, the affinities $K_A = 1.0$ and $K_B = 0.4$, and the feed fractions $c_A^F = 0.22$ and $c_B^F = 0.78$. Use Equations (7.7) and (7.8), and give three decimals.

(a) Compute κ_A, κ_B and the two roots $\omega_G > \omega_F$. *Checkpoint:* 0.220, 0.312; 1.053 and 0.357

(b) Compute the vertex w as the intersection of the edges wb and wf. *Checkpoint:* (0.421, 0.823)

(c) Compute m_3 on the upper edge at $m_2 = 0.62$ and m_2 on the left edge at $m_3 = 0.82$. The simulated unit reaches a recovery of 97 % at $m_3 = 0.8828$ for $m_2 = 0.62$, and an extract purity of 99.7 % only for $m_2 \geq 0.568$ at $m_3 = 0.82$. Compare. *Checkpoint:* 0.895 and 0.422

7.6 Theory and calculation · 7 min

Model each cycle-to-cycle PI as Equation (1.12) with a delay of $d = 2$ reports, one report per 12 min cycle. Ignore the proportional part and the tail of the real response, and use the final gains of the MPC's step responses as g: 9.35 decades of extract impurity per unit of m_2 and 15.26 decades of raffinate impurity per unit of m_3 (in magnitude). The integral part moves the handle by $K = K_c \Delta t / \tau_I$ per decade of error at each report.

(a) Compute K and gK for both loops at the shipped $K_c = 0.06$ per decade and $\tau_I = 24$ min. *Checkpoint:* $K = 0.03$; $gK = 0.281$ and 0.458

(b) Compute the magnitude of the closed-loop roots for each loop, and say where each sits against the double root and the edge of stability. *Checkpoint:* 0.530 and 0.677

(c) Repeat for both gains halved to 0.03. Which root is slowest, and what does the model predict for the recovery after the challenge's feed step? *Checkpoint:* 0.831 (extract) and 0.645 (raffinate)

7.7 App challenge · 8 min

Load case 1, Start-up: from a ring of desorbent to the cyclic steady state (Manual), at the default seed. Press Run at speed 2 (2 min of plant per wall second).

(a) Pause between 13 and 23 min and read the Paraxylene and Recovery (cycle) tiles. Pause again between 25 and 35 min and read them. *Checkpoint:* 1.81 t/h and 99.98 %; 12.67 t/h and 99.86 %

(b) The feed is 96 m^3/h of C_8 with 22 % paraxylene at 861 kg/m^3. What share of the paraxylene fed did the extract carry over cycle 1, and why does the Recovery tile read nearly 100 %? *Checkpoint:* 18.18 t/h fed; 10.0 %

(c) Run on at speed 10 and note the clock when CYCLIC STEADY STATE appears on the status strip. Pause at 130 min and read the Extract purity (cycle), Recovery (cycle) and Paraxylene tiles. Export the CSV and find the lowest and highest `truth_pe_now_pct` in the rows from 2.000 to 2.025 h (the file's time column is in hours for a run this long). Use a tolerance of 0.01 on every value. *Checkpoint:* 120 min; 99.893 %, 99.53 %, 18.10 t/h; 99.40 and 99.98 %

7.8 App challenge · 9 min

Load case 3, A richer feed and the cycle-to-cycle PI, at the default seed. It starts in Manual at the PIs' point, and the feed changes from 22 to 30 % paraxylene at 36 min.

(a) Run at speed 10 until the Recovery badge first reads below 97 %, near 98 min, and pause. Switch to Auto at any time before 110 min (Purity control is on Cycle-to-cycle PI), then run at speed 30 to 360 min. Read m_2 and m_3 on the two AIC cards and the Paraxylene tile. Export the CSV and list the cycles (`truth_cycle_of_record`) whose `truth_rec_cycle_pct` is below 97 %. *Checkpoint:* 0.526, 0.821, 36.11 t/h; cycles 8 to 11

(b) Reload the case with the reload button, and repeat (a) with Repetitive MPC chosen under Purity control before you switch to Auto. *Checkpoint:* 0.528, 0.821, 35.96 t/h; cycles 8 and 9

(c) Press Reset, which keeps your settings, so the run now starts in Auto with the MPC. Run to 360 min, export, and find the lowest `truth_rec_cycle_pct` from cycle 4 on. Choose Cycle-to-cycle PI, press Reset and repeat. Use a tolerance of 0.01 points. *Checkpoint:* MPC 98.41 %; PI 97.63 %

7.9 App challenge · 8 min

Load case 9, the Challenge (sixty cycles on an aging adsorbent), at the default seed. Run it as shipped at speed 300 (300 min of plant per wall second). The run pauses itself at 720 min.

(a) Read the So far tile and say which parts of the target the shipped fixed ratios meet. *Checkpoint:* 99.80 % · 82.6 % · 30.6 t/h

(b) Choose Cycle-to-cycle PI under Purity control, press Reset and run again. Read the So far tile and m_2 and m_3 on the AIC cards at the end. *Checkpoint:* 99.75 % · 97.5 % · 32.8 t/h; 0.500 and 0.769

(c) Set the Extract loop gain and the Raffinate loop gain to 0.030, press Reset and run again. Read the So far tile. Which part of the target fails? *Checkpoint:* 99.76 % · 96.8 % · 32.7 t/h

7.10 Application area · 4 min

Case 8 cuts the desorbent with the PIs on, at $t_s = 90$ s. At the PIs' point the desorbent ratio is 3.79 for 25.81 t/h of paraxylene. With m_1 lowered from 1.40 to 1.30 it is 3.59 for 24.96 t/h, and with m_4 then raised from 0.30 to 0.34 it is 3.73 for 22.91 t/h. The desorbent flow is $Q_D = (m_1 - m_4) \times 480\,\mathrm{m^3/h}$ at this switching time. Compute the desorbent per tonne of paraxylene at each step, say which cut an engineer should keep, and say what would lower the ratio toward the 1.20 by mass of an industrial unit.

7.11 Application area · 4 min

In case 6 the extract pump delivers 3 % more than its meter reads. In Auto the PIs restore the extract purity and command $m_2 = 0.590$ and $m_3 = 0.879$, while the beds see 0.565 and 0.855. The recovery estimate from the meters reads 97.78 % against a true 97.88 %. Every analyzer is now on its setpoint. Is the unit healthy? Check that the commanded offset fits a 3 % error on the extract flow, using $480\,\mathrm{m^3/h}$ per unit of m and $Q_E = (m_1 - m_2) \times 480\,\mathrm{m^3/h}$ with $m_1 = 1.40$, and say how an engineer could find the fault from plant data.

7.12 Review the agent · 5 min

An AI assistant asked how to make more paraxylene in case 2 writes: "At $m_3 = 0.82$ the left edge of the triangle is at $m_2 = 0.422$, and the design point is at 0.62. Lower m_2 to 0.44: the feed $m_3 - m_2$ rises from 0.20 to 0.38, the point stays inside the triangle, and equilibrium theory guarantees both outlets stay pure." Find the flaw and say how you would check it in the app.

CHAPTER 8

Fluidized-Bed Polyethylene Simulator for Cascade Control and Grade-Transition Economics



The controller moves 1 cooling-water valve, for the cooler outlet 2 ethylene valve, for the pressure 3 hydrogen valve, for melt index 4 hexene valve, for density 5 catalyst feeder, for the production rate 6 discharge valve, for the bed level

The controller reads 1 bed temperature 2 cooler outlet temperature 3 reactor pressure 4 bed level 5 gas analyzer, hydrogen and hexene ratios 6 quality laboratory, melt index and density 45 to 105 minutes late

What disturbs it 1 warmer cooling water 2 a fouling cooler 3 a recycle compressor that loses capacity

Notice that the gas mix can change in minutes while the bed of powder holds four hours of product and the laboratory reports long after its sample was drawn, and that in the chapter the catalyst enters the bed itself, not the space above it. Illustration, not to scale.

IN PLAIN TERMS A Fast Gas Steering a Slow Bed

The main idea. Plastic grains for film and molded parts grow in a tall vessel, kept floating by a stream of gas that also carries away the heat. The mix of gases decides what kind of plastic is being made, and that mix can be changed within minutes.

The challenge. The vessel holds about four hours of grains, all stirred together, so the plastic that leaves changes only slowly. A switch to a new product therefore makes a long run of plastic that fits neither order.

Why it matters. In the simulated plant, better tuning and a smaller bed cut that waste by more than half.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? and History boxes of the Background section, the Try it boxes in Sections 8.4 and 8.6, the Did you know? box in Section 8.7, and Key points in Section 8.10. No equations.
Study	65 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 8.7 with Figure 8.4, Section 8.8, and Section 8.9 with its Open challenge boxes.

The control question The gas in a fluidized-bed polyethylene reactor settles in minutes to hours, while melt index and density are averages over a four-hour bed of powder. How should cascades that span two decades of time scale be tuned, and what do tuning and bed preparation save in off-specification tonnes during a grade change?

Learning objectives

- Compute the time constants of the gas, thermal, catalyst and resin inventories from the model and rank them.
- Identify the inner cooler loop and the outer temperature loop from step tests, and explain why the outer plant is lead-lag.
- Derive the off-grade tonnage of a grade change and show that the production rate cancels from it.
- Reduce the off-grade tonnage of a transition with quality-loop tuning and bed level, and judge a rate cut by measurement.
- Read the cooling-water valve and the cooling utilization as early warnings of the heat-removal limit.

Prerequisites Sections 1.4.4 (heat removal), 1.4.5 (grade transitions), 1.4.6 (laboratory and bias update), 1.6.1 and 1.6.2 (engine and velocity-form PID). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/gas-phase-reactor/

See also Volume 3, Chapter 2 and Volume 3, Chapter 3 for cascade structure and the speed an inner loop needs. Volume 9, Chapter 4 for offset-free predictive control with a disturbance estimate. Volume 5, Chapter 3 and Volume 5, Chapter 4 for control through a slow analyzer and a late laboratory.

8.1 Introduction

Most teaching examples in process control carry one dominant time constant, and the controller is tuned against that number. A gas-phase polyethylene reactor carries two populations that respond about two decades apart. The gas above and inside the bed turns over in minutes, so ethylene inventory, hydrogen ratio and comonomer ratio are fast, measurable and reachable with a valve. The powder in the bed turns over in four hours, and the molecular weight and density of the resin leaving the discharge valve are the mass-weighted average of everything produced over the preceding residence time. No control action on the gas changes the resin already stored in the bed.

The consequence is a control problem that no amount of loop tuning removes. Melt index and density are the variables the customer buys and they cannot be manipulated directly. They are reached only through a gas-composition setpoint, which makes the quality loops the outer half of a cascade whose inner half settles between one and two orders of magnitude faster. The cooler outlet temperature in the simulator of this chapter answers its water valve in about 0.5 min, the 1-hexene inventory turns over in 4.0 min, the ethylene inventory in 6.7 min, the bed thermal lag is 11.3 min, the hydrogen inventory turns over in 142 min, the active catalyst inventory in 176 min and the resin property in 240 min. A melt-index controller tuned as though it acted on the hydrogen valve saturates a ratio the plant has already reached and then cycles for hours afterward.

The second consequence is economic rather than dynamic. A grade change is not a disturbance to be rejected. It is a scheduled excursion whose cost is the tonnage discharged while the resin sat outside its specification band. Integral absolute error is a poor objective for a transition, because the setpoint moved on purpose and hours of transient were always going to be spent. The relevant measure is the tonnage discharged at the wrong melt index, and the variables that reduce it differ from those regulatory tuning rules address.

The physics behind the model are the standard conservation statements of the process control curriculum. Material balances on the gas-phase species, an energy balance on the bed and a mixing balance on the solids inventory follow the forms of the introductory texts [1, 2], propagation follows Arrhenius kinetics, and each bed-average resin property is a first-order lag whose time constant is the solids residence time. The experiments built on that model show cascade structure, time-scale separation, integrating level dynamics, feedforward, anti-windup, override control and the difference

between an instantaneous and an inventory-averaged quality [1, 3]. Each of these behaviors appears in industrial polymer service, where off-grade production during transitions is a direct and recurring cost [4].

Gas-phase polyethylene reactors have carried a large modeling and control literature for exactly these reasons, covering the production processes and property correlations [5], the dynamics of a fluidized bed with solid catalyst [6], industrial copolymerization kinetics [7], the question of whether a two-phase description is needed at all [8], and the inferential and nonlinear control of resin properties in a running plant [9, 10]. The simulator of this chapter turns those questions into experiments. Nine loops can be opened one at a time, a grade change can be repeated with a single condition altered, and the off-specification tonnage is accumulated at the discharge valve during the run.

This chapter defines a fluidized-bed polyethylene reactor completely enough for it to be reconstructed, and measures what its two separated time scales cost during a grade transition. The plant is a thirteen-state model with nine independent controller stations, six non-defeatable overrides and a complete data export. Section 8.3 sets out the balances, the constants and the solved design point. Section 8.4 measures the seven process time constants that span the plant and shows that catalyst deactivation turns the outer temperature loop's plant from first order into lead-lag. Section 8.5 describes the nine loops, the cascades, the feedforward, the overrides and the optional predictive controller. Section 8.6 runs the rate, condensing-agent, cooling and impurity scenarios. Section 8.7 derives and tests the result that off-grade tonnage during a transition does not depend on production rate, runs one grade change in the ten ways of Table 8.9, measures its cost in tonnes of off-specification resin, and compares the cascades with the predictive controller. Section 8.8 names the hardware and the practice behind the model, and Sections 8.9 and 8.10 state the limitations and the conclusions.

8.2 Background

Polyethylene was made first at high pressure and then in solution and slurry, and each of those routes carries a diluent that has to be removed from the product. The gas-phase route eliminates the diluent. A vertical vessel holds a bed of polymer powder fluidized by its own recycled monomer, solid catalyst is metered into the side of the bed, and granules grow around the catalyst particles until they are withdrawn. Because nothing has to be filtered or distilled, the process took a large share of world polyethylene capacity, and its economics are dominated by the heat the recycle gas can carry away rather than by separation [5], which Section 1.4.4 works through for the running example of Chapter 1.

> **+ DID YOU KNOW? One Vessel, Hundreds of Thousands of Tonnes**
>
> The fluidized-bed reactor of a polyethylene plant is a vertical pressure vessel up to 40 m tall, and a 2007 European reference document gives 450 kt a year as the largest capacity of a gas-phase line for linear low-density polyethylene [11]. Fluidized-bed polymerization, introduced in 1968, accounts for 40 % of world polyethylene capacity [12].

The control problem that arrived with the technology was one of measurement rather than actuation. Bed temperature, pressure and level have had instruments since the first plants and melt index and density have not. Both are laboratory determinations on a sample pulled every hour or two and reported after the granules have been pressed, conditioned and run through an extrusion plastometer, so an operator holding melt index with hydrogen manipulated a four-hour inventory against a measurement that was several hours old.

Two developments closed the measurement gap, an inferential sensor and model-based control built on top of it, as the first History box below tells. Polymer units appear throughout the industrial model predictive control population that was inventoried a decade after the second of them [13].

> ⧖ HISTORY **Reading Melt Index Between Samples, 1991 and 1993**
>
> In 1991 an article named the lack of suitable on-line measurements of polymer properties as a major difficulty for quality control in industrial polymerization reactors. It predicted melt index and density in a fluidized-bed ethylene copolymerization reactor from the on-line temperature and gas composition, updated the models' adjustable parameters from the infrequent laboratory measurements by recursive estimation, and illustrated the scheme on operating data from an industrial reactor [9]. In 1993 a nonlinear model-based controller moved the hydrogen and butene feeds to regulate the instantaneous melt index and density and to carry them through grade changeovers, with an extended Kalman filter estimating the disturbances. Tested in simulation on a mechanistic model of the process, it handled both tasks, where an analogous linear controller was inadequate at changing operating conditions and during changeovers [10].

Grade transitions became the natural target once the properties could be estimated online. Optimal transition trajectories for polyolefin reactors were computed as dynamic optimization problems and showed that the shortest path is not the direct one [4], that the choice of control structure and the transition policy interact and should be selected together [14], and that the resulting policies can be commissioned on a plant with commercial real-time optimization software [15]. Alongside this, condensed-mode operation raised the ceiling on production by putting part of the recycle stream into the liquid phase and letting it evaporate in the bed, which requires the reactor model to carry a dew point and a condensed fraction rather than a sensible-heat balance alone [16]. The simulator of this chapter carries a reduced-order version of each of these pieces and exposes the loops that connect them. The second History box below tells where the fluid bed and its condensing cooler came from.

> ⧖ HISTORY **From a Low-Pressure Fluid Bed to Condensed Mode, 1975 to 1985**
>
> A patent application of 1975 copolymerized ethylene with other alpha-olefins in a fluid bed reactor at low pressure to make low- and medium-density polymers, which were then usually made by homopolymerizing ethylene at very high pressures with free-radical catalysts. Its recycle gas passed an external exchanger that stripped the heat of reaction before the gas returned to the bed [17]. A patent filed in 1984 and issued in 1985 started from the observation that the polymer a reactor of a given size can make is directly related to the heat that can be withdrawn. It cooled the recycle stream below its dew point and returned the two-phase mixture below the bed, raising the dew point with pressure or with an inert condensable fluid such as a butane, pentane or hexane, and keeping the liquid entrained in the gas until it vaporized [18].

8.3 Reactor Model

One fluidized-bed reactor with an external cycle gas loop is simulated, and Figure 8.1 carries every symbol the equations below use. Recycle gas leaves the top of the disengaging section, passes a compressor and a water-cooled exchanger, and returns below the distributor plate carrying whatever condensed in the cooler, with ethylene, hydrogen, 1-hexene and cyclohexane added to it as make-up.

8.3.1 Controlled Variables, Inputs and Disturbances

Nine variables are controlled, six valves and three inner setpoints are manipulated, and ten disturbances are available. Table 8.1 classifies them with their enforced ranges.

Only three of the nine loops produce performance statistics on the scorecard of Section 1.6.4, namely reactor temperature with a 1.0 °C band, bed level with a 3 % band and melt index with a band of $\pm 10\,\%$ of its setpoint, which is the specification band of Section 8.7. The other six are measured and logged but do not produce integral absolute error, overshoot or settling numbers.

8.3.2 Gas-Phase Balances and Pressure

The gas in the reactor, the disengaging section and the recycle loop is treated as one perfectly mixed volume. Four species are carried as molar holdups and nitrogen with the accumulated light ends is a fixed inert holdup n_I, so the balance for species $i \in \{C_2, H_2, C_6, ICA\}$ is Equation 8.1.

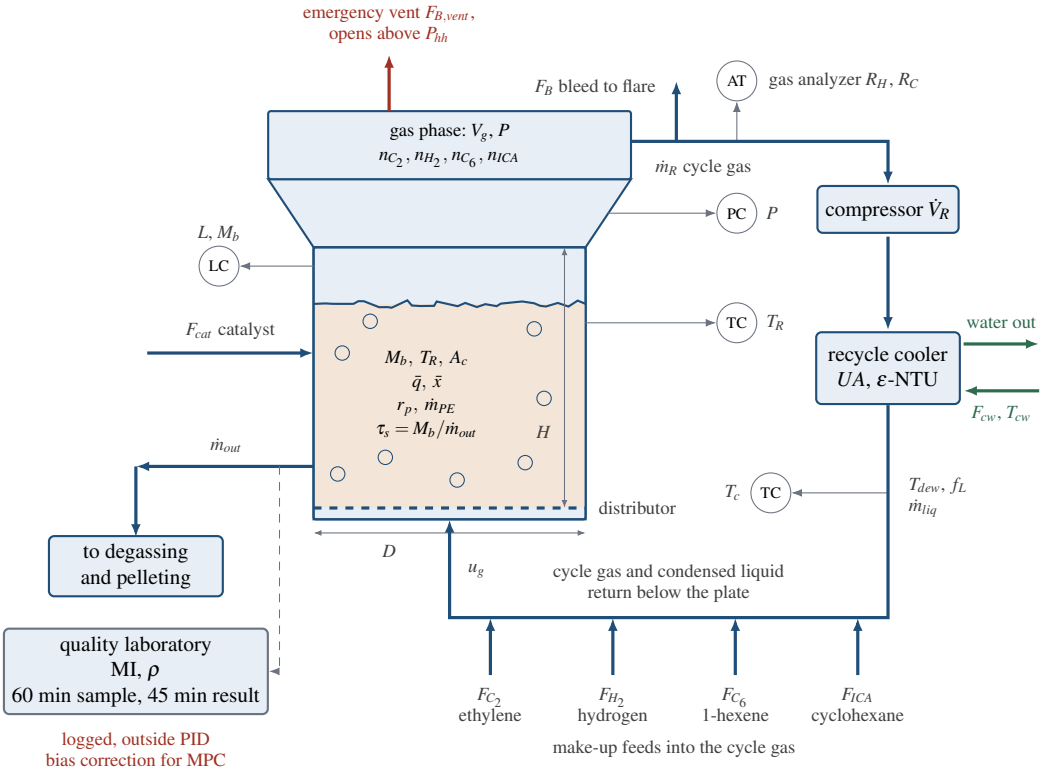

Figure 8.1 Fluidized-bed polyethylene reactor with its cycle gas loop, make-up feeds, instruments and quality laboratory. The vessel is drawn schematically and not to scale. Every symbol shown appears in (8.1)–(8.15), and the dashed path marks the laboratory, which is sampled, delayed, displayed and logged. It remains outside the PID loops and corrects the optional MPC estimate.

$$\frac{\mathrm{d}n_i}{\mathrm{d}t} = F_i - r_i - y_i F_B, \quad y_i = \frac{n_i}{n_g}, \quad n_g = n_I + \sum_j n_j \tag{8.1}$$

where F_i is the make-up feed, r_i is the consumption by reaction and F_B is the total gas leaving through the bleed. Only ethylene, hydrogen and 1-hexene react, so $r_{ICA} = 0$. Cyclohexane is an inert condensable whose only purpose is to raise the dew point of the recycle stream. The bleed carries all four away in proportion to their mole fractions, which is what makes the direction of a hydrogen move asymmetric.

The make-up flows are valve positions against fixed capacities, with the ethylene flow additionally scaled by the square root of the supply header pressure as shown in Equation 8.2.

$$F_{C_2} = F_{C_2}^{max} \frac{v_{C_2}}{100} \sqrt{\frac{\phi_h}{100}} \tag{8.2a}$$

$$F_{H_2} = F_{H_2}^{max} \frac{v_{H_2}}{100}, \qquad F_{C_6} = F_{C_6}^{max} \frac{v_{C_6}}{100} \tag{8.2b}$$

$$F_B = F_{B0} \frac{\phi_B}{100} + k_{vent} \max(0, P - P_{hh}) \tag{8.2c}$$

Table 8.1 Variable Classification and Panel Ranges

Symbol	Meaning	Unit	Range (default)
Controlled variables			
T_R	bed temperature	°C	78–100 (88)
T_c	cooler outlet temperature	°C	22–70 (40)
P	reactor pressure	bar	18–23 (21.7)
L	bed level	%	35–85 (65)
R	production rate	t/h	12–42 (30)
R_H	H_2/C_2 ratio $\times 100$	—	4–80 (20)
R_C	C_6/C_2 ratio $\times 100$	—	0.02–2.5 (1.5)
MI	melt index	g/10 min	0.2–12 (1.0)
ρ	resin density	g/cm^3	0.915–0.955 (0.918)
Manipulated variables			
v_{CW}	cooling-water valve	%	0–100 (50.0)
v_{C_2}	ethylene make-up valve	%	0–100 (57.6)
v_{dis}	product discharge valve	%	0–100 (41.7)
v_{cat}	catalyst feeder	%	0–100 (40.0)
v_{H_2}	hydrogen valve	%	0–100 (11.6)
v_{C_6}	1-hexene valve	%	0–100 (36.3)
$T_{c,\mathrm{SP}}$	cooler outlet setpoint	°C	22–70
$R_{H,\mathrm{SP}}$	H_2/C_2 setpoint	—	4–80
$R_{C,\mathrm{SP}}$	C_6/C_2 setpoint	—	0.02–2.5
Operator handle with no loop on it			
F_{ICA}	cyclohexane feed	mol/min	0–600 (216)

Symbol	Meaning	Unit	Range (default)
Disturbances			
T_{cw}	cooling-water inlet	°C	18–42 (30)
ϕ_f	cooler fouling	%	0–45 (0)
ϕ_K	recycle compressor capacity	%	55–110 (100)
ϕ_h	ethylene header pressure	%	40–115 (100)
ϕ_B	gas bleed to flare	%	20–320 (100)
z	feed impurity	ppm-eq	0–0.6 (0)
ϕ_a	catalyst activity bias	%	−30–30 (0)
ϕ_d	discharge restriction	%	0–60 (0)
ϕ_m	melt-index analyzer bias	%	−20–20 (0)
ϕ_n	measurement noise multiplier	—	0–4 (1)
States			
$n_{C_2}, n_{H_2}, n_{C_6}, n_{ICA}$	gas-phase molar holdups	mol	positive
A_c	active catalyst inventory	—	non-negative
M_b	bed mass	kg	2000–193798
T_R	bed temperature	°C	—
T_c	cooler outlet temperature	°C	—
$\bar{q}$	bed-average ln MI	—	—
$\bar{x}$	bed-average incorporated C_6	—	0–0.2
v_{CW}, v_{dis}, F_{cat}	actuator positions	%	0–100

where F_{ICA} is set directly in mol/min by the operator. The second term of (8.2c) is a hard-wired overpressure vent that no controller can inhibit, and because it discharges the whole gas mixture it moves every composition at once while it is open.

Pressure is not a state. It is an algebraic output of the ideal gas law applied to the moles present in the free volume, and that free volume shrinks as the bed grows, as shown in Equation 8.3.

$$P = \frac{n_g R T_R}{V_g}, \qquad V_g = V_{tot} - \frac{M_b}{\rho_s} \tag{8.3}$$

where R is the gas constant, T_R is absolute bed temperature, V_{tot} is the total gas and solids volume of the reactor and loop and ρ_s is the resin particle density. Equation 8.3 couples the level loop to the pressure loop, because moving 10 % of bed level moves 20.5 m^3 of free gas volume out of 400 m^3.

8.3.3 Catalyst, Kinetics and Production

Propagation is first order in ethylene concentration and in active catalyst inventory, with an Arrhenius temperature dependence referenced to 88 °C, a mild comonomer enhancement and a poisoning term, as shown in Equation 8.4.

$$r_p = k_{p,ref} \mathrm{e}^{-\frac{E_p}{R}\left(\frac{1}{T_R} - \frac{1}{T_{ref}}\right)} \frac{A_c n_{C_2}}{V_g} \frac{1 + k_{com} R_C}{1 + K_z z} \left(1 + \frac{\phi_a}{100}\right) \tag{8.4}$$

where $R_H = n_{H_2}/n_{C_2}$ and $R_C = n_{C_6}/n_{C_2}$ are the gas-phase mole ratios, T_R and $T_{ref} = 88\,^\circ\mathrm{C}$ are in kelvin inside the exponential, z is the feed impurity and ϕ_a is a catalyst activity bias. The reduced-order form follows the industrial gas-phase copolymerization kinetics of [7] with the site distribution collapsed into one lumped activity. Comonomer and hydrogen consumption and the polymer production rate are then given by Equation 8.5.

$$x_{inst} = \frac{K_C R_C}{1 + K_C R_C} \tag{8.5a}$$

$$r_{C_6} = r_p \frac{x_{inst}}{1 - x_{inst}}, \qquad r_{H_2} = \zeta\, r_p R_H \tag{8.5b}$$

$$\dot{m}_{PE} = \frac{M_{C_2} r_p + M_{C_6} r_{C_6}}{1000} \tag{8.5c}$$

where x_{inst} is the mole fraction of 1-hexene incorporated into the chain, K_C is the competitive incorporation constant, ζ is the chain-transfer hydrogen consumption per unit of propagation, M_{C_2} and M_{C_6} are molar masses in g/mol, and $\dot{m}_{PE}$ is in kg/min.

Catalyst does not act at the moment it is fed. It accumulates in the bed, deactivates thermally and leaves with the product, as shown in Equation 8.6.

$$\frac{\mathrm{d}A_c}{\mathrm{d}t} = k_a F_{cat} - A_c \frac{\dot{m}_{out}}{M_b} - k_d A_c, \quad k_d = k_{d0}\, \mathrm{e}^{\beta_d (T_R - 88)} \tag{8.6}$$

where F_{cat} is the feeder position in percent, k_a is a hidden activation factor and T_R is in °C inside the exponential. At the design point $\dot{m}_{out}/M_b = 4.168 \times 10^{-3}$/min and $k_d = 1.5 \times 10^{-3}$/min, so the active inventory has a time constant of 176 min. No controller gain shortens that number, so production is the slowest loop on the plant that is paired with its own valve.

8.3.4 Solids Inventory and Residence Time

The bed is a pure integrator. Nothing on the right-hand side of Equation 8.7 depends on how full the bed is, so the discharge rate is independent of head and any imbalance integrates without bound [3].

$$\frac{\mathrm{d}M_b}{\mathrm{d}t} = \dot{m}_{PE} - \dot{m}_{out} \tag{8.7a}$$

$$\dot{m}_{out} = \dot{m}_{dis}^{max} \frac{v_{dis}}{100} \left(1 - \frac{\phi_d}{100}\right) \tag{8.7b}$$

$$L = \frac{100\, M_b}{\rho_b A H}, \quad A = \frac{\pi D^2}{4}, \qquad \tau_s = \frac{M_b}{\dot{m}_{out}} \tag{8.7c}$$

where ρ_b is the fluidized bulk density, A is the bed cross-sectional area, H is the bed height at 100 % level and τ_s is the solids residence time. At the design point $M_b = 119970$ kg and $\dot{m}_{out} = 500$ kg/min, so $\tau_s = 240$ min. That single number sets the length of every grade change.

8.3.5 Energy Balance and the Condensing Cooler

The reaction heat leaves through the recycle gas and nowhere else. Sensible heat is carried by the gas warming from the distributor to the bed, latent heat by whatever condensed in the cooler and boils again in the bed, as shown in Equation 8.8.

$$\begin{aligned}(M_b c_{p,s} + C_{metal}) \frac{\mathrm{d}T_R}{\mathrm{d}t} = {} & (-\Delta H_p) r_p - \dot{m}_R c_{p,g} (T_R - T_c) \\ & - \dot{m}_{liq} \Delta H_{vap} - Q_{loss}\end{aligned} \tag{8.8}$$

where $c_{p,s}$ is the resin heat capacity, C_{metal} is the lumped heat capacity of the vessel steel, $-\Delta H_p$ is the heat of polymerization, $\dot{m}_R$ is the recycle mass flow, $c_{p,g}$ is the recycle gas heat capacity, $\dot{m}_{liq}$ is the condensed flow, ΔH_{vap} is its heat of vaporization, and $Q_{loss} = U A_{wall} (T_R - T_{amb})$. At the design point the three removal terms are 19.85 MW sensible, 5.90 MW latent and 0.32 MW of wall loss

against 26.07 MW of reaction heat. The dry recycle gas alone cannot carry the reaction heat, and that deficit is the motivation for condensed-mode operation.

Pause. Try Exercise 8.1 before reading on.

The recycle flow follows from the gas density and the compressor displacement, and the condensed fraction from how far the cooler outlet sits below the dew point, as shown in Equation 8.9.

$$\dot{m}_R = \frac{P\overline{M}}{RT_R}\dot{V}_R\frac{\phi_K}{100}, \qquad \overline{M} = \sum_i y_i M_i \tag{8.9a}$$

$$T_{dew} = T_{dew,0} + a_P(P-21.7) + a_I(100y_{ICA}-12) + a_6(100y_{C_6}-0.6) + a_2(100y_{C_2}-40) \tag{8.9b}$$

$$\Delta T_{sub} = w\ln\left(1+\mathrm{e}^{(T_{dew}-T_c)/w}\right) \tag{8.9c}$$

$$f_L = f_{L,max}\left(1-\mathrm{e}^{-\Delta T_{sub}/\Delta T_{ref}}\right), \quad \dot{m}_{liq} = f_L\dot{m}_R \tag{8.9d}$$

where $\overline{M}$ is the mixture molar mass including the inert, $\dot{V}_R$ is the compressor volumetric displacement, and w is a smoothing width that rounds the corner of the positive part in (8.9c) over about one kelvin. A hard $\max(0,\cdot)$ puts a kink in the condensation curve that a fourth-order Runge-Kutta step integrates straight through, and the rounded version is also more realistic because a real cooler condenses a small mist fraction before the bulk dew point is reached. Equation 8.9b couples the condensed fraction to the grade being made, because 1-hexene raises the dew point and low-density grades carry more of it.

The exchanger itself is a counterflow effectiveness-NTU calculation whose hot-side capacity rate is inflated to account for the latent duty, as shown in Equation 8.10.

$$UA = UA_0\left(1-\frac{\phi_f}{100}\right)\left(\frac{F_{cw}}{F_{cw,ref}}\right)^{n_{UA}} \tag{8.10a}$$

$$C_h = \dot{m}_R c_{p,g}\varphi, \quad \varphi = \mathrm{clamp}\left(1+\frac{Q_{vap}}{Q_{sens}},1,4\right) \tag{8.10b}$$

$$C_c = F_{cw}c_{p,w} \tag{8.10c}$$

$$\varepsilon = \frac{1-\mathrm{e}^{-N(1-C_r)}}{1-C_r\mathrm{e}^{-N(1-C_r)}}, \; N = \frac{UA}{C_{min}}, \; C_r = \frac{C_{min}}{C_{max}} \tag{8.10d}$$

$$T_c^* = T_R - \frac{\varepsilon C_{min}\max(0,T_R-T_{cw})}{C_h} \tag{8.10e}$$

$$\frac{\mathrm{d}T_c}{\mathrm{d}t} = \frac{T_c^*-T_c}{\tau_{HX}} \tag{8.10f}$$

where $n_{UA}=0.8$ is the flow exponent, $F_{cw}=F_{cw}^{max}v_{CW}/100$, $Q_{sens}=\dot{m}_R c_{p,g}(T_R-T_c)$, $Q_{vap}=\dot{m}_{liq}\Delta H_{vap}$, C_{min} and C_{max} are the smaller and larger of C_h and C_c, and τ_{HX} is the thermal lag between the solved outlet temperature and the measured one. The exchanger solution is clamped so that T_c can never fall below the water inlet temperature or rise above the bed temperature, C_h and C_c are floored at 500 kJ/(min K), ε reverts to $N/(1+N)$ when C_r exceeds 0.999, and F_{cw} is floored at 200 kg/min inside (8.10a) so that a closed valve does not produce a zero UA.

Two derived quantities are computed continuously. The first is the production rate the cooler could carry with the water valve wide open, obtained by evaluating (8.10) at $F_{cw}=F_{cw}^{max}$ and dividing the resulting duty by the heat of polymerization, as shown in Equation 8.11.

$$R_{max} = \frac{Q_{max}-Q_{loss}}{-\Delta H_p}\left(M_{C_2}+M_{C_6}\frac{x_{inst}}{1-x_{inst}}\right)\frac{0.06}{1000} \tag{8.11}$$

where the factor 0.06 converts kg/min to t/h. The second is the cooling utilization $100R/R_{max}$, which reads 88.3 % at the design point.

8.3.6 Quality Correlations and the Bed Average

Melt index is an inverse measure of molecular weight and hydrogen is the chain-transfer agent that shortens the chains. The instantaneous property of the resin currently being produced follows the reduced-order form used throughout the polyethylene inference literature [9], given in Equation 8.12.

$$\ln \mathrm{MI}_{inst} = a_T \left(\frac{1}{T_R} - \frac{1}{T_{ref}} \right) + m \ln(a_0 + a_H R_H + a_C R_C + a_z z) \tag{8.12a}$$

$$\rho_{inst} = \rho_{HD} - K_\rho x_{inst} + K_{MI} \ln \mathrm{MI}_{inst} \tag{8.12b}$$

where a_0, a_H, a_C and a_z weight the intercept, the hydrogen ratio, the comonomer ratio and the impurity, m is the exponent, a_T carries the temperature effect, ρ_{HD} is the homopolymer density base and K_ρ converts incorporated comonomer into a density decrement. The bracket in (8.12a) is floored at 0.05 to keep the logarithm defined. The design ratios $R_H = 0.20$ and $R_C = 0.015$ make the bracket exactly 1.000 and the temperature term exactly zero, so $\mathrm{MI}_{inst} = 1.00$ and $\rho_{inst} = 0.91801\,\mathrm{g/cm^3}$ without any fitting residual.

Neither of those is what leaves the reactor. The bed is one well-mixed solids volume, so the product of bed mass and bed-average property obeys a mixing balance, and expanding it against (8.7a) leaves the first-order form of Equation 8.13.

$$M_b \frac{\mathrm{d}\bar{q}}{\mathrm{d}t} = \dot{m}_{PE} (\ln \mathrm{MI}_{inst} - \bar{q}), \quad M_b \frac{\mathrm{d}\bar{x}}{\mathrm{d}t} = \dot{m}_{PE} (x_{inst} - \bar{x}) \tag{8.13}$$

where $\bar{q}$ is the bed-average logarithm of melt index and $\bar{x}$ the bed-average incorporated comonomer, from which $\mathrm{MI}_{bed} = \mathrm{e}^{\bar{q}}$ and $\rho_{bed} = \rho_{HD} - K_\rho \bar{x} + K_{MI} \bar{q}$. Equation 8.13 is a first-order lag of time constant $M_b/\dot{m}_{PE}$, the form of Equation (1.8), which is 240 min at design against gas time constants of 4 min to 142 min. This single lag is the source of the two-decade separation. Averaging the logarithm rather than the melt index itself follows the property that mixes additively in the reduced-order correlation, and it also reduces the transition analysis of Section 8.7 to one closed-form expression.

8.3.7 Sticking Limit and Fluidization

Two constraints bound the operating window and neither has a controller. Polyethylene softens over a range rather than melting at a point, and a bed whose particles have begun to soften agglomerates into sheets and then into one lump. The temperature at which that starts depends on the resin, as shown in Equation 8.14.

$$T_{stick} = \min\left(128,\; T_{s0} + k_\rho 10^3 (\rho_{bed} - 0.920) - k_I 10^2 y_{ICA} - k_{MI} \bar{q}\right) \tag{8.14}$$

where T_{s0} is the base sticking temperature, k_ρ, k_I and k_{MI} are sensitivities to density, cyclohexane content and molecular weight, and the reported sticking margin is $T_{stick} - T_R$. It reads 8.01 °C at the design point. Low-density film grades stick at the lowest temperatures, and the cyclohexane that adds heat-removal capacity lowers the limit further, so added condensing agent costs temperature margin.

The second constraint is fluidization. Below minimum fluidization the bed slumps and heat transfer collapses, so the superficial gas velocity and the minimum fluidization velocity are computed at every step from Equation 8.15.

$$u_{mf} = \frac{d_p^2(\rho_s - \rho_g)g}{1650\,\mu}, \qquad u_g = \frac{\dot{m}_R}{60\,\rho_g A} \tag{8.15}$$

where d_p is the mean particle size, ρ_g is the gas density, g is gravitational acceleration and $\mu = 1.4 \times 10^{-5}$ Pa s is the gas viscosity, with $\rho_s - \rho_g$ floored at 50 kg/m^3. At the design point $u_g = 0.583$ m/s and $u_{mf} = 0.182$ m/s, a ratio of 3.20.

8.3.8 Instruments, Actuators and the Quality Laboratory

Five measurements carry band-limited noise from a first-order-filtered Gaussian process whose innovation is scaled to give the quoted stationary standard deviation, all multiplied by the noise multiplier ϕ_n. Bed temperature uses $\tau = 1.5$ min and $\sigma = 0.06\,^{\circ}$C, pressure 1.0 min and 0.012 bar, level 2.0 min and 0.35 %, the hydrogen ratio 1.5 min and 0.09, and the comonomer ratio 1.5 min and 0.008. The cooler outlet temperature is read exactly. Production is not measured directly, because no instrument weighs 500 kg/min of powder in real time. The reported rate is the mass-balance inference $\dot{m}_{PE}$ passed through a first-order filter of time constant 3 min, and a second independent estimate from the heat balance, $(Q_{hot} + Q_{loss})/(-\Delta H_p)$ converted to tonnes per hour with $Q_{hot} = Q_{sens} + Q_{vap}$, is reported beside it. The two agree on a settled plant, both reading 30.00 t/h at the design point, so a disagreement between them indicates an instrument fault or a transient rather than a steady process change.

Three actuators carry their own first-order lags as states, namely the cooling-water valve at $\tau_V = 0.25$ min, the discharge system at $\tau_{Vd} = 0.4$ min and the catalyst feeder at $\tau_{feed} = 5.0$ min. The ethylene, hydrogen and 1-hexene valves are applied without lag.

The quality laboratory is modeled and corrects the optional MPC estimate, although it remains outside the nine PID loops. A sample is taken every $t_{lab} = 60$ min, holding the bed-average melt index and density at the instant of sampling, and the result is posted $\theta_{lab} = 45$ min later. The posted value is then held until the next result arrives, so the reported laboratory value is between 45 min and 105 min old. That value is displayed beside the inferred properties and exported with its age. It is not the measurement either cascaded quality controller uses. The melt-index controller reads $\mathrm{MI}_{bed}(1 + \phi_m/100)$ and the density controller reads ρ_{bed}, both taken from the plant states with no delay and no reconstruction. For the cascades the laboratory illustrates the measurement problem without participating in its solution, and the consequence is stated in Section 8.9.

8.3.9 Constants, Initialization and Numerics

Every constant appears in Table 8.2. Most are adjustable in the application, and the entries marked hidden appear only in the source code, which is deliberate because inferring a catalyst activation factor or a chain-transfer coefficient from a step test is part of the identification work.

The initial state is not hard-coded. Every scenario solves for the equilibrium of whatever setpoints and disturbances it declares, using two scalar bisections on monotone functions. The first finds the cooler outlet temperature at which $Q_{sens} + Q_{vap}$ equals the reaction heat less the wall loss, and the second finds the cooling-water valve position that produces exactly that outlet temperature through (8.10). Everything else follows in closed form. A built-in test evaluates all thirteen derivatives at the solved state and requires the largest relative rate to fall below 10^{-9}, so a scenario that starts at a different grade, level or production rate starts at a true equilibrium where the requested thermal line-up is feasible. If a learner chooses an infeasible disturbance before Run, the solve retains its closest bounded state and the app discloses the thermal mismatch. The disturbance remains a starting condition rather than becoming a hidden later step. A setpoint or disturbance chosen before Run is therefore part of the starting state, and the same setting carried through Reset gives the same run. A manual station set by hand before Run is treated differently. The start-up solve does not read it, so the plant starts lined out at the solved position and the edited station is applied when control begins, which survives Reset as well. A step meant to be watched on a setpoint is therefore made after Run, while a valve step on a station in Manual can be made before it.

Table 8.2 Model Constants. Hidden Constants Are Not Adjustable in the Application

Symbol	Meaning	Value	Unit
D	reactor diameter	5.0	m
H	bed height at 100% level	23.5	m
V_{tot}	gas and solids volume of reactor and loop	533.3	m^3
ρ_b	fluidized bulk density	400	kg/m^3
ρ_s	resin particle density	900	kg/m^3
d_p	mean particle size	700	µm
n_I	inert holdup	113899	mol
$k_{p,ref}$	propagation constant at 88 °C	54.79	m^3/min
E_p	propagation activation energy	14000	J/mol
k_{com}	comonomer rate enhancement	3.0	—
k_a	catalyst activation factor, hidden	1.41667×10^{-4}	1/(%·min)
k_{d0}	catalyst decay at 88 °C	0.0015	1/min
β_d	decay temperature sensitivity	0.05	1/K
K_z	poison sensitivity	1.6	—
ζ	chain-transfer H_2 use, hidden	0.005618	—
$-\Delta H_p$	heat of polymerization	94.5	kJ/mol
$F_{C_2}^{max}$	ethylene valve capacity	30000	mol/min
$F_{H_2}^{max}$	hydrogen valve capacity	1400	mol/min
$F_{C_6}^{max}$	1-hexene valve capacity	1200	mol/min
F_{B0}	bleed at 100% setting	1800	mol/min
$\dot{m}_{dis}^{max}$	discharge capacity	1200	kg/min
$\dot{V}_R$	compressor volumetric flow	686.4	m^3/min
$c_{p,g}$	recycle gas heat capacity	1.515	kJ/kg·K
ΔH_{vap}	condensate heat of vaporization	360	kJ/kg
UA_0	cooler UA at reference water flow	63139	kJ/min·K
F_{cw}^{max}	cooling water at 100% valve	80000	kg/min
$F_{cw,ref}$	reference water flow, hidden	40000	kg/min
n_{UA}	UA flow exponent, hidden	0.8	—
$c_{p,w}$	water heat capacity, hidden	4.19	kJ/kg·K
τ_{HX}	cooler thermal lag	0.4	min
UA_{wall}	vessel heat loss UA, hidden	300	kJ/min·K
T_{amb}	ambient temperature, hidden	25	°C
$c_{p,s}$	resin heat capacity, hidden	2.0	kJ/kg·K
C_{metal}	vessel steel heat capacity, hidden	40000	kJ/K

Symbol	Meaning	Value	Unit
$T_{dew,0}$	dew point at the design gas	43.287	°C
a_P	dew point per bar	0.9	°C/bar
a_I	dew point per mol% cyclohexane	1.3	°C/mol%
a_6	dew point per mol% 1-hexene	2.0	°C/mol%
a_2	dew point per mol% ethylene	−0.15	°C/mol%
$f_{L,max}$	maximum condensed fraction	0.25	—
ΔT_{ref}	subcooling scale	12	K
w	dew-point smoothing width, hidden	0.7	K
a_0	melt-index intercept	0.3287	—
a_H	melt index per H_2/C_2	3.244	—
a_C	melt index per C_6/C_2	1.5	—
a_z	melt index per impurity, hidden	0.8	—
a_T	melt-index temperature term	−5000	K
m	melt-index exponent, hidden	3.5	—
K_C	comonomer incorporation constant	1.709	—
ρ_{HD}	homopolymer density base	0.9585	g/cm^3
K_ρ	density per incorporated C_6	1.62	—
K_{MI}	density per ln MI, hidden	0.0025	—
T_{s0}	sticking base temperature	102.0	°C
k_ρ	sticking per 0.001 density	0.9	°C
k_I	sticking per mol% cyclohexane	0.35	°C
k_{MI}	sticking per ln MI	1.0	°C
P_{hh}	emergency vent pressure	23.5	bar
k_{vent}	vent capacity, hidden	9000	mol/min·bar
τ_V	cooling-water valve lag	0.25	min
τ_{Vd}	discharge system lag	0.4	min
τ_{feed}	catalyst feeder lag	5.0	min
t_{lab}	laboratory sample interval	60	min
θ_{lab}	laboratory turnaround	45	min
ε_{MI}	melt-index specification band	±10	%
ε_ρ	density specification band	±0.0015	g/cm^3

Solving rather than tabulating the design point produces the numbers in Table 8.3, and each of them can be checked by hand from Table 8.2.

Four training grades are provided and each pair of ratios was inverted out of (8.12), so loading a grade and letting the cascades line out reproduces its melt index and density to better than 2 % and 0.001 g/cm^3, which a built-in test enforces. Table 8.4 lists them. These are training grades calibrated to make the simulation behave like a polyethylene plant. They are not commercial specifications and no number in them is taken from any licensor's process.

A classical fourth-order Runge-Kutta step of $\Delta t = 0.05$ min advances the thirteen states, controllers execute every $\Delta t_c = 0.1$ min, and the historian logs four samples per minute of plant time. The order of work in one step and the separation of the clocks from the display are those of the shared engine in Section 1.6.1, with the post-step corrections below applied after the integration. The speed setting offers 1, 2, 5 and 10 simulated minutes per second of real time, with 5 as the default. Noise comes from a seeded mulberry32 generator with Box-Muller normal deviates, seed 20260825, so a scenario replays bit for bit.

Post-step corrections clamp the four molar holdups positive, the active catalyst inventory non-negative, bed mass to 2000 kg to 193798 kg which is 105 % level, the three actuator states to 0 % to 100 % and $\bar{x}$ to 0 to 0.2. The same pass integrates the off-grade accounting and services the laboratory sampler. Six annunciator conditions are raised and cleared there as well, namely sticking margin below 2 °C, the high-temperature interlock below 0.5 °C, the vent open, gas velocity below $1.6\,u_{mf}$, level above 92 % and level below 28 %.

Table 8.3 Solved Design Point for Film A at 65% Bed Level

Quantity	Value	Unit
Bed temperature, pressure	88.00, 21.70	°C, bar
Bed mass, level, height	119970, 65.0, 15.28	kg, %, m
Free gas volume V_g	400.0	m^3
Total gas holdup n_g	289083	mol
Gas $C_2/H_2/C_6$/ICA/inert	40.0/8.0/0.60/12.0/39.4	mol%
Propagation rate r_p	16550	mol/min
Production $\dot{m}_{PE}$	500.0 (30.00)	kg/min (t/h)
Residence time τ_s	240.0 (4.00)	min (h)
Recycle flow $\dot{m}_R$	16382	kg/min
Cooler outlet, dew point	40.00, 43.29	°C
Condensed fraction f_L	6.00	%
Reaction heat	26.07	MW
Sensible, latent, wall loss	19.85, 5.90, 0.32	MW
Heat-removal ceiling R_{max}	33.97	t/h
Cooling utilization	88.3	%
Melt index, density	1.000, 0.91801	g/10 min, g/cm^3
Sticking temperature, margin	96.01, 8.01	°C
u_g, u_{mf}, ratio	0.583, 0.182, 3.20	m/s
Valves CW/C_2/discharge	50.0, 57.6, 41.7	%
Valves catalyst/H_2/C_6	40.0, 11.6, 36.3	%

Table 8.4 Training Grade Slate. Not Commercial Specifications

Grade	MI g/10 min	ρ g/cm^3	Rate t/h	T_{SP} °C	$R_{H,SP}$	$R_{C,SP}$
A Film LLDPE	1.00	0.918	30	88	20.000	1.500
B Film LLDPE	2.00	0.925	32	88	26.843	1.301
C Molding MDPE	8.00	0.935	28	88	45.219	1.055
D HDPE	0.30	0.952	25	92	10.736	0.126

8.3.10 Modeling Assumptions

The gas phase is one perfectly mixed volume and the bed is another. Real fluidized beds have a bubble phase and an emulsion phase exchanging mass across a boundary, and a real reactor carries a vertical temperature and composition profile. The approximation is defensible because the recycle rate is large and the conversion per pass is a few percent, and it has been tested directly against two-phase models for industrial gas-phase conditions [8]. It stops holding near minimum fluidization, where the bed slumps and the two phases stop exchanging, which is exactly the region the fluidization override protects.

Molecular weight is represented by a single melt-index correlation rather than a population balance and catalyst is one lumped activity with one decay constant, so there is no molecular-weight distribution and no polydispersity. Anything that depends on the shape of the distribution, such as bimodal products, long-chain branching or rheology beyond melt index, is outside this model, and a real Ziegler-Natta or chromium catalyst has several site types with different reactivity ratios and different hydrogen responses.

Condensation is an effective dew-point correlation with a smooth subcooling curve rather than a rigorous flash. It gives the right qualitative couplings, because the dew point rises with pressure, with cyclohexane and with 1-hexene, but it will not give the right liquid composition and it cannot predict the onset of the entrainment and agglomeration problems that limit real condensed-mode operation [16].

Table 8.5 Dominant Time Constants Across the Plant, Evaluated at the Design Point

Path	Expression	Value min	Measured min
Cooler outlet	$\tau_V + \tau_{HX}$	0.65	0.5
1-hexene inventory	$n_{C_6}/(r_{C_6} + y_{C_6} F_B)$	4.0	—
Ethylene inventory	$n_{C_2}/(r_p + y_{C_2} F_B)$	6.7	—
Bed temperature	$(M_b c_{p,s} + C_{metal})/\dot{m}_R c_{p,g}$	11.3	12.0
Hydrogen inventory	$n_{H_2}/(r_{H_2} + y_{H_2} F_B)$	142	—
Active catalyst	$1/(\dot{m}_{out}/M_b + k_d)$	176	—
Resin property	$M_b/\dot{m}_{PE}$	240	—

The recycle loop has no transport delay, although a real reactor takes tens of seconds to move gas from the cooler outlet to the distributor plate, so a student identifying the inner loop from a step test here finds a pure lag where a plant would show a lag plus dead time. Nitrogen and the accumulated light ends are a fixed inert holdup replaced exactly as the bleed removes them, so inert accumulation never happens even though it is a real operating problem and part of why the bleed exists. Product discharge is a continuous valve rather than the pot-and-lock-hopper cycle that puts a saw-tooth on a real level measurement and a periodic disturbance on the pressure. Sensible heat carried in with cold feed gas and out with warm product is neglected, and both are small next to 26 MW of reaction heat but neither is zero. Every coefficient in the quality, sticking and condensation models was calibrated to make this simulation behave like a polyethylene plant, and none of them is plant data.

8.4 Dynamic Analysis

The separation of time scales is not asserted by the application. It is recoverable from step tests, and the scenario sequence is built so that a student measures the spread rather than reading it from a table. Table 8.5 collects the seven dominant time constants, each evaluated from Table 8.2 at the design point, with the two that a single step test can isolate reported beside them.

8.4.1 Inner Cooler Loop Identification

Scenario 2 opens with the cooler-outlet loop in MANUAL at 50 % on the water valve and its tuning deliberately detuned to $K_c = 1.0$ and $\tau_I = 8$ min. Stepping the valve from 50.0 % to 60.0 % at $t = 5$ min drives the cooler outlet from 40.00 °C to 38.54 °C within two minutes, a gain of $K_p = -0.146\,°\mathrm{C}/\%$. The response reaches 63.2 % of that change 0.51 min after the step, for any valve step from 1 to 20 points, faster than $\tau_V + \tau_{HX} = 0.65$ min suggests. A colder outlet condenses more liquid, raises φ and lowers the effectiveness, so T_c^* in (8.10e) rises 1.01 °C for each degree that T_c falls, which shortens the effective exchanger lag to about 0.2 min. This is by far the fastest loop on the plant, its gain is negative so the controller is direct acting and carries $a = -1$, meaning the water valve opens further as the cooler outlet rises, and both facts are why it is the inner half of a cascade. Those two numbers are also all a model-based rule needs to set the inner tuning, which is the route the scenario asks for [19].

The outlet temperature does not settle completely. It drifts a further 0.16 °C by $t = 20$ min because the bed temperature is falling underneath it, from 88.0 °C to 82.1 °C by $t = 20$ min, and the bed cools because the outer loop cannot act while the inner loop is in MANUAL. The slow tail is the outer loop's process appearing in the inner loop's measurement, and the step test must be read before the tail dominates the fit.

8.4.2 Outer Temperature Loop Identification

With the inner loop tightened to $K_c = 8.0$ and $\tau_I = 0.8$ min, the outer loop can be identified on its own. Taking the reactor temperature controller to MANUAL at $t = 11$ min leaves the cooler outlet setpoint where the controller last put it, at 39.93 °C, and an otherwise identical run without a step

shows the bed drifting from 88.00 °C to 87.66 °C by $t = 300$ min, so the gains below are read from the stepped run minus that unstepped run. Stepping the setpoint to 41.00 °C at $t = 11$ min drives the bed temperature through a peak of 93.47 °C at $t = 51.7$ min, a rise of 5.87 °C over the unstepped run and a peak gain of $+5.47\,°\text{C}/°\text{C}$. The gain then falls to $+4.40\,°\text{C}/°\text{C}$ at $t = 140$ min and recovers to $+4.75\,°\text{C}/°\text{C}$ at $t = 300$ min and $+5.38\,°\text{C}/°\text{C}$ at $t = 1500$ min. The response reaches 63.2 % of its peak rise 12.0 min after the step, close to the 11.3 min thermal lag of Table 8.5.

The dip is neither oscillation nor a numerical artifact. A hotter bed deactivates catalyst faster through k_d in (8.6), the active inventory falls over its 176 min time constant, and part of the temperature rise is given back. The production controller, still in AUTO, then raises the catalyst feeder over its 180 min integral time to restore 30 t/h and takes the gain back up. With the production loop also in MANUAL the catalyst path acts alone and the outer loop sees a lead-lag plant, with a peak gain of $+5.77\,°\text{C}/°\text{C}$ at $t = 58.8$ min that decays to $+4.47\,°\text{C}/°\text{C}$ at $t = 300$ min and $+4.63\,°\text{C}/°\text{C}$ at $t = 1500$ min, so the fast gain is 25 % larger than the slow one. A controller tuned on the peak is sluggish in the long run and one tuned on the slow gain overshoots in the short run.

The order in which the two loops are tuned matters more than either tuning. With the inner loop left at Scenario 2's $K_c = 1.0$ and $\tau_I = 8$ min, the suggested outer tuning of Table 8.6 overshoots a 1 °C setpoint step made at $t = 11$ min by only 0.29 °C, so the cascade appears well tuned. The warm cooling water of Scenario 8 then swings the bed from 86.6 °C to 90.4 °C, against 87.8 °C to 88.2 °C with the inner loop tightened to $K_c = 8.0$ and $\tau_I = 0.8$ min. A slow inner loop passes the cooling disturbance through to the bed before the outer loop can act on it, the same requirement on inner-loop speed that the fermenter cascade of Volume 3, Chapter 3 measures.

8.4.3 Integrating Level Dynamics

Scenario 4 puts the level loop in MANUAL and asks for a slope rather than a settling point. Stepping the discharge valve from 41.7 % to 47.0 % at $t = 5$ min raises the discharge from 500.0 kg/min to 564.0 kg/min, and between $t = 8$ min and 25 min the level falls at a constant −0.0351 %/min. The same run without the step drifts at +0.0003 %/min, so the difference gives an integrating gain of $K_I = -6.65 \times 10^{-3}\,\%/(\text{min}\,\%)$. The analytical value from (8.7b) and (8.7c) is $-\dot{m}_{dis}^{max}/(\rho_b A H) = -6.50 \times 10^{-3}\,\%/(\text{min}\,\%)$, and the difference is the small production shortfall while the bed drains, 498.4 kg/min rather than 500 kg/min by $t = 25$ min. There is no gain to read off a settling point because there is never a settling point.

Tight level control is not the objective. Downstream equipment is insensitive to whether the bed sits at 64 % or 66 % but is sensitive to variation in the discharge rate, so the loop is a classic averaging-level application and the shipped tuning is deliberately gentle at $K_c = 6$ and $\tau_I = 90$ min [1]. The scenario presents two upsets that produce similar level responses and demand different responses, a discharge valve partly plugged at $t = 25$ min, which is a sustained imbalance needing a permanent valve change, and a 4 t discharge surge at $t = 70$ min, which needs no valve change once the level has been returned to its setpoint.

8.4.4 Gas and Resin Time-Scale Separation

Scenario 5 demonstrates the time-scale separation that motivates the application, and it is the operating case shown in Figure 8.3. The melt-index controller is placed in MANUAL and its output, the hydrogen ratio setpoint, is stepped from 20.00 to 30.00 at $t = 20$ min. Four traces then separate on one time axis. The hydrogen valve opens immediately from 11.7 % to 23.7 %. The gas ratio reaches 26.90 by $t = 80$ min and 30.02 by $t = 200$ min, consistent with the 142 min hydrogen inventory. The instantaneous melt index follows the ratio almost without lag, because the instantaneous property depends only on the present gas composition, reaching 2.02 at $t = 80$ min and 2.68 at $t = 200$ min. The bed average moves far more slowly, reaching 1.15 at $t = 100$ min, 1.50 at $t = 200$ min and 1.99 at $t = 360$ min. The laboratory value lags further still, reporting 1.22 at $t = 200$ min and 1.83 at $t = 354$ min from a sample 54 min old.

The bed average is the trace the customer receives, and it drifted outside the $\pm 10\,\%$ band of the original melt index at $t = 81.8\,\text{min}$. By $t = 360\,\text{min}$ the run had discharged 136.8 t of off-specification resin out of 177.0 t made, 77 % of a six-hour shift, from one ratio move with no grade change ordered.

> **TRY IT IN THE APP** **5 · Fast Gas, Slow Resin**
>
> Load the scenario, press Run on the Quality tab and watch the three melt-index cards, Making now, In the bed and Laboratory, after the hydrogen ratio setpoint steps from 20 to 30 at 20 min. The resin being made reads 2.02 by 80 min, while at 200 min the bed average that leaves the reactor reads only 1.50 and the laboratory 1.22.

8.5 Controller Design

8.5.1 PID Implementation

All nine loops run the shared velocity-form PID of Section 1.6.2, Equations (1.14) to (1.17), with the derivative on the measurement and back-calculation on the integral alone [20]. Each executes every $\Delta t_c = 0.1\,\text{min}$ whatever mode the loop is in, with the error $e_k = a\,(\text{SP}_k - \text{PV}_k)$ and the loop action $a = \pm 1$ of Table 8.6. The measurement passes a first-order filter $\hat{y}_k = \hat{y}_{k-1}\alpha + \text{PV}_k(1-\alpha)$ with $\alpha = \mathrm{e}^{-\Delta t_c/\tau_f}$ and $\tau_f = 0.6\,\text{min}$ on all nine loops, and the tracking time is $T_t = \tau_I$. In the symbols of Chapter 1, the running sum S_k accumulates the proportional increment $K_c(e_k - e_{k-1})$ and the derivative increment $-aK_c\tau_D(\hat{y}_k - 2\hat{y}_{k-1} + \hat{y}_{k-2})/\Delta t_c$, the demanded output is $v_k = u_{bias,k} + S_k + I_{k-1}$, and u_k is v_k clipped to the output range. Because $T_t = \tau_I$, the proportional part of the demand cancels from the integral update while the output is clipped, and with no derivative action I relaxes toward $u_{lim} - u_{bias}$ with time constant τ_I whatever the error. Slew limits are unused, and derivative action is enabled only on the reactor temperature loop, where τ_D still defaults to zero.

MANUAL seeds the integral, $I_k = u_{man} - u_{bias,k} - S_k$, and in AUTO each station is set to its controller's output after every execution, which makes transfer bumpless in both directions as in Section 1.6.2. The header master switch throws all nine loops at once and, on the way to MANUAL, copies each controller's present output into its own station first. Every controller is initialized with u_{bias} equal to the output its own loop needs to hold the solved design point and with $I_0 = 0$, so a scenario opens with zero error and zero integral rather than with something quietly winding. Only the pressure loop's bias moves, because the feedforward below rewrites $u_{bias,k}$ at every execution. The temperature cascade carries an external reset. While the water valve sits at a stop and the outer controller asks for a cooler outlet beyond the one reached, its integral is set to $I_k = T_c - u_{bias,k}$, so its output is the measured outlet plus the proportional and derivative terms. In the cooling-fault run of Section 8.6 the setpoint therefore bottoms at 27.2 °C at $t = 271.9\,\text{min}$, when the bed is 7.4 °C above its setpoint, and the valve leaves its stop at $t = 342.4\,\text{min}$, as the bed returns to 88 °C. The two quality cascades carry none, and each of their outer controllers is held only by its own output range.

8.5.2 Loop Pairing and Cascade Structure

Figure 8.2 shows the nine pairings and the tier each belongs to. Two features of that map carry more weight than the count of loops.

The first is that pressure is held with ethylene and inventory with product discharge, and not the other way round. Both pairings are available on paper, because ethylene adds moles and therefore raises pressure, while discharge removes mass and therefore raises the free gas volume of (8.3) and lowers pressure. The chosen pairing points each loop at the handle with the direct, fast and large effect and leaves the cross-effects small, which is what a relative-gain analysis is for [21]. The interactions remain and they are visible. Bed level changes free gas volume and therefore pressure, bed mass changes thermal inertia and therefore the temperature response, and melt index enters the density correlation of (8.12b) directly.

The second is that, under Cascades, no polymer property is controlled directly with a valve. Melt index writes the hydrogen ratio setpoint and density writes the comonomer ratio setpoint, and the

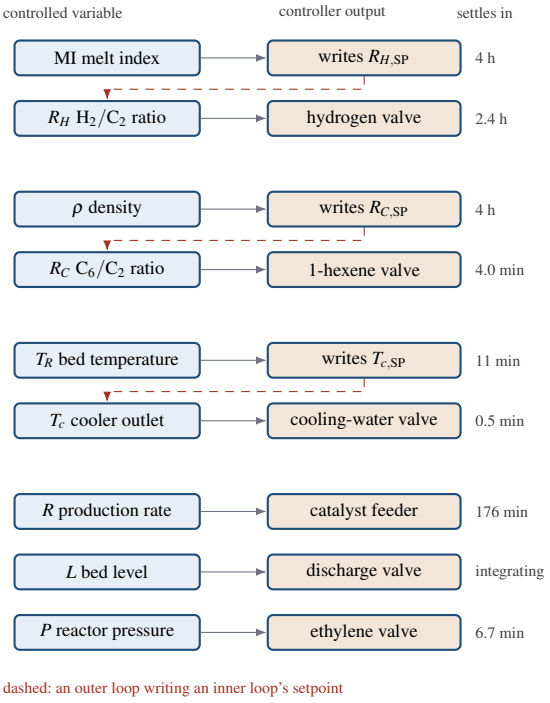

Figure 8.2 Loop pairing map. Three cascades are marked by the dashed links, in which the outer controller's output is the setpoint of the loop below it, and each pair spans roughly two decades of settling time.

Table 8.6 Controller Settings. Shipped Defaults and the Values the Graded Challenge Suggests

		Shipped		Suggested	
Loop	Action	K_c	τ_I min	K_c	τ_I min
Cooler outlet T_c	−1	4.0	1.2	8.0	0.8
Pressure P	+1	30	6	—	—
C_6/C_2 ratio	+1	60	5	—	—
Bed temperature T_R	+1	1.5	45	4.0	15
Bed level L	−1	6	90	—	—
H_2/C_2 ratio	+1	1.2	120	1.8	90
Production R	+1	2.5	180	5.0	120
Melt index MI	+1	7	400	18	260
Density ρ	−1	45	320	100	240

inner ratio loops track them, so a header pressure change, a bleed swing or a sticking valve is rejected before it reaches the quality loop [1]. This remains the default structure, because the property is a four-hour inventory and the chemistry can be reached only through the gas.

Pause. Try Exercise 8.2 before reading on.

Table 8.6 gives the shipped tuning and the tuning the graded challenge suggests. The suggestion also lowers the bed level setpoint from 65 % to 55 %, which Section 8.7 explains. The integral times range from 1.2 min on the cooler to 400 min on melt index, which restates the time-scale separation in tuning terms.

The reactor temperature loop is the one loop where derivative action is useful, because the process

Table 8.7 Non-Defeatable Overrides

Condition	Action	Label
$T_{stick} - T_R < 0.5\,^{\circ}\mathrm{C}$	$u_{cat} = 0,\ u_{CW} = 100$	high temperature
$T_{stick} - T_R < 2\,^{\circ}\mathrm{C}$	$u_{cat} \le 0.7\,u_{bias,cat}$	sticking margin
$u_g < 1.5\,u_{mf}$	$u_{cat} \le 20\,\%$	low fluidization
$L > 95\,\%$	$u_{dis} \ge 85\,\%$	bed level high
$L < 22\,\%$	$u_{dis} \le 10\,\%$	bed level low
$P > P_{hh}$	$v_{C_2} \le 5\,\%$	high pressure

is slow and the measurement noise is small at $\sigma = 0.06\,^{\circ}\mathrm{C}$. It ships with $\tau_D = 0$ and the suggested tuning sets 2.0 min. Raising the noise multiplier to four times the design level quantifies the cost of derivative action on a noisy measurement [20].

8.5.3 Feedforward on the Ethylene Valve

When the production target rises the reactor is about to consume more ethylene. Waiting for the pressure to fall before adding it leaves the pressure controller to find the extra ethylene by feedback alone. Instead the pressure controller's bias term follows the production inference, filtered with a three-minute time constant, as shown in Equation 8.16.

$$r_{p,ff} = \frac{10^6\,\widehat{R}/60}{M_{C_2} + M_{C_6}\,x_{inst}/(1 - x_{inst})} \tag{8.16a}$$

$$u_{bias,P} = \frac{100}{F_{C_2}^{max}}\left(r_{p,ff} + y_{C_2} F_{B0}\,\frac{\phi_B}{100}\right) \tag{8.16b}$$

where $\widehat{R}$ is the filtered production inference in tonnes per hour, the molar masses are in g/mol, and $r_{p,ff}$ is the propagation rate that inference implies in mol/min, which is 16550 mol/min at the design rate of 30 t/h. The pressure controller then trims the difference. The feedforward is on by default and can be disabled [1]. It follows production as the catalyst inventory changes rather than applying the eventual target's ethylene demand immediately. Pressure feedback corrects inference error. The three-minute filter is short relative to the catalyst's three-hour response.

8.5.4 Protective Overrides

Six protective actions sit between the controllers and the valves, listed in Table 8.7. They are not tuning and they cannot be disabled, because what they prevent ends a run rather than degrading it. When one acts, the loop whose handle it took is immediately told what was actually applied through the same tracking mechanism that MANUAL uses, so its integral does not wind up while the override holds. The active override is recorded in the exported data.

8.5.5 Model Predictive Control

The optional predictive controller follows the pattern of Section 1.6.3 and replaces only the quality and inventory loops. It moves the hydrogen, 1-hexene and discharge valves, and the pressure, temperature and production loops stay in charge of their own valves. Its six states are the gas ratios $h = 100 n_{H_2}/n_{C_2}$ and $c = 100 n_{C_6}/n_{C_2}$, the estimates $\hat{q}$ and $\hat{x}$ of the bed-average quality, the level L and the discharge valve position $\hat{v}_{dis}$. The local nominal balances are

$$\begin{aligned} \dot{\eta}_i &= F_i^{max} u_i/\hat{n}_2 - (k_i \hat{r}_p + F_{B0}\hat{y}_2)\eta_i/\hat{n}_2, \\ \dot{\hat{q}} &= a(q_{inst} - \hat{q}), \quad \dot{\hat{x}} = a(x_{inst} - \hat{x}), \\ \dot{L} &= (100\hat{m} - \dot{m}_{dis}^{max}\hat{v}_{dis})/M_{100}, \\ \dot{\hat{v}}_{dis} &= (u_{dis} - \hat{v}_{dis})/\tau_{Vd}. \end{aligned} \tag{8.17}$$

where $i = H, C$, $\eta_H = h$, $\eta_C = c$, $k_H = \zeta$, $k_C = K_C$, $a = \hat{m}/(M_{100}L/100)$, $M_{100} = \rho_b AH$, the u_i are valve percentages and hats denote estimates. The quality correlations receive $h/100$ and $c/100$ and assume zero impurity.

The ethylene storage $\hat{n}_2$ and its mole fraction $\hat{y}_2$ come from the measured pressure, temperature, ratios and level with the nominal geometry and the design inert and cyclohexane inventories. The controller's own ethylene command through the nominal valve capacity, less the nominal bleed and the change in storage, gives the propagation estimate $\hat{r}_p$ and from it the production $\hat{m}$ in kg/min, filtered over 3 min and bounded at 50–1000. The declared coefficients are frozen and exclude the hidden production and resin states. The quality maps are differentiated at the measured ratios with storage and rate held, the affine offsets are retained, and the model is discretized by an exact zero-order hold. Six noise-free ± 2 valve-point tests over 120 min give maximum RMS residuals across tests of 0.247 and 0.0128 ratio points, 0.00530 in log-MI, 0.0000586 g/cm^3 in density and 0.118 level points. These residuals cover local steps only.

The quality estimates propagate every 0.1 min. Each laboratory result corrects the log-MI and density biases against the estimate at the sample time, as in Equation (1.11), with gain 0.5 and bounds of 2 and 0.02 g/cm^3, respectively [22]. This is a simpler form of the disturbance estimate that the offset-free controller of Volume 9, Chapter 4 carries in a Kalman filter. The program decides every 5 min, including initialization, over 120 min ($N = 24$), with move blocks $[1, 1, 2, 4, 6, 10]$, and minimizes

$$J = \sum_{j=1}^{24} w_j \| y_j - r_j \|_Q^2 + 0.03 \sum_{b=1}^{6} \| \Delta u_b \|^2 + s_L^2 + s_H^2 + 100 s_C^2, \tag{8.18}$$

where $y = (\hat{q}, 1000(\hat{\rho} - \rho_{HD}), L)$ and r holds the corresponding targets. The weights are

$$Q = (5/240)\mathrm{diag}(\ln^{-2}(1.1), 1/1.5^2, 1/25),$$

with $w_{24} = 6$ (the stage weight plus a $5Q$ terminal weight) and $w_j = 1$ otherwise. The valves are bounded hard at 0–100%, with block-move caps of 35, 25 and 10 points. Three nonnegative slacks, each shared over the horizon, soften the level band 28–92%, h at 4–80 and c at 0.02–2.5, so the program has 18 valve variables and three slacks. The first move is applied [23].

With the grade preview On the plan reads the announced grade targets, and with it Off it holds the present targets. A 32-point sweep on the new case only selected the move weight $\lambda = 3$, the gentlest setting within 10% of the best off-grade result, 85.254 versus 81.016 t, and it was not retuned afterward. The five PI loops the controller displaces track the applied valves and the measured ratio setpoints. The protections and any loop the reader holds in Manual stay in force, and the controller stands by when it is not eligible. A plan that is not optimal, not finite or outside its bounds holds the first failed scan and then hands the valves to the tracked cascades until a plan succeeds.

8.6 Results

Table 8.8 lists the eleven shipped scenarios. Each one starts from a solved steady state and carries a guided sequence that asks for a prediction before the run and an explanation after it.

8.6.1 Production Rate Increase

Scenario 6 starts at 25 t/h and steps the production target to 34 t/h at $t = 30$ min. The catalyst feeder command moves immediately from 29.2 % to 51.7 %. The ethylene bias follows the filtered production inference, so the make-up valve remains near 48.3 % at the step, reads 47.9 % one minute later and 62.1 % at $t = 360$ min. Between $t = 25$ min and 150 min, pressure spans 21.68 bar to 21.72 bar. With feedforward switched off, the same window spans 21.65 bar to 21.71 bar. The previous target-based bias opened hours before the gas was consumed and widened that range to 21.65 bar to 22.08 bar. Everything downstream is slow. The active catalyst inventory climbs from 83.3 % to 100.5 % by $t = 120$ min and 108.7 % by $t = 360$ min, and the production rate reaches 30.1 t/h and 32.6 t/h at those two times without arriving at the target within the run.

Table 8.8 Shipped Scenarios and Their Timed Events

Id	Purpose	Events
lined-out	read a steady plant	none
tune-cooler	inner loop step test	loop starts in MANUAL, de-tuned
tune-temp	outer loop after inner	inner pre-tuned to 8.0/0.8
level-integrator	integrating dynamics	valve plugged 12% at 25 min, 4 t surge at 70 min
gas-vs-resin	the two time scales	$R_{H,\mathrm{SP}}$ 20 to 30 at 20 min
rate-up	rate into the heat limit	R_{SP} 25 to 34 t/h at 30 min
condensed	lose the condensing agent	F_{ICA} to 40 mol/min at 20 min, restored at 150 min
cw-upset	three cooling faults	water +6 °C at 40 min, fouling +15% at 140 min, compressor to 82% at 240 min
poison	diagnose an upset	impurity to 0.28 ppm-eq at 45 min, clean at 260 min
mpc-grade	MPC comparison	prepared 55% bed, grade A to C at 90 min
grade-change	the graded challenge	grade A to C at 90 min

The heat balance carries the larger changes. Reaction heat rises from 21.7 MW to 28.4 MW, the water valve opens from 37.0 % to 58.8 %, the cooler outlet falls from 42.2 °C to 38.7 °C, the condensed fraction rises from 2.5 % to 7.9 %, and the latent duty grows from 2.4 MW to 7.7 MW. The latent term, not the sensible term, supplies the added capacity. Cooling utilization moves from 83.3 % to 91.1 %, so the plant is still inside its ceiling. Pushing the target on to 45 t/h at $t = 150$ min finds the limit. The water valve reaches its stop at $t = 381$ min, the bed then warms to 94.0 °C, and at $t = 720.8$ min the sticking-margin override caps the catalyst feeder and production falls from 43.5 t/h to 34.4 t/h.

8.6.2 Loss of the Condensing Agent

Scenario 7 cuts the cyclohexane feed from 216 mol/min to 40 mol/min at $t = 20$ min and restores it at $t = 150$ min. Cyclohexane takes no part in the chemistry, so the response is entirely thermal.

The cyclohexane mole fraction falls from 12.0 % to 6.63 % by $t = 148$ min, because its only route out is the bleed and its inventory time constant is 161 min. The dew point falls with it from 43.29 °C to 35.76 °C. The condensed fraction barely moves, reading 6.00 % at the start and 6.23 % at $t = 148$ min, because the temperature controller drives the cooler outlet down from 40.00 °C to 32.33 °C to follow the falling dew point and condenses very nearly the same fraction at a lower temperature.

The cost appears in two quantities and neither is a controlled variable. The cooling-water valve climbs from 50.1 % just before the cut to a peak of 88.9 %, and the heat-removal ceiling drops from 33.97 t/h to 30.32 t/h. Production does not hold at 30 meanwhile. The cyclohexane leaving the gas raises the ethylene fraction from 40.0 % to 44.6 % at constant pressure, and production drifts up to 31.7 t/h although its controller cuts the catalyst feeder from 39.9 % to 34.5 %, so cooling utilization passes through 100 % and reaches 101.1 %. The bed temperature moves last and least, rising only from 88.0 °C to a peak of 89.7 °C. That ordering is the general signature of an approaching constraint on a well-controlled plant, because the controlled variable is the last place a constraint appears and the valve position is the leading indicator.

8.6.3 Unmeasured Cooling Faults

Scenario 8 applies three cooling faults in sequence at 30 t/h. None of the three has a measurement.

Warmer cooling water at $t = 40$ min takes the inlet from 30 °C to 36 °C, the water valve from 49.8 % to 77.5 %, the ceiling from 33.88 t/h to 30.36 t/h and utilization from 88.3 % to 98.8 %, read

just before this fault and the next. The bed temperature peaks only 1.1 °C high at $t = 43$ min, is within 0.11 °C of setpoint from $t = 60$ min, and the melt index stays between 0.998 and 1.000. Fifteen percent of fouling at $t = 140$ min repeats the pattern, taking the valve to 93.4 %, the ceiling to 29.63 t/h and utilization to 101.2 % by $t = 240$ min, with a bed temperature transient of 0.6 °C.

The compressor stepping down to 82 % at $t = 240$ min is different in kind. The water valve reaches its stop at $t = 243$ min and stays there until $t = 342$ min, saturated for 27.6 % of the run by the logged flag, so the bed temperature peaks at 95.4 °C at $t = 272.3$ min, the sticking margin closes from 8.0 °C to 0.54 °C, just above the 0.5 °C interlock (0.54 on the default seed and 0.61 to 0.72 °C over six additional noise seeds), and from $t = 258$ min the sticking-margin override caps the catalyst feeder and pulls production to 27.0 t/h. The same fault also takes the fluidization ratio from 3.20 to 2.62, which is the constraint that would have bind next. The valve position gives the warning the temperature does not. Lowering the production setpoint to 27 t/h when the fouling arrives keeps the margin above 5.7 °C, while the same cut made at $t = 240$ min lets it fall to 1.7 °C.

> **TRY IT IN THE APP** **8 · Warm Water and a Fouling Cooler**
>
> Press Run on the Heat removal tab and watch the Cooling water valve rather than the bed temperature. Warmer water at 40 min opens the valve from 49.8 to 77.5 % while the bed peaks only 1.1 °C high. After the compressor slows at 240 min the valve reaches its stop, and the sticking margin closes to 0.54 °C.

Pause. Try Exercise 8.3 before reading on.

8.6.4 Feed Impurity Diagnosis

Scenario 9 introduces 0.28 ppm equivalent of feed impurity at $t = 45$ min without saying so. Production falls at once from 29.9 t/h to 20.8 t/h and reaction heat from 26.0 MW to 18.0 MW, and within eight minutes the ethylene valve closes from 57.1 % to 39.0 % and the cooling-water valve from 49.8 % to 26.6 %, while the bed temperature dips to 85.0 °C at $t = 49.0$ min and then returns. Falling production with a closing ethylene valve is the distinguishing signature, because a supply problem starves the reactor and opens the valve while a catalyst problem stops the reactor asking for ethylene.

The quality response is less intuitive. The impurity term sits inside the melt-index correlation through a_z in (8.12a), so the instantaneous melt index rises from 1.00 to 2.01 with no hydrogen move behind it, and the melt-index controller responds by pulling the hydrogen ratio setpoint from 20.0 down to 16.6 to compensate for a disturbance that did not enter through hydrogen. Meanwhile the production controller has raised the catalyst feeder to 54 % and the active inventory to 137 % of design. When the feed cleans up at $t = 260$ min all of that catalyst acts at once. Production jumps from 28.9 t/h to 41.6 t/h, cooling utilization reads 124 % until the water valve catches up four minutes later, the bed reaches 93.1 °C at $t = 265.5$ min, and the hydrogen correction then has to be unwound. The run ends with 122.3 t of off-specification resin from an unannounced disturbance.

8.7 Grade Transition Case Study

8.7.1 Off-Grade Accounting

Resin is on grade when both properties sit inside their bands, as defined in Equation 8.19, and the accounting is integrated at the discharge valve on every physics step.

$$\text{on grade} \iff |\mathrm{MI}_{bed} - \mathrm{MI}_{\mathrm{SP}}| \le \varepsilon_{MI}\mathrm{MI}_{\mathrm{SP}} \quad \text{and} \quad |\rho_{bed} - \rho_{\mathrm{SP}}| \le \varepsilon_{\rho} \tag{8.19a}$$

$$W_{off} = \int \frac{\dot{m}_{out}}{1000}\left[1 - \mathbf{1}_{\text{on grade}}\right] \mathrm{d}t \tag{8.19b}$$



Figure 8.3 The gas-vs-resin scenario as shipped at $t = 354.4\,\text{min}$, on the Quality view. A single step of the hydrogen ratio setpoint from 20 to 30 at $t = 20\,\text{min}$ (4) has driven the instantaneous melt index (1) to 2.69 while the bed average that leaves the discharge valve (2) has reached only 1.98 and the laboratory (3), sampling a 54 min old specimen, reports 1.83. (5) The accumulated off-specification resin is 134.1 t out of 174.2 t made.

where $\varepsilon_{MI} = 10\,\%$ of the melt-index setpoint and $\varepsilon_\rho = 0.0015\,\text{g/cm}^3$, and W_{off} is in tonnes. Counting at the discharge valve rather than in the bed is a design decision and it changes the ranking of transition strategies. A metric evaluated on the bed average alone would reward emptying the reactor. The one that measures what actually reaches the silo does not.

8.7.2 Independence From Production Rate

The derivation repeats that of Section 1.4.5 in the symbols of this reactor. Take the bed mass and the production rate as constant through a transition and let q_{inst} be the logarithm of the instantaneous melt index that the gas composition is holding. Equation 8.13 is then a first-order lag of time constant τ_s, and the bed average crosses the specification edge q_{spec} at the time given by Equation 8.20.

$$\bar{q}(t) = q_{inst} - (q_{inst} - q_0)\mathrm{e}^{-t/\tau_s} \tag{8.20a}$$

$$t_{off} = \tau_s \ln \frac{q_{inst} - q_0}{q_{inst} - q_{spec}} \tag{8.20b}$$

$$W_{off} = \dot{m}_{out} t_{off} = M_b \ln \frac{q_{inst} - q_0}{q_{inst} - q_{spec}} \tag{8.20c}$$

where q_0 is the logarithm of the starting melt index, the last step uses $\tau_s = M_b/\dot{m}_{out}$ from (8.7c), and W_{off} in (8.20c) carries the units of M_b rather than the tonnes of (8.19b). The production rate cancels exactly. Halving the rate halves the tonnage at risk per hour and doubles the residence time in the same proportion, and the product is unchanged. Only two quantities are left. One is the bed mass, which the level setpoint controls. The other is how far past the target the resin being made can be driven and held, which is what the outer loop's gain and output limits decide. The loop reactor of Chapter 9 averages its product over an inventory that turns over in less than an hour, and there the hydrogen overshoot that shortens a transition is limited by the bubble-point margin it spends.

Equation 8.20c is a scaling argument rather than a prediction, because q_{inst} is not a step and the density band binds as well. It nevertheless gets close. For the Film A to Molding C change with

Table 8.9 Film A to Molding C at $t = 90$ min, Measured at $t = 900$ min. Off-Spec Is the Time From the Change to the Last Off-Specification Sample

Preparation	Off-grade t	Off-spec min	Peak R_H	Peak MI_{inst}
Shipped tuning, bed 65%	179.1	399.5	73.8	34.1
Retuned cascade, bed 65%	94.5	219.5	83.3	49.2
Retuned, 20 t/h until change	93.0	233.3	83.0	48.6
Retuned, 20 t/h throughout	128.7	404.5	83.4	49.4
Retuned, bed 55%	80.9	191.0	83.7	50.2
Retuned, bed 50%	84.9	235.0	84.0	50.6
Retuned, bed 40%	87.7	222.8	84.5	52.9
Retuned, drawn to 55%	77.5	189.0	84.2	51.0
Retuned, drawn to 50%	81.4	233.8	84.7	52.0
Retuned, drawn to 40%	83.6	223.5	86.0	56.4

$q_0 = 0$ and $q_{spec} = \ln 7.2 = 1.974$, and with the retuned quality cascade holding the instantaneous melt index near 47, the mean of $\ln \mathrm{MI}_{inst}$ from $t = 150$ min to 250 min being $q_{inst} = 3.85$, (8.20c) predicts 86.3 t at the 120.1 t of bed present at the change in the 65 % run and 70.9 t at the 98.6 t left by a drawdown to 55 %, against measured values of 94.5 t and 77.5 t. The first excursion of every retuned run in Table 8.9 costs between 0.77 and 0.84 t of off-grade resin per tonne of bed at the change, against the 0.72 that the logarithm gives, whatever the level or the production rate.

> **+ DID YOU KNOW? Running Slower Through the Change Costs More**
>
> From the simulator: with the retuned cascade, holding production at 20 t/h through the Film A to Molding C change leaves 128.7 t of off-grade resin, against 94.5 t at full rate. The slower bed carries the melt index past its band a second time (Table 8.9).

8.7.3 Ten Runs of One Transition

The graded challenge, Scenario 11, loads Film A, lets the plant line out, and changes the grade target to Molding C at $t = 90$ min. That is a factor of eight in melt index and 0.017 g/cm^3 in density, with ninety minutes of warning in which to prepare, and the grade change also writes the Molding C production target of 28 t/h. The challenge is judged at $t = 360$ min, with the model parameters as given. It passes when Molding C is on specification (MI 8.0 ± 10 %, density 0.935 ± 0.0015 g/cm^3), less than 90 t of off-grade resin has been discharged and the sticking margin never fell below 2 °C. Table 8.9 reports ten runs of the same event, each carried to $t = 900$ min so that the off-grade total has stopped accumulating. Every preparation is made before Run on the scenario as loaded, which makes it part of the starting state. A level chosen before Run starts the bed lined out at that level and a rate starts the plant at that rate, the same run that Reset gives with the setting carried. The three drawn rows instead lower the level setpoint at $t = 1$ min, so the bed is drawn down from 65 % during the warning. The retuned runs use the suggested tuning of Table 8.6 with the bed level setpoint at 65 % unless the table names another level, so the bed 55 % row is the suggestion exactly. Figure 8.4 draws the table's off-grade and off-specification columns as bars.

The shipped tuning costs 179.1 t and keeps the plant off specification for 399.5 min. Its failure is visible in the peak hydrogen ratio, which reaches only 73.8 against an available limit of 80. A melt-index controller with $K_c = 7$ and $\tau_I = 400$ min raises its output, the hydrogen ratio setpoint, no higher than 73.8, which it reaches at $t = 174$ min, 84 min after the change, and it never uses the limit it is permitted. The unused driving force is exactly the logarithm argument in (8.20c).

With the shipped settings the challenge fails at $t = 360$ min, with the bed-average melt index at 6.59, which is 0.61 below the lower band edge of 7.2, and 115.2 t of off-grade resin. The suggested

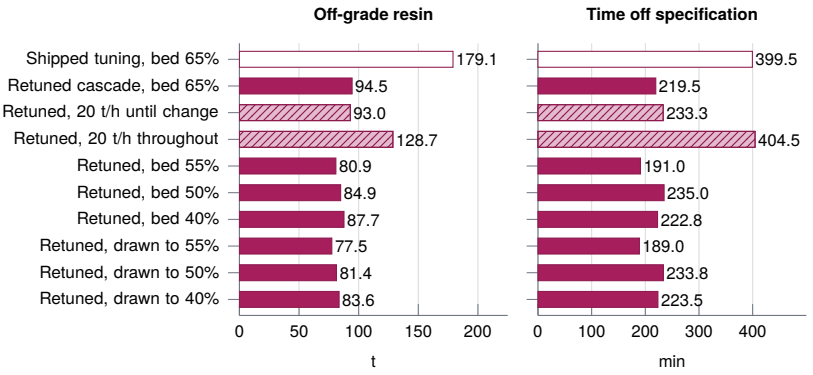

Figure 8.4 The ten runs of Table 8.9 as bars, off-grade resin on the left and the time from the change to the last off-specification sample on the right. Open bar: shipped tuning. Hatched bars: the two rate cuts. Solid bars: the retuned cascade at a chosen bed level. Notice that retuning removes about half of the loss, that every bed prepared below 65 % beats the retuned run at 65 %, and that holding the rate cut through the change is the worst of the retuned runs.

settings, the retuned cascade of Table 8.6 with the bed level setpoint at 55 %, start the bed lined out at 55 % and pass with MI 8.40, density 0.9352, 80.9 t of off-grade resin, 9.1 t inside the limit, and a lowest sticking margin of 7.8 °C. The run is the same whether the settings are applied before Run or carried through Reset.

Retuning the two quality controllers and the loops beneath them to the values in Table 8.6 takes the loss to 94.5 t, a reduction of 47 %, and shortens the excursion to 219.5 min. The retuned outer loop saturates the hydrogen ratio setpoint at its upper limit of 80 at its first execution after the change and holds it there until $t = 295.4$ min, for 3.4 h, and the density controller holds the comonomer ratio setpoint at its lower limit of 0.02 until $t = 124.0$ min. Both saturated outputs are productive, because an outer loop that can demand far more than the final grade requires is the mechanism that puts extra driving force into the transition.

Cutting production to 20 t/h before Run starts the plant at that rate only until the grade change writes 28 t/h, and it changes the loss from 94.5 t to 93.0 t, 2 % less. Holding 20 t/h through the change, which takes a second entry just after $t = 90$ min, leaves 92.9 t in the first excursion, as the cancellation predicts. The residence time is then about six hours, and the melt-index loop tuned for a four-hour bed carries the bed average to 8.89, past the upper band edge, so a second excursion adds 35.8 t and the total reaches 128.7 t. The rate reduction that operating intuition suggests first is the least effective preparation made before the change and the most costly one held through it.

A smaller bed at the change is effective, and the optimum level is not the lowest one available. With the bed started at 55 % the loss falls to 80.9 t and the excursion to 191.0 min, a 14 % improvement over retuning alone. Going further reverses the gain, with 50 % giving 84.9 t and 40 % giving 87.7 t even though the 40 % case has the shortest first excursion of the three, 148.8 min against 177.0 min and 191.1 min. A smaller bed turns over faster, and the melt-index loop tuned for the full bed carries the bed average past the upper band edge after reaching grade, to 8.87 at 50 % and 9.31 at 40 %, so second excursions add 10.5 t and 25.6 t. The 40 % bed also has less thermal inertia and catalyst, and cooling utilization reaches 122.4 %. Drawing the bed down in the first ninety minutes instead does slightly better at every level, 77.5 t, 81.4 t and 83.6 t, because the level overshoots its setpoint and leaves 98.6 t of bed rather than 101.6 t at the change in the 55 % case. The drawdown itself costs nothing, because it runs the discharge at up to the valve's 1200 kg/min capacity while the bed still holds on-grade Film A. Over five other noise seeds the bed started at 55, 50 and 40 % costs 81.0 t to 81.4 t, 84.2 t to 86.3 t and 86.0 t to 87.1 t, and drawn down to them 77.5 t to 78.1 t, 80.3 t to 82.4 t and

Table 8.10 Matched 900 min Comparisons, Seed 20260825. On/Off Denotes Announced-Target Preview. Costs Are Node Decision Times

Case	Control	Off-grade t	Before 90 t	Last off min	Valve travel total points	Outside bands min	Min. margin °C	Solves/fails	Mean cost ms
Shipped	Cascades	179.074	0.000	489.60	3978.7	0.00	7.917	0/0	—
Shipped	MPC On	93.845	23.975	261.75	217.7	95.35	8.008	180/0	0.257
Shipped	MPC Off	106.955	0.000	350.35	255.3	130.15	7.903	180/0	0.215
Prepared	Cascades	80.931	0.000	281.15	4124.7	121.50	7.755	0/0	—
Prepared	MPC On	85.254	24.956	226.80	202.9	38.15	7.938	180/0	0.311
Prepared	MPC Off	96.343	0.000	312.55	254.7	56.05	7.760	180/0	0.222
Drawdown	Cascades	77.498	0.000	279.20	4183.3	122.00	7.877	0/0	—
Drawdown	MPC On	103.944	32.009	235.60	206.5	26.15	7.946	180/0	0.228
Drawdown	MPC Off	106.066	0.000	316.80	270.8	78.05	7.826	180/0	0.230
Impurity	Cascades	122.258	0.000	342.90	3870.7	0.00	2.968	0/0	—
Impurity	MPC On	131.199	0.000	361.80	111.8	0.00	3.453	180/0	0.206
Impurity	MPC Off	131.199	0.000	361.80	111.8	0.00	3.453	180/0	0.206

82.0 t to 83.3 t, so the reversal is robust while the order of the two lower levels is within the noise. The best of the ten runs is a 57 % reduction from the shipped baseline, and both of the levers that produced it are tuning and preparation rather than equipment.

8.7.4 Predictive and Cascade Transition Control

Table 8.10 compares the cascades with the predictive controller of Section 8.5.5 with the grade preview On and Off. Within each triple of rows the initial state, seed, noise, instruments and retained loops are the same. The prepared case, Scenario 10, uses the suggested settings with the bed started at 55%. The drawdown case requests the level move from 65 to 55% at 1 min. Preview Off matches the cascades' grade information, although the cascades read bed quality without delay (Section 8.9).

On the shipped case the predictive controller reduces the loss from 179.074 to 93.845 t, and to 106.955 t with preview Off. It loses to the prepared cascades, 80.931 against 85.254 and 96.343 t, and to the drawdown, 77.498 against 103.944 and 106.066 t. Impurity is also a loss, 122.258 against 131.199 t with either preview setting. Preview reduces the prepared loss from 96.343 to 85.254 t but spends 24.956 t before 90 min, as Figure 8.5 shows: the plan starts moving the gas before the change, while the resin leaving is still Film A. Its earlier final excursion still costs more than the cascades, so an earlier last off-specification sample does not mean less off-grade resin.

Valve travel sums the scorecard travel totals of the physical valves at 0.25 min intervals. The time outside the soft gas and level envelopes falls from 121.50 to 38.15 min in the prepared case and from 122.00 to 26.15 min in the drawdown case with preview On. These envelopes are not protective limits. Paired production differs by less than 0.18 t. The shipped lined-out, rate-step and cooling-upset comparisons each produce zero off-grade under all three variants, and the minimum cooling-upset margin is 0.541/0.516/0.516 °C for Cascades/On/Off. An independent integration of the discharge reproduces the totals, and halving the physics step preserves the rankings within 0.014 t. In the Chrome browser a decision costs 1.386/0.706 ms at 1×/10× over 31/16 s, which occupies 0.0313/0.1455% of wall time with 7/33 plans, zero failures and 60 frames/s.

8.8 Technology in Practice

What Measures It

Figure 8.6 draws the loops of Section 8.5 with the hardware of a fluidized-bed reactor in place of the schematic's symbols. The pressure drop through a fluidized bed is equal to or slightly greater than the weight of the bed divided by its cross-sectional area [18], so a differential-pressure transmitter across the bed reads bed mass, as a liquid column's pressure reads its height [24]. The model's level, computed from bed mass with 0.35 % of noise, stands for that reading. Impulse lines prone to plugging are kept clear by a small, steady purge of clean fluid, often a gas such as nitrogen. Too much

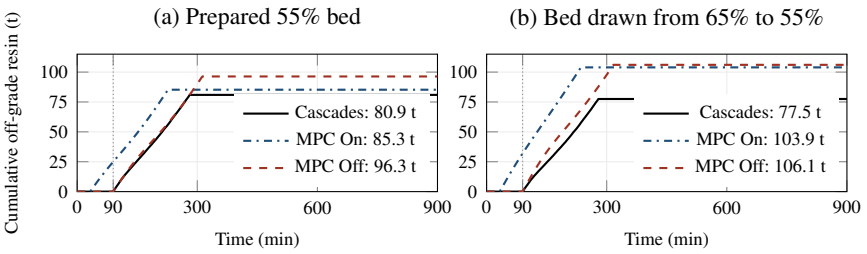

Figure 8.5 Full-duration economic comparisons over all 900 min. Cascades are the thick solid line, MPC with preview On the dash-dotted line and MPC with preview Off the dashed line. The dotted line marks the grade change at 90 min. Preview spends off-grade resin before the event and loses to both prepared cascade strategies over 900 min.

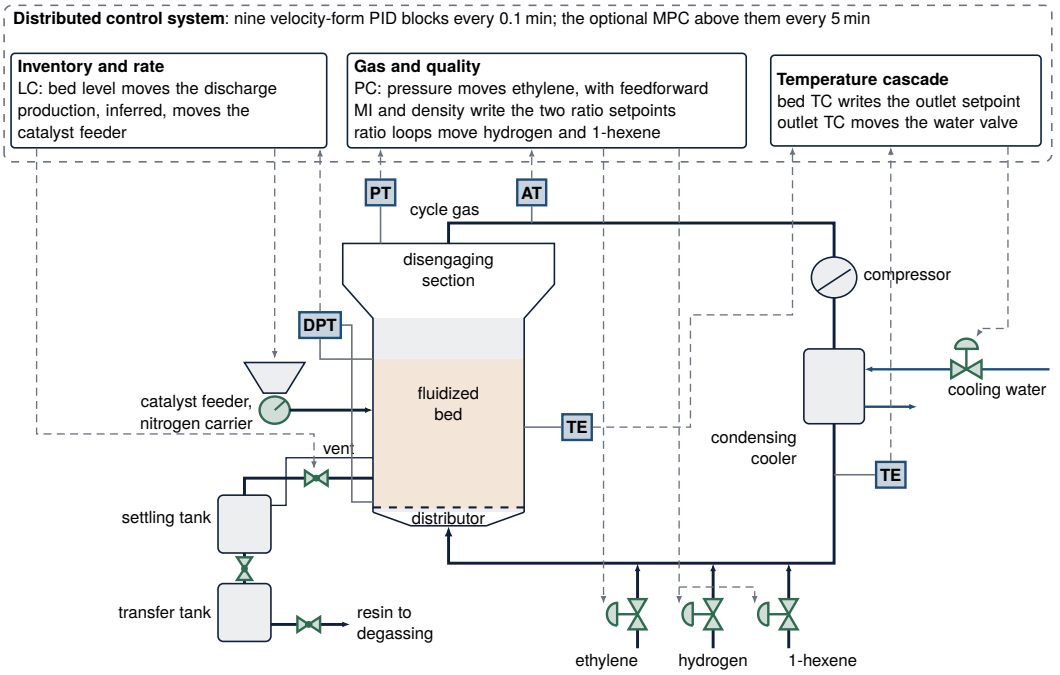

Figure 8.6 The reactor's main loops with their hardware. Blue tags are sensing devices: the bed and cooler-outlet temperature elements TE, the pressure transmitter PT, the differential-pressure transmitter DPT across the bed and the gas analyzer AT. Green symbols are final elements: the cooling-water, ethylene and hydrogen valves, the catalyst feeder and the on/off valves of the discharge tanks. Dashed arrows are signals. Notice that only the cooler outlet, the pressure and the ratio loops touch a valve directly. Melt index and bed temperature act through setpoints, and the laboratory reaches the plant only through the predictive controller's bias update. Not to scale, and not a piping diagram.

purge flow adds its own pressure drop and reads high, and a purge that stops lets the line plug [24]. A gas analyzer on the recycle stream sets the make-up composition that holds the gas steady [18], the instrument behind the model's two ratio measurements. Melt index itself comes from the laboratory, the subject of the box below.

> **TECHNOLOGY** **Sensor: The Melt Flow Test**
>
> Melt index is the mass of molten polymer extruded through a standardized capillary in ten minutes under a prescribed temperature and load, as ISO 1133-1 and ASTM D1238 define it. The test is simple and empirical, yet it classifies commercial grades and releases production lots. A sample has to be carried to the quality laboratory, conditioned, tested, checked and reported before an engineer sees the result, so the result describes the past [25]. A few units fit online melt flow analyzers to report sooner, but these need more equipment investment, read less accurately, and need their sampling molds replaced often, which interrupts production [26]. The chapter's laboratory, with a sample every 60 min and a result 45 min later, is that chain in short form.

What Moves It

Catalyst is carried into the bed by an inert gas such as nitrogen, at a point above the distributor plate, because catalyst injected below the plate could start polymerization there and plug it. The production rate is controlled by adjusting the rate of catalyst injection, and the temperature of the recycle stream entering the reactor is moved up and down to follow the rate of heat generation [18]. Those are the production loop and the temperature cascade of Figure 8.2. The model gives the feeder a 5.0 min lag and treats the discharge as a continuous valve with a 0.4 min lag, where a plant moves resin in batches.

> **TECHNOLOGY** **Final Element: The Discharge Valves and Their Tanks**
>
> Resin leaves a fluidized bed in batches. In one discharge process a settling tank below the bed is joined to it by a solids line near the distributor plate and by a vent line near the top of the bed. Opening the valves sends resin and gas into the tank, the gas vents back to the reactor, and the valves close. The resin then drops into a transfer tank, the pressures equalize, and an outlet valve passes it on at lower pressure. The tanks are sized by the discharge rate and the batch cycle time, and too large a transfer tank wastes reactor gas [27]. Ball, plug and gate valves suit such on/off duty because fully open they leave a nearly unimpeded path [24].

Where the Controller Runs

A distributed control system holds loops like these at a supervisory level that sends set points to field controllers, with loops nested or cascading and cycle times from milliseconds to minutes [28]. In condensed-mode polyethylene reactors, cascades of PID loops for the temperature are well established, and operators assist the critical loops by hand through that system [12]. In the app all nine loops are velocity-form PID blocks that execute every 0.1 min, the six overrides sit between the controllers and the valves, and the predictive controller is a supervisory layer that decides every 5 min and hands the valves back to the cascades when a plan fails.

> **IN PRACTICE** **Keep the Poisons Out of the Loop**
>
> The gas-phase process is highly sensitive to impurities, so purification beds remove sulfur compounds, acetylene and other impurities from the feed ethylene before it is compressed to reaction pressure. A small side stream of the recycle gas, mostly ethylene, is taken off to limit the build-up of impurities and burned as fuel gas or sent back to a cracker. Oxygen is kept very low because it poisons transition-metal catalysts, which are so effective that one gram of the metal makes more than 200 tonnes of product [11]. Scenario 9 shows what an impurity that passes the beds does to production and to melt index.

What Is Modernizing It

Soft sensors estimate product quality where online analyzers are not available [29], the line of work that the first History box traces to an industrial reactor. Online melt flow analyzers shorten the laboratory's delay at a cost in equipment and accuracy [26]. Among the industrial nonlinear predictive control applications counted in 2003, polymers were one of the areas with the largest number [13].

Keeping an inferred melt index right through a grade change, and handing the operator's supervisory decisions to a machine, are the subjects of the Open challenge boxes in Section 8.9.

8.9 Limitations

The laboratory is not in the default cascaded loops. It is sampled on a 60 min interval with a 45 min turnaround, displayed, exported and discussed, but both cascaded quality controllers read the true bed-average states directly as described in Section 8.3.8. Under Cascades the simulator therefore demonstrates what a delayed and infrequent measurement looks like beside a soft sensor, and it does not demonstrate control through one. For Cascades the melt-index analyzer bias disturbance is the only path by which a quality measurement can be made wrong, and it applies a constant multiplicative error rather than the state-dependent error a real inferential sensor shows during a transition. Closing either cascaded quality loop on the laboratory value, with or without a bias update between samples, would be a straightforward extension and it would change the transition numbers in Table 8.9. Control through a slow analyzer of that kind is the subject of Volume 5, Chapter 3.

OPEN CHALLENGE A Melt Index Estimate That Survives the Grade Change

A soft sensor loses accuracy when the process changes, through catalyst deactivation or a change in the feed, so it has to be maintained. Recursive updating with a bias correction has predicted melt index through grade changes of high-density polyethylene, but recursive methods cannot cope with abrupt changes [29]. In a 2026 study of 425 industrial polypropylene runs, melt index models that were accurate within one product family lost accuracy when carried across families [25]. Progress would be an estimate that stays right through a transition, the moment at which this chapter's cascades read the true bed average and a plant cannot.

The MPC instead uses each delayed report for sample-time bias correction. Its nominal model omits poison, catalyst activity, changed bleed/header capacities, pressure interactions and cooler dynamics. The squared quality objective does not minimize the off-grade indicator, and soft output envelopes provide no invariance guarantee. In the prepared, drawdown and impurity runs of Table 8.10 the MPC loses to the cascades.

Only three loops are scored. Overshoot and settling time appear only after a setpoint step larger than the loop's band and read as a dash otherwise, so a regulatory run such as the lined-out plant reports integral absolute error, manipulated-variable travel and the saturated fraction alone. Each cascade loop is scored on its own output, the cooler outlet setpoint for temperature and the hydrogen ratio setpoint, between 4 and 80, for melt index. The temperature tile counts saturation from the water valve beneath it, because the external reset holds the setpoint near the outlet rather than at its own limit, so in the cooling-fault run it reports 27.6 %. The melt-index tile reports 48.9 % by $t = 360$ min in the suggested transition, the time its output sat at 80.

The modeling simplifications set out in the assumptions bound what the reactor can be used to teach. There is no transport delay, no bubble phase, no molecular-weight distribution, no rigorous flash, no inert accumulation and no lock-hopper discharge cycle, and the sticking limit is a linear correlation rather than a measured softening curve evaluated on the bed-average resin, so the constraint itself moves during a grade change. The six overrides cannot be switched off, which is deliberate but means a student cannot run the experiment of removing them to see what the failure looks like.

OPEN CHALLENGE Supervision Still Rests With the Operator

In condensed-mode polyethylene reactors the temperature is held by cascades of PID loops, and operators still assist the critical loops by hand. A 2026 industrial study finds their intervention indispensable under rapid, strongly nonlinear disturbances, as in deciding when to take over the inner temperature loop. A machine framework that learns online is not yet established, and extremes such as a sudden large depressurization were not stress-tested [12]. Progress would be supervision that a machine can carry through such events.

8.10 Conclusion

This chapter studied a gas-phase polyethylene reactor simulation built around a separation of time scales of two decades. Thirteen states carry four gas-phase species balances, a lumped catalyst activity, an integrating solids inventory, a bed energy balance with an effectiveness-NTU condensing cooler, and two bed-average quality inventories, all initialized by solving for the starting line-up for the scenario setpoints. The dominant time constants span 0.5 min on the cooler outlet, as measured, to 240 min on the resin property, and with the catalyst feed held the outer temperature loop is shown to see a lead-lag plant whose fast gain of 5.77 °C/°C decays to about 4.5 °C/°C over the 176 min catalyst time constant. Nine proportional-integral-derivative loops run in three tiers with three cascades, one feedforward path and six non-defeatable overrides, each loop on its own manual station so that any one of them can be opened for a step test. The grade transition from Film A to Molding C costs 179.1 t of off-specification resin on the shipped tuning, 94.5 t after the quality cascade is retuned, 80.9 t with the bed also started at 55 % and 77.5 t with it drawn down to that level before the change. Production rate cancels out of the leading-order transition tonnage, and the measurements agree, since every retuned run spends 0.77 to 0.84 t per tonne of bed in its first excursion. Cutting production to 20 t/h before the change recovers only 1.5 t, and holding the cut through the change costs 34.3 t more than the full-rate run, because the slower bed makes the retuned melt-index loop overshoot into a second excursion.

The simulator is a reduced-order teaching model and not a plant replica. Its coefficients are calibrated rather than measured, its default cascades bypass the laboratory, and its recycle loop carries no transport delay. A fuller model would add a recycle transport delay, a two-site kinetic model and direct optimization of transition tonnage. The optional predictive controller costs 93.8 t on the shipped case (107.0 t without preview), but loses to the prepared and drawn-down cascades and on the impurity run. On one plant the reactor shows cascade control, constrained operation and the economics of a grade change.

KEY POINTS

- The gas changes in minutes, but the resin that leaves averages a bed that holds four hours of production.
- Melt index and density are reached only through gas-ratio setpoints, so each quality loop is the outer half of a cascade.
- Production rate cancels from the off-grade tonnage of a transition. Bed mass and the outer loop's driving force decide it.
- Retuning cut the Film A to Molding C loss from 179.1 to 94.5 t, and a bed started at 55 % cut it to 80.9 t.
- The cooling-water valve warns of the heat-removal limit before the bed temperature does.

Where else this applies. Quality held through a gas ratio over a slower inventory is also the problem of Chapter 9, and control through a late measurement that of Volume 5, Chapter 3.

References

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

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

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

[4] M. Takeda and W. H. Ray, "Optimal-grade transition strategies for multistage polyolefin reactors," *AIChE Journal*, vol. 45, no. 8, pp. 1776–1793, 1999.

[5] T. Xie, K. B. McAuley, J. C. C. Hsu, and D. W. Bacon, "Gas phase ethylene polymerization: Production processes, polymer properties, and reactor modeling," *Industrial & Engineering Chemistry Research*, vol. 33, no. 3, pp. 449–479, 1994.

[6] K.-Y. Choi and W. H. Ray, "The dynamic behaviour of fluidized bed reactors for solid catalysed gas phase olefin polymerization," *Chemical Engineering Science*, vol. 40, no. 12, pp. 2261–2279, 1985.

[7] K. B. McAuley, J. F. MacGregor, and A. E. Hamielec, "A kinetic model for industrial gas-phase ethylene copolymerization," *AIChE Journal*, vol. 36, no. 6, pp. 837–850, 1990.

[8] K. B. McAuley, J. P. Talbot, and T. J. Harris, "A comparison of two-phase and well-mixed models for fluidized-bed polyethylene reactors," *Chemical Engineering Science*, vol. 49, no. 13, pp. 2035–2045, 1994.

[9] K. B. McAuley and J. F. MacGregor, "On-line inference of polymer properties in an industrial polyethylene reactor," *AIChE Journal*, vol. 37, no. 6, pp. 825–835, 1991.

[10] ——, "Nonlinear product property control in industrial gas-phase polyethylene reactors," *AIChE Journal*, vol. 39, no. 5, pp. 855–866, 1993.

[11] European IPPC Bureau, "Reference document on best available techniques in the production of polymers," European Commission, Joint Research Centre, Institute for Prospective Technological Studies, Seville, Tech. Rep., Aug. 2007.

[12] Y. Liu, "Natural intelligence principle of supervisory-level control for condensed-mode polyethylene fluidized bed reactors: A proportional-integral-derivative perspective," *ACS Omega*, vol. 11, no. 39, pp. 60 064–60 072, 2026.

[13] 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.

[14] C. Chatzidoukas, J. D. Perkins, E. N. Pistikopoulos, and C. Kiparissides, "Optimal grade transition and selection of closed-loop controllers in a gas-phase olefin polymerization fluidized bed reactor," *Chemical Engineering Science*, vol. 58, no. 16, pp. 3643–3658, 2003.

[15] S. BenAmor, F. J. Doyle, III, and R. McFarlane, "Polymer grade transition control using advanced real-time optimization software," *Journal of Process Control*, vol. 14, no. 4, pp. 349–364, 2004.

[16] A. Mirzaei, A. Kiashemshaki, and M. Emami, "Fluidized bed polyethylene reactor modeling in condensed mode operation," *Macromolecular Symposia*, vol. 259, no. 1, pp. 135–144, 2007.

[17] I. J. Levine and F. J. Karol, "Preparation of low and medium density ethylene polymer in fluid bed reactor," U.S. Patent 4,011,382, assigned to Union Carbide Corp., 1977, filed March 10, 1975, issued March 8, 1977.

[18] J. M. Jenkins, III, R. L. Jones, and T. M. Jones, "Fluidized bed reaction systems," U.S. Patent 4,543,399, assigned to Union Carbide Corp., 1985, filed April 3, 1984, issued September 24, 1985.

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

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

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

[23] 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.

[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] J. Hernández-Fernández and J. Lopez-Martinez, "Reactor-grade non-additive polypropylene architecture defines the physicochemical limits of melt flow index model transferability: An industrial chemometric and machine learning study," *Polymers*, vol. 18, no. 15, p. 1880, 2026.

[26] M. A. Parvez and I. M. Mehedi, "An artificial intelligence-based melt flow rate prediction method for analyzing polymer properties," *Polymers*, vol. 17, no. 17, p. 2382, 2025.

[27] R. G. Aronson, "Fluidized bed discharge process," U.S. Patent 4,621,952, assigned to Union Carbide Corp., 1986, filed April 3, 1984, issued November 11, 1986.

[28] 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.

[29] 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.

8.11 Exercises

8.1 Concept check · 2 min

At the solved design point of Table 8.3, which statement about the bed energy balance of Equation (8.8) is right?

(a) The dry recycle gas alone carries the whole reaction heat, and the latent term is a spare margin

(b) The wall loss is the second largest removal term, ahead of the latent heat of the condensate

(c) The latent term removes 5.90 MW of 26.07 MW, and without it the dry gas falls short

(d) Most of the thermal inertia on the left side is the vessel steel rather than the resin in the bed

8.2 Concept check · 1 min

Under the default Cascades structure, how do the melt-index and density controllers act on the plant?

(a) They move the catalyst feeder, which sets how much resin of each kind is made

(b) They move the hydrogen and 1-hexene valves directly, one valve for each property

(c) They move the bed temperature setpoint, which shifts both correlations at once

(d) They write the H_2/C_2 and C_6/C_2 ratio setpoints, which inner loops hold with the valves

8.3 Concept check · 2 min

In the cooling-fault scenario the water warms, the cooler fouls and the compressor loses capacity, and none of the three is measured. Which reading warns first that the plant is approaching its heat-removal limit?

(a) The cooling-water valve position, together with the cooling utilization

(b) The bed temperature, because it is the variable the cascade is built to hold

(c) The bed-average melt index, because a hotter bed makes a different resin

(d) The reactor pressure, because the recycle compressor sets the gas inventory

8.4 Theory and calculation · 6 min

For the Film A to Molding C change, take $q_0 = 0$, the lower band edge $q_{spec} = \ln 7.2$, and a retuned cascade that holds the instantaneous property at $q_{inst} = 3.85$. Give each answer to three significant figures.

(a) Use Equation (8.20c) to estimate the off-grade tonnage with 120.1 t and with 98.6 t of bed at the change. *Checkpoint:* 86.3 t and 70.9 t.

(b) With 120.1 t of bed, find the off-specification time t_{off} at 28 t/h and at 20 t/h, and the tonnage in each case. *Checkpoint:* 185 min and 259 min, 86.3 t in both.

(c) Repeat (a) for 120.1 t if the cascade only drives the instantaneous melt index to the target, $q_{inst} = \ln 8$. *Checkpoint:* About 358 t.

8.5 Theory and calculation · 7 min

Use the constants of Table 8.2 at the design point, where $M_b = 119970\,\text{kg}$, $\dot{m}_R = 16382\,\text{kg/min}$ and $\dot{m}_{out}/M_b = 4.168 \times 10^{-3}/\text{min}$. Give each answer to three significant figures.

(a) Find the bed thermal lag $(M_b c_{p,s} + C_{metal})/(\dot{m}_R c_{p,g})$. *Checkpoint:* 11.3 min.

(b) The bed is held 5.5 K above 88 °C. Find k_d from Equation (8.6), the catalyst time constant, and the ratio of the new steady active inventory to the old at the same feeder position. *Checkpoint:* 1.97×10^{-3}/min, 163 min and 0.923.

(c) Find the factor by which the Arrhenius term of Equation (8.4) raises the propagation rate at the same temperature, and the product of the two factors. Explain what the product says about the temperature step test of Section 8.4. *Checkpoint:* 1.072 and 0.990.

8.6 Theory and calculation · 6 min

At the design point the recycle carries $\dot{m}_R = 16382\,\mathrm{kg/min}$ with $c_{p,g} = 1.515\,\mathrm{kJ/(kg\,K)}$ from 40.00 °C to the bed at 88.00 °C, 6.00 % of it condensed with $\Delta H_{vap} = 360\,\mathrm{kJ/kg}$, and the reaction runs at $r_p = 16550\,\mathrm{mol/min}$ with $-\Delta H_p = 94.5\,\mathrm{kJ/mol}$ while making 30 t/h. Give each answer to three significant figures, and the duties within 0.01 MW.

(a) Find the reaction heat, the sensible and latent duties and the wall loss $UA_{wall}(T_R - T_{amb})$ in MW. *Checkpoint:* 26.07, 19.85, 5.90 and 0.32 MW.

(b) Estimate the production the dry recycle could carry at the same temperatures. *Checkpoint:* About 23.2 t/h.

(c) Estimate the production each further percent of the recycle condensed would add, and compare it with the 0.525 t/h per percent of the running example in Section 1.4.4. *Checkpoint:* About 1.13 t/h.

8.7 App challenge · 8 min

Load scenario 2, "2 · Step test the inner cooler loop", at the default seed. The Cooler outlet temperature loop starts in Manual with the Cooling water valve at 50 %. Before Run, set the Cooling water valve to 60.0 %, so the step is applied at $t = 0$ to the plant lined out at the old position. Run at speed 1 (1 min per wall-clock second) until the clock passes 15 min, pause and export the CSV. Temperatures within 0.01 °C of the key (0.1 °C for the bed), the gain within 0.002 and the time within 0.05 min count as correct.

(a) Read `tc_C` in the rows at 0.0 and 2.0 min and find the process gain. *Checkpoint:* 40.00 °C and 38.54 °C, −0.146 °C per percent.

(b) Interpolate the `tc_C` rows to find when the outlet has made 63.2 % of the change it shows at 2.0 min, and compare it with $\tau_V + \tau_{HX}$. *Checkpoint:* About 0.51 min, against 0.65 min.

(c) Read `tc_C` and `temp_C` in the row at 15.0 min, and explain why the outlet is still moving. *Checkpoint:* 38.38 °C and 82.0 °C.

8.8 App challenge · 8 min

Load scenario 5, "5 · Fast gas, slow resin", at the default seed, with nothing changed. The Melt index loop starts in Manual, and the H_2/C_2 setpoint steps from 20 to 30 at 20 min. Run at speed 10 (10 min per wall-clock second) until the clock passes 360 min, pause and export the CSV. A run longer than an hour exports its time column as Time (h), in hours, so convert each time below from min. Values within 0.01 of the key (0.1 t for tonnages) count as correct.

(a) Read `h2c2` in the rows at 80.0 and 200.0 min, and `mi_inst` at 80.0 min. *Checkpoint:* 26.90, 30.02 and 2.02.

(b) Read `mi_bed` at 100.0, 200.0 and 360.0 min, `offgrade_t` and `produced_t` at 360.0 min, and the first row in which `on_spec` is 0. *Checkpoint:* 1.15, 1.50 and 1.99; 136.8 t of 177.0 t; 81.75 min.

(c) Use Equation (8.13) with $\tau_s = 240\,\mathrm{min}$ to estimate `mi_bed` at 360 min if the instantaneous melt index had jumped to its final 2.69 at 20 min. Explain the difference from (b). *Checkpoint:* About 2.12.

8.9 App challenge · 10 min

Load the graded challenge, "Challenge · Film A → Molding C", at the default seed. The grade target changes to Molding C at 90 min. Run at speed 10 (10 min per wall-clock second) throughout. A run longer than an hour exports its time column as Time (h), in hours, so convert each time below from min. Melt index within 0.01, density within 0.0001, tonnages and the margin within 0.1 and times within 0.5 min of the key count as correct.

(a) With nothing changed, run until the clock passes 360 min, pause and export the CSV. Read `mi_bed`, `density_bed` and `offgrade_t` in the row at 360.0 min, the largest `h2c2_sp` and its time, and say whether the challenge passes. *Checkpoint:* 6.59, 0.9321, 115.2 t; 73.8 at about 174 min; it fails on both properties.

(b) Continue to 900 min and export again. Read `offgrade_t` at 900.0 min and the last row in which `on_spec` is 0. *Checkpoint:* 179.1 t, at 489.5 min.

(c) Reload the case, press Reveal suggested tuning before Run, and repeat (a) and (b), reading also the smallest `stick_margin_C` up to 360 min. *Checkpoint:* 8.40, 0.9352, 80.9 t, margin 7.8 °C; 80.9 t at 900 min, last at 281.0 min.

8.10 Application area · 4 min

Use Equation (8.14) with the constants of Table 8.2, a gas with 12 mol% cyclohexane, and each bed at its grade's melt index and density (Table 8.4).

(a) Find the sticking temperature and margin of Film A at 88 °C, Molding C at 88 °C and HDPE D at 92 °C. *Checkpoint:* 8.0, 21.2 and 35.8 °C.

(b) A plant engineer proposes raising cyclohexane to 15 mol% on Film A to gain cooling. What does it cost in margin, and on which grades does the cost matter? *Checkpoint:* 1.05 °C, leaving 6.95 °C.

8.11 Application area · 3 min

In Scenario 5 the bed-average melt index leaves the band of 0.9 to 1.1 at 81.8 min. The laboratory draws a sample every 60 min from $t = 0$ and posts each result 45 min later. When does the first laboratory result outside the band arrive, how long after the bed left the band is that, and roughly how much off-grade resin has been discharged by then at 500 kg/min?

8.12 Review the agent · 5 min

An AI assistant asked how to cut the off-grade resin of the Film A to Molding C change writes: "Off-grade tonnage is the discharge rate times the time off specification. Keep the retuned cascade and the bed at 65 %, cut production from 28 to 20 t/h for the whole transition, and the loss falls by nearly a third, from about 94 to about 67 t." Find the flaw, and say how you would check it in the app.

CHAPTER 9

Polypropylene Loop Simulator to Teach Quality Control Against a Cavitation Limit

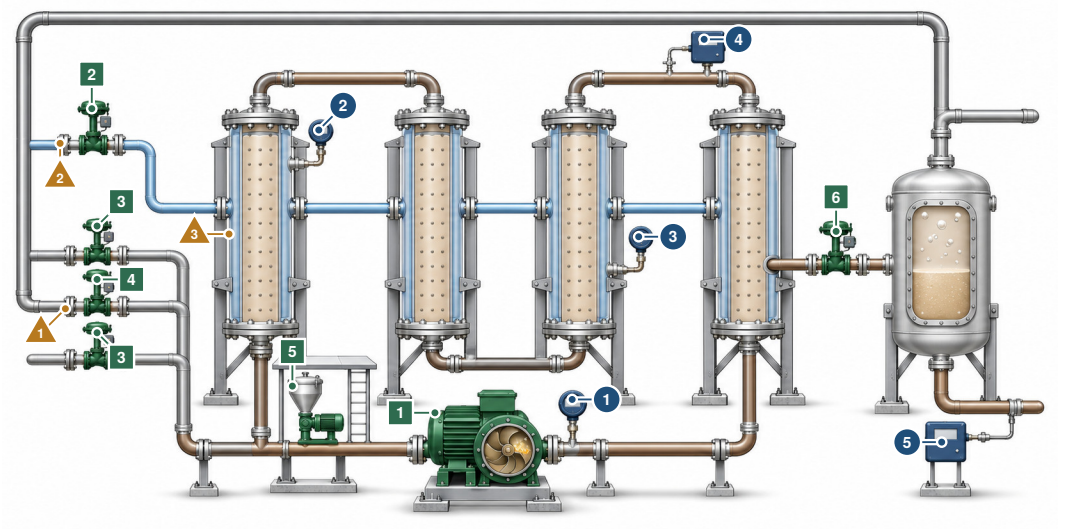

The controller moves 1 circulation pump speed, set by the operator 2 cold-water valve on the jackets 3 hydrogen and ethylene valves 4 fresh propylene valve, which holds the pressure 5 catalyst feeder 6 slurry take-off valve

The controller reads 1 reactor pressure 2 reactor temperature 3 slurry density 4 loop chromatograph: dissolved hydrogen, ethylene and inerts 5 melt flow rate and density analyzer

What disturbs it 1 light inerts in the propylene feed 2 cooling-water temperature 3 jacket fouling

Notice the bubbles in the pump's impeller: the hydrogen that valves 3 add for the grade also raises the liquid's bubble point toward the pressure the pump needs, and the chromatograph sees that hydrogen within minutes while the analyzer under the flash shows the grade about an hour later. Illustration, not to scale.

IN PLAIN TERMS **One Gas Sets the Product and Threatens the Pump**

The main idea. Plastic for cups and fibers grows as tiny grains in a pipe loop full of liquid. A large pump keeps the liquid moving, and a little dissolved hydrogen sets how easily the melted plastic flows.

The challenge. The same hydrogen brings the liquid closer to boiling, and bubbles in the pump stop it from pushing. The pump feels a change in minutes and the plastic only after about an hour, so chasing the plastic can wreck the pump.

Why it matters. A guard controller that limits the hydrogen saves the pump, but on its own it leaves the plant short of the plastic that was ordered.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the two History boxes in Section 9.2, the Did you know? boxes in Sections 9.3 and 9.6, the two Try it boxes in Section 9.6, and Key points at the end of the Conclusion. No equations.
Study	65 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 9.6 from Table 9.5 with Figure 9.3, Section 9.7, and Section 9.8 with its Open challenge box.

The control question In a liquid-full polypropylene loop the hydrogen that sets melt flow rate also raises the bubble point of the liquid toward the pressure the circulation pump needs. How should a quality controller and a margin controller share the one hydrogen setpoint, and what does each arrangement cost in pump damage and in off-specification polymer?

Learning objectives

- Compute the bubble point of the loop liquid and its gradients, and turn a grade request into a pressure requirement.
- Explain why one hydrogen move reaches the margin in minutes, the product in more than an hour and the pump in seconds.
- Describe the three details that make a low selector work: tracking of the loser, a range that lets the constraint win and parking of the idle controller.
- Compare the override and the predictive controller on one grade transition, and attribute the result to the plan or to the temperature.
- Plan a grade transition that buys margin with pressure before spending it on hydrogen.

Prerequisites Sections 1.4.4 (what spends the margin), 1.4.5 (inventory-averaged quality), 1.4.8 (one-directional guards), 1.6.2 (velocity-form PID) and 1.6.3 (the predictive option). Chapter 8 for a polyolefin grade change against a cooling limit.

The app https://apmonitor.com/apps/polypropylene-loop/

See also Volume 4, Chapter 2 for a protective override against surge and Volume 3, Chapter 4 for selectors that cross-limit fuel and air. Volume 9, Chapter 4 for offset-free predictive control with an estimated bias, and Volume 3, Chapter 2 for matching a cascade to where the disturbance enters.

9.1 Introduction

A constraint is usually drawn as a horizontal line on a trend and taught as a region to stay out of. A bulk polypropylene loop reactor refuses that picture. Its product specification is melt flow rate, the only effective handle on melt flow rate is dissolved hydrogen, and dissolved hydrogen is what raises the bubble point of the circulating liquid toward the pressure that keeps the circulation pump from cavitating. The line on the trend is therefore not a place the operator avoids. It is another controller wanting the same valve, in the opposite direction, and one of the two has to lose. A block diagram of the melt-flow cascade contains no term in which that collision can appear, because the hydrogen valve enters the quality transfer function once and the pump does not enter it at all. The misconception this simulator corrects is the belief that a constraint is tuned away from rather than arbitrated with.

The shape of the problem is general. Wherever the only effective actuator for a quality or throughput objective also consumes a mechanical or thermodynamic margin, a constraint controller and a quality controller share one manipulated variable and a selector decides which of them writes it. A centrifugal compressor raises throughput with the same guide vanes and speed that carry it toward surge, and the anti-surge controller is a genuine controller with its own measurement and its own authority [1], as Volume 4, Chapter 2 shows. A distillation column buys purity with reflux and pays for it in vapor loading toward the flooding limit. A pump on a hot suction gains flow at the cost of net positive suction head [2]. Each is arranged as a ratio, override or selective control problem rather than as a single loop [3], and each fails in the same way when the losing controller is allowed to wind up or is parked where it cannot act quickly. Polymer plants meet the pattern more often than most, because product quality there is inferred from a slow sampled analyzer and the mechanical constraints are fast [4, 5].

Loop reactor models in the literature answer polymerization questions. Population-balance treatments of the liquid-phase loop resolve the particle size distribution, the multiple site types and the recycle structure [6, 7], and commercial bulk processes have been matched at steady state and in transient [8]. The thermodynamic state of the circulating liquid receives far less attention even though it is what decides whether the pump survives a grade change. However, a student handed the melt-flow

transfer function and asked for proportional and integral settings never encounters the reason the resulting controller wrecks an impeller, because the pump is not in the model that was handed over. This chapter teaches override control on a bulk polypropylene loop reactor whose bubble point is computed from the same liquid composition the quality controllers manipulate, so that a well-tuned melt-flow loop and a wrecked circulation pump become one event. The kinetics, the bubble-point correlation and the selector are standard. What the app adds is that all of them act on one plant, in a browser, with every variable logged.

Section 9.2 traces how the bulk loop and its quality specifications came about. Section 9.3 carries every equation and every constant of the plant: nineteen states, a re-verified vapor-pressure surrogate, a Henry-law bubble point, an impeller-eye depression, an algebraic head-breakdown curve and a deposition-velocity settling limit. Section 9.4 prices each dissolved species in bubble point, turns a quality request into a pressure requirement and separates the three clocks of the plant. Section 9.5 describes the ten loops, the low selector, the interlocks and the predictive controller. Section 9.6 runs one grade transition twice, once to destroy the pump and once to abandon the grade, measures what the selector reports and what slow disturbances cost, and shows that the graded challenge is won only by raising the operating pressure and slowing the melt-flow controller together. It ends with the predictive controller on the same transition. Section 9.7 names the hardware and the practice behind the model, and Sections 9.8 and 9.9 state what the model cannot show and what a plant would add.

9.2 Polypropylene Process Technology

The catalysts of the bulk loop descend from Ziegler's low-pressure ethylene process [9] and from Natta's crystalline, stereoregular polymers of propylene [10], whose contested origin the first History box below tells. That stereoregularity is what made polypropylene a material rather than a curiosity, and it is the reason a polypropylene reactor has to control a variable an ethylene reactor does not.

> ⧖ HISTORY **A Crystalline Plastic From a Gas, 1951 to 1954**
>
> Near the end of 1953 Karl Ziegler's group found that ethylene polymerizes even at normal pressure on catalysts made by mixing an aluminum alkyl with titanium tetrachloride, and the catalysts became known to industry around the turn of the year [11]. Early in 1954 Giulio Natta's group polymerized propylene with them. Some of the product was crystalline, and Natta traced the crystallinity to a regular steric arrangement along the chain rather than to a higher molecular weight [12]. The two shared the 1963 Nobel Prize in Chemistry [11]. Yet in 1951 J. Paul Hogan and Robert L. Banks, trying to turn propylene into gasoline over a catalyst carrying chromium oxide, had already made crystalline polypropylene. Three laboratories filed patent applications on the discovery between 1951 and 1953, and the contest lasted three decades, until an appellate court ruled for Hogan and Banks in 1983 [13].

The process technology went through three stages, driven by the catalyst generations that the second History box below traces, until the reactor was almost the whole plant. The third stage removed the diluent. In a bulk or liquid-pool process the monomer is the reaction medium, which raises the propylene concentration at the active site, removes the solvent recovery train, and takes the reactor from a stirred vessel to a pipe loop that a single axial pump keeps in circulation. The Spheripol configuration pairs that liquid loop with a gas-phase reactor downstream for impact copolymer, and it is the arrangement most of the world's polypropylene is made in [14]. Metallocene catalysis arrived in the 1980s and became a commercial option in the 1990s, because a chiral single-site metallocene fixes stereochemistry at one well-defined center and therefore produces a narrow molecular weight distribution and a controllable comonomer sequence [15]. Single-site catalysts changed the grade slate rather than the reactor, and the bulk loop remained the vessel.

HISTORY From Washed Powder to Liquid Monomer, 1960s Onward

The first polypropylene catalysts, used in slurry processes from the 1960s, made about one tonne of polymer per kilogram of catalyst and 5 to 10 percent of atactic polymer, so the catalyst had to be washed out of the product and the atactic part removed. The catalysts of the 1970s reached about ten tonnes per kilogram, and supports with electron donors raised the isotacticity. The fourth generation, the present standard, sits on a spherical support made for liquid monomer, yields 30 to 50 tonnes per kilogram and made both cleaning steps obsolete. Modern suspension processes therefore use liquid propylene instead of a solvent, which the industry calls a bulk process [16].

The controlled quality variables followed from what a converter can measure. Molecular weight itself is a laboratory determination, and average molecular weight enters processing through melt viscosity, so the industry adopted melt flow rate, the mass of polymer extruded through a standard capillary in ten minutes under a standard load, as the operating specification. Hydrogen sets it. Hydrogen terminates a growing chain by transfer to the metal alkyl, and increasing the hydrogen concentration therefore shortens the chains and raises melt flow rate, with a response that depends on the catalyst and donor system and is a fitted quantity in every plant model that exists [17]. Density is the second specification and it follows ethylene incorporation, because the comonomer units interrupt the isotactic sequences and reduce crystallinity. Inference of both from composition and temperature, rather than waiting for the analyzer, became the standard route to closed-loop quality control in polyolefin plants [4], and grade transitions on a liquid-propylene loop are now posed as constrained optimal control problems [5]. What that literature treats sparingly is the arrangement below the quality loop, where the loop is liquid-full and the same dissolved hydrogen that sets the specification is also the largest avoidable term in the bubble point of the circulating liquid.

9.3 Process Model

Reactor R-101 is a 65 m^3 liquid-full loop of 0.65 m bore and 196 m developed length, built as four vertical legs of 40 m joined by 180° bends, with cooling jackets over 326.7 m^2 of the wall. Axial pump P-101 sits on the horizontal return and circulates the slurry at 5.50 m/s against the loop resistance. Catalyst and hydrogen and ethylene enter at separate points, fresh propylene make-up enters on pressure control, and product leaves through take-off valve XV-101 to a flash at 18 bar where the monomer is recovered and returned. The circuit, its instrument tags and each symbol carried by the equations of this section are drawn in Figure 9.1.

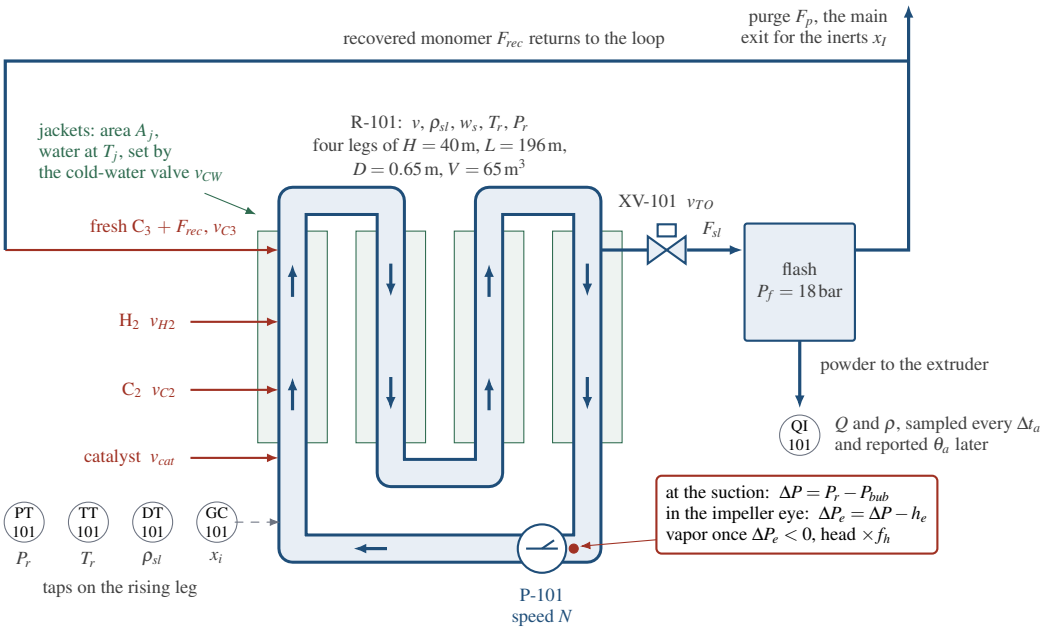

Figure 9.1 Loop reactor R-101 with its four jacketed legs, the axial circulation pump, the four injection points and the product take-off. Liquid propylene circulates at velocity v round a developed length L of bore D and volume V, carrying solids fraction w_s at slurry density ρ_{sl}, temperature T_r and pressure P_r. Catalyst, ethylene, hydrogen and fresh propylene enter separately on the leg leaving the pump, and the four instruments that the control system reads are PT-101, TT-101, DT-101 and the loop chromatograph GC-101 reporting the dissolved mole fractions $x_i = \{x_{H2}, x_{C2}, x_I\}$. The margin the application is about is measured at the pump suction as $\Delta P = P_r - P_{bub}$ and inside the impeller eye as ΔP_e, one eye depression h_e lower, where vapor appears and multiplies the delivered head by f_h. Slurry F_{sl} leaves through XV-101 to the flash at P_f, whose recovered monomer F_{rec} returns with the fresh propylene. The purge F_p and the 15 % of the flashed inerts that recovery does not return are the only exits for the accumulated inerts. Analyzer QI-101 reports melt flow rate Q and density ρ on a sample interval Δt_a with dead time θ_a.

✦ DID YOU KNOW? Pipe Loops Up to Fifty Meters Tall

Loop reactors are long jacketed pipes arranged as four or six vertical legs up to 50 m tall, and an in-line axial pump drives the slurry of a polyethylene loop round at 6 to 10 m/s. Liquid-propylene loops run near 70 °C and 4 MPa with about 40 wt% of polymer [16].

Everything the control system touches appears in Table 9.1, where melt flow rate is in g/10min, density in g/cm^3, slurry density in kg/m^3 and every valve position in %. Nineteen of those quantities are integrated. Units throughout are the ones on a control-room faceplate, temperature in °C, pressure in bar, mass in kg, amount in mol and velocity in m/s, and the clock is **in minutes**. Every rate below is therefore per minute, and every time constant is in minutes unless it is quoted in seconds.

Table 9.1 Classification of Variables. Nominal Values Are Those of the Solved Design Point at Which Every Scenario Opens.

Symbol	Role	Nominal	Range
Q melt flow rate	CV, sampled	3.00	0.5–40 SP
ρ density	CV, sampled	0.90550	0.888–0.907 SP
R production	CV, inferred	22.0	8–34 t/h SP
ΔP margin	CV, inferred	3.43	0.5–6 bar SP
T_r	CV, measured	70.0	58–84 °C SP
P_r	CV, measured	38.00	30–43 bar SP
ρ_{sl}	CV, measured	591	480–650 SP
r_H H_2/C_3	CV, measured	0.900	0.02–5.0
r_E C_2/C_3	CV, measured	0	0–6.0
v_{H2} hydrogen	MV	11.9	0–100 %
v_{C2} ethylene	MV	0.0	0–100 %
v_{C3} propylene	MV	43.3	0–100 %
v_{CW} cold water	MV	45.0	0–100 %
v_{TO} take-off	MV	45.0	0–100 %
v_{cat} catalyst	MV	44.7	0–100 %
N pump speed	MV, operator	100	40–110 %
T_{cw} water supply	DV, measured	22	4–34 °C
y_I feed lights	DV, measured	1.47	0–5 mol%
purge setting	DV, measured	100	0–160 %
header pressure	DV, measured	100	70–110 %
b_a batch activity	DV, unmeasured	0	±35 %
z poison	DV, unmeasured	0	0–0.40
jacket fouling	DV, unmeasured	0	0–60 %
analyzer bias	DV, unmeasured	0	±15 %
n_{H2} (mol)	state	394.7	≥ 0.01
n_{C2} (mol)	state	0	≥ 0
n_I (mol)	state	22280	≥ 1
M_s (kg)	state	19203	500–40950
A_c (—)	state	0.994	≥ 0
T_r, T_j (°C)	states	70.0, 41.2	20–94, 0–94
P_r (bar)	state	38.00	8–60
v (m/s)	state	5.50	0–14
$\bar{q}$ (—), $\bar{e}$ (wt%)	states	1.099, 0	—
ϕ settled layer	state	0	0–0.50
seven valve lags (%)	states	as above	0–110

9.3.1 Inventories and Composition

Solids and the three dissolved lights are separate inventories, as shown in Equations (9.1) to (9.4).

$$\frac{\mathrm{d}M_s}{\mathrm{d}t} = r_p - \dot{m}_{\text{out}} \tag{9.1}$$

$$\frac{\mathrm{d}n_{H2}}{\mathrm{d}t} = F_{H2} - \zeta\, \dot{n}_{C3}\, r_H - o_{H2} \tag{9.2}$$

$$\frac{\mathrm{d}n_{C2}}{\mathrm{d}t} = F_{C2} - r_{C2} - o_{C2} \tag{9.3}$$

$$\frac{\mathrm{d}n_I}{\mathrm{d}t} = \dot{n}_{\text{feed}}\, y_I - o_I \tag{9.4}$$

where M_s is the polymer inventory in kg, r_p is the polymerization rate in kg/min, $\dot{m}_{\text{out}}$ is polymer leaving with the take-off, n_i are dissolved amounts in mol, F_{H2} and F_{C2} are the gas feeds in mol/min,

$\zeta \dot{n}_{C3} r_H$ is the hydrogen consumed by chain transfer, r_{C2} is the ethylene built into chains, y_I is the light-inert fraction of the fresh propylene, and o_i is the loss of species i with the liquid drawn off and with the recycle purge. Equation (9.4) has no reaction term at all, which is the reason inerts accumulate until the purge and the unrecovered part of the flashed liquid remove them, the purge taking 53 % at the design point.

The liquid volume is what the solids leave behind, and the liquid composition follows from the mass that is not one of the three lights, as shown in Equation (9.5).

$$\begin{aligned} V_L &= V - M_s/\rho_p, \qquad M_L = \rho_L V_L \\ n_{C3} &= \frac{1000}{\mathrm{MW}_{C3}} \left(M_L - \tfrac{1}{1000} \textstyle\sum_i n_i \mathrm{MW}_i \right) \\ x_i &= n_i / n_L, \quad r_H = 1000 \frac{x_{H2}}{x_{C3}}, \quad r_E = 100 \frac{x_{C2}}{x_{C3}} \end{aligned} \tag{9.5}$$

where V is the loop volume, ρ_p is the particle density, ρ_L is the liquid density, n_L is the total liquid amount, r_H is the composition ratio in mol H_2 per kmol C_3 and r_E is mol C_2 per 100 mol C_3. Slurry density is $\rho_{sl} = (M_s + M_L)/V$ and the solids weight fraction is $w_s = M_s/(M_s + M_L)$. Liquid density is taken linear about the design point, $\rho_L = \rho_{L,0}\,[1 + \kappa(P_r - 38) - \alpha(T_r - 70)]$.

9.3.2 The Bubble Point and the Cavitation Margin

The constraint is one algebraic equation, and it is the reason the application exists. Raoult's law supplies the solvent term and Henry's law supplies the three solutes, as shown in Equation (9.6).

$$P_{bub} = x_{C3} P^{\mathrm{sat}}_{C3}(T_r) + \sum_{i \in \{H2, C2, I\}} x_i H_i(T_r) \tag{9.6}$$

where P^{sat}_{C3} is the propylene saturation pressure and H_i is the Henry constant of solute i. The saturation pressure is a three-constant Antoine surrogate, $\log_{10}(P^{\mathrm{sat}}/\mathrm{bar}) = A - B/(T[\mathrm{K}] + C)$ with $A = 4.23311$, $B = 932.239$ and $C = -4.262$, fitted by least squares to seven propylene saturation pressures between 0 °C and 90 °C. Re-checking that fit against all seven reference points gives a worst relative residual of 0.45 %, at 90 °C, and no more than 0.27 % anywhere from 0 °C to 80 °C, against the 0.5 % tolerance the application's own verification suite applies. Each Henry constant is fitted at 70 °C and slides with a power of the solvent's own vapor pressure, $H_i(T) = H_{i,0}\,[P^{\mathrm{sat}}_{C3}(T)/P^{\mathrm{sat}}_{C3}(70)]^{n_i}$, so that every solute becomes less volatile relative to the solvent as propylene approaches its critical point. The margin the pump lives on and the margin the impeller actually sees are Equation (9.7).

$$\Delta P = P_r - P_{bub}, \qquad \Delta P_e = \Delta P - h_e \tag{9.7}$$

$$h_e = \frac{\rho_{sl}}{2 \times 10^5} \left[\sigma_v v^2 + \sigma_N \left(\frac{N v_r}{100} \right)^2 \right] \tag{9.8}$$

where h_e is the impeller-eye depression in bar, a two-term net positive suction head surrogate carrying a flow term and a peripheral-speed term [2], N is pump speed in percent and v_r is the rated velocity. At the design point $h_e = 0.36$ bar, so the transmitter reports 3.43 bar of margin while the impeller has 3.08 bar. Vapor appears when ΔP_e goes negative, and the void fraction and the head it costs are Equation (9.9).

$$\alpha_v = \begin{cases} 0, & \Delta P_e \geq 0 \\ \alpha_{\max} \left(1 - \mathrm{e}^{\Delta P_e / \delta} \right), & \Delta P_e < 0 \end{cases} \tag{9.9}$$

$$f_h = \frac{1}{1 + (\alpha_v / \alpha_{50})^2} \tag{9.10}$$

where δ is the void-growth pressure scale and α_{50} is the void fraction at which the pump delivers half its head. The form reproduces the shape of a real head breakdown, which is nothing and then everything over a fraction of a bar, and nothing else about cavitation.

9.3.3 Circulation Momentum and the Settling Limit

The circulating slurry is a 196 m column of fluid between the pump and the loop resistance, so velocity obeys a momentum balance rather than an algebraic pump-curve intersection, as shown in Equation (9.11).

$$\rho_{sl} L \frac{\mathrm{d}v}{\mathrm{d}t} = 60\left(\Delta P_{\mathrm{pump}} f_h - \Delta P_{\mathrm{fric}}\right) \tag{9.11}$$

$$\begin{aligned} \Delta P_{\mathrm{pump}} &= \Delta P_{\mathrm{rated}}\left[R_{sh}\left(\tfrac{N}{100}\right)^2 - (R_{sh}-1)\left(\tfrac{v}{v_r}\right)^2\right] \\ \Delta P_{\mathrm{fric}} &= \tfrac{1}{2} K_f (1 + k_\phi \phi_S) \rho_{sl} v^2 (1 + 2.2\phi) \end{aligned} \tag{9.12}$$

where $K_f = fL/D + \sum K$ is the loop resistance coefficient, ϕ_S is the solids volume fraction, ϕ is the settled-layer state, R_{sh} is the pump shut-off head ratio, and the factor 60 converts an acceleration in m/s^2 to the per-minute clock the rest of the model keeps. Linearizing (9.11) about the design point gives a restoring slope of 79.7 kPa per m/s against 115.8 t/m^2 of inertia, so the circulation time constant is $\tau_v = 1.45$ s, confirmed by kicking the velocity state 0.30 m/s low and measuring 1.48 s to recover 63 %. That figure is the reason the application has two time scales, and it is the one place where reading the clock as seconds rather than minutes changes the answer by a factor of sixty.

Below a critical velocity the solids stop being carried and slide along the bottom of the horizontal runs [18]. The application uses a Durand form for that velocity and grows a settled layer whenever the circulation falls under it, as shown in Equation (9.13).

$$v_{\mathrm{dep}} = F_L \sqrt{2 g D \left(\rho_p / \rho_L - 1\right)} \tag{9.13}$$

$$\begin{aligned} \frac{\mathrm{d}\phi}{\mathrm{d}t} &= k_f \max\!\left(0, 1 - \tfrac{v}{v_{\mathrm{dep}}}\right) \frac{\phi_S}{0.328}\left(1 - \frac{\phi}{\phi_{\max}}\right) \\ &\quad - 6 \times 10^{-4} \phi \max\!\left(0, \tfrac{v}{v_{\mathrm{dep}}} - 1.25\right) \end{aligned} \tag{9.14}$$

Evaluating (9.13) at the design point gives 3.65 m/s for 900 kg/m^3 powder in 440 kg/m^3 liquid propylene through a 0.65 m bore, which is a ratio of 1.51 against the 5.50 m/s circulation. The erosion term in (9.14) removes a settled layer with a time constant of 108 h at the design velocity and still 68 h at 110 % pump speed, so a layer built during a transient does not come off again within any run the application makes.

9.3.4 Kinetics, Heat and Pressure

One lumped active-site inventory carries the kinetics, with the propagation rate first order in site concentration and in monomer, as shown in Equations (9.15) and (9.16).

$$r_p = k_p A_c C_{C3} f_H f_E f_z \left(1 + b_a/100\right) \tag{9.15}$$

$$\frac{\mathrm{d}A_c}{\mathrm{d}t} = k_a v_{cat} - \frac{A_c \dot{m}_{\mathrm{out}}}{M_s} - k_d A_c \tag{9.16}$$

with $k_p = k_{p,\mathrm{ref}} \exp[-(E_p/R)(1/T_r - 1/343.15)]$, $f_H = 1 + k_H r_H/(K_H + r_H)$, $f_E = 1 + k_E r_E$, $f_z = 1/(1 + K_z z)$ and $k_d = k_{d,0} \exp[\beta_d (T_r - 70)]$, where C_{C3} is the liquid monomer concentration in mol/m^3, A_c is the active-site inventory, v_{cat} is the feeder setting in percent, b_a is the batch activity bias and z is the poison level. Equation (9.16) has a decay time constant of 100 min at 70 °C, which makes production the slowest regulatory loop on the plant.

Reaction heat leaves through the jacket and with the cold feed, and the jacket coefficient carries a velocity exponent, as shown in Equation (9.17).

$$\begin{aligned} C_r \frac{\mathrm{d}T_r}{\mathrm{d}t} &= Q_{rxn} + Q_{\mathrm{feed}} - Q_j - Q_{\mathrm{loss}} \\ M_j c_{pw} \frac{\mathrm{d}T_j}{\mathrm{d}t} &= Q_j - F_{cw} c_{pw} (T_j - T_{cw}) \\ Q_j &= U_0 A_j \left(v / v_r\right)^{0.80} (1 - \varepsilon)\,(T_r - T_j) \end{aligned} \tag{9.17}$$

where $Q_{rxn} = (-\Delta H_p) r_p$, $C_r = M_s c_{pS} + M_L c_{pL} + M_{st} c_{p,st}$ is the total thermal capacitance including the loop steel, ε is the combined external and settled-layer fouling, and F_{cw} is the cold-water flow set by valve v_{CW}. The exponent 0.80 is what converts a cavitating pump into a thermal runaway, because the velocity and the heat-removal coefficient fall together.

The loop is liquid-full, so pressure is not a level. It is what the liquid's own compressibility says about the mass that will not fit, as shown in Equation (9.18).

$$\frac{\mathrm{d}P_r}{\mathrm{d}t} = \frac{\dot{M}_L + (\rho_L / \rho_p)\,\frac{\mathrm{d}M_s}{\mathrm{d}t}}{V_L\, \rho_{L,0}\, \kappa} + \frac{\alpha}{\kappa}\,\frac{\mathrm{d}T_r}{\mathrm{d}t} \tag{9.18}$$

where $\dot{M}_L$ is the net liquid mass rate into the loop and κ and α are the isothermal compressibility and the thermal expansivity. The second term is worth 4.6 bar per degree in a vessel with nowhere to go, which is why the pressure controller cannot be left in manual. The first term carries a sign that surprises people, because polypropylene at 900 kg/m^3 is denser than liquid propylene at 440 kg/m^3, so polymerizing shrinks the contents.

9.3.5 Product Quality and the Analyzer

Melt flow rate and comonomer incorporation are correlations for the polymer being made at this instant, as shown in Equation (9.19).

$$\begin{aligned} \ln Q_{\mathrm{inst}} &= \ln K_Q + a_T \left(\tfrac{1}{343.15} - \tfrac{1}{T_r}\right) \\ &\quad + m \ln(b_0 + r_H + b_E r_E - 0.6 z) \\ w_{C2} &= \frac{a_{E1} r_E}{1 + a_{E2} r_E} \\ \rho &= \rho_h - k_\rho w_{C2} + k_{\rho Q} \ln(Q / Q_{\mathrm{ref}}) \end{aligned} \tag{9.19}$$

where $m = 2$, so melt flow rate goes as the square of the chain-transfer group and the process gain is $\mathrm{d}Q/\mathrm{d}r_H = 2Q/(b_0 + r_H)$. Evaluating that derivative by re-solving the steady state gives 6.25 g/10min per unit of ratio at the design grade and 12.50 g/10min per unit at 12 g/10min, so the process gain doubles across one grade transition. What leaves the loop is not Q_{inst} but a first-order average over the polymer inventory, as shown in Equation (9.20).

$$\frac{\mathrm{d}\bar{q}}{\mathrm{d}t} = \frac{r_p\,(\ln Q_{\mathrm{inst}} - \bar{q})}{M_s}, \qquad \frac{\mathrm{d}\bar{e}}{\mathrm{d}t} = \frac{r_p\,(w_{C2} - \bar{e})}{M_s} \tag{9.20}$$

whose time constant is $M_s / \dot{m}_{\mathrm{out}} = 52.4$ min at the design point. Analyzer QI-101 samples that product every $\Delta t_a = 6.0$ min and reports $\theta_a = 4.0$ min later, so the effective dead time in front of the quality loop is about 7 min, being half the sample interval plus the transport.

9.3.6 Parameters and Initialization

Table 9.2 lists the constants that set the behavior reported in this chapter. The application exposes 95 parameters in all, of which 20 are hidden from the model panel because they are quantities a plant would have to infer rather than read.

Table 9.2 Model Parameters. The Seven Groups Below the Lower Rule Are Hidden From the Application's Model Panel Because They Are Quantities a Plant Would Have to Infer. The Rest Are Adjustable.

Symbol	Meaning	Value	Unit
V	loop volume	65.0	m^3
D	pipe bore	0.65	m
L	developed length	196	m
H	vertical leg height	40	m
A_j	jacketed area	326.7	m^2
ρ_p	particle density	900	kg/m^3
$\rho_{L,0}$	liquid density	440	kg/m^3
κ	compressibility	1.3×10^{-3}	bar
α	thermal expansivity	0.0060	1/K
$H_{H2,0}$	Henry, H_2 at 70 °C	2600	bar
$H_{C2,0}$	Henry, C_2 at 70 °C	130	bar
$H_{I,0}$	Henry, inerts at 70 °C	72	bar
v_r	rated velocity	5.5	m/s
K_f	loop resistance	10.22	—
R_{sh}	shut-off head ratio	1.90	—
$\alpha_{\max}$	maximum void fraction	0.25	—
α_{50}	void at half head	0.055	—
δ	void-growth scale	0.35	bar
F_L	Durand factor	1.00	—
k_f	settling growth rate	0.0040	1/min
E_p	propagation activation	45000	J/mol
$k_{d,0}$	site decay at 70 °C	0.0100	1/min
β_d	decay temperature term	0.045	1/K
$-\Delta H_p$	heat of polymerization	2046	kJ/kg
K_Q	melt-flow constant	3.2552	—
m	exponent on transfer group	2.00	—
b_0	spontaneous transfer	0.060	—
a_T	Arrhenius term on Q	4711	K
k_ρ	density per wt% C_2	0.00170	—
U_0	jacket coefficient	69.0	kJ/min m^2K
M_j	jacket water inventory	12000	kg
P_f	flash pressure	18	bar
$T_{\max}$	swelling limit	80.0	°C
P_{hh}	relief set pressure	44.0	bar
Δt_a	analyzer interval	6.0	min
θ_a	analyzer dead time	4.0	min
σ_v, σ_N	eye depression terms	2.0, 2.0	—
n_{H2}, n_{C2}, n_I	Henry exponents	0.35, 0.75, 0.85	—
$k_{p,\mathrm{ref}}$	propagation, 70 °C	0.0299873	kg m^3/mol min
k_a	catalyst activation	6.4646×10^{-4}	1/% min
ζ	H_2 chain-transfer use	7.13×10^{-4}	—
M_{st}	loop steel mass	78540	kg
μ_L	liquid viscosity	7.5×10^{-5}	Pa s

No scenario hard-codes a state. Every scenario solves for the state that makes its own setpoints an exact equilibrium, using three scalar solves for the inert inventory that balances its own feed against its own purge, the velocity where the pump curve crosses the loop resistance, and the jacket temperature that carries the reaction heat away. Evaluating all nineteen derivatives at the solved design point gives a worst relative residual of 2.7×10^{-15}, which is the numerical noise floor of the arithmetic that produced it. The design point is 70.0 °C and 38.00 bar with 591 kg/m^3 of slurry, which is 50.0 wt% solids and 19.20 t of powder in the 65 m^3 loop, making 22.0 t/h of a 3.00 g/10min homopolymer of density 0.90550 g/cm^3 at $r_H = 0.900\,\mathrm{mol/kmol}$. The residence time is 52.4 min on 9.1 g/min of catalyst, the circulation is 5.50 m/s at 263 kW of pump shaft power and a 35.6 s lap time, and reaction heat is 12.50 MW of which the cold monomer feed carries 1.46 MW and the jacket removes 10.83 MW with water at 41.2 °C through a 45.0 % cold-water valve.

The engine of Section 1.6.1 integrates the plant by fourth-order Runge-Kutta at a fixed step of

0.01 min (0.6 s), which places 2.4 integration steps inside one circulation time constant, well inside the stability limit of the method. Controller execution is every 0.05 min (3 s) and the historian keeps one sample every 0.17 min (10.2 s), every 17th integration step. The noise seed is 20260914.

A setting entered before Run gives the same run as the same setting carried through Reset. The start-up solves the steady state for the chosen setpoints and lines up only the stations the reader has not edited, and a loop switched to Manual or Auto stays where it was put. A setting the plant cannot hold, such as 84 °C or a negative margin, starts the run at the scenario's own steady state and acts on it as a step. A helper button beside the manual stations, "What does this grade cost in margin?", re-solves the steady state for the setpoints entered and reports the melt flow rate, the bubble point, the margin and the reactor pressure that the minimum margin would require.

9.3.7 Assumptions

Modeling effort went into the thermodynamic constraint and was withheld from almost everything else. Every simplification that resulted breaks down somewhere, and the paragraphs below name where.

The bubble point is a surrogate, not an equation of state. Equation (9.6) is linear in each solute mole fraction, carries no binary interaction parameters and treats the inert slate as one lumped species of molecular weight 34. A design calculation would use a cubic or associating equation of state and would find a genuinely nonideal liquid phase [19]. The Antoine fit is claimed only from 0 °C to 90 °C, and above 90 °C it extrapolates toward a critical point it does not represent, propylene's being 91.1 °C and 46.0 bar. The application clamps its temperature slider at 84 °C for that reason, and no number the model produces above 90 °C means anything. The Henry constants and their temperature exponents are fitted, and 2600 bar for hydrogen in a light hydrocarbon at 70 °C is the right order and nothing more.

One well-mixed volume. A real loop is a plug-flow circuit with the 35.6 s lap time the model itself reports, a measurable temperature rise across the pump, and composition gradients between the injection points and the take-off, which is precisely why catalyst enters at the pump discharge and hydrogen enters on a return leg. Residence-time distribution effects on the molecular weight distribution are the subject of a literature this model cannot enter [7].

One lumped active-site type with first-order decay. Ziegler-Natta catalysts have several site families with different hydrogen and comonomer responses, and that multiplicity is the origin of the broad molecular weight distribution that makes polypropylene processable [17]. Hydrogen activation of the catalyst appears in (9.15) as a fitted saturating multiplier rather than as a mechanism.

Quality is two correlations, not a chain-length distribution. There is no polydispersity, no xylene solubles, no isotacticity index and no bimodality in (9.19), so the model says nothing about the properties a converter cares about beyond melt flow rate and density.

Cavitation is a lumped void fraction with an algebraic breakdown curve. Equations (9.9) and (9.10) capture the shape of a real head breakdown and none of the acoustics, the damage rate, or the difference between incipient and fully developed cavitation. The eye depression in (9.8) is a two-term surrogate with fitted coefficients rather than a computed suction-specific speed.

Fouling is one scalar. Equation (9.14) represents settled powder in the horizontal runs. Real loop fouling includes a wall polymer film, grows preferentially at the bends, and is removed in a shutdown rather than by running faster.

The flash and the recycle are fixed recovery fractions. Inerts accumulate because recovery returns 85 % of what flashes and only the purge and the remaining 15 % remove them, which is the right mechanism on the right timescale, but the purge composition is assumed through enrichment factors rather than computed from a separation.

The state clamps bound a runaway rather than model it. Reactor temperature is clipped at 94 °C after every step, and the post-trip excursions reported in Section 9.6 run into both that clamp and the 90 °C validity limit of the Antoine fit. Those numbers are correct as statements about the simulator and meaningless as statements about a reactor.

9.4 The Constraint and Its Gradients

9.4.1 What Each Dissolved Species Costs

At the design point the 34.57 bar bubble point divides into 28.86 bar of propylene vapor pressure, 3.48 bar of accumulated light inerts at 4.83 mol% of the liquid, 2.23 bar of hydrogen at 856 mol ppm, and nothing from ethylene because the design grade is a homopolymer. The order of those two middle terms is the first surprise the application delivers, because the inerts nobody set on purpose are worth more bubble point than the hydrogen the quality controller is spending.

🕮 **Pause.** Try Exercise 9.1 before reading on.

Four gradients follow, each measured by re-solving the steady state on either side of the design point rather than by differentiating an expression. Raising r_H by one unit raises the bubble point by 2.44 bar, and the figure is almost exactly linear, at 2.439 bar per unit taken as a secant from 0.60 to 1.50 and 2.440 bar per unit taken locally. Raising r_E by one unit costs 0.87 bar locally at the homopolymer point and 0.86 bar taken as a secant out to $r_E = 2.0$. Each degree of reactor temperature is worth 0.629 bar. Reactor pressure is the only handle that buys margin, and it buys it at 1.004 bar per bar, because raising the pressure moves the bubble point itself by only −0.004 bar per bar through the liquid density.

The application reports the constraint as the sum Equation (9.6) makes it rather than as a single number. All four terms are trended separately on the pump view, and the hydrogen and ethylene terms appear again on the faceplate of the loop that moves them, beside the dissolved concentration in mol ppm and the valve position that produced it. Raising $r_{H,\mathrm{SP}}$ is therefore watched three ways at once: the hydrogen term of the sum rises, the margin gives up very nearly the same amount within the minute, and the melt flow rate the setpoint was asked for arrives an hour later. That is the whole collision, in three traces a student can read off one screen.

9.4.2 The Grade Envelope

Those numbers turn a quality request into a pressure requirement, as shown in Figure 9.2. The bubble point rises linearly with dissolved hydrogen while melt flow rate rises as the square of it, so the two lines that matter are the reactor pressure and the reactor pressure less the minimum margin. At 38.00 bar with a 2.00 bar minimum margin the bubble point may reach 36.00 bar, which is $r_H = 1.488$ and a melt flow rate of 7.80 g/10min. A 6 g/10min grade still leaves 2.46 bar. The 12 g/10min grade needs $r_H = 1.860$ and a bubble point of 36.91 bar, so it leaves 1.09 bar of margin at 38.00 bar and 4.11 bar at 41.00 bar. Raising the pressure to 41.00 bar moves the ceiling from 7.80 g/10min all the way to 25.2 g/10min, which is the whole 25 g/10min fiber grade, because melt flow rate is quadratic in the group that the linear margin constrains.

9.4.3 Three Time Scales in One Plant

The reason the collision cannot be tuned away is that the two controllers who want the hydrogen valve read clocks that differ by a factor of five to eight, and the equipment they are arguing about reads a third clock faster than either. Stepping r_H from 0.900 to 1.500 with the inventory loops closed and the quality loops open gives a liquid ratio that reaches 63 % of its change in 13.4 min, a bubble point that reaches 63 % of its change in 13.8 min, and a product melt flow rate that reaches 63 % of its change in 74.5 min. The margin is that bubble point read against a reactor pressure that wanders by a few tenths of a bar with the instrument noise, so its own 63 % crossing comes at 8.9 min in the shipped run and as late as 11.6 min on five other noise seeds. The same step moves melt flow rate from 3.00 to 7.92 g/10min and the steady-state margin from 3.43 to 1.97 bar. Against the 1.45 s circulation time constant of (9.11), the quality loop is roughly 3000 times slower than the equipment it can destroy. That ratio, and not a tuning constant, is why the constraint needs its own controller with its own measurement and its own authority.

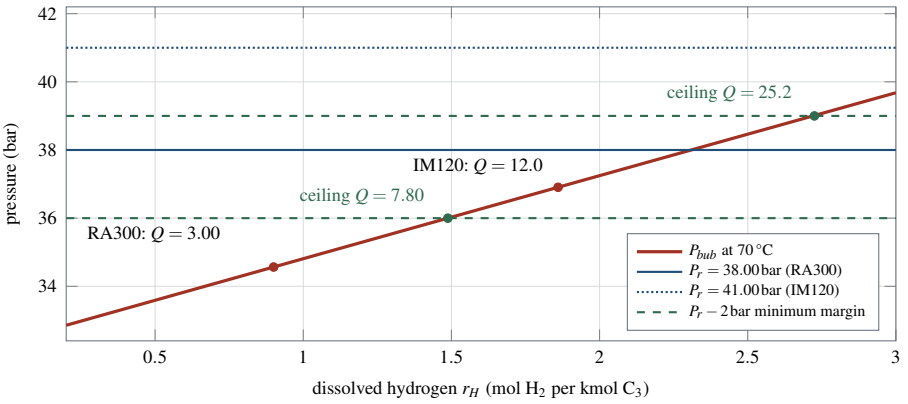

Figure 9.2 Bubble point of the loop liquid against dissolved hydrogen, with the reactor pressure and the minimum-margin line for two grade recipes. Every point is a separately solved steady state at 70 °C and the design inert level. The bubble point is the thick solid curve, the reactor pressure is a solid line at 38.00 bar and a dotted line at 41.00 bar, and the minimum-margin lines are dashed. Where the thick curve crosses a dashed line is the largest melt flow rate that pressure can support, 7.80 g/10min at 38.00 bar and 25.2 g/10min at 41.00 bar. The 12 g/10min grade sits above the 38 bar ceiling and comfortably below the 41 bar one.

9.5 Control Architecture

9.5.1 The Implemented PID

All ten loops run the velocity-form PID of Section 1.6.2 on a 0.05 min execution period, with the derivative on a measurement PV^f that a 0.4 min first-order filter smooths, advanced by its exact zero-order-hold recursion. Its limiting step carries two features of this app, as shown in Equation (9.21).

$$
\begin{aligned}
c_k &= u_{b,k} + s_k + I_{k-1} \\
u_k &= \mathrm{slew}\big(\mathrm{clamp}(c_k, u_{\min}, u_{\max,k}),\, u_{k-1},\, S\big) \\
I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t}\right] \Delta t_c
\end{aligned}
\tag{9.21}
$$

where e_k is the error of Equation (1.14), s_k is the never-clipped running sum of proportional and derivative increments of Equation (1.15), c_k is the demanded output (v_k in Equation (1.16)), $u_{b,k}$ is the output bias, $u_{\max,k}$ is an upper clamp that may move from scan to scan, S is the output slew rate, Δt_c is the execution period and $T_t = \tau_I$ is the tracking time, since no loop sets its own. The direction a of Table 9.3 is $+1$ on a loop whose process gain is positive and -1 where it is negative. Because the $(u_k - c_k)/T_t$ term acts on the difference between what was applied and what was demanded, output clamping and slew limiting both bleed the integral down over one integral time instead of resisting it. Manual operation, an interlock that overrides a loop, and a melt-flow controller that loses the low selector all set I_k directly so that c_k equals the applied output, which makes every transfer bumpless without any extra logic. The bias $u_{b,k}$ of each controller is preloaded with the actuator position that the solved equilibrium requires of that loop, which leaves both the error and the integral at zero when a scenario starts, and it stays constant on every loop except the constraint controller of Section 9.5.3.

9.5.2 Ten Loops in Three Tiers

Table 9.3 lists the loops, their default tunings and the response time each one has to live with. Three regulatory loops are fast, three are medium and two of the three quality loops carry the analyzer dead time.

Table 9.3 Implemented Loops With Their Default Tunings. Direction $a = +1$ Where the Process Gain Is Positive and -1 Where It Is Negative.

Loop	Actuator	K_c	τ_I (min)	a	Output range
Jacket water T_j	v_{CW}	8	0.6	-1	0–100 %
Reactor T_r	$T_{j,\mathrm{SP}}$	8	8	$+1$	8–62 °C
Pressure P_r	v_{C3}	4	2.5	$+1$	0–100 %
Slurry ρ_{sl}	v_{TO}	0.8	12	-1	0–100 %
Production R	v_{cat}	1.2	30	$+1$	0–100 %
Ratio r_H	v_{H2}	15	10	$+1$	0–100 %
Ratio r_E	v_{C2}	20	8	$+1$	0–100 %
Melt flow Q	$r_{H,\mathrm{SP}}$	0.55	45	$+1$	0.02–5.0
Density ρ	$r_{E,\mathrm{SP}}$	700	50	-1	0–6.0
Margin ΔP	$r_{H,\max}$	1.5	8	-1	0.02–clamp

The melt-flow and density controllers write composition setpoints rather than valve positions, and both are rate-limited at 0.05 units of ratio per minute. That limit is not a tuning constant in disguise. Hydrogen entering the loop leaves again only through a twenty-minute inventory, so a quality controller allowed to slam its setpoint spends bubble-point margin it cannot get back, and the rate limit is also what gives the constraint controller time to be a controller rather than a trip. Five architectures are offered as one selection. Manual leaves every valve to the operator. Inventory closes temperature, pressure, slurry density, production and both composition ratios and leaves melt flow rate and density to come out as they will, which is how loop reactors were run for years and is not unreasonable, because the ratios are what the catalyst sees. Quality adds the two outer quality cascades. Override adds the margin controller and the selector. The fifth keeps the override structure whole and places the predictive controller of Section 9.5.5 over it.

Each loop in Table 9.3 is presented as its own faceplate, carrying a MAN/AUTO switch, the setpoint it reads, the manual output station it writes and its two tuning constants, so every manipulated variable in the table is reachable directly and opening a cascade at any tier hands the student the signal that controller had been writing. The two composition ratios are the ones that matter pedagogically. Both $r_{H,\mathrm{SP}}$ and $r_{E,\mathrm{SP}}$ are outputs of the quality tier and setpoints to the regulatory tier, and taking the melt-flow controller to manual therefore turns $r_{H,\mathrm{SP}}$ into a direct operator handle on dissolved hydrogen, which is the same variable Equation (9.6) charges for. That is not a simulator convenience. A polyolefin plant, liquid loop or gas-phase fluidized bed (Chapter 8) alike, is operated through composition ratios rather than through feed valves, because the ratio is what the catalyst sees while the polymer property it sets is an inventory away. The ratio setpoint is the handle an operator is given and the flow controller beneath it is the handle the ratio controller is given.

9.5.3 The Low Selector

The melt-flow controller and the bubble-point controller both want the hydrogen setpoint and they want it in opposite directions, so a low selector decides, as shown in Equation (9.22).

$$r_{H,\mathrm{SP}} = \min\left(r_H^{Q},\; r_H^{B}\right) \tag{9.22}$$

where r_H^{Q} is the melt-flow controller's demand and r_H^{B} is the bubble-point controller's allowance. The acetylene converter of Chapter 3 takes its reagent away from a quality loop through the same kind of selector. Three details decide whether such an arrangement works [3]. The loser must not wind up, so whichever controller did not write the setpoint is told what was actually applied and its integral is tracked to it through the same track operation manual mode uses. The constraint controller must be able to win, so its output range reaches 0.02 mol/kmol, which is essentially no hydrogen at all.

And the idle controller must be parked where it can act, so its upper clamp and bias are set every execution to $\min(5.0, \max(r_{H,\mathrm{SP}}, r_H) + 0.5)$ rather than to the top of its range, with back-calculation holding it there. Section 9.6 measures what that third detail is worth. The implementation carries two separate flags out of the arrangement, one for the strict comparison and one for whether the constraint controller has left that clamp, because Section 9.6 shows that the two answer different questions and that only the second is fit to drive an operator indicator.

9.5.4 Hard Interlocks

Underneath every architecture are seven interlocks that no switch turns off, because each prevents an outcome that ends the run rather than degrading it. Hydrogen is slammed shut below 0.20 bar of margin at the pipe wall. Catalyst is limited to 15 % when the void fraction exceeds 2 % and to 30 % when the circulation falls below the deposition velocity. Catalyst is cut and the cold-water valve driven wide open above the 80.0 °C swelling limit, and above 82.0 °C the catalyst and every gas feed are latched shut until the unit is reset, because a loop full of tacky agglomerate is emptied with a hydro-jet rather than with a valve. The take-off is forced to at least 75 % above 60 wt% solids, and the propylene make-up is limited to 3 % above 43.5 bar. What matters pedagogically is which of those protects what. The hydrogen interlock defends the vessel against flashing its whole liquid-full contents. It is not a pump protection, and Section 9.6 shows it firing 2.15 min after the impeller has already been full of vapor.

9.5.5 Model Predictive Control

The fifth architecture puts a linear model predictive controller over the override structure [20, 21], an instance of the option of Section 1.6.3. It is never the application's default, and only scenario 11 loads it. Every $T_s = 2\,\mathrm{min}$, or 40 controller executions, it plans the three supervisory setpoints $u = [r_{H,\mathrm{SP}}, r_{E,\mathrm{SP}}, T_{\mathrm{SP}}]$ over $N = 90$ steps (180 min) in twelve move blocks of $\{1,1,1,1,2,2,3,4,6,9,15,45\}$ steps and applies the first move until the next plan, the first plan being solved at $t = 0$. The regulatory loops of Table 9.3 and the bubble-point controller stay closed beneath it. Its hydrogen request takes the place of r_H^Q in the low selector of (9.22), so the bubble-point controller and the seven interlocks keep the authority they have without it, and the displaced melt-flow and density controllers track the applied setpoints on every execution with no preload. The plan acts only while the header and the melt-flow, density, both ratio and reactor-temperature loops are in Auto.

The controller's model is its own. It sees the plant of Section 9.3 only through the instruments, the setpoints in force, the declared parameter values and its own state, so a parameter changed during a run is mismatch it is not told about. Its eight states are the two ratios with one intermediate lag each, the reactor temperature, the product log melt flow rate q, the product ethylene content e and a constant, as shown in Equation (9.23).

$$\begin{aligned} \tau_{H2}\dot{h}_H &= r_{H,\mathrm{SP}} - h_H, \qquad & \tau_{H1}\dot{r}_H &= h_H - r_H \\ \tau_{E2}\dot{h}_E &= r_{E,\mathrm{SP}} - h_E, \qquad & \tau_{E1}\dot{r}_E &= h_E - r_E \\ \tau_T \dot{T}_r &= T_{\mathrm{SP}} - T_r \\ \hat{\tau}\,\dot{q} &= \ln Q_{\mathrm{inst}}(r_H, r_E, T_r) - q, \qquad & \hat{\tau}\,\dot{e} &= w_{C2}(r_E) - e \end{aligned} \tag{9.23}$$

Here $\tau_{H1} = 9.5$, $\tau_{H2} = 2.35$, $\tau_{E1} = 16.3$, $\tau_{E2} = 1.7$ and $\tau_T = 1.5\,\mathrm{min}$, Q_{inst} and w_{C2} are the correlations of (9.19) at the declared constants with no poison, and the residence time $\hat{\tau}$ (52.4 min at design) is inferred from the measured slurry density and production. The outputs are $q + b_Q$, the density $\rho = \rho_h - k_\rho e + k_{\rho Q}(q + b_Q - \ln Q_{\mathrm{ref}}) + b_\rho$ and the margin $\Delta P = P_{r,\mathrm{SP}} - P_{bub}(r_H, r_E, T_r) + b_m$ from (9.6) at the design inert fraction of 4.83 mol%. The correlations and the bubble point are relinearized at each decision at the end of the previous plan, with the affine remainder kept, and discretized by an exact zero-order hold over T_s.

The ratio lags come from setpoint steps on scenario 5 with the inventory loops closed and no noise. They reproduce those responses with root-mean-square residuals of 3.7 % for hydrogen and 1.7 % for ethylene, and an independent two-lag fit gives 6.0 and 5.8 min for hydrogen, a different split of nearly the same total lag (11.9 against 11.85 min). The temperature loop reaches a setpoint step within about a minute and overshoots it by about 40 %, and its 1.5 min lag (residual 7.2 %) is slow on purpose because the plan moves T_{SP} at only 0.05 °C/min. At the design point the model gives 4.08 % of melt flow rate and −0.609 bar of margin per degree (against the plant's 0.629 bar of bubble point) and −2.44 bar per unit of r_H. The fixed inert fraction puts the model's margin 0.07 bar below the plant's at 66 °C, and over six decisions of scenario 11 the prediction 60 min ahead stays within 10 % of the melt flow rate and 0.36 bar of the margin.

The estimator runs on every execution whatever architecture is selected, integrating q and e from the measured ratios and temperature and running the lag states on the applied setpoints. The biases b_Q and b_ρ move by half of each analyzer residual, taken against the estimate held when the sample was drawn, and b_m filters over 5 min the difference between the inferred margin the override reads and the model's margin, which makes the prediction offset free [22], as the Kalman estimator of Volume 9, Chapter 4 does with a fuel bias. No event is shown to the plan in advance. Each plan minimizes Equation (9.24) subject to Equation (9.25).

$$
\begin{aligned}
J = {} & \sum_{k=1}^{N} w_k \left[\left(\frac{q_k + b_Q - \ln Q_{\mathrm{SP}}}{s_Q} \right)^2 + \left(\frac{\rho_k - \rho_{\mathrm{SP}}}{s_\rho} \right)^2 \right] \\
& + \sum_{j=1}^{12} \Delta u_j^{\top} S \, \Delta u_j + r_T \sum_{k=1}^{N} \left(T_{\mathrm{SP},k} - T_g \right)^2 + \sum_i \rho_i \varepsilon_i^2
\end{aligned}
\tag{9.24}
$$

$$
\begin{aligned}
& 0.02 \le r_{H,\mathrm{SP}} \le 5.0, \qquad 0 \le r_{E,\mathrm{SP}} \le 6.0 \\
& \max(58, T_g - \delta_T) \le T_{\mathrm{SP}} \le \min(84, T_g + \delta_T) \\
& |\Delta u_j| \le [0.1,\ 0.1,\ 0.1\,^{\circ}\mathrm{C}], \qquad \varepsilon_i \ge 0 \\
& |q_k + b_Q - \ln Q_{\mathrm{SP}}| \le s_Q (1 + \varepsilon_Q) \\
& |\rho_k - \rho_{\mathrm{SP}}| \le s_\rho (1 + \varepsilon_\rho) \\
& \Delta P_k \ge \Delta P_{\min} + \beta - \varepsilon_m
\end{aligned}
\tag{9.25}
$$

where $s_Q = \ln 1.1$ and $s_\rho = 0.0006\,\mathrm{g/cm^3}$ are the specification half-bands, $w_k = 1$ except $w_N = 6$, $S = \mathrm{diag}(3, 3, 30)$ weighs block moves measured from the command applied now, $r_T = 0.1$ per °C^2 draws the temperature toward the grade's own T_g within the freedom $\delta_T = 4\,^{\circ}\mathrm{C}$, $\Delta P_{\min} = 2.0$ bar is the override's setpoint and $\beta = 0.7$ bar the back-off above it. The input limits are the quality controllers' clamps, and the move caps are T_s times the rates at which the applied commands travel between decisions, 0.05 per minute on each ratio and 0.05 °C/min on the temperature. The bands are soft, each with one slack for the whole horizon, judged every third step and priced at $\rho_Q = \rho_\rho = 10$ per band squared and $\rho_m = 10^4$ per bar^2. The program has 39 variables. With one block and no constraints it agrees with independently formed normal equations to 1.7×10^{-13}, and a plan took 1.75 ms on average in Node.js over the 149 decisions of a five-hour scenario 11.

The setpoints follow a plan only when the solver reports it optimal, every number it returns (moves, predicted outputs, objective, slacks) is finite, and the moves respect their bounds. Any other plan leaves the setpoints where they stand, marks the run and hands the setpoints to the tracked melt-flow and density controllers under the override, and control returns with the first later plan that passes. With a failure forced from 100 to 110 min in scenario 11 and no noise, no execution moves a ratio setpoint by more than 0.0025, one execution of its slew, and the hydrogen valve moves by at most 0.049 %.

Table 9.4 Scenario Sequence and the Mechanism Each Shows

#	Scenario	What it shows
1	Manual at the design point	two clocks from one valve move
2	r_H step 0.90 to 1.50	gradient measurement, dead time
3	Production 22 to 27 t/h	activity dynamics, jacket ceiling
4	Inventory loops only, 7 h	what ratio control cannot see
5	Quality cascades lined out	tuning against the constraint
6	Grade transition, no override	cavitation, settling, trip
7	Same transition with override	the selector, and what it costs
8	Inerts creeping over a shift	a constraint spent by nobody
9	Cold night then fouled jacket	heat path into the bubble point
10	Pump speed 100 to 80 %	eye margin against settling
11	The override as a constraint	a predictive plan over the override
12	Graded challenge, 5 h	buy the margin before spending it

9.6 Results

Table 9.4 lists twelve scenarios in the order a class would work through them. A scenario pins the architecture, the grade recipe, the ten tunings, every disturbance setting and the timed events that arrive during the run.

9.6.1 The Transition That Destroys the Pump

Scenario 6 asks for 12 g/10min at $t = 15$ min with quality cascades in service, no override, and a melt-flow controller at $K_c = 1.6$ and $\tau_I = 20$ min, which is the tuning a step test at the starting grade produces. Nothing looks wrong for half an hour. The following sequence is read from the running model.

At $t = 51.35$ min, 36.35 min after the setpoint move, the impeller-eye margin goes negative while the transmitter at the pipe wall still reads 0.354 bar of margin and the liquid ratio has reached 2.169. Within 0.01 min the void fraction passes the 0.5 % at which the application counts cavitation, with the circulation still at 5.5 m/s. At $t = 53.4$ min the pump is delivering less than half its head and the velocity has fallen to 4.48 m/s. At the scan at $t = 53.5$ min the hydrogen interlock fires on 0.20 bar of wall margin, by which time the eye margin is −0.09 bar and has been negative for 2.15 min. At $t = 55.2$ min the velocity crosses the 3.67 m/s deposition velocity and the settled layer starts to grow. At $t = 67.6$ min the reactor passes the 80 °C swelling limit, and at the scan at $t = 68.25$ min the high-temperature trip latches the unit off at 82.1 °C with the circulation down to 1.61 m/s. Four hours after the start the run has spent 44.9 min cavitating, has 15.2 % of the bore under settled powder, and has made 44.2 t of polymer of which 38.7 t is off specification. Its worst reported margin of −20.1 bar is reached at $t = 78.5$ min, well after the trip and inside the region where the state clamp and the Antoine validity limit of Section 9.3.7 both apply, so the defensible statement is that the wall margin went negative at $t = 62.6$ min and stayed negative until the tripped reactor had cooled, at $t = 95.2$ min. On five other noise seeds the trip arrives between 64.5 and 125.3 min and the pump cavitates for 42.7 to 70.1 min.

Two things in that sequence are worth separating. The steady state alone would have survived, because 12 g/10min at 38.00 bar leaves 1.09 bar of margin. What cavitates the pump is the overshoot, since the liquid ratio reaches 2.29 against the 1.86 the grade needs while the analyzer has not yet registered any progress. And the eye margin, not the wall margin, is what decides, because the two differ by 0.36 bar at design flow and every alarm and interlock in the plant is written against the wall.

Table 9.5 Melt Flow Rate Step From 3 to 8 g/10min at $t = 20$ min, Read at $t = 300$ min. Both Columns Use the Default Tuning of Table 9.3. The Architecture Switch Is the Only Difference.

	Quality cascades	Plus override
Cavitation (min)	42.3	0.0
Worst margin (bar)	−19.82	+1.07
Minimum velocity (m/s)	1.51	5.50
Peak reactor temperature (°C)	94.0	70.2
Settled layer (% of limit)	15.2	0.0
Reactor still running	no	yes
Polymer made (t)	45.6	110.6
Off specification (t)	33.4	103.3
Melt flow rate at 300 min	1.71	6.55
Pump energy (MWh)	1.16	1.32
First cavitation (min)	56.5	—
First constraint selection (min)	—	45.7

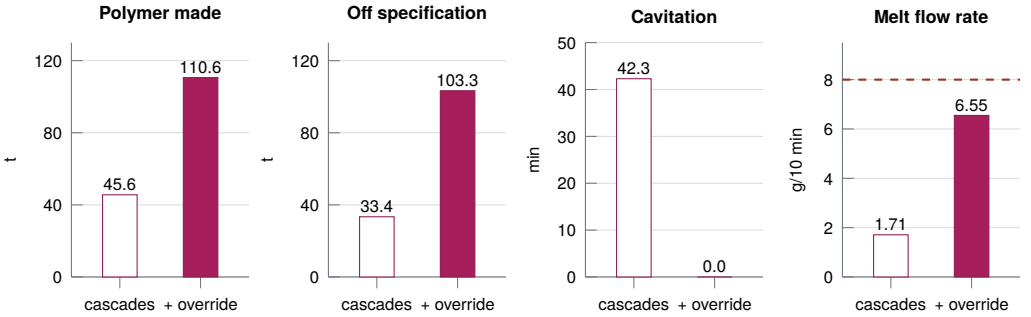

Figure 9.3 Four rows of Table 9.5 as bars, open for the quality cascades and solid for the cascades plus the override, each run read at 300 min. The dashed line in the last panel is the melt flow rate setpoint of 8. Notice that the override removes every minute of cavitation and more than doubles the polymer made, yet more of that polymer is off specification, because the margin it holds stops the melt flow rate short of the setpoint.

> **TRY IT IN THE APP** **6 · Grade Transition RA300 to IM120, No Override**
>
> Load the scenario, press Run and watch the Bubble-point margin strip on the Pump & bubble point view. After the melt-flow setpoint moves at 15 min, the dashed line at the impeller eye crosses zero at 51.35 min while the transmitter at the pipe wall still reads 0.354 bar. The circulation then slows and the reactor trips at 68.25 min.

Pause. Try Exercise 9.2 before reading on.

9.6.2 One Variable Changed

Table 9.5 compares the same run with and without the override, with the architecture switch as the only difference between the two columns. To keep the comparison free of the aggressive tuning that scenario 6 supplies, both runs start from scenario 5 at the design grade with the default melt-flow tuning of $K_c = 0.55$ and $\tau_I = 45$ min, and both receive a setpoint step from 3 to 8 g/10min at $t = 20$ min. The 8 g/10min target is chosen because its steady state at 38.00 bar does not cavitate the pump, at 1.95 bar of margin at the wall and 1.59 bar at the impeller eye, although it lies just above the 7.80 g/10min ceiling that a 2.00 bar minimum margin allows. Figure 9.3 draws four rows of the table as bars.

The override does exactly what a constraint controller is for and nothing more. It first takes the hydrogen setpoint at $t = 45.7$ min, 11 min before the unprotected run cavitates, and the pump then never leaves 5.50 m/s, the settled layer never grows, and the reactor is still running at five hours having made 110.6 t against 45.6 t. It also does not deliver the grade. Melt flow rate at 300 min is 6.55 g/10min against a setpoint of 8, which the 2.00 bar margin setpoint never allows at 38.00 bar, and 103.3 t of the 110.6 t is off specification, which is more off-specification polymer in absolute tonnes than the run that destroyed the pump produced. That is the honest trade and it is the centerpiece of the application. An override buys the equipment, not the product, and reading only the off-specification tonnage would rank the two runs the wrong way round. On five other noise seeds the protected run again never cavitates, while the unprotected run trips on four of them and on the fifth cavitates for 45.4 min without tripping.

The same comparison on scenario 6 and 7, where the melt-flow tuning is the aggressive $K_c = 1.6$ and $\tau_I = 20$ min and the target is 12 g/10min, gives 0.0 min of cavitation against 44.9 min, a worst margin of 1.21 bar against −20.1 bar, no settled layer against 15.2 %, and 88.6 t made against 44.2 t over four hours. Melt flow rate reaches only 7.49 g/10min at four hours against a setpoint of 12, and 83.2 t of the 88.6 t is off specification.

> ✦ DID YOU KNOW? **The Run That Saved the Pump Made More Off-Spec**
>
> From the simulator: in Table 9.5 the run without the override destroyed its pump and made 45.6 t of polymer, 33.4 t of it off specification. The run with the override kept the pump and made 110.6 t, of which 103.3 t was off specification, because the margin it protected never allowed the grade.

9.6.3 The Selector Chatters, and It Matters

A low selector whose loser is tracked to the winner's output has a diagnostic problem, and the application is where the problem and its remedy are both visible. Once the melt-flow controller's integral has been dragged down to the allowance, its own output is approximately equal to the allowance, so a strict inequality test reports that the constraint is not selected even while the constraint is what determines the answer. Counting the fraction of each of the first five twenty-minute windows after selection begins at $t = 41.8$ min in scenario 7 in which the constraint controller strictly wins gives 33.7 %, 19.8 %, 7.0 %, 19.0 % and 16.8 %, and the total is 39.9 min of the 198 min that follow. Over $70 \le t \le 240$ min the margin averages 1.99 bar against its 2.00 bar setpoint, and over the last hour of the run the melt flow rate is pinned between 7.2 and 7.5 g/10min against a setpoint of 12, so the constraint clearly governs the outcome for the whole period. An indicator built on the strict test alone therefore under-reports by a factor of about four and flickers at the execution rate.

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

The remedy is to ask a different question. Rather than testing which controller wrote this execution's setpoint, the application tests whether the constraint controller has come off the parking clamp of Section 9.5.3, which back-calculation holds it against for exactly as long as the margin is comfortable. That test reads true for 167.2 min of the same 198 min, and it is what colors the selector card, drives the badge and shades the hydrogen setpoint strip. The strict test is kept as a third state of the same panel and as its own exported column, so a student can see the difference between "the constraint wrote this setpoint" and "the constraint is holding the hydrogen down" rather than having one silently substituted for the other. The general lesson survives the fix. A selector indicator has to report authority and not arbitration, and the trend to read alongside either state is the separation between what the quality controller asked for and what the constraint allowed.

> **TRY IT IN THE APP** **7 · The Same Transition, With the Bubble-Point Override**
>
> Press Run on the Quality view and watch the low selector panel. Selection begins at 41.8 min. Over the 198 min that follow, the strict test of which controller wrote the setpoint is true for only 39.9 min, while the panel reports the constraint holding the ceiling for 167.2 min. The pump never cavitates.

The residual handover also leaves a real cycle. Averaging the inferred margin in three-minute blocks over $70 \le t \le 240$ min in scenario 7 gives a range of 1.70 to 2.37 bar, a peak-to-peak swing of 0.68 bar, and about 31 min between like extrema (0.36 to 0.64 bar and 23 to 31 min on five other noise seeds), which is the same order as the hydrogen inventory time constant the constraint loop is working against. Detuning the constraint controller from $\tau_I = 8$ min to 20 min cuts the swing to 0.29 bar and changes the worst margin only from 1.21 to 1.17 bar, while raising its gain from 1.5 to 3.0 raises the swing to 1.35 bar and costs 0.91 g/10min of melt flow rate averaged over the last hour. The shipped tuning is therefore a defensible compromise rather than a settled loop, and the application says so in the scenario note rather than claiming a settled margin.

The third detail of Section 9.5.3 is worth a number of its own. Forcing the constraint controller's clamp and bias back to the top of its 5.0 mol/kmol range on every execution, which is where an unconsidered implementation leaves it, changes the worst margin in scenario 7 from 1.21 to 0.48 bar, a loss of 0.74 bar out of a 2.00 bar setpoint, because the idle controller has to travel its whole range downward before it can act. That version does not cavitate in this scenario, and its ripple of 0.94 bar peak-to-peak is of the same order as the shipped 0.68 bar, so the cost of getting the parking wrong here is measured in margin rather than in a new oscillation.

9.6.4 Constraints Spent by Nobody in Particular

Scenario 4 runs the inventory loops alone for seven hours through a weak catalyst batch, a lighter feed specification and a trace of oxygenate. Melt flow rate holds 3.00 through the weak batch, reading 2.978 at 120 min and 3.003 at 240 min, while production dips from 22.0 to 16.9 t/h, because catalyst activity sets how fast chains are made and hydrogen sets how long they are. The oxygenate then moves melt flow rate to 2.65 g/10min, 11.7 % off target, with the H_2/C_3 ratio still at 0.901. Meanwhile the light inerts climb from 4.83 to 7.58 mol% and take the margin from 3.43 to 2.30 bar with every controlled variable on setpoint. Ratio control is an excellent inner loop and a poor outer one, which is the argument for the cascade, and the inert creep is invisible from every instrument except the loop chromatograph.

Scenario 8 makes that the whole story. A cheaper propylene cargo and a purge trimmed from 100 % to 30 % at $t = 20$ min take the inerts from 4.83 to 10.62 mol% by $t = 240$ min, raise their share of the bubble point from 3.48 to 7.66 bar, and walk the margin down onto the override setpoint at $t = 105$ min. The override takes its first hydrogen 91 min after the purge cut, which is the window a monitoring calculation on the loop chromatograph would have to flag the accumulation. The constraint controller then begins taking hydrogen away for a reason that has nothing to do with hydrogen. The liquid ratio falls from 0.900 to 0.36 at $t = 233$ min and the product melt flow rate to 1.30 g/10min at $t = 273$ min against a setpoint of 3.00, although the purge is restored to 130 % at $t = 240$ min, and 63.7 t is off specification by $t = 360$ min. Nothing alarms until the margin does.

Scenario 9 supplies the thermal route to the same place. Fourteen degrees of free cooling water takes the cold-water valve from 45.7 % to about 26 % and the jacket-limited production ceiling from 27.4 to 34.5 t/h. Removing 40 % of the jacket coefficient then drives the jacket water from 41.2 to a transient 17.5 °C and an average of 21.9 °C, with the cold-water valve reaching 100 % transiently and then working around 62 %, overshoots the reactor to 72.43 °C on the way, and leaves a ceiling of 23.8 t/h with 93 % of it in use on average. At 0.629 bar of bubble point per degree, the 2.43 °C overshoot alone is 1.53 bar of margin, and the worst margin over the run is 0.70 bar. The route from a fouled jacket to a cavitating pump is four steps long and there is no alarm on the first three.

Scenario 10 is the pump itself. Trimming the speed from 100 to 80 % saves 128 kW of shaft power, from 263 down to 135 kW, and improves the impeller-eye depression from 0.36 to 0.23 bar so that the steady-state eye margin actually rises from 3.08 to 3.21 bar. Everything else gets worse. The velocity falls from 5.50 to 4.40 m/s, which is 1.21 times the deposition velocity instead of 1.51, the jacket coefficient falls with $v^{0.80}$ so the heat-removal capacity in use climbs from 81 to 92 %, and the production ceiling falls from 27.4 to 24.2 t/h. Solving the steady state across the speed range puts the settling threshold at 66.4 % speed, so that at 66 % the circulation is already down to 3.63 m/s against the 3.65 m/s deposition velocity.

9.6.5 The Challenge

The challenge asks for 12 g/10min held within ±1.2 g/10min with the margin never below 2.00 bar and no cavitation at all, over five hours, starting from quality cascades with no override and a melt-flow controller at $K_c = 1.8$ and $\tau_I = 14$ min. Left alone that run cavitates the pump from $t = 44.3$ min for 115.5 min in all, takes the worst margin to −0.33 bar, puts 7.6 % of the bore under settled powder, and makes 72.8 t of which 47.2 t is off specification. It ends at 12.47 g/10min, inside the band, and it does not trip in the shipped run, but on five other noise seeds the same run trips the reactor between 59 and 109 min.

The answer is not a tuning constant alone. Raising the reactor pressure to 41.00 bar first is what the IM120 grade recipe in the application's own grade selector specifies, because on a liquid-full loop the operating pressure is part of the recipe in the same way the temperature is. Putting the override in service costs nothing when it is not needed. And slowing the melt-flow controller to $K_c = 0.15$ and $\tau_I = 80$ min recognizes both that the process gain doubles between 3 and 12 g/10min and that hydrogen leaves the loop only through a twenty-minute inventory, so an overshoot is paid for in bubble point long before it appears in melt flow rate. Done that way the five hours reach 11.88 g/10min, hold a minimum margin of 2.14 bar, never cavitate, and make 111.0 t of which 30.25 t is off specification. The override never acts in that run, and neither of the other two changes is enough alone, since 41.00 bar with the given tuning still trips the reactor after 45.2 min of cavitation and the slow tuning at 38.00 bar cavitates the pump for 66.9 min. The suggested settings give the same result whether they are entered before Run or carried through Reset. The general shape of that answer is to move the operating point until the constraint is not binding, keep the override as insurance, and stop asking a slow loop to move quickly.

The application is shown during that transition in Figure 9.4, at an instant when the strict selector comparison has momentarily handed the pen back while the constraint is still holding the ceiling. Each number quoted in that caption was read from the captured run and agrees with the same model stepped to $t = 51.01$ min outside the browser.

9.6.6 Comparison With Model Predictive Control

Scenario 11 repeats the transition of scenario 7, melt-flow tuning included, with the predictive architecture selected, and its own comparison is the same case reloaded with the override structure, which reproduces scenario 7 to the last digit. Table 9.6 runs the shipped structure at its shipped tuning and the plan at the settings of scenario 11 on that transition and on four other cases. Each pair shares its seed, instruments, noise, events and duration, the plan is not retuned between cases, and neither structure is shown anything in advance. Beside the headline the table carries the counterfactual runs that decide what the headline means.

Over four hours the plan never lets the override act, against 39.9 min of strict selection and 169.7 min off the parking clamp under the override structure, and the true margin never falls below 2.0 bar, with a worst value of 2.21 bar against 1.21 bar and 97.8 min below 2.0 bar. The product enters the 10.8 to 13.2 g/10min band at 136.8 min and stays, and 43.7 t of the 87.1 t made is off specification against 83.2 t of 88.6 t, as shown in Figure 9.5.

The comparison is not like-for-like, because the plan holds a temperature handle the override structure lacks. It lowers T_{SP} to 66.0 °C by 97.0 min, the floor of its freedom and 4 °C below the

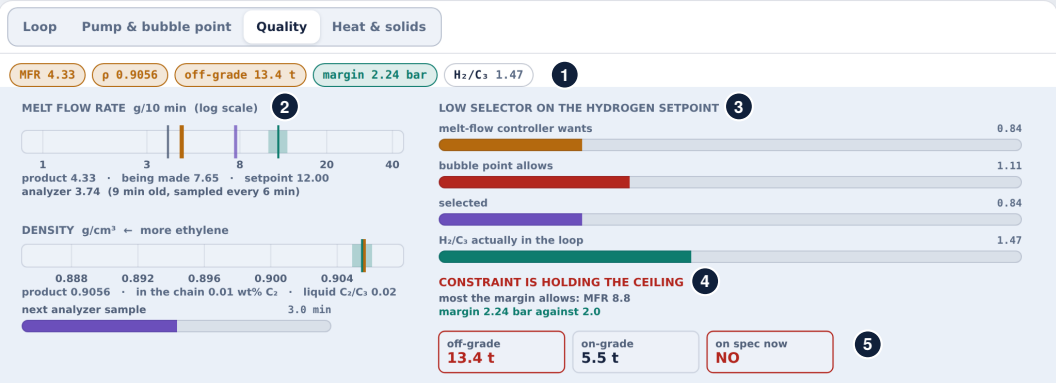

Figure 9.4 The Quality view of scenario 7 at $t = 51.01$ min, 36 min after a melt flow rate setpoint step from 3 to 12 g/10min with the bubble-point override in service. (1) The badges carry the product melt flow rate, the off-specification tonnage and the margin. (2) The melt flow rate scale marks the product, the polymer being made, the setpoint and the analyzer. (3) The low selector panel is the point of the view. The melt-flow controller is asking for 0.84, the bubble point allows 1.11, 0.84 is selected, and the liquid still holds 1.47 because the constraint took the pen at 41.8 min and hydrogen leaves only through a twenty-minute inventory. (4) The panel reads "constraint is holding the ceiling" rather than "constraint has the pen", because the strict comparison of the two signals is momentarily false at this instant while the constraint controller is still well off its parking clamp of 1.97. Product melt flow rate is 4.33 against 7.65 g/10min being made and an analyzer reading of 3.74 that is 9 min old, and the margin is 2.24 bar at the wall. (5) The reactor has made 13.43 t off specification against 5.49 t on grade, and the pump has not cavitated.

recipe, and spends the bubble point thereby freed on hydrogen, with $r_{H,\mathrm{SP}}$ peaking at 2.38 against the 1.86 that the grade needs at 70 °C. At 70 °C and 38.00 bar the grade leaves 1.09 bar of margin at steady state (Figure 9.2), so no structure can hold it there with 2.0 bar to spare. With no temperature freedom the plan makes 82.9 t off specification and ends at 6.17 g/10min, level with the override, although it still keeps the override idle and spends only 0.1 min below 2.0 bar.

The matched rows separate the temperature from the plan. Given the same ramp to 66 °C by hand, the override structure makes 62.4 t off specification and spends 54.5 min below 2.0 bar. The best such ramp, to 68 °C, makes 56.6 t (80.9 min below), and started at $t = 0$ with hindsight it makes 67.0 t. The plan held to the same path makes 43.8 t without a minute below 2.0 bar. Lowering the setpoint toward 66 °C by hand in steps of about 0.5 °C every 10 min instead of along the smooth ramp makes 68.3 t off specification under the override structure, with a worst margin of 0.82 bar and 40.3 min below 2.0 bar, against 1.21 bar for the smooth ramp, so each half-degree step evidently rings the margin through the pressure loop. On one temperature path, writing the margin as a constraint saves 18.6 t over the override and keeps the margin, while the colder reactor is what makes the grade reachable at all. Removing the back-off lowers the off-specification polymer to 36.4 t but lets the override select for 12.2 min, and pricing the margin at 1 instead of 10^4 per bar^2 returns the plan to the melt-flow controller's behavior, 83.7 t with 191.8 min of selection.

The colder reactor has a cost the controller's model does not see. Propagation is 17 % slower at 66 °C, so at steady state the production controller needs 49.1 % of the catalyst feeder against 42.9 % at 70 °C, about 14 % more, and the four hours make 1.6 t less polymer. Isotacticity, xylene solubles and the molecular weight distribution, which a real recipe protects with its temperature, are absent from the plant model as well, and the recipe for this grade asks for 70 °C at 41.0 bar. The plan has found an operating point the simulated plant permits, not one a recipe owner has approved.

The other cases divide the same way. Through the inert accumulation of scenario 8 the plan makes

Table 9.6 The Override Structure Against Model Predictive Control, Seed 20260914. Polymer in t, Margin in bar, Times in min

Case	Structure	Off spec	Q at end	Worst ΔP	Below 2.0 bar	Override selects, holds	$r_{H,\mathrm{SP}}$ travel	Lowest T_r (°C)	Mean v_{cat} (%)
11, 240 min	override structure	83.2	7.49	1.21	97.8	39.9, 169.7	17.89	69.8	43.55
	slow melt-flow tuning	83.0	6.65	1.21	51.9	11.4, 65.5	6.55	69.8	43.78
	hand ramp to 66 °C	62.4	9.76	1.21	54.5	20.2, 84.0	15.12	65.8	46.34
	hand ramp to 68 °C	56.6	11.23	1.32	80.9	31.5, 134.5	16.35	67.8	44.83
	the same from $t = 0$	67.0	10.27	1.24	67.4	28.0, 101.1	16.50	67.9	45.02
	predictive	43.7	11.89	2.21	0.0	0.0, 0.0	7.91	65.9	46.13
	no temperature freedom	82.9	6.17	2.00	0.1	0.0, 0.0	9.30	69.8	43.99
	the same, T_g ramped to 66 °C	43.8	12.00	2.32	0.0	0.0, 0.0	5.52	65.9	46.37
	no back-off	36.4	11.93	1.55	85.4	12.2, 155.5	9.11	66.5	45.18
	margin price 1	83.7	4.36	1.31	108.7	191.8, 207.7	95.00	69.9	42.54
8, 360 min	override structure	63.7	1.68	1.42	124.9	19.8, 184.0	7.58	69.8	46.54
	predictive	0.0	3.01	2.17	0.0	0.0, 0.0	7.72	67.5	47.36
	no temperature freedom	83.2	1.54	1.56	82.0	32.8, 111.6	16.48	69.8	46.60
9, 300 min	override structure	37.3	2.49	0.70	5.5	2.1, 6.2	4.24	69.8	44.89
	predictive	0.0	3.00	0.70	5.4	4.7, 6.2	3.16	69.8	44.64
5, 300 min	quality cascades	0.0	3.00	2.92	0.0	—	0.07	69.8	44.73
	predictive	0.0	3.00	2.94	0.0	0.0, 0.0	1.06	69.8	44.73
12, 300 min	quality cascades, as loaded	47.2	12.47	−0.33	220.2	—	6.11	66.7	20.27
	predictive	43.2	11.88	2.08	0.0	0.0, 0.0	9.86	65.9	46.30
	override structure, suggested	30.3	11.88	2.14	0.0	0.0, 0.0	3.30	69.8	42.85
	predictive, suggested	30.0	12.00	2.42	0.0	0.0, 0.0	4.94	69.5	42.67

Worst ΔP and the minutes below 2.0 bar use the true margin at the pipe wall. Selects counts the minutes of a true strict selector test and holds the minutes off the parking clamp, as in Section 9.6. Travel sums the absolute changes of the selected hydrogen setpoint. Hand ramps move T_{SP} at 0.05 °C/min from the 15 min setpoint step, and T_g ramped to 66 °C moves the plan's grade temperature along the same path. The slow tuning is the challenge's $K_c = 0.15$ and $\tau_I = 80$ min, and the suggested challenge runs at 41.0 bar. Only the challenge as loaded cavitates (115.5 min), no run trips, and the plan solves 120, 150 and 180 times in 240, 300 and 360 min without a rejected plan.

no off-specification polymer against 63.7 t by taking the reactor to 67.5 °C, and with its temperature held it loses, at 83.2 t. In the jacket case of scenario 9 it holds the grade but the override selects for 4.7 against 2.1 min, likely because the 40 % fouling step arrives faster than a plan made every 2 min can answer. On the lined-out scenarios 5 and 10 nothing is gained, and the plan moves the hydrogen setpoint 1.06 and 1.07 units over five hours against 0.07 and 0.12 under the shipped structures, most likely because it answers analyzer and margin noise that the slow melt-flow controller filters. The challenge as loaded is met by the plan at 11.88 g/10min and a worst margin of 2.08 bar, but with the reactor as cold as 65.9 °C, which the grade recipe does not specify, and with the suggested settings the two structures tie at 30.0 and 30.3 t. On five other noise seeds the five-hour transition gives 43.3 to 43.5 t under the plan against 105.1 to 105.2 t, and halving the integration step leaves scenario 11 at 43.7 t.

9.7 Technology in Practice

What Measures It

The variable the chapter is about has no sensor. As Section 9.9 says, the margin is computed from a pressure transmitter, a thermowell and the loop chromatograph, so an error in any of the three becomes an error in the margin. The temperature term alone is worth 0.629 bar per degree, and the chromatograph is the only instrument that sees the inerts creep of scenarios 4 and 8. The pressure

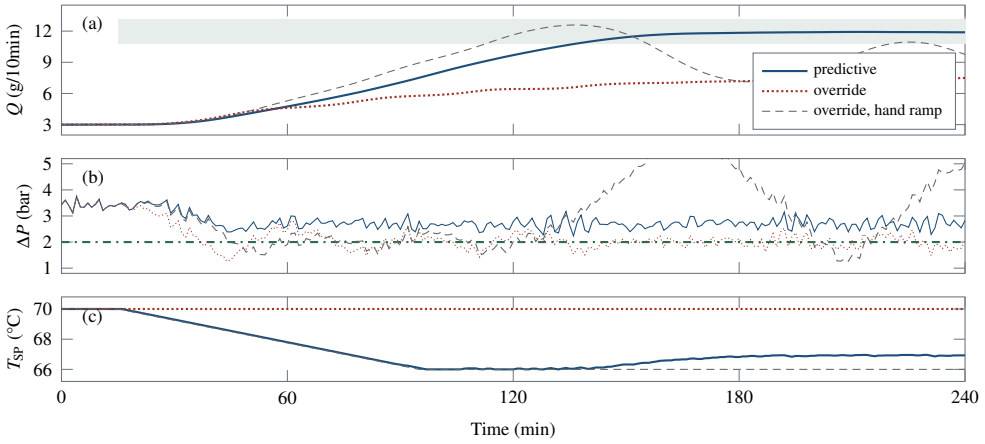

Figure 9.5 Scenario 11 over four hours under the predictive controller (solid), under the override structure (dotted), and under the override structure with the temperature setpoint ramped to 66 °C by hand (dashed). (a) Product melt flow rate and the 10.8 to 13.2 g/10min band (shaded) about the setpoint of 12 entered at 15 min. (b) True margin at the pipe wall and the 2.0 bar override setpoint (dash-dotted). (c) Reactor-temperature setpoint.

term is the subject of the box below. In the model PT-101 reads the pipe wall, where the reactor pressure wanders by a few tenths of a bar with the instrument noise, and the impeller sees 0.36 bar less than the wall.

> **TECHNOLOGY** **Sensor: A Pressure Transmitter Behind a Diaphragm Seal**
>
> Half of the margin is a pressure, and a tap on a slurry invites a plugged impulse line. A remote or chemical seal puts a flexible isolating diaphragm at the process connection and carries its pressure to the transmitter through a sealed, liquid-filled capillary, which removes the impulse line that clogs [23]. The fill must stay free of gas, because a bubble absorbs diaphragm motion and the reading falls short. A transmitter mounted below or above the tap reads high or low by the head of the fill liquid, heat near the capillary expands the fill and adds a false pressure, and the fill's viscosity in the narrow capillary slows the response [23]. On a trend these faults show as an offset or a sluggish line, and every bar of offset is a bar of margin that is not there.

What Moves It

Hydrogen reaches the loop through a valve that the ratio controller moves, so the selected ratio and not the valve is what the quality tier writes. A variable-speed motor is a good final element for a pump, because slowing the pump at low flow instead of throttling a valve saves a great deal of energy over time [23]. Scenario 10 shows the other side of that saving on a slurry loop, where the speed that saves 128 kW also brings the circulation toward the deposition velocity. In the model the pump speed is an operator setting, and its drive is the subject of the next box.

> **TECHNOLOGY** **Final Element: The Axial Circulation Pump and Its Drive**
>
> A single axial pump in each loop moves a large volume of slurry, which carries the heat to the water-cooled walls and keeps the particles from settling out of suspension [16]. A variable-frequency drive sets an induction motor's speed through the frequency of the power it supplies and its torque through voltage and current. The drive also limits speed, current, torque and the rate of acceleration, and shuts down on overcurrent or overvoltage [23]. The failure that matters is cavitation. Where liquid boils at a local pressure below its vapor pressure and the pressure then recovers, the bubbles collapse as microjets that pit any metal in time and crackle like gravel [23]. In the chapter's model, head and circulation velocity then fall together.

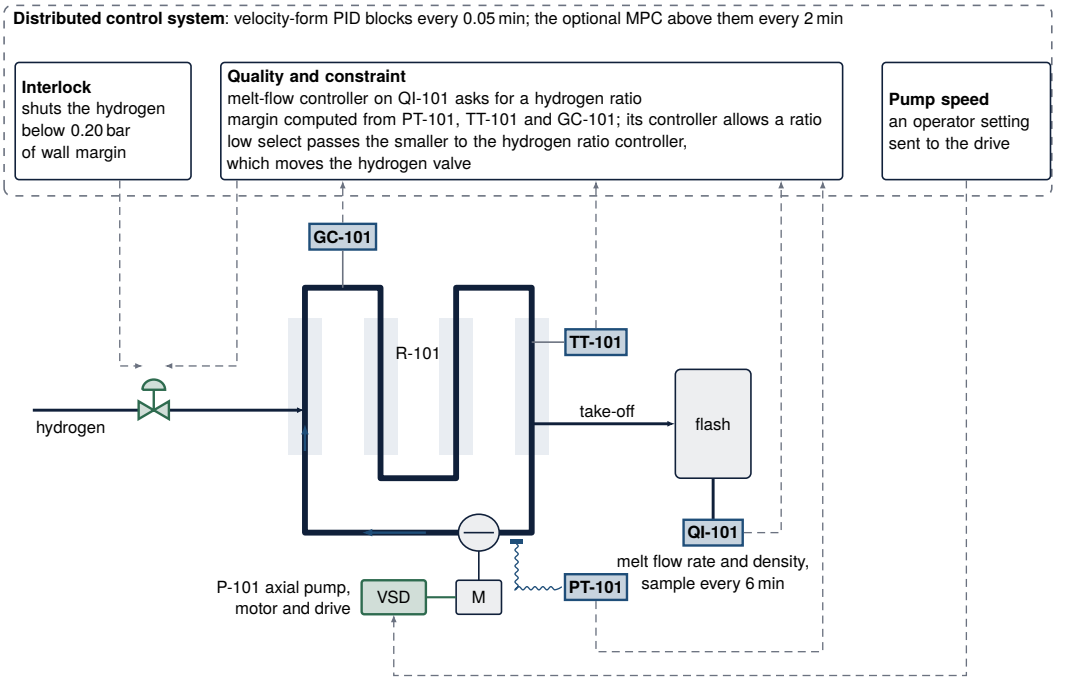

Figure 9.6 The hydrogen loops and the circulation pump with their hardware. Blue tags are sensing devices: the pressure transmitter PT-101 behind a diaphragm seal at the pump suction (its filled capillary drawn wavy), the temperature element TT-101, the loop chromatograph GC-101 and the product analyzer QI-101 under the flash. Green symbols are final elements: the hydrogen valve and the pump's variable-speed drive (VSD) with its motor (M). Dashed arrows are signals. Notice that the margin has no transmitter of its own: it is computed from three instruments, and the low selector then decides whether the melt-flow controller or the margin controller writes the hydrogen ratio setpoint. The interlock acts on the hydrogen valve directly, and the pump speed is an operator setting. The temperature, pressure, slurry-density and production loops are left out. Not to scale, and not a piping diagram.

Where the Controller Runs

Figure 9.6 draws the hydrogen loops of Section 9.5 with hardware in place of the schematic's symbols. A distributed control system holds loops like these at a supervisory level that sends set points to field controllers, with loops nested or cascading and cycle times from milliseconds to minutes [24]. An override is built there from a low-select function block between two controller outputs, and a de-selected controller must not wind up while it has no control. One remedy is the feedback input of the PID block, which stops integration once the block's output is no longer selected [23]. In the app all ten loops execute every 0.05 min, the loser of the low selector is tracked through the same operation manual mode uses, and the predictive controller is a supervisory layer that plans every 2 min.

What Is Modernizing It

Soft sensors built on industrial polypropylene data estimate quality between laboratory results, with the laboratory method kept as the certification reference [25]. A 2003 survey counted polymers among the areas with the most nonlinear predictive control applications and expected their number to grow where linear control had proven difficult, such as polymerization [20]. The EU reference document for polymer plants counts minimizing start-ups and stops as a best available technique. It

credits operating stability, assisted by computer monitoring and control systems, with fewer stops, and avoids an emergency stop by timely identification of deviating conditions followed by a controlled shutdown [16]. The override of this chapter is such an identification, made a controller. Whether a quality estimate can be trusted through a change of grade is the subject of the Open challenge box in Section 9.8.

> **IN PRACTICE** **Keep the Loop Liquid, in Two Layers**
>
> A liquid-propylene loop runs at a pressure high enough, typically 30 to 70 bar, to keep one continuous liquid phase with no separate gas. Gas in the loop can bring poor mixing, flow problems and mechanical damage to the loop pump [26]. A pump is best guarded in two layers: a soft override, where a low-select function block hands the final element to a constraint controller before the limit, and a hard shutdown behind it in case the soft layer is mistuned or fails [23]. The margin controller and the hydrogen interlock of this chapter are those two layers.

9.8 Limitations

The model is not a design tool for a loop reactor and no number in it belongs in a process design. The bubble point of (9.6) is a fit over a narrow band rather than thermodynamics, and the kinetic, quality, Henry, cavitation and fouling coefficients were chosen so that a teaching grade slate behaves like a polypropylene loop rather than to match any catalyst or any licensor's process. One well-mixed volume cannot show the composition and temperature gradients that decide where catalyst and hydrogen are actually injected [6], and one lumped site type cannot produce a molecular weight distribution. Cavitation is a lumped void fraction with an algebraic head-breakdown curve, so the application can say when a pump starts losing metal and never how fast. The post-trip excursions reported in Section 9.6 run into a 94 °C state clamp and past the 90 °C validity limit of the vapor-pressure fit, and are reported here only as the simulator's own behavior. Disturbances are entered as commanded steps and timed events, where a plant drifts. Three elements of the scenarios are stylized on purpose. No plant executes a transition from 3 to 12 g/10min by typing one setpoint, but a single step is the input whose consequences can be measured off a trend. A jacket that loses 40 % of its coefficient in one step stands for a season of operation, and it arrives as a step so that the four links from a dirty jacket to a cavitating pump can be counted. The poison and activity disturbances are separated cleanly, one moving melt flow rate and the other not, which a real feed impurity would not do. Fixing the integration step, the controller period and the random seed makes one run identical on every machine, and a distributed control system is not that repeatable. Five features are present in spite of all that trimming, the Henry-law bubble point, the eye depression, the two-time-scale split between the analyzer and the hydraulics, the nonlinear process gain and the selector, and every one of them is there because it changes the controller that has to be built. Section 9.3.7 names the condition under which every remaining assumption stops holding.

The predictive controller's model is narrower still. It fixes the inert fraction at its design value, which leaves its margin 0.07 bar below the plant's at 66 °C, an error the margin bias alone absorbs as the inerts move, and it omits the catalyst activity, the production, the poison, the pump, the hydrogen interlock and the pressure loop through which a temperature step rings the margin. Nothing in it makes 66 °C a poor temperature for the 12 g/10min grade, so the 4 °C freedom is a declared design choice rather than a property of the plant, and the extra catalyst and the polymer properties a recipe protects are charged to no one. Each output has one slack spanning all 90 steps, so a slack-free plan predicts that the line holds without promising it step by step, and in scenario 11 the margin slack is positive in 93 % of logged samples, at most 0.17 bar, which makes the 2.7 bar line a price rather than a wall. The plan loses where a disturbance outruns its 2 min decision, as in the jacket case, and on lined-out cases it moves the hydrogen setpoint nine to fifteen times as far as the shipped loops do.

OPEN CHALLENGE Quality Estimates That Survive a Change of Grade

Melt flow rate and the xylene-soluble fraction are laboratory results, the second by a gravimetric standard method (ASTM D5492), so neither gives continuous information on a reactor whose grade is moving [25]. Models that estimate them between samples exist. In a 2026 study of 425 industrial polypropylene runs, they lost accuracy when they crossed from one polymer architecture to another, and the study calls that gap between accuracy and transferability a central unresolved issue [27]. A companion study keeps the laboratory method as the certification reference and lets the model fill in between results [25]. Progress would be an estimate that holds through a transition and across homopolymer, random and impact grades.

9.9 Discussion

The organizational split around this constraint is as real as the technical one. The process engineer owns melt flow rate and density, the rotating-equipment engineer owns the circulation pump, and the bubble-point margin belongs to both and is trended by neither. That the margin is not measured at all makes it worse. It is computed from a pressure transmitter, a thermowell and a loop chromatograph through an equation-of-state routine in the control system, and the application does the same with (9.6) at its declared constants, an inert fraction reported 3 min after each 3 min sample and a 1.5 min filter. Getting that calculation trended, alarmed and put on a controller is most of the practical work, and the reason it is worth doing is that an axial pump in a 65 m^3 liquid-full loop is not a spared item.

Methods well beyond the ten implemented loops fit the same documented model. The selector of (9.22) is the regulatory answer to a problem that model predictive control states directly, by carrying the bubble-point margin as an inequality constraint on the same optimization that tracks melt flow rate and density, which removes the chatter measured in Section 9.6 and handles the doubling process gain through the model instead of through a selector [4, 5]. Grade transitions on this plant are naturally posed as optimal control problems in which the operating pressure is a decision variable rather than a recipe entry [5]. Moving horizon estimation on the same nineteen states would estimate the inert level and the catalyst activity bias from the pressure, temperature and analyzer records, so the invisible disturbance of scenario 8 becomes a soft sensor. Tuning the melt-flow controller by a rule that takes the destination gain rather than the origin gain is a small change with a measurable payoff, since the SIMC form for a 52 min lag with 7 min of effective dead time [28] returns different settings at the two ends of one transition.

9.10 Conclusion

This chapter has presented a bulk polypropylene loop reactor as an interactive laboratory for quality control against a thermodynamic constraint. The simulator integrates nineteen states and solves rather than assumes its design point, 70.0 °C and 38.00 bar with 50.0 wt% solids making 22.0 t/h of a 3.00 g/10min homopolymer, at which the bubble point of 34.57 bar divides into 28.86 bar of propylene, 3.48 bar of accumulated inerts and 2.23 bar of hydrogen, and all nineteen derivatives vanish to 2.7×10^{-15} relative. Hydrogen buys melt flow rate at a process gain that doubles from 6.25 to 12.50 g/10min per unit of ratio across one grade transition, and it spends bubble point at 2.44 bar per unit, against a design margin of 3.43 bar at the wall and 3.08 bar at the impeller eye. The plant therefore carries three clocks, 74.5 min for the product melt flow rate, 13.8 min for the bubble point and 1.45 s for the circulation. A melt flow rate step from 3 to 8 g/10min, whose steady state leaves 1.95 bar of margin, cavitates the pump for 42.3 min under quality cascades and not at all with a bubble-point override, and the override raises production from 45.6 to 110.6 t while leaving melt flow rate at 6.55 against a setpoint of 8. The application is deliberately explicit that this is a trade rather than a fix. Parking the idle constraint controller at the top of its range instead of just above the active demand costs 0.74 bar of worst-case margin, and the shipped constraint tuning leaves a 0.68 bar ripple at about a 31 min period that detuning the reset time to 20 min cuts to 0.29 bar. The predictive controller added as a fifth architecture never lets the override act on the grade change and makes 43.7 t off specification against 83.2 t, but from a reactor at 66.0 °C, and on the same temperature path

the margin written as a constraint is worth 18.6 t. Further work would add a second gas-phase reactor for impact copolymer, a plug-flow circuit with real composition gradients, a cavitation damage-rate correlation, and the operating pressure as a decision of the predictive controller. The loop reactor makes a constraint that has to become a controller into something a student can operate.

KEY POINTS

- The hydrogen that sets the grade also brings the loop liquid toward boiling.
- Hydrogen reaches the margin in minutes, the product in more than an hour and the pump in seconds, so the constraint needs its own controller.
- A low selector works when the loser cannot wind up, the constraint controller can win, and the idle controller is parked where it can act.
- An override buys the equipment, not the product: it removed 42.3 min of cavitation and left melt flow rate at 6.55 against 8.
- The challenge is won by raising the pressure before spending margin on hydrogen and by slowing the melt-flow controller.

Where else this applies. A fast limit and a slow quality loop sharing one actuator through a selector is also the anti-surge override of Volume 4, Chapter 2 and the hydrogen selector of Chapter 3.

References

[1] J. T. Gravdahl and O. Egeland, "Centrifugal compressor surge and speed control," *IEEE Transactions on Control Systems Technology*, vol. 7, no. 5, pp. 567–579, 1999.

[2] E. Grist, "Net Positive Suction Head (NPSH) and the pump operational range," in *Cavitation and the Centrifugal Pump*. Abingdon, UK: Routledge, 2023, pp. 51–69.

[3] C. A. Smith, "Ratio, override, and selective control," in *Automated Continuous Process Control*. New York, NY: John Wiley & Sons, 2002, pp. 74–126.

[4] M. Ohshima and M. Tanigaki, "Quality control of polymer production processes," *Journal of Process Control*, vol. 10, no. 2-3, pp. 135–148, 2000.

[5] M. Al-haj Ali, "Optimal grade transition control for liquid-propylene polymerization reactor," *Asian Journal of Control*, vol. 12, no. 3, pp. 413–425, 2010.

[6] J. J. Zacca and W. H. Ray, "Modelling of the liquid phase polymerization of olefins in loop reactors," *Chemical Engineering Science*, vol. 48, no. 22, pp. 3743–3765, 1993.

[7] A. S. Reginato, J. J. Zacca, and A. R. Secchi, "Modeling and simulation of propylene polymerization in nonideal loop reactors," *AIChE Journal*, vol. 49, no. 10, pp. 2642–2654, 2003.

[8] Z.-H. Luo, P.-L. Su, D.-P. Shi, and Z.-W. Zheng, "Steady-state and dynamic modeling of commercial bulk polypropylene process of Hypol technology," *Chemical Engineering Journal*, vol. 149, no. 1-3, pp. 370–382, 2009.

[9] K. Ziegler, E. Holzkamp, H. Breil, and H. Martin, "Das Mülheimer Normaldruck-Polyäthylen-Verfahren," *Angewandte Chemie*, vol. 67, no. 19-20, pp. 541–547, 1955.

[10] G. Natta, P. Pino, P. Corradini, F. Danusso, E. Mantica, G. Mazzanti, and G. Moraglio, "Crystalline high polymers of α-olefins," *Journal of the American Chemical Society*, vol. 77, no. 6, pp. 1708–1710, 1955.

[11] K. Ziegler, "Consequences and development of an invention," Nobel Lecture, Stockholm, 12 December 1963, 1963, nobelprize.org/uploads/2018/06/ziegler-lecture.pdf, read November 2026.

[12] G. Natta, "From the stereospecific polymerization to the asymmetric autocatalytic synthesis of macromolecules," Nobel Lecture, Stockholm, 12 December 1963, 1963, nobelprize.org/uploads/2018/06/natta-lecture.pdf, read November 2026.

[13] American Chemical Society, "Discovery of polypropylene and the development of a new high-density polyethylene," National Historic Chemical Landmark booklet, Office of Communications, Washington, DC, 1999, designated at Phillips Petroleum Company, Bartlesville, Oklahoma, 12 November 1999; acs.org/education/whatischemistry/landmarks/polypropylene.html, read November 2026.

[14] P. Galli and G. Vecellio, "Technology: driving force behind innovation and growth of polyolefins," *Progress in Polymer Science*, vol. 26, no. 8, pp. 1287–1336, 2001.

[15] H. H. Brintzinger, D. Fischer, R. Mülhaupt, B. Rieger, and R. M. Waymouth, "Stereospecific olefin polymerization with chiral metallocene catalysts," *Angewandte Chemie International Edition in English*, vol. 34, no. 11, pp. 1143–1170, 1995.

[16] European IPPC Bureau, "Reference document on best available techniques in the production of polymers," European Commission, Joint Research Centre, Institute for Prospective Technological Studies, Seville, Tech. Rep., Aug. 2007.

[17] J. B. P. Soares and A. E. Hamielec, "Kinetics of propylene polymerization with a non-supported heterogeneous Ziegler-Natta catalyst—effect of hydrogen on rate of polymerization, stereoregularity, and molecular weight distribution," *Polymer*, vol. 37, no. 20, pp. 4607–4614, 1996.

[18] M. C. Roco and C. A. Shook, "Critical deposit velocity in slurry flow," *AIChE Journal*, vol. 31, no. 8, pp. 1401–1404, 1985.

[19] D.-Y. Peng and D. B. Robinson, "A new two-constant equation of state," *Industrial & Engineering Chemistry Fundamentals*, vol. 15, no. 1, pp. 59–64, 1976.

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

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

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

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

[25] J. Hernández-Fernández, "Temporally validated and laboratory-reconciled soft sensing of xylene-soluble fraction in industrial polypropylene manufacturing," *Polymers*, vol. 18, no. 19, p. 2401, 2026.

[26] M. F. Bergstra, P. Denifl, M. Gahleitner, D. Jeremic, V. Kanellopoulos, D. Mileva, P. Shutov, V. Touloupidis, and C. Tranninger, "Polymerization in the Borstar polypropylene hybrid process: Combining technology and catalyst for optimized product performance," *Polymers*, vol. 14, no. 21, p. 4763, 2022.

[27] J. Hernández-Fernández and J. Lopez-Martinez, "Reactor-grade non-additive polypropylene architecture defines the physicochemical limits of melt flow index model transferability: An industrial chemometric and machine learning study," *Polymers*, vol. 18, no. 15, p. 1880, 2026.

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

9.11 Exercises

9.1 ○ Concept check · 2 min

At the design point the bubble point of the loop liquid is 34.57 bar. Which statement about the three dissolved lights is right?

(a) Hydrogen is worth the most after propylene, because its Henry constant is the largest of the three

(b) The accumulated inerts are worth more than the hydrogen, although no controller sets them

(c) Ethylene is worth the most after propylene, because the density loop keeps it in the liquid

(d) The three lights together are worth less than a bar, so the propylene term sets the margin

9.2 ○ Concept check · 2 min

In scenario 6 the melt-flow setpoint steps from 3 to 12 g/10min at 38.00 bar with no override, and the pump cavitates about 36 min later. The steady state of the new grade leaves 1.09 bar of margin at the pipe wall. Why does the pump cavitate?

(a) The steady margin is smaller than the eye depression, so the new grade cannot be held at all

(b) The inerts accumulate during the transition and take the margin that the hydrogen left

(c) The controller overshoots the liquid ratio to 2.29 against the 1.86 the grade needs, before the analyzer registers progress

(d) The pump speed falls as the slurry grows denser, so the eye depression grows during the transition

9.3 Concept check · 2 min

In scenario 7 the strict selector test, "the bubble-point controller wrote this scan's setpoint", is true for only 39.9 min of the 198 min after selection begins. Over the same period the margin averages 1.99 bar against its 2.00 bar setpoint and the melt flow rate is pinned far below its setpoint. Why does the strict test read false most of the time?

(a) The override is idle most of the time, and the ratio controller holds the margin by itself

(b) The analyzer dead time delays the selector, so the test lags the constraint by one sample

(c) The constraint controller sits on its parking clamp at the top of its range for most of the period

(d) The melt-flow controller's integral is tracked to the applied setpoint, so its demand stays close to the allowance and the two signals keep crossing

9.4 Theory and calculation · 7 min

At 70 °C with no ethylene and no poison, Equation (9.19) gives the instantaneous melt flow rate as $Q = K_Q\,(b_0 + r_H)^2$ with $K_Q = 3.2552$ and $b_0 = 0.060$. At the design point $r_H = 0.900$, the margin is 3.43 bar at 38.00 bar, the bubble point rises 2.44 bar per unit of r_H, and the margin rises 1.004 bar per bar of reactor pressure.

(a) Find the ratio r_H that the 12 g/10min grade needs. *Checkpoint:* 1.860.

(b) Find the steady margin of that grade at 38.00 bar, and the reactor pressure that would leave 2.00 bar. *Checkpoint:* 1.09 bar and 38.91 bar.

(c) Find the largest melt flow rate that 38.00 bar supports with a 2.00 bar minimum margin. *Checkpoint:* 7.78 g/10min, within 0.05 g/10min.

9.5 Theory and calculation · 7 min

The melt-flow loop runs from the ratio setpoint $r_{H,\mathrm{SP}}$ to the product melt flow rate Q. Treat it as first order with the inventory time constant $\tau = 52.4$ min, an effective dead time $\theta = 7$ min, and the process gain $k = \mathrm{d}Q/\mathrm{d}r_H$ of the correlation $Q = K_Q\,(b_0 + r_H)^2$ with $b_0 = 0.060$.

(a) Show that $k = 2Q/(b_0 + r_H)$, and evaluate it at 3.00 g/10min ($r_H = 0.900$) and at 12.0 g/10min ($r_H = 1.860$). *Checkpoint:* 6.25 and 12.5 g/10min per unit of ratio.

(b) The SIMC rule with a closed-loop time constant $\tau_c = \theta$ gives $K_c = \tau/[k(\tau_c + \theta)]$ and $\tau_I = \min[\tau, 4(\tau_c + \theta)]$. Find both settings at each gain, and compare them with the default ($K_c = 0.55$, $\tau_I = 45$ min) and with the tuning the challenge hands over ($K_c = 1.8$, $\tau_I = 14$ min). *Checkpoint:* $K_c = 0.60$ and 0.30, $\tau_I = 52.4$ min at both.

9.6 Theory and calculation · 6 min

The loop holds a slurry of density $\rho_{sl} = 591\,\mathrm{kg/m^3}$ over a developed length $L = 196$ m, and linearizing the momentum balance at the design point gives a restoring slope of 79.7 kPa per m/s. The powder has $\rho_p = 900\,\mathrm{kg/m^3}$, the liquid $\rho_L = 440\,\mathrm{kg/m^3}$, the bore is $D = 0.65$ m, $F_L = 1.00$ and $g = 9.81\,\mathrm{m/s^2}$. The circulation runs at 5.50 m/s at 100 % pump speed.

(a) Find the circulation time constant in seconds, and its ratio to the 74.5 min the product melt flow rate takes to cover 63 % of a hydrogen step. *Checkpoint:* 1.45 s, a ratio of about 3100.

(b) Find the deposition velocity of Equation (9.13) and its ratio to the circulation. *Checkpoint:* 3.65 m/s, a ratio of 1.51.

(c) Show from Equation (9.12) that, with no vapor and a fixed solids fraction, the steady velocity is proportional to the pump speed. Find the speed at which the circulation reaches the deposition velocity. *Checkpoint:* 66.4 %.

9.7 App challenge · 7 min

Load scenario 2, Hydrogen step test, and change no setting. The H_2/C_3 ratio setpoint steps from 0.90 to 1.50 at 20 min. Run to 116 min at speed 30 with the default seed and export the CSV file. For a run longer than an hour the time column is headed Time (h) and is in hours. After the step the steady state has a bubble point of 36.03 bar and a product melt flow rate of 7.92 g/10min.

(a) Read column `Pbub_bar` and column `rh_actual` in the last row before the step and in the row at 115 min. Compute the bubble point gained per unit of ratio. *Checkpoint:* 34.57 and 36.02 bar, 0.902 and 1.500, about 2.43 bar per unit, within 0.03 bar.

(b) Find the time after the step at which `Pbub_bar` first covers 63.2 % of its change, and the same for the product melt flow rate in column `mfr_prod`. Give each within 0.5 min. *Checkpoint:* 13.8 min and 74.5 min.

State what the ratio of the two times means for a melt-flow controller that is tuned to make the product respond faster.

9.8 App challenge · 8 min

Scenarios 6 and 7 run the same grade change, a melt-flow setpoint step from 3 to 12 g/10min at 15 min with $K_c = 1.6$ and $\tau_I = 20$ min on the melt-flow controller. Scenario 6, Grade transition RA300 to IM120, has no override, and scenario 7, The same transition, with the bubble-point override, adds the margin controller. Load each in turn, change no setting, run past 4 h at speed 30 with the default seed and export the CSV file. In the row at 4.00 h (the last row before it, at 3.998 h) read `metric_cav_min`, `metric_worst_margin_bar`, `metric_offgrade_t` and `metric_ongrade_t`.

(a) Give the four values for scenario 6. *Checkpoint:* 44.9 min, −20.14 bar, 38.7 t and 5.5 t.

(b) Give the four values for scenario 7. *Checkpoint:* 0.0 min, 1.21 bar, 83.1 t and 5.5 t.

Accept 0.1 min, 0.02 bar and 0.2 t. Explain why the on-grade tonnage is the same in both runs, and which run is better.

9.9 App challenge · 10 min

Load the challenge, Challenge: reach 12 g/10 min without spending the margin, the last case in the list. It steps the melt-flow setpoint from 3 to 12 g/10min at 10 min and is judged at 300 min: the product melt flow rate within ±1.2 g/10min of 12, the margin never below 2.0 bar and no cavitation. For each run use the default seed, run past 5 h at speed 30, export the CSV file and, in the row at 5.00 h (the last row before it, at 4.998 h), read `mfr_prod`, `metric_worst_margin_bar`, `metric_cav_min` and `metric_offgrade_t`. Accept 0.02 g/10min, 0.02 bar, 0.1 min and 0.2 t.

(a) Run the case as loaded. *Checkpoint:* 12.47 g/10min, −0.33 bar, 115.5 min, 47.2 t.

(b) Press the reload button and, before Run, set Reactor pressure setpoint to 41.00 bar, Control architecture to + bubble-point override, and on the melt-flow controller card Proportional gain to 0.15 and Integral time to 80 min. *Checkpoint:* 11.88 g/10min, 2.14 bar, 0.0 min, 30.25 t.

(c) Press the reload button and make the same changes except the architecture, which stays at Quality cascades. *Checkpoint:* The same four values as in (b).

Say which runs pass, and what result (c) says about the override.

9.10 Application area · 4 min

Every alarm and interlock on the loop reads the margin at the pipe wall, ΔP, while vapor forms in the impeller eye, one eye depression h_e lower. Use Equation (9.8) with $\rho_{sl} = 591\,\mathrm{kg/m^3}$, $\sigma_v = \sigma_N = 2.0$ and $v_r = 5.50\,\mathrm{m/s}$.

(a) Find h_e at the design circulation of 5.50 m/s and 100 % speed, and at 4.40 m/s and 80 % speed. *Checkpoint:* 0.36 bar and 0.23 bar.

(b) The hydrogen interlock closes the valve at 0.20 bar of wall margin. Find the eye margin at that instant at design flow, and say what the interlock protects. *Checkpoint:* −0.16 bar.

9.11 Application area · 4 min

An energy audit proposes running the circulation pump at 70 % speed instead of 100 %. At full speed the pump draws 263 kW at 5.50 m/s, shaft power scales with the cube of speed, the velocity is proportional to speed, the deposition velocity is 3.65 m/s and the jacket coefficient scales with $v^{0.80}$.

(a) Find the shaft power saved, and compare it with the saving at 80 % speed. *Checkpoint:* 173 kW against 128 kW.

(b) Find the velocity, its ratio to the deposition velocity and the fraction of the jacket coefficient that remains, and give a recommendation. *Checkpoint:* 3.85 m/s, 1.05, and 0.75.

9.12 Review the agent · 5 min

An AI assistant summarizes scenario 11 as follows. “On the same transition the predictive controller makes 43.7 t of off-grade polymer against 83.2 t under the override, and the override never has to act. Writing the bubble-point margin as a constraint in the plan is therefore worth about 40 t per transition. The plant should replace the override structure with the predictive controller.” Find the flaw in the reasoning and say how it could be checked on the app.

CHAPTER 10

Peak Temperature Control Against Ethylene Decomposition in a Tubular LDPE Simulator

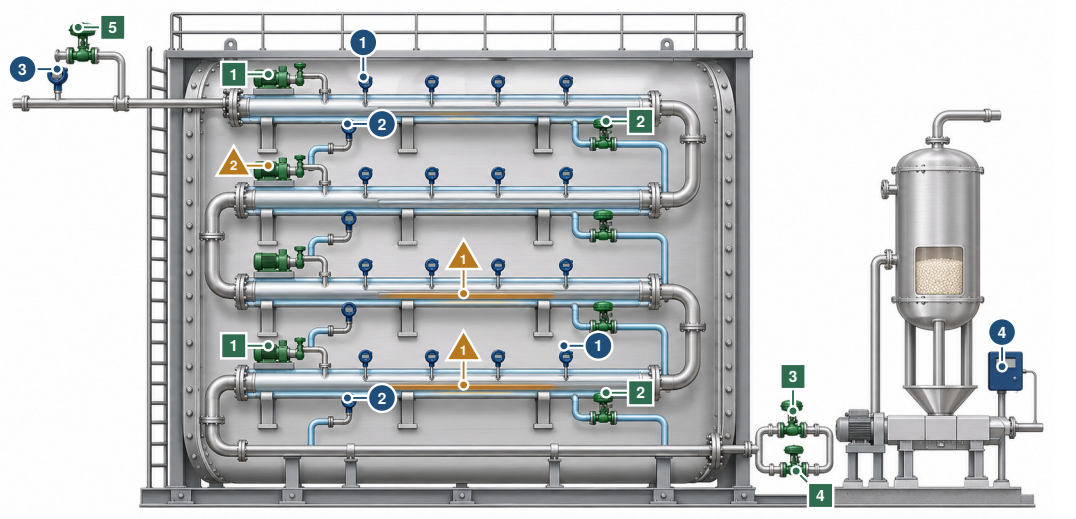

The controller moves 1 initiator pump at the start of each zone 2 cooler valve on each jacket's water 3 let-down valve that holds the pressure 4 pulse valve, opened on a timetable 5 chain-transfer agent feed

The controller reads 1 thermocouples along each zone, the hottest one read 2 jacket water inlet thermometers 3 reactor pressure 4 melt index of the pellets

What disturbs it 1 polymer layer on the wall 2 strength of the initiator lot

Notice the four injection points along one tube: each lights its own peak, the polymer layer on the later runs slows their cooling, and the pulse valve drops the pressure every three minutes to strip that layer. Illustration, not to scale.

IN PLAIN TERMS Four Hot Spots Kept Below a Danger Line

The main idea. Soft plastic for bags and film is made from a gas squeezed hard into a long, thin pipe. A chemical is pumped in at four places, and each one starts a short hot zone where some of the gas turns into plastic.

The challenge. Hotter zones make more plastic. Too hot, and the gas itself breaks apart. The thermometers see the heat a little late, and a planned drop in pressure every few minutes shakes every zone at once.

Why it matters. The gap kept below the danger line has to cover the worst swing the pipe ever makes, not only the steady temperature the controls aim for.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the two History boxes in Section 10.2, the Did you know? boxes in Sections 10.3 and 10.6, the two Try it boxes in Section 10.6, and Key points at the end of the Conclusion. No equations.
Study	60 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 10.6 with Figure 10.6, Section 10.7, and Section 10.8 with its Open challenge box.

The control question Four peroxide injections hold four temperature peaks a few tens of kelvin below the temperature at which ethylene decomposes, and the conversion rises with the peaks. How large must the margin to that line be when a scheduled pulse valve, late thermocouples and a changing initiator lot all make excursions that the setpoints do not show?

Learning objectives

- Compute the adiabatic rise per percent of conversion and explain why the tube converts in four separate peaks.
- Explain why a peak's gain falls as its pump strokes up, and why each zone is a disturbance to the zones after it.
- Use a feedforward keyed to the pulse schedule to stop the peak loops winding up, and measure the margin it buys.
- Size a peak setpoint from the largest excursion, the measurement error and a reserve, and explain why a trip below the line can still be too late.
- Tell a setpoint from a constraint by comparing the peak loops with a model predictive controller that carries a ceiling.

Prerequisites Sections 1.4.3 (runaway), 1.4.4 (heat removal), 1.4.8 (back-off from a limit), 1.6.1 and 1.6.2 (engine and velocity-form PID). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/ldpe-tubular-reactor/

See also Volume 4, Chapter 10 and Volume 4, Chapter 9 for mid-ranging and valve-position control, and Volume 4, Chapter 2 for a protective override that guards a limit. Volume 3, Chapter 6 for a feedforward that acts before feedback sees the disturbance, and Volume 3, Chapter 2 for matching loop structure to where a disturbance enters.

10.1 Introduction

Low-density polyethylene (LDPE) is made from ethylene by free-radical polymerization at very high pressure, either in stirred autoclaves or in long jacketed tubes [1, 2]. The product matters in volume. Polyethylene is the largest group of nonfiber plastics produced, 36 % of the total, and about 42 % of all nonfiber plastics have gone into packaging [3]. The tube of this chapter is 1500 m of 55 mm bore, fed 64.8 t/h of ethylene at 2500 bar, which carries the gas through in 109 s. Peroxide is injected at four points, and each injection lights a reaction front that heats the gas by 12.8 K for every percent of ethylene it converts. Four peaks at 290 °C convert 27.51 % and make 17.8 t/h of polymer. A single front could not, because converting that much in one rise would heat the gas by 353 K, from 165 °C to about 518 °C.

Above the peaks lies the temperature at which ethylene itself decomposes, to carbon, hydrogen and methane [4]. The model draws that line at 335 °C, 45 K above the design peaks, as a teaching value. A tempting reading of that gap holds that a peak kept on its setpoint by a well-tuned loop is as safe as the distance from its setpoint to the line, and the simulator is built to show why that reading fails. With all four loops at 290 °C, the pulse valve that the plant opens every three minutes takes the hottest thermocouple to 299.7 °C once the pressure has recovered, because the loops wound their pumps up while the pulse held the peaks down. During that recovery the hottest gas runs 10.9 K ahead of every thermocouple. An initiator lot twice as strong as the one the loops have adapted to sets off a 320 °C trip 25 s after it arrives, and a high-high trip set 5 K below the line fires 1.5 s before a decomposition it does not prevent, since it stops the pumps and not the initiator already in the gas, while a trip one 1 s scan earlier leaves 1.4 K. Excursions set the margin, and the setpoint is what they leave.

This shape recurs well beyond polyethylene. A tubular reactor with a steeply temperature-dependent reaction can be parametrically sensitive, so that a small change in the wall heat-transfer coefficient or a small dilution of the feed produces a large change in what leaves it [5], as Section 1.4.3 sets out and as the ethylene oxide tube of Chapter 6 shows on a catalyst bed. In the LDPE tube the

wall changes on its own, because polymer deposited on it lowers the heat-transfer coefficient and, in rare cases, the lost heat removal leaves hot spots that trigger decomposition [6]. Excess initiator, a feed impurity, a feed temperature disturbance, a failed controller and a poorly tuned one each drive an LDPE autoclave toward runaway [7]. A slow actuator that keeps a fast one inside its range is the mid-ranging structure, used wherever a small fast valve works beside a large slow one [8], as on the power block of Volume 4, Chapter 10, and a disturbance that arrives on a timetable is a case for feedforward rather than for feedback [9]. A sensitive hot spot, a fine valve kept in range by a coarse one, a disturbance on a schedule and an interlock that must act inside the process's own dead time are each familiar from other plants, and in the LDPE tube all four meet.

The published models of the tubular process were written to predict, fit and optimize a plant [1, 10–12], and none of them hands a student four pumps, a pulse valve and a trip and counts the cost of each choice in kelvin of margin. This chapter makes the margin to decomposition a measured quantity, on a tubular LDPE reactor built to run in a browser, whose gas is carried as parcels that ride the flow, whose thermocouples see that gas late and in pieces, and whose decomposition is an event that ends the run rather than a soft penalty. The kinetics are illustrative, and PI peak control, pulse feedforward, mid-ranging and a high-high trip are established practice. The optional model predictive controller follows the industrial form, a receding-horizon quadratic program on a model identified from plant step tests with a soft constraint on each output [13], as Section 1.6.3 describes.

The chapter runs as follows. Section 10.2 traces the process from the high-pressure laboratory to the pulsed tube. Section 10.3 states the model, a 1021-state model whose energy balance closes parcel by parcel and whose mass balance is its pressure equation, on a Lagrangian grid that keeps plug-flow kinetics exact where fixed well-mixed cells cannot, with every constant given. Section 10.4 measures how a peak forms and what moves it, among them a peak gain that falls from 1.9 to 0.6 K per percent of pump stroke while the peak moves toward its injector. Section 10.5 describes the thirteen loops, the pulse feedforward, the trip and the supervisor, and a model predictive controller on the four pumps whose model was fitted to step tests of the simulator and is given with its residuals. Section 10.6 reports nine cases and a four-hour challenge, which separate a peak from its measurement, feedback from a scheduled disturbance, a slow handle from the pump it serves, a setpoint from a margin, and a request from a constraint, and which measure the margin a scheduled pulse costs with and without feedforward. Section 10.7 names the hardware and the practice behind the model, Section 10.8 states the limitations and Section 10.9 sums up.

10.2 From a High-Pressure Laboratory to the Pulsed Tube

The high-pressure route to polyethylene came out of a laboratory that was studying how pressure changes organic reactions [14], a discovery told as a case history [15] and in the first History box below. The process that grew from it still polymerizes ethylene by free radicals at pressures of that order, in stirred autoclaves or in long jacketed tubes [1, 7], with the reaction started by oxygen, by organic peroxides or by both [12].

HISTORY Polythene Out of a High-Pressure Laboratory, 1930s

In 1935 three chemists reported how pressure changes reactions in solution, taking one reaction to 3000 kg/cm^2 and another to 8500 kg/cm^2 [14]. Polyethylene was invented in the same decade by a British chemical company, and its process needed extreme pressures of 20 to 30 thousand pounds per square inch [16]. When Karl Ziegler's group made polyethylene at normal pressure at the end of 1953, the British product had been known for some seventeen years and was made at 1000 to 2000 atmospheres and about 200 °C. In that process the chain grows with interruptions, so the molecule is strongly branched [17]. The branches make the polymer soft, tough and flexible, with a density between 915 and 935 kg/m^3. Its value was first established as insulation for electric cables, and film is its most popular use today [18].

Models of the tube arrived with the computer. A model published in 1976 assumed plug flow without axial mixing and computed the axial temperature profile, the conversion and the molecular weights and their distribution along the reaction coordinate [10], the plug-flow assumption this simulator keeps. The kinetics of olefin polymerization in high-pressure tubes were reviewed and updated in 1992 [2], and the modeling, optimization and quality control of these reactors were reviewed together in 1993 [1]. Models of industrial tubes then had to account for the wall. An improved model for industrial tubular reactors appeared in 1996 [11], and the same authors noted that the zones of such reactors exhibit different flow patterns and critical fouling behavior, so that neither a constant heat-transfer coefficient nor the usual Nusselt correlations predicts the temperature profiles accurately enough [19]. A reduced model built for the periodic parameter updates of a plant had to include the pressure pulse equations, and it treated each peroxide mixture as one fictitious initiator [12], which is the simplification the simulator makes zone by zone. Optimization followed the models, and a nonlinear model predictive control strategy for the operation of these tubular reactors was reported in 2009 [20].

The hazard was studied alongside, in a kinetic model of an autoclave that could run away [7] and in experiments on the decomposition boundary of ethylene and its comonomers [4], as the second History box tells. Models of a whole plant joined the reactor to the rest of the process [21], and the mechanisms by which polymer fouls the wall have been set out for a reactor that must remove about half of its heat of polymerization to the jacket water [6]. These studies describe the arrangement that the simulator implements, a tube divided into zones [19] with initiator injected along it and a pressure that is pulsed [12], though none of them can switch a piece off and read what it costs.

HISTORY Mapping Where Ethylene Gives Way, 1996 and 1999

The hazard behind this chapter's line was mapped in the 1990s. A kinetic model of a stirred autoclave published in 1996 added free-radical reactions for the decomposition of ethylene to the usual initiation, propagation and termination steps, so that a simulated reactor could run away, and its stability analysis gave safe operating limits for typical conditions [7]. In 1999 an experimental study filled ethylene mixed with comonomers into a preheated autoclave at up to 450 °C and 150 MPa. Depending on the starting pressure and temperature, each mixture either polymerized stably or polymerized with decomposition. Vinyl acetate and methyl methacrylate lowered the boundary between the two outcomes, while acrylic acid, 1-butene and 1-hexene raised it [4].

10.3 The Tube as Modeled

One tube of length $L = 1500\,\text{m}$ and bore $D = 55\,\text{mm}$ is simulated, fed $F_{in} = 18.0\,\text{kg/s}$ of ethylene by a hyper-compressor through a preheater at $T_{in} = 165\,^\circ\text{C}$, and discharged through a let-down valve and a pulse valve to a high-pressure separator at $P_{sep} = 250\,\text{bar}$. Its four zones begin at 0, 380, 750 and 1120 m, their injectors sit 15 m into each zone, and every symbol in the equations below appears in Figure 10.1.

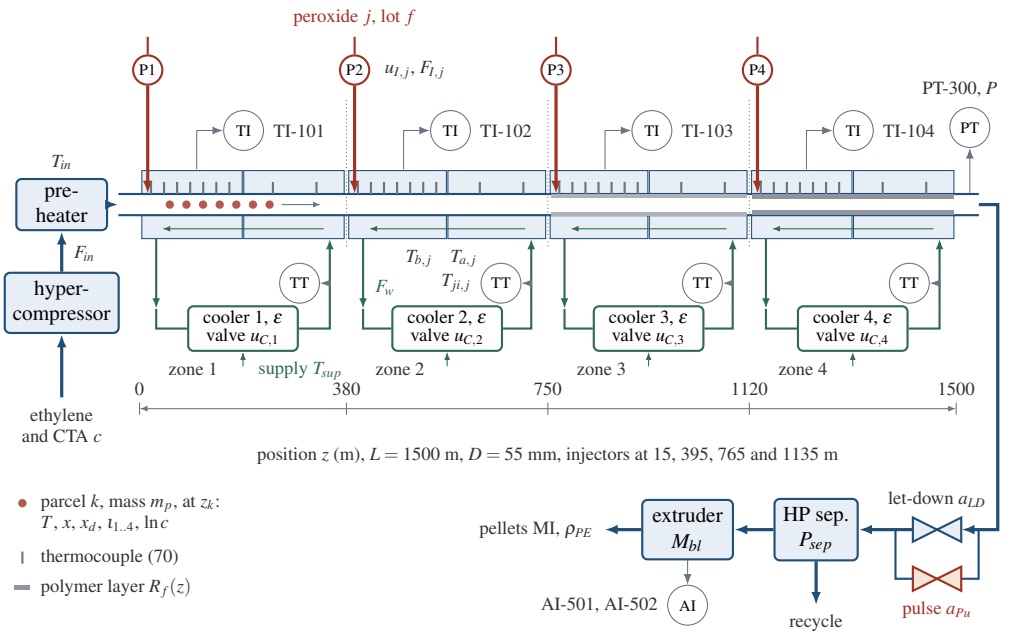

Figure 10.1 High-pressure tubular LDPE reactor as modeled. Ethylene with its chain-transfer agent (CTA) enters at F_{in} and T_{in} and crosses four zones of tube as parcels k of mass m_p at positions z_k. Pumps P1 to P4 inject peroxide $F_{I,j}$ at stroke $u_{I,j}$ from lots of efficiency f. Each zone's jacket water F_w enters its downstream section at $T_{ji,j}$, crosses sections a and b counter-current to the gas and returns through a cooler of effectiveness ε whose valve $u_{C,j}$ sets the share cooled by supply water at T_{sup}. Seventy thermocouples feed TI-101 to TI-104, the hottest of each zone, PT-300 reads the pressure P, TT-201 to TT-204 the jacket inlets, and AI-501 and AI-502 the pellets that leave the separators and extruder hold-up M_{bl}. The let-down valve a_{LD} and the pulse valve a_{Pu} discharge to the separator at P_{sep}. Labels drawn on zone 2 apply to every zone, and the polymer layer $R_f(z)$ is drawn thickest where the model grows it, in zones 3 and 4. The outlet line runs down the right edge to the valves, the separator and the extruder, drawn in a second row under the tube.

> **+ DID YOU KNOW? A Kilometer or Two of Thick-Walled Pipe**
>
> Commercial tubular reactors are typically 1000 to 2500 m long, built of high-pressure tubes 10 to 15 m long laid in a serpentine inside a concrete bay. The outside diameter of a tube is 2.1 to 2.5 times its inside diameter. Such a tube converts up to 36 % of its ethylene, an autoclave about 20 % [18].

Table 10.1 sorts the variables by role. The state vector holds 1021 entries, 12 states in each of 78 parcel slots, 20 lumped states, 5 ledgers and 60 wall segments. Inside the physics the units are SI, with temperatures in °C and pressures in bar, while the engine's clock runs in minutes.

10.3.1 Parcels That Ride the Gas

The tube is cut into $N = 70$ parcels of equal mass. Parcel k owns the slice of the cumulative feed between km_p and $(k+1)m_p$, so its center lies where Equation (10.1) puts it.

$$z_k(t) = \frac{M_{in}(t) - (k+\tfrac{1}{2})\,m_p}{\rho(P)\,A}, \qquad \frac{\mathrm{d}M_{in}}{\mathrm{d}t} = F_{in}$$
$$\rho(P) = \rho_0\left[1 + \kappa\,(P - P_{ref})\right], \qquad m_p = \frac{\rho_0 A L}{N} \tag{10.1}$$

Table 10.1 Variables by Role, With Their Values at the Design Point of Cases 2 and 3. Strokes, Valves and Ranges of MVs in Percent

Tag or symbol	Role	Design point	Range
TI-101 to 104	CV, hottest TC	290 °C SP	250–335
TT-201 to 204	CV, jacket inlet	175 °C SP	140–240
PT-300	CV, pressure	2500 bar SP	2000–3000
AI-501	CV, melt index	2.0 SP	0.5–8
Δ_{dec}	constraint, truth	44.7 K	≥ 0
$u_{I,1..4}$	MVs, pumps	36.3, 34.5, 36.5, 35.8	0–100
$u_{C,1..4}$	MVs, coolers	32.7, 35.5, 35.8, 36.3	0–100
a_{LD}	MV, let-down	50.0	0–100
c	MV, CTA (mol ppm)	1500	100–5000
a_{Pu}	scheduled, pulse	0, 0.25 open	—
f	DV, lot, unmeasured	100 % of 0.6	25–150
T_{sup}	DV, supply water	140 °C	100–200
T_{in}	DV, feed	165 °C	140–190
—	DVs, fouling, events	100 %, missed pulse, kick	0–500
$T, x, x_d, \iota_{1..4}$	states, parcel	see text	—
$\ln c, Q_r, Q_j, \Lambda, \Phi$	states, parcel	see text	—
P, M_{in}	states	2500 bar	—
$a_{LD}, a_{Pu}, s_{1..4}$	states, actuators	as above	—
$T_{a,j}, T_{b,j}, T_{ji,j}$	states, jackets	181, 188, 175 °C	—
$R_{f,1..60}$	states, wall	0 at reset	—

where $A = \pi D^2/4 = 2.376 \times 10^{-3}\,\mathrm{m}^2$, $\rho_0 = 550\,\mathrm{kg/m}^3$ at $P_{ref} = 2500\,\mathrm{bar}$ and $\kappa = 6.3 \times 10^{-5}\,\mathrm{bar}^{-1}$, so $m_p = 28.0\,\mathrm{kg}$ and each parcel is 21.4 m long. At the design flow the gas moves at 13.78 m/s and stays 108.9 s. When the pressure falls, ρ falls with it and every z_k grows in proportion, so parcels near the outlet move out fastest, which is the velocity profile of a compressible plug. Parcels are created at the inlet at T_{in} carrying the feed's CTA and leave once their rear has passed the outlet, and 78 slots hold them.

A fixed grid of well-mixed cells was tried and rejected. A cell of 21 m holds its radicals for its whole 1.5 s residence, while the initiator near a peak lives 25 to 190 ms, since its decomposition constant between 290 and 320 °C lies between 5.3 and 40 s^{-1}. At the initiator flows that give 290 °C peaks in plug flow, grids of 60, 70, 80 and 200 upwind cells have no steady state in their first cell short of decomposition. Parcels keep the plug-flow kinetics exact, because each one is a batch reactor that never mixes with its neighbors.

10.3.2 Initiators, Radicals and Heat

Each parcel carries four initiator concentrations, its conversion, the fraction of its ethylene decomposed, its temperature and the logarithm of its CTA, with the radicals in quasi-steady balance with termination, as shown in Equation (10.2).

$$\begin{aligned}
\frac{\mathrm{d}\iota_j}{\mathrm{d}t} &= -k_{d,j}\,\iota_j, \qquad j = 1,\ldots,4 \\
w = \frac{\mathrm{d}x}{\mathrm{d}t} &= G(1-x-x_d)\sqrt{2f\rho\textstyle\sum_j k_{d,j}\,\iota_j} \\
\frac{\mathrm{d}x_d}{\mathrm{d}t} &= k_{dec}\,(1-x-x_d) \\
c_p\frac{\mathrm{d}T}{\mathrm{d}t} &= \frac{-\Delta H_p}{M_w}\,w + \frac{-\Delta H_{dec}}{M_w}\,\frac{\mathrm{d}x_d}{\mathrm{d}t} - \frac{4U}{\rho D}\,(T-T_j) \\
\frac{\mathrm{d}\ln c}{\mathrm{d}t} &= -\frac{C_s\,w}{1-x-x_d}
\end{aligned} \tag{10.2}$$

where ι_j is initiator j in mol/kg, f its efficiency, x the conversion, x_d the fraction decomposed, c the CTA in mol/kg, $M_w = 28.05\,\mathrm{g/mol}$, $c_p = 2600\,\mathrm{J/(kg\,K)}$, $-\Delta H_p = 93.6\,\mathrm{kJ/mol}$, $-\Delta H_{dec} =$

120 kJ/mol, $C_s = 0.3$ and T_j is the jacket water the parcel passes. The rate constants follow Equation (10.3), with R the gas constant, T in kelvin and $P - P_{ref}$ in pascals inside the exponentials.

$$\begin{aligned} k_{d,j} &= \frac{\ln 2}{1\,\mathrm{s}} \exp\!\left[\frac{E_d}{R}\left(\frac{1}{T_{h,j}} - \frac{1}{T}\right) - \frac{\Delta V_d\,(P - P_{ref})}{RT}\right] \\ G &= G_{ref} \exp\!\left[\frac{E_G}{R}\left(\frac{1}{T_G} - \frac{1}{T}\right) - \frac{\Delta V_G\,(P - P_{ref})}{RT}\right] \\ k_{dec} &= k_{dec,0} \exp\!\left[\frac{E_{dec}}{R}\left(\frac{1}{T_D} - \frac{1}{T}\right)\right] \end{aligned} \tag{10.3}$$

Here $G = k_p/\sqrt{k_t}$ is one lumped constant, $G_{ref} = 0.30\,(\mathrm{m^3/(mol\,s)})^{1/2}$ at $T_G = 280\,^\circ\mathrm{C}$ with $E_G = 29\,\mathrm{kJ/mol}$ and $\Delta V_G = -29.5\,\mathrm{cm^3/mol}$, the four peroxides share $E_d = 90\,\mathrm{kJ/mol}$ and $\Delta V_d = 6\,\mathrm{cm^3/mol}$ and have one-second half-lives at $T_{h,j} = 212$, 228, 232 and 236 °C, and $k_{dec,0} = 5 \times 10^{-3}\,\mathrm{s^{-1}}$ at $T_D = 335\,^\circ\mathrm{C}$ with $E_{dec} = 450\,\mathrm{kJ/mol}$. Decomposition is left out below 260 °C, where it runs at less than 4×10^{-6} of its rate at 335 °C. Three consequences follow from these constants. The adiabatic rise $-\Delta H_p/(M_w c_p)$ is 12.83 K per percent of conversion. The square root rewards a slow supply of radicals, so a dose decomposed quickly makes less polymer per mole than the same dose decomposed slowly. And the negative activation volume makes a pulse that drops 300 bar slow the polymerization by 17 % at 290 °C.

The wall's coefficient is $U = [1/U_c + R_f(z)]^{-1}$, where the clean value $U_c = U_0\,[\varphi_{in} + (\varphi_{out} - \varphi_{in})\,z/L]$ follows the velocity with $\varphi = (F/F_{in,0})^{0.8}$ evaluated at the inlet and outlet flows, $F_{in,0} = 18.0\,\mathrm{kg/s}$ being the design feed, and $U_0 = 800\,\mathrm{W/(m^2\,K)}$. For a clean wall $4U/(\rho D) = 105.8\,\mathrm{W/(kg\,K)}$, which cools a parcel with a time constant of 24.6 s. Each parcel also integrates the heat it has made Q_r, the heat it has given the jacket Q_j, and the polymer-weighted integrals Λ and Φ of the quality correlations below, which makes 12 states.

10.3.3 Pressure, Valves and Pumps

The tube's mass balance is its pressure equation, as shown in Equation (10.4).

$$C\,\frac{\mathrm{d}P}{\mathrm{d}t} = F_{in} - C_v\,(a_{LD} + a_{Pu})\sqrt{P - P_{sep}} \tag{10.4}$$

where $C = \rho_0 A L \kappa = 0.1235\,\mathrm{kg/bar}$, and the valve coefficient $C_v = 0.759\,\mathrm{kg/(s\,bar^{1/2})}$ is set so that the let-down passes the design flow at P_{ref} with an opening of $a_{nom} = 0.5$. With the let-down held, the pressure answers a change in flow with a time constant of 30.9 s. The let-down follows its command through a 0.5 s lag, the pulse valve through 0.3 s to an extra opening of 0.25 while it is open, and the pump strokes s_j through 1.0 s. A pump delivers $F_{I,j} = s_j F_{I,max,j}$, with 13.0 mmol/s at full stroke for pump 1 and 3.8 mmol/s for pumps 2 to 4, and the panels convert moles at 216 g/mol. At the design point the pumps deliver 4.71, 1.31, 1.39 and 1.36 mmol/s.

10.3.4 Jackets and the Polymer Layer

Each zone's water runs counter-current through two well-mixed sections and returns through a cooler whose valve sets the share of it that is cooled, as shown in Equation (10.5).

$$\begin{aligned} m_s c_w\,\frac{\mathrm{d}T_{a,j}}{\mathrm{d}t} &= F_w c_w\,(T_{ji,j} - T_{a,j}) + Q_{a,j} \\ m_s c_w\,\frac{\mathrm{d}T_{b,j}}{\mathrm{d}t} &= F_w c_w\,(T_{a,j} - T_{b,j}) + Q_{b,j} \\ \tau_J\,\frac{\mathrm{d}T_{ji,j}}{\mathrm{d}t} &= T_{b,j} - u_{C,j}\,\varepsilon\,(T_{b,j} - T_{sup}) - T_{ji,j} \end{aligned} \tag{10.5}$$

where section a is the downstream half that the water enters first, $Q_{a,j}$ and $Q_{b,j}$ are the rates at which the parcels gave up heat in each section over the step, $m_s = 2000\,\mathrm{kg}$, $c_w = 4186\,\mathrm{J/(kg\,K)}$,

$F_w = 60\,\mathrm{kg/s}$, $\varepsilon = 0.8$ and $\tau_J = 20\,\mathrm{s}$. A parcel sees as T_j the time-weighted mean of the sections its center crosses during the step, and a parcel already past the outlet sees no jacket. At the design point the inlets hold 175 °C, the downstream sections run near 181 °C and the water returns at 187.4 to 189.3 °C.

A polymer layer grows on each of 60 wall segments of 25 m, where the wall is colder than the cloud point of the polymer in the gas, as shown in Equation (10.6).

$$\begin{aligned} \frac{\mathrm{d}R_{f,i}}{\mathrm{d}t} &= g_0 \,\frac{T_{cl} - T_w}{10\,\mathrm{K}}\, \min\!\left(2, \frac{x}{0.2}\right) \\ T_{cl} &= T_{c0} + T_{c1}\,x, \qquad T_w = \tfrac{1}{2}\,(T + T_j) \end{aligned} \tag{10.6}$$

applied while $T_w < T_{cl}$ and $x > 0.005$, once a second, with $g_0 = 4 \times 10^{-7}\,\mathrm{m^2\,K/(W\,s)}$ scaled by the fouling-rate disturbance, $T_{c0} = 150\,^{\circ}\mathrm{C}$ and $T_{c1} = 300\,\mathrm{K}$. Each completed pulse that lowered the pressure by ΔP_p multiplies every $R_{f,i}$ by $1 - \min(0.9,\, \phi_p \Delta P_p / 300\,\mathrm{bar})$ with $\phi_p = 0.3$. The law is illustrative.

10.3.5 Decomposition as an Event

A parcel held above $T_{dec} = 335\,^{\circ}\mathrm{C}$ for $t_{dwell} = 1.0\,\mathrm{s}$ decomposes. The event reports a local temperature of $T + (-\Delta H_{dec}/M_w c_p)(1 - x - x_d)$ and a pressure 700 bar above the operating one, both illustrative, bursts the rupture discs, shuts the plant down and stops the run. The threshold and the dwell are teaching values. The kinetics of Equation (10.3) give the choice a physical reading. At a conversion of 27 % the decomposition heats a spent parcel by 0.005 K/s at 290 °C, 0.63 K/s at 320 °C, 6.0 K/s at 335 °C and 12.4 K/s at 340 °C, and at 335 °C the clean jacket removes about 6 K/s from such a parcel, so above the line the decomposition wins. The limits of a real reactor come from its licensor's process-safety data, from the plant's own decomposition history and from measured decomposition boundaries such as those of ethylene and its comonomers [4], not from this model. The margin $\Delta_{dec} = T_{dec} - T_{hot}$ is judged on the hottest parcel in the tube, the truth, and the scorecard keeps its worst value over the run.

10.3.6 Product Quality and the Instruments

Melt index and density are illustrative correlations of the CTA, of the temperature at which the polymer is made and of the pressure, as shown in Equation (10.7).

$$\begin{aligned} \ln \mathrm{MI} &= \ln \mathrm{MI}_{ref} + a_C \ln \frac{c}{c_{ref}} + a_T\,(T - T_Q) \\ \rho_{PE} &= \rho_{ref} - b_T\,(T - T_Q) + b_P\,(P - P_{ref}) \end{aligned} \tag{10.7}$$

Each parcel accumulates $\mathrm{d}\Lambda/\mathrm{d}t = w \ln \mathrm{MI}$ and $\mathrm{d}\Phi/\mathrm{d}t = w\rho_{PE}$, so the polymer it carries has the melt index $\exp(\Lambda/x)$ and the density Φ/x of all it has made. The constants are $\mathrm{MI}_{ref} = 2.0\,\mathrm{g/10\,min}$ at $c_{ref} = 1500\,\mathrm{mol\,ppm}$ and $T_Q = 261.5\,^{\circ}\mathrm{C}$, $a_C = 1.0$, $a_T = 0.025\,\mathrm{K^{-1}}$, $\rho_{ref} = 0.9235\,\mathrm{g/cm^3}$, $b_T = 3.5 \times 10^{-4}\,\mathrm{g/(cm^3\,K)}$ and $b_P = 4 \times 10^{-6}\,\mathrm{g/(cm^3\,bar)}$. Parcels leaving the tube mix into one hold-up of $M_{bl} = 2500\,\mathrm{kg}$ of polymer for the separators and the extruder, each moving the pellets' ln MI and density toward its own by the fraction $\min(1, m_p x / M_{bl})$. At the design rate of 17.8 t/h the hold-up is 8.4 min of production, and the pellets are at 1.379 g/10 min and $0.92806\,\mathrm{g/cm^3}$.

No loop reads a parcel. Thermocouples sit every 12 m from 6 m to 150 m after each injector, then every 40 m from 190 m after it to the end of the zone, 70 in all. Each samples a parcel's temperature as that parcel's center crosses it, interpolated quadratically in time within the step, and adds a 1.5 s thermowell lag and 0.3 K of white noise at each scan, and TI-10j is the hottest of zone j. PT-300 has a 0.3 s lag and 1.5 bar of noise, and TT-201 to TT-204 have 3 s and 0.1 K. The melt indexer AI-501 and the laboratory density AI-502 are the framework's shared sampled analyzer. AI-501 draws from the hold-up every 6 min through a 2 min sample line and reports 6 min later with 0.02 g/10 min of noise, so its reading is at least 8 min old, and the laboratory reports density 30 min after an hourly sample with $0.0002\,\mathrm{g/cm^3}$ of noise.

Table 10.2 Model Constants. Visible Ones Can Be Moved in the Show Model Panel, Hidden Ones Only in the Source

Symbol	Meaning	Value	In panel
F_{in}	ethylene feed	64.8 t/h	visible
ρ_0, κ	density, compressibility	550, 6.3×10^{-5}	hidden
P_{ref}, P_{sep}	reference, separator	2500, 250 bar	hidden
a_{nom}	let-down at design	0.5	hidden
c_p	heat capacity	2600 J/(kg K)	visible
$-\Delta H_p$	heat of polymerization	93.6 kJ/mol	visible
f	initiator efficiency	0.6	visible
G_{ref}, E_G	$k_p/\sqrt{k_t}$ at 280 °C	0.30, 29	visible
ΔV_G	its activation volume	−29.5 cm^3/mol	visible
E_d	peroxide activation energy	90 kJ/mol	visible
$T_{h,1..4}$	one-second half-life	212, 228, 232, 236 °C	visible
ΔV_d	peroxide activation volume	6 cm^3/mol	hidden
C_s	chain-transfer constant	0.3	visible
T_{dec}, t_{dwell}	decomposition, dwell	335 °C, 1.0 s	visible
$k_{dec,0}$, E_{dec}	rate, activation	5×10^{-3}, 450 kJ/mol	visible
$-\Delta H_{dec}$	heat of decomposition	120 kJ/mol	hidden
U_0, F_w	clean wall, water	800, 60 kg/s	visible
m_s, τ_J	section water, cooler lag	2000 kg, 20 s	hidden
ε	cooler effectiveness	0.8	visible
g_0, ϕ_p	fouling growth, pulse strip	4×10^{-7}, 0.3	visible
T_{c0}, T_{c1}	cloud point	150 °C, 300 K	hidden
MI_{ref}, c_{ref}, T_Q	quality reference	2.0, 1500, 261.5	hidden
a_C, a_T	melt-index slopes	1.0, 0.025 K^{-1}	visible
ρ_{ref}	density reference	0.9235 g/cm^3	hidden
b_T, b_P	density slopes	3.5×10^{-4}, 4×10^{-6}	visible
M_{hl}	product hold-up	2500 kg	hidden
$F_{I,max}$	pump 1, pumps 2 to 4	13.0, 3.8 mmol/s	hidden
—	pump, let-down lags	1.0, 0.5 s	hidden
—	pulse stroke, longest	0.25, 10 s	both

Units where not shown: ρ_0 in kg/m^3, κ in bar^{-1}, G_{ref} in $(m^3/(mol\,s))^{1/2}$, E_G in kJ/mol, $k_{dec,0}$ in s^{-1}, U_0 in W/(m^2 K), g_0 in m^2 K/(W s), MI_{ref} in g/10 min, c_{ref} in mol ppm, T_Q in °C, b_T in g/(cm^3 K), b_P in g/(cm^3 bar).

10.3.7 Numerics, Initialization and Checks

The parameters are collected in Table 10.2. The tube advances on a fixed step of 0.25 s. Within a step the hydraulics of Equation (10.4) and the actuator lags are advanced first, with one classical fourth-order Runge-Kutta step, every parcel then moves between its exact start and end positions, and the jackets of Equation (10.5) are advanced last with the parcels' heat for the step. Near a peak the initiator decay is stiff, with $k_d\,\Delta t$ up to 10, so each parcel is advanced by Lawson's integrating-factor form of the Runge-Kutta method with $k_{d,j}$ frozen over a substep, and the substep is shortened until T moves at most 10 K across it at its starting rate. Case 3 needs at most four substeps. A parcel with no initiator below 300 °C is cooled by the exact exponential and its slow decomposition added explicitly. The initiator injected since the previous parcel passed an injector enters the next parcel at the instant its center passes, inside the step. The controllers run every 1 s, at the start of the physics step that begins each whole second, a scripted event acts at the start of the step at its time, the historian logs one row a second, and the random sequence is seeded at 20260927, so every run replays bit for bit, as Section 1.6.1 describes for every framework app.

Each scenario opens on the steady tube of its own configuration, and a setting changed before Run, or carried through Reset, rebuilds the steady tube for the settings in force, so a case started with its peak loops at 310 °C starts from 310 °C peaks. A setting no steady tube can hold, such as a closed cooler or a setpoint at the line, starts the tube from the scenario's own settings and then acts as a step. One parcel is marched from the inlet with the engine's own step, the two sections of each jacket are iterated to their balance, and in each zone under automatic control the initiator flow is found by regula falsi so that the zone's hottest thermocouple sits on its setpoint. Every parcel is

then marched from its own entry on the same step grid. Held there in MANUAL with the fouling off, no thermocouple moves by more than 0.18 K in a minute. The parcel, tube and jacket energy ledgers close to 3.8×10^{-15} relative through pulses, fouling and the decomposition of case 5, and $C\,\Delta P$ matches the integrated feed less the let-down to 5.6×10^{-16}. The same initiator flows on 60 or 80 parcels move the steady thermocouple peaks at 300 °C by at most 0.073 K (0.23 K at 290 °C) and the conversion by 0.007 points, halving the step moves the steady peaks by 0.021 K and both the dip and the overshoot of case 3's first pulse by 0.10 K, and a parcel with no cooling heats by 12.834 K per percent of conversion. These checks are among the 55 that the application passes. The speed selector runs from 0.05 to 60 simulated minutes per second, with 0.2 as the default, and the top rung is set by the 68 800 steps a second the browser sustains. At the default speed a pulse of 5.75 s passes in under half a second, and the four-hour challenge plays in four seconds at the top speed.

10.3.8 Assumptions

The tube is plug flow in parcels that are each well mixed, with no axial dispersion, and a parcel receives its injected initiator all at once. This breaks down wherever turbulence mixes neighboring gas over the length of a reaction front, and it makes the drawn profile's tip move by about 2 K from frame to frame as parcels pass the peak.

The kinetics are illustrative, with one lumped $k_p/\sqrt{k_t}$, radicals in quasi-steady balance, four peroxides that share one activation energy, and no transfer to polymer, long-chain branching, molecular-weight distribution or thermal initiation. The lot efficiency is carried with each dose, so a lot change reaches only the initiator pumped after it, and the switch between lots is instantaneous. The kinetics give peaks, conversions and initiator use of the right order and are calibrated to no plant, so a question about molecular weight, branching or a real peroxide's half-life lies outside the model.

The density follows the pressure but not the temperature, so the velocity varies along the tube only while the pressure moves, and the expansion cooling of a pulse is left out, although a 300 bar pulse expands the gas by 1.9 %. The polymer's separation from ethylene enters only as the cloud point that grows the layer, and the model stops being valid below about 2000 bar, above about 340 °C and above about 36 % conversion.

Each jacket is two well-mixed sections with a velocity-scaled clean coefficient, too coarse to resolve the water's profile along a section, and the fouling law is illustrative. The decomposition is a declared event with no model of what follows it, so nothing after the event is physical. The quality correlations are illustrative and hold only near the grade they were written for, the separators and extruder are one well-mixed hold-up, and the analyzers never fail. Pumps and valves are first-order lags without stiction, the hyper-compressor delivers a fixed mass flow whatever the pressure does, and the sizes and limits describe one teaching reactor rather than a design basis.

10.4 How a Peak Forms and What Moves It

A peak is a small runaway limited by its initiator and lit on purpose. At the design point the gas reaches injector 2 at 225.8 °C, where its peroxide decomposes at 0.63 s^{-1}. The polymerization heats the gas, the hotter gas decomposes the rest of the peroxide faster, 7.5 s^{-1} at 290 °C, and within 28.2 m, about 2 s of travel, the initiator is spent and the jacket takes over. Zone 1's gas arrives at only 167.9 °C, where its faster peroxide decomposes at 0.075 s^{-1}, and its peak forms 53.5 m after the injector. The conversion climbs in steps, 9.94, 15.63, 21.58 and 27.51 % at the ends of the four zones, and the gas leaves at 224.8 °C.

A larger dose heats the gas faster, so the rest of the initiator decomposes sooner and the front burns out earlier. The peak therefore rises and moves toward its injector, as shown in Figure 10.2. Between 5 and 80 % of stroke, zone 2's hottest thermocouple rises from 245.4 to 319.4 °C, and the hottest gas moves from 55.8 to 17.9 m after the injector. The gain falls as the stroke rises, 1.9 K per percent at 10 to 15 % and 0.6 K at 45 to 55 %, so one tuning is a compromise across a pump's range. The measurement adds its own distortion. The hottest thermocouple jumps from the one 54 m after

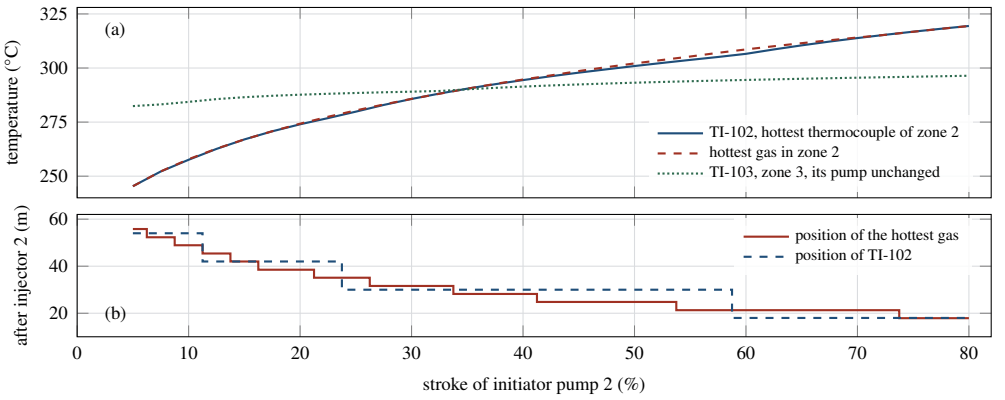

Figure 10.2 Steady peak of zone 2 against the stroke of its pump, from the application's steady solver with the other pumps at their design strokes and the jacket inlets at 175 °C. The measured peak rises with a slope that falls from 2.7 to about 0.5 K per percent while the hottest gas moves from 56 to 18 m after the injector, and the hottest thermocouple hops between readings 12 m apart. The position of the hottest gas is resolved to the 3.4 m a parcel moves in one step. In (a) TI-102 is solid, the hottest gas dashed and TI-103 dotted, and in (b) the hottest gas is solid and TI-102 dashed.

the injector to those at 42, 30 and 18 m as the peak moves, and at 55 % of stroke it reads 303.8 °C while the gas between two thermocouples reaches 305.3 °C.

🕮 **Pause.** Try Exercise 10.2 before reading on.

The zones are coupled one way. With its own pump unchanged, zone 3's hottest thermocouple rises from 282.4 to 296.4 °C as pump 2 goes from 5 to 80 %, because zone 2's hotter tail is the gas into which zone 3's initiator is injected. Open-loop steps in MANUAL from the design point measure the plant each loop sees, each read against an otherwise identical unstepped run. A 5 % step of pump 2 raises the lagged TI-102 by 4.06 K at 180 s, 0.81 K per percent, with the first 6 % of the change at the scan 3 s after the step, 41 % at 5 s and 63 % by 7 s, so the peak loop's plant is roughly 4 s of dead time and 2 s of lag. Zone 3 feels the same step only after 30 s, most of it the travel from zone 2's peak to injector 3, and gains 1.49 K by 180 s as jacket 2 warms.

A jacket acts mainly on the zone after it. Opening cooler 2 by 10 % lowers jacket 2's inlet by 8.9 K, 63 % of it in 2.0 min, and 10 min after the step, with the pumps held, zone 2's peak is 0.83 K lower but zone 3's 3.08 K lower. With the loops closed and the pulses off, a jacket 2 inlet setpoint 10 K warmer moves pump 3 from 36.5 to 31.1 % within 5 min and pump 2 only from 34.5 to 33.4 %. The jacket is a slow handle on the next zone's pump, not on its own.

The pulse moves everything at once. The gas expands, its velocity rises, the fronts are carried downstream, and the negative activation volume slows the polymerization while the pressure is low, so every peak dips. The thermocouples see the gas late when it moves fastest. Before the first pulse of case 3 the hottest gas runs at most 0.77 K ahead of the hottest thermocouple, while 42.5 s into the run, as the pressure recovers, the gas in zone 4 reaches 293.4 °C and no thermocouple reads above 282.5 °C, 10.9 K lower, because the peak is rising faster than a 1.5 s thermowell and a sample at each passing parcel can follow.

10.5 Implemented Controller

Thirteen PI loops and five pieces of logic run every $\Delta t_c = 1\,\text{s}$, four physics steps, in the order instruments, high-high trip, pulse scheduler, feedforward, peak loops, jackets with mid-ranging, pressure, quality trim and supervisor. In MANUAL the operator holds the four pump strokes, the four cooler valves, the let-down and the CTA, and every loop tracks what the operator sets. Within each

Table 10.3 The Thirteen PI Loops With Their Directions and Default Tunings

Tag	CV	MV	γ	K_c	τ_I, output
TIC-101..104	TI-10j	$u_{I,j}$	+1	0.5	6 s, 0–100
TIC-201..204	TT-20j	$u_{C,j}$	−1	3	30 s, 0–100
VPC-602..604	$u_{I,k}$	$T_{ji,k-1,\mathrm{SP}}$	−1	0.5	300 s, 140–240
PIC-300	PT-300	a_{LD}	−1	0.1	10 s, 0–100
AIC-501	AI-501	c	+1	300	900 s, 100–5000

Gains in % per K, % per K, K per %, % per bar and ppm per g/10 min.

physics step the controller runs before the integration, which runs before the post-step corrections, the outputs and the historian. An optional model predictive controller, Section 10.5.6, takes the four pumps in place of the peak loops' algorithm when it is selected, while every other loop and piece of logic runs as described here.

10.5.1 The PI Algorithm

Every loop runs the framework's shared velocity-form PID of Section 1.6.2, with its derivative term unused here, as shown in Equation (10.8).

$$\begin{aligned} e_k &= \gamma(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\, e_k \\ S_k &= S_{k-1} + (P_k - P_{k-1}), \qquad v_k = S_k + I_{k-1} \\ u_k &= \mathrm{clamp}(v_k, u_{min}, u_{max}) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I}\, e_k + \frac{u_k - v_k}{\tau_I}\right]\Delta t_c \end{aligned} \tag{10.8}$$

where $\gamma = +1$ where a larger output must follow a measurement below its setpoint and -1 where it must follow one above. These are Equations (1.14) to (1.17) with γ for the action a, no output bias, no derivative and a tracking time equal to τ_I. The running sum S_k is never clamped, and windup is handled on the integral alone by back-calculation, which the shared algorithm omits in a loop with no integral action. No loop here uses the derivative filter or a rate limit. The integral is seeded at reset with the output the steady tube needs, and in MANUAL, in fixed-stroke operation or while a feedforward window or a trip holds a loop, the loop tracks by setting $I_k = u_{op} - S_k$, which makes its output the held value, so every return to automatic is bumpless. Each loop's pairing, direction and default tuning appear in Table 10.3.

Under the ISA convention the peak loops and the melt-index trim are reverse acting and the jacket, pressure and mid-ranging loops direct acting. A cold peak raises its pump, a warm jacket inlet opens its cooler, a pressure above setpoint opens the let-down, a low melt index raises the CTA, and a pump above its target warms the jacket upstream of it.

10.5.2 Peak Loops and the Pulse Feedforward

TIC-101 to TIC-104 hold each zone's hottest thermocouple with that zone's pump, with one shared tuning or, in advanced mode, one per zone. In fixed-stroke operation the pumps hold the strokes on their cards while the jacket and pressure loops run. The pulse scheduler XV-310 opens the pulse valve every $t_p = 180$ s until PT-300 has fallen by the depth $\Delta P_p = 300$ bar from its value at the opening, or for at most 10 s, and holds PIC-300 at its last output until the valve has closed. It checks PT-300 once a second while the pressure falls about 60 bar a second, so a pulse can overrun its depth by up to a second of fall, to 317 bar in case 3 and 362 bar in case 4. Because the pulse is scheduled, FF-310 knows it in advance. From t_ℓ before each scheduled pulse to t_h after it, each peak loop tracks the output u_j^* it had when the window opened, and its pump receives that output reduced as shown in Equation (10.9).

$$u_{I,j} = (1 - \delta\,\omega)\,u_j^*, \qquad \omega = \min\left(1, \frac{t_p + t_h - t}{r}\right) \tag{10.9}$$

where the window runs from $t_p - t_\ell$ to $t_p + t_h$, $r = \min(5\,\text{s}, (t_\ell + t_h)/2)$ ramps the cut out at its end, and the defaults are $\delta = 20\,\%$, $t_\ell = 3\,\text{s}$ and $t_h = 30\,\text{s}$. The loops therefore neither fight the dip nor wind up against it, and they resume from u_j^* without a bump. The feedforward is keyed to the schedule, not to the valve, so it also cuts the initiator for a pulse that fails to open.

Pause. Try Exercise 10.3 before reading on.

10.5.3 Jackets and Mid-Ranging

TIC-201 to TIC-204 hold each jacket inlet with its cooler valve. Because a jacket mainly sets the temperature at which the next zone's initiator is injected, VPC-60k keeps pump k near a target stroke $u_{mid} = 50\,\%$ by moving the setpoint of jacket $k-1$, for $k = 2$, 3 and 4. Pump 1 has no mid-ranging, since its lever would be the feed preheater, which the model does not control. When a cooler is pinned in the direction mid-ranging would push it, wide open for a pump that wants colder water or shut for one that wants warmer water, the mid-ranging loop tracks its present output instead of integrating, which is conditional integration against a limit it cannot move [22].

10.5.4 Pressure, Trip and Restart

PIC-300 holds 2500 bar with the let-down valve. The high-high trip TSHH compares every zone's hottest thermocouple with T_{HH}, 325 °C by default, at each scan, and once it fires it latches, stops all four pumps, opens the let-down wide and suspends the pulses until it is reset. The permissive PSL-300 keeps the pumps at zero while PT-300 reads below 1500 bar, and when they are released each peak setpoint ramps back from the lower of its peak and its setpoint at 1 K/s, because in a counterfactual run without them a reset of case 6's trip decomposes the plant 18.25 s later.

10.5.5 Quality Trim and Supervisor

AIC-501 moves the CTA on each fresh melt-index report, with Δt_c of Equation (10.8) replaced by the 6 min report period, and holds while the analyzer is stale. The optional supervisor XY-700 walks each peak setpoint toward the highest value whose recent excursion stays a declared margin Δ_{set}, 20 K by default, below the line, as shown in Equation (10.10), which is Equation (1.13) with $\hat{e}_j$ as the excursion, $2\sigma_{TC}$ as the measurement error and Δ_{set} as the reserve.

$$\begin{aligned} \hat{e}_j &\leftarrow \max\left(\hat{e}_j\, \mathrm{e}^{-\Delta t_c/\tau_e},\ \mathrm{PV}_j - \mathrm{SP}_j\right) \\ \mathrm{SP}_j^{w} &= T_{dec} - \Delta_{set} - \max(0, \hat{e}_j) - 2\sigma_{TC} \\ \hat{\rho} &= \rho_{lab} - b_T\,(\bar{T}_{pk} - \bar{T}_{pk,lab}) \end{aligned} \tag{10.10}$$

where $\tau_e = \max(60\,\text{s}, 2t_p)$, $\sigma_{TC} = 0.3\,\text{K}$, and each setpoint moves toward SP_j^{w} by at most 1 K/min upward and four times that downward, between 250 °C and $T_{dec} - \Delta_{set}$. The density estimate $\hat{\rho}$ is the laboratory report moved by the change in the mean measured peak $\bar{T}_{pk}$ since its sample was drawn, a bias update of the kind Section 1.4.6 describes, and a rise is cut to the room left before $\hat{\rho}$ reaches the lowest density accepted, 0.9235 g/cm^3. While the melt-index trim is saturated with the melt index high, the supervisor lowers the setpoints instead.

10.5.6 Model Predictive Control

The control architecture offers a third setting for the initiator pumps beside fixed strokes and the four peak loops, one model predictive controller (MPC) that moves all four pumps every $T_s = 2\Delta t_c$ over a horizon of $N_p = 90$ decisions (180 s) and keeps every predicted peak under one ceiling, the trip setting less a margin, $T_{ceil} = T_{HH} - \Delta_{ceil}$ with $\Delta_{ceil} = 5\,\text{K}$ by default. It reads the signals TIC-101 to TIC-104 and FF-310 read, with TT-201 to TT-204, the trip setting and the strokes it applied, and nothing of the parcels, the kinetics or the margin. Every other loop and every protection runs above it as under the loops, and FF-310 is idle because the plan carries the pulse itself.

The controller's model was fitted to step tests of the simulator at the design point of cases 2 and 3 and is shipped as constants, so a plant the learner has changed is model mismatch, as on a real reactor.

From pump j to the hottest thermocouple of its own zone the path is one lag and a dead time with a gain that falls with the stroke, while the three cross paths, the pressure and the jacket inlets enter as output offsets simulated at the 1 s scan from the strokes already sent and the measured signals, as shown in Equation (10.11).

$$\begin{aligned}
\tau_j \frac{\mathrm{d}\hat{s}_j}{\mathrm{d}t} &= u_{I,j}(t-D_j) - \hat{s}_j, \quad \frac{\mathrm{d}F_j}{\mathrm{d}\hat{s}_j} = A_j\, \hat{s}_j^{\,-g_j} \\
\tau_x \frac{\mathrm{d}\hat{\zeta}_j}{\mathrm{d}t} &= u_{I,j}(t-D_x) - \hat{\zeta}_j, \quad \tau_P \frac{\mathrm{d}\hat{p}_j}{\mathrm{d}t} = \Delta P(t-D_P) - \hat{p}_j \\
\tau_w \frac{\mathrm{d}\hat{w}_j}{\mathrm{d}t} &= T_{ji,j}(t-D_w) - \hat{w}_j \\
\tau_w \frac{\mathrm{d}\hat{w}'_j}{\mathrm{d}t} &= T_{ji,j}(t-D'_w) - \hat{w}'_j \\
\hat{T}_j &= F_j(\hat{s}_j) + K_{P,j}\,\hat{p}_j + K_{w,j}\,\hat{w}_j + b_j \\
&\quad + K_{x,j-1}\,\hat{\zeta}_{j-1} + K'_{w,j-1}\,\hat{w}'_{j-1}
\end{aligned} \tag{10.11}$$

where a hat marks a state of the controller's model, $\hat{s}_j$ is the effective stroke in percent, F_j the static peak it gives, ΔP the measured PT-300 less its setpoint in bar, $T_{ji,j}$ the measured TT-20j, the terms in $j-1$ are absent for zone 1, and b_j is an output bias in kelvin. The own paths have $\tau_j = 2.8$, 1.1, 3.1 and 2.4 s, $D_j = 6$, 6, 4 and 6 s, $A_j = 28.12$, 19.37, 17.46 and 14.69 and $g_j = 0.890$, 0.863, 0.835 and 0.775. The cross paths have $K_{x,j} = 0.204$, 0.241 and 0.123 K per percent with $\tau_x = 4$ s and $D_x = 32$ s, the 24 to 25 s the gas takes from one peak to the next injector plus the own path, the pressure path has $K_{P,j} = 0.0713$, 0.0457, 0.0465 and 0.0523 K/bar with $\tau_P = D_P = 1$ s, and the jacket paths have $K_{w,j} = 0.149$, 0.098, 0.133 and 0.107 and $K'_{w,j} = 0.398$, 0.485 and 0.287 K/K with $\tau_w = 30$ s, $D_w = 35$ s and $D'_w = 27$ s. Refitted with the noise off, the own paths leave residuals of 0.7 to 2.2 % of the response at the design strokes and the other paths 0.5 to 8.7 %, and the design gains of 1.32, 0.88, 0.95 and 0.83 K per percent lie within 13 % of the schedule, with a delay plus lag of 8.8, 7.0, 7.0 and 7.6 s, the transport from each injector to its hottest thermocouple plus the pump lag and the thermowell.

The pulse enters through ΔP. With the preview on, as case 9 ships it, each scheduled pulse enters the prediction of ΔP as 91 one-second values of PT-300's measured deviation through a pulse of case 3, scaled to the depth setting, and what the template does not explain decays with PIC-300's recovery constant of 6.9 s, as the whole measured deviation does with the preview off. The bias b_j is the measured peak less the model's, filtered with $\tau_b = 5$ s at every scan under whichever structure acts, the output-disturbance form of offset-free control [23]. While $|\Delta P|$ exceeds 30 bar, and for 15 s after, the bias is not updated and the plan does not track its targets, which is FF-310's hold in the controller's own terms. After the pumps have been blocked by the trip or the permissive the model is put back on the measurements with the bias at zero.

Each decision solves the quadratic program of Equation (10.12) over the changes of the four strokes, taken as offsets from the strokes applied now, in twelve move blocks, four of one decision, then 2, 2, 3 and 4, then 6, 9, 15 and 45, 48 moves and four slacks on 15 states, the four lags and their delay chains.

$$\begin{aligned}
\min_{\Delta u,\,\varepsilon} \; & \sum_{k=1}^{N_p} q_k \sum_j \left(\hat{T}_{j,k} - r_j\right)^2 + \lambda \sum_i \sum_j \Delta u_{I,j,i}^2 \\
& + \sum_j \left(10^4\, \varepsilon_j^2 + 2\times 10^4\, \varepsilon_j\right) \\
\text{s.t. } & 0 \le u_{I,j,i} \le 100, \qquad \varepsilon_j \ge 0 \\
& \hat{T}_{j,k} \le T_{ceil} + \max(0,\, \hat{T}_{j,0} - T_{ceil})\, \mathrm{e}^{-kT_s/\tau_f} + \varepsilon_j
\end{aligned} \tag{10.12}$$

where $r_j = \min(\mathrm{SP}_j, T_{ceil})$ is the target, the request or the ceiling, whichever is lower, q_k is 1 except on the steps where the predicted $|\Delta P|$ exceeds 30 bar and for 15 s after them, where it is zero, $\lambda = 5$ per percent squared is the move price, i counts the move blocks, $\hat{T}_{j,0}$ is the model's present output and $\tau_f = 10\,\mathrm{s}$ is a funnel through which a peak already above the ceiling is asked back. The ceiling is a soft constraint priced 10^4 per K^2 and 2×10^4 per kelvin against 1 per K^2 of tracking, with one slack ε_j per zone over the whole horizon. The target is clipped rather than left at the request, because a request above the ceiling pulls the plan into that shared slack and lets the peaks ride above the ceiling. The stroke bounds are the peak loops' own and are hard. The framework's dual active-set solver [24] computes each plan with a ridge of $10^{-9}(1 + 2N_p K_{max}^2 + 2\lambda)$, K_{max} being the largest present gain, and only the first block's moves are applied.

A plan is accepted only when the solver reports an optimum and every move, prediction, cost and slack is finite, with the slacks nonnegative and the moves inside 0 to 100 %. Otherwise the historian is marked and the four peak loops, which track the applied strokes at every scan, take the pumps until the next valid plan. With the noise off a forced rejection at 120 s of case 9 moves no stroke at the hand-over, and one at 243 s, inside a pulse window, moves 0.06 points, because a window entered from a fallback holds the strokes without the feedforward's cut. In Chrome a plan takes 2.1 to 2.9 ms, 6.4 to 6.6 % of a core at the 1 min/s rung of case 9, and the 60 min/s rung slows to 17 to 21 min/s.

10.5.7 Simulation, Identified and Internal Models

Three models are kept apart here. The parcel model of Section 10.3 is what the plant does. Under the peak loops nothing inside the application fits a control model, and the open-loop steps of Section 10.4 serve as one, measuring the plants the loops see, 0.81 K per percent with a few seconds of dead time for a peak loop at mid-stroke and 0.89 K per percent with 2.0 min to 63 % for a jacket loop. Under the loops the controllers carry almost no internal model. The feedforward knows the schedule and a fixed cut, the scheduler knows the measured pressure, the supervisor knows T_{dec}, b_T and the laboratory, and none of them contains the parcels, the kinetics or the margin itself. The MPC's identified model, Equation (10.11), fitted once by the same kind of step test, contains none of them either.

10.6 Nine Cases and the Challenge, as Measured

Table 10.4 sets out the nine cases and the four-hour challenge, and Table 10.5 their measured outcomes. Each case starts from the four peak loops at 290 °C, the trip at 325 °C, pulses every 180 s from the first minute and the feedforward, mid-ranging, trim and supervisor off, and changes only what its row names. Every number below was read from the running engine at the default seed, times are simulated, and a controller action is dated by the scan at which it acts. Where a discrete outcome could turn on the measurement noise, the spread over six other seeds is given.

10.6.1 A Peak Lit by Hand

Case 1 starts with every station in MANUAL and pump 2 stopped. The tube converts 24.0 %, zone 2's gas passes its idle injector at 223.3 °C and only cools, and zone 3's peak is 267.4 °C because its gas arrives colder. Pump 2 starts at 20 % at 30 s, zone 2's thermocouples pass 260 °C 9.0 s later, and at 1.0 min the peak is 270.8 °C at the thermocouple 42 m after the injector. At 40 % from 90 s it is 293.0 °C at 3 min, at the one 30 m after it. Zone 3's peak, whose pump never moves, climbs to 280.2 °C by 2 min as zone 2's heat arrives with the gas and to 289.3 °C by 8 min as jacket 2's inlet warms from 159.0 to 173.3 °C. Zone 4's peak dips from 290.7 to 285.7 °C at 1.63 min because the zones before it have left it less ethylene, and the conversion rises to 28.3 % at 3 min and settles back to 27.8 % at 8 min.

10.6.2 Four Loops Coupled Downstream

Case 2 closes the four loops with the pulses off. At 30 s zone 2's setpoint steps from 290 to 300 °C, and TIC-102 moves pump 2 from 34.5 to 48.5 % by 3 min. TI-102 approaches 300 °C from below, stays within 1 K from 46 s after the step, and overshoots by no more than the noise of the hottest

Table 10.4 The Nine Cases and the Challenge

#	Scenario and events	Teaching point
1	MANUAL, pump 2 at 0, 20 at 0.5 min, 40 at 1.5 min	a peak forms and moves
2	pulses off, zone 2 to 300 °C at 0.5 min	one-way coupling
3	first pulse at 30 s, feedforward off	windup, feedforward
4	fixed strokes, jackets 160 °C, pulses every 15 min from 25 min	fouling
5	setpoints 315, 322, 328 °C at 1, 3, 5 min, trip 330 °C	the trip's dead time
6	300 °C, feedforward, trip 320 °C, lot 50 % at 1 min, 100 % at 10 min	gain change
7	feedforward and trim on, 310 °C at 2 min	two handles, late analyzers
8	300 °C, feedforward, mid-ranging at 55 %, supply 140 to 170 °C at 15 min	cooling runs out
9	MPC on the pumps, feedforward, 315, 322, 328 °C at 1, 3, 5 min	a request and a constraint
10	4 h, supply 165 °C at 60 min, lot 120 % at 150 min	the margin

Table 10.5 Measured Outcomes of Each Case and Variant

Case, configuration	Measured
1 as shipped	zone 2 peak 270.8 then 293.0 °C, 42 then 30 m after the injector
2 as shipped	in 1 K from 46 s, no overshoot above noise, zone 3 +2.06 K
3 feedforward off	dip 270.1, overshoot 299.7 and 299.9 °C, worst margin 33.1 K
3 feedforward on	dip 262.7, overshoot 292.2 and 292.8 °C, worst margin 41.9 K
3 hold 10 s	overshoot 294.2 and 295.8 °C, worst margin 38.2 K
4 fixed strokes	U/U_0 0.68, zone 4 peak 279.4 to 293.2 °C
4 peak loops	zone 4 at 290.0 °C, pump 4 at 29.4 %
5 trip 330 °C	trip at 7.350 min, decomposition at 7.375 min
5 trip 325 °C	trip at 4.283 min, worst margin 8.7 K
5 feedforward on	trip at 8.083 min, worst margin 4.8 K
6 trip on	trip 25 s after the lot, worst margin 12.9 K
6 trip off	decomposition 35.25 s after the lot
7 as shipped	in grade after 13.8 min, 4.54 t off grade
8 mid-ranging on	mean conversion 29.4 %, coolers 1 to 3 at 100 %
8 mid-ranging off	mean conversion 28.9 %, coolers 91 to 94 %
9 MPC, then loops	hottest 322.1 °C, margin 12.8 K, 31.54 %, no trip, then trip at 5.083 min, 19.25 %

of 17 thermocouples, its highest reading over the next 5 min being 300.8 °C. Its hotter tail reaches zone 3's injector and pushes TI-103 to 292.1 °C, 2.06 K above its setpoint, 37 s after the step, and TIC-103 cuts pump 3 from 36.5 to 33.3 % by 3 min. The conversion rises from 27.5 to 27.9 %. Each loop rejects what the zone before it sends, so a zone tuned hot becomes a disturbance to every zone after it.

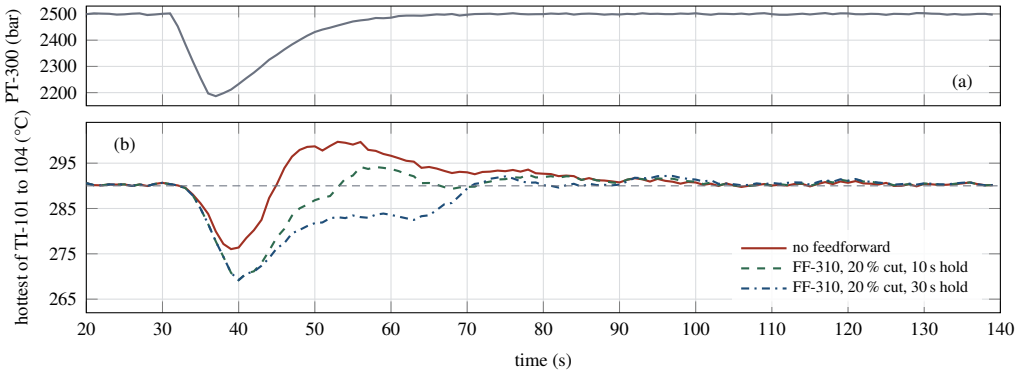

Figure 10.3 First pulse of case 3 ridden through three ways, one logged sample a second at the default seed. The pressure (a) is the same in all three runs. The hottest of the four peak measurements (b) overshoots to 299.7 °C without feedforward (solid), to 294.2 °C with a 20 % cut held 10 s (dashed) and to 292.2 °C with the default 30 s hold (dash-dotted), against the 290 °C setpoints marked by the thin gray dashed line. The lowest of the four, not drawn, dips to 270.1 °C without feedforward and to 262.7 °C with it.

10.6.3 Riding the Pulse

In case 3 the first pulse is scheduled at 30 s and opens at the scan at 30 s. The scheduler holds it open for 5 s, the valve is more than 2 % open for 5.75 s, the pressure falls 317.0 bar to 2184.0 bar, the gas leaves at 20.3 m/s instead of 13.8, and the mean velocity along the tube is 20.2 % higher while the valve is open. The pressure is back above 2490 bar 29.8 s after the scheduled start. The zones' hottest thermocouples move 12 to 72 m downstream, and the peaks dip to 270.1 °C. The loops see an error of 20 K and wind their pumps up, pump 1 to 53.5 %, and when the pressure returns that initiator takes TI-101 to 299.7 °C, 9.7 K above its setpoint, with the hottest gas at 301.6 °C. The conversion at the outlet swings between 26.05 and 28.16 %. The second pulse, at 210 s, repeats the pattern with an overshoot to 299.9 °C and the gas at 301.9 °C, the state shown in Figure 10.4, and it sets the worst margin of the two pulses at 33.1 K, against 33.4 K after the first.



Figure 10.4 The pulse valve, case 3, at $t = 3.90$ min (234 s) in the Profile view, 24 s after the second pulse was due, with no control changed. (1) The badges read a worst margin of 33.1 K, a conversion of 27.4 %, a melt index of 1.38 and a density of 0.9280, both flagged against the grade, U/U_0 of 0.98 and PT-300 at 2464 bar, with the next pulse 156 s away. The unrolled tube shows the fluid temperature of the parcels with four peaks, (2) the hottest thermocouple of each zone ringed, TI-101 at 299.9 °C at the top of its second overshoot, TI-102 at 293.4, TI-103 at 294.7 and TI-104 at 293.8 °C against setpoints of 290 °C, (3) below the dotted amber trip line at 325 °C and the dashed red decomposition line at 335 °C marked as a teaching value. Beneath them lie (4) the jacket water of the eight sections near 181 and 188 °C, the conversion climbing zone by zone on its own axis, and (5) the tube with its injectors, pump strokes and the let-down and pulse valves at the outlet.

TRY IT IN THE APP **3 · The Pulse Valve**

Load the scenario, press Run and watch the four ringed peaks on the Profile view. At 30 s the pulse valve opens, the pressure falls 317.0 bar and the peaks dip to 270.1 °C. Keep watching after the pressure returns: the wound-up loops carry TI-101 to 299.7 °C, almost 10 K above its setpoint.

The feedforward removes most of the overshoot, as shown in Figure 10.3. With a 20 % cut from 3 s before each pulse to 30 s after it, and the loops held meanwhile, the dip of the first pulse deepens to 262.7 °C, the overshoot stops at 292.2 °C, the hottest gas reaches only 292.8 °C, and over the two pulses the worst margin rises from 33.1 to 41.9 K. A hold of 10 s ends before the pressure has recovered, and the overshoot reaches 294.2 °C after the first pulse and 295.8 °C after the second, with a margin of 38.2 K. The margin is bought with conversion at fixed setpoints. Over 30 min the mean conversion is 27.41 % without the feedforward and 26.96 % with it, while the worst margin is 31.9 and 41.9 K, and the challenge shows what those 10 K are worth once the setpoints are raised to spend them. A missed pulse with the feedforward on shows its other side, because the cut still arrives on schedule and the peaks dip to 279.6 °C with no pulse at all.

10.6.4 Fouling Between Sparse Pulses

Case 4 runs the jackets at 160 °C, holds the pumps at fixed strokes and delays the first pulse to 25 min. The layer grows where the wall is below the cloud point, and by 24.9 min its mean resistance is 1.72×10^{-3} m^2 K/W in zone 4, 6.1×10^{-4} in zone 3, 5.0×10^{-5} in zone 2 and none in zone 1. U falls

to 0.82 of its clean value in 10 min, 0.76 in 15, 0.72 in 20 and 0.68 by 24.9 min. Less heat leaves zone 3's tail, so zone 4's initiator meets warmer gas and its peak creeps from 279.4 to 293.2 °C while zones 1 to 3 move by less than 2 K, and the conversion slips from 28.6 to 28.1 %. The pulse at 25 min stays open 6 s, because the scan 5 s into it read PT-300 1 bar above its closing value, so it falls 362 bar and strips 36.2 % of the layer, and U is back to 0.76 half a minute later. In case 3, with the jackets at 175 °C and a pulse every three minutes from 30 s, U is still 0.95 at 24.9 min. With the loops closed in case 4, zone 4 stays at 290.0 °C while pump 4 falls to 29.4 %. Parametric sensitivity to the wall coefficient [5] appears here as a peak that drifts toward the line or a pump that drifts toward its stop. The butyl rubber reactor of Chapter 11 fouls its jacket with its own product in the same way, without a pulse to strip it.

10.6.5 An Honest Decomposition

Case 5 raises the setpoints to 315, 322 and 328 °C at 1, 3 and 5 min, pulsing without feedforward, with the trip moved to 330 °C. At the steady tube each kelvin of peak buys about 0.2 points of conversion, from 27.51 % at 290 °C to 29.65 % at 300, 31.56 % at 310 and 33.54 % at 320 °C. After the pulse at 7 min zone 1's own thermocouple passes 330 °C and trips the plant at 7.350 min, but the gas behind it has already been dosed, and 1.5 s later, at 7.375 min, gas 64 m down zone 1 has been above 335 °C for a second, at 336.7 °C. It decomposes, the app reports about 1760 °C and 3010 bar, and the run stops. The outcome turns on one scan. A trip forced one scan earlier leaves 1.4 K, and over six other noise seeds the trip fires at 7.333 min three times, with 1.1 to 1.3 K to spare, and at 7.350 min three times, each followed by a decomposition. With the trip at 325 °C it fires at 4.283 min with the hottest gas 8.7 K below the line (8.7 to 9.1 K over the seeds), and with the feedforward and the 330 °C trip the plant trips at 8.083 min with 4.8 K to spare (3.9 to 5.1 K). A protection layer has a dead time of its own, and here it is the 1 s scan, the thermowell and the transport of initiator already injected. The surge override of Volume 4, Chapter 2 is another guard that a fast enough excursion crosses.

> **TRY IT IN THE APP 5 · Conversion Pushed Toward the Line**
>
> Press Run and watch the hottest parcel against the dashed decomposition line through each pulse. With the trip at 330 °C it fires at 7.350 min, and gas 64 m down zone 1 decomposes 1.5 s later. Reset, set the Trip temperature to 325 °C and run again: the trip fires at 4.283 min with the hottest gas 8.7 K below the line.

Pause. Try Exercise 10.1 before reading on.

10.6.6 A Weak Lot and the Trip

Case 6 holds 300 °C with the feedforward on and the trip at 320 °C, and the pumps start at 47.7, 46.3, 46.2 and 47.2 %. At 1 min the pumps begin to draw from a lot half as effective, and the loops hold the peaks, but by 9.9 min the strokes are 95.3, 93.0, 92.6 and 90.2 %, with half the gain per percent and almost no room. At 10 min the normal lot returns as a pulse opens, with the feedforward holding the commands a fifth lower, at 72.0 to 76.3 %, and the integrals frozen. The initiator already in the gas keeps the weak lot's efficiency, but every dose pumped after the change makes radicals twice as fast. As the pressure recovers, the four peaks climb past 313 °C on doses the frozen loops cannot trim, and the trip fires on zone 1 at the scan 25 s after the change, with the hottest gas at 322.1 °C, 12.9 K below the line. At 325 °C the same trip fires 34 s after the change, 1.25 s before a decomposition, which is why the case sets it at 320 °C. With the trip off, the feedforward hands the pumps back at 90 to 95 % and zone 1 decomposes 35.25 s after the change, at 61 m and 336.4 °C. Over six other seeds the trip fires 25 to 27 s after the change with 12.2 to 12.8 K left, and the decomposition comes 35.0 to 35.5 s after it. Excess initiator is the first runaway cause the autoclave study names [7], and here it arrives as an unmeasured gain change. A reset at 12 min brings the plant back, since the pumps wait until PT-300 reads 1500 bar, reached 11 s later, and by 22 min every peak reads within 0.5 K of 300 °C.

10.6.7 A Grade Transition

Case 7 moves from a grade made at 290 °C, melt index 1.379 g/10 min and density 0.92806 g/cm^3, to the film grade 2.0 ± 0.3 and 0.924 ± 0.001, with the feedforward and the trim on. At 2 min the setpoints go to 310 °C. Hotter peaks make polymer with more short branches, so the density falls, and a higher melt index. The pellets enter the grade box 13.8 min after the setpoints move and stay in it, and 4.54 t of off-grade polymer is made on the way. The trim raises the CTA from 1500 to 1638 ppm to bring the melt index to 1.942 g/10 min by 60 min, while AI-501 reads 1.904. The pellets are then at 0.92426 g/cm^3 while the laboratory still reports 0.92822, from the sample drawn at the start, and the conversion averages 30.62 % over the hour. The peaks carry the density and the CTA the melt index alone. Section 1.4.5 sets out why the hold-up and the analyzers set the off-grade tonnage, and Chapter 8 measures the same cost in off-specification tonnage on a fluidized bed.

10.6.8 Cooling That Runs Out

Case 8 holds 300 °C with mid-ranging on at a 55 % target. By 14.9 min mid-ranging has lowered jackets 1 to 3 to 166.8, 167.7 and 164.0 °C, their coolers run 43 to 49 % open, pumps 2 to 4 have risen to 52, 52 and 49 %, and the mean conversion is 29.6 % against 29.0 % without mid-ranging. At 15 min the supply water warms from 140 to 170 °C. Coolers 1 to 3 are wide open by 15.95 min, the inlets drift to 173.4 to 173.8 °C, several kelvin above the setpoints mid-ranging chose, mid-ranging holds instead of winding up, and pumps 2 to 4 fall back to 48, 47 and 44 %. By 29.9 min the mean conversion is 29.4 % against 28.9 % without mid-ranging, whose coolers run 91 to 94 % open and hold 175 °C. The setpoints are infeasible, and the Zones view reports the cooling as limited rather than met.

10.6.9 Four Peaks Under One Ceiling

Case 9 asks for the 315, 322 and 328 °C of case 5 at 1, 3 and 5 min, pulsing every three minutes from 1 min with the trip at 325 °C, and loads with the MPC on the pumps at a ceiling of 320 °C. Its comparison is the run after Reset with the peak loops and the feedforward on the same requests. Table 10.6 measures cases 9, 5 and 4 under the shipped structure and, with the settings case 9 ships and no retuning, under the MPC, with the same seed, instruments, noise and duration in every row, beside the stronger classical rows: the loops without the feedforward, the loops with the requests clipped by hand to the ceiling, and the best fixed tuning of a 35-point grid chosen in hindsight on each case. Figure 10.5 shows case 9 under both structures.

The loops follow the request into the trip. With the feedforward on, zone 4's hottest thermocouple passes 325 °C at the 5.08 min scan, 5 s after the 328 °C request reaches the pumps, and the run converts 19.25 % over its ten minutes. The MPC takes the 315 °C request, all four peaks within 1 K of it at 2.02 min, stops the 322 °C request at its ceiling, within 1 K of 320 °C at 3.52 min, and refuses the 328 °C one. The strokes hold through the pulses at 4 and 7 min, the peaks dip to 297.8 and 298.0 °C and come back to 320.8 °C after each, and from 7.7 min the four peaks average 320.0 to 320.2 °C. The hottest reading, 322.1 °C in zone 4 at 5.48 min, is the scatter of about 2 K between the parcels a thermocouple samples, the reason the ceiling sits 5 K below the trip. The hottest gas stays 12.8 K below the line, the run converts 31.54 %, and none of its 300 plans is rejected.

Against the operator's own remedy, the loops with the requests clipped by hand to 320 °C, the MPC's lead narrows to 0.45 points of conversion, 1.0 K of hottest reading and a third of the stroke travel, and retuning those loops in hindsight recovers only 0.07 points. Most of what the MPC adds is therefore the ceiling itself, and the rest is the strokes held still through the pulses where the loops integrate against the dip. The scorecard's IAE, 272.8 K min for the MPC, counts the refused request as error.

The MPC loses one comparison. A hindsight-tuned loop without the feedforward and with the requests clipped ($K_c = 1.5$, $\tau_I = 24$ s) converts 0.24 points more by running 0.4 K from the trip with 7.7 K of margin against 12.8 K, which is margin spent rather than control. By design the MPC also

Table 10.6 Shipped Structure Against the MPC on Cases 9, 5 and 4, Default Seed, Full Duration. IAE and Travel Are the Scorecard's, IAE_c Is Against the Lower of the Request and the Ceiling, and the Zone-Seconds Count Measured Peaks Near the Trip (Within 5 K) and Over the Ceiling

Case	Structure	IAE (K·min)	IAE_c (K·min)	Hottest (°C)	Near trip (zone·s)	Over ceiling (zone·s)	Margin (K)	Conv. (%)	Travel (%)	Outcome	Plans, rejected
9, 10 min	loops, FF-310 (shipped)	2963.3	2800.1	325.3	242	219	9.8	19.25	760	trip 5.08 min	—
trip 325,	loops, no FF	3345.3	3176.8	325.1	165	146	8.7	17.12	744	trip 4.28 min	—
ceiling 320	MPC as loaded	272.8	105.7	322.1	679	180	12.8	31.54	216	none	300, 0
	MPC, preview off	267.7	100.4	322.1	674	176	12.8	31.56	215	none	300, 0
	loops, FF, clipped to 320 °C	200.7	200.7	323.1	647	221	11.9	31.09	673	none	—
	clipped, hindsight K_c 1, τ_I 6 s	186.9	186.9	323.1	677	203	11.8	31.16	948	none	—
	no FF, clipped, hindsight 1.5, 24 s	75.9	75.9	324.6	1101	682	7.7	31.78	1577	none	—
5, 10 min	loops, FF-310	1124.9	1109.3	330.2	562	532	4.8	29.48	990	trip 8.08 min	—
trip 330,	loops, no FF (shipped)	85.4	96.1	331.2	560	506	−1.0	31.71	805	decomp. 7.37 min	—
ceiling 325	MPC	164.1	109.8	327.1	494	124	7.9	31.92	244	none	300, 0
	loops, FF, clipped to 325 °C	204.0	204.0	328.0	497	150	7.3	31.47	711	none	—
	clipped, hindsight K_c 0.5, τ_I 3 s	189.0	189.0	328.0	563	180	6.3	31.60	861	none	—
	no FF, clipped, hindsight 1, 12 s	77.7	77.7	329.6	748	342	4.2	32.19	996	none	—
4, 30 min	loops, FF-310	69.9	69.9	292.6	0	0	42.1	29.03	708	none	—
trip 325,	loops, no FF	53.3	53.3	302.0	0	0	30.7	29.07	752	none	—
ceiling 320	MPC	52.2	52.2	291.7	0	0	42.9	29.05	112	none	900, 0
	loops, FF, hindsight K_c 0.35, τ_I 6 s	67.6	67.6	292.3	0	0	42.3	29.03	497	none	—

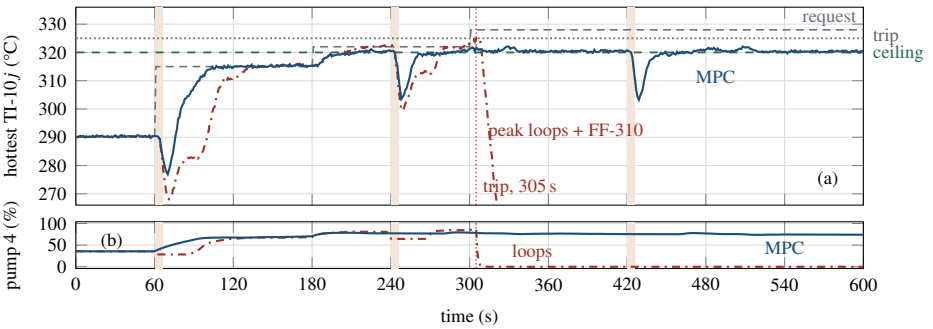

Figure 10.5 Case 9 as loaded under the MPC and, after Reset, under the peak loops with the feedforward, default seed, with the pulse openings shaded. (a) The hottest of TI-101 to TI-104 under the MPC (solid) and the loops (dash-dotted) against the requests of 315, 322 and 328 °C (gray, densely dashed), the ceiling at 320 °C (dashed) and the trip at 325 °C (dotted). The MPC stops at the ceiling, and the loops cross the trip at 5.08 min and leave the panel as the peaks collapse to 172.9 °C by 10 min. (b) Pump 4's stroke, in the same line styles, 76.6 under the MPC at 5.5 min against 87.3 % as the loops trip, and zero after it.

spends 679 zone-seconds within 5 K of the trip against 242 for the shipped loops, because it holds the peaks there while the loops trip and leave.

Case 5, with the trip at 330 °C and the ceiling at 325 °C, gives the same ranking. The MPC finishes at 31.92 % without a trip, where the loops trip with the feedforward and, as the case ships, decompose the plant without it, and again only the hindsight loop without the feedforward converts more, 0.27 points, at 0.4 K from the trip. In case 4, run under the peak controllers, the MPC's IAE is matched

Table 10.7 Challenge Configurations Over 240 min. MI and Density Columns Give the Pellets' Worst Deviation From the Grade After 60 min

Configuration	Conv. (%)	Margin (K)	MI dev.	ρ_{PE} dev.	Result
as shipped	27.72	35.31	0.638	0.00416	fail
suggested[a]	30.79	16.46	0.048	0.00057	pass
feedforward off	30.81	13.55	0.050	0.00062	fail
setpoints 312 °C	30.86	13.84	0.056	0.00052	fail
setpoints 305 °C	30.58	20.34	0.116	0.00152	fail
mid-ranging off	30.80	16.52	0.056	0.00029	pass
trim off	30.79	16.46	0.254	0.00057	pass
supervisor, 20 K	30.91	11.71	0.057	0.00074	fail
supervisor, 25 K	30.74	18.46	0.082	0.00112	fail
MPC on the pumps	30.80	16.46	0.048	0.00030	pass
MPC, else as shipped[b]	27.16	36.97	0.642	0.00419	fail

[a] peak loops at 310 °C, feedforward, mid-ranging at 55 %, trim. Other rows change the suggested configuration by one item, [b] the shipped one.

only by the loops without the feedforward (53.3, and 52.2 retuned in hindsight), which overshoot the pulse at 25 min to 302.0 °C (299.7 retuned), and it moves the pumps a sixth as far as the loops with the feedforward.

The pulse preview is credited with nothing, since switching it off moves the MPC's rows by at most 0.03 points and 0.1 K. Halving the physics step changes no verdict or ranking, and over four other noise seeds the loops trip at 5.08 to 5.10 min while the MPC never trips. After a trip's reset the pressure rule also holds the plan while PIC-300 restores the pressure, so the pumps stay at zero for 46 s after their release and then step by 11 %, the peaks topping out at 322.4 °C where the loops restart in 12 s and re-trip at 326.9 °C.

> ✦ **DID YOU KNOW? The Controller That Lives Near the Trip Never Trips**
>
> From the simulator: in case 9 of Table 10.6 the MPC holds the measured peaks within 5 K of the trip for 679 zone-seconds, against 242 for the shipped loops. Yet the loops trip at 5.08 min and convert 19.25 %, while the MPC never trips and converts 31.54 %.

10.6.10 The Four-Hour Challenge

The challenge runs 240 min with pulses every three minutes from 90 s, warms the supply water to 165 °C at 60 min, and brings a stronger initiator lot of 120 % at 150 min. Its target is a mean conversion of at least 30.5 %, a hottest gas never within 15 K of the line, pellets within 2.0 ± 0.3 in melt index and 0.924 ± 0.001 in density from 60 min, and no trip or decomposition, all scored on the truth. It ships with fixed strokes and the loops, feedforward, mid-ranging and trim off. The configurations measured against that target are collected in Table 10.7, and Figure 10.6 draws its conversion and margin columns against the target.

As shipped the plant stays 35.31 K from the line but converts only 27.72 %, and the pellets miss the grade by up to 0.638 in melt index and 0.00416 g/cm^3 in density. The suggested configuration converts 30.79 %, holds the melt index within 0.048 and the density within 0.00057, never trips, and comes closest to the line at 150.07 min, 4 s after the stronger lot arrives, leaving 16.46 K. That is 0.29 points above the conversion floor and 1.46 K more margin than required, and the run passes whether the configuration is set before Run or carried through Reset. The lot change decides the margin, and the feedforward is what lets the peaks sit at 310 °C. Without it the pulses take the margin to 13.55 K at 187.9 min, and at 312 °C the lot change takes it to 13.84 K. Over five other noise seeds every row keeps its verdict, and the suggested margin stays between 16.35 and 16.63 K. At 305 °C the margin

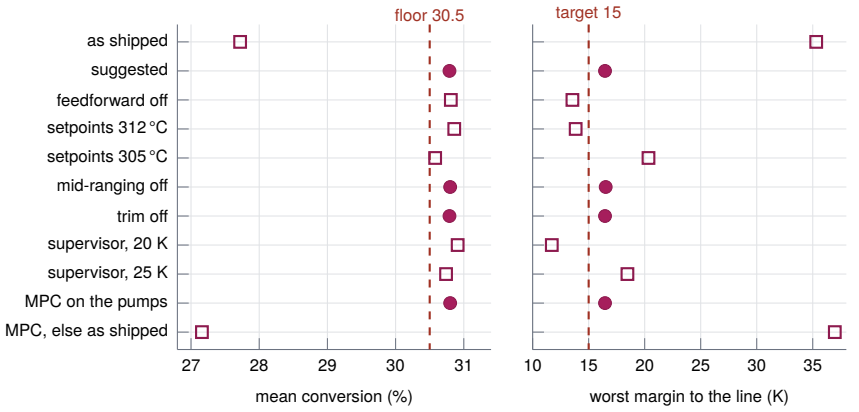

Figure 10.6 The conversion and margin columns of Table 10.7 as dots, filled for a configuration that passes the challenge and open for one that fails. The dashed lines are the conversion floor of 30.5 % and the 15 K the hottest gas must keep from the line. Notice that every passing configuration sits just right of both lines, and that the rows at 305 °C and with the supervisor at 25 K clear both lines yet fail, on density.

is ample but the density leaves the grade by 0.00152 g/cm^3, so the grade, not the line, sets the floor. Neither mid-ranging nor the trim decides the result, since the run passes without one, the other or both, although without the trim the melt index comes within 0.046 of its tolerance. The supervisor set at 20 K walks the peaks to 312 to 314 °C and leaves 11.71 K at the lot change, and at 25 K it keeps 18.46 K but lets the density leave the grade, so a margin chosen without the lot change in view fails either way.

The MPC in place of the suggested loops also passes with the same 16.46 K, because the lot change sets the margin 4 s after the lot arrives, before any controller acts, and it does so at a tenth of the loops' stroke travel, 1435 against 12 213 %. Of the unscored limits it reaches only the coolers' stops, for 14.2 % of the zone-scans against 23.4 % under the loops, and the pumps never. With every other setting as shipped it fails in the manner of the fixed strokes, holding 290 °C where they drift upward as the water warms and the stronger lot arrives.

10.7 Technology in Practice

What Measures It

Every peak loop of this chapter reads a row of thermocouples, and its measurement is the hottest of them. Section 10.3 gives each element a 1.5 s thermowell lag and 0.3 K of noise, so the reading hops from one element to the next as the peak moves, and Section 10.4 shows the gas running 10.9 K ahead of every thermocouple as the pressure recovers from a pulse. Figure 10.7 draws the element and its well, and the box below describes them.

TECHNOLOGY Sensor: Thermocouples in Thermowells Along the Tube

Temperature elements installed along the length of a tubular reactor follow the progress of the polymerization [18]. A thermocouple is the most rugged and wide-ranging contact sensor, and its accuracy depends on its type and its installation. Its pressure-tight thermowell lets it be replaced without depressurizing the process, but the well's mass adds a first-order lag, an air gap at the bottom of its bore adds a second, and a delayed reading can upset feedback. The most common failure is an open junction, called burnout, and the instrument's burnout mode decides whether a broken element then reads high or low [25]. On a trend a slow well reads a climbing peak late and low, as the model's 1.5 s lag does after a pulse.

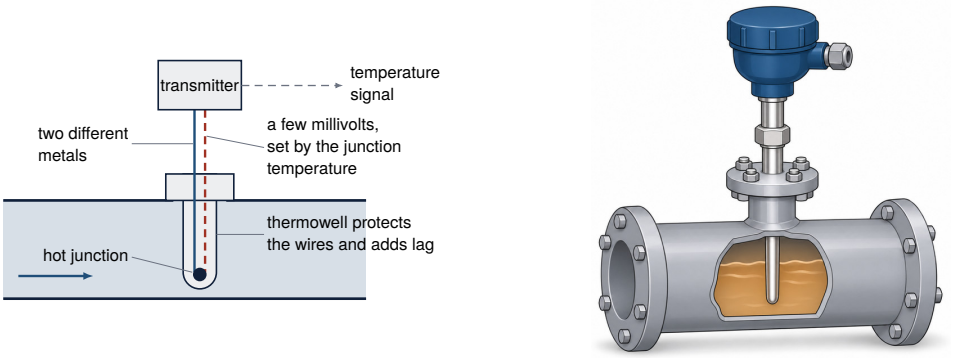

Figure 10.7 A thermocouple in a thermowell, the sensor behind TI-101 to TI-104. Notice the mass of metal between the gas and the junction: it is the 1.5 s lag of the model, and it is why the hottest gas can run ahead of every reading while a peak climbs.

What Moves It

Each peak has one final element, the initiator pump of its zone. The cooler valves move the jackets, the let-down valve holds the pressure, and the pulse valve opens on a timetable rather than on a controller's demand. The model gives every pump a stroke that delivers in proportion and follows its command through a 1.0 s lag, with no stiction.

> **TECHNOLOGY Final Element: High-Pressure Initiator Pumps**
>
> Where a tube runs on peroxides, the initiator added at each injection point is set by the speed of high-pressure pumps, and the pressure calls for reciprocating machines. The peroxide travels dissolved in a hydrocarbon solvent [18]. A positive-displacement pump discharges a fixed volume with every revolution, so at constant speed it delivers a fixed flow whatever the discharge pressure, and a blocked discharge locks it. A metering pump spans its flow from zero to a maximum set on the pump as its control signal goes from 0 to 100 %, and a plunger type delivers that flow as discrete pulses rather than a continuous stream [25]. The chapter's pumps, whose delivery is their stroke times a fixed maximum after a 1.0 s lag, are that behavior without the pulses.

Where the Controller Runs

Figure 10.8 draws one peak loop of Section 10.5 with hardware in place of the schematic's symbols. A distributed control system holds loops like these at a supervisory level that sends set points to field controllers, with loops nested or cascading and cycle times from milliseconds to minutes [26]. The hottest of a row of transmitters reaches the controller through a high-select function block, which passes only the greatest of its inputs. A textbook exothermic reactor holds its hottest measured temperature at setpoint in exactly this way [25]. In the app the thirteen PI loops, the feedforward and the trip logic execute every 1 s, four physics steps, the feedforward reads the pulse scheduler rather than the valve, and the MPC of case 9 plans every 2 s in place of the four peak blocks. Where the trip belongs is the subject of the In practice box below.

What Is Modernizing It

Model predictive control is accepted in industry for its multivariable control and its explicit handling of constraints, and for the tubular reactor a tracking nonlinear predictive controller was shown to keep the reactor stable through persistent cycles of fouling and cleaning while holding the melt index tightly [20, 27]. A 2003 survey counted polymers among the areas with the most nonlinear predictive control applications and expected their number to grow where linear control had proven difficult, such

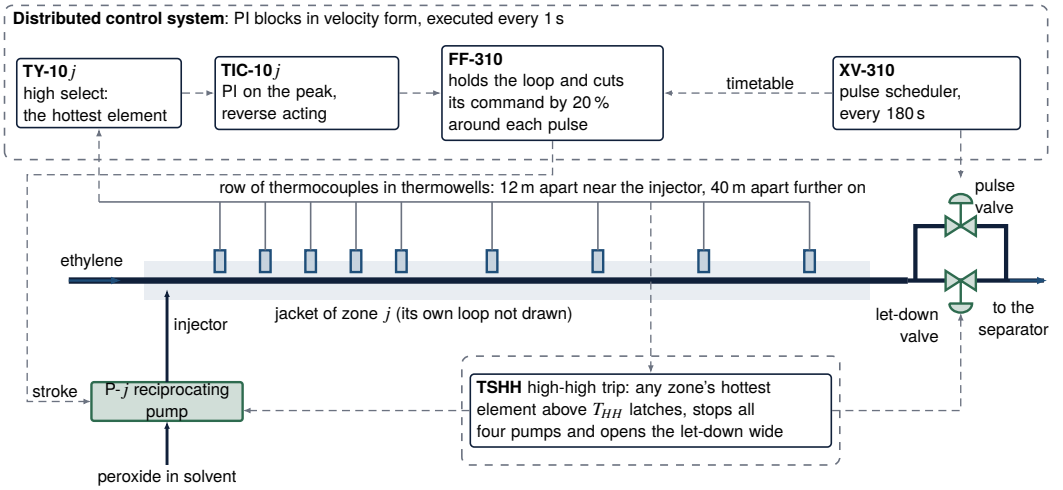

Figure 10.8 One peak loop and the trip with their hardware. Blue marks are sensing devices, the row of thermocouples in thermowells along zone j. Green symbols are final elements: the zone's reciprocating initiator pump, the let-down valve and the pulse valve at the outlet. Dashed arrows are signals. Notice that the controller never sees a single temperature: a high select passes the hottest element, and the feedforward stands between the peak controller and the pump around each scheduled pulse. The trip reads the same row of elements. The other three zones, the jacket loops with mid-ranging, PIC-300, the melt-index trim, the supervisor and the MPC are left out. Not to scale, and not a piping diagram.

as polymerization [13]. Plant data show the wall fouling, so the thickness of the layer can be inferred from the measured heat-transfer coefficient and a heat-transfer model, the basis of soft sensors for it [27]. The MPC of case 9 is the linear, identified form of the predictive idea, with the margin of this chapter written into its plan as a ceiling. How far a wall estimate can be trusted is the subject of the Open challenge box in Section 10.8.

> **IN PRACTICE** **A Trip With Its Own Sensors and a Relief Path**
>
> In service a high-high trip such as TSHH is a safety instrumented function. Its sensors may be transmitters separate from the regulatory control system, its logic solver is separate as well, and one piece of hardware should not both regulate and shut down, since a single fault would then disable both [25]. The standard for such systems is ANSI/ISA-84.00.01-2004, *Functional Safety: Safety Instrumented Systems for the Process Industry Sector*. Behind the trip, the reactor is protected by pressure relief devices that release its contents at once if ethylene runs away, and they vent to the air because the response must be fast [18]. In the chapter the trip reads the same thermocouples as the peak loops.

10.8 Limitations

The kinetics, the decomposition threshold and dwell, the fouling law and the quality correlations are illustrative or teaching values fitted to no plant, and nothing is modeled after a decomposition. The instruments misread only through lag, sampling and the spacing of the thermocouples, and the analyzers never fail. The supervisor is offered untuned and fails the challenge at its default margin, and case 2 runs without pulses, so its wall fouls over a long run. The weak lot of case 6 changes its strength in one step at a declared minute, which no plant would announce, so that a gain change can be told from everything else the tube does.

The MPC's model is four identified paths and a set of offsets. It leaves out the parcels, the kinetics, the margin, the decomposition and the fouling, which reach it only through the bias, the shortening of

the dead time as the stroke rises, the gains of the cross and jacket paths away from the design point, the lot, which arrives as an unmeasured gain change for the bias to carry, and the plant below 15 % of stroke, where the gain schedule is clamped. Its cross paths and jackets are offsets rather than states the plan can move, so a plan's own effect on the next zone is left to the next plan. The scatter of about 2 K between the parcels a thermocouple samples is not modeled and is the reason for the 5 K margin. After a trip's reset the pressure rule keeps the pumps at zero for 46 s after their release, and the one comparison the MPC loses is to a tuning chosen in hindsight that runs 0.4 K from the trip.

OPEN CHALLENGE Seeing the Wall Foul From Outside the Tube

The polymer layer on the inside wall of an LDPE tube is difficult to assess and has been described in only a few technical papers. Its prediction is not systematically formulated in practice, and when the product leaves its quality range the process is adjusted by hand, which is slow. Plants strip the layer online by pressure cleaning or thermal shock, and the reactor has to stay stable through every cycle of fouling and cleaning. A 2022 study pairs a soft sensor of the layer's thickness with a predictive controller that starts the cleaning when the layer passes a limit, tested on a process simulation [27]. Progress would be a wall estimate proved on a running tube.

10.9 Conclusion

A tubular LDPE reactor is run on a margin that its excursions spend, and the simulator of this chapter measures that margin on the truth while its loops see thermocouples, a pressure gauge and delayed analyzers. The tube is 70 parcels of 28 kg riding a 1500 m flow at 2500 bar, and at four 290 °C peaks it converts 27.51 % with 44.7 K to the teaching line at 335 °C. A peak's gain falls from 1.9 to 0.6 K per percent of stroke as it moves toward its injector, and each zone's tail is the next zone's feed.

The scheduled pulse winds the peak loops up and costs 9.7 K of overshoot, which a feedforward on the schedule reduces to 2.2 K, raising the worst margin over two pulses from 33.1 to 41.9 K. A trip set 5 K below the line fires 1.5 s before a decomposition that one scan earlier it would have prevented. Over four hours the suggested configuration converts 30.79 % with 16.46 K to spare against 27.72 % as shipped. An optional model predictive controller that carries the ceiling in its plan takes the request that trips the loops at 5.08 min and holds the four peaks at 320 °C, converting 31.54 % with 12.8 K to spare against 31.09 % for the loops with the same requests clipped by hand.

The model is plug flow in well-mixed parcels with illustrative kinetics and a declared decomposition. A fuller model would add axial dispersion, a temperature-dependent density with the expansion cooling of a pulse, a molecular-weight distribution, analyzer faults and a supervisor that anticipates the lot change. As it stands, the simulator is a place to learn to size a safety margin against the excursions a reactor actually makes.

KEY POINTS

- Each injection of initiator lights its own hot peak, and the gas it heats becomes the next zone's feed.
- A peak's gain falls from 1.9 to 0.6 K per percent of stroke as the peak moves toward its injector.
- The scheduled pulse winds the peak loops up and costs 9.7 K of overshoot, which a feedforward on the schedule reduces to 2.2 K.
- A trip set 5 K below the line fires 1.5 s before a decomposition it does not prevent, because it stops the pumps and not the initiator already in the gas.
- Excursions set the margin, and the setpoint is what they leave.

Where else this applies. A slow handle that keeps a fast one in range is also the mid-ranging of Volume 4, Chapter 10, and a guard that a fast excursion can cross is also the surge override of Volume 4, Chapter 2.

References

[1] C. Kiparissides, G. Verros, and J. F. MacGregor, "Mathematical modeling, optimization, and quality control of high-pressure ethylene polymerization reactors," *Journal of Macromolecular Science, Part C: Polymer Reviews*, vol. 33, no. 4, pp. 437–527, 1993.

[2] R. C. M. Zabisky, W.-M. Chan, P. E. Gloor, and A. E. Hamielec, "A kinetic model for olefin polymerization in high-pressure tubular reactors: a review and update," *Polymer*, vol. 33, no. 11, pp. 2243–2262, 1992.

[3] R. Geyer, J. R. Jambeck, and K. L. Law, "Production, use, and fate of all plastics ever made," *Science Advances*, vol. 3, no. 7, p. e1700782, 2017.

[4] J. Albert and G. Luft, "Thermal decomposition of ethylene–comonomer mixtures under high pressure," *AIChE Journal*, vol. 45, no. 10, pp. 2214–2222, 1999.

[5] O. Bilous and N. R. Amundson, "Chemical reactor stability and sensitivity: II. Effect of parameters on sensitivity of empty tubular reactors," *AIChE Journal*, vol. 2, no. 1, pp. 117–126, 1956.

[6] C. Kiparissides, "On the elucidation of polymer fouling mechanisms and ethylene decomposition in high-pressure LDPE tubular reactors," *Macromolecular Reaction Engineering*, vol. 16, no. 6, p. 2200023, 2022.

[7] S. X. Zhang, N. K. Read, and W. H. Ray, "Runaway phenomena in low-density polyethylene autoclave reactors," *AIChE Journal*, vol. 42, no. 10, pp. 2911–2925, 1996.

[8] B. J. Allison and A. J. Isaksson, "Design and performance of mid-ranging controllers," *Journal of Process Control*, vol. 8, no. 5–6, pp. 469–474, 1998.

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

[10] C. H. Chen, J. G. Vermeychuk, J. A. Howell, and P. Ehrlich, "Computer model for tubular high-pressure polyethylene reactors," *AIChE Journal*, vol. 22, no. 3, pp. 463–471, 1976.

[11] A. Brandolin, M. H. Lacunza, P. E. Ugrin, and N. J. Capiati, "High pressure polymerization of ethylene. An improved mathematical model for industrial tubular reactors," *Polymer Reaction Engineering*, vol. 4, no. 4, pp. 193–241, 1996.

[12] M. Asteasuain, S. Pereda, M. H. Lacunza, P. E. Ugrin, and A. Brandolin, "Industrial high pressure ethylene polymerization initiated by peroxide mixtures: A reduced mathematical model for parameter adjustment," *Polymer Engineering & Science*, vol. 41, no. 5, pp. 711–726, 2001.

[13] 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.

[14] R. O. Gibson, E. W. Fawcett, and M. W. Perrin, "The effect of pressure on reactions in solution I—Sodium ethoxide and ethyl iodide to 3000 kg/cm^2 II—Pyridine and ethyl iodide to 8500 kg/cm^2," *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 150, no. 869, pp. 223–240, 1935.

[15] E. W. Jenkins, "The discovery of polythene—a case history," *School Science and Mathematics*, vol. 68, no. 2, pp. 135–140, 1968.

[16] American Chemical Society, "Discovery of polypropylene and the development of a new high-density polyethylene," National Historic Chemical Landmark booklet, Office of Communications, Washington, DC, 1999, designated at Phillips Petroleum Company, Bartlesville, Oklahoma, 12 November 1999; acs.org/education/whatischemistry/landmarks/polypropylene.html, read November 2026.

[17] K. Ziegler, "Consequences and development of an invention," Nobel Lecture, Stockholm, 12 December 1963, 1963, nobelprize.org/uploads/2018/06/ziegler-lecture.pdf, read November 2026.

[18] European IPPC Bureau, "Reference document on best available techniques in the production of polymers," European Commission, Joint Research Centre, Institute for Prospective Technological Studies, Seville, Tech. Rep., Aug. 2007.

[19] M. H. Lacunza, P. E. Ugrin, A. Brandolin, and N. J. Capiati, "Heat transfer coefficient in a high pressure tubular reactor for ethylene polymerization," *Polymer Engineering & Science*, vol. 38, no. 6, pp. 992–1013, 1998.

[20] V. M. Zavala and L. T. Biegler, "Optimization-based strategies for the operation of low-density polyethylene tubular reactors: nonlinear model predictive control," *Computers & Chemical Engineering*, vol. 33, no. 10, pp. 1735–1746, 2009.

[21] M. Häfele, A. Kienle, M. Boll, and C.-U. Schmidt, "Modeling and analysis of a plant for the production of low density polyethylene," *Computers & Chemical Engineering*, vol. 31, no. 2, pp. 51–65, 2006.

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

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

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

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

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

[27] F. S. Rohman, M. R. Othman, D. Muhammad, A. Azmi, I. Idris, R. A. Ilyas, S. E. Elkhatif, and M. N. Murat, "Nonlinear control of fouling in polyethylene reactors," *ACS Omega*, vol. 7, no. 44, pp. 39 648–39 661, 2022.

10.10 Exercises

10.1 ⌕ Concept check · 2 min

In case 5 the peak setpoints are 328 °C and the high-high trip is set at 330 °C, 5 K below the decomposition line. After the pulse at 7 min the trip fires on zone 1, and 1.5 s later gas further down zone 1 decomposes. Why does the trip not prevent it?

(a) The trip stops the pumps, but the initiator already in the gas still reacts as it is carried to the peak

(b) The trip opens the let-down valve wide, and the falling pressure speeds up the polymerization in the peak

(c) The thermocouple noise of 0.3 K hides the approach of the peak until the scan that fires the trip

(d) The trip compares only the last zone's thermocouple with its setting, so zone 1 is not protected at all

10.2 ⌕ Concept check · 2 min

Between 10 and 15 % of stroke, zone 2's steady peak rises by about 1.9 K per percent of pump stroke, and between 45 and 55 % by only 0.6 K per percent. Which explanation is right?

(a) The pump's delivery flattens near full stroke, so each extra percent injects less peroxide

(b) A larger dose heats the gas faster, so the rest of the dose decomposes sooner and the front burns out earlier

(c) The decomposition of ethylene removes heat from the gas once the peak passes about 300 °C

(d) The pressure loop lowers the reactor pressure as the conversion rises, which slows the polymerization

10.3 Concept check · 2 min

The pulse feedforward FF-310 is on, and the pulse valve fails to open at its next scheduled time. What happens to the peaks?

(a) Nothing, because the feedforward watches the valve position and stays idle when the valve does not move

(b) They overshoot their setpoints, because the loops wind up against a pressure dip that never comes

(c) They dip with no pulse at all, because the cut is keyed to the timetable and still arrives on schedule

(d) The trip fires, because the pressure stays high while the pumps run at their full strokes

10.4 Theory and calculation · 5 min

The heat of polymerization is $-\Delta H_p = 93.6\,\mathrm{kJ/mol}$, the molar mass of ethylene $M_w = 28.05\,\mathrm{g/mol}$ and the heat capacity of the gas $c_p = 2600\,\mathrm{J/(kg\,K)}$. The tube is fed 64.8 t/h of ethylene at 165 °C and converts 27.51 % of it.

(a) Find the adiabatic temperature rise per percent of conversion. *Checkpoint:* 12.83 K.

(b) Find the outlet temperature if the whole conversion happened in one adiabatic front, and compare it with the decomposition line at 335 °C. *Checkpoint:* A rise of 353 K, to about 518 °C.

(c) Find the polymer production rate and how many minutes of production a product hold-up of 2500 kg holds. *Checkpoint:* 17.8 t/h and 8.4 min.

10.5 Theory and calculation · 6 min

The tube is $L = 1500\,\mathrm{m}$ long with a bore of 55 mm and holds $N = 70$ parcels. The gas density is $\rho_0 = 550\,\mathrm{kg/m^3}$ at $P_{ref} = 2500\,\mathrm{bar}$, its compressibility $\kappa = 6.3 \times 10^{-5}\,\mathrm{bar^{-1}}$, and the feed $F_{in} = 18.0\,\mathrm{kg/s}$. The separator is at 250 bar, and the let-down passes the design flow at P_{ref} with an opening of 0.5.

(a) Find the mass and length of one parcel, the gas velocity and the residence time. *Checkpoint:* 28.0 kg, 21.4 m, 13.78 m/s and 108.9 s.

(b) Find the capacity $C = \rho_0 A L \kappa$ and the valve coefficient C_v of Equation (10.4). *Checkpoint:* 0.1235 kg/bar and 0.759 kg/(s bar$^{1/2}$).

(c) Linearize Equation (10.4) about the design point with the let-down held, and find the time constant of the pressure. *Checkpoint:* 30.9 s.

10.6 Theory and calculation · 6 min

Ethylene decomposition in the model has $k_{dec,0} = 5 \times 10^{-3}\,\mathrm{s^{-1}}$ at $T_D = 335\,°\mathrm{C}$, an activation energy $E_{dec} = 450\,\mathrm{kJ/mol}$ and a heat of decomposition $-\Delta H_{dec} = 120\,\mathrm{kJ/mol}$. Take $M_w = 28.05\,\mathrm{g/mol}$, $c_p = 2600\,\mathrm{J/(kg\,K)}$, $R = 8.314\,\mathrm{J/(mol\,K)}$ and a spent parcel with $1 - x - x_d = 0.73$.

(a) Find the rate at which decomposition heats the parcel at 335 °C. *Checkpoint:* 6.0 K/s.

(b) Find the same rate at 320 °C. *Checkpoint:* 0.63 K/s, about 9.5 times slower.

(c) A clean wall cools a parcel with a time constant of 24.6 s. With the jacket water at 181 °C, find the cooling rate of a parcel at 335 °C and say what happens to a parcel held just above the line. *Checkpoint:* About 6.3 K/s.

10.7 App challenge · 9 min

Load scenario 3, “3 · The pulse valve”, at the default seed. The four peak loops hold 290 °C and the first pulse is scheduled at 30 s. Run at 1 min/s until the clock passes 5 min, pause and export the CSV. Temperatures and margins within 0.2 K of the key count as correct.

(a) With Pulse feedforward Off, find the highest value of the columns `Tpk1_meas_C` to `Tpk4_meas_C` between 0.5 and 3.4 min, and read the Worst decomposition margin tile at the end of the run. *Checkpoint:* 299.7 °C and 33.1 K.

(b) Pause, press Reset, set Pulse feedforward to On on the XV-310 card before Run, and repeat the run and the two readings. *Checkpoint:* 292.2 °C and 41.9 K.

(c) Take the supervisor's rule, Equation (10.10), with $\hat{e}_j$ equal to the overshoot above 290 °C that you found, $\Delta_{set} = 20$ K and $\sigma_{TC} = 0.3$ K. Find the highest peak setpoint it allows without and with the feedforward. *Checkpoint:* 304.7 °C and 312.2 °C.

10.8 App challenge · 8 min

Load scenario 5, "5 · Conversion pushed toward the line", at the default seed. The peak setpoints rise to 315, 322 and 328 °C at 1, 3 and 5 min, the pulses come every three minutes from 1 min without feedforward, and the trip is at 330 °C. Run at 1 min/s. Margins within 0.1 K and times within 0.05 min of the key count as correct.

(a) Run as shipped until the run stops. Read the Plant and Worst decomposition margin tiles. *Checkpoint:* Decomposed at 7.4 min, −1.0 K.

(b) Press Reset, set the Trip temperature (hottest thermocouple) to 325 °C on the Control architecture card, and run until the clock passes 10 min. Read the same two tiles, and find the first row of the exported CSV with `tripped` equal to 1. *Checkpoint:* Tripped, 8.7 K, at 4.30 min.

(c) The trip at 330 °C sat 5 K below the line. Explain why it was not enough and what the 325 °C setting buys. *Checkpoint:* The trip's dead time and the dosed gas.

10.9 App challenge · 10 min

Load the graded challenge, "Challenge · the most conversion inside the margin", at the default seed. It runs 240 min with pulses every three minutes, warmer jacket water at 60 min and a 120 % initiator lot at 150 min. Make every change before Run, then run at 60 min/s until the clock reaches 240 min. Conversions within 0.02 points and margins within 0.1 K of the key count as correct.

(a) Set Initiator pumps to Peak loops, all four peak setpoints to 310.0 °C, Pulse feedforward On, Mid-ranging (VPC) On with a Target pump stroke of 55 %, and Melt-index trim On. Read the Mean conversion, Worst decomposition margin and Challenge target tiles. *Checkpoint:* 30.79 %, 16.5 K, met.

(b) Press Reset, set Pulse feedforward Off, change nothing else, and run again. *Checkpoint:* 30.81 %, 13.5 K, not met.

(c) Press Reset, set Pulse feedforward On again and all four peak setpoints to 312.0 °C, and run again. From the CSV column `metric_worst_margin_K`, say when the worst margin is set in part (a)'s run and in this one. *Checkpoint:* 30.86 %, 13.8 K, not met. Both just after the lot change at 150 min.

10.10 Application area · 3 min

Tubular LDPE reactors open a pulse valve every few minutes on purpose.

(a) What does a pulse do for the wall, and what happens to the heat transfer when the pulses are sparse? Use case 4. *Checkpoint:* It strips the polymer layer. U falls to 0.68 of clean in 24.9 min without a pulse.

(b) Name two costs of pulsing, and say how the pulse period should be chosen. *Checkpoint:* A dip and an overshoot each time, and a pressure upset.

10.11 Application area · 3 min

The model draws the decomposition line at 335 °C with a dwell of 1.0 s, and every loop sees the peaks through thermocouples.

(a) Where should the decomposition limit and the trip setting of a real reactor come from? *Checkpoint:* Licensor data, plant history, measured boundaries.

(b) Give two reasons the hottest thermocouple can read well below the hottest gas, with a number from the chapter for one of them. *Checkpoint:* Spacing and lag. Up to 10.9 K after a pulse.

10.12 Review the agent · 5 min

An AI assistant is asked how close to the decomposition line the peaks can run. It answers: “In case 2 the zone 2 loop takes a step to 300 °C and its highest reading over the next five minutes is only 300.8 °C, so the loops hold the peaks within about 1 K. Set all four peaks to 328 °C and the trip to 330 °C. The peaks then stay 7 K below the line, the trip covers the last kelvin of loop error, and the conversion gains several points.” Find the flaw and say how you would check it in the app.

CHAPTER 11

A Butyl Rubber Reactor Simulator for Fouling, Actuator Margin and Wash Policy

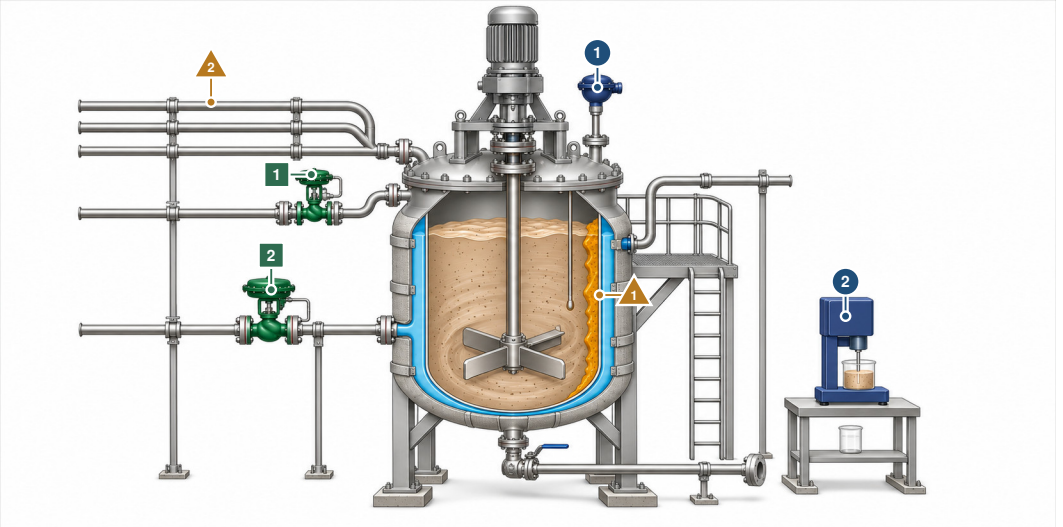

The controller moves 1 catalyst valve 2 refrigerant valve that sets the jacket temperature

The controller reads 1 reactor temperature 2 Mooney viscosity of the rubber, a laboratory test every two hours

What disturbs it 1 polymer coating the cold wall 2 isobutylene feed rate

Notice the amber layer between the slurry and the cold jacket: it insulates the wall, so the jacket must run ever colder to remove the same heat until it reaches its floor. Illustration, not to scale.

IN PLAIN TERMS A Cold Reactor That Coats Its Own Walls

The main idea. Butyl rubber is made in a stirred tank kept colder than any freezer. As it forms, some rubber sticks to the cold walls and blankets them, like frost on a freezer coil.

The challenge. The controls hide this by running the cooling ever colder, until it can go no colder and the plant must stop to wash. Running colder slows the coating but spends the cooling sooner, so how cold to run and when to wash are hard choices.

Why it matters. The warning to watch is how hard the cooling works, not the tank temperature the controls keep flat.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the two History boxes and the Did you know? box in Section 11.2, the two Try it boxes and the second Did you know? box in Section 11.6, and Key points at the end of the Conclusion. No equations.
Study	50 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 11.6 with Figure 11.3, Section 11.7, and Section 11.8 with its Open challenge boxes.

The control question A butyl rubber reactor fouls its own cooling surface, and its temperature loop holds the setpoint while the jacket walks toward a refrigeration floor it cannot pass. Which trend shows how much of the run is left, what reactor temperature gives the most saleable rubber per elapsed hour, and what should decide when to wash?

Learning objectives

- Compute the design duty, the jacket temperature and the cooling margin, and read the run length from the closed-form countdown.
- Explain why a flat controlled variable and a drifting controller output together mean that the loop is spending a finite resource.
- Explain why colder operation shortens the run, and find the grade window that bounds the temperature optimum.
- Compare a fixed wash schedule with a wash on the cooling margin by the yield per elapsed hour, and name the measurement that ends a run.
- Judge when a constraint on the cooling margin, by override or by predictive control, pays on a reactor whose end of run creeps.

Prerequisites Sections 1.4.2 (run length against a limit), 1.4.4 (capacity set by heat removal), 1.6.2 (velocity-form PID), 1.4.6 (laboratory bias update) and 1.6.3 (the predictive option). Chapter 2 for the same countdown in coke.

The app https://apmonitor.com/apps/butyl-rubber/

See also Volume 4, Chapter 2 for a protective override and Volume 4, Chapter 8 for a run against a cash floor. Volume 3, Chapter 2 for matching a cascade to where the disturbance enters and Volume 3, Chapter 3 for the speed of an inner loop. Volume 8, Chapter 9 for corrections from late, sampled results.

11.1 Introduction

A control course teaches that a well tuned loop holding its setpoint is a loop doing its job. On the reactor modeled here it is a loop counting down. Butyl rubber is polymerized at −95 °C in a methyl chloride slurry, and the polymer deposits on every cooled surface from the first minute of a run. The overall heat transfer coefficient falls from 2400 to 976 W/(m^2 K) over forty-eight hours. A temperature controller absorbs that almost perfectly. Measured, the reactor temperature never leaves its setpoint by more than 0.01 °C while the jacket walks from −99.4 °C to −105.9 °C at hour 48 and reaches its −106 °C floor at hour 48.9, the temperature difference the duty requires grows from 4.4 K to 11.0 K, and the fraction of the jacket's capability still spare falls from 0.598 at hour zero to zero at hour 49.0. The controlled variable is a flat line for two days and then the loop has nothing left.

The misconception the simulator corrects is that the controlled variable is where the information is. On any plant with a slowly dying gain, the controlled variable stays put by construction and the manipulated variable is the diagnostic. Heat exchanger fouling, catalyst decay, valve wear and filter blinding all have this shape, and in every case the trend worth an alarm is the controller output rather than its input. The coke in a cracking furnace coil is the same countdown, read in tube metal temperature. Section 1.4.2 gives its run length in closed form, and Chapter 2 works it on a four-pass furnace. The class is large enough that scheduling cleaning against a fouling model is a research subject in its own right [1], and the reactor here adds the feature that makes it a control problem rather than a maintenance one: the actuator has a hard floor, so the end of the run is a saturation event, abrupt in the controller even though its cause was gradual.

Two further results in this application are not the intuitive ones and both are measured. Colder operation gives a higher molecular weight and slower fouling and a shorter run, because the refrigeration floor does not move when the setpoint does, and the measured optimum in tonnes of rubber per hour of elapsed time is near −94 °C rather than at the coldest reachable temperature. And a wash triggered on the cooling margin, the obvious feedback replacement for the fixed schedule most plants

run, yields less than the schedule when the feed fouls faster than the schedule assumed, because the margin runs out about fifteen hours before the product leaves grade.

This chapter works with an interactive butyl rubber reactor whose fouling rate, temperature acceleration and refrigeration floor are hidden, so that the countdown must be read from the actuator, the temperature optimum found by experiment, and a wash policy judged by the yield it delivers rather than by whether it uses feedback. The polymerization kinetics, the fouling model and cascade control are standard, and the app exports its complete time history (Section 1.6.4). Section 11.2 gives the history of the product and of its control. Section 11.3 writes out the slurry reactor with its fouling dynamics and its instrument chain in the detail needed to rebuild it. Section 11.4 derives the run length in closed form and the behavior at the refrigeration floor. Section 11.5 describes the four control architectures, the two wash policies and the predictive option. Section 11.6 reports the scenario sequence, whose temperature and scheduling conclusions contradict the obvious answers and are measured rather than argued, and compares the predictive controller with the shipped structures. Section 11.7 names the hardware and the practice behind the model, and Sections 11.8 to 11.10 give the limitations, the lessons that transfer to other plants and the conclusions.

11.2 Background

Butyl rubber is a copolymer of isobutylene with a small amount of isoprene, made by cationic polymerization at very low temperature, and the History box below tells how it was found [2, 3]. The account of its development given by one of its inventors thirty years later [4] makes clear that the low temperature was not a design preference but a requirement: the chain transfer reaction that limits molecular weight is strongly temperature dependent, so useful molecular weight is only reachable below about −90 °C. The diluent is methyl chloride because the polymer is insoluble in it and precipitates as a slurry rather than thickening the reactor [5].

> ⧖ HISTORY **A Rubber With Just Enough Double Bonds, 1937 to 1944**
>
> In 1941 a team of five chemists described butyl rubber as a new hydrocarbon product. Their work on polymers of simple olefins had then run for ten years. Long chains made from the butylenes behaved much like rubber, but a fully saturated polymer could not be vulcanized in the usual way [2]. The answer, in an application first filed in December 1937 by Robert M. Thomas and William J. Sparks, was to copolymerize isobutylene with a small amount of a diolefin such as isoprene, at low temperature, with aluminum chloride dissolved in ethyl or methyl chloride as the catalyst. The polymer kept a low, controlled unsaturation that sulfur could cure, and the patent was granted in 1944 [6]. Because polymers of isobutylene are so impermeable, they are widely used in inner tubes and in layers of tires [7].

That slurry is the source of the control problem. Polymer particles adhere to the cooled surfaces, and the deposit is an insulator between the reaction and the refrigerant. Butyl plants have therefore always run in cycles, with a scheduled shutdown for hydroblasting, and the run length between washes is one of the numbers a plant is managed by.

HISTORY Washing the Reactor Clean, 1946

A patent application filed in 1946 describes the continuous butyl reactor: a central draft tube with an agitator, and return tubes around it, all cooled by liquid ethylene. A run proceeded smoothly for 10 to 150 hours while a layer of polymer built up on the vessel, the stirrer and the shaft, until the layer interfered seriously with the transfer of heat. The tough, rubbery coating was very hard to remove by hand, and light naphtha dissolved too little of it. The inventors drained the refrigerant, warmed the reactor with ethylene gas and filled it with warm liquid isobutylene. That softened the polymer and loosened its hold enough for the stirrer to break it away [8].

Control of the reactor developed in the order the instruments allowed. The first requirement is temperature, because molecular weight follows it directly, represented here by −1.40 Mooney units per degree. A jacket temperature loop is straightforward and the refrigeration system is the constraint underneath it, which is why split range and model-based algorithms for jacketed reactor temperature remain a subject even on much simpler services [9]. Product quality entered as a laboratory Mooney viscosity with a cycle measured in hours, which supports a cascade that corrects over shifts rather than minutes. Inferential estimation of polymer properties from reactor conditions became common later and carries the maintenance problems any soft sensor has [10]. Constraint handling above all of it, which is what decides whether a conversion loop should be allowed to recover production on a reactor that is nearly out of cooling, is the layer a modern plant delegates to a multivariable controller [11].

DID YOU KNOW? Seven Percent of the Synthetic Rubber

In the world statistics for 2002, isobutylene-isoprene rubber, butyl and halobutyl together, made up 7 % of the production volume of the major synthetic rubbers, the same share as polyisoprene. Emulsion styrene-butadiene rubber led with 28 % and polybutadiene followed with 24 % [12].

The safety literature on polymerization reactors concentrates on thermal runaway, and the mechanisms are well catalogued [13]. This reactor is not a runaway problem, and the distinction is worth making explicitly. Its conversion is limited by catalyst rather than by temperature, so the heat release does not accelerate with temperature the way an adiabatic exotherm does. What accelerates with temperature here is the fouling, at 0.1386 per degree, so a warming reactor fouls faster and reaches its stop sooner. The failure is economic rather than energetic, and it is the more common of the two on operating plants.

11.3 Process Model

Reactor R-201 is a 30 m^3 slurry reactor with 190 m^2 of cooled surface, charged with 6.0 t/h of isobutylene, 18 t/h of methyl chloride and 0.12 t/h of isoprene. At 85 % conversion it makes 5.10 t/h of polymer at 21.1 % slurry solids, releasing 2.018 MW. Figure 11.1 shows the arrangement and every symbol used in the equations, and Section 1.4.4 places the reactor among the polymer reactors whose capacity is set by heat removal.

11.3.1 Reaction and Energy Balance

Conversion is set by catalyst and by temperature, Equation (11.1), and the heat released follows the polymer made, Equation (11.2).

$$X^{*} = \frac{k\tau}{1 + k\tau}, \qquad k = k_p \, a \, F_{cat} \, e^{\,b(T - T_{ref})} \tag{11.1}$$

$$Q = \Delta H_p \, F_{IB} \, X + \dot{m} \, c_p \, (T_{feed} - T) \tag{11.2}$$

where a is the catalyst activity, $k_p = 1.35$ per hour per kg/h of catalyst, τ is the residence time of 0.871 h, $b = -0.012/^{\circ}\mathrm{C}$ and $T_{ref} = -95\,^{\circ}\mathrm{C}$. The conversion X approaches X^{*} with the residence

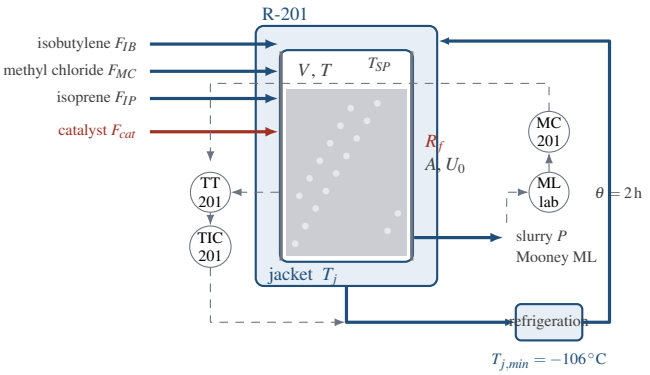

Figure 11.1 Butyl rubber slurry reactor R-201. Isobutylene F_{IB}, methyl chloride F_{MC} and isoprene F_{IP} enter with catalyst F_{cat} and polymerize at temperature T in volume V. Heat leaves through cooled surface A of clean coefficient U_0, carrying a growing fouling resistance R_f, to a jacket at T_j whose refrigeration cannot pass $T_{j,min}$. The slurry product P is characterized by its Mooney viscosity ML, reported by a laboratory on a cycle θ. TIC-201 manipulates the jacket temperature and MC-201 trims the reactor setpoint T_{SP} from the laboratory result.

Table 11.1 Classification of Variables at the Start of a Clean Run

Symbol	Role	Nominal	Range
T (TT-201)	CV, continuous	−95 °C	−102 to −86 °C SP
ML	CV, laboratory	51.0	46–56 spec
X conversion	CV	0.85	0.50–0.95 SP
T_j	MV	−99.42 °C	$\geq -106\,^{\circ}\mathrm{C}$
F_{cat}	MV	4.82 kg/h	0–8 kg/h
F_{IB}	DV, measured	6.0 t/h	3–9 t/h
catalyst activity	DV, unmeasured	1×	0.4–1.3×
fouling tendency	DV, unmeasured	1×	0.4–3×
R_f	state	0	growing
U	derived	2400 W/(m^2 K)	falling
margin	derived	0.598	0 ends the run

time as its time constant, and the catalyst reaches the reactor through a 0.08 h lag. The heat of polymerization is $\Delta H_p = 930\,\mathrm{kJ/kg}$, and the second term chills $\dot{m} = 24.1\,\mathrm{t/h}$ of feed from $T_{feed} = -40\,^{\circ}\mathrm{C}$ at $c_p = 1.9\,\mathrm{kJ/(kg\,K)}$, so the design duty is 1.32 MW of reaction and 0.70 MW of feed chilling. The temperature dependence of conversion is deliberately weak, which is what makes this a fouling problem rather than a runaway problem.

11.3.2 Fouling, and the Actuator Margin

The fouling resistance accumulates and accelerates with temperature, Equation (11.3), and the overall coefficient follows from it, Equation (11.4).

$$\frac{\mathrm{d}R_f}{\mathrm{d}t} = k_f \frac{P}{P_0}\,\phi\,\exp\!\left[c_f (T - T_{ref})\right] \tag{11.3}$$

$$U = \frac{1}{1/U_0 + R_f} \tag{11.4}$$

with $k_f = 1.29 \times 10^{-5}\,\mathrm{m^2\,K/(W\,h)}$, $c_f = 0.1386/^{\circ}\mathrm{C}$, which is a doubling of the fouling rate for every five degrees of warming, ϕ the fouling tendency of the feed, and P the polymer rate against a reference $P_0 = 5.2\,\mathrm{t/h}$, so the design production of 5.10 t/h fouls at 0.98 of k_f.

Table 11.2 Model Parameters

Symbol	Meaning	Value	Class
V	reactor volume	30 m^3	adjustable
A	cooled surface	190 m^2	adjustable
U_0	clean coefficient	2400 W/(m^2 K)	adjustable
k_f	fouling rate	1.29×10^{-5}	adjustable
c_f	fouling temperature rise	0.1386/°C	adjustable
$T_{j,min}$	refrigeration floor	−106 °C	adjustable
k_{ML}	Mooney against temperature	−1.40/°C	adjustable
θ	Mooney laboratory cycle	2.0 h	adjustable
τ_{MW}	product mixing time	0.60 h	adjustable
τ_j	refrigerant loop lag	0.035 h	adjustable
wash time	shutdown and restart	8 h	adjustable
ML band	specification	46 to 56	visible

The jacket temperature the duty requires is Equation (11.5), and the cooling margin is the fraction of the available temperature difference still unused, Equation (11.6).

$$T_j = T - \frac{Q}{UA} \tag{11.5}$$

$$m = 1 - \frac{Q/(UA)}{T - T_{j,min}} \tag{11.6}$$

The margin is the clock. It is readable at any instant, it falls smoothly, and it reaches zero at the moment the jacket reaches its floor. At the design condition it starts at 0.598 and reaches zero at hour 49.0.

11.3.3 Product Quality

Mooney viscosity follows the reactor temperature directly, Equation (11.7), with the mixing time of the reactor between the temperature and the product.

$$\tau_{MW} \frac{\mathrm{dML}}{\mathrm{d}t} = \mathrm{ML}_{ref} + k_{ML}\,(T - T_{ref}) - \mathrm{ML} \tag{11.7}$$

with $k_{ML} = -1.40$ Mooney units per degree Celsius and $\tau_{MW} = 0.60$ h. The specification band is 46 to 56, so the reactor temperature has about seven degrees of usable range before the product leaves grade at one end or the other.

11.3.4 Assumptions

The reactor is a single well mixed volume. A real butyl reactor has an axial circulation driven by a large impeller and the fouling is not uniform, so the coefficient in this model is an average of something that varies over the surface.

The fouling resistance grows linearly in time at a given temperature and production rate and is removed completely by a wash. Real deposits have an induction period, an asymptotic component and a removal term, and a hydroblast does not always return the surface to its clean condition.

Conversion is a function of catalyst and temperature with no molecular weight distribution. Mooney viscosity is a correlation on reactor temperature rather than a moment calculation, so nothing about the breadth of the distribution can be studied.

Isoprene incorporation and the unsaturation specification are represented by a linear relationship to the feed ratio, which is a simplification of a copolymerization that depends on temperature as well.

The refrigeration floor is a hard limit on jacket temperature with no dynamics of its own. A real refrigeration system loses capacity as its own condensers foul and as ambient temperature rises.

The wash takes a fixed eight hours and restores the surface exactly. Both are convenient rather than true.

11.4 Dynamic Analysis

11.4.1 The Countdown Is Readable in Closed Form

With the loop holding the reactor and the duty roughly constant, the required temperature difference is $\Delta T = Q/(UA)$, and substituting Equation (11.4) gives Equation (11.8).

$$\begin{aligned} \Delta T(t) &= \frac{Q}{A}\left(\frac{1}{U_0} + R_f(t)\right) \\ &= \Delta T_0 + \frac{Q}{A}\, k_f\, \frac{P}{P_0}\, \phi\, e^{c_f (T - T_{ref})}\, t \end{aligned} \tag{11.8}$$

At a fixed reactor temperature the required difference grows linearly in time, which is why the jacket trend in scenario 2 is a straight line and why its slope is a direct reading of the fouling rate. Setting $\Delta T = T - T_{j,min}$ gives the run length, Equation (11.9).

$$t_{run} = \frac{A\, P_0}{Q\, k_f\, P\, \phi}\, e^{-c_f (T - T_{ref})} \left[(T - T_{j,min}) - \frac{Q}{A U_0} \right] \tag{11.9}$$

This is the constant-rate limit of Equation (1.2): the fouling rate here does not grow during a run held at one temperature. At the design condition Equation (11.9) gives 48.9 h, and the simulated run of scenario 2 spends its margin at 49.0 h. Both temperature effects are visible in that expression and they pull opposite ways. The exponential prefactor lengthens the run as the reactor is cooled, because the fouling slows. The bracket shortens it, because the distance to the refrigeration floor is what the run is spending, and it goes to zero at $T = T_{j,min} + Q/(AU_0)$, which is −101.6 °C at the design duty and −101.4 °C once the larger feed-chilling duty at that temperature is included. From a clean surface held at −101 °C the expression gives a run of 6.7 h. The measured run at −101 °C in Table 11.3 is shorter still, 2.9 h, because that setpoint arrives after two hours of fouling at −95 °C, and the wash follows 0.8 h after the setpoint change. Both are that bracket almost closed.

The maximum of t_{run} is not the operating optimum, because a longer run at a lower rate is not necessarily more rubber. The quantity to maximize is on-specification production per hour of elapsed time including the wash, which for a run that stays in grade is Equation (11.10), the on-stream rate of Equation (1.3) with the wash as the downtime.

$$Y = \frac{P\, t_{run}}{t_{run} + t_{wash}} \tag{11.10}$$

which is why the eight-hour wash appears in every economic comparison in Section 11.6 and why the answer to scenario 5 depends on it.

11.4.2 Behavior at the Refrigeration Floor

Nothing in Equation (11.8) is discontinuous. The discontinuity is the actuator. While $T_j > T_{j,min}$ the loop is closed and the reactor is held. At the floor the loop is open and the reactor obeys Equation (11.11).

$$\rho V c_p \frac{\mathrm{d}T}{\mathrm{d}t} = Q(T) - U(t)\, A\, (T - T_{j,min}) \tag{11.11}$$

The reactor holds 39.9 MJ/K, so a sustained shortfall of 0.1 MW would warm it 9 °C per hour, but no such shortfall builds up. Each degree of warming lowers the feed-chilling duty and raises the duty the floor can remove by UA, 0.18 MW/K when the floor is first reached, so within a few minutes the reactor settles onto the temperature at which the two terms balance and then creeps with that temperature as the fouling grows. Warming raises $\mathrm{d}R_f/\mathrm{d}t$ by Equation (11.3), which lowers U, which

moves the balance point faster, so the creep accelerates. The loss of control is abrupt and the loss of product is slow, and the measured consequence in Section 11.6 is that the margin gives about a day of warning before off-grade rubber and the reactor temperature gives sixteen hours once it has moved a degree.

11.5 Controller Design

11.5.1 Velocity Form PID

Every loop uses the shared velocity-form PID of Section 1.6.2 with the derivative on the measurement, written in this chapter's symbols as Equations (11.12) to (11.16).

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k) \tag{11.12}$$

$$P_k = K_c e_k, \qquad D_k = -K_c \tau_D\, a\, \frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{11.13}$$

$$s_k = s_{k-1} + \big(P_k - P_{k-1}\big) + \big(D_k - D_{k-1}\big)$$

$$c_k = u_{bias} + s_k + I_{k-1} \tag{11.14}$$

$$u_k = \min\big(\max(c_k, u_{min}), u_{max}\big) \tag{11.15}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\, e_k + \frac{u_k - c_k}{\tau_t}\right]\Delta t \tag{11.16}$$

where a is the controller direction, s_k is the never-clipped running sum of proportional and derivative increments (S_k of Equation (1.15)), c_k is the demanded output (v_k of Equation (1.16)), and the integral state I is seeded with the initial output so that $u_{bias} = 0$. Anti-windup is the back-calculation term of Equation (11.16), with tracking time $\tau_t = \tau_I$, so only the integral is corrected at a limit [14]. The derivative acts on the measurement through a 0.05 h filter, and $\tau_D = 0$ on every loop as shipped. In manual, and for any loop not in service, the integral tracks the value applied, so the transfer to automatic is bumpless. The controller scan Δt is 0.02 h (72 s) for all three loops and integration is fixed-step fourth-order Runge-Kutta at 0.005 h.

The low limit of the temperature loop is the refrigeration floor itself, $u_{min} = T_{j,min}$, and the high limit is −80 °C, so the back-calculation of Equation (11.16) holds the integral at the floor. Once the jacket reaches −106 °C the command stays there rather than winding on, and a warmer setpoint or a rate cut acts on the jacket at the next scan. In manual the operator may ask for as low as −120 °C, and the floor is imposed on the applied jacket temperature after the station.

11.5.2 Four Architectures and Two Wash Policies

Off leaves the jacket in manual. *Temperature* closes TIC-201 from the reactor temperature to the jacket temperature with $a = +1$, because the process gain is positive and a reactor above its setpoint calls for a colder jacket, with $K_c = 4.0$ degrees of jacket per degree and $\tau_I = 0.30$ h. *Cascade* adds MC-201 from the two-hour Mooney laboratory result to the reactor temperature setpoint, with $a = -1$ because a warmer reactor makes a lower Mooney, $K_c = 0.30$ °C per Mooney unit, $\tau_I = 8$ h, and the setpoint confined to −102 °C to −86 °C. The laboratory result is held for its two-hour cycle and passed through a 0.05 h filter. The fourth architecture, *MPC*, replaces MC-201 and the conversion loop with a receding-horizon plan and is described in Section 11.5.3.

The outer integral time is long on purpose. A Mooney result up to two hours old, with a laboratory repeatability of 0.4 units (1σ) in scenario 6, drives a setpoint through a plant gain of 1.40 units per degree, so $K_c = 0.30$ °C per unit is a loop gain of 0.42. At that gain a 2 h integral time only makes the setpoint busier, 0.75 °C peak to peak over two days against 0.70 °C, and no cycle appears. Raising K_c to 1.0 °C per unit as well, a loop gain of 1.4, gives a sustained cycle with a 4.0 h period, twice the laboratory cycle, in which the setpoint swings between its −102 °C limit and −86.8 °C and the Mooney between 42.8 and 56.4. The sample age is what sets the outer loop speed, not the tuning

preference [10]. The quality cascade of Chapter 8 and the per-head correction of Volume 8, Chapter 9 also act on results that describe product made some time before.

A conversion loop manipulates the catalyst rate between 0 and 8 kg/h to hold conversion, with $K_c = 12$ kg/h per unit of conversion and $\tau_I = 1.0$ h. Whether it should be in automatic is the subject of scenario 7.

The wash policy is an operator decision, a fixed elapsed time, or a threshold on the cooling margin. The two automatic policies are implemented as a shutdown, an eight-hour wash and a restart in which the catalyst is ramped in over 1.5 h and the quality trim is held for 3.0 h. Whatever the policy, a high-temperature trip on the reactor transmitter at −80 °C kills the reaction and starts the same wash, so a run that nobody stops ends in a trip. The threshold is an illustrative simulation choice, eight degrees above the warm-slurry alarm, rather than a limit derived from a plant safety study. This model cannot establish protection adequacy.

11.5.3 Model Predictive Control

The predictive option treats the cooling margin of Equation (11.6) as a constraint on a prediction rather than as a wash trigger, in the receding-horizon form surveyed by Mayne et al. [15] and set up in general in Section 1.6.3. Every $T_s = 0.1$ h, five controller scans, it moves two handles: the reactor temperature setpoint T_{SP}, which TIC-201 follows exactly as it follows the output of MC-201, and the catalyst command u_c, which is applied as the conversion loop's command is. MC-201 and the conversion PI follow the commands actually applied, scan by scan and without a preload, so a hand-back to the cascade is bumpless. With the conversion loop in manual the catalyst stays with the operator and the plan moves the setpoint alone. TIC-201, the refrigeration floor, the wash policy, the catalyst ramp after a restart and the high-temperature trip remain in force beneath and above the plan, and the plan never decides a wash.

Controller model. The internal model uses only the declared defaults of Table 11.2 and the design feeds, never from the run's parameters, the plant state or the disturbances, Equations (11.17) to (11.22).

$$\tau_T \frac{\mathrm{d}T}{\mathrm{d}t} = T_{SP} + b_T - T \tag{11.17}$$

$$\tau_{MW} \frac{\mathrm{d}M}{\mathrm{d}t} = \mathrm{ML}_{ref} + k_{ML}(T - T_{ref}) - M \tag{11.18}$$

$$\tau_c \frac{\mathrm{d}c}{\mathrm{d}t} = u_c - c, \qquad \tau \frac{\mathrm{d}X}{\mathrm{d}t} = X^*(\hat{a}\,c, T) - X \tag{11.19}$$

$$R_f(t+s) = \hat{R}_f + \hat{r}\,s + z(s) \tag{11.20}$$

$$\frac{\mathrm{d}z}{\mathrm{d}s} = \hat{r}\left(c_f\,\delta T + \frac{\delta X}{\bar{X}}\right) \tag{11.21}$$

$$T_{j,need} = T - \frac{Q(X,T)}{A}\left(\frac{1}{U_0} + R_f\right) \tag{11.22}$$

Here τ_T is the closed-loop lag of TIC-201, b_T its offset filtered over 0.5 h, M the Mooney of the polymer being made, corrected to $\widehat{\mathrm{ML}} = M + b_M$ by the laboratory bias b_M, c the delivered catalyst with the 0.08 h lag τ_c, X^* the conversion law of Equation (11.1) with the activity estimate $\hat{a}$, Q the declared duty of Equation (11.2), $\hat{R}_f$ and $\hat{r}$ the inferred fouling resistance and its trend, δ a deviation from the present point and $\bar{X}$ the present conversion. The model is linearized at the measured point at every decision (the resistance taken at mid-horizon for the slopes) and sampled exactly at T_s, giving six states and four outputs: the Mooney, the margin row, the conversion and the temperature. The margin row

$$T_{j,need} - m\,T \;\geq\; (1-m)\,T_{j,min} \tag{11.23}$$

is Equation (11.6) rearranged into a form linear in the states, with m the margin the plan keeps, 5 % as shipped, and $T_{j,min}$ read as TIC-201 reads it.

The lag τ_T was identified with TIC-201 closed at its shipped tuning, because an open-loop step test removes the loop the plan relies on. Setpoint steps of $\pm 1\,^{\circ}\mathrm{C}$ at the design point with noise off give a gain of 1.00 in both directions and a ringing response, a peak of 1.56 at 0.055 h and a period of 0.12 h, that no low-order model fits to better than a 4 to 9 percent residual. The transmitter's mean delay by the area method is 0.080 h and a first-order fit at the decision instants gives 0.085 h, so $\tau_T = 0.08\,\mathrm{h}$. Through the 0.60 h mixing lag the model predicts −1.095 Mooney units per degree at one hour against a measured −1.099 and −1.342 against −1.346 at two hours, but only −0.094 against −0.215 at 0.1 h, which is the ring the model leaves out. At the design point the conversion law gives $\partial X/\partial c = 0.0265$ per kg/h, the production is $6.0\,\mathrm{t/h} \times X$, and the duty rises by 1.55 MW per unit of conversion.

Fouling as a trend of its own future. The resistance is never read. At every scan the controller computes $1/U = A(T - T_j)/Q(X,T)$ from the transmitter, the jacket and the declared duty at the measured conversion, filters it over 0.1 h, keeps a sample only while the reactor is within 0.1 °C of its setpoint and 0.3 h after any setpoint move of 0.05 °C or more, and regresses the last 4 h for the trend $\hat{r}$. The ratio of $\hat{r}$ to the declared rate law of Equation (11.3) at the measured temperature and production is a tendency $\hat{f}$, bounded between 0.2 and 5, which supplies the trend when less than an hour of samples exists after a start or a wash. Equation (11.20) extrapolates that trend in a straight line over the horizon and adds the first-order effect of the plan's own temperature and production on the rate. Against the simulated truth, sampled every half hour outside washes, the inferred resistance has a mean relative error of 0.1, 0.3 and 0.6 % under the predictive controller on cases 7, 8 and 9, and the trend a mean relative error of 2.2, 2.7 and 4.0 %. The worst tendency errors, 0.56 on case 8 and 0.69 on case 9, fall in the hours after a wash while the window refills. Under the shipped loop the jacket on its floor leaves no quiet samples, and on case 7 the tendency freezes at 0.90 against a true 1.00.

Estimators. The design is offset-free in the sense of Pannocchia and Rawlings [16], with one disturbance estimate per channel. The catalyst activity multiplier comes from a steady-state Kalman filter on the sampled pair $[X, \ln \hat{a}]$ with process covariances 10^{-8} and 2×10^{-5} and measurement covariance 10^{-6} per scan, whose gain is $L = [0.156,\ 4.11]^{\mathsf{T}}$. It updates only while the delivered catalyst exceeds 1 kg/h and the conversion exceeds 0.3, within bounds of 0.3 and 1.6, and in scenario 10 it reads 0.650 two hours after the poisoning to 0.65. The Mooney bias takes half of each laboratory residual, within ±5 units, and the residual is formed against the model's Mooney at the instant the sample was taken, held in a 20 h history, and read 0.2 h after the sample once the laboratory filter has settled. Results during the 3.0 h restart hold are skipped, the rule the shipped trim follows. On the challenge, with a laboratory repeatability of 0.5 units, the bias stays within 0.69 units of zero.

Objective and constraints. The horizon is $N = 80$ steps (8 h) in ten move blocks, their lengths running 1, 1, 2, 3, 5, 8, 10, 10, 20 and 20 steps, so the quadratic program has 20 move variables and three slacks. The objective, Equation (11.24), prices production as a linear cost in the manner of economic MPC [17].

$$
\begin{aligned}
J = \sum_{k=1}^{N} \Big[& - w_P\,(6.0\,\mathrm{t/h})\,X_k T_s + q_M \big(\widehat{\mathrm{ML}}_k - \mathrm{ML}_{SP}\big)^2 \Big] \\
& + \sum_{j} \Big[s_T\,\Delta T_{SP,j}^2 + s_c\,\Delta u_{c,j}^2 \Big] + \sum_{i=1}^{3} \rho_i \big(\varepsilon_i^2 + \varepsilon_i\big)
\end{aligned}
\tag{11.24}
$$

with $w_P = 1$ per tonne, $q_M = 0.005$ per Mooney unit squared, $s_T = 2$ per $^{\circ}\mathrm{C}^2$ and $s_c = 0.5$ per $(\mathrm{kg/h})^2$ at each block boundary. Three outputs are soft, judged at every second step, each with one slack $\varepsilon_i \geq 0$ over the whole horizon: the Mooney within $\mathrm{ML}_{SP} \pm 2.5$ at $\rho_1 = 100$, the margin row of Equation (11.23) at $\rho_2 = 10^4$ per kelvin, and the conversion at or below its target at $\rho_3 = 10^3$, so the plan never runs the reactor above the rate the operator commissioned. Where the present point already violates a limit, the limit re-enters through a funnel with a 1 h time constant. The hard limits

are the ranges the shipped loops clamp to, $T_{SP} \in [-102, -86]\,°\mathrm{C}$ and $u_c \in [0, 8]\,\mathrm{kg/h}$, with rate limits of 0.3 °C and 0.8 kg/h per decision.

Acceptance and fallback. A plan stands only while the solver's status is optimal, every planned command lies inside its range, the plan's first command matches the move being applied, and the cost, the slacks and every published prediction are finite. Otherwise the historian is marked, MC-201 and the conversion loop take the two handles from the commands they have been tracking, and the plan is retried at every scan. Six forced failures injected at hour 20 of scenario 10 with noise off were each rejected, the setpoint moved −0.0175 °C and the catalyst 0.0019 kg/h at the hand-over, and the next plan was accepted one scan later. A plan costs a mean of 0.57 to 0.59 ms in Chrome over the 480 plans of scenario 10, the largest 7.3 ms for the first. That load is 0.6 % of a single core when one simulated hour passes each second and 35 % at the top speed of sixty.

Three models kept apart. The plant is Equations (11.1) to (11.7) with the true activity, fouling tendency and exponential temperature acceleration, the refrigerant lag and the loop's ring. The identified model is the closed-loop lag τ_T from the step tests above. The controller's model is Equations (11.17) to (11.22) at the declared defaults, linearized at each decision, with the inferred fouling trend, the estimated activity and the two biases. The controller reads the transmitter, the jacket temperature, the conversion, the laboratory result, the commands applied and $T_{j,min}$, and nothing else, so a parameter changed by the learner is model mismatch rather than knowledge.

11.6 Results

Eleven scenarios run from watching the coefficient fall by hand to a graded four-day challenge. Table 11.4 lists them with the point each one teaches and what completes it. The app runs at 0.5 to 60 simulated hours per second, 5 h per second by default, so a two-day cycle takes about ten seconds. Reset restarts a case from time zero and keeps the reader's own settings, which is how the comparisons in scenarios 4 to 8 are made (Section 1.6.4).

11.6.1 The Controlled Variable Is Not Where the Information Is

With the temperature loop closed in scenario 2 the reactor stays within 0.010 °C of its setpoint for forty-eight hours. That residual is the steady error of a PI loop following a ramp, the jacket's 0.134 °C per hour divided by K_c/τ_I, and otherwise the loop is working perfectly. Over the same period the jacket walks from −99.4 °C to −105.9 °C, the required temperature difference grows from 4.4 K to 10.9 K, and the cooling margin falls in a straight line, 0.0122 per hour, from 0.598 to 0.012. The margin reaches zero, and the jacket its floor, at hour 49.0. With the loop open in scenario 1 and the jacket held at −99.4 °C, the reactor warms instead, to −83.3 °C at hour 48, and the coefficient falls further, to 600 $\mathrm{W/(m^2\,K)}$, because warmer slurry fouls faster. The Mooney follows at 1.40 units per degree, leaves the grade at hour 22.7 and reaches 35.1 by hour 48.

> **TRY IT IN THE APP 2 · The Temperature Loop, and What It Is Spending**
>
> Load the scenario, press Run and watch the jacket trace walk toward the dashed refrigeration floor while the reactor trace stays flat. Over 48 hours the reactor stays within 0.010 °C of its setpoint, the jacket walks from −99.4 °C to −105.9 °C, and the cooling margin falls from 0.598 to 0.012.

The end of a run is abrupt in the controller and slow in the process. Scenario 3 starts the reactor with the fouling of about 43 hours at the design rate and the jacket at the −105.2 °C that the fouled surface needs, so the reactor holds its setpoint within 0.01 °C until the jacket reaches its floor and the margin is spent, both at 6.3 h. The reactor is 1 °C above its setpoint at 13.8 h, the Mooney leaves the grade at 29.9 h, and the laboratory result confirms it at 30.2 h. The creep accelerates with the fouling, from 0.13 °C per hour just after the floor to 0.29 °C per hour by hour 43, and the reactor passes −88 °C at 43.6 h. The margin trend gives about a day of warning before off-grade rubber, the



Figure 11.2 The butyl rubber application at hour 48 of scenario 3, the end of a run, cropped to the plant view, with the temperature loop in automatic and the reactor started at the fouling resistance of about 43 hours at the design rate. The jacket and the temperature controller output have been at the −106 °C floor since hour 6.3, as (1) the badges show with the margin at zero, while (2) the reactor has crept to −86.6 °C and the Mooney viscosity has fallen to 39.5, with the laboratory result at 40.4. (3) The fouling layer lines the cooled wall, (4) the bar beside the vessel is the temperature difference the duty needs against the floor, and (5) the cooling margin through the run has reached the point marked no control. The run has made 153 t on specification and 18.1 h of off-grade rubber.

reactor temperature gives sixteen hours once it has moved a degree, and the laboratory gives none. Figure 11.2 shows the case at hour 48.

🕮 **Pause.** Try Exercise 11.1 before reading on.

11.6.2 The Temperature Optimum Is Warm

Table 11.3 reports run length, Mooney viscosity and yield at seven reactor setpoints, measured in scenario 4. Each setpoint is entered 0.1 h after the scripted step at hour 2, the wash is triggered at a cooling margin of 0.03, and the run, the Mooney and the tonnage are read when the wash starts. The two yields divide the tonnes made, and the tonnes on specification, by the run plus its eight-hour wash.

> **TRY IT IN THE APP 4 · Temperature Does Three Things, and They Disagree**
>
> Run to the first automatic wash after the step to −99 °C at hour 2: the run lasts 27.0 h and the Mooney ends at 56.6, above the grade. Press Reset, run past the step, set the Reactor temperature setpoint to −93 °C and run to the wash again. The run now lasts 47.6 h, with the Mooney at 48.2.

Colder is shorter, and the arithmetic is not close. The fouling does slow as the reactor gets colder and it is beaten by the collapse of the temperature difference. At −101 °C there are only five degrees between the reactor and the refrigeration floor, the fouling laid down in the first two hours nearly exhausts them, and the wash follows 0.8 h after the setpoint change, before the product has lined out at its Mooney of 59.4.

Table 11.3 Reactor Setpoint Against Run Length, Product and Yield, at 85 % Conversion

Setpoint (°C)	Run (h)	Mooney at wash	Made (t/h)	On spec (t/h)
−101	2.9	57.1	1.37	1.25
−99	27.0	56.6	3.94	0.51
−97	40.6	53.8	4.26	4.26
−95	46.5	51.0	4.35	4.35
−93	47.6	48.2	4.37	4.37
−91	45.8	45.4	4.34	0.34
−89	42.4	42.6	4.29	0.28

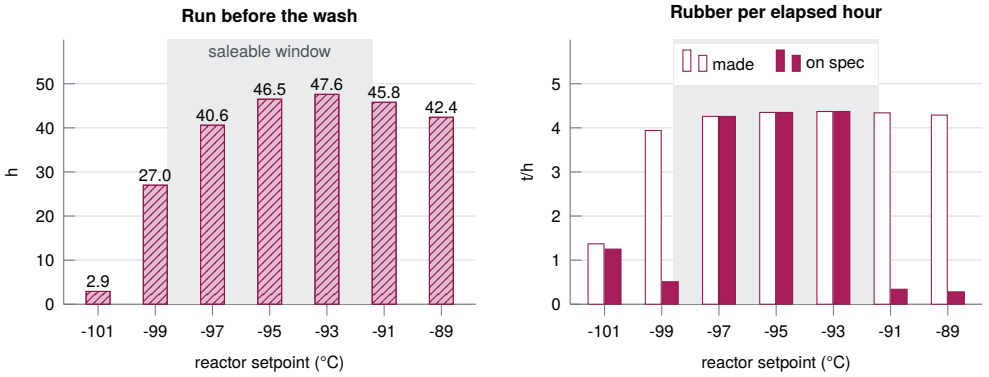

Figure 11.3 The run and yield columns of Table 11.3 against the reactor setpoint. Left, the run before the wash. Right, the tonnes made (open bars) and on specification (solid bars) per elapsed hour, counting the eight-hour wash. The shaded band is the saleable window of −98.6 °C to −91.4 °C. Notice that the run collapses on the cold side, that the tonnes made are nearly flat from −97 °C to −89 °C, and that only the setpoints inside the band keep their rubber on specification.

🕮 **Pause.** Try Exercise 11.2 before reading on.

The yield optimum is near −94 °C and the peak is very flat, 4.35 t/h at −95 °C, 4.37 t/h at −93 °C and 4.34 t/h at −91 °C counting every tonne made. What stops the plant running warmer is not the physics but the grade. At −91 °C the Mooney is 45.4, below its 46 minimum, and only 18.3 t of the 233.7 t made is on specification. The grade binds on the cold side as well, where the Mooney of 56.6 at −99 °C is above the 56 maximum. The saleable window is −98.6 °C to −91.4 °C, so the economic optimum is a constrained one and the binding constraint is a number on a contract. Figure 11.3 draws the run and yield columns of Table 11.3 against the setpoint.

11.6.3 The Case for a Constraint Override

Scenario 7 starts the reactor with the fouling of about 38 hours at the design rate and poisons the catalyst to an activity of 0.65 at hour 2, with the conversion loop in automatic. The loop does exactly what it was asked to do and raises the catalyst from 4.82 kg/h to 7.2 kg/h by hour 12 to recover the production rate. Recovering the production recovers the heat load, which on a reactor with a nearly saturated jacket is the resource the run is spending. The jacket reaches its floor at 11.6 h against 16.2 h with the loop in manual, and the Mooney leaves grade at 35.0 h against 43.7 h. That is the standard argument for a constraint override, and nothing in a conventional cascade knows about it. The protective override of Volume 4, Chapter 2 is the same idea where the limit is reached in seconds.

On this reactor the arithmetic goes the other way. Manual makes 4.72 t/h of polymer instead of 5.07 t/h, and because the end of a run creeps rather than runs away, the hours that automatic loses cost

less than the tonnes that manual loses. At hour 36 the yield is 4.92 t/h in automatic against 4.74 t/h in manual, and stopping each run when its rubber leaves grade, with an eight-hour wash, gives 4.12 t/h against 4.00 t/h. An override that gave way to the cooling margin has to be justified by that arithmetic, and it pays only where the end of a run is abrupt or the wash is long.

11.6.4 A Schedule Is a Controller With No Feedback

Scenario 8 runs a fixed wash at hour 44 with the feed fouling 1.6 times faster than the schedule was written for. The margin is gone at hour 31.3, so the jacket sits on its floor for the last 12.7 hours of the first run. The reactor creeps rather than runs away, that run stays in grade, and over four days the Mooney is out of grade for 0.9 h, at the end of the second run. The schedule yields 4.54 t/h. A wash on a margin of 0.10 makes no off-grade rubber and yields 3.99 t/h, because it washes 26.2 hours into every run and pays for three washes in four days. Lower thresholds of 0.05 and 0.02 both give 4.09 t/h. The schedule is lucky rather than right. Moved to hour 60 it yields 3.87 t/h with 13.6 h off grade, because the rubber of the first run leaves the grade at hour 46.4.

The general form transfers to other plants. A schedule is right on the day it is set and wrong every day the plant is different, and the cure is to find the measurement the schedule was standing in for and use that instead [1]. On this reactor that measurement is not the cooling margin, which runs out about fifteen hours before the product leaves grade, but the quantity that ends the run, the Mooney viscosity or the reactor temperature that sets it.

🕮 **Pause.** Try Exercise 11.3 before reading on.

11.6.5 The Graded Challenge

The challenge runs ninety-six hours with the fouling already 1.35 times the design rate, a loss of catalyst activity to 0.78 at hour 18, a rate increase to 6.8 t/h at hour 40 and a dirtier feed at 2.1 times the design fouling at hour 66. It starts with cautious tuning ($K_c = 1.0$, $\tau_I = 1.2$ h on the temperature loop), the conversion loop in manual and the wash left to the operator, and it is judged by the criterion in the last row of Table 11.4.

As shipped, with no wash chosen, the jacket reaches its floor at 38.6 h, the reactor passes −88 °C at 58.2 h, and once the feed fouls at 2.1 times the design rate the creep of the floor turns into a runaway that the −80 °C trip ends at 68.5 h with an eight-hour wash. The catalyst station keeps its 4.82 kg/h through that wash, so the restart at hour 76.5 adds 100 t of rubber within grade and the yield is 3.65 t/h, which the 15.00 °C excursion still fails. The suggested configuration, the Mooney cascade with the tuned loops, the conversion loop in automatic and a wash on a margin of 0.10, passes with a worst error of 0.83 °C, all of it at start-ups, a lined-out error of 0.21 °C, and a yield of 3.83 t/h from three washes and 367 t on specification. The cascade settles the setpoint near −93.5 °C, on the flat top of Table 11.3. Holding −97 °C with the temperature loop alone makes the same 367 t with a worst error of 3.4 °C. With the conversion loop switched to manual at hour 66 the same configuration still passes, at 3.86 t/h and a worst error of 1.07 °C, and with the conversion loop in manual from the start it passes at 3.90 t/h from two washes instead of three. The same configuration on the 44-hour clock yields 4.09 t/h but lets the reactor run 22 °C above its setpoint, with 11.5 h off grade, until the trip at hour 93.9. Passing needs the margin wash and the tighter temperature tuning, without which a lined-out run sits 1.1 °C off setpoint, but neither the conversion loop in automatic nor the cascade: the temperature loop alone at −95 °C passes at 1.46 °C.

11.6.6 The Predictive Controller Against the Shipped Structures

Table 11.5 runs cases 7, 8 and 9 and the challenge twice, once per structure, from identical seeds, durations, instruments and noise. The shipped structure keeps its pinned tuning, with the suggested configuration on the challenge, and the predictive controller runs at the settings of scenario 10 on every case with no retuning. The application has no preview, so neither structure receives one. Three further rows are measured beside them. The best fixed setting of the shipped structure is the highest yield with no off-grade hour found in hindsight over the conversion target (0.60 to 0.90) under the

Table 11.4 Scenarios and Completion Requirements

#	Scenario	Kind	Teaching point, and what is required to complete
1	Watch the coefficient fall	Demonstration	*A gain that dies.* With the jacket held in manual at –99.4 °C, run until the Mooney leaves the 46 to 56 grade (hour 22.7), then on to hour 48 with nothing changed.
2	The temperature loop	Experiment	*The output is the diagnostic.* With the temperature loop in automatic, run until the cooling margin falls below 0.10 (hour 40.8), then on to hour 48.
3	The end of a run	Demonstration	*Saturation is a discontinuity.* From the fouling of about 43 hours with the temperature loop in automatic, run until the jacket reaches its floor after the first hour, the reactor rises 1 °C above setpoint, and the laboratory Mooney falls below 46 (hour 30.2).
4	Temperature does three things	Experiment	*A constrained optimum.* Run to the first margin-triggered wash after the step to –99 °C at hour 2, then Reset, set –93 °C after the step and run to a wash more than 40 hours into the run.
5	Rate against run length	Experiment	*Fixed downtime arithmetic.* Run to the first wash after the isobutylene push to 7.2 t/h at hour 2, then Reset, raise isoprene to 0.14 t/h to 0.15 t/h after the push and run to a wash with more than 100 t on specification.
6	The Mooney cascade	Experiment	*Sample age sets the speed.* Run 48 hours on the cascade with $\tau_{I,M} \geq 6$ h, then Reset, set $\tau_{I,M} \leq 3$ h before the activity drop at hour 3, and run 48 hours again.
7	The loop that helps then hurts	Experiment	*The arithmetic of an override.* Run 12 hours with the conversion loop in automatic, then Reset, put it in manual before the poisoning at hour 2, and run 12 hours again.
8	The clock against the margin	Experiment	*Which measurement ends a run.* Run 96 hours with the wash on a 44-hour clock, then Reset, switch to the margin policy before the dirtier feed at hour 2, and run 96 hours again.
9	A dirtier feed	Exploration	*Early warning.* Run past the fouling step to 2.2 times the design rate at hour 1 until the margin policy at 0.08 starts a wash (hour 19.8), and confirm that the Mooney, the reactor temperature and the conversion stayed on target.
10	The margin as a constraint	Experiment	*What a margin costs.* From case 7's fouled reactor and poisoning, run 48 hours under the MPC at a 5 % margin, then Reset under the temperature loop alone and run 48 hours, then return to the MPC at 2 %, Reset and run 48 hours again.
11	Challenge	Challenge	*Operating, not only controlling.* Over 96 hours, hold the reactor within ±1.5 °C of setpoint from hour 63.4, start-ups included, and within ±0.3 °C once a run is more than four hours old, and finish with a yield above 3.4 t/h, with the model parameters and disturbances as given.[a]

[a] With the shipped configuration (temperature loop at $K_c = 1.0$ and $\tau_I = 1.2$ h, conversion loop in manual, no wash chosen) the attempt fails by a wide margin: the reactor runs away after the jacket reaches its floor until the high-temperature trip shuts it down at hour 68.5, and the reported worst error is 15.00 °C. The yield, 3.65 t/h, clears the target only because the restart at hour 76.5 resumes at the operator's catalyst rate. With the suggested configuration (Mooney cascade, conversion loop in automatic, wash on a margin of 0.10, tuning of Section 11.5) it passes with a worst error of 0.83 °C, 0.67 °C inside the band, a lined-out error of 0.21 °C, 0.09 °C inside, and a yield of 3.83 t/h, 0.43 t/h above the target.

case's own wash policy and over a grid of margin washes. The margin override on case 7 is a classical structure built in the measurement script only: a PI loop on the app's cooling margin with a setpoint of 5 %, low-selected against the conversion loop, with the reactor setpoint raised after the first step to –93.21 °C, the warm edge of the plan's Mooney band. The last row of each case removes the fouling trend from the plan by holding the resistance at its present value over the horizon.

Scenario 10 is case 7's reactor, started at about 38 hours of fouling with the cooling margin at 13.4 % and poisoned at hour 2, with the predictive controller in charge, the conversion loop in automatic and the wash left to the operator. The plan recovers the conversion to its 0.85 target with the catalyst at 7.59 kg/h by hour 4.3, then buys margin cheaply first by walking the reactor up its Mooney band, reaching –93.25 °C at hour 27.0. From about hour 13, when the margin has fallen to 5 %, it brings the catalyst down, 7.49, 4.01, 2.27 and 1.55 kg/h at hours 12, 24, 36 and 48, and the margin then stays between 4.79 and 5.03 %, as shown in Figure 11.4. The jacket never reaches its

Table 11.5 The Predictive Controller Against the Shipped Structures, Measured on the Engine

Case	Structure	Yield (t/h)	First wash (h)	Washes	Floor (h)	Off grade (h)	Mooney IAE	Catalyst travel (kg/h)	Plans, fallbacks
7, 48 h	temperature loop, conversion loop automatic	3.689	none	0	36.36	13.01	144.4	3.11	–
	MPC	4.412	none	0	0	0	85.6	8.90	480, 0
	best fixed: conversion 0.75	4.484	none	0	27.64	0	59.2	1.53	–
	margin override at 5 %, warm setpoint	4.319	none	0	0	0	118.6	8.70	–
	MPC without the fouling trend	4.485	none	0	0	0	83.4	9.06	480, 0
8, 96 h	cascade, wash on a 44-hour clock	4.545	44.01	1	26.24	0.95	104.2	10.35	–
	MPC	4.473	44.01	1	0	0	58.9	13.65	880, 0
	best fixed: conversion 0.84	4.542	44.01	1	24.82	0	97.0	9.73	–
	MPC without the fouling trend	4.529	44.01	1	0	0	50.4	13.71	880, 0
9, 48 h	cascade, wash on a margin of 0.08	4.085	19.83	2	0	0	34.1	9.93	–
	MPC	4.065	21.78	1	0	0	8.2	10.64	400, 0
	best fixed: conversion 0.90, wash at 0.03	4.206	19.25	2	0	0	33.8	15.23	–
	MPC without the fouling trend	4.067	19.83	1	0	0	3.0	10.54	401, 0
Challenge, 96 h	suggested cascade, wash on a margin of 0.10	3.827	30.28	3	0	0	74.6	38.40	–
	MPC with the suggested regulatory settings	4.139	33.04	3	0	0	24.5	31.51	765, 0
	best fixed: conversion 0.85, wash at 0.03	4.143	33.18	3	0.02	0	80.3	26.89	–
	MPC without the fouling trend	3.805	30.30	3	0	0	22.0	37.39	721, 0

Yield is tonnes on specification per elapsed hour including washes. Washes counts the washes started. Floor is the hours the jacket spends at −106 °C outside washes. Mooney IAE accumulates the absolute error of the true Mooney against its target outside washes, and catalyst travel sums the absolute changes of the catalyst command, both by the scorecard's definitions. A best fixed setting that allows off-grade rubber yields more on two cases, 4.517 at a conversion target of 0.76 on case 7 (0.28 h off grade) and 4.548 at 0.87 on case 8 (2.77 h off grade).

floor, the reactor never leaves its setpoint by more than 0.02 °C, the Mooney never falls below 48.51, and all 211.8 t are on specification, a yield of 4.412 t/h. The loop of case 7 makes 243.3 t but only 177.1 t on specification, 3.689 t/h, after 36.36 h on the floor with the reactor up to −88.20 °C. The comparison has a crossover. The yield at hours 12, 24, 36 and 48 is 5.08, 4.96, 4.71 and 4.41 t/h under the plan and 5.00, 5.04, 4.92 and 3.69 t/h under the loop, so the loop leads from some point between hours 12 and 24 until after its rubber leaves grade. The margin the plan keeps is a price: at 2 % the plan yields 4.530 t/h and at 8 % 4.284 t/h, with no hour on the floor and no off-grade hour at either.

Most of the case 7 win belongs to the margin row, not to the plan. The classical override at the same 5 % with the warm setpoint yields 4.319 t/h, so it carries 0.63 of the 0.72 t/h gained, and the prediction and coordination of the two handles add 0.09 t/h. At a 2 % margin the two yield 4.423 and 4.530 t/h. Neither ingredient works alone: the override at −95 °C yields 3.970 t/h, and the warm setpoint without an override yields 3.266 t/h, less than the shipped loop, because its floor comes

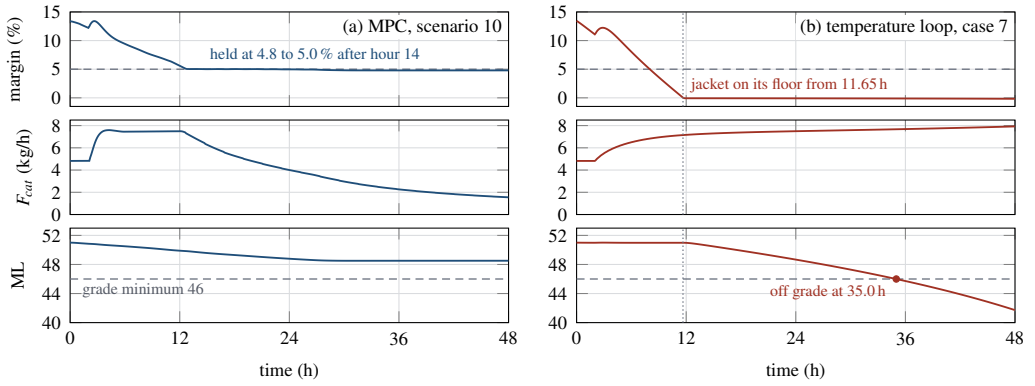

Figure 11.4 The same fouled reactor, poisoned to an activity of 0.65 at hour 2, over two days under (a) the predictive controller of scenario 10 and (b) the temperature loop with the conversion loop in automatic, case 7. Top, the cooling margin with the 5 % the plan keeps. Middle, the catalyst command. Bottom, the Mooney viscosity of the polymer being made against the grade minimum. The plan cuts catalyst from about hour 13 and the jacket never reaches its floor, while the loop holds production until the jacket stops at hour 11.65 and its rubber leaves grade at hour 35.0.

later but its reactor then creeps further. A fixed conversion target of 0.75 with no wash beats the plan at 4.484 t/h, all of it on specification, by sitting 27.64 h on the refrigeration floor and ending the two days with a Mooney of 46.31, a third of a unit inside the grade. Counted over a fixed 48 hours, rubber made after the loop's product leaves grade weighs against the loop, which is why these margin structures pay here while the per-run arithmetic of scenario 7 favors holding production.

The predictive controller loses cases 8 and 9, and the mechanism is the one scenario 8 already teaches. The end of a run on this reactor creeps, so the tonnes the plan gives up to protect the margin are worth more than the grade they protect. On case 8 the shipped cascade spends 26.24 h on its floor before the wash at hour 44 and loses only 0.95 h of grade, while the plan has cut the catalyst to 1.75 kg/h by hour 44 and yields 4.473 t/h against 4.545 t/h. On case 9 the margin wash at 0.08 ends the run in either structure, and the plan's cut only stretches the first run from 19.83 to 21.78 h at a lower rate, 4.065 t/h against 4.085 t/h. The best fixed setting beats the plan on both, 4.542 and 4.206 t/h. On the challenge the plan yields 4.139 t/h against the suggested cascade's 3.827 t/h, 397.3 t against 367.4 t, because its cut delays the first margin wash from hour 30.28 to 33.04, and the best fixed setting of the cascade in hindsight ties it at 4.143 t/h. On the variant of case 8 with the wash moved to hour 60, which scenario 8 itself measures as a mistake, the plan yields 4.248 t/h against 3.872 t/h with 13.61 h off grade.

The fouling trend is what keeps the margin the operator set. Without it the plan holds 3.38 and 2.59 % at its lowest on cases 7 and 8 instead of 5 % and makes more, 4.485 and 4.529 t/h, but on the challenge it yields 3.805 t/h, below the suggested cascade, because it no longer stretches the first run, which reaches its wash at hour 30.30. In no case does the predictive controller reach the floor or make off-grade rubber, and no plan was rejected in 2525 solves across the four cases.

DID YOU KNOW? A Tighter Grade, Less Rubber

From the simulator: on case 9 of Table 11.5 the predictive controller holds the Mooney to an IAE of 8.2, against 34.1 for the cascade with a margin wash. Yet it yields 4.065 t/h against 4.085 t/h. Holding the grade more tightly did not make more saleable rubber.

11.7 Technology in Practice

What Measures It

Two measurements carry the chapter's loops. TT-201 reads the reactor temperature for TIC-201 on every 72 s scan, and in the setpoint steps of Section 11.5.3 its mean delay is 0.080 h. The thermocouple and the resistance thermometer that real reactors use are described in Chapter 10 and Chapter 4. The quality loop MC-201 reads the Mooney viscosity, the measurement that decides whether a run's rubber can be sold, and the box below describes how that number is made.

TECHNOLOGY Sensor: The Mooney Viscometer

Mooney viscosity is the resistance of the uncured rubber to flow. The standard test, ISO 289 or ASTM D1646, preheats a sample for one minute, then turns a large rotor in it at 2 revolutions per minute for four minutes and records the torque. At a test temperature of 100 °C the result is reported as ML(1+4) 100 °C [18, 19]. It is a laboratory test, not an online instrument. MC-201 of the chapter reads its result as a staircase, each step held for the two-hour cycle and describing rubber made earlier, with a repeatability of 0.4 units in scenario 6.

What Moves It

TIC-201 moves the jacket temperature, and the conversion loop moves the catalyst rate between 0 and 8 kg/h. The first is the handle the fouling consumes. In a real reactor it is the temperature of a refrigerant boiling against the cooled surface, as the box below explains.

TECHNOLOGY Final Element: Boiling Refrigerant Against the Wall

The butyl reactor of a 1946 patent was cooled by liquid ethylene, circulated around a central draft tube and its return tubes or held in a refrigerating jacket [8]. A boiling liquid holds its temperature at its boiling point until it has boiled away, and that boiling point rises and falls with the pressure on the liquid [20]. The refrigerant side of the wall can therefore be held at a steady temperature set by pressure, and what changes over a run is the polymer layer on the slurry side. TIC-201 of the chapter moves no valve directly: its output is the jacket temperature itself, reached through a 0.035 h refrigerant lag and stopped at the −106 °C floor that stands for the coldest the refrigeration can deliver.

Where the Controller Runs

Figure 11.5 draws the loops of Section 11.5 with hardware in place of the schematic's symbols. A distributed control system holds loops like these at a supervisory level that sends set points to field controllers, with loops nested or cascading and cycle times from milliseconds to minutes [21]. In the app the three PI blocks run in velocity form every 72 s, MC-201 writes the setpoint of TIC-201 as a cascade, and the laboratory result enters MC-201 held for its two-hour cycle. The MPC of Section 11.5.3 replaces MC-201 and the conversion loop with a plan every 0.1 h, five scans, and the wash policy and the −80 °C trip stay in force above every layer. The lag of TIC-201 that the plan uses was identified with the loop closed at its shipped tuning, and the box below describes how industry tests a plant for such a controller.

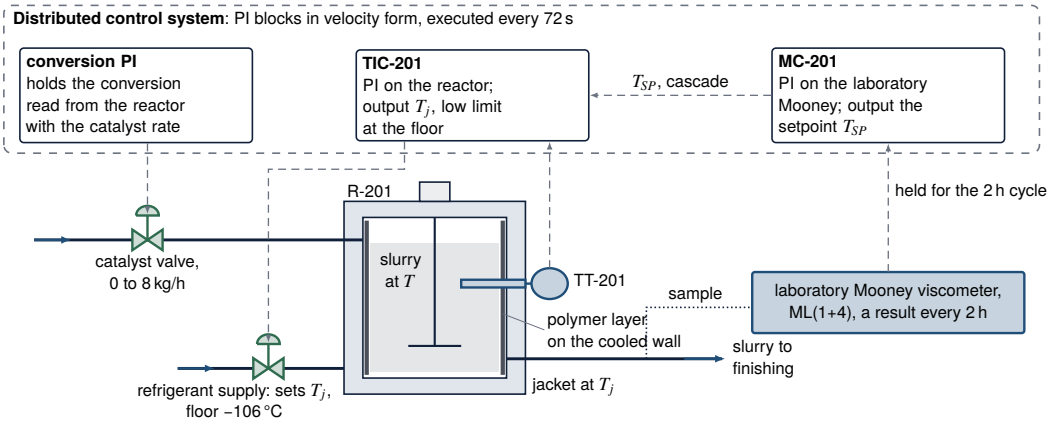

Figure 11.5 The loops of R-201 with their hardware. Blue marks are sensing devices: the temperature element and transmitter TT-201 and the laboratory Mooney viscometer. Green symbols are final elements: the refrigerant supply that sets the jacket temperature and the catalyst valve. Dashed arrows are signals. Notice that the only measurement of product quality reaches MC-201 from a laboratory, held for two hours, and that MC-201 acts only by moving the setpoint of TIC-201, whose output walks toward the −106 °C floor as the wall fouls. The feeds, the trip and the MPC are left out. Not to scale, and not a piping diagram.

> **IN PRACTICE** **Commissioning a Predictive Controller**
>
> An industrial predictive controller gets its model from a test on the plant. A pre-test steps each manipulated variable to adjust the instruments and PID controllers beneath it and to find each output's settling time. The plant test runs around the clock with engineers watching. Each handle is stepped 8 to 15 times, with an output signal-to-noise ratio of at least six, and the test can take 5 to 15 days. No PID configuration or tuning may change during the test, and operators may step in to avoid critical situations but make no correlated moves. If the PID tuning later changes much, the model may need rebuilding [11]. The chapter's closed-loop step tests on TIC-201 follow the same rule.

What Is Modernizing It

Three developments bear on this reactor. Soft sensors trained on data from the distributed control system and the laboratory estimate Mooney viscosity between laboratory results and can cut how often the laboratory must measure it [22]. In-line fouling sensors, applied in dairy plants, detect the deposit as it forms and help trigger cleaning on the actual buildup rather than on fixed schedules [23]. A 2003 survey counted polymers among the areas with the most nonlinear predictive control applications [11]. The cooling margin of this chapter is an inferred measurement of the deposit, and the MPC writes it into its plan as a constraint rather than a wash trigger. How far each development can be trusted is the subject of the Open challenge boxes in Section 11.8.

11.8 Limitations

The reactor is a single well mixed volume with a uniform fouling layer. Real fouling is not uniform, and the practical consequence is that a real plant loses circulation in one region before the average coefficient predicts it.

The fouling model is linear in time with an exponential temperature acceleration and no induction period, no asymptote and no removal term. That shape makes the countdown clean and it is simpler than the fouling literature supports.

Molecular weight is a correlation on reactor temperature rather than a calculation from chain transfer kinetics, so the distribution breadth, the gel content and the branching that a real butyl plant manages are all absent [3].

> **OPEN CHALLENGE** **Mooney Viscosity Without the Laboratory**
>
> In industrial rubber mixing, no reliable online instrument measures Mooney viscosity. The laboratory reports it every 4 to 6 hours, against batches that take 2 to 5 minutes, so an abnormal batch is found long after it was made. Soft sensors trained on data from the distributed control system could fill the gap, but laboratory results are too costly and slow to give them enough labeled data. Semi-supervised methods that also learn from the unlabeled process data are one answer under study [22]. Progress would be an estimate that a quality loop like MC-201 could close on between samples, and that stays right as the plant drifts.

The refrigeration floor is fixed. A real system loses capacity in summer and as its own surfaces foul, which means the countdown on a real plant has a second term this model does not carry.

The wash is a fixed eight hours restoring the surface exactly. A real hydroblast varies, and the fraction of the original coefficient recovered is one of the numbers a plant argues about.

> **OPEN CHALLENGE** **Washing on the Deposit, Not the Clock**
>
> Fouling of heated and cooled surfaces is unsolved well beyond rubber. A 2025 review of fouling in dairy heat treatment finds a notable gap in evaluating in-line measurement and prediction of fouling, in the laboratory and in plants. Computational flow models are still far from simulating, optimizing or scheduling whole heat exchangers over complete cleaning cycles, and the learned models that predict fouling well are hard to interpret, which slows their industrial use [23]. Progress would be a deposit sensor or model trusted enough to trigger a wash on the actual buildup rather than on a fixed schedule, the decision that scenario 8 shows the cooling margin cannot make alone.

Finally, the shipped structures deliberately do not include the multivariable layer that a real plant uses to arbitrate between production rate, quality and cooling margin [11]. Scenario 7 poses the question an override has to answer, and on this reactor the measured arithmetic favors holding production. No classical override is implemented, and the predictive option of Section 11.5.3 carries the margin as a constraint instead.

The predictive controller's fouling model is a trend of its own future: a straight line at the trend inferred over the last four hours, with only the first-order effect of the plan's temperature and production on the rate. It leaves out the exponential acceleration of Equation (11.3) along the plan, it freezes while the jacket is on its floor and refills after a wash, and it infers the coefficient from the duty at the design feeds, so a changed feed lands in the coefficient. The margin settles at 4.8 % rather than the 5 % asked for, 4.79 on scenario 10 and 4.78 on case 8, because each soft output carries one slack over the whole horizon and the funnel re-grants the present shortfall at every decision. The margin prediction is conservative while the plan cuts catalyst: a plan applied open loop for eight hours realizes up to 3.06 points more margin than it predicted, while its Mooney is within 0.04 units and its temperature within 0.16 °C. The plan does not decide the wash, and it loses cases 8 and 9 to the shipped structure.

11.9 Discussion

Three results generalize past butyl rubber.

The first is that a monotonically drifting controller output is a measurement. On this reactor the controlled variable is flat for two days while the plant runs out of cooling, and the only trend that says so is the jacket temperature. Trending controller outputs and alarming on their rate of change is cheap, and it is the difference between planning a shutdown and being surprised by one.

The second is that a fixed limit turns a monotone trade into an optimum. Colder gives a higher molecular weight and slower fouling, and the run length collapses anyway because the distance to the refrigeration floor is what the run is spending. Any time a decision variable improves two things at once, it is worth asking what fixed limit it is also consuming.

The third is that a schedule is an open-loop controller, and that replacing it with feedback helps only when the feedback measures what ends the run. The fixed wash at hour 44 is right for one fouling rate and lucky or wrong for every other. The cooling margin is already measured, but on this reactor it runs out about fifteen hours before the product leaves grade, so a wash triggered on it comes early and gives up 0.55 t/h against the schedule. Replacing a schedule with the measurement of the binding constraint it approximates is among the cheapest improvements available on a cyclic plant.

11.10 Conclusion

A butyl rubber reactor insulates itself with the product it is making. The simulated reactor of this chapter holds −95 °C to within 0.01 °C for two days while its overall coefficient falls from 2400 to 964 W/(m^2 K), its jacket walks from −99.4 °C to its −106 °C floor, and its cooling margin falls from 0.598 to zero at hour 49, so the run ends at an actuator stop rather than at a measurement. Colder operation is better for the molecular weight and for the fouling rate and worse for the plant: measured yields of on-specification rubber per hour of elapsed time are 1.25 t/h at −101 °C and 4.37 t/h at −93 °C, with the optimum near −94 °C and the binding constraint the Mooney grade of 46 to 56 rather than anything in the equipment. A wash scheduled at a fixed hour is correct only for the fouling rate it was written for, and a wash triggered on the cooling margin is not the cure on this reactor, because the margin runs out long before the product leaves grade. A fuller model would add non-uniform fouling, a refrigeration system with capacity of its own, and a molecular weight distribution model. A model predictive controller that keeps a predicted cooling margin of 5 % never reaches the floor or makes off-grade rubber in any case measured, yet it wins on yield only on the poisoned reactor and the challenge, loses the scheduled wash and the dirtier feed by at most 0.08 t/h, and owes most of its poisoned-reactor win to the margin row that a classical override also carries. The reactor is a place to learn to read a plant from its controller outputs.

KEY POINTS

- A butyl rubber reactor insulates itself with its own product, so its run ends when the jacket reaches its floor.
- The controlled variable stays put and the manipulated variable is the diagnostic: the cooling margin is the clock.
- Colder operation shortens the run, because the floor does not move with the setpoint. The optimum is near −94 °C, bounded by the Mooney grade.
- A fixed wash schedule suits only the fouling rate it was written for, and a margin wash comes about fifteen hours early.
- A predictive controller keeping a 5 % margin never reaches the floor, yet wins on yield only on the poisoned reactor and the challenge.

Where else this applies. A run spent against a limit is also the coke countdown of Chapter 2 and the fouling preheat train of Volume 11, Chapter 3.

References

[1] J. Tian, Y. Wang, and X. Feng, "Simultaneous optimization of flow velocity and cleaning schedule for mitigating fouling in refinery heat exchanger networks," *Energy*, vol. 109, pp. 1118–1129, 2016.

[2] R. M. Thomas, I. E. Lightbown, W. J. Sparks, P. K. Frolich, and E. V. Murphree, "Butyl rubber. a new hydrocarbon product," *Rubber Chemistry and Technology*, vol. 14, no. 1, pp. 177–195, 1941.

[3] J. L. White, T. D. Shaffer, C. J. Ruff, and J. P. Cross, "Incorporation of isoprene in isobutylene/isoprene copolymers: Nmr identification of branching in butyl rubber," *Macromolecules*, vol. 28, no. 9, pp. 3290–3300, 1995.

[4] R. M. Thomas, “Early history of butyl rubber. charles goodyear medal address—1969,” *Rubber Chemistry and Technology*, vol. 42, no. 4, pp. G90–G96, 1969.

[5] O. Sofronova, E. Markina, S. Chelnokova, and A. Sakhabutdinov, “The effect of alkyl halides on the stability of a butyl rubber suspension in methyl chloride,” *International Polymer Science and Technology*, vol. 40, no. 3, pp. 11–14, 2013.

[6] R. M. Thomas and W. J. Sparks, “Mixed olefinic polymerization process and product,” U.S. Patent 2,356,128, filed 20 October 1939 (continuation in part of an application filed 29 December 1937), granted 22 August 1944, 1944.

[7] I. E. Nifant’ev, S. A. Korchagina, M. S. Chinova, and A. N. Tavtorkin, “Polyisobutylenes with controlled molecular weight and chain-end structure: Synthesis and actual applications,” *Polymers*, vol. 15, no. 16, p. 3415, 2023.

[8] L. R. Gould and G. G. Wanless, “Method for washing butyl reactors with isobutylene,” U.S. Patent 2,580,019, filed 9 October 1946, granted 25 December 1951, 1951.

[9] M. G. Balaton, L. Nagy, and F. Szeifert, “Model-based split-range algorithm for the temperature control of a batch reactor,” *Engineering*, vol. 04, no. 09, pp. 515–525, 2012.

[10] 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.

[11] 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.

[12] European IPPC Bureau, “Reference document on best available techniques in the production of polymers,” European Commission, Joint Research Centre, Institute for Prospective Technological Studies, Seville, Tech. Rep., Aug. 2007.

[13] A. Kummer and T. Varga, “What do we know already about reactor runaway? a review,” *Process Safety and Environmental Protection*, vol. 147, pp. 460–476, 2021.

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

[15] 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.

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

[17] M. Ellis, H. Durand, and P. D. Christofides, “A tutorial review of economic model predictive control methods,” *Journal of Process Control*, vol. 24, no. 8, pp. 1156–1178, 2014.

[18] A. Aggarwal, F. Grunert, S. Ilisch, and A. Blume, “A quantitative analysis of influencing parameters for cure torque in silica-filled SSBR compounds,” *Polymers*, vol. 18, no. 20, p. 2596, 2026.

[19] H. Yangthong, P. Nun-Anan, A. Krainoi, B. Chaisrikhwun, S. Karrila, and S. Limhengha, “Hybrid alumina–silica filler for thermally conductive epoxidized natural rubber,” *Polymers*, vol. 16, no. 23, p. 3362, 2024.

[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.

[22] Y. Zhang, H. Jin, H. Liu, B. Yang, and S. Dong, “Deep semi-supervised just-in-time learning based soft sensor for Mooney viscosity estimation in industrial rubber mixing process,” *Polymers*, vol. 14, no. 5, p. 1018, 2022.

[23] J. Tarapata, T. R. Murphy, E. W. Finnegan, T. F. O’Callaghan, and J. A. O’Mahony, “Approaches for measuring and predicting fouling during thermal processing of dairy solutions,” *Comprehensive Reviews in Food Science and Food Safety*, vol. 24, no. 4, p. e70209, 2025.

11.11 Exercises

11.1 ○ Concept check · 2 min

In scenario 2 the temperature loop holds the reactor within 0.01 °C of −95 °C for two days while the surface fouls. Which reading tells the operator how many hours of the run are left?

(a) The jacket temperature against the −106 °C floor, or the cooling margin computed from it

(b) The reactor temperature deviation, which grows in step with the fouling resistance

(c) The laboratory Mooney result, because the molecular weight follows the heat transfer coefficient

(d) The integrated absolute error on the scorecard, which accumulates the fouling over the run

11.2 Concept check · 1 min

Fouling slows by half for every five degrees of cooling, yet in Table 11.3 the run at −99 °C is much shorter than the run at −95 °C. Why?

(a) Colder polymerization is faster, so the reaction duty rises enough to shorten the run

(b) The run spends the distance to a fixed refrigeration floor, and cooling takes that distance away faster than it slows the fouling

(c) The Mooney viscosity rises above its grade limit and the wash policy stops the run early

(d) The temperature loop saturates at its high limit and stops correcting the jacket temperature

11.3 Concept check · 2 min

In scenario 8 the feed fouls 1.6 times faster than the 44-hour wash schedule assumed. A wash on a cooling margin of 0.10 yields 3.99 t/h against the schedule's 4.54 t/h. What is the reason?

(a) The margin is computed from noisy measurements, so it triggers washes at random times

(b) The schedule washes later and makes no off-grade rubber, while the margin policy makes some

(c) The margin runs out well before the rubber leaves grade, so the margin policy washes early and pays for an extra wash

(d) The schedule happens to wash at the moment the margin reaches zero, which is the optimum

11.4 Theory and calculation · 6 min

Reactor R-201 is fed 6.0 t/h of isobutylene, 18 t/h of methyl chloride and 0.12 t/h of isoprene at $T_{feed} = -40\,^{\circ}\mathrm{C}$, and holds $T = -95\,^{\circ}\mathrm{C}$ at a conversion of 0.85. Use $\Delta H_p = 930\,\mathrm{kJ/kg}$ of polymer, $c_p = 1.9\,\mathrm{kJ/(kg\,K)}$, $A = 190\,\mathrm{m^2}$, $U_0 = 2400\,\mathrm{W/(m^2\,K)}$ and $T_{j,min} = -106\,^{\circ}\mathrm{C}$.

(a) Compute the heat of reaction, the feed-chilling duty and the total duty Q of Equation (11.2), in MW to three decimals. *Checkpoint:* 1.318 MW, 0.700 MW and 2.018 MW.

(b) Compute the temperature difference the duty requires on a clean surface, the jacket temperature of Equation (11.5) and the cooling margin of Equation (11.6). *Checkpoint:* 4.42 K, −99.42 °C and 0.598.

(c) Compute the overall coefficient at which the margin reaches zero with the same duty, and compare it with the 964 W/(m^2 K) the chapter reports when the jacket reaches its floor. *Checkpoint:* about 965 W/(m^2 K).

11.5 Theory and calculation · 7 min

Use Equation (11.9) with the values of the previous exercise, $k_f = 1.29 \times 10^{-5}\,\mathrm{m^2\,K/(W\,h)}$, $c_f = 0.1386/^{\circ}\mathrm{C}$, $P = 5.10\,\mathrm{t/h}$, $P_0 = 5.2\,\mathrm{t/h}$, $\phi = 1$ and $T_{ref} = -95\,^{\circ}\mathrm{C}$. The run starts from a clean surface and the reactor is held at one temperature throughout.

(a) Compute the prefactor $AP_0/(Qk_fP\phi)$ and the run length at −95 °C. *Checkpoint:* 7.44 h per kelvin and 48.9 h.

(b) Repeat at −93 °C and at −101 °C, with the duty recomputed for the feed chilled to each temperature. *Checkpoint:* 1.992 MW and 49.3 h, then 2.094 MW and 6.7 h.

(c) Holding Q at its design value, find the temperature that maximizes t_{run}, and check it against the Mooney grade window of 46 to 56, with ML = 51.0 at −95 °C and $k_{ML} = -1.40$ per degree. *Checkpoint:* −94.4 °C, inside the window −98.6 to −91.4 °C.

11.6 Theory and calculation · 6 min

The reactor holds 30 m^3 of slurry at 700 kg/m^3 with $c_p = 1.9\,\mathrm{kJ/(kg\,K)}$. When the jacket reaches its floor the coefficient is 964 W/(m^2 K) on $A = 190\,m^2$, and 24.12 t/h of feed is chilled from −40 °C.

(a) Compute the heat capacity of the contents and the rate at which a sustained shortfall of 0.1 MW would warm them. *Checkpoint:* 39.9 MJ/K and 9.0 K per hour.

(b) Each degree of warming raises the duty the floor removes by UA and lowers the feed-chilling duty. Find how far the reactor warms before a 0.1 MW shortfall is absorbed, and the time constant of that approach. *Checkpoint:* 0.51 K and about 3.4 min.

(c) Before the floor, TIC-201 with $K_c = 4.0$ and $\tau_I = 0.30\,h$ moves the jacket down a ramp of 0.134 K per hour. Compute the steady error of the PI loop on that ramp. *Checkpoint:* 0.010 °C.

11.7 App challenge · 8 min

Load scenario 2, The temperature loop, and what it is spending, and change no setting. Run a little past hour 48 at speed 20 with the default seed and export the CSV file. The time column `time_h` is in hours and a row is logged every 0.035 h.

(a) Read column `margin` in the first row and in the last row at or before hour 48 (at 47.985 h), and compute its mean slope per hour. *Checkpoint:* 0.598 and 0.012, a slope of −0.0122 per hour within 0.0003.

(b) Find the first row in which `margin` is below 0.10, and extrapolate the straight line of part (a) to zero. *Checkpoint:* hour 40.8 within 0.1 h, and about hour 49.0.

(c) Find the largest absolute difference between columns `temp_C` and `temp_sp_C`, and explain its size. *Checkpoint:* 0.010 °C within 0.002 °C.

11.8 App challenge · 9 min

Load scenario 4, Temperature does three things, and they disagree. A script steps the reactor setpoint from −95 to −99 °C at hour 2, and the wash starts when the cooling margin falls below 0.03. Use speed 5 and the default seed, and export the CSV file after each run.

(a) Change no setting and run until the wash has started. In the last CSV row before column `washing` becomes 1, read `run_hours`, `metric_onspec_t` and `mooney`. Compute the tonnes on specification per elapsed hour, counting the run and its eight-hour wash. *Checkpoint:* 27.0 h, 17.8 t, a Mooney of 56.6, and 0.51 t/h within 0.02.

(b) Press Reset. Run past hour 2 until the step to −99 °C has appeared on the trend, pause before hour 3, set Reactor temperature setpoint to −93.0 °C, and run until the wash has started. Read the same three columns and compute the same yield. *Checkpoint:* a run of 47.6 to 48.2 h, 242.7 to 245.6 t, a Mooney of 48.2, and 4.37 t/h within 0.01.

Explain why the colder run is shorter and why almost none of its rubber is on specification.

11.9 App challenge · 8 min

Load scenario 8, The clock against the margin. The fouling tendency steps to 1.6 at hour 2, and the wash is on the clock at 44 h. Use speed 20 and the default seed. Run each case a little past hour 96, export the CSV file, and read it in the last row at or before hour 96 (the row at 95.97 h). The yield is the tonnes on specification divided by the 96 elapsed hours.

(a) Change no setting. Read columns `metric_onspec_t`, `metric_offspec_h` and `downtime_h`, compute the yield, and find the first hour at which column `margin` is zero or below. *Checkpoint:* 436.3 t, 0.9 h, 8.0 h, 4.54 t/h within 0.01, and hour 31.3 to 31.4.

(b) Press Reset, set When to wash to On margin before hour 2, keeping the threshold at 0.10, and run again. Read the same three columns, compute the yield, and find the hours at which column

`washing` turns to 1. *Checkpoint:* 383.4 t, 0.0 h, 17.8 h, 3.99 t/h within 0.01, washes at about 26.2, 60.2 and 94.2 h.

State which measurement the schedule was standing in for.

11.10 Application area · 3 min

The Limitations section notes that a real refrigeration system loses capacity in summer. Suppose the floor rises from −106 to −105 °C while everything else stays at the design point of the earlier exercises.

(a) Use the prefactor of 7.44 h per kelvin to find the new run length. *Checkpoint:* 41.5 h, 7.4 h shorter.

(b) Use Equation (11.10) with the eight-hour wash to compare the yield before and after, for runs that stay in grade. *Checkpoint:* 4.38 against 4.28 t/h.

11.11 Application area · 3 min

Section 11.2 argues that this reactor is not a runaway problem, and Section 11.5 calls the −80 °C trip an illustrative threshold eight degrees above the warm-slurry alarm at −88 °C.

(a) State what does not accelerate with temperature on this reactor, and what does. *Checkpoint:* the heat release does not, the fouling does, doubling for every five degrees.

(b) In scenario 3 the creep reaches 0.29 °C per hour. At that rate, how long would the reactor take from the alarm to the trip? On the challenge as shipped, the reactor passes −88 °C at 58.2 h and trips at 68.5 h. What does the difference say about warning time? *Checkpoint:* about 28 h by extrapolation against 10.3 h measured.

11.12 Review the agent · 5 min

A learner asks an AI assistant to make the Mooney cascade of scenario 6, The Mooney cascade, respond faster to the two-hour laboratory result. The assistant replies: “The result is two hours old, so the outer loop must act harder on each one. Raise Mooney loop gain from 0.30 to 1.0 and cut Mooney integral time from 8 h to 2 h. The Mooney will then be corrected within one laboratory cycle.” Find the flaw, and say how you would check the proposal in the app before using it.

CHAPTER 12

Loss-in-Weight Additive Feeders and the Holdup Their Scales Cannot Weigh

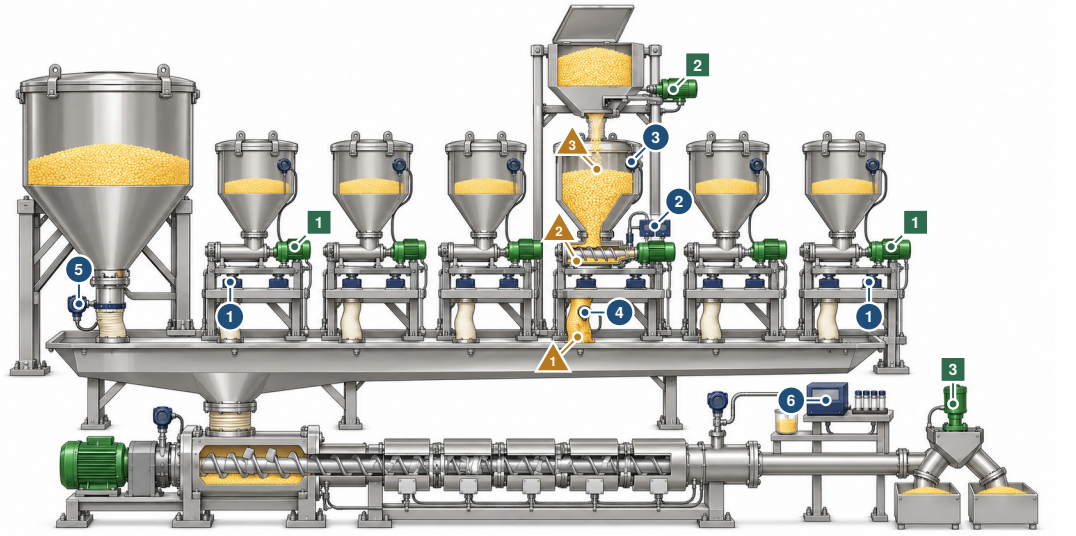

The controller moves 1 speed of each feeder's screw 2 refill valve above the hopper 3 diverter to prime or to the diversion bin

The controller reads 1 load cells under each feeder 2 screw speed and motor torque 3 hopper pressure 4 sock high-level switch 5 resin flow 6 laboratory assay of a sample, hours late

What disturbs it 1 powder held in the sock below the scale 2 buildup where the screw discharges 3 a refill, which can pressurize the hopper

Notice that each discharge sock hangs below its feeder's load cells, so powder held in it has left the weighed feeder without reaching the extruder. Illustration, not to scale.

IN PLAIN TERMS **A Scale That Cannot See the Last Step**

The main idea. Plastic pellets get small doses of powders that protect the plastic and help it work. Each powder sits in a hopper standing on a scale. A screw turns faster or slower so that the weight drops at the right rate, even when the screw starts to deliver less on each turn.

The challenge. The scale weighs only the hopper. Powder can pile up in the soft sleeve below it, so the pellets run short while every screen reads normal, and the laboratory finds out hours later.

Why it matters. A scale that reads right does not prove that the pellets are right.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? box in Section 12.1, the History box in Section 12.2, the Try it boxes in Sections 12.3 and 12.4, the History box in Section 12.5, the Did you know? box in Section 12.6, and Key points at the end of the Conclusion. No equations.
Study	45 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 12.6 with Figure 12.3, Section 12.7, and Section 12.8 with its Open challenge box.

The control question A loss-in-weight feeder holds the flow its scale weighs, but the product sees only what leaves the discharge sock below the scale and a laboratory reports it hours late. Which faults does the weighing loop correct, which can it not see, and what do a laboratory trim, a moving horizon estimator and a predictive controller add?

Learning objectives

- Convert a recipe in ppm of finished compound into feeder flows for neat and masterbatch ingredients, and account for the carrier in the resin rate.
- Compute the speed a feed-factor loss demands and the time a sock restriction takes to trip the high-level switch.
- Explain why the sock holdup is not observable from hopper weight, and name the measurements that reveal it.
- Decide when a delayed assay may move the laboratory trim, judged against its own collection interval and lot.
- Compare ratio control, MPC, the simple estimate and the MHE on matched runs, and state where each changes the product.

Prerequisites Sections 1.4.5 (grade changes) and 1.4.6 (laboratory bias update), 1.6.2 (velocity-form PID), 1.6.3 (MPC and MHE) and 1.6.4 (scenarios and the exported file).

The app https://apmonitor.com/apps/weigh-feeder/

See also Volume 5, Chapter 6 for loss-in-weight feeders with refill feedforward ahead of a tablet press, and Volume 5, Chapter 7 for diversion timed by residence. Volume 3, Chapter 6 for a dose computed from an upstream measurement, with the slow downstream measurement demoted to a ratio trim.

12.1 Introduction

A specification of 1,000 ppm of a phosphite antioxidant is one kilogram per tonne of compound, so a compounding line making 1,000 kg/h of pellets needs 1 kg/h of that powder, or 5.6 g every twenty seconds. The feeder that delivers it is usually a loss-in-weight (LIW) screw feeder, a hopper, motor and screw standing on load cells whose flow is taken as the rate at which the weighed assembly loses mass [1, 2]. The same principle meters active ingredients and excipients into continuous tablet lines [3, 4], the subject of Volume 5, Chapter 6, and the additive packages it meters into polyolefins enhance the polymer's properties and prolong its life [5]. A published polypropylene composition carries 0.1 % of a phosphite antioxidant, 0.05 % of a hindered phenolic antioxidant, 0.11 % of calcium stearate and 0.2 % of a sorbitol clarifier, compounded in a twin-screw extruder [6], and the simulator's grades sit at those magnitudes.

> **DID YOU KNOW? Fractions of a Percent**
>
> A 2018 review of additives in plastics gives typical contents of 0.05 to 3 % by weight for stabilizers and antioxidants, and 0.1 to 3 % for slip agents. Phenolic antioxidants are used in low amounts and phosphites in higher ones, and polyolefins take the lowest amounts [5].

This simulator targets a single belief, that a feeder holding its weighed flow on target holds the product on specification. The scale weighs the hopper, the motor and the screw. Whatever leaves the screw and stays in the flexible discharge sock below it has left the weighed system without reaching the product. In case 4 of the app the sock narrows to a fifth of its healthy capacity. The phosphite feeder's LIW reading then averages 1.00 kg/h from 15 to 30 min and still reads 1.03 kg/h when a level switch trips at 36.6 min, while the phosphite at the sample tap sits at 600 ppm against 1,000 ppm. The class of problem is general. A flow measured upstream of an unmeasured inventory, a product property seen only by a laboratory, and a model updated from infrequent assays describe polymer reactors [7] as well as feeders, and grade changes on such plants are judged by the off-specification material they leave [8], as Section 1.4.5 sets out and Chapter 8 measures on a fluidized bed.

However, the feeder literature is mostly pharmaceutical and treats the feed rate as the controlled quantity, with refill transients, feed-factor prediction and the coordination of several feeders as its subjects [4, 9, 10], while polymer quality control treats the reactor rather than the additive lanes. This chapter measures what ratio control, a laboratory trim, a short moving horizon estimator (MHE) and a supervisory model predictive controller (MPC) each add to the product, on a browser simulator that records each lane's weighed flow, its delivered flow, the off-scale holdup and the product concentration together. The balances, the proportional-integral (PI) loop, the MHE [11, 12] and the MPC [13] are standard methods. The results are measurements on one teaching plant, and three of them are negative findings, that the off-scale holdup is not observable from hopper weight, that the MHE is quieter than the simple estimate but not faster, and that the MPC's measured product benefit is confined to grade transitions. The model's total mass balance closes to 10^{-13} of the input, the recipe solver reports infeasible ingredient sets instead of clipping them, the feeder MHE solves a five-minute window in 0.2 ms, and no controller or estimator crosses the measurement boundary to read the simulator's truth. Ten cases run from recipe arithmetic to a graded twelve-hour shift.

Section 12.2 follows the domain, from gravimetric feeding to delayed laboratory assays. Section 12.3 builds the line, its instruments, the recipe and the laboratory. Section 12.4 works out by hand what the weighing boundary implies for the two restrictions and for washout. Section 12.5 describes the ratio and LIW loops, the two estimators, the laboratory trim, the supervisor and the MPC. Section 12.6 measures the ten cases and compares the four control architectures on matched runs. Section 12.7 names the hardware and the practice behind the model. Sections 12.8 and 12.9 state the limitations and the conclusions, and the Exercises put the cases to work.

12.2 Weighing Powder, Dosing Additives and Waiting for the Laboratory

A screw feeder delivers a volume of powder per revolution, so its mass flow moves with bulk density, hopper level and anything that changes how the screw fills. Gravimetric feeding answers this by weighing. In the loss-in-weight arrangement the hopper, the drive and the screw stand on load cells, the controller differentiates the net weight to obtain the flow, and it trims the screw speed to hold a setpoint [1]. The weight is a derivative measurement, so the flow estimate needs a window and is only as good as the scale beneath it. A hopper cannot be weighed while it is being refilled, which forces the feeder to run on its learned speed until the weight settles, and refills have been shown to cause measurable feed-rate deviations [9]. Characterization methods for these feeders measure their feed-rate variability and the feed factor, the delivery per revolution, as the hopper empties [2], and the feed factor has been predicted from the properties of the powder [10].

> ⧖ HISTORY **A Feeder Standing on a Scale, 1945**
>
> In a patent application filed in March 1945, a weigh hopper discharged through a star feeder driven by a variable-speed motor. The hopper, the feeder, the motor and the other accessories all stood on a scale platform as tare weight, so the scale showed the net weight of material left in the hopper [14]. A later application described how such systems held a rate. A counterpoise, moved along the scale by a motor-driven lead screw, set a weight that fell at the desired rate, and any unbalance of the beam adjusted the flow. The reservoir had to be refilled by hand and the scale balanced again after every emptying, and a change in density could set the flow oscillating about its target [15].

Continuous pharmaceutical manufacturing made the feeder the first unit of a line whose composition matters at every instant [3]. One study treated a set of LIW feeders as a control system in its own right [4], and a hybrid of MPC and PID control was implemented on a direct compaction pilot plant [16]. Polymer compounding has the same structure at larger scale and with sparser measurement. Additives of these kinds are found throughout plastic products [5], and assays of a stabilizer in a polymer commonly begin by extracting the additive from the matrix [17], so the product's composition is known from a laboratory rather than from an online instrument.

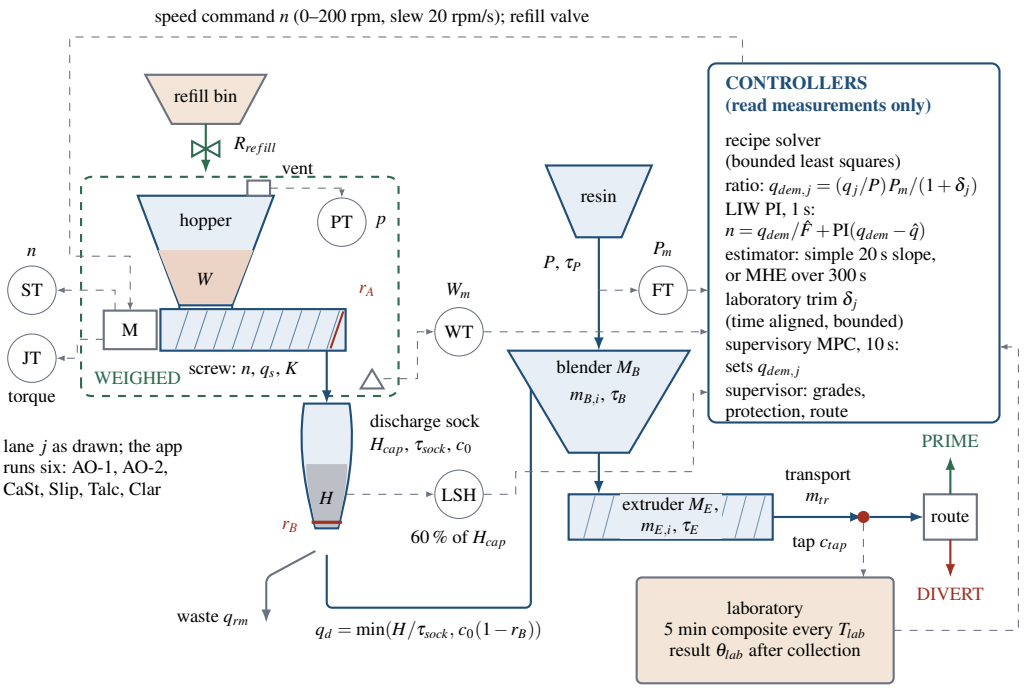

Figure 12.1 Additive compounding line as modeled. One lane is drawn in full: the hopper weight W, the screw at speed n with its flow q_s, the feed-factor loss r_A of a restriction that backs into the screw, the vent and hopper pressure p, and the refill R_{refill} lie inside the dashed weighing boundary, while the sock holdup H, its restriction r_B and its delivered flow q_d lie outside it. The resin P, the blender (M_B, τ_B), the extruder (M_E, τ_E), the transport m_{tr} to the tap c_{tap}, the laboratory and the diverter follow. Instruments: WT net weight W_m, ST speed, JT torque, PT pressure, LSH sock high level, FT resin flow. Not to scale.

Infrequent laboratory measurements were brought into polymer control through models whose parameters are updated each time a result arrives [7], the bias update of Section 1.4.6, and dynamic optimization of grade changes showed how much off-specification polymer a transition policy costs [8]. Moving horizon estimation, a least-squares fit of a model to a sliding window of data with constraints on states and disturbances [11, 12] (Section 1.6.3), is the estimation counterpart of the receding-horizon controllers that became standard in the process industries [13]. A robust penalty on the residuals in the sense of Huber [18] limits the influence of a single bad weight reading on the estimate. The simulator puts these pieces on one compounding line and asks which of them the product notices.

12.3 The Compounding Line as Modeled

The simulated line makes 1,000 kg/h of finished polyolefin compound, selectable from 500 to 2,000 kg/h, as shown in Figure 12.1. Base resin enters through its own gravimetric feeder, the master stream. Six additive lanes, one per active, each comprise a weighed hopper, motor and screw on load cells, a refill bin above, and a discharge sock below the weighing boundary. The lanes carry a phenolic antioxidant of the AO-1010 type (AO-1), a phosphite antioxidant of the AO-168 type (AO-2), calcium stearate (CaSt), an erucamide slip agent, a fine talc antiblock and a sorbitol clarifier (Clar). Everything discharges into a blender of 20 kg that feeds an extruder of 30 kg, and 16.7 kg of product downstream a sample tap feeds the laboratory before a diverter sends the product to prime or to a diversion bin.

Table 12.1 Variables by Role for the AO-2 Lane at the Design Point (PP-A grade)

Variable	Role	Design	Range
$c_{tap,i}$	CV, product at tap (lab only)	1,000 ppm	band ±10 %
q_{dem}	CV of the LIW loop, flow	1.000 kg/h	0–2.5 kg/h
n	MV, auger speed	80 rpm	0–200 rpm
$q_{dem,j}$	MV of the MPC, lane demand	1.000 kg/h	0–0.95$\hat{F}n_{max}$
Q	DV, production setpoint	1,000 kg/h	500–2,000 kg/h
r_A	DV, feed-factor loss	0	0–60 %
r_B	DV, sock restriction	0	0–100 %
s	DV, scale span error	0	±10 %
x_{MB}, c_{res}	DV, potency, resin AO-1	5 %, 0	3–6, 0–800 ppm
W	state, hopper	1.5 kg	0–3 kg
H	state, sock (unmeasured)	2.8 g	0–300 g
a, p	states, aeration, pressure	0, 0	0–0.06, Pa
P	state, resin flow	997.5 kg/h	—
M_B, M_E	states, mixing masses	20, 30 kg	—
W_m, torque	measurements	1.5 kg, 25 %	—

Table 12.1 assigns each variable its role. Per lane five states are integrated, the hopper weight W, the sock holdup H, the feed-factor loss r_A, the aeration a after a refill and the hopper pressure p. The line adds the resin flow P, the blender and extruder masses with six active masses each, and ledgers of resin and refill in and of product and waste out, so that conservation can be checked to round-off. The model integrates in seconds, and flows written in kg/h are divided by 3,600 where they enter a derivative. Concentrations are mass fractions of the finished compound, displayed in mass ppm, with $1\,\% = 10{,}000\,\text{ppm}$.

12.3.1 Hopper, Screw and Two Restrictions

The weighed hopper loses what the screw takes and gains what a refill adds, as shown in Equation (12.1).

$$\begin{aligned} \frac{\mathrm{d}W}{\mathrm{d}t} &= R_{refill} - \frac{q_s}{3600} \\ q_s &= K\,n\,f(W)\,(1-a)\,(1-r_A)\,\mathrm{acc}(H) \\ f(W) &= \left[1+\beta\left(\frac{W}{W_{max}}-\frac{1}{2}\right)\right]\min\left(1,\frac{W}{0.1\,W_{max}}\right) \end{aligned} \tag{12.1}$$

where K is the commissioned feed factor in kg/h per rpm, n the auger speed, $\beta = 0.08$ the hopper-level effect (±4 % from empty to full), and the second factor of f starves the screw below 10 % of the hopper. The lanes are sized so that the design rate $q_{des} = 80K$ sits at 80 rpm, with $K = 0.00625$, 0.0125, 0.0125, 0.0125, 0.0625 and 0.025 kg/h/rpm for AO-1, AO-2, CaSt, slip, talc and clarifier. The AO-2 value delivers 1 kg/h at 80 rpm, the example that opens the chapter. The feed-factor loss r_A represents a restriction that backs into the screw, so it is inside the weighing boundary, and it ramps toward its disturbance value at a buildup rate of 25 % per 30 min and clears with a 5 s time constant. The backpressure factor $\mathrm{acc}(H)$ is 1 below 70 % of the sock capacity H_{cap} and falls by a smoothstep to 0 at H_{cap}, so a full sock stops the screw instead of losing mass.

12.3.2 The Discharge Sock Below the Scale

The sock receives the screw flow and releases it with an availability-limited outflow, as shown in Equation (12.2).

$$\begin{aligned} \frac{\mathrm{d}H}{\mathrm{d}t} &= \frac{q_s - q_d - q_{rm}}{3600} \\ q_d &= \min\left(\frac{3600\,H}{\tau_{sock}},\; c_0\,(1-r_B)\right), \qquad c_0 = 3\,q_{des} \end{aligned} \tag{12.2}$$

where $\tau_{sock} = 10\,\mathrm{s}$ is the transit time of a healthy sock, c_0 its capacity, and r_B the sock restriction, with 1 a full plug. The outflow can never exceed what is in the sock, so H stays nonnegative when material enters an empty sock. Cleaning locks the lane out for 300 s, sets $q_d = 0$ and sends the holdup to waste at $q_{rm} = 3600H/30\,\mathrm{s}$, while knocking the sock clear removes r_B so that its contents fall into the blender at the healthy capacity c_0. Neither action deletes mass. The microfeeder socks hold $H_{cap} = 0.30\,\mathrm{kg}$ and the talc and masterbatch socks 2 kg.

12.3.3 Refill, Aeration and Hopper Pressure

The feeder's own sequence opens the refill valve when the measured weight falls below 25 % of the hopper and closes it at 85 % or after 40 s, with $R_{refill} = 0.6W_{max}/20\,\mathrm{s}$. Fresh powder is aerated and the displaced air must leave through the vent filter, as shown in Equation (12.3).

$$\frac{\mathrm{d}a}{\mathrm{d}t} = \frac{0.06\,R_{refill}}{0.6\,W_{max}} - \frac{a}{\tau_{aer}}, \qquad \frac{\mathrm{d}p}{\mathrm{d}t} = u_{refill} - \frac{p}{\tau_v} \tag{12.3}$$

where u_{refill} is 1 while the valve is open and 0 otherwise, p is in pascals, $\tau_{aer} = 120\,\mathrm{s}$, a full refill costs up to 6 % of the delivery per revolution, and the vent time constant τ_v is 2 s when healthy and 90 s with the filter plugged. A plugged filter therefore leaves about 18 Pa in the hopper at the end of a 20 s refill. The term represents hopper pressure reaching the load cells through the flexible connections, a declared teaching relation.

> **TRY IT IN THE APP** **6 · Refill With a Plugged Vent Filter**
>
> Run 15 min on the Feeder view and watch the AO-2 speed after the refill that starts at 220 s. With the timer alone the speed falls to 65.1 rpm and AO-2 at the tap dips to 911 ppm. Set Weight validity to Timer + pressure, press Reset and run again: the speed stays at 75.2 rpm or above and the tap holds 995 to 1,021 ppm.

12.3.4 What the Instruments Read

Nothing reaches the controllers except the instruments of Equation (12.4), each read once per second.

$$\begin{aligned} W_m &= (1+s)\,(W + \alpha H) + k_p\,p + v_W \\ \tau_q &= 15 + 25\,\frac{n}{200} + 40\,\frac{H}{H_{cap}} + 20\,r_A + d_\tau + v_\tau \\ \mathrm{LSH} &= [H > 0.6\,H_{cap}] \end{aligned} \tag{12.4}$$

where s is a span error, $\alpha = 0$ the share of the sock carried by the scale (a commissioning value), $k_p = 0.5\,\mathrm{g/Pa}$ the pressure coupling, v_W colored noise of 0.2 g (1 g on the talc and masterbatch lanes) with a 1 s correlation time, τ_q the motor torque in percent with a slow drift d_τ of 2 % over 2 h and noise of 1.5 %, and LSH a discrete sock high-level switch. The drive's speed, the hopper pressure (σ 0.2 Pa), the refill valve, the maintenance log of cleanings and knocks, the resin flow (σ 0.3 %) and the blender level (σ 0.2 kg) complete the packet. The sock holdup itself is measured by nothing.

12.3.5 Resin, Blender, Extruder and Tap

Ingredient j carries A_{ij} kilograms of active i per kilogram, 1 for a neat additive, 0.05 for the slip masterbatch whose other 0.95 is carrier polymer, and 1/3 and 2/3 for the optional antioxidant premix. The blender and the extruder are level-regulated ideal mixers, as shown in Equation (12.5).

$$\begin{aligned} Q_{in} &= P + \sum_j q_{d,j}, \qquad \dot{m}_{in,i} = P\,c_{res,i} + \sum_j A_{ij}\,q_{d,j} \\ \frac{\mathrm{d}m_{B,i}}{\mathrm{d}t} &= \frac{\dot{m}_{in,i}}{3600} - \frac{m_{B,i}}{\tau_B}, \qquad \frac{\mathrm{d}m_{E,i}}{\mathrm{d}t} = \frac{m_{B,i}}{\tau_B} - \frac{m_{E,i}}{\tau_E} \\ c_{tap}(t) &= c_E\left(M_{prod}^{-1}(M_{prod}(t) - m_{tr})\right) \end{aligned} \tag{12.5}$$

where the total masses follow the same pair with Q_{in}, $\tau_B = 72\,\text{s}$ and $\tau_E = 108\,\text{s}$ correspond to 20 and 30 kg at 1,000 kg/h, and the tap reads the extruder outlet concentration of the product that left $m_{tr} = 16.7\,\text{kg}$ earlier, so its delay follows the throughput. The resin follows its command with $\tau_P = 15\,\text{s}$, and its command moves at most 20 kg/h per second.

12.3.6 The Recipe and Its Feasibility

For a finished-compound demand Q, targets c_i and a declared resin background $c_{res,i}$, the solver seeks nonnegative ingredient flows that satisfy Equation (12.6).

$$\begin{aligned} P + \sum_j q_j &= Q, \qquad q_j \geq 0 \\ P\,c_{res,i} + \sum_j A_{ij}\,q_j &= c_i\,Q \quad \forall i \end{aligned} \tag{12.6}$$

Eliminating P leaves one row per active, and the rows are scaled by each active's tolerance and solved as a bounded least-squares problem with Goldfarb and Idnani's dual active-set method [19]. The recipe is feasible only when every residual is below $10^{-9}Q$. Otherwise the solver names the active and the reason, a premix whose fixed ratio differs from the targets or a resin that already carries more than the target, and never clips one ingredient on its own. For PE-B at 1,000 kg/h it returns 0.5, 1, 0.5, 1 and 5 kg/h of neat actives and 992 kg/h of resin. With a 5 % slip masterbatch the slip lane carries 20 kg/h, of which 19 kg/h is carrier, and the resin falls to 973 kg/h. Both agree with hand arithmetic to 3×10^{-12} kg/h.

12.3.7 The Laboratory

A five-minute composite is drawn from the tap every $T_{lab} = 3\,\text{h}$ (2 to 4 h) from 1 h, weighted by the product mass that passes it, and reported $\theta_{lab} = 45\,\text{min}$ (20 to 90 min) after collection ends. Each species carries a relative noise of 3 % for the chromatographic assays, 4 % for calcium stearate assayed as calcium and 5 % for talc assayed as ash, with a 10 ppm floor and a 20 ppm reporting limit. A scheduled event can lose a sample or multiply the next result by 1.35. Each record keeps its collection interval and the grade requested when collection began, so a late result belongs to its own lot.

12.3.8 Parameters, Initialization and Numerics

The values of Table 12.2 are teaching assumptions sized for the line's operating envelope, except where a source is named. The line starts lined out on the controller's recipe, with the declared potency and declared resin background, while the product already carries the true composition of those flows. Hoppers start half full, socks at their healthy holdup $q\tau_{sock}$, and the blender and extruder at their steady masses. With this start the derivatives of the sock, blender and extruder states are zero, and a hidden fault reaches the controllers only through a measurement. The engine of Section 1.6.1 takes fourth-order Runge–Kutta steps of 1 s, runs the controllers every second, the MPC every 10 s and each MHE every 5 s, and logs a row every 10 s. Refills, the sock-full mark, cleaning, sampling and reports are events at step boundaries with historian marks, and the derivative function is pure.

12.3.9 Assumptions

The actives are conserved through the extruder and the assay measures the intact parent compound, so a low phosphite result is always a dosing error here, which on a real line it need not be. The blender and extruder are ideal mixers with constant residence, with no segregation and no axial dispersion in the transport. Calcium stearate is assayed as calcium and talc as ash through exact mappings with larger noise. A grade change between polypropylene and polyethylene requires a resin purge that is not simulated, and the grade selector refuses it. The torque relation, the pressure coupling, the sock capacity and the bands are declared teaching relations, not supplier correlations or standards, and a real grade's bands come from its customer specification and its assay validation.

Table 12.2 Parameters (V visible, A adjustable, H hidden truth)

Symbol	Meaning	Value	
K	feed factor, AO-2	0.0125 kg/h/rpm	V
W_{max}	hopper, micro (talc, MB)	3 (20, 25) kg	V
β, starve	level effect, starving	0.08, 10 %	V
τ_{aer}, a_{max}	aeration after refill	120 s, 0.06	V
τ_{sock}	sock transit	10 s	A
c_0	healthy sock capacity	$3q_{des}$	A
H_{cap}	sock capacity, micro (large)	0.30 (2) kg	V
$\dot{r}_A$	buildup rate	25 %/30 min	A
α	sock load on scale	0	A
k_p	pressure coupling	0.5 g/Pa	A
τ_v	vent, healthy (plugged)	2 (90) s	A
σ_W	load-cell noise, micro (large)	0.2 (1) g	V
τ_P	resin feeder	15 s	A
τ_B, τ_E	blender, extruder residence	72, 108 s	A
m_{tr}	product, extruder to tap	16.7 kg	A
T_{lab}, θ_{lab}	sample interval, turnaround	3 h, 45 min	A
σ_{rel}, σ_0	assay noise, floor	3 %, 10 ppm	A
bands	nonzero target, zero target	$\pm$10 %, $\leq$ 50 ppm	V
x_{MB}	masterbatch potency	5 %	H
c_{res}	resin AO-1 background	0 ppm	H

12.4 What the Weighing Boundary Implies

The two restrictions sit on opposite sides of the scale and have opposite consequences. A feed-factor loss r_A removes delivery per revolution from the weighed flow, so the scale sees it and speed restores it. At a half-full hopper the AO-2 lane needs $1/(0.0125 \times 0.75) = 106.67$ rpm for 1 kg/h after a 25 % loss, 33.3 % more than the healthy 80 rpm, and the 200 rpm limit allows losses up to 60 %. A sock restriction r_B removes nothing from the weighed flow until the sock pushes back. Its capacity $3(1 - r_B)q_{des}$ falls below the demand once r_B exceeds two thirds, and no speed can then deliver the demand, while a full plug delivers nothing at any speed. At 200 rpm with the sock plugged the delivered flow after an hour is 0.0000 kg/h and the sock holds 299 g of its 300 g.

The partial sock restriction of case 4 can be followed by hand. The sock starts at $1\,\text{kg/h} \times 10\,\text{s} = 2.8\,\text{g}$ and fills at $1 - 0.6 = 0.4\,\text{kg/h}$, so it reaches the switch point of 180 g after 26.6 min, at $t = 36.6$ min. The engine trips the switch at 36.62 min. Over the 15 to 30 min window the LIW reading averages 0.9998 kg/h while the sock delivers 0.600 kg/h, and the phosphite at the tap is 608 ppm at 20 min. Figure 12.2(a) shows the three flows parting at 10 min and joining again only after the sock passes 70 % of its capacity, when backpressure throttles the screw. Released, the holdup is a slug. With the speed fixed and the sock restricted 80 % from 600 s, the screw still leaves the hopper at 0.98 to 0.99 kg/h from 1,200 to 2,400 s while 0.600 kg/h is delivered, the sock holds 247.9 g when it is knocked clear at 3,990 s, and the phosphite at the tap peaks at 2,871 ppm.

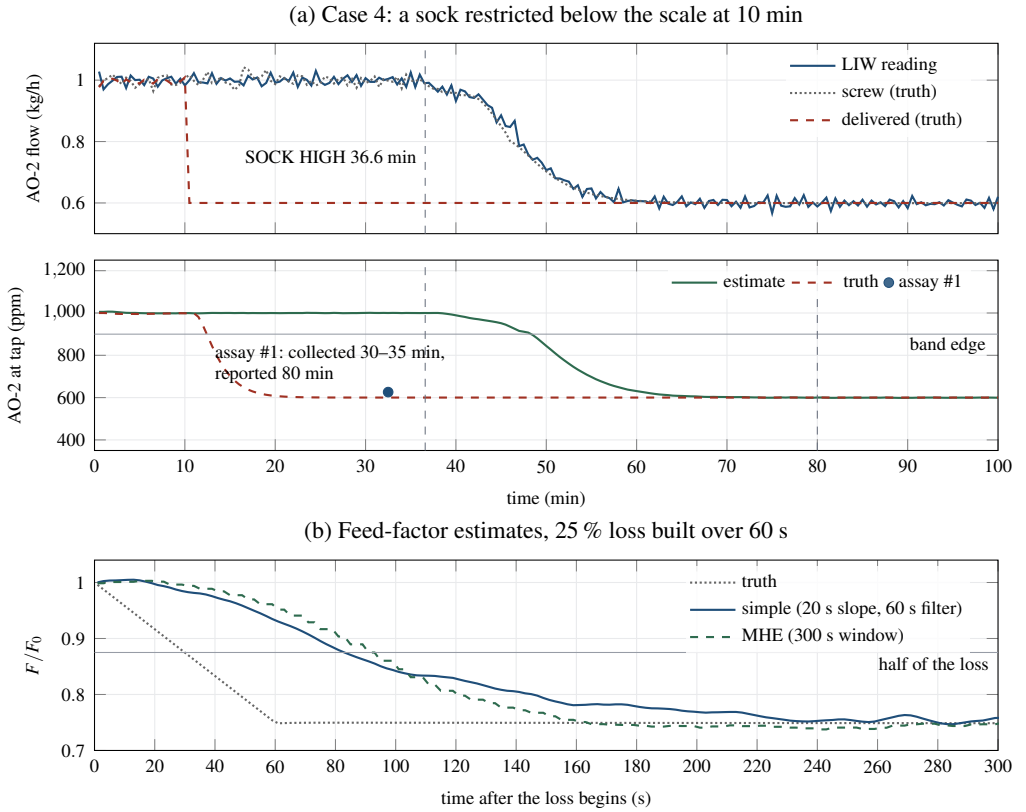

Figure 12.2 Two blind spots of the weighing loop. (a) Case 4 as shipped: the LIW reading follows the screw while the delivered flow and the phosphite at the tap fall, and the estimate stays near target until backpressure throttles the screw and pulls it to 599 ppm against a true 600 ppm by 100 min. The assay reported at 80 min is accepted as evidence of holdup and leaves the trim at zero. (b) A 25 % feed-factor loss on the AO-2 lane in Manual: the MHE is smoother but reaches half the loss later than the simple estimate.

> **TRY IT IN THE APP** **4 · Hidden Holdup: The Scale Says All Is Well**
>
> Load the scenario, keep the Line view and press Run. Watch the AO-2 feeder's card: its LIW flow keeps reading about 1 kg/h. Now watch AO-2 at the tap, the dashed true line. It is below 900 ppm from 12.4 min, and nothing on the feeder changes until the sock high-level switch trips at 36.6 min.

The line's residence decides when a recipe change reaches the product. The blender, extruder and transport hold 66.7 kg, which is 240 s at 1,000 kg/h, so the scheduled feed switch leads a grade change by that time. Removal is slower than the mean residence suggests. After the clarifier feeder stops, its concentration passes three lags of 10, 72 and 108 s, and integrating them by hand gives 520.5 s to fall by the factor of 40 from 2,000 ppm to 50 ppm. With the 60 s transport the clarifier should leave the band 9.67 min after the feed stops at 26 min, at 35.67 min, and the engine reports 35.7 min.

12.5 Control and Estimation as Implemented

Every controller, estimator and supervisor reads only the instrument packet of Equation (12.4), the released laboratory results, the controls and its own state. The exported file marks the simulator's

truth columns with the prefix `truth_`, so a controller replayed from the file cannot read them. Each control execution, once per second, reads the instruments, forms the recipe, updates the estimator and the laboratory trim, runs the supervisor, solves the MPC if it is due, and updates the six LIW loops, in that order.

12.5.1 Ratio Control and the LIW Loop

The recipe flows of Equation (12.6) at the production setpoint, computed with the declared ingredient matrix and the declared resin background, become ratios to the resin, and each lane's demand is that ratio times the measured resin flow P_m, divided by the laboratory trim, as shown in Equation (12.7).

$$\begin{aligned} q_{dem,j} &= \frac{q_j}{P}\,\frac{P_m}{1+\delta_j} \\ n &= \frac{q_{dem}}{\hat{F}} + K_c\,e + \frac{K_c}{\tau_I}\int e\,dt, \qquad e = 100\,\frac{q_{dem}-\hat{q}}{q_{des}} \end{aligned} \tag{12.7}$$

where $\hat{F}$ is the estimated feed factor, the feedforward inverse, $\hat{q}$ the estimated flow, and the error is in percent of the lane's design rate so that one tuning reads the same on every lane. The shipped tuning is $K_c = 0.2$ rpm per percent and $\tau_I = 60$ s. The PI is the velocity-form PID of Section 1.6.2 with no derivative action, a 0 to 200 rpm range and a slew limit of 20 rpm/s, reduced to 0.5 rpm/s for 5 min after a cleaning. The feedforward enters as the loop's moving bias. While the weight is invalid, during a refill, the 10 s settling after it, and optionally while the hopper pressure is outside 1.5 Pa, the lane runs volumetrically at $q_{dem}/\hat{F}$ and the PI tracks that speed, so the return to gravimetric operation is bumpless.

> ⧖ HISTORY **A Load Cell, a Derivative and a Hold for the Refill, 1967**
>
> An application filed in 1967 replaced the moving weight with electronics. A transducer, which could be a load cell built on strain gauges, turned the weight of the reservoir and its contents into a voltage. A circuit differentiated it with respect to time, and a controller compared that mass flow with the desired rate and adjusted the outflow. Differentiation magnifies anything that changes faster than the weight, so frequencies above about one cycle per second were filtered out to remove the vibration of nearby machinery. Two level settings opened and closed a refill valve. While the weight rose, its derivative would have driven the controller to the wrong flow, so the controller held its last output or switched to manual until the refill ended [15]. The volumetric mode described above does the same.

12.5.2 The Simple Estimate

The classical LIW reading is the least-squares slope of the last 20 s of net weight, valid only when no sample in the window is masked. Its feed factor is that flow divided by the speed, bounded to 0.2 to 1.5 times K, filtered with a 60 s time constant and updated only while valid and while the lane is not in alarm.

12.5.3 The Feeder MHE

The moving horizon estimator works at the scale of the feeder, not of the product: it replaced an eight-hour composition window considered first. A feeder empties gradually and fills repeatedly, its flow is the change of weight over time, and the estimator's job is to detect a change of delivery per revolution against speed for the controller, so its horizon is about five minutes. Over the last $N = 300$ s of 1 s samples it chooses the starting weight W_0, one feed factor F_b per 10 s block and one

free mass J_g per refill that ends in the window, as shown in Equation (12.8).

$$\begin{aligned} \min \sum_{t\in\mathcal{V}} \rho_H\Big(\frac{W_{m,t}-W_t}{\sigma_W}\Big) + \frac{(W_0-\bar{W}_0)^2}{\sigma_{W0}^2} + \frac{(F_1-\bar{F})^2}{\sigma_{F0}^2} \\ + \sum_b \frac{(F_{b+1}-F_b)^2}{\sigma_q^2} + \varepsilon \sum_g J_g^2 \\ W_t = W_0 - \sum_{\tau<t} \frac{F_{b(\tau)}\, n_\tau}{3600} + \sum_{g\le t} J_g \\ 0.2K \le F_b \le 1.5K \end{aligned} \tag{12.8}$$

where $\mathcal{V}$ holds the valid samples, so refill and settling samples are masked, ρ_H is the Huber penalty with threshold 2, $\bar{W}_0$ and $\bar{F}$ are the previous window's trajectory at this window's start, $\sigma_{W0} = 2\,\mathrm{g}$, $\sigma_{F0} = 0.05K$, the random-walk step is $\sigma_q = 0.002K$ per block, and ε is a small regularization that keeps the problem strictly convex. The estimate is carried across a refill by the random walk and is never reset. The weights are linear in the decisions, so each window is a dense convex quadratic program of 31 to 33 variables and 60 bounds, solved by the same dual active-set method [19]. The Huber penalty is applied by iteratively reweighted least squares, two solves per update with weights $\min(1, 2\sigma_W/|r|)$ from the first. A solve is rejected when its status is anything but optimal, when a value is infinite or when a block leaves its bounds, and the lane then keeps its previous estimate, as it also does for a window with fewer than 30 valid samples. The lane's flow is then $\hat{q} = F_B\, n$ from the last block.

Three things are deliberately outside the estimator. A constant scale offset is absorbed by W_0, but a span error or drift cannot be told from a feed-factor change by the weight alone and is left to the laboratory trim. Laboratory assays do not enter the MHE. The sock holdup is not observable from the hopper weight at all, because the weight records what the screw removed and the holdup is what happened to it afterward. Its witnesses are the torque, the high-level switch and the laboratory.

🕮 **Pause.** Try Exercise 12.1 before reading on.

The same problem written for the shared estimator of the framework, with one noise vector per one-second step, has 602 decisions and 903 rows and costs 161 ms per update at $N = 300$. The block structure and one mass per refill reduce it to the 31 to 33 decisions above, which cost 0.205 ms per solve in Node, about 785 times less.

12.5.4 The Product Between Assays and the Laboratory Trim

Between assays the product concentration is a model driven by the estimated lane flows, each times its declared composition and its trim $1+\delta_j$, plus the measured resin times the declared background, passed through the declared sock, blender and extruder lags at the measured throughput and the 16.7 kg transport. A released result is used once, against this model's prediction averaged over its own collection interval, never against the present state, as shown in Equation (12.9). It is the bias update of Equation (1.11) written as a gain on each lane's delivery and time-aligned to the collection interval.

$$\delta_j \leftarrow \delta_j + \lambda\Big(\frac{y_{lab}-b}{\hat{y}_{raw}(t_s,t_e)-b} - 1 - \delta_j\Big), \qquad |\delta_j| \le 0.2 \tag{12.9}$$

where $\lambda = 0.5$, b is the background's share and $\hat{y}_{raw}$ the untrimmed prediction for the interval $[t_s, t_e]$. A result within 1.5 assay standard deviations of the present trim moves nothing. A result more than 25 % from its prediction is held as a possible outlier unless the next result confirms it on the same side or the lane's own diagnostics, the switch or torque above the healthy relation, were raised after its collection began. The estimate's standard deviation combines the assay noise, an aging term of 2 % of target per hour since the last accepted assay, and a holdup term of 30 % of target per unit of torque and switch evidence, the switch alone counting one half.

🕮 **Pause.** Try Exercise 12.2 before reading on.

12.5.5 Supervisor

The supervisor keeps three grade states apart from the laboratory, namely old grade, transition, and new grade predicted on specification once the predicted product has held every band for 10 min, with laboratory confirmation only from an accepted result collected after that instant. With Preview on, the feed recipe switches ahead of a scheduled change by the declared holdup over the measured throughput, 240 s at 1,000 kg/h. Protection acts per lane without waiting for the laboratory. AT LIMIT means a speed at 98 % of maximum with the flow more than 10 % short for 30 s, SOCK HIGH is the switch, and PLUGGED is the switch with the flow under half its demand for 30 s or with the torque above 75 %, or AT LIMIT with either. Under SOCK HIGH or PLUGGED the speed may not rise above its present value, because more speed only fills the sock. Auto routing diverts during a transition, while a lane reads SOCK HIGH or PLUGGED, for 15 min after a knock recorded in the maintenance log, while the recipe is infeasible and while the predicted product is outside a band.

12.5.6 Supervisory MPC

The MPC replaces the ratio demands every 10 s in Auto and leaves the LIW loops underneath it. It has the general form of Section 1.6.3. One small problem per lane uses the declared sock, blender and extruder lags at the measured throughput, with the estimator's own lag states as the initial condition and the trim as the gain, as shown in Equation (12.10).

$$\begin{aligned} \min_{u} \; & \sum_{k} \left\| \frac{y_k - r_k}{tol} \right\|^2 + S \sum \left(\frac{\Delta u}{q_{des}} \right)^2 \\ & + R \sum \left(\frac{u - u_{ratio}}{q_{des}} \right)^2 + \rho \|\sigma\|^2 \\ & 0 \le u \le 0.95 \, \hat{F} \, n_{max}, \qquad |\Delta u| \le q_{des} \end{aligned} \tag{12.10}$$

where y_k is the predicted concentration of each active the lane carries, r_k the destination grade when that material reaches the tap (Preview on), $S = 2$, $R = 0.02$ and u_{ratio} the ratio demand. The bands are soft at $0.8\,tol$ about target with a price $\rho = 50/tol^2$ judged every ninth step with one slack per judged step. The horizon is 15 min, 90 steps, with twelve graduated move blocks per lane, four of one step each and the rest lengthening through 2, 2, 3, 4, 6, 9, then 15, to a final 45. While the supervisor holds a lane, u may not exceed its present flow. The lane takes the ratio demand for a decision whenever the plan's status is not optimal, a move is infinite or a move leaves its box. The MPC never reads a fault, the true sock or an unannounced event, and its capacity is what the estimated feed factor allows, so it cannot overcome a plug.

12.5.7 Plant, Estimator and Controller Models

Three models are kept apart in the app, the simulator's model of Section 12.3, the estimator's declared lags and compositions, and the MPC's three-state linear model per lane. Neither of the last two contains the sock holdup, the fill-level effect, aeration or the true potency. Twin runs that differ only in a hidden truth confirm the boundary. With a 25 % feed-factor loss on AO-2 the first measured difference appears at 2 s and the first speed difference at 5 s. With an 80 % sock restriction the torque differs at 2 s, but no speed differs until 1,602 s, one second after the high-level switch trips and the supervisor holds the lane. A resin background of 300 ppm AO-1 in case 1 and a masterbatch potency of 4.0 % on the shift as shipped change no speed in two hours. A source scan finds no read of the plant state, a disturbance, a run parameter, the case script, the scorecard or a laboratory record's hidden flags in the five controller-side files.

Table 12.3 The Ten Cases

#	Setup	Lesson
1	PE-A, slip neat or as 5 % masterbatch, 1 h	carrier displaces resin
2	PP-A, Manual, production 1,000 → 1,500 kg/h at 10 min	ratio, not speed
3	AO-2 feed factor −25 % over 30 min from 10 min	inside the scale
4	AO-2 sock restricted 80 % at 10 min, sample at 30 min	outside the scale
5	AO-2 sock plugged at 10 min, clean or knock	infeasible, slug
6	AO-2 refill at 220 s with the vent plugged	weight validity
7	PE-A → PE-B at 0.5 h, back at 2.5 h	transition mass
8	PP-B → PP-A at 0.5 h	clarifier washout
9	PP-A → PP-B at 1 h, samples every 2 h, one lost, one botched	late lots
10	12 h shift: two grade changes, 4.3 % masterbatch, AO-2 loss 20 %	challenge

Table 12.4 Measured Outcomes at the Default Seed

#	Result
4	Tap below 900 ppm at 12.4 min, switch and first diversion at 36.6 min. 403 kg off-spec prime. Assay #1 626 ppm against 1,000 ppm predicted, accepted as holdup evidence at 80 min, trim unchanged
5	Switch and diversion at 20.7 min, PLUGGED at 27 min, speed held at 77.8 rpm. Cleaned at 50 min: 293 g to waste, in band at 71.3 min, 148 kg off-spec prime. Knocked: peak 2,902 ppm, 169 kg off-spec prime on Auto, 903 kg on Prime
6	Timer only: speed down to 65.1 rpm, tap down to 911 ppm (MHE: 68.2, 929). Timer and pressure: 281 s volumetric, speed ≥ 75.2 rpm, tap 995–1,021 ppm
7	Transition mass ratio + Preview 410 kg (1.6, 3.0 min). MPC 361 kg (0.7, 1.0 min). Ratio without Preview 541 kg (5.5, 6.9 min)
8	Clarifier above 50 ppm until 35.7 min, 262 kg transition, predicted on spec at 45.8 min. Without Preview 39.7 min. MPC 34.3 min, 239 kg
9	#1 (PP-A lot, 0.90–0.98 h) reported at 1.73 h, clarifier 104 ppm against 89 ppm predicted. #2 lost. #3 reads 1,392 ppm AO-2 against 1,000, rejected. No trim moved
10	Shipped 0 kg on spec, 9,000 kg transition. Suggested (Lab trim and Preview on) 9,574 kg on spec, 450 kg transition, 1,896 kg off-spec prime before the first assay

12.6 Ten Cases, Measured

The ten cases are listed in Table 12.3. Every number in this section comes from headless runs of the engine at each case's default seed and shipped settings, or from three-seed means where so stated.

12.6.1 Recipes, Rates and the Two Restrictions

The measured outcomes of the cases at their default seed are collected in Table 12.4. In case 1 the slip lane moves from 0.5 kg/h neat to 10 kg/h as masterbatch, the resin falls from 995 to 985.5 kg/h, and the slip at the tap stays at 495 to 501 ppm and 498 to 501 ppm. Asked for 1,500 ppm of AO-2 the 1:2 premix is reported infeasible, its closest point being 577 and 1,154 ppm. In case 2 fixed speeds leave AO-2 at 660 ppm by 30 min and Auto routing diverts 1,176 kg of the hour, while ratio control takes the lane to 118.8 rpm with a worst error of 0.07 of a tolerance. In case 3 the speed averages 104.6 rpm between 50 and 90 min (100.6 to 108.5 rpm as the hopper cycles), the tap stays at 997 to 1,003 ppm, and the simple and MHE feed factors end at 0.00941 and 0.00944 kg/h/rpm against a true 0.00947 kg/h/rpm. The same buildup in Manual leaves AO-2 at 764 ppm by 90 min.

In case 4 the scale reads normal for the whole time the product is wrong. The phosphite at the tap is below 900 ppm from 12.4 min, the switch trips at 36.6 min with the LIW reading 1.03 kg/h, and Auto routing diverts on the switch in the same scan. By then 403 kg of off-specification product has

gone to prime, all of it made while every feeder screen read normal. The first assay, collected at 30 to 35 min and reported at 80 min, reads 626 ppm against a prediction of 1,000 ppm. It is accepted although it lies outside the 25 % gate, because the switch corroborates it, and for the same reason it is accepted as evidence of holdup and does not move the trim. The full-width view of this case at 45 min is shown in Figure 12.4. In case 5 the switch trips and diverts the product at 20.7 min, and the protection holds the speed at 77.8 rpm, so the lane never climbs toward 200 rpm. Cleaning sends the 293 g in the sock to waste. Knocking it clear with the route on Prime instead releases a slug that peaks at 2,902 ppm of phosphite, and the case ends with 903 kg of off-specification prime.

12.6.2 Refills, Transitions and Late Results

A plugged vent leaves 18 Pa in the hopper, which at 0.5 g/Pa is a 9 g weight error decaying with 90 s, and the derivative reads the decay as extra flow. With the timer alone the AO-2 speed dips to 65.1 rpm and the tap to 911 ppm, while the pressure check keeps the lane volumetric for 281 s and holds the tap at 995 to 1,021 ppm. The MHE with the timer alone dips less, to 68.2 rpm and 929 ppm. Grade changes in case 7 cost 410 kg of transition material with ratio control and Preview, 541 kg without Preview, and 361 kg with the MPC, which overfeeds within capacity on the way up and cuts early on the way down. In case 8 the clarifier feeder stops at 26 min, ahead of the 30 min change, and the clarifier stays above its 50 ppm limit until 35.7 min. No controller can remove it faster than the line washes out, because a feeder cannot deliver a negative flow. In case 9 the first result belongs to the PP-A lot although it arrives after the change to PP-B, the lost sample moves nothing, and the botched analysis is rejected.

12.6.3 The Four Architectures on Matched Runs

The comparison runs each case with ratio control or MPC, each with the simple estimate or the MHE, on identical measurements, schedules, fault injections and sample availability, at three paired seeds, as shown in Table 12.5. The conventional reference matches or beats the advanced options on the nominal, feed-factor-loss and holdup cases. On the nominal hour the mass-weighted error is 0.021 of a tolerance for ratio control with the simple estimate against 0.032 for MPC with the MHE, and on the feed-factor ramp it is 0.018 against 0.046. The MPC's product benefit is confined to grade transitions, 414 kg of transition material for ratio control against 370 and 365 kg for the two MPC variants, and to the delayed-laboratory case's transition, 219 against 192 kg. Over the twelve-hour shift with the suggested settings the four architectures make 9,576 kg (ratio, simple), 9,570 kg (ratio, MHE), 9,616 kg (MPC, simple) and 9,613 kg (MPC, MHE) on specification, within 0.5 % of each other. The MHE estimates the nominal feed factor with 0.28 % mean error against 0.47 % for the simple estimate, but it changes no product outcome measurably. Its visible effect is on the actuator. With the MPC, the AO-2 speed travels 895 rpm in the nominal hour against 1,463 rpm with the simple estimate, and 12,294 rpm against 20,428 rpm over the shift. Nothing beats the off-scale holdup, which no estimator can see: all four rows of case 4 end with 403 kg of off-specification prime, set by when the sock switch trips. Figure 12.3 draws the error and transition rows of Table 12.5.

> **✦ DID YOU KNOW? The Plainest Pair Is the Most Accurate**
>
> From the simulator: in Table 12.5, ratio control with the simple estimate has the smallest error of the four architectures on the nominal hour, 0.021 of a tolerance, and on the feed-factor loss, 0.018. The MPC with the MHE scores 0.032 and 0.046 on the same runs.

12.6.4 How Long the Feeder MHE Should Look

The horizon was chosen by an open-loop test on the AO-2 lane in Manual, in which the feed factor falls 25 % over 60 s and detection is the first time an estimate falls past half the loss and stays there for 30 s, as shown in Table 12.6. A 60 s window detects in 64 s, faster than the simple estimate's 83 s, but with 0.88 % steady noise against 0.48 %. From 120 s up the noise falls to 0.36 to 0.38 % and

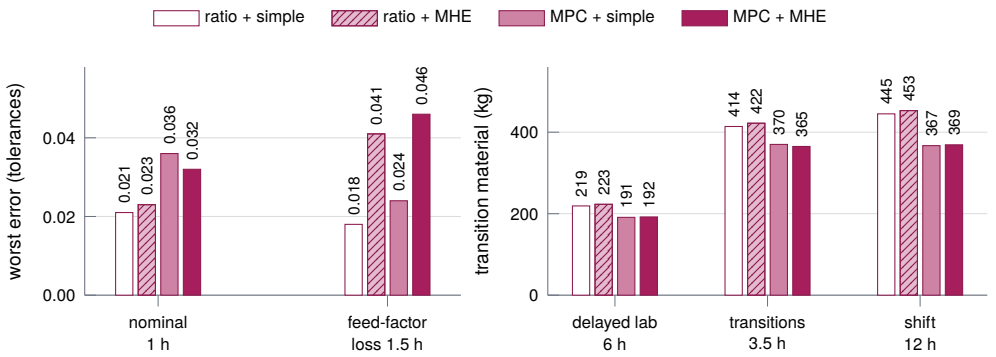

Figure 12.3 The error and transition rows of Table 12.5, three-seed means. Left, the worst active's mass-weighted error in tolerances. Right, the transition material in kilograms. Open and hatched bars are ratio control with the simple estimate and the MHE, medium and dark bars the MPC with the same two. Notice that ratio control with the simple estimate has the smallest errors, that the MPC's gain is confined to the transitions, and that the holdup and plug rows, not drawn, read 403 kg and 148 kg in every column.

Table 12.5 Four Architectures, Three Paired Seeds (Means)

Case, metric	R+S	R+M	M+S	M+M
nominal 1 h, error (tol)	0.021	0.023	0.036	0.032
nominal, AO-2 travel (rpm)	1,522	1,397	1,463	895
delayed lab 6 h, transition (kg)	219	223	191	192
feed-factor loss 1.5 h, error (tol)	0.018	0.041	0.024	0.046
holdup 2 h, off-spec prime (kg)	403	403	403	403
plug 2 h, off-spec prime (kg)	148	148	148	148
transitions 3.5 h, transition (kg)	414	422	370	365
shift 12 h, on spec (kg)	9,576	9,570	9,616	9,613
shift, transition (kg)	445	453	367	369
shift, AO-2 travel (rpm)	20,524	19,652	20,428	12,294

R ratio + PI, M MPC (first letter). S simple estimate, M MHE (second letter). Error is the worst active's mass-weighted error in tolerances. MPC: 0 fallbacks, 0.57 to 0.65 ms per decision for all lanes in Node.

Table 12.6 MHE Horizon Against the Simple Estimate (Three Seeds)

N (s)	noise MHE	noise simple	half-loss MHE	half-loss simple
60	0.88 %	0.48 %	64 s	83 s
120	0.38 %	0.48 %	87 s	83 s
300	0.36 %	0.48 %	91 s	83 s
600	0.36 %	0.48 %	91 s	83 s

detection slows to 87 to 91 s, and beyond 300 s nothing changes, because the random-walk step, not the window, sets the estimator's memory. At the shipped 300 s the MHE is quieter than the simple estimate but not faster, 91 s against 83 s, as shown in Figure 12.2(b). The 300 s window was kept because it spans a refill and its settling, up to 281 s of invalid weight in case 6, with valid data on both sides.

🕮 **Pause.** Try Exercise 12.3 before reading on.

12.6.5 Holdup Is Not a Delivery Loss

Following case 4 past its end tests how the laboratory trim treats a holdup. The trim corrects a lasting error in delivery, while the sock's shortfall is transient and leaves with the sock's contents. A result

taken while a lane's holdup alarms are raised is therefore accepted as evidence and leaves that lane's trim at zero, and a cleaning or a knock recorded in the maintenance log resets the trim. Cleaned at 90 min, the lane returns to its recipe flow and the phosphite at the tap averages 990 ppm between 110 and 170 min, with a minimum of 875 ppm in the restart and a maximum of 1,001 ppm. Knocked clear instead, it averages 1,000 ppm. At 150 min the estimate and the truth both read 1,000 ppm. By 3 h the off-specification prime is 403 kg after a cleaning and 407 kg after a knock, with the trim on or off. The app's suite checks the diversion at the switch, the unmoved trim and the 1,100 ppm band edge after the cleaning.

12.6.6 The Shift Challenge

The shift runs PE-A, PE-B from 3 h and PE-A again from 8 h, with a slip masterbatch that is really 4.3 % active and an AO-2 feed factor that falls 20 % from 1.5 h. Every scale reads correctly while the slip runs 14 % low, so as shipped, with Lab trim and Preview off, no product is on specification and all 9,000 kg made after the first change count as transition. The target is 9,200 kg on specification and at most 500 kg of transition by 12 h. With both settings on, the three seeds make 9,574 kg, 9,587 kg and 9,566 kg on specification with 450, 430 and 455 kg of transition, and all three pass. Lab trim alone makes 9,564 to 9,586 kg but 562 to 589 kg of transition and fails, and Preview alone makes none.

12.6.7 Verification

The total mass balance, inventories against the resin and refill in and the product and waste out, closes to 5.8×10^{-14} of the input over case 4 and to 8.9×10^{-14} through a sock restriction and a knock, and the app's own test holds it below 10^{-9} through refills, a plug, cleaning to waste and a knock. Halving the step to 0.5 s with the controllers still at 1 s moves case 7's transition mass from 409.72 to 409.31 kg, its recoveries from 93 and 180 s to 93 and 179.5 s and its peak from 5.853 to 5.869 tolerances, case 4's off-specification prime from 402.90 to 402.76 kg, and the slug's peak not at all at 2,871 ppm. The four-way benchmark reproduces the app's recorded file in every non-timing field of all 32 rows. In Node a simulated hour of the shift costs 0.86 s of wall time with ratio control and 1.05 s with MPC and MHE, and the browser plays the twelve-hour shift in 72 s at its top speed of 600 times real time. The app's own suite passes 53 of 53 checks.

12.7 Technology in Practice

What Measures It

Every lane's controlled flow comes from one measurement, the net weight of a hopper, its motor and its screw, read once per second as W_m in Equation (12.4). A real feeder obtains it from load cells under the weighing platform, and its controller divides each change of weight by the time between readings to obtain the flow [3]. The box below describes the cell and the forces that can pass for powder, and Figure 12.5 shows its principle.

> **TECHNOLOGY Sensor: Strain-Gauge Load Cells Under the Feeder**
>
> A load cell is a metal element of precisely known elasticity with strain gauges bonded to it. Under load the element compresses on a microscopic scale and the gauges change their resistance. A vessel usually stands on several cells whose signals are summed, because its weight is not shared equally among its supports. The cells read every force on the vessel: a pipe without a flexible coupling, vibration from nearby machinery, or a current through a cell from a welder grounded to the frame all look like a change in contents [20]. A flow taken from the change of weight is a net flow, so it cannot separate an outflow from a refill [20]. WT of the chapter stands for the summed cells, with 0.2 g of noise, a span error s and the hopper pressure reaching it at 0.5 g/Pa.

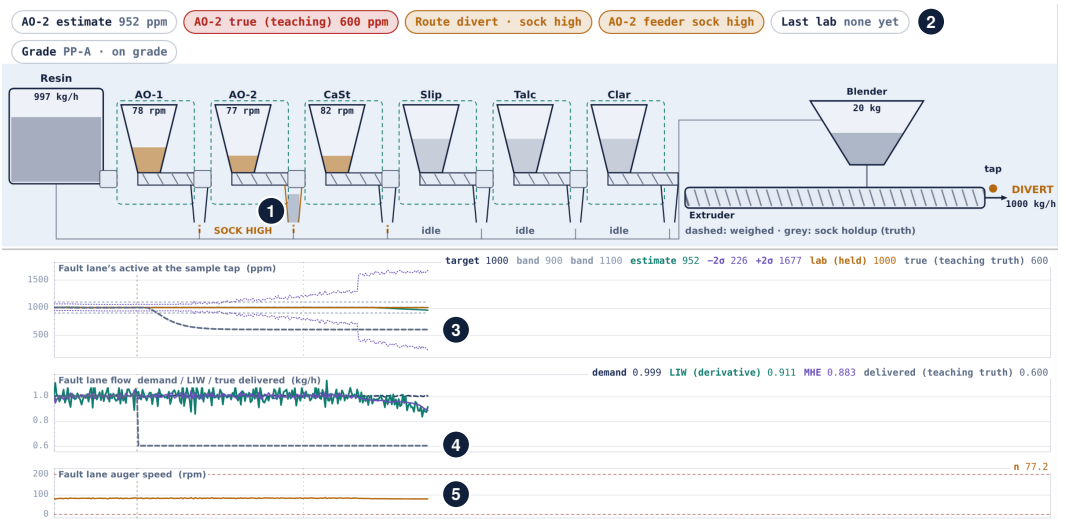

Figure 12.4 Hidden holdup of case 4 at $t = 45$ min on the Line view, cropped to its badges, its plant view and its first three trends. (1) The AO-2 lane reads SOCK HIGH with its sock drawn below the dashed weighing boundary. (2) The badges: the phosphite estimate is 952 ppm (± 725 ppm, two standard deviations) against a true 600 ppm and a target of 1,000 ppm, and the switch has sent the product to diversion. (3) The phosphite at the sample tap, the dashed true line below the estimate and its band. (4) The flow trend shows a demand of 0.999 kg/h, an LIW reading of 0.911 kg/h, an MHE flow of 0.883 kg/h and 0.600 kg/h delivered, (5) at 77.2 rpm.

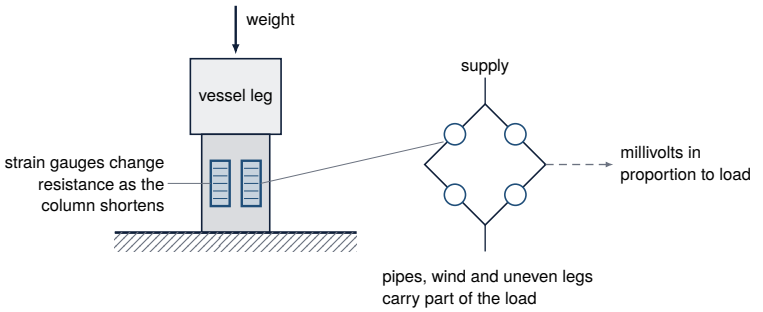

Figure 12.5 The strain-gauge load cell from the book's instrument library. Notice that the cell measures every force that reaches it through the leg. In the chapter's model the hopper pressure enters W_m this way, through the flexible connections, and the share α of the sock that the scale carries is set to 0, which is why the holdup of case 4 never reaches the weight.

What Moves It

The manipulated variable of each LIW loop is the auger speed n, and the torque JT is the lane's only continuous witness of a sock that pushes back. In a real feeder the speed is set by a motor drive turning a metering screw, as the box below explains. Neither the screw nor its drive can see the sock: a full sock throttles the screw through the backpressure factor $\mathrm{acc}(H)$, and the supervisor then forbids more speed.

TECHNOLOGY Final Element: A Metering Screw on a Variable-Speed Drive

A loss-in-weight feeder is a volumetric feeder, a hopper with a metering screw, mounted on a weighing platform under a gravimetric controller that corrects the screw speed whenever the weight lost per unit time departs from the set rate. Screws come single or twin, with various thread depths, as spirals or augers [3]. A variable-frequency drive sets the speed of an alternating-current motor through the frequency of its supply, and drives commonly limit speed, torque and acceleration [20]. A powder of low density can compact in the flights until the motor reaches its maximum torque and the feeder shuts down [3]. The chapter's screw runs from 0 to 200 rpm under a slew limit of 20 rpm/s, and JT reads its torque.

Where the Controller Runs

Figure 12.6 draws one lane of the chapter with hardware in place of the schematic's symbols. The gravimetric controller is part of the feeder itself, together with the volumetric feeder and the weighing platform [3]. The recipe, the ratio to the resin and the laboratory trim sit a level higher, in a control system whose supervisor sends set points to field controllers, with loops nested or cascading and cycle times from milliseconds to minutes [21]. In the app the six PI blocks run in velocity form every second, the MHE every 5 s and the MPC every 10 s, and the lane falls back to volumetric operation whenever its weight is invalid. The settings of that fallback are a matter of practice.

IN PRACTICE Choosing When to Refill

During a refill a loss-in-weight feeder runs volumetrically and is blind to density: the incoming powder can compress the bed below it or aerate it. The number of refills is therefore chosen so that the feed-rate deviation stays acceptable to the operations downstream. In the study a 2018 review reports, refilling a hopper at a low fill level such as 20 % gave larger deviations, while refilling at 60 % fed more accurately but demanded more frequent, smaller refills. A hopper designed for the feed rate and the powder also reduces the disturbance [3]. The chapter's lanes start a refill below 25 % of the hopper and stop at 85 %.

What Is Modernizing It

Three developments bear on this line. In-line Raman and near-infrared spectroscopy at the extruder reads the active content of the melt itself. In one study a model built on in-line Raman spectra saw a powder accumulation in the barrel that a soft sensor built on the feeders' data could not, because that sensor knew only what the feeders had dosed [22]. Near-infrared probes at a feeder outlet have been used to follow the bulk density of the powder and the screw speed that compensates for it [3]. A hybrid of predictive and PID control, informed by process analytical tools, has been implemented on a continuous direct-compaction pilot plant for tablets [16]. The first development is the one this chapter's holdup asks for, and the Open challenge box in Section 12.8 says why it is not yet routine.

12.8 Limitations

The equipment sizes, the sock capacity and backpressure, the torque relation, the pressure coupling and the bands are teaching values sized for one self-consistent line at 1,000 kg/h, not data from a plant or a supplier. The blender and extruder are ideal mixers, the actives are conserved, and the assays map exactly onto the actives with added noise, so the model cannot represent antioxidant consumption, segregation, axial dispersion or an ambiguous ash result. A polypropylene to polyethylene changeover is refused rather than modeled. The faults arrive as scheduled steps and ramps.

The MHE estimates only the feed factor and the weighed flow. It cannot see a span error, a potency error or the sock, and no confidence interval is claimed for it. The MPC plans each lane separately, has no resin-rate decision, and carries no stability guarantee beyond its linear prediction rebuilt every 10 s. The laboratory trim separates holdup from delivery only through the torque and the switch: a holdup that raises neither would still be learned as a lasting efficiency loss. Protection and routing are

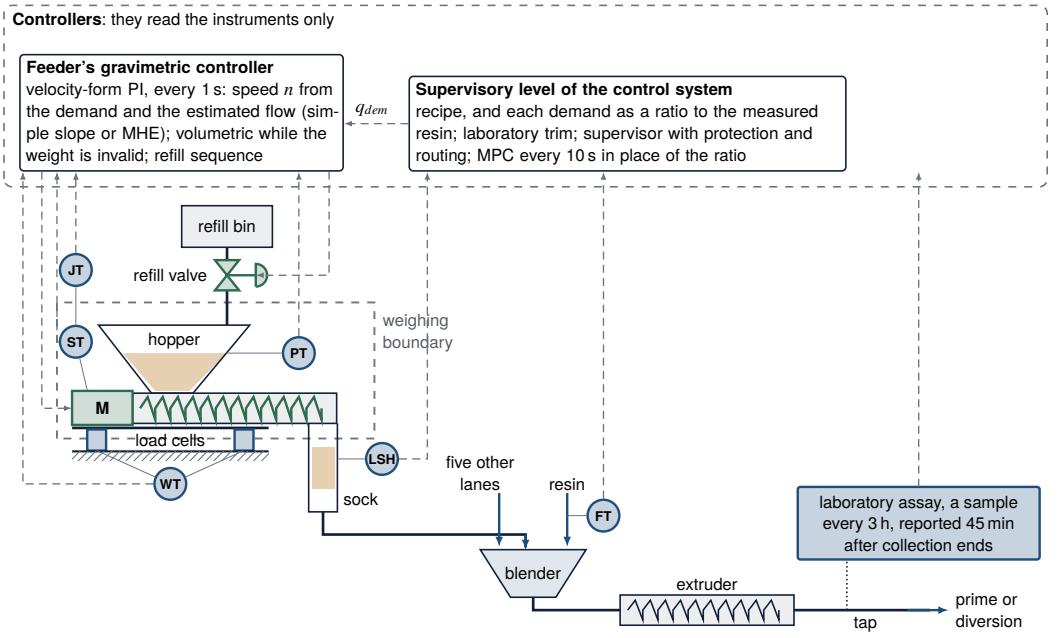

Figure 12.6 One additive lane with its hardware. Blue marks are sensing devices: the load cells WT under the weighing platform, the speed and torque of the drive, the hopper pressure, the sock high-level switch and the resin flow FT. Green symbols are final elements: the screw with its variable-speed drive and the refill valve. Dashed arrows are signals. Notice that the dashed weighing boundary encloses the hopper, the drive and the screw but not the sock, and that the only measurement of the product reaches the controllers from a laboratory, 45 min after each sample. The other five lanes and the diverter are left out. Not to scale, and not a piping diagram.

an educational supervisor, not certified interlocks, and a simulated pass is not evidence of compliance with any product specification.

> ⚠ **OPEN CHALLENGE** **Reading the Additive in the Melt**
>
> In-line spectroscopy could report an additive in the melt while the product is made, not hours later. Near-infrared and Raman measurements of molten polymer have been researched intensively, yet industry still applies them sparsely because their measurements are unpredictable and unstable [23]. In one study of polymer additive content in an extruder, a calibration built on near-infrared spectra from one time lost its accuracy on spectra collected at another [22]. Progress would be an in-line reading that keeps its calibration across days, grades and machines, so that a holdup like the sock of case 4 shows in the product long before the next assay.

12.9 Conclusion

A loss-in-weight feeder measures what leaves its hopper, and this simulator logs that reading beside the flow that reaches the blender, the holdup between them and the product at the tap. A restriction inside the weighing boundary is corrected by speed, from 80 to about 105 rpm for a 25 % loss, while one in the sock below it keeps the reading near 1 kg/h as the product falls to 600 ppm, and 403 kg of it reaches prime before the sock switch diverts it. On matched runs ratio control matches the advanced options except on grade transitions, where the MPC saves 44 to 49 kg of 414 kg, and the five-minute MHE estimates the feed factor with less error and less speed travel but reaches half a step loss in 91 s against 83 s. The holdup is unobservable from the weight, so the laboratory trim must not learn

from a result taken under the holdup alarms. The model is a teaching model with declared relations. Natural extensions are a measured sock load ($\alpha > 0$), a holdup estimate from the torque, antioxidant consumption in the extruder and a resin purge between polymer families. The simulator shows that a feeder's scale vouches for the hopper, never for the pellets.

KEY POINTS

- A loss-in-weight feeder weighs what leaves its hopper, never what reaches the pellets.
- A restriction inside the weighing boundary is corrected by speed. One in the sock below it keeps the reading near 1 kg/h while the product falls to 600 ppm.
- The holdup is not observable from the weight. Its witnesses are the torque, the high-level switch and the laboratory, so the trim must not learn from a result taken under the holdup alarms.
- On matched runs ratio control matches the advanced options except on grade transitions, where the MPC saves 44 to 49 kg of 414 kg.

Where else this applies. A dose set from an upstream measurement, with the slow laboratory result demoted to a trim, is also the coagulant feedforward of Volume 3, Chapter 6, and loss-in-weight feeders meter the dose ahead of the tablet press of Volume 5, Chapter 6.

References

[1] M. Hopkins, "Loss in weight feeder systems," *Measurement and Control*, vol. 39, no. 8, pp. 237–240, 2006.

[2] W. E. Engisch and F. J. Muzzio, "Method for characterization of loss-in-weight feeder equipment," *Powder Technology*, vol. 228, pp. 395–403, 2012.

[3] C. A. Blackshields and A. M. Crean, "Continuous powder feeding for pharmaceutical solid dosage form manufacture: a short review," *Pharmaceutical Development and Technology*, vol. 23, no. 6, pp. 554–560, 2018.

[4] J. Hanson, "Control of a system of loss-in-weight feeders for drug product continuous manufacturing," *Powder Technology*, vol. 331, pp. 236–243, 2018.

[5] J. N. Hahladakis, C. A. Velis, R. Weber, E. Iacovidou, and P. Purnell, "An overview of chemical additives present in plastics: Migration, release, fate and environmental impact during their use, disposal and recycling," *Journal of Hazardous Materials*, vol. 344, pp. 179–199, 2018.

[6] J. Pezzutti, A. Benito, G. Cassano, L. Roth, W. Schoene, H. Siebert, A. Winter, A. Dimeska, and V. Galiatsatos, "Supertransparent high impact strength random block copolymer," U.S. Patent 8,076,429, Dec. 13, 2011.

[7] K. B. McAuley and J. F. MacGregor, "On-line inference of polymer properties in an industrial polyethylene reactor," *AIChE Journal*, vol. 37, no. 6, pp. 825–835, 1991.

[8] ——, "Optimal grade transitions in a gas phase polyethylene reactor," *AIChE Journal*, vol. 38, no. 10, pp. 1564–1576, 1992.

[9] 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.

[10] F. Tahir, J. Palmer, J. Khoo, J. Holman, I. K. Yadav, G. Reynolds, E. Meehan, A. Mitchell, and G. Bajwa, "Development of feed factor prediction models for loss-in-weight powder feeders," *Powder Technology*, vol. 364, pp. 1025–1038, 2020.

[11] D. G. Robertson, J. H. Lee, and J. B. Rawlings, "A moving horizon-based approach for least-squares estimation," *AIChE Journal*, vol. 42, no. 8, pp. 2209–2224, 1996.

[12] 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.

[13] 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.

[14] R. J. Harkenrider, "Apparatus for feeding material at a predetermined rate," U.S. Patent 2,544,155, filed 27 March 1945, granted 6 March 1951, 1951.

[15] H. H. Marhauer, "Mass flow controller," U.S. Patent 3,481,509, filed 18 October 1967, granted 2 December 1969, 1969.

[16] 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.

[17] H. J. Vandenburg, A. A. Clifford, K. D. Bartle, J. Carroll, I. Newton, L. M. Garden, J. R. Dean, and C. T. Costley, "Analytical extraction of additives from polymers," *The Analyst*, vol. 122, no. 9, pp. 101R–116R, 1997.

[18] P. J. Huber, "Robust estimation of a location parameter," *The Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73–101, 1964.

[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.

[22] N. Munir, M. Nugent, D. Whitaker, and M. McAfee, "Machine learning for process monitoring and control of hot-melt extrusion: Current state of the art and future directions," *Pharmaceutics*, vol. 13, no. 9, p. 1432, 2021.

[23] S. Zhu, Z. Song, S. Shi, M. Wang, and G. Jin, "Fusion of near-infrared and raman spectroscopy for in-line measurement of component content of molten polymer blends," *Sensors*, vol. 19, no. 16, p. 3463, 2019.

12.10 Exercises

12.1 Concept check · 2 min

In case 4 the AO-2 discharge sock is restricted to a fifth of its healthy capacity, below the weighing boundary. Which measurements can reveal the holdup before a laboratory result does?

(a) The hopper weight, once the MHE replaces the simple 20-second slope

(b) The LIW derivative, once its window is lengthened from 20 to 60 seconds

(c) The auger speed, which climbs as the loop makes up the missing flow

(d) The motor torque and the sock high-level switch, read outside the weight

12.2 Concept check · 1 min

Assay #1 of case 4 reads 626 ppm of AO-2 against a prediction of 1,000 ppm, and the AO-2 sock switch was raised after its collection began. Why does the app accept the result without moving the AO-2 trim?

(a) The shortfall sits in the sock and leaves with it, so a trim would overdose after the repair

(b) The result is more than 25 % from its prediction, so it is held as a possible outlier

(c) The result arrives 45 min after collection, too late to describe the present product

(d) The MPC sets the AO-2 demand in Auto, so the laboratory trim is not used

12.3 Concept check · 2 min

Table 12.6 shows that lengthening the feeder MHE's window from 300 s to 600 s changes neither its noise nor its detection time. Why?

(a) The Huber penalty discards the older samples of a longer window as outliers

(b) The random-walk step of the feed factor, not the window, sets how far back the estimate remembers

(c) The solver stops after two reweighted solves, before the extra feed-factor blocks can move

(d) A refill falls in every 600 s window, and its samples are masked from the fit

12.4 Theory and calculation · 5 min

The PE-B grade at 1,000 kg/h of finished compound asks for 500 ppm of AO-1, 1,000 ppm of AO-2, 500 ppm of calcium stearate, 1,000 ppm of slip and 5,000 ppm of talc, and no clarifier. The resin carries none of them.

(a) With every additive fed neat, find each feeder's flow and the resin flow. *Checkpoint:* 0.5, 1, 0.5, 1 and 5 kg/h, and 992 kg/h of resin.

(b) The slip is fed instead as a 5 % masterbatch. Find the slip feeder's flow, the carrier it brings and the new resin flow. *Checkpoint:* 20 kg/h, 19 kg/h of carrier, 973 kg/h of resin.

(c) The masterbatch is really 4.3 % active while the recipe assumes 5 %. Find the slip in the product and say whether the ±10 % band holds it. *Checkpoint:* 860 ppm, outside the band.

12.5 Theory and calculation · 8 min

The AO-2 lane has the feed factor $K = 0.0125\,\mathrm{kg/h/rpm}$, a design rate of 1 kg/h at 80 rpm, a speed limit of 200 rpm, a sock of capacity $H_{cap} = 0.30\,\mathrm{kg}$ with transit time $\tau_{sock} = 10\,\mathrm{s}$ and outflow capacity $c_0 = 3q_{des}$, and a high-level switch at $0.6\,H_{cap}$. Take the hopper half full ($f = 1$) and no aeration.

(a) A buildup at the screw discharge removes 25 % of the delivery per revolution. Find the speed that restores 1 kg/h, and the largest loss the speed limit can make up. *Checkpoint:* 106.67 rpm; 60 %.

(b) The sock is restricted instead, with $r_B = 0.8$. Find its outflow capacity, and the smallest r_B at which no speed can deliver 1 kg/h. *Checkpoint:* 0.6 kg/h; $r_B > 2/3$.

(c) The screw keeps sending 1 kg/h into that sock from $t = 10\,\mathrm{min}$, starting from its healthy holdup. Find when the switch trips. *Checkpoint:* After 26.6 min, at $t = 36.6\,\mathrm{min}$.

12.6 Theory and calculation · 6 min

The blender holds 20 kg, the extruder 30 kg and the transport to the tap 16.7 kg, at 1,000 kg/h. A stopped feeder's concentration passes three first-order lags, the sock (10 s), the blender (72 s) and the extruder (108 s), and then the transport delay.

(a) Find the lead that Preview gives the feeders ahead of a scheduled grade change, and the transport delay alone. *Checkpoint:* 240 s; 60 s.

(b) The response of the three lags to a unit step down is $y(t) = \sum_i a_i e^{-t/\tau_i}$ with $a_i = \tau_i^2 / \prod_{j \neq i}(\tau_i - \tau_j)$. Find the three a_i, and show that $y = 1/40$ at $t = 520.5\,\mathrm{s}$. *Checkpoint:* $a_i = 0.0165$, -2.323 and 3.306; $y(520.5) = 0.0250$.

(c) The clarifier feeder of case 8 stops at 26 min from 2,000 ppm. Find when the clarifier at the tap falls below its 50 ppm limit. *Checkpoint:* 35.67 min.

12.7 App challenge · 9 min

Load scenario 3, "Gradual sock restriction: the speed rises a third". Despite its name, the fault is a buildup at the AO-2 screw discharge, inside the weighing boundary: the feed factor falls 25 % over 30 min from 10 min. Use the default seed and the speed 300× (300 simulated seconds per wall-clock second). The case pauses at 1.5 h. Export the CSV file, whose time column is in hours.

(a) Run the case as loaded, in Auto. Find the mean of `rpm_1` over the rows from 50 to 90 min, and the range of `truth_ppm_1` over the same rows. *Checkpoint:* 104.5 rpm, within 0.3 rpm; 997 to 1,004 ppm, within 1 ppm.

(b) Read `ff_simple_1`, `ff_mhe_1` and `truth_ff_1` in the last row. *Checkpoint:* 0.00941, 0.00944 and 0.00947 kg/h/rpm, within 0.00001.

(c) Press the reload button, switch to Manual before Run, and run to 1.5 h. Read `truth_ppm_1` in the last row. *Checkpoint:* 764 ppm, within 2 ppm.

Compare (a) with the 106.67 rpm of a half-full hopper.

12.8 App challenge · 8 min

Load scenario 4, "Hidden holdup: the scale says all is well", and change no setting. The AO-2 sock is restricted 80 % at 10 min. Use the default seed and the speed 300×. Run until the case pauses at 2 h and export the CSV file, whose time column is in hours.

(a) Find the mean of `liw_kgh_1` over the rows from 15 to 30 min, and read `truth_del_kgh_1`, `truth_ppm_1` and `est_ppm_1` in the row at 30 min. *Checkpoint:* 1.00 kg/h, within 0.01 kg/h; 0.600 kg/h, 600 ppm and 1,000 ppm, each within 0.5 %.

(b) Find the first row whose `route` reads `divert`, and its `route_why`. Read the Off-spec prime tile at 2 h. *Checkpoint:* 36.7 min, within 0.2 min, "sock high"; 403 kg.

(c) Find the first row in which `lab_ppm_1` differs from 1,000 ppm, its value, and `trim_1` in the last row. *Checkpoint:* 80.2 min, 626 ppm within 1 ppm; a trim of 0.

12.9 App challenge · 9 min

Load scenario 7, "PE-A ↔ PE-B on a known schedule". The grade changes to PE-B at 0.5 h and back to PE-A at 2.5 h. Use the default seed and the speed 600×, and run each case until it pauses at 3.5 h. Read the second mass and the last recovery on the Diverted · transition tile.

(a) Run the case as loaded, with Ratio + PI and Preview On. *Checkpoint:* 410 kg; last recovery 3.0 min.

(b) Set Preview to Off, press Reset and run again. *Checkpoint:* 541 kg; 6.9 min.

(c) Set Preview to On and Control structure to MPC, press Reset and run again. *Checkpoint:* 361 kg; 1.0 min.

Accept a mass within 2 kg. Explain what Preview and the MPC each change.

12.10 Application area · 4 min

A plugged vent filter leaves 18 Pa in a hopper at the end of a 20 s refill, and the pressure decays with a time constant of 90 s. The load cells read the pressure at 0.5 g/Pa, and the feeder treats the weight as valid again 10 s after the refill ends.

(a) Find the weight error at the end of the refill, and the extra flow the weight derivative reads 10 s after it. *Checkpoint:* 9 g; about 0.32 kg/h.

(b) Say what the AO-2 loop does with that reading, and what a plant can do about it. *Checkpoint:* It cuts the speed. Keep the weight invalid while the hopper pressure is out of its band, and keep the vent filter clean.

12.11 Application area · 3 min

In case 4, all four control architectures send 403 kg of off-specification product to prime, and the AO-2 at the tap falls below 900 ppm 24 min before the sock switch trips. A plant engineer is asked what would shorten that gap. Name two changes to the equipment or the instruments, and say why a sample requested at 10 min would not help.

12.12 Review the agent · 5 min

An AI assistant reviewing case 4 writes: "The simple LIW estimate is a 20 s slope of a noisy weight, which is why the restriction went unnoticed for 26 min. The MHE fits the last 5 min of weight against speed, and its feed-factor error on the nominal hour is 0.28 % against 0.47 %. Setting Estimator to MHE will detect the sock restriction sooner and cut the 403 kg of off-specification prime." Find the flaw and say how you would check it.

CHAPTER 13

Summary and Outlook

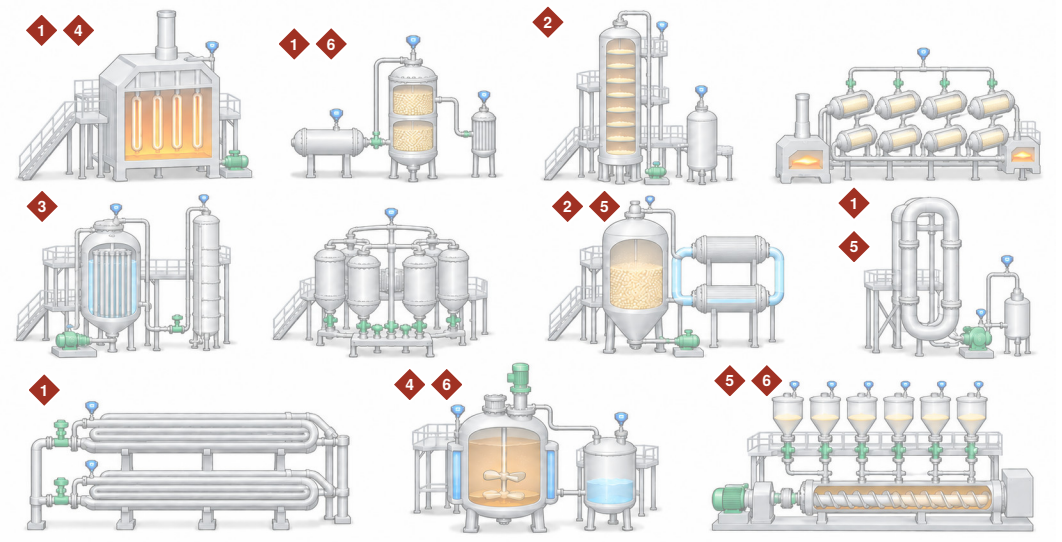

Open gaps 1 limit where it is 2 inference validity 3 flat optimum 4 run length 5 whole grade changes 6 unweighed holdups

Notice that each number stands beside every unit where a chapter met that gap, so the first gap appears four times and only the dehydrogenation beds and the moving bed carry none. Illustration, not to scale.

IN PLAIN TERMS **What Eleven Petrochemical Units Taught**

The main idea. Eleven chapters followed gas on its way to plastic pellets, from the furnace that cracks it to the feeders that add the last ingredients. Each unit had to keep a product quality that is checked only now and then close to a limit it must never cross.

The challenge. The limit is often measured in the wrong place, so the true value can run ahead of the reading. The quality comes late, and a tester that has stopped can still look normal.

Why it matters. A unit that trusts a wrong reading makes plastic nobody can sell, or shuts down, while every display looks fine.

Eleven chapters followed the chain from the cracking furnace to the additive-dosed pellet. Table 1.2 listed, for each unit, the quality it holds and how often that quality is seen, the limit that binds and the structure that guards it. This chapter adds what worked, what failed and the measured result, and reads the comparison for the question of the volume: how a quality that is rarely measured is held against a safety limit that must never be crossed. It then gives a short guide for choosing a design, 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.

13.1 What the Chapters Showed

Tables 13.1 and 13.2 compare the chapters, split between the units that make and purify monomers and intermediates and the units that polymerize and dose. The second column condenses Table 1.2

Table 13.1 How Each Unit Held a Rarely Measured Quality Against Its Limit, Chapters 2 to 7

Chapter	Quality, how often, limit	What worked	What failed	Measured result
2, cracking furnace	Conversion on a 21 min analyzer. Tube metal behind the coke, which makes a coked pass read cold.	Balancing the passes on tube metal. A decoke at 1090 °C with the tuned outlet loop.	Balancing on coil outlet temperature. Nobody decoking, as shipped.	Runs of 910, 888 and 841 h balanced on metal, not balanced and balanced on outlet (Table 2.4). Challenge 20.67 t/h with no time over the limit, against 858.2 h over (Section 2.6).
3, acetylene converter	Outlet under 1 ppm on a 5 min chromatograph. A window open on one side. Bed rise trip at 70 K.	Feedforwards from the inlet analyzer with a trim on log ppm at 0.1 ppm. A balance inference when the analyzer fails.	A fixed ratio, which breaks through at 1.5 and spends ethylene at 2.0.	Worst sample 0.53 ppm at a loss of 0.297 %, against 185 of 288 samples over at 1.5 and 0.566 % at 2.0 (Table 3.6).
4, propylene splitter	Purity on a 20 min chromatograph against a response of hours. A tray temperature that reads pressure first.	The analyzer on the reflux ratio while it is healthy. The cascade with its validity hold armed from the start.	A cascade that integrates a frozen analyzer until a late hold. An uncompensated tray.	Peak departure 0.063, 0.143 and 0.373 mol% for analyzer, cascade and inferential alone. Frozen analyzer: 98.90 with the late hold, 99.17 with it from the start (Table 4.4).
5, dehydrogenation beds	Conversion once a 3 min slot. Bed limit 680 °C. The handle acts on the next cycle's stored heat.	A cycle-to-cycle PI on the regeneration, with two beds on line after one is lost.	Three beds on line under any control. Fixed regeneration, as shipped.	Worst hour 1.73 points from target, against 3.19 shipped, 8.14 with three on line under the PI and 9.14 under the MPC (Table 5.5).
6, ethylene oxide	Selectivity on a 30 min analyzer. Hot spot, 4.2 °C of coolant from the end of the cool state. Flammable oxygen.	A retuned oxygen loop, all four composition loops closed, the override on the rate.	A chloride PID started above the optimum. The cautious oxygen loop, as shipped.	From 3.30 ppm the PID trips at 381.8 min, from 1.50 it creeps (Table 6.5). Challenge oxygen error 0.105 against 0.314 mol% (Section 6.5).
7, paraxylene moving bed	Both purities as 12 min cycle composites. Edges of the triangle, which move as the adsorbent ages.	Two PIs once a cycle on the logarithm of each impurity. The repetitive MPC.	Fixed ratios at the fresh point, as shipped. Every fixed point on a 0.01 grid.	PIs pass at 32.78 t/h with recovery no lower than 97.47 %. Shipped recovery falls to 82.62 %, best fixed point 27.00 t/h (Table 7.5).

into the quality, how often it is known, and the limit. The last column gives the measured result with the table or section of the chapter that reports it.

The columns make five points about the question the volume asked. Figure 13.1 draws one like pair of results from each chapter's row.

The limit is measured in the wrong place, and the guard is only as good as that reading. On the cracking furnace the coil outlets agreed and the tube metal did not. Balancing on the outlet took feed away from the coked pass that needed it, and the run fell to 841 h against 888 h with no balancing at all. On the acetylene converter the ethylene loss passed one percent 1.58 min after the hydrogen was raised by hand, and the bed rise reached its trip only at 3.95 min (Table 3.5). The selectivity was lost long before the limit was. On the LDPE tube a trip set 5 K below the decomposition line fired 1.5 s before a decomposition that one scan earlier it would have prevented (Section 10.9). The back-off of Equation (1.13) has a term for how far the truth runs ahead of its measurement, and each of these units needed that term.

Table 13.2 How Each Unit Held a Rarely Measured Quality Against Its Limit, Chapters 8 to 12

Chapter	Quality, how often, limit	What worked	What failed	Measured result
8, gas-phase polyethylene	Melt index and density from a laboratory 45 to 105 min old. Sticking temperature and cooling capacity.	The retuned quality cascade, with the bed drawn down to 55 % before the grade change.	The shipped melt-index tuning, which never used its hydrogen limit. A rate cut held through the change.	Off-grade 77.5 against 179.1 t as shipped, and 128.7 t at 20 t/h throughout (Table 8.9). The cascades read the bed quality without delay.
9, polypropylene loop	Melt flow rate on a 6 min analyzer behind a 52 min average. Pump cavitation as hydrogen spends the bubble point.	A bubble-point controller at a low selector. For the challenge, 41 bar first and a slower quality loop.	Quality cascades alone. Pressure alone or the slow tuning alone in the challenge.	Cavitation 42.3 against 0 min, polymer 45.6 against 110.6 t, of which 33.4 and 103.3 t off specification (Table 9.5).
10, tubular LDPE	Melt index every 6 min, density from the laboratory. Decomposition at 335 °C, gas ahead of every thermocouple.	Peak loops at 310 °C with the pulse feedforward. A predictive controller that carries the ceiling.	The feedforward off. Peaks at 312 °C. A supervisor reserve of 20 K blind to the lot change.	Conversion 30.79 % with 16.46 K to spare, against 27.72 as shipped. Without feedforward 13.55, with the supervisor 11.71 K (Table 10.7).
11, butyl rubber	Mooney from the laboratory every 2 h. A jacket walking to its refrigeration floor as the wall fouls.	The Mooney cascade with tuned loops and a wash on a cooling margin of 0.10. A setpoint near the warm optimum.	No wash, as shipped. A wash on a clock written for another fouling rate.	Challenge 3.83 t/h, worst error 0.83 °C, against a trip at 68.5 h (Section 11.6). On specification 4.37 at −93 °C, 1.25 at −101 °C (Table 11.3).
12, additive feeders	Additive content from a composite every 3 h, 45 min late. A sock below the scale that the weight cannot see.	Ratio control with the laboratory trim and the grade preview. Diversion on the sock switch.	Either setting alone. Any estimator against the holdup.	Shift 9574 kg on specification with 450 kg of transition, against none on specification as shipped (Table 12.4). Holdup 403 kg off specification in all four architectures (Table 12.5).

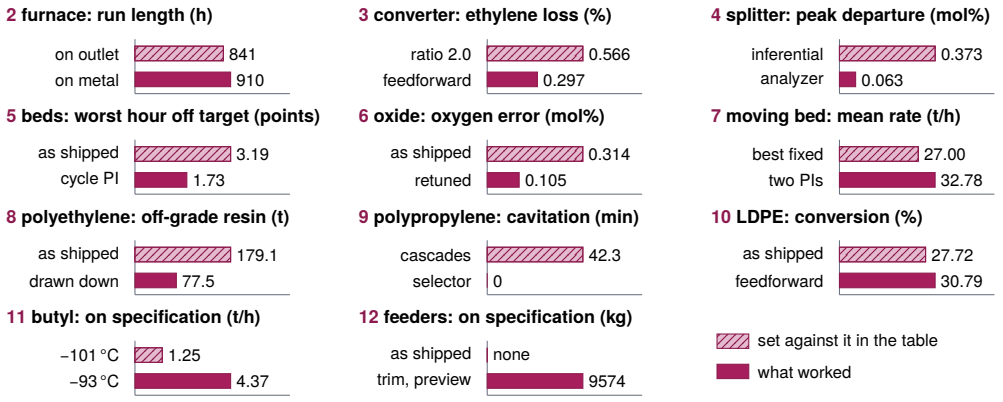

Figure 13.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 13.1 and 13.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 2, 7, 10, 11 and 12. Notice that on the polypropylene pump and the feeders the structure that worked changed all of the result.

A guard buys the equipment, not the product. The polypropylene override kept the pump out of cavitation and the reactor running, and made 110.6 t against 45.6 t. It also left the melt flow rate at 6.55 against a setpoint of 8, so 103.3 t of its polymer was off specification, more than the run that destroyed the pump. Read by off-specification tonnage alone, the two runs rank the wrong way round. The butyl reactor shows the converse. With its conversion loop in manual, as an override that gave way to the cooling margin would leave it, the poisoned reactor lost tonnes faster than it saved hours, 4.00 t/h against 4.12 t/h in automatic (Section 11.6). Whether a guard pays depends on whether the end of a run is abrupt or creeping.

A late quality is held by what is measured early, and lost to a measurement that has stopped. The acetylene converter passed its day only with feedforward from the inlet analyzer, ahead of a 20 min drum. The LDPE pulse feedforward raised the worst margin over two pulses from 33.1 to 41.9 K (Table 10.5), and the grade preview cut the feeders' transition material from 541 to 410 kg (Table 12.4). Where the analyzer is fast against the plant, as a 20 min chromatograph is against a splitter that settles in about twelve hours, it can close the loop directly. The danger is the analyzer that stops. The splitter's cascade that kept integrating a frozen result settled at 98.90 mol%, worse than the tray temperature alone, and the feeders' scale read normal for the whole time that 403 kg of wrong product went to prime.

An optimum is not a setpoint. The chloride of the ethylene oxide reactor is flat in selectivity and steep in hot spot, 22.6 °C across the same two parts per million (Table 6.4). A PID with one tuning creeps toward its setpoint from below the optimum and trips the reactor from above it. The furnace's best severity, near 850 °C at 20.31 t/h per elapsed hour (Table 2.5), is visible to no continuous instrument, and the butyl reactor's warm optimum is bounded by the Mooney grade and not by the equipment. On the LDPE tube the grade set the floor as well: peaks at 305 °C kept ample margin and let the density leave its band. Each of these optima is found by running the unit to the end of a campaign or a grade, which no loop does.

Cycles and grade changes are prepared, not chased. Prediction paid where it carried a ceiling. With three dehydrogenation beds on line no controller met the target, because the regeneration could not reach it under the bed limit. Drawing the polyethylene bed down before the grade change cut the off-grade resin from 179.1 to 77.5 t, and the polypropylene challenge passed by raising the pressure first. The predictive controllers were mixed. On the LDPE tube the plan that carried the ceiling converted 31.54 % with 12.8 K to spare against 31.09 for the loops with the requests clipped by hand (Table 10.6). On the coking furnace it cut the time above the tube limit from 260.5 to 228.5 h (Table 2.6). It lost to the prepared polyethylene cascades, 85.3 against 80.9 t (Table 8.10), and tied the dehydrogenation PI. On the polypropylene loop most of its gain came from a colder reactor, and the margin written as a constraint was worth 18.6 t on the same temperature path (Section 9.6.6). Volume 9 treats that choice in full.

> ⧖ HISTORY **1991 to 2026: The Laboratory Still Corrects the Model**
>
> The timeline of Chapter 1 closed in 1991 with a scheme that predicted the melt index and density of an industrial polyethylene reactor from its temperatures and gas compositions, and corrected the model with infrequent laboratory results [1]. A 2013 review of soft sensors found model maintenance the most critical issue in their industrial use. It listed a recursive model with a bias update that predicted melt index through grade changes, and noted that recursive methods cannot cope with abrupt changes [2]. In 2026 a study of 425 industrial polypropylene runs found that models lost accuracy when moved between product families [3], and a second study kept the laboratory method as the certification reference, with the model estimating between results [4]. The design of 1991 is still the design. Gap 2 of Section 13.5 asks what it lacks.

13.2 Choosing a Method

The chapters suggest an order of questions for a unit that holds a rarely measured quality near a limit. Each names the chapter where its answer decided the result. Figure 13.2 sets the same questions out as a flowchart.

1. *Where is the limit measured, and how far can the truth run ahead of it?* Size the setpoint from Equation (1.13) with that distance in it, and balance or protect on the variable nearest the limit (Chapter 2, Chapter 10).
2. *Is the guard one-way, faster than the quality loop, and told what was applied?* A low selector with tracking keeps the losing loop from winding up, and its indicator should say whether the constraint holds the ceiling, not only whether it won the last scan (Chapter 3, Chapter 9).
3. *Does the process gain keep its sign over the range?* Near an optimum it does not. Hold a fixed policy, search slowly, or keep the handle on the side where the sign is known (Chapter 6).
4. *How late is the quality against the plant's own time scale?* An analyzer fast against the plant can close the loop. A slow one needs the load fed forward from what is measured early (Chapter 4, Chapter 3).
5. *What happens when the analyzer stops or lies?* Arm a validity hold, carry an inference through the outage, and never let a laboratory trim learn from a sample that the diagnostics already distrust (Chapter 4, Chapter 12).
6. *Does the unit run in cycles?* Correct once a cycle on the cycle's report, with the gain the delay of reports allows (Section 1.4.7), and first check that the cycle can reach the target at all (Chapter 5, Chapter 7).
7. *Is the binding constraint the grade or the equipment?* Find the optimum per elapsed hour, wash or decoke on the quantity that ends the run, and expect the grade to bound it (Chapter 2, Chapter 11).
8. *Can preparation do what control cannot?* Draw the bed down, move the pressure, preview the grade (Chapter 8, Chapter 9, Chapter 12). Then ask what a predictive controller adds over the prepared structure at the same target and ceiling.

13.3 Where Practice Is Heading

Inferred quality learns from data. Soft sensors are widely built in the process industries for quality prediction and monitoring, and a 2021 survey reviews deep learning for them, the demands it meets in practical applications and the problems still open [5]. In the editor's view the inference will grow more capable faster than the logic that decides when to trust it. The chapters of this volume failed at that logic, not at the inference: the splitter's cascade integrated a stale analyzer, and the feeders' trim had to be kept from learning a holdup as a loss.

Nonlinear prediction reaches polymer reactors. Nonlinear model predictive control has been reported on an industrial polymerization process, where it controlled polymer quality across production rates better than reactor temperature targets computed offline from the same model [6]. For tubular LDPE reactors, optimization-based strategies combining nonlinear predictive control with fast state estimation were proposed to regulate the process and its economics under fouling [7]. In the editor's view such controllers will carry more of the limits of this volume in their plans, as the LDPE controller of Chapter 10 carried its ceiling. The comparisons of this volume suggest testing each against a structure given the same preparation.

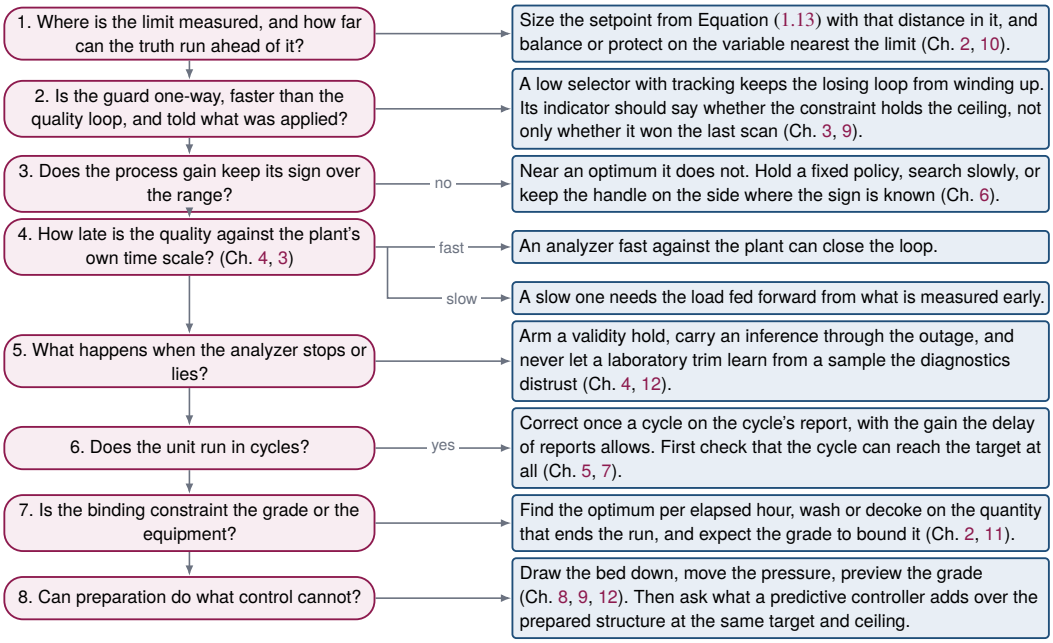

Figure 13.2 The eight questions of this section in order, with the list's answer and the chapters where it decided the result. Question 4 branches on the list's two cases. Notice that questions 1, 4 and 5 ask about a measurement before any structure is chosen.

Cyclic units are operated by optimization. Optimization-based control of a reactive simulated moving bed was reported in 2004 [8], and a 2007 article treats feedback control for optimal process operation [9]. The paraxylene ring of Chapter 7 found no fixed point that passed sixty cycles of aging, while two PIs on cycle reports did. In the editor's view the correction from one cycle to the next will stay the core of such control, with an optimizer setting what it corrects toward.

The cracker's fuel and feed change. Steam cracking emits large quantities of carbon dioxide, and a 2020 review names enhanced heat transfer in the radiation section, green energy and novel furnace designs as the keys to reducing it [10]. Pyrolysis oils from plastic waste are expected to become a steam cracker feed, but their contaminants drive corrosion, fouling and downstream catalyst poisoning, and the crackers rely on a stable and predictable feed [11]. In the editor's view both changes act on the limits of this volume. A new firing pattern moves the tube metal that Chapter 2 balances, a fouling feed shortens the run whose optimum it finds, and a poisoned converter moves the window of Chapter 3.

Electric heat reaches the cracker. Industrial steam crackers use gas-fired furnaces, which emit about 366 Mt of carbon dioxide a year worldwide. A 2022 review finds that improvements to the present designs are limited and cannot reach the reductions required. Electric heating avoids the flue gas losses, and resistance heating, promising for naphtha, is relatively easy to scale up and can retrofit existing crackers. It still has to be tested and demonstrated, and a cracker's heat is highly integrated: its steam drives the compressor turbines [12]. In the editor's view the tube metal limit of Chapter 2 would stay, held by power instead of fuel.

Spectroscopy moves into the melt. Near-infrared and Raman spectroscopy of molten polymer, measured in line, have been researched intensively and are still sparsely applied in industry, because

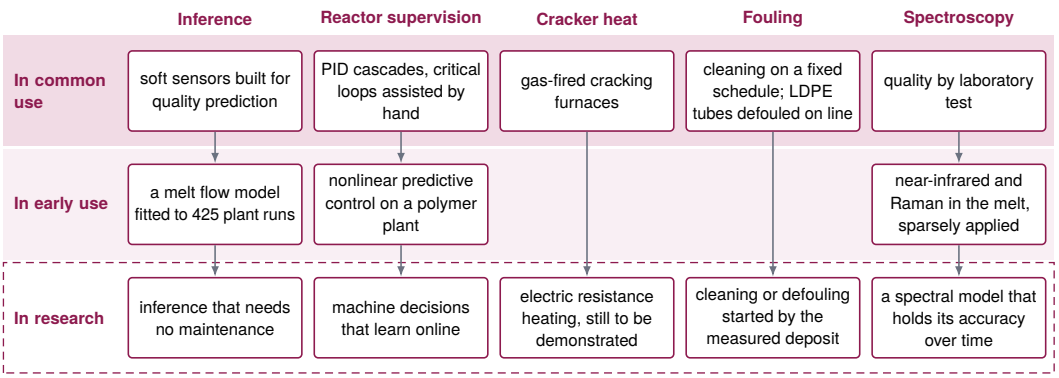

Figure 13.3 The developments of this section in five threads. The top lane holds what the cited source calls standard, established or current practice, the middle lane what it reports in service at some sites or on plant data, and the bottom lane what it reports from simulation or as a need. An empty cell has no cited development. The placing is the editor's reading of the sources, and no box carries a date. Sources by thread: inference [2, 3, 5], reactor supervision [6, 17], cracker heat [12], fouling [15, 16] and spectroscopy [4, 13, 14]. Notice that no thread has a box in common use that reads the product or the deposit in line.

their readings are unpredictable and unstable [13]. A 2021 review of extrusion monitoring reports a study of polymer additive content in which a model trained on near-infrared spectra from one time did not keep good accuracy on spectra taken at other times [14]. In the editor's view such a reading would shorten the three-hour wait of Chapter 12, once it passes the validity test of gap 2.

Cleaning follows the deposit. A 2025 review of fouling in thermal processing describes sensors that detect the deposit as it forms and can trigger cleaning on the actual buildup rather than on the conventional fixed schedule [15]. For tubular LDPE reactors, which are defouled on line by pressure cleaning or thermal shock, a 2022 study let a soft sensor of the layer's thickness start the defouling when the layer passed its limit, under a nonlinear predictive controller, on a process simulation [16]. The butyl wash of Chapter 11 applies the idea to a deposit inferred from the jacket. Figure 13.3 sets the developments of this section on three lanes, by how far each has come into use.

13.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [18]. In the editor's view, agents can now carry out much of the work that used to fill the hours of a control project. They derive a model from balances or fit one to plant data. They write the simulator and its tests, code the controller, tune it against stated criteria, script commissioning checks, write the documentation and watch running loops for degradation. The apps and chapters of this series were produced that way under the editor's direction, and the account that follows this section reports what the agents got wrong in this volume and how it was caught.

In the editor's view, this changes where an engineer's value lies. When a model, a controller and a report can each be produced in minutes, producing them is no longer the scarce skill. The work moves to five things.

Quality control. An agent's output arrives as a claim: this model fits, this tuning settles in four minutes, this code implements the equation. The engineer checks claims against the plant or the

simulator. A reported number is measured again on a fresh run at the stated seed. The code is read for what it clamps, what it initializes and what it reads, because a display that shows the controller's model in place of the plant looks correct until the two differ. A tuning is tried on a disturbance the agent did not test. The most useful checks are independent of the author, made by a second agent or a person who has not seen the first answer, and they are checks that can fail.

Supervision. The engineer sets the scope of each task and decides what an agent may touch. Reading a historian and proposing a tuning is one level of access. Writing a setpoint, downloading a configuration or changing an interlock is another, and it passes through the plant's management of change with a named approver. Supervision also fixes the order of work: a baseline before a change, one change at a time and a comparison that shows what moved.

Creativity. An agent answers the question it is given. Choosing the question remains the engineer's task: which loop is worth improving, whether a disturbance calls for feedforward, a cascade or a change to the process, which step test the operation can tolerate, and which comparison would show that a predictive controller earns its place over a retuned PID loop. The experiment that separates two explanations is designed by someone who can imagine both.

Process insight. A model contains what its author put in. The equipment does more. A valve sticks, the dead time of an analyzer grows when its sample flow drops, a thermowell lags, a heat exchanger fouls and an operator intervenes at the same hour each night. The engineer who knows the unit recognizes a result that is too good, asks what the model left out and knows which measurement to distrust first.

Accountability. Responsibility does not transfer to a tool. Functional safety standards assign the activities of the safety lifecycle to persons and organizations of stated competence, and they require an assessment that includes someone independent of the design team [19]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

These five are the professional form of the habit introduced in Section 1.7: check before trusting, and know enough to check.

13.4.1 How This Volume Was Made

Each of the eleven apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace every number to a run, and a referee report was written on each paper. Four of the eleven reports are titled self-referee reports. The others state a reviewer's stance without saying who wrote them. Seven papers had a further referee pass when a predictive controller was added. Defects went into a log for each app with a severity, a cause, a suggested fix and a status. For this book one agent per chapter reran the numbers its exercises use. The editor set the scope of each round and the rules on what each agent could change, and statements that runs contradicted were left for the editor to decide.

The defect logs. Ten of the eleven logs hold 99 entries with a severity. Of these, 37 are wrong behavior or output of an app, 47 are documentation, 12 are cosmetic or display and 3 are other kinds, such as a fragile test. The splitter's log is written under headings without a severity field and is not counted. Two logs hold a single entry each, and four apps keep a second log in their own folders, which is not counted either. Five cases show the errors that matter for a volume about quality and limits.

An inferred margin that read the truth. The polypropylene override and its predictive controller acted on a bubble-point margin described in the code, the app and the paper as inferred from the pressure transmitter, the thermowell and the chromatograph. Only the pressure was a measurement. The bubble point came from the plant's true composition. An audit that compared two plants with identical readings, one with a quarter more light inerts, found the margins 0.29 bar apart after 1.5 min and the hydrogen valve 16.8 % lower within 12 min. The margin is now computed from instruments, and every number of a run in which the override or the predictive controller acts moved. An audit of the same kind found the acetylene converter's predictive model built from the run's own kinetic settings. Doubling a rate constant moved its ratio setpoint from 1.827 to 1.906 with nothing measured changing. Three more entries, on the dehydrogenation and moving-bed apps, record start-up values and an operator aid taken from the plant rather than its instruments, with no shipped run moved.

A laboratory trim that learned a holdup. On the feeders a powder holdup below the scale made the first assay read 626 ppm against a prediction of 1000 ppm, and the trim took the shortfall as a lasting loss of delivery. After the sock was cleaned the lane was asked for 1/0.813 of its recipe, the product averaged 1217 ppm against a band edge of 1100 ppm, and the estimate still read 1000 ppm while the product stayed on prime. By 3 h the off-specification prime reached 1741 kg against 598 kg with the trim off. The paper agent found it on a second reading, reported it with the fix, and the owner approved the change. A result taken under the holdup alarms now leaves the trim at zero.

A selector that looked idle, and a cycle that did not exist. The polypropylene referee report made three findings that bear on the selector. The indicator that said the constraint "has the pen" was true for 34.8 of the 198 minutes after selection began, while the margin held a mean of 1.99 against its 2.00 bar setpoint, because tracking pulls the losing controller onto the winning value. The app's source, its guide and the brief all claimed a forty-minute limit cycle when the idle constraint controller was parked at the top of its range. The referee emulated it and found none. The draft's central comparison changed the architecture and the tuning at once, and the comparison with one variable changed became Table 9.5.

Protections that did not count what they guarded. The furnace's tube metal could jump past its limit at the start of a decoke, to 1109.0 °C on one noise seed, while the over-limit tile and the alarm skipped decoking time and the run passed. The feed is now withdrawn no faster than the firebox cools, and decoking time counts. On the LDPE tube the decomposition of case 5 turned on one thermocouple scan, so a change in the order of random draws could have made the case uncompletable. The text now says so, and a test fails if the case stops decomposing.

A steady state that had not converged. The ethylene oxide app computes "the steady state with every feed fixed" with a helper that stops after 220 iterations. The chapter agent found that off the design point the helper had not converged. The oxygen margin printed at a coolant of 241.0 °C in the fixed-feed column of Table 6.6, 0.39 mol%, is −0.50 mol% when the solve is run to convergence, inside the trip guard, and a long run of the app with every station in manual trips on oxygen at 1279.5 min. The slow part is the purge, with a time constant near 2000 min. The printed values stay as the paper has them, under the rule that leaves a contradicted statement to the editor, and the defect is logged for the editor and the app owner. The fixed-feed column is therefore not a measured result, and this chapter uses none of its values. The chapter's lesson holds and is stronger: cooling with the feeds fixed eventually trips the unit. The chapter agents that followed were told to check the convergence of any helper behind a number they used. Five of their records report the check, and in each the helper converged.

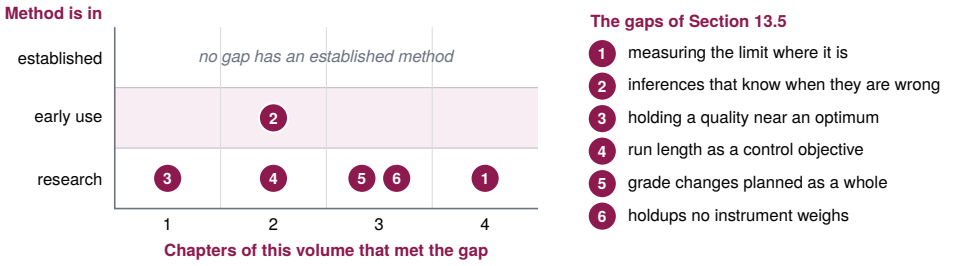

Figure 13.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: four for gap 1 (Chapters 2, 3, 9 and 10), two for gap 2 (4 and 8), one for gap 3 (6), two for gap 4 (2 and 11), three for gap 5 (8, 9 and 12) and three for gap 6 (3, 11 and 12). Up is the maturity of the method available, which is the editor's reading of the sources cited in each item and not a measurement. Notice that only the gap of inferred qualities has a method in use anywhere.

The book round. Each chapter agent first confirmed that its paper had not changed since conversion, and none of the eleven had. The eleven chapter records list 170 rows of numeric verification, each a stated value or group of values set beside its reproduction. Ten records report that every rerun number reproduced, and the eleventh is the ethylene oxide record above. The book round's defect list holds ten entries for this volume. Beside the ethylene oxide helper they are mostly wording, display and stale numbers in the apps' own texts. One briefing calls a speed of 200 min/s "200 times", which is 12 000 times real time, and one exported column holds ten times the chloride its name gives. Three chapter statements were aligned with their own tables. The acetylene conclusion had the loss passing one percent 1.58 min before the trip, where the trip came 3.95 min after the change. The agent writing Chapter 1 first claimed that each chapter meets at least two of the industry's defining problems, and the splitter meets one.

Every error above was caught by a check that could fail: two plants with the same readings, a selector flag set beside the margin it should explain, a cycle emulated instead of quoted, a solver run to convergence and then against the dynamic app. Four working rules follow. Ask what every guard reads, and from which instrument. Count the time a protection is meant to prevent, including the time it is busy. Do not let a quality correction learn from a sample its own diagnostics distrust. Run every helper to convergence before printing what it returns.

13.5 Gaps in Petrochemical and Polymer Control

Six gaps belong to the control of petrochemical and polymer units, and the gaps common to agent-era practice follow them in Section 13.5.1. In the editor's view none is closed. A sentence on difficulty or progress that carries no citation is the editor's view. Figure 13.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the units in the picture that opens this chapter.

1. *Measuring the limit where it is.* The furnace's tube metal sits behind coke, the LDPE gas runs ahead of every thermocouple and the pump's margin is inferred from a composition (Chapter 2, Chapter 10, Chapter 9). On the acetylene converter the selectivity was lost minutes before the bed rise reached its trip (Chapter 3). Methods that bound the distance between the measured and the true limit, and size the back-off from that bound, are not routine. The Open challenge box after this list takes the gap further.

2. *Inferred qualities that know when they are wrong.* The splitter's cascade trusted a frozen analyzer until a hold was armed (Chapter 4), and the polyethylene cascades of Chapter 8 read the bed quality without delay, a luxury no plant has. A validity layer that ranks the analyzer, the laboratory and the inference, and switches between them without a bump, is still designed by hand for each

unit. A 2026 study of industrial polypropylene notes that most published soft sensors are validated within a single grade family or a limited operating region [3]. It is hard because recursive updates cannot cope with abrupt changes [2], and a grade change is one. Progress would be an inference that says when it has left the region of its data.

3. *Holding a quality near an optimum.* The ethylene oxide chloride defeats a feedback loop at its optimum, and a search hunts across the top (Chapter 6). A method that holds a flat optimum next to a steep limit, with a guarantee on the limit, remains open. A 2022 review of optimization by feedback explains why: near the optimum the steady-state relation between input and cost is typically flat, so a gradient fitted to plant data is least accurate exactly there. Search without a model converges in hours to days on chemical processes, one of the main reasons it is not used widely there [20]. Progress would be a search that keeps a stated margin on the steep side.

4. *Run length as a control objective.* The furnace's best severity and the butyl reactor's wash are decided over a campaign that no loop sees (Chapter 2, Chapter 11). Joining the decoke and wash schedule to the loops that spend the margin, with the coking or fouling rate estimated as it changes, is an engineering problem with little routine practice behind it. Across reported coke properties that differ by orders of magnitude, a 2020 study found that the predicted run of an industrial ethane furnace ranged from 20 to 227 days [21]. A 2025 review finds models still far from scheduling whole exchangers over their cleaning cycles [15]. The gap is hard because the rate that ends the run is the least known number in the model. Progress would be a run planned on a fouling rate estimated from the running unit and revised at each decoke or wash.

5. *Grade changes planned as a whole.* The polyethylene bed drawdown and the polypropylene pressure were preparation decided by hand, and the preview of the feeders cut transition material by about a quarter (Chapter 8, Chapter 9, Chapter 12). A planner that chooses inventory, pressure and the quality trajectory together, within the limits of every unit downstream, is not yet standard. A 1992 study optimized the grade changes of a gas-phase polyethylene reactor on a model: lowering the bed level and the catalyst feed during a change cut the off-specification polymer significantly, and the policy needed feedback, because disturbances and model mismatch move the trajectory [22].

6. *Holdups no instrument weighs.* The feeders' sock, a bed's green oil and a jacket's fouling are inventories inferred from their effects (Chapter 12, Chapter 3, Chapter 11). Estimating them from torque, temperature or a laboratory result, and saying when the estimate cannot see them, is needed before a correction is allowed to learn. A 2025 review finds a notable gap in the in-line measurement and prediction of fouling [15], and on LDPE reactors fouling of the inner wall is difficult to assess and its prediction is not systematically formulated in normal practice [16]. It is hard because a holdup shows only through effects that a change of feed also produces.

⚠ OPEN CHALLENGE A Limit Read Behind the Truth

Gap 1 is the one most chapters met. A 2021 review of reactor runaway finds that knowledge of the phenomenon has improved a lot but is not fully integrated into practice, and it names as a research direction the handling of parameter uncertainty in detecting a runaway during operation [23]. A margin kept from a constraint to avoid violating it in transients costs a loss that grows linearly with the margin [20]. Progress would be a stated bound on how far the truth can run ahead of each reading, carried into the setpoint and the trip.

13.5.1 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 [19]. In the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

4. *Security of agent access to control systems.* An agent with credentials to a control system is a new path into it. Industrial security standards state the requirements a control system must meet at each security level [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 limit is often read in the wrong place, and a guard is only as good as that reading.
- A guard buys the equipment, not the product.
- A late quality is held by what is measured early, and lost to a measurement that has stopped.
- Cycles and grade changes are prepared, not chased, and an optimum is not a setpoint. Prediction paid where it carried a ceiling.

References

[1] K. B. McAuley and J. F. MacGregor, "On-line inference of polymer properties in an industrial polyethylene reactor," *AIChE Journal*, vol. 37, no. 6, pp. 825–835, 1991.

[2] 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.

[3] J. Hernández-Fernández and J. Lopez-Martinez, "Reactor-grade non-additive polypropylene architecture defines the physicochemical limits of melt flow index model transferability: An industrial chemometric and machine learning study," *Polymers*, vol. 18, no. 15, p. 1880, 2026.

[4] J. Hernández-Fernández, "Temporally validated and laboratory-reconciled soft sensing of xylene-soluble fraction in industrial polypropylene manufacturing," *Polymers*, vol. 18, no. 19, p. 2401, 2026.

[5] 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.

[6] R. Bindlish, "Nonlinear model predictive control of an industrial polymerization process," *Computers & Chemical Engineering*, vol. 73, pp. 43–48, 2015.

[7] V. M. Zavala and L. T. Biegler, "Optimization-based strategies for the operation of low-density polyethylene tubular reactors: Nonlinear model predictive control," *Computers & Chemical Engineering*, vol. 33, no. 10, pp. 1735–1746, 2009.

[8] A. Toumi and S. Engell, "Optimization-based control of a reactive simulated moving bed process for glucose isomerization," *Chemical Engineering Science*, vol. 59, no. 18, pp. 3777–3792, 2004.

[9] S. Engell, "Feedback control for optimal process operation," *Journal of Process Control*, vol. 17, no. 3, pp. 203–219, 2007.

[10] I. Amghizar, J. N. Dedeyne, D. J. Brown, G. B. Marin, and K. M. Van Geem, "Sustainable innovations in steam cracking: CO_2 neutral olefin production," *Reaction Chemistry & Engineering*, vol. 5, no. 2, pp. 239–257, 2020.

[11] M. Kusenberg, A. Eschenbacher, M. R. Djokic, A. Zayoud, K. Ragaert, S. De Meester, and K. M. Van Geem, "Opportunities and challenges for the application of post-consumer plastic waste pyrolysis oils as steam cracker feedstocks: To decontaminate or not to decontaminate?" *Waste Management*, vol. 138, pp. 83–115, 2022.

[12] M. E. H. Tijani, H. Zondag, and Y. Van Delft, "Review of electric cracking of hydrocarbons," *ACS Sustainable Chemistry & Engineering*, vol. 10, no. 49, pp. 16 070–16 089, 2022.

[13] S. Zhu, Z. Song, S. Shi, M. Wang, and G. Jin, "Fusion of near-infrared and raman spectroscopy for in-line measurement of component content of molten polymer blends," *Sensors*, vol. 19, no. 16, p. 3463, 2019.

[14] N. Munir, M. Nugent, D. Whitaker, and M. McAfee, "Machine learning for process monitoring and control of hot-melt extrusion: Current state of the art and future directions," *Pharmaceutics*, vol. 13, no. 9, p. 1432, 2021.

[15] J. Tarapata, T. R. Murphy, E. W. Finnegan, T. F. O'Callaghan, and J. A. O'Mahony, "Approaches for measuring and predicting fouling during thermal processing of dairy solutions," *Comprehensive Reviews in Food Science and Food Safety*, vol. 24, no. 4, p. e70209, 2025.

[16] F. S. Rohman, M. R. Othman, D. Muhammad, A. Azmi, I. Idris, R. A. Ilyas, S. E. Elkhatif, and M. N. Murat, "Nonlinear control of fouling in polyethylene reactors," *ACS Omega*, vol. 7, no. 44, pp. 39 648–39 661, 2022.

[17] Y. Liu, "Natural intelligence principle of supervisory-level control for condensed-mode polyethylene fluidized bed reactors: A proportional-integral-derivative perspective," *ACS Omega*, vol. 11, no. 39, pp. 60 064–60 072, 2026.

[18] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, W. X. Zhao, Z. Wei, and J. Wen, "A survey on large language model based autonomous agents," *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024.

[19] International Electrotechnical Commission, "Functional safety – Safety instrumented systems for the process industry sector – Part 1: Framework, definitions, system, hardware and application programming requirements," IEC 61511-1:2016, Edition 2.0, 2016.

[20] D. Krishnamoorthy and S. Skogestad, "Real-time optimization as a feedback control problem – a review," *Computers & Chemical Engineering*, vol. 161, p. 107723, 2022.

[21] M. Geerts, S. H. Symoens, P. A. Reyniers, G. B. Marin, M.-F. Reyniers, and K. M. Van Geem, "Steam cracking coke properties and their influence on furnace run length predictions: Experimental and modeling study," *Industrial & Engineering Chemistry Research*, vol. 59, no. 52, pp. 22 460–22 472, 2020.

[22] K. B. McAuley and J. F. MacGregor, "Optimal grade transitions in a gas phase polyethylene reactor," *AIChE Journal*, vol. 38, no. 10, pp. 1564–1576, 1992.

[23] A. Kummer and T. Varga, "What do we know already about reactor runaway? a review," *Process Safety and Environmental Protection*, vol. 147, pp. 460–476, 2021.

[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.

13.6 Exercises

13.1 Discussion · 8 min

An agent proposes a pass-balancing scheme for a new cracking furnace that equalizes the coil outlet temperatures, and reports that its simulation holds the four outlets within half a degree. Using Chapter 2 and Chapter 10, explain why that report does not show the scheme is safe, and list what you would ask for before accepting it.

13.2 Discussion · 7 min

A plant manager ranks two runs of the polypropylene loop of Chapter 9 by off-specification tonnage and concludes that the override made things worse. Use Table 9.5 to show why the ranking is wrong, and propose three measures a scorecard for a quality loop under a safety limit should report beside off-specification tonnage.

13.3 Discussion · 8 min

Three chapters show a quality measurement failing quietly: the frozen chromatograph of Chapter 4, the assay taken during a holdup in Chapter 12, and the cascades of Chapter 8 that read the bed quality without delay. For each, say what the loop believed, what was true, and what logic would have protected it. Then state a general rule for when a quality correction may learn from a measurement.

13.4 Discussion · 7 min

The chloride of Chapter 6 has a flat selectivity optimum near 2.5 ppm and a steep hot spot. Explain why a PID controller on selectivity cannot hold that optimum, why a step-and-search optimizer hunts across it, and argue for an operating policy that respects the hot spot limit.

13.5 Discussion · 10 min

Section 13.4.1 reports that an app's steady-state helper stopped after 220 iterations, so a printed oxygen margin of 0.39 mol% was −0.50 mol% when run to convergence. Explain why such an error passes ordinary review, describe the checks that caught it, and write a procedure of five to eight steps for an agent that computes steady states or limits that a person will rely on.

13.6 Discussion · 8 min

A vendor claims its predictive controller cut off-specification polymer on a polypropylene grade change from 83.2 to 43.7 t. Using Table 9.6, Chapter 8 and Chapter 10, describe the comparison you would require before crediting the controller, and say what each chapter's fair comparison showed.

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) $t^{*} = 1000\,\mathrm{h}$, run 917.9 h. (b) 2500 h, and 20.8 K after 500 h.

1.3 The hotter policy, 15.342 t/h against 15.200 t/h, about 0.9 % more.

1.5 (d). Breakthrough puts all of the product off grade, while each extra mole costs a little ethylene and moves the bed toward light-off.

1.7 25.7 kg/s, 17.1 % of the recycle.

1.9 0.625 K per point for the double root at 0.5, and 2.5 K per point at the edge.

Chapter 2. Ethane Cracking Furnace Pass Balancing

2.1 (c). Coke cuts the flux that reaches the gas and raises the metal above the gas, and a cold outlet calls for less feed.

2.3 (a). With the flux fixed, Equation (2.9) holds with $\mathscr{S} = q$, and its solution is $\dot{R}_0/(1 - t/t^{*})$.

2.5 (a) 2.51×10^{-7} m^2 K/(W h) and 0.0268 °C/h. (b) 1120 h and 1073 h. (c) 3550 h, 3.3 times the true run.

2.7 (a) 8.4 h, 16.1 °C. (b) 20.9 °C and 34.0 °C, while the outlets stay within 0.48 °C. (c) Equation (2.8) predicts about 16 °C. The spread lies behind the coke, and the hotter pass cokes faster.

2.9 (a) MPC 2.176 to 2.225, Loops 2.986 to 3.040 (2.176 and 2.986 at 120 h). (b) MPC: 0 rows, 1087.3 °C. Loops: 153 rows, 1094.8 °C. (c) The plan's ceiling sits one degree under the override's start, and it moves feed onto pass 4 before it gives up severity.

2.11 (a) 1087 °C and 1093 °C, with true metal at 1092 °C and 1098 °C. (b) The excursion of the loops in the override's band, the thermocouple's error and a reserve. Keep 12 °C unless the thermocouples are verified.

Chapter 3. Acetylene Converter and Its Hydrogen Window

3.1 (c). The two failures are priced very differently, so the cheapest safe point is the least hydrogen that clears the edge.

3.3 (a). The heat is released near the inlet, and the bed's thermal wave takes about 2.6 min to carry it to the outlet.

3.5 (a) 1.00. (b) 1.96 min, against 2.0 min measured. (c) 113.4 kmol/h. The same ratio on more acetylene is more spare hydrogen.

3.7 (a) 14.9, 0.30 and 0.000 16 ppm, with losses of 0.137, 0.200 and 0.470 %. (b) 4.97 decades and 3.43 times, and the step from 1.50 to 1.90. (c) 7.2 t/d, against all 2680 t/d off grade at 1.30.

3.9 (a) 1079 ppm, 185 of 288 over, 0.192 %: fails. (b) 1.36 ppm, 12 over, 0.242 %: fails. (c) 0.53 ppm, none over, 0.297 %: passes.

3.11 (a) 40.3 more days on specification, at about 3.0 t/d more ethylene lost on day 150. (b) Green oil takes the activity until no inlet temperature or ratio meets the limit, and the catalyst is regenerated.

Chapter 4. Propylene Splitter Inferential Control

4.1 (b). Against a plant that needs hours to respond, a sample even an hour old is recent.

4.3 (d). Pressure enters the column twice, through the bubble point and through the relative volatility, and the difference cancels only the first.

4.5 (a) $D = 494.0\,\mathrm{kmol/h}$, $L = 7000\,\mathrm{kmol/h}$. (b) 424.0 kmol/h, −14.2 %. (c) 489.4 kmol/h, −0.93 %. A ratio move changes the product in proportion, while a flow move is multiplied by fourteen.

4.7 (a) 715 min (±2). (b) 99.560 mol% and 99.586 mol% (±0.002), 29 and 53 percent of the change. (c) The fit sees a ramp. It returns a time constant of about 3.1 h against the 5.2 h at which the step is 63 percent complete, and a gain 43 percent low.

4.9 (a) 0.98 kmol (±0.01), 942.3 min (±0.5), lowest 99.459 (±0.002) at 361 min. (b) 8.07 kmol, 355.8 min, lowest 99.364 at 366 min. (c) 4.26 kmol, 111.0 min, lowest 99.491, so no minute is more than 0.01 mol% below the specification.

4.11 Any two of: the hydraulic capacity limit (flooding, weeping, entrainment), the energy balance, and tray efficiency that falls near flooding. The app accepts any boilup up to its range and keeps separating.

Chapter 5. Cyclic Propane Dehydrogenation Control

5.1 (a). No move reaches a bed on line: the hotter air charges beds that come on line later, and a second rise follows a cycle on.

5.3 (c). The bed stays hotter on line and comes back hotter from its regeneration, which makes up most of the lost activity.

5.5 (a) 0.334 bar. (b) 63.3 and 67.5 %. (c) 62.2 %, about one point lower. (d) About 0.08 of the rise: the heat stored in the beds caps the conversion.

5.7 (a) 630.8 °C, 86.9 %, 0.75 to 0.76 t/h, 2.0 y and 60.73 %. (b) 0.293 points per K of air and 0.357 per K of start temperature. (c) 0.647 and 0.755 t/h, within the tile's flicker.

5.9 (a) Shipped: worst hour ±3.2 points, hottest 660 °C, ripple 58.6 %, fuel 0.59 t/h; it fails. PI with two on line: ±1.7, 658 °C, 6.0 %, 0.50 t/h; it passes. (b) ±0.1 then ±1.7: hour 4, the close-up after the lost bed. (c) The price is throughput, 11.49 t/h of propylene on average against 13.66 t/h.

5.11 720, 480 and 360 a day. The model counts the swing's cost in selectivity, ripple and catalyst life, and leaves out valve wear, purge steam and vacuum.

Chapter 6. Ethylene Oxide Selectivity and Runaway

6.1 (d). The reaction that releases most of the heat is also the one that accelerates faster, so the heat grows with an effective activation between the two.

6.3 (b). Cooling relieves the hot spot and spends the flammability margin, while less oxygen feed relieves both.

6.5 (a) $7.606 - 1.20 \times 0.0720 = 7.520\,\mathrm{kmol/min}$, and $0.5 \times 8.00 + 3 \times 1.173 = 7.520\,\mathrm{kmol/min}$, with 0.940 mol of oxygen per mol of product. (b) 7.91 kmol/min (±0.01). (c) 1.064 kmol/min per kmol/min and 7.66 min.

6.7 (a) 286.6 °C (±0.1) at 25.9 min, and 6.9 min (±0.1) above the alarm. (b) 267.6 °C, 5.68 mol% and 7.592 kmol/min (±0.1, ±0.01 and ±0.002). The recycle composition did. (c) 283.2 °C (±0.1) at 24.9 min and 4.6 min (±0.1) above the alarm.

6.9 (a) 0.314 mol%, fails, margin 1.73 mol%. (b) 0.077 mol%, passes, margin 1.56 mol%. (c) 0.105 mol%, passes, margin 1.75 mol% (±0.002 on the errors, ±0.01 on the margins). No run trips. Their hot-spot peaks are 272.6, 272.6 and 274.4 °C, respectively.

6.11 A ceiling of 247.47 °C. Any two of these: one tube stands for 9200, so the worst tube is missing, twelve cells understate the peak, the boundary was found at the design composition and catalyst activity, and the gas loop has no transport delay.

Chapter 7. Paraxylene Simulated Moving Bed Cycle Control

7.1 (d). At the cyclic steady state the profile repeats itself shifted by one bed, and only the cycle average is steady.

7.3 (b). A larger extract flow lowers Q_{II} and Q_{III}, so the extract becomes less pure and the raffinate purer.

7.5 (a) $\kappa_A = 0.220$, $\kappa_B = 0.312$, $\omega_G = 1.053$, $\omega_F = 0.357$. (b) $w = (0.421, 0.823)$. (c) 0.895 and 0.422: the simulated unit loses the recovery a little below the upper edge and needs a large margin from the left edge.

7.7 (a) 1.81 t/h at 99.98 %, then 12.67 t/h at 99.86 %. (b) 18.18 t/h is fed, so the extract carried 10.0 %: the rest filled the ring, and almost none reached the raffinate. (c) 120 min, the end of cycle 10. The tiles read 99.893 %, 99.53 % and 18.10 t/h, and the extract swings from 99.40 to 99.98 % inside the first switch of cycle 11.

7.9 (a) 99.80 %, 82.6 %, 30.6 t/h: the purity and the production pass, the recovery fails. (b) 99.75 %, 97.5 %, 32.8 t/h, a pass, with $m_2 = 0.500$ and $m_3 = 0.769$. (c) 99.76 %, 96.8 %, 32.7 t/h: the recovery fails.

7.11 No: the fault is hidden in the setpoints. The offset of 0.024 to 0.025 in m_2 and m_3 matches 11.7 $\mathrm{m^3/h}$, 3 % of the extract flow. Compare the commanded ratios with those the same purities needed before, and check the extract and raffinate flows against a second measurement downstream.

Chapter 8. Fluidized-Bed Polyethylene Cascade Control

8.1 (c). Sensible 19.85 MW plus wall loss 0.32 MW leave about 5.9 MW of the reaction heat for the condensate to carry.

8.3 (a). The temperature loop hides the approach by opening the valve, so the valve position and the utilization move first.

8.5 (a) 11.3 min. (b) $k_d = 1.97 \times 10^{-3}/\mathrm{min}$, a catalyst time constant of 163 min and an active inventory 0.923 of its old value. (c) 1.072, and a product of 0.990.

8.7 (a) 40.00 °C and 38.54 °C (±0.01), a gain of -0.146 ± 0.002 °C per percent. (b) 0.51 min (±0.05), faster than the 0.65 min of the two lags. (c) 38.38 °C (±0.01) and 82.0 °C (±0.1). The bed is cooling under the outlet.

8.9 (a) MI 6.59, density 0.9321 and 115.2 t (±0.1); the peak $R_{H,\mathrm{SP}}$ is 73.8 at about 174 min; it fails, with both melt index and density outside their bands. (b) 179.1 t (±0.1), last off specification at 489.5 min. (c) MI 8.40, density 0.9352, 80.9 t and a lowest margin of 7.8 °C, a pass; still 80.9 t at 900 min, last off specification at 281.0 min.

8.11 The sample drawn at 120 min, posted at 165 min, about 83 min after the bed left the band. Roughly 42 t has gone off grade by then.

Chapter 9. Polypropylene Loop Against a Cavitation Limit

9.1 (b). The inerts are worth 3.48 bar and the hydrogen 2.23 bar.

9.3 (d). Tracking drags the loser onto the winner, so the strict comparison flickers while the constraint governs.

9.5 (a) 6.25 and 12.5 g/10min per unit. (b) $K_c = 0.60$ at the starting grade and 0.30 at the destination, with $\tau_I = 52.4$ min. The default is close to the first, and the challenge's tuning is six times the second.

9.7 (a) 34.57 to 36.02 bar for a ratio change from 0.902 to 1.500, about 2.43 bar per unit. (b) 13.8 min for the bubble point and 74.5 min for the product, a factor of about 5.4.

9.9 (a) 12.47 g/10min, –0.33 bar, 115.5 min, 47.2 t: it fails. (b) 11.88 g/10min, 2.14 bar, 0.0 min, 30.25 t: it passes. (c) The same values: the override never acted in (b).

9.11 (a) 173 kW, against 128 kW at 80 %. (b) 3.85 m/s, 1.05 times the deposition velocity, with 75 % of the jacket coefficient. Not recommended.

Chapter 10. Tubular LDPE Reactor Peak Temperature Control

10.1 (a). A trip removes future initiator, not the dose that is already reacting downstream of the injector.

10.3 (c). The feedforward knows the schedule, not the valve, so it cuts the initiator for a pulse that never opens.

10.5 (a) 28.0 kg, 21.4 m, 13.78 m/s, 108.9 s. (b) $C = 0.1235\,\mathrm{kg/bar}$ and $C_v = 0.759\,\mathrm{kg/(s\,bar^{1/2})}$. (c) $\tau_P = 2C(P - P_{sep})/F_{in} = 30.9\,\mathrm{s}$.

10.7 (a) 299.7 °C and 33.1 K. (b) 292.2 °C and 41.9 K. (c) 304.7 °C without the feedforward and 312.2 °C with it.

10.9 (a) 30.79 %, 16.5 K, met. (b) 30.81 %, 13.5 K, not met. (c) 30.86 %, 13.8 K, not met. In both runs the worst margin is set within seconds of the stronger lot's arrival at 150 min.

10.11 (a) From the licensor's process-safety data, the plant's own decomposition history and measured decomposition boundaries for ethylene and its comonomers, not from a teaching model. (b) The thermocouples sample the gas only where they stand, and they lag through the thermowell. As the pressure recovers from a pulse the hottest gas runs 10.9 K ahead of every thermocouple.

Chapter 11. Butyl Rubber Reactor Fouling and Wash Policy

11.1 (a). The loop holds the controlled variable flat, so the fouling shows only in the controller output, the jacket, and in the margin.

11.3 (c). The margin runs out about fifteen hours before the product leaves grade, so the margin wash gives away run time that was still saleable.

11.5 (a) 7.44 h/K, 48.9 h. (b) 49.3 h at −93 °C, 6.7 h at −101 °C. (c) −94.4 °C, inside the grade window from −98.6 to −91.4 °C.

11.7 (a) 0.598 and 0.012, −0.0122 per hour. (b) Hour 40.8, and zero near hour 49.0. (c) 0.010 °C, the PI error on the jacket ramp.

11.9 (a) 436.3 t and 4.54 t/h, 0.9 h off grade, 8.0 h of downtime, margin gone at hour 31.3. (b) 383.4 t and 3.99 t/h, no off-grade hour, 17.8 h of downtime, washes at 26.2, 60.2 and 94.2 h. The schedule stands in for the Mooney or the reactor temperature that ends the run, not for the margin.

11.11 (a) The heat release is limited by the catalyst. The fouling rate grows by $c_f = 0.1386$ per degree. (b) About 28 h against the 10.3 h of the challenge: the creep accelerates, so a linear extrapolation overstates the warning.

Chapter 12. Loss-in-Weight Additive Feeding

12.1 (d). The weight records what the screw removed, and the holdup is what happened to it afterward. Only the torque and the switch respond to H.

12.3 (b). With $\sigma_q = 0.002K$ per block, data older than a few minutes no longer constrain the present feed factor.

12.5 (a) 106.67 rpm, losses up to 60 %. (b) 0.6 kg/h, and $r_B > 2/3$. (c) At 36.6 min, after 26.6 min of filling.

12.7 (a) 104.5 rpm (the chapter's one-second mean is 104.6), and 997 to 1,004 ppm. (b) 0.00941, 0.00944 and 0.00947 kg/h/rpm. (c) 764 ppm. The loop makes up the loss with speed. Without the loop the product falls a quarter short.

12.9 (a) 410 kg, 3.0 min. (b) 541 kg, 6.9 min. (c) 361 kg, 1.0 min. Preview moves the recipe ahead by the line's residence, and the MPC also overfeeds on the way up and cuts early on the way down.

12.11 Weigh the sock with the hopper ($\alpha > 0$), or estimate the holdup from the motor torque. A sample answers 45 min after it is drawn, later than the switch.

Chapter 13. Summary and Outlook

13.1 A good answer notes that the outlet is not the limit. In Chapter 2 outlets matched within half a degree carried 16.1 °C of tube metal spread, and on a furnace with one degraded burner outlet balancing ran 841 h against 888 unbalanced and 910 balanced on metal, because a coked pass reads cold and loses feed. The LDPE gas runs ahead of every thermocouple, and a trip 5 K below the line fired 1.5 s too late. It asks for the run length and tube metal spread on a degraded case, the variable nearest the limit, the distance the truth can run ahead of the measurement, and a back-off sized with that term.

13.3 A good answer says the splitter's cascade believed a stale result and walked the tray setpoint from 46.80 to 46.99 °C, settling at 98.90 against 99.17 with the hold armed from the start, so a validity hold is the

protection. The feeders' trim believed a shortfall of 626 against 1000 ppm was a lasting loss when it was a holdup, so results taken under holdup diagnostics must not move the trim and a cleaning must reset it. The polyethylene cascades read a value no plant has, so their tonnages are optimistic and a delayed laboratory with a bias update is the honest loop. The rule: learn only from a measurement that is fresh, valid and taken when nothing else explains the error.

13.5 A good answer notes that the helper returned a plausible number in the expected direction with no error flag, and that the lesson of the table (two margins moving opposite ways) still read correctly, so rereading could not catch it. It was caught by checking the residual of the oxygen balance and by running the dynamic app at the same setting, which tripped at 1279.5 min. The procedure: report the residual and iteration count of every solve, fail on a cap rather than return, confirm the result with a long dynamic run or an independent solver, record the time scales (here a purge near 2000 min), keep the value as printed and log it for a person to decide, and have the next agents check every helper they use.

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.

acetylene converter A fixed-bed reactor that removes the last traces of acetylene from an ethylene stream by adding hydrogen over a catalyst. *Chapter 3*.

activation energy How steeply a reaction rate rises with temperature, written in this volume as an activation temperature in kelvin. *Chapter 1*.

adsorbent A porous solid that holds some molecules on its surface more strongly than others, so that a bed of it can separate them. *Chapter 7*.

analyzer An instrument that takes a sample of a stream and reports its composition or a property, usually some minutes later. *Chapter 1*.

analyzer validity A check that an analyzer reading is fresh and plausible before a controller is allowed to act on it. *Chapter 1*.

averaging level control Level control tuned loosely, so that a drum absorbs upsets instead of passing them downstream as changes of flow. *Chapter 4*.

back-calculation A way to stop windup: when a controller's output is limited, its integral is pulled toward the output actually applied. *Chapter 1*.

back-off The margin kept between a setpoint and a limit, so that excursions and measurement error do not carry the plant over the limit. *Chapter 1*.

bias update A correction to a model's prediction made each time a laboratory or analyzer result shows how far off the model was. *Chapter 1*.

bubble point The temperature at which a liquid starts to boil at a given pressure, or the pressure at which it boils at a given temperature. *Chapter 9*.

butyl rubber A synthetic rubber made from isobutylene with a little isoprene, polymerized as a slurry at very low temperature. *Chapter 11*.

cascade control Two loops in series, in which the outer controller sets the setpoint of a faster inner one. *Chapter 2*.

catalyst deactivation The slow loss of a catalyst's activity under deposits such as green oil or coke, which changes how the plant responds over months. *Chapter 3*.

cavitation Vapor bubbles that form and collapse in a pump when the pressure at its inlet falls too low, which damages the pump. *Chapter 9*.

chain-transfer agent A chemical added to a polymerization to end growing chains early, which shortens them and raises the melt index. *Chapter 10*.

chloride moderator A trace of organic chloride dosed into an ethylene oxide reactor to improve selectivity; too much lowers the catalyst's activity. *Chapter 6*.

coil outlet temperature (COT) The temperature of the cracked gas where it leaves a furnace coil, the variable the firing loop holds. *Chapter 1*.

coke Solid carbon that hot hydrocarbons lay down on tube walls or on a catalyst, where it insulates the wall or blocks the catalyst. *Chapter 1*.

compounding Mixing additives into a polymer, usually in an extruder, before the product is cut into pellets. *Chapter 12*.

condensed-mode cooling Cooling a gas-phase reactor by condensing part of its recycle gas, so that the liquid boils again in the bed and carries extra heat. *Chapter 1*.

constrained variable A variable that need not sit at a setpoint but must stay on one side of a limit, such as a tube metal temperature. *Chapter 2*.

cooling margin How much more cooling a jacket or cooler could still give before it reaches its limit. *Chapter 11*.

cycle-to-cycle control A loop that moves its handle once a cycle or once a report, from the results of earlier cycles; also called run-to-run control. *Chapter 1*.

cyclic steady state The state of a cycling unit in which every cycle repeats the one before it. *Chapter 5*.

dead time The time between a change in the process and the first sign of it in a measurement. *Chapter 1*.

decoke Taking a cracking furnace out of service to remove the coke from its coils, after which a new run begins. *Chapter 1*.

desorbent The liquid that pushes the held components off an adsorbent, so that the bed can be used again. *Chapter 7*.

distributed control system (DCS) The plant's control computers: field controllers that run the regulatory loops, and the stations from which operators watch and act. *Chapter 1*.

ethylene The smallest olefin, a gas made by cracking ethane or naphtha and the feed of polyethylene and ethylene oxide. *Chapter 1*.

ethylene oxide A reactive gas made by oxidizing ethylene over a silver catalyst, while part of the ethylene burns to carbon dioxide. *Chapter 1.*

extract and raffinate The two outlets of a moving bed: the extract carries the strongly held component, here paraxylene, and the raffinate the rest. *Chapter 7.*

feed factor The mass a screw feeder delivers per revolution, which can change with the powder and with how full the hopper is. *Chapter 12.*

feedforward Control that moves a handle when a disturbance is measured, before the disturbance reaches the controlled variable. *Chapter 3.*

final control element The device that changes a flow of material or energy, such as a valve, a pump or a feeder. *Chapter 1.*

flammability limit The oxygen content above which the gas of an ethylene oxide loop could burn; it moves with the ethylene and carbon dioxide in the gas. *Chapter 6.*

floating pressure Lowering a column's pressure setpoint while the condenser still has room, which can save heat at the reboiler. *Chapter 4.*

fluidized bed A bed of particles held up by gas blown upward through it, so that the bed behaves like a boiling liquid. *Chapter 8.*

fouling The slow buildup of deposits, such as coke or polymer, on a surface that transfers heat. *Chapter 1.*

gas chromatograph (GC) An analyzer that separates the components of a sample in a column and measures each as it leaves, giving one result per cycle. *Chapter 1.*

gas-phase polymerization Making polymer from a gaseous monomer on catalyst particles held in a fluidized bed. *Chapter 8.*

giveaway Product made purer than its specification, which costs energy and earns nothing. *Chapter 4.*

grade transition The change of a polymer plant from one product grade to the next, during which off-grade material is made. *Chapter 1.*

green oil A heavy liquid of small polymers that forms on the catalyst of an acetylene converter and slowly lowers its activity. *Chapter 3.*

heat of polymerization The heat released for each kilogram of polymer made, which the reactor's cooling must carry away. *Chapter 1.*

hot spot The hottest point of a cooled reactor tube or bed, where a runaway would begin. *Chapter 6.*

inferential control Control of a quality computed from fast measurements, such as temperatures and gas ratios, between slow analyzer or laboratory results; also called a soft sensor. *Chapter 1.*

initiator A chemical, usually an organic peroxide, that breaks into radicals and starts the chains of a free-radical polymerization. *Chapter 10.*

interlock Logic that forces a valve or a unit to a safe state when a limit is crossed, whatever the controllers ask. *Chapter 9.*

inventory averaging The product leaving a reactor averages everything its inventory holds, so a property changes slowly after the conditions change. *Chapter 1.*

light-off A sudden, self-accelerating rise of a bed's temperature when an unwanted reaction takes over. *Chapter 3.*

load cell A sensor that turns a force into an electrical signal through strain gauges; under a hopper it weighs the hopper and its contents. *Chapter 12.*

loop reactor A pipe bent into a closed loop, through which a pump circulates polymer particles suspended in liquid monomer. *Chapter 9.*

loss-in-weight feeder A feeder standing on a scale, whose controller measures the delivery as the rate at which the hopper's weight falls. *Chapter 12.*

low selector A block that passes the lower of two controller outputs, so that the more cautious controller wins. *Chapter 9.*

low-density polyethylene (LDPE) A branched polyethylene made from ethylene at very high pressure, with chains started by radical initiators. *Chapter 10.*

melt index (MI) A laboratory measure of how easily a molten polymer flows through a standard die; shorter chains give a higher index. *Chapter 1.*

mid-ranging Using a slow second handle to bring a fast one back toward the middle of its range. *Chapter 10.*

model predictive control (MPC) Control that predicts the plant with a model and chooses future moves of several handles at once, within their limits. *Chapter 1.*

Mooney viscosity A rubber's resistance to shearing in a standard instrument, the quality that decides whether a run of butyl rubber can be sold. *Chapter 11.*

moving horizon estimation (MHE) Estimation that fits a model to a window of recent measurements to find quantities that no instrument reads. *Chapter 12.*

multitubular reactor A reactor made of many catalyst-filled tubes inside one shell, cooled by a coolant outside the tubes. *Chapter 6.*

naphtha A light liquid cut of crude oil, one of the feeds that a cracker breaks into olefins. *Chapter 1.*

net positive suction head (NPSH) The margin by which the pressure at a pump's inlet exceeds the pressure at which the liquid would boil. *Chapter 9.*

off-grade Product made outside its specification band, as during a grade transition. *Chapter 1*.

olefin A hydrocarbon with a double bond, such as ethylene or propylene, the building block of most plastics. *Chapter 1*.

on-stream rate The production averaged over a run and the shutdown that follows it, the quantity a decoke or wash schedule should maximize. *Chapter 1*.

override control A structure in which a limit controller takes a handle from the normal controller when the limit comes near. *Chapter 1*.

parametric sensitivity The state of a reactor in which a small change of an input, such as the coolant temperature, gives a large change of its temperature. *Chapter 1*.

paraxylene One of the three xylenes, separated from the others and from ethylbenzene by adsorption in a simulated moving bed. *Chapter 7*.

pass balancing Trimming the feed of each pass of a furnace so that the passes run alike. *Chapter 2*.

PID control The proportional, integral and derivative controller that runs most of the loops in a plant. *Chapter 1*.

polymer A material made of long chains of small molecules joined together. *Chapter 1*.

polyolefin A polymer made from olefins, mainly polyethylene and polypropylene. *Chapter 1*.

pressure-compensated temperature A tray temperature corrected for the column's pressure, so that what is left follows the composition. *Chapter 1*.

propylene splitter A distillation column with many trays that separates propylene from propane. *Chapter 4*.

radiant coil The tubes in a furnace's firebox through which the feed flows and cracks. *Chapter 2*.

ratio control Holding one flow at a set ratio to another, such as hydrogen to the gas it treats. *Chapter 3*.

reflux ratio The liquid returned to the top of a column divided by the product drawn off there. *Chapter 4*.

regeneration Restoring a catalyst bed by burning off its coke with hot air, which also stores heat in the bed for the next reaction step. *Chapter 5*.

relative volatility How much more readily one component boils off than another; near one, as for propylene and propane, separation needs many trays. *Chapter 4*.

run length The time a furnace, reactor or exchanger can run before fouling forces a shutdown. *Chapter 1*.

runaway A loss of the balance between heat made and heat removed, after which a reactor's temperature rises faster and faster. *Chapter 1*.

safety instrumented system (SIS) A protective system of sensors, logic and final elements, kept apart from the control system, that brings the plant to a safe state. *Chapter 1*.

selectivity The share of the converted feed that becomes the wanted product. *Chapter 1*.

severity How hard a furnace cracks its feed, set through the coil outlet temperature or the conversion. *Chapter 2*.

simulated moving bed (SMB) A ring of fixed adsorbent beds whose inlets and outlets step around the ring, so that the solid seems to move against the liquid. *Chapter 7*.

slurry reactor A reactor in which solid polymer particles are suspended in a liquid. *Chapter 11*.

steam cracking Heating hydrocarbons with steam in furnace tubes until they break into olefins such as ethylene and propylene. *Chapter 1*.

sticking temperature The temperature at which polymer powder in a fluidized bed softens and begins to stick together into sheets and lumps. *Chapter 8*.

thermocouple A temperature sensor made of two different metals joined at a junction, which gives a small voltage that depends on temperature. *Chapter 1*.

thermowell A closed metal tube that holds a temperature sensor in the process, so it can be replaced without opening the process; it slows the reading. *Chapter 1*.

tray temperature The temperature on one tray of a column, read to infer the composition of the product. *Chapter 4*.

triangle theory A rule that draws the region of flow ratios inside which a simulated moving bed can make both outlets pure. *Chapter 7*.

tube metal temperature (TMT) The temperature of a furnace tube's wall, which rises as coke builds up inside and so ends the run. *Chapter 1*.

wash Stopping a fouled reactor to clean the polymer off its cooled walls, on a fixed schedule or when the cooling margin runs low. *Chapter 11*.

Ziegler-Natta catalyst A catalyst made from a transition-metal compound and an aluminum alkyl that polymerizes olefins at low pressure. *Chapter 9*.

Index of Control Topics

active-set method, 34
actuator saturation, 257
  open loop at the limit, 263
AI agents
  engineering practice, 311
  open problems, 315
  reproduction, 314
  truth audit, 313
  verification record, 312
  while studying, 18
analyzer, 10, 31
  sampled and delayed, 55, 102, 171, 233, 264
analyzer delay, 82
analyzer validity, 10, 79, 308
averaging control, 105
averaging level control, 79, 175

back-calculation, 31, 176, 237, 264
back-off, 11, 33, 86, 226, 238
  design procedure, 309
bias update, 10, 33, 80, 179
bumpless transfer, 16, 176
burner management system, 111

capacity limit, 179
cascade control, 31, 71, 79, 162, 264
  inner-loop speed, 174
catalyst deactivation, 54
constrained variable, 24, 35, 180, 306
  measurement, 315
constraint conflict, 128
constraint control, 194
cycle-to-cycle control, 95, 103, 148
cyclic operation, 10, 95, 259
  simulated moving bed, 139
cyclic steady state, 102, 145

data export, 17
differential temperature, 79
distillation
  history, 73

economic optimum, 308
electrification, 43
evolutionary operation, 127

external reset, 31, 176

feedforward
  from an analyzer, 57
  scheduled disturbance, 237
feedforward control, 308
  model-based, 178
Fenske equation, 73
final control element, 12
finite-time blow-up, 30
floating pressure, 79, 83
flow scheme, 4
fouling, 4, 233
  condition-based cleaning, 311
  monitoring, 251, 276
  reactor washing, 260

giveaway, 81
  in practice, 88
grade transition, 9, 162, 210, 245, 293, 308

heat removal, 8

in-line spectroscopy, 300, 310
infeasibility, 61
inferential control, 10, 57, 71, 164
  history, 74, 165
  tray selection, 78, 82
integrating process, 168
integrator windup, 242
interlocks, 208, 238
inventory averaging, 202
iterative learning control, 98, 140

Kalman filter, 17, 124, 266

laboratory bias update, 266, 292
  holdup, 313
  late results, 295
Langmuir–Hinshelwood kinetics, 52
load cell, 297
loop pairing, 123, 147, 176
loss-in-weight feeding
  refill hold, 291
low selector, 207
  indicator, 212

mid-ranging, 238
minimum reflux, 73

model predictive control, 32, 57, 80, 178, 265
  against constraint control, 308
  economic, 124, 266
  history, 32
  in the apps, 16
  model maintenance, 66
  move blocking, 149, 239, 266, 293
  nonlinear, 309
  offset-free, 266
  plant test, 275
  repetitive, 149
  soft constraint, 33, 208, 238, 266
  step response model, 103
Mooney viscometer, 274
moving horizon estimation, 17, 291

observability, 292
on-stream rate, 5, 36, 263
operating window
  one-sided, 48
override control, 11, 57, 117, 178, 194, 269
  cost, 308
  integral tracking, 207
  measurement, 306
  on the production rate, 123
  soft and hard, 219
  tube metal, 31

parametric sensitivity, 6, 27, 49, 118, 227
pass balancing, 24, 31
  on coil outlet temperature, 30
peak temperature control, 226
petrochemical control
  open problems, 314
PID control
  anti-windup, 16, 79, 123
  velocity form, 16, 30, 78, 123, 148, 176, 237, 264
plug flow
  parcel model, 230
positive feedback, 28, 127
pressure measurement
  remote seal, 217
pressure-compensated temperature, 10, 79, 81
process gain
  sign change, 122
process identification
  slow plant, 78

ratio control, 57, 88, 282
  analyzer trim, 57
real-time optimization, 120
recycle dynamics, 119
relative gain array, 147
run length, 5, 26, 30, 257, 263
  prediction, 42
run-to-run control, 10, 95, 96, 140
runaway, 6, 60, 121, 227, 260
  criteria, 120

safety instrumented system, 61, 250
scenario
  sequence, 81
scenarios, 17, 125, 267
  completion rules, 17
scorecard, 17
selectivity, 6, 29, 48, 117
  history, 50
  open problem, 66
selector control, 57
  indicator, 313
self-inhibition, 53
sequencing, 101
  on timers, 97
severity optimization, 36
simulated moving bed
  optimizing control, 310
simulation engine, 15
simulator framework, 15
soft sensor, 72, 309
  history, 308
  maintenance, 89, 188
  Mooney viscosity, 276
  transferability, 220
steady state
  convergence, 313
step response model, 149
step test, 72, 174
supervisory control
  operator, 188

temperature measurement
  thermowell, 248
time-scale separation, 163, 205

wash scheduling, 265

Index of Applications

absolute pressure transmitter, 109
acetylene converter, 48
    front-end, 50
    model, 51
    tail-end, 50
acetylene hydrogenation, 48
    history, 51
adsorbent
    aging, 151
    new materials, 156
adsorption
    competitive Langmuir isotherm, 142
autoclave reactor, 227

ballast gas, 133
bubble point, 200
bulk polymerization, 196
    history, 197
burner
    degraded, 36
butyl rubber, 257
    history, 259

carbon monoxide moderator, 50
catalyst deactivation, 111, 168
catalyst life, 101
catalyst promoters, 133
catalyst regeneration, 65, 100
cationic polymerization, 259
Catofin process, 97
cavitation, 200
chain-transfer agent, 233
chloride moderator, 120
chromia catalyst, 97
    history, 98
coil outlet temperature, 25
coke
    burn, 100
coke resistance, 28
coking, 4, 26
compounding, 283
condensed-mode cooling, 8, 165
    history, 165
condenser
    cooling water, 77
control valve
    pass flow, 41
cooling margin, 262
cyclic fixed-bed reactor, 95

decoking, 26
    policy, 31
    practice, 41
deposition velocity, 201
desorbent, 142
    quality, 156
    ratio to feed, 146
dilution steam, 26
discharge sock, 286
distillation
    stage model, 74
    superfractionator, 71, 74

ethylene decomposition, 227
    history, 229
    threshold, 233
ethylene furnace, 24
ethylene oxide, 7, 117
    history, 119
ethylene purification, 48
extract, 142

feed drum, 51
feed factor, 284
feed purification, 187
firebox, 27
flammability limit, 120
fluidization, 170
fluidized-bed reactor, 162
fouling, 258
free-radical polymerization, 227

gas chromatograph
    process, 63
gas-phase polymerization, 162
green oil, 50

Houdry process, 97
hydroblasting, 259
hyper-compressor, 229

inert accumulation, 213
initiator injection pump, 249

instrument, 12
intercooler, 51
isobutylene, 259
isoprene, 259

light-off, 60
liquid-full reactor, 202
loop reactor
  circulation pump, 217
loss-in-weight feeding, 282
  history, 284
low-density polyethylene, 226
  history, 229

masterbatch, 287
melt flow rate, 197
melt flow rate test, 187
melt index, 9, 170, 233
membrane separation, 89
methyl chloride, 259
mixing cells, 142
Mooney viscosity, 262
multitubular reactor, 117

naphtha cracking, 42
net positive suction head, 200

off-grade resin, 181
on/off valve, 110
organic peroxide initiator, 231
oxygen service, 132

packed bed
  adiabatic, 53
  cell stack, 99
palladium catalyst, 50
paramagnetic oxygen analyzer, 130
paraxylene, 139
petrochemicals, 1
  history, 2
pinch, 75
polyethylene, 162
polymer additives, 282
polymer fouling, 233
polymer production, 1
polymerization reactor, 8
polyolefin
  history, 3
polypropylene
  history, 196
polypropylene loop reactor, 194
product compressor, 101
product discharge system, 187
product diversion, 293
propane dehydrogenation, 95
propylene splitter, 71
pulse valve, 229

quality laboratory, 171

radiant coil, 27
raffinate, 142
Raman analyzer, 155
reactor jacket, 232
recycle gas cooler, 168
refill, 284
  fill level, 299
reflux ratio, 76
refrigeration
  liquid ethylene, 274
regeneration air, 101
regenerative heat storage, 95
relative volatility, 73
resistance thermometer, 87
reverse-flow reactor, 96
rotary valve, 155

screw feeder, 299
selective hydrogenation, 48
severity, 26
shutdown valve, 64
silver catalyst, 120
simulated moving bed, 139
  history, 141
skin thermocouple, 26
slurry polymerization, 258
slurry reactor, 260
steam cracking, 4, 24
  decarbonization, 310
  electric heating, 310
  history, 3, 26
sticking temperature, 170
surge volume, 101

thermocouple, 233
transfer line exchanger, 42
tray temperature, 76
triangle theory, 145
  history, 142
tube metal temperature, 5
  measurement, 40
tubular reactor
  LDPE, 226

Ziegler-Natta catalyst, 196

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.