INTERACTIVE CONTROL ENGINEERING

Control in Chemical Plants

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 13

Control in Chemical Plants

Recycle, equilibrium and energy in continuous chemical processes

John D. Hedengren
Editor

Advanced Process Solutions, LLC

Interactive Control Engineering: Principles and Applications
Volume 13: Control in Chemical Plants
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/v13-chemical-plants/

How to cite this volume. John D. Hedengren, Ed., *Control in Chemical Plants*, vol. 13 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026.

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

Preface

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

How to read this book with the apps

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

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

How AI agents were used

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

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

About this volume

This volume is about continuous chemical plants shaped by recycle, equilibrium and the cost of energy. Chapters 2 to 5 cover synthesis loops and inorganic chemicals: ammonia, urea, sulfuric acid and chlor-alkali. Chapters 6 to 9 cover separations whose columns are coupled by heat or by reaction. Chapters 10 and 11 cover a specialty product and the optimization of a plant whose model is known to be incomplete.

Each chapter names the technology volume that teaches its main technique. Volume 6 develops the plantwide ideas that recur here.

John D. Hedengren
2026

Contents

Preface v

Notation x

1 Introduction: Recycle, Equilibrium and Energy in Chemical Plants 1
- 1.1 The Question This Volume Answers . . . 1
- 1.2 A Short History . . . 3
- 1.3 Notation and Conventions . . . 4
- 1.4 Recycle, Equilibrium and Energy on One Site . . . 4
- 1.5 Instruments and Final Elements in This Volume . . . 13
- 1.6 The Simulator Framework . . . 14
- 1.7 Working With AI Agents While Studying . . . 18
- 1.8 Chapter Roadmap . . . 19
- 1.9 Exercises . . . 21

2 Ammonia Synthesis Loop Simulator for Recycle Amplification and Inventory Pairing 24
- 2.1 Introduction . . . 25
- 2.2 The Synthesis Loop and How It Came To Be Controlled . . . 26
- 2.3 Process Model . . . 28
- 2.4 What the Recycle Does . . . 34
- 2.5 Controller Design . . . 36
- 2.6 Results . . . 38
- 2.7 Technology in Practice . . . 45
- 2.8 Limitations . . . 47
- 2.9 Discussion . . . 48
- 2.10 Conclusion . . . 48
- 2.11 Exercises . . . 50

3 A Urea Synthesis Loop Whose Controlling Ratio Is Never Measured 53
- 3.1 Introduction . . . 54
- 3.2 From Carbamate to a Total-Recycle Loop . . . 55
- 3.3 The Loop as Modeled . . . 56
- 3.4 What the Loop Does Before It Is Controlled . . . 62
- 3.5 Control as Implemented . . . 63
- 3.6 Nine Cases on the Running Loop . . . 65
- 3.7 Technology in Practice . . . 69
- 3.8 Limitations . . . 70
- 3.9 Conclusion . . . 71
- 3.10 Exercises . . . 73

4 Four Catalyst Beds, an Ignition Floor and Double Absorption in a Contact Plant 76
- 4.1 Introduction . . . 77
- 4.2 Contact Process, Melt Catalyst and the Stack Limit . . . 78
- 4.3 The Plant as Modeled . . . 79
- 4.4 Lines, a Floor and Fronts . . . 84
- 4.5 Seven Loops and a Predictive Layer . . . 86
- 4.6 Nine Cases on the Running Plant . . . 88
- 4.7 Technology in Practice . . . 91

4.8 Limitations . . . 94
4.9 Conclusion . . . 94
4.10 Exercises . . . 96

5 Chlor-Alkali Membrane Electrolyzer Simulator Fed in Ratio to Its Current 99
5.1 Introduction . . . 100
5.2 Mercury, Asbestos and the Ion-Exchange Membrane . . . 101
5.3 The Electrolyzer as Modeled . . . 102
5.4 Instruments and Implemented Controllers . . . 108
5.5 Riding Through a Compressor Trip . . . 110
5.6 Measured Cases . . . 112
5.7 Technology in Practice . . . 116
5.8 Limitations . . . 119
5.9 Conclusion . . . 119
5.10 Exercises . . . 121

6 Air Separation Simulator for Thermally Coupled Columns and Interacting Purity Loops 124
6.1 Introduction . . . 125
6.2 Background . . . 126
6.3 Air Separation Unit Model . . . 127
6.4 Coupling and Interaction Analysis . . . 136
6.5 Implemented Control Structure . . . 138
6.6 Results . . . 141
6.7 Technology in Practice . . . 148
6.8 Limitations . . . 150
6.9 Discussion . . . 150
6.10 Conclusion . . . 151
6.11 Exercises . . . 153

7 Amine Capture Plant Simulator for Recycle Gain and the Scheduling of Regeneration Energy 156
7.1 Introduction . . . 157
7.2 Amine Scrubbing and the Loops Added to It . . . 158
7.3 Process Model . . . 159
7.4 What the Recycle Does to the Gain . . . 166
7.5 Controller Design . . . 168
7.6 Results . . . 172
7.7 Technology in Practice . . . 178
7.8 Limitations . . . 180
7.9 Discussion . . . 180
7.10 Conclusion . . . 181
7.11 Exercises . . . 183

8 A Dividing-Wall Column Whose Spare Handle Moves Only the Energy 187
8.1 Introduction . . . 188
8.2 Two Columns in One Shell . . . 189
8.3 The Shell as Modeled . . . 190
8.4 The Spare Handle and Its Valley . . . 193
8.5 Loops as Implemented . . . 197
8.6 Nine Cases, Measured . . . 198
8.7 Technology in Practice . . . 202
8.8 Limitations . . . 204

8.9 Conclusion . . . 204
8.10 Exercises . . . 206

9 MTBE Reactive Distillation With Multiple Steady States and a Branch-Dependent Temperature Loop 210
9.1 Introduction . . . 211
9.2 Etherification on a Distillation Tray . . . 212
9.3 The Column as Modeled . . . 214
9.4 The Branches and Their Stability . . . 218
9.5 Loops as Implemented . . . 221
9.6 Nine Cases, Measured . . . 222
9.7 Technology in Practice . . . 226
9.8 Limitations . . . 228
9.9 Conclusion . . . 228
9.10 Exercises . . . 230

10 A Synthetic Base Oil Simulator for Mid-Ranging, Inferential Control and Constraint Back-Off 233
10.1 Introduction . . . 234
10.2 Background . . . 235
10.3 Process Model . . . 236
10.4 Controller Design . . . 238
10.5 Results . . . 241
10.6 Technology in Practice . . . 247
10.7 Limitations . . . 249
10.8 Discussion . . . 249
10.9 Conclusion . . . 250
10.10 Exercises . . . 252

11 Real-Time Optimization of a Williams–Otto Reactor Whose Model Lacks a Reaction 255
11.1 Introduction . . . 256
11.2 From a Benchmark Plant to Modifier Adaptation . . . 257
11.3 The Reactor as Modeled . . . 258
11.4 The Optimizer's Model and the Mismatch . . . 261
11.5 Loops and Optimization Layer as Implemented . . . 262
11.6 Ten Cases, Measured . . . 266
11.7 Technology in Practice . . . 274
11.8 Limitations . . . 276
11.9 Conclusion . . . 276
11.10 Exercises . . . 278

12 Summary and Outlook 282
12.1 What the Chapters Showed . . . 282
12.2 Choosing a Method . . . 285
12.3 Where Practice Is Heading . . . 286
12.4 Engineering With AI Agents . . . 288
12.5 Gaps in Chemical Plant Control . . . 291
12.6 Exercises . . . 295

Answers to Selected Exercises 296

Glossary 300

Index of Control Topics **303**

Index of Applications **306**

About the Editor **308**

Volumes in This Series **309**

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: Recycle, Equilibrium and Energy in Chemical Plants

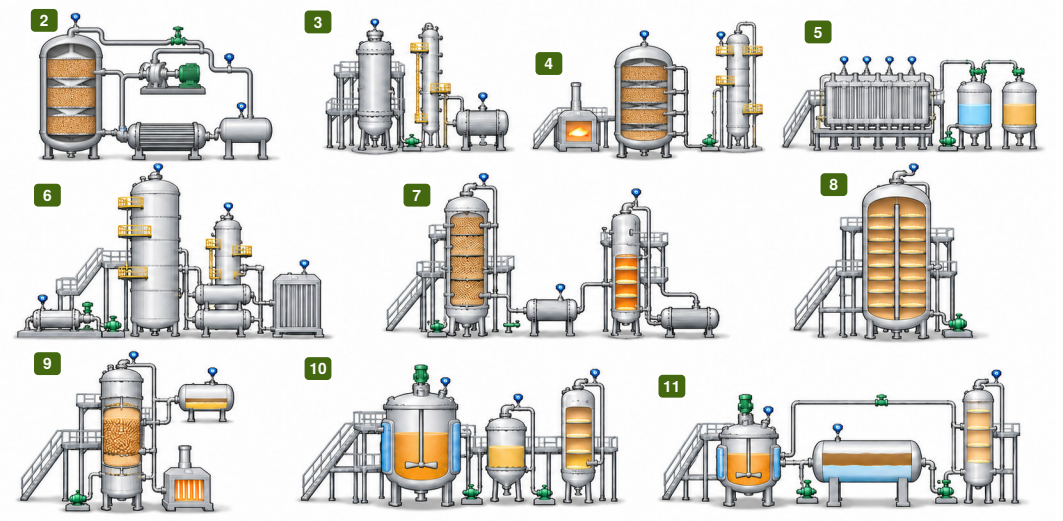

The chapters 2 ammonia loop 3 urea loop 4 contact acid plant 5 membrane electrolyzer 6 air separation 7 amine capture 8 dividing-wall column 9 reactive column 10 base oil plant 11 reactor and optimizer

Notice that the top row holds the reactors and cells of the heavy chemicals, the middle row three separations, and the bottom row a reactive column and two reactors, each followed by its separation. Illustration, not to scale.

IN PLAIN TERMS **Plants that send most of their material around again**

The main idea. A chemical plant turns air, natural gas, salt water and sulfur into goods such as fertilizer, acid and chlorine. In a synthesis loop, most of what enters the reactor has been there before, because the part that did not react is sent back around.

The challenge. Many reactions stop partway, at a balance that heat and pressure can shift. Sending material around again makes the plant slow and sensitive. What a stream holds is often known only from a late laboratory test.

Why it matters. Good control decides how much product each unit of energy buys and keeps the plant inside its limits.

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, 1.4.3 and 1.4.5, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, working each number of the example S-1, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.8 with Tables 1.1 and 1.2, Section 1.5 with Figures 1.7 and 1.10, and Section 1.8.

1.1 The Question This Volume Answers

A chemical plant turns a few raw materials into a product by reaction and separation, and it rarely does so in one pass. Unconverted reactants return through a recycle, a solvent circulates between an absorber and a stripper, and the heat of one column boils another. The reactions stop at an equilibrium that temperature and pressure move, and the energy that drives the separations is bought at a price

that changes by the hour. This volume answers one question: how do recycle, equilibrium and energy cost shape the control structure of a chemical plant, and which techniques recur from one plant to the next?

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

- *Recycle sets the gains and the time constants.* A loop that converts a fifth of its feed per pass, or a solvent that returns after an hour, answers a handle with a gain and a speed that no single vessel has.
- *Equilibrium caps the conversion.* A reversible exothermic reaction runs faster when hot and stops sooner, so temperature acts on conversion in both directions.
- *Some compositions are inventories.* A reactant fed in slight excess to a loop it cannot leave accumulates, and its ratio integrates every feed error.
- *Energy is shared or priced.* One exchanger may be the reboiler of one column and the condenser of another, a spare handle may move only the energy, and the power may cost more in one hour than the next.
- *Composition is inferred.* Loop ratios, solvent loadings and product purities have no fast transmitter, and they are carried between laboratory results by densities, temperatures and models.
- *Feeds follow the throughput.* Stoichiometry or Faraday's law fixes how much of each feed a given rate needs, so the feeds are set in ratio to the throughput and trimmed.
- *The steady state is not unique, and the optimum is known only through a model.* A converter can go out, a column can sit on more than one branch, and an optimizer moves the plant along slopes its model may have wrong.

Ten chapters answer the question on ten plants, each with an app that the reader runs in a browser. Chapters 2 and 3 are the two synthesis loops of a fertilizer site, ammonia and urea. Chapters 4 and 5 are a contact sulfuric acid plant and a membrane electrolyzer. Chapters 6 and 7 separate air and capture carbon dioxide from a flue gas. Chapters 8 and 9 are a dividing-wall column and a reactive column. Chapter 10 makes a synthetic base oil, and Chapter 11 closes the volume with an optimizer above a reactor whose model is wrong.

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

- Find where an adiabatic bed stops on the conversion and temperature diagram, and why its inlet temperature has an optimum.
- Predict the gain and the time constant that a recycle or a purge gives a composition, and recognize a composition that integrates.
- Say what a lean loading loop does to the gain of an absorber, and what lowering the loading costs in energy.
- Count the handles of a column, find the one that moves only the energy, and estimate the energy a thermally coupled arrangement saves.
- Set feeds in ratio to the throughput with a trim, and correct an inferred composition from a late laboratory result.
- Explain why a refitted model can stop an optimizer short of the plant's best point, and how a gradient correction removes the gap.

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, overrides and mid-ranging from Volume 4, dead time and estimation from Volume 5, the relative gain array and plantwide inventory control from Volume 6, multiplicity and nonlinear analysis from Volume 7 and predictive control and optimization from Volume 9. The refinery and the petrochemical complex that supply several of these plants are Volume 11 and Volume 12.

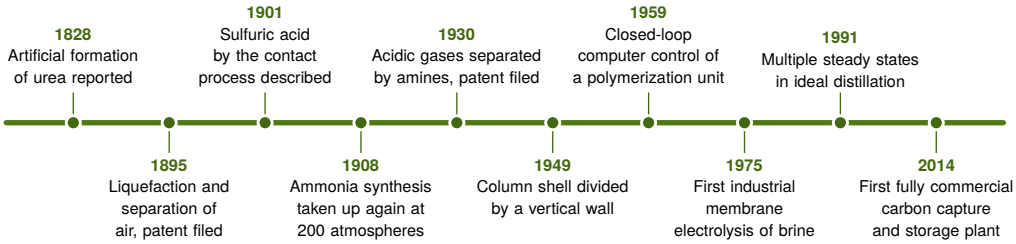

Figure 1.1 Milestones of the chemical plants of this volume and their control. The dates are from [1] (1828), [2] (1895), [3] (1901), [4] (1908), [5] (1930), [6] (1949), [7] (1959), [8] (1975), [9] (1991) and [10] (2014). A patent is dated by its filing where the card says so. Notice that most processes of the volume are older than the computers that now run them. Events are evenly spaced, not to a time scale.

1.2 A Short History

The milestones of Figure 1.1 follow the plants of this volume, from the first heavy chemicals to the separations, the columns and the computers that came to run them.

> ⧖ HISTORY **1905 to 1932: The First Synthesis Loops**
>
> In 1905 Fritz Haber judged ammonia from its elements not worth pursuing, because from the start of red heat no catalyst gave more than a trace of it at normal pressure. He took the problem up again in 1908 at 200 atmospheres, the limit of the compressor he had. His process made and removed the ammonia in a cycle at a constant high pressure, and it led the heat of the exhaust gas to the fresh gas, so that the process supplied its own heat. By 1920 industry had kept the pressure, the circulation and the heat exchange [4]. Urea followed. A mixture of carbon dioxide and ammonia heated in a closed vessel stops at about 40 percent of the theoretical yield, so the excess gases were drawn off at high pressure by heated compressors and condensed again [11]. Chapter 2 and Chapter 3 run both loops.

Air separation began with a patent filed in 1895. Air let down through a valve from a high pressure comes out colder, and the cold air was led back through a long countercurrent coil, about one hundred meters of pipe, to cool the compressed air arriving until it liquefied [2]. A second patent, filed in 1903, rectified the liquid in a column. Compressed gas condensed in a coil in the liquid chamber at the column's foot and boiled that liquid, until pure nitrogen left the top and pure liquid oxygen left the chamber [12]. A coil that condenses one stream and boils another is the shared exchanger of Chapter 6. In 1930 a patent was filed for separating acidic gases with a small group of amines that unite with them when cool and give them up when hot, the solvent circulating in a closed cycle with heat interchangers [5]. That is the solvent loop of Section 1.4.4 and of Chapter 7.

> ⧖ HISTORY **1851 to 1986: Sulfuric Acid and Chlorine**
>
> In 1901 a paper of 47 pages described sulfuric acid and its manufacture by the contact process [3], and Haber later likened the difficulties of scaling up ammonia to those overcome in it [4]. Of the catalysts tested for oxidizing sulfur dioxide only vanadium compounds, platinum and iron oxide proved satisfactory, and vanadium pentoxide on silica, a liquid melt at operating conditions, is now used almost exclusively [13]. Chlorine took longer. An English patent of 1851 made chlorine and caustic soda by electrolysis of brine, and from 1890 the electrodes were separated by a porous diaphragm or replaced by a mercury cathode, unchanged for decades. Stable membranes came at the end of the 1960s, the first industrial membrane plant ran in 1975, and in 1986 Japan converted all its mercury cells [8]. Chapter 4 and Chapter 5 model the plants that came of both.

The columns came next. A column shell divided by a vertical wall through a set of trays was proposed in 1949 for a three-product split, and it is usually named after a theoretical study of 1965 [6]. Reaction and separation in one column have been developed since the early 1920s [14], and in 1991 multiple steady states were shown in a distillation column with ideal thermodynamics [9]. Chapter 8 and Chapter 9 build on both. Computers arrived in 1959, when one achieved closed-loop control of an industrial polymerization unit [7]. The first process computers were supervisory. They adjusted the setpoints of the controllers below them, so the plant kept working when the computer was off [15]. The optimizer of Chapter 11 sits in the same place. The newest milestone is about energy and carbon. In October 2014 the world's first fully commercial carbon capture and storage plant opened at a coal-fired power station [10].

1.3 Notation and Conventions

Temperatures are in °C, differences of temperature in kelvin, and absolute temperatures in a rate or an equilibrium constant in kelvin. Molar flows in the running example are in kmol/h and pressures in bar, written p. A conversion is X, an equilibrium conversion X^*, a gas mole fraction y and a solvent loading x. In the synthesis loop F_M is the makeup, F_P the purge, F the gas fed to the converter and N the gas held in the loop. A recycle is described by a forward gain g_F and the gain g_R of the path that returns. The absorption factor is $\mathscr{A}$ and the stripping factor $\mathscr{S}$. The framework's h is its integration step and its a the controller action, so neither is used for anything else here.

The chapters keep their own symbols, and several collide. λ is a latent heat in Chapter 6 and Chapter 8, an adiabatic rise in Chapter 2, a bias update gain in Chapter 3 and Chapter 10, and a gradient modifier in Chapter 11. The loading of Chapter 7 is α, which Chapter 8 and Chapter 9 use for relative volatility and Chapter 10 for a chain growth probability. r is the hydrogen to nitrogen ratio in Chapter 2, while Chapter 3 writes its ammonia to carbon dioxide ratio as N/C. Chapter 8 splits its liquid by β and its vapor by ρ, and K is an equilibrium constant in some chapters and a filter gain in Chapter 11.

1.4 Recycle, Equilibrium and Energy on One Site

1.4.1 The Plants of This Volume

Figure 1.2 places each chapter's plant among the others. Natural gas is reformed and shifted into a synthesis gas whose hydrogen and nitrogen an ammonia loop converts [16]. Nitrogen can also come from an air separation unit. Ammonia and carbon dioxide make urea, and an amine plant can recover the carbon dioxide from the reformer's flue gas. Sulfur burned in air becomes sulfuric acid, and brine electrolyzed with power becomes chlorine, caustic soda and hydrogen. Among the organic plants, a reactive column makes an ether from a refinery C4 cut and methanol, a dividing-wall column splits aromatics into three products, and a reactor and column make a base oil from decene. The last block is a reactor with an optimizer above it, the benchmark on which Chapter 11 tests real-time optimization.

The chapters take the plants one at a time, and the techniques they share are developed below on one running example, a small synthesis site called S-1, drawn in Figure 1.3. A solvent unit removes an impurity I from a raw gas, and the cleaned gas is the makeup of a loop that makes C from A and B in a two-bed converter. A separator condenses C, a purge removes some of the loop gas, and the rest returns. A column splits the crude C from a light and a heavy byproduct. Every number of S-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 Equilibrium-Limited Reactors

S-1's reaction A + B $\rightleftharpoons$ C runs in the gas at $p = 100\,\text{bar}$ and releases heat. Its equilibrium constant falls as the temperature rises, $K(T) = \exp[B_K(1/T - 1/T_h)]$ in bar^{-1}, with $B_K = 10\,000\,\text{K}$ and $K = 1\,\text{bar}^{-1}$ at $T_h = 525\,\text{K}$. For a feed of equal A and B with an inert fraction y_I, the equilibrium at a conversion

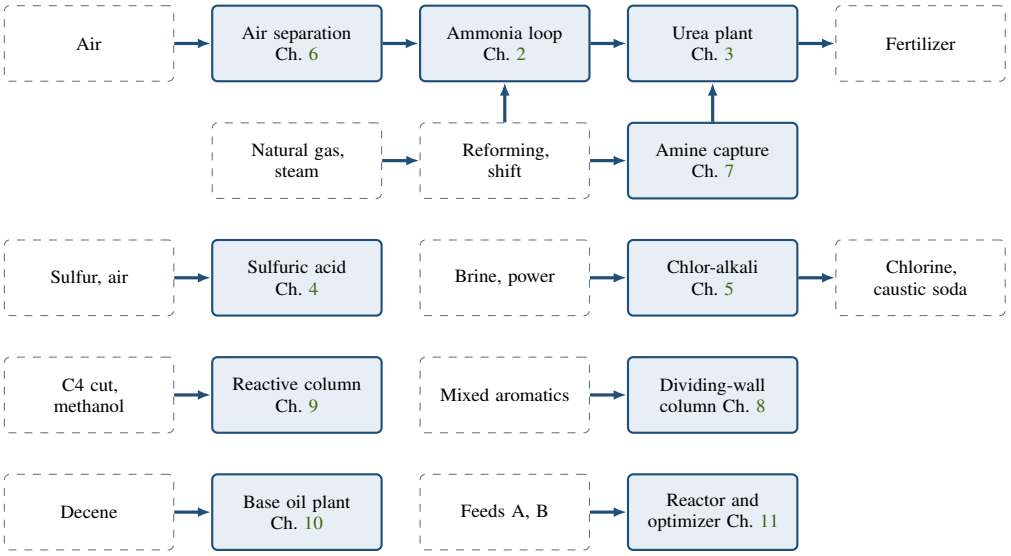

Figure 1.2 The Plants of the Volume. Filled Blocks Are the Plants of Chapters 2 to 11, Dashed Blocks Lie Outside the Volume. The Arrows Show Typical Connections: Nitrogen From Air Separation and Synthesis Gas From the Reformer Feed the Ammonia Loop, and Ammonia With Carbon Dioxide Captured From the Reformer's Flue Gas Feeds the Urea Plant. A Site Need Not Have Every Unit.

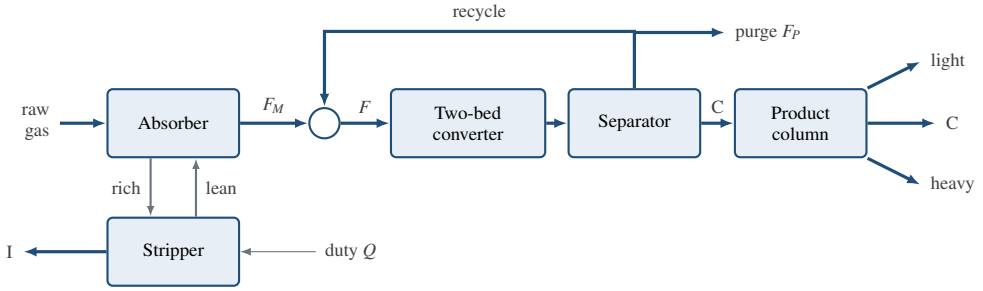

Figure 1.3 The Running Example S-1. A Four-Stage Absorber Removes the Impurity I, Which Is Inert in the Loop, With a Solvent That a Four-Stage Stripper Regenerates. The Loop Converts A + B to C in Two Adiabatic Beds With an Interbed Cooler, Condenses C in the Separator, Purges F_P and Recycles the Rest.

X of A is Equation (1.1).

$$K(T)\,p = \frac{y_C}{y_A y_B} = \frac{X(1-aX)}{a(1-X)^2}, \qquad a = \frac{1-y_I}{2} \tag{1.1}$$

It is a quadratic in X, whose root below one is the equilibrium conversion of Equation (1.2).

$$X^* = \frac{1+2\kappa-\sqrt{(1+2\kappa)^2-4\kappa(a+\kappa)}}{2(a+\kappa)}, \qquad \kappa = K(T)\,p\,a \tag{1.2}$$

At 400 °C with 10 % inert, $K = 0.0151\,\text{bar}^{-1}$, $\kappa = 0.680$ and $X^* = 0.345$. A higher pressure raises κ, and more inert lowers it. Chapter 2 solves the same kind of quadratic in closed form for the ammonia equilibrium.

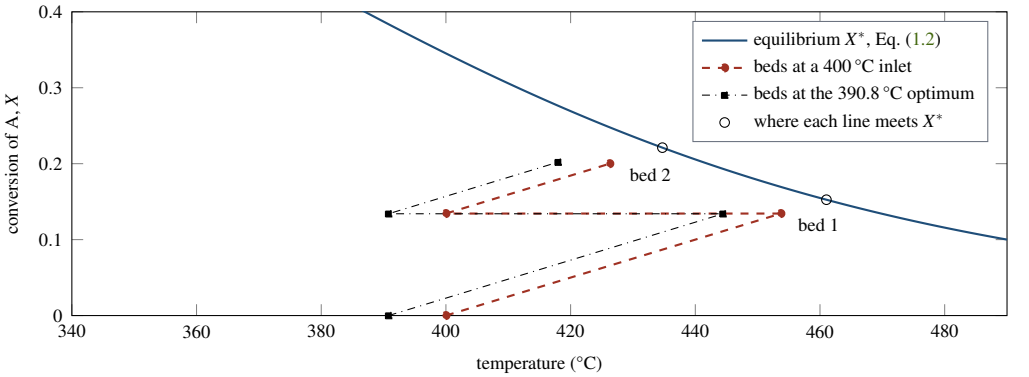

Figure 1.4 S-1's Converter on the Conversion and Temperature Diagram. The Equilibrium Curve (Solid) Falls With Temperature. Each Bed Climbs an Adiabatic Line Toward It and Stops Short (Dashed, Dots), the Cooler Returns the Gas to the Inlet Temperature, and Open Circles Mark Where Each Line Meets the Curve. At the Optimum Inlet (Dash-Dotted, Squares) Both Beds Start Cooler and End Higher.

An adiabatic bed cannot hold its temperature. Each unit of conversion heats the gas by the adiabatic rise ΔT_{ad}, 400 K in S-1, so a bed moves along the straight line of Equation (1.3) and stops where the line meets the equilibrium curve, or short of it.

$$T = T_{in} + \Delta T_{ad}\,(X - X_{in}), \qquad X_{out} = X_{in} + \eta\,(X^*_{line} - X_{in}), \qquad \eta = 1 - \mathrm{e}^{-Da(T_m)} \tag{1.3}$$

where X^*_{line} is the intersection, η the fraction of the way the catalyst gets there, and Da a Damköhler number that grows with the mean bed temperature T_m with an activation temperature of 15 000 K, 1.9 at 700 K. Figure 1.4 draws S-1's converter on this diagram [17]. Both beds enter at 400 °C. Bed 1 would stop on the curve at $X^* = 0.153$ and 461 °C, gets 88 % of the way and leaves at $X = 0.134$ and 454 °C. The interbed cooler returns the gas to 400 °C, and bed 2 climbs a parallel line to $X = 0.200$. One bed alone would give 0.134, and two beds without the cooler 0.152, because the second bed would start next to the curve.

The inlet temperature has an optimum because temperature acts both ways. A hotter inlet speeds the catalyst and lowers the curve it is heading for. S-1 converts 0.193 per pass at 410 °C, 0.170 at 430 °C, and at most 0.2020 at 390.8 °C, 1.0 % more than at the design inlet. Below the optimum the catalyst is too slow, above it the equilibrium too low. Pressure moves the whole picture: at 90 bar the converter gives 0.189 and at 110 bar 0.210. So does the inert. At 12 % inert the per-pass conversion falls to 0.1975 and at 20 % to 0.187, which is how the purge of the next section reaches the reactor.

Chapter 2 runs three quench beds on the Haber equilibrium. Its first bed reaches 82.8 % of the way to equilibrium, and its converter has an optimum near 480 °C of hot spot, worth 11 % of production, against an operating 505 °C. Chapter 4 draws four beds live on this diagram, at 3.20 K per point of conversion, and its first bed stops on the curve at 65.77 %. Removing the product between beds starts the next bed on a new curve. S-1 would reach 0.25 instead of 0.20 if C were taken out after bed 1, and the interpass absorber of Chapter 4 raises its conversion from 97.910 to 99.837 % the same way. The urea reactor of Chapter 3 approaches an equilibrium conversion that rises with its N/C ratio and peaks in temperature, and Chapter 9 removes its ether by distillation as it forms.

> ✦ DID YOU KNOW? **Most Sulfur Becomes Acid**
>
> The world produced an estimated 85 million metric tons of sulfur in 2024, most of it as a result of processing fossil fuels. In the United States about 90 percent of the sulfur consumed that year was in the form of sulfuric acid [18].

A bed that heats its own feed through an exchanger adds a feedback the diagram does not show. The ignited state can then be lost, as in Chapter 2, where a quench of 15.3 % keeps the converter lit and 15.4 % puts it out [19]. Chapter 4 finds an ignition floor near a 395 °C inlet. The loop gain of such a heat exchanger is derived in Volume 11.

1.4.3 Synthesis Loops and Recycle

S-1's loop takes $F_M = 1000\,\mathrm{kmol/h}$ of makeup carrying 1 % of I and purges $F_P = 100\,\mathrm{kmol/h}$. Each mole of C uses one mole each of A and B, so with the purge fixed the loop makes $(F_M - F_P)/2 = 450\,\mathrm{kmol/h}$ of C. At a per-pass conversion of 0.20 the converter needs 2250 kmol/h each of A and B. The separator returns 4000 kmol/h of gas, of which 3900 kmol/h is recycled, 3.9 times the makeup. The loop holds $N = 500\,\mathrm{kmol}$ of gas, seven and a half minutes of the separator's flow. The converter of Section 1.4.2 was computed with the loop gas's inert, although the makeup dilutes it to 8.2 % at the converter inlet, a difference the example leaves out.

An inert can leave only through the purge. With the loop well mixed at constant inventory, its balance is Equation (1.4).

$$N\frac{\mathrm{d}y_I}{\mathrm{d}t} = F_M z_I - F_P y_I, \qquad y_I^{\infty} = \frac{F_M}{F_P} z_I, \qquad \tau_I = \frac{N}{F_P} \tag{1.4}$$

where z_I is the inert in the makeup. Both results are ratios of a large number to a small one. S-1's loop holds 10 % inert from a 1 % makeup, and its time constant is 5 h, forty passes. Chapter 2 measures the same structure: a makeup with 1 % inert and a purge of 8.4 % of the makeup hold 11.90 % in the loop with a time constant of 5.1 h. If the makeup carries 1.2 % inert, S-1 either runs at 12 % and loses 1.3 % of its per-pass conversion, or purges 120 kmol/h to stay at 10 % and makes 2.2 % less C.

> ✦ DID YOU KNOW? **A Fertilizer Made From the Air**
>
> The world's ammonia plants produced an estimated 150 million metric tons of nitrogen in 2024. In the United States natural gas was the main feedstock, and about 88 percent of the ammonia made went to fertilizer [18].

The same result holds for any recycle. If a variable answers a handle with a gain g_F in one pass, and a fraction g_R of the change returns to the inlet, the loop settles at the gain and the time constant of Equation (1.5) [20].

$$K_{ss} = \frac{g_F}{1 - g_R}, \qquad \tau_{ss} = \frac{\tau}{1 - g_R} \tag{1.5}$$

where τ is the time constant of one pass, for a first-order forward path and a return path that is a pure gain. For S-1's inert, the separator gas is the pass, $g_F = F_M/4000 = 0.25$ and $g_R = 1 - F_P/4000 = 0.975$, so both the gain and the time constant are multiplied by forty, to 10 and 5 h. Three regimes recur in the volume. A return gain near one makes a variable large and slow, the snowball effect of reactor and separator plants [21]. A return gain of exactly one, with no exit at all, turns the variable into an integrator. A negative return gain, a recycle that undoes part of the change, makes the settled gain smaller than the gain seen before the recycle returns.

A reactant fed in excess is the integrator case. Suppose the front end feeds B by a meter that reads 1 % high, so the makeup brings 495 kmol/h of A and 490.05 kmol/h of B. The reaction uses them equally, and the excess of A in the loop gas, $\Delta = y_A - y_B$, obeys Equation (1.6).

$$N\frac{\mathrm{d}\Delta}{\mathrm{d}t} = (F_{M,A} - F_{M,B}) - F_P\,\Delta, \qquad \Delta^{\infty} = \frac{F_{M,A} - F_{M,B}}{F_P} \tag{1.6}$$

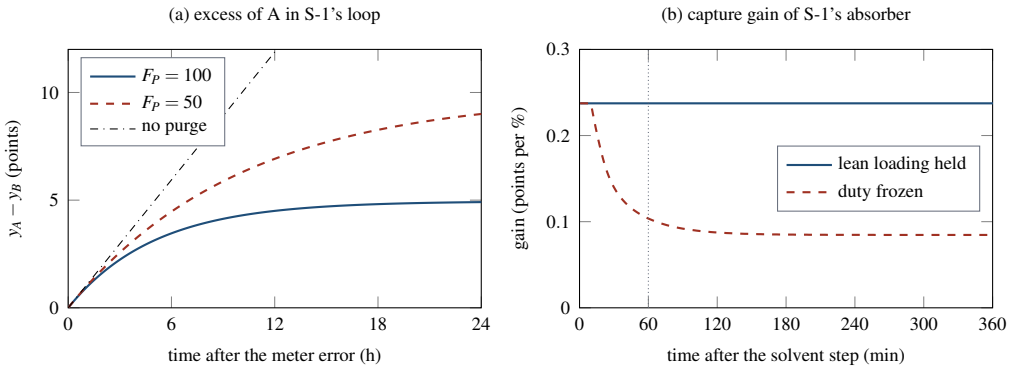

Figure 1.5 Two Recycles of S-1. (a) The Excess of A in the Loop After the B Meter Starts to Read 1 % High, With the Design Purge (Solid), Half the Purge (Dashed) and No Purge (Dash-Dotted). (b) The Apparent Gain of the Absorber's Capture to a 10 % Solvent Step, With the Stripper Duty Frozen (Dashed) and With the Lean Loading Held (Solid). The Dotted Line Marks One Lap of the Solvent.

The excess settles at 4.95 points with the same 5 h time constant as the inert, as shown in Figure 1.5a. The loop gas then holds A and B in the ratio 1.116 against the makeup's 1.010, an error 11.5 times as large. Halving the purge doubles both the error and its time constant. With no purge the excess grows by 0.99 points an hour and never stops. Chapter 2 measures an amplification of 8.0 for a 5 % error in the hydrogen to nitrogen ratio and finds that it is not symmetric, because the real reaction responds to the excess. In Chapter 3 the excess ammonia has no exit, and the N/C ratio integrates at 0.0180 an hour for a fresh ratio 0.02 too high. A ratio held between two measured feeds is not a ratio held in the loop.

The makeup is the only flow that can set the loop's throughput, since every mole of C made and every mole purged entered through it. The design decision is which variable the makeup valve holds [22]. If it holds the loop pressure, the inventory is regulated and a disturbance appears as a change in production. If it holds production, the pressure carries the disturbance, and an override must take the valve back before the pressure reaches its limit. Both arrangements regulate the inventory with a flow that enters it [23]. Chapter 2 offers both. A loss of catalyst activity there costs 17 t/d in one arrangement and raises the pressure to 158.3 bar, 6.7 bar inside the limit of its low selector, in the other.

1.4.4 Absorption and Stripping

S-1's absorber washes I out of 1000 kmol/h of raw gas holding 5 % of it. For a dilute gas with a straight equilibrium line $y = mx$, the outlet of an n-stage counter-current absorber follows the Kremser relation of Equation (1.7) [24].

$$y_{out} = mx_L + (y_{in} - mx_L)\,\varphi(\mathscr{A}), \qquad \varphi(\mathscr{A}) = \frac{\mathscr{A} - 1}{\mathscr{A}^{n+1} - 1}, \qquad \mathscr{A} = \frac{L}{mG} \tag{1.7}$$

where x_L is the lean loading of the solvent entering at the top, L the solvent flow and G the gas flow. The stripper obeys the same relation with the roles of the phases exchanged. It leaves the fraction $\varphi(\mathscr{S})$ of the rich loading in the lean solvent, where the stripping factor $\mathscr{S}$ grows with the vapor the reboiler duty raises. At the design, $\mathscr{A} = \mathscr{S} = 1.25$ with four stages each, so $\varphi = 0.122$, the solvent flow is 1250 kmol/h, the lean loading 0.0044 and the rich loading 0.0365, and the makeup leaves with 1.0 % of I. The absorber removes 80 % of it.

The solvent is a recycle. A larger solvent flow raises $\mathscr{A}$ and captures more at once, but with the duty fixed the same vapor must strip more solvent, so $\mathscr{S}$ falls and the lean loading rises a lap later. Figure 1.5b shows a 10 % solvent step on S-1, whose lap is 60 min. For ten minutes the capture rises

0.237 points per percent of solvent, and with the duty frozen the gain decays to 0.085 as the leaner solvent stops arriving. By Equation (1.5) the return gain is $g_R = 1 - 0.237/0.085 = -1.8$. A lean loading loop that moves the duty to hold x_L cuts the return path, $g_R = 0$, and the gain stays at 0.237. It asks for 7.4 % more vapor, 4.3 % more for each mole of I removed.

Chapter 7 measures this behavior on an amine plant whose lap is 70 min [25]. With the reboiler duty frozen its capture gain falls from 0.177 at ten minutes to 0.040 after eight hours, $g_R = -3.4$, and with the lean loading loop in service it stays at 0.177. A capture loop tuned on the early gain and run with the duty frozen is therefore 4.4 times too weak, and one tuned on the settled gain is 4.4 times too strong once the inner loop is closed. The inner loop is there because it makes one gain out of two.

A leaner solvent costs more than its share of energy. Lowering S-1's lean loading in four equal steps of a fifth of the design value takes 140, 188, 280 and 533 kmol/h more stripping vapor, and the last step buys the same 1.6 points of capture as the first. Chapter 7 pays 38, 43, 51, 65, 90 and 137 MW for six equal steps of loading, because the stripper boils more water to carry less carbon dioxide. The Kremser relation also solves the top of the high-pressure column in Chapter 6, where the heavy key is a few tens of parts per million.

1.4.5 Heat-Integrated and Reactive Columns

A column with a condenser and a reboiler has four flows to move: reflux, distillate, boilup and bottoms. Two hold the drum and sump levels, and two remain for two product purities. A column with a side draw and a liquid split, such as a dividing-wall column, has six handles for three purities and two levels, so one handle is spare. Once the purities are held, the material balances fix the products, and the spare handle can only redistribute the internal flows. It moves the energy and nothing else, and the energy it moves has a minimum.

The minimum is bounded by Underwood's method [26]. For a saturated liquid feed with relative volatilities α_j and mole fractions z_j, the roots θ of Equation (1.8) give the minimum vapor of each sharp split.

$$\sum_j \frac{\alpha_j z_j}{\alpha_j - \theta} = 0, \qquad V_{min} = F \sum_{j \in D} \frac{\alpha_j z_j}{\alpha_j - \theta} \tag{1.8}$$

where the second sum runs over the components in the distillate and θ is the root between the two key volatilities. S-1's crude product is 100 kmol/h of a light, C and a heavy in equal parts with $\alpha = 4$, 2 and 1. The roots are 2.756 and 1.244, the light split needs 107.2 kmol/h of vapor and the heavy split 136.6 kmol/h. Two columns in the direct sequence need $107.2 + 100.0 = 207.2\,\mathrm{kmol/h}$. A thermally coupled arrangement in one shell needs only the larger of the two minima [27], 136.6 kmol/h, a saving of 34.1 %. Chapter 8 finds 30.2 % for its benzene, toluene and o-xylene feed with the same bound. With finite stages its boilup reaches a minimum of 5.83 MW at one liquid split and costs 17.7 % or 21.4 % more at splits a little above and below, and the minimum moves with the feed.

Thermal coupling can also join two columns through one exchanger. In Chapter 6 the nitrogen condensing at the top of the high-pressure column boils the oxygen in the sump of the low-pressure column, across 1.894 K. The driving force is fixed by the duty, so a pressure change in one column must be matched in the other. With the Clausius–Clapeyron slope $\mathrm{d}T/\mathrm{d}p = RT^2/(\lambda_v p)$ on each side, where λ_v is the latent heat, the two pressures move as Equation (1.9) states.

$$\frac{\mathrm{d}p_{HP}}{\mathrm{d}p_{LP}} = \frac{T_{reb}^2\, \lambda_{N_2}\, p_{HP}}{T_{cond}^2\, \lambda_{O_2}\, p_{LP}} \tag{1.9}$$

With the chapter's temperatures and latent heats, 95.36 K and 5000 J/mol at 5.6 bar against 93.47 K and 6700 J/mol at 1.4 bar, the ratio is 2.87. The chapter measures 3.11 bar for each bar of low-pressure column with the bubble points of the real compositions, so pinching a valve on the low-pressure column loads a compressor that feeds the other one.

> **DID YOU KNOW? The Electricity Behind a Tonne of Oxygen**
>
> A typical cryogenic air separation unit uses more than 200 kilowatt-hours of electricity per tonne of oxygen, against a minimum thermodynamic work of around 51 [28]. In Chapter 6 that electricity drives the main air compressor, whose power the controllers cap.

A reactive column removes its product as it forms, which pushes an equilibrium-limited reaction past the limit a reactor would reach. The coupling of reaction and separation can also give one set of handles more than one steady state [9, 29]. Chapter 9 finds three stable states at a duty of 42.5 MW, converting 98.22, 90.05 and 94.89 %. Its stripping-tray temperature peaks at 143.36 °C at 43.15 MW, so 143.13 °C is read at both 43.02 and 43.4 MW. That is an input multiplicity [30]. The process gain from duty to the tray is positive below the peak and negative above it, and a loop whose action was chosen for one side drives the duty the wrong way on the other. The chapter shows such a loop raising the duty without limit on the low-conversion branch.

1.4.6 Ratio Control and Inferred Compositions

When stoichiometry fixes the feeds, the throughput sets each of them. A ratio station multiplies a measured rate by a ratio and adds the output of a trim that reads the composition, as in Equation (1.10).

$$F_{B,SP} = R\,F_{A,m} + u_{trim} \tag{1.10}$$

The ratio is a feedforward that moves the feed the instant the rate moves, and the trim's integral removes what the ratio got wrong. S-1's meter error of Section 1.4.3 is such an unknown. The trim must add 4.95 kmol/h of B, and it sees its own moves through the loop's gain of 0.01 per kmol/h and its 5 h time constant, so it is a slow loop on a slow plant. The membrane electrolyzer of Chapter 5 shows the same split on a fast plant. Faraday's law fixes every product from the current, and with the feeds in ratio to the current and trimmed, a load change keeps the anolyte within 0.28 g/L of its setpoint. With the ratios alone it settles 2.0 g/L low, and with the trims alone it swings 4.88 g/L. Chapter 4 sets its air and its dilution water in ratio to the metered sulfur, and Chapter 9 sets its methanol in ratio to the isobutene fed.

The compositions that matter most are rarely measured fast. A model carries them between laboratory results from secondary measurements [31, 32], and each result corrects the model's bias as in Equation (1.11).

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

where $\hat{y}(t_{s,k})$ is the model's value at the instant the sample was drawn and $0 < \lambda_b \leq 1$. During S-1's meter error a sample of loop gas drawn after one hour and reported an hour later would, if compared with the model at the report, put 0.73 points into the bias of a model that is right, 45 % of the excess at that time. Chapter 3 infers its N/C ratio from the reactor's density and temperature and corrects it every 4 h from a laboratory that reports 45 min late, with $\lambda_b = 0.7$. Chapter 8 turns tray temperatures into purities with analyzer trims, and Chapter 10 corrects a fouling viscometer from a laboratory every 75 min. Estimation in general is the subject of Volume 5.

1.4.7 Steady-State Optimization Under Model Mismatch

A real-time optimizer waits for a steady state, fits its model to it and moves the setpoints to the model's optimum. This two-step method is Equation (1.12).

$$\theta_k = \arg\min_{\theta} \left\| \bar{y}_k - \hat{y}(u_k, \theta) \right\|^2, \qquad u_{k+1} = \arg\max_{u} \hat{\phi}(u, \theta_k) \tag{1.12}$$

where u are the setpoints, $\bar{y}_k$ the measured steady state, θ the fitted parameters and $\hat{\phi}$ the model's profit. If the model's structure is right, the method converges to the plant's optimum. If it is not, the method stops where the fitted model's slope is zero, which need not be where the plant's is [33].

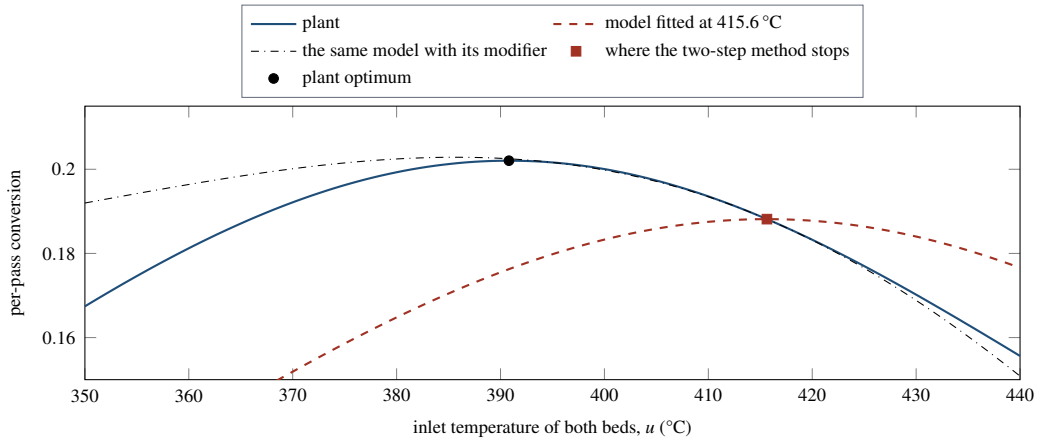

Figure 1.6 Two-Step Optimization on S-1. The Plant's Per-Pass Conversion (Solid) and the Model Fitted Where the Two-Step Method Stops (Dashed) Agree in Value There, but the Model Is Flat and the Plant Falls. Adding the Measured Slope Difference (Dash-Dotted) Moves the Model's Maximum Toward the Plant's Optimum (Dot).

S-1's optimizer sets the inlet temperature of both beds to maximize the per-pass conversion, which at a fixed circulation sets how much makeup the loop can take. Its model is the converter of Section 1.4.2 with an equilibrium constant 43 % too high at every temperature, and it fits only the catalyst activity. As shown in Figure 1.6, from the design inlet of 400 °C, where the plant makes 99.0 % of its best, the method walks to 415.6 °C and stays there, at 93.1 %. The fitted model matches the measured conversion there and predicts a slope of zero, while the plant's slope is −0.0011 per kelvin. A model that reproduces the measurements is not good enough to optimize with, because the optimizer moves along the slopes.

Modifier adaptation keeps the model and corrects it at each run with the measured error and the measured slope [34], as in Equation (1.13).

$$\phi_m(u) = \hat{\phi}(u) + \varepsilon_k + \lambda_k^{\mathsf{T}}(u - u_k), \qquad \lambda_k = \nabla\phi_p(u_k) - \nabla\hat{\phi}(u_k) \tag{1.13}$$

where ε_k is the difference in value and λ_k the difference in gradient, often filtered. The corrected model has the plant's value and slope at u_k, so a point where its slope vanishes is a point where the plant's does. From 415.6 °C with exact plant slopes, S-1's optimizer reaches 99.7 % of the best at its first run and 99.99 % by its fourth, close to the optimum at 390.8 °C. The cost is the plant gradient, which a real plant gives only through deliberate moves and noisy differences.

Chapter 11 measures exactly this on a reactor whose model lacks one of the plant's three reactions. Its two-step optimizer stops at 92.83 % of the best profit while its model predicts more than the plant earns, and modifier adaptation with exact gradients reaches 100.000 % from its tenth run. The other chapters optimize in smaller ways. Chapter 5 plans the current on the hourly price of power with a linear program, Chapter 7 schedules the regeneration energy on the same kind of price, and Chapter 8 moves its spare handle toward the valley's minimum. Chapter 10 backs its reactor off a viscosity index limit that the laboratory sees late. Optimization in general is the subject of Volume 9.

1.4.8 Recycle, Equilibrium and Energy in Each Plant

The three themes of this chapter appear in every plant, though not every plant has all three. Table 1.1 collects what Equation (1.5) predicts and the chapters measure for the plants with a recycle. Table 1.2 lists, for each chapter, what recycles, what equilibrium limits, where the energy goes and the control structure the chapter implements [35]. Chapter 12 adds the measured results and asks how recycle, equilibrium and energy cost shaped each structure.

Table 1.1 What a Recycle Does to a Gain and a Time Constant, in S-1 and in the Chapters

Plant and variable	What returns, how it leaves	Effect
S-1, inert	loop gas, purge 2.5 % of the separator gas	$g_R = 0.975$: 10 % in the loop from 1 %, $\tau_l = 5$ h, forty passes
S-1, excess of A	loop gas, purge	1 % feed error, 11.5 times in the loop ratio, 5 h
S-1, capture	solvent, lean loading	$g_R = -1.8$: 0.237 points per percent at once, 0.085 settled with the duty frozen
2, inert	loop gas, purge 8.4 % of makeup	11.90 % from 1 %, 5.1 h
2, H:N ratio	loop gas, purge	5 % low in the makeup, 8.0 times in the loop, 63 % in 3.2 h
3, N/C ratio	excess ammonia, no exit	$g_R = 1$: integrates 0.0180 an hour for a fresh ratio 0.02 high
7, capture	solvent, lean loading	$g_R = -3.4$: 0.177 at ten minutes, 0.040 after eight hours, 0.177 with the inner loop

Table 1.2 For Each Chapter, What Recycles, What Equilibrium Limits, Where the Energy Goes, and the Structure Implemented

Ch.	Recycle	Equilibrium	Energy	Structure
2	Loop gas, circulation 5.95 times the makeup	Haber, three quench beds	Converter heats its own feed	Makeup on pressure or on production with a low selector, hot spot on quench, slow inert and ratio loops or MPC
3	Unconverted ammonia and carbon dioxide	Conversion rises with N/C	Condenser raises low-pressure steam	Ratio on the flows or on an inferred N/C with laboratory bias, MPC
4	None in the gas	Four beds, double absorption	Bypasses set the bed inlets	Four inlet loops, air and water in ratio to sulfur with trims, MPC on the inlets
5	None	None, Faraday's law	Power on an hourly price	Feeds in ratio to the current with trims, pressure difference cascade with feedforward, price schedule
6	None, one shared exchanger	Vapor and liquid	Compressor near its motor rating	Pressure on the guide vanes with a power override, averaging level, purity loops, MPC
7	Solvent, 70 min lap	Loading against partial pressure	Reboiler duty on a price	Capture on circulation over a lean loading loop, economic supervisor, MPC
8	None	Vapor and liquid	Spare split moves only the duty	Three tray-temperature loops with analyzer bias, split loop, feed feedforward, MPC
9	None	Etherification on catalytic trays	Duty decides the branch	Tray-temperature loop on duty, methanol in ratio, branch monitor, supervisor
10	None	None	Yield against viscosity index, not energy	Viscosity on the cut, mid-ranging to the reactor, bias update, one-way index cap, MPC
11	None	None, irreversible	Jacket, not priced	Three regulatory loops, two-step or modifier adaptation, economic MPC

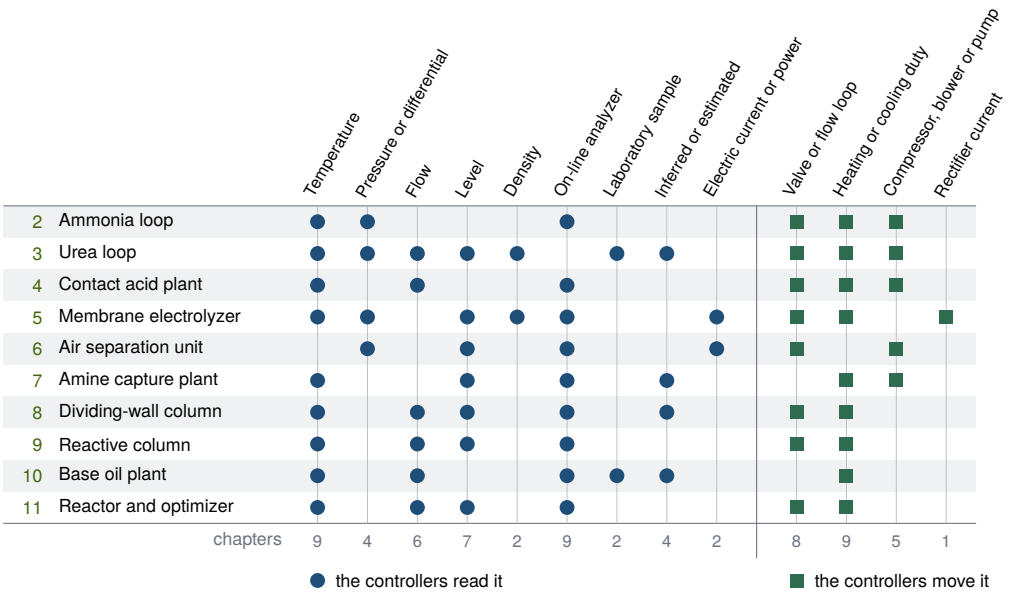

Figure 1.7 Sensors and final elements by chapter. A round mark means that the chapter's controllers or operator read the device, a square mark that they move it. A mark is set only where the chapter's variables table, its instrument text or its section on technology in practice names the device or its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

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 and an online analyzer are each read in nine of the ten chapters, while a laboratory sample is read in two and a composition or property inferred from other measurements is held in four. A heating or cooling duty is moved in nine chapters and a valve or flow loop in eight. Five chapters move a compressor, a blower or a pump, and one moves the current of a rectifier.

The analyzers set the pace of the composition loops. A chromatograph separates the species of a sample in time, as drawn in Figure 1.8, so cycles of minutes are not uncommon and some dead time is inevitable [36]. The chromatograph of Chapter 2 reports the inerts and the ratio every 15 min, and the analyzers of Chapter 8 sample every 10 min and report a cycle later. Chapter 3 has no analyzer for its ratio at all. It infers the ratio from a density and a temperature and corrects it from a laboratory that reports 45 min after each sample, drawn every 4 h.

Most final elements are valves, and many of them sit on steam, cooling water or a jacket. A positioner compares the stem position with the controller's signal and applies as much air to the actuator as it needs, a small cascade of its own, while a valve with excess friction is unresponsive to small reversals of its signal, most of all without a positioner [36]. Figure 1.9 draws both parts. Large machines carry the energy of several plants. The pressure loop of Chapter 6 acts on the guide vanes of the main air compressor under a power cap, and the rectifier of Chapter 5 sets the current from which Faraday's law fixes every product.

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 [37, 38]. A distributed control system hosts the loops, and its supervisory level sends setpoints to the field controllers [39]. The apps run their loops every 5 s in Chapter 5 up to every 0.5 min in Chapter 7, their

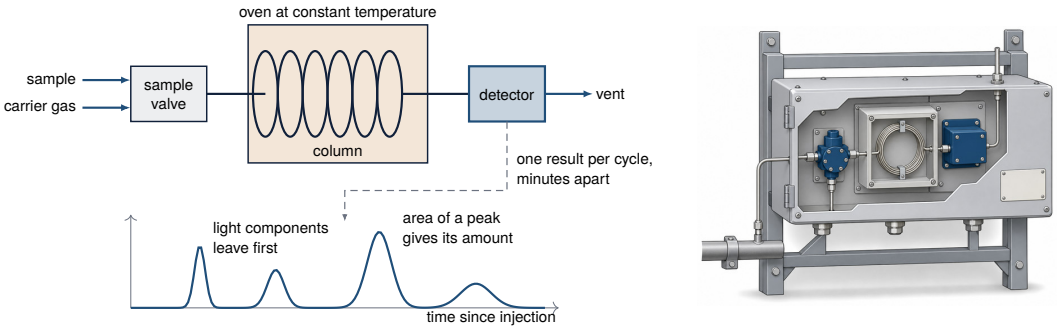

Figure 1.8 Principle of a process gas chromatograph, with a view of the device as installed. Notice that one sample gives one result at the end of the cycle, so a loop on this instrument acts on a value that is always old.

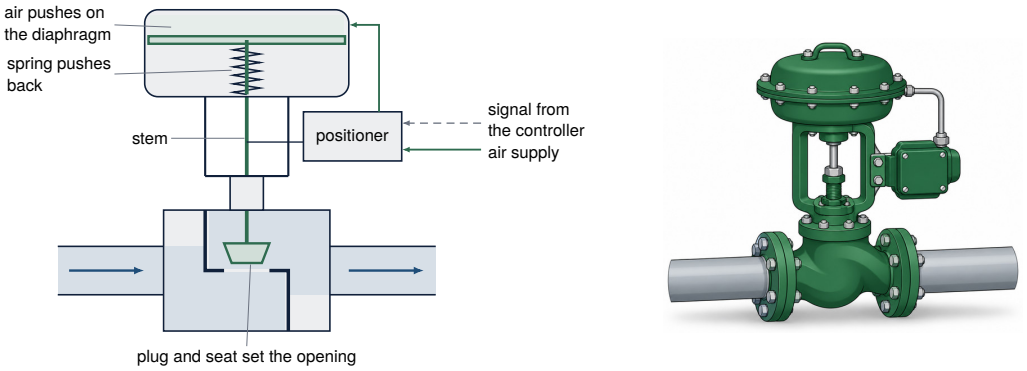

Figure 1.9 Principle of a globe control valve with its spring-and-diaphragm actuator and positioner, with a view of the device as installed. Notice that the positioner closes a loop on the stem position, so the controller's output sets a position, not an air pressure.

predictive controllers every 1 min in Chapter 6 up to every 15 min in Chapter 2, and the optimizer of Chapter 11 no sooner than an hour after its last move.

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

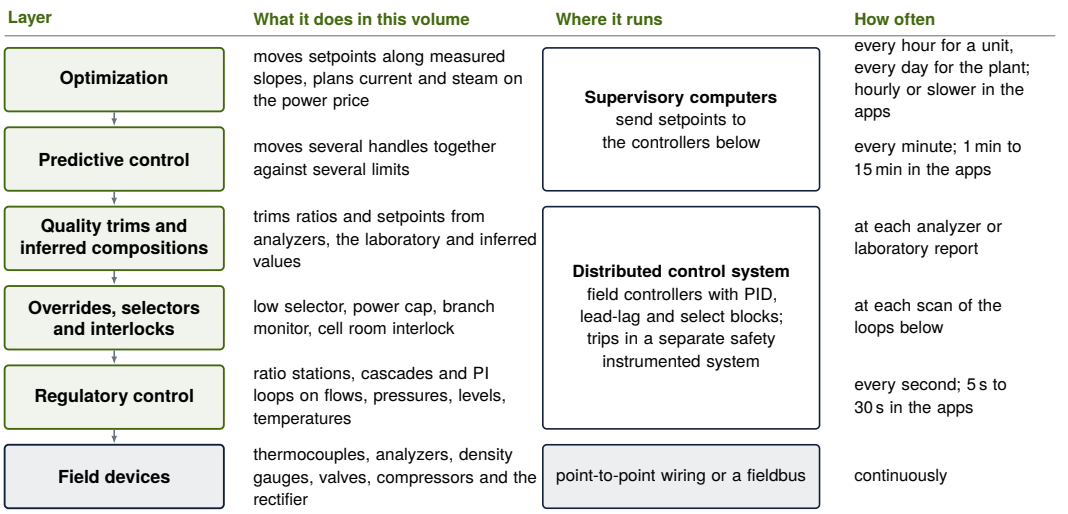

Figure 1.10 The control layers of this volume, the platform that hosts each and how often each acts. The periods of the regulatory, predictive and optimization layers are those of [37, 38], the platforms those of [37, 39], and the ranges in the apps those of the chapters. Notice that each layer sends setpoints to the layer below and acts less often than it.

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.

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.

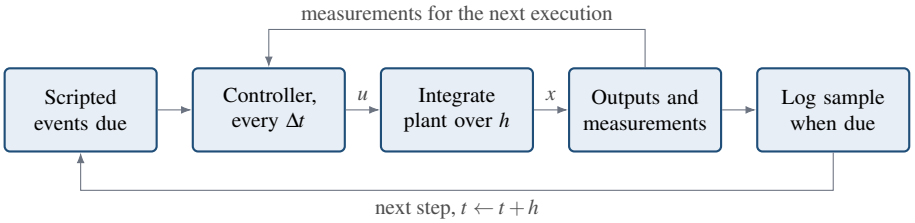

Figure 1.11 Work of the Engine in One Integration Step

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 [40]. 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 [41]. 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 [42]. 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 [43].

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.

Table 1.4 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Ammonia synthesis loop	Recycle sets the gains, and makeup holds pressure or rate.	advanced	3 h
3	Urea synthesis loop	The loop's ratio integrates every feed error and is inferred.	introductory	2 h
4	Contact sulfuric acid plant	The coolest inlet converts most and sits nearest the floor.	intermediate	2 h
5	Membrane electrolyzer	The current is the load, so every feed follows it, and the gas side is faster than feedback.	introductory	2 h
6	Air separation unit	One exchanger couples two columns, so no valve is local.	intermediate	3 h
7	Amine capture plant	An inner loop makes the gain a constant, not the loop faster.	advanced	3.5 h
8	Dividing-wall column	The spare handle moves only the energy, along a valley.	intermediate	2 h
9	MTBE reactive column	The same handles hold more than one steady state.	advanced	2 h
10	Synthetic base oil plant	Two handles reach one viscosity with two different oils.	introductory	2.5 h
11	Reactor with an optimizer	A fitted model stops short, and the slopes decide.	advanced	2.5 h

How the exercises train the habit. One exercise in each chapter is marked *Review the agent*. It presents a tuning, a model or a claim as an AI assistant might propose it, with a flaw, and asks the reader to find the flaw in about five minutes. The flaws are ordinary ones: a unit, a sign, a steady state that does not close, a result measured on the wrong scenario or at a different run length, or a limit that was ignored. Each is found with one of the tests above. The purpose is not distrust. An engineer who can say why an answer is right can hand an assistant much larger tasks and still stand behind the result. That larger picture, in which agents do most of the routine work and the engineer supervises and checks, is the subject of Chapter 12.

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 3, 5 and 10 form a first course: a ratio that integrates, feeds that follow a current, and two handles on one quality with a late laboratory. They rest on Sections 1.4.3 and 1.4.6. Chapters 4, 6 and 8 add an equilibrium-limited reactor train, thermal coupling and a column's energy valley. Chapters 2, 7, 9 and 11 combine a recycle that moves gains and time scales, more than one steady state, or an optimizer with a wrong model, and are best left until last. A reader who follows the fertilizer chain can read Chapters 6, 2, 7 and 3 in that order.

KEY POINTS

- In a synthesis loop most of what enters the reactor has been there before, so the recycle sets how large and how slow each response is.
- A ratio held between two measured feeds is not a ratio held in the loop, because a reactant that cannot leave integrates every feed error.
- A lean loading loop does not make the absorber faster. It makes one gain out of two.
- The spare handle of a column with a side draw moves only the energy, and the energy it moves has a minimum.
- A model that reproduces the measurements is not good enough to optimize with, because the optimizer moves along the slopes.

Where else this applies. Recycle and inventory control return in Volume 6, ratio control in Volume 3, estimation in Volume 5 and optimization in Volume 9.

References

[1] F. Wöhler, “Ueber künstliche bildung des harnstoffs,” *Annalen der Physik*, vol. 88, no. 2, pp. 253–256, 1828.

[2] C. Linde, “Process of producing low temperatures, the liquefaction of gases, and the separation of the constituents of gaseous mixtures,” U.S. Patent 727,650, filed July 9, 1895, granted May 12, 1903, 1903. [Online]. Available: https://patents.google.com/patent/US727650A/en

[3] R. Knietsch, “Ueber die Schwefelsäure und ihre Fabrication nach dem Contactverfahren,” *Berichte der deutschen chemischen Gesellschaft*, vol. 34, no. 3, pp. 4069–4115, 1901.

[4] F. Haber, “The synthesis of ammonia from its elements,” Nobel Lecture, 2 June 1920, 1920, nobelprize.org/uploads/2018/06/haber-lecture.pdf, read November 2026.

[5] R. R. Bottoms, “Process for separating acidic gases,” U.S. Patent 1 783 901, Dec. 2, 1930, filed October 7, 1930.

[6] E. A. Wolff and S. Skogestad, “Operation of integrated three-product (Petlyuk) distillation columns,” *Industrial & Engineering Chemistry Research*, vol. 34, no. 6, pp. 2094–2103, 1995.

[7] Computer History Museum, “Computers in control,” CHM Revolution, online exhibition, section Real-Time Computing, accessed 18 November 2026. [Online]. Available: https://www.computerhistory.org/revolution/real-time-computing/6/130

[8] M. Paidar, V. Fateev, and K. Bouzek, “Membrane electrolysis—history, current status and perspective,” *Electrochimica Acta*, vol. 209, pp. 737–756, 2016.

[9] E. W. Jacobsen and S. Skogestad, “Multiple steady states in ideal two-product distillation,” *AIChE Journal*, vol. 37, no. 4, pp. 499–511, 1991.

[10] M. Bui, C. S. Adjiman, A. Bardow, E. J. Anthony, A. Boston, S. Brown, P. S. Fennell, S. Fuss, A. Galindo, L. A. Hackett, J. P. Hallett, H. J. Herzog, G. Jackson, J. Kemper, S. Krevor, G. C. Maitland, M. Matuszewski, I. S. Metcalfe, C. Petit, G. Puxty, J. Reimer, D. M. Reiner, E. S. Rubin, S. A. Scott, N. Shah, B. Smit, J. P. M. Trusler, P. Webley, J. Wilcox, and N. Mac Dowell, “Carbon capture and storage (CCS): the way forward,” *Energy & Environmental Science*, vol. 11, no. 5, pp. 1062–1176, 2018.

[11] C. Bosch, “The development of the chemical high pressure method during the establishment of the new ammonia industry,” Nobel Lecture, 21 May 1932, 1932, nobelprize.org/uploads/2018/06/bosch-lecture.pdf, read November 2026.

[12] C. Linde, “Apparatus for producing pure nitrogen and pure oxygen,” U.S. Patent 795,525, filed October 22, 1903, granted July 25, 1905, 1905. [Online]. Available: https://patents.google.com/patent/US795525A/en

[13] European Commission, “Reference document on best available techniques for the manufacture of large volume inorganic chemicals: Ammonia, acids and fertilisers,” European Commission, Integrated Pollution Prevention and Control, Tech. Rep., Aug. 2007, eippcb.jrc.ec.europa.eu, read November 2026.

[14] A. A. Kiss and M. Jobson, “Taking reactive distillation to the next level of process intensification,” *Chemical Engineering Transactions*, vol. 69, pp. 553–558, 2018.

[15] M. Pitt, “Taking a look back at control: Part 2,” *The Chemical Engineer*, Oct. 2023, published online 26 October 2023. [Online]. Available: https://www.thechemicalengineer.com/features/taking-a-look-back-at-control-part-2/

[16] M. Appl, “Ammonia, 2. production processes,” in *Ullmann’s Encyclopedia of Industrial Chemistry*. Wiley-VCH, 2011.

[17] M. J. King, W. G. Davenport, and M. S. Moats, “Oxidation of SO_2 to SO_3—equilibrium curves,” in *Sulfuric Acid Manufacture*. Elsevier, 2013, pp. 125–133.

[18] U.S. Geological Survey, *Mineral Commodity Summaries 2025*. Reston, VA: U.S. Geological Survey, 2025, nitrogen (fixed)–ammonia, pp. 128–129, and Sulfur, pp. 172–173.

[19] J. C. Morud and S. Skogestad, “Analysis of instability in an industrial ammonia reactor,” *AIChE Journal*, vol. 44, no. 4, pp. 888–895, 1998.

[20] M. M. Denn and R. Lavie, “Dynamics of plants with recycle,” *The Chemical Engineering Journal*, vol. 24, no. 1, pp. 55–59, 1982.

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

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

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

[24] M. Souders and G. G. Brown, "Fundamental design of absorbing and stripping columns for complex vapors," *Industrial & Engineering Chemistry*, vol. 24, no. 5, pp. 519–522, 1932.

[25] G. T. Rochelle, "Amine scrubbing for CO_2 capture," *Science*, vol. 325, no. 5948, pp. 1652–1654, 2009.

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

[27] I. J. Halvorsen and S. Skogestad, "Minimum energy consumption in multicomponent distillation. 2. three-product Petlyuk arrangements," *Industrial & Engineering Chemistry Research*, vol. 42, no. 3, pp. 605–615, 2003.

[28] M. Variny, D. Jediná, M. Rimár, J. Kizek, and M. Kšiňanová, "Cutting oxygen production-related greenhouse gas emissions by improved compression heat management in a cryogenic air separation unit," *International Journal of Environmental Research and Public Health*, vol. 18, no. 19, p. 10370, 2021.

[29] S. Hauan, T. Hertzberg, and K. M. Lien, "Why methyl tert-butyl ether production by reactive distillation may yield multiple solutions," *Industrial & Engineering Chemistry Research*, vol. 34, no. 3, pp. 987–991, 1995.

[30] L. B. Koppel, "Input multiplicities in nonlinear, multivariable control systems," *AIChE Journal*, vol. 28, no. 6, pp. 935–945, 1982.

[31] 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.

[32] 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.

[33] J. F. Forbes, T. E. Marlin, and J. F. MacGregor, "Model adequacy requirements for optimizing plant operations," *Computers & Chemical Engineering*, vol. 18, no. 6, pp. 497–510, 1994.

[34] A. Marchetti, B. Chachuat, and D. Bonvin, "Modifier-adaptation methodology for real-time optimization," *Industrial & Engineering Chemistry Research*, vol. 48, no. 13, pp. 6022–6033, 2009.

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

[36] 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.

[37] 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.

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

[39] 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.

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

[41] 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.

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

[43] 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 gas-phase reaction A + B $\rightleftharpoons$ C is fed with equal A and B, and at the bed temperature $K(T)\,p = 2$, as in Section 1.4.2.

(a) Find the equilibrium conversion of A with no inert in the feed. *Checkpoint:* $a = 0.5$, $\kappa = 1$, $X^* = 0.423$.

(b) Find it again with 10 % inert at the same $K(T)\,p$. *Checkpoint:* $a = 0.45$, $\kappa = 0.9$, $X^{*} = 0.398$.

1.2 Concept check · 2 min

A two-bed adiabatic converter for a reversible exothermic reaction runs both inlets at 400 °C. Raising both inlets to 420 °C lowers the per-pass conversion. Why?

(a) the catalyst slows down when it is hotter, so each bed gets less far along its line

(b) the adiabatic rise per unit of conversion grows with temperature, so each bed heats faster

(c) each bed starts closer to the falling equilibrium curve, so the line meets it at a lower conversion

(d) the interbed cooler can no longer return the gas to the inlet temperature

1.3 Theory and calculation · 4 min

A synthesis loop holds 600 kmol of gas. Its makeup is 1200 kmol/h with 0.8 % of an inert that leaves only through a purge of 80 kmol/h.

(a) Find the settled inert fraction in the loop and its time constant. *Checkpoint:* 12.0 % and 7.5 h.

(b) How long after a step in the makeup's inert does the loop reach 95 % of its new level? *Checkpoint:* About 22.5 h.

1.4 Theory and calculation · 5 min

In S-1 the meter on B starts to read 2 % high, so the makeup brings 495 kmol/h of A and 485.1 kmol/h of B. The purge stays at 100 kmol/h, the makeup's inert at 10 kmol/h and the loop at 500 kmol.

(a) Find the settled excess $y_A - y_B$ and the time constant. *Checkpoint:* 0.099 and 5 h.

(b) With the loop's C production $(F_M - F_P)/2$, find y_A, y_B and the loop ratio y_A/y_B, and compare its error with the makeup's. *Checkpoint:* 0.4995, 0.4005, 1.247 against 1.020, about 12 times.

1.5 Concept check · 2 min

A urea loop sets its ammonia in ratio to the measured carbon dioxide, and the excess ammonia returns through the recycle with no purge. The ammonia meter starts to read 1 % high. What does the ammonia to carbon dioxide ratio in the reactor do?

(a) it falls and keeps falling until something outside the ratio station corrects it

(b) it falls by about 1 % and settles there, as the feeds do

(c) it falls by about ten times 1 % and settles with a time constant of hours

(d) it does not change, because the ratio station holds the measured ratio

1.6 Theory and calculation · 3 min

A circulation step raises an absorber's capture with a gain of 0.20 points per unit before the returning solvent changes, and the solvent loop returns a fraction $g_R = -3$ of the change once it has gone around.

(a) Find the settled gain. *Checkpoint:* 0.050 points per unit.

(b) A capture loop is tuned for the settled gain, and then an inner lean loading loop is closed. By what factor is the capture loop now too aggressive? *Checkpoint:* Four.

1.7 Theory and calculation · 4 min

A three-stage absorber treats a dilute gas with 4 % of a solute. The equilibrium line is $y = mx$, the absorption factor is $\mathscr{A} = 1.5$, and the lean solvent's back pressure is $mx_L = 0.002$. Find the outlet mole fraction and the fraction of the solute removed.

1.8 Theory and calculation · 5 min

S-1's crude product is 100 kmol/h of saturated liquid with a light, C and a heavy in equal parts, at relative volatilities 4, 2 and 1. The Underwood roots are $\theta_1 = 2.756$ and $\theta_2 = 1.244$.

(a) Find the minimum vapor of the light split, with only the light in the distillate, and of the heavy split, with the light and C in the distillate. *Checkpoint:* 107.2 kmol/h and 136.6 kmol/h.

(b) The second column of the direct sequence needs 100.0 kmol/h. Compare the direct sequence with a thermally coupled column. *Checkpoint:* 207.2 kmol/h against 136.6 kmol/h, 34.1 % less.

1.9 Concept check · 3 min

A real-time optimizer refits its model at every steady state, and after a few runs it stops moving. The fitted model reproduces the measured outputs at the final point. What can be concluded?

(a) the plant is at its optimum, because a model that fits the measurements has the plant's slopes
(b) the plant is at its optimum only if the steady-state detector was strict enough
(c) the model's structure must be wrong, because a refitted model never stops
(d) the fitted model's slope is zero there, which says nothing certain about the plant's slope

1.10 Theory and calculation · 4 min

A plant's profit is $\phi_p(u) = -(u-5)^2$, and its model is $\hat{\phi}(u) = -(u-3)^2 + c$, whose only fitted parameter is the offset c.

(a) Where does the two-step method stop? *Checkpoint:* $u = 3$.

(b) From $u_k = 3$, find the gradient modifier with exact plant gradients and the next setpoint of modifier adaptation. *Checkpoint:* $\lambda_k = 4$, $u_{k+1} = 5$.

CHAPTER 2

Ammonia Synthesis Loop Simulator for Recycle Amplification and Inventory Pairing

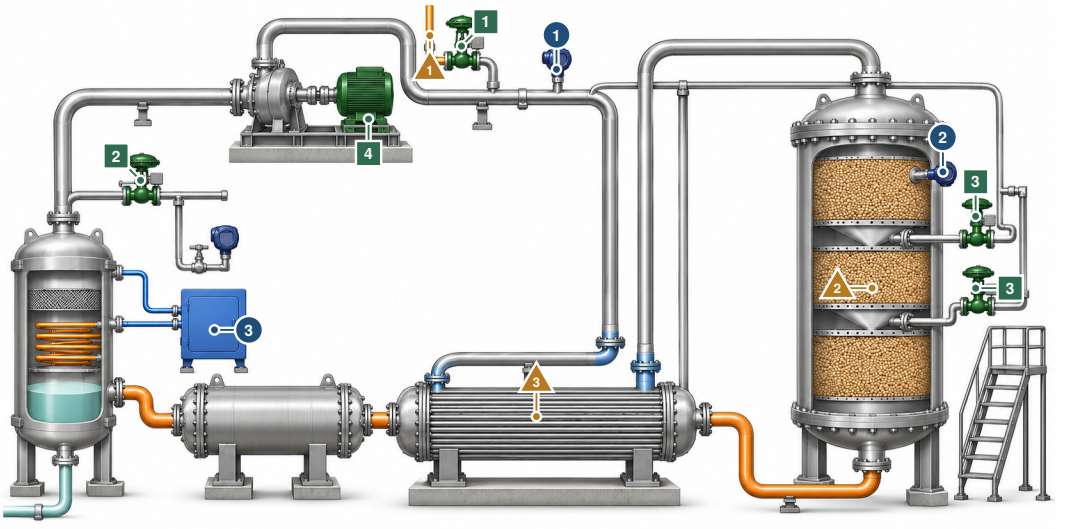

The controller moves 1 makeup gas valve 2 purge valve 3 quench gas between the beds 4 circulator speed

The controller reads 1 loop pressure 2 hot spot in the beds

3 chromatograph: inerts and the hydrogen to nitrogen ratio, every quarter hour

What disturbs it 1 makeup gas that is dirtier, off ratio or scarcer 2 aging catalyst

3 an exchanger that loses effectiveness

Follow the gas around: it leaves the separator, passes the circulator, takes on makeup, is heated in the exchanger and reacts in the converter, so most of it goes around again and only the small purge lets the inerts out. Illustration, not to scale.

IN PLAIN TERMS **A Loop That Magnifies Small Mistakes**

The main idea. An ammonia plant turns only a little of its gas into ammonia on each pass through the reactor. So it sends the rest around again in a loop, and a small bleed lets out the gases that cannot react.

The challenge. A small error in the fresh gas grows much larger inside the loop, and it takes hours to arrive. The valve that tops up the loop can hold the pressure or the output steady, and the trouble then shows up in the other one.

Why it matters. No alarm sounds while the plant quietly makes less ammonia or uses up its safety margin.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page; the history and fact boxes of Sections 2.1 and 2.2; the try-it boxes of Sections 2.4 and 2.6; the fact box of Section 2.6; the key points
Study	60 min	the whole chapter with its pause points and the exercises
Practice	20 min	the first page; Section 2.6 with its result chart; Section 2.7; Section 2.8 with its open challenge

The control question An ammonia converter reacts about four percent of the gas on each pass, so the plant is a recycle with a converter inside it. How far and how slowly does the loop amplify a small error in the makeup gas, and which variable should the makeup valve hold: the loop pressure or the production rate?

Learning objectives

- Compute the steady inert level and the composition time constant of the loop from its holdup, its makeup and its purge.
- Explain why an exothermic equilibrium converter has a best temperature and, because it heats its own feed, a second steady state in which it is out.
- Measure the amplification of a front-end ratio error and explain why it is larger for a hydrogen-rich error than for a nitrogen-rich one.
- Predict where a disturbance appears under each pairing of the makeup valve, and read constraint margin from a valve position.
- Compare two slow PI loops with one predictive plan on the purge and the trim, in integrated error and in margin.

Prerequisites Sections 1.4.2 (equilibrium-limited beds), 1.4.3 (inert balance, recycle gain, feed ratio error, throughput manipulator), 1.6.2 (velocity form) and 1.6.4 (scenarios and exported files). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/ammonia-loop/

See also Volume 6, Chapter 7 for the snowball effect of a reactor and column recycle, and Volume 6, Chapter 8 for a gas recycle with an inert and a purge. Volume 7, Chapter 2 for ignition and extinction folds in an exothermic reactor. Volume 9, Chapter 4 for offset-free predictive control with a Kalman disturbance estimate.

2.1 Introduction

A recycle is normally drawn as an arrow that leaves the separator and returns to the front of the block diagram, and it is normally treated as bookkeeping. An ammonia synthesis loop makes that treatment untenable. The extent of reaction over the iron catalyst is 0.0386 per pass at the design point of the simulator, so 6.00 kmol/s of gas circulates against 1.01 kmol/s of makeup, a recycle ratio of 5.95 (circulator flow over makeup, as the application reports it), and the properties a student would predict from the converter alone are not the properties the plant has. The converter answers a temperature change in minutes and the loop answers a composition change in five hours. A front-end error of five percent in the hydrogen to nitrogen ratio arrives inside the loop as an error of forty percent. The misconception the simulator corrects is that a recycle plant is a unit operation with a return line, when in fact the return line sets the gains, the time constants and the place a disturbance finally appears.

DID YOU KNOW? World Ammonia Output, 2024

The world's ammonia plants produced an estimated 150 million metric tons in 2024, counted as the nitrogen the ammonia contains. In the United States 37 plants ran that year, natural gas was the dominant feedstock, and about 88 percent of the output went to fertilizer [1].

The shape of the problem is general and it does not require a catalyst. Whenever a component enters a recycle loop and can leave it only through a small purge, the steady-state loop concentration is the ratio of two small flows and the loop amplifies any persistent imbalance in proportion to the inverse of the purged fraction. Reactor and separator systems joined by a recycle carry that amplification as a documented structural property [2], the snowball effect of Volume 6, Chapter 7, and the loop inventory divided by the purge is the time constant that comes with it. A second and separate question follows from conservation of mass. One valve on a recycle loop sets throughput, and the plant designer decides which variable that valve is allowed to hold, which fixes where every disturbance lands and which measurements are informative [3, 4]. Both questions are the ordinary subject matter of plantwide control rather than special features of an ammonia plant [5], and

hydrotreaters, reforming units, polymerization loops and gas plants with recovery all present them in the same form.

Ammonia loop models in the literature answer process questions. The converter has been simulated at steady state and in transient since 1965 [6, 7], the loop variables have been mapped by steady-state simulation [8], and a full control structure has been designed for the process by self-optimizing methods [9]. However, a student who is handed the loop as a set of transfer functions never performs the experiment that shows a five percent front-end error becoming forty percent, because the amplification lives in a mass balance that a transfer function has already absorbed into a gain. This chapter works with an interactive ammonia synthesis loop whose per-pass conversion is small enough that the recycle governs the plant, so that recycle amplification and the choice of what the makeup valve holds become experiments rather than assertions. The simulator carries fourteen states, a three-bed quench converter solved against the Gillespie-Beattie equilibrium, a feed and effluent exchanger that gives the vessel two steady states, a refrigerated separator, a sampled loop analyzer, a low selector on the makeup valve and a linear model predictive controller for the two slow handles. Thirteen scenarios run the same plant through a converter optimum, an extinction, two inventory pairings, three slow composition disturbances and a day on which two PI loops and one plan meet the same pair of drifts. The plant runs in a browser and exports every variable it computes. The equilibrium correlation, the kinetic form and the plantwide arguments are standard and are cited to their sources. Chapter 1 treats the same balances on its running example in Sections 1.4.2 and 1.4.3.

Section 2.2 traces how the loop came to be controlled. Section 2.3 carries every equation and constant the plant is built from. Section 2.4 derives the steady inert level and its time constant, measures the amplification of a front-end ratio error and states the inventory question. Section 2.5 describes the four PI loops, the selector and the predictive controller. Section 2.6 reports the scenarios, from the converter outward, and compares the two slow loops with one plan. Section 2.7 names the hardware and the practice behind the model. Sections 2.8 and 2.9 state where the model stops holding and what it supports beyond the implemented controllers, and Section 2.10 collects the measured results.

2.2 The Synthesis Loop and How It Came To Be Controlled

Nitrogen fixation at industrial scale began with the high-pressure catalytic synthesis commercialized in 1913, and the process is credited with sustaining roughly half of the present world population [10, 11]. The reaction that made that possible is also the reason the plant is a loop. Ammonia synthesis is exothermic, so its equilibrium constant falls steeply with temperature, while the rate over promoted iron is negligible until roughly 400 °C. The catalyst therefore has to be run where the equilibrium is poor, and the equilibrium ammonia fraction at any temperature the catalyst tolerates is small enough that the unconverted gas has to be recovered and returned [12]. Recovery by refrigerated condensation and return by a circulator is the arrangement that every modern plant uses, and it is what turns a reactor design problem into an inventory problem.

> **HISTORY A Loop Under Constant Pressure, 1905 to 1920**
>
> In his Nobel lecture of 1920, Fritz Haber traced the loop to his own laboratory. In work set aside in 1905 he had found that at red heat no catalyst gives more than a trace of ammonia at normal pressure, and had described a circulation that brought the gas to the hot metal and then washed the ammonia out, converting the gas stage by stage. Taking the problem up again in 1908 with Robert le Rossignol, he worked at 200 atmospheres, the limit of his compressor. Their small plant made and removed ammonia in one cycle at constant high pressure and led the heat of the outgoing gas to the fresh gas, so the process supplied its own heat [13]. Industry, he said, kept both features: the circulator and the feed and effluent exchanger of Figure 2.1.

The engineering attention followed the vessel first. Converter modeling was posed as a steady-state simulation of an autothermal reactor whose feed is preheated by its own effluent, and the transient behavior of the same reactor was analyzed in the same year [6, 7]. The structural fact those studies inherited had been established a decade earlier. An autothermal process whose feed is heated by its product has a heat generation curve and a heat removal line that can intersect more than once, so it possesses an ignited state, an extinguished state and an ignition condition between them [14]. The cooled stirred tank of Volume 7, Chapter 2 has the same ignition and extinction folds, and the adiabatic beds of Chapter 4 have an ignition floor of the same origin. The practical consequence is that a converter which has gone out does not come back when the operator returns the quench valve to where it was, and that the same feedback can produce sustained oscillation on an industrial machine rather than a clean approach to either branch [15].

Pause. Try Exercise 2.2 before reading on.

Loop questions were answered later and mostly by simulation. The purge rate, the inert level, the loop pressure and the hydrogen to nitrogen ratio were mapped against each other by steady-state calculation, which established the now familiar picture of a purge that buys conversion with hydrogen [8]. Before an on-line analyzer existed, the loop composition was a laboratory result rather than a measurement, the purge was a valve set by a standing instruction and revised on a shift report, and the converter was managed by the temperature profile because the temperature profile was the only thing measured continuously. Gas chromatography on the loop replaced the standing instruction with a controller, and a controller pointed at a five-hour plant is a different design problem from the two temperature loops next to it.

> **HISTORY The Leaks Were the Purge, 1932**
>
> Carl Bosch's Nobel lecture of 1932 shows how the high-pressure plants he built watched their loops. An automatic recorder of the gas density followed the composition of the circulating gas, and it also controlled the proportion of hydrogen in the mixture. Of the thousands of flanges and valves he wrote that everything had to be gas-tight, and that the measures had been so successful that the gas losses were just enough to keep the argon from the air and the accumulating methane from rising too high in the circulating system [16]. The leaks did the work of a purge. In the chapter's plant that work belongs to valve FV-102, watched by chromatograph GC-101.

The current arrangement became standard when plantwide control was posed as a subject in its own right rather than as a collection of loops. The recycle amplification of a persistent imbalance was named and quantified for reactor and separator systems [2], the choice of which valve sets throughput was identified as the decision that fixes the rest of the control structure [3], and the requirement that every inventory in a plant be held by a flow on one side or the other of it was stated as a consistency condition [4, 5]. The Tennessee Eastman process of Volume 6, Chapter 8, with its gas recycle, its inert and its purge, became the common test of those ideas. Applied to this process the result is four loops, a pressure controller and a converter temperature controller that answer in minutes, and an inert controller and a synthesis gas ratio controller that answer in hours, with the throughput manipulator

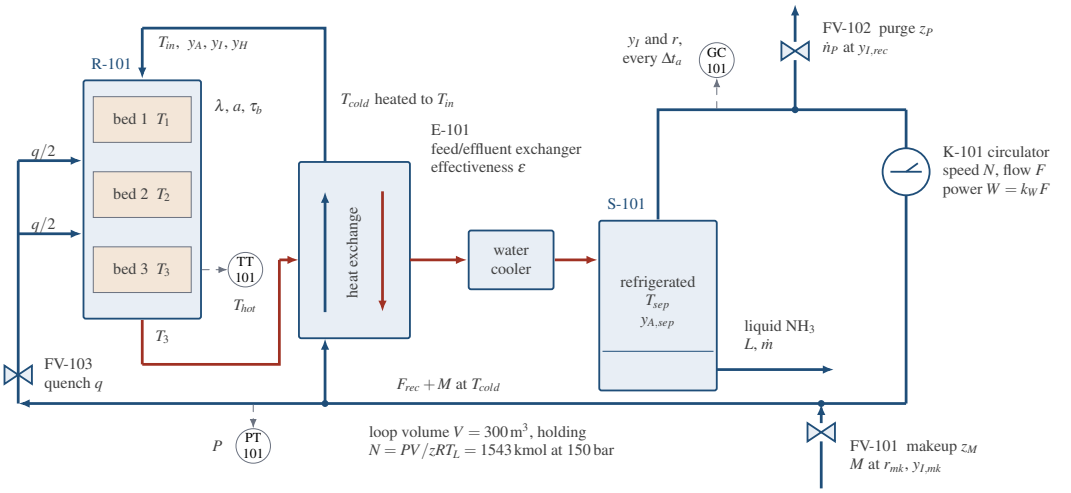

Figure 2.1 Synthesis loop L-101 with its three-bed quench converter, the feed and effluent exchanger that preheats the converter charge with the converter product, the refrigerated separator and the circulator. Gas at T_{cold} leaves the circulator, is heated to T_{in} in E-101 by the bed three effluent at T_3, and passes down through beds at T_1, T_2 and T_3 with a fraction q of the cold gas split evenly between the two quench points. Cooled effluent enters S-101, where liquid ammonia L is condensed at T_{sep} and leaves as product $\dot{m}$. The overhead gas splits into purge $\dot{n}_P$ through FV-102 and recycle F_{rec} back to the circulator, where makeup M at ratio r_{mk} and inert fraction $y_{I,mk}$ joins it through FV-101. Composition y_A, y_I and y_H is carried at the converter inlet mixing point, and the loop holds N kilomoles of gas in volume V at pressure P. PT-101 and TT-101 are continuous, and chromatograph GC-101 reports the loop inert fraction and the hydrogen to nitrogen ratio once every Δt_a.

chosen deliberately rather than by default [9]. Interest in ammonia as an energy carrier rather than only as a fertilizer intermediate has renewed attention to the efficiency of exactly this loop [17], and the efficiency number the loop owns is the syngas consumed per tonne of product, 128.6 kmol/t at the design purge here against a stoichiometric 117.4 kmol/t.

2.3 Process Model

Loop L-101 circulates 6.00 kmol/s of synthesis gas at 150 bar through a three-bed quench converter R-101, a feed and effluent exchanger E-101, a water cooler, and a refrigerated separator S-101 that condenses the ammonia out at –4.7 °C. Makeup syngas enters through valve FV-101 on the circulator discharge, and the purge leaves through FV-102 on the separator overhead. Quench valve FV-103 takes cold gas around the exchanger and injects half of it between beds one and two and half between beds two and three. The circuit, its instrument tags and every symbol carried by the equations below are drawn in Figure 2.1.

Table 2.1 lists everything the control system touches. Pressure is in bar, temperature in °C, flow in kmol/s, valve position in percent, and the clock is in minutes, so every rate below is per minute and every flow that enters one is converted with the factor of 60 written out. Fourteen quantities are integrated. Eight of them are the plant, and six are first-order actuator positions.

2.3.1 Loop Inventory and Pressure

The gas held in the loop is an ideal inventory with a compressibility correction at a fixed mean temperature, so pressure and moles are the same variable, as shown in Equation (2.1).

$$N = \frac{PV}{zRT_L}, \qquad \frac{\mathrm{d}P}{\mathrm{d}t} = 60\,\frac{P}{N}\,(M + F_{rec} - F) \tag{2.1}$$

Table 2.1 Classification of Variables. Nominal Values Are Those of the Solved Design Point at Which Every Scenario Opens.

Symbol	Role	Nominal	Range
P loop pressure	CV, measured	149.97	120–168 bar SP
T_{hot} hot spot	CV, measured	505.2	465–525 °C SP
y_I loop inerts	CV, sampled	10.07	5–18 % SP
r loop H_2:N_2	CV, sampled	3.000	2.6–3.4 SP
z_M makeup valve	MV	59.3	0–100 %
z_P purge valve	MV	37.4	0–100 %
q quench valves	MV	10.0	0–45 %
r_{mk} front-end trim	MV	3.00	2.4–3.6
N circulator speed	MV, operator	100	60–110 %
refrigeration	MV, operator	63	0–100 %
a catalyst activity	DV, unmeasured	1.00	0.50–1.05
makeup inert factor	DV, unmeasured	1.00	0.30–3.00
Δr_{mk} ratio drift	DV, unmeasured	0	±0.30
makeup availability	DV, measured	1.00	0.40–1.05
exchanger condition	DV, unmeasured	1.00	0.85–1.03
refrigeration shortfall	DV, unmeasured	0	−10–20 °C
P (bar)	state	149.97	20–260
y_A ammonia	state	0.0316	0.0001–0.45
y_I inerts	state	0.1007	0–0.60
y_H hydrogen	state	0.6508	0.05–0.92
T_1, T_2, T_3 (°C)	states	497/505/498	100–700
T_{sep} (°C)	state	−4.65	−40 to 60
six actuator positions	states	as above	as above

where N is the loop holdup in kmol, V is the loop gas volume, z is the compressibility, T_L is the mean loop gas temperature, M is the makeup flow, F_{rec} is the recycle returning from the separator and F is the circulator throughput, all in kmol/s. At the design point $N = 1543$ kmol. Pressure is inventory, and nothing else on this plant is, which is the fact the whole control structure is arranged around.

The circulator is a fixed volumetric machine, so its molar throughput rises with speed and with density, and its shaft power is proportional to throughput, as shown in Equation (2.2).

$$F = F_0 \frac{N_{sp}}{100} \frac{P}{P_{ref}}, \qquad W = k_W F \tag{2.2}$$

where N_{sp} is the speed demand in percent, F_0 is the rated throughput and k_W is the power coefficient. Equation (2.2) couples pressure back into contact time, because a loop running at higher pressure passes gas through the beds faster.

2.3.2 Composition at the Mixing Point

The loop is treated as one well-mixed gas volume whose composition is taken at the converter inlet, where the recycle and the makeup meet. Writing the species balance as $\mathrm{d}(N y_k)/\mathrm{d}t$ and substituting the inventory balance of (2.1) removes the total flow terms and leaves Equation (2.3).

$$\frac{\mathrm{d}y_k}{\mathrm{d}t} = \frac{60}{N}\left[M\left(y_{k,mk} - y_k\right) + F_{rec}\left(y_{k,rec} - y_k\right)\right] \tag{2.3}$$

where y_k is the mole fraction of ammonia, inerts or hydrogen at the mixing point, $y_{k,mk}$ is that species in the makeup and $y_{k,rec}$ is that species in the recycle. Nitrogen is the balance, $y_N = 1 - y_A - y_I - y_H$, and the loop ratio is $r = y_H / y_N$. Makeup carries no ammonia, so $y_{A,mk} = 0$, and the other two follow

from the front-end specification, $y_{I,mk}$ set by the reformer and $y_{N,mk} = (1 - y_{I,mk})/(1 + r_{mk})$ with $y_{H,mk} = r_{mk} y_{N,mk}$.

Equation (2.3) is the whole recycle argument in one line. A species that the converter does not consume and the separator does not condense appears in the makeup term and in the recycle term with almost the same concentration, so its only net removal is the purge, and it accumulates until the two balance.

2.3.3 Equilibrium and One Adiabatic Bed

The equilibrium constant is the Gillespie-Beattie correlation for the Haber equilibrium, written for ammonia over the roots of its reactants and returning K_p in inverse atmospheres, as shown in Equation (2.4) [18].

$$\begin{aligned} \log_{10} K_p = {} & \frac{2250.322}{T} - 0.85340 - 1.51049 \log_{10} T \\ & - 2.58987 \times 10^{-4}\, T + 1.48961 \times 10^{-7}\, T^2 \end{aligned} \tag{2.4}$$

where T is in kelvin. Because the reaction is exothermic, K_p falls by a factor of 1.49 for every 30 °C of rise at 400 °C and by 1.36 at 500 °C, and its square, the constant for the whole reaction, by about two, so every degree that speeds the reaction up worsens the destination it is heading for. With $y_N = (1 - y_A - y_I)/(1 + r)$ and $y_H = r y_N$, the equilibrium condition $y_A = K_p P\, y_N^{1/2} y_H^{3/2}$ collapses to a quadratic in the unreacted fraction, which is solved in closed form rather than by bisection, as shown in Equation (2.5).

$$y_A^{eq} = s - u, \quad u = \frac{-1 + \sqrt{1 + 4\kappa s}}{2\kappa}, \quad \kappa = \frac{K_p P\, r^{3/2}}{(1 + r)^2} \tag{2.5}$$

where $s = 1 - y_I$ and u is the combined hydrogen and nitrogen fraction at equilibrium. It is the same kind of closed form as Equation (1.2) of Chapter 1, written in mole fractions instead of conversion. Equation (2.5) is evaluated inside the bed solve, which is evaluated inside the derivative, which the fixed-step integrator calls four times per step, so the closed form matters for the speed of the steady-state solver.

A bed is adiabatic, so its state moves along a straight line in the ammonia against temperature plane, and the extent it can reach is where that line meets the equilibrium curve, as shown in Equation (2.6).

$$\begin{aligned} T(x) &= T_{in} + \lambda x, \qquad y_A(x) = \frac{y_A + 2x}{1 - 2x} \\ x_{eq} &: \; y_A(x) = y_A^{eq}\big(T(x), P, r, \tfrac{y_I}{1-2x}\big) \\ x &= \eta\, x_{eq}, \qquad \eta = 1 - \exp(-k(T_m)\, \tau_b^*\, a) \end{aligned} \tag{2.6}$$

where x is the nitrogen extent per mole of bed feed, λ is the adiabatic temperature rise for complete conversion of the nitrogen present, η is the fractional approach to equilibrium, a is the catalyst activity multiplier and $\tau_b^* = \tau_b F_0 / \max(0.3, F)$ is the contact time scaled by the actual circulation. The equilibrium extent is found by bisection in forty steps between zero and $0.98(1 - y_A - y_I)/(1 + r)$, which is the nitrogen the feed actually contains. The rate group is one Arrhenius term evaluated at the mean bed temperature $T_m = T_{in} + \tfrac{1}{2}\lambda x_{eq}$ with a square-root hydrogen coverage, as shown in Equation (2.7).

$$k(T_m) = k_0\, \mathrm{e}^{-E_a / R T_m} \sqrt{y_H / 0.652} \tag{2.7}$$

where the reference fraction 0.652 is the design loop hydrogen content, so the coverage term is unity at the design ratio. A single lumped group stands in for a Temkin-Pyzhev form with a diffusion correction [19], and it is calibrated so the first bed reaches 82.8 % of its equilibrium at the design inlet while beds two and three reach 98.2 and 98.1 %. Bed one converts $x_1 = 0.0263$, bed two 0.0101 and bed three 0.0050.

Quench mixes cold gas into the effluent between beds, half of the total fraction at each of the two points, and the converter extent is the flow weighted sum, as shown in Equation (2.8).

$$\begin{aligned} T^{+} &= \frac{T + \tfrac{q}{2} T_{cold}}{1 + \tfrac{q}{2}}, \qquad y_A^{+} = \frac{y_A^{out} + \tfrac{q}{2} y_A}{1 + \tfrac{q}{2}} \\ x_{tot} &= \frac{x_1 + \left(1 + \tfrac{q}{2}\right) x_2 + (1+q)\, x_3}{1+q} \end{aligned} \tag{2.8}$$

where q is the quench fraction of circulator throughput and the weights account for the quench gas that never saw the beds ahead of it. At the design line-up $x_{tot} = 0.03862$, the converter outlet carries 11.79 % ammonia against 3.16 % at the inlet, and the hot spot is 505.2 °C against a sintering limit of 530 °C.

2.3.4 The Converter Heats Its Own Feed

One algebraic equation gives this vessel two steady states, and it is Equation (2.9).

$$T_{in} = T_{cold} + \varepsilon\, c_{hx}\, (T_3 - T_{cold}) \tag{2.9}$$

where ε is the exchanger effectiveness and c_{hx} is a condition multiplier that represents a leaking bypass. At design, (2.9) lifts gas leaving the circulator at 40 °C to 419.1 °C. The operating point is a fixed point of the loop formed by (2.6) and (2.9) rather than a consequence of a valve position, which is the structure that produces an ignited branch and an extinguished branch [14, 15]. The beds carry thermal capacity so the gas answers immediately and the catalyst does not, as shown in Equation (2.10).

$$C_{cat} \frac{\mathrm{d}T_i}{\mathrm{d}t} = 60\, F\, c_{pg} \left(T_i^{out} - T_i\right), \qquad i = 1,2,3 \tag{2.10}$$

where C_{cat} is the catalyst and vessel heat capacity per bed in kJ/°C, c_{pg} is the molar heat capacity of the gas in kJ per kmol per °C and T_i^{out} is the algebraic bed outlet of (2.6). Both sides carry kJ per minute, and the resulting bed time constant is close to two minutes.

2.3.5 Separator, Purge and Makeup

The separator condenses ammonia to its vapor pressure with a departure factor, and its temperature follows the refrigeration demand through a first-order lag, as shown in Equations (2.11) and (2.12).

$$y_{A,sep} = k_{sep}\, \frac{\mathrm{e}^{A_s - B_s/(T_{sep} + 273.15)}}{P} \tag{2.11}$$

$$\tau_{sep} \frac{\mathrm{d}T_{sep}}{\mathrm{d}t} = -\frac{R_{duty}\, z_{ref}}{100} + 30 + \Delta T_{cool} - T_{sep} \tag{2.12}$$

where A_s and B_s are vapor pressure constants, k_{sep} is the departure from ideal separation at loop pressure, z_{ref} is the refrigeration demand in percent, R_{duty} is the refrigeration span and ΔT_{cool} is a shortfall disturbance. Equation (2.11) is clamped to the interval $[10^{-4}, 0.5]$. At –4.65 °C it leaves 3.79 % ammonia in the recycle, where an ideal Raoult calculation would predict about 2.4 %, and the ammonia going back round is the strongest single brake on an equilibrium-limited converter.

The material balance round the converter and separator gives the recycle and the product, as shown in Equation (2.13).

$$\begin{aligned} \Xi &= F x_{tot}, \qquad F_{out} = F - 2\Xi, \qquad y_A^{out} = \frac{F y_A + 2\Xi}{F_{out}} \\ L &= \frac{F_{out}\left(y_A^{out} - y_{A,sep}\right)}{1 - y_{A,sep}}, \qquad F_{gas} = F_{out} - L \\ F_{rec} &= \max\left(0,\, F_{gas} - \dot{n}_P\right), \qquad \dot{m} = 17.03\, L \end{aligned} \tag{2.13}$$

where Ξ is the nitrogen converted in kmol/s, F_{out} accounts for four moles of reactant becoming two of product, L is the liquid ammonia in kmol/s and $\dot{m}$ is the product in kg/s, reported on the scorecard as $86.4\,\dot{m}$ in t/d. Recycle compositions are the converter outlet corrected for the mole change and for the ammonia removed, as shown in Equation (2.14).

$$\begin{aligned} y_{I,rec} &= \frac{y_I}{1-2x_{tot}}\,\sigma, \qquad y_{H,rec} = \frac{y_H - 3x_{tot}}{1-2x_{tot}}\,\sigma \\ \sigma &= \frac{1-y_{A,sep}}{1-y_A^{out}} \end{aligned} \tag{2.14}$$

The purge is a square-root valve on the loop pressure and the makeup is a linear valve whose delivery falls as the loop it feeds gets harder to push into, as shown in Equation (2.15).

$$\begin{aligned} \dot{n}_P &= k_P \frac{z_P}{100}\sqrt{P} \\ M &= M_{max}\frac{z_M}{100}\,\alpha_{mk}\,\mathrm{clamp}\left(1 - 0.35\frac{P-P_{ref}}{P_{ref}},\, 0.3,\, 1.4\right) \end{aligned} \tag{2.15}$$

where k_P is the purge valve capacity in kmol/s per $\sqrt{\text{bar}}$, M_{max} is the makeup at a fully open valve and α_{mk} is the availability the front end is actually delivering. The pressure term in (2.15) is what gives the loop any self-regulation at all, because without it the only outlet that responds to pressure is the purge. Separating α_{mk} from z_M is what allows the front end to be taken away independently of the valve, and scenario 10 turns on exactly that distinction.

Each of the six actuators follows its command through a first-order lag, $\tau_j\,\mathrm{d}u_j/\mathrm{d}t = u_j^{cmd} - u_j$, with the makeup, purge, quench, speed, refrigeration and front-end trim constants listed in Table 2.2. Every state is clipped after each step, the pressure to 20–260 bar, the bed temperatures to 100–700 °C and the actuator positions to their travel.

2.3.6 Initialization and Numerics

Every scenario opens at a solved equilibrium. The application marches the derivative set from a fixed starting guess at a step of 0.25 min for 16 000 iterations, which is 67 hours of plant time, and successive substitution converges because every pass through the converter removes part of the reactant inventory. Evaluated at that line-up, the largest relative state derivative is 9.6×10^{-12}, which the application's own verification suite checks. The engine of Section 1.6.1 integrates by the classical fourth-order Runge-Kutta method at a fixed 0.05 min (3 s) step, the controller executes every 0.25 min (15 s), and the historian logs four samples per simulated minute. The exported file holds 104 columns, among them the true inert level beside the last analysis and each commanded valve position beside the actual one. The challenge scenario is the one exception to the solved start, because it is handed over lit and already drifting with a purge valve at 22 %, and solving that line-up would present a converter that is already out.

2.3.7 Assumptions

The model is detailed on the loop and thin on the vessel internals, and each simplification fails in a stated place.

Three lumped adiabatic beds. Each bed carries one temperature and one extent rather than an axial profile, so the hot spot is a bed outlet, while a real hot spot sits perhaps a third of the way into the bed and moves as the catalyst ages. Mixing at the quench points is instantaneous.

One kinetic group instead of a rate expression. The approach to equilibrium in (2.6) is a single Arrhenius term with a square-root hydrogen coverage, calibrated to a bed-one approach of 82.8 %. A Temkin-Pyzhev rate integrated down the bed with a diffusion correction is the honest form [19], and it carries a stronger ammonia inhibition than this one does.

Table 2.2 Model Parameters. Every Constant Above the Second Rule Is Shown in the Application and Adjustable by the User. The Six Below It Are Hidden.

Symbol	Meaning	Value	Unit
λ	adiabatic rise per unit extent	2958	°C
k_0	rate pre-exponential	1.46×10^{12}	—
E_a	apparent activation energy	165 000	J/mol
τ_b	contact time per bed	0.55	s
T_{max}	catalyst sintering limit	530	°C
T_{cold}	gas leaving the circulator	40	°C
ε	exchanger effectiveness	0.828	—
c_{pg}	loop gas heat capacity	35.5	J/mol/°C
V	loop gas volume	300	m^3
F_0	circulator rated throughput	6.00	kmol/s
P_{ref}	design loop pressure	150	bar
P_{max}	loop design pressure	172	bar
M_{max}	makeup at a full-open valve	1.70	kmol/s
k_P	purge valve capacity	0.01850	$\text{kmol s}^{-1}\text{bar}^{-1/2}$
$y_{I,mk}$	inerts in the makeup	0.0100	—
r_{mk}	makeup H_2:N_2	3.00	—
k_{sep}	separator departure factor	1.60	—
τ_{sep}	separator thermal lag	2.5	min
τ_{mk}	makeup valve travel	0.40	min
τ_{pg}	purge valve travel	0.25	min
τ_q	quench valve travel	0.30	min
τ_N	circulator speed response	0.50	min
τ_{ref}	refrigeration response	3.0	min
τ_{rmk}	front-end response to a trim	12	min
$y_{I,max}$	inert operating limit	0.16	—
W_{max}	circulator driver rating	12 500	kW
k_W	power per unit throughput	1650	kW per kmol/s
τ_{Tm}	thermocouple response	0.25	min
τ_{Pm}	pressure transmitter response	0.10	min
Δt_a	loop gas analysis cycle	15	min
—	ammonia value	0.42	$/kg
—	makeup syngas cost	2.50	$/kmol
—	electricity cost	0.075	$/kWh
Hidden constants, never displayed in the application			
C_{cat}	bed heat capacity	26 000	kJ/°C
z	loop compressibility	1.12	—
T_L	mean loop gas temperature	40	°C
A_s	vapor pressure constant	11.76	—
B_s	vapor pressure constant	2817	K
R_{duty}	refrigeration span	55	°C

Mole fractions rather than fugacities. Equation (2.4) is a fit to real gas data, but it is applied here to mole fractions. At 150 bar the fugacity coefficients depart from unity by tens of percent, and this is the largest single approximation in the model, so the equilibrium ammonia fractions are optimistic by a few points.

One well-mixed loop volume. Equation (2.3) places the whole inventory at one point. A real loop is a train of vessels with their own holdups and the analyzer looks at one place in it, so the true composition response carries a transport delay this model does not have.

Four components and one lumped inert. Argon and methane are combined into a single species that neither reacts nor condenses. The two behave differently in the separator, and methane in particular is

not entirely insoluble in liquid ammonia at loop conditions.

A separator that is a vapor pressure and a factor. Equation (2.11) is not a flash calculation, ammonia dissolved in the liquid drawn off is ignored, and k_{sep} is fitted rather than derived.

A circulator with no map. Equation (2.2) has no head curve, no surge line and no efficiency variation, so the machine cannot be operated into a region a real one would refuse.

Catalyst deactivation as a slider. Activity is one multiplier applied uniformly. Real deactivation is worst at the front of bed one, takes years, and changes the shape of the profile rather than only its height.

No front end and no hydrogen recovery. Everything upstream of the makeup valve is a supply and a ratio, and the purge goes to fuel rather than through the recovery unit most modern plants install [12]. The recovery unit is left out because it hides the economics the application is about.

2.4 What the Recycle Does

2.4.1 Two Small Flows Set the Steady Composition

Setting (2.3) to zero for the inert species and using the fact that the purge is drawn from the recycle stream gives the steady inert level directly, as shown in Equation (2.16).

$$y_{I,rec}^{\infty} = \frac{M\, y_{I,mk}}{\dot{n}_P}, \qquad \tau_I = \frac{N}{\dot{n}_P} \tag{2.16}$$

where τ_I is the inert time constant in seconds. This is Equation (1.4) of Chapter 1 written for the purge stream. Both are ratios of two small numbers. At the design line-up $M y_{I,mk}/\dot{n}_P = 1.0082 \times 0.01/0.08473 = 0.1190$, and the application reports 11.90 % inerts in the purge stream against 10.07 % at the converter inlet, the difference being the makeup dilution and the mole change through the beds. The time constant is $1543/0.08473 = 18\,209$ s, which is 303 min or 5.1 h. Cutting the purge valve from 37.4 to 20 % moves the inert level from 10.07 toward 15.34 % and covers 63 % of that distance in 5.12 h, so the measured response runs on the five-hour clock of (2.16). The smaller purge alone would set a slower clock, but the makeup falls with production as the inerts rise, and the smaller inert feed speeds the approach.

The same structure governs the hydrogen to nitrogen ratio, which is why both of the loop's slow variables are slow for the same reason. What the loop cannot consume in the proportion it is supplied has to leave through the same small purge, and the purge removes it at the loop concentration rather than at the concentration it arrived in.

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

2.4.2 Amplification, and Why It Is Not Symmetric

Figure 2.2 is the settled loop ratio against the ratio the front end delivers, measured by running the plant to steady state with the pressure controller holding 150 bar. A once-through plant would lie on the identity line. This one does not come close to it.

Defining the amplification as the relative loop error divided by the relative makeup error, a front end 2.0 % low is amplified 9.9 times, one 5.0 % low is amplified 8.0 times and one 10.0 % low is amplified 5.9 times. The gain falls as the error grows because a loop that is already nitrogen rich purges nitrogen faster, so the removal term catches up with the imbalance sooner. The curve is also not symmetric. The same relative errors high are amplified 13.5, 17.2 and 23.9 times, because hydrogen leaves through a purge less readily than nitrogen does once the loop is hydrogen rich. Production is a maximum at the stoichiometric ratio, 679 t/d at $r_{mk} = 3.00$, falling to 553 t/d at 10 % low and 516 t/d at 10 % high.

Section 1.4.3 derives the settled excess of one reactant after a feed ratio error, Equation (1.6), which is the mechanism behind this curve. The amplification takes as long as it is large. Following the loop ratio after the front end drifts from 3.00 to 2.85, the settled value is 1.804 and the trajectory covers 63 % of that change in 3.2 h, 95 % in 10.7 h and 97.8 % in 14 h. Production falls from 678 to

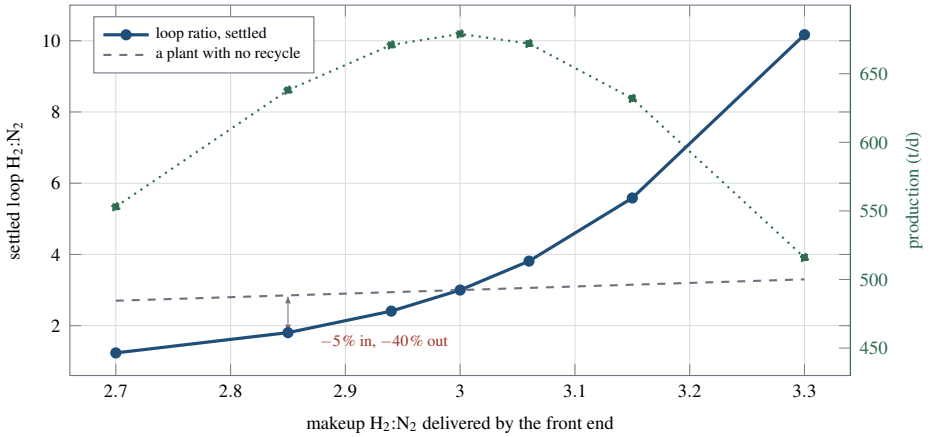

Figure 2.2 Recycle amplification of a front-end composition error, measured at steady state with the pressure controller holding 150 bar and the purge valve fixed at 37.4 %. The solid curve is the loop ratio, the dashed line is where a plant without a recycle would sit, and the dotted curve read against the right axis is ammonia production. Against the 679 t/d made at the stoichiometric ratio, a makeup ratio 5 % low settles the loop at 1.80 and costs 41 t/d, while the same error high settles it at 5.58 and costs 47 t/d.

638 t/d over the same period and the loop pressure never leaves 150.0 bar, because a controller is holding it. The whole disturbance therefore arrives as a production number that drifts down across half a day on a plant where nothing has alarmed.

> **TRY IT IN THE APP** **6 · A Small Error at the Front End**
>
> Load scenario 6 and press Run. From 30 min the front end delivers 2.85 to 1 instead of 3 to 1. Watch the loop H_2:N_2 trend head for 1.80, covering 63 % of the way in 3.2 h, while production falls from 678 toward 638 t/d and the pressure trend stays flat at 150 bar.

2.4.3 Which Variable Is the Inventory

At steady state the makeup entering the loop equals the reactant leaving it as ammonia plus the reactant leaving it through the purge, so the makeup valve is the throughput manipulator and there is no other candidate. Section 1.4.3 poses the same choice for the running example. The design decision is which variable it is allowed to hold, and the application offers both answers as a segmented control. With the pressure controller on the makeup valve, pressure is the controlled inventory and production is whatever holding it requires. With the operator on the makeup valve, production is set and pressure is whatever the converter needs in order to consume the gas, with the pressure controller retained behind a low selector as a high limit. Each arrangement is consistent in the sense that the loop inventory is regulated by a flow entering it [4], and each is used in practice [3, 9].

The two differ in where a disturbance appears, which Section 2.6 measures. What matters for the operator is that only one of the two failure modes announces itself. A plant losing catalyst activity with the pressure controller in charge quietly makes less ammonia, and the same plant with the operator in charge quietly consumes the margin to its relief valves.

Pause. Try Exercise 2.1 before reading on.

Table 2.3 The Four Loops as Implemented, With the Default Tunings Every Automatic Scenario Loads

CV	MV	K_c	τ_I (min)	Action	Plant
P	z_M	2.4	9	reverse	minutes
T_{hot}	q	0.14	8	direct	minutes
y_I	z_P	2.5	300	direct	5.1 h
r	r_{mk}	0.09	260	reverse	3.2 h

2.5 Controller Design

2.5.1 The Implemented PI Algorithm

All four loops use the shared velocity-form law of Section 1.6.2, Equations (1.14) to (1.17), executed every $\Delta t_c = 0.25$ min. The action is reverse on a loop where a rising measurement calls for a falling output and direct where it calls for a rising one, and the measurement passes through a first-order filter advanced by its exact zero-order-hold recursion. The clamp acts on the emitted output only, back-calculation unwinds the integral with a tracking constant T_t that defaults to τ_I, and Manual operation solves the integral so that the output equals the operator's valve position, which makes transfer in both directions bumpless. All four loops run with $\tau_D = 0$, so the derivative path is present in the implementation and inactive in every shipped tuning. No loop carries an output rate limit, because the actuator lags of Table 2.2 already limit how fast a valve moves.

2.5.2 Four Loops on Two Time Scales

Table 2.3 lists the loops, their default tunings, their direction and the plant each has in front of it. The pressure and hot-spot loops read continuous transmitters filtered at 0.10 and 0.25 min. The inert and ratio loops read one chromatograph on the loop that reports both numbers every 15 min and holds them between analyses, so both slow loops move in steps. The sample period is not the difficulty on those two loops. A 15 min analysis on a 5.1 h plant is twenty times faster than the process it measures, and the difficulty is the process.

Gains carry the units of their pairing, 2.4 % of makeup valve per bar, 0.14 % of quench per °C, 2.5 % of purge valve per percent of inerts, and 0.09 of makeup ratio per unit of loop ratio. Output limits are 0 to 100 % on the makeup and purge valves, 0 to 45 % on the quench, and 2.4 to 3.6 on the front-end trim. The two reset times in the hundreds of minutes are conservative by a wide margin, and scenario 9 measures how wide. Running the doubled makeup inerts of scenario 9 with the shipped $K_c = 2.5$ and $\tau_I = 300$ min leaves the inert level at 10.52 % after 48 h and still approaching its 10 % setpoint, while $K_c = 10$ with $\tau_I = 60$ min first reaches 10.00 % 6.3 h after the change and moves the purge valve through 79 percentage points of travel against 55. Neither tuning oscillates, and even the slider maxima of $K_c = 30$ and $\tau_I = 20$ min settle within eight hours with no sustained cycle, after one undershoot to 9.81 %, 12 min with the purge valve wide open and 125 points of travel. The reason is in the model rather than in the tuning rule. The inert level is a single lag of 5.1 h behind a 15 min sample hold, so the loop has almost no phase lag beyond first order and gain alone cannot destabilize it. What limits the purge controller in practice is the cost of overshooting, because every excess percent of valve position is hydrogen leaving through the stack for hours before the analyzer reports that it was not needed.

2.5.3 Four Structures and One Selector

The application offers manual operation, a basic-loops structure in which the pressure and hot-spot controllers are in service and nobody is watching the composition, a supervised structure that adds the inert loop and, on a separate selection, the ratio loop, and a predictive structure in which one plan replaces the two slow loops (Section 2.5.4). The basic-loops setting is where most of the scenarios

live, because a plant with the fast loops closed and the slow ones open works perfectly for a shift and poorly for a month.

When the operator owns the makeup valve, the pressure controller is retained as an override against the high limit rather than against the normal setpoint, and a low selector passes whichever signal asks for less gas, as shown in Equation (2.17).

$$z_M = \min\left(z_M^{op},\ \mathrm{PI}\left(P_{hi}, \hat{P}\right)\right) \tag{2.17}$$

where z_M^{op} is the operator demand and $P_{hi} = 165\,\mathrm{bar}$ is the high limit against a design pressure of 172 bar. Whenever the operator demand is selected, the pressure controller is tracked to the applied output through the same expression used for manual operation, so it takes over without a jump. Pointing the override at the normal setpoint instead of the high limit is the mistake worth naming, because it makes the production structure into a slower version of the pressure structure. The override then retakes the valve the moment the pressure reaches the number the other structure was already holding, and the two arrangements become indistinguishable.

2.5.4 Model Predictive Control

Section 1.6.3 sets out the predictive option of the framework in general, and this subsection gives what is specific to the loop. Under the predictive structure the pressure and hot-spot loops run exactly as in Table 2.3, and a quadratic program chooses the purge valve z_P and the front-end trim r_{mk} together [20, 21]. A plan is made every $T_s = 15\,\mathrm{min}$, on each new loop analysis at the declared cycle, and its filter update reads the last analysis reported. The horizon is $N_h = 48$ periods (12 h, adjustable from 6 to 24 h), grouped into twelve blocks of $\{1,1,1,1,2,2,3,4,6,9,15,3\}$ periods. Every decision is expressed as a change from the commands in force, and the first block alone is sent to the valves and kept until the next analysis.

The plan predicts with the equations of Section 2.3, evaluated at the declared parameter values of Table 2.2 and the nominal disturbances of Table 2.1, closed around the pressure and hot-spot PI loops at their shipped tuning with the 0.10 and 0.25 min transmitter lags, and linearized at the design line-up by central differences. Its eighteen states are the fourteen of the plant, the two transmitter filters and the two loop integrals, and an exact zero-order hold over T_s gives Equation (2.18).

$$\begin{aligned} \tilde{x}_{k+1} &= A_d \tilde{x}_k + B_d\left(\tilde{u}_k + d_k\right) + E_d \tilde{w}_k, \qquad d_{k+1} = d_k \\ \tilde{y}_k &= C\tilde{x}_k + D_w \tilde{w}_k \end{aligned} \tag{2.18}$$

where tildes denote departures from the design point, $\tilde{u} = (z_P, r_{mk})$ holds the decided inputs, d one unmeasured disturbance per decided input, $\tilde{w}$ the known inputs held over the horizon (the pressure and hot-spot setpoints, the operator's makeup demand, the circulator speed and the refrigeration demand), and $\tilde{y} = (y_I, r, z_M)$ the loop inerts in percent, the loop ratio and the makeup valve command, which alone reads a known input directly through D_w. Two versions are built, one for each owner of the makeup valve. Step-response fits were not used, because a step made in manual opens the two fast loops that stay in service under the plan.

The settled gains are -0.194 percent of inerts per percent of purge (-0.244 with the operator on the makeup valve) and 11.48 of loop ratio per unit of front-end ratio. With the pressure loop on the makeup valve the inerts cover 45.5, 78.4 and 97.5 % of a purge move's settled effect after 2, 5 and 12 h, close to a 3.3 h time constant. Against the engine's mean response to ± 1 percent of purge and ± 0.01 of trim from the line-up of scenario 4 the model errs by at most 0.26 % of the largest response between 15 and 720 min, while a single 5 percent opening lowers the inerts by 0.866 % in 12 h against the model's 0.945 %, the nonlinearity the model leaves out.

A steady-state Kalman filter estimates $\tilde{x}$ and d at every decision under every structure, so a switch to the plan starts warm, and it reads only what the slow loops read, the analyzed y_I and r, with assumed errors $\delta_I = 0.01\,\%$ and $\delta_r = 0.002$. Each disturbance carries process noise $\varphi\,\delta_m^2 / G_m^2$, with $\varphi = 10^4$ and G_m the settled gain of its input on its own analysis, which makes the prediction offset

free [22]. The makeup valve's offset from the model, b_M, is read from the pressure loop's command at each decision. Replacing the filter by a constant output bias, on a diagnostic copy that is not part of the application, raises the inert IAE of scenario 12 from 56.6 to 539.3 % · min and its ratio IAE from 8.5 to 68.5 min, because a bias corrects a drift only as fast as the five-hour plant reveals it.

Each plan minimizes Equation (2.19), written in dollars per decision period.

$$\begin{aligned} J = & \sum_{k=1}^{N_h} \left[q_r \left(\frac{r_k - r_{\mathrm{SP}}}{0.05} \right)^2 + c_P z_{P,k} \right] \\ & + 50 \lambda_u \sum_{j=1}^{12} \left[\left(\frac{\Delta z_{P,j}}{5} \right)^2 + \left(\frac{\Delta r_{mk,j}}{0.05} \right)^2 \right] \\ & + \sum_{i \in \{I,M\}} \left(\rho_i s_i^2 + \mu_i s_i \right) \end{aligned} \tag{2.19}$$

Here $q_r = 20$, $\lambda_u = 5$, Δ is the move of block j, and $c_P = \$5.10$ per percent of purge prices the synthesis gas the valve vents in one period at the declared \$2.50 per kmol and the design pressure, so the plan purges only as much as the inert ceiling demands. The slacks s_I and s_M of the soft rows of Equation (2.20) cost $\rho_I = 10^5$ per square percent and $\mu_I = 10^4$ per percent of inerts, and $\rho_M = \mu_M = 10^4$ for the makeup valve, so a reachable limit costs nothing.

$$\begin{aligned} & 0 \le z_{P,k} \le 100, \qquad 2.4 \le r_{mk,k} \le 3.6 \\ & y_{I,k} \le y_{I,\mathrm{SP}} + \max(0,\, \hat{y}_{I,0} - y_{I,\mathrm{SP}})\, \mathrm{e}^{-kT_s/\tau_f} + s_I \\ & z_{M,k} + b_M \le 98 + s_M, \qquad s_I,\, s_M \ge 0 \end{aligned} \tag{2.20}$$

where $\hat{y}_{I,0}$ is the filter's present inert level and $\tau_f = 120\,\mathrm{min}$ lets a level already above its setpoint come back along a lag instead of pricing the whole horizon. The input rows are hard and are the limits the slow loops clamp to. The makeup row and a purge hold act only with the pressure controller on the makeup valve, and while the pressure reads more than 2 bar under its setpoint the plan may close the purge but not open it. The model has the pressure loop on its setpoint, and a loop being drained concentrates its inerts, which the filter would otherwise take for purge that has lost its worth. The compressor limit is a soft row rather than a hard constraint, because the makeup valve is an output of the pressure loop, so a hard row on it would be infeasible in every front-end cut. The only hard rows are the input bounds, which always contain the present commands, so a plan cannot be infeasible.

If the solver returns a plan that is suboptimal, nonfinite or beyond the valve limits, the supervised inert and ratio loops, which have followed the applied commands throughout, take over both handles bumplessly and give them back at the next acceptable plan. Nothing in the plan comes from the run's parameters or from the plant's states. Its inputs are the analysis, the pressure transmitter, the makeup command the pressure loop sends, the setpoints, the operator's settings and the commands applied, so an edit a learner makes to the plant is, for the plan, model error that the filter absorbs into its disturbance estimates. A dual active-set method [23] solves a plan in a mean of 0.77 ms in a browser at the 12 h horizon and 1.19 ms at 24 h over three runs, under 0.01 percent of one core at real time.

2.6 Results

Thirteen scenarios are arranged from the converter outward, listed in Table 2.4. Each fixes the starting mode, the control structure, the four tunings, the disturbance settings and any timed events, and each carries the briefing and the completion rule of Section 1.6.4. The exercises at the end of the chapter run several of them.

2.6.1 The Converter Has an Optimum and a Cliff

Scenario 2 runs the quench by hand. Closing it from 10 to 4 % takes the beds from 497/505/498 to 528/528/525 °C an hour later and the ammonia leaving bed three falls from 11.79 to 10.33 %.

Table 2.4 Scenario Sequence and the Mechanism Each One Exercises

#	Scenario	What it exercises
1	Four valves by hand	two actuators answer in minutes and two in hours
2	The converter has a best temperature	rate against equilibrium on one vessel
3	Blow it out and get it back	two steady states from (2.9)
4	Pressure and hot spot on automatic	a held variable stops carrying information
5	Which variable is the inventory	throughput manipulator pairing
6	A small error at the front end	recycle amplification over 14 hours
7	Inerts and the five hours they take	the time constant of (2.16)
8	What a purge valve is buying	conversion against hydrogen loss
9	The makeup gets dirty	a controller that converts one cost into another
10	The front end cuts back	self-regulation until an actuator saturates
11	Five years of catalyst	constraint margin spent by the first disturbance
12	Five hours ahead: Supervised or MPC	one plan against two PI loops on two slow drifts
13	A day on the loop (challenge)	structure first, then the slow loops early

Opening it to 14 % runs the beds at 471/485/479 °C after the next hour and the outlet rises to 12.55 %. Hotter is faster and worse, because the equilibrium curve of (2.4) has come down further than the rate has climbed.

Scenario 4 puts the pressure controller on the makeup valve and the hot-spot controller on the quench and then moves the temperature setpoint. Stepping the hot spot from 505 to 498 °C half an hour into the run takes production from 678 to 726 t/d two hours later, an improvement of 7.0 %, while the loop pressure stays within 0.4 bar of 150 bar. Waiting a day gives 719 t/d, because the extra production draws more makeup and the inert level rises from 10.07 to 10.52 % at a fixed purge valve, which is the two time scales appearing inside a single step response. Continuing downward in five-degree steps, each setpoint held from the start of its own run and read at 1440 min, gives 738, 748, 753 and 754 t/d at 493, 488, 483 and 478 °C, with the quench valve opening from 10.10 to 16.26 %, so the converter's own optimum is near 480 °C and worth 11 % of production.

Nobody operates there, and scenario 3 is the reason. Opening the quench to 22 % for seventy minutes and then returning it to exactly 10 % does not relight the converter, which goes out 6.6 min after the valve opens. The inlet temperature falls from 419 to 90 °C and stays there, while the loop pressure runs from 150 to 214 bar within twenty minutes of the opening, and to the 260 bar ceiling of the model by 51 min, because the makeup gas is still arriving and nothing is consuming it. Held from the design state in manual, a quench of 15.3 % keeps the converter lit and 15.4 % puts it out, so past that point the converter is on the wrong side of the ignition condition, and the recovery is a startup heater rather than a valve position.

2.6.2 Where the Disturbance Lands

Scenario 5 runs the same plant twice. With the pressure controller on the makeup valve the loop sits at 149.99 bar and makes 678 t/d. Handing the valve to the operator and asking for 66 % takes production to 749 t/d and lets the pressure float to 155.5 bar, which is what the converter needs in order to consume the extra gas. Dropping the catalyst activity to 0.75 then separates the two arrangements. Production holds at 744 t/d and the pressure carries the loss instead, reaching 157.6 bar two hours after the change and settling at 158.3 bar, 6.7 bar inside the 165 bar high limit at which (2.17) would take the valve back. Under the other pairing the same disturbance appears as a production number and nothing else, a loss of 17 t/d two hundred minutes after the change with the pressure on 150 bar.

Scenario 10 is the self-regulation the arrangement buys. Removing 45 % of the makeup gas availability with the pressure controller in charge drives the makeup valve from 59.4 to 100.0 % within ten minutes, where it stops being a controller. The pressure falls to 140.6 bar eight minutes after the cut, recovers only to 146 bar, and production settles at 631 t/d against 679 t/d. Nothing trips

Table 2.5 Measured Outcomes of the Composition and Constraint Scenarios. Every Reading States the Time Since the Disturbance at Which It Was Taken, Because a Loop Composition Is Still Moving Hours After the Disturbance.

#	Reading	y_I (%)	r	Rate (t/d)
—	design line-up	10.07	3.000	677
6	front end 3.00 to 2.85, after 4.25 h	10.01	2.128	659
6	the same, after 13.5 h	9.73	1.834	639
6	the same, settled after 24 h	9.65	1.807	638
7	purge 37.4 to 20 %, after 0.5 h	10.55	3.000	675
7	the same, after 3 h	12.40	3.000	629
7	the same, after 5.5 h	13.54	3.000	598
7	the same, after 9.5 h	14.51	3.000	572
8	purge shut to 8 %, after 11.2 h	20.81	3.000	471
8	purge open to 85 %, after 11.7 h	5.44	3.000	791
9	makeup inerts doubled, after 24 h	11.38	3.000	647
9	the same, after 72 h	10.21	3.000	677
11	activity to 0.65, after 4.5 h	9.88	3.000	650
11	exchanger −8 %, 20 min later	8.92	3.000	0

and nothing alarms. The hot spot dips 6.9 °C below setpoint as the pressure falls and is back within 1 °C twenty minutes later, and it overshoots by 9.5 °C when the gas returns at 400 min and recovers as quickly, because a recycle absorbs a large feed disturbance by changing its own inventory. The useful part of that behavior ends exactly where the valve travel does.

2.6.3 Slow Composition and the Cost of a Purge

Table 2.5 collects the measured outcomes of the scenarios that turn on composition, each with the observation window it was read over. Scenario 7 halves the purge and is the clearest of them. Thirty minutes after the cut the inert level has moved from 10.07 to 10.55 %, which on a trend is indistinguishable from noise, and three hours after the cut it is 12.4 %. Production falls from 678 to 629, then 598, then 572 t/d while the pressure stays within 0.6 bar of 150 bar and the circulator power, after a rise of 0.04 MW in the first twenty minutes, stays within 0.01 MW of 9.90 MW. Sixteen percent of the plant output has gone and the only instrument that reports it is the analyzer.

Scenario 8 takes both extremes of the purge valve. Nearly shut at 8 %, the inert level reaches 20.81 % after 11.2 h and production falls to 471 t/d, so a fifth of the loop is gas that does nothing but occupy partial pressure and get compressed, and the saving is 72 kg/h of hydrogen against 389 kg/h at design. Wide open at 85 %, the loop is clean at 5.44 % inerts and makes 791 t/d, and the stack is carrying 939 kg/h of hydrogen while the makeup valve has gone from 59.4 to 75.0 %. That is 26 % more syngas bought for 17 % more ammonia.

Table 2.6 sweeps the valve properly, with each point run to a settled state after moving the valve five minutes into a design-start run. This protocol follows the lit converter branch; a cold start at the lowest purge can instead converge to an extinguished converter. Read down the table the trade is visible in two places at once. Syngas per tonne of product rises monotonically with purge, from 123.7 to 139.3 kmol/t, so feedstock efficiency wants the valve shut, while production rises monotonically from 484 to 792 t/d, so plant output wants it open.

Which of the two wins depends on the prices and pairing, a consequence of Section 2.4. With pressure control on makeup and the illustrative syngas price of $2.50 per kmol, the settled margin peaks near 42 % purge at $2 049 per hour, compared with $2 041 per hour at the design 37.4 % and $1 628 per hour at 85 %. Sampled valve positions from 37.5 to 47.5 % remain within 1 % of the

Table 2.6 Purge Valve Sweep on the Lit Branch, Each Move Made After a Design Start and Settled Over 96 h With the Pressure Controller Holding 150 bar. Margin, Quoted in the Text, Is the Application's Own Metric, Ammonia Sold Less Makeup Gas and Circulator Power, Evaluated at the Settled State.

z_P (%)	y_I (%)	Rate (t/d)	z_M (%)	H_2 purged (kg/h)	Syngas (kmol/t)
15	18.15	484	40.8	140	123.7
25	13.32	599	51.4	249	125.9
37.4	10.08	679	59.4	389	128.6
42	9.27	699	61.7	441	129.7
55	7.58	740	66.8	591	132.5
85	5.43	792	75.1	939	139.3

sampled best. Both extremes cost money: dilution loses production at low purge, while the additional syngas bought to hold pressure costs more than its additional ammonia earns at high purge. The flat optimum explains why a valve setting can look correct on every screen and still lose money over a year. The application's price slider extends to \$5 per kmol; reducing it to \$0.145 per kmol instead makes the margin rise as the purge opens. This optimum belongs to the stated prices and pairing, rather than to a universal purge setting.

> ✦ **DID YOU KNOW?** **Venting More Gas Makes More Ammonia**
>
> From the simulator: opening the purge valve from 15 to 85 % sends 939 kg/h of hydrogen out of the loop instead of 140 kg/h, yet production rises from 484 to 792 t/d, because the loop then carries 5.43 % inerts instead of 18.15 % (Table 2.6).

Scenario 9 supplies the cleanest argument for treating the inert setpoint as an economic decision. Doubling the methane slipping through into the makeup with the inert controller in service takes the purge valve from 37.6 to 69.9 % in the twenty-four hours after the change and to 81.4 % after seventy-two, still opening, while the hydrogen leaving through it rises from 391 to 845 kg/h and the syngas consumed per tonne rises from 128.7 to 141.9 kmol/t for 677 t/d of product against 678 before the change. The controller does the right thing and the plant is worse off anyway, because an inert controller does not solve an inert problem. It converts an upstream problem into a hydrogen bill, and the answer is at the reformer.

2.6.4 Margin, Measured Before It Is Spent

Scenario 11 ages the catalyst to 0.65 and the plant absorbs it. The beds cool, the hot-spot controller closes the quench from 10.08 to 8.45 % to put the temperature back on 505 °C, and production settles at 650 t/d. Every indication on the display is normal. Four and a half hours later the feed and effluent exchanger loses 8 % of its effectiveness, and the quench controller now needs more travel than the 8.45 % it has. The valve is shut four minutes after the event, T_{in} falls, and the converter is out six minutes after it, with the loop pressure passing 160 bar two minutes later. Fresh catalyst survives the same exchanger loss, with the quench driven shut and the inlet holding at 392 °C, which is the comparison the scenario asks for after a Reset.

The first disturbance did not damage the plant. It removed the plant's ability to survive the second one, and no controlled variable said so, because the hot spot was on setpoint at every instant until it was not. Valve position is therefore the trend worth its own strip on this plant, and the application reports the quench command beside the hot spot for that reason.

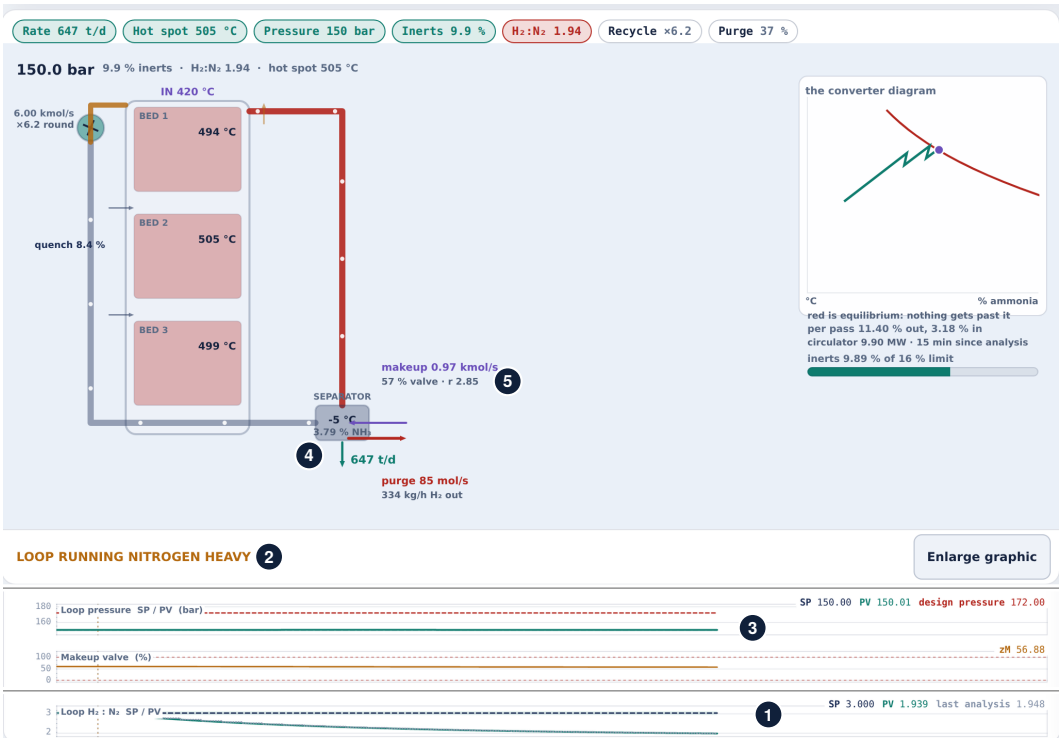

Figure 2.3 Scenario 6 at $t = 480$ min on the 12 h trend window, seven and a half hours after the front-end ratio drifted from 3.00 to 2.85, with the pressure and hot-spot loops in service and the composition loops open, cropped to the badges, the plant view and the loop pressure, makeup valve and loop H_2:N_2 strips of the trends. (1) The loop ratio has fallen to 1.939 against a last analysis of 1.948 that is 14.75 min old, and it is still moving toward 1.80, which is what the status line (2) calls a loop running nitrogen heavy. The point of the view is the pressure trend (3) read against production (4). Loop pressure is 150.01 bar against a 150 bar setpoint and has stayed within 0.04 bar of it since the drift, while production has drifted from 678 to 647 t/d and the makeup valve (5) has closed from 59.4 to 56.9 % to keep the pressure there. The hot spot is 504.9 °C on a 505 °C setpoint, the quench sits at 8.41 %, the loop carries 9.89 % inerts, and the recycle badge reads 6.2, up from 5.94 before the drift, because the same circulation is now serving less makeup.

> **TRY IT IN THE APP** **11 · Five Years of Catalyst**
>
> Load scenario 11 and press Run. At 30 min the catalyst activity falls to 0.65 and the quench closes from 10.08 to 8.45 % to hold the hot spot. Watch the quench strip at 300 min, when the exchanger loses 8 %: the valve is shut four minutes later, and the converter is out six minutes after the event.

2.6.5 The Challenge

The graded scenario hands over a fifteen-hour shift with the purge valve already at 22 %, makeup inerts 30 % above specification and rising to 80 % above at two hours, a front end running 0.06 low on ratio, catalyst at 0.90, and a reformer that removes 45 % of the makeup availability from the fifth hour to the eighth and a half. The case ships with weak tunings on the basic loops, and with those the converter goes out at 302.7 min, 2.7 min into the front-end cut. That run fails, with a worst pressure error of 3.15 bar over the last five hours, 7.7 % inerts and a ratio of 2.70 at the end. The target, judged at 900 min on the case as given, is the loop pressure within ±3 bar of 150 over the last five hours, the

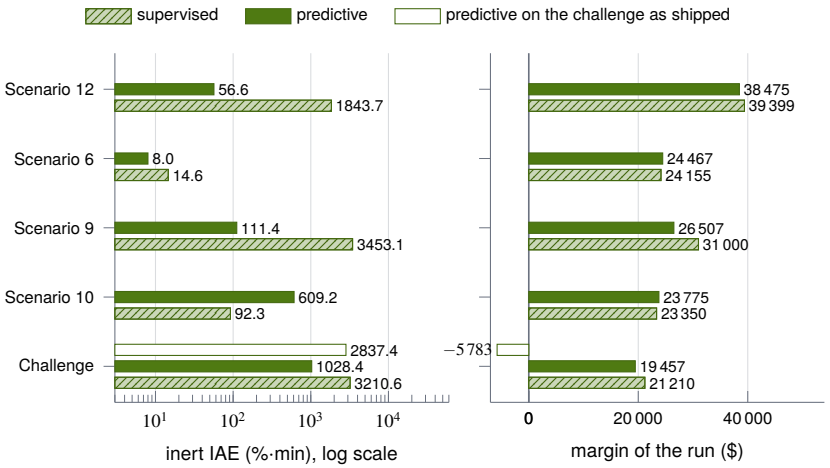

Figure 2.4 Table 2.7 as bars, digit for digit: the inert IAE on a logarithmic axis (left) and the margin of each run (right), with the supervised loops hatched, the predictive controller solid and the predictive controller on the challenge as shipped open. Notice that the plan cuts the inert error in every case but scenario 10, yet earns less margin than the supervised loops on scenario 12, scenario 9 and the challenge, and that on the challenge as shipped it ends with a negative margin.

hot spot under 530 °C with the converter alight all day, and at the end inerts under 13 % and the ratio within 0.15 of three. The suggested configuration, the default tunings of Table 2.3 with both slow loops in service from the start and the purge and quench stations left where the shift found them, meets it with a worst pressure error of 0.04 bar, a peak hot spot of 512.7 °C, 12.47 % inerts, 0.53 point inside the limit, and a ratio of 2.99, and the application verifies that on itself.

The answer is a structure and a schedule rather than a tuning constant. On the default tunings the two fast loops hold the plant for the whole shift, alight and within 0.004 bar of setpoint over the last five hours, but on their own they let the inerts cross the 16 % operating limit at 3.1 h and finish at 21.0 % with the ratio at 2.22. The inert loop and the ratio loop decide where the plant finishes, and both need three time constants to get there. With the fast loops tuned from the start, putting the slow loops in service at twenty minutes finishes inside the target at 12.89 % inerts and a ratio of 2.93, waiting two hours finishes at 13.21 % and fails, and waiting six hours finishes at 14.73 %. The front-end cut in the middle is a distraction and treating it as one is correct, because the loop absorbs it by depressuring to 141.7 bar and recovers on its own before the scored window opens, while a purge valve moved for a three-hour disturbance is still moving five hours after the disturbance has gone. Figure 2.3 shows the application during scenario 6, which is the same lesson without the clock.

2.6.6 Two Slow Loops Against One Plan

Table 2.7 runs the supervised structure, with its shipped tuning, the ratio loop on the analyzer and the suggested settings on the challenge, against the predictive structure with the settings of scenario 12 on every case and only the structure changed. Both face the same events, analyzer, transmitters and duration, the application has no measurement noise, and neither receives a preview. Figure 2.4 draws the table’s inert error and margin columns side by side.

Scenario 12 starts on the design line-up and at 60 min raises the methane slipping into the makeup by half and turns the front end 0.10 nitrogen heavy, two drifts that each take hours to reach the loop. On the first analysis that shows them, at 75 min, the plan opens the purge from 37.9 to 56.2 %, reaching 59.5 by 120 min, while it takes the trim to 3.12 by 120 min and eases it back to the 3.10 the drift needs by the end of the day (Figure 2.5). The inerts peak at 10.36 % at 81.5 min and stay

Table 2.7 Supervised Loops Against Model Predictive Control

Case (run length)	Structure	IAE y_I (% · min)	Peak, end y_I (%)	Purge travel (%)	IAE r (min)	Trim travel	IAE P (bar · min)	H_2 purged (t)	NH_3 (t)	Margin ($)	Plans, mean (µs)
Scenario 12 (1440 min)	supervised	1843.7	11.98, 10.73	22.8	269.95	0.157	26.0	11.74	648.6	39 399	—
	predictive	56.6	10.36, 10.00	22.0	8.47	0.147	13.3	14.58	680.2	38 475	96, 529
Scenario 6 (720 min)	supervised	14.6	10.07, 10.01	0.3	321.60	0.185	5.0	4.45	337.0	24 155	—
	predictive	8.0	10.07, 10.00	0.3	13.50	0.243	1.4	4.72	340.1	24 467	48, 525
Scenario 9 (1440 min)	supervised	3453.1	13.65, 11.42	42.5	0.00	0.000	43.4	14.28	623.3	31 000	—
	predictive	111.4	10.70, 10.00	46.0	0.03	0.000	26.5	20.44	679.9	26 507	96, 764
Scenario 10 (720 min)	supervised	92.3	10.64, 10.07	8.3	0.00	0.000	1680.3	4.61	328.5	23 350	—
	predictive	609.2	11.41, 9.64	117.8	0.08	0.002	708.8	4.56	331.9	23 775	48, 502
Challenge (900 min)	supervised, suggested	3210.6	15.11, 12.47	40.7	164.31	0.111	370.7	6.59	368.9	21 210	—
	predictive, suggested	1028.4	13.10, 10.10	183.4	39.34	0.285	862.8	9.70	400.5	19 457	60, 489
	predictive, as shipped	2837.4	13.10, 5.59	122.8	37.52	0.274	33 197.0	4.81	150.7	−5 783	60, 455

Errors and travel accumulate from $t = 0$ as on the scorecard, end values are read at the run length, no plan was rejected, and solve times are wall-clock time in Node. The last row switches only the structure on the challenge as shipped.

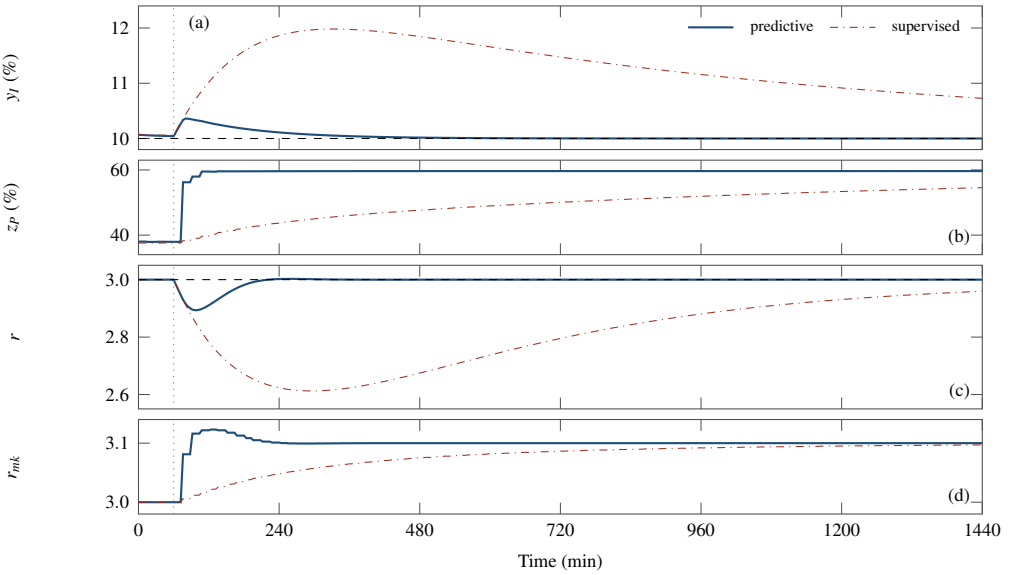

Figure 2.5 Scenario 12 over its 1440 min day under the predictive structure (thick solid lines) and the supervised structure (dash-dotted lines), with methane slip raised by half and the front end turned 0.10 nitrogen heavy at 60 min (dotted). (a) True loop inerts against the 10 % setpoint (dashed). (b) Purge valve command. (c) True loop H_2:N_2 against its setpoint of 3 (dashed). (d) Front-end trim command.

within 0.1 % of the setpoint from 252 min, and the ratio, lowest at 2.89, stays within 0.02 of 3 from 178.5 min. Under the supervised loops the purge creeps from 37.6 to 54.6 % over the day, the inerts peak at 11.98 % at 333 min and still read 10.73 % at 1440 min, and the ratio falls to 2.61. The plan cuts the inert IAE from 1843.7 to 56.6 % · min and the ratio IAE from 269.95 to 8.47 min with less travel, 22.0 against 22.8 percent of purge and 0.147 against 0.157 of trim, because it moves each handle sooner rather than further. The bill is the purge. Holding 10 % on dirtier gas vents 14.58 against 11.74 t of hydrogen, and the day's margin is $924 lower. Neither the makeup ceiling nor the purge hold acts on this day, with the makeup valve at most at 62.7 % and the pressure never below 149.4 bar, and switching the ceiling off changes no plant variable.

Scenario 6, the front end alone 0.15 nitrogen heavy, is a clear win on the ratio, with a ratio IAE of 13.50 against 321.60 min and a largest error of 0.163 against 0.596, for 0.243 against 0.185 of trim

travel and \$312 more margin.

Scenario 9 is a loss in money. The plan holds the doubled makeup inerts at 10.00 %, against 11.42 % under supervision at 1440 min, with an inert IAE of 111.4 against 3453.1 % · min and 56.6 t more ammonia, but purging 16 039 against 11 547 kmol of gas costs \$4 493 of margin. The loss is the price of the 10 % setpoint rather than of the controller. Supervised with a fast inert loop, $K_c = 10$ and $\tau_I = 60$ min, which lies outside the shipped tuning, also ends the day at 10.00 % and earns \$26 753, \$246 more than the plan, with 79.1 against 46.0 percent of purge travel, and the shipped loop is cheaper only because it is still 0.52 % above its setpoint after two days.

Scenario 10 is a loss in inerts. The plan holds the pressure better, with an IAE of 708.8 against 1680.3 bar · min and \$425 more margin, and the lowest pressure is the same 140.6 bar at 48.25 min under both, before any plan has closed the purge. Its inert IAE is 609.2 against 92.3 % · min and its purge travel 117.8 against 8.3 percent, because while the makeup valve is at its stop, from 50.25 to 400.25 min, the makeup ceiling and the hold close the purge, to 20.2 % at 90 min, trading inerts, which can wait five hours, for pressure, which cannot. A replay outside the application, through the basic loops, of the supervised purge with the plan's purge substituted only while the makeup is at its stop gives an inert IAE of 645.1 and a pressure IAE of 715.0, so the closure during the cut is the whole trade.

The graded challenge keeps its scoring and is given here for information only. Switching only the structure does not meet it. The plan moves only the purge and the trim, and the weak fast tunings still lose the converter in the front-end cut, at 334.6 min against 302.7 min for the shipped line-up. The loop then spends 548.4 min above the 172 bar design pressure, peaking at 228.8 bar. With the suggested settings the plan meets the target with a worst pressure error of 0.06 bar, 10.1 % inerts and a ratio of 3.00, and it breaks no unscored limit, with a peak pressure of 160.9 bar and no minute over temperature, over the design pressure or over 16 % inerts. Its pressure IAE is 862.8 against 370.7 bar · min and its margin \$1 753 lower, because inerts held near 10 % mean more purge and more makeup during the cut.

2.7 Technology in Practice

What Measures It

Figure 2.6 draws the loops of Section 2.5 with their devices. Loop pressure and the hot spot are read continuously. The two slow variables come from one chromatograph on the loop gas, which separates the species of a sample in time, so its cycle takes minutes and some dead time is inevitable [24]. GC-101 reports every 15 min and holds its last result in between, and the model's single well-mixed loop leaves out the transport delay of a real sample line (Section 2.8).

> **TECHNOLOGY** **Sensor: The Loop Gas Chromatograph**
>
> A process chromatograph injects a small fixed sample into a long column, where each species travels at its own speed, so the detector at the end need only tell the carrier gas from anything else. A thermal conductivity detector senses the passing gas by the cooling of a heated element, and the other common type, the flame ionization detector, cannot see ammonia. Most process units recalibrate themselves from bottles of calibration gas on a schedule. The sample valve matters most for accuracy: if the injected amount is not repeatable, the readings change from cycle to cycle while the gas does not, and a leaking valve shifts the baseline or corrupts the analysis [24].

What Moves It

Four handles move the loop: the makeup valve, the purge valve, the quench valve and the trim sent to the front end, which sets the ratio of the gas the reformer delivers. The synthesis gas is raised to loop pressure by centrifugal compressors, usually driven by steam turbines on steam raised from the plant's surplus heat, and the catalyst is divided into layers cooled between them, either directly by

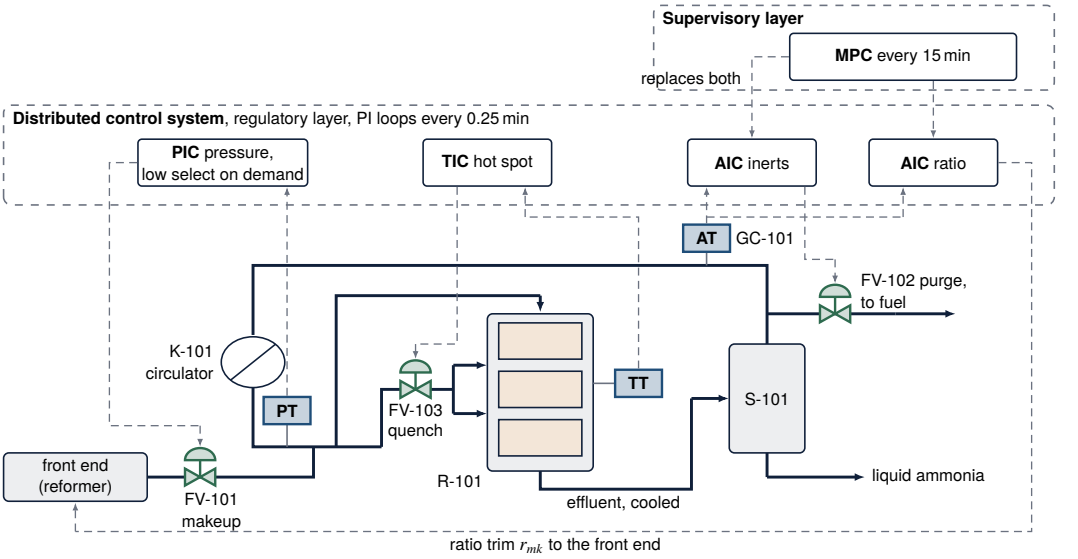

Figure 2.6 The loops of the synthesis loop with their hardware. Lettered boxes are sensing devices: the loop pressure transmitter (PT), the bed temperature element (TT) and the loop gas chromatograph (AT), which reports the inerts and the ratio together. Valve symbols are final elements, and dashed arrows are signals. Notice that one analyzer feeds both slow loops, that the predictive controller in the supervisory layer replaces those two loops on the same two handles, and that the operator's makeup demand passes a low selector against the pressure override. The exchanger, the cooler and the circulator's drive are left out. Not to scale, and not a piping diagram.

cold synthesis gas or by raising steam [25]. The model keeps the direct form, one quench valve split evenly between two bed inlets, and a circulator with no map, which cannot surge.

> **TECHNOLOGY** **Final Element: A Control Valve With a Positioner**
>
> A positioner compares the stem position of a pneumatic control valve with the controller's signal and adjusts the actuator air until the two agree, a cascade inside the valve. A smart positioner also measures the stem position, the actuator and supply air pressures and the temperature. Without a positioner, or after poor maintenance, friction makes the stem move in stair steps and ignore small reversals of the signal. On loss of air an air-to-open valve fails closed, and the risk of the process should decide which way each valve fails [24]. The model's actuator lags of 0.25 to 0.40 min stand in for this hardware.

Where the Controller Runs

A plant runs such loops in a distributed control system, in which a supervisory level oversees local controllers and sends them set points, loops are sometimes nested, and cycle times range from milliseconds to minutes [26]. The app executes its four PI loops every 0.25 min and its predictive controller every 15 min, on each new analysis, so both sit inside that range. The PI loops and the low selector belong to the regulatory layer. The plan belongs to the supervisory layer above it and writes to the purge and the trim, the same two handles the slow loops hold when it is not in service.

IN PRACTICE **Purge, Startup and the Long Run**

A 2007 European reference document on ammonia plants describes the practice the model leaves out. A continuous purge keeps argon and methane from accumulating, and its size holds the loop at roughly 10 to 15 percent inerts. A converter is brought up to temperature with a fired startup heater, and a full startup takes one or two days. Plants run for long periods, and shutdowns for technical failure average 5.7 a year [25]. The design point here, 10.07 % inerts, sits inside that band, and the converter of scenario 3 stays out because the model has no such heater.

What Is Modernizing It

Three developments bear on this loop. Recent designs recover the hydrogen from the purge by cryogenic separation, membranes or pressure swing adsorption and return it to the loop, where older designs burned the purge as reformer fuel [25]. Recovery changes the price of the purge that Table 2.6 trades against production, which is why the model leaves it out. Model-based advanced control with weighted and hierarchical optimization was reported in service on an ammonia plant in 2004 [25], the industrial counterpart of the plan of Section 2.5.4. And a loop fed with hydrogen from electrolysis will have to change its rate with the renewable supply, which calls for plants that start, stop and adjust production quickly [27]. The open challenge after the Limitations takes this up.

2.8 Limitations

The model is not a design tool and no number in it should be used to size anything. The kinetics are one fitted group rather than a mechanism, the equilibrium of (2.4) is applied to mole fractions at a pressure where fugacity coefficients are not unity, and the separator is a correlation with a factor on it. The three-bed converter cannot show an axial profile, so it cannot show where a hot spot actually sits or how deactivation moves it. The loop is one well-mixed volume, so the transport delay a real analyzer sees between the converter and the sample point is absent, and the measured composition responses here are therefore slightly optimistic. There is no compressor map, so nothing in the application can surge. Startup, catalyst reduction and the fired heater are outside the model, which means a blown-out converter in this application stays out and the recovery that a real plant performs cannot be practiced. Economics are three illustrative prices, rather than a market quotation or a complete plant cost model. At the shipped \$2.50 per kmol, the design feedstock costs about \$322 per tonne against \$420 of revenue. Hydrogen recovery, purge fuel credit, upstream production costs and capital are absent, so the economic optimum is an exercise in the stated assumptions. Disturbances arrive as commanded steps and timed events where a plant drifts. Fixing the integration step, the controller period and the random seed reproduces one trajectory on every machine, and no distributed control system is that repeatable. Five features survived all of that trimming, the recycle amplification of (2.3), the two steady states of (2.9), the two time scales, the sampled composition measurement and the selector of (2.17), and each one is present because it changes the controller that has to be built. Section 2.3.7 names where each remaining assumption stops holding.

The plan's model is linear about a single design point, with the fast loops closed at their shipped tuning, so it leaves out the weaker effect of a large purge move (0.866 % of inerts for a 5 percent opening against its 0.945 %), a retuned fast loop, as in the challenge, and a makeup valve at its stop, which only the priced ceiling and the measured hold stand in for. It prices the purge at the declared gas price whatever the slider says, and it holds whatever inert setpoint it is given, which is why it gives up \$4 493 of margin on scenario 9.

OPEN CHALLENGE **A Loop That Follows the Wind**

Ammonia made from hydrogen split out of water with renewable electricity must follow an intermittent supply. A 2020 review finds that processes with chains of high-pressure compressors, heat integration and sensitive catalysts cannot leave steady state without large stores of hydrogen and electricity [27]. A 2018 study of

a three-bed autothermal converter found its widest steady range in the inert fraction and the hydrogen to nitrogen ratio, the two slow variables of this chapter, and noted that a pilot plant had been studied only at steady state. Outside that range the converter goes out or overheats its catalyst [28]. Progress would be a loop that changes rate on demand, proven in dynamic operation, without losing its converter.

2.9 Discussion

The organizational split around this plant matches the technical one. The converter belongs to the reaction engineer, the front end belongs to a different unit and often a different department, and the two variables that decide whether the loop is economic, the inert level and the hydrogen to nitrogen ratio, belong to nobody in particular and are trended by nobody in particular. Both of them are five-hour variables measured every fifteen minutes and neither shows on the pressure trace, and a loop left for a month at the 1.80 ratio of scenario 6 gives away 40 t/d, which is 1200 tonnes of ammonia that no alarm reported. The quantity worth putting in front of an operator is not on the standard display at all, and it is the purge-corrected efficiency, 128.6 kmol of syngas per tonne against a stoichiometric 117.4, with the difference attributable to the purge.

The documented model supports methods well beyond the four implemented loops and the linear plan over two of their handles. Recycle amplification is exactly the structure that steady-state real-time optimization, Section 1.4.7 and Chapter 11, exists to handle, because the optimum purge rate of Table 2.6 moves with the makeup inert content, the catalyst activity and the price of hydrogen, and it is flat enough that an operator will not find it by hand. A nonlinear predictive controller on the same equations would carry the sintering limit, the design pressure and the quench travel as explicit constraints and would coordinate the makeup and purge valves as one two-input problem rather than as two loops with a five-hour interaction between them [9]. Moving horizon estimation on the fourteen states would recover the catalyst activity and the exchanger condition from the temperature and pressure records, which turns the invisible margin loss of scenario 11 into a soft sensor that reports before the second disturbance arrives rather than after. The equations transfer directly to an algebraic modeling language, with the six hidden constants of Table 2.2 recast as estimated parameters.

2.10 Conclusion

This chapter has used an ammonia synthesis loop as an interactive laboratory for the behavior a recycle imposes on a plant. The simulator integrates fourteen states and solves rather than assumes its design point, 149.97 bar and a 505.2 °C hot spot with 10.07 % inerts at a hydrogen to nitrogen ratio of 3.000, making 677 t/d from 1.01 kmol/s of makeup against 6.00 kmol/s of circulation, at which line-up the largest relative state derivative is 9.6×10^{-12}. Single-pass extent is 0.0386, the recycle ratio is 5.95, and the purge is 8.4 % of the makeup, which fixes the inert level at 11.90 % in the purge stream and the composition time constant at 5.1 h. A front end delivering hydrogen and nitrogen 5 % below three to one settles the loop at 1.80, an amplification of eight that costs 40 t/d against the 678 t/d the plant was making, and it takes 14 h to arrive with the loop pressure never leaving 150.0 bar. The same error high is amplified 17 times, because hydrogen leaves through a purge less readily than nitrogen does. Handing the makeup valve to the operator rather than to the pressure controller changes nothing about the plant and everything about the result, since a catalyst aged to 0.75 costs 17 t/d of production under one pairing and 2.3 bar of pressure margin under the other. A hot-spot setpoint seven degrees lower is worth 7.0 % of production, and the quench travel that buys it is the same travel a converter needs in order not to go out, and scenario 11 spends that travel four and a half hours before it is needed. A linear model predictive controller over the purge and the trim holds the inerts within 0.1 % of setpoint from 192 min after a two-drift upset that two PI loops have not corrected a day later, for $924 of margin in purged hydrogen. Further work would add a hydrogen recovery unit on the purge, a compressor map with a surge line, a fired startup heater so that an extinguished converter can be recovered, and an axial converter profile. The synthesis loop is

an accessible browser platform for learning what a recycle does to a small error and what it costs to choose the wrong inventory.

KEY POINTS

- An ammonia loop sends most of its gas around again, so a small, steady error in the makeup gas grows inside the loop over hours.
- The inert level and its five-hour time constant are both set by the small purge.
- Holding pressure turns a catalyst loss into lost production, and holding production turns it into lost pressure margin.
- The quench valve position, not the hot spot, shows the margin left.
- A predictive controller on purge and trim corrects slow drifts sooner than two PI loops and pays in purged hydrogen.

Where else this applies. Volume 6, Chapter 7 shows the same amplification, and Volume 6, Chapter 8 a gas recycle with an inert and a purge.

References

[1] U.S. Geological Survey, *Mineral Commodity Summaries 2025*. Reston, VA: U.S. Geological Survey, 2025, nitrogen (fixed)–ammonia, pp. 128–129.

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

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

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

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

[6] R. F. Baddour, P. L. T. Brian, B. A. Logeais, and J. P. Eymery, "Steady-state simulation of an ammonia synthesis converter," *Chemical Engineering Science*, vol. 20, no. 4, pp. 281–292, 1965.

[7] P. L. T. Brian, R. F. Baddour, and J. P. Eymery, "Transient behaviour of an ammonia synthesis reactor," *Chemical Engineering Science*, vol. 20, no. 4, pp. 297–310, 1965.

[8] L. D. Gaines, "Ammonia synthesis loop variables investigated by steady-state simulation," *Chemical Engineering Science*, vol. 34, no. 1, pp. 37–50, 1979.

[9] A. Araújo and S. Skogestad, "Control structure design for the ammonia synthesis process," *Computers & Chemical Engineering*, vol. 32, no. 12, pp. 2920–2932, 2008.

[10] V. Smil, *Enriching the Earth: Fritz Haber, Carl Bosch, and the Transformation of World Food Production*. Cambridge, MA: MIT Press, 2001.

[11] J. W. Erisman, M. A. Sutton, J. Galloway, Z. Klimont, and W. Winiwarter, "How a century of ammonia synthesis changed the world," *Nature Geoscience*, vol. 1, no. 10, pp. 636–639, 2008.

[12] M. Appl, "Ammonia, 2. Production processes," in *Ullmann's Encyclopedia of Industrial Chemistry*. Weinheim, Germany: Wiley-VCH, 2011.

[13] F. Haber, "The synthesis of ammonia from its elements," Nobel Lecture, 2 June 1920, 1920, nobelprize.org/uploads/2018/06/haber-lecture.pdf, read November 2026.

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

[15] J. C. Morud and S. Skogestad, "Analysis of instability in an industrial ammonia reactor," *AIChE Journal*, vol. 44, no. 4, pp. 888–895, 1998.

[16] C. Bosch, "The development of the chemical high pressure method during the establishment of the new ammonia industry," Nobel Lecture, 21 May 1932, 1932, nobelprize.org/uploads/2018/06/bosch-lecture.pdf, read November 2026.

[17] D. R. MacFarlane, P. V. Cherepanov, J. Choi, B. H. R. Suryanto, R. Y. Hodgetts, J. M. Bakker, F. M. Ferrero Vallana, and A. N. Simonov, "A roadmap to the ammonia economy," *Joule*, vol. 4, no. 6, pp. 1186–1205, 2020.

[18] L. J. Gillespie and J. A. Beattie, "The thermodynamic treatment of chemical equilibria in systems composed of real gases. I. an approximate equation for the mass action function applied to the existing data on the Haber equilibrium," *Physical Review*, vol. 36, no. 4, pp. 743–753, 1930.

[19] D. C. Dyson and J. M. Simon, "Kinetic expression with diffusion correction for ammonia synthesis on industrial catalyst," *Industrial & Engineering Chemistry Fundamentals*, vol. 7, no. 4, pp. 605–610, 1968.

[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] 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.

[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] European Commission, "Reference document on best available techniques for the manufacture of large volume inorganic chemicals: Ammonia, acids and fertilisers," European Commission, Integrated Pollution Prevention and Control, Tech. Rep., Aug. 2007, eippcb.jrc.ec.europa.eu, read November 2026.

[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] C. Smith, A. K. Hill, and L. Torrente-Murciano, "Current and future role of Haber–Bosch ammonia in a carbon-free energy landscape," *Energy & Environmental Science*, vol. 13, no. 2, pp. 331–344, 2020.

[28] I. I. Cheema and U. Krewer, "Operating envelope of Haber–Bosch process design for power-to-ammonia," *RSC Advances*, vol. 8, no. 61, pp. 34 926–34 936, 2018.

2.11 Exercises

2.1 ⌕ Concept check · 2 min

The loop runs on the basic loops with the pressure controller on the makeup valve and the hot-spot controller on the quench. The catalyst activity falls from 1.00 to 0.75. What does the operator see on the display once the loops have settled?

(a) The loop pressure climbs until the high limit takes the makeup valve back
(b) The hot spot falls below its setpoint and stays there with the quench shut
(c) Production falls while the loop pressure stays on its setpoint
(d) The loop inerts climb, because the converter consumes less of the makeup

2.2 ⌕ Concept check · 1 min

In scenario 3 the quench is opened to 22 % with the plant in Manual, and the converter goes out. At 90 min the quench returns to exactly 10 %, the value at which the converter ran lit at the start. What happens next?

(a) The converter relights within minutes, as the exchanger restores its inlet temperature
(b) The converter relights slowly, on the five-hour clock of the loop composition
(c) The converter relights once the pressure has risen far enough to speed up the reaction
(d) The converter stays out, because a dead bed cannot heat its own feed back to ignition

2.3 ⌕ Concept check · 2 min

The loop pressure answers a makeup valve move in minutes, while the loop inert level answers a purge valve move over hours. What sets the inert time scale?

(a) The loop gas holdup divided by the small purge flow that is the only exit for inerts
(b) The 15 min cycle of the loop gas analyzer that measures the inerts
(c) The 300 min integral time of the inert controller in its shipped tuning
(d) The heat capacity of the catalyst beds, which delays every change in conversion

2.4 Theory and calculation · 8 min

At the design line-up the loop pressure is 149.97 bar, the purge valve is at 37.4 % with capacity $k_P = 0.01850\,\mathrm{kmol/s}$ per $\sqrt{\mathrm{bar}}$, and the makeup is 1.0082 kmol/s carrying 1.00 % inerts. The loop holds $V = 300\,\mathrm{m^3}$ of gas with $z = 1.12$ at a mean $T_L = 40\,^{\circ}\mathrm{C}$. Use $R = 8.314\,\mathrm{J/mol/K}$.

(a) Find the purge flow and the share of the makeup it removes. *Checkpoint:* 0.08473 kmol/s, 8.4 % of the makeup.

(b) Find the loop holdup N. *Checkpoint:* 1543 kmol.

(c) Find the steady inert fraction in the purge stream and the inert time constant, in seconds and in hours. *Checkpoint:* 11.90 %, about 18 210 s or 5.06 h.

(d) The purge valve is cut to 20 % with the pressure held at 150 bar. Compute $N/\dot{n}_P$ at the new purge. Scenario 7 measures 63 % of the change in 5.12 h. Why is the loop faster than this ratio? *Checkpoint:* About 9.5 h. The makeup falls with production.

2.5 Theory and calculation · 6 min

The equilibrium constant of Equation (2.4) is used in Equation (2.5) with the pressure written as the model writes it, $P = 150$.

(a) Evaluate K_p at 400 °C and at 430 °C and find the factor by which it falls over those 30 °C. *Checkpoint:* 1.291×10^{-2} and 8.671×10^{-3}, a factor of 1.49.

(b) For a loop gas with $r = 3$ and $y_I = 0.10$ at 150 bar, find κ, u and the equilibrium ammonia fraction at 450 °C. *Checkpoint:* $\kappa = 0.3295$, $u = 0.7262$, $y_A^{eq} = 0.1738$.

(c) Repeat (b) at 505 °C, the design hot spot, and say what the two results mean for a bed that runs hotter. *Checkpoint:* $\kappa = 0.1770$, $y_A^{eq} = 0.1103$.

2.6 Theory and calculation · 4 min

At the design line-up the circulator moves $F = 6.00\,\mathrm{kmol/s}$, the converter extent is $x_{tot} = 0.03862$ per mole of feed, the separator returns $y_{A,sep} = 0.0379$ in its overhead gas, the purge is 0.08473 kmol/s drawn from that gas, and the makeup is 1.0082 kmol/s. Ammonia is 17.03 kg/kmol.

(a) Find the ammonia made in the converter and the liquid ammonia drawn from the separator at steady state, and express the product in tonnes per day. *Checkpoint:* 0.4634 kmol/s made, 0.4602 kmol/s drawn, 677 t/d.

(b) Find the recycle ratio, the circulator flow over the makeup as the app reports it, and the makeup gas consumed per tonne of product. *Checkpoint:* 5.95 and 128.6 kmol/t.

2.7 App challenge · 8 min

Load scenario 6, A small error at the front end, and change no setting. At 30 min the front-end ratio drifts 0.15 low, from 3.00 to 2.85, with the pressure and hot-spot loops in service. Run at 60x with the default seed. Pause between 720 and 740 min, set Front-end ratio drift to +0.15, and run 240 min more. Export the CSV file, whose time column is in hours.

(a) Read the true loop ratio, column `ratio`, and the production at 720 min, and the largest departure of the true pressure, column `press_bar`, from 150 bar between 30 and 720 min. *Checkpoint:* 1.853 within 0.005, 640 t/d within 1 t/d, and under 0.05 bar.

(b) Taking the amplification as the relative loop ratio error divided by the relative makeup error, find its value at 720 min, and compare it with the settled value of 1.804 that Figure 2.2 reports. *Checkpoint:* 7.6 at 720 min, against 8.0 settled.

(c) Read the loop ratio 240 min after the drift was set to +0.15. How far is it from the settled value of Figure 2.2 for a makeup ratio of 3.15? *Checkpoint:* 3.40 within 0.01, against 5.58 settled.

2.8 App challenge · 8 min

Load scenario 11, Five years of catalyst, and change no setting. The catalyst activity falls to 0.65 at 30 min and the feed and effluent exchanger loses 8 % of its effectiveness at 300 min. Run to 400 min at 60x with the default seed and export the CSV file, whose time column is in hours.

(a) Read the quench command, column `quench_cmd_pct`, and the production at 299.75 min, the last row before the exchanger loss. *Checkpoint:* 8.45 % within 0.02 %, and 650 t/d.

(b) Find the first row after 300 min at which the quench command reaches zero and the first at which the hot spot, column `hotspot_C`, is below 380 °C. *Checkpoint:* 304.25 min and 306.0 min.

(c) Press Reset. Once the activity event has fired at 30 min, and before 300 min, set Catalyst activity to 1.00. Run to 360 min. Does the converter stay lit? Read the converter inlet, column `tin_C`, and the quench command at 360 min. *Checkpoint:* It stays lit, with the inlet at 392.0 °C within 0.3 °C and the quench at zero.

2.9 App challenge · 8 min

Load scenario 12, Five hours ahead: Supervised or MPC, and change no setting. At 60 min the methane slip rises by half and the front end turns 0.10 nitrogen heavy. Run to 1440 min at 180x with the default seed under the MPC structure the scenario loads, and export the CSV file, whose time column is in hours. Then press Reset, set Control structure to Supervised before pressing Run, run the same 1440 min, and export again.

(a) From the MPC run, read the peak of `inert_pct` and its time, and the purge command, column `purge_cmd_pct`, at 76 min. *Checkpoint:* 10.36 % at 81.5 min, and 56.2 %.

(b) From the supervised run, read the peak of `inert_pct` and its time, and its value at 1440 min. *Checkpoint:* 11.98 % at 333 min, and 10.73 %.

(c) Read the IAE of the Loop inerts row of the scorecard and the column `metric_profit` at 1440 min for both runs. What does the plan's better control cost? *Checkpoint:* IAE 56.6 against 1843.7, within 1, and $38 475 against $39 399, $924 less margin.

2.10 Application area · 3 min

A plant manager has to choose which variable the makeup valve holds. Plant A sells to a customer who takes a fixed tonnage every day. Plant B sends its product into a pipeline that takes whatever it makes. Which pairing suits each, and what must the operators of each watch that the other need not?

2.11 Application area · 4 min

In scenario 9 the methane slip from the reformer doubles the makeup inerts. The inert controller holds its setpoint by opening the purge, and once settled the gas per tonne rises from 128.7 to 141.9 kmol/t at about 677 t/d. At the shipped syngas price of $2.50 per kmol, estimate what the dirtier gas costs per day. Where does the fix belong, and what can the loop engineer do in the meantime?

2.12 Review the agent · 5 min

An AI assistant reviews Table 2.6 and writes: "Production rises monotonically with the purge valve, from 484 to 792 t/d. Opening the purge from the design 37.4 % to 85 % adds 17 % to production with the loop still at 150 bar, so the plant should run with the purge at 85 %." Find the flaw and say how you would check it.

CHAPTER 3

A Urea Synthesis Loop Whose Controlling Ratio Is Never Measured

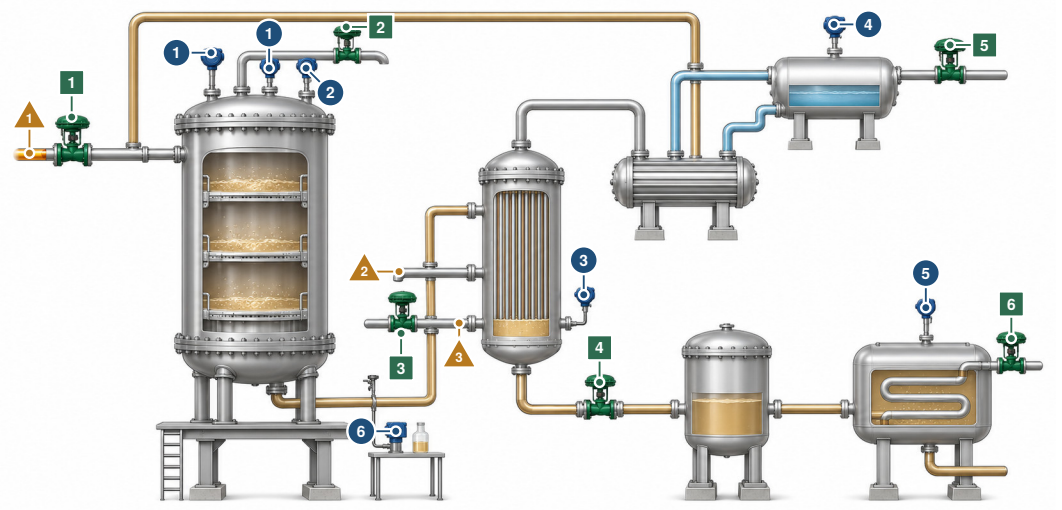

The controller moves 1 fresh ammonia, held in ratio to the carbon dioxide 2 inert vent valve 3 stripper steam valve 4 stripper level valve 5 steam drum valve, the reactor's heat 6 evaporator steam valve

The controller reads 1 density and temperature at the reactor top, from which the ratio is inferred 2 reactor pressure 3 stripper level 4 steam drum pressure 5 evaporator temperature 6 laboratory sample every four hours

What disturbs it 1 wetter ammonia 2 a carbon dioxide turndown 3 a dip in the stripper steam

Nothing in the picture reads the ratio of ammonia to carbon dioxide that decides the run, only a density and a temperature on the reactor and a laboratory sample taken hours apart. Illustration, not to scale.

IN PLAIN TERMS A Mixing Ratio Nobody Can See

The main idea. A urea plant joins ammonia and carbon dioxide under high pressure and sends what does not react around the loop again. The amount of spare ammonia in the reactor decides how much urea forms and whether the steel lining stays safe.

The challenge. No instrument can read that amount, so it can drift while both feed meters look normal. The plant guesses it from how heavy and how hot the liquid is, and a slow lab test fixes the guess hours later.

Why it matters. Too little ammonia lets the hot liquid eat the lining, and too much drives up the pressure.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page; the history and fact boxes of Sections 3.2 and 3.3; the try-it box of Section 3.4; the try-it and fact boxes of Section 3.6; the key points
Study	45 min	the whole chapter with its pause points and the exercises
Practice	10 min	the first page; Section 3.6 with its result chart; Section 3.7; Section 3.8 with its open challenge

The control question The molar ratio of ammonia to carbon dioxide in a urea reactor sets the conversion, the pressure the liquid needs and whether the stainless liner stays passive, and no online instrument measures it. Can a ratio held on the two feed flows hold it, and what do an inference from the reactor's density and temperature, corrected by a laboratory that reports 45 min late, and a predictive controller add?

Learning objectives

- Explain why the reactor N/C is an inventory that integrates every unmeasured error in the fresh feed, and estimate how fast wet ammonia drains it.
- Compute the inferred N/C from density and temperature, and the laboratory bias update made against the inference at the sampling instant.
- Evaluate the corrosion index and say how the reactor temperature and the passivation air move the floor.
- Explain how the carbamate condenser sets the reactor temperature, and why a cascade onto the steam drum restores it at the cost of biuret.
- Compare ratio control on the flows, a PI on the inference and a predictive controller over a twelve-hour shift, including the margin each one leaves.

Prerequisites Sections 1.4.3 (a recycle whose excess has no exit), 1.4.6 (ratio stations, inferred compositions and the bias update), 1.6.2 (velocity form) and 1.6.4 (scenarios and exported files). Section 1.6.3 is needed only for the predictive controller, and Chapter 2 gives a second synthesis loop with the same recycle arithmetic.

The app https://apmonitor.com/apps/urea-plant/

See also Volume 5, Chapter 10 for a soft sensor whose bias a sampled and delayed laboratory corrects, and Volume 5, Chapter 3 for composition control through a slow analyzer. Volume 3, Chapter 7 for a ratio blend trimmed on a composition measurement, and Volume 3, Chapter 6 for a feedforward dose demoted to a slow ratio trim. Volume 6, Chapter 7 for the snowball effect of a reactor and column recycle, and Volume 6, Chapter 8 for a plant whose inventories must close before anything else.



3.1 Introduction

The variable a urea control room cares most about has no transmitter. In the CO_2-stripping loop simulated here, 45.8 t/h of carbon dioxide and 35.83 t/h of ammonia make 62.5 t/h of urea at 140 bar a, and the reactor outlet carries ammonia and carbon dioxide in the molar ratio $N/C = 3.000$ on the initial-mixture basis. Ammonia beyond the two moles each urea molecule takes stays in the loop. When the ammonia arrives with 2.0 instead of 0.2 weight percent water at 1 h, a ratio held on the measured flows lets the reactor's N/C fall to the 2.80 corrosion floor at 7 h 49 min, while the laboratory's first sample, drawn at 4 h, reports 2.918 at 4 h 45 min. A ratio inferred from the reactor's density and temperature had already crossed 2.93 at 3 h 42 min, because free ammonia is the lightest component of the synthesis liquid.

The simulator is aimed at a common misconception, that a ratio held between two measured feeds is a ratio held in the process. A ratio station is a feedforward on the feeds, as in Section 1.4.6, and the quantity downstream of a recycle is an inventory that integrates whatever the feeds get wrong. The general remedy is inferential control, estimating an unmeasured product quality from secondary measurements [1] and maintaining that estimate against slow, accurate analyses, which is the role soft sensors play across the process industries [2]. The milk powder dryer of Volume 5, Chapter 10 corrects one such soft sensor from a sampled and delayed laboratory. Recycles also add dynamics of their own, among them the snowball effects of reactor and separator processes [3] that Volume 6, Chapter 7 measures. The same structure recurs wherever a reactor's feed ratio is held in a loop of unconverted reactants, among them ammonia synthesis (Chapter 2), methanol synthesis, hydrocracker hydrogen and polymerization monomer recovery.

Urea synthesis has been modeled for decades, in a simulation program for urea plants [4], as a thermodynamic framework and reactor model [5, 6], as the synthesis section of an industrial plant [7] and as a high-pressure loop [8], and its phase and reaction equilibria have been brought together in

one thermodynamic model of the synthesis [9]. However, those models are built to answer design and rating questions, and none of them hands a student an N/C that drifts while every flow meter reads its setpoint. This chapter works with a browser simulator of a CO_2-stripping urea loop whose true ratio, inferred ratio and laboratory reports are logged side by side, and it measures what ratio control on the flows, a laboratory-corrected inference and a predictive controller each hold through drift, turndown, a stripper steam dip and a rate cut. The balances, the correlations, the velocity-form controller and the linear predictive controller are standard. The measurements on this one synthetic plant show, among other things, that the corrosion floor is reached by wet ammonia rather than by a CO_2 turndown, which raises the ratio. The model has 39 states whose nitrogen, carbon, hydrogen and oxygen ledgers close to 3×10^{-11} kmol over twelve hours, and every controller reads a measurement packet only. Nine graded cases judge both what the learner does and what the plant does. Chapter 1 treats the recycle and the inferred composition in general terms in Sections 1.4.3 and 1.4.6.

Section 3.2 follows the process from carbamate to the total-recycle loop. Section 3.3 sets out the balances, the teaching correlations, the instruments and the inference. Section 3.4 measures what the loop does before it is controlled: the ratio integrates, the condenser sets the reactor's heat and the recycle stores a disturbance. Section 3.5 describes the loops and the three ways to hold the ratio, and Section 3.6 measures the nine cases and the twelve-hour challenge. Section 3.7 names the hardware and the practice behind the model, Section 3.8 states where the model stops holding, and Section 3.9 collects the measured results.

3.2 From Carbamate to a Total-Recycle Loop

In the industrial route, whose origin the history box at the end of this section tells, ammonia and carbon dioxide combine to ammonium carbamate, which then loses water to give urea. The first step is fast, exothermic and nearly complete above about 135 °C at pressures above the carbamate's decomposition pressure, while the second is slow, slightly endothermic and never complete, so every process separates the unconverted carbamate and returns it to the reactor. Because the carbamate is volatile at these temperatures the synthesis must run at high pressure, typically 170 °C to 200 °C and 130 to 300 bar, and both the corrosion of the liquid and the rate of urea formation rise roughly exponentially with temperature [10].

The equilibrium was measured and modeled in steps. Equilibrium pressures of the CO_2–NH_3–urea–water system were measured as a function of the loading density of the vessel in 1951 [11]. A single overall equilibrium constant for the liquid, proposed in 1948, remained the most widely used model, and a later model with free ammonia and carbon dioxide, carbamate, urea and water in the liquid and a vapor phase was presented in 1972 [10]. The densities of the synthesis solution in chemical equilibrium were measured for N/C from 2.04 to 6.3 and fitted by an empirical equation [12], which is the physical basis for inferring the ratio from a density. A phase model of the same system places minimum bubble points and a ternary azeotrope in the composition plane [13]. The Gorlovskii–Kucheryavyi equilibrium conversions tabulated with a 1993 reactor model [5] and the conversions and bubble pressures of a 2016 thermodynamic model [9] are the published values this chapter compares the simulator's correlations with.

Process design moved from once-through plants to total recycle, in which the reactor effluent passes a decomposer or stripper and the recovered ammonia and carbon dioxide return to the synthesis [10]. In a CO_2-stripping plant both reactants return together with water as carbamate, so the concentration of urea in the reactor liquid, rather than the CO_2 conversion alone, decides the optimum, and that yield passes through a maximum near $N/C = 3$. The same review observes that the reactor is a bubble column of boiling liquid whose endothermic reaction is heated by condensing gas, that the total pressure is a system pressure plus the pressure of inerts that include the passivation air, and that the temperature rise along the reactor indicates the conversion [10]. The worldwide use of urea as a fertilizer has escalated far enough to be studied as a global change that contributes to coastal eutrophication [14]. None of the sources consulted here documents how a commercial plant's

Table 3.1 Variables, Their Roles and Design Values (Case 1, $t = 0$)

Variable	Role	Design	Range
N/C	CV, true (truth only)	3.000	SP 2.7–3.4
$\widehat{N/C}$	CV, inferred	3.014	—
R	MV, fresh NH_3/CO_2 ratio	2.0173	1.80–2.30
S	MV, stripper steam ratio	1.15 t/t	0.8–1.5
p, v	CV, MV: pressure, vent	140.0 bar a, 39.7 %	SP 130–150
L	CV, stripper level	50 %	SP 20–80
p_s	CV, steam drum	4.50 bar a	SP 3.6–6.5
T	CV of the cascade, reactor top	185.02 °C	SP 175–192
T_e	CV, evaporator	138.0 °C	SP 130–145
—	DV, CO_2 load	100 %	50–115
—	DV, ammonia purity	99.8 wt%	96–100
O_2	DV, passivation air	0.6 vol%	0–1
—	DV, HP steam supply	80 t/h	30–100 %
—	DV, CO_2 meter, density gauge	0, 0	±5 %, ±10 kg/m^3
X, H_2O/C	outputs, reactor outlet	61.11 %, 0.406	—
ρ, p_b	outputs	970.6 kg/m^3, 127.6 bar	—
B, w	outputs, product	0.748, 0.300 wt%	—
I_c	output, corrosion index	0.018	stop at 1

control system is configured, so the simulator's loops are a study of the structure above rather than a description of practice.

> ⧖ HISTORY **From Wöhler's Urea to a Loop Under Pressure, 1828 to 1932**
>
> Wöhler reported the artificial formation of urea in 1828 [15], but the industrial route is a different one, through ammonium carbamate. It was first proposed by Basaroff and carried into manufacture by German chemists about 1920 [10]. In his Nobel lecture of 1932, Carl Bosch described the difficulty. Carbon dioxide and ammonia heated to 130 to 140 °C in a closed vessel give urea and water, but stop at about 40 percent of the theoretical yield, and on slow cooling the water would turn most of the urea back, and the use of heat would be very poor. His company therefore separated the mixture at high pressure, drew off the excess ammonia and the carbon dioxide with heated compressors and condensed them again, making about 40,000 tons of urea a year [16]. The total-recycle loop of Figure 3.1 returns, in the same way, what its reactor leaves unconverted.

3.3 The Loop as Modeled

The plant is a generic 1500 t/d CO_2-stripping loop with no licensor's data, as shown in Figure 3.1. The reactor overflows into a falling-film stripper heated by high-pressure (HP) steam, where the fresh CO_2 rises as stripping gas. The stripper gas condenses to carbamate in the HP carbamate condenser, whose shell raises low-pressure (LP) steam in a drum, and the liquid and the uncondensed gas flow into the bottom of the reactor, where the gas finishes condensing and supplies the reactor's heat. The fresh ammonia and the carbamate recycle from the LP section enter the condenser. The stripper's liquid passes the LP section, a urea solution tank and an evaporator that makes the melt. The engine's clock is the minute, species flows are in kmol/min, heat in MJ/min, pressures in bar absolute, and feed flows are reported in t/h.

Table 3.1 lists each variable's role and its value at the design point of case 1. The integrated state holds four apparent species and a temperature in each reactor cell (15), the stripper's bottoms holdup (4), the LP recycle holdup (3), the urea solution tank (2), the inerts in the reactor gas space, the drum pressure, the evaporator temperature and its biuret, seven actuator positions and four element ledgers, 39 in all.

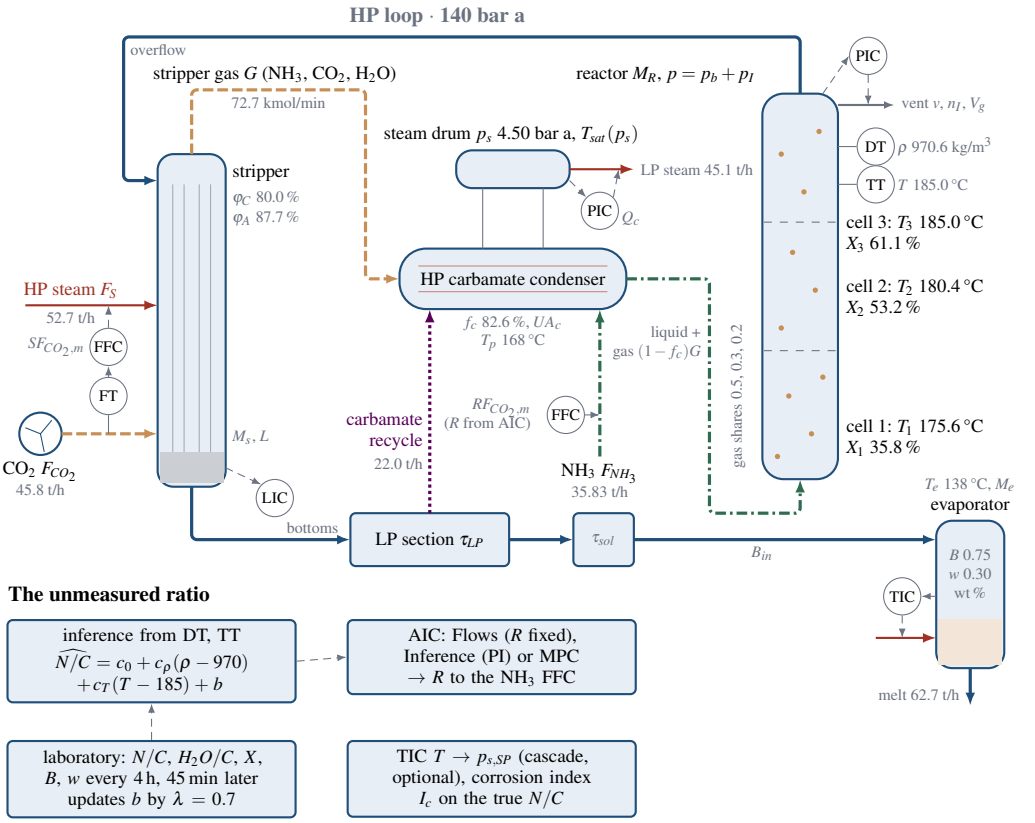

Figure 3.1 The CO_2-stripping synthesis loop as modeled, redrawn from the Loop view of case 1 at $t = 0$, the design point. The stripper gas G (amber, dashed) goes to the HP carbamate condenser, which condenses a fraction f_c and raises LP steam in a drum at p_s. The liquid and the uncondensed gas $(1 - f_c)G$ (green, dash-dotted, as is the fresh ammonia) enter cell 1 of the three-cell reactor, and the overflow (navy, solid) returns to the stripper, which strips fractions φ_C and φ_A with the fresh CO_2 and HP steam F_S. The carbamate recycle (violet, dotted) returns through the LP section with lag τ_{LP}, and the urea solution passes a tank τ_{sol} to the evaporator at T_e. DT and TT on the reactor top feed the inference (lower left), which the laboratory corrects. Process model only, not a piping diagram. Not to scale.

3.3.1 Species, Ratios and the Reactor Cells

Every holdup carries apparent ammonia A, carbon dioxide C, urea U and water W in kmol. Carbamate formation is taken as instantaneous and complete while ammonia is in excess, so the carbamate is $\min(C, A/2)$, and the ratios the industry quotes are those of the initial mixture [10], as shown in Equation (3.1).

$$N/C = \frac{A + 2U}{C + U}, \quad \mathrm{H_2O/C} = \frac{W - U}{C + U}, \quad X = \frac{U}{C + U} \tag{3.1}$$

where X is the CO_2 conversion. Dehydration changes neither ratio, so they describe the feed a cell received. The reactor is three equal cells of 45 t in series, and each balances its species and its heat,

as shown in Equation (3.2).

$$\begin{aligned}
\frac{\mathrm{d}A_k}{\mathrm{d}t} &= A_k^{in} + g_{A,k} - f_k A_k - 2r_k - \delta_{k3}\, n_v \\
\frac{\mathrm{d}C_k}{\mathrm{d}t} &= C_k^{in} + g_{C,k} - f_k C_k - r_k \\
\frac{\mathrm{d}U_k}{\mathrm{d}t} &= U_k^{in} - f_k U_k + r_k \\
\frac{\mathrm{d}W_k}{\mathrm{d}t} &= W_k^{in} + g_{W,k} - f_k W_k + r_k \\
r_k &= k(T_k)\,(C_k + U_k)\,\big(X_{eq,k} - X_k\big) \\
k(T) &= k_{185}\exp\Big[-\frac{E_d}{R}\Big(\frac{1}{T} - \frac{1}{T_{185}}\Big)\Big] \\
M_k c_{pL}\frac{\mathrm{d}T_k}{\mathrm{d}t} &= \dot{m}_k^L c_{pL}(T_k^{in} - T_k) + \dot{m}_k^G c_{pG}(T_p - T_k) \\
&\quad + q_G(g_k) - \Delta H_d\, r_k
\end{aligned} \tag{3.2}$$

where $g_k = s_k(1-f_c)G$ is the share $s_k = 0.5, 0.3, 0.2$ of the condenser's uncondensed gas that condenses in cell k, f_k is the outflow per unit holdup, $\dot{m}_k^L$ and $\dot{m}_k^G$ are the liquid and gas mass flows entering, $T_p = 168\,^{\circ}\mathrm{C}$ is the condenser's process-side temperature, and n_v is the ammonia lost with the vent. The condensation heat is $q_G = \Delta H_c \min(C, A/2) + 20A_{free} + 40W$ in MJ/kmol with $\Delta H_c = 100\,\mathrm{kJ/mol}$, and the dehydration takes $\Delta H_d = 15.5\,\mathrm{kJ/mol}$. Cells 1 and 2 overflow what they receive, and cell 3 overflows $\dot{m}_{out} = \bar{m} + (M_3 - \bar{M}_3)/\tau_{ov}$ with $\tau_{ov} = 0.5\,\mathrm{min}$. At design the cells run at 175.58, 180.40 and 185.02 °C with conversions of 35.80, 53.22 and 61.11 % against an equilibrium of 66.26 % at the top, and the liquid stays 47.7 minutes in the reactor. The rising temperature along the reactor reflects the conversion, as the industrial reactor's does [10].

> **+ DID YOU KNOW? Urea, the Leading Solid Nitrogen Fertilizer**
>
> In 2007 urea was the most used solid nitrogen fertilizer in the world, with an output above 100 million metric tons a year and new plants of typically 2000 tonnes a day, some up to 3500 [17]. In the United States in 2024 it was the derivative of ammonia made in the largest quantity [18]. This chapter's plant makes 1500.

3.3.2 Equilibrium Conversion, a Teaching Correlation

The equilibrium conversion is a polynomial in $a = N/C$, $b = \mathrm{H_2O/C}$ and $\tau = T/100$ with T in °C, as shown in Equation (3.3).

$$\begin{aligned}
X_{eq} = {} & 0.2616a - 0.01945a^2 + 0.0382ab - 0.1160b \\
& - 0.02732a\tau - 0.1030b\tau + 1.640\tau \\
& - 0.1394\tau^3 - 1.869
\end{aligned} \tag{3.3}$$

clamped to $2 \le a \le 5$, $0 \le b \le 1.5$ and 140 °C–210 °C, with a declared shape for a from 2.2 to 4.5. It is a teaching correlation of the published shape, not a fit to any data set. Its conversion rises with N/C (64.10, 66.37 and 75.39 % at 2.8, 3.0 and 4.0 for $\mathrm{H_2O/C} = 0.4$ and 185 °C), falls with water (62.53 % at $\mathrm{H_2O/C} = 0.6$) and peaks in temperature near 190 °C (190.4 °C at $N/C = 3$). Against the 34 Gorlovskii–Kucheryavyi equilibrium conversions tabulated by Isla, Irazoqui and Genoud for N/C 3.5 to 5 at 190 °C and $N/C = 4$ from 180 to 210 °C [5], it runs from 1.5 points below to 2.8 points above, 1.3 points above on average. Against eleven points read off the conversion curves of a process paper whose sources include the Ullmann's encyclopedia [10], for N/C from 2.2 to 5.0 and 165 to 210 °C, it runs 0.6 to 3.1 points above, and 2.4 to 3.1 points above for N/C up to 3. Against Voskov and Voronin's model [9] at N/C 2.5 and 3, 180 and 200 °C and $\mathrm{H_2O/C}$ up to 0.6 it runs 4.4 to 9.1

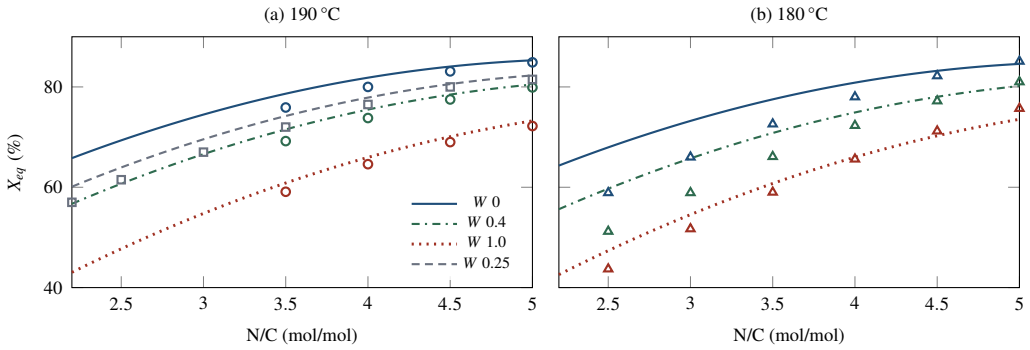

Figure 3.2 The equilibrium conversion correlation (lines) against published values (marks) at $H_2O/C = W$. The water ratios differ by line style and mark fill as well as by color, from a solid line and open marks for the dry liquid to a dotted line and filled marks for the wettest, as in the key. (a) At 190 °C, circles are the Gorlovskii–Kucheryavyi conversions tabulated by Isla et al. and squares the curve reproduced by Brouwer at $W \approx 0.25$ (gray dashed line). (b) At 180 °C, triangles are the Voskov–Voronin model. The correlation's declared range ends at $N/C = 4.5$.

points above, and that model itself lies about 3 to 5 points below the other two near $N/C = 3$. At the design point the correlation gives 66.4 % where the curve reads about 64.5 and Voskov–Voronin 59.3. The comparison is drawn in Figure 3.2.

The shape matters more than the level for the lessons. The plan for the application asked for a conversion maximum near N/C 3 to 3.5, but none of the three sources has one, since the published conversion rises with N/C throughout. The correlation rises too, and the maximum near 3 appears where the published literature puts it, in the urea made per kilogram of synthesis liquid, $X_{eq}M_U/(aM_A + M_C + bM_W)$. It peaks at $N/C = 2.98$ at $H_2O/C = 0.4$ and 185 °C and at 2.56 at $H_2O/C = 0.25$ and 190 °C, where the published yield curve peaks near 3 [10]. The maximum is flat, so its location is sensitive to the correlation. The water slope at $N/C = 3.5$ and 190 °C, −0.178 per unit of H_2O/C, is close to the curve's −0.17, and the temperature maximum at $N/C = 3.5$, $H_2O/C = 0.25$ is 190.6 °C against about 193 °C on the published curve.

3.3.3 Stripper, Condenser and Steam Drum

The stripper decomposes fractions of the reactor overflow's CO_2, ammonia and water from the CO_2 and the steam it receives per unit of liquid, relative to design, and the condenser condenses a fraction of the stripper gas set by the duty its shell can take, as shown in Equation (3.4).

$$\begin{aligned}
\varphi_C &= 1 - \mathrm{e}^{-\eta(k_g r_g + k_q r_q)} \\
\varphi_A &= 1 - (1 - \varphi_C)^{1.3}, \qquad \varphi_W = 0.08\sqrt{r_q} \\
f_c &= 1 - \mathrm{e}^{-2z} \\
z &= \frac{UA_c\big(T_p - T_{sat}(p_s)\big) + Q_f - Q_g}{q_G(G)} \\
Q_c &= f_c\, q_G(G) + Q_g - Q_f \\
C_s \frac{\mathrm{d}p_s}{\mathrm{d}t} &= \frac{Q_c}{\lambda_{LP}} - K_s v_s \sqrt{p_s^2 - p_h^2}
\end{aligned} \tag{3.4}$$

where r_g and r_q are the CO_2 and the steam per kilomole of liquid on the CO_2 basis over their design values, $k_g = 0.5$, $k_q = 1.109$, η is the stripper efficiency factor, Q_f is the heat that warms the fresh ammonia from 60 and the recycle from 100 °C to T_p, Q_g the heat the gas gives up cooling from

185 °C, T_{sat} the IAPWS-IF97 saturation temperature, $\lambda_{LP} = 2.12\,\mathrm{MJ/kg}$, v_s the LP steam valve and $p_h = 3.5\,\mathrm{bar\,a}$ the header. At design the stripper strips 80.0 percent of the CO_2 and 87.7 percent of the ammonia with 52.67 t/h of steam, the gas is 72.72 kmol/min, and the condenser condenses 82.6 percent of it and raises 45.06 t/h of steam at 4.50 bar a ($T_{sat} = 147.91\,^{\circ}\mathrm{C}$). The conductance $UA_c = 84.8\,\mathrm{MJ/(min\,K)}$ was sized so that the top cell sits at 185 °C. The exponential condensed fraction is a teaching closure in the manner of a number of transfer units, which replaced a clipped linear fraction under which the reactor and the condenser oscillated together. The gas the condenser leaves heats the reactor, so the drum pressure is the reactor's temperature handle, 11.6 K per bar from 4.2 to 4.8 bar a and 15.1 K per bar from 4.5 to 5.0 at the lined-out design.

3.3.4 Pressure, Recycle and Evaporator

The reactor pressure is the liquid's bubble pressure plus the inerts that collect in the gas space at the top [10], as shown in Equation (3.5).

$$\begin{aligned} p &= p_b + \frac{n_I R T_3}{V_g} \\ p_b &= 124\,\mathrm{e}^{0.018(T_3-185)+0.40(a-3)-0.30(b-0.5)} \\ \frac{\mathrm{d}n_I}{\mathrm{d}t} &= F_C\left(y_I + \frac{\mathrm{O_2}}{21}\right) - K_v v\sqrt{p}\,\frac{p-p_b}{p} \end{aligned} \tag{3.5}$$

where $V_g = 3\,\mathrm{m^3}$, $y_I = 0.005$, $K_v = 1.4$ and the vent loses 0.3 mole of ammonia per mole of inert. At design $p_b = 127.60\,\mathrm{bar}$ and the inerts hold 12.40 bar. The bubble pressure is a teaching correlation. Against Voskov and Voronin's model, fitted to measured bubble points, it gives 127.8 bar where theirs gives about 135.6 bar at $N/C = 3$, $\mathrm{H_2O/C} = 0.4$ and 185 °C [9]. Its temperature slope, 0.018 per kelvin in $\ln p$, is below their 0.024 (160–180 °C) and 0.028 (180–200 °C), its water slope of -0.30 is three times their -0.10, and its N/C slope of 0.40 is six times their 0.064 from 3.0 to 3.5. Below 3 their bubble pressure falls with N/C (-0.36 from 2.5 to 3.0), consistent with the minimum bubble points of the phase model [13]. The model therefore exaggerates the ammonia's share of a pressure rise and understates the temperature's.

The stripper's liquid holdup drains through a level valve. Its unstripped ammonia (less a 2 percent slip to the evaporator), carbon dioxide and enough water to make a 45 weight percent carbamate solution return to the condenser through a first-order holdup with $\tau_{LP} = 15\,\mathrm{min}$, 21.95 t/h at design. The urea and the remaining water pass a solution tank of $\tau_{sol} = 30\,\mathrm{min}$, a teaching choice that keeps the level loop's transients out of the evaporator. The evaporator balances its heat and its biuret, as shown in Equation (3.6).

$$\begin{aligned} M_e\frac{\mathrm{d}B}{\mathrm{d}t} &= (B_{in} - B)\,\dot{m}_{melt} + M_e k_b \\ k_b &= 0.0585\,\mathrm{e}^{-\frac{E_b}{R}\left(\frac{1}{T_e}-\frac{1}{T_{138}}\right)}\,\frac{1.5}{1+0.5\,n} \\ B_{in} &= 0.30\sqrt{U_0/U}, \qquad w = 0.30\,\mathrm{e}^{-(T_e-138)/4} \end{aligned} \tag{3.6}$$

where B and w are the biuret and moisture in weight percent, $M_e = 8\,\mathrm{t}$, $E_b = 100\,\mathrm{kJ/mol}$, n is the ammonia reaching the evaporator per unit of urea relative to design, and U/U_0 is the urea flow relative to design, with B_{in} capped at three times its design value. The square-root law for the biuret formed upstream is a teaching law, chosen so that the quality window can close at low rate. At design the melt stays 7.66 min and carries 0.748 wt% biuret and 0.300 wt% moisture.

3.3.5 Corrosion Index, Instruments and the Inference

The liner's corrosion is a teaching index on the true N/C at the reactor top, as shown in Equation (3.7).

$$I_c = \mathrm{e}^{(2.8-N/C)/0.05}\sqrt{0.6/\mathrm{O_2}}\;\mathrm{e}^{(T-185)/15} \tag{3.7}$$

The index is 1 at the floor of 2.8 with design passivation air and temperature, and it stops the run once it has stayed at 1 or above for 20 min (the app reports an exercise stop). The floor, the hold and the alarm at $N/C = 3.3$ are teaching values, and real limits come from the licensor's and the liner's corrosion data. The instruments read the CO_2 and ammonia flows, the reactor top's density (lag 0.5 min, noise 0.6 kg/m^3) and temperature (lag 0.3 min, noise 0.08 K), the pressure, the level, the steam flows, the drum pressure and the evaporator temperature. A laboratory samples the reactor outlet's N/C, H_2O/C and conversion and the product's biuret and moisture every 4 h and reports 45 min later, with a standard deviation of 0.01 on the N/C. The density follows the partial molar volumes of free ammonia, carbamate, urea and water, 970.6 kg/m^3 at design, in the spirit of the measured equilibrium densities [12]. The inference and its laboratory correction are given in Equation (3.8), the bias update of Equation (1.11) with $\lambda = \lambda_b$.

$$\begin{aligned} \widehat{N/C} &= c_0 + c_\rho(\rho_m - 970) + c_T(T_m - 185) + b \\ b &\leftarrow b + \lambda\left[N/C_{lab} - \left(\widehat{N/C}_{raw}(t_s) + b\right)\right] \\ \hat{X} &= 0.6041 + 0.1544(\widehat{N/C} - 3) + 0.00431(T_m - 185) \\ &\quad - 0.25(\mathrm{H_2O/C}_{lab} - 0.40) + b_X \end{aligned} \tag{3.8}$$

where $c_0 = 3.018$, $c_\rho = -0.00697$ per kg/m^3, $c_T = -0.01006$ per kelvin, $\lambda = 0.7$, $|b| \le 0.4$, and $\widehat{N/C}_{raw}(t_s)$ is the unbiased inference at the instant t_s the sample was drawn, not at the report. The coefficients are a least-squares fit on 35 lined-out plants of this model at design load and water with N/C from 2.7 to 3.3 and the drum from 3.9 to 5.1 bar a. The fit's worst error is 0.026 over its own plants, and lined-out plants at 70, 85 and 110 % load, with 98 or 99 weight percent ammonia or with steam ratios of 0.95 and 1.30 read within 0.016 (0.0153). At design it reads 3.0137 against a true 3.000, an offset the bias removes, and a gauge 5 kg/m^3 low reads 0.035 high, since the fit gives 14.3 kg/m^3 for each 0.1 of N/C at constant temperature. The fit data reach 215.7 °C, past the correlation's clamp, although no case runs hotter than 192.4 °C.

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

3.3.6 Parameters, Initialization and Numerics

Table 3.2 lists the parameters. Each case starts from the lined-out plant for its own settings, found by running the model with the fast loops held at their setpoints for 2400 min of pseudo-time, integrating the loop's ammonia toward the N/C target, and then by Newton's method with the conserved modes pinned. Every derivative at the design state is at most 4.2×10^{-16} of its state. The integrator is the engine's fourth-order Runge–Kutta (Section 1.6.1) with a step of 2 s, control executes at 6 s intervals and a log row is written at 10 s intervals. Halving the step changes the N/C at 8 h of the challenge by 6.1×10^{-11} as shipped and 2.8×10^{-10} with the suggested settings. Over the challenge's twelve hours the four element ledgers close to better than 2.9×10^{-11} kmol, and the molar masses are built from the atomic weights so that the overall reaction conserves mass exactly.

3.3.7 Assumptions

Carbamate formation is complete in excess ammonia, so free CO_2 appears only below $N/C = 2$, and carbonate, bicarbonate and the non-ideal electrolyte solution are left out. The reactor is three stirred cells with fixed gas shares, while a real reactor has trays, back-mixing and a vapor phase throughout. The stripper is a decomposition law whose steam enters no energy balance, and the condenser is one duty at a fixed process-side temperature. The LP section and the solution tank are first-order holdups. Every correlation (conversion, bubble pressure, density volumes, stripping, condensing, biuret, moisture, corrosion) is a teaching value of the published shape, and the inference is fitted to this model, so its error comes only from what the fit did not vary and from the instruments. The heat of carbamate formation, 100 kJ/mol, is below the 117 kJ/mol quoted at 160 °C and 110 atm, while

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

Symbol	Meaning	Value	
M_R	reactor holdup, three cells	135 t	A
k_{185}, E_d	dehydration rate, activation	0.12 min^{-1}, 90 kJ/mol	A
ΔH_d, ΔH_c	dehydration, carbamate heats	15.5, 100 kJ/mol	A
c_{pL}, τ_{ov}	liquid heat capacity, overflow	2.7 kJ/(kg K), 0.5 min	H
UA_c, T_p	condenser conductance, process side	84.8 MJ/(min K), 168 °C	A
C_s, K_s	drum capacity, LP valve	900 kg/bar, 1200	H
p_h	LP header	3.5 bar a	A
η	stripper efficiency factor	1.00	A
k_g, k_q	stripping exponents	0.5, 1.109	H
—	HP steam supply limit	80 t/h	A
τ_{LP}, τ_{sol}	LP recycle, solution tank	15, 30 min	A
—	water in LP carbamate	45 wt%	A
ε_N	ammonia slip to evaporator	0.02	H
y_I, V_g, K_v	inerts in CO_2, gas space, vent	0.005, 3 m^3, 1.4	H
$p_{b,0}$	bubble pressure at design	124 bar	A
—	ln p_b slopes T, N/C, H_2O/C	0.018, 0.40, −0.30	H
M_e	evaporator melt holdup	8 t	A
B_{in}	biuret upstream at design	0.30 wt%	A
$k_{b,0}$, E_b	biuret rate at 138 °C, activation	0.0585 wt%/min, 100 kJ/mol	A
—	biuret, moisture limits	0.90, 0.50 wt%	A
—	floor, hold, N/C HIGH	2.8, 20 min, 3.3	A

the dehydration's 15.5 kJ/mol is the quoted value [10]. The condenser was sized on the lower value, so changing it would move every run.

3.4 What the Loop Does Before It Is Controlled

3.4.1 The Ratio Integrates

The loop's ammonia is an inventory. A fresh ratio 0.02 too high raises the reactor N/C by 0.0180 per hour at design load, an integrating gain of 0.014 99 per minute per unit of fresh ratio. In case 1 a 2.85 percent step in the ammonia at constant CO_2 takes the N/C from 3.000 to 3.030 in the first 50 min and 3.083 by 110 min, still climbing, as Section 1.4.3 expects of a recycle with no exit for the excess. The density falls with it, from 970.60 to 965.70 and then 956.12 kg/m^3, and the inference follows within 0.024 of the truth over two hours, while the laboratory's first sample is still four hours away. Wet ammonia is the same error with the opposite sign. Ammonia of 98.0 instead of 99.8 weight percent is a shortfall of 0.0363 in the fresh ratio that no flow meter sees.

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

A CO_2 turndown does not drain the ratio. With the fresh ratio fixed on the flows, a turndown to 80 % raises the N/C from 3.000 to 3.021 by 8 h in case 5, because only an error in the fresh ratio integrates into the inventory. That measurement decided the challenge's design. The application's plan expected the turndown to carry the shipped controller to the floor, but it does not, so the challenge adds wet ammonia at 1 h. Without it the shipped flows hold the N/C between 2.989 and 3.034 for twelve hours, and with it the N/C crosses the floor.

> **TRY IT IN THE APP** **3 · Ratio on the Flows Drifts**
>
> Load scenario 3 and press Run, at a faster playback speed if the clock is too slow. From 1 h the ammonia arrives wetter while both flow meters read as before. Watch the inferred N/C on the Ratio view fall below 2.93 at 3 h 42 min, the laboratory's first report arrive at 4 h 45 min reading 2.918, and the run stop at 8 h 26 min.

3.4.2 The Condenser Sets the Reactor's Heat

The reactor has no heater. Its heat is the gas the condenser does not condense, so at a fixed drum pressure a smaller gas flow loses a larger share to the condenser. In case 5 the turndown raises the

condensed share from 82.6 to 86.1 percent, the reactor top falls from 185.0 to 176.6 °C at 1 h 54 min and the conversion from 61.1 to 58.1 percent at 8 h. The HP steam per tonne of urea stays at 0.84 t/t because the stripper steam follows the CO_2, while the LP steam raised per tonne rises from 0.721 to 0.825 t/t, heat sent to the drum instead of the reactor. The plan's expectation that energy per tonne would climb did not happen in the model, and the case teaches the coupling instead.

Pause. Try Exercise 3.2 before reading on.

3.4.3 Bubble Pressure and a Recycle That Stores a Disturbance

The high side of the ratio costs pressure. Left running after case 1's step, the inert vent is wide open 1 h 58 min after the step and PRESSURE HIGH (measured pressure above 143 bar a) shows 2 h 9 min after it, at $N/C = 3.100$ with the top at 187.28 °C. That alarm is half ammonia and half temperature in this model. The bubble pressure rises from 127.59 to 137.98 bar, of which the N/C alone accounts for 5.2 bar and the 2.3 K warmer reactor alone for 5.3 bar. With the bubble-pressure slopes refitted near design to Voskov and Voronin (0.026, 0.07 and −0.10), the alarm would come 3 h 12 min after the step, the suggested challenge's peak pressure would rise from 142.93 to 144.35 bar a, still under its limit, and no verdict would change, so the shipped slopes were kept with the comparison stated.

A stripper steam dip shows the recycle storing a disturbance. In case 6 the supply falls to 40 % of 80 t/h from 30 to 90 min, so the stripper gets 32.0 t/h and strips 64.5 instead of 80.0 percent of the CO_2. The recycle peaks at 46.9 t/h at 1 h 34 min, the conversion bottoms at 54.68 % at 1 h 36 min, and the reactor's H_2O/C rises from 0.406 to 0.619 at 2 h 17 min, all after the steam's return. On the return the stripper decomposes the backlog at once. The gas to the condenser jumps from 73.3 to 85.2 kmol/min, the condenser at its fixed drum pressure condenses 77.1 instead of 83.7 percent, and the heat the gas brings the reactor rises from 376 to 640 MJ/min. The top reaches 192.37 °C at 2 h 26 min and the measured pressure 143.59 bar a at 2 h 18 min, after the vent saturates. The overshoot is physical. With the ratio on the flows it is 191.76 °C, a pressure loop four times faster leaves it at 192.37, a drum loop five times slower raises it to 194.10 °C, and the drum valve on manual to 194.90 °C.

3.5 Control as Implemented

Every controller reads one measurement packet built by the instrument code, holding the meters, the gauges, the laboratory's reports and their quality, and an automated check rejects any packet holding another key. The header's Manual and Auto switch belongs to the N/C loop alone (Manual hands the operator the ammonia feed), while the pressure, level, drum and evaporator loops each have their own switch, so that an integrating level and a fast pressure are not left unattended in case 1. Every controller not acting tracks the one that is, so mode, architecture and cascade changes are bumpless.

3.5.1 The Shared PI Law

All six loops use the framework's velocity-form PID of Section 1.6.2, Equations (1.14) to (1.17), with the derivative off. Each executes every $\Delta t = 6$ s. The loops marked reverse in Table 3.3 have a measurement that falls as their output rises and run with $a = -1$. The running sum is never clamped, and back-calculation pulls the integral toward the limited output with τ_I as the tracking time. An integral time at its slider's maximum removes integral action. The defaults are listed in Table 3.3.

3.5.2 Three Ways to Hold the Ratio

The ammonia command is always a ratio R to the measured CO_2, $F_{NH_3,SP} = R F_{CO_2,m} M_A / (M_C(1 - 0.002))$, where the ratio station assumes the design purity. With *Flows*, R is the operator's number, 2.0173 at design (2.000 for the urea, 0.0072 for the ammonia slip and 0.0101 for the vent), and nothing looks at the reactor. With *Inference*, the AIC of Table 3.3 moves R on $\widehat{N/C}$. With *MPC*, a receding-horizon plan sets R and the stripper steam ratio S every 10 min over 4 h on the model of

Table 3.3 Loops, Pairings and Default Tunings

Loop	PV → MV	K_c	τ_I	Output range
AIC	$\widehat{N/C} \to R$	1.5 per unit	120 min	1.80–2.30
PIC	$p \to$ vent, reverse	2 %/bar	3 min	0–100 %
LIC	$L \to$ bottoms, reverse	1.5 %/%	30 min	0–100 %
PIC	$p_s \to$ LP valve, reverse	20 %/bar	3 min	0–100 %
TIC	$T \to p_{s,SP}$ (cascade)	0.05 bar/K	20 min	3.6–6.5 bar a
TIC	$T_e \to$ heating steam	1.0 %/K	5 min	0–100 %

Equation (3.9).

$$
\begin{aligned}
\frac{\mathrm{d}a}{\mathrm{d}t} &= g\,(R - R_{bal}) \\
\frac{\mathrm{d}X}{\mathrm{d}t} &= \left[X_0 + k_N(a - a_0) + k_S(S - S_0) - X\right]/\tau_X \\
\frac{\mathrm{d}B}{\mathrm{d}t} &= \left[B_0 + k_B(S - S_0) - B\right]/\tau_B \\
\min\ & \textstyle\sum_k \left[w_N(a_k - \mathrm{SP})^2 - 100 w_X X_k + w_S(S_k - 1.15)^2\right] \\
& + w_{\Delta u} \textstyle\sum_j \left[(\Delta R_j/0.01)^2 + (\Delta S_j/0.05)^2\right] \\
& + 10^4 \textstyle\sum \sigma^2 \\
\text{s.t. } & 2.90 - \sigma \le a_k \le 3.20 + \sigma \\
& B_k \le B_{lim} + \sigma, \qquad \sigma \ge 0
\end{aligned}
\tag{3.9}
$$

where a_0 is the inferred N/C, X_0 the conversion soft sensor of Equation (3.8), B_0 the last laboratory biuret, S_0 the present steam ratio, $g = 0.0149 F_{CO_2,m}/45.8$ per minute, and R_{bal}, the ratio that would hold the N/C still, is estimated from the inferred N/C's motion and filtered over an hour. The step-identified constants are $k_N = 0.154$, $k_S = 0.10$ per t/t, $k_B = 0.27$ wt% per t/t, $\tau_X = 60$ min and $\tau_B = 30$ min. A lined-out step of 0.1 t/t in steam gives 1.00 point of conversion and 0.0273 wt% of biuret, which matches k_S and k_B. The model is converted to discrete time exactly under a zero-order hold, the horizon is 24 steps in six move blocks, two of one step and then 2, 4, 8 and 8 steps, R is bounded by 1.80–2.30 and moves at most 0.04 per step, S lies between 0.8 and the lesser of 1.5 and the supply's limit and moves at most 0.05, and the defaults are $w_N = 200$, $w_X = 1$, $w_S = 50$ and $w_{\Delta u} = 1$. The quadratic program goes to the framework's shared solver of Section 1.6.3, a dual active-set algorithm due to Goldfarb and Idnani [19], and the first move is applied, as in industrial practice [20]. A plan that is not optimal, not finite or out of range is rejected, and the tracked inference PI takes the ratio without a bump. The conversion reward pulls the N/C above its setpoint, and the MPC buys conversion with steam until its biuret constraint binds. Before the first laboratory report the MPC's biuret starts from the shipped default of B_{in} plus 0.45 rather than from the engine's parameter, which moves no shipped run.

3.5.3 Cascade and the Three Models

The optional reactor temperature cascade lets a TIC on the measured top temperature set the drum's setpoint between 3.6 and 6.5 bar a, the handle of Section 3.4. The stripper steam follows the measured CO_2 in the ratio S unless the MPC sets it. The plant is the model of Section 3.3. The control models are the inference fitted to that plant, the conversion soft sensor and the MPC's three-state linear model, none of which knows the ammonia's water, a gauge's error or the corrosion index. The exported file carries every measurement beside its truth, and the columns holding the simulator's truth carry a prefix (`truth_`) that keeps them out of any replayed controller. The corrosion stop reads the truth and is a teaching guardrail outside every controller.

Table 3.4 The Nine Cases

#	Case and events	Teaching point
1	Manual, learner steps the ammonia to 36.85 t/h	the ratio integrates
2	Inference, gauge -5 kg/m^3, lab every 2 h	the laboratory bias
3	Flows, ammonia 98.0 wt % from 1 h	drift on the flows
4	Flows from N/C 2.90, learner cuts R to 1.95	the corrosion floor
5	Flows, CO_2 to 80 % at 30 min	the condenser's heat
6	Inference, HP steam 40 % from 30 to 90 min	a recycle stores it
7	Inference, learner moves T_e	biuret against moisture
8	Inference, CO_2 to 55 % at 30 min	a window that closes
9	Challenge, 12 h, four upsets, Flows as shipped	the whole loop

Table 3.5 Case Outcomes Measured on the Engine

Case and condition	Result
1, step at 5 min	N/C 3.030 / 3.083 at 50 / 110 min after; ρ 956.12 kg/m^3; PRESSURE HIGH 2 h 9 min after
2, bias update on	$\widehat{N/C}$ 3.045 at $t = 0$, truth 2.962 lowest; bias -0.035 at 2 h 45 min; truth 2.997 at 4 h 45 min
2, bias update off	truth 2.952 at 8 h, $\widehat{N/C}$ 3.001
3, Flows	$\widehat{N/C} < 2.93$ at 3 h 42 min; lab 2.918 at 4 h 45 min; floor 7 h 49 min; stop 8 h 26 min
3, Inference at 3 h 42 min	R 2.0955 an hour later; truth $\geq$ 2.913; within 0.03 at 5 h 45 min
3, Inference from $t = 0$	truth 2.969–3.020 for 12 h
4, R 1.95 at 10 min	N/C 2.808 at 2 h; $I_c = 1$ at 2 h 21 min; stop 2 h 41 min
5, operator drum	top 176.56 °C; X 58.12 %; biuret 0.858 wt %
5, cascade	drum SP 5.17 bar a; top $\geq$ 181.02 °C; X 63.23 %; biuret 0.919 wt %
6, steam dip	recycle 46.9 t/h; $X \geq$ 54.68 %; top 192.37 °C; p 143.59 bar a
7, T_e 142.5 / 135 / 138.5	biuret 0.918 / 0.660 / 0.763; moisture 0.097 / 0.638 / 0.267 wt %
8, 55 % load	$X \geq$ 51.26 %; residence 13.38 min; biuret 1.163 wt % at 138 °C
8, best T_e	136.0 °C: biuret 0.968, moisture 0.497 wt %

3.6 Nine Cases on the Running Loop

The cases are listed in Table 3.4 and their main measurements in Table 3.5. Every number comes from the application's own engine, run headless with its 2 s step and its default seed and started from each case exactly as it loads, with any setting the text names entered at the time it names (Section 1.6.4 describes the cases, their completion rules and the export). Two independent measuring scripts reproduce every value with identical output. Times are from $t = 0$ unless a step is named.

3.6.1 An Inferred Ratio Against a Late, Accurate One

Case 2 runs the inference with a density gauge that reads 5 kg/m^3 low, so the inference starts at 3.045 against a true 3.000 and the AIC cuts the ammonia until the truth reaches 2.962. The first report, sampled at 2 h, reads 2.970 at 2 h 45 min and moves the bias to -0.0347, and the truth is 2.997 at the second report. With the update off the error stays and the truth ends at 2.952 at 8 h. The update also inherits the laboratory's noise. A third report 0.018 below the truth moves the bias to -0.0575 at 6 h 45 min. In case 3 the flows let the inferred N/C fall below 2.93 at 3 h 42 min and the truth below 2.90 at 4 h 32 min, before the first report. The truth crosses the floor at 7 h 49 min, but the index reaches 1 only at 8 h 6 min ($N/C = 2.791$) because the wet reactor runs at 182.5 °C, and the run stops at 8 h 26 min. Switched to the inference at 3 h 42 min, the AIC raises R to 2.0955 within the hour and holds the lowest truth at 2.913, and run on the inference from the start the truth stays between 2.969 and 3.020. The state of case 3 at 500 min is shown in Figure 3.5.

3.6.2 Heat, Biuret and a Window That Closes

With the cascade in case 5 the TIC raises the drum to 5.17 bar a, the reactor dips only to 181.02 °C and is back within 1 K of 185 °C at 1 h 54 min, and the conversion settles at 63.23 %. The biuret then rises to 0.919 wt%, over its limit, against 0.858 wt% with the operator's drum. The condenser does

Table 3.6 The Challenge Over 12 h at the Default Seed

Configuration	Samples	Low	X (%)	B (wt%)	p (bar a)	Verdict
Flows, operator, 138 (shipped)	2.949, 2.857	2.774	56.53	0.872	140.75	stop
Inference, cascade, 137 (sugg.)	2.965, 3.021, 3.033	2.964	61.32	0.882	142.93	pass
MPC, cascade, 137	2.976, 3.049, 3.070	2.973	61.98	0.896	142.59	pass
Inference, operator, 138	2.976, 3.012, 3.021	2.967	58.40	0.870	140.98	fail
Inference, cascade, 138	2.965, 3.021, 3.033	2.964	61.32	0.922	142.93	fail
Inference, drum 4.8, 137	2.968, 3.009, 3.022	2.967	60.71	0.858	148.94	fail
Flows, cascade, 137	2.937, 2.839	2.775	60.46	0.880	141.20	stop

not reach the evaporator. The turndown alone lengthens the melt's stay from 7.66 to 9.57 min and raises B_{in} through the square-root law, which moves the biuret from 0.748 through 0.783 to 0.895 in the absence of any ammonia effect. With the operator's drum the cooler reactor converts less and 0.1218 kmol/min of unconverted ammonia reaches the evaporator and holds the biuret back to 0.858. With the cascade the hotter reactor converts more and only 0.0881 kmol/min arrives, so the suppression weakens and the biuret reaches 0.919. With the suppression switched off both settings give 0.895, so the cascade's whole difference is that ammonia term.

> **TRY IT IN THE APP** **5 · A CO_2 Turndown and the Reactor's Heat**
>
> Load scenario 5 and press Run. The carbon dioxide drops to 80 % at 30 min, and the reactor top falls from 185.0 to 176.6 °C. Then set Steam drum setpoint from to Reactor T on the TIC card, press Reset and run again: the drum setpoint climbs to 5.17 bar a, the reactor dips only to 181.02 °C, and the biuret rises to 0.919 instead of 0.858 wt%.

Case 7 finds the evaporator's window at design between the 0.50 wt% moisture limit and the 0.90 wt% biuret limit, 136.00 to 142.10 °C. At 142.5 °C the biuret crosses its limit 19 minutes after the change. Case 8 cuts the CO_2 to 55 %. The urea falls from 62.50 to 35.76 t/h, the reactor top to 163.88 °C and the conversion to 51.26 %, and the fixed melt holdup keeps the melt 13.38 min. On the lined-out plants the window narrows to 136.00–136.60 °C at 65 % load and is gone at 60 percent (smallest excess 0.008 wt%) and 55 percent (0.039 wt%). The case is the honest infeasible one. No evaporator setpoint meets both limits, 136.0 °C giving 0.968 and 0.497 after five hours, while a 5 t melt holdup reopens the window (0.772 and 0.438 at 136.5 °C).

> **DID YOU KNOW?** **Making Less Urea Makes Worse Urea**
>
> From the simulator: with the carbon dioxide cut to 55 %, the melt stays 13.38 min in the evaporator and its biuret reaches 1.163 wt% at 138 °C. The best evaporator temperature, 136.0 °C, gives 0.968 biuret with 0.497 wt% moisture, so no setting meets both limits (Table 3.5).

3.6.3 A Twelve-Hour Shift

The challenge brings 98.0 weight percent ammonia at 1 h, the CO_2 to 80 % at 2 h, the stripper steam to 40 % from 6 to 7 h and the CO_2 back at 9 h. It is met when the true N/C lies within 2.90–3.20 at the samples of 4, 8 and 12 h and never falls below 2.80, the average true conversion is at least 59 %, the biuret never exceeds 0.90 wt% and the reactor pressure never exceeds 145 bar a, judged every step on the truth. The 59 percent and 145 bar a are teaching values chosen by measurement. At the plan's 58 percent the inference alone passed by 0.4 points, so the condenser handle did not matter, and without a pressure limit a drum held at 4.8 bar a all shift passes on conversion (60.71 %) while it reaches 148.94 bar a when the load returns. The configurations are compared in Table 3.6 and the trajectories in Figure 3.4. Figure 3.3 sets the table's last four columns against the limits of the challenge.

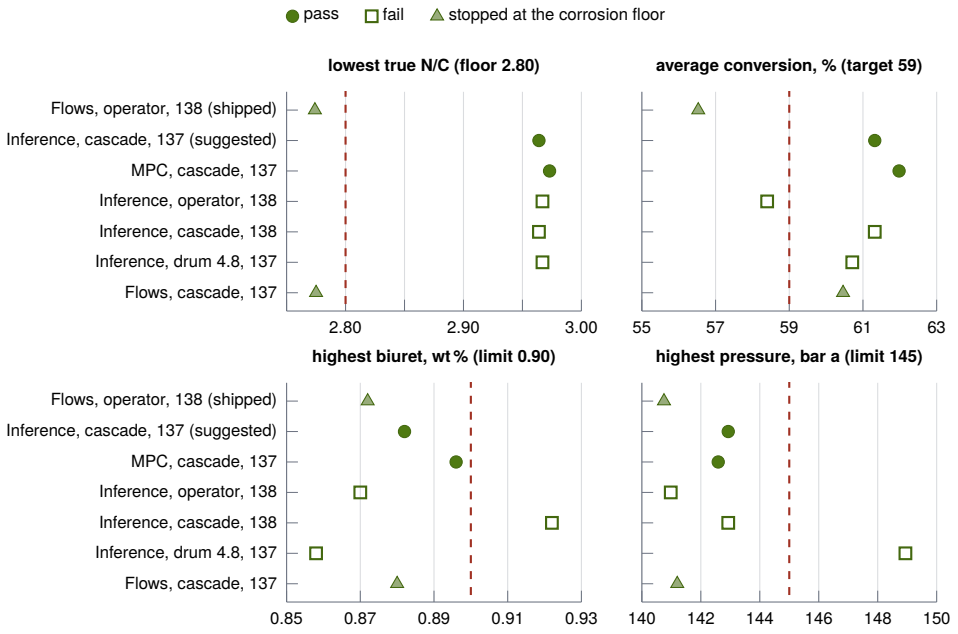

Figure 3.3 Table 3.6 as dots, digit for digit: the lowest true N/C, the average conversion, the highest biuret and the highest pressure of each configuration over the shift, with the floor, target or limit of the challenge dashed. Passing runs are filled circles, failing runs open squares and runs stopped at the corrosion floor filled triangles. Notice that each failing configuration misses on exactly one axis, the conversion, the biuret or the pressure, and that the two runs on the flows are the only ones below the floor.

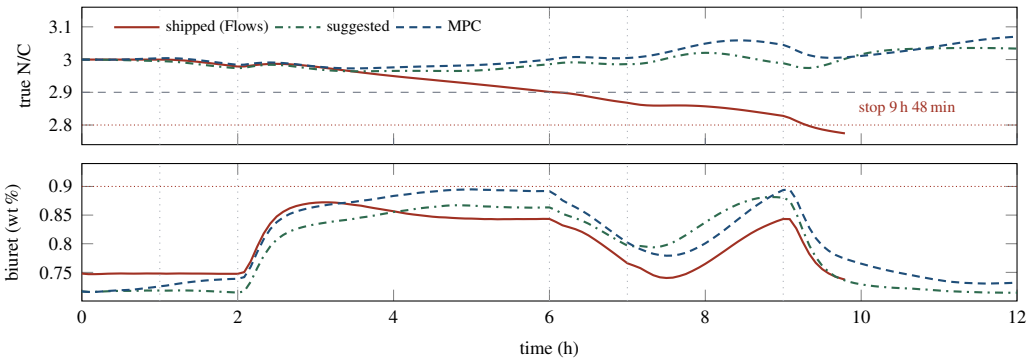

Figure 3.4 The challenge's true reactor N/C and product biuret, shipped (Flows, operator's drum, 138 °C, solid), suggested (Inference, cascade, 137 °C, dash-dotted) and MPC (cascade, 137 °C, densely dashed). Dotted verticals mark the events at 1, 2, 6, 7 and 9 h, dashed lines the 2.90–3.20 band, dotted horizontals the 2.80 floor and the 0.90 wt% limit. Rows every 5 min of the log.

As shipped, the wet ammonia drains the inventory for nine hours, the samples read 2.949 and 2.857, and the run stops at 9 h 48 min. The inference holds the N/C, but with the operator's drum the reactor cools to 174.42 °C through the turndown and the conversion averages only 58.40 %. The cascade raises the drum only while the load is down, which keeps the conversion at 61.32 % and the pressure under 145 bar a when the load returns, and the turndown's longer residence then decides the



Figure 3.5 Case 3 at $t = 500$ min on the Ratio view with the teaching overlay on Show truth and an eight-hour trend window, in Auto with the ratio on the flows, cropped to the Ratio view and two trend strips. Wet ammonia since 1 h has drained the loop. The badges (1) read the inferred N/C 2.777, the last laboratory report 2.918 (sampled at 4 h), the corrosion index 1.14, an estimated conversion of 55.5 %, the reactor top at 182.3 °C and 0.73 wt% biuret, and the status strip (4) N/C LOW, CORROSION, 14 OF 20 MIN. On the conversion plot (2) the inferred point and the truth's × (2.784) sit inside the shaded floor band, and the inference plot (3) places the gauge reading of 1007.5 kg/m^3 beyond the 2.8 line with a bias of −0.007. The trends show the inferred and true N/C falling together through 2.80, the laboratory's single report (5) flat and late, and the fresh ratio unmoved at 2.0173.

evaporator setpoint, 0.922 wt% at 138 °C and 0.882 wt% at 137 °C. Each handle answers one upset. The suggested settings pass at seeds 0 to 19, the default seed and half the step, 22 of 22 runs, with the highest biuret between 0.8807 and 0.8857 wt%, a margin of at least 0.014 wt%. The MPC also passes all 22, but only just. Its highest biuret lies between 0.8946 and 0.8999 wt%, a margin of 0.0001 to 0.0054 wt%, because the biuret is a soft constraint on its own linear model fed by a laboratory reading, and the conversion reward spends stripper steam until that constraint binds. Its samples run higher, 3.049 and 3.070 at the default seed, as the reward pulls the N/C up, and it solves 71 plans with none rejected at seeds 0 to 3, at the default seed and with the step halved. The MPC was left as built, with its margin stated, because holding a soft constraint at its limit is what a reward on the other side of it produces.

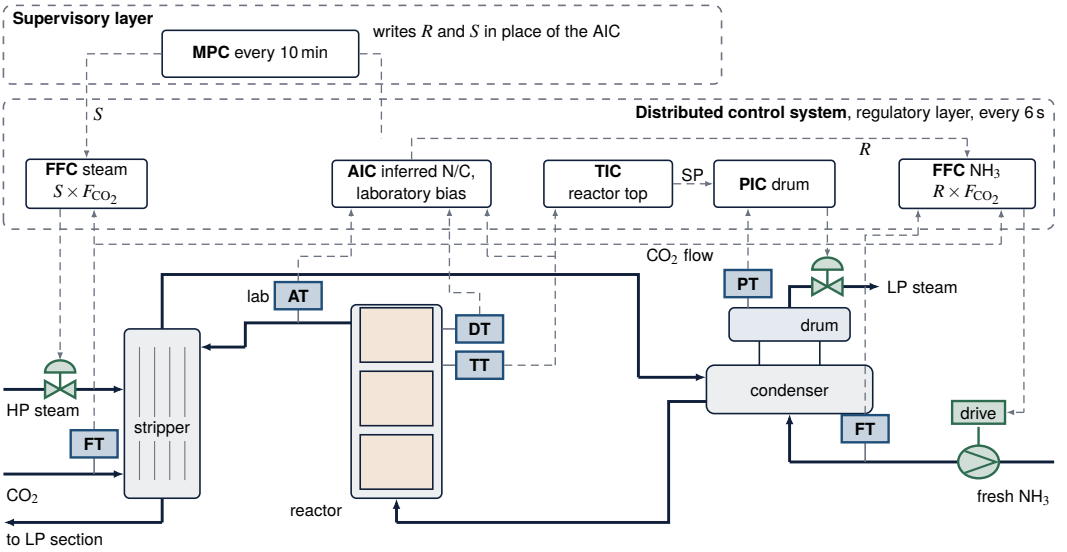

Figure 3.6 The ratio loop and the reactor's heat loop with their hardware. Lettered boxes are sensing devices: flow transmitters on the carbon dioxide and the fresh ammonia (FT), the density and temperature at the reactor top (DT, TT), the drum pressure (PT) and the laboratory sample (AT). Valve symbols and the pump with its drive are final elements, and dashed arrows are signals. Notice that nothing measures the ratio itself: the AIC reads a density, a temperature and a late laboratory result, and the ratio station multiplies its R by the measured carbon dioxide. The reactor's temperature is moved through the drum pressure, a cascade, and the predictive controller writes R and S in place of the AIC. The vent, level and evaporator loops, the LP section and the carbon dioxide compressor are left out. Not to scale, and not a piping diagram.

3.7 Technology in Practice

What Measures It

Figure 3.6 draws the loops of Section 3.5 with their devices. The feed flows, the pressures, the level and the temperatures are read continuously. The ratio that matters has no transmitter. The inference reads a density and a temperature at the reactor top, and the laboratory corrects it every 4 h, 45 min after each sample. Offline analysis is accurate but late, and drawing a sample of hot liquid at synthesis pressure is itself a hazard, which is why plants prefer in-line methods [21]. The model's density lag of 0.5 min and noise of 0.6 kg/m^3 stand for the gauge below.

> **TECHNOLOGY** **Sensor: A Density Gauge on the Reactor Outlet**
>
> No instrument reads the reactor's ratio directly, so plants infer it. A 2026 study lists vibrating-element density sensors, Raman spectroscopy, gas chromatography and laboratory samples, and notes that sensors in contact with the carbamate corrode and foul, so their calibration drifts. A gamma-ray gauge avoids contact: a source and a detector sit on opposite sides of the outlet pipe, and the density of the liquid sets how much radiation gets through. The same study recalibrates it with process water at shutdowns, to remove drift in the pipe, the detector and the decaying source [21]. On a trend, a sensor drift is a slow offset that looks like a change in the process [2]. The gauge 5 kg/m^3 low of case 2 is one.

What Moves It

Two ratios and one pressure move this loop. The ratio station sets the fresh ammonia in proportion to the measured carbon dioxide, and the stripper steam follows the carbon dioxide in its own ratio. In a

stripping plant of this kind the fresh carbon dioxide enters through the high-pressure stripper [17], as it does in Figure 3.1. The reactor has no heater, so its temperature is moved by the pressure of the steam drum on the carbamate condenser, through the low-pressure steam valve. The same handle appears in the reference document's pool reactor, a design in which the pressure of the steam raised controls the rate of condensation [17].

> **TECHNOLOGY Final Element: A Pump at Synthesis Pressure**
>
> The fresh ammonia reaches synthesis pressure through high-pressure pumps [17]. A pumped flow can be moved in two ways. A control valve throttles it, or a variable-frequency drive changes the speed of the motor by changing the frequency of its power, which saves energy and removes a component. A variable-speed machine, however, often cannot change the flow as fast as a valve, cannot stop the flow positively and offers fewer fail-safe states [22]. Either way, the ratio station of Section 3.5 only hands the ammonia flow its setpoint, and the seven actuator positions among the model's 39 states stand in for such hardware.

Where the Controller Runs

A plant runs such loops in a distributed control system, in which a supervisory level oversees local controllers and sends them set points, loops are sometimes nested or cascaded, and cycle times range from milliseconds to minutes [23]. The app executes its six PI loops every 6 s and its predictive controller every 10 min, so both sit inside that range. The ratio stations, the AIC with its laboratory bias and the cascade from the reactor top onto the drum belong to the regulatory layer. The plan belongs to the supervisory layer above it and writes the ammonia ratio and the steam ratio. When a plan is rejected, the tracked inference PI takes the ratio back without a bump.

> **IN PRACTICE The Operating Window and the Biuret Limit**
>
> A 2007 European reference document gives the window urea reactors work in: 140 to 250 bar, 180 to 210 °C and an ammonia to carbon dioxide ratio of 2.8 to 4, for a conversion of 50 to 80 percent of the carbon dioxide. Stripping with carbon dioxide, as this chapter's plant does, is now the most frequent type. Fertilizer urea should preferably carry at most 1.2 percent biuret, which can damage crops sprayed on their leaves, and urea solutions low in ammonia should not stay long at high temperature [17]. The chapter's biuret limit of 0.90 wt% and the evaporator window of case 7 are teaching values of the same trade.

What Is Modernizing It

Three developments bear on this loop. Revamps of older plants have added continuous monitoring of the ratio, along with relined reactors and stripping equipment [17], and gamma-ray gauges, which never touch the carbamate, have gained attention for reading it [21]. Combining condensation and reaction in one vessel, a pool reactor, makes the synthesis much less sensitive to deviations of the ratio and removes the loop's inverse response [17], which would change the dynamics that Section 3.4 measures. And data-driven soft sensors, of which the chapter's least-squares inference is a simple case, are spreading through the process industries, with adaptive versions that update their models as the process drifts [2]. The open challenge after the Limitations takes up what they still lack.

3.8 Limitations

The sizes, rates and limits were chosen to teach, for a generic plant. The conversion correlation has the published shape and lies within about 3 points of two published sources but 4 to 9 points above a third near $N/C = 3$, the bubble pressure overstates the N/C slope sixfold and understates the temperature's, the heat of carbamate formation is 17 kJ/mol below the quoted value, and the stripping, condensing, biuret, moisture and corrosion laws are closures chosen for their shape. The square-root biuret law and the 30 min solution tank were added so that the evaporator's lessons work, and the condensed fraction's exponential form was chosen to remove an oscillation that the clipped form produced. The urea-per-kilogram maximum moves with the correlation, 2.98 at the design's water

and temperature but 2.56 at the conditions of the published curve that peaks near 3. The cases also idealize their upsets. A clean step in ammonia purity or a gauge error that appears at once would not happen on a plant, but each isolates one path into the inventory.

The inference is fitted to the same model it estimates, so its 0.026 worst error is a floor that a real gauge on a real solution would not reach. The reactor is three stirred cells with no vapor phase, the stripper has no energy balance, and the corrosion stop is a guardrail that reads the truth, not a model of a corroding liner. The MPC's challenge pass rests on a biuret margin of thousandths of a weight percent, and its model is identified once on this plant and never updated.

OPEN CHALLENGE **An Inference That Stays Right**

A 2009 review names maintenance as a main open issue of soft sensors in the process industry. After launch, the predictions gradually deteriorate as the process changes, until the model is retrained or rebuilt. Drift comes from the process, such as the purity of the input materials, or from the measuring devices, and telling the two apart is hard and still needs expert judgment. Most soft sensors had no automated maintenance and often no objective measure of their quality [2]. Here one bias, moved by a laboratory report every four hours, does all of that work. Progress would be an inference that separates a drifting gauge from a changing process on its own and states how far it can be trusted.

3.9 Conclusion

A CO_2-stripping urea loop holds its most important variable, the N/C in the reactor, as an inventory that no instrument reports, and this simulator logs the true ratio, the inference from density and temperature and the laboratory's reports side by side from element balances that close to 3×10^{-11} kmol. Two weight percent of water in the ammonia carries a ratio held on the flows from 3.000 to the 2.80 floor in 6 h 49 min, while a CO_2 turndown raises it slightly. A PI on the inferred ratio with a laboratory bias holds the truth between 2.969 and 3.020, and the bias removes a 5 kg/m^3 gauge error within two reports. The condenser, not a heater, sets the reactor temperature, 11.6 K per bar of drum pressure near design, so a turndown cools the reactor from 185.0 to 176.6 °C and a cascade onto the drum restores it, at the cost of biuret through the ammonia that no longer reaches the evaporator. The suggested settings pass the twelve-hour shift at 22 of 22 runs and the MPC passes only by thousandths. A fuller model would fit the teaching correlations to the published equilibria and give the reactor a vapor phase. The simulator shows that a ratio held between two flow meters is a ratio of the meters, not of the process.

KEY POINTS

- No instrument reports the reactor's ratio of ammonia to carbon dioxide, and a ratio held between two flow meters is a ratio of the meters.
- Two weight percent of water in the ammonia carries a ratio held on the flows to the 2.80 corrosion floor.
- A PI on the ratio inferred from density and temperature, with a laboratory bias, holds the truth between 2.969 and 3.020.
- The condenser, not a heater, sets the reactor temperature, and a cascade onto the steam drum restores it at the cost of biuret.
- The suggested settings pass the twelve-hour shift in 22 of 22 runs, the predictive controller only by thousandths.

Where else this applies. Volume 5, Chapter 10 corrects a soft sensor from a delayed laboratory in the same way, and Chapter 2 holds a feed ratio in another loop of unconverted reactants.

References

[1] 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.

[2] 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.

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

[4] M. Dente, S. Pierucci, A. Sogaro, G. Carloni, and E. Rigolli, “Simulation program for urea plants,” *Computers & Chemical Engineering*, vol. 12, no. 5, pp. 389–400, 1988.

[5] M. A. Isla, H. A. Irazoqui, and C. M. Genoud, “Simulation of a urea synthesis reactor. 1. Thermodynamic framework,” *Industrial & Engineering Chemistry Research*, vol. 32, no. 11, pp. 2662–2670, 1993.

[6] H. A. Irazoqui, M. A. Isla, and C. M. Genoud, “Simulation of a urea synthesis reactor. 2. Reactor model,” *Industrial & Engineering Chemistry Research*, vol. 32, no. 11, pp. 2671–2680, 1993.

[7] M. Hamidipour, N. Mostoufi, and R. Sotudeh-Gharebagh, “Modeling the synthesis section of an industrial urea plant,” *Chemical Engineering Journal*, vol. 106, no. 3, pp. 249–260, 2005.

[8] X. Zhang, S. Zhang, P. Yao, and Y. Yuan, “Modeling and simulation of high-pressure urea synthesis loop,” *Computers & Chemical Engineering*, vol. 29, no. 5, pp. 983–992, 2005.

[9] A. L. Voskov and G. F. Voronin, “Thermodynamic model of the urea synthesis process,” *Journal of Chemical & Engineering Data*, vol. 61, no. 12, pp. 4110–4122, 2016.

[10] M. Brouwer, “Thermodynamics of the urea process,” UreaKnowHow.com Process Paper, Jun. 2009.

[11] S. Kawasumi, “Equilibrium of the CO_2-NH_3-urea-H_2O system under high temperature and pressure. I. equilibrium pressure of reaction in urea-synthesis from ammonia and carbon dioxide,” *Bulletin of the Chemical Society of Japan*, vol. 24, no. 3, pp. 148–151, 1951.

[12] J. N. Broers, S. M. Lemkowitz, and P. J. van den Berg, “Densities of urea–ammonia–water–carbon dioxide solutions in chemical equilibrium at and above urea synthesis conditions,” *Journal of Applied Chemistry and Biotechnology*, vol. 25, no. 10, pp. 769–779, 1975.

[13] S. M. Lemkowitz, G. A. M. Diepen, and P. J. van den Berg, “A phase model for the gas–liquid equilibria in the ammonia–carbon dioxide–water–urea system in chemical equilibrium at urea synthesis conditions. I. Theory,” *Journal of Applied Chemistry and Biotechnology*, vol. 27, no. 1, pp. 327–334, 1977.

[14] P. M. Glibert, J. Harrison, C. Heil, and S. Seitzinger, “Escalating worldwide use of urea – a global change contributing to coastal eutrophication,” *Biogeochemistry*, vol. 77, no. 3, pp. 441–463, 2006.

[15] F. Wöhler, “Ueber künstliche bildung des harnstoffs,” *Annalen der Physik*, vol. 88, no. 2, pp. 253–256, 1828.

[16] C. Bosch, “The development of the chemical high pressure method during the establishment of the new ammonia industry,” Nobel Lecture, 21 May 1932, 1932, nobelprize.org/uploads/2018/06/bosch-lecture.pdf, read November 2026.

[17] European Commission, “Reference document on best available techniques for the manufacture of large volume inorganic chemicals: Ammonia, acids and fertilisers,” European Commission, Integrated Pollution Prevention and Control, Tech. Rep., Aug. 2007, eippcb.jrc.ec.europa.eu, read November 2026.

[18] U.S. Geological Survey, *Mineral Commodity Summaries 2025*. Reston, VA: U.S. Geological Survey, 2025, nitrogen (fixed)–ammonia, pp. 128–129.

[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] 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] S. Z. Islami Rad and R. Gholipour Peyvandi, “A novel calibration method for online determination of nitrogen-to-carbon ratio in urea synthesis,” *Scientific Reports*, vol. 16, no. 1, p. 36571, 2026.

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

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

3.10 Exercises

3.1 Concept check · 2 min

In case 3 the ratio is held on the flows: the fresh ammonia follows the measured carbon dioxide in the fixed ratio 2.0173, and the ratio station assumes ammonia of 99.8 wt%. From 1 h the ammonia arrives at 98.0 wt%, and both flow meters keep reading their usual values. What does the true N/C at the reactor top do?

(a) It falls by about two percent within the hour the liquid stays in the reactor, then holds

(b) It falls steadily for hours, because the loop's ammonia integrates the shortfall

(c) It stays near 3.000, because the ratio station already corrects for the water it meters

(d) It rises, because the water dilutes the carbon dioxide more than the ammonia

3.2 Concept check · 2 min

In case 5 the carbon dioxide compressor is turned down to 80 % with the ratio on the flows and the steam drum held at 4.50 bar a by the operator's setpoint. The reactor top falls from 185.0 to 176.6 °C. Why does the reactor cool?

(a) The stripper steam is cut with the carbon dioxide, so less HP steam per tonne of urea reaches the loop

(b) The N/C falls with the turndown, and a leaner liquid converts less and releases less heat

(c) The condenser at a fixed drum pressure condenses a larger share of a smaller gas flow, leaving less gas to heat the reactor

(d) The liquid stays longer in the reactor, so the endothermic dehydration takes more of its heat

3.3 Concept check · 1 min

The laboratory samples the reactor at 2 h and reports at 2 h 45 min. Against which value of the inference does the bias update compare the laboratory result?

(a) The inference at the report, since that is when the result is known

(b) The average of the inference over the 45 min the analysis takes

(c) The inference at the next control execution after the report

(d) The inference at the instant the sample was drawn

3.4 Theory and calculation · 6 min

At design load a fresh ratio 0.02 too high raises the reactor N/C by 0.0180 per hour. In case 3 the ratio station holds the fresh ratio at 2.0173 and assumes ammonia of 99.8 wt%, and from 1 h the ammonia is 98.0 wt%. Give each answer to two significant figures.

(a) Find the integrating gain of the loop in N/C per minute per unit of fresh ratio. *Checkpoint:* 0.015 per minute.

(b) Find the shortfall in the fresh ratio, in moles of ammonia per mole of carbon dioxide, that the wet ammonia causes. *Checkpoint:* About 0.036.

(c) Estimate how long after 1 h the N/C takes to fall from 3.000 to the 2.80 floor. The run measures 6 h 49 min. Give one reason from the chapter's measurements why the estimate is earlier. *Checkpoint:* About 6.1 h. The holdups delay the start of the ramp.

3.5 Theory and calculation · 6 min

The corrosion index of Equation (3.7) stops the run once it has stayed at 1 or above for 20 min. Give N/C values to three decimals.

(a) Evaluate the index at N/C 2.90 with the design passivation air of 0.6 vol% and a reactor top at 185 °C. *Checkpoint:* 0.135.

(b) In case 3 the wet reactor runs at 182.5 °C. At what N/C does the index reach 1 with design air? Compare with the chapter. *Checkpoint:* 2.792, against 2.791 measured at 8 h 6 min.

(c) The passivation air is halved to 0.3 vol% with the top at 185 °C. Find the index at N/C 2.80 and the N/C at which it reaches 1. *Checkpoint:* 1.41, and 2.817.

3.6 Theory and calculation · 7 min

The inference of Equation (3.8) has $c_0 = 3.018$, $c_\rho = -0.00697$ per kg/m^3 and $c_T = -0.01006$ per kelvin, with $\lambda = 0.7$.

(a) At the design point the top reads 970.6 kg/m^3 and 185.02 °C. Find the unbiased inference to three decimals, and its offset from the true 3.000. *Checkpoint:* 3.014, an offset of 0.014.

(b) In case 2 the density gauge reads 5 kg/m^3 low. Predict the inference at $t = 0$, and find the density change the fit expects for a change of 0.1 in N/C at constant temperature. *Checkpoint:* 3.048, and 14.3 kg/m^3.

(c) The first report in case 2 reads 2.970. When that sample was drawn the unbiased inference read 3.0195, and the bias was 0. Find the new bias, and the bias a gain of $\lambda = 1$ would give. *Checkpoint:* -0.0347, and -0.0495.

3.7 App challenge · 9 min

Load case 1, Step the ammonia (Manual), which opens with the header switch on Manual, and change no other setting. Run at 10x with the default seed, pause at about 10 min, set the Ammonia feed on the AIC card from 35.83 to 36.85 t/h and run to 150 min. Export the CSV file, whose time column is in hours, and find the step in column `nh3_cmd_tph`.

(a) Read the true N/C, column `truth_nc`, and the true density, column `truth_rho_kgm3`, 50 min and 110 min after the step. *Checkpoint:* 3.030 and 3.083 within 0.002, 965.7 kg/m^3 and 956.1 kg/m^3 within 0.2 kg/m^3.

(b) The step is 0.0574 in the fresh ratio. From the slope between the two readings, find the integrating gain in N/C per minute per unit of fresh ratio, and compare it with the chapter's 0.01499. *Checkpoint:* 0.0156 within 0.0003, about 4 % above.

(c) Find how long after the step the measured pressure, column `p_meas_bar`, first exceeds 143 bar a, which is when the status strip shows PRESSURE HIGH. *Checkpoint:* 2 h 9 min within 2 min.

3.8 App challenge · 8 min

Load case 2, The inference and the laboratory: the ratio on the inference, the density gauge 5 kg/m^3 low and the laboratory every 2 h. Set the Teaching overlay on the AIC card to Show truth, run to 8 h at 60x with the default seed and export the CSV file, whose time column is in hours.

(a) Read the lowest true N/C, column `truth_nc`, before the first report, and the bias, column `nc_bias`, on the row where `lab_reports` first reads 1. *Checkpoint:* 2.962 within 0.002, and -0.0347 within 0.0005.

(b) Read the true N/C on the row where `lab_reports` first reads 2. *Checkpoint:* 2.997 within 0.002.

(c) Switch to the Advanced level, set Laboratory bias update to Off, press Reset, run to 8 h again and export. Read the true and the inferred N/C at 8 h, and explain their difference. *Checkpoint:* 2.952 and 3.001 within 0.002.

3.9 App challenge · 9 min

Load case 9, the twelve-hour challenge, and keep the default seed. Run each configuration for 12 h at 60x, and read the verdict from the So far tile, whose hover text lists the samples, the lowest N/C, the average conversion, the highest biuret and the highest pressure, and from the note at 12 h.

(a) Run as shipped (ratio on Flows, steam drum setpoint from the Operator, evaporator at 138 °C). When does the run stop, and what do the samples at 4 h and 8 h read? *Checkpoint:* A corrosion stop at 9 h 48 min, samples 2.95 and 2.86.

(b) Set Ratio control to Inference, Steam drum setpoint from to Reactor T and the Evaporator temperature setpoint to 137.0 °C, press Reset and run. Read the three samples, the average conversion, the highest biuret and the highest pressure. *Checkpoint:* 2.97, 3.02 and 3.03; 61.3 %; 0.88 wt%; 142.9 bar a; target met.

(c) Raise the evaporator setpoint to 138.0 °C, press Reset and run again. Which part of the target fails, and why does the evaporator setpoint matter only with the cascade on? *Checkpoint:* The biuret, 0.92 wt%.

3.10 Application area · 4 min

In case 8 the carbon dioxide load is cut to 55 %, the melt holdup of the evaporator stays at 8 t, and no evaporator setpoint meets both the 0.50 wt% moisture limit and the 0.90 wt% biuret limit.

(a) The moisture law is $w = 0.30\,\mathrm{e}^{-(T_e-138)/4}$ in weight percent. Find the lowest evaporator temperature that meets the moisture limit, and say why it is the same at every load. *Checkpoint:* 136.0 °C, because the law has no flow in it.

(b) The melt stays 13.38 min at 55 % load. How long would it stay with a 5 t holdup, and why does that reopen the window? *Checkpoint:* About 8.4 min.

3.11 Application area · 4 min

The ratio runs on the inference with the laboratory every 4 h. A plant manager proposes sampling every 8 h, since the inference already holds the N/C. Using the chapter's measurements, say what the laboratory does for the loop, what the change would risk, and what you would recommend.

3.12 Review the agent · 5 min

An AI assistant reviews Table 3.6 and writes: "Use the MPC for the shift. It passes like the suggested settings, its average conversion is higher (61.98 against 61.32 %) and its samples sit higher in the band (3.049 and 3.070), so it has more margin on every part of the target." Find the flaw, and say how you would check the claim.

CHAPTER 4

Four Catalyst Beds, an Ignition Floor and Double Absorption in a Contact Plant

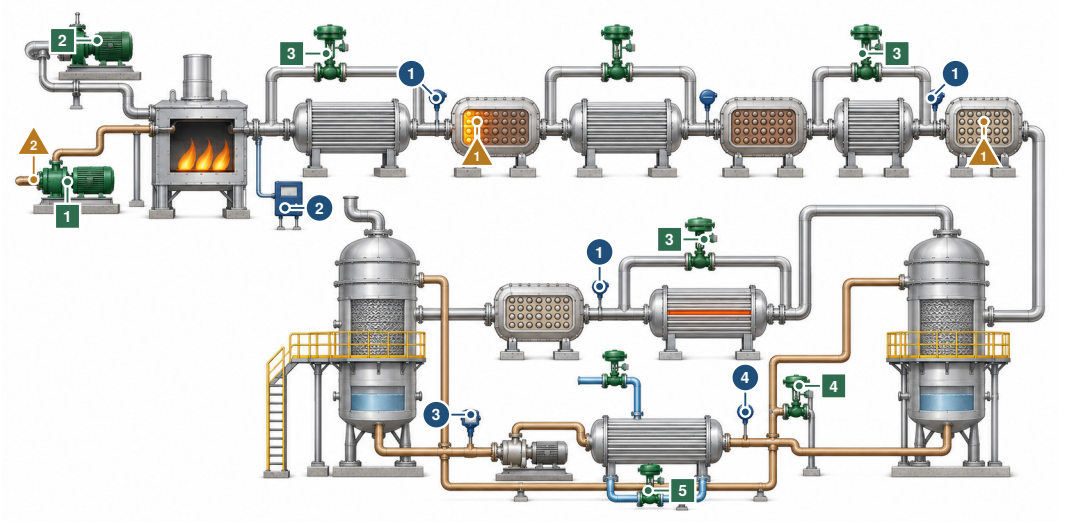

The controller moves 1 sulfur pump, which sets the rate 2 air blower, in ratio to the sulfur 3 gas bypass valves around the coolers, one per bed 4 dilution water valve 5 acid cooler water valve

The controller reads 1 temperature at each bed inlet 2 sulfur dioxide analyzer on the gas to the beds 3 acid strength analyzer 4 acid temperature

What disturbs it 1 catalyst that ages 2 sulfur richer than its meter counts

Notice that each catalyst bed has a cooler with a bypass and a temperature head ahead of it, and that the absorber after the third bed lets the fourth start on fresh gas. Illustration, not to scale.

IN PLAIN TERMS The Coolest Gas That Still Lights

The main idea. A sulfuric acid plant burns sulfur and passes the gas through four layers of catalyst, with coolers in between. Each layer heats the gas as it reacts, and hot gas reacts less fully, so each layer stops short.

The challenge. The coolest gas the catalyst can still light makes the most acid, but a few degrees colder the layer goes out. When the plant slows down, its coolers chill the gas more, and layers go out one by one unless hot gas skips the coolers.

Why it matters. Gas that does not react leaves by the stack, and the law limits how much may leave.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	the first page; the fact box of Section 4.1; the history box of Section 4.2; the try-it boxes of Sections 4.4 and 4.6; the fact box of Section 4.6; the key points
🕮	**Study**	45 min	the whole chapter with its pause points and the exercises
⛰	**Practice**	10 min	the first page; Section 4.6 with its result chart; Section 4.7; Section 4.8 with its open challenge

The control question Each adiabatic bed of a contact plant climbs a straight line until it meets the equilibrium curve, so the coolest inlet that still lights the catalyst converts the most. How close to that ignition floor can inlet loops, an air ratio, an analyzer trim and a predictive layer hold four beds through a turndown, a richer sulfur and an aged catalyst?

Learning objectives

- Compute the adiabatic line of a bed from its feed and heat-capacity flow, and read where it meets the equilibrium curve.
- Explain how an ignition floor emerges from a two-regime rate law and the bed's heat balance, and why the break of the rate law is not the floor.
- Predict the wrong-way rise and the creeping front from the heat stored in the catalyst, and estimate how fast a front crosses a bed.
- Configure the inlet loops and the strength trim to pass a turndown day, and say what each loop protects.
- Relate a stack reading in ppmv to a limit per tonne of acid.

Prerequisites Sections 1.4.2 (adiabatic lines, interbed cooling, removing the product), 1.4.6 (ratio with a trim), 1.6.2 (velocity form) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive layer, and Chapter 2 shows a converter of quench beds.

The app https://apmonitor.com/apps/sulfuric-acid-plant/

See also Volume 7, Chapter 2 for ignition and extinction in a cooled stirred tank. Volume 9, Chapter 8 for a supervisory layer that chooses the setpoints PI loops hold. Volume 6, Chapter 8 for a plant whose loops must be closed before production can be set, with sampled analyzers.

4.1 Introduction

Every point of conversion in the first bed of a contact plant heats its gas by 3.20 K. In the generic plant simulated here, 13.66 t/h of sulfur burned with 1076 mol/s of dry air gives a gas of 11 % sulfur dioxide that enters bed 1 at 420 °C and leaves it at 630.6 °C with 65.77 % of its SO_2 oxidized, which is where that straight line meets the equilibrium curve. Three more beds, with cooling between them and an absorber that removes the trioxide after the third, raise the conversion to 99.837 % and leave 214.5 ppmv of SO_2 in the stack. Without the interpass absorber the same four beds stop at 97.910 %.

The process matters because of its size and its stack. More sulfuric acid is produced than any other chemical, about 150 million tonnes worldwide in 1997 by one estimate, and the European Union had 95 producing plants in 2004 [1]. Anthropogenic SO_2 emissions peaked worldwide in the early 1970s and then fell until 2000 [2], the same years in which the United States began to regulate acid plants built after August 17, 1971 [3]. The converter is analyzed on a diagram of conversion against temperature, on which equilibrium is a curve and each adiabatic bed a heatup path that climbs toward it [4, 5]. That diagram is the reason a converter has several beds, and the reason its inlet temperatures are its most watched variables. Section 1.4.2 draws it for the running example of Chapter 1, and the ammonia converter of Chapter 2 climbs the same kind of lines with quench gas between its beds.

DID YOU KNOW? World Sulfur, 2024

World sulfur production was an estimated 85 million metric tons in 2024, most of it a result of processing fossil fuels, as in the refinery unit of Volume 11, Chapter 13. In the United States about 90 percent of the sulfur consumed that year was in the form of sulfuric acid [6].

The simulator corrects the belief that a hotter catalyst always converts more. On an exothermic equilibrium the coolest inlet that still lights the catalyst converts the most, so the operating point sits near a floor below which the bed goes out, and the floor is not a number printed on the catalyst. In this model bed 1 lined out at the design gas converts 62.8 % with a 395 °C inlet and 3.7 % at 392.5 °C, while halving its catalyst raises that edge to about 400 °C. A second belief is that cooling a bed's

inlet cools the bed. Cooled to 385 °C, bed 1's outlet first rises to 653.4 °C, the wrong-way behavior of packed beds [7], and the bed goes out 32 minutes later. The class is wider than sulfuric acid. Any adiabatic packed bed that carries a reversible exothermic reaction shares the lines, the steps and the extinction, the transient rise after a cut in feed temperature is a property of packed beds in general [7], and a fixed bed whose inlet falls below ignition can keep converting on the heat stored in its catalyst while a reaction zone creeps toward the outlet [8].

🕮 **Pause.** Try Exercise 4.1 before reading on.

However, the sulfuric acid literature treats these parts apart. Plant models are built to optimize operation, dynamically or under uncertainty [9, 10], kinetic studies measure rates on an industrial vanadium catalyst [11], and the creeping and wrong-way transients are analyzed on idealized single beds [7, 8]. This chapter puts the parts together on a browser simulator of a four-bed double-absorption plant, which logs its cell temperatures, conversions and stack beside every instrument. It measures how inlet temperature loops, an air ratio, a strength analyzer trim and a predictive controller hold conversion and emissions through a turndown, a richer sulfur, an aged catalyst and a warm start-up. The rate law is a teaching law of the published shape, and the thermochemistry and the controllers are standard. What the chapter adds is measurement on one synthetic plant, among it the finding that an ignition floor near 395 °C and a creeping front at a 289 °C inlet both follow from the same heat balance. The model has 117 states, its sulfur and oxygen balances close to 2×10^{-15}, it is built on a packed-bed module that other applications share, and no controller sees anything but the instruments. Nine graded cases credit a learner's actions only when the plant responds to them.

Section 4.2 traces the process, its catalyst and its stack limit. Section 4.3 builds the plant: the rate law, the equilibrium, the beds as cells, the burner, exchangers and absorbers, and the instruments. Section 4.4 explores it uncontrolled, from what double absorption buys to the ignition floor, a creeping front and the wrong-way rise. Section 4.5 adds the seven loops and the predictive layer, and Section 4.6 runs the cases, ending with the graded turndown day. Section 4.7 names the hardware and the practice behind the model. Section 4.8 states the limitations and Section 4.9 concludes. The exercises at the end of the chapter run the same cases.

4.2 Contact Process, Melt Catalyst and the Stack Limit

The history box at the end of this section traces the contact process from 1901 to the catalyst used today, vanadium pentoxide in a melt at operating temperature. Cesium sulfate lowers the melting point so that the catalyst can be used cooler, and the lower range of industrial practice is 410 to 440 °C for conventional catalysts and 380 to 410 °C for cesium-promoted ones, against an upper range of 600 to 650 °C above which activity can be lost for good [1].

Design moved from one absorption to two, and the change from single to double absorption is one of the techniques the European reference document lists [1]. In a double contact plant the first catalysis stage converts 85 to 95 % of the SO_2 before an intermediate absorber, the SO_3 formed so far is absorbed in 98.5 to 99.5 % acid, and the removal shifts the equilibrium so that one or two further beds convert most of the rest [1]. The four-bed arrangement with three beds before the intermediate absorber is the 3+1 configuration of the European reference document, and five beds in 3+2 is described there as the normal one [1]. The equilibrium curves and the heatup paths of the stages are the subject of whole chapters of a modern text on the manufacture [4, 5].

The stack limit came with the regulation of new sources. The United States standard of performance applies to every sulfuric acid production unit that commenced construction or modification after August 17, 1971, and limits the SO_2 discharged to 2 kg per metric ton of acid produced as 100 % H_2SO_4 [3]. The rule converts a stack monitor's reading to that unit with a factor computed from the SO_2 entering the converter, $CF = k(1-0.015r)/(r-s)$ with $k = 0.0653$ and r and s the inlet and stack percentages [3], which puts 2 kg/t at 402.0 ppmv for an 11 % gas. The European reference document of 2007 associates best available techniques for an existing sulfur-burning double-absorption plant

with 99.8 to 99.92 % conversion and 30 to 680 mg/Nm^3 of SO_2 as daily averages, and new plants with 99.9 to 99.92 % and 30 to 340 mg/Nm^3 [1].

Models followed the plant into dynamics and optimization. A dynamic model of an industrial plant has been used to optimize its operation [9], and an industrial-scale plant has been optimized under uncertainty [10]. The transients of the beds themselves belong to the theory of fixed-bed reactors, which has represented a bed by a computational model extended to reacting systems since 1960 [12], has analyzed a creeping reaction zone with a model of cells [8], and has explained the transient temperature rise that follows a sudden reduction in feed temperature [7]. None of the sources read for this chapter documents a commercial converter's control configuration, so the simulator's loops are built on the structure above and claim to describe no plant's practice.

> ⧖ HISTORY **A Contact Process and Its Vanadium Melt, 1901 to the 1970s**
>
> In 1901 the contact process received a 47-page account by Knietsch in the *Berichte* of the German Chemical Society, on sulfuric acid and its manufacture by the contact method [13]. Of the catalysts tested for the oxidation of sulfur dioxide, only vanadium compounds, platinum and iron oxide proved technically satisfactory, and vanadium pentoxide is now used almost exclusively. A commercial catalyst carries 4 to 9 weight percent of it with alkali sulfate promoters on silica, and at operating temperature these form a liquid melt in which the reaction is thought to take place [1]. In the 1970s the catalyst was studied as a supported liquid phase: its kinetics and effectiveness factor were measured on an industrial vanadium catalyst, and its liquid dispersion and mass transport in their own right [11, 14]. The steep fall of the model's rate below its 420 °C break stands for that melt solidifying.

4.3 The Plant as Modeled

The plant is a generic sulfur-burning double-absorption unit of 1000 t/d, sized without any licensor's data (Figure 4.1). Liquid sulfur burns with dried air, a waste-heat boiler (WHB) cools the gas, and beds 1, 2 and 3 follow with an exchanger ahead of each of the last two. The reformer of Volume 11, Chapter 9 divides an endothermic reaction among beds in the same way, with a fired heater ahead of each. Each of these three coolers has a hot-gas bypass, so opening it warms the next bed's inlet. An economizer cools the gas leaving bed 3 to the interpass absorber, a reheater warms the returning gas for bed 4, and its bypass on the cold side cools bed 4's inlet when opened. A final absorber precedes the stack, and both absorbers drain into one acid loop held at 98.5 wt% by dilution water. The X–T diagram of the same instant is Figure 4.2. The engine's clock is the minute, flows in the converter are in mol/s, heat in watts and acid in kg/s.

Table 4.1 lists each variable's role and its design value. The integrated state holds the temperatures of the 96 bed cells, four exchanger metals, four bypass positions, the sulfur and air flows behind their pump and blower, the acid strength and temperature, the dilution water and the cooling water openings, 110 physical states in all, and seven ledgers that integrate the sulfur burned, the SO_2 and SO_3 leaving, the SO_3 absorbed, the water added, the acid made and its SO_3.

4.3.1 The Reaction and a Rate Law With a Break

Sulfur dioxide and oxygen react reversibly on the catalyst at a rate per kilogram that falls with the SO_3 it adsorbs, as shown in Equation (4.1).

$$\begin{aligned} &\mathrm{SO_2} + \tfrac{1}{2}\mathrm{O_2} \rightleftharpoons \mathrm{SO_3} \\ &r = a\,k(T)\,\frac{p_{SO_2}\sqrt{p_{O_2}} - p_{SO_3}/K_p(T)}{1 + K_s(T)\,p_{SO_3}} \\ &K_s = K_{s,0}\exp\!\left[\frac{Q_s}{R}\left(\frac{1}{T} - \frac{1}{T_{500}}\right)\right] \end{aligned} \tag{4.1}$$

where r is in mol/(s kg), the partial pressures $p_i = P_b F_i / F_t$ are in bar at the bed pressure P_b, a is the catalyst activity, $K_{s,0} = 2\,\mathrm{bar}^{-1}$ at $T_{500} = 773.15\,\mathrm{K}$ and $Q_s = 60\,\mathrm{kJ/mol}$. The rate constant is two

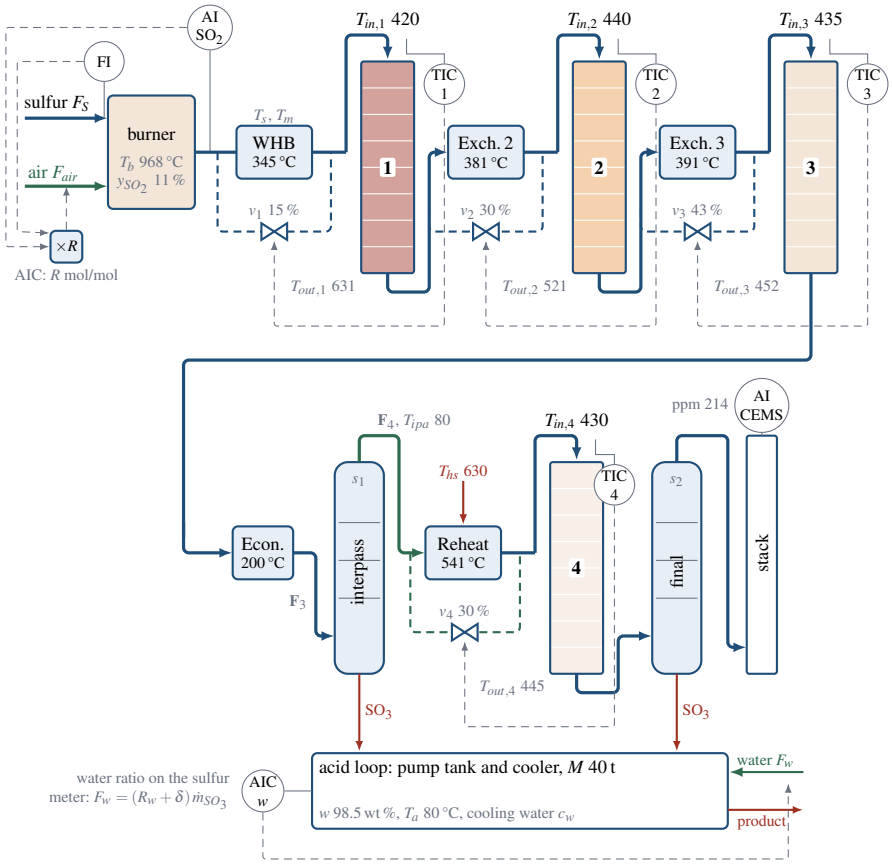

Figure 4.1 The double-absorption contact plant as modeled, case 1 at $t = 0$, after the app's Plant view, with the gas leaving each cooler and the bypass openings v_1 to v_4. The upper row runs from the burner to bed 3, and bed 3's outlet continues on the lower row. Sulfur F_S burns with air $F_{air} = R F_{S,m}$ to 11 % SO_2 at T_b, the WHB and exchangers 2 and 3 (sink T_s, metal T_m) set the inlets $T_{in,b}$ of beds 1 to 3, the interpass absorber passes a share s_1 of the SO_3 in $\mathbf{F}_3$, and the reheater, held on a stream at T_{hs}, heats $\mathbf{F}_4$ for bed 4 with its bypass on the cold side. TIC 1 to 4 move the bypasses, the AIC on the SO_2 analyzer sets the air ratio, and dilution water F_w follows the metered sulfur with a trim from the strength analyzer on the acid loop (w, T_a). A model drawing, not piping and not a safety layout. Not to scale.

Arrhenius terms in series, as shown in Equation (4.2).

$$\frac{1}{k} = \frac{1}{k_{hi}} + \frac{1}{k_{lo}}, \qquad k_{hi} = k_{500}\,\mathrm{e}^{-\frac{E_{hi}}{R}\left(\frac{1}{T}-\frac{1}{T_{500}}\right)}$$
$$\frac{k_{lo}}{k_{hi}} = \mathrm{e}^{-\frac{E_{lo}-E_{hi}}{R}\left(\frac{1}{T}-\frac{1}{T_{br}}\right)} \tag{4.2}$$

where $k_{500} = 0.5\,\mathrm{mol/(s\,kg\,bar^{1.5})}$, $E_{hi} = 80\,\mathrm{kJ/mol}$, $E_{lo} = 600\,\mathrm{kJ/mol}$ and the break $T_{br} = 420\,^\circ\mathrm{C}$. Above the break k follows k_{hi}, below it the steep k_{lo}, and at the break $k = k_{hi}/2$. The constant is 5.94×10^{-2} at 420 °C, 5.05×10^{-3} at 400 °C and 2.02×10^{-4} at 380 °C, the shape of a melt catalyst solidifying [1]. The law is a teaching law, not a vendor's kinetics. It was given a break because, while the model was built, a single Arrhenius term could not place the floor and keep the late beds active at once. With $E_{lo} = E_{hi} = 80\,\mathrm{kJ/mol}$ bed 1 lights at 380 °C, the bottom of the range searched (Section 4.4).

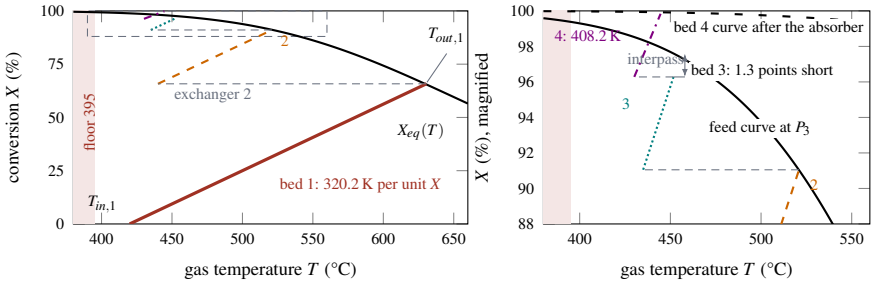

Figure 4.2 The X–T diagram of case 1 at $t = 0$, drawn from the engine's 96 cells: the feed's equilibrium $X_{eq}(T)$ at bed 2's pressure (solid), the floor shaded below 395 °C, the beds as adiabatic lines (bed 1 thick solid, bed 2 dashed, bed 3 dotted, bed 4 dash-dotted) and the exchangers as steps (thin dashed). The right panel magnifies the dashed box, with the feed's curve at bed 3's pressure (solid), which bed 3 ends 1.3 points short of, and bed 4's own curve after the interpass absorber (loosely dashed).

Table 4.1 Variables, Their Roles and Design Values (Case 2 at $t = 0$)

Variable	Role	Design	Range
$T_{in,b}$	CV, bed inlets	420, 440, 435, 430 °C	SP 370–480
v_b	MV, bypasses	15.0, 29.6, 42.6, 30.0 %	0–100
y_{SO_2}	CV, SO_2 to the converter	11.0 vol%	SP 7.5–12.5
R	MV, air-to-sulfur ratio	9.091 mol/mol	7–14
w	CV, acid strength (analyzer)	98.50 wt%	SP 97.5–99.5
F_w	MV, dilution water	2.305 kg/s	0–5
T_a, φ	CV, MV: acid, cooling water	80 °C, 63.0 %	SP 60–100
—	MV, production rate	100 % of 1000 t/d	30–110
π	DV, sulfur against its meter	100 %	90–110
a_b	DV, catalyst activity	100 %	20–120
—	DV, bypass 2 stuck, water supply	free, 100 %	0–100
$T_{out,b}$	outputs, bed outlets	630.6, 520.9, 451.7, 444.5 °C	—
X_b	outputs, after each bed	65.77, 91.05, 96.28, 99.837 %	—
—	output, stack SO_2, SO_3	214.5, 2.4 ppmv	limit 400
T_b	output, burner gas	967.9 °C	—
$T_{i,b}$	states, 96 cells	see Fig. 4.2	—

4.3.2 Equilibrium From NIST-JANAF

The equilibrium constant is a two-constant fit, as shown in Equation (4.3).

$$\ln K_p = -11.2583 + \frac{11839.8}{T}$$
$$\Delta H = -R \cdot 11\,839.8\,\mathrm{K} = -98.44\,\mathrm{kJ/mol} \tag{4.3}$$

with K_p in $\mathrm{bar}^{-1/2}$. It was fitted over 400 to 650 °C to the constant computed from the Shomate parameters of the NIST-JANAF tables [15] for SO_2, O_2 and SO_3. Recomputed by an independent route, $\Delta_r G^\circ = \Delta_r H^\circ - T\Delta_r S^\circ$, the fit gives 166.37 against 166.29 $\mathrm{bar}^{-1/2}$ at 450 °C (+0.05 %) and 9.990 against 9.991 at 600 °C, and its worst error from 350 to 700 °C is −0.91 %, at 350 °C. The tables' heat of reaction runs from −98.93 to −97.75 kJ/mol over that range against the constant −98.44. The same ΔH is the heat released in the energy balance, so the equilibrium curve and the adiabatic lines are consistent by construction. The heat capacities at 500 °C from the same tables are 52.1, 33.6, 72.2 and 31.3 J/(mol K) for SO_2, O_2, SO_3 and N_2.

4.3.3 Beds as Cells on the Catalyst's Heat Capacity

Each bed is 24 cells of equal catalyst mass on a bed module that several applications share. Gas spends a fraction of a second in a cell, so its composition is solved algebraically as a stirred tank at

the cell's temperature, and only that temperature, carried by the catalyst, is integrated, as shown in Equation (4.4).

$$\begin{aligned} \mathbf{F}_{out} &= \mathbf{F}_{in} + \nu\,\xi, \qquad \xi = W_c\, r(T, \mathbf{F}_{out}) \\ W_c\, c_{cat}\, \frac{\mathrm{d}T}{\mathrm{d}t} &= \dot{F}C_p\,(T_{prev} - T) + \xi\,(-\Delta H) \end{aligned} \tag{4.4}$$

where $\nu = (-1, -\frac{1}{2}, 1, 0)$ over SO_2, O_2, SO_3 and N_2, W_c is a twenty-fourth of the bed's 22, 28, 38 or 70 t, $c_{cat} = 1000\,\mathrm{J/(kg\,K)}$, T_{prev} is the previous cell or the bed inlet, and $\dot{F}C_p = \sum_i F_i c_{p,i}$ is the stream's heat-capacity flow, taken at the converter feed and held along beds 1 to 3, with bed 4 a second train on the gas the interpass absorber returns. The extent ξ is found by six damped Newton iterations from zero, a fixed count with no convergence test, so the right-hand side is a pure function of the state and runs replay bit for bit. Building this plant carried one change to the shared module. A species at zero flow whose Newton step is negative is now clipped at zero rather than allowed to set the damping to zero, which had frozen cells with inhibited rates at zero extent. The three applications then on the module, and two column applications checked alongside, logged byte-identical runs before and after this one guard. A bed-1 cell holds 917 kJ/K against an $\dot{F}C_p$ of 36.39 kW/K, a time constant of 25.2 s, and the beds' time constants $C_{bed}/\dot{F}C_p$ are 10.1, 12.8, 17.4 and 40.9 min. Pressures are fixed at 1.40, 1.36, 1.32 and 1.22 bar absolute.

Holding $\dot{F}C_p$ makes every bed a straight line on the diagram, as shown in Equation (4.5), the adiabatic line of Equation (1.3) in this plant's symbols.

$$T_{out} - T_{in} = \frac{F_{SO_2,0}\,(-\Delta H)}{\dot{F}C_p}\,(X_{out} - X_{in}) \tag{4.5}$$

where $X = 1 - F_{SO_2}/F_{SO_2,0}$ is the conversion of the SO_2 fed to the converter. By hand, $118.363 \times 98\,442/36\,389 = 320.20$ K per unit conversion for beds 1 to 3, 3.202 K a point, and bed 4's gas after the interpass absorber, with an $\dot{F}C_p$ of 28.545 kW/K, climbs 408.19 K per unit. The engine's beds lie on these slopes to 2.4×10^{-11} relative and every cell of beds 1 to 3 on its line to 9.8×10^{-10} K.

4.3.4 Burner, Exchangers and Absorbers

Each mole of sulfur takes one mole of oxygen, so the burner gas and its temperature follow from the two flows, as shown in Equation (4.6).

$$\begin{aligned} S &= \min(\pi F_S,\ 0.21 F_{air}) \\ \mathbf{F}_0 &= (S,\ 0.21 F_{air} - S,\ 0,\ 0.79 F_{air}) \\ T_b &= T_{air} + q_c\, S/(\dot{F}C_p)_{hot} \end{aligned} \tag{4.6}$$

where π is the sulfur burned per kilogram the meter counts, $T_{air} = 80\,^\circ$C, $q_c = 279\,\mathrm{kJ/mol}$ is the heat of formation of SO_2, 296.8 kJ/mol in the tables, less a teaching furnace loss, and $(\dot{F}C_p)_{hot}$ uses heat capacities near 600 °C. At design $T_b = 967.9\,^\circ$C. Each cooler passes a share $\beta = 0.8\nu$ of the gas around itself, as shown in Equation (4.7).

$$\begin{aligned} G &= (1-\beta)\dot{F}C_p, \qquad \varepsilon = 1 - \mathrm{e}^{-UA_g/G} \\ T_x &= T_h - \varepsilon(T_h - T_m), \qquad T_{in} = \beta T_h + (1-\beta) T_x \\ C_m \frac{\mathrm{d}T_m}{\mathrm{d}t} &= G\varepsilon(T_h - T_m) - UA_s(T_m - T_s) \end{aligned} \tag{4.7}$$

where T_h is the burner gas or the previous bed's outlet. The reheater is the mirror image, with the cold gas at $T_{ipa} = 80\,^\circ$C heated by metal held near a stream at $T_{hs} = 630\,^\circ$C, $T_x = T_{ipa} + \varepsilon(T_m - T_{ipa})$ and $C_m \dot{T}_m = UA_s(T_{hs} - T_m) - G\varepsilon(T_m - T_{ipa})$, so opening its bypass cools bed 4. At design the coolers' gas leaves at 345.2, 380.9, 390.5 and 540.6 °C before the bypassed gas rejoins it, and one point of bypass moves the four inlets by 4.21, 1.32, 0.39 and –2.76 K.

The absorbers pass a share of the SO_3 that is least at 98.5 wt% and cool acid, and the acid loop balances its strength and its temperature, as shown in Equation (4.8).

$$\begin{aligned} s &= \frac{s_0(1+g)}{1+s_0 g/s_{max}} \\ g &= k_w\left(\frac{w-98.5}{0.2}\right)^2 + k_T(T_a-85)^+ \\ M\frac{\mathrm{d}\sigma}{\mathrm{d}t} &= \dot{m}_{SO_3} - \sigma(\dot{m}_{SO_3}+F_w) \\ Mc_a\frac{\mathrm{d}T_a}{\mathrm{d}t} &= \dot{n}_{abs}q_{abs} + Q_{gas} + F_w c_w(25-T_a) \\ &\quad - UA_c\varphi(T_a-T_{cw}) \end{aligned} \tag{4.8}$$

where s_0 is 0.1 % for the interpass and 0.05 % for the final absorber, $k_w = 3$, $k_T = 0.5\,\mathrm{K}^{-1}$, $s_{max} = 10\,\%$, σ is the acid's SO_3-equivalent mass fraction (0.804 at 98.5 wt%), $M = 40\,\mathrm{t}$, $\dot{m}_{SO_3}$ the SO_3 absorbed, $q_{abs} = 130\,\mathrm{kJ/mol}$, Q_{gas} the heat of the gas cooled to 200 and 180 °C in the two absorbers, the dilution water enters at 25 °C with $c_w = 4.186\,\mathrm{kJ/(kg\,K)}$, and φ is the cooling water opening. The interpass absorber returns $F_{4,SO_3} = s_1 F_{3,SO_3}$, and the stack's SO_2 is $10^6 F_{o,SO_2}$ over the gas leaving the final absorber. At 98.5 % the acid takes 0.244 kg of water per kilogram of SO_3. The bypass valves follow their commands with a lag of 0.15 min, the sulfur pump and the blower 0.5 min, and the water and cooling-water valves 0.2 min.

4.3.5 Instruments and the Lit Flag

Thermocouples on every bed inlet and outlet have a lag of 0.2 min and 0.3 K of noise, the SO_2 analyzer on the converter feed and the stack monitor (CEMS) a lag of 1 min with noise of 0.02 vol% and 3 ppm, and the acid thermocouple 0.3 min and 0.1 K. The sulfur meter reads the flow the pump delivers, not the sulfur in it, and the flow meters carry 0.2 % noise. The strength analyzer reports a sample every 5 min after a 1 min sample line and 1 min of analysis, with 0.02 wt% of noise. A bed counts as lit while it covers at least a fifth of the way from its inlet conversion to the equilibrium at its inlet temperature, which sets the BED EXTINCT status.

4.3.6 Parameters, Initialization and Numerics

Table 4.2 lists the parameters. Every case begins from the plant lined out for the settings it ships with, each bypass found by bisection so that its inlet sits on its setpoint, the exchanger metals at their balance, the beds lined out by the module's steady solver, and the acid loop solved for the water and the cooling water. The largest derivative of a physical state at the design point is 8.6×10^{-10}. Classical fourth-order Runge–Kutta advances the state in 2 s steps, control acts each 6 s, and a row is logged each 20 s (the clocks of Section 1.6.1). Over a disturbed three-hour run the sulfur and oxygen balances close to 1.3×10^{-15} and 7.9×10^{-16} at every step, and the acid's SO_3 to 1.7×10^{-15}. Halving the step through the whole engine moves no logged conversion by more than 2.2×10^{-6} points.

4.3.7 Assumptions

The plant has one reaction on one catalyst, and the rate law of Equations (4.1) and (4.2) is a teaching law of the published shape. Cesium-promoted catalysts that light lower are not modeled, and the break temperature and activation energies are not fitted to any measured catalyst. The gas holds no heat and no composition, so a cell's composition lines out at every instant and its heat sits in the catalyst. Heat losses, the change of $\dot{F}C_p$ with conversion along a train and pressure drop that varies with rate are left out, and each pressure is fixed. Each exchanger is one metal temperature, and the reheater's hot side is held at 630 °C where a plant would take it from bed 1's outlet through a hot interpass exchanger, so the beds are not coupled through it. The two absorbers share one acid loop of one strength and one temperature, and their slip against strength and temperature is a teaching

Table 4.2 Parameters. A: Adjustable in the Model Dialog. H: Hidden and Fixed

Symbol	Meaning	Value	
k_{500}	rate constant at 500 °C	0.5 mol/(s kg bar$^{1.5}$)	A
E_{hi}, E_{lo}	activation above, below the break	80, 600 kJ/mol	A
T_{br}	break of the Arrhenius plot	420 °C	A
$K_{s,0}$, Q_s	SO_3 inhibition, its heat	2 bar^{-1}, 60 kJ/mol	H
W_b	catalyst in beds 1–4	22, 28, 38, 70 t	H
c_{cat}	catalyst heat capacity	1000 J/(kg K)	H
P_b	bed pressures	1.40, 1.36, 1.32, 1.22 bar	H
T_{air}, q_c	air, net combustion heat	80 °C, 279 kJ/mol	H
β_{max}	largest bypassed share	0.8	H
UA_g	gas side, WHB, X2, X3, reheater	96, 52, 21, 50 kW/K	H
UA_s	sink side, same order	319, 195, 66, 261 kW/K	H
T_s, T_{hs}	sinks; reheater's hot stream	250, 300, 250; 630 °C	H
C_m	exchanger metal	30, 20, 20, 20 MJ/K	H
T_{ipa}	gas from the interpass absorber	80 °C	H
$s_{0,1}$, $s_{0,2}$	SO_3 slip at best	0.10, 0.05 %	A
k_w, k_T, s_{max}	slip growth, largest slip	3, 0.5 K^{-1}, 10 %	H
M, c_a	acid holdup, heat capacity	40 t, 1.45 kJ/(kg K)	H
q_{abs}	heat per mol SO_3 absorbed	130 kJ/mol	H
UA_c, T_{cw}	acid cooler, cooling water	700 kW/K, 30 °C	H
$F_{w,max}$	dilution water at full valve	5 kg/s	H
—	stack SO_2 limit	400 ppmv	A
—	ACID HOT alarm	100 °C	H

law, not a design correlation. The sulfur pump's flow loop is ideal. The stack limit of 400 ppmv, the floor marked at 395 °C, the strength band of 98.3 to 98.7 % and the ACID HOT alarm at 100 °C are teaching values. A real stack limit comes from the plant's permit (the United States standard above gives 402.0 ppmv at 11 % SO_2 [3]), a real floor from the catalyst supplier, and the strength band and the acid temperature from the absorbers' design, for which the European reference document lists 98.5 to 99.5 % acid normally at 70 to 120 °C [1].

4.4 Lines, a Floor and Fronts

4.4.1 What Double Absorption Buys

Figure 4.2 shows the design point on the diagram. At the design inlets bed 1 climbs to its curve and stops on it, at 65.773 % and 630.6 °C, bed 2 stops 0.09 point short of its own outlet's equilibrium at 91.053 %, and bed 3 ends at 96.282 % and 451.7 °C, where the feed's curve at its pressure reads 97.599 %, 1.3 points away, because its catalyst is too cool and its gas too lean to go further. The interpass absorber then removes all but 0.1 % of the SO_3, bed 4 sees a gas of 4.40 mol/s SO_2, 50.6 mol/s O_2 and 0.114 mol/s SO_3, and it converts 95.62 % of what it receives to finish at 99.837 % overall. With the interpass absorber's SO_3 left in the gas, the same inlets give 97.910 %. Removing the product starts bed 4 on a new equilibrium curve, the effect Section 1.4.2 shows for the running example of Chapter 1.

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

4.4.2 An Ignition Floor From the Heat Balance

Nothing in the model clamps a temperature, a rate or a conversion. The floor appears because a bed lights only if the reaction near its inlet warms the gas fast enough to light the catalyst behind it before the gas leaves. Lined out at the design gas, bed 1 converts 62.80 % at a 395.0 °C inlet, 17.48 % at 394.0 °C and 3.04 % at 392.0 °C, as shown in Figure 4.3a. The inlet at which the bed passes 30 % is 395.12 °C in an independent plug-flow integration of the transcribed law, and 393.44, 394.33 and 394.74 °C at 12, 24 and 48 cells. Started lit at 420 °C and held ten hours at each inlet, the 24-cell bed ends where it does from cold to within 1×10^{-4} points between 384 and 404 °C, so the floor shows no hysteresis at this resolution. The rate constant at 395 °C is 3.93 % of its value at the 420 °C break,

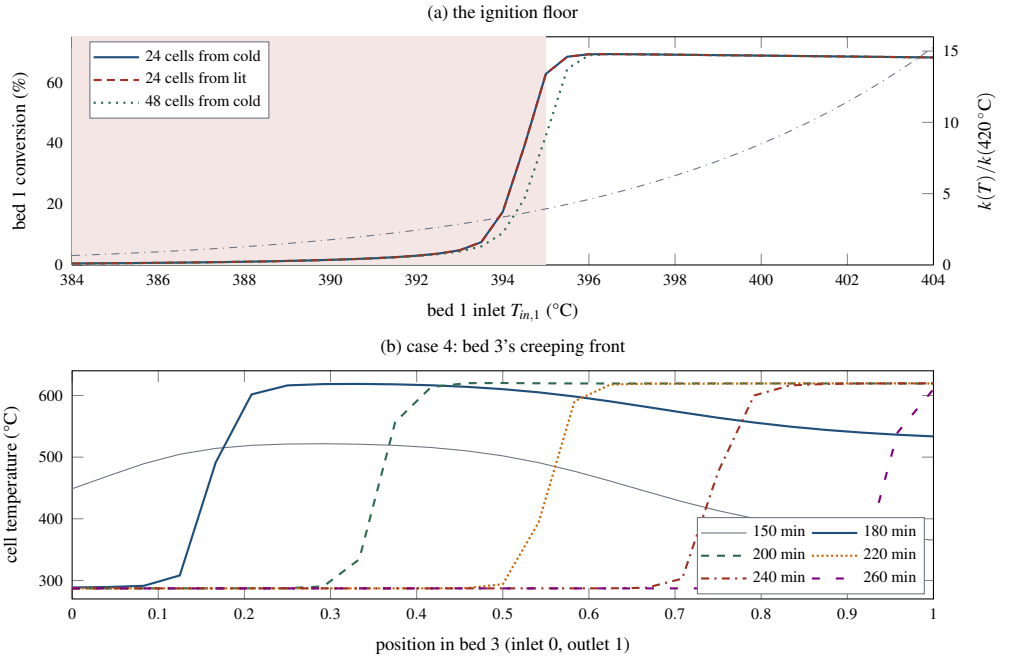

Figure 4.3 The floor and the front, from the engine's bed module. (a) Bed 1 lined out at the design gas against its inlet temperature, from cold (solid) and from a lit bed (dashed) at 24 cells and from cold at 48 (dotted), with the rate constant relative to its value at 420 °C (dash-dot, right axis). The shaded band is the 395 °C floor. (b) Case 4 as shipped, bed 3's 24 cell temperatures from inlet to outlet at six instants, each in its own line style as the legend shows. From 180 min the inlet is at 289 °C and the bed converts on the heat stored in its catalyst.

and the bed at 420 °C still converts 65.77 %, so the break is where the catalyst slows, not where it goes out. The floor belongs to the bed as well as the catalyst. It moves to 386.5 or 402.2 °C with the break at 410 or 430 °C, to 389.2 or 400.3 °C with k_{500} doubled or halved, and to 400.3 °C with half of bed 1's catalyst, while a single Arrhenius term at 80 kJ/mol lights the bed at 380 °C. Six cells per bed, the count first planned, kept a lit bed 1 burning down to a 370 °C inlet and lit it from cold only at 392 °C, a heat-smearing artifact that set the count at 24. The stirred-tank reactor of Volume 7, Chapter 2 ignites and goes out at folds of the same contest between heat released and heat carried away, there through a cooling jacket.

4.4.3 Turndown and a Front That Creeps Out

Half the gas through the same exchangers is cooled further, because each exchanger then has twice the transfer units. In case 4 the rate falls to 50 % at 20 min with the four bypasses held, and by 60 min the inlets are 367.7, 400.5, 393.7 and 477.3 °C. The stack first falls to 89 ppmv at 45 min, then bed 3, which has little conversion left and so little heat of its own, goes out at 1 h 9 min, the CEMS passes 400 ppmv at 1 h 22 min, bed 1 goes out at 1 h 53 min and bed 2 at 2 h 48 min. In between, bed 2, taking over bed 1's work on hot catalyst, relights bed 3 at 2 h 11 min. At 3 h the inlets are 367.8, 320.6, 288.6 and 476.7 °C, yet bed 3 adds 88.33 points of conversion with a 288.6 °C inlet, 106 K under the floor, and the stack reads 2128 ppmv.

Bed 3 is lit by the heat in its own catalyst. Its hottest cell is the 8th of 24 at 618.7 °C at 180 min, the 12th at 200 min, the 21st at 220 min and the last at 240 min, as shown in Figure 4.3b, and the bed is out at 4 h 26 min. This is the creeping reaction zone of fixed-bed theory [8]. At 50 % rate a thermal wave would cross bed 3 in 34.8 min (38 MJ/K of catalyst over 18.18 kW/K), and the reaction

heat behind the front slows it by $1 - \Delta T_{rx}/(T_{hot} - T_{in}) = 1 - 68.4 \times 3.20/(620 - 287) = 0.34$, about 102 min to cross, against the 80 to 90 min measured. Run with the four loops in Auto instead, the bypasses open to 25.2, 48.2, 79.6 and 42.0 %, every bed stays lit and the stack reads 74 ppmv at 3 h.

> **TRY IT IN THE APP** **4 · Turndown to 50 % on Fixed Bypasses**
>
> Load scenario 4 and press Run. At 20 min the rate falls to half with the four bypasses held. Watch the stack trend fall to 89 ppmv at 45 min, then climb past the 400 ppmv limit at 1 h 22 min, after bed 3 has gone out. Reload, switch the header to Auto and run again: every bed stays lit, and the stack reads 74 ppmv at 3 h.

4.4.4 Outlets Above Equilibrium and the Wrong-Way Rise

The same stored heat explains a transient that looks wrong on the diagram. A bed's outlet can sit above the equilibrium of its own temperature, by at most 2.60 points (bed 4 in the shipped challenge at 355.7 min) and 1.12 points (bed 4 in case 3). In every such instant the outlet cell is the bed's hottest, the conversion peaked in a cooler cell upstream, and the last cells run backward, five of them in the challenge, because gas converted in cooler cells meets catalyst still hot from before. Six and forty Newton iterations give the same outlet to 1×10^{-8} points, so the excess is physics of the model and not a solver artifact. Lined out, every outlet lies under its curve. In case 3 bed 1's inlet setpoint drops from 420 to 385 °C at 5 min, the cold gas cools the catalyst at the front first, the reaction zone moves down the bed into catalyst that is still hot, and the outlet rises to 653.4 °C at 15.2 min after the change before the bed goes out 31.8 min after it.

Pause. Try Exercise 4.2 before reading on.

4.4.5 How Many Cells

The cell count does not move the floor by more than 1.3 K from 12 to 48 cells, but it does set the speed of a front and the last ppm. At 48 cells bed 1 goes out 21.6 min after case 3's change instead of 31.8 min, with a peak of 677.8 °C instead of 653.4, case 4's bed 1 goes out at 1 h 14 min instead of 1 h 53 min, and the design stack is 198.6 ppmv instead of 214.5 ppmv (conversion 99.849 against 99.837 %). Twelve cells lose the wrong-way rise altogether. The lessons hold at 24 and 48 cells, and 24 were kept because 48 nearly double the cost, 14.66 against 7.76 ms per simulated minute headless in Node, on which the fastest speed of 60 min/s depends. Front times and the stack's last ppm are therefore this model's numbers.

4.5 Seven Loops and a Predictive Layer

The instrument code builds a single packet of readings for the controllers, and the application's tests whitelist its keys. Each loop carries a separate Auto/Manual switch on its card, and the header's Manual and Auto is the master of the four bed loops alone, so the header's Manual followed by its Auto throws all four while the SO_2, strength and acid temperature loops keep their own switches. Per-loop switches were chosen because the challenge ships with fixed bypasses that the suggested settings must undo, and because case 3 moves one loop while the others hold. Every controller not acting tracks what acts, so changes of mode and structure are bumpless (Section 1.6.2).

4.5.1 The Shared PI Law

The seven loops share the framework's velocity-form controller of Section 1.6.2, Equations (1.14) to (1.17), with the derivative term switched off in every loop here. Each loop executes every 6 s except the strength trim, only the emitted output is limited while the running sum runs free, and the integral is corrected toward the limited output with a tracking time equal to τ_I. Setting τ_I to the top of its slider turns integral action off. The integral reaches the output one execution after the error. The pairings and defaults are listed in Table 4.3. Bed 4's loop acts in the opposite sense to the other bed loops, direct acting in the convention of Section 1.6.2, because its bypass is on the reheater's cold

Table 4.3 Loops, Pairings and Default Tunings

Loop	PV → MV	K_c	τ_I	Output
TIC 1	$T_{in,1} \to v_1$	0.2 %/K	1 min	0–100 %
TIC 2	$T_{in,2} \to v_2$	0.5 %/K	1 min	0–100 %
TIC 3	$T_{in,3} \to v_3$	1.0 %/K	1 min	0–100 %
TIC 4	$T_{in,4} \to v_4$, direct	0.3 %/K	1 min	0–100 %
AIC	$y_{SO_2} \to R$, direct	0.4 per %	2 min	7–14 mol/mol
AIC	$w \to \delta$ (trim), direct	0.05 per wt %	15 min	±0.08 kg/kg
TIC	$T_a \to \varphi$, direct	8 %/K	4 min	0–100 %

side, and the production rate setpoint ramps at 2 % of design a minute, a teaching value in place of a step.

4.5.2 Air by Ratio, Water by Ratio and a Trim

The air is the ratio R times the metered sulfur, and the SO_2 loop moves R on the analyzer, so a sulfur richer than its meter shows first as a stronger gas. Both feeds follow the ratio station with a trim of Equation (1.10), and the Claus unit of Volume 11, Chapter 13 sets its air the same way, on a stoichiometric feedforward with a trim from a tail gas analyzer. The dilution water is a feedforward on the SO_3 the meters imply, as shown in Equation (4.9).

$$\begin{aligned} F_w &= \big(R_w(w_{SP}) + \delta\big)\, F_{S,m} X_m M_{SO_3}, \qquad R_w = \frac{1-\sigma_{SP}}{\sigma_{SP}} \\ X_m &= \frac{y_{in} - y_{out}}{y_{in}(1 - 1.5 y_{out})} \end{aligned} \tag{4.9}$$

where $F_{S,m}$ is the sulfur meter, y_{in} and y_{out} are the SO_2 analyzer and the CEMS as mole fractions, and δ is the trim. With Strength control on Ratio only, $\delta = 0$ and nothing looks at the acid. On Ratio + analyzer, the trim of Table 4.3 acts once per analyzer report with the analyzer's own 5 min period as its step and holds between reports. The ratio uses the conversion the analyzers report, so a change in conversion does not move the strength.

4.5.3 The MPC on the Inlet Setpoints

With Inlet setpoints from set to MPC, a receding-horizon controller moves the four inlet setpoints every 10 min for the least stack SO_2, and the four TICs hold what it sets. It is the framework's predictive option of Section 1.6.3 with an economic cost in place of a tracking error: the stack, as a steady copy of the beds predicts it. Its model is a coarser copy of the beds, 12 cells per bed with fresh catalyst, lined out at the inlets around the present setpoints u_0 and fed the gas the analyzers report (air meter times SO_2 analyzer). By central differences of ±2 K it fits $\text{ppm}(u_0 + \delta) \approx p_0 + \sum_b g_b \delta_b + \frac{1}{2}\sum_b H_b \delta_b^2$, and it plans over one hour as shown in Equation (4.10).

$$\begin{aligned} & x_{b,k+1} = a_b x_{b,k} + (1 - a_b) u_{b,k}, \quad a_b = \mathrm{e}^{-T_s/\tau_b} \\ & \min \ \textstyle\sum_{k=1}^{6} \sum_b \left[\frac{1}{2} H_b (x_{b,k} - x_b^*)^2 \text{ or } g_b x_{b,k} \right] \\ & \qquad + w_{\Delta u} \textstyle\sum \Delta u^2 \\ & \text{s.t. } 395 + \Delta T_{fl} \le u_b \le 470\,^\circ\text{C}, \qquad |\Delta u_b| \le 3\,\text{K} \end{aligned} \tag{4.10}$$

where $T_s = 10\,\text{min}$, $\tau_b = C_{bed,b}/\dot{F}C_p$, a bed with curvature $H_b > 0.02\,\text{ppm}/\text{K}^2$ tracks its local optimum $x_b^* = u_{0,b} - g_b/H_b$, a bed without it, bed 1 on its equilibrium curve, gets its gradient as a linear cost, $w_{\Delta u} = 1$ per K^2, the margin ΔT_{fl} is 15 K, and the moves are blocked one, one and four steps. Goldfarb and Idnani's dual active-set algorithm [16] in the framework's shared module solves each program, and only the first move is implemented, as in industrial MPC [17]. If a plan is not optimal or not finite, the setpoints stay where they are and the loops keep holding them. Nine steady solves of the

Table 4.4 The Nine Cases

#	Case and events	Teaching point
1	Manual, learner opens bypass 2 from 29.6 to 39.5 %	reading the diagram
2	Auto, rate to 80 % at 15 min	the inlet loops
3	Auto, learner sets bed 1 to 385 then 420 °C	the ignition floor
4	Manual, rate to 50 % at 20 min	turndown, fronts
5	Auto, learner sets SO_2 to 9 %	gas strength
6	Ratio only, sulfur +5 % at 20 min	ratio and analyzer
7	Auto, activity 60 % at 15 min	aging, infeasibility
8	Warm beds 430/430/400/400 °C, rate 30 to 100 %	start-up emissions
9	Challenge, 8 h, bed loops Man, Ratio only	the whole plant

copy make a decision. Over the eight-hour challenge it made 48 decisions with no failure. The model knows nothing of an aged catalyst or a richer sulfur except through the analyzers, and its masses, heat capacity, rate law, pressures and interpass slip are the declared defaults, so a change to the plant's parameters is a mismatch it does not see. The supervisory layer of Volume 9, Chapter 8 also chooses the setpoints that PI loops hold, against an optimum that sits on a moving constraint.

4.5.4 Plant, Identified and Controller Models

The simulated plant is the 24-cell model of Section 4.3, an identified model is whatever a learner fits from a run, such as a bypass's inlet gain, and the controllers' models are the PI gains, the water ratio and the MPC's 12-cell fresh-catalyst copy on the declared parameters, none of which sees the activity, the sulfur's purity, the true acid or a change to the plant's parameters. The status strip and the verdict read the truth and act on nothing.

4.6 Nine Cases on the Running Plant

The cases are listed in Table 4.4 and their outcomes in Table 4.5. Every number was measured with the application's engine run headless at its 2 s step and default seed, each case loaded as it ships, with any setting a row names entered at the time it names. An independent measuring script, run again for this chapter, reproduced its earlier output line for line, and every verdict is unchanged for seeds 0, 1, 2 and 3 and with the step halved. The exported file logs each reading beside its true value, in columns marked as truth that no controller reads (Section 1.6.4).

> **DID YOU KNOW? A Weaker Gas Gives a Cleaner Stack**
>
> From the simulator: burning the same sulfur with more air, to 9 instead of 11 percent sulfur dioxide, raises the conversion from 99.837 to 99.867 % and lowers the stack from 215 to 138 ppmv. The price is 22.2 % more air and a burner at 815 instead of 968 °C (Table 4.5, case 5).

4.6.1 Inlet Loops Through a Rate Change

Opening exchanger 2's bypass by ten points in case 1 moves bed 2's line 12.4 K to the right, and because the line then meets the curve lower, the conversion after bed 2 falls from 91.05 to 89.85 %, while bed 3, whose bypass is held, sees a 5.8 K warmer inlet and the stack rises from 214 to 218 ppmv. In case 2 the rate cut opens every bypass, 15.1 to 20.4, 29.6 to 40.0, 42.5 to 66.0 and, through bed 4's direct-acting loop, 30.0 to 36.7 %, and less gas per kilogram of catalyst takes the stack from 214 to 133 ppmv.

4.6.2 The Gas, the Acid and an Aged Catalyst

Lowering the SO_2 setpoint from 11 to 9 % in case 5 takes 22.2 % more air, cools the burner from 968 to 815 °C, shortens the temperature rise to 2.650 K a point, and raises the conversion from 99.837 to 99.867 % while the stack falls from 215 to 138 ppmv. In case 6 a sulfur cargo 5 % richer than its

Table 4.5 Case Outcomes Measured on the Engine

Case and condition	Result
1, bypass 2 at 1 min	at 46 min bed 2 in 440.0 → 452.4, out 521.0 → 529.4 °C; after bed 2 91.05 → 89.85 %; bed 3 in +5.8 K; stack 214 → 218 ppmv
2, rate 80 %	inlet error at most 2.77 K, within 1.5 K 11.7 min after; bypasses 20.4/40.0/66.0/36.7 %; stack 214 → 133 ppmv
3, bed 1 at 385 °C	outlet 653.4 °C 15.2 min after; BED 1 EXTINCT 31.8 min after; bed 2 out at 74.6 min; 7242 ppmv at 125 min
3, back to 420 °C	relit 3.2 min after; stack under 220 ppmv 56.5 min after
4, Manual	bed 3 out 1 h 9 min, CEMS > 400 at 1 h 22 min, bed 1 out 1 h 53 min; 2128 ppmv at 3 h
4, Auto	every bed lit, 74 ppmv at 3 h
5, 9 % SO_2	conversion 99.837 → 99.867 %; stack 215 → 138 ppmv; burner 968 → 815 °C; air +22.2 %
6, Ratio only	analyzer past 98.7 % 21 min after; true 99.381 % at 4 h; stack 215 → 348 ppmv
6, trim at 41 min	analyzer back under 98.7 % 30 min after
7, nothing done	762 ppmv at 60 min
7, bed 4 at 460 °C	under 400 ppmv 23.4 min after; 388 ppmv at 3 h (342 with bed 3 at 453 °C)
8, as shipped	CEMS peak 3948 ppmv at 18.3 min; 44.5 min over 400
8, bed 4 at 460 °C	2288 ppmv at $t = 0$, then at most 1817; 37.1 min over 400

meter arrives at 20 min with the water on the ratio alone. The analyzer reads above 98.7 % 21 min later and the true strength reaches 99.381 % at 4 h. The stack rises too, from 215 to 348 ppmv of SO_2, because off-strength acid lets the interpass absorber pass 3.745 % of the SO_3 instead of 0.100 %, which lowers bed 4's conversion of its own feed from 95.62 to 93.10 %, and the stack's SO_3 climbs from 2.4 to 216.5 ppmv. Switched to Ratio + analyzer at 41 min, the trim brings the reading back inside 98.7 % 30 min later.

> **TRY IT IN THE APP** **6 · Acid Strength: The Ratio and the Analyzer**
>
> Load scenario 6 and press Run on the Acid view. At 20 min a sulfur 5 % richer than its meter arrives with the dilution water on the ratio alone, and the analyzer reads above 98.7 % 21 min later. Set Strength control on the AIC · acid strength card to Ratio + analyzer. Switched at 41 min, the trim brings the reading back inside 98.7 % 30 min later.

Case 7 ages every bed to 60 % activity at 15 min, and the stack reaches 762 ppmv by 60 min. At 3 h bed 2 stops 2.0 points short of its curve and bed 4 converts 89.37 of its own feed's 97.67 %. Bed 4 at 460 °C brings the stack under the limit 23.4 min later and to 388 ppmv at 3 h, and bed 3 at 453 °C as well gives 342 ppmv. Below some activity no inlet answers. A search over inlets from 400 to 480 °C on the lined-out plant finds a best stack of 150.2 ppmv fresh, 333.2 ppmv at 60 % activity and 443.3 ppmv at 50 %, at 404.5, 453.5, 457.5 and 466.0 °C and 99.663 %, and the dynamic plant held five hours at those setpoints reads 443 ppmv. At half activity the 400 ppmv limit cannot be met by any inlet, which is why the case ages the catalyst to 60 % and the aging button to 50.

4.6.3 A Warm Start-Up

Case 8 starts with beds preheated to 430, 430, 400 and 400 °C and the rate ramping from 30 to 100 % from 1 min. At $t = 0$ the gas reaches beds 2 and 3 at 358.9 and 374.4 °C and the CEMS reads 2288 ppmv. As the rate rises the stack peaks at 3948 ppmv at 18.3 min, bed 3 lights at 32.4 min, and the stack is over the limit for 44.5 min in all. The lever the case offers is bed 4's inlet. With bed 4 at 460 °C the later peak falls to 1817 ppmv and the time over the limit to 37.1 min, while the reading at $t = 0$, fixed by the start, stays at 2288 ppmv. Over seeds 0 to 3 that first reading lies between 2281 and 2288 ppmv, and bed 4's setpoint moves only the later peak.

4.6.4 A Turndown Day

The challenge runs eight hours. The rate falls to 60 % at 1 h and returns at 5 h, and a sulfur 5 % richer than its meter arrives at 4 h. It is met when the average true conversion is at least 99.7 %, the true

Table 4.6 The Challenge Over 8 h at the Default Seed

Configuration	$\bar{X}$ (%)	Stack (ppmv)	Coldest (°C)	Acid (wt %)	
loops Man, Ratio only (shipped)	95.969	48 468	292.4	98.47–99.38	fail
loops Auto, Ratio + analyzer (sugg.)	99.878	260	417.2	98.45–98.63	pass
MPC, Ratio + analyzer	99.904	215	409.4	98.45–98.63	pass
loops Auto, Ratio only	99.852	350	417.2	98.47–99.38	fail
loops Man, Ratio + analyzer	96.042	48 481	292.4	98.45–98.63	fail

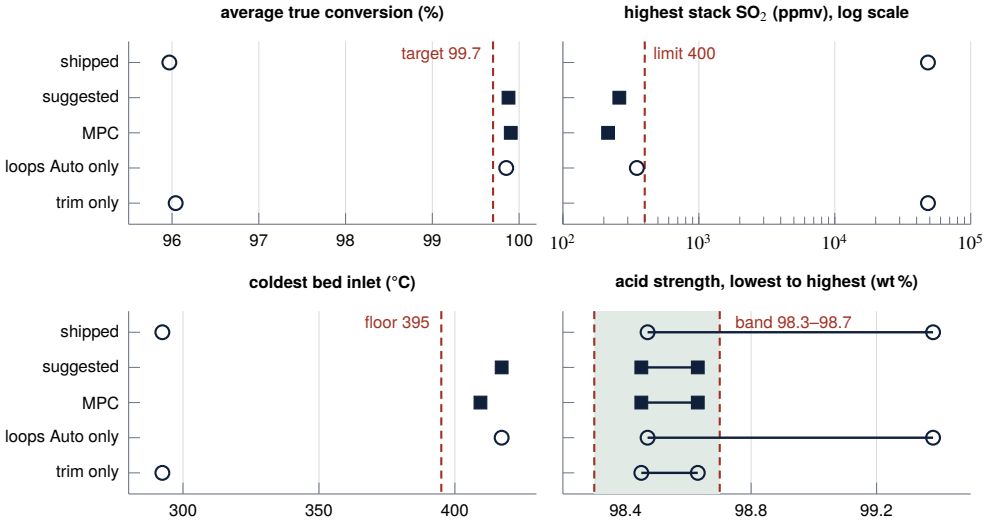

Figure 4.4 Table 4.6 as dots, digit for digit: the average true conversion against its target, the highest stack reading on a logarithmic axis against the limit, the coldest bed inlet against the floor, and the range of the acid strength against its band. The rows are the table's: shipped (loops Man, Ratio only), suggested (loops Auto, Ratio + analyzer), MPC (Ratio + analyzer), loops Auto only (Ratio only) and trim only (loops Man, Ratio + analyzer). Filled squares are the configurations that pass, open circles those that fail. Notice that each half of the fix fails on its own: the loops alone miss only the acid band, and the trim alone loses the converter.

stack SO_2 never exceeds 400 ppmv, no bed inlet falls below 395 °C and the acid stays within 98.3 to 98.7 wt%, scored every step against the simulator's truth. It ships with the four bed loops in Manual and the water on the ratio alone. The configurations are compared in Table 4.6, and Figure 4.4 draws its four measured columns against the target of each.

As shipped, the 60 % rate cools bed 1's inlet to 375.9 °C by 1 h 40 min, bed 1 goes out at 2 h 17 min, the stack passes the limit at 2 h 21 min and the beds go out one after another, while the richer sulfur takes the acid to 99.38 %. Each half of the fix fails alone. The loops without the trim pass on conversion, stack and inlets but not on the acid, and the trim without the loops keeps the acid but loses the converter. With both, the inlets never fall below 417.2 °C, the stack peaks at 260 ppmv and the conversion averages 99.878 %. The MPC with the same two changes averages 99.904 % with a stack of at most 215 ppmv, taking bed 1's inlet down to 409.4 °C, within a kelvin of its lower bound of 410 °C. Over four further seeds and a halved step the suggested settings and the MPC always pass and the shipped settings always fail. The turndown case at 3 h is shown in Figure 4.5.

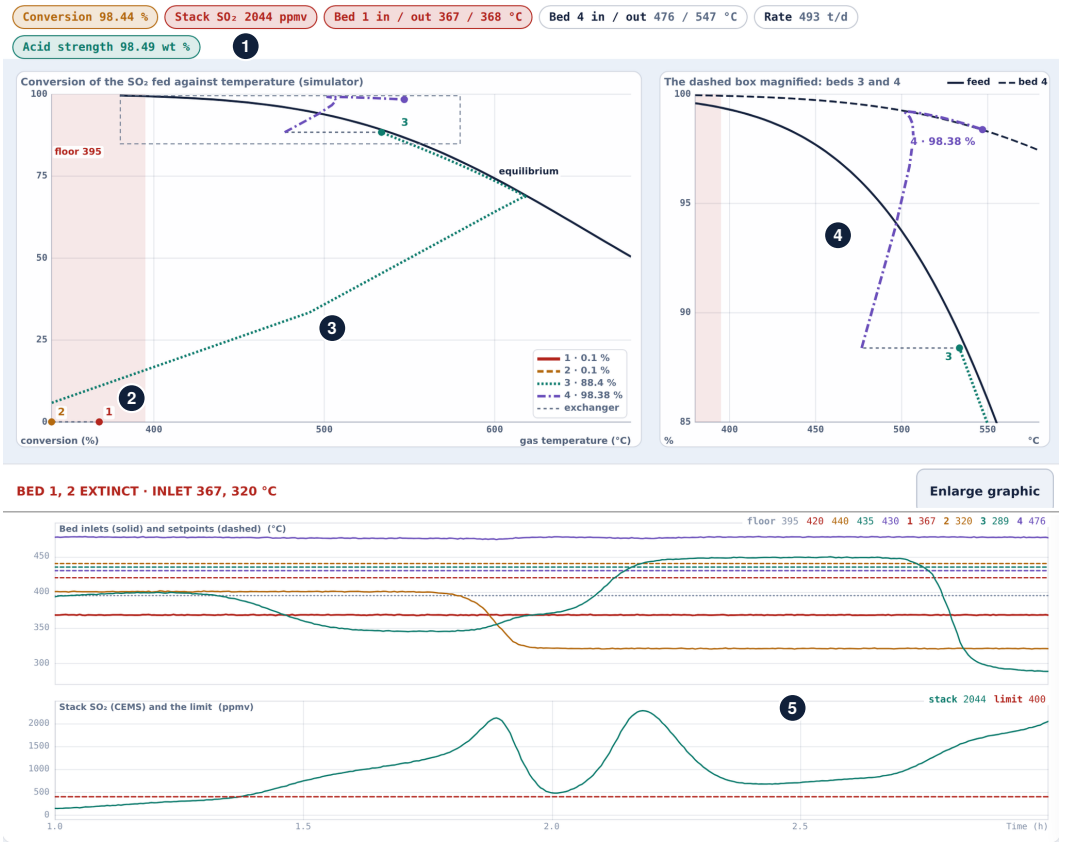

Figure 4.5 Case 4 at $t = 180$ min on the X–T diagram view, in Manual with the four bypasses held, more than two hours after the rate fell to 50 %, cropped to the plant view above the bed inlet and stack trends. Numbered callouts mark (1) the stack badge, (2) beds 1 and 2 on the floor, (3) bed 3's bent line, (4) bed 4 in the magnifier and (5) the CEMS trend above its limit. The badges read a conversion of 98.44 % from the analyzers, a stack of 2044 ppmv on the CEMS against 2128 ppmv true, bed 1 in and out at 367 and 368 °C, bed 4 at 476 and 547 °C, 493 t/d and acid at 98.49 wt%, and the status strip BED 1, 2 EXTINCT, INLET 367, 320 °C. On the diagram beds 1 and 2 lie flat on the shaded floor at 0.1 %, and bed 3 climbs a bent line from a 289 °C inlet to 88.4 % on its stored heat, while the magnifier, which has followed the plant down to bed 4's start, shows bed 4 reaching 98.38 %. The trends show bed 1's inlet flat at 367 °C below the 395 °C floor line, bed 3's inlet rising while bed 2 burned and falling once bed 2 went out, and the CEMS above its 400 ppmv limit since 1 h 22 min.

4.7 Technology in Practice

What Measures It

Figure 4.6 draws the loops of Section 4.5 with their devices. The variable the chapter turns on, each bed's inlet temperature, is a gas temperature read by an element in a well, and a few kelvin of it decide whether a bed stays lit. The gas itself is read twice, by the analyzer on the converter feed and by the stack monitor, and the acid by a strength analyzer that reports a sample every 5 min. No case of Table 4.4 breaks a thermocouple, so the model's trends show the gas, never a failed element. Figure 4.7 shows the element in its well.

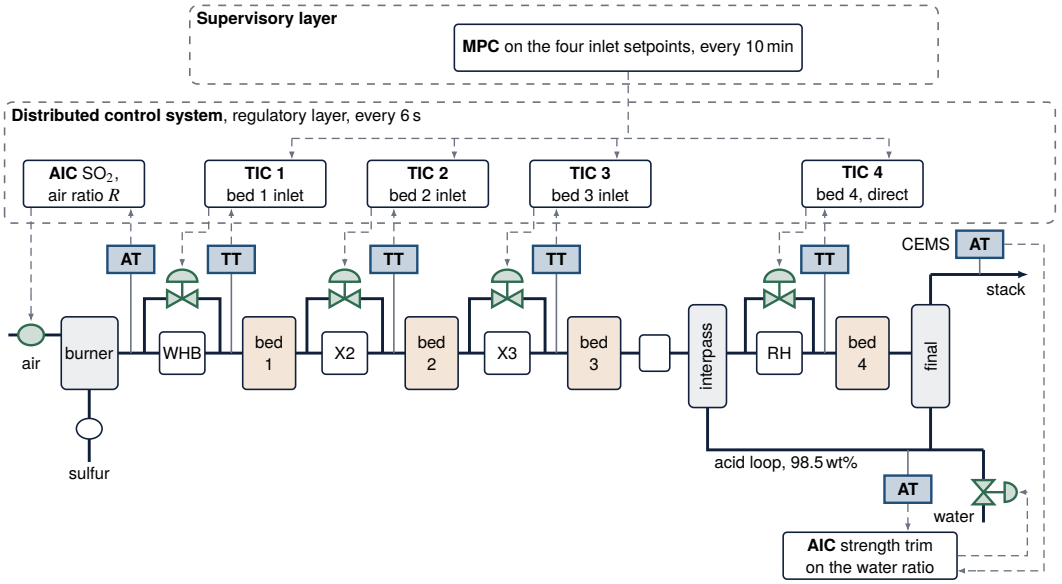

Figure 4.6 The loops of the contact plant with their hardware. Lettered boxes are sensing devices: a temperature element (TT) at each bed inlet, the SO_2 analyzer (AT) on the converter feed, the stack monitor (CEMS) and the strength analyzer (AT) on the acid loop. Valve symbols and the blower are final elements, dashed arrows are signals, and WHB, X2, X3 and RH are the waste-heat boiler, exchangers 2 and 3 and the reheater, and the small box after bed 3 is the economizer. Notice that each bed has its own loop on its own bypass, that bed 4's bypass sends cold gas around the reheater, which makes its loop direct acting, and that the predictive layer writes the four setpoints the same loops hold. The sulfur meter, the water ratio's other inputs and the acid temperature loop are left out. Not to scale, and not a piping diagram.

> **TECHNOLOGY** **Sensor: A Thermocouple in a Thermowell**
>
> Two wires of dissimilar metals joined at one end produce a voltage at the other that is roughly proportional to the junction's temperature. Thermocouples are less accurate than resistance thermometers but more rugged, with wider spans, and the pair of alloys is chosen for the gas: type K, for one, is attacked by sulfur and by reducing gases. A thermowell lets the element be replaced without opening the process, at the price of a slower first-order lag, and an air gap at its bottom adds a second lag. The most common failure is an open circuit. Without burnout detection the reading then wanders wildly with electrical noise, and with it the transmitter drives the reading high or low by design [18]. The model's 0.2 min lag and 0.3 K noise stand in for the element and its well.

What Moves It

Besides the sulfur pump, which sets the rate, seven handles move the plant: four gas bypasses, the air through the blower's ratio station, the dilution water and the acid cooler's water. A 2019 study of a smelter's acid plant describes the same structure in service, heat exchangers equipped with bypass valves to regulate the temperatures throughout the plant [19]. Three bypasses here send hot gas around a cooler and warm the next bed, and the fourth sends cold gas around the reheater.

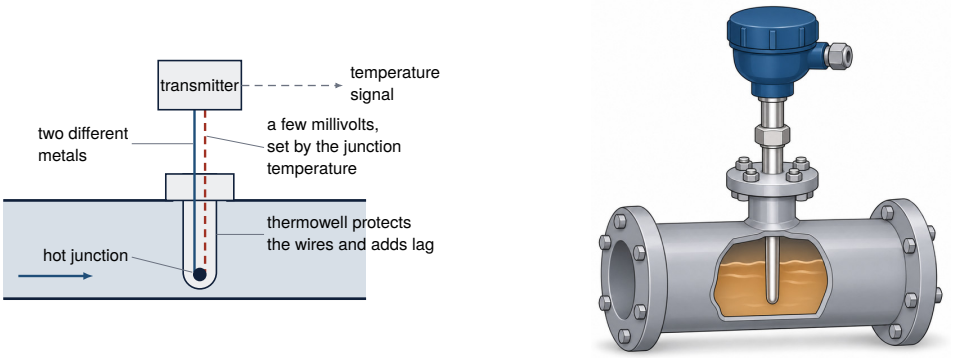

Figure 4.7 A thermocouple in a thermowell, the element behind TT 1 to TT 4. Notice the metal of the well between the gas and the junction: it is the 0.2 min lag of the model, and an air gap at the bottom of the well would add a second lag [18].

TECHNOLOGY **Final Element: A Rotary Valve on a Gas Bypass**

Large flows of air at low pressure are throttled by dampers, whose vanes turn together, as in furnace and boiler draft. A butterfly valve is a single disk turning across the pipe: edge-on it barely obstructs the flow, square to it it blocks the flow, and tight shutoff is hard to obtain unless its seat is lined with a soft material. It passes far more flow than a globe valve of the same pipe size. Installed, a valve sees less pressure drop as it opens, so its flow droops against its position, and a valve with excess friction ignores small reversals of the signal [18]. The model's bypass passes a share of the gas in proportion to its opening, none when shut, and follows its command with a 0.15 min lag.

Where the Controller Runs

A plant runs such loops in a distributed control system, in which a supervisory level oversees local controllers and sends them set points, loops are sometimes nested, and cycle times range from milliseconds to minutes [20]. The app executes its PI loops every 6 s, the strength trim once per analyzer report and the predictive layer every 10 min, all inside that range. The PI loops and the two ratio stations belong to the regulatory layer. The predictive controller is supervisory: it writes the four inlet setpoints the TICs hold, and when a plan fails the setpoints stay where they are.

IN PRACTICE **Two Analyzers and a Factor Three Times a Day**

A plant is judged by its conversion and its stack sulfur dioxide. A 2007 European reference document has both rest on sulfur dioxide measured at the converter inlet and after the final absorber, with the conversion determined daily and the levels monitored continuously [1]. The United States rule for new acid plants requires a continuous monitor with its span set at 1000 ppm, a conversion factor from the gas entering the converter at least three times a day, and it counts as excess emissions every three-hour period whose average exceeds the standard [21]. The chapter's feed analyzer and CEMS are those two instruments, and its challenge judges every step, a stricter test than a three-hour average.

What Is Modernizing It

Three developments bear on this converter. Cesium-promoted catalysts are active at 380 to 400 °C and can be run from 380 to 620 °C, in the first bed to lower its inlet or in the final bed at a low inlet, and at double-contact plants burning sulfur they reach 99.9 % conversion [1]. Heat recovery has made the plant an exporter of energy: up to 67 % of the waste heat of a double-contact plant burning sulfur is recovered as high-pressure steam, and a modern one can export about 6 GJ per tonne of acid [1]. And models identified from plant data now inform the temperature control: the 2019 study fitted dynamic models to industrial operating data and explored a feedforward and feedback scheme [19].

4.8 Limitations

The plant is generic, and its sizes, rates and limits serve teaching. The rate law has the published shape but its break, activation energies and inhibition are not fitted to a measured catalyst, so the floor of about 395 °C is this model's, below the 410 to 440 °C quoted for conventional catalysts and inside the 380 to 410 °C quoted for cesium-promoted ones [1]. The equilibrium constant is a two-constant fit within 0.91 % of NIST-JANAF from 350 to 700 °C, and holding $\dot{F}C_p$ and ΔH constant makes the lines exactly straight where real ones bend slightly. The gas carries no heat, so a front's speed is the catalyst's, and front times and the stack's last ppm depend on the 24 cells. The reheater's hot side is held, and the absorbers' slip law is a teaching closure that sets how strongly off-strength acid reaches the stack.

The MPC plans on a steady 12-cell copy with fresh catalyst, so after aging its model is wrong in a way only the analyzers report. The PI loops are not scheduled with the rate, and no case models a trip, a blower failure or the start-up preheat itself.

> **OPEN CHALLENGE** **A Converter Fed by a Smelter**
>
> Non-ferrous metal smelting accounts for 39 percent of sulfuric acid production, and its gas, unlike a sulfur burner's, varies in strength and flow. A 2007 European reference document puts 99.8 percent conversion within reach of four beds on gas of low variability and 99.7 percent on smelter gas, and tabulates a larger loss the more the sulfur dioxide varies in an hour [1]. A 2019 study of a smelter's acid plant found its feed typically between 5 and 15 percent sulfur dioxide and noted that most models assume a burner's steady feed, so continued research is required [19]. Progress would be loops that keep every bed lit through such swings, proven on plant data.

4.9 Conclusion

A sulfuric acid contact plant is a train of adiabatic beds on an exothermic equilibrium, and this simulator draws its beds live on the conversion–temperature diagram from balances that close to 2×10^{-15}. Each point of conversion heats the gas 3.20 K, and four beds with an interpass absorber reach 99.837 % where the same four beds without the absorber would leave 97.910 %. A floor near 395 °C emerges from a two-regime rate law and the heat balance with no clamp, while the 420 °C break of the rate law is not the floor. A turndown on fixed bypasses puts the beds out one by one, leaving bed 3 converting 88.33 points from a 289 °C inlet on stored heat. Four inlet loops and an analyzer trim together pass an eight-hour turndown day with 99.878 %, and an MPC on the inlet setpoints reaches 99.904 %, while at half activity no inlet meets the limit. A fuller model would fit the rate law to a measured catalyst, add a cesium-promoted bed 4 and couple the reheater to bed 1. The simulator shows that the best inlet temperature on an exothermic equilibrium is the coolest one the catalyst can still light.

KEY POINTS

- A contact plant converts sulfur dioxide in several catalyst beds with cooling between them, because each bed climbs only until it meets equilibrium.
- The interpass absorber starts bed 4 on a new curve: 99.837 % with it, 97.910 % without.
- The coolest inlet that still lights the catalyst converts the most, and the floor near 395 °C comes from the heat balance, not from the 420 °C break of the rate law.
- On fixed bypasses a turndown puts the beds out one by one. Inlet loops and the strength trim pass the turndown day at 99.878 %, and the MPC reaches 99.904 %.

Where else this applies. Chapter 2 climbs the same adiabatic lines with quench gas between its beds, and Volume 7, Chapter 2 ignites and goes out on the same heat balance.

References

[1] European Commission, "Reference document on best available techniques for the manufacture of large volume inorganic chemicals: Ammonia, acids and fertilisers," European Commission, Integrated Pollution Prevention and Control, Tech. Rep., Aug. 2007.

[2] S. J. Smith, J. van Aardenne, Z. Klimont, R. J. Andres, A. Volke, and S. Delgado Arias, "Anthropogenic sulfur dioxide emissions: 1850–2005," *Atmospheric Chemistry and Physics*, vol. 11, no. 3, pp. 1101–1116, 2011.

[3] U.S. Environmental Protection Agency, "Standards of performance for sulfuric acid plants, 40 CFR part 60, subpart H," Electronic Code of Federal Regulations, §§60.80–60.85, 2026, read October 2026. https://www.ecfr.gov/current/title-40/chapter-I/subchapter-C/part-60/subpart-H.

[4] M. J. King, W. G. Davenport, and M. S. Moats, "Oxidation of SO_2 to SO_3—equilibrium curves," in *Sulfuric Acid Manufacture*. Elsevier, 2013, pp. 125–133.

[5] ——, "SO_2 oxidation heatup paths," in *Sulfuric Acid Manufacture*. Elsevier, 2013, pp. 135–150.

[6] U.S. Geological Survey, *Mineral Commodity Summaries 2025*. Reston, VA: U.S. Geological Survey, 2025, sulfur, pp. 172–173.

[7] P. S. Mehta, W. N. Sams, and D. Luss, "Wrong-way behavior of packed-bed reactors: 1. The pseudo-homogeneous model," *AIChE Journal*, vol. 27, no. 2, pp. 234–246, 1981.

[8] H.-K. Rhee, D. Foley, and N. R. Amundson, "Creeping reaction zone in a catalytic, fixed-bed reactor: a cell model approach," *Chemical Engineering Science*, vol. 28, no. 2, pp. 607–615, 1973.

[9] A. A. Kiss, C. S. Bildea, and J. Grievink, "Dynamic modeling and process optimization of an industrial sulfuric acid plant," *Chemical Engineering Journal*, vol. 158, no. 2, pp. 241–249, 2010.

[10] M. Tejeda-Iglesias, J. Szuba, R. Koniuch, and L. Ricardez-Sandoval, "Optimization and modeling of an industrial-scale sulfuric acid plant under uncertainty," *Industrial & Engineering Chemistry Research*, vol. 57, no. 24, pp. 8253–8266, 2018.

[11] H. Livbjerg and J. Villadsen, "Kinetics and effectiveness factor for SO_2 oxidation on an industrial vanadium catalyst," *Chemical Engineering Science*, vol. 27, no. 1, pp. 21–38, 1972.

[12] H. A. Deans and L. Lapidus, "A computational model for predicting and correlating the behavior of fixed-bed reactors: II. extension to chemically reactive systems," *AIChE Journal*, vol. 6, no. 4, pp. 663–668, 1960.

[13] R. Knietsch, "Ueber die Schwefelsäure und ihre Fabrication nach dem Contactverfahren," *Berichte der deutschen chemischen Gesellschaft*, vol. 34, no. 3, pp. 4069–4115, 1901.

[14] H. Livbjerg, B. Sorensen, and J. Villadsen, "Liquid dispersion, mass transport, and chemical reaction in supported liquid phase catalysts for SO_2 oxidation," in *Chemical Reaction Engineering—II*, ser. Advances in Chemistry. American Chemical Society, 1974, pp. 242–258.

[15] M. W. Chase, Jr., "NIST-JANAF Thermochemical Tables, Fourth Edition," *Journal of Physical and Chemical Reference Data, Monograph*, vol. 9, pp. 1–1951, 1998.

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

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

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

[19] S. Caron, "Temperature modeling and control of the sulfuric acid plant in an industrial smelter," M.A.Sc. thesis, Laurentian University, Sudbury, ON, 2019, https://laurentian.scholaris.ca/handle/10219/3303.

[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] U.S. Environmental Protection Agency, "Standards of performance for sulfuric acid plants: Standard for sulfur dioxide and emission monitoring, 40 CFR 60.82 and 60.84," Code of Federal Regulations, Title 40, Volume 7, edition of July 1, 2025, 2025, https://www.govinfo.gov/app/collection/cfr/2025/title40/vol7.

4.10 Exercises

4.1 Concept check · 2 min

In case 1 the Exchanger 2 bypass is opened from 29.6 to 39.5 % with the other bypasses held, and bed 2's inlet rises by 12.4 K. Bed 2 was already at its equilibrium curve. What happens to the conversion after bed 2?

(a) It rises, because a hotter catalyst reacts faster and climbs further along its line

(b) It falls, because a line that starts further right meets the falling curve lower

(c) It stays the same, because the bypass does not change the slope of bed 2's line

(d) It rises, because bed 3's warmer inlet is credited to the conversion after bed 2

4.2 Concept check · 1 min

In case 3 the bed 1 inlet setpoint is lowered from 420 to 385 °C, below the ignition floor. What does bed 1's outlet temperature do first?

(a) It falls at once by about the 35 K of the inlet change

(b) It holds steady until the bed goes out, and then it falls

(c) It rises for about a quarter of an hour, and then the bed goes out

(d) It falls at first, and then the inlet loop brings it back

4.3 Concept check · 2 min

With the same four inlets the plant converts 99.837 % with the interpass absorber and 97.910 % with its SO_3 left in the gas. Why does removing SO_3 after bed 3 help bed 4?

(a) The absorber cools the gas, so bed 4 runs further below its floor and converts more

(b) The absorber takes out part of the SO_2, so less is left for bed 4 to convert

(c) The absorber puts oxygen back into the gas, which bed 4's rate law needs

(d) Bed 4 starts on a new equilibrium curve, because less SO_3 slows the back reaction

4.4 Theory and calculation · 6 min

The converter is fed 118.363 mol/s of SO_2. The heat of reaction is −98.442 kJ/mol, and the heat-capacity flow is held at 36.389 kW/K along beds 1 to 3 and at 28.545 kW/K for bed 4 after the interpass absorber. Conversion X is counted on the SO_2 fed to the converter.

(a) Find the temperature rise of beds 1 to 3 per unit conversion and per point of conversion. *Checkpoint:* 320.20 K and 3.202 K.

(b) Bed 1 enters at 420.0 °C with no conversion and stops on its curve at 65.77 %. Find its outlet temperature. *Checkpoint:* 630.6 °C.

(c) Bed 4 enters at 430.0 °C with $X = 96.28$ % and leaves at 99.837 %. Find its rise per point and its outlet temperature, and say why its line is flatter on the diagram. *Checkpoint:* 4.082 K a point and 444.5 °C.

4.5 Theory and calculation · 6 min

Use the rate constant of Equation (4.2) with $k_{500} = 0.5\,\mathrm{mol/(s\,kg\,bar^{1.5})}$, $E_{hi} = 80\,\mathrm{kJ/mol}$, $E_{lo} = 600\,\mathrm{kJ/mol}$, $T_{500} = 773.15\,\mathrm{K}$, the break at 420 °C and $R = 8.314\,\mathrm{J/(mol\,K)}$.

(a) Find k_{hi} and k at the break. *Checkpoint:* 0.1189 and 0.0594.

(b) Find k at 400 °C. *Checkpoint:* 5.05×10^{-3}.

(c) Find k at 395 °C as a share of its value at 420 °C, and the same share for a single Arrhenius term at 80 kJ/mol. *Checkpoint:* 3.93 % against 59.5 %.

4.6 Theory and calculation · 7 min

A thermal front in a bed moves at the speed at which the gas can change the temperature of the catalyst. The catalyst holds 1000 J/(kg K), and beds 1 to 4 hold 22, 28, 38 and 70 t.

(a) Find the time constant $C_{bed}/\dot{F}C_p$ of bed 1 at an $\dot{F}C_p$ of 36.39 kW/K, and of bed 4 at 28.545 kW/K. *Checkpoint:* 10.1 min and 40.9 min.

(b) In case 4 the rate is 50 % and $\dot{F}C_p$ is 18.18 kW/K. Find the time a thermal wave takes to cross bed 3. *Checkpoint:* 34.8 min.

(c) Behind the front bed 3 converts 68.4 points at 3.20 K a point, from a 287 °C inlet to a 620 °C hot spot. The reaction slows the front by the factor $1 - \Delta T_{rx}/(T_{hot} - T_{in})$. Find the factor and the crossing time, and compare it with the 80 to 90 min measured. *Checkpoint:* 0.34 and about 102 min.

4.7 App challenge · 8 min

Load scenario 2, The inlet temperature loops, and change no setting. The production rate setpoint goes to 80 % at 15 min. Use the default seed and run at 20 min/s.

(a) Run to 135 min and read the four bypass openings on the TIC cards. Accept 0.3 % on each. *Checkpoint:* 20.4, 40.0, 66.0 and 36.7 %.

(b) Export the CSV file. Its time column is in hours. Read the true stack SO_2, column `truth_stack_so2_ppm`, in the rows at 0 and 2.25 h. Accept 1 ppmv. *Checkpoint:* 214.5 and 132.8 ppmv.

(c) Raise the Bed 4 inlet setpoint on TIC 4 to 440 °C and run 40 min more. Say which way the reheater bypass moves and read where it settles, within 0.4 %. *Checkpoint:* It closes, to about 33.7 %.

Explain why all four bypasses opened for the rate cut.

4.8 App challenge · 8 min

Load scenario 7, An aged catalyst: higher inlets. Every bed falls to 60 % activity at 15 min. Use the default seed, run each case to 180 min at 20 min/s, export the CSV file (time column in hours), and read the true stack SO_2, column `truth_stack_so2_ppm`. Accept 2 ppmv.

(a) Run the scenario as shipped and read the rows at 1 and 3 h. *Checkpoint:* 762 and 766 ppmv.

(b) Press the reload button, set the Bed 4 inlet setpoint on TIC 4 to 460 °C before Run, and read the row at 3 h. *Checkpoint:* 388 ppmv.

(c) Also set the Bed 3 inlet setpoint to 453 °C before Run, and read the row at 3 h. *Checkpoint:* 342 ppmv.

Explain why bed 4 is the inlet to raise first.

4.9 App challenge · 9 min

Load scenario 9, Challenge, a turndown day. The rate goes to 60 % at 1 h, a sulfur 5 % richer than its meter arrives at 4 h, and the rate returns at 5 h. Use the default seed and run at 60 min/s until the case

pauses at 8 h. For each run read the So far tile: average conversion, highest stack SO_2, coldest bed inlet and the acid strength range. Accept one unit in the last digit the tile shows.

(a) Run it as shipped. *Checkpoint:* 95.97 %, 48 468 ppm, 292 °C, 98.47 to 99.38 %.

(b) Press the reload button. Before Run, put the four bed loops in Auto (the header's Manual and then Auto throws all four), keep Strength control on Ratio only, and run again. *Checkpoint:* 99.85 %, 350 ppm, 417 °C, 98.47 to 99.38 %.

(c) Press the reload button, put the bed loops in Auto and set Strength control to Ratio + analyzer before Run, and run again. *Checkpoint:* 99.88 %, 260 ppm, 417 °C, 98.45 to 98.63 %.

Say which limit each run breaks and which loop protects it.

4.10 Application area · 4 min

The United States standard limits a sulfuric acid plant to 2 kg of SO_2 per metric ton of acid. A stack reading C in ppmv converts to that unit by multiplying it by $CF = k(1 - 0.015r)/(r - s)$, with $k = 0.0653$ and r and s the SO_2 percentages entering the converter and in the stack.

(a) Find the reading in ppmv that equals the limit for an 11 % gas. *Checkpoint:* 402.0 ppmv.

(b) Find the same for the 9 % gas of case 5, and the emission of the design point, 214.5 ppmv at 11 %, in kilograms per ton. *Checkpoint:* 317.5 ppmv and 1.07 kg/t.

4.11 Application area · 4 min

Industrial vanadium catalysts are quoted with a lower operating range of 410 to 440 °C for conventional and 380 to 410 °C for cesium-promoted grades, and an upper range of 600 to 650 °C, above which activity can be lost for good.

(a) Compare the design inlets and bed 1's design outlet with these ranges.

(b) In case 3 bed 1's outlet peaks at 653.4 °C after its inlet is cooled, and at 677.8 °C in the 48-cell model. What does this mean for an operator who tests the floor on a real converter?

(c) Where should a real plant take its floor, and what does the MPC's margin of 15 K above 395 °C give against the conventional range?

4.12 Review the agent · 5 min

After case 7 ages the catalyst to 60 % activity, an AI assistant advises: "On an exothermic equilibrium the coolest inlet that still lights the catalyst converts the most. Lower the four inlet setpoints to 410, 430, 425 and 420 °C, all above the 395 °C floor, and the extra conversion will bring the stack back under 400 ppmv." Find the flaw and say how you would check it.

CHAPTER 5

Chlor-Alkali Membrane Electrolyzer Simulator Fed in Ratio to Its Current

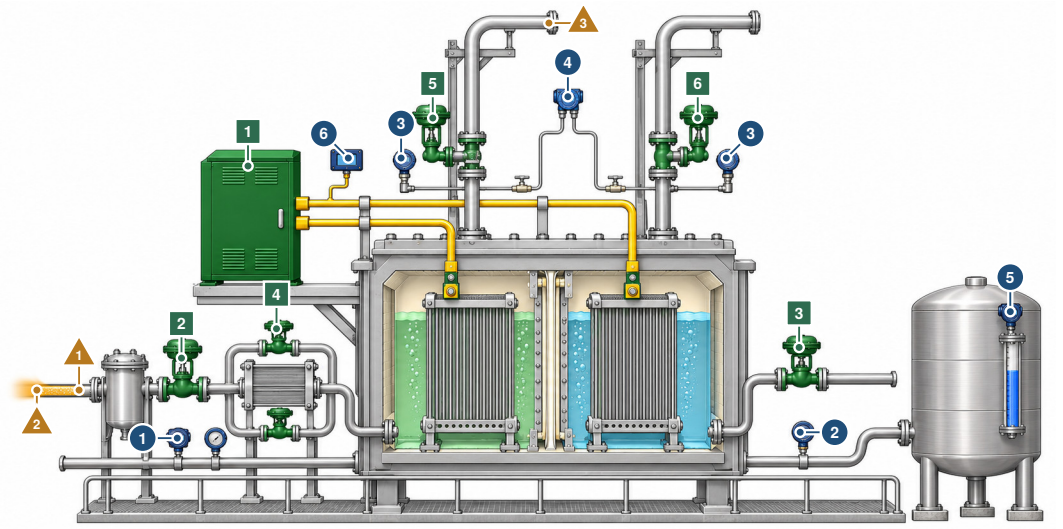

The controller moves 1 rectifier, which sets the current and so the production 2 brine feed valve, in ratio to the current 3 water feed valve, in ratio to the current 4 heater and cooler on the brine 5 chlorine valve to the compressor 6 hydrogen valve

The controller reads 1 brine strength (density) and cell temperature 2 caustic strength (conductivity) 3 chlorine and hydrogen pressures 4 hydrogen-over-chlorine pressure difference 5 caustic tank level 6 voltage of every cell

What disturbs it 1 hardness the ion exchange lets through 2 weaker or colder brine 3 chlorine compressor trip

Notice the two gas pipes and the transmitter between them: the hydrogen must stay a few tens of millibars above the chlorine while every product follows the one current from the rectifier. Illustration, not to scale.

IN PLAIN TERMS Feeding the Cells in Step With the Current

The main idea. Electricity splits salt water into chlorine, hydrogen and caustic soda. Every bit of product is made by the electric current, so the salt water and the fresh water must be fed in step with it.

The challenge. A fixed feeding rule is never quite right, so measured strengths must correct it. The hydrogen must also stay a little above the chlorine in pressure, and when the chlorine stops being drawn away, a controller that waits to see the gap shrink acts too late.

Why it matters. Too small a gap stops the plant to protect the thin sheet between the two sides. Some harm to that sheet can never be undone.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page; the fact box of Section 5.1; the history box of Section 5.2; the fact box of Section 5.5; the try-it boxes of Section 5.6; the key points
Study	50 min	the whole chapter with its pause points and the exercises
Practice	15 min	the first page; Section 5.6 with its result chart; Section 5.7; Section 5.8 with its open challenge

The control question A membrane electrolyzer makes every product in proportion to one current, and its hydrogen must stay a few tens of millibars above its chlorine. What do feeds in ratio to the current, a chlorine pressure feedforward on the hydrogen header and a load plan on the electricity price each buy, and which losses can no controller recover?

Learning objectives

- Compute the products, the oxygen in the chlorine and the specific energy of a membrane electrolyzer from Faraday's law and its current efficiencies.
- Explain why a ratio station needs an analyzer trim, and what the ratio alone and the trim alone each get wrong.
- Estimate how long a feedback-only pressure-difference loop rides through a compressor trip, and explain why a chlorine pressure feedforward removes the race.
- Relate the hydrogen header volume to the fastest rise the hydrogen can make, and find the volume above which no tuning keeps pace.
- Evaluate what a load scheduler saves on an hourly price, what limits it, and what permanent membrane damage costs a day.

Prerequisites Sections 1.4.6 (ratio stations with a trim), 1.6.2 (velocity-form PI) and 1.6.4 (scenarios, scorecard and exported files) of Chapter 1. Chapter 3 holds a second ratio on measured feeds.

The app https://apmonitor.com/apps/chlor-alkali/

See also Volume 3, Chapter 4 for an air-to-fuel ratio trimmed on an oxygen analyzer, Volume 3, Chapter 7 for a ratio blend trimmed on a refractometer, and Volume 3, Chapter 6 for a dose computed from a measured cause. Volume 9, Chapter 11 for a pumping schedule on the electricity price, and Volume 9, Chapter 6 for heat priced across a network. Volume 4, Chapter 2 for a protective loop on a compressor, and Volume 4, Chapter 6 for a safety limit that bounds the throughput.



5.1 Introduction

Chlorine and caustic soda are made together by the electrolysis of brine, in an industry that has run since the nineteenth century on three technologies, the diaphragm cell, the mercury cell and the membrane cell, of which the membrane cell is the most energy efficient and the only one free of asbestos and mercury [1]. Brine electrolysis is the predominant industrial use of ion-selective membranes [2]. The membrane passes sodium from a brine anolyte to a caustic catholyte and keeps the two gases apart, and Faraday's law ties every product to the one current that crosses it. In the electrolyzer of this chapter, 16.2 kA through 150 cells in series makes 3.45 t/h of NaOH, 3.18 t/h of chlorine and 91 kg/h of hydrogen for 7.54 MW of direct current, 2188 kWh/t for every tonne of caustic.

> **+ DID YOU KNOW? World Chlorine Around 2016**
>
> Around 2016 the world could make roughly 76 million tonnes of chlorine a year, and more than half of it came from membrane electrolysis [2]. Each of those tonnes comes with its caustic soda and hydrogen in the fixed proportions of Faraday's law, which Section 5.3 applies to one electrolyzer.

The simulator is aimed at a common belief about cell rooms, that the plant is under control once its strengths sit on their setpoints. Because the current is the throughput, a feed that is not tied to it drifts whenever the load moves. In case 1 a morning load-up from 5.0 to 6.5 kA/m^2, with the feeds left where they were, takes the anolyte from 210 to 177.9 g/L and sends 81.0 t of the day's 86.8 t of caustic off specification. The gas side is faster than any analyzer. When the chlorine compressor stops, the chlorine header rises 1.11 mbar/s, and a hydrogen loop that learns of it only through the pressure difference loses its 15 mbar margin to the interlock in 20 s. Some losses are permanent. In case 6 three hours of hardness raise the cell voltage by 75.4 mV for the rest of the membrane's life while every loop is back on its setpoint. None of this appears in a transfer function from a brine valve

to a densitometer, which knows nothing of the load, of the race between the two headers, or of a loss that no setpoint recovers.

This class of problem reaches well beyond chlorine. A plant whose production rate is a measured and exactly known quantity can feed its reagents forward in ratio to it and trim the ratios on a composition, which is textbook ratio control [3] and the subject of Section 1.4.6. The fired heater of Volume 3, Chapter 4 trims its air-to-fuel ratio on an oxygen analyzer in the same way. The same electrochemical structure recurs in water electrolysis, fuel cells and the smaller electromembrane processes [2]. Every two-compartment device whose membrane must not see a reversed pressure difference shares the gas-side problem, and every power-intensive process that can move its load within hours is a candidate for demand response [4], as the pumping station of Volume 9, Chapter 11 is. The chemistry sets the limits the controllers work inside. Oxygen evolution competes with chlorine at the anode [5], impurities in the brine degrade the membrane [6], and the cell voltage and current efficiency depend on the process conditions [7].

The published work on the membrane cell relates its voltage and efficiency to the operating conditions [7], and the demand-response studies schedule the plant as a load over hours [8–11]. Neither gives a plant in which the ratio stations, the millibar loop, the interlock, the permanent damage and the price schedule interact, and in which each can be taken away to see what fails. This chapter audits, in tonnes, dollars and millibars, what each feedforward buys in a membrane electrolyzer, on a browser simulator in which every feed follows the current, the hydrogen follows the chlorine and the load follows the price. Faraday's law, ratio control with an analyzer trim, cascade pressure control, an interlock and a scheduling linear program (LP) are established practice. The analysis gives a closed-form estimate of how long the pressure-difference loop can ride through a compressor trip on feedback alone, 21.5 s at the shipped tuning against 20 s on the engine, and a ratio of the two headers' fastest rise rates, 1.207 at every current, which fixes a hydrogen header volume of 301.8 m^3 above which the hydrogen cannot keep pace with the chlorine even with its valve shut. The model has 17 physical states and 33 ledgers that close to 4.0×10^{-13}, and the controls use the framework's shared sampled analyzer, an interlock on its own transmitter and a bounded-variable simplex that plans the day in about a millisecond, all written down so that the model can be rebuilt from this text. The eight cases and the graded day set a ratio apart from its trim, a feedforward from a faster feedback, a loss the loops can correct from one they cannot, and a cheap hour from a tank with no room left.

Section 5.2 traces the industry from mercury and asbestos to the ion-exchange membrane and to load management on the electricity price. Section 5.3 sets out the balances, the voltage and efficiency correlations, the heat balance, the gas headers and the tank. Section 5.4 describes the instruments, the ratio stations and their trims, the pressure cascade with its feedforward, the interlock, the rectifier and the load scheduler. Section 5.5 derives how long the hydrogen can follow a tripped chlorine compressor on feedback alone and checks the estimate on the engine. Section 5.6 measures the eight cases and the graded day, Section 5.7 names the hardware and the practice behind the model, Section 5.8 states where the model stops holding, and Section 5.9 collects the results.

5.2 Mercury, Asbestos and the Ion-Exchange Membrane

Electrolytic chlorine is older than the membrane by most of a century. The diaphragm and mercury processes carried the industry for more than a hundred years and remain the most used worldwide, and both proved the most harmful to the environment through their asbestos and their mercury, while the membrane process was developed within the last sixty years [1]. The history of the industry, its brine preparation and treatment, the design of its cell rooms and their instrumentation and control systems are collected in one handbook [12].

The history box at the end of this section traces the membrane from the first patent to its adoption. The membrane brought constraints of its own. Brine for membrane cells is polished by ion exchange after its primary treatment [13], because the impurities that reach the membrane cost the cell its

performance [6]. The anode evolves some oxygen beside its chlorine, and the selectivity between the two has been reviewed at length [5], while the effects of the process conditions on the cell voltage, the current efficiency and the voltage balance of a membrane cell have been studied in their own right [7]. Oxygen-depolarized cathodes, which reduce oxygen at the cathode instead of evolving hydrogen and so lower the cell voltage, have a history, a present status and prospects of their own [14].

Control reached beyond the cell room through the electricity bill. The electrolytic process is highly energy intensive, and load management under a time-of-use tariff was posed for a caustic-chlorine plant as a mixed-integer nonlinear program that cut its peak demand by about 19 % and its electricity cost by about 3.9 % [15]. Later work treated chemical plants as grid-level batteries that bid into short-term electricity markets [8], built dynamic models of membrane chlor-alkali plants for demand response [9], let a plant switch between operating modes with an oversized cell room [10], and integrated the scheduling with the control under price and demand uncertainty [11], inside a wider literature on industrial demand-side management [4]. Those schedules sit on regulatory loops that the scheduling papers take as given. In the simulator the two layers share one plant, and the true strengths, efficiency and pressure difference are logged beside every instrument that reports them.

> ⧖ HISTORY **From Diaphragms and Mercury to the Membrane, 1851 to 1986**
>
> Making chlorine, caustic soda and hypochlorite by the electrolysis of brine was patented in England by Watt in 1851. In 1890 production began in a cell whose anode and cathode were separated by a porous cement diaphragm, and porous separators, asbestos among them, or a mercury cathode then kept the technology unchanged for decades. In 1970 a safer and more efficient technology was sought. Chemically stable perfluorinated sulfonated membranes of low resistance had appeared at the end of the 1960s, and industrial electrolysis with perfluorinated membranes began in Japan in 1975 [2]. An account of the ion-exchange membrane chlor-alkali process followed in 1976 [16]. In 1986 Japan changed all of its mercury cell production to membranes, and the process spread quickly to other industrial countries. A 2016 review called brine electrolysis the predominant industrial use of ion-selective membranes [2].

5.3 The Electrolyzer as Modeled

The simulated plant is one bipolar electrolyzer of $N = 150$ cells of $A = 2.7\,\mathrm{m}^2$ each at a design current density of $6\,\mathrm{kA/m}^2$, making 32 % caustic at 88 °C from brine of 305 g/L that leaves at 210 g/L. Figure 5.1 shows the loops, the headers, the tank, every instrument and the states and flows used below. Table 5.1 classifies the variables. The application integrates in minutes and quotes flows as a plant quotes them, while the equations below are written in consistent units. Each balance is integrated together with a ledger of what entered and left, so every inventory can be closed from the state itself.

5.3.1 Faraday's Law and the Products

All N cells carry the same current $I = iA$, so the electrons through the stack, $\dot{n}_e = NI/F$, fix every product, as shown in Equation (5.1).

$$\begin{aligned} \dot{n}_{H_2} &= \frac{\dot{n}_e}{2}, \qquad \dot{n}_{Cl_2} = \eta_{Cl}\frac{\dot{n}_e}{2}, \qquad \dot{n}_{O_2} = (1-\eta_{Cl})\frac{\dot{n}_e}{4} \\ \dot{m}_{NaOH} &= \eta\,\dot{n}_e M_{NaOH}, \qquad y_{O_2} = \frac{\dot{n}_{O_2}}{\dot{n}_{Cl_2}+\dot{n}_{O_2}} = \frac{1-\eta_{Cl}}{1+\eta_{Cl}} \end{aligned} \tag{5.1}$$

where $F = 96\,485.332\,\mathrm{C/mol}$ is Faraday's constant, the current efficiency η is the fraction of the current that the membrane gives to sodium, and $1-\eta_{Cl}$ is the fraction of the anode current that oxidizes water to oxygen. The brine is taken as acidified, so the hydroxide that migrates back through the membrane is neutralized in the anolyte and makes no oxygen, which gives chlorine and caustic separate efficiencies. The design point has $\eta = 95.00\,\%$ and $\eta_{Cl} = 98.8\,\%$, which puts 0.60 % oxygen by volume in the dry chlorine, and the stack makes 3.445 t/h of NaOH, 3.18 t/h of chlorine and 91.4 kg/h of hydrogen.

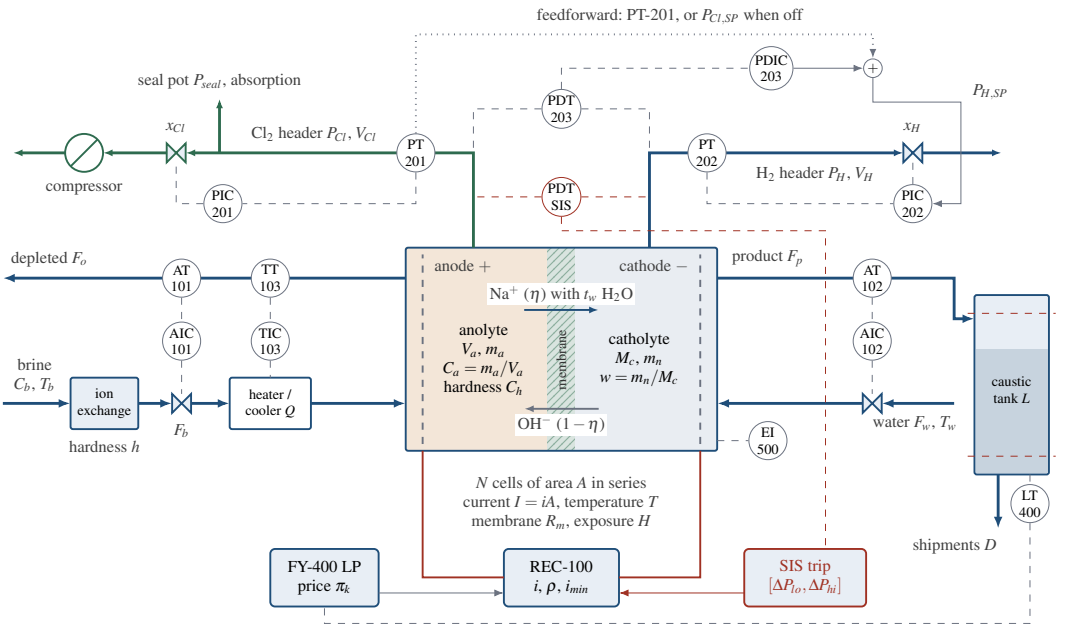

Figure 5.1 Membrane electrolyzer as modeled. Brine of strength C_b and temperature T_b passes the ion exchange, which lets through the hardness h, the brine valve F_b and the heater or cooler Q into the anolyte (V_a, m_a, C_h), and depleted brine F_o overflows past AT-101 and TT-103. Water F_w enters the catholyte (M_c, m_n), whose product F_p overflows past AT-102 to the caustic tank L and its shipments D. The membrane passes a share η of the current as sodium with t_w waters each and returns the rest as hydroxide. Chlorine and its oxygen leave through header P_{Cl} to the compressor's suction valve x_{Cl}, with a seal pot to the emergency absorption, and hydrogen through header P_H and its valve x_H. PDT-203 feeds the master PDIC-203, whose output is added to PT-201 (feedforward on) or to the chlorine setpoint (off) to form the setpoint of PIC-202, while the interlock's own transmitter trips the rectifier REC-100. FY-400 plans the current on the price and the tank gauge LT-400, and EI-500 scans the cell voltages. Measurement lines are dashed.

5.3.2 Anolyte and Catholyte

The anolyte is a well-mixed loop fed with brine and emptied over a weir, as shown in Equation (5.2).

$$\begin{aligned}
\frac{dV_a}{dt} &= F_b - F_o - \frac{\dot m_d + \dot m_{v,a}}{\rho_w} - v_s \dot m_{tr} \\
\frac{dm_a}{dt} &= F_b C_b - F_o C_a - \dot m_{tr}, \qquad \frac{d(V_a C_h)}{dt} = F_b h - F_o C_h \\
F_o &= k_{oa} \max(0, V_a - V_{a,0}), \qquad \dot m_{tr} = \eta\, \dot n_e M_{NaCl} \\
\dot m_d &= t_w\, \eta\, \dot n_e M_{H_2O}, \qquad t_w = t_{w,0} + t_{w,a}(C_{ref} - C_a)
\end{aligned} \tag{5.2}$$

where $C_a = m_a/V_a$ is the NaCl in kg/m^3 (numerically g/L), $\dot m_{tr}$ is the salt the membrane removes as sodium, counted as NaCl, $\dot m_d$ the water it drags with each sodium ion, $\dot m_{v,a}$ the vapor that leaves with the chlorine, $v_s = 0.33 \times 10^{-3}\,\mathrm{m^3/kg}$ the apparent volume of dissolved salt, $\rho_w = 997\,\mathrm{kg/m^3}$, and C_h the hardness of the anolyte in ppb. The weir overflows with a 2 min time constant ($k_{oa} = 0.5\,\mathrm{min^{-1}}$), and at the design point the anolyte stands $0.92\,\mathrm{m^3}$ above it. The drag term above $t_{w,0}$ is an illustrative osmotic correction that grows as the anolyte weakens against the strong caustic, and it couples the two loops. The brine density is $\rho_b(C) = \rho_w(1 - v_s C) + C$. The catholyte gains the sodium and its water, loses one water per electron at the cathode, of which only the half hydrogen leaves, and loses

Table 5.1 Variables of the Electrolyzer and Their Design Values

Variable	Role, instrument	Design	Range
C_a	CV, AT-101 (density)	210 g/L SP	170–240
w	CV, AT-102 (conductivity)	32 % SP	28–36
T	CV, TT-103	88 °C SP	80–92
P_{Cl}	CV, PT-201	150 mbar(g) SP	80–250
$\Delta P = P_H - P_{Cl}$	CV, PDT-203	35 mbar SP	22–48
P_H	CV, PT-202 (slave)	185 mbar(g)	0–450
i	CV and MV, REC-100	6.0 kA/m^2	2–7
L	CV of FY-400, LT-400	30 t	10–90 %
F_b, F_w	MVs, feeds (m^3/h)	35.46, 3.869	0–60, 0–8
Q	MV, heater (+), cooler	0.093 MW	±1.5
x_{Cl}, x_H	MVs, header valves	52.6, 50.0 %	0–100
C_b	DV, feed brine	305 g/L	270–315
h	DV, hardness in the feed	10 ppb	0–2000
T_b, T_{cw}	DVs, brine, cooling water	65, 25 °C	30–75, 15–40
π	DV, price, spike ($/MWh)	18–165	+0–500
—	DVs, events	compressor, breakthrough, cell	
V_a, m_a, C_h	states, anolyte	18.92, C_aV_a, 12.9	—
M_c, m_n	states, catholyte	20.36 t, wM_c	—
T, R_m, H	states, cell, membrane	88, 0.06375, 0	—
P_{Cl}, P_H, x_{Cl}, x_H	states, headers, valves	as above	—
F_b, F_w, Q	states, actuator lags	as above	—
L, S	states, tank NaOH, solution	30, 93.75 t	—

Units: m^3 for V_a, ppb for C_h, mΩ m^2 for R_m and ppb h for H. The tank starts at 24 of 30 t in case 8 and the challenge.

the hydroxide that migrates back and the vapor carried by the hydrogen, as shown in Equation (5.3).

$$\begin{aligned}\frac{\mathrm{d}M_c}{\mathrm{d}t} &= \rho_w F_w + \dot{n}_e\left[\eta M_{Na} - (1-\eta)M_{OH} - \tfrac{1}{2}M_{H_2}\right] \\ &\quad + \dot{m}_d - \dot{m}_{v,c} - F_p \\ \frac{\mathrm{d}m_n}{\mathrm{d}t} &= \eta\,\dot{n}_e M_{NaOH} - wF_p \\ F_p &= k_{oc}\max(0, M_c - M_{c,0})\end{aligned} \tag{5.3}$$

where $w = m_n/M_c$. The catholyte's water closes exactly, because $\frac{1}{2}M_{H_2} + M_{OH} = M_{H_2O}$ and $M_{Na} + M_{OH} = M_{NaOH}$. At the design point the membranes carry 5.03 t/h of salt out of the anolyte and drag 5.43 t/h of water with it, the brine feed is 35.46 m^3/h and the depleted brine 27.53 m^3/h, and the catholyte needs 3.869 m^3/h of water to stay at 32 %. The anolyte turns over in 41 min and the catholyte in 113 min, and every one of these flows is proportional to the current. The catholyte weir has the same constant and holds 0.36 t above its level at the design point.

5.3.3 Membrane Resistance and Hardness

Calcium and magnesium above a precipitation threshold damage the membrane in proportion to the exposure and the current, and irreversibly, as shown in Equation (5.4).

$$\frac{\mathrm{d}H}{\mathrm{d}t} = \max(0, C_h - h_{thr})\,\frac{i}{i_{nom}}, \qquad \frac{\mathrm{d}R_m}{\mathrm{d}t} = \frac{r_{age}}{t_{yr}} + k_{HR}\frac{\mathrm{d}H}{\mathrm{d}t} \tag{5.4}$$

where H is the exposure in ppb h, t_{yr} is one year, so that the creep adds only 6.8×10^{-6} mΩm^2 a day, and R_m is the membrane's area resistance referred to 88 °C and 32 %, which starts at $R_{m,0} + r_{age}\,\mathrm{age} = 0.063\,75$ mΩm^2 for an 18-month-old membrane. The threshold and both coefficients are illustrative.

5.3.4 Cell Voltage and Current Efficiency

The cell voltage is an illustrative correlation with area resistances in $\mathrm{m\Omega\,m^2}$, so that i in $\mathrm{kA/m^2}$ times a resistance is volts, as shown in Equation (5.5).

$$\begin{aligned} U &= U_{rev} + b\log_{10}\frac{\max(i,0.05)}{i_0} \\ &\quad + i\big(R_m f_T(1+k_m\Delta w) + R_{el} f_T g_{el} + R_{hw}\big) \\ U_{rev} &= U_0 + \beta_T\Delta T + k_{Uw}\Delta w + k_{Ua}(C_{ref} - C_a) \\ g_{el} &= s_a\Big(\frac{C_{ref}}{\max(50,C_a)}\Big)^{0.9} + (1-s_a)(1+k_{cw}\Delta w) \end{aligned} \tag{5.5}$$

where $\Delta T = T - 88$, $\Delta w = w - 32$ and $f_T = \mathrm{e}^{-\alpha_T\Delta T}$. At the design point the 3.103 V of a cell divide into 2.200 V reversible, 0.180 V at the electrodes, 0.383 V in the membrane, 0.270 V in the electrolyte and bubbles and 0.070 V in the hardware, and the membrane and electrolyte resistances fall about 1.5 % per kelvin. The current efficiency is an illustrative correlation with two knees, as shown in Equation (5.6).

$$\begin{aligned} \eta_{max} &= \eta_0 - a\,\mathrm{age} - k_1\max(0,w-w_k)^2 \\ &\quad - k_2 H - k_3\max(0,C_k - C_a) \\ \eta &= \mathrm{clamp}\big(\eta_{max} - j_d/\max(i,0.1),\ 0.5,\ 0.999\big) \\ \eta_{Cl} &= 1 - \min\big(0.5,\ \phi_0\,\mathrm{e}^{(C_{ref}-C_a)/\phi_s}\big) \end{aligned} \tag{5.6}$$

where the knees sit at $w_k = 32\,\%$ and $C_k = 200\,\mathrm{g/L}$, and j_d/i is hydroxide back-diffusion that does not scale with the current, which makes the efficiency fall at low load, from 95.00 % at 6 $\mathrm{kA/m^2}$ to 94.50 % at 3 $\mathrm{kA/m^2}$. The specific energy follows as $SE = NUI/\dot m_{NaOH} = UF/(\eta M_{NaOH})$, 2188 kWh/t at the design point, 1927 kWh/t at 3 $\mathrm{kA/m^2}$ and 2276 kWh/t at 7 $\mathrm{kA/m^2}$.

5.3.5 Heat Balance

The cells generate heat above the thermoneutral voltage $U_{tn} = \Delta H/2F = 2.31\,\mathrm{V}$, from $\Delta H = 446\,\mathrm{kJ}$ per mole of chlorine, and the feeds, the vapor, the room and the heater or cooler take it away, as shown in Equation (5.7).

$$\begin{aligned} C_{th}\frac{\mathrm{d}T}{\mathrm{d}t} &= I\,(NU - \Delta U_{fs} - NU_{tn}) + F_b\rho_b(C_b)c_{p,b}(T_b - T) \\ &\quad + \rho_w F_w c_{p,w}(T_w - T) - (\dot m_{v,a} + \dot m_{v,c})h_{fg}(T) \\ &\quad - UA\,(T - T_{amb}) + Q \end{aligned} \tag{5.7}$$

where ΔU_{fs} is the voltage a failing cell loses (below), $c_{p,w} = 4.18\,\mathrm{kJ/(kg\,K)}$ and $h_{fg}(T) = 2501 - 2.37T$ in kJ/kg. Both gases leave saturated with water, as shown in Equation (5.8).

$$\begin{aligned} \dot m_{v,a} &= (\dot n_{Cl_2} + \dot n_{O_2})M_{H_2O}\frac{p_a}{P_{Cl}^{abs} - p_a} \\ p_a &= \min\big(a_b\,p^{sat}(T),\ 0.95P_{Cl}^{abs}\big) \\ \log_{10}\frac{p^{sat}}{0.133\,322\,\mathrm{kPa}} &= 8.07131 - \frac{1730.63}{233.426+T} \end{aligned} \tag{5.8}$$

with $\dot m_{v,c}$ likewise from $\dot n_{H_2}$, P_H^{abs} and a_c, the illustrative water activities $a_b = 1 - 0.00047\,C_a$ and $a_c = 1 - 0.0145\,w$ (with w in percent), and absolute pressures from $P_{atm} = 1013.25\,\mathrm{mbar}$. At the design point 1.93 MW of the 7.54 MW, a quarter, is heat, of which the brine takes 0.90 MW, the vapor 0.73 MW, the water 0.26 MW and the room 0.13 MW, while the heater adds 0.09 MW. The vapor is 820 kg/h with the chlorine and 334 kg/h with the hydrogen, and its cooling grows by 51 kW

per kelvin of cell temperature, against 39 for the brine, while the heat generated falls by 26 kW per kelvin as the voltage drops. At fixed strengths these slopes give the heat balance a time constant of 20 min.

5.3.6 Gas Headers

Each header is an ideal-gas volume on a dry basis after the coolers, filled by the cells and emptied through its valve, as shown in Equation (5.9).

$$\begin{aligned} \frac{V_{Cl}}{RT_g}\frac{\mathrm{d}P_{Cl}}{\mathrm{d}t} &= \dot{n}_{Cl_2} + \dot{n}_{O_2} - \kappa C_{v,Cl} x_{Cl}\sqrt{(P_{Cl} - P_s)} \\ &\quad - C_{abs}\max(0, P_{Cl} - P_{seal}) \\ \frac{V_H}{RT_g}\frac{\mathrm{d}P_H}{\mathrm{d}t} &= \dot{n}_{H_2} - C_{v,H} x_H \sqrt{(P_H - P_d)} \\ &\quad - C_{rel}\max(0, P_H - P_{rel}) \end{aligned} \tag{5.9}$$

where the pressures are gauge in mbar, $\sqrt{(d)}$ is $\sqrt{d}$ above 5 mbar, $d/\sqrt{5}$ between 0 and 5 mbar and zero below, and κ is the compressor's share of its suction capacity, zero while it is tripped and recovering linearly to one over 1 min after a restart. At the design point the chlorine header is 150 mbar(g) with its valve at 52.6 %, the hydrogen header 185 mbar(g) with its valve at 50.0 %, and the difference $\Delta P = P_H - P_{Cl}$ is 35 mbar.

5.3.7 Feeds, Actuators, Tank and Events

The brine and water flows respond to their commands as first-order lags of 20 s, the heater duty as one of 1 min and the two valves as lags of 3 s. The caustic tank keeps its NaOH L and its solution S, as shown in Equation (5.10).

$$\frac{\mathrm{d}L}{\mathrm{d}t} = wF_p - D, \qquad \frac{\mathrm{d}S}{\mathrm{d}t} = F_p - \frac{D}{L/S} \tag{5.10}$$

Shipments $D = 3.0\,\mathrm{t/h}$ stop when the tank holds 0.2 t and resume at 1 t, and product reaches the tank on specification while $|w - 32| \le 0.5$ in percent. Hard events are handled after each integration step. The interlock reads its own transmitter, ΔP with 0.2 mbar of noise, once a second, and opens the rectifier when two consecutive scans fall outside $[\Delta P_{lo}, \Delta P_{hi}] = [20, 50]$ mbar. A restart is refused for 30 min and while ΔP is outside the band, and the rectifier then ramps up from zero. The band and its two-scan delay are teaching values that follow the handbook's treatment of the cell room [12]. Real limits come from the electrolyzer's vendor and are implemented as a safety instrumented function designed under IEC 61511. A failing cell, cell 87, opens a pinhole whose severity s grows linearly from 0 to 1 over 120 min. The cell then loses $\Delta U_f s = s \times 0.25\,\mathrm{V}$ and a share $\delta_\eta s = 0.30\,s$ of its efficiency, which lowers the stack's efficiency by $\delta_\eta s/N$, 0.2 % at full failure.

5.3.8 Initial State, Step and Replay

Table 5.2 collects the parameters. Every scenario starts from the explicit steady state of its own configuration, including any setting changed before the run starts. Given the current density, the temperature, both strengths, the brine and the header pressures, the balances give the brine and water flows, the heater duty and the two valve positions that hold the plant still, with no iteration because the temperature and the strengths are given. At that state no derivative of the loop inventories, the temperature, the headers or the actuator lags exceeds 1×10^{-16} relative, where the application's own check asks for 1×10^{-10}. Only the tank, which the cells fill faster than the shipments draw it at the design point, and the slow creep of the membrane with age are not stationary. The ratio stations as shipped are not the ones the steady state needs, 2.15 m³/h of brine and 0.251 m³/h of water per kA against 2.189 and 0.239, errors of −1.8 and +5.1 %, so the trims start out carrying the difference.

Table 5.2 Model Parameters. Class M Entries Can Be Changed in the Show Model Dialog, While Class H Entries Stay Hidden

Symbol	Meaning	Value	Class	Symbol	Meaning	Value	Class
N, A	cells, area per cell	150, 2.7 m^2	H	$V_{a,0}$, k_{oa}	anolyte weir holdup, overflow	18 m^3, 0.5 min^{-1}	H
i_{nom}, i_{max}	design, rectifier limit	6.0, 7.0 kA/m^2	H, M	$M_{c,0}$, k_{oc}	catholyte weir holdup, overflow	20 t, 0.5 min^{-1}	H
P_{max}	rectifier rating (trend scale)	10 MW	H	T_w	water temperature	30 °C	H
age	membrane age	18 months	M	$F_{b,max}$, $F_{w,max}$	feed capacities	60, 8 m^3/h	M
U_0, U_{tn}	reversible, thermoneutral	2.20, 2.31 V	M	—	flow, heater, valve lags	20 s, 1 min, 3 s	H
b, i_0	Tafel slope, exchange current	0.080, 0.033 75 kA/m^2	M, H	C_{th}	heat capacity	150 MJ/K	M
$R_{m,0}$, r_{age}	new membrane, creep per year	0.060, 0.0025 mΩ m^2	M	UA, T_{amb}	loss to the room	2 kW/K, 25 °C	H
R_{el}, R_{hw}	electrolyte, hardware	0.045, 0.0117 mΩ m^2	M	$c_{p,b}$	brine heat capacity	3.3 kJ/(kg K)	H
α_T, β_T	temperature coefficients	0.015 K^{-1}, –0.9 mV/K	H	Q_{heat}, Q_{cool}	heater, cooler at 25 °C	1.5, 1.5 MW	M
k_{Uw}, k_{Ua}	U_{rev} per %, per g/L	5, 0.3 mV	H	T_{max}	temperature alarm	92 °C	H
k_m, k_{cw}, s_a	per % caustic, anolyte share	0.03, 0.03, 0.4	H	V_{Cl}, V_H	header gas volumes	300, 250 m^3	M
η_0, a	new efficiency, loss per year	95.8, 0.2 %	M, H	T_g, τ_{comp}	gas temperature, reload	320 K, 1 min	H
k_1, w_k	caustic knee	0.004 $\%^{-2}$, 32 %	M, H	$C_{v,Cl}$, $C_{v,H}$	valves, mol/s per $\sqrt{\text{mbar}}$	1.505, 1.853	H
k_2	loss per exposure	1×10^{-5} $(\text{ppbh})^{-1}$	M	P_s, P_d	suction, hydrogen outlet	−100, 0 mbar(g)	H
k_3, C_k	anolyte knee	0.0006 $(\text{g/L})^{-1}$, 200 g/L	M, H	P_{seal}, C_{abs}	seal pot, mol/s per mbar	300 mbar(g), 0.6	M, H
j_d	back-diffusion	0.03 kA/m^2	M	P_{rel}, C_{rel}	hydrogen relief	450 mbar(g), 0.6	H
ϕ_0, ϕ_s, C_{ref}	oxygen share, scale, reference	1.2 %, 30, 210 g/L	M, H, H	ΔP_{lo}, ΔP_{hi}	interlock band (teaching)	20, 50 mbar	M
$t_{w,0}$, $t_{w,a}$	drag per Na^+, per g/L	3.5, 0.01	M	n_{SIS}, t_{rst}	scans outside, restart	2, 30 min	M
h_{thr}, k_{HR}	threshold, damage per ppb h	30 ppb, 7×10^{-6} mΩ m^2	M	L_{cap}, D	tank capacity, shipments	60 t[a], 3.0 t/h	M
σ_a, σ_w	analyzer repeatability	0.6 g/L, 0.03 %	M	L_{lo}, L_{hi}, L_0	tank band, start	10, 90, 50 %[a]	M
θ_a, Δt_a	sample loops, update	1 min, 30 s	H	δ_L, δ_m	scheduler margins	0.5, 0.3 t	H
—	TT-103 lag, noise	30 s, 0.05 °C	H	ΔU_f, δ_η, τ_f	failing cell 87	0.25 V, 0.30, 120 min	M, H, M
—	PT, PDT, SIS, LT noise	0.1, 0.1, 0.2 mbar, 0.05 t	H	—	EI-500 scan, noise, spread	30 s, 2, 4 mV	H, H, M

[a] 30 t starting at 80 % in case 8 and the challenge. U_0 is taken at 88 °C, 32 % and 210 g/L. Constants: $F = 96\,485.332\,\text{C/mol}$, $R = 8.314\,462\,\text{J/(mol K)}$, and molar masses of NaOH, NaCl, Cl_2, H_2, O_2 and H_2O of 39.997, 58.443, 70.906, 2.016, 31.998 and 18.015 g/mol.

The engine of Section 1.6.1 advances the electrolyzer by fourth-order Runge-Kutta in steps of one second. The controllers run every 5 s and the historian writes a row every 30 s, while a ring buffer keeps the true ΔP once a second for the scene's scope. The speed selector spans 1 to 1200 minutes of plant time per second, with 10 as the default, and the seed is 20261127. With every instrument noise-free, halving the step moves the lowest ΔP of the compressor trip with the feedforward on by 0.064 mbar and the challenge's electricity cost by less than 1×10^{-6} %. Over the challenge day the salt, volume, hardness, catholyte mass, NaOH, water, tank and header ledgers close to 4.0×10^{-13} relative, far inside the 1×10^{-9} that one of the application's 42 checks demands.

5.3.9 Assumptions

One lumped cell stands for all 150, with the same current, strengths and temperature, so end cells, gradients along the stack and uneven flows are absent, and only the monitor sees the cells apart, through a fixed cell-to-cell fingerprint of 4 mV and the one failing cell. The voltage, efficiency, water-activity, oxygen and hardness correlations are illustrative teaching choices fitted to one design point, valid between about 2 and 8 kA/m^2, 80 and 92 °C, 28 and 36 % and 150 and 240 g/L. The application itself notes that real damage depends on calcium, magnesium, strontium, barium, iodide and silica together and on the membrane type, and the impurities of the brine set the performance of a membrane cell in general [6]. The brine is acidified without the hydrochloric acid dosing being modeled, and chlorate, hypochlorite, dechlorination and resaturation are left out, so the saturators appear only as the strength and temperature of the brine they deliver.

The heat balance uses the thermoneutral voltage at 25 °C without a heat-capacity correction and one lumped heat capacity. The headers are ideal gases on a dry basis, the compressor is a suction capacity, and the emergency absorption has no limit. Chlorine and hydrogen are taken by their consumers without limit, so the caustic tank is the only storage the scheduler sees, while in many real plants the chlorine outlets limit flexibility more than the caustic does. Electricity is priced on the direct current with rectifier losses neglected, and the steam is not priced. The price profiles are synthetic.

Table 5.3 Loops, Directions and Shipped Tunings

Tag	CV → MV	γ	K_c	τ_I	Output
AIC-101	$C_a \to F_b$, bias $r_b I$	+1	1.0	30 min	0–60 m³/h
AIC-102	$w \to F_w$, bias $r_w I$	−1	2.0	60 min	0–8 m³/h
TIC-103	$T \to Q$	+1	0.25	20 min	$-Q_{cool}(T_{cw})$–1.5 MW
PIC-201	$P_{Cl} \to x_{Cl}$	−1	2	2 min	0–100 %
PDIC-203	$\Delta P \to u_{\Delta P}$	+1	0.5	1 min	0–150 mbar
PIC-202	$P_H \to x_H$	−1	6	1 min	0–100 %

Gains in m³/h per g/L, m³/h per %, MW/K, % per mbar, mbar per mbar and % per mbar, in the order of the rows.

5.4 Instruments and Implemented Controllers

5.4.1 Instruments

No controller reads the plant's state. AT-101 reports the anolyte strength that a densitometer infers from $\rho_b(C_a)$, and AT-102 the caustic strength from conductivity, through the framework's shared sampled analyzer, each with a 1 min sample loop, a result every 30 s describing a sample that left the cell at least 1.5 min earlier, repeatabilities of 0.6 g/L and 0.03 %, and a STALE flag when a result has not been refreshed for four cycles. TT-103 lags the cell by 30 s with 0.05 °C of noise, the two header transmitters and PDT-203 carry 0.1 mbar of noise, LT-400 carries 0.05 t, and the rectifier knows its own current exactly and measures the stack voltage. EI-500 scans every cell every 30 s with 2 mV of noise.

5.4.2 The Shared PI Law

All six loops run the framework's velocity-form PI of Section 1.6.2, written for this chapter as Equation (5.11).

$$\begin{aligned} e_k &= \gamma(\mathrm{SP}_k - \mathrm{PV}_k), \quad P_k = K_c e_k, \quad D_k = -K_d \gamma \dot{\mathrm{PV}}_k \\ S_k &= S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ v_k &= b_k + S_k + I_{k-1}, \qquad u_k = \mathrm{clamp}(v_k, u_{min}, u_{max}) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - v_k}{\tau_I}\right]\Delta t_c \end{aligned} \tag{5.11}$$

where $\gamma = \pm 1$ is the loop's direction, b_k an output bias and $\Delta t_c = 5$ s. Derivative action is switched off in every loop of the cell room ($K_d = 0$), so $S_k = P_k$. The back-calculation corrects the integral alone, with τ_I as the tracking time, and the clamp bounds nothing but the emitted value. No PI output is rate limited, and the only slew in the cell room is the rectifier's ramp of Equation (5.14). A loop in MANUAL recomputes its integral as $I_k = u_{op} - b_k - S_k$ at every scan, so the return to AUTO starts from wherever the operator left the actuator. Table 5.3 lists what each loop measures and moves, its direction and its shipped tuning.

5.4.3 Ratio Stations and Trims

AIC-101 and AIC-102 are ratio stations whose bias is the feedforward, the station with a feedback trim of Equation (1.10), as shown in Equation (5.12).

$$\begin{aligned} F_{b,cmd} &= \mathrm{clamp}\big(r_b I + u_a,\ 0,\ F_{b,max}\big) \\ F_{w,cmd} &= \mathrm{clamp}\big(r_w I + u_w,\ 0,\ F_{w,max}\big) \end{aligned} \tag{5.12}$$

where I is the stack current in kA and u_a, u_w are the PI trims of Equation (5.11) on AT-101 and AT-102. The ratio moves the flow the instant the rectifier does, and the integral carries the correction. Each loop has three strategies. Ratio plus trim is Equation (5.12), ratio only sends $r_b I$ with the trim tracking, and trim only drops the bias so that feedback sets the whole flow. A switch of strategy moves

the integral by exactly the feedforward that appears or disappears, which keeps the flow continuous. While its analyzer is out of service or STALE, a trim holds what it has learned and the ratio carries on alone. More brine raises the anolyte, so AIC-101 is direct acting, and more water dilutes the caustic, so AIC-102 is reverse acting. TIC-103 is one PI over a split range, the steam heater up to 1.5 MW and the cooler down to $-Q_{cool}(T_{cw}) = -1.5(78 - T_{cw})/53$ MW, whose capacity shrinks as the cooling water warms.

5.4.4 The Gas Side and the Interlock

PIC-201 holds the chlorine header with the compressor's suction valve. PDIC-203 is the master of a cascade, and its output $u_{\Delta P}$ is the difference it asks the hydrogen slave for, as shown in Equation (5.13).

$$P_{H,SP} = \mathrm{clamp}(B + u_{\Delta P},\, 0,\, 450), \quad B = \begin{cases} \text{PT-201} & \text{FF on} \\ P_{Cl,SP} & \text{FF off} \end{cases} \tag{5.13}$$

in mbar(g), with $u_{\Delta P}$ between 0 and 150 mbar and FF the chlorine pressure feedforward. With the feedforward on, the hydrogen setpoint follows the chlorine header as PT-201 measures it, and the master only trims the difference. With it off, the setpoint is built on the chlorine setpoint and every rise of the chlorine header has to appear as a PDIC-203 error first. With PIC-202 on its local setpoint or in MANUAL the cascade is open and the master tracks. The interlock of Section 5.3 sits behind the cascade on its own transmitter, so a controller that is too slow loses the day's production but not the membrane.

5.4.5 Rectifier

REC-100 takes the scheduler's plan when FY-400 is on and the plant is in AUTO, and the operator's setpoint otherwise, as shown in Equation (5.14).

$$\begin{aligned} i_{tgt} &= \mathrm{clamp}(i_{want},\, i_{min},\, i_{max}) \\ i_{tgt} &\leftarrow \max(i_{min},\, i_{sh}) \quad \text{if } L \geq 0.97 L_{cap} \text{ and } i_{tgt} > i_{sh} \\ i_k &= \mathrm{slew}(i_{tgt},\, i_{k-1},\, \rho,\, \Delta t_c) \end{aligned} \tag{5.14}$$

where the ramp limit is $\rho = 0.03\,\mathrm{kA/m^2}$ per minute, the minimum load is $i_{min} = 3.0\,\mathrm{kA/m^2}$, and $i_{sh} = (D/c + j_d)/0.95$ is about the current density whose production matches the shipments, with $c = NA\,M_{NaOH}/F = 0.6044\,\mathrm{t/h}$ per kA/m^2 at full efficiency. After a trip the current is zero until the interlock is reset. The I key requests an interlock reset; R keeps its shared simulation Reset action.

5.4.6 The Load Scheduler

FY-400 plans the current density v_k at every hour boundary from now to midnight with a linear program, and between boundaries the rectifier ramps linearly, so a plan that respects the ramp row is executed exactly. The day's electricity cost is minimized subject to the ramp, the load limits, the tank band and the target, as shown in Equation (5.15).

$$\begin{aligned} \min\ & \textstyle\sum_{k=1}^{K} \pi_k \tfrac{1}{2}\big(\hat{P}(v_{k-1}) + \hat{P}(v_k)\big)\Delta t_k \\ & \quad + M_T\,\sigma + M_L(\zeta_{lo} + \zeta_{hi}) \\ \text{s.t. } & |v_k - v_{k-1}| \leq \rho\,\Delta t_k, \qquad i_{min} \leq v_k \leq i_{max} \\ & L_{lo} + \delta_L \leq \hat{L}_k + \zeta_{lo} - \zeta_{hi} \leq L_{hi} - \delta_L \\ & \textstyle\sum_k \hat{m}_k + \sigma \geq m_{day} + \delta_m - m_{made} \\ & \hat{m}_k = c\,\big(\hat{\eta}\,\tfrac{1}{2}(v_{k-1} + v_k) - j_d\big)\Delta t_k \\ & \hat{L}_k = L_0 + \textstyle\sum_{j \leq k}(\hat{m}_j - D\,\Delta t_j) \end{aligned} \tag{5.15}$$

where v_0 is the present current, π_k the price of hour k, $m_{day} = 72\,\mathrm{t}$ the daily target, m_{made} the NaOH made since midnight as the tank gauge and the shipments' meter measure it, and L_0 the tank gauge.

The power $\hat{P}(v) = NAv\hat{U}(v)$ uses the voltage correlation with a new membrane at the measured temperature and strengths, shifted by a bias so that it passes through the measured stack voltage, so a damaged membrane raises the price of every tonne the scheduler plans. $\hat{P}$ is convex, and it is replaced by its chords over four equal segments between i_{min} and i_{max}, whose slopes rise from 1.312 to 1.613 MW per kA/m^2, so the LP fills the cheap segments first without integer variables and overstates the power by at most 12.7 kW. The efficiency $\hat{\eta}$ is the model's η_{max} at the measured strengths with no hardness loss, and the back-diffusion term enters exactly, because j_d/i times the current is a constant. The margins $\delta_L = 0.5$ t and $\delta_m = 0.3$ t absorb the model's error and the gauge's noise, and the slacks σ, ζ_{lo} and ζ_{hi} carry prices of $M_T = \$10^5$ and $M_L = \$10^6$ per tonne, so the LP always returns a plan and reports a shortfall rather than relaxing a limit. With four segment variables per hour the full day has 99 variables and 49 rows. A bounded-variable simplex solves it, pricing by the largest reduced cost and falling back to Bland's rule [17] after 30 pivots without progress, in 172 pivots and bound flips and about 1 ms in Node.js. The plan is re-solved on the hour, when the present hour's price moves by more than 20 %, at a new day, after a restart and whenever a limit or the target changes. The day-ahead price is known and a real-time spike is seen only once it has begun.

5.4.7 Cell-Voltage Monitor

EI-500 subtracts the stack mean from each cell's scan, and a cell whose deviation exceeds ±40 mV for two consecutive scans raises a cell fault. The fingerprint is a fixed, zero-mean draw of 4 mV standard deviation, the same in every case, so the monitor must find a failing cell against a stack that is never uniform.

5.4.8 Three Models Kept Apart

Section 5.3 describes the simulated plant. The application fits no control model, and manual steps from the design point with the instrument noise switched off show instead the plants the loops act on. A 1 m^3/h step of brine raises the anolyte by 2.97 g/L, 63 % of the way in 36 min. A 0.1 m^3/h step of water lowers the caustic by 0.28 %, 63 % in 108 min. A 0.1 MW step of the heater warms the cells by 0.82 K, 63 % in 22 min. A one-percent closure of either header valve raises its header by 7.6 (hydrogen) or 9.8 mbar (chlorine), 63 % in 289 s or 469 s, because a root-law valve makes each header self-regulating. Linearized, a header at P above a downstream pressure P_o (P_d or P_s) rises by $2(P - P_o)/x$ per percent of closure, with x in percent, and responds with the time constant $2(P - P_o)/r$, with r the fill rate of Equation (5.16), which gives 7.4 and 9.5 mbar and 276 and 450 s for the two headers. The controllers' internal models are the ratio coefficients, which approximate the stoichiometry in m^3/h per kA, the unit-gain chlorine feedforward of Equation (5.13), and the scheduler's voltage correlation, efficiency, chords and tank integrator of Equation (5.15). The PI trims carry no model at all.

5.5 Riding Through a Compressor Trip

When the chlorine compressor stops, its suction valve can draw nothing and the chlorine header fills with everything the anodes make. Both header rates follow from Equations (5.1) and (5.9), as shown in Equation (5.16).

$$r_{Cl} = \frac{(1+\eta_{Cl})\,NIRT_g}{4FV_{Cl}}, \qquad r_H = \frac{NIRT_g}{2FV_H}$$
$$\frac{r_H}{r_{Cl}} = \frac{2V_{Cl}}{(1+\eta_{Cl})V_H} \tag{5.16}$$

where r_{Cl} is the chlorine header's rise with the compressor stopped and r_H the fastest the hydrogen header can rise, with its valve shut. At 6 kA/m^2 they are 1.110 mbar/s and 1.340 mbar/s. Both scale with the current, so their ratio, 1.207, is the same at every load, and the hydrogen can outrun the chlorine by only 21 % however it is controlled. A hydrogen header larger than $2V_{Cl}/(1+\eta_{Cl}) = 301.8$ m^3 could not keep pace at all. The engine agrees. With the feedforward on, a 300 m^3 hydrogen

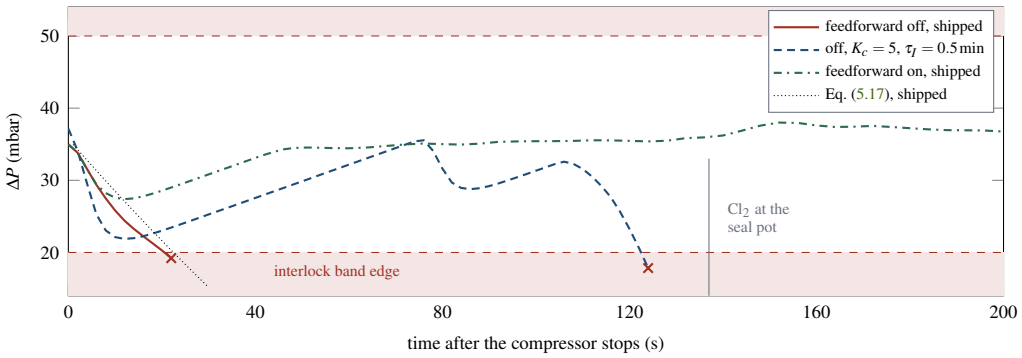

Figure 5.2 Hydrogen-over-chlorine difference after the chlorine compressor stops in case 5, true values from the engine. The shipped tuning without the feedforward (solid) trips at 22 s, ten times its gain with half its integral time ($K_c = 5$, $\tau_I = 0.5$ min, dashed) holds until the master's output limit and trips at 124 s, and the shipped tuning with the feedforward (dash-dotted) never leaves 27.4 to 38.9 mbar. The dotted curve is Equation (5.17) for the shipped tuning, the shaded bands lie outside the interlock band, and crosses mark the trips.

header rides through case 5 with ΔP no lower than 26.1 mbar, while 320 m^3 trips the rectifier 108 s after the compressor stops with the hydrogen valve fully shut, and 400 and 500 m^3 trip after 32 and 24 s.

With the feedforward off and a slave fast enough to hold $P_H = P_{Cl,SP} + u_{\Delta P}$, a chlorine header that starts to rise at r_{Cl} from equilibrium gives the master an error $e = \Delta P_{SP} - \Delta P = r_{Cl}t - (u_{\Delta P} - u_0)$, with $u_{\Delta P} - u_0 = K_c e + (K_c/\tau_I)\int e\,dt$. Differentiating gives Equation (5.17).

$$(1 + K_c)\frac{de}{dt} + \frac{K_c}{\tau_I}e = r_{Cl}$$
$$e(t) = \frac{r_{Cl}\tau_I}{K_c}\left(1 - e^{-t/\tau_e}\right), \qquad \tau_e = \frac{(1 + K_c)\tau_I}{K_c} \tag{5.17}$$

The error settles toward $r_{Cl}\tau_I/K_c$, and the margin $m = \Delta P_{SP} - \Delta P_{lo} = 15$ mbar is gone after the time given by Equation (5.18).

$$t_m = -\tau_e \ln\left(1 - \frac{K_c m}{r_{Cl}\tau_I}\right), \qquad K_c m < r_{Cl}\tau_I \tag{5.18}$$

At the shipped tuning $\tau_e = 180$ s and the settled error would be 133.2 mbar, nearly nine times the margin, so $t_m = 21.5$ s. On the engine the header starts to rise 1 s after the scheduled event, once the integration step in which the compressor stops is complete, the true difference crosses 20 mbar 20 s later, and the interlock opens the rectifier 22 s after the compressor stopped, as shown in Figure 5.2. The estimate assumes a perfect slave and a continuous master, so it overstates the ride-through, by 1.5 s at 6 kA/m^2. On the challenge's day, where the trip comes at 5.3 or 7.0 kA/m^2, it predicts 24.7 and 18.3 s against 23 and 15 s on the engine, whose interlock trips 25 and 17 s after the compressor stops.

Equation (5.18) says the loop holds on feedback alone only if $K_c/\tau_I > r_{Cl}/m$, 4.44 min^{-1} at 6 kA/m^2. A second limit then binds. The master may ask for at most 150 mbar above the chlorine setpoint, so the hydrogen setpoint stops at 300 mbar(g) while the chlorine climbs toward its seal pot at the same pressure, and the difference falls through 20 mbar once the chlorine passes about 280 mbar(g), 117 s after the rise begins. The two runs with $K_c/\tau_I = 10\,min^{-1}$ cross at 122 and 123 s and trip at 124 and 125 s, their hydrogen header coasting 4 to 5 mbar above the capped setpoint, as Table 5.4 shows. Faster feedback also costs noise. In normal operation the standard deviation of

Table 5.4 PDIC-203 Tuning Against the Compressor Trip of Case 5

			Feedforward off			FF on
K_c	τ_I (min)	t_m, Eq. (5.18) (s)	crossing (s)	trip (s)	σ (mbar)	lowest ΔP
0.5	1	21.5	20	22	0.09	27.4
1	1	30.6	28	30	0.12	28.3
2	1	53.9	48	50	0.29	28.5
2	0.5	104.0	93	95	0.35	28.4
5	0.5	holds	122	124	2.81	22.5
10	1	holds	123	125	2.82	24.0
10	0.1	holds	11	13	6.99	trips

Crossing: true ΔP below 20 mbar, seconds after the chlorine header starts to rise. Trip: seconds after the compressor stops. σ: standard deviation of the true ΔP from 20 to 60 min, before the trip. Lowest: the least ΔP from 1:00 to 2:00 with the feedforward (FF) on, where the last tuning trips at 13 s, inside its limit cycle.

ΔP grows from 0.09 mbar at the shipped tuning to 2.8 mbar at 10 min^{-1}, and at the extreme of the tuning range, $K_c = 10$ and $\tau_I = 0.1$ min, the cascade limit-cycles with a 40 s period between 24.9 and 47.9 mbar, close to both trip limits, with the hydrogen valve commanded to one of its stops 97 % of the time. The compressor trip then catches a swing near its bottom, and that tuning trips the rectifier 13 s after the compressor stops with the feedforward on or off.

> **DID YOU KNOW? Faster Feedback Trips Sooner**
>
> From the simulator: the fastest pressure-difference tuning of Table 5.4, a gain of 10 with an integral time of 0.1 min, trips the rectifier 13 s after the compressor stops, sooner than the 22 s of the shipped tuning without the feedforward, and with the feedforward on or off. Before the trip it already swings between 24.9 and 47.9 mbar.

The feedforward removes the ramp from the master's error, because the hydrogen setpoint now moves with PT-201 from the first second. What is left is the slave's transient and the master's 5 s cycle. The difference dips to 27.4 mbar 12 s after the trip, the chlorine reaches its seal pot 137 s after it, 1.48 t of chlorine go to the emergency absorption before the restarted compressor has reloaded, and the difference stays between 27.4 and 38.9 mbar through the restart. A feedback loop acts on an error it has already seen, while the feedforward acts on the cause, which is measured, and that is the whole lesson of the gas side.

Pause. Try Exercise 5.2 before reading on.

5.6 Measured Cases

The nine scenarios are summarized in Table 5.5, and the Exercises at the end of the chapter work through them. Every case lasts 24 h and starts from the shipped tunings of Table 5.3, with the chlorine feedforward on, both feeds at ratio plus trim and the scheduler off, except for the items its row names, which are set before the run starts. Each number was read from the engine, stepped without the browser through the same step function at the default seed, and a run that trips is continued to midnight. Table 5.6 gathers each day's scorecard, which is kept on the truth at every physics step.

5.6.1 The Current Is the Load

Case 1 starts steady at 5.0 kA/m^2 with 29.51 $\text{m}^3\text{/h}$ of brine and 3.222 $\text{m}^3\text{/h}$ of water set by hand. At 1:00 the rectifier ramps to 6.5 kA/m^2 and reaches it at 1:50, so the membranes carry 5.46 t/h of salt out of the anolyte against 4.19 t/h before, and the brine does not follow. The anolyte passes 200 g/L at 1:43 and 190 g/L at 2:09 and reaches 177.9 g/L by 8:00. The caustic leaves its specification at 1:52 and rises from 32.00 to 33.11 %, less than the current alone would make it because the weaker anolyte

Table 5.5 The Cases, Their Upsets and Their Lessons

#	Scenario and events	Teaching point
1	feeds in MANUAL, 5.0 to 6.5 kA/m^2 at 1:00	the current is the load
2	6.0 to 4.0 at 2:00, to 6.5 kA/m^2 at 12:00	a ratio and its trim
3	caustic SP 33, 34, 35 % at 2:00, 8:00, 14:00	an optimum on a curve
4	ratio only, brine 305 to 285 g/L at 1:00	depletion, oxygen
5	compressor trip 1:00 to 1:30, feedforward off	feedforward, interlock
6	500 ppb hardness 2:00 to 5:00	damage no loop undoes
7	brine 65 to 45 °C at 1:00, heater held	the heat balance
8	FY-400 on, 30 t tank from 24 t	load following on price
9	graded day, three upsets	challenge

Table 5.6 Truth-Log Scorecard of Each Case and Variant Over Its 24 h Day

Case and configuration	Made (t)	On spec (t)	Off spec (t)	Electricity ($)	*SE* (kWh/t)	ΔP (mbar)	Anolyte lowest (g/L)	η lowest (%)	Other
1 feeds in MANUAL (shipped)	86.8	5.6	81.0	13 724	2300	34.5–35.6	177.9	93.22	O_2 1.78 %
1 both feed loops in AUTO	88.4	88.4	0.0	13 508	2226	34.5–35.6	209.7	94.90	
2 ratio plus trim (shipped)	74.2	74.2	0.0	11 318	2157	34.5–35.6	209.7	94.75	caustic 31.98–32.02 %
2 trim only	74.2	74.2	0.0	11 319	2158	34.5–35.6	205.0	94.75	caustic 31.64–32.40 %
2 ratio only	74.2	9.4	65.0	11 278	2150	34.5–35.6	208.0	94.75	caustic 31.36–32.00 %
3 caustic 33 to 35 %[a]	81.0	7.5	72.9	12 557	2265	34.5–35.6	209.7	91.24	
4 ratio only (shipped)	82.0	82.0	0.0	12 447	2220	34.5–35.6	185.8	94.15	O_2 1.36 %
4 ratio plus trim	82.7	82.7	0.0	12 373	2189	34.5–35.6	205.0	95.00	
5 feedforward off (shipped)	3.5	3.6	0.0	394	2189	19.8–35.3	210.0	95.00	trip, dry 772 min
5 feedforward on	82.7	82.7	0.0	12 373	2189	27.4–38.9	209.6	95.00	1.48 t Cl_2 absorbed
6 breakthrough (shipped)	81.4	81.4	0.0	12 636	2268	34.5–35.6	209.7	93.21	$534/day lasting
6 with cell 87 failing at 12:00	81.3	81.3	0.0	12 633	2270	34.5–35.6	209.7	93.01	flagged 12:22
7 heater held (shipped)	82.7	82.7	0.0	12 684	2243	34.5–35.6	209.5	95.00	80.7 °C
7 TIC-103 in AUTO	82.7	82.7	0.0	12 374	2189	34.5–35.6	209.7	95.00	
8 scheduler (shipped)	72.3	72.3	0.0	9 008	2163	34.5–35.6	209.7	94.50	tank 21.4–26.7 t
8 constant 5.3 kA/m^2	73.0	73.0	0.0	10 617	2127	34.5–35.6	209.7	94.93	tank 24.0–25.0 t

All cases are priced on the sunny profile. Made is NaOH at the membrane, on and off specification are what reached the tank inside and outside 32 ± 0.5 %, and after the trip of case 5 the catholyte above its weir still drains to the tank, so more reaches it than the membranes made. [a] Case 3 moves the caustic off the 32 ± 0.5 % specification on purpose, to trace the efficiency curve.

pushes more water through the membrane. The efficiency falls from 94.90 to 93.22 %, the oxygen in the chlorine rises from 0.60 to 1.78 %, the specific energy climbs from 2101 to 2314 kWh/t, and 81.0 t of the day's 86.8 t reach the tank off specification. With both feed loops in AUTO the anolyte stays within 0.28 g/L and the caustic within 0.016 % of their setpoints all day.

Case 2 turns the load down to 4.0 kA/m^2 at 2:00 and up to 6.5 kA/m^2 at 12:00, which asks the brine for 35.46, 23.57 and 38.43 m^3/h in turn. With ratio plus trim the anolyte stays within 0.28 g/L of 210 and the caustic within 0.022 % of 32. With the trims alone the flows move only after the analyzers have seen the strengths move, and the anolyte swings 4.88 g/L after the turndown and the caustic 0.40 % after the load-up, inside the specification only because the 0.03 kA/m^2 per minute ramp stretches each step over more than an hour. With the ratios alone the flows follow at once but keep the stations' own errors, the anolyte settles 2.0 g/L low and the caustic 0.64 % weak, off specification from 2:47, and 65.0 t of 74.2 t go off specification. Feedforward does the known part at once, and feedback does the part nobody knew.

🕮 **Pause.** Try Exercise 5.1 before reading on.

5.6.2 Two Curves in the Efficiency

Case 3 raises the caustic setpoint one point at a time. Averaged over the hour before each step, the efficiency is 95.00, 94.60, 93.40 and 91.40 % at 32, 33, 34 and 35 %, losses of 0.40, 1.20 and 2.00 points, so the third point costs five times the first. The cell voltage rises 64 mV over the three steps and the specific energy climbs from 2188 to 2213, 2257 and 2322 kWh/t, 6.1 % in all, for 1.00 m^3/h

less water to evaporate downstream (the water added falls from 3.87 to 3.52, 3.18 and 2.87 m^3/h). In case 4 the brine weakens from 305 to 285 g/L under the ratio alone. The anolyte, at 208.4 g/L from the station's own error, falls through 190 g/L at 2:00 and reaches 185.8 g/L by 8:00, the dry chlorine carries 1.36 instead of 0.64 % oxygen, and below the 200 g/L knee the efficiency falls to 94.15 %, while the caustic trim holds 32 % by cutting the water from 3.86 to 3.48 m^3/h. With the densitometer trim the brine rises from 35.5 to 44.9 m^3/h, the anolyte dips only to 205.0 g/L and is back within 1 g/L of 210 by 2:26. A ratio is a model of the feed's strength, and only a measurement can tell it that the strength has changed.

> **TRY IT IN THE APP** **4 · A Weak Brine and Oxygen in the Chlorine**
>
> Load scenario 4 and press Run. At 1:00 the brine weakens from 305 to 285 g/L while its feed follows the current by the ratio alone. Watch the anolyte strip on the Cell view fall through 190 g/L at 2:00 to 185.8 g/L by 8:00. Reset, set the Brine feed strategy to Ratio + trim, and run again: the anolyte dips only to 205.0 g/L.

5.6.3 Damage No Loop Undoes

In case 6, 500 ppb of hardness reaches the cells from 2:00 until a fresh ion-exchange column comes on line at 5:00. The anolyte passes the 30 ppb threshold at 2:01, peaks at 636 ppb at 5:00 and falls back below it at 7:31, an exposure of 1794 ppb h. The cell voltage rises from 3.103 to 3.178 V, 75.4 mV, and the efficiency falls from 95.00 to 93.21 %, so the specific energy is 2285 kWh/t instead of 2188 while the anolyte, the caustic and the temperature are back on their setpoints. At the profile's average price of $68.38/MWh the scorecard prices the damage at $534 a day for as long as this membrane stays in. The monitor stays silent, because hardness raises every cell alike, and the worst of the 150 cells never reads more than 17.2 mV from the mean, against a threshold of 40 mV. In the variant with cell 87 failing from 12:00, the monitor flags cell 87 after 22.1 min, 42 mV below the mean, while the stack mean has moved by 0.3 mV. A threshold of 20 mV flags it after 11.1 min and one of 100 mV after 50.1 min, while 10 mV raises a false alarm on a healthy cell 35 s into the day, at its second scan, inside the stack's own spread.

Pause. Try Exercise 5.3 before reading on.

5.6.4 Cold Brine and the Heat Balance

In case 7 the brine arrives at 45 °C instead of 65 from 1:00 with the heater held at 0.09 MW. The cells cool 63 % of the way in 24.5 min and reach 80.7 °C, the voltage rises by 82 mV to 3.185 V by 8:00, and the specific energy goes from 2188 to 2246 kWh/t, 0.199 MW or 2.6 % more electricity for the same caustic. With TIC-103 in AUTO the heater rises from 0.09 to 0.87 MW, the temperature dips by 1.70 K and is back within 0.3 K of 88 °C by 1:58. The 0.78 MW of steam buys back 0.199 MW of electricity, a trade the price decides, while keeping the membrane at its design temperature keeps it where its correlations were fitted.

5.6.5 Load Following on the Price

Case 8 gives the scheduler a sunny day's price from $18 to $165/MWh, a 30 t tank starting at 24 t, customers drawing 3.0 t/h and a 72 t target. Its first plan, solved at midnight, reaches the rectifier limit at 3:00, from 12:00 to 15:00 and from 23:00, and the minimum load through both price peaks. The rectifier runs within 0.01 kA/m^2 of 7.0 kA/m^2 from 2:59 to 3:06, from 11:59 to 15:01 and after 22:59, and of 3.0 kA/m^2 from 6:58 to 8:09 and from 18:00 to 20:00. What stops it running harder in the cheap hours is the tank, which reaches the scheduler's ceiling of 26.5 t, 0.5 t inside the tank's own 90 %, by 5:00 and again at 16:00, and never leaves 21.4 to 26.7 t. No hourly re-solve plans to pass that ceiling by more than 0.01 t, and the tank still passes it by up to 0.16 t from 16:05 to 17:11, an error of the scheduler's model that its margin absorbs. The day makes 72.3 t for $9 008 at 2163 kWh/t, buying its 156.3 MWh at an average of $57.63/MWh. A constant 5.3 kA/m^2, which starts the day steady at that current, makes 73.0 t for $10 617, 18 % more, although each of its tonnes takes 1.6 % less energy

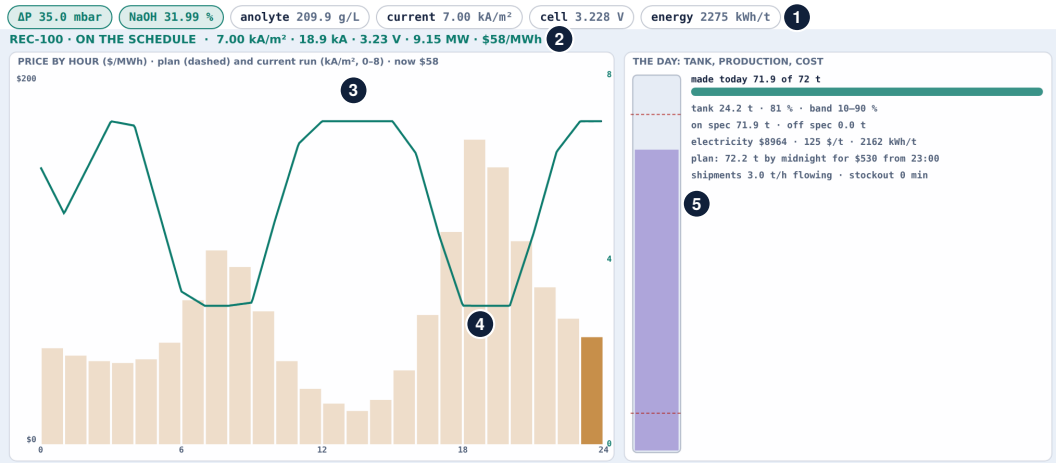

Figure 5.3 The electrolyzer in case 8, load following on the price, at $t = 23.92$ h (1435 min) in the Plant view, with every control as shipped, cropped to the plant view. Numbered callouts mark (1) the badges, (2) the rectifier's status line, (3) the current at the rectifier's limit through the midday dip in price, (4) the minimum load through the evening peak and (5) the tank and the day's tally. The price panel shows the sunny day's hourly price with the present hour at $58/MWh highlighted, and on its right-hand scale the current density actually run, 5.0 kA/m^2 at 1:00, close to the rectifier's 7.0 kA/m^2 at 3:00 and at it from 12:00 to 15:00 and from 23:00, and the 3.0 kA/m^2 minimum load through both price peaks. The plan for the last hour (dashed) lies under the run at 7.0 kA/m^2. The day panel gives 71.9 of the 72 t target made, the tank at 24.2 t (81 %) inside its 10 to 90 % band, 71.9 t on specification and none off, $8964 of electricity at $125/t and 2162 kWh/t, a plan of 72.2 t by midnight, and shipments of 3.0 t/h with no stockout. The header reads REC-100 on the schedule at 7.00 kA/m^2, 18.9 kA, 3.23 V and 9.15 MW, and the badges read 35.0 mbar, 31.99 % NaOH, 209.9 g/L of anolyte and 2275 kWh/t.

(2127 kWh/t), because it pays the profile's average of $68.38/MWh. On the scheduler's own model the marginal energy of one more tonne is 2185 kWh/t at the minimum load and 2881 kWh/t at the rectifier limit, a ratio of 1.32 against a price ratio of 9.2 between the dearest and cheapest hours, which is why the plan sits on its limits. The flexibility is bought with storage. A 60 t tank starting at 24 t, set before the day is restarted, saves another $220 and a 15 t tank starting at 12 t costs $158 more, while a ramp limit of 0.1 kA/m^2 per minute saves $40 and a minimum load of 2.0 kA/m^2 saves $164. A two-hour spike of $250/MWh at 17:00 adds $1 900 to the scheduled day and $3 235 to the constant one. The last minutes of the day are shown in Figure 5.3.

> **TRY IT IN THE APP** **8 · Load Following on the Price**
>
> Load scenario 8 and press Run on the Plant view. Watch the current against the price bars: it runs at the rectifier's 7.0 kA/m^2 from 12:00 to 15:00 and at the 3.0 kA/m^2 minimum through both price peaks. The tank reaches the scheduler's ceiling of 26.5 t by 5:00, which is what stops the plan running harder in the cheap hours.

5.6.6 The Graded Day

The challenge runs case 8's day with three upsets, brine at 290 g/L from 6:00, a compressor trip at 14:00 for 15 min, and a $250/MWh spike at 17:00 for two hours. The target asks for at most $11 500 of electricity, at least 71.5 t of NaOH at 32 ± 0.5 %, an anolyte never below 190 g/L, an efficiency never below 94 %, ΔP always inside the band, no trip and no stockout. The cost ceiling

Table 5.7 Challenge Configurations Over the 24 h Day

Configuration	On spec (t)	Off (t)	Cost ($)	ΔP (mbar)	Fails on
shipped[a]	9.8	33.0	4 880	19.9–35.4	trip, ΔP, short, dry
suggested[b]	72.3	0.0	10 908	26.4–39.0	passes
constant, trims, FF	73.0	0.0	13 853	28.5–38.7	cost
sched., FF, ratios	10.8	61.4	10 868	26.5–38.9	short, η
sched., trims, no FF	42.8	0.0	4 288	19.98–35.4	trip, ΔP, short, dry

[a] constant 5.3 kA/m^2, ratio-only feeds, feedforward (FF) off.
[b] ratio plus trim on both feeds, scheduler on, feedforward on.
Sched.: scheduler. Short: less than 71.5 t on specification.

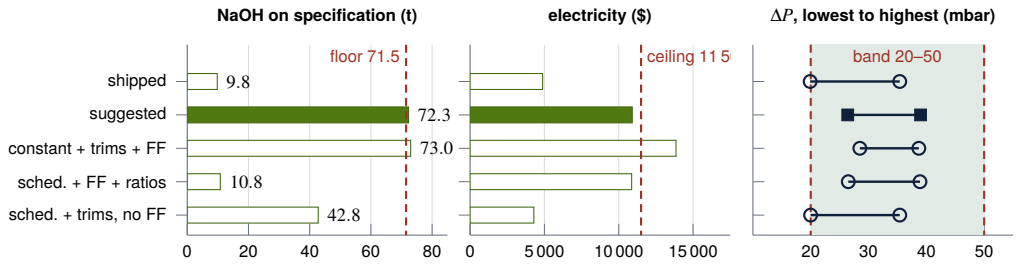

Figure 5.4 Table 5.7 as bars and ranges, digit for digit: the NaOH on specification against the 71.5 t floor, the electricity cost against the $11 500 ceiling, and the lowest to highest true ΔP against the 20 to 50 mbar band. The rows are the table's, and the filled bar and squares mark the one configuration that passes. Notice that the two cheapest days are the two that trip, and that the constant current fails on cost, the ratio-only feeds on tonnage, and the day without the feedforward on ΔP.

is the suggested configuration's cost plus 5.4 %, and the tonnage floor its tonnage less 1.0 %, so the suggested configuration passes with margins of $592 and 0.76 t. Changing a model parameter or a disturbance voids an attempt, so a larger tank, which case 8 prices, cannot be bought here. Table 5.7 compares the measured configurations, and Figure 5.4 draws three of its columns against the targets.

As shipped the caustic leaves its specification at 3:12 and falls to 30.65 %, the anolyte falls to 191.3 g/L, the interlock opens the rectifier 25 s after the compressor stops, and the tank runs dry for 107 min from 22:13, so of 42.5 t made only 9.8 t reach the tank on specification. The suggested configuration holds every limit, with 72.3 t on specification for $10 908, the anolyte never below 207.7 g/L, the efficiency never below 94.50 %, ΔP between 26.4 and 39.0 mbar through a compressor trip at the full 7.0 kA/m^2, and the tank between 21.5 and 26.7 t. Leaving out any one element fails a different line of the scorecard. A constant current with the trims and the feedforward is safe and on specification but costs $13 853. The scheduler and the feedforward on ratio-only feeds cost $10 868 but put 61.4 t off specification with the caustic down to 30.60 % and the efficiency down to 93.998 %, just under its floor. The trims and the scheduler without the feedforward trip 17 s after the compressor stops and run the tank dry for 109 min.

5.7 Technology in Practice

What Measures It

Figure 5.5 draws the loops of Section 5.4 with their devices. The two strengths are read on the liquids leaving the cells, the anolyte by a densitometer and the caustic by conductivity. An electrodeless conductivity cell induces a current in the liquid with one toroidal coil and measures it with a second, a voltage in proportion to the conductivity, and such cells are preferred wherever they can be used for

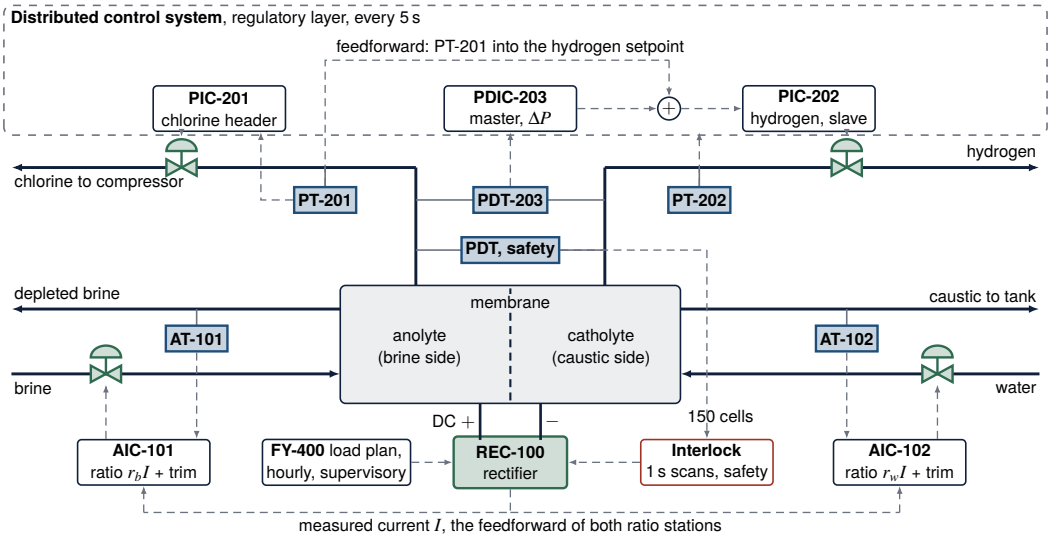

Figure 5.5 The loops of the electrolyzer with their hardware. Lettered boxes are sensing devices: the densitometer (AT-101) on the depleted brine, the conductivity analyzer (AT-102) on the caustic, the two header transmitters (PT), the difference transmitter of the loop (PDT-203) and the interlock's own transmitter (PDT, safety). Valve symbols and the rectifier are final elements, dashed arrows are signals. Notice that the measured current enters both ratio stations before any analyzer has seen a change, that PT-201 enters the hydrogen setpoint in the same way, and that the interlock reaches the rectifier by a path of its own. The temperature loop, the tank gauge and the cell-voltage monitor are left out. Not to scale, and not a piping diagram.

their ruggedness and near immunity to fouling [18]. The gas side, the lesson of Section 5.5, is read by difference transmitters, one for the loop and one for the interlock, of the kind shown in Figure 5.6.

> **TECHNOLOGY** **Sensor: A Differential Pressure Transmitter**
>
> A differential pressure transmitter has two ports, high and low, and a diaphragm between them that flexes with any difference of pressure. It ignores the pressure common to both ports and sends a standard signal, such as 4 to 20 mA or a digital code, for the difference alone. Where the process fluid is corrosive, isolating diaphragms and a filled capillary, a chemical seal, keep it out of the instrument, and on a gas service a bleed fitting lets condensed liquid out. An impulse line that plugs makes the transmitter respond too slowly, or not at all, so its trend turns sluggish or flat while the process moves [18]. The model's PDT-203 carries 0.1 mbar of noise and the interlock's own transmitter 0.2 mbar, and neither ever plugs.

What Moves It

The rectifier moves everything at once: the chapter's REC-100 sets the current, and Faraday's law turns it into chlorine, hydrogen and caustic. The design point matches a typical bipolar membrane plant of a 2014 European reference document, 2.70 m^2 membranes at 6.0 kA/m^2 and 16.2 kA per electrolyzer [19]. The brine and water flows follow their ratio stations as first-order lags of 20 s, and the two header valves follow their controllers as lags of 3 s.

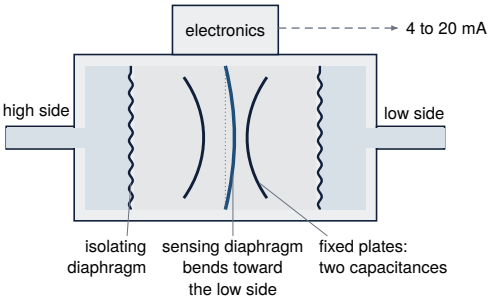

Figure 5.6 A differential pressure transmitter, the device behind PDT-203 and the interlock's transmitter. Notice that the sensing diaphragm responds to the difference between its two sides and ignores the pressure common to both [18].

> **TECHNOLOGY** **Final Element: The Rectifier**
>
> A cell room needs huge quantities of direct current, usually taken from a high-voltage alternating supply. Step-down transformers lower the voltage, silicon diode or thyristor rectifiers convert it to direct current, and busbars carry it to the cells. In 2010 a rectifier and transformer unit had an efficiency between about 94 percent (older units) and 98 percent, and the units are cooled by air or by water circuits [19]. A controlled rectifier sets its output by timing the trigger pulses of its silicon-controlled rectifiers against the alternating cycle: the earlier each one fires, the longer it passes current [18]. The model's REC-100 knows its current exactly, ramps at 0.03 kA/m^2 per minute, and after a trip stays at zero until the interlock is reset.

Where the Controller Runs

A plant runs such loops in a distributed control system, in which a supervisory level sends set points to local controllers, loops are sometimes nested, and cycle times range from milliseconds to minutes [20]. The app executes its six PI loops every 5 s, inside that range. The ratio stations and the pressure cascade with its feedforward belong to the regulatory layer, the load scheduler is supervisory and writes the rectifier's target at every hour, and the interlock scans its own transmitter once a second, apart from both.

> **IN PRACTICE** **A Trip With Its Own Transmitter**
>
> When a mercury cell room is converted to membranes, the main concern is gas pressure control: the membrane needs a steady differential pressure of hydrogen over chlorine, so a differential pressure controller is added. Behind it a chlorine absorption unit is sized to take the full cell-room production in an emergency [19]. A safety instrumented system reads its own transmitters and decides in its own logic solver, apart from the regulatory controls, and a plant built for automatic shutdown will probably shut down more often than it needs to [18]. IEC 61511-1 sets its requirements from specification to maintenance [21]. The chapter's interlock reads its own transmitter and waits for two scans.

What Is Modernizing It

Three developments bear on this cell room. Oxygen-depolarized cathodes reduce oxygen to hydroxide instead of making hydrogen, which lowers the cell voltage by about 1 V at 4 kA/m^2 and the electricity by about 30 %, and the first industrial installation went into operation in 2011 [19]. Plants already run at varying current densities to follow product demand and the electricity price [19], and a 2023 study used a digital model of a cell room to shift its load on the day-ahead and balancing markets, finding that the flexibility depends largely on intermediate storage or on downstream processes [22], as the tank of case 8 shows. And bipolar electrolyzers make faulty cells easier to find by monitoring every cell voltage [19], the job of EI-500.

5.8 Limitations

One lumped cell stands for all 150, and the voltage, efficiency, water transport, oxygen and hardness correlations are illustrative fits through one design point rather than a vendor's data. The damage law is a threshold and a linear exposure that knows nothing of the individual impurities [6]. The brine circuit stops at the cells, with no acid dosing, chlorate, dechlorination or resaturation, and the gases leave for consumers that never limit the plant.

The analyzers are sampled, delayed and noisy but never biased, and the scheduler knows the day-ahead price exactly. The interlock band of 20 to 50 mbar and its two-scan delay are teaching values. The scheduler plans on hourly boundaries with one lumped tank and a fixed customer draw, where a real plant would also schedule its chlorine consumers [10]. The heater's steam is never weighed against the electricity it saves.

OPEN CHALLENGE Caustic at the Strength Customers Buy

A membrane cell makes caustic of up to about 35 weight percent, a limit set by the water each sodium ion carries through the membrane, close to four molecules, and hydroxide leaking back costs about five percent of the charge [2]. Membranes that would make the commercial 50 percent directly were abandoned around 2000 for technical problems, and their higher cell voltage made them pay only where steam is dear and electricity cheap [19]. Cutting the back-migration of hydroxide to raise the current efficiency remains a target [2]. Progress would be a membrane that keeps its efficiency in stronger caustic, moving the knee that case 3 steps into.

5.9 Conclusion

A membrane chlor-alkali electrolyzer is a plant whose throughput is a current, and this simulator shows what follows from that at each layer. Feeds tied to the current hold the anolyte within 0.28 g/L and the caustic within 0.022 % through two load steps, feeds left in MANUAL send 81.0 t of a day's 86.8 t off specification, and each ratio needs its trim because a station 1.8 % wrong or a brine 20 g/L weak leaves an offset no feedforward can see. On the gas side a chlorine header rising 1.110 mbar/s takes the margin from a feedback-only hydrogen loop in 20 s, against the 21.5 s of the closed-form estimate, and the rectifier trips 22 s after the compressor stops, while the chlorine pressure feedforward holds the difference above 27.4 mbar. The two headers' fastest rise rates stand in a ratio of 1.207 at every current. Hardness costs 75.4 mV and $534 a day that no controller recovers, and the scheduler makes the day's caustic for $9 008 against $10 617 at a constant current, limited mainly by its tank. The model is one lumped cell with illustrative correlations and teaching limits. A fuller model would add stack gradients and individual membranes, the brine's acid dosing and chlorate, a limited chlorine outlet, analyzer bias and an uncertain price. The application needs only a browser, and its lesson is that a measured cause, fed forward, does in time what no feedback can.

KEY POINTS

- In an electrolyzer the current is the production rate, so a feed that is not tied to it drifts whenever the load moves.
- Feeds tied to the current hold both strengths through two load steps, and each ratio still needs its analyzer trim.
- A chlorine header rising 1.110 mbar/s takes the margin from a feedback-only hydrogen loop in 20 s, while the chlorine pressure feedforward holds the difference above 27.4 mbar.
- Hardness costs 75.4 mV and $534 a day that no controller recovers, and the scheduler makes the day's caustic for $9 008 against $10 617 at a constant current, limited mainly by its tank.

Where else this applies. Volume 3, Chapter 4 trims an air-to-fuel ratio on an oxygen analyzer in the same way, and Volume 9, Chapter 11 schedules its load on the electricity price.

References

[1] J. Crook and A. Mousavi, "The chlor-alkali process: A review of history and pollution," *Environmental Forensics*, vol. 17, no. 3, pp. 211–217, 2016.

[2] M. Paidar, V. Fateev, and K. Bouzek, "Membrane electrolysis—history, current status and perspective," *Electrochimica Acta*, vol. 209, pp. 737–756, 2016.

[3] 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.

[4] Q. Zhang and I. E. Grossmann, "Enterprise-wide optimization for industrial demand side management: Fundamentals, advances, and perspectives," *Chemical Engineering Research and Design*, vol. 116, pp. 114–131, 2016.

[5] R. K. B. Karlsson and A. Cornell, "Selectivity between oxygen and chlorine evolution in the chlor-alkali and chlorate processes," *Chemical Reviews*, vol. 116, no. 5, pp. 2982–3028, 2016.

[6] J. T. Keating and K.-J. Behling, "Brine, impurities, and membrane chlor-alkali cell performance," in *Modern Chlor-Alkali Technology*, N. M. Prout and J. S. Moorhouse, Eds. Dordrecht: Springer Netherlands, 1990, pp. 125–139.

[7] A. A. Jalali, F. Mohammadi, and S. N. Ashrafizadeh, "Effects of process conditions on cell voltage, current efficiency and voltage balance of a chlor-alkali membrane cell," *Desalination*, vol. 237, no. 1–3, pp. 126–139, 2009.

[8] J. I. Otashu and M. Baldea, "Grid-level "battery" operation of chemical processes and demand-side participation in short-term electricity markets," *Applied Energy*, vol. 220, pp. 562–575, 2018.

[9] ——, "Demand response-oriented dynamic modeling and operational optimization of membrane-based chlor-alkali plants," *Computers & Chemical Engineering*, vol. 121, pp. 396–408, 2019.

[10] L. C. Brée, K. Perrey, A. Bulan, and A. Mitsos, "Demand side management and operational mode switching in chlorine production," *AIChE Journal*, vol. 65, no. 7, p. e16352, 2019.

[11] J. M. Simkoff and M. Baldea, "Stochastic scheduling and control using data-driven nonlinear dynamic models: Application to demand response operation of a chlor-alkali plant," *Industrial & Engineering Chemistry Research*, vol. 59, no. 21, pp. 10 031–10 042, 2020.

[12] T. F. O'Brien, T. V. Bommaraju, and F. Hine, *Handbook of Chlor-Alkali Technology*. Boston, MA: Springer US, 2005.

[13] I. F. White and T. F. O'Brien, "Secondary brine treatment: Ion-exchange purification of brine," in *Modern Chlor-Alkali Technology*, N. M. Prout and J. S. Moorhouse, Eds. Dordrecht: Springer Netherlands, 1990, pp. 271–289.

[14] I. Moussallem, J. Jörissen, U. Kunz, S. Pinnow, and T. Turek, "Chlor-alkali electrolysis with oxygen depolarized cathodes: history, present status and future prospects," *Journal of Applied Electrochemistry*, vol. 38, no. 9, pp. 1177–1194, 2008.

[15] C. A. Babu and S. Ashok, "Peak load management in electrolytic process industries," *IEEE Transactions on Power Systems*, vol. 23, no. 2, pp. 399–405, 2008.

[16] M. Seko, "The ion-exchange membrane, chlor-alkali process," *Industrial & Engineering Chemistry Product Research and Development*, vol. 15, no. 4, pp. 286–292, 1976.

[17] R. G. Bland, "New finite pivoting rules for the simplex method," *Mathematics of Operations Research*, vol. 2, no. 2, pp. 103–107, 1977.

[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] T. Brinkmann, G. Giner Santonja, F. Schorcht, S. Roudier, and L. Delgado Sancho, "Best available techniques (BAT) reference document for the production of chlor-alkali," European Commission, Joint Research Centre, Luxembourg, JRC Science and Policy Report EUR 26844 EN, 2014.

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

[21] International Electrotechnical Commission, "IEC 61511-1:2016, functional safety – safety instrumented systems for the process industry sector – part 1: Framework, definitions, system, hardware and application programming requirements," Scope page, webstore.iec.ch/en/publication/24241, 2016, accessed 2026-11-18.

[22] P. Lerch, F. Scheller, and T. Bruckner, "Potential of demand response for chlor-alkali electrolysis processes," in *2023 19th International Conference on the European Energy Market (EEM)*, 2023, pp. 1–7.

5.10 Exercises

5.1 Concept check · 2 min

In case 2 both feeds are set to Ratio only before Run, and the load is turned down from 6.0 to 4.0 kA/m^2 at 2:00. The stations send 2.15 m^3/h of brine and 0.251 m^3/h of water per kA, where the design point needs 2.189 and 0.239. What does the anolyte do?

(a) It settles about 2 g/L below its setpoint and stays there, since nothing corrects the station's own error
(b) It returns to 210 g/L after the turndown, since the ratio moves the brine in proportion to the current
(c) It keeps falling while the load stays low, since the anolyte integrates the brine shortfall
(d) It swings about 5 g/L after the turndown and then recovers, as the analyzer catches up

5.2 Concept check · 2 min

In case 5 the chlorine compressor stops and the chlorine header rises about 1.1 mbar/s. With the chlorine pressure feedforward off, ten times the shipped PDIC-203 gain with half its integral time holds the difference for about two minutes and then trips the rectifier, while the shipped tuning with the feedforward on never leaves 27.4 to 38.9 mbar. Why does the feedforward succeed?

(a) It doubles the master's gain while the chlorine header rises, so the hydrogen reacts faster
(b) The hydrogen setpoint follows the measured chlorine pressure, so the rise never becomes an error
(c) It throttles the chlorine suction valve less, so the chlorine header rises more slowly
(d) It widens the interlock band while the compressor is stopped, so a lower difference is allowed

5.3 Concept check · 1 min

In case 6 a hardness breakthrough raises the cell voltage by 75.4 mV for the rest of the membrane's life. Why does the cell-voltage monitor EI-500 raise no alarm?

(a) The rise takes three hours, slower than the monitor's two-scan delay can follow
(b) The monitor watches only cell 87, the one cell the model allows to fail
(c) It compares each cell with the stack mean, and hardness raises every cell alike
(d) The trims return the anolyte and caustic to setpoint, which takes the rise back out

5.4 Theory and calculation · 6 min

The electrolyzer has 150 cells of 2.7 m^2 in series at 6.0 kA/m^2, with $F = 96\,485.332\,\mathrm{C/mol}$, $\eta = 0.950$ and $\eta_{Cl} = 0.988$. Molar masses are 39.997 (NaOH), 70.906 (Cl_2) and 2.016 g/mol (H_2). Give each answer to three significant figures.

(a) Find the moles of electrons through the stack per second and the NaOH made in t/h. *Checkpoint:* 25.2 mol/s and 3.45 t/h.

(b) Find the chlorine in t/h, the hydrogen in kg/h and the oxygen as a percentage by volume of the dry chlorine. *Checkpoint:* 3.18 t/h, 91.4 kg/h and 0.604 %.

(c) The cell voltage is 3.103 V. Find the DC power and the specific energy in kWh per tonne of NaOH. *Checkpoint:* 7.54 MW and 2190 kWh/t.

5.5 Theory and calculation · 7 min

With the feedforward off and a perfect slave, the PDIC-203 error after a compressor trip follows Equation (5.17). At 6.0 kA/m^2 the chlorine header rises at $r_{Cl} = 1.110\,\mathrm{mbar/s}$, the setpoint is 35 mbar and the interlock acts at 20 mbar. Give times to 0.1 s.

(a) For the shipped $K_c = 0.5$ and $\tau_I = 1\,\mathrm{min}$, find τ_e, the error the master would settle at, and the time t_m until the margin is gone. *Checkpoint:* 180 s, 133.2 mbar and 21.5 s.

(b) Find t_m for $K_c = 2$ and $\tau_I = 1\,\text{min}$, and the smallest K_c/τ_I for which the loop holds on feedback alone. *Checkpoint:* 53.9 s, and $4.44\,\text{min}^{-1}$.

(c) The rise scales with the current. Find r_{Cl} and t_m at the shipped tuning when the compressor trips at $7.0\,\text{kA/m}^2$, and say why the engine's crossing, 15 s after the rise begins, comes earlier. *Checkpoint:* 1.295 mbar/s and 18.3 s.

5.6 Theory and calculation · 6 min

Use Equation (5.16) with $R = 8.314\,\text{J/(mol K)}$, $T_g = 320\,\text{K}$, 150 cells, $I = 16.2\,\text{kA}$, $\eta_{Cl} = 0.988$, $V_{Cl} = 300\,\text{m}^3$ and $V_H = 250\,\text{m}^3$. Give rates to 0.001 mbar/s.

(a) Find the chlorine header's rise with the compressor stopped and the fastest rise of the hydrogen header, with its valve shut. *Checkpoint:* 1.110 mbar/s and 1.340 mbar/s.

(b) Find their ratio, and the largest hydrogen header volume that can still keep pace with the chlorine. *Checkpoint:* 1.207, and $301.8\,\text{m}^3$.

(c) Find both rates and their ratio at $3.0\,\text{kA/m}^2$. Then estimate how long the chlorine header takes at $6.0\,\text{kA/m}^2$ to climb from its 150 mbar(g) setpoint to the seal pot at 300 mbar(g), and compare with the 137 s measured after the compressor stops. *Checkpoint:* 0.555 mbar/s, 0.670 mbar/s, 1.207, and about 135 s.

5.7 App challenge · 8 min

Load case 1, Manual feeds: the electrolyzer starts at $5.0\,\text{kA/m}^2$ with the Anolyte loop and the Catholyte loop in Manual, and the current is raised to $6.5\,\text{kA/m}^2$ at 1:00. Keep the default seed, run at 300x and let the run pause at 24:00. Read the scorecard tiles NaOH made, Off-spec caustic, Anolyte lowest and Efficiency lowest.

(a) Report the four tiles for the day as shipped. *Checkpoint:* 86.8 t made with 5.6 t on specification, 81.0 t off, anolyte at least 177.9 g/L, efficiency at least 93.22 %.

(b) Press Reset, set the Anolyte loop and the Catholyte loop to Auto on the AIC-101 and AIC-102 cards before pressing Run, and run the day again. Report the same tiles. *Checkpoint:* 88.4 t made, all on specification, anolyte at least 209.7 g/L, efficiency at least 94.90 %.

(c) Both days run the same current. Explain why the second makes 1.6 t more caustic. *Checkpoint:* The efficiency stays near 94.90 %.

5.8 App challenge · 8 min

Load case 5, Chlorine compressor trip, which opens with Chlorine pressure feedforward Off on the PDIC-203 card. The compressor stops at 1:00 and restarts at 1:30. Keep the default seed and run at 60x. Read the scorecard tiles ΔP worst and Trips · stockout.

(a) Run as shipped. The run pauses when the interlock opens the rectifier. Report the two tiles. *Checkpoint:* 19.8 to 35.3 mbar, 1 trip.

(b) Press Reset, set the feedforward On before pressing Run, run to 1:40 and export the CSV file, whose time column is in hours. Report the ΔP worst tile and the chlorine sent to the emergency absorption, column `cl2_absorbed_t`, in the last row, within 0.01 t. *Checkpoint:* 27.4 to 38.9 mbar, no trip, 1.48 t.

(c) Press Reset, set the feedforward back to Off, the ΔP gain to 5.00 and the ΔP integral time to 0.5 min, and run until the run pauses. Report the two tiles, and say why this tuning holds the ramp and still trips. *Checkpoint:* 19.3 to 40.0 mbar, 1 trip.

5.9 App challenge · 9 min

Load the challenge, a day on the price, which starts at a constant $5.3\,\text{kA/m}^2$ on ratio-only feeds with the chlorine pressure feedforward Off. Keep the default seed and run each day to 24:00 at 1200x,

pressing Run again if the run pauses at a trip. Read the tiles Electricity, NaOH made, ΔP worst, Trips · stockout and Day target.

(a) Run as shipped. Which parts of the target fail? *Checkpoint:* \$4880, 9.8 t on specification of 42.5 t, ΔP 19.9 to 35.4 mbar, 1 trip, 107 min of stockout: not met.

(b) Press Reset, set Brine feed strategy and Water addition strategy to Ratio + trim, the Load scheduler On and the Chlorine pressure feedforward On, all before pressing Run, and run the day. *Checkpoint:* \$10 908, 72.3 t on specification, ΔP 26.4 to 39.0 mbar, no trip or stockout: met.

(c) Press Reset, set only the feedforward back to Off and run the day. Which parts fail, and why does the ΔP worst tile turn red although its lowest value reads 20.0 mbar? *Checkpoint:* \$4288, 42.8 t on specification, ΔP 20.0 to 35.4 mbar, 1 trip, 109 min of stockout.

5.10 Application area · 4 min

In case 7 the brine arrives at 45 °C instead of 65 from 1:00. With the heater held at 0.09 MW the cells settle near 80.7 °C and draw 0.199 MW more electricity for the same caustic. With TIC-103 in Auto the heater rises to 0.87 MW and holds 88 °C. The sunny profile averages \$68.38/MWh.

(a) What does the extra electricity cost a day, and at what price of heat, in dollars per MWh, would the steam cost the same? *Checkpoint:* About \$327 a day, and about \$17.4/MWh of heat.

(b) Give one reason from the chapter to hold 88 °C even when the steam costs more than that.

5.11 Application area · 4 min

In case 6, three hours of 500 ppb hardness leave the membranes with a lasting cost of \$534 a day, priced at the sunny profile's average. The membranes are 18 months old. A plant manager asks what one breakthrough costs if the membranes are replaced at four years of age, and whether the cell-voltage monitor will warn of the next one. Answer both, and recommend how to catch a breakthrough.

5.12 Review the agent · 5 min

An AI assistant is asked how to keep case 5 from tripping and writes: "A larger hydrogen header buffers its pressure. Raise V_H from 250 to 500 m^3 in the Show Model dialog. Twice the gas volume halves how fast the difference can change after the trip, so the interlock has twice the time, with or without the feedforward." Find the flaw, and say how you would check the claim.

CHAPTER 6

Air Separation Simulator for Thermally Coupled Columns and Interacting Purity Loops

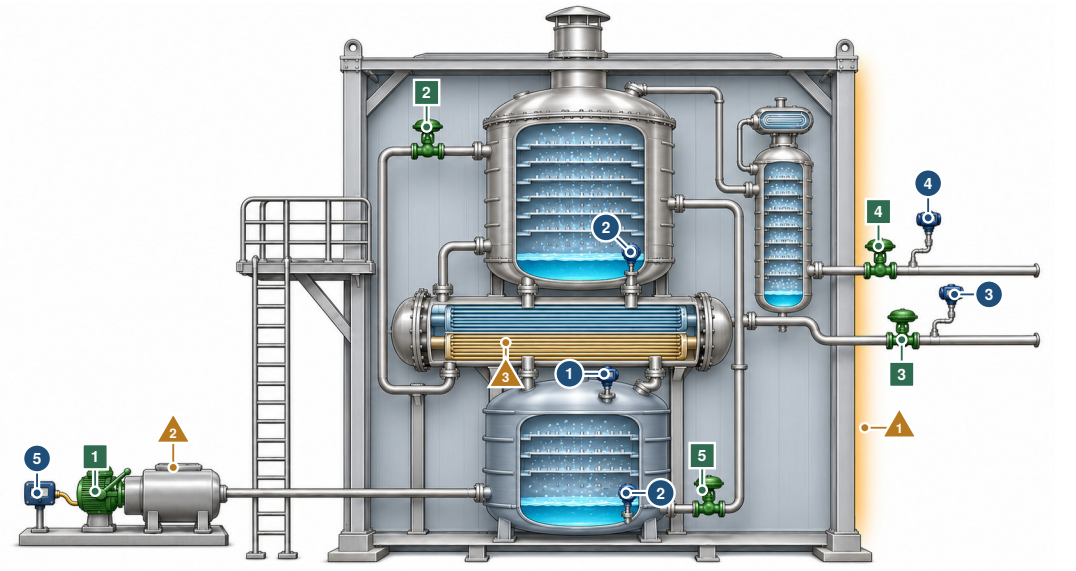

The controller moves 1 compressor guide vanes 2 reflux split on the liquid nitrogen 3 liquid oxygen product valve 4 crude argon draw valve 5 rich liquid valve

The controller reads 1 high-pressure column pressure 2 levels in the two sumps 3 oxygen purity analyzer 4 crude argon analyzer 5 compressor power

What disturbs it 1 heat leak into the cold box 2 warm air at the compressor suction 3 fouling of the shared exchanger

Notice that the one exchanger between the two columns is the only condenser of the lower column and the only reboiler of the upper one, so a valve on either column moves both. Illustration, not to scale.

IN PLAIN TERMS **One Exchanger Shared by Two Columns**

The main idea. An air separation plant cools air until it turns to liquid, then boils it apart into oxygen and nitrogen in two tall columns. One heat exchanger between them cools the top of one column and heats the bottom of the other.

The challenge. So a valve on one column changes the pressure in the other and the power of the machine that squeezes the air. Some faults keep every reading on target while the plant quietly makes less.

Why it matters. Squeezing the air is nearly the whole energy bill, so a few percent of lost output decides whether a plant pays.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page; the history box and the fact box of Section 6.2; the history box of Section 6.3; the try-it boxes of Sections 6.4 and 6.6; the fact box of Section 6.6; the key points
Study	65 min	the whole chapter with its pause points and the exercises
Practice	20 min	the first page; Section 6.6 with its result chart; Section 6.7; Section 6.8; the open challenge of Section 6.9

The control question One shell is the condenser of one column and the reboiler of the other, so all reflux and boil-up come from one duty. Which handle should each purity loop own when the strongest one belongs to the argon loop, and how are failures that hold every setpoint seen?

Learning objectives

- Explain why a valve on the low-pressure column moves the high-pressure column and the compressor at a fixed driving force.
- Compute the oxygen product from the sump balance, and the oxygen lost for each mole of liquid nitrogen sold.
- Rank the four handles by their effect on oxygen purity, and compute the relative gain array of the two purity loops.
- Tune the averaging level loop on the integrating sump by what it does to the product flow.
- Diagnose exchanger fouling and cold-box heat leak from the driving force, the recovery and the specific power.

Prerequisites Sections 1.4.5 (thermal coupling and Equation (1.9)), 1.4.4 (the Kremser relation), 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/air-separation/

See also Volume 6, Chapter 2 for pairing by the relative gain array. Volume 6, Chapter 4 for units heated by one another's vapor. Volume 6, Chapter 8 for a plantwide control structure. Volume 9, Chapter 4 for offset-free predictive control with a disturbance estimator.

6.1 Introduction

A distillation column normally owns its reboiler and its condenser. Steam to one and cooling water to the other are utilities, and a utility can be turned up for one column without asking permission from another. A double-column air separation unit has neither. The nitrogen condensing at the top of the high-pressure column is what boils the oxygen in the sump of the low-pressure column, in the same shell and across the same tubes, and the exchanger transfers heat only because one column sits at 5.600 bar and the other at 1.400 bar, which buys 1.894 K of driving force and nothing more. That one temperature difference carries 546.0 kW. It condenses 109.2 mol/s of nitrogen, which is the reflux for both columns, and it boils 81.5 mol/s of oxygen, which is the only boil-up the low-pressure column has. There is no second condenser and no second reboiler, so nothing on the unit can be unloaded alone.

The consequence is a plant on which the unit boundary an engineer draws is fictitious. Pinching the waste nitrogen valve on top of the *low*-pressure column to 78 % raises the *high*-pressure column from 5.600 to 6.303 bar and the main air compressor from 722 to 767 kW, past its 752 kW motor rating, for exactly the same oxygen, with no pipe joining the two column tops and no controller involved. What the plant corrects is the belief that a valve belongs to the vessel it is bolted to, and behind that a second belief, that a badly behaved purity loop wants better tuning, when the crude argon draw valve moves the oxygen purity by 378 ppm per percent against the reflux split's 70 ppm and the strong handle is already spoken for.

The class of problem is wider than cryogenics. Any two units joined by one exchanger that is the only reboiler of one and the only condenser of the other inherit the same structure, and heat-integrated distillation trains, double-effect evaporators, reactors working against their own feed preheaters and refrigeration cascades are all built that way. Chapter 8 meets the other kind of thermal coupling, two column sections that share one condenser and one reboiler in a single shell. Dynamic comparisons of heat-integrated column configurations established that the integration changes the control problem and not merely the energy balance [1], interaction between candidate loops is measured with the relative gain array [2], and which variables a coupled plant should hold constant is the plantwide control question [3]. Inventories on such a plant are held by averaging loops rather than tight ones [4].

Cryogenic air separation is where these ideas earn money rather than marks. Air compression is essentially the whole energy bill of the unit, oxygen is worth so little per cubic meter that a few percent of specific power decides whether a site is competitive, and published assessments put the compression work at the center of the economics [5, 6]. Published control studies of this plant work on tray-by-tray models or on commercial simulator cases, and evaluate structures through steady-state loss or through model predictive control on reduced-order models [7–9]. However, none of those models can be opened, perturbed and re-run in a classroom hour, so a student meets the coupling as a claim rather than as an experiment. This chapter works with a browser-hosted double-column air separation unit with a crude argon sidearm, written out closely enough that its eighteen states, its shared exchanger and its five loops can be rebuilt from the text, and uses it to show that the interaction on this plant is a conservation law rather than a matrix of gains.

Section 6.2 traces how the double column and its operation developed. Section 6.3 sets out the model: the shared exchanger, the pressures and the compressor, a reduced-order column in which each lumped section is marched along its own operating line rather than treated as one equilibrium stage, the argon tracer and the sidearm, and every constant. Section 6.4 measures the coupling. The high-pressure column rises 3.11 bar for every bar of low-pressure column at constant duty, the duty setpoint meets the compressor rating at 5.699 bar and 84.10 mol/s of boil-up, and a relative gain array on the two purities returns $\lambda_{11} = 0.471$ for the pairing the plant actually uses. Section 6.5 describes the five loops, the compressor power override and the predictive controller. Section 6.6 produces the three failures that hold every setpoint, exchanger fouling, cold-box heat leak and a recovery pinch, diagnoses them from variables that are not controlled, and compares full cascade with model predictive control. Section 6.7 names the real devices behind the loops, and Sections 6.8 to 6.10 give the limitations, a discussion and the conclusion.

6.2 Background

Air separation began as a liquefaction problem and became a distillation problem. Commercial air liquefaction arrived at the end of the nineteenth century on the Joule-Thomson effect with regenerative heat exchange, and the expansion engine that followed cut the pressure and the work required, so that liquid air became a feedstock [10]. A single rectifying column answered the separation badly, because it can make pure nitrogen overhead or oxygen-enriched liquid in the sump but not both, the only reflux available being what its own condenser makes and the condenser needing a colder fluid than the column contains. The double column solved it by putting two columns at two pressures and letting the high-pressure overhead condense against the low-pressure bottoms, which supplies reflux to both from one duty. Exergy analyses show why that arrangement persists, since the temperature difference between the two ends is the entire thermodynamic penalty of the separation and the double column keeps it near the minimum a real exchanger allows [11]. Argon was added later as a sidearm hung off the low-pressure column where argon accumulates, and it has no reboiler either, because the vapor drawn from the argon belly is its boil-up.

> ⧖ HISTORY **Liquid Air and the First Rectifier, 1895 to 1905**
>
> In July 1895 Carl Linde, a professor in Munich, applied for a United States patent on a process for producing very low temperatures, liquefying air and separating its constituents. Air throttled through a valve from high to low pressure comes out colder, and Linde let the expanded air cool the incoming compressed air in a countercurrent coil about one hundred meters long until it liquefied [12]. His application of October 1903 added a rectifying column. Compressed air condensed in a coil in the liquid chamber at the column's foot, boiling the liquid there, and nitrogen drawn from the top was compressed again and condensed in a second coil to return as reflux, until pure nitrogen left the top and pure liquid oxygen collected below [13]. One stream condensing to boil another is the job the shared exchanger of Section 6.3 does between two columns.

Operation followed the hardware slowly. Early plants ran on fixed line-ups with the valve positions written on a card, a shift chemist who analyzed the product once or twice a day, and an operator who judged the column from pressure gauges and from the frost line on the cold box. The inventory is what made that workable, because twenty cubic meters of liquid oxygen and a sidearm holding two hundred kilomoles absorb hours of imbalance, so a plant lined up correctly stays lined up and a plant lined up incorrectly takes most of a shift to say so. What made it expensive was that nobody could see the compressor bill against the separation being bought with it.

Analysis of the columns themselves matured in the 1980s. Simulation of the upper column together with the crude argon sidearm as one connected problem was published in 1989 and established the belly draw and the sidearm reflux ratio as the quantities that decide argon recovery [14]. First-principles dynamic models of cryogenic distillation with multivariable control designed on them appeared around 2000 [7], and the obstacle to using such models online was removed by compartmental reduction, in which a column of many trays is collapsed to a few lumped sections whose dynamics reproduce the full model closely enough for control [8]. Nonlinear model predictive control on those reduced models followed on high-purity air separation columns [9], and that combination, regulatory loops for pressure and inventory with a multivariable controller above them, is the current standard answer.

The economics moved again in the last fifteen years. An air separation unit is a large, flexible electrical load with liquid storage attached, and the literature has since asked how much of its rate can be moved on an hourly electricity price given those inventories and dynamics [15, 16]. Both arguments rest on the physics this simulator carries, since what limits how fast the plant can be moved is the shared duty at one end and the liquid inventories at the other. The model of this chapter is in the spirit of the compartmental work, lumped where the control problem allows it and detailed where the coupling lives.

> ✦ DID YOU KNOW? **Electricity Is the Feedstock**
>
> In 2010 the industrial gas sector of the United States used 19.4 TW h of electricity, 2.5 percent of what all of manufacturing used [17]. A typical cryogenic unit uses more than 200 kW h of electricity per tonne of oxygen, against a thermodynamic minimum of around 51 kW h [18].

6.3 Air Separation Unit Model

The application simulates a double column taking 124.1 mol/s of air (10 014 m^3/h at normal conditions) through a main air compressor at 722.0 kW, with 12 % of it diverted to an expansion turbine and admitted to the low-pressure column as vapor. It delivers 23.70 mol/s of liquid oxygen (1912 m^3/h) at 99.622 %, 1.000 mol/s of crude argon (81 m^3/h) at 96.02 %, 2.18 mol/s of liquid nitrogen (176 m^3/h) and 97.2 mol/s of waste nitrogen carrying 2.53 % of the heavy key, for an oxygen recovery of 90.8 %, an argon recovery of 83.2 % and a specific power of 0.379 kW h per normal cubic meter of contained oxygen. Figure 6.1 carries the arrangement together with every symbol that appears below. Flows are in mol/s, hold-ups in mol, pressures in bar, temperatures in K and the clock runs in minutes, so every accumulation term carries a factor of 60.

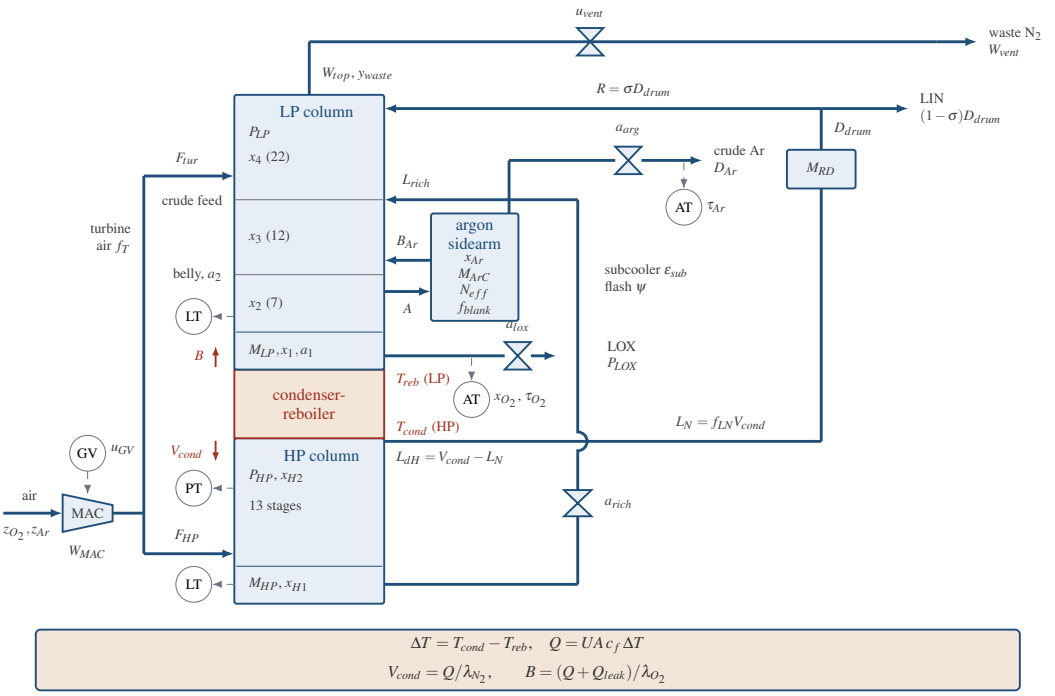

Figure 6.1 Double column with the condenser-reboiler drawn as the single wall between the two columns. The shared exchanger condenses V_{cond} of nitrogen on the high-pressure side and boils B of oxygen on the low-pressure side from one duty Q across one driving force $\Delta T = T_{cond} - T_{reb}$. The low-pressure column carries three marched sections of 22, 12 and 7 theoretical stages with the crude argon sidearm hung off its belly, and the high-pressure column is one 13-stage rectifying section solved as an absorber.

6.3.1 Variable Classification

Every signal a loop or the scorecard reads is classified in Table 6.1. The reflux split σ never touches the compressor, so it moves purity, recovery and liquid nitrogen production at constant power, and the argon in the liquid oxygen a_1 is the oxygen purity written from the other end, because the nitrogen left in the sump is 14 ppb.

6.3.2 The Shared Exchanger

One duty and one driving force set every internal flow on the plant, as shown in Equation (6.1).

$$\begin{aligned} Q &= UA\, c_f\, (T_{cond} - T_{reb}), \\ V_{cond} &= \frac{Q}{\lambda_{N_2}}, \qquad B = \frac{Q}{\lambda_{O_2}} + \frac{Q_{leak}}{\lambda_{O_2}} \end{aligned} \tag{6.1}$$

where UA is the exchanger conductance in W/K, c_f is a fouling multiplier, λ_{N_2} and λ_{O_2} are the latent heats in J/mol, Q_{leak} is cold-box heat leak in W taken from a slider calibrated in kW, and V_{cond} and B are the nitrogen condensed and the oxygen boiled in mol/s, with the duty floored at zero. Heat leak appears in B and not in V_{cond}, which is why a warm cold box costs product rather than energy.

Both saturation temperatures are Raoult bubble points of the actual local compositions at the actual local pressures, as shown in Equation (6.2).

$$\begin{aligned} T_{cond} &= T_{bub}(P_{HP}, x_{H2}, 0), \\ T_{reb} &= T_{bub}(P_{LP}, x_1 - a_1, a_1) \end{aligned} \tag{6.2}$$

Table 6.1 Signal Classification. Nominal Values Belong to the Design Line-Up Every Run Starts From.

Symbol	Role	Nominal	Range
x_{O_2}	CV, measured	99.622 %	98.5–99.9 % SP
P_{HP}	CV, measured	5.600 bar	4.6–6.6 bar SP
L (LOX sump)	CV, measured	50.0 %	20–80 % SP
x_{Ar}	CV, measured	96.02 %	85–99 % SP
L_{HP}	CV, measured	50.0 %	30–70 % SP
W_{MAC}	constraint	722.0 kW	cap at 752 kW
u_{GV}	MV	100.0 %	45–110 %
σ	MV	95.0 %	55–98 %
u_{lox}	MV	52.66 %	0–100 %
u_{arg}	MV	40.0 %	5–80 %
u_{rich}	MV	53.29 %	5–100 %
f_T	DV, set	12.0 %	0–25 %
u_{vent}	DV, set	100 %	40–100 %
u_A	DV, set	100 %	50–130 %
T_{suc}	DV, measured	30.0 °C	5–50 °C
c_f, c_a	DV, unmeas.	1.00×	0.50–1.05, 0.80–1.02
Q_{leak}	DV, unmeas.	0 kW	0–60 kW
ε_{sub} factor	DV, unmeas.	1.00×	0.40–1.00
P_{HP}, P_{LP}	states, bar	5.600, 1.400	1.05–12, 1.02–6
M_{HP}, M_{LP}, M_{RD}	states, mol	120000, 700000, 30000	0–240000, 0–1400000
x_{H1}, x_{H2}	states	0.3646, 4.68×10^{-5}	0–1
x_1–x_4	states	1.000, 1.000, 0.99995, 0.5365	0–1
a_1, a_2, x_{Ar}	states	0.003778, 0.10457, 0.9602	0–0.5, 0–0.6, 0–1
a_{GV}, a_{rich}, a_{lox}, a_{arg}	actuators	as MV	as MV
ΔT, V_{cond}, B	outputs	1.894, 109.2, 81.5	K, mol/s
a_1 in product	output	3778 ppm	—
y_{N_2} at belly	output	143 ppm	alarm 320 ppm
w_{spec}	output	0.379	kW h/m^3

where $T_{bub}(P,x_{O_2},x_{Ar})$ solves $x_{N_2}p^{sat}_{N_2}+x_{O_2}p^{sat}_{O_2}+x_{Ar}p^{sat}_{Ar}=P$ with $x_{N_2}=1-x_{O_2}-x_{Ar}$ and Antoine vapor pressures $\log_{10}(p^{sat}/\mathrm{bar})=A-B/(T+C)$ taken from the NIST WebBook, with (A,B,C) equal to $(3.7362,264.651,-6.788)$ for nitrogen, $(3.9523,340.024,-4.144)$ for oxygen and $(3.29555,215.24,-22.233)$ for argon. The solve is exactly ten Newton iterations with an analytic derivative and no early exit, because the duty is proportional to the difference of two of these temperatures that are only 1.894 K apart. Equation (6.2) makes the coupling run in both directions, since a purer sump boils hotter, closes the driving force and takes away its own boil-up.

The ratio of the two latent heats decides what the plant can make. Nitrogen at 5.6 bar gives up 5000 J/mol and oxygen at 1.4 bar needs 6700 J/mol, so one mole condensed boils only 0.746 of a mole and the rest has to leave the low-pressure column as liquid. The steady sump balance $L_2=B+P_{LOX}$, written per mole of air fed to the high-pressure column, is Equation (6.3).

$$\frac{P_{LOX}}{F_{HP}}=\sigma f_{LN}(1-\psi_R)+(1-f_{LN})(1-\psi)+\frac{B_{Ar}}{F_{HP}} -\frac{\lambda_{N_2}}{\lambda_{O_2}}-\frac{Q_{leak}}{\lambda_{O_2}F_{HP}} \tag{6.3}$$

where σ is the reflux split, f_{LN} is the fraction of the high-pressure condensate drawn off as liquid, ψ and ψ_R are the flash fractions of the two liquid feeds and B_{Ar} is the sidearm bottoms returned to the low-pressure column. At the design line-up the five terms are 0.37052, 0.49200, 0.10073, −0.74627 and zero, which sum to 0.21697 against the 0.2095 of oxygen the air contains, and multiplying by $F_{HP}=109.21\,\mathrm{mol/s}$ recovers the 23.70 mol/s of product. Equation (6.3) reproduces the model's own liquid oxygen draw to better than 2×10^{-9} over every control and disturbance setting the application

allows, and two of its terms are the whole of the plant's economics. Every mole of liquid nitrogen sold instead of returned as reflux costs very nearly a mole of oxygen, and the last term is cold-box heat leak, which subtracts product without touching any other term.

> ⧖ HISTORY **Reboiler Explosions, 1964 and 1997**
>
> The exchanger at the center of this chapter is also where an air separation unit is most dangerous. Hydrocarbons and other contaminants in the intake air boil higher than oxygen and collect in the oxygen-rich reboilers, above all in the sump of the low-pressure column. An industry code records three violent explosions that began in a reboiler, one in 1964 and two in 1997, and one of the 1997 explosions destroyed its unit and seriously damaged the neighboring plant that used the oxygen. All three are thought to have run the same course: contaminants accumulated in the exchanger passages ignited, the aluminum of the exchanger itself burned, the liquid oxygen flashed to vapor, and the column and the cold box burst. The code concludes that any hydrocarbon accumulation in a liquid oxygen reboiler must be prevented [19].

6.3.3 Pressures, Air and the Compressor

Each column pressure is a vapor inventory that integrates a flow imbalance, as shown in Equation (6.4).

$$\begin{aligned} C_{HP}\,\frac{\mathrm{d}P_{HP}}{\mathrm{d}t} &= 60\,(F_{HP} - V_{cond}) \\ C_{LP}\,\frac{\mathrm{d}P_{LP}}{\mathrm{d}t} &= 60\,(W_{top} - W_{vent}) \end{aligned} \tag{6.4}$$

where C_{HP} and C_{LP} are vapor capacitances in mol/bar, W_{top} is the vapor leaving the top of the low-pressure column and W_{vent} is what the waste valve passes. The high-pressure column therefore pressurizes until the nitrogen at its top condenses as fast as the air arrives, which fixes the condensing temperature, the driving force and the boil-up in the other column, so putting P_{HP} on the guide vanes puts that boil-up on the guide vanes. Differentiating V_{cond} of (6.1) with respect to P_{HP} gives 132.6 mol/s per bar, so (6.4) is a lag of $C_{HP}/(60 \times 132.6) = 2.51$ min on its own and about 11 min once the low-pressure column has moved with it. The waste valve is a square-root device, as shown in Equation (6.5).

$$W_{vent} = C_{vent}\,\frac{u_{vent}}{100}\,\sqrt{\max(0,\, P_{LP} - P_{vent})} \tag{6.5}$$

Air is a volumetric flow through the guide vanes of the main air compressor, so a warm morning is fewer moles through the same opening and more work for each of them, as shown in Equation (6.6).

$$\begin{aligned} F_{air} &= F_{air,0}\,\frac{a_{GV}}{100}\,\frac{T_{suc,0}}{T_{suc}}\,c_a, \qquad F_{tur} = F_{air}\,\frac{f_T}{100}, \\ F_{HP} &= F_{air} - F_{tur}, \\ W_{MAC} &= \frac{F_{air}\,R\,T_{suc}}{\eta_{MAC}}\,\ln\max\left(1.02,\; \frac{P_{HP} + \Delta P_{mac}}{P_{amb}}\right) \end{aligned} \tag{6.6}$$

where a_{GV} is the vane position in percent, T_{suc} is the suction temperature in K, c_a is the filter and molecular sieve condition, f_T is the turbine share of the air, R is the gas constant, η_{MAC} is the isothermal efficiency and ΔP_{mac} is the exchanger and piping drop above the column. The shaft power of (6.6) comes out in watts and is reported in kilowatts. Both ambient effects act the same way, which is why a twelve-degree afternoon costs 29 kW at a fixed duty.

6.3.4 Sections Marched Along Their Own Operating Lines

Relative volatilities are referred to the heavy key, which is oxygen and argon together, and both fall as pressure rises, as shown in Equation (6.7).

$$\alpha_{N_2}(P) = \alpha_{N_2}^{ref}\left(\frac{P_{ref}}{P}\right)^{n_{N_2}}, \quad \alpha_{Ar}(P) = \alpha_{Ar}^{ref}\left(\frac{P_{ref}}{P}\right)^{n_{Ar}} \tag{6.7}$$

At the design pressures α_{N_2} is 3.100 in the high-pressure column and 3.450 in the low-pressure column, and α_{Ar} is 1.105. Equilibrium between a liquid at heavy-key fraction x and its vapor is Equation (6.8).

$$y = \frac{x}{\alpha - (\alpha - 1)x}, \qquad x = \frac{\alpha y}{1 + (\alpha - 1)y} \tag{6.8}$$

Within a lumped section the flows are constant, so the net heavy key traveling downward is the same across every tray in it, and the section can be marched one stage at a time instead of solved as a boundary value problem, as shown in Equation (6.9).

$$\begin{aligned} W &= Lx - Vy \quad \text{(constant in a section)}, \\ y_k &= \max\left(\frac{Lx_k - W}{V},\ \frac{x_k}{\alpha - (\alpha - 1)x_k}\right), \\ x_{k+1} &= \frac{\alpha y_k}{1 + (\alpha - 1)y_k} \end{aligned} \tag{6.9}$$

where L and V are the liquid and vapor rates in the section, W is the net heavy key in mol/s and the march starts from the liquid entering the section's top. Equation (6.9) is exact under constant molar overflow and constant relative volatility and costs one multiply per stage. One equilibrium stage per lump reaches 97 % oxygen and stops, and α^n reaches ppm purities but only because that is the total-reflux answer, whereas this column runs at L/V near unity. The equilibrium floor in the second line of (6.9) states that liquid must get richer in the heavy key at every stage going down, so a section asked to send down more heavy key than its inlet liquid carries pinches and passes its inlet composition through instead of driving the operating line negative and producing pure nitrogen in one step.

The top of the high-pressure column is the exception, because its heavy key is 46.8 ppm and marching it would need the difference of two numbers near 24 mol/s whose answer is near 0.003 mol/s. It is solved as an absorber by the Kremser relation instead, the relation of Section 1.4.4 written for the heavy key, as shown in Equation (6.10).

$$\begin{aligned} A &= \frac{\alpha L}{V}, \qquad \phi = \frac{A - 1}{A^{n+1} - 1}, \\ y_{top} &= \frac{x_{in}}{\alpha} + \left(g_{in} - \frac{x_{in}}{\alpha}\right)\phi, \\ x_{bot} &= x_{in} + \frac{V}{L}\left(g_{in} - y_{top}\right) \end{aligned} \tag{6.10}$$

where g_{in} is the heavy key in the vapor entering the bottom of the section, n is the stage count and ϕ is replaced by $1/(n+1)$ when A is within 10^{-6} of unity.

6.3.5 Internal Flows and the Net Heavy Key

The liquid streams follow from the valve positions and the condensate, as shown in Equation (6.11).

$$\begin{aligned} L_{rich} &= C_{rich}\frac{a_{rich}}{100}\sqrt{\max(0.02,\Delta P)}\min\left(1,\frac{M_{HP}}{8000}\right), \\ L_N &= f_{LN}V_{cond}, \qquad L_{dH} = V_{cond} - L_N, \\ D_{drum} &= k_{drum}\frac{M_{RD}}{M_{RD,0}}, \\ R &= \sigma D_{drum}, \qquad P_{LIN} = (1-\sigma)D_{drum}, \\ P_{LOX} &= C_{lox}\frac{a_{lox}}{100}\min\left(1,\frac{M_{LP}}{20000}\right), \\ D_{Ar} &= D_{Ar}^{max}\frac{a_{arg}}{100}, \\ A &= A_0\frac{u_A}{100}, \qquad B_{Ar} = \max(0, A - D_{Ar}) \end{aligned} \tag{6.11}$$

where A is the argon belly vapor draw and the two min terms starve a valve whose vessel has emptied. Both liquid feeds are throttled from 5.6 to 1.4 bar and would flash without the subcooler, so the flash fractions of Equation (6.12) carry its effectiveness.

$$\begin{aligned} \psi &= \frac{c_{pL}\left(T_{rich} - T_{rich}^{LP}\right)}{\lambda_{O_2}}\left(1 - \varepsilon_{sub}c_{sub}\right) + \frac{B_{Ar}\lambda_{Ar}}{\lambda_{O_2}L_{rich}}, \\ \psi_R &= \frac{c_{pL}\left(T_{cond} - T_{LN}^{LP}\right)}{\lambda_{N_2}}\left(1 - \varepsilon_{sub}c_{sub}\right) \end{aligned} \tag{6.12}$$

where $\Delta P = P_{HP} - P_{LP}$, c_{pL} is the liquid heat capacity, c_{sub} is the subcooler condition disturbance, the temperatures are bubble points of the same liquid at the two pressures, and the second term of ψ folds in the crude argon condenser, which has no coolant of its own and boils against the rich liquid. Each of the three terms is clipped to $[0, 0.45]$ and their sum to $[0, 0.70]$, and at design $\psi = 0.1800$ and $\psi_R = 0.0250$. The internal traffic then follows as shown in Equation (6.13).

$$\begin{aligned} R_l &= R\left(1 - \psi_R\right), \qquad L_f = L_{rich}\left(1 - \psi\right), \\ L_3 &= R_l + L_f, \qquad L_2 = L_3 + B_{Ar}, \\ V_3 &= \max(0, B - A), \qquad V_4 = V_3 + L_{rich}\psi + F_{tur}, \\ W_{top} &= V_4 + R\,\psi_R \end{aligned} \tag{6.13}$$

with design values $R_l = 40.46\,\text{mol/s}$, $L_3 = 94.19\,\text{mol/s}$, $L_2 = 105.19\,\text{mol/s}$, $V_3 = 69.50\,\text{mol/s}$, $V_4 = 96.19\,\text{mol/s}$ and $W_{top} = 97.2\,\text{mol/s}$.

The net heavy key crossing each section plane is written as what leaves below that plane plus what accumulates below it, which is what makes the whole column conserve, as shown in Equation (6.14).

$$\begin{aligned} W_2 &= (L_2 - B)\,x_1, \qquad W_3 = W_2 + D_{Ar}, \\ W_4 &= W_3 - L_{rich}x_{H1} - F_{tur}\left(z_{O_2} + z_{Ar}\right) \end{aligned} \tag{6.14}$$

Below the argon belly the only inventory is the sump, so the accumulation term collapses the product draw out of W_2 and at steady state $W_2 = P_{LOX}x_1$. That form is a conditioning choice. The net nitrogen rising between the crude feed and the belly is $V_3 - L_3 + W_3$, terms of 69.50, -94.19 and 24.70 mol/s whose design value is -3.4×10^{-7} mol/s, the 14 ppb of nitrogen in the product, and (6.14) telescopes it analytically to $-(L_2 - B)(1 - x_1)$ so that nothing is left to cancel. In the product-draw form a liquid oxygen valve 0.01 points from balance would add 0.0045 mol/s to it, as much as the nitrogen crossing

the belly in each direction. At design $W_4 = -2.456\,\mathrm{mol/s}$, so above the crude feed the heavy key travels net upward, and dividing that by W_{top} gives the 2.53 % of oxygen the waste nitrogen carries away. No stage count appears in it, which is why recovery is decided at the reflux split and not at the top of the column.

The three low-pressure sections are marched downward by (6.9) from the liquid entering each of them, using $(x_R, W_4, R_l, V_4, \alpha_{N_2}, 22)$ above the crude feed, $(x_{in,3}, W_3, L_3, V_3, \alpha_{N_2}, 12)$ between the crude feed and the belly, and $(x_{in,2}, W_2, L_2, B, \alpha_{N_2}, 7)$ below the belly, with the mixed inlets $x_{in,3} = (R_l x_4 + L_f x_{fl})/L_3$ and $x_{in,2} = (L_3 x_3 + B_{Ar} x_{B_{Ar}})/L_2$ and the flashed feed compositions $x_{fl} = (x_{H1} - \psi y(x_{H1}))/(1-\psi)$ and $x_R = (x_{H2} - \psi_R y(x_{H2}))/(1-\psi_R)$. When the top section's operating line starts below the equilibrium curve the march has nothing to say, so the model reports a drain rate $\Gamma = \max(0, W_4 - R_l x_R)$ instead and lets the section lose heavy key from its own hold-up at exactly the rate the balance is short by.

6.3.6 Argon as a Tracer and the Sidearm

Argon rides on the heavy key rather than appearing as a rigorous third component, and it leaves the belly three ways, to the sidearm, downward into the oxygen product and upward with the waste nitrogen. The two slips are set by the argon stripping factor in the section concerned, as shown in Equation (6.15).

$$\begin{aligned} S_{dn} &= \frac{\alpha_{Ar} B}{L_2}, \qquad S_{up} = \frac{\alpha_{Ar} V_4}{R_l}, \\ k_{dn} &= \frac{k_{dn,0}}{S_{dn}^{m_{dn}}}, \qquad k_{up} = k_{up,0} S_{up}^{m_{up}} \end{aligned} \tag{6.15}$$

with S_{dn} clamped to $[0.30, 3]$ and S_{up} to $[0.30, 8]$. At the design line-up $S_{dn} = 0.856$ and $S_{up} = 2.627$. The sixth power on S_{dn} is what makes anything that raises the liquid traffic below the belly put argon into the product, which is the mechanism behind the reflux split's effect on purity.

The sidearm has no reboiler, so the belly vapor is its boil-up and whatever is not drawn off condenses and returns as reflux, as shown in Equation (6.16).

$$\begin{aligned} f_{blank} &= \frac{1}{1 + \left(y_{N_2}/y^*_{N_2}\right)^4}, \quad R_{Ar} = f_{blank} \frac{B_{Ar}}{\max(0.02, D_{Ar})}, \\ N_{eff} &= N_{Ar} \frac{R_{Ar}}{R_{Ar} + R_{1/2}}, \qquad \frac{x^*_{Ar}}{1 - x^*_{Ar}} = \alpha_{Ar}^{N_{eff}} \frac{z^F_{Ar}}{1 - z^F_{Ar}} \end{aligned} \tag{6.16}$$

where y_{N_2} is the nitrogen in the belly vapor, $y^*_{N_2}$ is the blanketing threshold, N_{Ar} is the stage count at total reflux, $R_{1/2}$ is the reflux ratio at which the column has half its stages and z^F_{Ar} is the argon fraction of the argon-plus-oxygen entering the sidearm. Nitrogen will not condense at the argon condenser temperature, so it collects in the top of the shell and covers the tubes, and the fourth power makes that a threshold rather than a trend. At design $R_{Ar} = 10.58$, $N_{eff} = 52.4$, $f_{blank} = 0.9615$ and the belly nitrogen is 143 ppm against a threshold of 320 ppm. The sidearm relaxes toward x^*_{Ar} with the hold-up time constant $\tau_{Ar} = M_{ArC}/(60A)$, which is 277.8 min at the design belly draw.

6.3.7 State Equations

Eighteen states are integrated, as shown in Equations (6.4), (6.17), (6.18), (6.19) and (6.20). The three liquid inventories are plain balances.

$$\begin{aligned} \frac{\mathrm{d}M_{HP}}{\mathrm{d}t} &= 60\,(L_{dH} - L_{rich})\,, \\ \frac{\mathrm{d}M_{RD}}{\mathrm{d}t} &= 60\,(L_N - D_{drum})\,, \\ \frac{\mathrm{d}M_{LP}}{\mathrm{d}t} &= 60\,(L_2 - B - P_{LOX}) \end{aligned} \tag{6.17}$$

The liquid oxygen sump is the integrator the level loop lives on. Its inventory at 50 % is 700 000 mol, which is 19.6 m^3, and differentiating (6.17) against valve position gives an integrating gain of 0.001 93 %/min of level per percent of valve.

The compositions relax toward what each section's own profile asks for at the rate its liquid turns over, as shown in Equation (6.18).

$$\begin{aligned}
\frac{dx_{H1}}{dt} &= \frac{60\,L_{dH}\left(x_{bot}^{HP} - x_{H1}\right)}{\max(4000, M_{HP})},\\
\frac{dx_{H2}}{dt} &= \frac{60\,L_{dH}\left(y_{top}^{HP} - x_{H2}\right)}{M_{H2}},\\
\frac{dx_1}{dt} &= \frac{60\left[L_2\left(x_2 - x_1\right) - B\left(y(x_1) - x_1\right)\right]}{\max(20000, M_{LP})},\\
\frac{dx_2}{dt} &= \frac{60\,L_2\left(x_{bot,2} - x_2\right)}{M_{st}}, \quad \frac{dx_3}{dt} = \frac{60\,L_3\left(x_{bot,3} - x_3\right)}{M_{st}},\\
\frac{dx_4}{dt} &= \frac{60\left[\max(4, R_l)\left(x_{bot,4} - x_4\right) - \Gamma\right]}{M_{st}}
\end{aligned} \tag{6.18}$$

where $x_{bot,k}$ is the liquid leaving the bottom of section k from the march, frozen at the state itself when that section is infeasible. The sump is written as differences of compositions rather than of flows so that a sump which is filling or draining cannot make its own composition drift by arithmetic alone, and the argon states follow the same pattern, as shown in Equation (6.19).

$$\begin{aligned}
\frac{da_2}{dt} &= \frac{60\left[Ar_{in} - D_{Ar}x_{Ar} - P_{LOX}k_{dn}a_2 - W_{top}k_{up}a_2\right]}{M_{belly}},\\
\frac{da_1}{dt} &= \frac{60\left[L_2\left(k_{dn}a_2 - a_1\right) - B\left(\alpha_{Ar}a_1 - a_1\right)\right]}{\max(20000, M_{LP})},\\
\frac{dx_{Ar}}{dt} &= \frac{x_{Ar}^{*} - x_{Ar}}{\max(1, \tau_{Ar})}
\end{aligned} \tag{6.19}$$

where $Ar_{in} = F_{HP}z_{Ar}\eta_{Ar}^{HP} + F_{tur}z_{Ar}$ is the argon arriving at the low-pressure column. The exchange terms in the sump argon balance are L_2 and B rather than the product draw, which is why the argon in the liquid oxygen lines out in about 111 min rather than the eight hours the draw rate alone would suggest. The four actuator positions each lag their own command, as shown in Equation (6.20).

$$\tau_j \frac{da_j}{dt} = u_j - a_j, \qquad j \in \{GV,\ rich,\ lox,\ arg\} \tag{6.20}$$

The outputs the scorecard reports are the oxygen purity, the oxygen and argon recoveries and the specific power, as shown in Equation (6.21).

$$\begin{aligned}
x_{O_2} &= x_1 - a_1, \qquad & w_{spec} &= \frac{W_{MAC}}{\dot{V}_{O_2}x_{O_2}},\\
r_{O_2} &= \frac{P_{LOX}x_{O_2}}{F_{air}z_{O_2}}, \qquad & r_{Ar} &= \frac{D_{Ar}x_{Ar}}{F_{air}z_{Ar}}
\end{aligned} \tag{6.21}$$

where $\dot{V}_{O_2}$ is the product volumetric rate at normal conditions and r_{O_2} and r_{Ar} are the two recoveries.

6.3.8 Parameters, Initialization and Numerics

Every constant appears in Tables 6.2 and 6.3. Nine are hidden and never displayed, and the four stage counts are structure rather than parameters, so anybody who wants the volatilities, the air composition or the vessel spans has to infer them. The remainder are visible in the parameter panel and adjustable.

Table 6.2 Air, Exchanger, Pressure and Separation Parameters

Symbol	Meaning	Value	Unit	Visibility
$F_{air,0}$	air at 100 % vane	124.1	mol/s	adjustable
z_{O_2}	oxygen in air	0.2095	—	hidden
z_{Ar}	argon in air	0.0093	—	hidden
$T_{suc,0}$	reference suction	303.15	K	hidden
η_{MAC}	isothermal efficiency	0.767	—	adjustable
P_{amb}	atmospheric pressure	1.013	bar	hidden
ΔP_{mac}	drop above the column	0.35	bar	adjustable
W_{max}	compressor power limit	752	kW	adjustable
UA	exchanger conductance	288300	W/K	adjustable
λ_{N_2}, λ_{O_2}	latent heats, N_2 and O_2	5000, 6700	J/mol	adjustable
λ_{Ar}	argon latent heat	6400	J/mol	hidden
ΔT_{min}	unloading threshold	1.10	K	adjustable
C_{HP}, C_{LP}	vapor capacitances	20000, 25000	mol/bar	adjustable
C_{vent}	waste valve size	164.33	mol/s bar$^{-1/2}$	adjustable
P_{vent}	downstream of the vent	1.05	bar	adjustable
C_{rich}	rich liquid valve size	60.0	mol/s bar$^{-1/2}$	adjustable
$\alpha_{N_2}^{ref}$, α_{Ar}^{ref}	volatilities at 1.4 bar	3.45, 1.105	—	adjustable
P_{ref}	volatility reference	1.4	bar	hidden
n_{N_2}, n_{Ar}	pressure exponents	0.0771, 0.060	—	adjustable
f_{LN}	condensate drawn as liquid	0.400	—	adjustable
ε_{sub}	subcooler effectiveness	0.85	—	adjustable
c_{pL}	liquid heat capacity	56	J/mol K	hidden

Table 6.3 Inventory, Argon, Actuator, Instrument and Economic Parameters

Symbol	Meaning	Value	Unit	Visibility
$M_{HP,0}$, $M_{LP,0}$	sumps at 50 %	120000, 700000	mol	adjustable
M_{HP}^{max}, M_{LP}^{max}	sumps at 100 %	240000, 1400000	mol	hidden
$M_{RD,0}$	reflux drum hold-up	30000	mol	adjustable
M_{st}	hold-up per LP section	70000	mol	adjustable
M_{H2}	HP rectifying hold-up	45000	mol	adjustable
M_{belly}	argon belly hold-up	30000	mol	adjustable
M_{ArC}	sidearm hold-up	200000	mol	adjustable
A_0	belly draw at 100 %	12.0	mol/s	adjustable
D_{Ar}^{max}	argon draw, valve open	2.5	mol/s	adjustable
N_{Ar}	sidearm stages at total reflux	67	—	adjustable
$R_{1/2}$	reflux ratio for half the stages	2.95	—	adjustable
$y_{N_2}^*$	half-blanketing nitrogen	3.2×10^{-4}	—	adjustable
$k_{dn,0}$	downward argon slip	0.01538	—	adjustable
$k_{up,0}$	upward argon slip	1.240×10^{-3}	—	adjustable
m_{dn}, m_{up}	slip exponents	6.0, 2.0	—	adjustable
η_{Ar}^{HP}	air argon into the crude liquid	0.99	—	adjustable
N_{HP}, $N_{2..4}$	stages per section	13, 7, 12, 22	—	hidden
C_{lox}	LOX draw, valve open	45.0	mol/s	adjustable
k_{drum}	drum drain at normal level	44.0	mol/s	adjustable
τ_{GV}, τ_{rich}, τ_{lox}, τ_{arg}	actuator travel	0.30, 0.25, 0.35, 0.40	min	adjustable
τ_{O_2}, τ_{Ar}^m	the two analyzers	6.0, 9.0	min	adjustable
τ_{Pm}, τ_{Lm}	transmitter response	0.10, 0.50	min	adjustable
$x_{O_2}^{spec}$, x_{Ar}^{spec}	product specifications	0.995, 0.93	—	adjustable
L_{lo}, L_{hi}	LOX level alarms	25, 80	%	adjustable
p_W	electricity	0.062	\$/kW h	adjustable
p_{O_2}, p_{Ar}, p_{LIN}	product values	0.055, 1.20, 0.030	\$/m^3	adjustable

Every run opens on a genuine steady state. Two of the eighteen states are pure integrators, so the plant has no steady state unless the rich liquid valve and the liquid oxygen draw are chosen to close their balances. The initializer therefore runs the real derivative function forward at a fixed 0.4 min step for exactly 9000 iterations, with each of those two valves driven by a feedforward term asking for the flow that closes its own balance plus a small level trim, and it returns the state and the two positions. At the design controls it returns 53.29 % on the rich liquid valve and 52.66 % on the liquid oxygen valve, which is the line-up every quoted result starts from. Feedforward is used rather than a proportional level controller because any gain slow enough to be stable took longer to line out than the sidearm did, and any gain fast enough pushed the stripping section through its pinch.

The engine of Section 1.6.1 advances the state vector on a fixed 0.05 min (3 s) step under classical fourth-order Runge-Kutta, the five controllers execute on a 0.25 min (15 s) scan and four samples of every simulated minute reach the historian. The noise generator is seeded at 20260914, and the fastest speed is 60 minutes of plant time per second.

6.3.9 Assumptions

Detail is concentrated where the coupling and the inventories live and taken out everywhere else. Every simplification listed here has an operating region in which it stops being true.

Eighteen states stand in for several hundred. Each column section is a lump of theoretical stages marched along its own operating line, so there is no tray hydraulics, no flooding or weeping and no temperature profile of the kind an operator reads. A theoretical stage here is worth roughly two real trays, and the stage counts were calibrated against the design purities.

Relative volatility is corrected for pressure and not for composition. The real volatility of argon against oxygen runs from about 1.08 to 1.13 with temperature, and taking 1.105 throughout gets the difficulty of the separation right and the exact stage requirement wrong.

Argon is a tracer with two empirical slip coefficients. The exponents in (6.15) were calibrated to put 3778 ppm of argon in the liquid oxygen at design and to give the reflux split the sensitivity a real unit shows. The sidearm is one steady-state relation with a hold-up lag and its condenser blanketing is a fitted quartic sigmoid, so the shape of (6.16) is an argument and the 320 ppm threshold is a calibrated number.

The compressor is isothermal work at a constant efficiency. There is no surge line, no stage matching, no intercooler train and no real turndown limit, and the guide vane characteristic is linear in volume.

Refrigeration is not modeled. The expansion turbine appears only as air diverted from the high-pressure column and the cold-box heat leak is a slider, whereas on a real unit those two and the liquid production are one balance, and that balance is the most important slow constraint on the plant.

The arithmetic is too repeatable. One frozen step, one frozen scan and a bubble-point solve with no early exit hand back the identical run everywhere, whereas a real control system delivers scan jitter and an execution order that moves. The heavy key balance closes to machine precision except at the sidearm condenser vent, which carries under 0.002 mol/s of nitrogen at design.

6.4 Coupling and Interaction Analysis

6.4.1 A Valve on One Column Is a Valve on the Other

The waste nitrogen valve sits on top of the low-pressure column and touches no part of the high-pressure column, no compressor setting and no air rate. Pinching it from 100 to 78 % in scenario 2, with every handle in MANUAL, raises the low-pressure column from 1.400 to 1.626 bar by the time the plant has settled at 220 min, so the oxygen in its sump boils at 95.01 K instead of 93.47 K. Neither the duty nor the conductance has changed, so (6.1) still demands the same 1.894 K and the nitrogen at the top of the high-pressure column must condense at 96.91 K instead of 95.36 K, which it will only do at 6.303 bar. The ratio is the useful number, because the high-pressure column rises 3.11 bar for every bar of low-pressure column, the two bubble point curves having different slopes with a

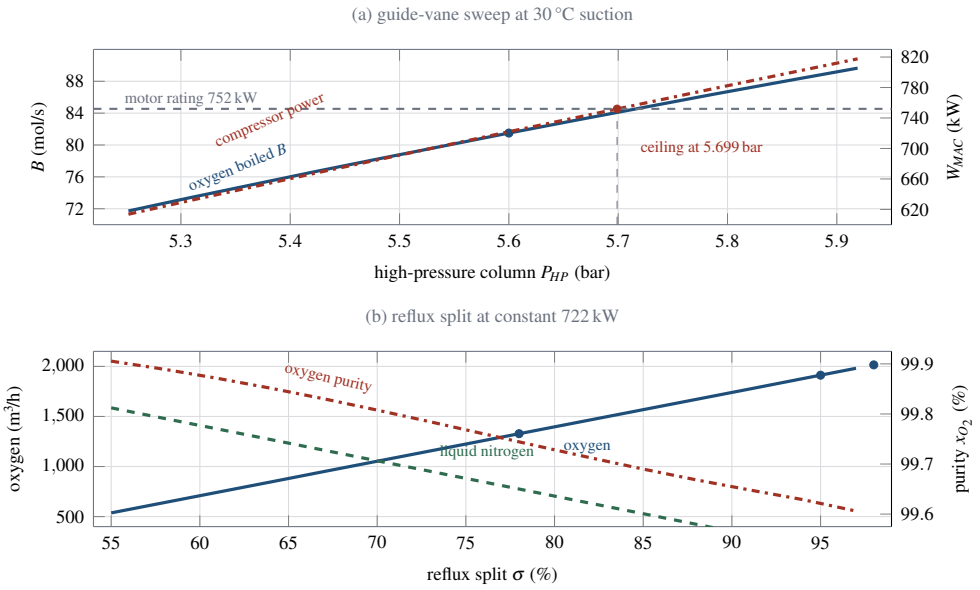

Figure 6.2 Two steady-state sweeps of the shipped model. In (a) the guide vane runs from 88 to 110 % and both low-pressure boil-up (solid, left axis) and compressor power (dash-dotted, right axis) are plotted against the pressure the loop holds, which crosses the motor rating (thin dashed line) at 5.699 bar. In (b) the reflux split runs from 55 to 97 %, with oxygen production solid, liquid nitrogen dashed and oxygen purity dash-dotted on the right axis, while the compressor power varies by less than half a kilowatt.

fixed driving force between them, so one tenth of a bar of low-pressure column is worth 20.2 kW continuously. The low-pressure column reaches 63 % of its move in 13 min and the high-pressure column 71 % of its own in 20 min, settling in about an hour, held back by nothing but the two vapor volumes of (6.4).

> **TRY IT IN THE APP** **2 · An LP-Column Valve Raises the HP-Column Pressure (MANUAL)**
>
> Load scenario 2 and press Run. At 40 min the waste nitrogen valve on top of the low-pressure column is pinched to 78 %. Watch the high-pressure column, which nobody touches: it settles at 6.303 bar instead of 5.600 bar, the gap across the condenser-reboiler stays at 1.894 K, and the compressor climbs to 767.5 kW, past its 752 kW rating.

Pause. Try Exercise 6.1 before reading on.

6.4.2 The Pressure Setpoint Is a Duty Setpoint With a Ceiling

Sweeping the guide vane at steady state produces Figure 6.2(a), which is the characteristic the pressure loop rides on. Boil-up in the low-pressure column is very nearly linear in the pressure of a column it does not enter, at 26.9 mol/s per bar over the range swept, and compressor power is linear at 306 kW per bar. Neither of those curves belongs to one unit. The motor rating of 752 kW is reached at 5.699 bar and 84.10 mol/s of boil-up, and that intersection is the ceiling on the duty setpoint. Asking for 5.780 bar instead needs the vane at 105.7 % and 776.3 kW, which is 54 kW on top of 722 and past the rating.

Dropping the setpoint instead is instructive because the two halves of the answer arrive on different clocks. In scenario 3 a move from 5.600 to 5.420 bar on the pressure and level loops brings the transmitter within 0.01 bar of the new setpoint 8 min after the step, undershoots to 5.385 bar and

stays within 0.01 bar from 24 min, and it takes the driving force from 1.894 to 1.780 K, the nitrogen condensed from 109.2 to 102.6 mol/s and the oxygen boiled from 81.5 to 76.6 mol/s. Less liquid arriving in the sump then lowers the level, the level controller closes the product valve, and the oxygen walks from 1912 to 1795 m^3/h five hours after the step, so the pressure loop settles in minutes and the plant in hours for a saving of 55 kW against 6.1 % of the production.

6.4.3 One Knob, Three Answers, at Constant Power

The reflux split decides how much of the liquid nitrogen drawn off the high-pressure column goes back as reflux and how much is sold, and Figure 6.2(b) is the trade it buys. At 78 % the plant makes 1328 m^3/h of oxygen at 99.744 % with 775 m^3/h of liquid nitrogen and 63.1 % recovery, at 95 % it makes 1912 at 99.622 % with 176 and 90.8 %, and at 98 % it makes 2015 at 99.594 % with 70 and 95.7 %, for 721.9, 722.0 and 722.0 kW. Scenario 5 makes the same two moves on the pressure and level loops, and its production reaches those settled values only through the sump level loop, undershooting to 1237 m^3/h six hours after the cut to 78 % and overshooting to 2126 m^3/h ten hours after the move to 98 %. The air, the duty and the driving force did not change, so the split is not a purity handle with a production side effect. It is a decision about how to divide a fixed amount of separation between two products and one specification.

6.4.4 Ranking the Handles, and What the Array Says

At steady state with the sump level held, three of the four manual handles move the oxygen purity, and only one of them was supposed to. Measured as central differences of one percent at the design line-up, the crude argon draw valve moves it by 378 ppm per percent, the guide vane by −174 ppm per percent and the reflux split by −70 ppm per percent. The argon valve wins because the nitrogen left in the liquid oxygen is 14 ppb, so the only impurity that matters is the 3778 ppm of argon the sidearm did not take out and the argon valve is the only handle that acts on it directly. Taking the two purities as controlled variables and the split and the argon draw as manipulated variables gives the steady-state gain matrix of Equation (6.22).

$$K = \begin{bmatrix} -0.0070 & 0.0378 \\ -0.1151 & -0.5536 \end{bmatrix} \%/\%, \qquad \lambda_{11} = 0.471 \tag{6.22}$$

where the rows are oxygen purity and crude argon purity and the columns are the reflux split and the argon valve. The relative gain array returns 0.471 for the diagonal pairing the application implements and 0.529 for the off-diagonal one, so it marginally prefers putting oxygen purity on the argon valve. The application pairs the other way anyway, for availability rather than sensitivity, because the guide vane belongs to the duty loop and the argon valve belongs to the argon specification, which leaves the reflux split as the only free handle. A λ_{11} of 0.471 is close enough to 0.5 that neither pairing decouples, and that is why both purity loops are tuned at 200 minutes of reset rather than made fast. The array is a steady-state measure, and Volume 6, Chapter 2 shows its other limit, a relative gain at the value for no interaction on two loops that still interact.

🕮 **Pause.** Try Exercise 6.3 before reading on.

6.5 Implemented Control Structure

6.5.1 Five Loops, Four Architectures and One Override

Five loops are implemented and a control architecture switch selects how many are in service. In MANUAL the four handles are whatever the operator typed and only the high-pressure sump level loop runs. The inner-loops setting closes column pressure on the guide vanes and sump level on the product valve, and leaves both purities to whatever the day gives, which is how a great many air separation units are actually run. Full cascade adds both purity loops and enables a compressor power override that caps the guide vane. The fourth setting hands the guide vanes, the product valve and the

reflux split to one model predictive controller, described in Section 6.5.5, while the argon loop, the override and the high-pressure sump loop run as they do in full cascade.

6.5.2 The Discrete Algorithm

One implementation serves all five loops. It is the velocity form of Section 1.6.2, and one execution is Equation (6.23), in which k counts scans and $\Delta t_c = 0.25\,\text{min}$.

$$e_k = a\,(\text{SP}_k - \text{PV}_k)\,, \qquad P_k = K_c e_k \tag{6.23a}$$

$$D_k = -K_c \tau_D\, a \left(\frac{\text{dPV}^f}{\text{d}t} \right)_k$$

$$s_k = s_{k-1} + \Delta P_k + \Delta D_k \tag{6.23b}$$

$$c_k = u_{bias} + s_k + I_{k-1} \tag{6.23c}$$

$$u_k = \min(u_{max},\ \max(u_{min},\ c_k)) \tag{6.23d}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t} \right] \Delta t_c \tag{6.23e}$$

Here $a = \pm 1$ is the action, $\Delta P_k = P_k - P_{k-1}$ and $\Delta D_k = D_k - D_{k-1}$ are the proportional and derivative increments, s_k accumulates them without limit, c_k is the demanded output and u_k is what reaches the actuator lag of (6.20). In the symbols of Equations (1.14) to (1.17), s_k is S_k, c_k is v_k and u_{bias} is u_b. The last term of (6.23e) is the back-calculation of Chapter 1, with tracking time $T_t = \tau_I$ on every loop [20], so the integral alone absorbs the clipping. Derivative action is switched off everywhere, since $\tau_D = 0$ on all five loops makes D_k vanish, and none of the loops carries a rate limit, the actuator lags already supplying one. MANUAL keeps every loop executing with $I_k = u_{man} - u_{bias} - s_k$, the bumpless transfer of Section 1.6.2. Output bias is zero throughout, and every initialization, whether a scenario load or a Reset, seeds each integral from the actuator position the initializer returns for the settings in force, so the opening execution hands back the line-up it started from.

The actions follow Chapter 1's convention, in which $a = +1$ marks a reverse-acting loop on which a rising measurement lowers the output and $a = -1$ a direct-acting one. Table 6.4 lists the five loops with their pairings, default tunings, actions and output limits. Only the pressure loop is reverse acting. The level, purity and sump loops are direct acting, because raising each of their outputs lowers its own measurement.

Table 6.4 Default Tuning of the Five Loops. Action +1 Is Reverse Acting.

CV	MV	K_c	τ_I (min)	a	Limits
P_{HP}	u_{GV}	40	3	+1	45–110 %
L	u_{lox}	3	300	−1	0–100 %
x_{O_2}	σ	30	200	−1	55–98 %
x_{Ar}	u_{arg}	2.5	200	−1	5–80 %
L_{HP}	u_{rich}	2.0	30	−1	5–100 %

Three of the five settings in Table 6.4 are worth explaining. The pressure loop carries 40 % of vane per bar with 3 min of reset against a process gain of 0.0306 bar per percent and about 11 min of coupled column dynamics, which makes it the fastest thing on the plant. The level loop carries 3 % of valve per percent with 300 min of reset, which looks slack and is the correct answer, because the process gain is 0.001 93 %/min per percent of valve and a loop tight enough to hold the level throws the whole flow imbalance into the product line [4]. That is averaging level control, and Volume 2, Chapter 5 tunes a loop on another integrating level from a measured slope. The oxygen purity loop carries 30 % of split per percent with 200 min of reset, slow because it faces 111 min of sump inventory behind a handle worth 70 ppm per percent, so the 6 min analyzer is not what makes it

slow. The argon loop faces 278 min of sidearm hold-up, so anything responsive on that trend is far too fast.

6.5.3 The Compressor Power Override

The power limit is enforced on the actuator rather than as a controlled variable of its own, as shown in Equation (6.24).

$$u_{max} = \mathrm{clamp}\big(u_{max} + K_W\,(W_{max} - \hat{W}_{MAC})\Delta t_c,\, 45,\, 110\big) \tag{6.24}$$

where $K_W = 0.06\,\%$ per kW min and $\hat{W}_{MAC}$ is the filtered power measurement. Equation (6.24) is an integral-only constraint controller whose state is a valve position, so it is seeded at the top of its range and falls only when the limit is real. Written as a high limit on the pressure loop's output rather than as a second controller into a low selector, the pressure loop's own back-calculation in (6.23e) does the anti-windup, so when the limit lets go the loop comes off it without a bump. The cap is deliberately slow, because the guide vane is also the boil-up handle for a column with twenty cubic meters of inventory underneath it, and snatching it back at the speed a power trip would prefer costs an hour of recovery on the argon. A protective override that must act in seconds, on a compressor near surge, is the subject of Volume 4, Chapter 2.

6.5.4 Instruments

The oxygen analyzer is a first-order filter of 6.0 min standing for a paramagnetic analyzer on a sample line and the crude argon chromatograph is a filter of 9.0 min, while the pressure transmitter and the power transducer are 0.10 min and the two level transmitters are 0.50 min. Optional colored noise is zero by default and is scaled by 1, 8 and 12 on the oxygen analyzer, the argon chromatograph and the level transmitters.

6.5.5 Model Predictive Control

The fourth architecture replaces the pressure, level and oxygen purity loops by a plan of the kind set out in Section 1.6.3. Every $T_s = 1$ min, four controller scans, a quadratic program chooses the guide vane u_{GV}, the product valve u_{lox} and the reflux split σ over $N = 120$ steps in twelve move blocks of $\{1,1,1,1,2,2,3,4,6,9,15,75\}$ steps, each a move from the input applied now, and the opening move is held for T_s [21, 22]. The argon loop keeps its valve, whose applied position enters the plan as a measured input held over the horizon, and the cap of (6.24) stays the vane's upper bound.

The controller's model is the plant of Section 6.3 at the declared parameter values and nominal disturbances (30 °C suction, $c_f = 1$, no heat leak), linearized by central differences at the design line-up (100 % vane, 95 % split, 40 % argon valve). The integral of the high-pressure sump loop at its shipped tuning, which runs under every architecture, and the 6.0 min analyzer lag join the eighteen plant states, and an exact zero-order hold over T_s gives Equation (6.25).

$$\begin{aligned} \tilde{x}_{k+1} &= A_d\,\tilde{x}_k + B_d\,(\tilde{u}_k + d_k) + E_d\,\tilde{u}_{arg,k}, \qquad d_{k+1} = d_k \\ \tilde{y}_k &= C\,\tilde{x}_k \end{aligned} \tag{6.25}$$

Here a tilde marks a deviation from the design point, $\tilde{x}$ holds the twenty states, $\tilde{u}$ the three decided inputs, d one unmeasured disturbance per decided input, $\tilde{u}_{arg}$ the applied argon valve and $\tilde{y}$ the outputs P_{HP}, L, x_{O_2}, W_{MAC} and the analyzer reading. Step-response fits were not used, because one percent of vane raises the purity 0.000 78 % at 15 min and lowers it 0.0191 % by 960 min, while one percent of split delivers 42, 65 and 94 percent of its −0.0096 % answer by 60, 120 and 480 min. Against the engine's mean response to ±1 percent of vane, ±1.5 of product valve, ±0.5 of split and ±2 of argon valve from scenario 1 without noise, the model errs by at most 0.73 % of the largest response between 15 and 960 min, the worst input being the argon valve, and by 0.20 percent for the split.

A steady-state Kalman filter estimates $\tilde{x}$ and d once a minute under every architecture from the readings the loops use, P_{HP}, L and the analyzer, with assumed noise of 0.002 bar, 0.05 % and 0.002 %. Each disturbance carries process noise $\rho\, r_m^2 / G_m^2$, with $\rho = 100$, r_m the noise of its measurement and G_m the effect of one percent of its input on it (0.0306 bar, 0.116 % of level in an hour, 0.0097 % of purity), which makes the prediction offset free [23]. Volume 9, Chapter 4 builds the same kind of estimator around an unmeasured fuel bias. The compressor power is corrected by its measured offset. A constant output bias instead of the filter, tried on a diagnostic copy outside the application, loses scenario 7 by an oxygen IAE of 21.25 against 16.38 % · min and 22.0 against 0.1 min above the motor rating.

The plan minimizes Equation (6.26).

$$\begin{aligned} J = \sum_{k=1}^{N} \Big[& q_P \Big(\frac{P_{HP,k} - P^{SP}}{0.05} \Big)^2 + q_L \Big(\frac{L_k - L^{SP}}{3} \Big)^2 \\ & + q_O \Big(\frac{x_{O_2,k} - x_{O_2}^{SP}}{0.05} \Big)^2 \Big] + \lambda \sum_{j=1}^{12} \sum_{m=1}^{3} \Big(\frac{\Delta u_{m,j}}{2} \Big)^2 \\ & + \sum_{i \in \{O,L,W\}} \left(\rho_s s_i^2 + \mu_s s_i \right) \end{aligned} \tag{6.26}$$

Here pressure is in bar, level, purity and inputs in percent, $q_P = 1$, $q_L = q_O = 0.1$, $\lambda = 3$, $\Delta u_{m,j}$ is the move of input m in block j, and the slacks of the oxygen, level and power bands, in units of 0.05 %, 3 % and 10 kW, are priced at $\rho_s = 10^4$ and $\mu_s = 10^3$, so a reachable band costs nothing. The tracked purity is the model's true x_{O_2}, and its low weight makes the target a preference while the specification is a limit, as shown in Equation (6.27).

$$\begin{aligned} & 45 \le u_{GV} \le \min(110, u_{max}), \qquad 0 \le u_{lox} \le 100 \\ & 55 \le \sigma \le 98, \qquad x_{O_2,k} \ge 100\, x_{spec} + 0.01 - s_O \\ & L_{lo} + 4 - s_L \le L_k \le L_{hi} - 4 + s_L \\ & W_{MAC,k} \le W_{max} - 1 + s_W, \qquad s_O,\ s_L,\ s_W \ge 0 \end{aligned} \tag{6.27}$$

where u_{max} is the cap of (6.24), $x_{spec} = 0.995$, $L_{lo} = 25\,\%$ and $L_{hi} = 80\,\%$ are the level alarms and $W_{max} = 752$ kW, all read from the run. The input rows are hard, the output rows soft and judged at every second step, and a band already broken when a plan starts is regained along a first-order lag of 60 min (5 min for the power). The back-offs of 0.01 % and 1 kW keep a plan that rides a limit from crossing it on its estimate's error, and without the kilowatt a diagnostic copy counts 754.6 min of scenario 7 above the rating. Energy carries no price, because with the pressure on its setpoint the vane has no steady-state freedom and a price on kilowatts could buy only a pressure offset.

A rejected plan, not optimal, not finite or outside its bounds, hands the three handles without a bump to the full cascade loops, which track the applied commands while the plan acts. As Section 1.6.3 describes, the model never reads the run's parameters or the plant's states, so a learner's change to the plant reaches it only as a disturbance for the filter. A dual active-set method [24] solves a plan in a mean of 2.06 ms in a browser, 0.21 percent of one core at real time, and in 7.88 ms at a horizon of 480 min.

6.6 Results

The application carries fourteen scenarios arranged as a laboratory sequence, each of which loads a starting mode, a control architecture, eleven tuning constants, a set of disturbance positions and its own timed events. In order they are four handles on one plant (1), an LP-column valve that raises the HP-column pressure (2), the pressure loop as the duty loop (3), twenty cubic meters of memory in the sump (4), purity, recovery and liquid nitrogen on one knob (5), full cascade held steady (6), a

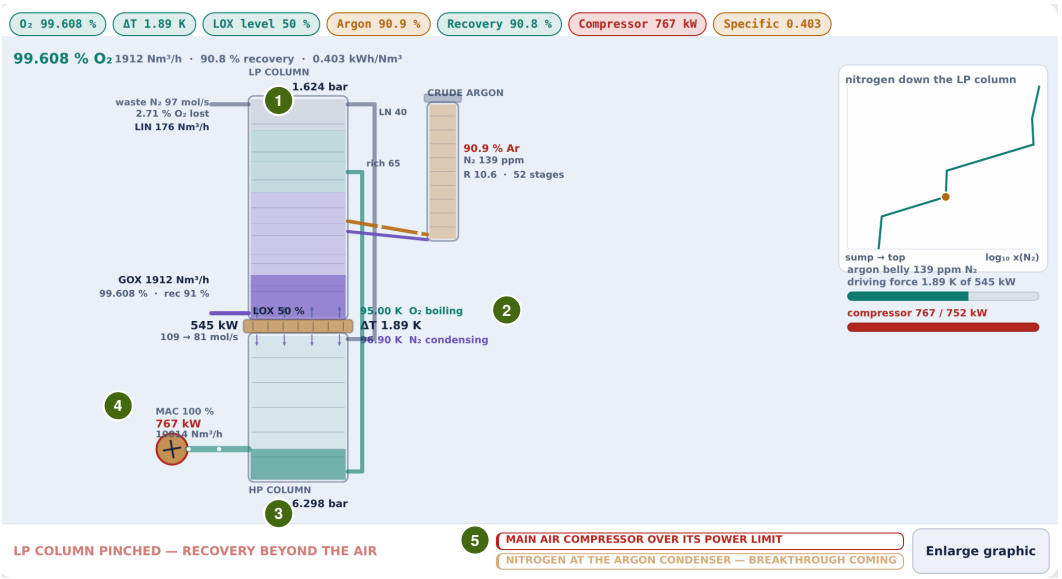

Figure 6.3 Scenario 2 at $t = 120$ min, eighty minutes after the waste nitrogen valve on top of the low-pressure column was pinched to 78 %, in MANUAL with only the high-pressure sump loop in service. (1) The low-pressure column sits at 1.624 bar. (2) Both saturation temperatures at the two ends of the condenser-reboiler are translated upward, 95.00 K boiling against 96.90 K condensing, with the gap between them still 1.89 K. (3) The high-pressure column has arrived at 6.298 bar. (4) The guide vane has not moved from 100 %, and the compressor takes 767 kW against its 752 kW rating for the same 1912 m³/h of oxygen. (5) The limit banner names the overload.

warm afternoon and the power limit (7), the shared exchanger fouling (8), the argon lost and found (9), nitrogen at the argon condenser (10), the cold box warming (11), the last two percent of recovery (12), a purity under full cascade and under model predictive control (13) and the graded challenge, a shift on a badly lined-up plant (14). Each is credited by a completion rule of the kind described in Section 1.6.4, and the challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt. Integrated absolute error, overshoot, settling time, actuator travel and saturated fraction appear on the scorecard for each loop, and the integrated error and the travel keep accruing after a transient has ended. Beside the loops, the scorecard carries the oxygen and argon purities and recoveries, the driving force, the specific power, the compressor power, the minutes off specification and the margin. One hundred and forty-five columns leave the historian, among them the two saturation temperatures, the driving force, the duty, the argon in the liquid oxygen, the belly nitrogen and the running specific power. An integrated error from a plant whose slowest loop has a 278 min time constant is meaningless without the window it was taken over, so every comparison below is made between two settled conditions, at a named instant, or as an elapsed time to a named fraction of the total move.

6.6.1 Open-Loop Moves and Settled Comparisons

Scenario 1 makes the pairing argument with four open-loop moves, each held for an hour from the design line-up with only the high-pressure sump loop in service, and each measured from the controller scan at which it took effect. Within the hour five percent more guide vane is 0.157 bar in the high-pressure column, 47 kW at the compressor and 4.06 mol/s more boil-up in a column the vane does not enter, but only 25 ppm off the oxygen purity, which keeps falling to 917 ppm off at 12 h and 968 ppm at 24 h. Six points more liquid oxygen valve is 2117 m³/h one minute later and 2130 an hour

later, after which the level falls at 0.69 % per hour and the purity creeps up only 173 ppm in 12 h, so what that valve really moves is inventory. With the product valve fixed, one percent more reflux split leaves the production at 1912 m³/h and takes 41 ppm off the purity in the hour and 103 ppm in a day, against the 70 ppm of the level-held steady state. One percent more argon valve moves the oxygen purity only 23 ppm in the hour but 370 ppm by 12 h, and takes 0.52 points off the crude argon in 12 h and 0.60 in a day. The strongest handle on the purity is therefore also the slowest to show it. Table 6.5 collects the closed-loop results.

Table 6.5 Measured Results. Each Block Compares the Condition Before a Move With the Settled Condition After It, Unless the Block Names an Instant.

Quantity	Before	After	#
Waste valve 100 to 78 %, MANUAL			
P_{LP} (bar)	1.400	1.626	2
P_{HP} (bar)	5.600	6.303	2
ΔT (K)	1.894	1.894	2
W_{MAC} (kW)	722.0	767.5	2
Exchanger conductance $1.00 \rightarrow 0.90 \rightarrow 0.78$*, inner loops*			
ΔT (K)	1.894	2.253	8
Guide vane (%)	100.0	92.8	8
Oxygen (m³/h)	1912	1762	8
W_{MAC} (kW)	722	670	8
w_{spec} (kW h/m³)	0.379	0.381	8
Then the setpoint raised to 5.780 bar			
ΔT (K)	2.253	2.386	8
Guide vane (%)	92.8	98.3	8
Oxygen (m³/h)	1762	1874	8
W_{MAC} (kW)	670	722	8
Heat leak 0 to 45 kW, full cascade, at 1000 min			
Oxygen purity (%)	99.622	99.763	11
Oxygen (m³/h)	1912	1295	11
Recovery (%)	90.8	64.9	11
w_{spec} (kW h/m³)	0.379	0.531	11
Argon valve 40 to 58 %, inner loops, at the argon minimum			
Sidearm reflux ratio	10.58	6.72	9
Crude argon (%)	96.02	68.8	9
Oxygen purity (%)	99.622	99.931	9
Argon in the product (ppm)	3778	690	9

6.6.2 A Failure That Holds Every Setpoint

Scenario 8 fouls the shared exchanger to 0.90 and then 0.78 with the pressure and level loops in service. The driving force widens from 1.894 to 2.253 K because the pressure controller allows it, but $Q = UA\,c_f\,\Delta T$ and the conductance has fallen further than the driving force has risen, so the condenser takes less nitrogen away and the controller holds 5.600 bar by *closing* the vane from 100.0 to 92.8 %. Oxygen production falls from 1912 to 1762 m³/h, 7.8 % of the plant, the compressor power goes *down* from 722 to 670 kW, specific power moves only from 0.379 to 0.381 kW h/m³ and the purity rises to 99.75 % by 1400 min because there is less to separate. Every controlled variable is on setpoint, nothing alarms, and the only trend that says the plant has become eight percent smaller is an exchanger temperature difference of two kelvin.

🕮 **Pause.** Try Exercise 6.2 before reading on.

Raising the pressure setpoint to 5.780 bar is what an operator watching that trend would have done hours earlier, and it widens the driving force to 2.386 K, brings the vane back to 98.3 % and the production to 1874 m³/h at 722 kW. That is where the bill arrives, because the same 52 kilowatts the fouling saved now buy back production that used to be free, which at 0.062 $/kWh is about twenty-eight thousand dollars a year.

> **TRY IT IN THE APP** **8 · The Shared Exchanger Fouls, and the Plant Quietly Unloads**
>
> Load scenario 8 and press Run. The exchanger fouls at 60 and 600 min. Watch the driving force widen from 1.894 to 2.253 K, the guide vane close from 100.0 to 92.8 % and the oxygen fall from 1912 to 1762 m^3/h with nothing alarming. At 1400 min the pressure setpoint goes to 5.780 bar and buys the production back to 1874 m^3/h.

6.6.3 Purity Is Not Performance

Scenario 11 puts 25 kW and then 45 kW of heat leak into the low-pressure sump in full cascade. Through (6.1) the 45 kW boils 6.7 mol/s more liquid oxygen and condenses nothing anywhere, which is the last term of (6.3) acting alone, so those moles leave with the waste nitrogen instead of as product. By 1000 min production has fallen from 1912 to 1295 m^3/h, recovery from 90.8 to 64.9 % and specific power has risen from 0.379 to 0.531 kW h/m^3, while the purity analyzer has moved from 99.622 to 99.763 % because what is left in the sump is cleaner. The purity loop cannot pull it back, because it has held the reflux split against its 98 % stop since 218 min, and production recovers only to 1373 m^3/h at 68.8 % recovery by 2880 min, once the level loop has refilled the sump. The two measurements that would have caught it within minutes, recovery and specific power, are not controlled variables.

Scenario 12 is the other blind spot. At a 98 % reflux split the plant already recovers 95.7 % of the oxygen in the air, and cutting the expansion turbine from 12 to 3 % sends more air to the high-pressure column and pushes for the rest. It does not arrive, and Equation (6.14) is the reason, because a column cannot send more of a component out of the bottom than came in at the feed, so when W_4 crosses zero the section above the crude feed runs out of heavy key to send upward. Its profile pinches, nitrogen floods down to the argon belly to a peak of 17 700 ppm, the crude argon falls from 95.3 to 38.7 % and the oxygen purity follows it to 98.81 %, bottoming 3.4 h behind the argon and 4.8 h *after* the turbine was put back. Both failures share a shape worth carrying away, because on a plant with several products any single quality measurement can be improved by making less of the product.

> **DID YOU KNOW?** **A Warmer Cold Box, a Purer Product**
>
> From the simulator: with 45 kW of heat leak under full cascade, the oxygen purity at 1000 min has risen from 99.622 to 99.763 %, while the production has fallen from 1912 to 1295 m^3/h and the recovery from 90.8 to 64.9 % (Table 6.5). The number the product is sold on improves as the product disappears.

6.6.4 Inventory and the Argon

Scenario 4 starts the liquid oxygen sump at 42 % instead of 50. The level controller closes the product valve by twenty-four points at once, from 52.7 to 28.7 %, and the oxygen production drops from 1912 to 1042 m^3/h with it. The level then returns on the inventory's own clock, reaching 44.6 % at one hour, 46.7 % at two, setpoint at four and a quarter, a peak of 51.8 % at 8.6 h and within 0.2 points of setpoint from 18 h. That is averaging level control working as designed, because the alternative throws the whole flow imbalance into a product line whose customer is a steel plant.

Scenario 9 opens the crude argon draw from 40 to 58 %, and the crude argon reaches a minimum of 68.8 % five hours later, then climbs back to 75.6 % by the time the valve closes at 600 min, close to the 75.7 % it would line out at. Argon recovery reads 120.5 % three minutes after the draw opens and 119 % an hour later, which is not a fault, because the sidearm is taking argon out faster than the air brings it in and the difference is the belly emptying. Coming back is slower than going, at 84.7 % three hours after the valve is closed, 93.5 % after nine and 95.5 % after fourteen, still short of where it started. Scenario 10 blankets the sidearm condenser instead by pinching the waste valve with a heavy belly draw, which takes the belly nitrogen from 267 to a peak of 3375 ppm 23 min after the pinch and settles it at 344 ppm over the 320 ppm threshold, collapsing the sidearm reflux ratio from 9.85 to 0.001 and the crude argon to a minimum of 79.9 % 67 min after the pinch.

6.6.5 What the Constraint Controller Decides

Scenario 7 warms the compressor suction from 30 to 42 °C. Guide vanes pass a volume, so the same opening passes four percent fewer moles, and the isothermal work of (6.6) is four percent higher because it is proportional to suction temperature, so holding 5.600 bar takes the vane from 100.0 to 104.0 % and the power from 722 to 751 kW, already touching the rating. Raising the pressure setpoint to 5.720 bar for more oxygen is the interesting half. On inner loops the pressure controller takes what it needs, driving the vane to 108.0 % and the compressor to 788 kW, thirty-six over the rating and continuously, while in full cascade Equation (6.24) caps the vane at 104.1 % and the trend sits at 5.605 bar against 5.720 bar for the rest of the afternoon. The cap moves at the slow rate the vane allows, so the request itself still takes the compressor to 786 kW one minute later and keeps it more than half a kilowatt above the rating for 9 min before the vane is held. A constraint controller does not solve the constraint, it decides which variable gives way, and here it is the duty and therefore the production, because the alternative is a tripped motor.

Cooling the suction to 28 °C in the evening lets the same 752 kW buy more moles, the cap relaxes to 103.0 % and the setpoint is met at 5.715 bar. What arrives is duty rather than oxygen, because boil-up rises from 81.5 to 84.5 mol/s and the purity controller spends every mole of it on liquid nitrogen, taking the reflux split from 95.0 to 92.7 % and the liquid nitrogen from 176 to 267 m³/h while the oxygen settles at 1898 m³/h, below the 1912 it started at. More separation is not more oxygen unless somebody decides it is.

6.6.6 The Graded Challenge

Scenario 14 hands over a plant that came off the night shift with the crude argon valve at 33 %, weak tunings loaded and the exchanger already at 0.88, and then warms the suction to 39 °C at 90 min, fouls the exchanger to 0.74 at 300 min and cools the suction again at 520 min, in a run judged over twenty hours. Every attempt starts from this prescribed handover state. Learner valve settings made before Run are actuator requests applied when control begins, rather than inputs to a new steady starting plant. Reset keeps those requests and restores the same handover state. The case passes when the oxygen stays within ±0.06 % of 99.62 % and the high-pressure column within ±0.05 bar of 5.60 bar over the last five hours, with the crude argon at or above its 93 % specification and the compressor at or below 780 kW throughout. At handover the oxygen reads 99.310 %, off the 99.5 % sales specification, while the crude argon analyzer reads a comfortable 98.32 % and the argon recovery is 70.3 %. Almost all of it is the argon valve, because at 33 % the sidearm draws 0.83 mol/s against 1.14 mol/s arriving, so 30 % of the argon in the air escapes the crude product, half of it into the liquid oxygen at 6902 ppm. Left as shipped, the plant still has the oxygen 0.191 % short of its setpoint during the last five hours and fails, against the 0.06 % allowed, while the other three terms pass, with the high-pressure column within 0.024 bar, the crude argon at least 98.13 % and a compressor peak of 727 kW.

The suggested configuration is full cascade with the tuned settings, the manual valves left at the handover positions. The argon loop opens the valve from 33 to 40 % within the hour and to a peak of 43.9 %, the oxygen is back on specification at 99.50 % after 159 min, and the argon recovery is 83.8 % at the end of the twenty hours with the crude argon never below 94.5 %. Holding 5.600 bar through a quarter of the exchanger’s conductance means closing the vane from 99.4 % two hours in to 90.5 %, so the power falls from 730 to 654 kW and never approaches the 780 kW the target forbids, and what the plant gives up instead is production, 1944 down to 1837 m³/h at the end of the run after a trough of 1544 m³/h at 400 min. The case passes, with the oxygen within 0.049 %, the high-pressure column within 0.000 03 bar, the crude argon at least 94.52 % and a compressor peak after the handover instant of 722 kW. Buying the duty back with the pressure setpoint is the wrong answer, because at 5.85 bar the compressor peaks at 786 kW and (6.24) is far too slow to catch it.

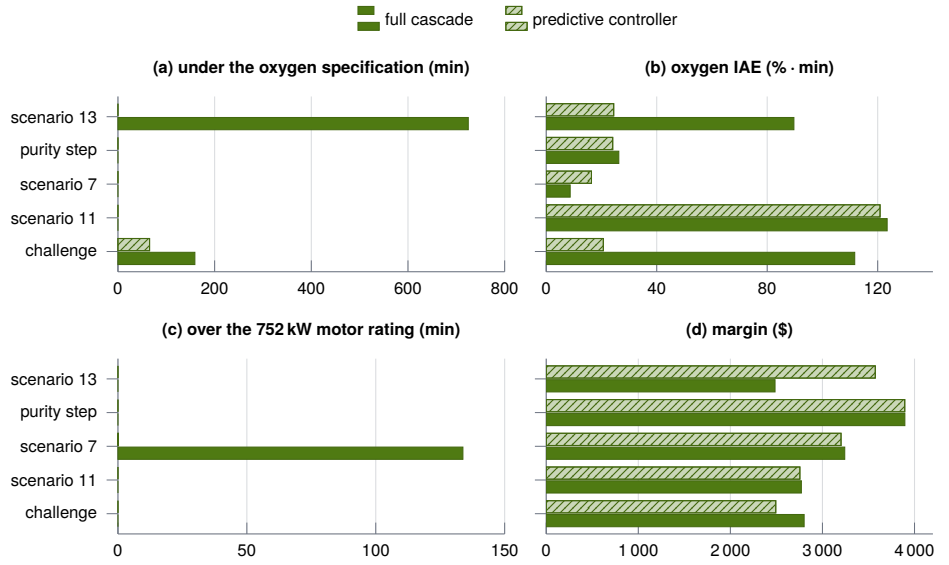

Figure 6.4 Four columns of Table 6.6 as bars, full cascade dark and the predictive controller hatched. Notice that the plan wins where time under the specification or over the motor rating is counted, in scenario 13, scenario 7 and the challenge, while the oxygen IAE of scenario 7 and the margins of the step and the heat leak barely move or go the other way.

6.6.7 Comparison With Model Predictive Control

Table 6.6 runs full cascade, with its shipped tuning and the suggested line-up on the challenge, against the predictive controller with the settings of scenario 13 on every case and only the architecture changed. Both face the same events, instruments and duration without noise, neither receives a preview, and the purity step is a scripted event on the line-up of scenario 6 because no shipped scenario contains one. Figure 6.4 draws four columns of the table as bars.

Scenario 13 lowers the oxygen target from 99.62 to 99.55 % at 60 min, five hundredths above the sales specification, and raises the crude argon target from 96 to 97.5 % at 480 min. Both structures answer the first step by driving the split to its 98 % stop. The argon request then closes the argon valve from 38.4 to 33.8 % within the hour and leaves more argon in the liquid oxygen. Under full cascade the purity loop, on its 200 min reset, brings the split off its stop only as the error accumulates, to 97.1 % at 600 min and 91.1 at 900 min, so the oxygen is under specification from 663.0 to 1387.5 min, at lowest 99.441 %, and 21 400 m^3 of it is vented. The plan lowers the split as the sag begins, to 94.3 % at 540 min, 90.1 at 600 min and 81.0 at 961 min, and the purity never falls to the specification, its lowest value being 99.545 %, as shown in Figure 6.5. The margin is $3574 against $2485, paid for with 25.2 and 25.5 percent of split and product valve travel against 14.7 and 16.6.

The oxygen floor never binds on this case. With the floor switched off every logged variable except the slack, the solver iterations and the solve time is identical, and with it on the largest oxygen slack is 4.3×10^{-19} percent, so the plan's tracking of the 99.55 % target holds the specification. The floor binds only on the challenge, which starts at 99.310 %, where it shortens the time under specification from 71.4 to 65.4 min.

The purity step and the heat leak are draws. On the step the oxygen IAE is 24.07 against 26.27 % · min and the margin $3894 against $3895, with a pressure IAE of 0.228 against 0.007 bar · min. On scenario 11 the oxygen IAE is 120.86 against 123.42, the level IAE halves from 1518 to 725 % · min and the pressure IAE rises from 0.535 to 1.790 bar · min.

Scenario 7 is a loss. The plan holds the motor rating, with a peak of 752.0 kW against 786.4

Table 6.6 Full Cascade Against Model Predictive Control

Structure	Under spec (min)	Lowest x_{O_2} (%)	IAE x_{O_2} (% min)	IAE P_{HP} (bar min)	IAE L (% min)	Travel u_{GV}, u_{lox}, σ (%)	Peak W_{MAC} (kW)	Over rating (min)	Margin ($)	Plans, fall-backs
Scenario 13 (1440 min)										
full cascade	724.7	99.441	89.66	0.024	1122	0.3, 16.6, 14.7	722.5	0.0	2485	—
predictive	0.0	99.545	24.43	0.204	218	0.3, 25.5, 25.2	722.5	0.0	3574	1440, 0
Purity step (1440 min)										
full cascade	0.0	99.544	26.27	0.007	319	0.1, 4.7, 3.9	722.0	0.0	3895	—
predictive	0.0	99.549	24.07	0.228	153	0.1, 6.5, 5.1	722.1	0.0	3894	1440, 0
Scenario 7 (1200 min)										
full cascade	0.0	99.591	8.67	60.269	90	17.9, 2.4, 2.4	786.4	133.8	3242	—
predictive	0.0	99.578	16.38	63.614	244	5.6, 21.8, 16.6	752.0	0.1	3202	1200, 0
Scenario 11 (1200 min)										
full cascade	0.0	99.622	123.42	0.535	1518	7.6, 17.9, 2.9	722.0	0.0	2771	—
predictive	0.0	99.621	120.86	1.790	725	6.7, 29.1, 3.0	722.2	0.0	2755	1200, 0
Challenge (1200 min), full cascade on the suggested line-up										
full cascade	159.1	99.310	111.68	2.038	1218	25.9, 21.5, 14.3	724.6	0.0	2801	—
predictive	65.4	99.310	20.67	1.743	286	25.4, 75.7, 71.6	731.7	0.0	2493	1200, 0

Errors, travel and the minutes of true purity under the 99.5 % specification accumulate from $t = 0$ as on the scorecard, and the challenge's lowest purity is its starting value. Over rating counts minutes above 752 kW. Margin is the value of on-specification oxygen, crude argon and liquid nitrogen at the prices of Table 6.3 less the electricity. The purity step lowers the oxygen target from 99.62 to 99.55 % at 60 min.

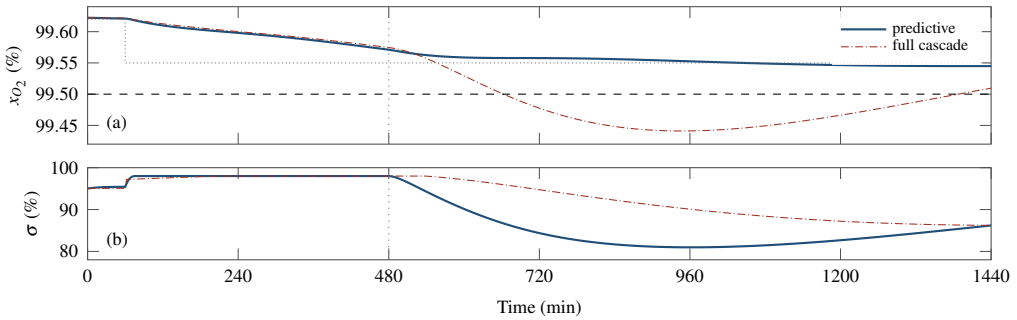

Figure 6.5 Scenario 13 under the predictive controller (thick solid line) and under full cascade (dash-dotted line). (a) True oxygen purity against the 99.5 % specification (dashed) and the target (dotted). (b) Reflux split. The vertical dotted line marks the argon request at 480 min.

and 0.1 min above it against 133.8, but its oxygen IAE is 16.38 against 8.67 % · min, its level IAE 244 against 90, its product valve and split travel 21.8 and 16.6 percent against 2.4 and 2.4, and its margin $3202 against $3242, because the 5.72 bar request is out of reach. Under full cascade the vane is within 0.05 percent of the override's cap on 3598 of 4801 logged rows and under the plan on none, since the plan's own ceiling 1 kW under the rating holds the compressor in 1152 of its 1200 decisions, and the quadratic objective spends the split, at its stop for 10.2 percent of the run, and the product valve on the hundredths of a bar they can still buy. A diagnostic copy outside the application that clips the pressure target to the present pressure whenever a plan is made within 0.5 kW of that ceiling reduces the oxygen IAE to 0.77 and the level IAE to 28, which confirms the mechanism. The application has no steady-state target calculation that would remove the loss.

The challenge is reported for information, because its scoring is unchanged. Both structures meet the target, full cascade with the suggested line-up holding the oxygen within 0.049 percent and the plan, with only the architecture switched, within 0.001 percent with a compressor peak of 732 kW.

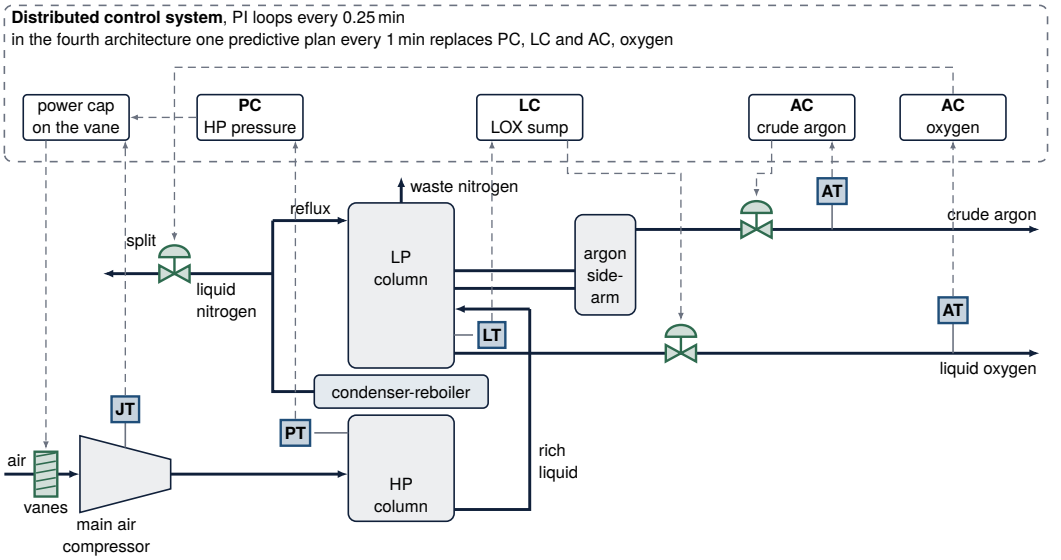

Figure 6.6 The loops of the air separation unit with their hardware. Lettered boxes are sensing devices: the pressure transmitter (PT) on the high-pressure column, the level transmitter (LT) on the low-pressure sump, the oxygen analyzer (AT) on the liquid oxygen, the chromatograph (AT) on the crude argon and the power transducer (JT) of the compressor motor. The vanes and the valve symbols are final elements, dashed arrows are signals. Notice that the pressure loop acts through the compressor's vanes, that the power cap limits the same vanes, and that the oxygen loop reaches across the plant to the split on the liquid nitrogen. The high-pressure sump loop on the rich liquid is left out. Not to scale, and not a piping diagram.

Unscored, the oxygen inherited at 99.310 % stays under specification for 159.1 min under full cascade and 65.4 min under the plan. No plan was rejected on any case of Table 6.6.

6.7 Technology in Practice

What Measures It

Figure 6.6 draws the five loops of Section 6.5 with their devices. The loop the chapter turns on reads the oxygen purity, and on a real unit that reading comes from an analyzer fed by a sample line. The product here is liquid, and a cryogenic liquid sample has to be taken so that it does not partly vaporize, which would misread heavy components such as hydrocarbons [19]. The pressure and the two levels are read by ordinary transmitters, and the level of the low-pressure sump matters for safety as well as for the product flow, because control of the reboiler's level is important in preventing dry boiling [19].

TECHNOLOGY Sensor: A Paramagnetic Oxygen Analyzer

A paramagnetic analyzer is selective because no other industrial gas is nearly as paramagnetic as oxygen, although nitric oxide interferes, and one that reads above the oxygen of ambient air is spanned with pure oxygen [25]. In the classic cell two nitrogen-filled glass spheres hang on a quartz thread in the field of a permanent magnet, and oxygen drawn into the field turns them out of equilibrium. Such cells answer in 3 to 10 s, limited by the gas exchange in the cell, need a small and stable sample flow and are sensitive to vibration [26]. A sample line that plugs leaves an analyzer reading old, stale sample with no way of knowing it [25]. The model's 6.0 min lag stands for the analyzer and its sample line together.

What Moves It

The pressure loop's final element is not a valve. It is the main air compressor's guide vane, which through the shared exchanger sets the boil-up of a column the air never enters (Section 6.4). The other handles are cryogenic valves on the product, the argon draw and the rich liquid, each lagging its command as in Equation (6.20), and the reflux split, which the model applies without a lag. The compressor power override writes a high limit on the vane rather than a second controller into a selector.

TECHNOLOGY Final Element: Compressor Inlet Guide Vanes

The main air compressor of an air separation unit delivers between 4 and 10 atm. On a centrifugal machine the capacity is typically set by variable inlet guide vanes on at least the first stage, which change it more efficiently than a throttling valve on the suction, and on an axial machine by rows of variable stator blades, whose actuating mechanism can bind from rust and dirt [27]. A control valve with excess friction moves in jerks, and integral action turns that into a cycle that detuning slows but does not remove [25]. Every dynamic compressor also needs an automatic anti-surge system [27]. The model's vane is linear in volume, moves between 45 and 110 % and meets no surge line.

Where the Controller Runs

An industry guideline divides the instrumented systems of an air separation unit into critical safety systems, operational safety systems and routine plant control. The critical systems may share components with the routine controls only if a failure of the routine system cannot compromise them [27]. Routine control runs in a distributed control system, in which a supervisory level sends set points to field controllers, loops may be nested, and cycle times range from milliseconds to minutes [28]. The app executes its five PI loops every 0.25 min and its predictive controller every 1 min, inside that range. The loops and the power override belong to the regulatory layer and the predictive controller to the layer above it. The shutdowns that protect a real reboiler belong to a safety layer the model does not have.

IN PRACTICE Keep the Reboiler Under Liquid

A bath-type main reboiler is kept submerged to the depth its designer specifies, or fully when that is unknown. Below it the liquid stops recirculating and evaporates completely in the passages, concentrating the heavier impurities by extremely high factors, and that accumulation cannot be detected by analyzing the sump. The level shown on the instruments is not the submergence, and level instruments and their alarms are verified at least once a year. A plant that makes little liquid oxygen purges 0.1 to 0.2 percent of its air as liquid, while one that makes more than 2 percent dilutes the contaminants by itself [19]. The chapter's plant sells all of its oxygen as liquid.

What Is Modernizing It

Three developments bear on this plant. Linear model predictive control is already the technology of the air separation industry, and nonlinear predictive control on a reduced compartmental model of a high-purity column has been solved within a 2 min sampling interval in closed-loop simulation [9].

Production scheduled on day-ahead electricity prices, with product stored in the cheap hours, cut the operating cost and the peak power demand of an air separation unit in simulation studies [17]. And computer-based control systems let such plants run unattended, with minimal staff or from a remote site, provided that every response an operator would give at an attended plant is designed into the controls [27].

6.8 Limitations

Nobody should size a column or an exchanger with this model. Lumped sections, constant relative volatilities, an argon tracer and an absent refrigeration balance together rule out any statement about the tray count, recovery, argon yield or turndown of a particular plant. Three lumped sections place every composition the model knows at a section boundary, so nothing can be said about where inside a section a pinch sits or what temperature profile an operator would read, and the argon results are calibrated in magnitude rather than predicted. Refrigeration is the most consequential omission, because on a real unit the expansion turbine, the cold-box heat leak and the liquid production are one balance and that balance decides whether a plant can choose its liquid make at all. Bit-exact repeatability is itself unlike practice. Two of the scenarios are stylized on purpose. Fouling that steps the conductance to 0.90 and then 0.78 in one move is faster than the weeks a real carbon dioxide film takes, but it separates the loss of conductance from every other cause. Cutting the expansion turbine to 3 % would not be done on a real unit, but it is the shortest path to a rectifying section with no heavy key left to send upward. What survived the simplification, one duty shared by two columns, two inventories measured in hours, an argon balance read from both ends and a power limit that binds on a warm afternoon, survived because each of them has to be designed around, and Section 6.3.9 bounds every simplification individually.

The predictive controller’s model is linear at one line-up and at the declared parameter values. It omits every gain that moves with the operating point, such as the pinch near the 98 % split stop, the nitrogen blanketing of the argon condenser, the vane’s moles and kilowatts at another suction temperature and the exchanger conductance after fouling, it holds the argon valve over the horizon although the argon loop moves it, and it has no steady-state target calculation, which is why it loses scenario 7.

6.9 Discussion

The economic argument on an air separation unit is a ratio rather than a rate. Fouling moves the production by eight percent while leaving specific power at 0.381 kW h/m^3, and heat leak moves it to 0.531 kW h/m^3 while the purity analyzer improves, so neither failure is visible on a controlled variable. Fouling shows in the driving force and heat leak in recovery and specific power, which is why studies of air separation operation evaluate candidate structures by steady-state loss rather than by closed-loop response [3, 15]. The same equations carry methods beyond the linear predictive controller of Section 6.5.5, which plans three handles from one design point and leaves the argon to its own loop. Eighteen states with an active compressor constraint, two purity specifications and two large inventories are a natural nonlinear model predictive control problem, which is where the published work on this plant went [8, 9], and a nonlinear controller of that kind with a steady-state target calculation would trade oxygen purity, argon recovery and compressor power directly instead of leaning on a relative gain array of 0.471 to pick a pairing that cannot decouple. The inventories are the second opportunity, because twenty cubic meters of liquid oxygen and a sidearm holding two hundred kilomoles are the storage that makes an air separation unit a demand-response asset [16], and the same equations with an electricity price attached become a scheduling problem. Rebuilding the eighteen equations in the Gekko Optimization Suite, with the nine hidden constants declared as parameters to be estimated [29], is the shortest route to either.

OPEN CHALLENGE Scheduling the Plant on the Price of Power

Electricity is in effect the only paid feedstock of an air separation unit, and products stored as liquid let it make more when power is cheap and less when it is dear [17, 30]. Prices can change from every several hours to every five minutes, the time scale of the plant's own dynamics, so a schedule has to carry a model of them, and such models are large, nonlinear and stiff [17]. Linear models fitted to a year of industrial data served over a scheduling horizon, but their accuracy drifted unpredictably and had to be updated online [30]. Progress would be a schedule that knows how far the sump, the sidearm and the compressor rating let the plant move as the plant changes.

6.10 Conclusion

This chapter studied a double-column cryogenic air separation unit as an interactive laboratory for thermal coupling and the loop interaction that follows from it. The simulator integrates eighteen states covering two column pressures, three inventories, six composition lumps solved by marching each section along its own operating line, two argon tracers, a sidearm relaxing with a 278 min hold-up constant and four actuator positions, all driven by one duty of 546.0 kW across 1.894 K that condenses 109.2 mol/s of nitrogen and boils 81.5 mol/s of oxygen. Pinching a valve on the low-pressure column raises the high-pressure column by 3.11 bar per bar at constant driving force and costs 45.5 kW for the same oxygen. The duty setpoint meets the 752 kW motor rating at 5.699 bar and 84.10 mol/s of boil-up, and the reflux split moves production from 1328 to 2015 m^3/h and purity from 99.744 to 99.594 % while the compressor stays within 0.2 kW of 722 kW. The crude argon valve moves the oxygen purity by 378 ppm per percent against the reflux split's 70 ppm, and the relative gain array of the two purity loops returns 0.471 for the pairing the plant is obliged to use. Three failures hold every setpoint, with fouling taking eight percent of the production while the compressor power falls, 45 kW of heat leak taking recovery from 90.8 to 64.9 % while the purity analyzer improves, and a recovery pinch that gives no warning until the belly nitrogen reaches 17 700 ppm. A linear model predictive controller on the vanes, the product valve and the split keeps the oxygen above its specification through a day that full cascade spends 724.7 min under, and loses the oxygen and level comparison when the pressure setpoint lies beyond the motor rating. A fuller model would add a refrigeration balance so that liquid make becomes a constraint rather than a slider, a rigorous third component in place of the argon tracer, and a steady-state target calculation that keeps the predictive controller from pursuing an unreachable setpoint. The app puts a plant on which nothing can be moved alone within reach of anyone with a browser.

KEY POINTS

- One exchanger is the only condenser of the high-pressure column and the only reboiler of the low-pressure column, so nothing on the plant can be moved alone.
- A valve on the low-pressure column raises the high-pressure column by 3.11 bar per bar at constant driving force, and the duty setpoint meets the 752 kW motor rating at 5.699 bar.
- The crude argon valve moves the oxygen purity by 378 ppm per percent against the reflux split's 70 ppm, and the relative gain of the pairing the plant uses is 0.471.
- Fouling, heat leak and a recovery pinch hold every setpoint, and show only in variables nobody controls.

Where else this applies. Volume 6, Chapter 4 heats units with one another's vapor, and Chapter 8 shares one condenser and one reboiler between two column sections.

References

[1] T. P. Chiang and W. L. Luyben, "Comparison of the dynamic performances of three heat-integrated distillation configurations," *Industrial & Engineering Chemistry Research*, vol. 27, no. 1, pp. 99–104, 1988.

[2] 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.

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

[4] D. Sbarbaro and R. Ortega, "Averaging level control: An approach based on mass balance," *Journal of Process Control*, vol. 17, no. 7, pp. 621–629, 2007.

[5] A. Ebrahimi, M. Meratizaman, H. Akbarpour Reyhani, O. Pourali, and M. Amidpour, "Energetic, exergetic and economic assessment of oxygen production from two columns cryogenic air separation unit," *Energy*, vol. 90, pp. 1298–1316, 2015.

[6] A. R. Smith and J. Klosek, "A review of air separation technologies and their integration with energy conversion processes," *Fuel Processing Technology*, vol. 70, no. 2, pp. 115–134, 2001.

[7] B. Roffel, B. H. L. Betlem, and J. A. F. de Ruijter, "First principles dynamic modeling and multivariable control of a cryogenic distillation process," *Computers & Chemical Engineering*, vol. 24, no. 1, pp. 111–123, 2000.

[8] S. Bian, S. Khowinij, M. A. Henson, P. Belanger, and L. Megan, "Compartmental modeling of high purity air separation columns," *Computers & Chemical Engineering*, vol. 29, no. 10, pp. 2096–2109, 2005.

[9] Z. Chen, M. A. Henson, P. Belanger, and L. Megan, "Nonlinear model predictive control of high purity distillation columns for cryogenic air separation," *IEEE Transactions on Control Systems Technology*, vol. 18, no. 4, pp. 811–821, 2010.

[10] W. F. Castle, "Air separation and liquefaction: recent developments and prospects for the beginning of the new millennium," *International Journal of Refrigeration*, vol. 25, no. 1, pp. 158–172, 2002.

[11] R. Agrawal and D. W. Woodward, "Efficient cryogenic nitrogen generators: An exergy analysis," *Gas Separation & Purification*, vol. 5, no. 3, pp. 139–150, 1991.

[12] C. Linde, "Process of producing low temperatures, the liquefaction of gases, and the separation of the constituents of gaseous mixtures," U.S. Patent 727,650, filed July 9, 1895, granted May 12, 1903, 1903. [Online]. Available: https://patents.google.com/patent/US727650A/en

[13] ——, "Apparatus for producing pure nitrogen and pure oxygen," U.S. Patent 795,525, filed October 22, 1903, granted July 25, 1905, 1905. [Online]. Available: https://patents.google.com/patent/US795525A/en

[14] Y. Hirose, T. Kitazawa, and I. Funada, "Simulation of upper and crude-argon columns in air separation process," *Kagaku Kogaku Ronbunshu*, vol. 15, no. 4, pp. 695–702, 1989.

[15] Y. Cao, C. L. E. Swartz, M. Baldea, and S. Blouin, "Optimization-based assessment of design limitations to air separation plant agility in demand response scenarios," *Journal of Process Control*, vol. 33, pp. 37–48, 2015.

[16] R. C. Pattison, C. R. Touretzky, T. Johansson, I. Harjunkoski, and M. Baldea, "Optimal process operations in fast-changing electricity markets: Framework for scheduling with low-order dynamic models and an air separation application," *Industrial & Engineering Chemistry Research*, vol. 55, no. 16, pp. 4562–4584, 2016.

[17] M. T. Kelley, R. C. Pattison, R. Baldick, and M. Baldea, "An MILP framework for optimizing demand response operation of air separation units," *Applied Energy*, vol. 222, pp. 951–966, 2018.

[18] M. Variny, D. Jediná, M. Rimár, J. Kizek, and M. Kšiňanová, "Cutting oxygen production-related greenhouse gas emissions by improved compression heat management in a cryogenic air separation unit," *International Journal of Environmental Research and Public Health*, vol. 18, no. 19, p. 10370, 2021.

[19] European Industrial Gases Association, "Safe operation of reboilers/condensers in air separation units," European Industrial Gases Association, Brussels, EIGA Doc 65/20, 2020.

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

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

[22] 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.

[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. 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.

[27] European Industrial Gases Association, "Guideline for safe practices for cryogenic air separation plants," European Industrial Gases Association, Brussels, EIGA Doc 147/22, 2022.

[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] L. D. R. Beal, D. C. Hill, R. A. Martin, and J. D. Hedengren, "GEKKO optimization suite," *Processes*, vol. 6, no. 8, p. 106, 2018.

[30] C. Tsay, Y. Cao, Y. Wang, J. Flores-Cerrillo, and M. Baldea, "Identification and online updating of dynamic models for demand response of an industrial air separation unit," *IFAC-PapersOnLine*, vol. 54, no. 3, pp. 140–145, 2021.

6.11 Exercises

6.1 Concept check · 2 min

In scenario 2 the waste nitrogen valve on top of the low-pressure column is pinched to 78 % with every handle in Manual, and the high-pressure column rises from 5.600 to 6.303 bar. No pipe joins the two column tops. Why does the high-pressure column rise?

(a) Waste nitrogen backs up through the main exchanger into the high-pressure column

(b) The pressure loop opens the guide vanes to keep the oxygen production up

(c) The condenser-reboiler duty rises, so more nitrogen condenses at a higher pressure

(d) The sump oxygen boils hotter, and the same duty needs the same driving force

6.2 Concept check · 1 min

In scenario 8 the condenser-reboiler loses 22 % of its conductance while the pressure and level loops hold their setpoints. What does the main air compressor power do?

(a) It falls, because the pressure loop closes the vanes as less nitrogen condenses

(b) It rises, because the pressure loop opens the vanes to restore the lost duty

(c) It stays at 722 kW, because the pressure and so the pressure ratio are held

(d) It rises, because the wider driving force needs a higher column pressure

6.3 Concept check · 2 min

At steady state with the sump level held, which handle moves the oxygen purity most per percent?

(a) The reflux split, the handle the oxygen purity loop is paired with

(b) The crude argon draw valve, since the impurity left in the product is argon

(c) The guide vanes, because they set the boil-up of the low-pressure column

(d) The liquid oxygen valve, because it sets how fast the product leaves the sump

6.4 Theory and calculation · 8 min

Equation (6.3) gives the oxygen product per mole of air fed to the high-pressure column. At design $F_{HP} = 109.21\,\text{mol/s}$, $f_{LN} = 0.400$, $\psi_R = 0.0250$, and one mol/s of product is 80.69 m³/h at normal conditions.

(a) The five terms are 0.37052, 0.49200, 0.10073, −0.74627 and zero. Compute P_{LOX}/F_{HP} and P_{LOX}, and say where the fourth term comes from. *Checkpoint:* 0.21698 and 23.70 mol/s; $-5000/6700$.

(b) Hold every other term and raise the reflux split by one percentage point. Compute the gain in oxygen product in mol/s and in m³/h, and the oxygen lost for each mole of liquid nitrogen sold instead of returned. Give three significant figures. *Checkpoint:* 0.426 mol/s, 34.4 m³/h; 0.975 of a mole.

(c) Compute the last term and the product lost for 45 kW of cold-box heat leak, in mol/s and m³/h. *Checkpoint:* 0.0615; 6.72 mol/s, 542 m³/h.

6.5 Theory and calculation · 6 min

Equation (6.22) gives the steady-state gains of the two purities on the reflux split and the argon valve: $K_{11} = -0.0070$, $K_{12} = 0.0378$, $K_{21} = -0.1151$ and $K_{22} = -0.5536$ (%/%).

(a) Compute $\lambda_{11} = 1/(1 - K_{12}K_{21}/(K_{11}K_{22}))$ and λ_{12} to three decimals. *Checkpoint:* 0.471 and 0.529.

(b) Compute the gain of the reflux split on the oxygen purity when the argon loop is closed and holds the crude argon purity, and its ratio to K_{11}. Give three significant figures. *Checkpoint:* –0.0149 %/%, 2.12 times.

(c) Say what (b) means for tuning the oxygen purity loop with the argon loop in Manual and then in Auto. *Checkpoint:* The loop gain doubles when the argon loop closes.

6.6 Theory and calculation · 8 min

The liquid oxygen sump holds 1 400 000 mol at 100 % level, and the product valve passes $C_{lox} = 45.0$ mol/s fully open (Equation (6.11)). The level loop has action -1 (direct acting) and the shipped tuning of Table 6.4.

(a) Compute the integrating gain of the level on the valve, in percent of level per minute per percent of valve, to three significant figures. *Checkpoint:* 0.001 93 %/min per percent.

(b) Scenario 4 starts the sump at 42 % against a 50 % setpoint with the valve at 52.66 %, and the integral is seeded so that the first execution asks for that position plus the proportional term. Use Equation (6.23) to find the valve it asks for at $K_c = 3$ and at $K_c = 9$. *Checkpoint:* 28.66 %; -19.34, held at 0 %.

(c) Ignore the integral. Compute the closed-loop time constant $1/(K_cK)$ for both gains, in minutes, and the level after 60 min at $K_c = 3$, to one decimal. *Checkpoint:* 173 min and 58 min; 44.3 %.

6.7 App challenge · 8 min

Load scenario 2, An LP-column valve raises the HP-column pressure (MANUAL), and change no setting. The waste nitrogen valve is pinched to 78 % at 40 min. Use the default seed, run at 20 min/s to 215 min and export the CSV file. Its time column is in hours.

(a) In the rows at 0.5 and 3.5 h, read `hp_bar`, `lp_bar`, `Tcond_K`, `Treb_K` and `mac_kW`. Accept 0.002 bar, 0.02 K and 0.2 kW. *Checkpoint:* 5.600, 1.400, 95.36, 93.47 and 722.0 kW; 6.303, 1.626, 96.91, 95.01 and 767.5 kW.

(b) Compute the rise of the high-pressure column per bar of low-pressure column and the driving force before and after. Compare the ratio with the 2.87 that Equation (1.9) gives. *Checkpoint:* 3.11; 1.894 K both times.

(c) The vane stays at 100 %, so $F_{air} = 124.1$ mol/s. With Equation (6.6), $T_{suc} = 303.15$ K, $\eta_{MAC} = 0.767$, $\Delta P_{mac} = 0.35$ bar, $P_{amb} = 1.013$ bar and $R = 8.314$ J/(mol K), compute W_{MAC} at both high-pressure column pressures. *Checkpoint:* 722.0 and 767.5 kW, 45.5 kW apart.

6.8 App challenge · 8 min

Load scenario 8, The shared exchanger fouls, and the plant quietly unloads, and change no setting. The conductance falls to 0.90 at 60 min and to 0.78 at 600 min, and the pressure setpoint goes to 5.78 bar at 1400 min. Use the default seed, run at 60 min/s to 1700 min and export the CSV file. Its time column is in hours.

(a) In the rows at 0.5 and 23.3 h, read the guide vane, the oxygen, the compressor, the driving force and the specific power, columns `gv_pct`, `gox_Nm3h`, `mac_kW`, `dTcr_K` and `spec_power_`

`kWh_Nm3`. Accept 0.2 %, 3 m³/h, 0.5 kW, 0.005 K and 0.001. *Checkpoint:* 100.0, 1912, 722.0, 1.894, 0.379; 92.8, 1762, 669.8, 2.253, 0.381.

(b) In the row at 28.3 h, read the same columns. Accept 0.2 %, 5 m³/h, 0.5 kW, 0.005 K and 0.001. Compute the power the setpoint move costs and its price over a year of 8760 h at 0.062 $/kWh. Accept $600. *Checkpoint:* 98.3, 1877, 721.4, 2.386, 0.386; 51.6 kW, $28 000.

(c) Say which controlled variable showed the fouling before 1400 min, and which trend did. *Checkpoint:* None; the driving force.

6.9 App challenge · 8 min

Load scenario 4, Twenty cubic meters of memory, and change no setting. The sump starts at 42 % against a 50 % setpoint. Use the default seed, run at 20 min/s to 480 min and export the CSV file. Its time column is in hours.

(a) Read `lox_cmd_pct` in the first row, the lowest `gox_Nm3h`, and `lox_level_pct` at 1 and 8 h. Accept 0.05 %, 2 m³/h and 0.05 %. *Checkpoint:* 28.66, 1042, 44.56 and 51.79.

(b) Pause, press Reset, set LOX level loop gain to 9 before Run, run to 480 min again and export. Read `lox_cmd_pct` in the first row, from which row to which `gox_Nm3h` stays below 1 m³/h (accept 0.01 h), and the level at 1 h. *Checkpoint:* 0; from 0.046 to 0.450 h; 47.37 %.

(c) Compare the two runs as the steel plant that buys the oxygen would. *Checkpoint:* The shipped gain keeps the product flowing.

6.10 Application area · 4 min

A plant manager wants an alarm that would have caught both the exchanger fouling of scenario 8 and the cold-box heat leak of scenario 11 while every loop held its setpoint.

(a) From the chapter's measured values, say what each failure does to the oxygen purity, the production and the specific power. *Checkpoint:* Purity up in both; production down in both; specific power flat for fouling, up for heat leak.

(b) Propose what to trend and alarm, and say why the oxygen analyzer cannot be one of them. *Checkpoint:* Driving force, recovery and specific power.

6.11 Application area · 4 min

Table 6.3 prices oxygen at 0.055 $ and liquid nitrogen at 0.030 $ per normal cubic meter. At a 95 % reflux split the plant makes 1912 of oxygen and 176 of liquid nitrogen (m³/h), and at 78 % it makes 1328 and 775, with the compressor unchanged.

(a) Compute the hourly value of the two products at each split and the difference. *Checkpoint:* $110.44 and $96.29 an hour; $14.15 less at 78 %.

(b) Find the liquid nitrogen price at which the two splits are worth the same, and relate it to part (b) of the first theory exercise. *Checkpoint:* 0.0536 $ per cubic meter.

(c) Name the limitation of the model that makes selling liquid nitrogen look cheaper here than on a real unit. *Checkpoint:* No refrigeration balance.

6.12 Review the agent · 5 min

An AI assistant proposes a plan for scenario 14, the challenge A shift on a badly lined-up plant: "Press Reveal suggested tuning, then raise the HP column pressure setpoint to 5.85 bar before Run to buy back the duty the fouling takes, and return it to 5.60 bar before 900 min. The compressor power override caps the guide vanes at the 752 kW rating, so the compressor can never reach the 780 kW the target forbids." Find the flaw and say how you would check it in the app.

CHAPTER 7

Amine Capture Plant Simulator for Recycle Gain and the Scheduling of Regeneration Energy

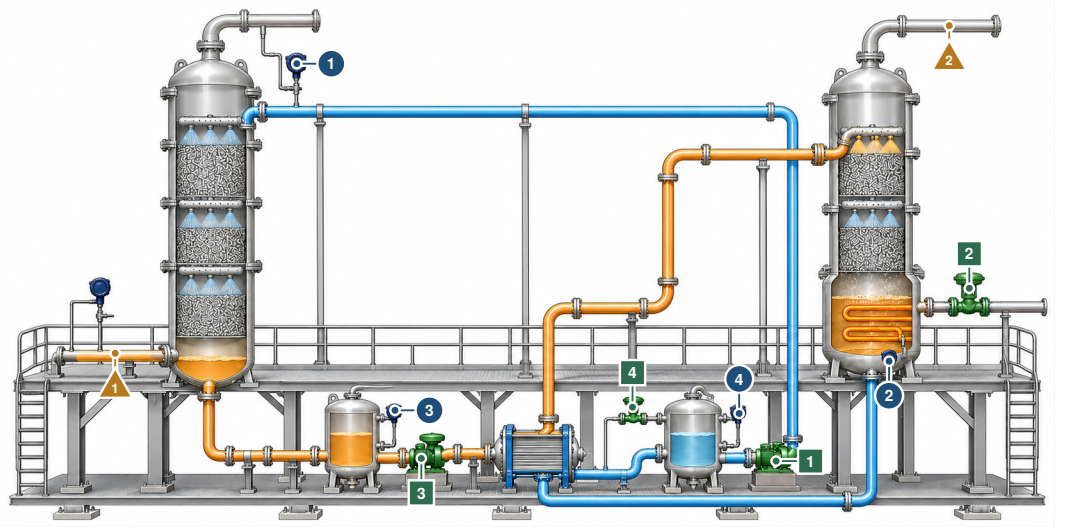

The controller moves 1 lean solvent circulation pump 2 reboiler steam valve 3 stripper feed pump 4 amine makeup valve

The controller reads 1 capture analyzer on the treated gas 2 reboiler temperature 3 rich tank level 4 lean tank level

What disturbs it 1 host load, which sets the flue gas 2 stripper pressure drift

Notice that the solvent leaving the absorber passes the stripper and both tanks before it returns, so a change of circulation comes back to the absorber about seventy minutes later. Illustration, not to scale.

IN PLAIN TERMS **A Cleaning Liquid That Comes Back Round**

The main idea. A capture plant washes the smoke of a coal power station with a liquid that soaks up carbon dioxide. Steam from the station then boils the gas out, and the liquid flows back to wash again about seventy minutes later.

The challenge. More liquid catches more gas at first, but unless the steam rises too, the liquid returns less clean and the gain fades. Nobody measures how clean the liquid is, and the reading used in its place can be fooled.

Why it matters. Steam spent on the liquid is power the station cannot sell, so when to spend it is worth real money.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page; the history boxes of Section 7.2; the fact box and the try-it box of Section 7.3; the fact box and the try-it box of Section 7.6; the key points
Study	65 min	the whole chapter with its pause points and the exercises
Practice	20 min	the first page; Section 7.6 with its result chart; Section 7.7; Section 7.8; the open challenge of Section 7.9

The control question A capture plant's solvent takes seventy minutes to come back round its circuit, so every move of the circulation acts through a recycle. What sets the gain of the capture loop, and when should the regeneration energy, bought out of the host station's saleable output, be spent?

Learning objectives

- Explain why the capture gain decays with the reboiler duty frozen and why a lean loading loop underneath holds it constant.
- Compute the transport delays and the lap of a solvent circuit from its holdups and the circulation.
- Diagnose an inferred lean loading that a drift in stripper pressure misleads, and name the defense.
- Tune a surge tank level loop for averaging, and read a level setpoint as an energy schedule bounded by the other tank.
- Compare flat operation, an economic supervisor and a predictive controller on margin, captured carbon dioxide and solvent left in store.

Prerequisites Sections 1.4.3 (recycle gain), 1.4.4 (the solvent loop), 1.4.6 (inferred compositions), 1.6.2 (velocity form) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive controller, and Chapter 2 shows another large recycle.

The app https://apmonitor.com/apps/co2-capture/

See also Volume 6, Chapter 7 for a recycle whose gain depends on which flow is held. Volume 9, Chapter 6 for a flow that sets its own transport delay, and a price of heat that coordinates the consumers. Volume 5, Chapter 10 for a product quality inferred from a temperature that a second variable misleads.

7.1 Introduction

A cascade is usually justified by speed. The inner loop is described as fast, the outer loop as slow, and the benefit as a disturbance intercepted before it reaches the primary measurement. A post-combustion capture plant offers a different justification, and offers it in numbers. The circuit in this chapter holds 9100 m^3 of 30 wt% monoethanolamine and circulates 130 m^3/min, so a parcel of solvent needs 70 min to return to the point it left. Raising the circulation by 20 m^3/min with the reboiler duty held fixed lifts the capture rate by 3.54 points within ten minutes, an apparent gain of 0.177 points per m^3/min, and 77 % of that rise is handed back over the following eight hours to a settled gain of 0.040. Closing a lean loading loop underneath, so that the reboiler duty follows the solvent flow, removes the decay entirely and holds the gain at 0.177 from ten minutes to eight hours. The inner loop did not make the outer loop faster. It made the outer loop's gain a constant.

What the simulator corrects is the habit of reading a large recycle as a slowness problem. A block diagram absorbs the recycle into a lumped transfer function whose steady-state gain is a single number, and a student who is handed that number never watches it fall by a factor of 4.4 according to what else in the circuit happens to be regulated. The shape is general and it does not require an amine. Whenever a stream returns to the front of a process and something inside the return path is left on manual, the settled gain of the forward loop is set by the unregulated variable rather than by the unit the actuator is attached to, which is the structural property that recycle amplification describes for reactor and separator systems [1] and that Volume 6, Chapter 7 measures on a reactor and column. Refinery hydrogen loops, ammonia synthesis loops such as the one of Chapter 2, hydrometallurgical leach circuits and paper machine broke systems all present it, and so does any solvent regeneration circuit in a gas plant [2]. The second property the plant carries is transport delay that is not a time constant. A pipe holds a fixed number of cubic meters rather than a fixed number of minutes, so the 9.00 min delay in each 1170 m^3 line becomes 7.55 min at 155 m^3/min and a controller tuned around it at one circulation is mistuned at another.

Capture plants have been modeled dynamically for about fifteen years [3, 4], control structures have been proposed for them by plantwide reasoning and by self-optimizing selection [5, 6], and the modeling, identification and control literature has been surveyed in detail [7, 8]. Those studies are written for the engineer who already accepts that the circuit governs the plant, and the measurement

that establishes it is a transient test on a plant nobody will hand to a student. This chapter works on a capture plant, built to be operated in a browser, whose recycle is slow enough and whose inventory is large enough that the change in loop gain can be measured rather than asserted, and whose regeneration energy is bought out of the host station's saleable output so that the timing of that energy becomes a control decision. The plant is a sixteen state absorber, stripper and inventory model with two volume-keyed plug flow lines, an inferential reboiler measurement and an economic supervisor, given here in the detail a reconstruction of it would need. Its fifteen scenarios run from a hand-operated absorber to a sixteen-hour price day against a rolling permit, and every variable the simulator computes can be exported. The solvent chemistry, the plantwide arguments and the tuning arithmetic are standard, and each is cited to the work that produced it. Section 1.4.4 treats the solvent loop as a recycle with a negative return gain on the small solvent unit of S-1, and this chapter measures the same effect on a full-size plant.

Section 7.2 traces amine scrubbing and the loops added to it. Section 7.3 gives the constants and balances the plant is assembled from and the assumptions behind them. Section 7.4 measures what the recycle does to the gain of the capture loop and times the lap item by item. Section 7.5 describes the three PI loops, the tank availability factors, the economic supervisor and the predictive controller. Section 7.6 reports loop performance, the surge tank, the cost and timing of the energy, two silent faults, the graded challenge and the predictive controller against the supervisor. Section 7.7 names the real devices behind the loops. Sections 7.8, 7.9 and 7.10 give the limitations, a discussion and the conclusions, and the exercises in Section 7.11 close the chapter.

7.2 Amine Scrubbing and the Loops Added to It

Washing an acid gas with an aqueous alkanolamine and regenerating the solution with steam is a 1930s technology, and it has been the standard route to sweetening natural gas and to treating refinery and ammonia plant gas ever since [2]. The same solvent applied to a power station flue gas is a much larger and much colder duty, and the reason it stayed a laboratory proposition for decades is the reboiler rather than the absorber [9]. An absorber is a well-behaved contactor. A stripper that has to reverse an 85 kJ/mol heat of absorption on hundreds of tonnes of carbon dioxide an hour is a heat sink comparable with a boiler's own reheat load, and the cost of capture is dominated by what that heat is worth as electricity [10].

> ⧖ HISTORY **Amines in a Closed Cycle, 1930**
>
> On October 7, 1930 Robert Roger Bottoms applied for a United States patent on a process for separating acidic gases, and named carbon dioxide, sulfur dioxide and hydrogen sulfide as the main ones. He had found that a small group of amines unite chemically with such gases at a low temperature and give them up as gas at a higher temperature, regenerated in the same step. A liquid with that property can circulate continuously in a closed cycle between a cool absorption stage and a hot regeneration stage, with heat interchangers to save heat. Monoethanolamine, the solvent of this chapter's plant, is among the amines the patent lists, although it preferred triethanolamine for its higher boiling point [11]. Amine scrubbing has separated carbon dioxide from natural gas and hydrogen ever since [9].

The operating problem was answered by hand for a long time. A gas treating unit ran a fixed circulation and a fixed steam rate chosen from a design case, the reboiler was managed on its own temperature because the temperature was the only continuous measurement on the regeneration side, and the lean loading that the plant actually cared about arrived on a laboratory report the following day [2]. The equilibrium data that would let a temperature be read as a loading were measured for the 30 mass percent solution over a wide loading range only much later [12]. Two consequences of the manual arrangement survived into the power station application. The first is that the loading is still inferred rather than measured on almost every plant, and an inference is a function of the measurement and of something else, which is the classical situation that secondary measurements

were introduced to handle [13]. The second is that the solvent degrades thermally at a rate that roughly doubles for every eight degrees of reboiler temperature, so running the stripper harder is paid for twice, once in steam and again in amine [14].

Dynamic analysis of the coupled plant is recent. Absorber column models written for transient rather than steady-state duty appeared around 2009 [3], and the first studies aimed at the energy question showed that the lean loading and the stripper pressure are the variables that decide the reboiler bill and therefore the penalty [4]. Control structure followed. Plantwide selection of pairings for the absorber and stripper together was published in 2011 [5], self-optimizing selection of controlled variables for the same process appeared in the same year [6], and a survey of the field a decade later found the arrangement now regarded as standard, which is capture rate on solvent circulation with an inner loop on reboiler temperature or an equivalent inferential loading measurement, and a deliberately slack level controller on the surge tank between the two columns [7, 8].

Flexibility arrived as an economic argument rather than a control one. Part load operation of a capture plant on a supercritical coal unit was evaluated early enough to establish that the steam extraction, not the absorber, sets what the host can sell [15], and the observation that the carbon dioxide does not care when it is regenerated turned solvent storage into a scheduling problem against a volatile electricity price [16]. Rolling-average emissions obligations are what make the scheduling legal, because they permit the instantaneous capture rate to move while the total does not. The application implements that argument in full, including the part that is easy to get wrong, which is what a supervisor is allowed to do with the permit when the price is high enough.

> ⧖ HISTORY **Capture Reaches Coal Power Stations, 2014 to 2019**
>
> Amine absorption had removed carbon dioxide from natural gas for decades before it treated a power station’s flue gas at commercial scale. In October 2014 a capture plant on a 110 MW coal-fired unit in Canada opened as the world’s first fully commercial plant of its kind [17]. On December 29, 2016 a second began commercial operation in Texas, treating a slipstream of flue gas equivalent to 240 MW and capturing at least 90 percent of its carbon dioxide, about 5,200 short tons a day at full capacity. Instead of drawing steam from the coal unit, as this chapter’s plant does, it raised most of its steam from the exhaust heat of a separate gas turbine. Over its three-year demonstration, 2017 to 2019, it captured 3,542,537 metric tons of carbon dioxide [18].

7.3 Process Model

Capture plant U-400 treats the flue gas of a 580 MW coal unit. Absorber C-401 receives 21.0 kmol/s of gas carrying 12.4 % carbon dioxide, which is 412.6 t/h presented, and removes 90.0 % of it into 130 m³/min of solvent at the design line-up. Rich solvent leaves the bottom of the column at 0.4472 mol/mol of carbon dioxide per mole of amine, travels 900 m of DN1300 pipe to surge tank TK-401, and is fed through cross exchanger E-401 to stripper C-402, whose reboiler takes 373.3 MW of low-pressure steam from the host turbine crossover and returns the solvent to 0.2375 mol/mol at 115.87 °C and 190 kPa. Lean solvent collects in TK-402 and returns through a second 900 m line to the top of the absorber. Figure 7.1 draws that circuit with its instrument tags and with every symbol the equations below use.

Table 7.1 sorts those quantities by the role each one plays. Units throughout are the ones a utility operator reads, volumetric flow in m³/min, duty in MW, loading in mol/mol, level in percent of tank capacity, and time in minutes. Sixteen states are integrated.

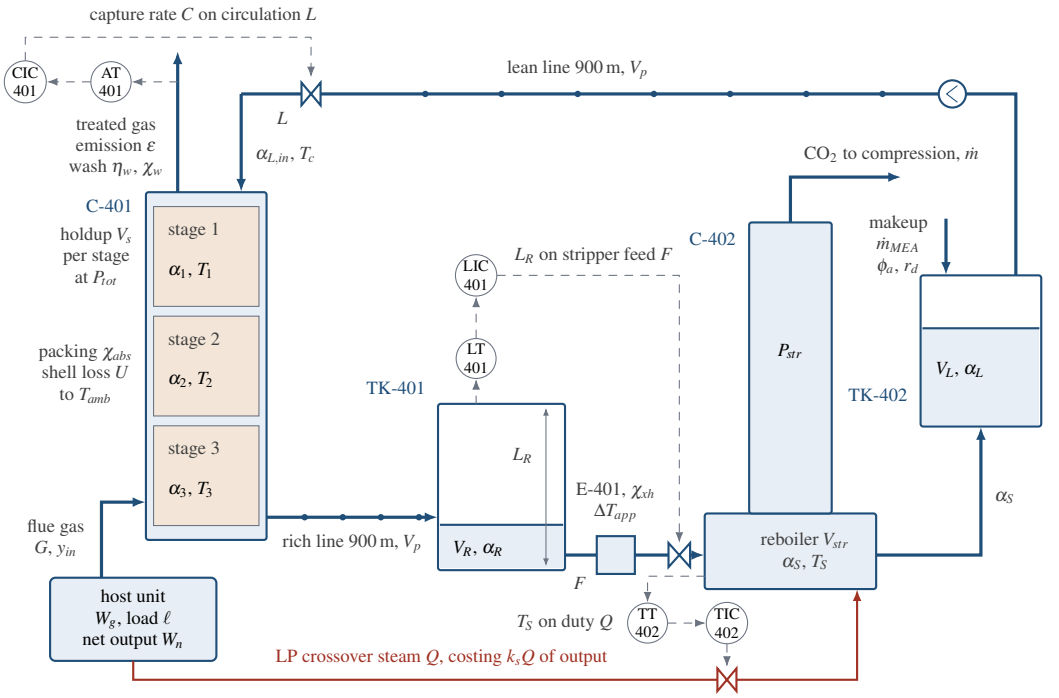

Circuit inventory $\sum V = 9100\,\mathrm{m}^3$: two lines at $1170\,\mathrm{m}^3$ each, lean tank $5100\,\mathrm{m}^3$, rich tank $800\,\mathrm{m}^3$, reboiler $500\,\mathrm{m}^3$, absorber $360\,\mathrm{m}^3$, so one lap at $L = 130\,\mathrm{m}^3/\mathrm{min}$ is 70.0 min.

Figure 7.1 Capture plant U-400 drawn as the circuit it is. Flue gas G of composition y_{in} meets lean solvent of loading $\alpha_{L,in}$ at temperature T_c in the three stages of absorber C-401, each holding V_s at loading α_i and temperature T_i. Rich solvent crosses the 900 m rich line of volume V_p to surge tank TK-401, whose level L_R is held by stripper feed F through cross exchanger E-401 of approach ΔT_{app}. Reboiler duty Q is crossover steam worth $k_s Q$ of the host's saleable output, and it sets the lean loading α_S at bubble point T_S and stripper pressure P_{str}. Lean solvent returns through TK-402 at circulation L, the manipulated variable of the capture loop. Dashed lines are signals and the markers in the two solvent lines are equally spaced parcels whose spacing is the transport delay.

7.3.1 Vapor–Liquid Equilibrium

Every quantity in the plant follows from the equilibrium partial pressure of carbon dioxide over the loaded solvent, which is written as a reference value times three factors in Equation (7.1).

$$P^* = P_0^* \left(\frac{\alpha}{\alpha_0}\right)^n \left(\frac{\alpha_{max} - \alpha_0}{\alpha_{max} - \alpha}\right)^m f_T$$
$$f_T = \exp\left[-\frac{\Delta H_{abs}}{R}\left(\frac{1}{T} - \frac{1}{T_0}\right)\right] \tag{7.1}$$

where α is the loading in mol of carbon dioxide per mol of amine, T is the liquid temperature in K, $n = 2$ carries the carbamate stoichiometry of one carbon dioxide to two amines, $m = 1.6$ is the depletion exponent, $\alpha_{max} = 0.58$ is the loading at which the free amine runs out, and $\Delta H_{abs} = 85\,\mathrm{kJ/mol}$ is the heat of absorption. The loading is clamped to $0.004 \le \alpha \le \alpha_{max} - 0.004$ before the depletion term is evaluated, because a transient that overshoots the rich loading otherwise divides by zero and the integrator carries the infinity into every other state within one step. The shape of (7.1) is taken from the measured solubility of the 30 mass percent solution [12] and the constants are fitted rather than regressed. The Arrhenius factor f_T is the separation. At 0.2375 mol/mol it is worth

Table 7.1 Everything the Control System Touches, With Units Carried in the Symbol Column. Each Nominal Is Read at the Design Line-Up Where Every Scenario Begins, and a Range Marked SP Spans the Setpoint Slider.

Symbol and unit	Role	Nominal	Range
C, AT-401 (%)	CV, measured	90.00	60–98 SP
T_S, TT-402 (°C)	CV, measured	115.87	100–126 clamp
L_R, LT-401 (%)	CV, measured	10	5–90 SP
L_L, LT-402 (%)	measured, MPC	63.75	0–100
α_S (mol/mol)	inferred CV	0.2375	0.15–0.34 SP
L (m^3/min)	MV	130	55–170
Q (MW)	MV	373.3	0–560
F (m^3/min)	MV	130	0–175
makeup (kg/min)	MV	0	0–120
host load (pu)	DV, measured	1.00	0.40–1.00
price ($/MWh)	DV, measured	60	0–260
T_c (°C)	DV, measured	40	25–50
P_{str} (kPa)	DV, unmeasured	190	150–240
r_d (%/day)	DV, unmeasured	0	0–8
gas bias (×)	DV, unmeasured	1.00	0.85–1.15
χ_{abs} (×)	DV, unmeasured	1.00	0.50–1.05
χ_{xh} (×)	DV, unmeasured	1.00	0.60–1.05
χ_w (×)	DV, unmeasured	1.00	0.50–3.00
analyzer noise, 1σ (%)	DV, unmeasured	0	0–2
α_1 (mol/mol)	state	0.2841	0.01–0.57
α_2 (mol/mol)	state	0.3680	0.01–0.57
α_3 (mol/mol)	state	0.4472	0.01–0.57
T_1, T_2, T_3 (°C)	states	52.8/60.4/58.5	0–110
V_c (m^3)	state	0	unbounded
V_R (m^3)	state	800	0–8000
α_R (mol/mol)	state	0.4472	0.01–0.57
α_S (mol/mol)	state	0.2375	0.01–0.57
T_S (°C)	state	115.87	40–170
V_L (m^3)	state	5100	0–8000
α_L (mol/mol)	state	0.2375	0.01–0.57
Q_a (MW)	state	373.3	0–900
ϕ_a (×)	state	1.000	0.30–1.05
ℓ_a (pu)	state	1.00	0.05–1.15

a factor of 582 between the 40 °C of the absorber and the 115.87 °C of the reboiler, and no other term in the model moves by anything like that much.

7.3.2 Absorber Stages

The gas side is solved algebraically because it is quasi-steady, with two seconds of gas residence against three minutes of liquid holdup. Within a stage the gas is treated as plug flow over a well-mixed liquid, which integrates in closed form to Equation (7.2).

$$n_{out} = n_{eq} + (n_{in} - n_{eq})\,\mathrm{e}^{-k_i P_{tot}/G}, \qquad n_{eq} = \frac{P^*(\alpha_i, T_i)\,G}{P_{tot}} \tag{7.2}$$

where n is the molar flow of carbon dioxide in kmol/s, G is the total gas flow, $P_{tot} = 101.3\,\mathrm{kPa}$ is the absorber pressure and k_i is the stage transfer coefficient of Equation (7.3). The closed form is worth the algebra. A well-mixed gas stage drives the outlet flow negative once the stage efficiency passes one, which happens as soon as the column runs warm because the kinetics are Arrhenius, and a negative partial pressure propagates into the next stage's driving force. Equation (7.2) is bounded by n_{eq} from both sides whatever the coefficient does.

$$k_i = K_G a\,\chi_{abs}\,\phi_a^{0.8}\left(\frac{L}{L_0}\right)^{n_w}\exp\left[-\frac{E_a}{R}\left(\frac{1}{T_i}-\frac{1}{T_0}\right)\right] \tag{7.3}$$

where $K_G a$ is the design coefficient in kmol/(s kPa), χ_{abs} is the packing condition, ϕ_a is the solvent strength relative to design and floored at 0.2 before the power is taken, $n_w = 0.30$ is the wetting exponent on a liquid load ratio floored at 0.05, $L_0 = 130\,\mathrm{m^3/min}$ and $E_a = 21\,\mathrm{kJ/mol}$. Stage 1 is the top of the column and stage 3 the bottom, and the gas passes from 3 to 1 against liquid passing from 1 to 3. The two state equations per stage are the amine balance and the energy balance of Equations (7.4) and (7.5).

$$H_a \frac{\mathrm{d}\alpha_i}{\mathrm{d}t} = L_a\,(\alpha_{i-1} - \alpha_i) + N_i \tag{7.4}$$

$$\begin{aligned} C_l \frac{\mathrm{d}T_i}{\mathrm{d}t} &= W_{cp}\,(T_{i-1} - T_i) + N_i \Delta H_{abs} - G c_{p,g}\,(T_i - T_{i+1}) \\ &\quad - G\lambda_w\,[x_s(T_i) - x_s(T_{i+1})] - U\,(T_i - T_{amb}) \end{aligned} \tag{7.5}$$

where N_i is the carbon dioxide absorbed on stage i in kmol/s, $H_a = V_s c_a$ is the amine held on a stage in kmol with c_a the amine concentration, L_a is the amine molar flow, $C_l = V_s \rho c_p$ is the stage heat capacity in kJ/K, W_{cp} is the solvent heat capacity flow in kW/K, λ_w is the latent heat of water, $x_s(T)$ is the saturation mole fraction of water from the Buck equation, and U is the shell loss to an ambient at $T_{amb} = 15\,^{\circ}\mathrm{C}$. The upstream liquid for stage 1 is the lean solvent arriving from the pipe and the downstream gas for stage 3 is the flue gas at $T_c + 5\,\mathrm{K}$. Every derivative is written per second and multiplied by 60 in one place, so that the whole model advances on the minute clock the operator reads. The evaporation term in (7.5) is what bounds the bulge, since the gas is water-saturated everywhere and its water content is a steep function of temperature. At the design line-up the profile is 52.8, 60.4 and 58.5 °C from the top down, with the maximum in the middle where the absorption is fastest.

7.3.3 Transport as Volume, Not as Time

The two 900 m lines are true plug flow keyed on cumulative volume rather than on elapsed time, as shown in Equation (7.6).

$$\frac{\mathrm{d}V_c}{\mathrm{d}t} = L, \qquad \alpha_{in}(t) = \alpha_{out}\left(V_c^{-1}[V_c(t) - V_p]\right) \tag{7.6}$$

where V_c is the cumulative volume pumped in $\mathrm{m^3}$ and $V_p = 1170\,\mathrm{m^3}$ is the line volume. A parcel entering the line leaves it when V_p more cubic meters have passed, which is 9.00 min at $130\,\mathrm{m^3/min}$ and 7.55 min at $155\,\mathrm{m^3/min}$. Writing the delay this way needs no variable-delay filter and gets the stretching right for free, which a time-keyed delay does not. The pipes of a district heating network behave the same way, with the delay set by the flow the consumers draw (Volume 9, Chapter 6). The history is held in a ring buffer with linear interpolation and a new sample is pushed only after the cumulative volume has advanced by $V_p/240$, because pushing on every step lets a stalled pump fill the buffer with identical entries and silently empty the line at exactly the moment the plant is slowest.

7.3.4 Tanks, Reboiler and the Duty Split

The two tanks are one inventory seen from opposite ends, which Equation (7.7) makes explicit.

$$\frac{\mathrm{d}V_R}{\mathrm{d}t} = L - F, \qquad \frac{\mathrm{d}V_L}{\mathrm{d}t} = F - L \tag{7.7}$$

Each tank is well mixed in loading, with the holdup floored at $20\,\mathrm{m^3}$ so that a tank running dry cannot turn a finite inflow into an infinite derivative, as shown in Equation (7.8).

$$\max(20, V_R)\, c_a \frac{\mathrm{d}\alpha_R}{\mathrm{d}t} = L_a\,(\alpha_{R,in} - \alpha_R) \tag{7.8}$$

with the identical expression for the lean tank on F and α_S. Equation (7.9) carries the reboiler, which holds a loading, a temperature and the delivered duty.

$$V_{str}c_a\frac{\mathrm{d}\alpha_S}{\mathrm{d}t} = F_a\,(\alpha_R-\alpha_S) - n_{str} \qquad (7.9)$$
$$\frac{\mathrm{d}T_S}{\mathrm{d}t} = \frac{T_b(P_{str},\alpha_S)-T_S}{\tau_{reb}}, \qquad \frac{\mathrm{d}Q_a}{\mathrm{d}t} = \frac{Q-Q_a}{\tau_Q}$$

where n_{str} is the carbon dioxide released, T_b is the bubble point of Equation (7.11), τ_{reb} is the reboiler thermal lag and τ_Q is the steam valve lag. The stripping rate is written causally, duty in and carbon dioxide out, in Equation (7.10).

$$n_{str} = \frac{Q_a - q_{sens}}{\Delta H_{abs} + \lambda_w \varphi \dfrac{P_{str}-P^*}{P^*}}, \qquad q_{sens} = \frac{\dot{m}c_p\Delta T_{app}}{\chi_{xh}} \qquad (7.10)$$

where the equilibrium pressure is floored at 0.3 kPa and the water-to-product ratio at zero before either is used, φ is the fraction of the reboiler vapor that survives to the overhead instead of condensing against the cold rich feed, $\dot{m}$ is the solvent mass flow, ΔT_{app} is the cross exchanger approach and χ_{xh} is the exchanger condition. The numerator of (7.10) is floored at zero. Three properties of that expression organize the economics. The sensible term is paid on the whole solvent flow whether or not there is carbon dioxide in it, which is what punishes over-circulation. The desorption term is a constant of the chemistry at 1.931 MJ/kg and no column goes below it. The stripping term is the molar ratio of water to carbon dioxide in the reboiler vapor, and because P^* collapses as the solvent is stripped the plant boils more and more water to carry away less and less product. Six equal steps of 0.03 mol/mol from 0.34 down to 0.16 cost 38, 43, 51, 65, 90 and 137 MW, so the last increment is 3.6 times the first. The Manual helper that sets the duty for a nominated lean loading solves this stripper alone, at the present feed and rich loading, and at 0.16 mol/mol it reports the 640 MW required and applies the 560 MW steam limit instead. Its answer holds only until the changed lean solvent has gone round the loop and moved the rich loading the stripper is fed. Raising the duty lowers the loading until the rate has climbed to meet it, which is a stable fixed point and the reason a reboiler sits still with no controller of its own.

7.3.5 The Bubble Point and What the Plant Can Measure

Nothing on the plant measures loading. The reboiler thermocouple stands in for it through the bubble point of the loaded solvent, as shown in Equation (7.11).

$$a_w P_w^{sat}(T_b) + P^*(\alpha, T_b) = P_{str} \qquad (7.11)$$

where $a_w = 0.85$ is the water activity in a solution full of carbamate salts, which is the boiling point elevation written as the thing that causes it. Equation (7.11) is solved by 24 bisections over a bracket from 20 to 180 °C, giving 1×10^{-5} K of resolution, which is four orders finer than the thermocouple and cheap enough to evaluate inside every integration stage. At the design pressure 0.2375 mol/mol boils at 115.87 °C. The control system inverts (7.11) at the assumed reference pressure of 190 kPa rather than at the actual pressure, which is the single most important thing about the second loop. 25 kPa of drift makes the same solvent boil about three degrees hotter, so a controller holding the old temperature holds 0.274 mol/mol while the number reported on the screen is still 0.238.

> **TRY IT IN THE APP 6 · The Lean Loading Nobody Measures**
>
> Load scenario 6 and press Run. At 90 min the stripper pressure drifts 25 kPa high. Watch the reboiler temperature return to its setpoint with nothing alarming, while the loading leaving the reboiler settles near 0.274 and the inference still reports 0.238. Only the circulation shows it, climbing by 24 m^3/min.

7.3.6 Host Station, Emission and Degradation

The gas presented to the absorber follows the host load in both its quantity and its strength, and the load itself is a state because a boiler cannot change output instantly. Equation (7.12) carries all three.

$$\begin{aligned} \frac{\mathrm{d}\ell}{\mathrm{d}t} &= \frac{\ell_{cmd} - \ell}{\tau_\ell}, \qquad G = \max(0.2,\, G_{max}\,\ell\, b_g) \\ y_{in} &= y_0\,[1 - y_t\,(1-\ell)] \end{aligned} \tag{7.12}$$

where ℓ_{cmd} is the load the operator or the scenario asks for, $\tau_\ell = 4.0\,\mathrm{min}$ is the boiler ramp lag, b_g is the gas bias disturbance, $y_0 = 0.124$ is the design carbon dioxide fraction and $y_t = 0.10$ is the fall in that fraction per unit fall in load, so a station at 60 % presents 40 % less gas at 11.9 % carbon dioxide rather than 12.4. What the host is then left with is one algebraic expression, Equation (7.13).

$$W_n = W_g\,\ell - k_s Q_a - w_c \dot{m}_{str} - W_p \frac{L}{L_0} - W_b\,\ell_c^3 \tag{7.13}$$

where $W_g = 580\,\mathrm{MW}$ is the gross output at full load $\ell = 1$, ℓ_c is the load clamped to the interval 0.05 to 1.15, $k_s = 0.175$ is the electricity lost per unit of extracted heat, $w_c = 0.378\,\mathrm{MJ/kg}$ is the compression and drying work charged on the stripper outlet mass flow $\dot{m}_{str}$ rather than on what the absorber took out, $W_p = 7.5\,\mathrm{MW}$ is the solvent pumping at design circulation and $W_b = 4.5\,\mathrm{MW}$ is the flue gas blower. At the design line-up the four terms are 65.3, 39.0, 7.5 and 4.5 MW, so the capture plant takes 116.3 MW off the station and leaves 463.7 MW, which is 0.3132 MWh/t of lost output per tonne of carbon dioxide captured, of which 0.1759 MWh/t is the steam alone. Amine leaving the top of the absorber is volatility times what the water wash does not catch, as shown in Equation (7.14).

$$\varepsilon = \frac{x_M\, p_M(T_1)}{P_{tot}} \frac{M_{MEA} \times 10^6}{V_{Nm}}\,(1 - \eta_w)\, f_{aero}\, \chi_w \tag{7.14}$$

where p_M is 0.0067 kPa at 40 °C and carries a 62 kJ/mol enthalpy of vaporization, so it roughly doubles every ten degrees, $M_{MEA} = 61.08\,\mathrm{kg/kmol}$, x_M is the effective liquid activity, $\eta_w = 0.90$ is the wash removal, $V_{Nm} = 22.414\,\mathrm{m^3/kmol}$, $f_{aero} = \max(0.2,\, 1 + 0.3[(G/G_{max})^2 - 1])$ is a velocity-dependent aerosol allowance and χ_w is the wash condition. Solvent strength falls with thermal and oxidative degradation and is restored by makeup, as shown in Equation (7.15).

$$\frac{\mathrm{d}\phi_a}{\mathrm{d}t} = -\frac{r_d}{100}\,\frac{2^{(T_S - T_d)/E_d}}{1440} + \frac{\dot{m}_{MEA}}{\sum V \rho\, w_{MEA}} \tag{7.15}$$

where r_d is the degradation rate in percent per day at the reference reboiler temperature $T_d = 120\,^\circ\mathrm{C}$, $E_d = 8\,\mathrm{K}$ is the rise that doubles it [14], and the denominator of the second term is the mass of amine in the whole circuit.

> **+ DID YOU KNOW? The Thermodynamic Floor of Capture**
>
> The minimum work to separate carbon dioxide from coal flue gas and compress it to 150 bar is 0.11 MWh per metric ton, and better processes and solvents are expected to bring the energy used down to about 0.2 MWh [9]. At its design line-up the plant of this chapter costs its host 0.3132 MWh of output per tonne captured.

7.3.7 Constants, the Opening Line-Up and the Numerics

Table 7.2 lists the constants. Those marked hidden are used by the model and never displayed, and they are the quantities a learner has to infer from a trend rather than read from a panel.

No scenario opens at a guessed steady state. The initializer bisects the whole recycle on one scalar, the lean loading leaving the reboiler, until it equals the lean loading arriving at the top of the absorber,

Table 7.2 Model Parameters. Visible Constants Appear in the Model Panel and Are Adjustable There. A Hidden Constant Reaches No Panel.

Symbol	Meaning	Value	Unit	Visibility
W_g	host gross output	580	MW	visible
G_{max}	flue gas at full load	21.0	kmol/s	visible
y_{in}	flue gas CO_2 fraction	0.124	—	visible
τ_ℓ	host load ramp lag	4.0	min	visible
k_s	output lost per unit duty	0.175	MWe/MWth	visible
w_c	compression work	0.378	MJ/kg	visible
W_p	pumping at design flow	7.5	MW	visible
W_b	blower at full load	4.5	MW	visible
c_{MEA}	cost of makeup amine	2.40	\$/kg	visible
V_s	holdup per absorber stage	120	m^3	visible
$K_G a$	transfer coefficient	0.21746	kmol/(s kPa)	visible
w_{MEA}	MEA mass fraction	0.30	—	visible
P_0^*	reference CO_2 pressure	0.30	kPa	visible
α_{max}	free amine exhausted	0.58	mol/mol	visible
n, m	loading exponents	2.0, 1.6	—	visible
ΔH_{abs}	heat of absorption	85 000	J/mol	visible
V_{str}	stripper and reboiler holdup	500	m^3	visible
φ	vapor surviving overhead	0.2144	—	visible
ΔT_{app}	cross exchanger approach	12.0	K	visible
P_{ref}	assumed stripper pressure	190	kPa	visible
τ_{reb}	reboiler temperature lag	4.0	min	visible
τ_Q	steam valve lag	3.0	min	visible
V_p	each pipe run volume	1170	m^3	visible
V_{tk}	tank capacity, each	8000	m^3	visible
—	lean tank level at design	63.75	%	visible
C_{obl}	rolling capture obligation	88	%	visible
W_{roll}	averaging window	240	min	visible
ε_{max}	amine emission limit	4.0	mg/Nm^3	visible
—	pump suction limit	5	%	visible
—	tank full limit	95	%	visible
$T_{s,max}$	degradation limit	127	°C	visible
c_{CO_2}	price of an emitted tonne	95	\$/t	visible
c_{obl}	penalty per tonne short	400	\$/t	visible
τ_{cm}	stack analyzer response	2.0	min	visible
τ_{tm}	thermocouple response	0.8	min	visible
τ_{lm}	level transmitter response	0.5	min	visible
P_{tot}	absorber pressure	101.3	kPa	hidden
n_w	wetting exponent	0.30	—	hidden
E_a	mass transfer activation	21 000	J/mol	hidden
$c_{p,g}$	gas heat capacity	33.0	kJ/(kmol K)	hidden
λ_w	latent heat of water	43.5	kJ/mol	hidden
U	shell loss per stage	12	kW/K	hidden
ρ	solvent density	1050	kg/m^3	hidden
c_p	solvent heat capacity	3.60	kJ/kg/K	hidden
α_0	reference loading	0.35	mol/mol	hidden
T_0	equilibrium reference	313.15	K	hidden
a_w	water activity	0.85	—	hidden
T_d, E_d	degradation reference	120, 8.0	°C, K	hidden
η_w	water wash removal	0.90	—	hidden
x_M	MEA liquid activity	0.056	—	hidden
y_t	CO_2 fraction tilt with load	0.10	—	hidden

and then fills the pipes and the tanks with whatever last passed through them. In manual operation at the design line-up, the largest relative state derivative is 1.8×10^{-15}. The initializer reads the manual circulation, the manual duty and the level setpoint in force, so a value chosen before Run sets where the plant starts, and the same value carried through Reset gives the same run. A second Manual helper sets the duty that holds a nominated capture rate at the present circulation. Both helpers solve the declared design plant, so a plant the learner has changed shows as a missed target. Integration uses

Runge-Kutta of fourth order at a fixed 0.25 min (15 s) step, controller execution falls every 0.5 min (30 s), and the historian keeps two samples per simulated minute (Section 1.6.1). Each run exports 121 columns, with the true lean loading beside the inferred one and every commanded flow beside the flow actually delivered. A generator seeded at 20260913 supplies every random sequence, so a noisy scenario replays identically. The fastest setting of the speed selector runs 120 simulated minutes per second of wall time, so a sixteen-hour price scenario finishes in eight seconds, and the setting changes nothing about the trajectory.

7.3.8 Assumptions

The circuit is where the detail went, and the chemistry is thin, so every simplification below breaks down somewhere.

Three absorber stages rather than a packed height. Three stages are enough for a genuine interior temperature bulge, because the evaporation term in (7.5) is strongly nonlinear and pins the profile. They are not enough to place the bulge in height, and they leave the top stage several degrees hotter than a real twenty meter packing would be, so the amine emission predicted by (7.14) is pessimistic.

The equilibrium is a correlation. Equation (7.1) has the right stoichiometry, the right runaway near $\alpha = 0.58$ and the right heat of absorption, and the design point is anchored to it, but it is not an electrolyte activity model and the bicarbonate route real monoethanolamine takes above $\alpha = 0.5$ is absent [12]. Absolute pressures away from the design region are indicative rather than predictive.

The stripper is one lumped reboiler. Everything between the reboiler and the condenser appears only as φ in (7.10), calibrated so that the design point gives 3.619 MJ/kg and not derived from a column. A model that had to predict the effect of stripper pressure or feed location on specific duty would need the stages [3].

The water balance is not closed. A real plant loses water from the absorber top and gains it in the stripper overhead, and holding 30 wt% needs a wash section and a makeup loop. Solvent strength here moves only through (7.15), so the water balance excursions that are a routine operating problem cannot appear.

The host is a fixed gross output and a fixed conversion factor. A real low-pressure turbine loses efficiency as it is throttled, so the penalty per megawatt of extracted steam rises with the extraction and there is a maximum extraction beyond which the turbine cannot run at all [15]. Both omissions make (7.13) slightly optimistic about very high capture rates and about very deep storage.

The tanks are perfectly mixed and the pipes perfect plug flow. There is no dispersion in either, so a step arrives more sharply here than it would in the field. That is the honest direction to err for this application, because it makes the delay visible rather than hiding it inside a lag.

Aerosol emission is one allowance. Real amine aerosol depends on sulfuric acid mist nuclei in the flue gas and can be an order of magnitude worse than volatility alone, which is why real plants fit demisters that have no equivalent in (7.14).

7.4 What the Recycle Does to the Gain

7.4.1 One Scalar Closes the Plant

Everything downstream of the absorber is transport of loading. The pipes and the tanks change when a parcel of solvent arrives and never what is in it, so the steady state of the whole circuit is fixed by one scalar, the requirement that the lean loading leaving the reboiler equal the lean loading arriving at the top of the absorber. Whether that fixed point moves when the circulation moves is the entire question, and the answer depends on what else is being held. With the reboiler duty frozen the same 373 MW has to regenerate whatever flow is sent to it, so a circulation 15 % higher settles at a lean loading of 0.2630 instead of 0.2375. The solvent arriving at the top of the absorber is a worse absorbent than it was, and it takes back most of what the extra flow gained.

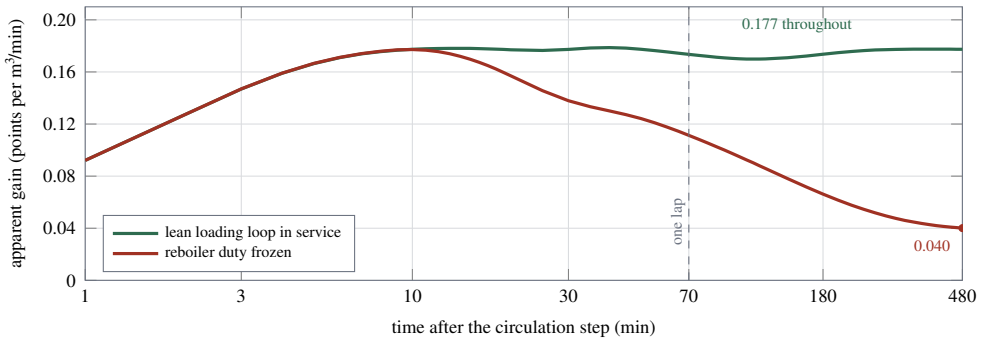

Figure 7.2 Apparent gain of the capture loop against observation time, measured on the running model for a circulation step of 130 to 150 m³/min at $t = 20$ min, with the stripper feed moved with it. The two runs are the same step and differ only in what is regulated underneath. With the reboiler duty frozen at 373.3 MW the gain peaks at 0.177 at ten minutes and decays to 0.040 by eight hours as the returning solvent grows less lean. With the lean loading loop closed underneath the reboiler duty rises from 372 to 401 MW and the gain is 0.177 at ten minutes and 0.177 at eight hours, never leaving the band 0.170 to 0.179 after ten minutes. Both curves are in points of capture per unit of circulation.

7.4.2 The Measured Gain and Its Decay

Figure 7.2 is the apparent gain of the capture loop against the time at which it is read, for the same 20 m³/min step under two conditions.

With the duty frozen the gain reaches 0.177 at ten minutes and falls to 0.040 at eight hours, a loss of 77 %, and the decay is not complete until several laps have passed. With the lean loading loop in service the reboiler duty climbs from 372 to 401 MW as the circulation rises, the returning solvent is the same solvent it always was, and the gain is 0.177 at ten minutes and 0.177 at eight hours. The two curves agree to four figures for their first two minutes, 0.0921 and then 0.1268, because for that long the plant is only an absorber, and they separate from the moment the first affected solvent reaches the reboiler. Scenario 1 makes the frozen-duty step by hand, and its capture readings at 10, 30, 90, 180 and 480 min after the step trace the same fall from 0.177 to 0.040. The circulation station is locked on basic loops, so the held arm cannot be stepped by hand, but in scenario 3 the ratio of capture change to circulation change stays between 0.166 and 0.170 from 180 to 420 min with the lean loading loop in service.

Two consequences follow for the controller. Tuning the capture loop against the gain read at ten minutes and running it with the duty frozen produces a loop whose true gain is 4.4 times smaller than the one it was tuned for, so it never reaches setpoint and the integral term walks the circulation to its stop. Scenario 3 run with the reboiler loop gain at zero shows exactly that, with the capture rate peaking at 93.43 % against a setpoint of 94 and the circulation on its 170 m³/min stop from $t = 316$ min, 256 minutes after the step. Tuning against the settled gain of 0.040 and then closing the inner loop produces the opposite, a loop 4.4 times too aggressive on a plant with a seventy minute memory. The inner loop is not there because it is fast. It is there because it makes one number out of two.

🕮 **Pause.** Try Exercise 7.1 before reading on.

7.4.3 The Lap, Item by Item

The lap time is the circuit inventory divided by the circulation, and the inventory is a sum of holdups no single unit dominates. At design the two lines contribute 9.00 min each, the lean tank 39.2 min, the rich tank 6.2 min, the reboiler 3.8 min and the absorber 2.8 min, which totals 70.0 min. Only the two lines behave as delays. Stepping the circulation from 130 to 155 m³/min moves the loading leaving

Table 7.3 Loop Pairings, Shipped Tunings and the Plant Each Controller Faces

CV	MV	Action	K_c	τ_I (min)	Plant
C	L	reverse	3.5	20	70 min lap
T_S	Q	reverse	30	12	7 min lag
L_R	F	direct	1.5	240	integrator

the absorber within 0.5 min, moves what reaches the rich tank not at all until 8.0 min and then sharply, reaches the reboiler at 9.5 min and reaches the top of the absorber at 21.5 min. Those are the times a student reads off the half-minute trend, taking a move of 2×10^{-4} mol/mol as the first sign. At the 0.25 min integration step the same four moves fall at 0.25, 7.75, 9.25 and 21.25 min, and 7.75 is the first step after the 7.55 min pipe transit. The extra 1.5 min between the pipe and the reboiler is the rich tank, which is stirred and begins to move as soon as its inlet does. A step test on the absorber alone finds 0.5 min and concludes that the plant is fast.

7.5 Controller Design

7.5.1 The Implemented PI Algorithm

All three loops run the velocity-form PI law of Section 1.6.2, Equations (1.14) to (1.17), executed every $\Delta t_c = 0.5$ min. Its action takes the value $+1$ for a reverse-acting pairing and -1 wherever the output has to rise with the measurement. The optional first-order derivative filter is off in every loop, so the rate term would see the raw measurement, and every loop ships with $\tau_D = 0$, leaving the derivative path present and idle. The back-calculation time T_t is left at its default of τ_I, and the integral used at execution k is the one accumulated through execution $k-1$. Saturation reaches the emitted output and stops there, never the running sum of increments, so the emitted output equals the positional law only away from the limits, and nothing else touches the integral except back-calculation, which is dormant unless there is integral action to unwind [19]. Manual solves the same relation for the integral and matches the output to the operator's station, while in automatic mode the three manual stations are set to the controller outputs after every execution, so a transfer in either direction lands without a bump. Rate limiting is absent from all three loops, because the steam valve lag and the two pipe runs already bound how fast anything can move.

7.5.2 Three Pairings, Three Time Scales

Table 7.3 lists the pairings, the shipped tunings and the plant each controller has in front of it.

Each gain carries the units of its own pairing, 3.5 m³/min of circulation per point of capture, 30 MW of duty per °C of reboiler, and 1.5 m³/min of stripper feed per percent of level. Output limits are 55 to 170 m³/min on the circulation, 0 to 560 MW on the duty and 0 to 175 m³/min on the feed, and the tuning sliders span $0 \le K_c \le 25$ with $2 \le \tau_I \le 400$ min on the capture loop, $0 \le K_c \le 200$ with $1 \le \tau_I \le 200$ min on the reboiler, and $0 \le K_c \le 20$ with $10 \le \tau_I \le 2000$ min on the level. Measurements are filtered at 2.0, 0.8 and 0.5 min, and the stack analyzer is the only fast instrument in the capture loop. Only the level loop is direct acting, because it is the only one of the three whose process gain is negative, and getting that sign wrong is invisible until a tank overflows.

The capture loop is tuned against the constant 0.177 of Figure 7.2 rather than against the 0.040 the plant settles at with the duty frozen, which is legitimate only because the loop underneath is in service. Breaking that assumption is a scenario rather than a footnote. Setting the reboiler gain to 3 MW/°C with a 120 min reset and repeating the same setpoint step leaves the lean loading drifting out to 0.2715 while the capture controller chases it with circulation until the pump sits on its 170 m³/min stop, and the capture rate is still only 93.03 % against a setpoint of 94 when the setpoint returns. Nothing is wrong with the loop that is failing.

The reboiler loop is where the inference lives. A lean loading setpoint is converted to a temperature setpoint by inverting (7.11) at the assumed reference pressure and clamping the result to 100 to

126 °C, and the loop then compares two temperatures. The level loop is deliberately slack. One cubic meter per minute of imbalance moves the level 0.0125 %/min, so a tank of this size is an integrator with a very small gain, and there is no reason to hold it tightly.

7.5.3 Two Availability Factors

What the controllers ask for is not always what the plant can do, and the limit is a tank end rather than a valve. Both factors are smooth, because a hard cut-off at a level alarm gives the fixed-step integrator a discontinuous right-hand side and turns the level trace into a staircase, as shown in Equation (7.16).

$$\begin{aligned} s(x) &= \operatorname{clamp}\left(x / x_{lo}, 0, 1\right) \\ r(x) &= \operatorname{clamp}\left(\frac{100 - x}{100 - x_{hi}}, 0, 1\right) \\ L' &= L_{cmd}\, s(L_L), \qquad F' = F_{cmd}\, s(L_R) \\ L_a &= F' + (L' - F')\, r(L_R) \\ F_a &= L_a + (F' - L_a)\, r(L_L) \end{aligned} \tag{7.16}$$

where L_L and L_R are the two tank levels in percent, $x_{lo} = 5$ is the level at which the pumps lose suction and $x_{hi} = 95$ is a full tank. A lean tank on its floor cuts the circulation whatever the capture controller has to say about it, which is the constraint that ends every storage maneuver. The factors act after the controller output, and the output limits of the PI law, Equation (1.16), are the pump's 55 to 170 m³/min rather than the delivered flow, so the capture integral keeps winding while a starved pump falls short. In scenario 11 the command first reaches its 170 m³/min stop at $t = 82$ min and spends 119 min there before the price breaks at 420 min, while the pump delivers 105 to 148 m³/min.

7.5.4 The Economic Supervisor

Above the three loops sits an optional supervisor whose entire content is that the capture rate is a law and the energy is a price. It moves two setpoints and never the capture setpoint, as shown in Equation (7.17).

$$\begin{aligned} z &= \operatorname{clamp}\left(\frac{p - p_{lo}}{p_{hi} - p_{lo}}, 0, 1\right) \\ u &= \operatorname{clamp}\left(\frac{b - (\bar{C} - C_{obl})}{b}, 0, 1\right) \\ z_e &= z\,(1 - u) \\ L_{R,\mathrm{SP}} &= \ell_c + z_e\,(\ell_d - \ell_c) \\ \alpha_{\mathrm{SP}} &= \alpha_c + z_e\,(\alpha_d - \alpha_c) \end{aligned} \tag{7.17}$$

where p is the electricity price, $p_{lo} = 55$ \$/MWh and $p_{hi} = 175$ \$/MWh bound the price index z, $\bar{C}$ is the rolling-average capture over the last 240 min, C_{obl} is the obligation, $b = 1.2$ points is the credit the supervisor is allowed to spend, and the four targets are $\ell_c = 6\,\%$, $\ell_d = 56\,\%$, $\alpha_c = 0.218$ and $\alpha_d = 0.252$. The rolling average is computed the way a regulator computes it, as carbon dioxide kept out of the stack over carbon dioxide presented across a fixed window, and both z and u are published on every architecture so that a run on basic loops shows what a supervisor would have done.

The product in the third line of (7.17) is the design decision worth arguing about. Multiplying the price index by the remaining credit removes the price signal's authority as the obligation comes under pressure, rather than outweighing it. A layer that added the two would carry an implied exchange rate between dollars and permit compliance, and at a high enough price it would sell the permit.

🕮 **Pause.** Try Exercise 7.2 before reading on.

7.5.5 Model Predictive Control

A fourth control setting replaces the supervisor and the two outer loops with a receding-horizon optimization [20] of the economic kind that industrial practice places above the regulatory layer [21]. Section 1.6.3 describes the shared option, and this one differs from it in its economic objective, its six-state model and its obligation rows. Every $T_s = 10\,\text{min}$, twenty controller executions, it chooses the circulation L, the lean loading setpoint α_{SP} and the stripper feed F and holds them until the next decision. The reboiler temperature loop of Table 7.3 keeps acting beneath the plan with its shipped tuning, turning α_{SP} into a temperature setpoint through (7.11) at 190 kPa as it does under the supervisor. The capture and level controllers track the applied circulation and feed without a preload, so a transfer either way is bumpless, and the factors of (7.16) and every alarm stay in force.

Controller model. The plan predicts with six states written from the declared defaults of Table 7.2 and the design disturbances, never from the running plant, as shown in Equation (7.18).

$$\begin{aligned}
\tau_T \frac{\mathrm{d}\hat{\alpha}_S}{\mathrm{d}t} &= \alpha_{\text{SP}} - \hat{\alpha}_S \\
\frac{V_L}{F_m}\frac{\mathrm{d}\hat{\alpha}_L}{\mathrm{d}t} &= \hat{\alpha}_S - \hat{\alpha}_L, \qquad \hat{\alpha}_{in}(t) = \hat{\alpha}_L(t - T_s) \\
\tau_C \frac{\mathrm{d}\hat{C}}{\mathrm{d}t} &= C_m + g_L\,(L - L_m) + g_\alpha\,\delta\hat{\alpha}_{in} - \hat{C} \\
\frac{\mathrm{d}\hat{L}_R}{\mathrm{d}t} &= \frac{100\,(L - F)}{V_{tk}} \\
\hat{Q} &= Q_m + q_F\,\delta F + q_\alpha\,\delta\hat{\alpha}_S + q_R\,(r_\alpha\,\delta\hat{\alpha}_{in} + r_L\,\delta L)
\end{aligned} \tag{7.18}$$

where the subscript m marks a value measured or estimated at the decision, δ a deviation from it, $\hat{\alpha}_S$, $\hat{\alpha}_L$ and $\hat{\alpha}_{in}$ the lean loadings at the reboiler, in TK-402 and at the absorber inlet, and $\hat{Q}$ the duty the temperature loop will demand. The lag $\tau_T = 9.4\,\text{min}$ folds the 6.4 min lag and 3.0 min dead time identified from α_{SP} to the inferred loading with the temperature loop closed, and $\tau_C = 2.1\,\text{min}$ is the capture lag identified from circulation steps with the cascade closed. A second identification with the outer loops pinned gives 5.2 to 6.8 min plus 3.0 to 3.25 min and 1.3 to 2.0 min, with residuals of 2.9 to 4.5 % of the response. The tank constant V_L/F_m is recomputed at each decision (39.2 min at design), and the lean line is one decision of transport. The gains g_L and g_α are re-read at every decision from the declared absorber at the circulation where it gives the measured capture, 0.241 points per m^3/min and -139 points per mol/mol at design, with a chord from 130 to 150 of 0.179 against the 0.177 of Figure 7.2. The rich tank is the integrator of (7.7). The duty is the declared stripper of (7.10) linearized at design, $q_F = 2.87\,\text{MW}$ per m^3/min, $q_\alpha = -2114\,\text{MW}$ and $q_R = 1311\,\text{MW}$ per mol/mol of lean and rich loading, with a curvature $q_{\alpha\alpha} = 1.86 \times 10^4\,\text{MW}$ per $(\text{mol/mol})^2$ and a rich loading that follows the lean one at $r_\alpha = 0.676$ and falls by $r_L = 0.00105\,\text{mol/mol}$ per m^3/min. There are no recycle states, for the reason of Section 7.4. With the temperature loop acting, a 20 m^3/min step reads a capture gain of 0.1775 at 10 min and 0.1791 at 480 min, and pricing the rich leg's duty 10 or 20 min later changes scenario 10's margin by \$6 and \$13. Both gains are read with the inlet loading capped at 0.32 mol/mol.

Measurements and offset. Each plan starts from the stack analyzer, the thermocouple converted to loading by the same inverse of (7.11) at P_{ref} that the temperature loop uses, so a pressure drift misleads the plan as it misleads every loop, and the two level transmitters, LT-401 and the LT-402 of Table 7.1, both through the 0.5 min lag and without noise. A soft sensor carries $\hat{\alpha}_L$ and $\hat{\alpha}_{in}$ through the declared tank and line. Offset is removed by an output bias, not a Kalman filter. Because C_m anchors the capture equation, a host load, cooling water temperature, packing condition or solvent strength the model omits appears within the absorber's two minutes as a step that the bias holds over the horizon [22].

Horizon and objective. The horizon is $N = 144$ decisions (24 h) in twelve move blocks, two of a single decision followed by 2, 2, 4, 6, 6, 12, 12, 24, 24 and 50, with three handles and one obligation

slack per block, so the quadratic program has 48 variables. The objective is the next day's cost in dollars, Equation (7.19).

$$\begin{aligned} J = \sum_{k=1}^{N} \Big[\pi_k \Big(k_s \hat{Q}_k + k_c \dot{n}_k + W_p \frac{L_k}{L_0} \Big) + \frac{\pi_k k_s q_{\alpha\alpha}}{2} (\delta_0 \hat{\alpha}_{S,k})^2 \\ + w_C (\hat{C}_k - C_{\mathrm{SP}})^2 + w_I (\hat{L}_{R,k} - L_{R,\mathrm{SP}})^2 \Big] \\ + P (\hat{L}_{R,N} - L_{R,\mathrm{SP}})^2 + \sum_j \| \Delta u_j \|_S^2 + \sigma s \\ + \sum_b \rho_b (\varepsilon_b^2 + \varepsilon_b) \end{aligned} \tag{7.19}$$

Here $\pi_k = p_k T_s / 60$ is the published price of decision k in dollars per megawatt held for one decision, k_s, W_p and L_0 are those of (7.13), $k_c \dot{n}_k$ is the compression power on the carbon dioxide the plan predicts the stripper releases, and δ_0 is a deviation from the design lean loading of 0.2375. The weights are $w_C = \$300$ per point^2 per decision, $w_I = \$0.02$ per point^2 per decision and $P = \$2$ per point^2, and the move prices S are \$0.05 per $(\mathrm{m}^3/\mathrm{min})^2$ on L and F and \$3 per $(0.01\,\mathrm{mol/mol})^2$ on α_{SP}, times a move suppression of 1. The inventory weight treats stored rich solvent as a loan, and without it the plan ends scenario 10 at 61.4 % rich instead of 11.0, earning \$1,278 more by leaving that regeneration to the next day.

Constraints. The handles are hard-bounded by the shipped clamps, 55 to 170 m³/min, 0.15 to 0.34 mol/mol and 0 to 175 m³/min. The rich level is soft between 7 % and the smaller of 90 % and the level that leaves the lean tank at 10 %, at $\rho = \$200$ per point^2 and per point, and the duty is soft 15 MW below its 560 MW maximum at \$20 per MW^2 and per MW, each with one slack over the horizon judged at every second decision. The obligation is a hard row for every four-hour window ending inside the horizon, Equation (7.20),

$$\frac{100\, n_j^{cap} + r\, T_s \sum_{k=k_0}^{j-1} \hat{C}_{k+1}}{n_j^{in} + r\, T_s\, (j - k_0)} \;\geq\; C_{obl} + m - s \tag{7.20}$$

where n_j^{in} and n_j^{cap} are the carbon dioxide presented and captured over the part of window j already past, from the regulator's record the supervisor reads, r is the arrival rate over the last decision, $k_0 = \max(0, j - W_{roll}/T_s)$, $m = 0.3$ points is the margin kept, and the slack $s \geq 0$ costs $\sigma = \$2 \times 10^4$ per point-decision (0.69 t at full load), so the row holds whenever it can be met and the problem stays feasible when it cannot. The rows never bind on scenarios 9, 10 and 11 or the challenge (3.375 point-decisions to spare at the tightest on scenario 10), and act when scenario 11's obligation is raised to 89.7 %, holding a rolling minimum of 89.993 % against 89.884.

Preview. The price over the horizon is the scenario's scripted schedule, read as a day-ahead market's hourly prices published the day before, and a price set by hand is not on it. Switching the preview off leaves only the present price over the whole horizon, which is all the supervisor reads. A real plant knows the day-ahead clearing but not the real-time departures from it [23], and nothing announces a dispatch, so there is no load preview.

Acceptance and fallback. A plan is applied only when the solver reports an optimum, everything it returns is finite, no planned command leaves its bounds and its first command is the one being applied. Otherwise the historian is marked, the supervisor takes over with the loops that have been tracking, its two setpoints glide from the plan's values at 1 % and 0.001 mol/mol per minute, until a valid plan reclaims the three handles. Six forced failures across scenario 14 moved the feed by at most 0.75 m³/min at the hand-over, where the supervisor's level target applied at once would have stepped it by up to 84 m³/min. The first plan is solved at $t = 0$. A plan takes 2.71 to 2.77 ms in a warmed-up Chrome page, 0.03 % of one processor core at the slowest playback, a simulated minute each second, and 3.3 % at the fastest, 120, and a frame guard that ends a frame's stepping after 25 ms never engaged at the top speed.

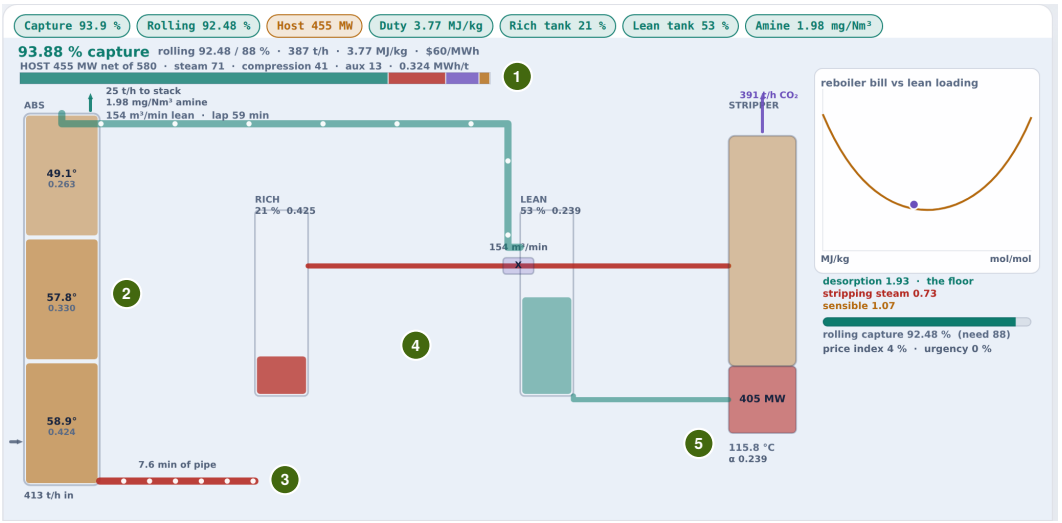

Figure 7.3 Scenario 3 at $t = 240$ min, three hours after the capture setpoint stepped from 90 to 94 % with both inner loops in service, cropped to the plant view. (1) The capture rate has followed the setpoint to 93.88 %, settling into its one-point band 42.5 min after the step, and the host output of 454.6 MW at a specific duty of 3.77 MJ/kg is on the way from 463.9 MW at 90 % capture to the 453.6 MW reached by $t = 420$ min, which is the price of four more points of capture. (2) The three absorber stages, each with its temperature and loading. (3) The rich line, whose markers are equally spaced parcels of solvent. (4) The rich and lean tanks, one inventory seen from opposite ends. (5) The reboiler, whose duty has followed the circulation from 372 to 405.8 MW while the circulation climbed from 130 to 153.6 m^3/min, which is the second loop paying for what the first one spent.

Three models kept apart. The plant is the sixteen-state model of Section 7.3 with its recycle, the true stripper pressure and the solvent strength. The identified model is the set of closed-loop lags read from step tests on it. The controller's model is (7.18) at the declared defaults, relinearized at each decision, so a parameter the learner changes is model mismatch rather than knowledge.

7.6 Results

The simulator carries fifteen scenarios in the order a laboratory would take them: the absorber by hand (scenario 1), the circulation step timed on four loadings (2), capture on circulation with and without the loop underneath (3), the energy penalty on the host (4), lean loading (5), the inferential measurement (6), the surge tank (7), a host load ramp (8), a price spike (9), a day of prices, flat and shifted (10), the rolling obligation (11), solvent degradation (12), a cold night and a hot afternoon (13), the same price day under the predictive controller (14) and the graded challenge. Each one pins the starting mode, the architecture, all thirteen tuning constants (six for the loops and seven for the supervisor), the disturbances it names and any timed event, and it declares an objective, a predicted observation and a completion rule ahead of the run (Section 1.6.4).

7.6.1 Capture Loop Performance

Table 7.4 reports the measures the application computes for itself on four runs of the same three loops. Integrated absolute error accumulates from the start of each run, so the observation window is stated with every column and the columns are read against each other rather than in isolation.

The first two columns separate a controller fault from a structural one. Both carry identical capture tuning and both meet the same four point setpoint step, and both report a settling time of 42.5 min

Table 7.4 Measured Capture-Loop Performance. Every Measure Accumulates From $t = 0$ to the Stated Read Time, and All Four Runs Use the Same Capture Loop Tuning of $K_c = 3.5$ and $\tau_I = 20$ min.

	SP step shipped	SP step weak loop 2	P_{str} drift	load ramp
Scenario	3	3	6	8
Read at (min)	420	420	500	780
K_c reboiler (MW/°C)	30	3	30	30
τ_I reboiler (min)	12	120	12	12
IAE (% min)	128.5	259.2	140.0	709.1
Max. deviation (%)	3.08	3.07	1.20	7.83
Circulation travel (m^3/min)	31.4	47.3	24.6	125.2
Time at a stop (%)	0.0	15.5	0.0	0.0
Settling (min)	42.5	42.5	—	—
Final capture (%)	94.01	93.03	90.01	90.03
Final L (m^3/min)	153.9	170.0	153.9	130.5
Final α_S (mol/mol)	0.2379	0.2715	0.2740	0.2379
Host output (MW)	453.6	456.7	462.5	461.4

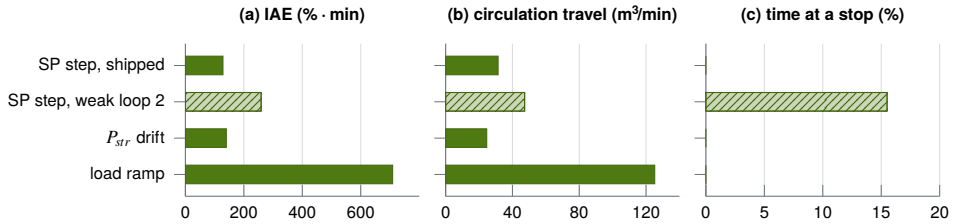

Figure 7.4 Three rows of Table 7.4 as bars, the run with the weak reboiler loop hatched. Notice that the second run carries the same capture tuning as the first and still doubles its error, travels further and is the only run that sits against the pump stop, while the load ramp makes the largest error and the most travel without ever reaching a stop.

because the first approach to the band is the same. What differs afterward is that the second column never holds. Its integrated error is twice the first at the same read time, 259.2 against 128.5 % min, its circulation travel is 47.3 m^3/min against 31.4, it spends 15.5 % of the run against the pump stop, and it finishes almost a point short of setpoint with the lean loading at 0.2715 instead of 0.2379. No adjustment to the capture controller repairs that, because the variable that moved is not in its loop. Figure 7.4 draws three rows of the table as bars.

The third column is the inferential failure. The pressure rise lifts the reboiler temperature by 2.1 °C, the reboiler loop has it back within 0.1 °C of its 115.83 °C setpoint 19 min after the rise, nothing alarms, and the capture controller does its job, dipping to 88.80 % and returning to setpoint. The true lean loading nonetheless settles at 0.274 mol/mol while the inference on the screen still reads 0.238, and the only evidence anywhere on the plant is 24 m^3/min of extra pumping and 1.4 MW off the host. That shape, in which the loop holds its measurement perfectly while the inferred quantity walks away, is the standard failure of an inference built on a secondary measurement [13], the same shape Volume 5, Chapter 10 meets when an outlet temperature stands in for powder moisture, and the defense is a pressure-compensated calculation or an alarm on the pressure rather than a better controller.

🕮 **Pause.** Try Exercise 7.3 before reading on.

The fourth column is the load ramp, and it is the largest error in the table because feedback alone

Table 7.5 Level Loop Gain Against What the Reboiler Sees, Scenario 7, Measured Over $60 \le t \le 480$ min With $\tau_I = 240$ min Throughout

K_c (m³/min per %)	0.5	1.5	4	8
Stripper feed range (m³/min)	105–130	55–133	0–136	0–142
Reboiler duty range (MW)	304–372	176–380	2–391	0–427
Time to reach setpoint (min)	223	110	55	44

has to reject it through the lap. The host returns to full load in four minutes, the capture rate falls to 82.17 % within 6.0 min because the solvent standing in the absorber was sized for sixty percent load, and the rate climbs back above 89.5 % 108 min after the return to full load. In the other direction the plant is over-circulated rather than under-circulated and the capture rate peaks at 96.90 %, which costs nothing except steam. Feedforward from the host load signal is the first thing a real plant adds, and it is deliberately absent here so that the recycle is visible.

7.6.2 The Surge Tank Is Not a Setpoint

Moving the rich tank level setpoint from 10 to 60 % cuts the stripper feed, and the reboiler duty follows it down because the duty needed is very nearly proportional to the solvent being regenerated. Half an hour into the move the feed is 81.1 m³/min, the duty is 236.8 MW and the host is making 501.6 MW instead of 463.9. That is 38 MW for as long as the tank is filling, and it is the mechanism the whole economics rests on. It is also strictly temporary. The level crosses its setpoint 110 min after the move, overshoots to 67.4 %, and when the setpoint returns to 10 % the feed runs hard and the output falls below where it started while the backlog is worked off.

Table 7.5 is the level gain sweep over the filling period. The tighter tunings track the setpoint faster and pass the whole of the imbalance to the reboiler while doing it.

Raising the gain from 1.5 to 8 cuts the time to setpoint from 110 to 44 min, and widens the reboiler duty swing from 204 MW to 427 MW. A surge tank exists to absorb imbalance so that the downstream unit does not have to see it, and a level controller tight enough to hold a setpoint has thrown away the only thing the tank was bought for. The loop must still keep both ends off the tank. As the rich tank fills the lean tank empties by exactly the same volume and reaches 6.3 % at the worst of the shipped tuning, against a pump suction limit of 5 %.

> **+ DID YOU KNOW? A Level Controller That Shuts Off the Reboiler**
>
> From the simulator: at the tightest level gain of 8, scenario 7's setpoint move swings the reboiler duty from 0 to 427 MW, shutting the reboiler off on the way. At a gain of 0.5 the duty stays between 304 and 372 MW, though the tank takes 223 min instead of 44 to reach its new level (Table 7.5).

7.6.3 What the Energy Costs and When to Spend It

Capture rate is bought in megawatts and the price rises with the rate. Moving the setpoint from 90 to 95 % takes the host from 463.9 to 450.7 MW, a cost of 13.2 MW, moves the circulation from 130 to 164 m³/min and the specific duty from 3.61 to 3.86 MJ/kg. Coming down from 90 to 85 % returns only 7.8 MW, so the last five points cost nearly twice what the five below them do. The convexity in the other variable is stronger, since the six equal loading steps of Section 7.3 run from 38 to 137 MW. Both directions away from the design lean loading are worse for the host, because a leaner solvent needs less circulation and a richer one needs more, and the sum has a shallow optimum near the design point that neither column can find alone.

The only genuine freedom is timing [16]. Scenario 9 puts the supervisor on a five-hour peak at $210 a megawatt-hour. It raises the rich tank setpoint to 72 %, the level controller cuts the feed, the reboiler falls to 209.4 MW and the host output goes from 451.6 to 508.8 MW half an hour into the peak. The carbon dioxide is still being captured, because the absorber has not changed at all. The

limit then arrives, and it is not the tank an operator watches. The rich tank never passes 70.5 %, while the lean tank is at 8 % 90 min into the peak and on its 5 % floor by 98 min, after which the circulation is starved and the feed has to come back up to whatever the absorber is sending. Storage bought ninety good minutes out of a five-hour peak, and the swing available is the lean tank's 4700 m^3.

> **TRY IT IN THE APP 9 · A Price Spike, and What the Tanks Are Actually For**
>
> Load scenario 9 and press Run. At 60 min power jumps to $210 a megawatt-hour for five hours. Watch the supervisor raise the rich tank setpoint, the stripper feed fall and the host output climb from 451.6 to 508.8 MW. Then watch the lean tank: it is on its 5 % floor 98 min into the peak, and from there the starved circulation ends the storage.

Over the 960 min price day of scenario 10 the supervisor sells $580,549 of electricity against $575,925 on basic loops, a gain of $4,624 or 0.80 %, and once the carbon price on the escaped carbon dioxide is charged its margin is $5,063 or 0.99 % higher. It does that while capturing 90.068 % of the carbon dioxide presented against 89.998 % and finishing with the rich tank at 4.4 % rather than 10.0 %. Both halves of the fairness test pass, so the money is real. It is also thin, and an arbitrage that finished with a full tank would have borrowed the money rather than made it.

Scenario 11 asks what happens when the price outlasts the tanks. Seven hours at $200 against an 88 % obligation on a four-hour rolling average puts the lean tank on its floor within 73 min, drives the instantaneous capture rate down to 77.96 % and the rolling average to 88.51 %. The obligation urgency of (7.17) comes off zero at 82 min, the price begins losing authority, and the plant settles into an uncomfortable equilibrium with the rich tank near 70 % and the lean tank near 4 %. The run finishes 0.00 t short of the obligation, which is the only acceptable answer. The timing problem underneath is worth naming. A four-hour average on a plant with a seventy minute lap means that by the time the average has moved, the decision that moved it is hours old and the correction will take hours more, and the credit band exists to buy that lag back.

7.6.4 Two Faults That Hold Every Setpoint

Degrading the solvent at 6 % per day moves no controlled variable at all. Over eleven and a half hours the capture rate stays at 89.98 % and the reboiler temperature at 115.83 °C, while the strength falls from 1.000 to 0.981, the circulation rises from 130.0 to 132.6 m^3/min and the duty from 372.2 to 374.8 MW, because the sensible term of (7.10) is paid on the whole flow. The host loses 0.7 MW with nothing to show for it on any trend an operator watches. Holding the strength takes 90 kg/min of fresh amine, which at $2.40 a kilogram is $12,960 an hour and an argument for fixing the degradation rather than out-running it [14].

The weather fault is a permit rather than a cost. At 30 °C cooling water the equilibrium pressure over the solvent falls, the driving force widens all the way up the column, and 90 % capture costs 114.7 m^3/min instead of 130, worth 3.8 MW on the host's meter. At 49 °C the same setpoint costs 159.0 m^3/min and 455.3 MW, and the absorber profile lifts from 52.8, 60.4 and 58.5 °C to 59.3, 65.5 and 64.5 °C. Amine volatility roughly doubles every ten degrees, so the emission rises from 1.33 to a peak of 4.60 mg/Nm^3 and the 4.0 mg/Nm^3 permit is breached 2.0 min after the water warms. The capture rate dips to 82.66 % and is back within half a point of setpoint 129 min after the warming, but the emission stays between 3.99 and 4.16 mg/Nm^3 from 500 min to the end of the hot spell and finishes at 4.01 mg/Nm^3, on the permit, with the capture rate held. On a hot day the binding constraint on a capture plant is the amine limit rather than the carbon dioxide one.

7.6.5 The Graded Challenge

The challenge runs sixteen hours with the price from $30 to $195, a five degree warming of the cooling water, the host backed down to 80 % across the start of the peak, and gentle degradation throughout. It starts on basic loops with a capture reset of 250 min on a plant whose lap is seventy, and it asks for the capture rate held within ±2 points through the last third, the rolling average at or

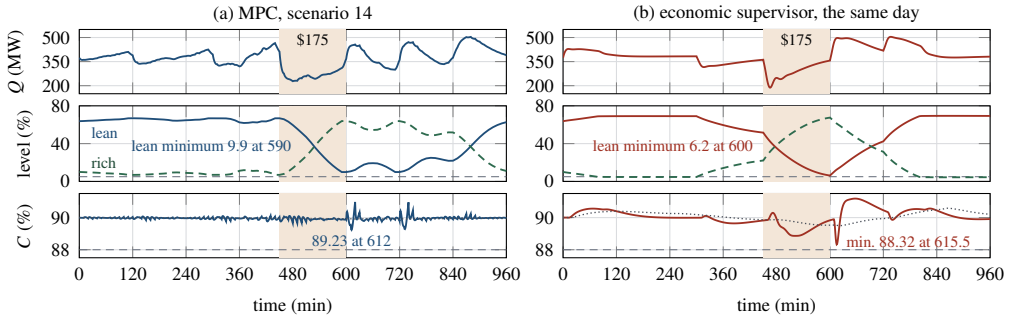

Figure 7.5 The price day of scenarios 10 and 14 over its 960 min under (a) the predictive controller of scenario 14 and (b) the economic supervisor. The prices are \$30, 55, 95, 175, 105, 42 and 24 from 0, 120, 300, 450, 600, 720 and 840 min, with the evening peak shaded. From the top, the reboiler duty, the lean (solid) and rich (dashed) tank levels with the 5 % pump suction limit, and the capture rate with its four-hour rolling average (dotted) above the 88 % obligation. The plan strips deeper before the peak and stops at its 10 % lean margin, while the supervisor empties its lean tank further and its capture dips below 89 % after the peak.

above 88 %, the emission under 4.0 mg/Nm^3 and the rich tank drawn back down to 15 % or less. The shipped settings are basic loops with a capture gain of 0.25 and a 250 min reset, a reboiler loop of 3 MW/°C and 120 min, and a level loop of 0.2 and 1800 min. They fail two of the four parts, with a worst last-third deviation of 3.19 points and a rolling average that sinks to 86.96 %, leaving the plant 17.24 t short with a \$6,897 penalty, although the amine peak of 3.70 mg/Nm^3 and the final rich tank of 11.1 % pass. The suggested configuration puts the loop tunings of Table 7.3 under the economic supervisor, with the dear-power level at 50 %, $\alpha_d = 0.248$, $p_{hi} = 180$ \$/MWh and $b = 1.5$ points. It meets all four parts, holding the last third within 1.11 points, 0.89 inside the band, and never letting the rolling average below 88.93 %, with the emission peaking at 3.80 mg/Nm^3 and the rich tank finishing at 4.7 %. Integrated absolute error over the whole sixteen hours falls from 2414 to 677 % min, and the circulation travel rises from 7 to 130 m^3/min, which is the loop doing work rather than watching.

7.6.6 The Predictive Controller Against the Supervisor

Table 7.6 runs scenarios 9, 10, 11, 8 and 13 and the challenge for 960 min from each scenario's own settings, with the shipped seed, instruments and noise in every row. The plan keeps scenario 14's settings with no retuning, and on the challenge both structures carry the suggested tunings. Two supervisors exist only in the measurement script. The first is given the published prices and replaces the price index z of (7.17) by the share of the next eight published hours priced below the present hour, eight being the best of four windows tried over scenarios 9 to 11. The second, the third row of the comparison, is that ranked rule with one setting chosen in hindsight over the same three days from 280 candidates that keep the obligation and end at or under 15 % rich (a four-hour window, levels of 6 and 72 % and lean targets of 0.228 and 0.275). A copy used only by the script decides every 5 min.

Scenario 14 is the price day of scenario 10 under the predictive controller. The plan strips deeper than design before the evening, to a lean loading setpoint of 0.2229 and a duty of 426.2 MW at 420 min while the price is \$95. Through the \$175 hours it holds a mean duty of 268.6 MW against the supervisor's 290.0 and spends the lean tank to 9.89 % at 590 min, its own margin, regenerating the stored solvent overnight, as shown in Figure 7.5. Its capture never falls below 89.23 %, while the supervisor empties its lean tank to 6.21 % and its capture dips to 88.32 %. The plan makes \$523,372 against \$518,267 and the flat plant's \$513,204, with \$74,243 of steam against \$78,648 and \$81,088.

Table 7.6 The Predictive Controller Against the Supervisor Over 960 Minutes, Measured on the Engine

Case	Structure	Margin ($)	Steam ($)	Capture min (%)	Capture <85 (min)	Rolling min (%)	Lean min (%)	Rich max (%)	Rich end (%)	Plans
9	supervisor, shipped	664,951	94,456	79.71	21.5	88.80	3.3	70.5	4.4	–
	MPC	667,086	91,838	89.01	0	89.87	9.9	63.9	10.1	96
	supervisor given the published prices	666,899	93,453	84.61	6.5	88.96	3.5	70.3	4.5	–
	hindsight supervisor (third row)	666,813	92,934	78.77	9.5	89.13	3.4	70.4	4.6	–
	MPC, preview off	653,831	100,146	89.84	0	89.89	63.4	10.3	10.1	96
10	supervisor, shipped	518,267	78,648	88.32	0	89.50	6.2	67.5	4.4	–
	MPC (scenario 14)	523,372	74,243	89.23	0	89.91	9.9	63.9	11.0	96
	supervisor given the published prices	521,841	76,089	88.93	0	89.57	11.9	61.9	4.5	–
	hindsight supervisor (third row)	525,838	73,052	80.74	6.5	88.88	3.5	70.2	14.9	–
	MPC, preview off	513,364	80,456	89.85	0	89.92	63.4	10.4	10.1	96
	basic loops	513,204	81,088	89.98	0	89.99	63.6	10.1	10.0	–
	MPC, 4 h plan	523,327	74,030	89.19	0	89.90	9.9	63.9	11.2	96
11	supervisor, shipped	742,909	105,226	77.96	23.5	88.51	3.1	70.6	4.3	–
	MPC	743,462	104,104	89.00	0	89.88	9.9	63.9	10.1	96
	supervisor given the published prices	743,600	105,674	83.27	18.5	88.82	3.3	70.4	4.4	–
	hindsight supervisor (third row)	744,484	104,820	78.54	10.0	89.09	3.4	70.4	4.6	–
	MPC, preview off	733,460	110,689	89.78	0	89.90	63.5	10.2	10.1	96
Challenge	supervisor, suggested	600,292	90,961	85.46	0	88.93	10.2	63.5	4.7	–
	MPC, suggested loop tunings	604,273	86,942	80.80	11.5	89.17	9.9	63.8	12.6	96
	supervisor given the published prices	602,374	89,732	85.85	0	88.90	16.5	57.2	4.6	–
	hindsight supervisor (third row)	605,113	87,393	85.65	0	88.59	6.5	67.3	4.6	–
	MPC, preview off	594,064	93,297	80.86	12.0	89.15	57.3	16.5	10.1	96
	MPC, 4 h plan	604,012	86,944	80.82	7.0	89.27	9.9	63.8	16.2	96
8	basic loops, shipped	317,837	50,639	82.17	26.5	88.20	54.0	19.8	12.2	–
	MPC	317,330	50,151	64.96	24.5	88.29	52.9	20.9	10.1	96
	MPC deciding every 5 min	318,131	50,220	74.30	13.0	89.22	54.7	19.1	10.1	192
13	basic loops, shipped	381,202	63,529	82.66	5.5	88.92	44.7	29.0	4.7	–
	MPC	380,731	62,621	77.27	16.0	89.11	56.4	17.3	10.2	96
	MPC deciding every 5 min	381,391	62,691	77.40	9.0	89.57	55.8	18.0	10.1	192

Margin is the application's own figure, the revenue at each minute's price less the carbon price on escaped carbon dioxide, makeup and the obligation penalty, c_{obl} for each tonne short of holding every four-hour window at C_{obl}, charged once and paid by no row. Steam is the reboiler duty's cost in electricity at each minute's price. No row spends a minute under the rolling obligation, so the rolling minimum gives the depth. Capture <85 counts minutes below 85 %. The tanks start at 63.75 (lean) and 10 % (rich). No plan was rejected. On scenarios 8 and 13 the price is flat and the plan with its preview off equals the plan.

Most of that gain is knowledge of the prices, and the plan adds the rest. With its preview off the plan makes $513,364, within $160 of the flat plant, because a plan that sees one price has no reason to store. The supervisor given the same published schedule makes $521,841, which apportions the plan's $5,105 lead as $3,574 (70 %) of price knowledge and $1,531 (30 %) of optimization. The split is 91 to 9 on scenario 9 and 52 to 48 on the challenge, and on scenario 11 the ranked supervisor beats the plan by $138. The hindsight supervisor makes more than the plan on scenario 10 ($2,466 more), scenario 11 ($1,022) and the challenge ($840) and $273 less on scenario 9, by letting the capture fall to 80.74, 78.54 and 78.77 % on scenarios 10, 11 and 9, where the plan never lets it below 89.0 %. The capture weight and the lean margin are a decision not to starve the absorber, and the hindsight row prices it. Nor is the horizon the source, since a four-hour plan is within $261 of the 24-hour plan on all four priced cases. With scenario 10's schedule published and other prices realized (the peak an hour late or early, at $130 or $230, or absent), the plan still beats the supervisor by $2,073 to $5,998, and the ranked supervisor on the same wrong schedule ends between $1,057 ahead of the plan and $2,767 behind it.

The plan loses both flat-price days to basic loops, and the cause is its decision period. On scenario 8 it makes $317,330 against $317,837. The decision at 480 min reads the analyzer before the returning gas reaches the absorber and holds the circulation for ten minutes, so the capture falls to 64.96 % against 82.17 % under a capture loop that acts every 0.5 min. On scenario 13 it makes $380,731 against $381,202, with the capture down to 77.27 % at the warming against 82.66. It recovers faster on both, with capture integrated absolute errors of 465 against 711 and 380 against 521 % min, and deciding every 5 min it makes $318,131 and $381,391 and beats basic loops on both days. Scenario

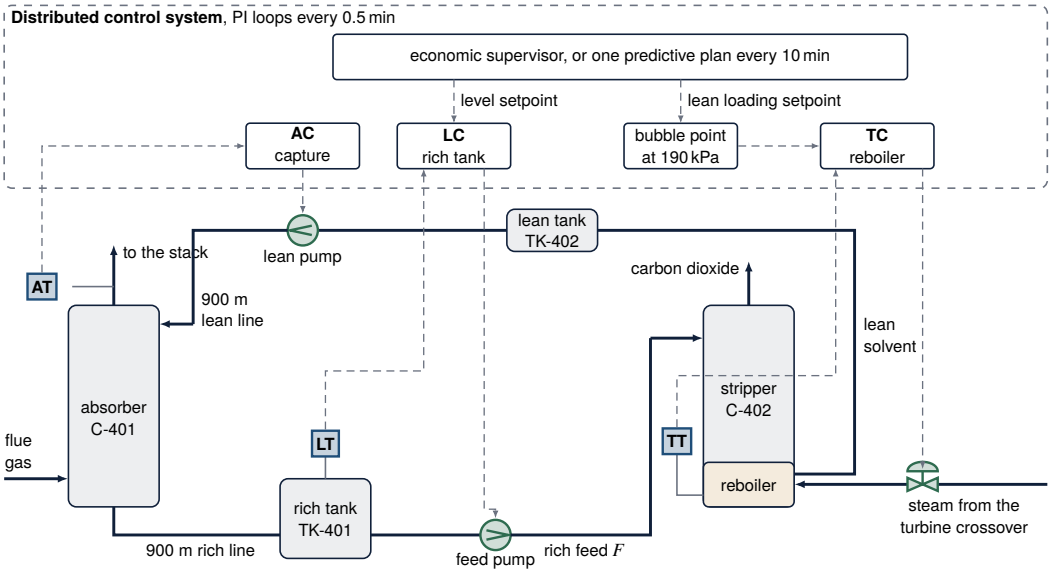

Figure 7.6 The three loops of capture plant U-400 with their hardware. Lettered boxes are sensing devices: the stack analyzer (AT-401) on the treated gas, the thermocouple (TT-402) in the reboiler and the level transmitter (LT-401) on the rich tank. The pump symbols and the steam valve are final elements, and dashed arrows are signals. Notice that the capture controller moves the lean pump at one end of the circuit while the lean loading it depends on is set by steam at the other end, and that the reboiler loop receives its setpoint as a temperature computed at an assumed stripper pressure. The cross exchanger and the makeup are left out. Not to scale, and not a piping diagram.

10's two fairness tests, which the supervisor passes, both fail narrowly for the plan. It captures 0.042 points less carbon dioxide than the flat plant (2.77 t, already charged in the margin) and ends 1.03 points fuller in the rich tank (83 m^3, \$17 to \$41 of steam). On the challenge with the suggested loop tunings the plan meets all four parts of the target, with a worst last-third deviation of 1.48 points, a rolling minimum of 89.17 %, an amine peak of 3.90 mg/Nm^3 and the rich tank at 12.6 %, while the four-hour plan ends at 16.2 % and misses the last part. On the inherited tunings the plan fails, with a worst deviation of 6.73 points and 205 min under the obligation (54.87 t short, a \$21,949 penalty), because the reboiler loop at 3 MW/°C and a 120 min reset takes 312.5 min to reach 63 % of a lean setpoint step where the model assumes 9.4 min.

7.7 Technology in Practice

What Measures It

Figure 7.6 draws the three loops of Section 7.5 with their devices. The capture loop reads the carbon dioxide left in the treated gas, and at a commercial capture unit that reading comes from certified analyzers on the absorber stack, which also measure the stack flow [18]. The reboiler loop reads a thermocouple and the level loop a level transmitter. Lean loading is usually found by titration, an analysis that can take up to two hours while the open sample loses carbon dioxide and the amine degrades [24], which is why the chapter's plant infers it from the reboiler temperature.

TECHNOLOGY Sensor: A Non-Dispersive Infrared Analyzer

Carbon dioxide in a combustion exhaust is the standard example of the non-dispersive infrared analyzer, developed in 1937 and 1938. Sample gas flows through a windowed cell, and a heat-sensitive detector behind it measures how much infrared light the gas did not absorb. Water vapor, carbon monoxide and sulfur dioxide absorb too, so a selective analyzer uses a detector filled with the gas of interest. A sealed reference cell cancels the zero shift of an aging lamp but not its span shift, which calls for recalibration. The sample must arrive clean and dry, since liquid or dust in the analyzer causes anything from delay to gross error [25]. On a trend a span drift is a slow walk of the reading that no move of the plant explains. The model's AT-401 is a 2.0 min lag with optional noise.

What Moves It

The capture loop moves the lean circulation through pumps, the level loop moves the stripper feed, and the reboiler loop moves low-pressure steam. On a coal unit that steam is extracted from the crossover pipe between the intermediate- and low-pressure sections of the steam turbine, and in a reference design its flow is proportional to the rich amine sent to the stripper and depends also on the stripper top temperature [26]. That extraction is the term $k_s Q$ of Equation (7.13). The model gives the steam valve a lag of 3.0 min and bounds the pumps by their stops and by the tank ends of Equation (7.16).

TECHNOLOGY Final Element: A Solvent Pump on a Variable-Speed Drive

Setting a pumped flow by the speed of the pump uses less energy than running it at full speed and throttling its discharge or recycling to its suction, and it leaves the control valve out of the line. A variable-speed machine often cannot change flow as fast as a valve and has limited fail-safe options, and its drive limits speed, current and acceleration and shuts down on overcurrent [25]. A pump held at such a limit cannot follow its command, so on a trend the commanded and the delivered flow part company. The model's circulation is bounded at 55 and 170 m^3/min and falls once the lean tank drops below its 5 % pump suction limit, and every run exports the commanded flow beside the delivered one.

Where the Controller Runs

At the commercial capture unit that began operating in 2016, the capture process and its balance of plant were run from a distributed control system in their own control room, and the carbon dioxide compressor had a dedicated programmable logic controller. The operator changed the plant's load by changing a carbon dioxide production setpoint, and automatic control held it anywhere from 50 to 100 percent load, although the plant ran best at full load [18]. A distributed control system runs nested and cascaded loops under a supervisory level, with cycle times from milliseconds to minutes [27]. The app executes its three PI loops every 0.5 min and its predictive controller every 10 min, inside that range. The three loops belong to the regulatory layer, and the economic supervisor and the predictive controller to the layer above it.

> **IN PRACTICE** **Keeping a Stack Analyzer Certified**
>
> A United States coal unit under the federal acid rain rules measures its carbon dioxide emissions with a certified analyzer and a flow monitor whose data are recorded automatically [28]. On a boiler the analyzer is spanned between 14 and 20 percent carbon dioxide. Every operating day it is checked with zero and upscale calibration gases, and an error above 1.0 percent carbon dioxide puts its data out of control. A linearity check follows each quarter, the relative accuracy against a reference method must stay within 10 percent, or within 1.0 percent carbon dioxide, and the cycle time within 15 min [29, 30]. The commercial capture unit of 2016 carried certified analyzers for flow and carbon dioxide on its absorber stack [18].

What Is Modernizing It

Three developments bear on this plant. Near-infrared spectrometry calibrated on a gas plant's solvent reads amine strength, carbon dioxide loading and water content, and has been proposed to replace titration [24]. An online loading measurement of that kind would expose the pressure-drift failure of Section 7.6. Scheduling amine regeneration for times of surplus power, so that the station's output can rise when it is needed, is proposed as a service capture plants can offer the grid [17]. And economic predictive control that joins scheduling and dynamic control, together with coordinated control of the power station and the capture plant, is named as the next step for the field [7].

7.8 Limitations

No number printed here will size a piece of equipment, because the model was never built as a design tool. The equilibrium of (7.1) is a three-factor correlation anchored at one point rather than an electrolyte activity model, the stripper is one lumped reboiler with a calibrated φ, and three absorber stages cannot place a temperature bulge in height. Specific duty, solvent inventory and emission performance are therefore internally consistent and not transferable to a real plant [7]. The water balance is open, so the excursions that occupy real operators are absent. The host is a fixed gross output and a fixed steam-to-electricity factor, which flatters very deep storage and very high capture rates [15]. There is no feedforward from the host load signal and no analyzer on the lean loading, both withheld deliberately, so scenario 8 takes longer than a real plant would and scenario 6 has something to break. A real plant drifts, while every disturbance here arrives as a commanded step or a timed event, and freezing the random seed alongside the integration step and the controller period returns a trajectory no distributed control system could repeat. Three elements are stylized on purpose. A stripper pressure that rises 25 kPa at an instant is not how a condenser fouls, but it isolates the inferential error from everything a season of fouling would change alongside it. Degradation at 6 % per day compresses months into a shift so the actuator drift reads off one trend. The price arrives as a staircase because a staircase is the disturbance whose response can be timed, while a real pool price moves every five minutes. Five features survived the trimming, the volume-keyed transport of (7.6), the loop-closing scalar of Section 7.4, the convex reboiler bill of (7.10), the inference of (7.11) and the tank ends of (7.16), each because it changes the controller a reader would have to build. Section 7.3.8 says where the remaining assumptions stop holding.

The predictive controller's model omits the host load, the cooling water, the stripper pressure, the solvent strength and the packing, which reach it only through the analyzer bias, and it has no rich-leg recycle state. Its preview is an exact copy of the schedule, and it knows nothing of a dispatch. Its 10 min decision loses scenarios 8 and 13 to basic loops by \$507 and \$471, the hindsight supervisor out-earns it on three of four priced cases by letting capture fall to 78.54 %, and it fails the challenge on the inherited tuning of the temperature loop it relies on.

7.9 Discussion

The plant makes an unusual demand on the control engineer, which is that the manipulated variable of the most important loop is bought out of the product the owner is paid for. Every megawatt of reboiler duty is 0.175 MW the station does not sell, so the design duty of 373 MW is 65.3 MW of lost

revenue before the compressors are counted [10]. Neither end of that trade is negotiable, because the capture rate is written on a permit and the energy per tonne is set by the heat of absorption and the water boiled with it, so the only decision left is when to spend it and the tanks are the only instrument for making it. A level setpoint on the rich tank is therefore an energy schedule worth 38 MW while the tank is filling, bounded not by the tank being filled but by the one being emptied.

The same reasoning explains why the supervisor multiplies rather than adds. An economic layer that traded permit compliance against price at any exchange rate would eventually find a price at which selling the permit is the optimal action, and (7.17) is written so that no such price exists. That is a control structure decision rather than an optimization one, and it is the same distinction that separates a self-optimizing choice of controlled variables from a numerically optimal one [6].

The model supports methods beyond the three implemented loops. A sixteen state balance with two hard tank limits and a convex reboiler bill is what model predictive control is for, and the option of Section 7.5.5 keeps both tank ends inside their alarm levels explicitly and coordinates circulation, lean loading and feed as one three-input problem with the price in the objective rather than as three loops with an hour of interaction between them [5, 16]. The same equations under moving horizon estimation would recover solvent strength and packing condition out of the circulation and duty records, turning the invisible fault of scenario 12 into a soft sensor. The pressure-compensated inference scenario 6 asks for is one line of arithmetic once (7.11) is available, and it is the cheapest improvement anywhere in this plant.

> ⚠ **OPEN CHALLENGE** **Running the Power Station and the Capture Plant as One**
>
> A power station that follows the electricity market changes its flue gas flow and the steam drawn from its turbine often and over a wide range, while a capture plant can take more than two hours to settle after a change in reboiler duty. Most control studies treat the capture plant on its own, although the two are strongly coupled through the flue gas and the extracted steam, and a steady-state optimization cannot see the transients. Loading, flows and concentrations are hard to measure accurately in real time, and the plant itself changes as the solvent degrades and the equipment corrodes [7]. Progress would be coordinated control of station and capture plant, with the schedule and the dynamic control solved together [7].

7.10 Conclusion

This chapter has presented a post-combustion amine capture plant as an interactive laboratory for what a large material recycle does to the gain of a control loop. The simulator integrates sixteen states, keys its two 1170 m^3 transport lines on cumulative volume so that the 9.00 min delay stretches to 7.55 min at 155 m^3/min, and solves rather than assumes a design line-up of 90.00 % capture, 371.3 t/h of carbon dioxide, loadings of 0.2375 and 0.4472 mol/mol, a reboiler at 115.87 °C, 3.619 MJ/kg of specific duty and 463.7 MW left of the host’s 580 across a 9100 m^3 circuit whose lap is 70.0 min. With the reboiler duty frozen a 20 m^3/min circulation step is worth 0.177 points of capture per m^3/min at ten minutes and 0.040 at eight hours, while the same step with the lean loading loop in service is worth 0.177 at both times, paid for by the reboiler moving from 372 to 401 MW. Weakening that inner loop to 3 MW/°C leaves the capture controller on its 170 m^3/min stop almost a point short of setpoint with the lean loading at 0.2715, and raising the surge tank gain from 1.5 to 8 shortens the time to setpoint from 110 to 44 min while widening the reboiler duty swing to 427 MW, which is the opposite of what the tank was bought for. Time-shifting regeneration across a sixteen-hour price day sells 0.80 % more electricity and earns a 0.99 % larger margin at the same captured tonnage, and seven hours of expensive power finish 0.00 t short of the obligation because the price signal is multiplied by the remaining credit rather than added to it. On that day the predictive controller earns $5,105 more than the supervisor, and 70 % of it is recovered by a supervisor given the same prices. The next additions to the model would be feedforward from the host load, a pressure-compensated loading inference, a throttled low-pressure turbine and an axial absorber profile. Why an inner loop

earns its cost when the outer loop is wrapped around a recycle is the lesson the capture plant was built to carry, and a browser is all it asks for.

KEY POINTS

- A capture plant's solvent comes back to the absorber about seventy minutes after it leaves, so every move of the circulation acts through a recycle.
- With the reboiler duty frozen the capture gain falls from 0.177 at ten minutes to 0.040 at eight hours. A lean loading loop underneath holds it at 0.177, so the inner loop makes the outer loop's gain a constant.
- An inference built on the reboiler temperature at an assumed pressure holds its measurement perfectly while the true lean loading walks away.
- A surge tank level loop must stay slack, and a rich tank setpoint is an energy schedule bounded by the tank being emptied.

Where else this applies. Volume 6, Chapter 7 measures the same recycle amplification on a reactor and column, and Volume 5, Chapter 10 infers a product quality from a temperature that a second variable misleads.

References

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

[2] A. L. Kohl and R. B. Nielsen, *Gas Purification*, 5th ed. Houston, TX: Gulf Publishing, 1997.

[3] H. M. Kvamsdal, J. P. Jakobsen, and K. A. Hoff, "Dynamic modeling and simulation of a CO_2 absorber column for post-combustion CO_2 capture," *Chemical Engineering and Processing: Process Intensification*, vol. 48, no. 1, pp. 135–144, 2009.

[4] S. Ziaii, G. T. Rochelle, and T. F. Edgar, "Dynamic modeling to minimize energy use for CO_2 capture in power plants by aqueous monoethanolamine," *Industrial & Engineering Chemistry Research*, vol. 48, no. 13, pp. 6105–6111, 2009.

[5] Y.-J. Lin, T.-H. Pan, D. S.-H. Wong, S.-S. Jang, Y.-W. Chi, and C.-H. Yeh, "Plantwide control of CO_2 capture by absorption and stripping using monoethanolamine solution," *Industrial & Engineering Chemistry Research*, vol. 50, no. 3, pp. 1338–1345, 2011.

[6] M. Panahi and S. Skogestad, "Economically efficient operation of CO_2 capturing process part I: Self-optimizing procedure for selecting the best controlled variables," *Chemical Engineering and Processing: Process Intensification*, vol. 50, no. 3, pp. 247–253, 2011.

[7] X. Wu, M. Wang, P. Liao, J. Shen, and Y. Li, "Solvent-based post-combustion CO_2 capture for power plants: A critical review and perspective on dynamic modelling, system identification, process control and flexible operation," *Applied Energy*, vol. 257, p. 113941, 2020.

[8] T. Nittaya, P. L. Douglas, E. Croiset, and L. A. Ricardez-Sandoval, "Dynamic modelling and control of MEA absorption processes for CO_2 capture from power plants," *Fuel*, vol. 116, pp. 672–691, 2014.

[9] G. T. Rochelle, "Amine scrubbing for CO_2 capture," *Science*, vol. 325, no. 5948, pp. 1652–1654, 2009.

[10] E. S. Rubin, J. E. Davison, and H. J. Herzog, "The cost of CO_2 capture and storage," *International Journal of Greenhouse Gas Control*, vol. 40, pp. 378–400, 2015.

[11] R. R. Bottoms, "Process for separating acidic gases," U.S. Patent 1 783 901, Dec. 2, 1930, filed October 7, 1930.

[12] F.-Y. Jou, A. E. Mather, and F. D. Otto, "The solubility of CO_2 in a 30 mass percent monoethanolamine solution," *The Canadian Journal of Chemical Engineering*, vol. 73, no. 1, pp. 140–147, 1995.

[13] R. Weber and C. Brosilow, "The use of secondary measurements to improve control," *AIChE Journal*, vol. 18, no. 3, pp. 614–623, 1972.

[14] J. Davis and G. Rochelle, "Thermal degradation of monoethanolamine at stripper conditions," *Energy Procedia*, vol. 1, no. 1, pp. 327–333, 2009.

[15] H. Chalmers and J. Gibbins, "Initial evaluation of the impact of post-combustion capture of carbon dioxide on supercritical pulverised coal power plant part load performance," *Fuel*, vol. 86, no. 14, pp. 2109–2123, 2007.

[16] S. M. Cohen, G. T. Rochelle, and M. E. Webber, "Optimizing post-combustion CO_2 capture in response to volatile electricity prices," *International Journal of Greenhouse Gas Control*, vol. 8, pp. 180–195, 2012.

[17] M. Bui, C. S. Adjiman, A. Bardow, E. J. Anthony, A. Boston, S. Brown, P. S. Fennell, S. Fuss, A. Galindo, L. A. Hackett, J. P. Hallett, H. J. Herzog, G. Jackson, J. Kemper, S. Krevor, G. C. Maitland, M. Matuszewski, I. S. Metcalfe, C. Petit, G. Puxty, J. Reimer, D. M. Reiner, E. S. Rubin, S. A. Scott, N. Shah, B. Smit, J. P. M. Trusler, P. Webley, J. Wilcox, and N. Mac Dowell, "Carbon capture and storage (CCS): the way forward," *Energy & Environmental Science*, vol. 11, no. 5, pp. 1062–1176, 2018.

[18] Petra Nova Power Holdings LLC and G. Kennedy, "W.A. Parish post-combustion CO_2 capture and sequestration demonstration project (final technical report)," U.S. Department of Energy, Office of Scientific and Technical Information, Tech. Rep. DOE-PNPH-03311, Mar. 2020.

[19] 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.

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

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

[22] 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] Z. He, M. H. Sahraei, and L. A. Ricardez-Sandoval, "Flexible operation and simultaneous scheduling and control of a CO_2 capture plant using model predictive control," *International Journal of Greenhouse Gas Control*, vol. 48, pp. 300–311, 2016.

[24] M. Yusop Nurida, D. Norfadilah, M. R. Siti Aishah, C. Zhe Phak, and S. M. Saleh, "Monitoring of CO_2 absorption solvent in natural gas process using Fourier transform near-infrared spectrometry," *International Journal of Analytical Chemistry*, vol. 2020, p. 9830685, 2020.

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

[26] T. Schmitt, S. Leptinsky, M. Turner, A. Zoelle, C. White, S. Hughes, S. Homsy, M. Woods, H. Hoffman, T. Shultz, and R. James III, "Cost and performance baseline for fossil energy plants, volume 1: Bituminous coal and natural gas to electricity," National Energy Technology Laboratory, Pittsburgh, PA, Tech. Rep., Oct. 2022.

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

[28] U.S. Environmental Protection Agency, "40 CFR 75.10, general operating requirements, continuous emission monitoring," Code of Federal Regulations, Title 40, Part 75, edition of July 1, 2024, govinfo.gov, 2024.

[29] ——, "40 CFR part 75, appendix A, specifications and test procedures," Code of Federal Regulations, Title 40, Part 75, edition of July 1, 2024, govinfo.gov, 2024.

[30] ——, "40 CFR part 75, appendix B, quality assurance and quality control procedures," Code of Federal Regulations, Title 40, Part 75, edition of July 1, 2024, govinfo.gov, 2024.

7.11 Exercises

7.1 ⌕ Concept check · 2 min

With the reboiler duty frozen at 373.3 MW, a circulation step from 130 to 150 m^3/min raises the capture rate by 0.177 points per m^3/min at ten minutes and by only 0.040 after eight hours. Why does the gain decay?

(a) The absorber packing wets less evenly at the higher flow, so its transfer coefficient falls

(b) The pipe delay grows with the flow, so the extra solvent reaches the absorber later and later

(c) The same duty must regenerate more solvent, so the solvent that returns is less lean

(d) The stack analyzer's lag hides the early response, so the first reading overstates the gain

7.2 ⌕ Concept check · 2 min

The economic supervisor of Equation (7.17) multiplies the price index z by $1 - u$, where u is the obligation urgency. Why a product and not a sum?

(a) A product makes the supervisor answer a price change faster than a sum would
(b) A product gives the price a fixed exchange rate against the permit, which a sum lacks
(c) A product keeps the rich tank below its full limit when the price stays high
(d) A product strips the price of its authority as the credit runs out, so no price can sell the permit

7.3 Concept check · 2 min

In scenario 6 the stripper pressure rises from 190 to 215 kPa at 90 min. The reboiler temperature loop has its temperature back within 0.1 °C of setpoint 19 min later and nothing alarms. Which statement is true once the plant has settled?

(a) The true lean loading is near 0.274 while the inferred one still reads 0.238, and the only sign is extra circulation
(b) The inferred lean loading reads 0.274, so the operator sees the richer solvent at once
(c) The capture rate settles about a point below its setpoint because the solvent is richer
(d) The temperature loop is mistuned, and a higher gain would bring the lean loading back

7.4 Theory and calculation · 6 min

The circuit holds two lines of 1170 m^3 each, a lean tank of 5100 m^3, a rich tank of 800 m^3, a reboiler of 500 m^3 and an absorber of 360 m^3.

(a) Compute the residence time of each holdup at the design circulation of 130 m^3/min and the lap. *Checkpoint:* lap 70.0 min

(b) Scenario 2 steps the circulation to 155 m^3/min. Compute the transit time of one line and the lap at that flow. The trend shows the solvent leaving the absorber moving after 0.5 min and the solvent reaching the top of the absorber after 21.5 min. Why is the second time longer than the absorber's half minute plus two line transits? *Checkpoint:* 7.55 min and 58.7 min

(c) At part load the capture loop cuts the circulation to about 70 m^3/min. Compute the line transit and the lap, and say what this does to a capture loop tuned at design. *Checkpoint:* 16.7 min and 130 min

7.5 Theory and calculation · 7 min

Use the recycle relation of Equation (1.5), settled gain $g_F/(1-g_R)$, with the gains of Figure 7.2: 0.177 points per m^3/min at ten minutes, before any changed solvent has come back, and 0.040 settled with the duty frozen.

(a) Compute the return gain g_R of the solvent loop with the duty frozen. *Checkpoint:* $g_R = -3.4$

(b) Scenario 3 steps the capture setpoint from 90 to 94 %. Estimate the settled circulation with the lean loading loop in service and with the duty frozen, starting from 130 m^3/min. The pump's stop is 170 m^3/min. *Checkpoint:* about 153 and 230 m^3/min

(c) The capture loop has $K_c = 3.5\,m^3/\text{min}$ per point. Compute the loop gain K_c times the process gain for both cases and their ratio, and say which case the shipped tuning was chosen for. *Checkpoint:* 0.62 and 0.14, ratio 4.4

7.6 Theory and calculation · 6 min

At the design line-up the reboiler takes 373.3 MW, the plant captures 371.3 t/h of carbon dioxide, and the host makes 580 MW gross. Use Equation (7.13) with $k_s = 0.175$, $w_c = 0.378\,\text{MJ/kg}$, $W_p = 7.5\,\text{MW}$ and $W_b = 4.5\,\text{MW}$, and take the stripper outlet equal to the capture.

(a) Compute the steam and compression terms, the total the capture plant takes and the net output. *Checkpoint:* 116.3 MW and 463.7 MW

(b) Express the total and the steam alone in MWh of lost output per tonne captured. Give three decimals. *Checkpoint:* 0.313 and 0.176

(c) Scenario 4 measures 13.2 MW for the step from 90 to 95 % and 7.8 MW returned by the step from 90 to 85 %. Compute the cost per point on each side and their ratio. *Checkpoint:* 2.64 and 1.56 MW per point, ratio 1.7

7.7 App challenge · 8 min

Open scenario 1, *The absorber by hand (MANUAL)*, default seed, as loaded.

(a) Run at 1 min/s and pause when the clock reads between 3 and 30 min. Note the time t_s, set Solvent circulation and Stripper feed rate both to 150 m^3/min, and run at 60 min/s to at least $t_s + 480$ min. From the exported CSV (time column `Time (h)`), read `capture_pct` at $t_s + 10$ and $t_s + 480$ min and `a_lean_reb_m` at $t_s + 480$ min, and compute the apparent gain at both times against the design 90.00 %. Give the capture rates to ± 0.02 points and the loading to ± 0.001. *Checkpoint:* 93.55 and 90.81 %

(b) Press the reload button, run a few minutes as in (a), and set Reboiler steam duty to 420 MW. Read `capture_pct` 20 and 60 min after the step (± 0.02 points) and `W_net_MW` just before the step and 60 min after it (± 0.2 MW). *Checkpoint:* 90.25 and 91.73 %

7.8 App challenge · 8 min

Open scenario 3, *Capture on circulation, and the loop underneath that makes it work*, default seed. The capture setpoint steps from 90 to 94 % at 60 min and returns at 420 min.

(a) Run as loaded (Basic loops) at 60 min/s past 420 min. From the exported CSV row at `Time (h)` = 7.000, read `capture_pct`, `L_cmd_m3min`, `a_lean_reb_m` and `W_net_MW`. *Checkpoint:* 94.01, 153.9, 0.2379, 453.6

(b) Set Reboiler temperature loop gain to 3 MW per °C and Reboiler temperature integral time to 120 min, press Reset, and repeat the run. Read the same four values, and the first row with `L_saturated` equal to 1. Give the capture to ± 0.02 points, the loading to ± 0.0005, the host output to ± 0.2 MW and the time to ± 1 min. *Checkpoint:* 93.03, 170.0, 0.2715, 456.7

(c) The capture tuning is the same in both runs. Explain why no change to the capture loop repairs the second run.

7.9 App challenge · 9 min

Open the challenge, *A day of prices, a load ramp and a permit*, default seed. Run every attempt to 960 min at 120 min/s without stopping, and read the result line of the challenge in the progress panel and the CO_2 short and Margin tiles.

(a) Run it as shipped. *Checkpoint:* worst 3.19 points, rolling minimum 86.96 %

(b) Set the six loop tunings to those of Table 7.3 (capture gain 3.5 with 20 min, reboiler 30 with 12 min, level 1.5 with 240 min), stay on Basic loops, press Reset and run again. Give the four values of the result line. *Checkpoint:* worst 0.58 points, rolling minimum 89.10 %

(c) Keep those tunings and set Control to Economic supervisor. In the Economic supervisor group of the advanced controls, set Rich level setpoint when power is dear to 50 %, Lean loading target when power is dear to 0.248, Price the supervisor calls dear to \$180 and Obligation credit the supervisor may spend to 1.5 points. Press Reset and run again. Compare the margins of (b) and (c) over the day. *Checkpoint:* worst 1.11 points, rich tank 4.7 %

Give the result-line values as printed and the margins to $\pm$\$10 per hour.

7.10 Application area · 4 min

In scenario 9 the price goes to \$210 for five hours and the supervisor raises the rich tank setpoint to 72 %. Half an hour into the peak the stripper feed is about 64 m^3/min against a circulation of about

130 m^3/min, and the host makes 508.8 MW against 451.6 MW before the peak. The lean tank starts at 63.75 % of 8000 m^3, and the pumps lose suction at 5 %.

(a) Compute the lean tank's usable swing and how long it would last at the imbalance above. *Checkpoint:* 4700 m^3, about 71 min

(b) What is the extra output worth per hour at \$210, and why is the binding tank the lean one rather than the rich one an operator watches?

7.11 Application area · 4 min

Two faults in the chapter hold every setpoint. Use Section 7.6 and Equation (7.15).

(a) Degradation doubles for every 8 K above 120 °C. By what factor is the rate at 128 °C larger than at the design 115.87 °C? What does the makeup of 90 kg/min that holds the strength cost per hour at \$2.40 a kilogram? *Checkpoint:* 2.9 times, \$12,960 per hour

(b) On a 49 °C afternoon the capture loop holds 90 %. Which limit binds, and what would an operator have to give up to stay inside it?

7.12 Review the agent · 5 min

An AI assistant reads Table 7.6 and writes: "On the price day of scenario 10 the predictive controller earns \$523,372 against the supervisor's \$518,267. Its optimizer is worth \$5,105 a day, so the predictive controller should replace the supervisor on every day the plant runs." Find the flaw, and say how you would check it in the app.

CHAPTER 8

A Dividing-Wall Column Whose Spare Handle Moves Only the Energy

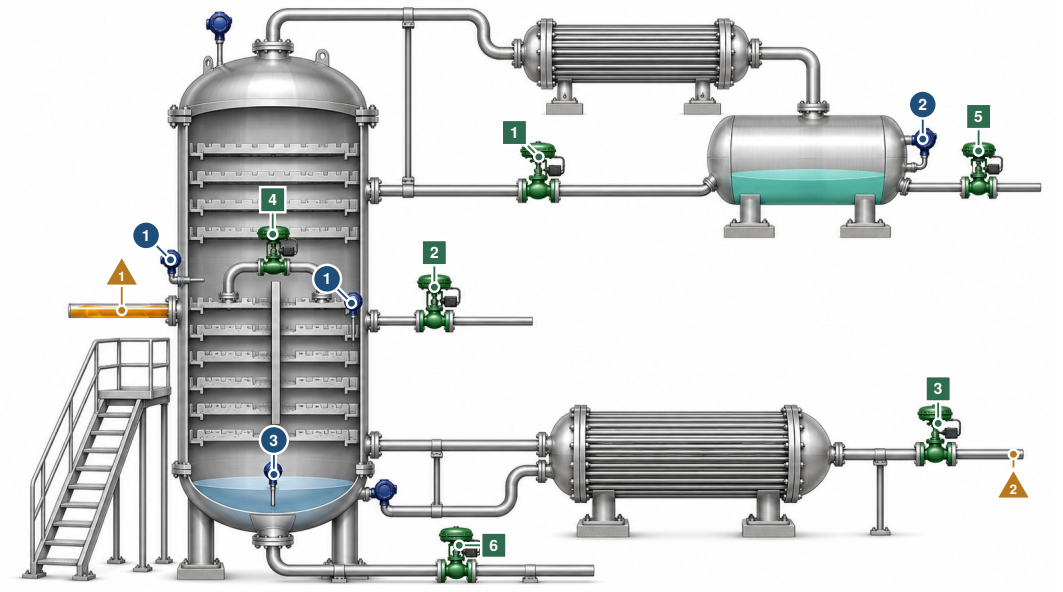

The controller moves 1 reflux valve 2 side draw valve 3 reboiler steam valve, which sets the boilup 4 liquid splitter above the wall 5 distillate valve, on the drum level 6 bottoms valve, on the base level

The controller reads 1 tray temperatures beside the wall 2 reflux drum level 3 base level

What disturbs it 1 a heavier or larger feed 2 a limit on the steam supply

Notice the wall up the middle of the tower and the splitter at its top: with the three products held pure, that one setting decides how much steam the reboiler at the lower right uses. Illustration, not to scale.

IN PLAIN TERMS One Tower Doing the Work of Two

The main idea. One tall tower with a wall up its middle splits a liquid mixture into three products. Three controllers keep the products pure, and a fourth setting decides how much of the liquid falling from the top goes to each side of the wall.

The challenge. With the products held pure, that setting changes only the steam used to boil the liquid at the bottom. The best value moves when the feed changes, and too low a value means no amount of steam makes the middle product.

Why it matters. A column with every product on target can still burn far more steam than it needs.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page; the history boxes of Section 8.2; the fact box of Section 8.4; the try-it and fact boxes of Section 8.6; the key points
Study	45 min	the whole chapter with its pause points and the exercises
Practice	15 min	the first page; Section 8.6 with its result chart; Section 8.7; Section 8.8 and its open challenge

The control question Three purity loops use three of a dividing-wall column's four handles, and the fourth, the liquid split above the wall, moves only the reboiler duty along a valley whose minimum shifts with the feed. How is the column kept in that valley when the feed changes, and what does each control structure cost on the way?

Learning objectives

- Count the degrees of freedom of a dividing-wall column and explain why the spare handle moves only the energy.
- Compute the Underwood minimum vapor of each split, the thermally coupled saving, and the bound a fixed vapor split sets.
- Explain the steep side of the energy valley from the liquid-to-vapor ceilings on the two sides of the wall.
- Compare a fixed split, a split loop on the prefractionator, a feed-rate feedforward and an MPC through a feed change and a rate change.
- Judge when an inferred purity or a tray temperature can stand for the quantity a loop must hold.

Prerequisites Sections 1.4.5 (degrees of freedom and Underwood's method), 1.4.6 (inferred compositions and the bias update), 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/dividing-wall-column/

See also Volume 6, Chapter 2 for pairing by the relative gain array and Volume 6, Chapter 7 for choosing a control structure. Volume 9, Chapter 8 for an optimum that moves with a disturbance and Volume 9, Chapter 4 for offset-free predictive control with a disturbance estimator. Volume 5, Chapter 3 for composition control through a slow analyzer.

8.1 Introduction

Three products from one feed normally need two distillation columns, two reboilers and two condensers. A dividing-wall column does the two splits in one shell. In the generic column simulated here, 300 kmol/h of equimolar benzene, toluene and o-xylene enters the feed side of a vertical wall, benzene leaves overhead at 99.2 mol %, toluene leaves by a side draw at 98.2 mol % and o-xylene leaves the base at 99.2 mol %, with one reboiler boiling 649 kmol/h (5.83 MW). That duty is a minimum, and it holds only while 31.4 % of the liquid leaving the section above the wall is sent down the feed side. With 40 % sent there instead, the same three purities need 17.7 % more steam, and with 28 % they need 21.4 % more. Below 28 % no boilup under the flooding rate makes the side product at all.

The simulator is built around one misconception, that holding every product on specification leaves nothing more for the operator to decide. Here three purities take three handles (reflux, side draw and boilup), and the fourth, the liquid split above the wall, is left free. Once the purities are held it cannot change any product flow or composition, so it moves only the internal flows and with them the energy. The same structure, more handles than specifications and a free one that sets the cost, appears in parallel trains that share a load, in heat-integrated units and wherever an optimum lies inside the constraints, and choosing what to hold constant there is the problem of self-optimizing control [1]. Chapter 6 couples two columns the other way, through one condenser-reboiler, so that one duty serves both. A control course meets the relative gain array [2] on a two-product column and rarely on a plant where a spare handle has a minimum to find and lose. Two more lessons come with the dividing wall. Its three purity loops interact strongly, and the column is ill-conditioned in the way high-purity distillation is [3]. Its minimum also moves with the feed, so a split fixed at commissioning wastes energy after the next feed change, which is the problem a measurement on the prefractionator was proposed to solve [4].

However, a student who reads the dividing-wall literature sees the energy valley as a figure from a steady-state solver and the control structures as trends from a commercial simulator [4, 5]. This chapter works with that column in a browser, with the liquid split as a handle and the valley computed

from the same equations the plant integrates. It measures what a fixed split, a split loop on the prefractionator, a feed-rate feedforward and a model predictive controller (MPC) do with one feed change and one rate change. The stage model, the shortcut methods of Fenske, Underwood and Gilliland [6–8], the four-point structure of Ling and Luyben [4] and the proportional-integral (PI) law are all established methods. The results are measured on one generic design. The valley computed by the app agrees with an independent steady-state solver to 0.01 kmol/h at twenty splits, its steep side is a ceiling on the prefractionator's liquid-to-vapor ratio, and an MPC on a fixed step-response model reaches the valley but lets the benzene fall below its specification on a feed change. The model has 138 states, is built on a distillation module that other apps share, closes its balances to 8.7×10^{-10} of the feed, and is given here with every equation, constant and default needed to build it again.

Section 8.2 traces the domain, from two columns in sequence to the wall in one shell. Section 8.3 sets out the shell as modeled: its stages, the wall, the drum and sump, the instruments and the parameters. Section 8.4 counts the degrees of freedom, computes the energy valley, checks it against an independent solver and the shortcut bounds, and explains why one side is steep. Section 8.5 describes the loops, the inferential measurements, the split loop, the feedforward and the MPC. Section 8.6 measures the nine cases of the app, from walking the valley by hand to the graded challenge, Section 8.7 names the real devices behind the loops, and Section 8.8 states the limitations. The exercises at the end of the chapter run several of the cases.

8.2 Two Columns in One Shell

Separating a ternary mixture by distillation takes two splits, and the classic answer is a sequence of two columns. In the direct sequence the first column takes the light component overhead and the second splits the other two, so the middle component is boiled in both. The thermally coupled alternative sends the feed to a prefractionator that splits it roughly, light and middle up, middle and heavy down, and returns both ends to a main column that draws the middle component from its side. That arrangement is the Petlyuk column, whose origin and first operating studies are told in the history boxes of this section. For infinite stages, constant relative volatility and constant molar flows, its minimum vapor rate equals that of the more difficult of the two binary splits done in a single column [9], the analytical form of Underwood's minimum reflux [7].

> ⧖ HISTORY **A Wall Through the Middle Trays, 1937 to 1965**
>
> In 1949 Wright proposed an alternative to columns in series for a three-component feed: an ordinary column shell in which a vertical wall through a set of trays divides the feed from the side draw, saving investment, with one shell and two heat exchangers, as well as operating cost. Brugma had proposed a similar design in 1937, with its own condenser and reboiler for the prefractionator. The arrangement carries the name of Petlyuk and coworkers, who studied it theoretically in 1965 to avoid the losses of mixing streams of different composition at the feed tray. Many authors predicted large savings, yet by 1995 few such columns had been built, and the only industrial one reported stood in Germany. A probable reason was that the column has many more degrees of freedom than an ordinary one, which complicates both its design and its control [10].

Putting both columns of the Petlyuk arrangement in one shell, with a vertical partition between the feed side and the draw side, gave the dividing-wall column. A 1987 account of columns with vertical partitions reported energy use 20 to 35 % lower than other distillation arrangements and production columns that were easily controlled [11]. Operation and control of the dividing wall were studied by the late 1990s [12], and reviews in 2010 and 2011 surveyed its design, its industrial record and its current activity [13, 14]. By 2009 about forty columns were reported in service [4].

Control arrived after design. Most published structures held the three purities with the reflux, the side draw and the boilup, while the vapor split was fixed when the wall was placed and the liquid split was left as a handle for some further objective [4]. Ling and Luyben proposed the four-point structure, which keeps the energy near its minimum implicitly by controlling the heaviest component

Table 8.1 Variables, Their Roles and Design Values (Case 4 at $t = 0$)

Symbol	Role	Design	Range
x_D	CV, benzene in the distillate (mol %)	99.20	SP 97–99.9
x_S	CV, toluene in the side draw (mol %)	98.20	SP 94–99.8
x_B	CV, o-xylene in the bottoms (mol %)	99.20	SP 97–99.9
x_P	CV (four-point), o-xylene on the prefractionator's top stage (mol %)	1.57	SP 0.2–8
ℓ_D, ℓ_B	CV, drum and base levels (%)	50, 50	SP 20–80
L	MV, reflux (kmol/h)	548.9	0–1500
S	MV, side draw (kmol/h)	100.2	0–300
V	MV, boilup (kmol/h)	649.3	0–V_{cap}
β	MV, liquid split to the prefractionator	0.314	0.15–0.70
D, B	MV, distillate, bottoms (kmol/h)	100.3, 99.5	0–400
F	DV, feed rate (kmol/h)	300	200–400
z_1, z_3	DV, benzene, o-xylene in feed (mol %)	33.3, 33.3	15–50
ΔT_{sub}	DV, feed subcooling (K)	0	0–40
V_{steam}	DV, boilup the steam allows (kmol/h)	1500	400–1500
M_i, x_{ij}	states, stage holdup (kmol), fractions	1 or 2, profile	40 stages
M_D, M_B	states, drum, sump holdup (kmol)	20, 30	cap 40, 60
Q	output, reboiler duty (MW)	5.83	—

in the prefractionator, and tested it on feed rate and feed composition changes [4]. Their study of the same benzene, toluene and o-xylene separation on temperatures found that one temperature on the prefractionator side and three in the main column handled feed rate changes, while product purities deviated for composition changes greater than about 10 % [5]. Multivariable model predictive control, by then the industrial standard for interacting units [15], became the other candidate.

> ⧖ HISTORY **The Energy Valley, 1995 and 1999**
>
> In 1995 a study of operating the three-product column found that the boilup could have two solutions at one liquid split, and that at a higher side-draw purity a hole opened where no split met the specifications. It warned of serious problems, at least at high purity [10]. A 1999 study noted that industrial use had been limited, with difficulties in operation reported as one reason. With three purities controlled, two degrees of freedom remained for optimization, and the optimal surface proved narrow, very steep in one direction and moved by disturbances, its optimum lying along the bottom of a valley. The authors proposed holding the column there by feedback on a variable that characterizes optimal operation, and called the heavy key at the top of the prefractionator very interesting, though probably in need of a composition analyzer [16]. The four-point loop of this chapter feeds back exactly that measurement.

8.3 The Shell as Modeled

No plant stands behind the simulated column, whose layout is shown in Figure 8.1. Forty theoretical stages are arranged as four sections, 8 above the wall, 12 on each side of it and 8 below. The feed enters stage 7 of the prefractionator, counted from the top of the wall, and the side draw leaves stage 7 of the main side. A total condenser feeds a reflux drum and a partial reboiler with its sump closes the shell. The liquid leaving the top section is divided above the wall by a splitter, and the vapor leaving the bottom section is divided at the foot of the wall by the wall's position. Sizes, volatilities and limits are teaching values picked so that the design point hangs together, and each export is labeled a synthetic educational model.

Table 8.1 lists the variables by role. Time runs in minutes while flows are in kmol/h, so the code multiplies each balance by $1/60$, a factor the equations below leave out. The model integrates 138 states, a holdup and two mole fractions on each of the 42 stages (drum, 40 stages, sump), six flow states and six ledgers that count the moles, the benzene and the o-xylene entering and leaving.

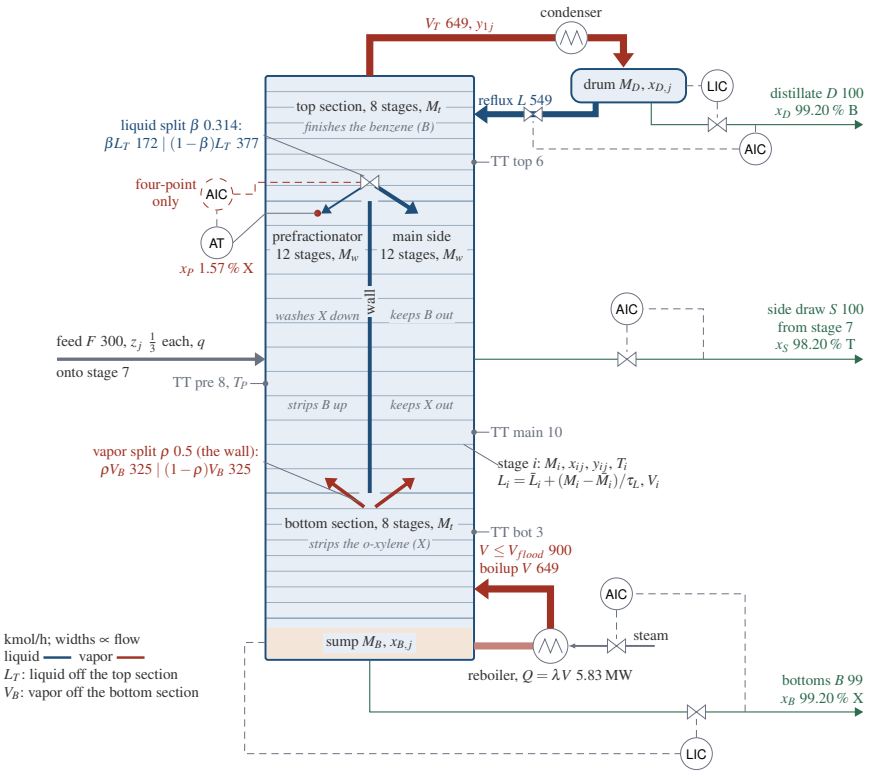

Figure 8.1 The dividing-wall column at its design point, following the application's Column view of case 4 at $t = 0$, where the split sits at the valley's minimum. Each section carries its job, the liquid split β above the wall and the vapor split ρ at its foot are drawn as the flows they send to the two sides (arrow widths proportional to flow), and the loops of the three-point structure are drawn with the fourth, four-point loop dashed in red, its controller bubble outlined dashed as well. Liquid runs down and vapor up, and every stream is labeled by name. TT marks the thermowells the inferential loops read. Flows in kmol/h. Process model only. Not to scale.

8.3.1 Stages, Equilibrium and Temperature

Each stage is an equilibrium stage with a liquid holdup, constant molar overflow and a linearized weir, as shown in Equation (8.1).

$$\begin{aligned} \frac{\mathrm{d}M_i}{\mathrm{d}t} &= L_{i-1} + V_{i+1} + F_i - L_i - V_i - S_i \\ \frac{\mathrm{d}(M_i x_{ij})}{\mathrm{d}t} &= L_{i-1}x_{i-1,j} + V_{i+1}y_{i+1,j} + F_i z_j \\ &\quad - (L_i + S_i)\,x_{ij} - V_i\,y_{ij} \\ L_i &= \bar{L}_i + \frac{M_i - \bar{M}_i}{\tau_L}, \qquad V_i = V_{i+1} + (1-q)\,F_i \end{aligned} \tag{8.1}$$

where i counts the stages of a section from its top, $j = 1,2,3$ are benzene, toluene and o-xylene, F_i is the feed on the feed stage and zero elsewhere, S_i is the side draw on the draw stage, q is the liquid fraction of the feed, and $\bar{L}_i$ and $\bar{M}_i$ are the weir's reference flow and holdup. The reference flows are those of a 500 kmol/h reflux, a split of 0.305, a 300 kmol/h feed and a 100 kmol/h draw, and $\tau_L = 0.1$ min. The reference holdup is $M_t = 2$ kmol on each stage above and below the wall and $M_w = 1$ kmol beside it. A draw takes at most the liquid there is on its stage. The vapor composition

follows constant relative volatility and the temperature a closed-form bubble point, as shown in Equation (8.2).

$$y_{ij} = \frac{\alpha_j x_{ij}}{\sum_k \alpha_k x_{ik}}, \qquad T_i = \frac{\Lambda}{\ln\left(\frac{P_0}{P}\sum_k x_{ik}\, e^{\Lambda/T_{b,k}}\right)} \tag{8.2}$$

where $\alpha = (2.4, 1, 0.44)$ relative to toluene, $\Lambda = 3890\,\mathrm{K}$ is one common latent heat over the gas constant, $P = P_0 = 1.013\,25\,\mathrm{bar}$, and the normal boiling points $T_{b,k}$ follow from α and toluene's 110.6 °C as 80.1 and 144.4 °C for benzene and o-xylene. One latent heat is the constant-molar-overflow assumption itself, and it makes α exactly constant. The latent heat $\lambda = R\Lambda = 32.3\,\mathrm{kJ/mol}$ turns the boilup into the duty $Q = \lambda V$, so 1 kmol/h of boilup is 0.008 984 MW. The stage stacks, the drum and the sump are a distillation module that the apps share, written for a reactor, column and recycle plant and moved into the framework when this column became its second user.

8.3.2 The Wall

The four sections are joined across the wall, as shown in Equation (8.3).

$$\begin{aligned} L_{P,0} &= \beta L_T, & L_{M,0} &= (1-\beta)\, L_T \\ V_{P,13} &= \rho V_B, & V_{M,13} &= (1-\rho)\, V_B \\ L_{bot,0} &= L_{P,12} + L_{M,12}, & V_{top,9} &= V_{P,1} + V_{M,1} \end{aligned} \tag{8.3}$$

where P and M are the prefractionator and the main side, index 0 marks the liquid entering a section's top stage and index $n+1$ the vapor entering its bottom stage, L_T is the liquid leaving the top section, V_B the vapor leaving the bottom section, the entering streams carry the composition of the stage they come from, and the two mixed streams carry the flow-weighted composition of their parts. The code computes the vapor bottom-up and the liquid top-down before any composition, and then evaluates the four stacks bottom-up. The vapor split $\rho = 0.5$ is equipment, not a handle.

8.3.3 Drum, Sump and Actuators

The drum is a total condenser's holdup and the sump a partial reboiler's equilibrium stage, as shown in Equation (8.4).

$$\begin{aligned} \frac{\mathrm{d}M_D}{\mathrm{d}t} &= V_T - L - D \\ M_D \frac{\mathrm{d}x_{D,j}}{\mathrm{d}t} &= V_T\,(y_{1j} - x_{D,j}) \\ \frac{\mathrm{d}M_B}{\mathrm{d}t} &= L_N - V - B \\ \frac{\mathrm{d}(M_B x_{B,j})}{\mathrm{d}t} &= L_N x_{N,j} - V y_{B,j} - B x_{B,j} \end{aligned} \tag{8.4}$$

where V_T is the vapor leaving the top stage at y_{1j}, L_N the liquid off the last stage at $x_{N,j}$, and $y_{B,j}$ is in equilibrium with the sump. The product purities are $x_D = 100\,x_{D,1}$, $x_S = 100\,x_{M7,2}$ on the draw stage, $x_B = 100\,x_{B,3}$, and x_P is the o-xylene on the prefractionator's top stage. Each manipulated flow follows its setpoint through a first-order lag, 0.2 min for the liquid flows, 1 min for the reboiler and 0.5 min for the splitter. As the drum or the base runs dry its pump delivers only the flow times $\min(1, M/(0.02\,M_{cap}))$. A subcooled feed enters with $q = 1 + c_{p,L}\Delta T_{sub}/\lambda$.

8.3.4 Instruments

A measurement packet carries everything the controllers see, and the application's test fails if it ever carries anything else. Level transmitters lag 0.1 min with 0.05 % of noise, flow meters lag 0.05 min with 0.3 kmol/h, and thermowells on four candidate stages per section lag 0.5 min with 0.05 K. Four analyzers on the shared analyzer module measure the three product purities and x_P. Each takes a sample every 10 min, after a 1 min transport delay, and publishes the result at the end of the following

Table 8.2 Parameters, Marked A When the Learner Can Change Them, H When the Panel Hides Them and S When Fixed in the Source

Symbol	Meaning	Value	Unit	
α_1, α_3	benzene, o-xylene volatility	2.4, 0.44	—	A
Λ	latent heat over R	3890	K	H
$T_{b,2}$	toluene normal boiling point	110.6	°C	H
ρ	vapor split to the prefractionator	0.5	—	A
M_t, M_w	stage holdup off, beside the wall	2, 1	kmol	A
τ_L	stage hydraulic lag	0.1	min	H
$M_{D,cap}, M_{B,cap}$	drum, base capacity	40, 60	kmol	A
V_{flood}	boilup at flooding	900	kmol/h	A
$c_{p,L}$	liquid heat capacity	0.16	kJ/(mol K)	H
—	specifications of x_D, x_S, x_B	99, 98, 99	mol %	A
$\tau_Q, \tau_V, \tau_\beta$	flow, reboiler, splitter lags	0.2, 1, 0.5	min	H
N_T, N_P, N_M, N_B	stages per section	8, 12, 12, 8	—	S
f, s	feed, draw stage (wall's top = 1)	7, 7	—	S
P	column pressure	1.01325	bar	S
$\bar{L}, \bar{\beta}, \bar{F}, \bar{S}$	weir references	500, 0.305, 300, 100	kmol/h	S
E	stage efficiency	1	—	S

cycle, with 0.01 mol % of noise (0.02 mol % on x_P). Every execution draws the same twelve normal deviates, so a seeded run replays bit for bit (Section 1.6.1).

8.3.5 Parameters, Initialization and Numerics

Table 8.2 gives every parameter. Each case opens lined out at its own split and setpoints, a state found by the Newton solver of Section 8.4 at the case's split and purity setpoints, with the holdups placed on the weirs and the drum and base at 50 %. The design point is an exact equilibrium, its largest derivative 4.5×10^{-15} per minute. Where no steady state exists at the case's split (case 8), the plant starts from the steady state at that split with the design's reflux and boilup and the side draw at the toluene in the feed. A fourth-order Runge-Kutta integrator advances the plant in 2 s steps, a third of the weir lag, the controllers act every 6 s, and one row is logged per 5 min. Halving the step moves case 4 at 4 h by 7.9×10^{-9} relative, and over a disturbed run the ledgers close the moles, the benzene and the o-xylene to 8.7×10^{-10} of the feed.

8.3.6 Assumptions

Constant molar overflow and constant relative volatility with one latent heat stand for the thermodynamics, so no energy balance is solved per stage and the volatilities do not drift with temperature. Benzene, toluene and o-xylene are close to ideal, and the error lies mostly there. The volatilities are teaching values. The Antoine fits of Willingham and coauthors, as the NIST Chemistry WebBook lists them [17], give 2.60 for benzene over toluene at 80 °C and 0.31, 0.34 and 0.37 for o-xylene over toluene at 70, 90 and 110 °C, so the 0.44 used is 19 to 42 % high and the toluene and o-xylene split is easier here than in a real column. The model's boiling points, 80.1, 110.6 and 144.4 °C, match the same fits' 80.1, 110.62 and 144.41 °C. The stages are theoretical with an efficiency of one, the shell has one pressure, and the vapor split is fixed, while a real wall's vapor split moves with the two sides' pressure drops. Flooding is a ceiling on the boilup, so the column has no flooded state. The flow loops, the condenser and the splitter are ideal lags, and the pressure and feed temperature loops are not shown.

8.4 The Spare Handle and Its Valley

8.4.1 Degrees of Freedom

The count follows Section 1.4.5. The drum and base levels take the distillate and the bottoms. Four handles remain, L, S, V and β, and three purities use three of them. With the purities fixed, the

component balances fix the three product flows, so the fourth handle can only redistribute the internal flows, as shown in Equation (8.5).

$$Q(\beta) = \lambda\, V(\beta)\big|_{x_D, x_S, x_B \text{ held}}, \qquad \beta^* = \arg\min_{\beta} Q(\beta) \tag{8.5}$$

The application computes $V(\beta)$ by Newton's method on every stage's two fractions and on L, S and V, with the three purities as equations, a numerical Jacobian, a step capped at 0.25 in any fraction and 30 % in any flow, and each stage projected into the simplex. The curve is walked from the minimum in steps of 0.005 until no solution exists or the boilup passes 1.6 times the flooding rate, and the minimum is refined by a parabola through its neighbors. It is computed at the purity setpoints in force, because those are what the loops hold, with the feed sampled at most every 10 min.

🕮 **Pause.** Try Exercise 8.1 before reading on.

At the design feed and the setpoints 99.2 mol %, 98.2 mol % and 99.2 mol %, the minimum is 649.3 kmol/h (5.83 MW) at $\beta^* = 0.314$, 72 % of flooding, with 548.9 kmol/h of reflux. The boilup is 788.5 kmol/h at 0.28 (21.4 % more) and 763.9 kmol/h at 0.40 (17.7 % more), and stays within 5 % of the minimum from 0.295 to 0.35. All three products can be made below flooding from 0.28 to 0.445. At the specifications the minimum would be 601.1 kmol/h at 0.305, so the 0.2 mol % margins cost 8.0 % of the duty. The heavier feed of cases 4, 5 and 9 (benzene 25, o-xylene 41.7 mol %) moves the minimum to 0.339 and 668.1 kmol/h (6.00 MW), as shown in Figure 8.3a. The plant lands on the curve when it lines out. In case 2 the split is moved every 3 h, and just before each move the duty is 7.31 against the curve's 7.31 (+0.03 %), 6.21 against 6.20 (+0.30 %) and 5.83 against 5.84 MW (−0.22 %), as shown in Figure 8.2.

8.4.2 An Independent Check and the Shortcut Bounds

A steady-state solver written for the technical review from the stage equations alone, by the bubble-point method at fixed flows with Newton's method on (D, S, V) for the purities, shares no code with the application. It reproduces the curve to 0.01 kmol/h at twenty splits from 0.275 to 0.50, puts the minimum at 0.3139 and 649.26 kmol/h against the application's 0.314 and 649.28 kmol/h, and returns the design purities at the application's design flows to six decimals, as Table 8.3 lists.

Underwood's method [7] at infinite stages, sharp splits and a saturated-liquid feed gives the roots of Equation (8.6) and the minimum vapor of each binary split. For $q = 1$ it is Equation (1.8) of Chapter 1.

$$\sum_j \frac{\alpha_j z_j}{\alpha_j - \theta} = 1 - q, \qquad V_{min,k} = F \sum_{j \le k} \frac{\alpha_j z_j}{\alpha_j - \theta_k} \tag{8.6}$$

With $\theta_1 = 1.4667$ and $\theta_2 = 0.5625$, the benzene and toluene split needs 257.1 kmol/h and the toluene and o-xylene split 359.2 kmol/h. The Petlyuk minimum is the larger of the two [9], here 359.2 kmol/h against 514.3 kmol/h for the direct sequence, a saving of 30.2 %. The wall fixes the vapor split, so the prefractionator can reach its preferred split (195.9 kmol/h of vapor) only when $\rho V \ge 195.9$, which puts the fixed-wall minimum at 391.8 kmol/h (3.52 MW) and a split of 0.231. The application's 649 kmol/h is 1.66 times that because its sections are short. Fenske's minimum [6] at the design compositions is 11.1 stages for the toluene and o-xylene split and 11.6 for the benzene and toluene split, against 14 stages for each (ratios of 1.26 and 1.21), and Gilliland's correlation [8] for the controlling split gives 894 kmol/h, an estimate rather than a bound. The prefractionator also needs a liquid-to-vapor ratio of 0.53 above the feed with its six stages, against the 0.344 of the preferred split, so the finite-stage minimum sits at a larger split, 0.314 against 0.231.

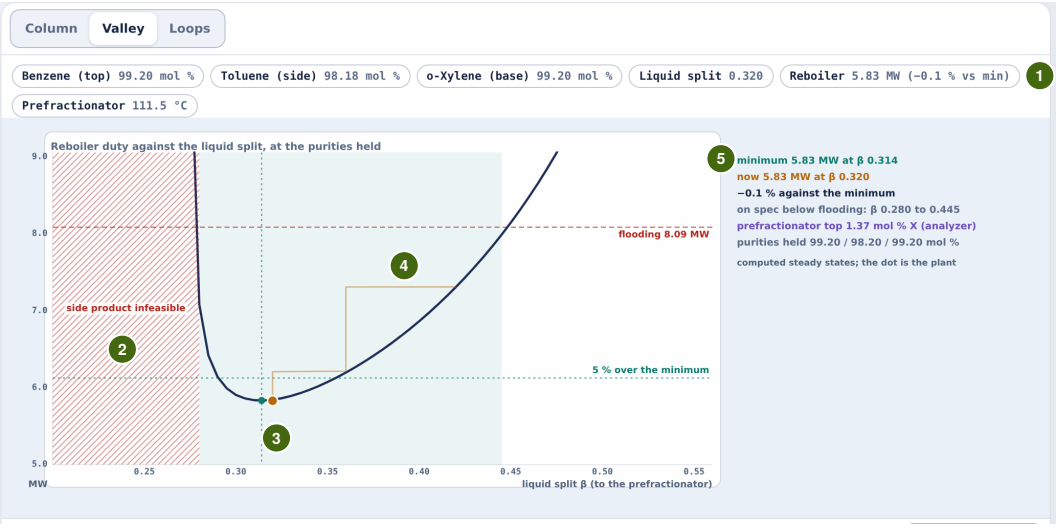

Figure 8.2 The energy valley of case 2 at $t = 9$ h on the Valley view, the three-point structure in Auto, light theme, cropped to the plot and its status chips. (1) The chips: benzene at 99.20 mol %, toluene at 98.18 mol % and still settling, o-xylene at 99.20 mol %, and the reboiler at 5.83 MW. (2) The zone below 0.28, hatched, where no boilup up to flooding (8.09 MW) makes the side product. The splits from 0.280 to 0.445 are on specification below flooding, and the dotted line lies 5 % over the minimum. (3) The computed curve holds 99.20 mol %, 98.20 mol % and 99.20 mol % at the design feed, with its minimum at 5.83 MW and a split of 0.314 (left dot). The plant's dot reads 5.83 MW at 0.320, 0.1 % under the minimum and 0.22 % under the curve at that split. (4) The plant's path down the gentle side through the case's moves from 0.42 to 0.36 at 3 h and 0.32 at 6 h. (5) The readout, with the analyzer's o-xylene on the prefractionator's top stage at 1.37 mol %.

Table 8.3 The Valley Against an Independent Solver and the Shortcut Bounds

Quantity	Application	Check
minimum split β^*	0.314	0.3139
minimum boilup (kmol/h)	649.28	649.26
boilup at 0.28, 0.40 (kmol/h)	788.5, 763.9	788.50, 763.89
largest difference on 20 splits (kmol/h)	—	0.01
RGA diagonal at the design	12.7, 0.44, 24.9	12.68, 0.44, 24.86
Underwood, Petlyuk V_{min} (kmol/h)	—	359.2
Underwood, $\rho = 0.5$ fixed (kmol/h, β)	—	391.8, 0.231
Fenske N_{min} against 14 stages	—	11.1, 11.6
vapor split for equal pressure drops	0.5	0.502

✦ DID YOU KNOW? Thirty Percent, If the Column Is Run Right

Studies of the three-product column report typical savings of about 30 percent in both energy and capital cost against two columns in series [16]. A four-product column with two side draws can save up to 40 percent of the energy, a saving that is lost when the column runs away from its optimum vapor split [18].

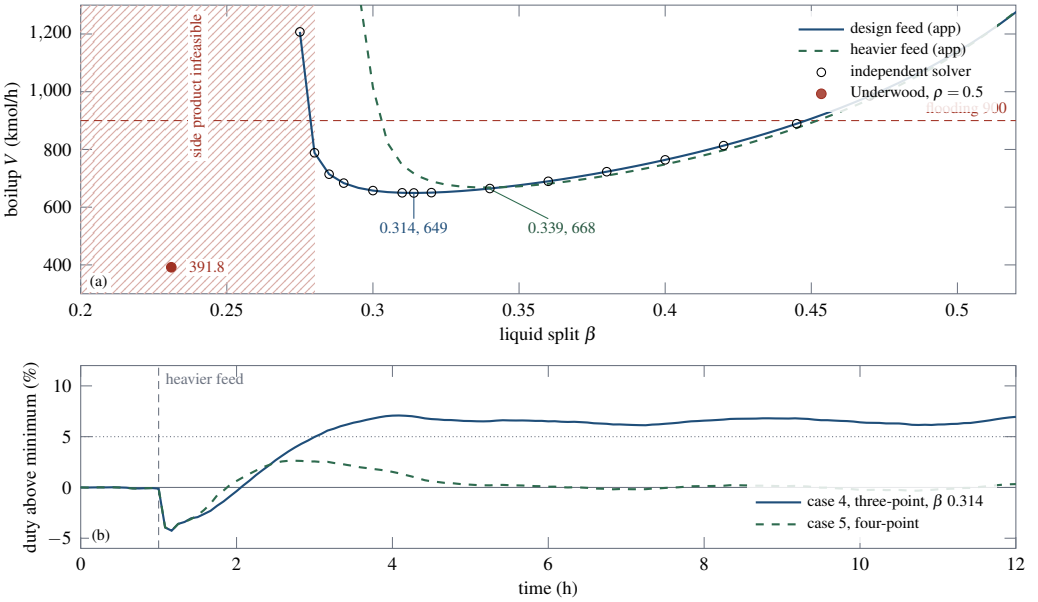

Figure 8.3 The valley and what the four-point loop does with it. (a) Boilup against the liquid split at the setpoints from the application at the design feed (solid) and at the heavier feed (dashed), the independent solver's points on the design curve (open circles), and the Underwood minimum with the vapor split fixed at 0.5 (filled dot). The hatched zone is the design feed's, and the flooding rate is the thin dashed line. (b) The duty against the valley's minimum in cases 4 (solid) and 5 (dashed) as shipped, every logged row, with the feed made heavier at 1 h.

8.4.3 Why One Side Is Steep

However hard the reboiler boils, each side of the wall has a ceiling on its liquid-to-vapor ratio set by the two splits, as shown in Equation (8.7).

$$\begin{aligned} \left(\frac{L}{V}\right)_{P,\text{above feed}} &= \frac{\beta L}{\rho V} < \frac{\beta}{\rho} \\ \left(\frac{L}{V}\right)_{M,\text{below draw}} &= \frac{(1-\beta)L - S}{(1-\rho)V} < \frac{1-\beta}{1-\rho} \end{aligned} \tag{8.7}$$

because $L < V$ whenever a distillate is drawn. The prefractionator needs a ratio of about 0.5 above the feed to wash the o-xylene down, 0.53 at the minimum against a ceiling of 0.63 and 0.49 at a split of 0.28 against 0.56. As the split falls the ceiling closes on the need, and the boilup climbs steeply, 788.5 kmol/h at 0.28, 1207 at 0.275, and about five times the flooding rate (4543 kmol/h) at 0.25. The main side below the draw needs 0.65 to 0.85, and its ceiling comes down to that only near a split of 0.675 (the boilup is 98 029 kmol/h at 0.674, and no steady state exists at 0.6765), so the high side rises gently over a far wider range. Below the minimum the o-xylene on the prefractionator's top stage rises, from 1.57 to 3.77 and 4.56 mol % at 0.28 and 0.275, and that o-xylene can leave only with the toluene. The asymmetry is in the stage balances, not in a correlation. The vapor split of 0.5 agrees with a wall on the centerline. With 12 trays on each side and equal pressure drops, taking a tray's drop as 40 % dry and the liquid head's crest share as 40 %, the two sides' weir-crest loads at the design are 550 and 567 (3 % apart) and equal drops need $\rho = 0.502$, between 0.488 and 0.515 across every split that runs below flooding.

🕮 **Pause.** Try Exercise 8.2 before reading on.

8.5 Loops as Implemented

8.5.1 Pairing and the Shared PI Law

The three-point structure pairs the reflux with the benzene in the distillate, the side draw with the toluene in the side draw and the boilup with the o-xylene in the bottoms, while the drum level moves the distillate and the base level the bottoms. At the design the relative gain array [2] of reflux, side draw and boilup against the three purities, each gain taken with the other two handles held and the levels closing the balance, has the diagonal 12.7, 0.44 and 24.9 and off-diagonal elements down to −24.4. Of the six pairings two have every element positive, the one in force and the swap of reflux and side draw, whose element of 0.006 for the side draw on the benzene rules it out. On the distillate and the bottoms the gain of the reflux alone is 476 and 124 times the gain of reflux and boilup moved together, the ill-conditioning of high-purity distillation [3]. All loops use the framework's velocity-form PID of Section 1.6.2, here without derivative action. In the loops' own symbols it reads as Equation (8.8).

$$\begin{aligned} e_k &= \sigma\,(r_k - y_k), \qquad s_k = s_{k-1} + K_c\,(e_k - e_{k-1}) \\ \hat{u}_k &= u_{b,k} + s_k + I_{k-1} \\ u_k &= \min(u_{max}, \max(u_{min}, \hat{u}_k)) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - \hat{u}_k}{\tau_I}\right]\Delta t_k \end{aligned} \tag{8.8}$$

where r_k is the setpoint, y_k the measurement, $u_{b,k}$ the feedforward term (zero when it is off, the starting flow being carried in the integral), and s_k, $\hat{u}_k$ and σ are the S_k, v_k and action a of Equations (1.14) to (1.17), with the tracking time equal to τ_I. The limit acts on the output only, and the integral alone carries the anti-windup. The action is $\sigma = +1$, reverse acting in Chapter 1's convention, for the reflux and boilup loops, and -1, direct acting, for the side draw (more draw, dirtier toluene), the levels and the split loop. Integral action is switched off once the integral time reaches 600 min. The levels move 10 kmol/h per percent with $\tau_I = 20$ min. On tray temperatures the purity loops move 50, 20 and 30 kmol/h per mol % with integral times of 10, 20 and 10 min and execute every 0.1 min. On the analyzers they run at 2, 1 and 2 with 60, 120 and 60 min and execute on each report with Δt_k the time since the last. Each measurement brings its own tuning, because on the analyzers the temperature gains run the side draw to zero 1.52 h into case 7. The boilup's upper limit is $V_{cap} = \min(V_{flood}, V_{steam})$, and the valve ranges of Table 8.1 limit the rest. Switching structure or mode is bumpless because every idle controller tracks its flow.

8.5.2 Tray Temperatures With an Analyzer Trim

By default each purity loop reads one tray temperature converted by a straight line fitted on the steady-state model at the design point, corrected by each analyzer report, as shown in Equation (8.9). The update is the laboratory bias update of Equation (1.11) with a first-order filter on the bias.

$$\begin{aligned} \hat{x} &= x_0 + m\,(T_k - T_{k,0}) + b_f, \qquad \tau_b \frac{\mathrm{d}b_f}{\mathrm{d}t} = b - b_f \\ b &\leftarrow b + \lambda_b \left(x_{an} - (x_0 + m(T_k - T_{k,0}))\big|_{t_{sample}} - b\right) \end{aligned} \tag{8.9}$$

where the slope m is −0.139, −0.750 and +0.118 mol %/K on top-section stage 6, main-side stage 10 and bottom-section stage 3, $\lambda_b = 0.5$ on the shared bias updater, which clamps the bias to ±3 mol %, $\tau_b = 5$ min, and the report is compared with the inferential's value when its sample was drawn. The bias starts lined out. Each loop can instead read its analyzer alone or the uncorrected temperature, which case 7 compares. Volume 12, Chapter 4 shows how a tray temperature that stands in for a product purity also answers to pressure, and Volume 5, Chapter 3 what an analyzer's delay does to a composition loop.

8.5.3 The Split Loop, the Feedforward and the MPC

The four-point structure adds the split on x_P, the o-xylene on the prefractionator's top stage read by its analyzer, at 0.015 per mol % with $\tau_I = 40$ min. More liquid on the feed side washes the o-xylene down [4], so the split must rise when x_P rises. The loop is direct acting ($\sigma = -1$) in Chapter 1's convention. Its setpoint, 1.57 mol %, is the value at the design valley's minimum. In the steady-state model, holding x_P there keeps the boilup within 0.31 % of each feed's minimum for six feed changes, while the best prefractionator tray temperature held the same way costs from 29.2 to 219 % when the benzene or the o-xylene changes, and a split fixed at 0.314 costs up to 13.97 %. In the language of self-optimizing control [1], x_P is a variable whose constant setpoint keeps the loss small. Volume 9, Chapter 8 meets an optimum that moves with the weather and follows it with a supervisory search instead. The temperature remains an Advanced option. The feed-rate feedforward adds $K(F_m - F_{ref})$ to the reflux, side draw and boilup loops' bias, with $K = 1.669$, 0.333 and 2.004 (the flows per unit of feed at the specifications and a split of 0.305) and F_{ref} the measured feed when it is switched on or the run starts.

The MPC runs every 5 min over a 5 h horizon of 60 steps on the apps' shared linear MPC module (Section 1.6.3) and moves L, S, V and β together. Its outputs are the logarithms of the three impurities and of x_P, and its model is the plant's step response at the design point fitted to first order plus dead time per pair, with the measured feed as a disturbance and an output bias held over the horizon. Reflux alone and reflux with boilup are identified apart because the column is ill-conditioned. The objective is given by Equation (8.10).

$$\begin{aligned} \min \; & \sum_{k,i} q_i \left(y_{ik} - r_i\right)^2 + \sum_i \left(400\, s_i^2 + 40\, s_i\right) \\ & + w_Q\, 10^{-3} \sum_k V_k + \sum_{j,k} \frac{w}{\delta_j^2}\, \Delta u_{jk}^2 \\ \text{s.t. } & y_{ik} \le r_i + s_i, \quad s_i \ge 0 \end{aligned} \tag{8.10}$$

where $q = (0.2, 0.2, 0.2, 2)$, the band applies to the three impurities only, $\delta = (10, 2, 10, 0.01)$ is a typical move of each handle, $w_Q = w = 1$, the moves are capped at 30, 6 and 30 kmol/h and 0.01 per execution, the split stays within 0.15 and 0.70, and a failed solve hands the handles back to the PI loops. On case 5 it solves 96 times with no fallback.

8.5.4 Annunciators and Model Separation

FLOODING marks the boilup request at V_{flood}, STEAM LIMIT at a lower steam supply, OFF SPEC a true purity below its specification, and SIDE PRODUCT INFEASIBLE a split below the lowest one at which the valley has a point under flooding. None of them trips the column. The flooding rate and the specifications are teaching values, and on a plant they would come from a tray or packing rating and from the buyer's contract. The application keeps three models apart. The plant is Section 8.3, the valley is its steady state at the setpoints, read only by the display, the infeasibility flag, the challenge's score and the helper that moves the split to the minimum, and the MPC's model is a fixed set of step responses.

8.6 Nine Cases, Measured

Table 8.4 names the nine cases. The figures in this section were taken headless at the default seed, advancing the engine one physics step at a time exactly as the page does (Section 1.6.1), each case opened the way the page opens it and each move of the learner applied at the time the text gives. Table 8.5 collects the scorecard tiles at the instants the cases are read. A case gives the same run whether it is opened first, picked after another case or restarted with Reset. At the fastest speed of 300 min per second a 12 h case plays in 2.4 s. The exported CSV file carries the structure and measurement in force, every setpoint, purity reading and handle, the split, the prefractionator's

Table 8.4 The Nine Cases

#	Settings and events	Teaching point
1	three-point, split 0.40, learner moves it to 0.34, 0.30, 0.28	the spare handle
2	split 0.42, to 0.36 at 3 h, 0.32 at 6 h, learner 0.29 at 9 h	the curve against the plant
3	purity loops in Manual, bottoms loop, then two loops	interaction, RGA
4	three-point at 0.314, heavier feed at 1 h	the valley moves
5	four-point, the same feed, then the MPC	a loop that holds it
6	four-point, steam 740 kmol/h, feed +20 % at 1 h	a reboiler limit
7	tray temperatures only, heavier feed at 1 h	inferential offset
8	split 0.25, side draw and boilup on stations	infeasible side product
9	challenge, 12 h, three-point at 0.40	the whole problem

Table 8.5 Scorecard Tiles, Purities and Split at the Instants Each Case Is Read

Case	Settings and instant	Duty (MW)	Against minimum (%)	Split, β^*	x_D / x_S / x_B (mol %)	Also
1	as shipped, 1 h	6.86	+17.5	0.400, 0.314	99.20 / 98.18 / 99.21	
1	split 0.34 at 1 h, 4 h	5.98	+2.6	0.340, 0.314	99.22 / 98.18 / 99.23	toluene low 98.11
1	split 0.30 at 4 h, 7 h	5.86	+0.5	0.300, 0.314	99.19 / 98.18 / 99.19	toluene low 97.95
1	split 0.28 at 7 h, 12 h	6.86	+17.5	0.280, 0.314	99.14 / 98.22 / 99.10	3.2 % under the curve
2	as shipped, 9 h	5.83	−0.1	0.320, 0.314	99.20 / 98.18 / 99.20	−0.22 % against the curve
4	as shipped, 5 h	6.39	+6.5	0.314, 0.339	99.21 / 98.18 / 99.22	x_P 3.22
4	as shipped, 8 h	6.40	+6.6	0.314, 0.339	99.19 / 98.20 / 99.18	
5	as shipped, 8 h	6.01	+0.1	0.345, 0.339	99.19 / 98.21 / 99.19	x_P 1.59
5	MPC, 8 h	6.05	+0.7	0.345, 0.339	99.15 / 98.24 / 99.20	benzene low 98.12
6	as shipped, 3 h	6.65	−5.0	0.312, 0.314	99.19 / 98.26 / 98.62	STEAM LIMIT from 95 min
6	setpoints 99.05 / 98.05 / 99.05 at 3 h, 7 h	6.57	−0.3	0.312, 0.307	99.04 / 98.02 / 99.06	boilup 731
7	as shipped, 3 h	5.82	−3.1	0.314, 0.339	99.36 / 98.01 / 99.02	loops read 99.20 / 98.17 / 99.21
8	as shipped, 2 h	5.86	+0.5	0.250, 0.314	99.21 / 97.31 / 98.09	x_P 6.38, INFEASIBLE

analyzer and temperature, the duty and the valley's minimum, the status flags, the disturbances, parameters and tunings. The true compositions and the balances sit in columns marked as truth, apart from the measurements a controller may use.

8.6.1 Walking the Valley

Case 1 starts at a split of 0.40 with the purities on their setpoints and the duty 6.86 MW, 17.5 % above the minimum. Moved to 0.34 at 1 h, the duty is 5.98 MW at 4 h, and moved to 0.30 it is 5.86 MW at 7 h. Moved to 0.28 it climbs, 5.95, 6.37, 6.62 and 6.72 MW at 8, 9, 10 and 11 h, toward the curve's 7.08 MW, and it is still 3.2 % under the curve at 12 h. On the curve the three moves are worth −0.89, −0.06 and +1.17 MW, so the smallest move changes the duty most. After each of the first two moves the duty is within 1 % of the curve for good after 2.31 and 2.35 h. Each move down dips the toluene first, to 98.11 mol % 83 minutes after the move to 0.34 and to 97.95 mol % half an hour after the move to 0.30, while the side loop catches up. In case 2 the learner's move to 0.29 at 9 h leaves the plant at 6.11 MW against the curve's 6.14 at 12 h, and within 1 % of the curve from 11.76 h.

> **TRY IT IN THE APP** **1 · The Fourth Handle: Move the Liquid Split**
>
> Load scenario 1 and press Run. At 1 h the three purities sit on their setpoints and the reboiler duty reads 6.86 MW, 17.5 % above the valley's minimum. Type 0.34 into the Liquid split box and run on. By 4 h the duty is 5.98 MW, while the purities have hardly moved.

8.6.2 Three Loops That Fight

Case 3 starts with the purity loops in Manual. With only the bottoms loop closed, its setpoint raised from 99.2 to 99.4 mol % reaches 99.35 mol % 9 min later on 0.4 kmol/h more boilup. The extra vapor leaves as distillate and takes toluene with it, so the benzene falls below 99.0 mol % 12 min after the move and stands at 98.86 mol % at 6 h. With the distillate loop closed as well, after Reset, the same move takes 75 min, 8.3 times as long, and by 8 h reflux and boilup have both risen about 45 kmol/h (the duty 6.24 MW). The relative gain of 24.9 for the boilup on the bottoms says the same thing, because with the other loops closed the boilup moves the bottoms about 25 times less than it does alone.

8.6.3 A Valley That Moves

Case 4 makes the feed heavier at 1 h. The distillate falls to 74.3 kmol/h and the bottoms rises to 125.2 kmol/h by 1.5 h as the levels pass the change through, and the toluene dips to 97.99 mol % and the o-xylene to 98.95 mol % before the loops recover them. The split stays at 0.314, now on the steep side of the new curve, and the duty is 6.39 MW at 5 h, 6.5 % above the new minimum, and 6.6 % above at 8 h, with x_P at 3.22 mol % from 1.57, as shown in Figure 8.3b. The helper that moves the split to the computed minimum, pressed at 5 h, takes it to 0.339, and by 10 h the duty is 5.97 MW. The helper uses the model and the true feed. Chapter 11 shows where an optimizer stops when its model is not exact. Case 5 repeats the feed change on the four-point structure. Without being told where the minimum went, the split loop takes the split to 0.345 by 8 h, a little on the gentle side of 0.339, the duty is 6.01 MW, 0.1 % above the minimum, and x_P is back at 1.59 mol %. The toluene dips to 97.98 mol % and the o-xylene to 98.86 mol % on the way.

The MPC on the same case also reaches the valley, with the split at 0.350 at 5 h, 0.345 and 0.7 % above the minimum at 8 h, but lets the benzene fall to 98.12 mol % and the toluene to 97.31 mol %, where the four-point PI loops hold the benzene above 99.08 mol %. It corrects an unmeasured feed change only through its output bias, and its first-order fits, with lags up to 290 min, describe responses that have two time scales.

> **TRY IT IN THE APP** **5 · Four-Point Structure: A Loop on the Prefractionator**
>
> Load scenario 5 and press Run. At 1 h the feed turns heavier. Watch the split on the Valley view: the loop, which has no model of the valley, takes it to 0.345 by 8 h, against the new minimum at 0.339, and the duty ends 0.1 % above it. Then choose MPC, press Reset and run again. It reaches the valley too, but the benzene falls to 98.12 mol %.

8.6.4 A Reboiler Limit and Inferential Measurement

Case 6 caps the steam at 740 kmol/h of boilup and raises the feed to 360 kmol/h at 1 h, where the valley's minimum at the setpoints is 779 kmol/h. STEAM LIMIT comes on at 95 min and the bottoms loop loses its handle, the o-xylene falling to 98.62 mol % by 3 h while the benzene and toluene hold. Lowering the setpoints to 99.05, 98.05 and 99.05 mol % lowers the minimum to 734 kmol/h, and by 7 h the boilup is 731 kmol/h with all three products on specification. Case 7 runs the loops on tray temperatures alone through the heavier feed. At 3 h the loops read 99.20, 98.17 and 99.21 mol % while the true purities are 99.36, 98.01 and 99.02 mol %, because the calibration made at the design feed misreads the heavier one. With the analyzer bias switched on at 3 h, the true purities are 99.17, 98.20 and 99.16 mol % by 5 h. On the analyzers alone, with their own tuning, the benzene falls to 86.0 mol % and the toluene to 81.2 mol % before the reports arrive.

Table 8.6 The Challenge Over 12 h, Default Seed

Settings	Low x_D (mol %)	Low x_S (mol %)	Low x_B (mol %)	Worst hour (%)	Verdict
three-point, 0.40 (shipped)	98.63	98.15	98.13	+12.4	fails
four-point and feedforward	99.11	98.07	99.07	+0.3	passes
four-point alone	98.61	98.15	98.07	+3.6	fails
feedforward alone	99.09	98.03	99.02	+12.4	fails
three-point, 0.339, feedforward	99.11	98.08	99.07	+0.0	passes
MPC	98.29	98.15	99.15	+16.7	fails

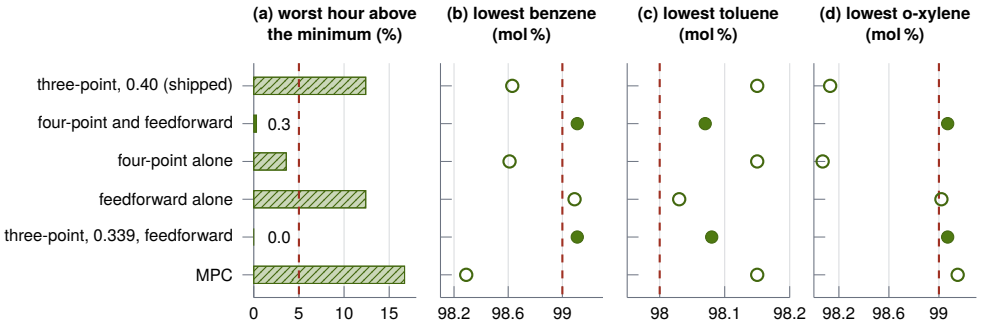

Figure 8.4 The challenge of Table 8.6 as a chart. (a) The worst judged hour against the valley's minimum, with the 5 % pass line dashed. (b) to (d) The lowest sample of each purity, with its specification dashed. Dark bars and filled dots are settings that pass, hatched bars and open dots settings that fail. Notice that the four-point structure alone holds the energy but lets the benzene and the o-xylene fall below their specifications, that the feedforward alone holds the purities but not the energy, and that only the settings with both a split on the minimum and the feedforward pass on every count.

🕮 **Pause.** Try Exercise 8.3 before reading on.

8.6.5 A Split That Cannot Be Fixed With Steam

Case 8 starts at a split of 0.25, where the curve has no point under flooding. With the side draw at 100 kmol/h and the boilup raised by hand, the side product holds 97.32 mol % toluene at the starting boilup, 97.31 at 750 kmol/h, 97.30 at 850 and 97.31 at 900 (flooding), while x_P reaches 7.2 mol %. The steady-state sweep at that split, with the benzene held, gives 97.32 to 97.34 mol % at every boilup from 650 to 900 kmol/h. Moved to 0.31, the side product reaches 98.83 mol % 2 h later. In Auto at that split, the side loop runs the draw to zero.

8.6.6 The Challenge

The challenge runs 12 h on the three-point structure at a split of 0.40. The feed turns heavier over 30 min from 2 h and rises 20 % over 30 min from 7 h. It passes when, from 2 h, every 10 min sample of the true purities is at or above 99 mol %, 98 mol % and 99 mol %, and when each of the six whole hours that start 2 h after a feed change has a mean duty within 5 % of the valley's minimum at the feed and setpoints in force. The changes arrive as ramps, the way a tank blend or a rate move arrives, although measured as steps the suggested settings also pass, with the toluene at 98.02 mol %. Table 8.6 lists the settings tried, and Figure 8.4 draws them as a chart. As shipped, the first judged hour is 12.4 % above the minimum, and the rate change drops the o-xylene to 98.13 mol %. The four-point structure fixes the energy and the feedforward fixes the rate change, and neither alone passes. Typing the new minimum, 0.339, on the three-point structure with the feedforward also passes, but only because the day's feed is known in advance.

The suggested settings pass at all 32 seeds from 0 to 31, with the lowest toluene between 98.040 and 98.107 mol % (mean 98.070, standard deviation 0.015). At the default seed and seeds 0 to 3 the lowest toluene sample falls between 7.33 and 7.5 h, in the half hour after the rate change starts. Halving the step leaves it at 98.073 mol %. The shipped settings fail at all 32. The toluene margin is thin, but it is set by the rate change and not by the noise.

DID YOU KNOW? Two Fixes That Fail Alone

From the simulator: on the twelve-hour challenge the four-point structure alone fails, with a worst hour 3.6 % above the minimum and the o-xylene down to 98.07 mol %, and the feedforward alone fails at 12.4 %. Together they pass at 0.3 %, while the MPC fails at 16.7 % (Table 8.6).

The MPC fails for its model's structure, not for a code defect. At the default seed and seeds 0 to 3, the benzene falls to between 98.27 and 98.33 mol % after the feed change and the worst hour is 15.8 to 17.0 % above the minimum. Its hourly excess falls from 16.7 through 14.5, 12.8, 10.6 and 8.0 to 4.5 %. The split is not the cause, since at 5 h it sits at 0.354 against the minimum's 0.339. The purities are, with the toluene at 98.61 mol % and the o-xylene at 99.33 mol % above their setpoints while the benzene, at 98.93 mol %, is still under its own. The slow directions of its step-response model take hours to bring all three back.

The side draw carries a visible fuzz, 0.58 to 0.78 kmol/h root mean square from one logged row to the next on lined-out stretches of cases 1, 2 and 4. It is the side loop answering thermowell noise. The loop's gain through the calibration is 15 kmol/h per kelvin, and 0.05 K of noise gives 0.75 kmol/h, against a measured standard deviation of the request of 0.82 kmol/h. The fuzz vanishes with every random draw set to zero and does not change when the step is halved.

8.7 Technology in Practice

What Measures It

Figure 8.5 draws the loops of Section 8.5 with their devices. Each purity loop reads a tray temperature, and an analyzer above it corrects the conversion. Monitoring and controlling distillation columns is a very common use of process chromatographs, which quantify the purity of the separated products [19], and the chapter's four analyzers, sampling every 10 min and reporting a cycle later, stand for such instruments. The tray temperatures come from thermocouples or resistance thermometers in thermowells, and nearby elements on a column can read differently where the flow patterns differ [19].

TECHNOLOGY Sensor: A Resistance Thermometer in a Tray Thermowell

A resistance temperature detector is a winding of pure metal wire, usually platinum, whose resistance rises with temperature. It is relatively insensitive but very linear [19]. Platinum elements cover about -200 to 800 °C and are chosen for their accuracy and precision [20]. On a column the element sits in a thermowell, so that it can be replaced without depressurizing, and the well adds a first-order lag. An air gap at the bottom of the well adds a second lag that can push a loop toward oscillation, and transmitters close together on a column can disagree where the flow patterns differ [19]. On a trend a badly seated element shows as a reading that trails its neighbors. The model's thermowells are a 0.5 min lag with 0.05 K of noise.

What Moves It

The reflux, the side draw, the distillate and the bottoms are liquid flows, and the boilup is set by the steam to the reboiler. In the model each manipulated flow follows its setpoint through a first-order lag, 0.2 min for the liquid flows and 1 min for the reboiler, which stands for a flow loop around a valve. The fourth handle is the liquid split above the wall, the one final element a conventional column does not have, and the model keeps the vapor split at its foot fixed at $\rho = 0.5$.

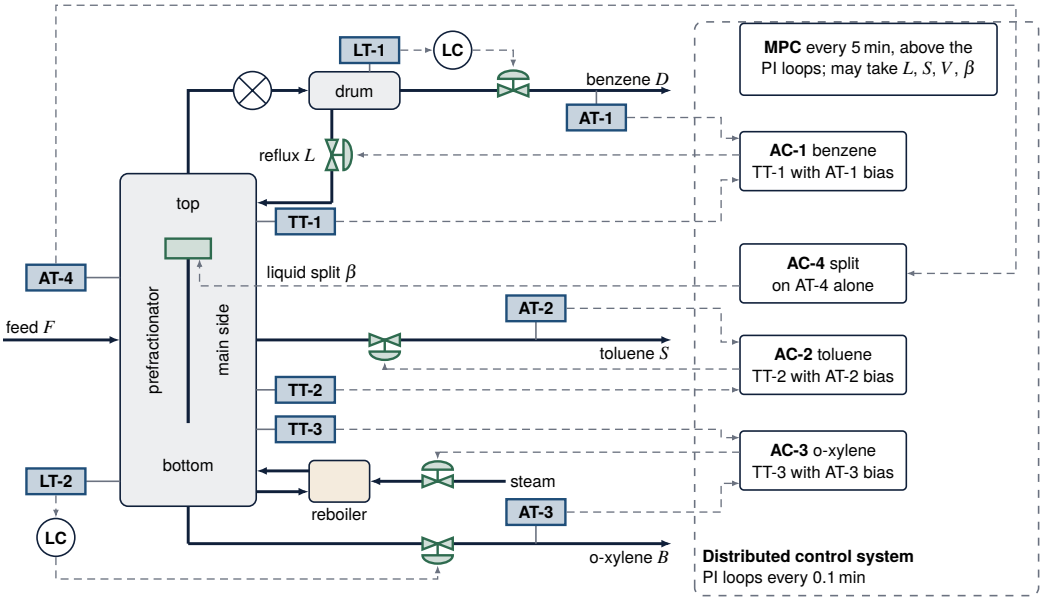

Figure 8.5 The loops of the dividing-wall column with their hardware. Lettered boxes are sensing devices: a tray temperature (TT) and a product analyzer (AT) behind each purity loop, the analyzer on the prefractionator's top stage (AT-4) and the two level transmitters, whose level controllers (LC) are drawn beside their vessels. Valve symbols and the splitter are final elements, and dashed arrows are signals. Notice that each purity loop reads its tray temperature through the analyzer's bias update, that the split loop alone reads an analyzer directly, and that the splitter above the wall is the one handle a conventional column lacks. The feed-rate feedforward and the flow loops are left out. Not to scale, and not a piping diagram.

TECHNOLOGY **Final Element: The Liquid Splitter Above the Wall**

The liquid split is made outside the shell. All the liquid leaving the trays above the wall can be drawn into an external accumulator and returned to the two sides under flow control [16], and controlling one of the two streams is enough, because tray overflow regulates itself [10]. A laboratory column split its reflux with a swinging funnel driven by solenoids, an on-off device that a PI controller set through pulse-width modulation [18]. A return valve with excess stem friction moves in jerks and ignores small reversals [19], so on a trend the split moves in steps. The vapor split is fixed by the wall, since vapor in an open column cannot be throttled like flow in a pipe [10]. The model's splitter follows its setpoint through a 0.5 min lag between 0.15 and 0.70.

Where the Controller Runs

A distributed control system runs nested and cascaded loops under a supervisory level, with cycle times from milliseconds to minutes [21]. The app executes its temperature loops every 0.1 min, its analyzer loops on each report and its predictive controller every 5 min, inside that range. The PI loops, the bias update and the feedforward sum belong to the regulatory layer, and the predictive controller to the layer above it, which takes the four handles when it runs. Industrial predictive controllers skip the bias update of an output whose analyzer has not reported, and they must cope with ill-conditioning such as that of adjacent tray temperatures in a column, where small errors call for large moves [15]. Section 8.5 measures the same property with the relative gain array.

IN PRACTICE **Temperature Control and Purity Give-Away**

Columns are widely controlled on a tray temperature instead of a product analyzer, because chromatographs and similar analyzers carry large delays and high costs. The practice assumes the product is on specification while the temperature is held, which is hard to keep when the feed composition or rate changes in a multicomponent column, and pressure moves the temperatures too [22]. Plants also tend to purify beyond need. Energy guidance asks for steam use to be compared with reflux and purity daily, and the reflux lowered in small steps until the purity is just met, which lowers the reboiler duty [23]. In this chapter the 0.2 mol % margins over the specifications cost 8.0 % of the duty.

What Is Modernizing It

Three developments bear on this column. A vapor split valve on a laboratory column has been shown to work as a handle under closed-loop temperature control, with feedback correcting the valve's inaccuracy [18]. It would add the fifth handle that the model keeps fixed. Soft sensors, which estimate a composition from several measurements, have distillation as their main target, and keeping their models up to date is the most critical issue that engineers report [22]. Such a sensor would replace the single-temperature lines of Equation (8.9), which case 7 shows misreading a heavier feed. And multivariable predictive control, with a solid foundation in refining and petrochemicals [15], has been suggested for this column in a nonlinear form built on a physical model [24], where the chapter's predictive controller uses one fixed linear model.

8.8 Limitations

Every size and setting is a generic teaching value. Constant molar overflow, constant relative volatility and one latent heat stand for the thermodynamics, and the o-xylene volatility of 0.44 is 19 to 42 % above the measured vapor-pressure data, so the toluene and o-xylene split here is easier than on a real column and the valley's position would move on a rigorous model. The stages are ideal, the pressure is uniform, and the vapor split is fixed at 0.5, while a real wall's vapor split moves with the liquid loads. Flooding is a ceiling on the boilup, and the column has no hydraulic limit of its own.

The four-point loop's setpoint is the value at the design minimum, which commissioning would have to find by stepping the split. The feedforward's gains are the flows per unit of feed at the specifications rather than at the setpoints, which moves the challenge's lowest toluene by 0.006 mol %. The MPC has a fixed model identified at one point and an output-bias disturbance model, and an industrial MPC would be identified on the plant and would often sit above temperature loops. When the reflux stops, the side draw is capped at the liquid on its stage while the side flow meter still reports the request, a state no shipped case reaches. The disturbances arrive as steps and ramps. A feed that turns heavier in one step would not happen on a real column, but it separates the valley's movement from the loops' response.

OPEN CHALLENGE **Holding the Valley With a Fixed Vapor Split**

The window of minimum energy is narrow and moves with the feed composition and the feed's vapor fraction [18]. With the vapor split fixed, as it normally is, the two ends of the prefractionator cannot both be controlled, and for some disturbances the column must use more than its minimum energy. Industry prefers decentralized PI loops with selectors where they work, but tuning them can be difficult, and a linear predictive controller would need models from several operating regions, so nonlinear predictive control on a physical model may be the better route [24]. Progress would be a vapor split that can be moved reliably on a full-scale column, with controllers that hold the minimum across feeds.

8.9 Conclusion

This chapter modeled a benzene, toluene and o-xylene dividing-wall column on four coupled stage stacks whose balances close to 8.7×10^{-10} of the feed, with the energy valley computed from the same equations and drawn under the running plant. With the three purities held, the liquid split

moves only the duty, from 5.83 MW at 0.314 to 17.7 % more at 0.40 and 21.4 % more at 0.28, and an independent solver puts the minimum at 0.3139 and 649.26 kmol/h. The Underwood minimum with the vapor split fixed at 0.5 is 391.8 kmol/h, and the column needs 1.66 times that because its 14 stages per split sit close to the Fenske minimum. The steep side is the prefractionator reaching its ceiling on liquid-to-vapor ratio. A heavier feed moves the minimum to 0.339 and leaves a fixed split 6.6 % above it, while the split loop on the prefractionator's o-xylene ends 0.1 % above. With the feedforward the four-point structure passes the challenge at 32 of 32 seeds, and the MPC fails it with the benzene at 98.29 mol %. The volatilities are teaching values and the hydraulics a ceiling. A fuller model would add rigorous thermodynamics, a vapor split that follows the pressure drops, temperature-based split control with a differential measurement, and an MPC whose model includes an input disturbance. The app shows that a column holding every specification can still be burning 17.7 % more steam than it needs.

KEY POINTS

- A dividing-wall column makes three products in one shell, and with the purities held its fourth handle, the liquid split, moves only the energy.
- At the design feed the duty is 5.83 MW at a split of 0.314, 17.7 % more at 0.40 and 21.4 % more at 0.28. The steep side is the prefractionator's ceiling on liquid-to-vapor ratio.
- A heavier feed moves the minimum to 0.339 and leaves a fixed split 6.6 % above it, while the split loop on the prefractionator ends 0.1 % above.
- With the feedforward the four-point structure passes the challenge at 32 of 32 seeds, and the MPC fails it.

Where else this applies. Volume 9, Chapter 8 meets an optimum that moves with a disturbance, and Chapter 11 shows where an optimizer stops when its model is not exact.

References

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

[2] 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.

[3] S. Skogestad, M. Morari, and J. C. Doyle, "Robust control of ill-conditioned plants: high-purity distillation," *IEEE Transactions on Automatic Control*, vol. 33, no. 12, pp. 1092–1105, 1988.

[4] H. Ling and W. L. Luyben, "New control structure for divided-wall columns," *Industrial & Engineering Chemistry Research*, vol. 48, no. 13, pp. 6034–6049, 2009.

[5] ——, "Temperature control of the BTX divided-wall column," *Industrial & Engineering Chemistry Research*, vol. 49, no. 1, pp. 189–203, 2010.

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

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

[8] E. R. Gilliland, "Multicomponent rectification estimation of the number of theoretical plates as a function of the reflux ratio," *Industrial & Engineering Chemistry*, vol. 32, no. 9, pp. 1220–1223, 1940.

[9] I. J. Halvorsen and S. Skogestad, "Minimum energy consumption in multicomponent distillation. 2. three-product Petlyuk arrangements," *Industrial & Engineering Chemistry Research*, vol. 42, no. 3, pp. 605–615, 2003.

[10] E. A. Wolff and S. Skogestad, "Operation of integrated three-product (Petlyuk) distillation columns," *Industrial & Engineering Chemistry Research*, vol. 34, no. 6, pp. 2094–2103, 1995.

[11] G. Kaibel, "Distillation columns with vertical partitions," *Chemical Engineering & Technology*, vol. 10, no. 1, pp. 92–98, 1987.

[12] M. I. A. Mutalib and R. Smith, "Operation and control of dividing wall distillation columns," *Chemical Engineering Research and Design*, vol. 76, no. 3, pp. 308–318, 1998.

[13] I. Dejanović, L. Matijašević, and Ž. Olujić, "Dividing wall column—a breakthrough towards sustainable distilling," *Chemical Engineering and Processing: Process Intensification*, vol. 49, no. 6, pp. 559–580, 2010.

[14] Ö. Yildirim, A. A. Kiss, and E. Y. Kenig, "Dividing wall columns in chemical process industry: A review on current activities," *Separation and Purification Technology*, vol. 80, no. 3, pp. 403–417, 2011.

[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] I. J. Halvorsen and S. Skogestad, "Optimal operation of Petlyuk distillation: steady-state behavior," *Journal of Process Control*, vol. 9, no. 5, pp. 407–424, 1999.

[17] C. B. Willingham, W. J. Taylor, J. M. Pignocco, and F. D. Rossini, "Vapor pressures and boiling points of some paraffin, alkylcyclopentane, alkylcyclohexane, and alkylbenzene hydrocarbons," *Journal of Research of the National Bureau of Standards*, vol. 35, no. 3, p. 219, 1945.

[18] D. Dwivedi, J. P. Strandberg, I. J. Halvorsen, H. A. Preisig, and S. Skogestad, "Active vapor split control for dividing-wall columns," *Industrial & Engineering Chemistry Research*, vol. 51, no. 46, pp. 15 176–15 183, 2012.

[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] 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.

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

[23] 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] D. Dwivedi, I. J. Halvorsen, and S. Skogestad, "Control structure selection for three-product Petlyuk (dividing-wall) column," *Chemical Engineering and Processing: Process Intensification*, vol. 64, pp. 57–67, 2013.

8.10 Exercises

8.1 ○ Concept check · 2 min

In scenario 1 the three purity loops hold 99.2, 98.2 and 99.2 mol % while the liquid split is moved from 0.40 to 0.34. The distillate, side draw and bottoms stay near 100 kmol/h each, but the reboiler duty falls from 6.86 to 5.98 MW. Why can the split change the duty but not the product flows?

(a) The level loops hold the distillate and bottoms at fixed flows, so only the boilup is left free

(b) The vapor split at the wall follows the liquid split, so each side of the wall keeps its own flows

(c) With the feed and three purities fixed, the balances fix the products, so the split only moves internal flows

(d) The duty belongs to the bottoms loop alone, and the liquid split does not reach the boilup

8.2 ○ Concept check · 1 min

Below the valley's minimum the duty climbs steeply, and at a split of 0.25 no boilup up to flooding makes the side product. What makes that side steep?

(a) The main side runs short of liquid, so benzene reaches the side draw

(b) The boilup reaches flooding, so the bottoms loop loses its handle

(c) The side loop closes the draw, so the toluene backs up into the bottoms

(d) The prefractionator's ratio L/V cannot exceed β/ρ, so o-xylene climbs over the wall

8.3 Concept check · 2 min

In scenario 7 the purity loops read tray temperatures converted by straight lines fitted at the design feed. At 3 h, two hours after the feed turns heavier, the loops read 99.20, 98.17 and 99.21 mol % while the true purities are 99.36, 98.01 and 99.02 mol %. What causes the gap?

(a) A tray temperature depends on the whole composition, and the calibration fitted at the design feed misreads the heavier one

(b) The thermowells lag the trays, so the loops still read the compositions from before the feed change

(c) The analyzer bias update has reached its clamp and pushes each reading off by the same amount

(d) The loops run without integral action, so each settles with a proportional offset

8.4 Theory and calculation · 8 min

Use Equation (8.6) with $\alpha = (2.4,\ 1,\ 0.44)$, $z_j = 1/3$ each, $F = 300\,\mathrm{kmol/h}$ and a saturated-liquid feed ($q = 1$).

(a) Check that $\theta_1 = 1.4667$ and $\theta_2 = 0.5625$ are roots, and compute the minimum vapor of the benzene and toluene split (θ_1, benzene in the distillate) and of the toluene and o-xylene split (θ_2, benzene and toluene in the distillate). Give four significant figures. *Checkpoint:* 257.1 kmol/h and 359.2 kmol/h.

(b) In the direct sequence the second column receives 200 kmol/h of equimolar toluene and o-xylene. Find its root and minimum vapor, the direct sequence's total, and the saving of the thermally coupled column, whose minimum is the larger of the two values in part (a). *Checkpoint:* $\theta = 0.6111$, 257.1 kmol/h; 514.3 kmol/h; 30.2 %.

(c) The wall fixes $\rho = 0.5$, and the prefractionator reaches its preferred split only with 195.9 kmol/h of vapor. Compute the fixed-wall minimum boilup, its duty at 0.008 984 MW per kmol/h, and the ratio of the app's 649.3 kmol/h to it. *Checkpoint:* 391.8 kmol/h, 3.52 MW; 1.66.

8.5 Theory and calculation · 6 min

The latent heat is $\lambda = R\Lambda$ with $\Lambda = 3890\,\mathrm{K}$ and $R = 8.314 \times 10^{-3}\,\mathrm{kJ/(mol\,K)}$. On the design valley the boilup is 649.3 kmol/h at the minimum ($\beta^* = 0.314$), 763.9 kmol/h at 0.40 and 788.5 kmol/h at 0.28, and $\rho = 0.5$.

(a) Compute λ, the duty of 1 kmol/h of boilup in MW, the duty at the three splits, and how far the two outer points lie above the minimum. *Checkpoint:* 32.34 kJ/mol, 0.008 984 MW; 5.83, 6.86 and 7.08 MW; 17.7 % and 21.4 %.

(b) Compute the prefractionator's ceiling β/ρ at splits of 0.314, 0.28 and 0.25. It needs a ratio of 0.53 above the feed at the minimum and 0.49 at 0.28. Say what the ceiling at 0.25 means for scenario 8. *Checkpoint:* 0.628, 0.56, 0.50.

(c) The main side below the draw needs a ratio of 0.65 to 0.85. At what split does its ceiling $(1-\beta)/(1-\rho)$ fall to the lower value? *Checkpoint:* 0.675.

8.6 Theory and calculation · 6 min

The feed-rate feedforward adds $K(F_m - F_{ref})$ to the reflux, side draw and boilup. Its gains are the flows per unit of feed at the specifications and a split of 0.305, where the valley gives $L = 500.6\,\mathrm{kmol/h}$, $S = 100.0\,\mathrm{kmol/h}$ and $V = 601.1\,\mathrm{kmol/h}$ at $F = 300\,\mathrm{kmol/h}$. At the setpoints the minimum is $V = 649.3\,\mathrm{kmol/h}$ with $L = 548.9\,\mathrm{kmol/h}$.

(a) Compute the three gains. *Checkpoint:* 1.669, 0.333, 2.004.

(b) The feed rises from 300 to 360 kmol/h. Compute what the feedforward adds to each flow, and the duty its boilup adds. *Checkpoint:* 100.1, 20.0 and 120.2 kmol/h; 1.08 MW.

(c) With constant molar overflow the steady state at a fixed split and fixed purities scales with the feed. Compute the boilup and reflux the setpoints need at 360 kmol/h, and how much the purity loops must add to what the feedforward supplies. *Checkpoint:* 779.2 and 658.7 kmol/h, about 9.6 kmol/h each (9.7 from the rounded values).

8.7 App challenge · 8 min

Open scenario 1, *The fourth handle: move the liquid split*, at the default seed and the default speed of 60 min per second. Read the *Reboiler duty* and *Duty against the valley* tiles of the scorecard.

(a) Run as shipped (split 0.40) to 1 h and read both tiles. *Checkpoint:* 6.86 MW, +17.5 %.

(b) Reload the scenario, type 0.34 into the *Liquid split* box before Run and run to 4 h. Then reload, type 0.28 before Run and run to 4 h. Read both tiles each time, to within 0.02 MW and 0.2 points. *Checkpoint:* 5.98 MW, +2.5 %; 6.86 MW, +17.5 %.

(c) Run the 0.28 case on to 12 h and read the tiles again. The valley curve gives 7.08 MW at 0.28 and 6.86 MW at 0.40. Explain why 0.28 and 0.40 read the same at 4 h and what the run at 0.28 shows about the two sides of the valley. *Checkpoint:* 7.02 MW, +20.3 %.

8.8 App challenge · 8 min

Run scenario 4, *Three-point structure: a heavier feed*, and then scenario 5, *Four-point structure: a loop on the prefractionator*, each as shipped at the default seed to 8 h. The feed turns heavier at 1 h in both.

(a) For each run read the *Reboiler duty* and *Duty against the valley* tiles, and on the Liquid split card the split and the valley's minimum. *Checkpoint:* Scenario 4: 6.40 MW, +6.6 %, 0.314 and 0.339. Scenario 5: 6.01 MW, +0.1 %, 0.345 and 0.339.

(b) Export the CSV file of each run and read column `truth_xPX_pct`, the true o-xylene on the prefractionator's top stage, in the row at `Time (h)` 8, to within 0.02 mol %. *Checkpoint:* 3.22 mol % and 1.59 mol %.

(c) The split loop's setpoint, 1.57 mol %, is the value at the design feed's minimum. At the heavier feed's minimum, 0.339, the steady-state model puts x_P at 1.80 mol %. Explain where the loop settles and what that costs. *Checkpoint:* A little above 0.339, on the gentle side; about 0.1 % of duty.

8.9 App challenge · 8 min

Open scenario 9, the challenge *twelve hours in the valley*, at the default seed. Each run stops at 12 h, and the verdict gives the lowest sample of each true purity and the worst judged hour of duty against the valley's minimum.

(a) Run as shipped (three-point, split 0.40, feedforward off). Read the verdict and say which half of the target fails. *Checkpoint:* 98.63 / 98.15 / 98.13 mol %, +12.4 %; both halves fail.

(b) Reload, set *Control structure* to Four-point and *Feed-rate feedforward* to On before Run, and run again. *Checkpoint:* 99.11 / 98.07 / 99.07 mol %, +0.3 %; passes.

(c) Reload, keep the three-point structure, set the feedforward On and type the new minimum, 0.339, into the *Liquid split* box before Run. Say what the box shows, read the verdict, and explain why the setting of part (b) is still the one to keep. *Checkpoint:* 0.340; 99.11 / 98.08 / 99.07 mol %, +0.0 %; passes.

8.10 Application area · 3 min

The model's o-xylene volatility relative to toluene is 0.44, while vapor-pressure data give 0.31 to 0.37 between 70 and 110 °C.

(a) Does a real benzene, toluene and o-xylene column find the toluene and o-xylene split easier or harder than the model, and what does that ask of its stages or its boilup? *Checkpoint:* Harder; more stages or more boilup for the same purities.

(b) What would a plant team do before trusting the model's minimum or the four-point setpoint of 1.57 mol %? *Checkpoint:* Find both on the plant by stepping the split.

8.11 Application area · 4 min

At the design feed the setpoints sit 0.2 mol % above the specifications of 99, 98 and 99 mol %, and the valley's minimum is 649.3 kmol/h at the setpoints against 601.1 kmol/h at the specifications.

(a) Compute the cost of the margins as a share of the duty and in MW. *Checkpoint:* 8.0 %, 0.433 MW.

(b) In scenario 6 the steam supports 740 kmol/h of boilup and the feed rises to 360 kmol/h, where the minimum at the setpoints is 779 kmol/h. Lowering the setpoints to 99.05, 98.05 and 99.05 mol % lowers it to 734 kmol/h. What does the operator trade, and what must stay true for the trade to be safe? *Checkpoint:* Margin for throughput; the disturbances must stay inside the smaller margin.

8.12 Review the agent · 5 min

An AI assistant reviews the four-point structure and writes: *"The prefractionator analyzer reports only every 10 min and is costly to maintain. Set* Split loop measurement *to the prefractionator tray temperature at its shipped setpoint, 111.7 °C, the value near the design minimum. A tray temperature answers in a minute and follows the composition, so the split loop will find the valley after a feed change just as well."* Find the flaw, and say how you would check the claim in the app.

CHAPTER 9

MTBE Reactive Distillation With Multiple Steady States and a Branch-Dependent Temperature Loop

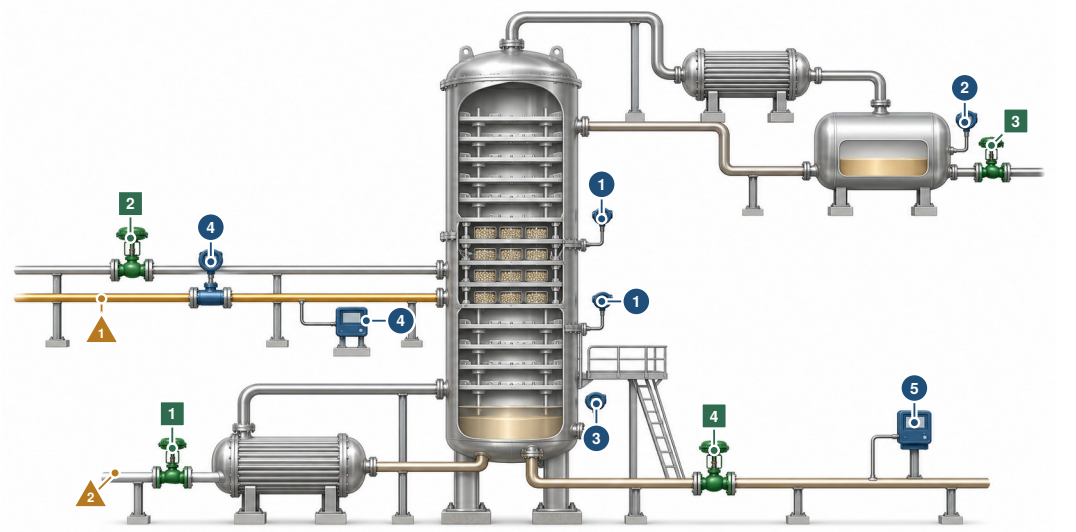

The controller moves 1 reboiler steam valve, which sets the duty 2 methanol valve, in ratio to the C4 3 distillate valve, on the drum level 4 bottoms valve, on the base level

The controller reads 1 tray temperatures: catalytic tray 8 for the branch monitor, tray 12 for the loop 2 reflux drum level 3 base level 4 C4 flow meter and feed analyzer 5 bottoms analyzer

What disturbs it 1 a larger or leaner C4 feed 2 a dip in the steam supply

Notice the band of catalyst trays where the methanol must stay: with the same feeds and the same steam, the column can hold it there or lose it below. Illustration, not to scale.

IN PLAIN TERMS **One Column, Three Ways to Settle**

The main idea. This tall column makes a chemical called an ether. Two feeds react on trays that hold a catalyst, and the column separates the product as it forms. With the same feeds and the same heat, the column can settle in three very different ways, and which one it reaches depends on how it got there.

The challenge. The controller adds or removes heat by watching one tray low in the column. In one bad way to settle, that tray's temperature stops rising, so the controller keeps adding heat and makes things worse.

Why it matters. A loop that works well in one way can quietly push the column into another.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page, the history boxes of Sections 9.1 and 9.2, the fact boxes of Sections 9.2 and 9.6, the two try-it boxes of Section 9.6 and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises
Practice	15 min	the first page, Section 9.6, Section 9.7 with its three boxes, Section 9.8 and the open challenge

The control question At one set of feeds and one reboiler duty, an MTBE column can convert most of its isobutene, lose its methanol into the stripping trays, or fail to strip its C4. Which steady state does the column reach, and what does a stripping-tray temperature loop tuned on one of them do on the others?

Learning objectives

- Explain why the methanol's position decides the branch, and read the branches, folds and Hopf point of an S-curve against the duty.
- Compute the boilup, the vapor made by the heat of reaction and the temperature dependence of the equilibrium constant and the rate.
- Predict the hysteresis of a slow duty ramp from the S-curve and measure it on the running column.
- Identify an input multiplicity in a tray temperature and decide which setpoints a loop can hold on each branch.
- Evaluate a branch monitor and a duty-cutting supervisor, and choose a control tray and setpoint with margin on both branches.

Prerequisites Sections 1.4.5 (multiplicity in a reactive column), 1.4.2 (equilibrium-limited reactions), 1.4.6 (ratio control), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios and the challenge). Chapter 8 uses the same stage stack.

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

See also Volume 7, Chapter 2 for folds and a Hopf point in a stirred tank, and Volume 7, Chapter 4 for a setpoint at the peak of the plant characteristic. Volume 5, Chapter 3 for composition control through a slow analyzer. Volume 6, Chapter 7 for the same stage stack in a recycle plant.

9.1 Introduction

Fifteen trays, eight of them carrying an acid resin, turn 775.8 kmol/h of methanol and 1976.4 kmol/h of a C4 cut holding 35.68 mol % isobutene into ether in the simulated column. At a reboiler duty of 42.5 MW the column can sit in three stable states. In one it converts 98.22 % of the isobutene and sends 99.94 mol % MTBE to the bottoms. In another the methanol has fallen through the catalyst into the stripping trays, only 90.05 % is converted, and the bottoms carry 16.91 mol % methanol. In the third the vapor is too weak to strip the C4, 94.89 % is converted and the bottoms hold 6.94 mol % C4. The feeds, the reflux ratio of 7 and the duty are the same in all three.

The simulator is built to correct a belief that most control courses leave in place, that a plant's handles fix its steady state, so that a controller tuned at one operating point meets the same plant wherever the handles return to it. Distillation alone can contradict that belief, since ideal two-product columns can have several steady states [1], and a reactor example shows that more than one set of inputs can give the same outputs, which decides whether a loop pairing is stable [2]. Reactive distillation combines both, and its multiplicity and oscillations are reviewed as part of its control problem [3]. In the column simulated here a duty ramped up at 1 MW per hour drops the conversion at 44.82 MW, and ramped down it recovers only at 40.88 MW. A tray-temperature loop that holds the high-conversion state through a feed step raises the duty without limit on the low-conversion state, and one tray temperature is reached at two duties.

> ⧖ **HISTORY** **1960 to 1991: Can a Plain Column Have Two States?**
>
> Multiple steady states in distillation were studied for three decades before 1991. Rosenbrock proved in 1960 and 1962 that a binary column with constant relative volatility and constant molar flows cannot have them, and a 1982 review concluded the same for any multistage column separating a binary mixture with constant molar flows. In 1991 a study showed, for the first time, multiplicity in columns with ideal thermodynamics. One kind came from reflux and boilup set in mass or volume units, as industrial columns usually set them, the other from the energy balance. In both, more reflux made the top of the column separate worse, at an unstable operating point [1]. The MTBE column adds a reaction to the trays and finds three stable states.

However, the literature presents these effects as steady-state analyses and simulation studies [3–5], and a student rarely meets a column that can be pushed across a fold and watched while its branches are drawn underneath. This chapter describes an MTBE column simulated in the browser, whose branches an independent steady-state solver computes and draws under the running plant, and measures what a duty ramp, a feed change, methanol excess, two start-up charges and a stripping-tray temperature loop do to the branch the column is on. The equilibrium-stage model, the Wilson equation [6], the activity-based kinetics [7, 8], continuation and the proportional-integral (PI) law are all standard. The measurements are made on one column at one seed. They show that the multiplicity emerges from parameters read from published sources, that the plant lines out within 0.007 percentage points of the solver's states in the comparisons of Table 9.4, and that the high branch is lost through a complex pair of eigenvalues at 43.50 MW rather than at a fold. The model is a 79-state column on the series' shared distillation module, with an integrator whose fixed points are exact steady states, and the chapter gives the equations, constants and defaults from which the column can be rebuilt. Nine cases come with completion rules that credit what a learner does.

Section 9.2 traces etherification on a distillation tray and the history of its multiplicity. Section 9.3 sets out the column as modeled, its trays, activity model, rate, instruments and integrator. Section 9.4 computes the branches with an independent solver, classifies their stability and explains why the methanol decides the branch and why one tray temperature is read at two duties. Section 9.5 describes the loops, the branch monitor and the supervisor. Section 9.6 measures the nine cases, from the hysteresis loop of a duty ramp to the graded challenge, and Section 9.7 names the real devices behind the loops. Section 9.8 states the limitations and Section 9.9 sums up. The exercises at the end of the chapter run the cases of Section 9.6.

9.2 Etherification on a Distillation Tray

An ether made from isobutene and methanol over an ion-exchange resin is limited by equilibrium, and in a fixed bed the conversion stops short of complete. Boiling the heavy ether away from the catalyst as it forms removes that limit, and etherification became one of the reaction classes of industrial reactive distillation. How the technique grew to more than 150 industrial columns is told in the history box below. The kinetics behind such designs were written in liquid-phase activities on the resin [7], and rate constants of that form are still used in reactor models [8].

> ⧖ HISTORY **From the 1920s to More Than 150 Columns**
>
> Reaction and separation in one column have been developed since the early 1920s, and the technique counts among the best success stories of process intensification. It overcomes the equilibrium limit of reactions such as esterification and etherification [9]. A 2007 survey reported more than 150 industrial reactive distillations operating worldwide at 100 to 3000 kilotonnes per year, most of them started up less than fifteen years before. Their reactions include hydrogenations, hydrodesulfurization, esterifications and etherification, and several commercial ways of packing a solid catalyst into a column had become available. The designs were scaled up from pilot plants by brute force and by modeling [10].

The column's second life began as a puzzle. Multiple solutions of an MTBE column were reported in 1993 [11], and a 1995 study asked why MTBE production by reactive distillation may yield them [12]. The analysis of columns at infinite length and infinite internal flows then predicted output multiplicities from the reactive vapor-liquid equilibrium alone, extended it to columns with a reactive core between nonreactive sections, and studied the physical causes of the MTBE multiplicities [4]. Theoretical analysis and experimental verification of steady-state multiplicities in MTBE and TAME columns appeared the same year [5], as did a nonequilibrium, rate-based model that showed multiple steady states in MTBE synthesis [13]. A 1996 dynamic simulation of a seventeen-stage column found the reported low- and high-conversion states together with unstable states between them, and its base

Table 9.1 Variables, Their Roles and Design Values (Case 3 at $t = 0$)

Symbol	Role	Design	Range
T_{12}	CV, control tray temperature (°C)	137.4	SP 70–155
x_B	CV (trim), bottoms MTBE (mol %)	99.94	SP 95–99.9
ℓ_D, ℓ_B	CV, drum and base levels (%)	50, 50	SP 20–80
Q	MV, reboiler duty (MW)	42.5	0–60
F_M	MV, methanol (kmol/h), $r_M F_{C4} z_{IB}$	775	0–1600
D, B	MV, distillate, bottoms (kmol/h)	1364, 694	0–3000, 0–2000
L	reflux, RD (kmol/h)	9546	—
r_M, R	settings, methanol ratio, reflux ratio	1.100, 7	0.9–1.3, 5–10
F_{C4}	DV, C4 feed (kmol/h)	1976	1600–2400
z_{IB}	DV, isobutene in the C4 (mol %)	35.68	28–42
s	DV, steam available (% of normal)	100	60–110
M_i, x_{ij}	states, holdups (kmol), fractions	15, profile	15 trays
M_D, M_B	states, drum, sump (kmol)	100, 150	cap 200, 300
X	output, isobutene converted (%)	98.22	—
$x_{D,M}$	output, methanol in distillate (mol %)	6.08	—
T_8	monitor tray temperature (°C)	77.6	saddle 93

case of feeds, reflux ratio and reactive stages [14] is the one later work, and this application, start from [15].

Control followed the nonlinear analysis. Two-point control of a reactive column for composition and conversion was studied in 1999 [16], a review chapter set the nonlinear dynamics of reactive distillation, including MTBE and TAME synthesis, beside control structure selection and control algorithms [3], and a monograph on reactive distillation gave a chapter to the control of MTBE and ETBE columns [17]. The application takes the simplest regulatory choice, one stripping-tray temperature on the reboiler duty, and shows what the branches do to it.

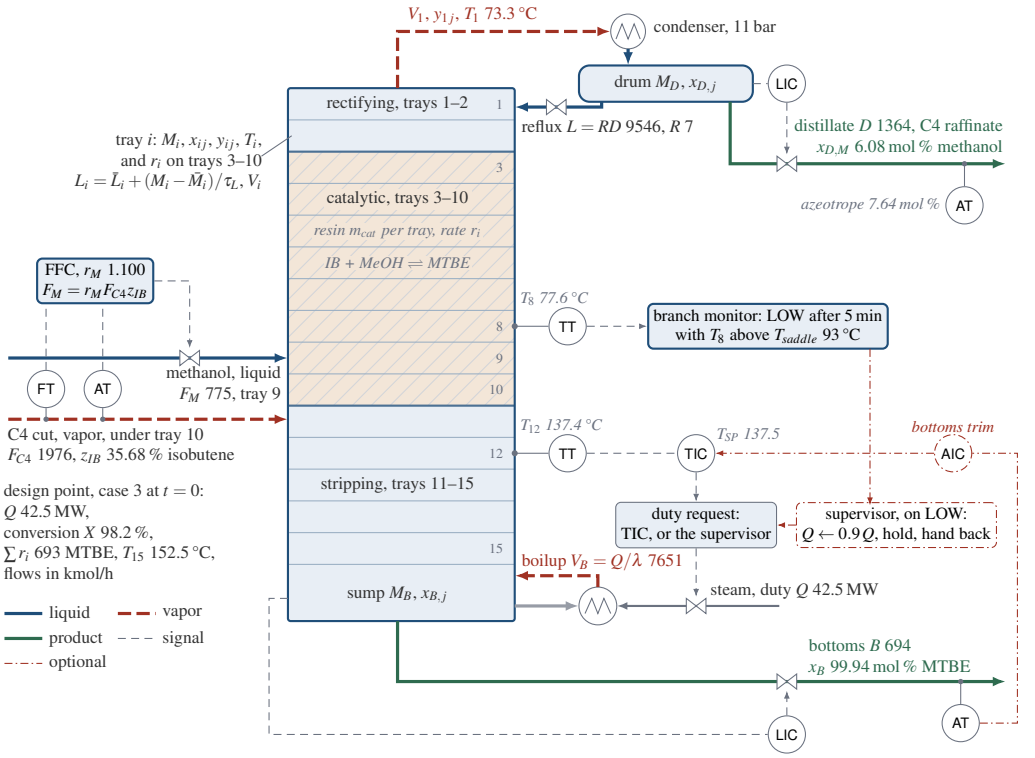

Figure 9.1 The MTBE column drawn after the Column view of case 3 at $t = 0$, the design point on the high branch at 42.5 MW. Methanol enters tray 9 as liquid and the C4 cut enters under tray 10 as vapor, the catalytic trays 3 to 10 are hatched, and the tray callout lists the states and flows of Equations (9.1) to (9.3). The temperature loop reads tray 12 and sets the reboiler duty, the methanol follows the measured C4 flow and the feed analyzer's isobutene, the two level loops set the distillate and bottoms, and the dash-dotted red elements (the bottoms trim, the monitor's supervisor and its duty override) are optional. Liquid and product streams are solid, vapor streams dashed and signals thin dashed. Flows in kmol/h. Process model only. Not to scale.

> ✦ **DID YOU KNOW? A Fifth or More Saved**
>
> In every industrial case of the 2007 survey, reactive distillation cut the variable cost, the capital and the energy by 20 percent or more against a reactor followed by a column [10]. A 2018 comparison put the energy saving of the classic reactive column at 15 to 80 percent [9].

9.3 The Column as Modeled

The simulated column has no plant behind it, but its shape and feeds are a published base case. A doctoral thesis on reactive distillation dynamics tabulates the base case of Schrans and coworkers [14], a total condenser, two rectifying, eight reactive and five stripping trays and a partial reboiler, with 215.5 mol/s of liquid methanol, 549 mol/s of C4 vapor at 35.68 % isobutene on stage 11, reactive stages 4 to 11, a reflux ratio of 7 and feeds at 11 atm, and its multiplicity figures put the methanol on stage 10 [15]. The application keeps that shape, numbers the trays 1 to 15 below the condenser, so that methanol enters tray 9 and the C4 vapor enters under tray 10, and holds every tray at 11 bar. The layout is shown in Figure 9.1.

Table 9.1 lists the variables by role. Flows are in kmol/h on an engine clock in minutes, so the code divides each balance by 60, a factor the equations below omit. The model integrates 79 states,

a holdup and three mole fractions on each of 17 stages (drum, 15 trays, sump), five flow lags, and six ledgers that count the inert, the isobutene group (isobutene plus MTBE) and the methanol group (methanol plus MTBE) in and out, the three quantities the reaction conserves.

9.3.1 Trays With a Reaction

Each tray holds liquid in phase equilibrium with the vapor it sends up. On a catalytic tray the reaction adds its stoichiometry to the component balances, as shown in Equation (9.1).

$$\begin{aligned} \frac{\mathrm{d}M_i}{\mathrm{d}t} &= L_{i-1} + V_{i+1} + F_i - L_i - V_i - r_i \\ \frac{\mathrm{d}(M_i x_{ij})}{\mathrm{d}t} &= L_{i-1}x_{i-1,j} + V_{i+1}y_{i+1,j} + F_i z_{ij} \\ &\quad - L_i x_{ij} - V_i y_{ij} + \nu_j r_i \end{aligned} \tag{9.1}$$

where j runs over isobutene, n-butene (the inert), methanol and MTBE, the last closing the fractions to one, $\nu = (-1, 0, -1, +1)$, r_i is the net rate of MTBE formation on tray i (zero off trays 3 to 10), and F_i is a feed on tray 9 or 10. One mole of liquid is lost per mole of ether formed. Constant molar overflow with one latent heat λ moves the flows, and the heat of reaction boils liquid on its own tray, as shown in Equation (9.2).

$$\begin{aligned} V_i &= V_{i+1} + (1 - q_i)F_i + \nu_V r_i, \qquad V_{16} = V_B = Q/\lambda \\ L_i &= L_{i-1} + q_i F_i - (1 + \nu_V) r_i, \qquad \nu_V = -\Delta H_r/\lambda \end{aligned} \tag{9.2}$$

where $q = 1$ for the liquid methanol and $q = 0$ for the C4 vapor. With $\Delta H_r = -37.7\,\mathrm{kJ/mol}$ [8] and $\lambda = 20\,\mathrm{kJ/mol}$, each mole of MTBE boils $\nu_V = 1.885$ moles, and the boilup is 180 kmol/h per MW. In the dynamic model these two equations are steady values about which a linearized weir moves the liquid, as shown in Equation (9.3).

$$L_i = \bar{L}_i + \frac{M_i - \bar{M}_i}{\tau_L} \tag{9.3}$$

where $\bar{L}_i$ and $\bar{M}_i$ are the flows and holdups of a reference set once from the base case's reflux and feeds and $\tau_L = 0.1\,\mathrm{min}$. The trays, the drum and the sump come from the stage stack the series shares, first written for its reactor-column-recycle plant and promoted to the framework once its dividing-wall column needed the same code (Volume 6, Chapter 7 and Chapter 8). This column is the stack's third user, and it plugs in through the stack's own interfaces. The activity model is the stack's bubble point. The reaction enters as a vapor feed of the tray's liquid with an equal liquid draw for the heat of reaction, a liquid draw of r_i for the mole lost, and one stoichiometric term added to the derivatives. The flows run the other way when MTBE decomposes.

9.3.2 Liquid Activity and the Bubble Point

Every tray's temperature and vapor follow from the bubble point with Wilson activity coefficients [6], as shown in Equation (9.4).

$$\begin{aligned} P &= \sum_j \gamma_j x_j P_j^{sat}(T), \qquad y_j = \gamma_j x_j P_j^{sat}/P \\ \ln \gamma_i &= 1 - \ln \sum_j x_j \Lambda_{ij} - \sum_k \frac{x_k \Lambda_{ki}}{\sum_j x_j \Lambda_{kj}} \\ \Lambda_{ij} &= \frac{V_j}{V_i}\, \mathrm{e}^{-a_{ij}/T} \end{aligned} \tag{9.4}$$

where $\ln P^{sat} = C_1 + C_2/T + C_3 \ln T + C_4 T^{C_5}$ in pascals (the DIPPR equation 101 form), solved for T by Newton's method from the fraction-weighted boiling points. The vapor-pressure coefficients, the

Table 9.2 Property Constants Read From the Sources (Fixed in the Source)

	isobutene	methanol	MTBE
C_1	102.5	81.768	55.875
C_2 (K)	−5021.8	−6876	−5131.6
C_3	−13.88	−8.7078	−4.9604
C_4	0.020296	7.1926×10^{-6}	1.9123×10^{-17}
C_5	1	2	6
V_j (cm^3/mol)	93.33	44.44	118.8
$a_{IB,j}$ (K)	0	49.4802	8.87958
$a_{M,j}$ (K)	1243.186	0	711.7
$a_{E,j}$ (K)	−0.01494	−197.9	0

rate: $a_0 = 6.3\times10^{12}$ mol h^{-1} g^{-1}, $E_a = 85.4$ kJ/mol, $\Delta H_r = -37.7$ kJ/mol, equilibrium $K(350\,\text{K}) = 28.5$, van 't Hoff with ΔH_r ($K = 271.2$ at 25 °C)

Wilson energies a_{ij} and the molar volumes V_j of isobutene, methanol and MTBE are those tabulated in a 2021 thesis [18] and appear in Table 9.2. The inert takes isobutene's properties.

The thesis prints the coefficient with V_i/V_j. The standard form [6] with V_j/V_i is used because only that reading of the tabulated energies gives the azeotropes the sources describe. With it the infinite-dilution coefficients at 350 K are 29.9 for methanol in isobutene, 4.02 for methanol in MTBE and 2.91 for MTBE in methanol, methanol and MTBE form a minimum-boiling azeotrope at 1 atm at 29.4 mol % methanol and 50.8 °C, and the C4 and methanol form an azeotrope at 11 bar at 7.64 mol % methanol and 73.1 °C. The printed reading gives 0.50 and 0.55 for the methanol and MTBE pair, a negative deviation with no azeotrope. At 11 bar the pure components boil at 75.2, 140.5 and 152.8 °C.

9.3.3 The Rate on the Resin

The net rate on a catalytic tray is written in activities $a = \gamma x$, as shown in Equation (9.5).

$$\begin{aligned} r_i &= m_{cat}\, k(T_i)\, \frac{a_{IB}\, a_M - a_E/K(T_i)}{(a_M+\varepsilon)^2} \\ k(T) &= a_0\, \mathrm{e}^{-E_a/RT} \\ K(T) &= K_{350}\, \exp\!\left[\frac{-\Delta H_r}{R}\left(\frac{1}{T}-\frac{1}{350\,\text{K}}\right)\right] \end{aligned} \tag{9.5}$$

where $m_{cat} = 1000$ kg of resin per tray, E is MTBE and M methanol. As $\varepsilon \to 0$ the fraction becomes the methanol-rich form $a_{IB}/a_M - a_E/(K a_M^2)$ with the constants a_0 and E_a of Table 9.2 [8]. $K_{350} = 28.5$ is the activity-based equilibrium constant of the 1990 kinetics [7] as tabulated in the thesis that supplies the column [15], and K falls as T rises because the reaction is exothermic.

The term $\varepsilon = 0.05$ keeps the rate finite where the methanol runs out, but it is not a small guard. Written out, the denominator is the adsorption form $(1 + K_{ads} a_M)^2$ with $K_{ads} = 1/\varepsilon = 20$ and the rate constant scaled by $1/\varepsilon^2$. At the design point it slows the rate on the catalytic trays to 0.48 to 0.81 of the $\varepsilon \to 0$ form, because the methanol activity there is only 0.11 to 0.44. The folds depend on it. At $\varepsilon = 0.02$ all six lie 0.9 to 3.0 MW lower (41.26, 38.34, 39.40, 37.23, 43.19 and 40.89 MW), while at $\varepsilon = 0.1$ only the under-boiled state remains at 42.5 MW (93.3 %). The catalyst mass, ε and λ are the three teaching values of the reaction and the flows. The folds survive each being changed but move with it. With 300 kg of resin per tray the curve folds at 48.10, 36.09, 53.46 and 52.80 MW, and with $\lambda = 16$ kJ/mol its six folds lie between 29.28 and 35.55 MW.

Table 9.3 Parameters With Their Status: A, Adjustable by the Learner. H, Hidden in the Parameter Panel. S, Set in the Source Code

Symbol	Meaning	Value	Unit	
m_{cat}	resin on each catalytic tray	1000	kg	A
Q_{max}	reboiler duty at full steam	60	MW	A
—	bottoms specification, conversion	98, 95	mol %, %	A
—	distillate methanol alarm	7.0	mol %	A
λ	latent heat (constant molar overflow)	20	kJ/mol	H
ε	methanol term, rate denominator	0.05	—	H
M_t	tray liquid at the weir reference	15	kmol	H
$M_{D,cap}$, $M_{B,cap}$	drum, base capacity	200, 300	kmol	H
τ_L, τ_Q	tray hydraulic lag, reboiler response	0.1, 2	min	H
P	pressure on every stage	11	bar	S
N	trays (rectifying, catalytic, stripping)	2, 8, 5	—	S
—	methanol, C4 feed trays	9, under 10	—	S
$\bar{L}$	weir reference reflux	9399.6	kmol/h	S
τ_F	feed, reflux and product flow lags	0.2	min	S

9.3.4 Drum, Sump, Actuators and Instruments

The drum holds what the total condenser returns, so its vapor in is y_1, and the sump boils at equilibrium as the partial reboiler, as written in Equation (9.6).

$$\begin{aligned} \frac{\mathrm{d}M_D}{\mathrm{d}t} &= V_1 - L - D \\ M_D \frac{\mathrm{d}x_{D,j}}{\mathrm{d}t} &= V_1\,(y_{1j} - x_{D,j}) \\ \frac{\mathrm{d}M_B}{\mathrm{d}t} &= L_{15} - V_B - B \\ \frac{\mathrm{d}(M_B x_{B,j})}{\mathrm{d}t} &= L_{15}\,x_{15,j} - V_B\,y_{B,j} - B\,x_{B,j} \end{aligned} \tag{9.6}$$

where $L = RD$. The delivered duty lags the request times the steam available with $\tau_Q = 2\,\mathrm{min}$, and the methanol, reflux, distillate and bottoms follow their setpoints with 0.2 min lags. A pump delivers only $\min(1, M/(0.02\,M_{cap}))$ of its flow as its vessel empties. The controllers see the plant only through a measurement packet, and a test in the application fails whenever the packet holds anything more. Thermowells on every tray lag 0.5 min with 0.05 K of noise, levels lag 0.1 min with 0.05 %, flows lag 0.05 min with 0.2 % of the flow, and the duty carries 0.02 MW. The series' analyzer module supplies three analyzers, one on the bottoms (MTBE, methanol, C4), one on the distillate's methanol and one on the feed's isobutene. Each samples every 10 min through a 1 min sample line and reports a cycle later with 0.02 mol % of noise. Each execution consumes a fixed count of random draws, which makes a seeded run repeat exactly.

9.3.5 Initialization and the Integrator

Table 9.3 gives the parameters. Each case opens lined out on the branch it names, at its own duty, feeds and ratios. The state comes from the steady-state solver of Section 9.4, started from a stored profile on that branch and walked to the case's settings, each tray holding its weir's steady liquid and the drum and base half full. The design point is a steady state of the dynamic model to 1.8×10^{-13} per minute.

The engine's classical fourth-order Runge-Kutta method (Section 1.6.1) at a 1 s step runs case 3 but fails one minute into the methanol-wetted start-up, where the reaction on a methanol-lean tray relaxes faster than the explicit step can follow. The application therefore supplies its own step. The transport is explicit, and on each catalytic tray the reaction is implicit through the rank-one Jacobian

of the rate, as shown in Equation (9.7).

$$\Delta\mathbf{x}_i = h\,\mathbf{f}_i + h^2\,\mathbf{u}_i\,\frac{\mathbf{g}_i^{\mathsf{T}}\mathbf{f}_i}{1 - h\,\mathbf{g}_i^{\mathsf{T}}\mathbf{u}_i}, \qquad u_{ij} = \frac{\nu_j + x_{ij}}{M_i} \tag{9.7}$$

where $\mathbf{f}_i$ is the tray's explicit derivative, $\mathbf{g}_i = \partial r_i/\partial \mathbf{x}_i$ at the tray's temperature and $h = 1\,\mathrm{s}$. This is the Sherman-Morrison solution of $(I - h\,\mathbf{u}\mathbf{g}^{\mathsf{T}})\Delta\mathbf{x} = h\mathbf{f}$, so $\Delta\mathbf{x} = 0$ exactly when $\mathbf{f} = 0$, and a fixed point of the step is a steady state of the equations. Where the linearized step would carry a fraction out of $[0, 1]$, the extent over the step is found by bisection on the interval that keeps every fraction inside. The stages are updated in holdup and component amount, so the inert and the two groups are conserved. Over the application's disturbed test run the ledgers close to 9.0×10^{-14} of what was fed, and one step from the solver's high, low and under-boiled states moves no state by more than 5×10^{-12}.

A step that would still carry a fraction more than 0.001 outside $[0, 1]$ is retaken as two half steps, split again while needed down to $h/32$. The first minutes of the methanol-wetted start-up need this, and there the lowest fraction is -7.1×10^{-4}, the methanol group closes to 3.6×10^{-3} kmol of 11 008 kmol fed over 4 h, and the column lands on the same state at $h/2$, $h/4$ and $h/10$. A step inside the tolerance is the single step, bit for bit. On the low branch at high duty (cases 4 and 8, and case 5 in Auto) fractions dip to at most -3.7×10^{-5}, inside the tolerance. Halving h leaves the jumps of case 1 at 44.816 and 40.884 MW.

The conversion is logged as the extent the reaction actually ran at over each step, summed over the catalytic trays and divided by the isobutene fed. Because the reaction relaxes within a fraction of a step, the rate re-evaluated at the stepped state jitters about that extent. In case 8 run at 143.5 °C, 300 min in, the applied extent gives 62.21 % at the 1 s step and 62.01 % at a quarter of it, while the re-evaluated rate reads 62.57 %.

9.3.6 Assumptions

The trays are equilibrium stages with an efficiency of one, and the reaction runs in the liquid of the catalytic trays with no resistance between vapor, liquid and resin, so the model is least reliable where the rate is fast, which is the case rate-based models address [13]. Constant molar overflow with one latent heat replaces an energy balance per tray, and the feeds are saturated. The inert is one component with isobutene's volatility, and no side reaction (dimethyl ether, isobutene dimers, tert-butyl alcohol) is modeled. The base case's inert is n-butane, its methanol enters at 320 K and its C4 at 350 K, it fixes the bottoms flow at 197 mol/s where the application fixes the duty, and it lists its feeds at 11 atm and its column at 10 atm [15]. The model holds 11 bar everywhere and takes both feeds as saturated. The pressure, the condenser and the flow loops are ideal. Its fifteen trays depart from the twelve first planned for the application because the fifteen-tray base case could be read, and a twelve-tray variant also folds in a prototype model built during development.

9.4 The Branches and Their Stability

9.4.1 An Independent Steady-State Solver

The branches are found by a second program. It writes the 51 steady balances (three fractions on 17 stages) again from Equations (9.1), (9.2) and (9.4) at constant holdup, with $x_D = y_1$ for the drum and the sump as an equilibrium stage, and it imports only the bubble point, the rate and an eigenvalue routine. It never calls the dynamic model's derivatives or the shared stage stack. The curve $\mathbf{f}(\mathbf{x}, Q) = 0$ is followed in the duty by pseudo-arclength continuation, as shown in Equation (9.8).

$$\begin{bmatrix} \partial_{\mathbf{x}}\mathbf{f} & \partial_Q\mathbf{f} \\ \mathbf{t}^{\mathsf{T}} & \end{bmatrix} \Delta\mathbf{z} = -\begin{bmatrix} \mathbf{f} \\ \mathbf{t}^{\mathsf{T}}(\mathbf{z} - \mathbf{z}_{pred}) \end{bmatrix}, \qquad \mathbf{z} = (\mathbf{x}, Q) \tag{9.8}$$

where $\mathbf{t}$ is the unit tangent and $\mathbf{z}_{pred}$ the predictor step along it. The Jacobian is numerical, the walk runs from 25 to 56 MW in arclength steps of 0.03, each fold's duty is refined by a walk at a tenth of

Table 9.4 The Dynamic Plant Against the Independent Solver

Comparison	Plant	Solver
41.0 MW, high, mean of the hour after 5 h (%)	98.237	98.230
45.5 MW, low, mean of the hour after 5 h (%)	81.548	81.548
42.5 MW, C4 start-up, 10–11 h (%)	98.218	98.219
42.5 MW, methanol start-up, 10–11 h (%)	90.039	90.041
derivatives at 8 solver states (per min)	$\leq 5.3 \times 10^{-11}$	
Hopf on the high branch (MW)	43.57	43.50
low branch held 10 h at 41.80 / 41.70 MW	low / gone	fold 41.75
high branch held 10 h at 46.10 / 46.20 MW	high / low	fold 46.15

the step and a parabola through the extreme, and each point is classified by the largest real part of the eigenvalues of the balances linearized with the dynamic model's holdups (15 kmol per tray, 100 kmol in the drum, 150 kmol in the sump). The curve takes 2.9 s headless and about 3 s in a browser with the run paused. Nothing logged, controlled or scored depends on it.

The solver and the plant agree where they can be compared without circularity, as listed in Table 9.4. Evaluated at the solver's states, the dynamic derivatives are zero to between 7×10^{-14} and 5×10^{-11} per minute. Moved off one solver state, the plant lines out on another. Started from a column refilled with the C4 cut or with methanol, neither of which is a solver state, it reaches the solver's high and low states at 42.5 MW.

9.4.2 Six Folds and a Hopf Point

At the published feeds the walk folds at 42.90, 40.46, 40.77, 38.35, 46.15 and 41.75 MW, as shown in Figure 9.2a. Three stable families result. The under-boiled branch, with too little vapor to strip the C4, runs up to its fold at 42.90 MW (a short stable sliver lies between the 40.46 and 40.77 MW folds). The high branch begins at its fold at 38.35 MW, and the low branch, with the methanol in the stripping trays, begins at its fold at 41.75 MW and continues to the end of the walk. Between 41.75 and 43.50 MW the high and the low branch coexist, the output multiplicity. At 42.5 MW the unstable saddles between the three stable states convert 95.24 and 95.43 %.

The high branch does not end at a fold. A complex pair of eigenvalues crosses the imaginary axis at 43.50 MW in the solver's constant-holdup linearization, with a period of 82 min, and between 43.55 and 43.60 MW in the full dynamic model's, whose real part goes from -2.8×10^{-3} to $+6.9 \times 10^{-3}$ per minute, so about 43.57 MW. A 10 h hold oscillates visibly only from 43.7 MW, where the conversion swings between 96.11 and 99.29 % over the last two hours. A short stable stretch of the high branch returns from 44.99 to 46.15 MW. Every fold reproduces at arclength steps of 0.05, 0.03, 0.02 and 0.005. At 0.01 one step straddles the 0.31 MW sliver between 40.46 and 40.77 MW and reports four folds. The exothermic stirred tank of Volume 7, Chapter 2 has the same ingredients in one vessel, ignition and extinction folds and a Hopf point on its hot branch.

9.4.3 Why the Methanol Decides the Branch

Methanol boils between the C4 and the MTBE, so the vapor rate decides where it sits. On the high branch it stays on the catalytic trays, where it reacts, and the bottoms are nearly pure ether. More boilup pushes it down. Once it is below the catalyst it cannot react, less MTBE is made, the heat of reaction boils less liquid, and the flows that put it there persist. On the low branch at 42.5 MW the bottoms carry 16.91 mol % methanol, and more duty pushes more down, so the conversion falls from 93.0 to 53.9 % between the low branch's fold and 56.17 MW. The branch monitor reads this on catalytic tray 8, which stays below 87.3 °C everywhere on the stable high branch and above 100.0 °C everywhere on the low branch.

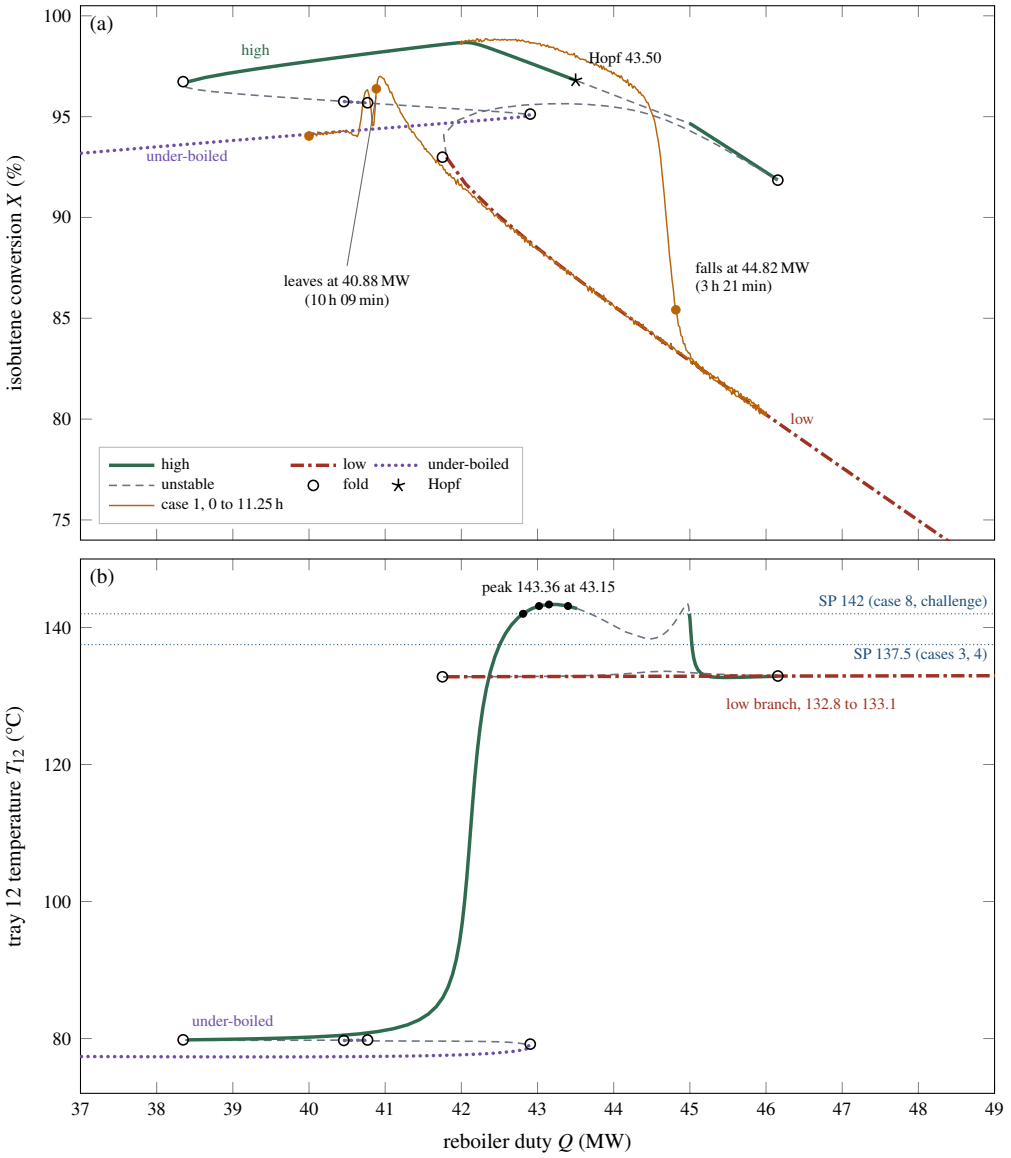

Figure 9.2 The S-curve at the published feeds from the application's own continuation. (a) Conversion against duty, stable branches heavy (high solid, low dash-dotted, under-boiled dotted), unstable states thin dashed, folds circled, the Hopf point starred, and the plant's path in case 1 from 42 MW up to 46 MW and down to 40 MW at 1 MW per hour, with the two jumps and the state at 11.25 h marked. (b) Tray 12 against duty on the same branches, with the two setpoints the cases use and the two duties at which tray 12 reads 143.13 °C on the high branch.

9.4.4 One Temperature at Two Duties

The stripping trays carry the temperature a loop can hold. On the high branch tray 12 rises with the duty from the C4's 75 °C toward the ether's 153 °C at 11 bar, peaks at 143.36 °C at 43.15 MW and falls after it, so 143.13 °C is read at both 43.02 and 43.4 MW, the input multiplicity of Figure 9.2b. On the low branch tray 12 holds methanol and MTBE and reads 132.81 to 133.07 °C at every duty from the fold to 56.17 MW, 133.10 at 58 and 133.13 °C at the reboiler's 60 MW. Tray 11 reads

128.87 to 130.43 °C there. A setpoint above that plateau cannot be reached on the low branch at all. The anti-lock brake of Volume 7, Chapter 4 holds its setpoint at the top of a friction curve, where the process gain passes through zero and reverses in the same way.

9.5 Loops as Implemented

9.5.1 The Tray-Temperature Loop and Its Helpers

Every loop runs the velocity-form PID law of Section 1.6.2, Equations (1.14) to (1.17), with its derivative term off and no output bias. The output is clamped after the sum, and the integral is pulled back toward the clamped value with a tracking time T_t equal to τ_I. The temperature loop (TIC) reads the thermowell of one stripping tray, 12 by default and 10 to 14 on offer, and sets the reboiler duty with action $a = +1$, because more duty heats the tray. In the convention of Section 1.6.2 that is a reverse-acting controller, although the app's text for case 8 calls it direct acting. Its defaults are $K_c = 0.15\,\mathrm{MW/K}$ and $\tau_I = 45\,\mathrm{min}$, it executes every $\Delta t = 0.1\,\mathrm{min}$, and its output runs from 0 to Q_{max}. Open loop at 42.5 MW, 0.2 MW more duty raises tray 12 by 3.42 K, two thirds of it in about 45 min. Integral action is off when τ_I reaches 600 min. In Manual the duty station is the operator's and the TIC tracks it, so a transfer in either direction is bumpless. When the control tray changes, a setpoint that sat within 1 K of the old tray's temperature moves to the new tray's, and any other is kept.

The level loops move the distillate and the bottoms with $a = -1$ (direct acting) at 40 kmol/h per percent and $\tau_I = 30\,\mathrm{min}$, the reflux follows as $L = RD$, and the methanol is set in ratio to the isobutene fed, $F_M = r_M F_{C4} z_{IB}$, with the measured C4 flow and the feed analyzer's isobutene or, as an option, the design 35.68 mol %. The optional bottoms trim (AIC) moves the TIC's setpoint by up to $\pm 15\,\mathrm{K}$ from the bottoms analyzer at 0.5 K per mol % and $\tau_I = 60\,\mathrm{min}$, executing on each report with the time since the last report as its control period. The analyzer is too slow to move the duty itself, which is why the TIC reads a tray, and Volume 5, Chapter 3 measures what such a delay costs a composition loop.

9.5.2 The Branch Monitor, the Supervisor and the Branch Names

The monitor reads catalytic tray 8 against a saddle temperature of 93 °C, between the high branch's 87.3 °C and the low branch's 100.0 °C. Five minutes above it reads LOW, and five minutes below it reads high again. With the supervisor on and the TIC in Auto, a LOW reading takes the duty from the TIC and cuts it to $f_{cut} = 0.9$ of the TIC's request, cuts it by 0.9 again after each hour the monitor still reads LOW, holds it until the monitor has read high for 30 min, and hands the duty back to the TIC, which has tracked it throughout. A path through methanol excess, duty and ratio was considered first. In a prototype a ratio cut left the column on the low branch for 10 h, so the supervisor cuts the duty instead.

A single cut lands below the low branch's fold at 41.75 MW only for a request under 46.4 MW, and two for one under 51.5 MW. Switched on at the start of case 4, at a duty of 42.50 MW, it cuts the TIC's first request, already raised by the proportional term, to 38.9 MW. The monitor reads high 13 min later and the column is on the high branch at 17 min. Switched on at 5 h, when the TIC asks for 47.79 MW, it cuts to 43.0 MW, above the fold, and an hour later to 38.7 MW, after which the monitor reads high at 6 h 14 min and the column is on the high branch at 6 h 18 min.

Three names for a branch appear in the application, and they are kept apart. The simulator's truth, used by the alarms, the challenge and the CSV, is low when the bottoms carry more than 6 mol % methanol and high otherwise. The readouts name the branch from the S-curve, low by the truth, otherwise the stable state at the plant's duty nearest it in conversion (high or under-boiled), with "past its stable end" and "between branches" when no such state exists, because the C4 in the bottoms cannot tell the high branch from the under-boiled one (the high branch carries 4.67 mol % C4 at 41 MW). The monitor is what an operator has and reads two states. The three can disagree. In case 5, 2 h after the methanol ratio is raised to 1.200, the bottoms carry 9.41 mol % methanol and 6.34 mol %

Table 9.5 The Nine Cases

#	Settings and events	Teaching point
1	Manual from 42 MW, ramp to 46 and back to 40 MW at 1 MW/h	hysteresis
2	Manual at 42.5 MW, learner moves to 41 MW	curve against plant
3	TIC on tray 12 at 137.5 °C, C4 feed +5 % at 1 h	the loop on the high branch
4	the same loop, started on the low branch	the loop on the low branch
5	Manual, learner raises the methanol ratio to 1.200	azeotrope, methanol down
6	Manual, C4 cut to 33 mol % isobutene at 1 h	a fold moves
7	Manual, two start-up charges	the path picks the branch
8	Manual at 43.4 MW, TIC setpoint 142 °C	input multiplicity
9	challenge, 8 h, C4 +5 % at 2 h, steam 85 % for 20 min at 5 h	the whole problem

C4, the truth and the readout say low, and the monitor still reads high with tray 8 at 75.3 °C, because the excess methanol reached the bottoms while the C4 was not stripped.

🕮 **Pause.** Try Exercise 9.3 before reading on.

9.5.3 Three Models

The plant is Section 9.3. The S-curve is a separate solution of its steady equations, read by the Branches view, the branch readout and the scenario texts. The controllers carry no process model. The monitor's saddle temperature and the supervisor's cut are the only places where knowledge of the branches enters the control system, and both are fixed numbers. The alarms, the specification of 98 mol % MTBE, the target of 95 % conversion, the reboiler's 60 MW and the distillate alarm at 7.0 mol % are teaching values. The real ones come from the unit's licensor, its steam system and the buyer's contract.

9.6 Nine Cases, Measured

Table 9.5 names the nine cases and Table 9.6 reads the app's own scorecard at the instant each is read. The cases, their completion rules and the challenge follow Section 1.6.4. Three cases carry most of the lesson. Case 1 draws the hysteresis loop before any controller is closed, case 4 shows a correct loop driving the column the wrong way, and case 7 shows that the start-up charge decides where the column lands. The measurements reported here were made headless at the default seed. The engine advanced through the same one-second physics steps as the page, each case was loaded through the page's own loading path, and every learner's move was applied at the time the text gives. Every case runs identically when loaded first, picked from the list after another case, or restarted with Reset.

9.6.1 A Hysteresis Loop on a Duty Ramp

Case 1 opens on the high branch at 42 MW, 98.67 % converted, and ramps the duty at 1 MW per hour from 30 min. The conversion eases to 95.53 % at 44.47 MW by 3 h, then falls, 92.29 at 44.63 MW and 85.76 % at 44.80 MW, and the column is on the low branch at 44.82 MW, 3 h 21 min into the run. On the way down it stays there, 82.82 % at 45.03 MW at 6 h, and it leaves at 40.88 MW at 10 h 9 min, so the two jumps are 3.93 MW apart. The ramp carries the column 1.3 MW past the Hopf point, because the oscillation that loses the high branch grows slowly, and 0.9 MW past the low branch's fold, as shown in Figure 9.2a and Figure 9.3. Both jumps are unchanged at half the step, and at seeds 1 to 3 the upward jump lies between 44.816 and 44.833 MW.

🕮 **Pause.** Try Exercise 9.1 before reading on.

Case 2 compares the plant with the curve. At 42.5 MW the curve lists under-boiled 94.9, high 98.2 and low 90.1 %. Moved to 41.00 MW, the plant holds within 0.1 percentage points of the solver's 98.230 % for 20 min by 87 min after the move, and its mean over the seventh hour is 98.220 %. The high branch at 41 MW reaches 4.71 mol % C4 in its bottoms by 7 h, and C4 IN BOTTOMS lights 53 min after the move and OFF SPEC 72 min after it. High conversion and an on-specification product are different things on this column.

Table 9.6 The Application's Scorecard Tiles at the Instants Each Case Is Read

Case	Settings and instant	Duty (MW)	Conversion (%)	Branch tile	Bottoms MTBE (mol %)	Control tray (°C)	Also
1	as shipped, 3 h	44.47	95.5	high (past end)	99.90	138.5	
1	as shipped, 3 h 21 min	44.82	85.4	low	93.79	132.9	monitor LOW
1	as shipped, 6 h	45.03	82.8	low	75.26	132.9	methanol 24.74 mol %
1	as shipped, 10 h 9 min	40.88	96.4	high	94.20	132.6	
1	as shipped, 11 h 15 min	40.00	94.0	under-boiled	85.89	77.4	C4 11.73 mol %
2	41 MW from 10 min, 7 h	41.00	98.3	high	95.29	81.2	C4 4.71 mol %
3	Auto, 8 h	44.75	98.1	high	99.94	137.4	
3	Manual, 8 h	42.50	97.8	high	92.93	80.4	C4 7.07 mol %
4	Auto, 5 h	47.76	75.6	low	69.58	132.9	monitor LOW
4	cut to 40 MW at 2 h, 8 h	40.00	94.1	under-boiled	84.55	77.3	C4 13.86 mol %
5	ratio 1.200 from 30 min, 2.5 h	42.50	94.8	low	84.25	77.3	monitor high
5	1.200 at 30, 1.100 at 50 min, 6.83 h	42.50	94.9	under-boiled	92.67	77.8	C4 6.94 mol %
6	fixed duty, 3 h	42.50	93.5	under-boiled	82.78	76.7	C4 16.55 mol %
6	44 MW from 91 min, 7.52 h	44.00	98.4	high	96.59	81.4	C4 3.41 mol %
7	C4 start-up, 2 h	42.50	98.8	high	99.75	112.5	
7	methanol start-up, 2 h	42.50	89.9	low	83.12	132.8	methanol 16.88 mol %
8	Auto, SP 142 °C, 8 h	42.81	97.8	high	99.96	141.9	
8	Auto, SP 143.5 °C, 6 h	54.97	57.7	low	53.21	133.3	



Figure 9.3 The hysteresis loop of case 1 at $t = 11.25$ h on the Branches view, in Manual, light theme, cropped to the plot, its status chips and its alarm strip. (1) The high branch, green (the same branches carry line styles in Figure 9.2a), with the six folds circled and the unstable states dashed. (2) The low branch, red. (3) The plant's path in amber, its chevrons pointing up the high branch, down onto the low branch at 44.82 MW, back along it and off it at 40.88 MW. (4) The plant at 40.00 MW on the under-boiled branch, violet, 94.0 % converted with 85.89 mol % MTBE and 11.73 mol % C4 in the bottoms. (5) The readout, with the curve's under-boiled 94.1 % and high 97.7 % at that duty.

TRY IT IN THE APP **1 · Ramp the duty up and back down**

Load the scenario, open the Branches view and press Run. Watch the plant's amber dot against the S-curve: it drops to the low branch at 44.82 MW on the way up and leaves it only at 40.88 MW on the way down.

9.6.2 One Loop, Two Branches

Case 3 runs the TIC on tray 12 at 137.5 °C and raises the C4 feed 5 % at 1 h. The loop takes the duty from 42.5 to 44.76 MW within the hour, tray 12 dips to 128.8 °C and is back within 1 K from 3 h 22 min, and the bottoms never fall below 99.90 mol %. On the fixed duty tray 12 falls to 80.4 °C and the bottoms to 92.93 mol % MTBE with 7.07 mol % C4, while the conversion hardly moves (97.83 against 98.09 % with the loop). The reaction section does not see the feed step, and the stripping section does.

Case 4 starts the same loop on the low branch, where tray 12 reads 132.8 °C, below its setpoint. The loop raises the duty past 45 MW at 2 h and to 47.76 MW at 5 h, tray 12 stays at 132.8 to 132.9 °C, and the conversion falls from 90.0 to 75.6 %. No duty on the low branch reaches 137.5 °C, which makes this the honest infeasible case of the set. Cut to 40 MW in Manual at 2 h, below the low branch's fold, the column leaves it, the monitor reading high 24 min later and the truth 27 min later, and it settles on the under-boiled branch.

Pause. Try Exercise 9.2 before reading on.

Case 8 starts on the high branch at 43.4 MW, past tray 12's peak, where the tray reads 143.13 °C. Switched to Auto at 142 °C, the loop cuts the duty, and the tray first gets hotter, 143.20 °C at 1 h and 143.35 °C at most, as the column climbs back over the peak. It then cools, and the TIC holds 142 °C at 42.81 MW on the rising side, within 0.5 K for 30 min by 2 h 57 min. At a setpoint of 143.5 °C, above the peak, the loop raises the duty, the tray cools, and the column falls onto the low branch 1 h 36 min in with the duty still climbing, 48.77 MW at 3 h and 54.97 MW at 6 h. Across seeds 0 to 3, and with the step halved, the holding run meets its rule between 2 h 57 min and 3 h 19 min, and the lost run falls between 1 h 36 min and 1 h 44 min.

> **TRY IT IN THE APP** **8 · One tray temperature, two duties**
>
> Load the scenario and switch the header to Auto. The loop cuts the duty, yet tray 12 first warms, to 143.35 °C at most, before it settles at 142 °C with 42.81 MW. Reload with the circular arrow, set the tray temperature setpoint to 143.5 °C and switch to Auto again. Now the loop raises the duty and the column falls to the low branch 1 h 36 min in.

9.6.3 Methanol, the Feed and the Start-Up

Case 5 raises the methanol ratio from 1.100 to 1.200 at 30 min. The distillate passes its alarm 3 min later and sits at the azeotrope, 7.63 mol % methanol, and the rest falls, past 3 mol % of the bottoms 20 min after the move. Left at 1.200, the column is on the low branch 2 h after the move. Put back to 1.100 at 50 min, the methanol leaves the bottoms within 55 min, but 6 h later the column is on the under-boiled branch, 94.90 % converted with 6.94 mol % C4, at the duty and ratio it started from. With the TIC in Auto the same excess takes the duty to 49.08 MW and the conversion to 72.11 % 5 h after the move.

Case 6 turns the C4 cut leaner, 33 mol % isobutene, at 1 h. The feed analyzer reports it and the methanol ratio uses it 20.1 min later. On the new curve the under-boiled branch reaches up to 43.47 MW, the high branch begins at 38.77 MW and the low branch at 44.95 MW, so 42.5 MW lies inside the under-boiled range. C4 in the bottoms passes 5 mol % at 1 h 31 min, and on the fixed duty it reaches 16.55 mol % by 3 h. Raised to 44 MW then, past the new fold, the column is above 98 % converted 97 min later and at 98.40 % after 6 h, with 3.41 mol % C4 still in the bottoms. The conversion dips to 84.2 % in the minute rows about 9 min after the feed change. The dip is physical rather than numerical, since it reappears at a quarter of the step (84.17 %) and at seeds 1 to 3 (82.81 to 84.01 %).

Case 7 refills every tray at 42.5 MW. Charged with the C4 cut, the column is on the high branch at 2 h, 98.8 % converted with 99.75 mol % MTBE in the bottoms. Charged with methanol, the resin wetted first, it is on the low branch at 2 h, 89.9 % converted with 16.88 mol % methanol in the bottoms, and it stays there. The case was first planned with the methanol-rich path landing high. The model

Table 9.7 The Challenge Over 8 h, Default Seed

Settings	Low x_B (mol %)	Low X (%)	Low branch	Verdict
tray 12, 142 °C (shipped)	67.58	73.0	from 6 h 32	fails
tray 11, 120 °C, supervisor (suggested)	99.70	97.2	never	passes
tray 11, 120 °C, no supervisor	99.70	97.2	never	passes
tray 12, 142 °C, supervisor	97.18	92.4	never	fails
tray 12, 137.5 °C	99.64	97.3	never	passes

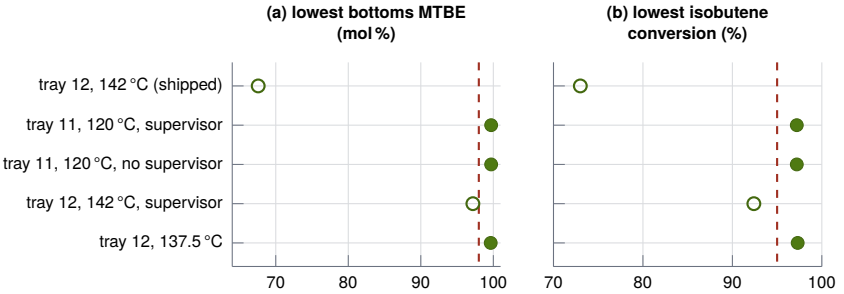

Figure 9.4 The challenge of Table 9.7 as a chart: (a) the lowest bottoms MTBE of each setting against the 98 mol % specification and (b) the lowest conversion against the 95 % target, both dashed. Filled dots pass and open dots fail. Notice that the supervisor keeps tray 12 at 142 °C off the low branch but still short of both lines, while the same loop at 137.5 °C, and tray 11 at 120 °C with or without the supervisor, clear both.

lands it low, and the case teaches what the model shows. At half, a quarter and a tenth of the step the methanol start-up lands on the same state.

> ✦ **DID YOU KNOW? Two Start-Ups, Two Columns**
>
> From the simulator: at the same 42.50 MW, feeds and reflux ratio, the column charged with the C4 cut is 98.8 % converted after 2 h, with 99.75 mol % MTBE in its bottoms. Charged with methanol it is 89.9 % converted, with 16.88 mol % methanol in its bottoms (Table 9.6).

9.6.4 The Challenge

The challenge runs 8 h under the TIC with tray 12 at 142 °C and no supervisor. The C4 feed rises 5 % at 2 h and the steam dips to 85 % for 20 min at 5 h. It passes when, at every 10 min sample of the true column from 30 min, the column is on the high branch with at least 98 mol % MTBE in the bottoms and at least 95 % converted. As shipped the column holds through the feed step, but after the steam dip the loop's catch-up overshoots tray 12's peak and the column falls to the low branch at 6 h 32 min. The loop then raises the duty to 51.2 MW by 8 h, and the bottoms fall to 67.58 mol %. The settings tried are listed in Table 9.7, and Figure 9.4 draws them as a chart.

The shipped setting fails on its margin, not on its tuning. Tray 12 peaks at 143.36 °C on the high branch, 1.4 K above the 142 °C setpoint, and the same gains on the same tray at the default 137.5 °C pass with 99.64 mol % and 97.3 % at their lowest. The suggested settings move the loop to tray 11 at 120 °C, which the high branch reaches and which lies below the low branch's 128.87 to 130.43 °C on that tray, so a loop on the low branch cuts the duty and returns the column by itself. They pass at the default seed as well as seeds 1 through 6, with 99.70 to 99.71 mol % and 97.2 to 97.3 % at their lowest, and pass at half the step. The supervisor never acts in a suggested run. With tray 12 at 142 °C the supervisor cuts the duty to 43.3 MW at 6 h 22 min and keeps the column off the low branch, but not on specification. No cut fraction from 0.7 to 0.98 passes on that setting. The shipped run fails at

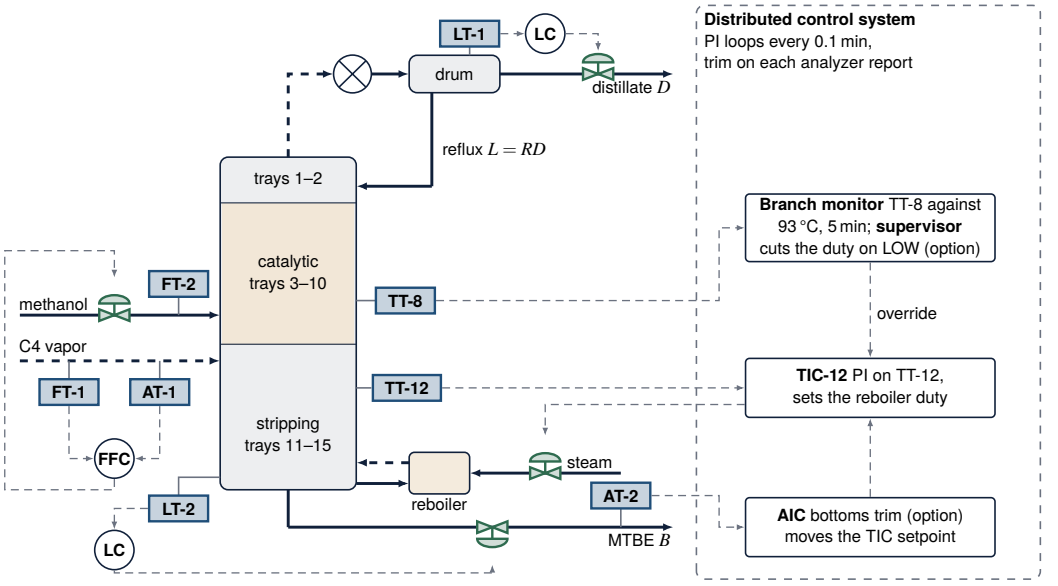

Figure 9.5 The loops of the MTBE column with their hardware. Lettered boxes are sensing devices: the thermowells on trays 8 and 12 (TT), the C4 flow meter (FT-1) and feed analyzer (AT-1), the methanol flow meter (FT-2), the bottoms analyzer (AT-2) and the two level transmitters, whose level controllers (LC) and the methanol ratio station (FFC) are drawn beside their devices. Valve symbols are final elements, thick dashed lines carry vapor and thin dashed arrows are signals. Notice that the temperature loop moves the steam valve on the reboiler, that the branch monitor reads a catalytic tray rather than the control tray, and that the methanol follows the C4 that is measured, not the C4 that is wanted. The distillate analyzer and the alarms are left out. Not to scale, and not a piping diagram.

seeds 1 to 6 as well, its lowest bottoms between 67.66 and 67.75 mol %. The scored samples begin to read low at 6 h 40 min, the first 10 min sample after the fall.

9.7 Technology in Practice

What Measures It

Figure 9.5 draws the loops of Section 9.5 with their devices. The temperature loop and the branch monitor read tray thermowells, modeled with a 0.5 min lag and 0.05 K of noise. They hold thermocouples or resistance thermometers (Chapter 4 and Chapter 8), and transmitters close together on a column can disagree where the flow patterns differ [19]. The methanol ratio, which decides the branch, rests on two flow meters and the feed analyzer.

TECHNOLOGY Sensor: A Coriolis Meter on the Methanol Feed

A Coriolis meter shakes one or more tubes that carry the fluid at their resonant frequency. The frequency reads the density, and the twist that the moving fluid adds, seen as a phase shift between two points on the tube, is proportional to the mass flow. The meter is very accurate but costly, has difficulty with gases and with mixed liquid and vapor, and errs when outside vibration reaches its tubes or gas bubbles collect in its bent ones [19]. Under a flow loop a meter that reads low makes the valve deliver more methanol than the ratio asks for, the excess of case 5. The model's flow meters lag 0.05 min with 0.2 % noise.

What Moves It

The reboiler duty is the temperature loop's handle, and in a plant it is the steam to the reboiler. The methanol, the distillate and the bottoms are liquid flows behind control valves, and the reflux follows the distillate in the ratio $L = RD$. In the model each of these flows follows its setpoint through a 0.2 min lag.

TECHNOLOGY Final Element: The Steam Valve on the Reboiler

A temperature loop on a heat exchanger usually throttles the heating steam with a control valve and leaves the process flow to the rest of the plant. An air-to-open steam valve closes as its signal falls, and that choice sets the controller's action. If the steam system fails, integral action drives the valve wide open in vain, and when the steam returns the saturated output sends too much and the temperature overshoots, unless the operator switches the loop to manual and brings the output down [19]. The model lumps valve and reboiler into a duty that follows the request times the steam available with a 2 min lag. Its challenge dips the steam to 85 % for 20 min, and the loop's catch-up then overshoots tray 12's peak.

Where the Controller Runs

A distributed control system runs nested and cascaded loops under a supervisory level, with cycle times from milliseconds to minutes [20]. The app executes its temperature and level loops every 0.1 min and its trim on each analyzer report, every 10 min, inside that range. The PI blocks, the ratio station and the trim belong to the regulatory layer. The branch monitor and the supervisor's duty cut are logic above it that overrides the temperature loop's output, and the loop's tracking of that output is what an industrial controller calls output tracking, which lets a loop pick up from a value it did not set [19].

IN PRACTICE One Temperature Loop to Hold the Profile

A column is first stabilized by its level and pressure loops. Without composition control its profile behaves almost like an unstable system that must be watched continuously, a heavy burden on the operators, so at least one composition loop is closed, for example on the boilup. The operator then tends to overpurify the uncontrolled product, which costs energy and capacity. A tutorial on column control recommends at least three temperature sensors in each section, so that pressure changes can be corrected and fast temperatures can feed a cascade or estimate the product compositions, and a reboiler that can change the boilup fast [21]. The tray-12 loop on the duty is this one-point control.

What Is Modernizing It

Model predictive control has a solid foundation in refining and petrochemicals [22]. For distillation a tutorial of 1997 noted that the cost of a dynamic model and of tuning had kept industry from a predictive controller on all five handles of a column, and that a first-principles model is probably better than an identified one [21]. The supervisor that knows where the fold lies, which the Conclusion names as a next addition, would rest on such a model. Soft sensors that estimate compositions from tray temperatures have distillation as their main target, and keeping their models up to date is the most critical issue their users report [23]. Reactive dividing-wall columns, which put the reaction of

this chapter inside the column of Chapter 8, promise further energy savings and are discussed as an open challenge after the limitations.

9.8 Limitations

The column's shape, feeds and pressure come from a published base case and its properties and kinetics from read sources, but the catalyst mass, the rate term ε and the latent heat are teaching values on which the fold duties depend, so the duties at which this column jumps are not predictions for a real one. The trays are equilibrium stages with constant molar overflow, the reaction has no mass-transfer resistance, the inert has isobutene's volatility, and no side reaction or energy balance per tray is modeled. The Wilson parameters are fitted values read at second hand, in the reading that gives the azeotropes the sources describe rather than the one printed.

The controllers carry no model and the supervisor's path is one declared move, where a real one would come from the unit's procedures. The monitor reads one tray and two states, and case 5 shows it reading high while methanol has reached the bottoms. Within a walk step of a fold the readout can name the wrong family (at 38.40 MW a plant still at 96.7 % reads under-boiled), a state no case reaches. On the low branch at high duty some fractions leave $[0, 1]$ by less than the step-splitting tolerance (cases 4 and 8 at 143.5 °C late in their runs). The disturbances arrive as steps. A C4 cut that turns leaner in one step is not how a real feed changes, but it separates the fold's movement from the loop's response.

> **OPEN CHALLENGE** **A Reactive Column Behind a Dividing Wall**
>
> Putting the reaction of this chapter inside the dividing-wall column of Chapter 8 promises 15 to 75 percent less energy. By 2018, however, experiments covered only five chemical systems, the models were uncertain, and complex control such as model predictive control had not been demonstrated on such a column. Progress would be experimental and pilot-plant studies with dynamic and rate-based models checked against them [9].

9.9 Conclusion

An MTBE column of fifteen trays, built on the series' shared stage stack with a Wilson activity model and an activity-based rate, has at 42.5 MW a high state converting 98.22 %, a low state converting 90.05 % with the methanol below the catalyst, and an under-boiled state converting 94.89 %. A solver that shares only the thermodynamics with the plant finds six folds from 38.35 to 46.15 MW and a Hopf point at 43.50 MW, and the plant lines out on its states within 0.007 percentage points in the comparisons of Table 9.4. Ramped at 1 MW per hour the column jumps at 44.82 and 40.88 MW. A tray-12 loop that holds the high branch through a feed step raises the duty to 47.76 MW on the low branch, where tray 12 never passes 133.13 °C, and a setpoint 1.4 K under the high branch's peak loses the challenge where tray 11 at 120 °C passes at seven seeds. The fold duties rest on three teaching values. The next additions to the model would be a rate-based tray model, side reactions, an energy balance per tray and a supervisor that knows where the fold lies. The simulator shows that a correct loop on a correct measurement can still drive a column away from the state it was tuned for.

KEY POINTS

- With the same feeds, reflux and duty, the column has three steady states.
- At 42.5 MW the high, low and under-boiled states convert 98.22, 90.05 and 94.89 %. Ramped at 1 MW per hour, the column jumps at 44.82 and 40.88 MW.
- On the low branch tray 12 never passes 133.13 °C, and the tray-12 loop raises the duty to 47.76 MW.
- Tray 12 at 142 °C loses the challenge on its margin, not its tuning.

Where else this applies. Tray 12's peak returns in Volume 7, Chapter 4 and Volume 2, Chapter 9.

References

[1] E. W. Jacobsen and S. Skogestad, "Multiple steady states in ideal two-product distillation," *AIChE Journal*, vol. 37, no. 4, pp. 499–511, 1991.

[2] L. B. Koppel, "Input multiplicities in nonlinear, multivariable control systems," *AIChE Journal*, vol. 28, no. 6, pp. 935–945, 1982.

[3] A. Kienle and W. Marquardt, "Nonlinear dynamics and control of reactive distillation processes," in *Reactive Distillation: Status and Future Directions*, K. Sundmacher and A. Kienle, Eds. Wiley, 2002, pp. 241–281.

[4] T. E. Güttinger and M. Morari, "Predicting multiple steady states in equilibrium reactive distillation. 2. Analysis of hybrid systems," *Industrial & Engineering Chemistry Research*, vol. 38, no. 4, pp. 1649–1665, 1999.

[5] K.-D. Mohl, A. Kienle, E.-D. Gilles, P. Rapmund, K. Sundmacher, and U. Hoffmann, "Steady-state multiplicities in reactive distillation columns for the production of fuel ethers MTBE and TAME: theoretical analysis and experimental verification," *Chemical Engineering Science*, vol. 54, no. 8, pp. 1029–1043, 1999.

[6] G. M. Wilson, "Vapor-liquid equilibrium. XI. A new expression for the excess free energy of mixing," *Journal of the American Chemical Society*, vol. 86, no. 2, pp. 127–130, 1964.

[7] A. Rehfinger and U. Hoffmann, "Kinetics of methyl tertiary butyl ether liquid phase synthesis catalyzed by ion exchange resin—I. Intrinsic rate expression in liquid phase activities," *Chemical Engineering Science*, vol. 45, no. 6, pp. 1605–1617, 1990.

[8] Z. Nawaz, "Methyl-tert-butyl-ether synthesis reactor modelling and optimization using an Aspen Custom Modeler," *Hungarian Journal of Industry and Chemistry*, vol. 45, no. 2, pp. 1–7, 2017.

[9] A. A. Kiss and M. Jobson, "Taking reactive distillation to the next level of process intensification," *Chemical Engineering Transactions*, vol. 69, pp. 553–558, 2018.

[10] G. J. Harmsen, "Reactive distillation: The front-runner of industrial process intensification," *Chemical Engineering and Processing: Process Intensification*, vol. 46, no. 9, pp. 774–780, 2007.

[11] R. Jacobs and R. Krishna, "Multiple solutions in reactive distillation for methyl tert-butyl ether synthesis," *Industrial & Engineering Chemistry Research*, vol. 32, no. 8, pp. 1706–1709, 1993.

[12] S. Hauan, T. Hertzberg, and K. M. Lien, "Why methyl tert-butyl ether production by reactive distillation may yield multiple solutions," *Industrial & Engineering Chemistry Research*, vol. 34, no. 3, pp. 987–991, 1995.

[13] A. P. Higler, R. Taylor, and R. Krishna, "Nonequilibrium modelling of reactive distillation: Multiple steady states in MTBE synthesis," *Chemical Engineering Science*, vol. 54, no. 10, pp. 1389–1395, 1999.

[14] S. Schrans, S. de Wolf, and R. Baur, "Dynamic simulation of reactive distillation: An MTBE case study," *Computers & Chemical Engineering*, vol. 20, pp. S1619–S1624, 1996.

[15] A. D. Estrada-Villagrana, "On the modeling of dynamics and operation of reactive distillation," Ph.D. dissertation, University of London (University College London), 2002. [Online]. Available: https://discovery.ucl.ac.uk/id/eprint/1317706/

[16] M. G. Sneesby, M. O. Tade, and T. N. Smith, "Two-point control of a reactive distillation column for composition and conversion," *Journal of Process Control*, vol. 9, no. 1, pp. 19–31, 1999.

[17] W. L. Luyben and C.-C. Yu, "Control of MTBE and ETBE reactive distillation columns," in *Reactive Distillation Design and Control*. Wiley, 2008, pp. 407–427.

[18] D. Porcelli, "Comparative techno-economic assessment between reactive and reactive-cyclic distillation for process intensification: theoretical framework and application to MTBE and MeOAc reactive distillation," Master's thesis, Politecnico di Milano, 2021. [Online]. Available: http://hdl.handle.net/10589/182157

[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] S. Skogestad, "Dynamics and control of distillation columns: A tutorial introduction," *Chemical Engineering Research and Design (Transactions of the Institution of Chemical Engineers, Part A)*, vol. 75, no. 6, pp. 539–562, 1997.

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

[23] 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.

9.10 Exercises

9.1 Concept check · 2 min

In case 1 the duty ramps up at 1 MW per hour, and the column falls to the low branch at 44.82 MW. The S-curve says the high branch turns unstable at its Hopf point, 43.50 MW. Why does the fall come 1.3 MW later than the curve suggests?

(a) The reboiler responds with a 2 min lag, so the plant sees an older and lower duty
(b) The branch is lost through an oscillation that grows slowly, and the ramp outruns its growth
(c) The solver holds the tray holdups constant, which puts its Hopf point about a megawatt too low
(d) The low branch does not exist below 44.82 MW, so the column has no other state to go to

9.2 Concept check · 1 min

In case 4 the TIC holds tray 12 at a setpoint of 137.5 °C, and the column starts on the low branch with tray 12 at 132.8 °C. By 5 h the loop has raised the duty from 42.5 to 47.76 MW while tray 12 stays at 132.8 to 132.9 °C. Why?

(a) The integral has wound up against the reboiler's 60 MW limit and cannot come back
(b) The tray's gain to the duty is negative on the low branch, so the loop is in positive feedback
(c) On the low branch tray 12 holds methanol and MTBE near 133 °C at any duty
(d) The thermowell's 0.5 min lag hides the duty's effect from a loop that runs every 6 s

9.3 Concept check · 2 min

In case 5 the methanol ratio is raised to 1.200 at 30 min and left there. Two hours later the bottoms carry 9.41 mol % methanol and 6.34 mol % C4. The truth flag reads low, while the branch monitor still reads high with tray 8 at 75.3 °C. Which explains the disagreement?

(a) The monitor waits 5 min above its saddle before it reads LOW, and that time has not passed
(b) The truth flag is misled by the azeotrope, which carries the methanol out with the distillate
(c) The saddle temperature of 93 °C is set too high for a column fed at a ratio of 1.200
(d) Tray 8 stays near the C4's temperature because the C4 is not stripped, while methanol still falls

9.4 Theory and calculation · 6 min

The column runs at the design point, 42.5 MW, with a latent heat $\lambda = 20\,\mathrm{kJ/mol}$ and a heat of reaction $\Delta H_r = -37.7\,\mathrm{kJ/mol}$. It is fed 1976.4 kmol/h of C4 cut at 35.68 mol % isobutene, and the high branch makes 693 kmol/h of MTBE.

(a) Use Equation (9.2) to find the boilup V_B in kmol/h and the moles of vapor ν_V boiled by each mole of MTBE. *Checkpoint:* 7650 kmol/h and 1.885.

(b) Find the vapor the reaction adds in the catalytic section, and the reboiler duty that would boil the same amount. *Checkpoint:* 1306 kmol/h, the vapor of 7.26 MW.

(c) At the same duty the low branch converts 90.05 % of the isobutene and the high branch 98.22 %. How much less vapor does the reaction make on the low branch, in kmol/h and as reboiler duty? *Checkpoint:* 108.6 kmol/h less, the vapor of 0.60 MW.

9.5 Theory and calculation · 6 min

The rate of Equation (9.5) uses the equilibrium constant $K_{350} = 28.5$ with $\Delta H_r = -37.7\,\mathrm{kJ/mol}$ and $R = 8.314\,\mathrm{J/(mol\,K)}$, and the term $\varepsilon = 0.05$ in its denominator.

(a) Find K at 25 °C and at 100 °C. *Checkpoint:* 271.2 and 12.76.

(b) Write the denominator $(a_M + \varepsilon)^2$ as $\varepsilon^2(1 + K_{ads}a_M)^2$ and give K_{ads} and the factor by which the rate constant is scaled. *Checkpoint:* $K_{ads} = 20$, scaled by 400.

(c) Find the ratio of the rate to its $\varepsilon \to 0$ form, $a_M^2/(a_M+\varepsilon)^2$, at methanol activities of 0.11 and 0.44, for $\varepsilon = 0.05$ and for $\varepsilon = 0.02$. Give two decimals. *Checkpoint:* 0.47 and 0.81 at $\varepsilon = 0.05$, 0.72 and 0.91 at $\varepsilon = 0.02$.

9.6 Theory and calculation · 6 min

The supervisor cuts the duty to $f_{cut} = 0.9$ of the TIC's request when the monitor reads LOW, and again to 0.9 of the held value after each hour it still reads LOW. At the published feeds the low branch begins at its fold at 41.75 MW and the high branch at its fold at 38.35 MW.

(a) Find the largest request that one cut brings below 41.75 MW, and the largest that two cuts bring below it. *Checkpoint:* 46.39 MW and 51.54 MW.

(b) In case 4 the supervisor is switched on at 5 h, when the TIC asks for 47.79 MW. Find the duty after the first cut and after the second. Say for each whether the low branch and the high branch exist there. *Checkpoint:* 43.01 MW, where both exist, then 38.71 MW, where only the high branch of the two does.

(c) The cut is set to 0.80 instead. Find the duty after one cut from the same request, and the range of requests that one cut at 0.80 brings between the two folds. *Checkpoint:* 38.23 MW, below both folds. Requests from 47.94 to 52.19 MW.

9.7 App challenge · 8 min

Load scenario 1, Ramp the duty up and back down, and change no setting. With the default seed, run to 11 h 15 min (at the top speed of 120 min per second this takes about 6 s) and export the CSV file. Its time column is in hours, one row per minute.

(a) Find the first row in which the column `truth_branch` reads low. Give its time and the duty in `duty_MW`, within 0.02 MW. *Checkpoint:* 3.35 h (3 h 21 min) and 44.82 MW.

(b) Find the first later row in which `truth_branch` reads high again, its duty, and the width of the hysteresis loop. *Checkpoint:* 10.15 h (10 h 9 min), 40.88 MW, a loop 3.93 MW wide.

(c) The S-curve ends the high branch at its Hopf point, 43.50 MW, and the low branch at its fold, 41.75 MW. By how much does the ramp pass each before the column jumps? *Checkpoint:* 1.32 MW and 0.87 MW.

9.8 App challenge · 9 min

Load scenario 8, One tray temperature, two duties. The column starts on the high branch in Manual at 43.4 MW, with the tray temperature setpoint at 142.0 °C. Use the default seed and export the CSV file after each run.

(a) Switch the header to Auto before pressing Run, and run to 8 h. Read the highest value of `truth_T12_C` in the file, within 0.05 K, and the duty `duty_MW` at 8 h, within 0.02 MW. *Checkpoint:* 143.35 °C, 27 min in, and 42.81 MW.

(b) Press the reload button, set the Tray temperature setpoint to 143.5 °C, switch to Auto and run to 6 h. Find the first row in which `truth_branch` reads low, and read the duty at 3 h and at 6 h, within 0.05 MW. *Checkpoint:* 1.60 h (1 h 36 min), then 48.77 and 54.97 MW.

(c) Use Figure 9.2b to explain both runs. *Checkpoint:* One side of the tray-12 peak has a positive gain, the other a negative one, and 143.5 °C is above the peak.

9.9 App challenge · 10 min

Load the challenge, eight hours on the high branch. A C4 feed step comes at 2 h and a steam dip at 5 h, and the run stops itself at 8 h. Use the default seed. After each run, read the three So far tiles (on the low branch, lowest bottoms MTBE, lowest conversion) and the verdict.

(a) Run as shipped, with tray 12 at 142.0 °C and the supervisor off. *Checkpoint:* Low branch from 6 h 40 min, lowest bottoms 67.58 mol % MTBE, lowest conversion 73.0 %. Fails.

(b) In the TIC card choose Control tray 11, then set the Tray temperature setpoint to 120.0 °C. Leave the supervisor off, press Reset and run again. *Checkpoint:* Never on the low branch, 99.70 mol %, 97.2 %. Passes.

(c) Choose Control tray 12, then set the setpoint to 137.5 °C, press Reset and run again. What do the three runs say about the cause of the shipped failure? *Checkpoint:* Never on the low branch, 99.64 mol %, 97.3 %. Passes.

9.10 Application area · 4 min

The methanol ratio station computes the isobutene fed from the measured C4 flow and either the feed analyzer or the design 35.68 mol %. The analyzer samples every 10 min through a 1 min sample line and reports a cycle later.

(a) After a step in the isobutene of the C4 cut, what is the longest the station can run on the old value with the analyzer? *Checkpoint:* About 21 min. Case 6 measured 20.1 min.

(b) With the design fraction instead, the C4 cut turns to 33 mol % isobutene as in case 6. What methanol ratio does the column then receive at the shipped setting of 1.100? Compare it with case 5. *Checkpoint:* 1.189, close to the 1.200 that put methanol in the bottoms in case 5.

9.11 Application area · 4 min

An operations engineer proposes to copy the simulator's numbers into the procedure of a real MTBE column: "Keep the reboiler below 43.5 MW. If the conversion drops, cut the duty below 41.75 MW and wait."

(a) Which model values set these duties, and how far do the folds move when they change? *Checkpoint:* Catalyst mass, ε and λ. The folds move by megawatts.

(b) What would you watch on the real column instead of a duty to know that the high branch is near its end, and that the column has left it? *Checkpoint:* Tray 12's temperature levelling off as the duty rises, and the catalytic tray 8 temperature.

9.12 Review the agent · 5 min

An AI assistant reviews the failed challenge run and replies: "The column holds through the feed step and is lost after the steam dip, when the loop catches up too hard and overshoots. The setpoint of 142 °C is fine, since it is below the 143.36 °C peak. Halve the TIC gain to 0.075 MW/K and double the integral time to 90 min, and the column will pass." Find the flaw and say how you would check it in the app.

CHAPTER 10

A Synthetic Base Oil Simulator for Mid-Ranging, Inferential Control and Constraint Back-Off

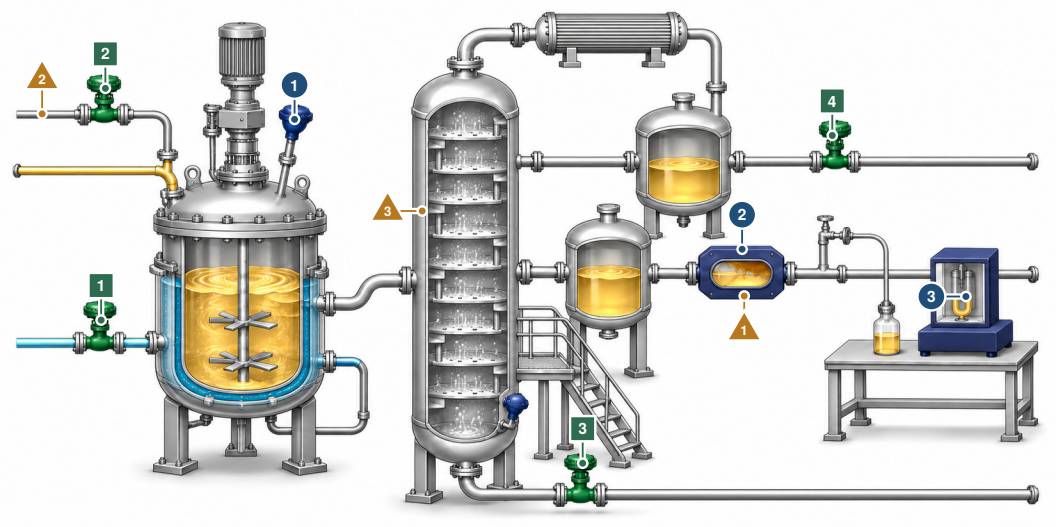

The controller moves 1 jacket water valve, which sets the reactor temperature (the slow handle) 2 catalyst dosing valve 3 bottom cut point, the fast handle (heavy product valve) 4 top cut point (light product valve)

The controller reads 1 reactor temperature 2 in-line viscometer on the motor oil 3 laboratory viscosity and viscosity index, every 75 minutes

What disturbs it 1 fouling that makes the viscometer read high 2 a more active catalyst lot 3 a column that loses cut sharpness

Notice the viscometer on the middle product line, which can foul and read high, and the laboratory bench that checks it only every 75 minutes. Illustration, not to scale.

IN PLAIN TERMS **Two Handles for One Oil**

The main idea. A synthetic motor oil is made by growing chains of molecules in many lengths and keeping the middle ones. The plant sets the thickness by changing which chains it keeps, and the slow reactor brings that choice back to the middle.

The challenge. Two routes to one thickness give oils that thin differently when hot, and only a slow laboratory test shows it. The thickness meter slowly gets dirty, so the controller can hold its reading while the oil goes wrong.

Why it matters. Failed oil sells for less, so the plant must decide how far to stay from a limit it cannot see.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	13 min	The first page, the history boxes in Section 10.2, the two try-it boxes and the fact boxes in Sections 10.5 and 10.8, and the key points
🕮	**Study**	45 min	The whole chapter with its pause points and the twelve exercises
⛰	**Practice**	20 min	The first page, Section 10.5, Section 10.6, Section 10.7 with its open challenge

The control question A reactor and a column cut point can both put a base oil on its viscosity, one in about an hour and one in a quarter of an hour, and the two routes make different oils. How should the two handles be arranged, how is a drifting meter corrected from a late laboratory, and how far should the plant stay from a viscosity index limit that only the laboratory can see?

Learning objectives

- Explain why two routes to the same viscosity give different viscosity index and yield, from the distribution and the cut.
- Derive the reactor and cut point gains from the model and work out the action of a mid-ranging loop whose gain runs through another controller.
- Predict the offset a laboratory bias update leaves behind a steadily drifting meter, for a given update gain.
- Recognize a loop in perfect control while the plant goes off specification, and name the layer that catches it.
- Price the back-off from a laboratory-measured constraint and compare classical mid-ranging with a predictive controller that prices the cut.

Prerequisites Sections 1.4.6 (the laboratory bias update), 1.6.2 (velocity-form PI and its action) and 1.6.4 (running a scenario). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/pao-lubricant/

See also Volume 4, Chapter 10 for mid-ranging on a power cycle and Volume 4, Chapter 9 for valve-position control in a trigeneration plant. Volume 5, Chapter 10 for a soft sensor corrected from a delayed laboratory and Volume 5, Chapter 3 for composition control through a slow analyzer.

10.1 Introduction

When one controlled variable has two manipulated variables, the textbook question is which to pair it with. On the plant modeled here the answer is both, and the interesting work is arranging them. Decene oligomerizes over a metallocene catalyst into a distribution of chain lengths, and a column cuts motor oil base stock out of the middle of it with a lighter grade off the top and a bright stock for gear oils and grease off the bottom. The viscosity can be moved by the reactor, which changes what the molecules are and takes most of an hour, or by the cut point, which changes only which tank each molecule goes into and takes a quarter of one.

Measured on the model, both reach 6.30 cSt from the design point and the oils are not the same. Cooling the reactor eleven degrees gives a viscosity index of 138.5 at 51.0 % yield, and opening the bottom cut by four carbons gives 134.8 at 60.9 %. Ten points of yield against nearly four of viscosity index, at an identical number on the certificate.

The misconception the simulator corrects is that hitting the specification is the job. The specification is one functional of a distribution and the customer buys three, so two routes to the same viscosity are two different products, and a viscosity controller on its own cannot say which one was made. The measured consequence is severe. A column that loses cut sharpness leaves nine percent dimer in the product, the viscosity controller corrects the viscosity to exactly 6.000 cSt by opening the bottom cut six carbons, and the viscosity index ends below specification with a hundred and fifty tonnes downgraded in a day. The loop is in perfect control throughout.

Two structural answers follow and both generalize. The first is mid-ranging: control with the fast handle and use a slow outer loop, whose measurement is the fast loop's output, to return that output to the middle of its range [1]. The second is inferential correction: the fast on-line viscometer fouls and reads high, and a laboratory result up to eighty-seven minutes old corrects it through a bias, which is the structure most polymer and refinery quality loops run on [2, 3]. Above both is a constraint that no continuous measurement can see, and the distance the plant keeps from it is set by how late the news is rather than by how good the control is.

This chapter works on an interactive base oil plant whose chain growth probability, cut sharpness and property correlations are hidden, so that the two handles must be compared by experiment, a

mid-ranging loop signed correctly, and a back-off chosen against a measured cliff. The app runs in a browser and exports its complete time history. Section 10.2 gives the domain history. Section 10.3 builds the model, a Schulz-Flory reactor with a two-cut column, three product specifications, a fouling viscometer and a laboratory, in the detail needed to rebuild it. Section 10.4 describes the implemented loops, the laboratory bias update, the viscosity index cap and an optional predictive controller. Section 10.5 reports the measured scenarios, among them a sign error that produces no oscillation and a controlled variable that stays perfect while the plant goes off specification. Section 10.6 names the real devices behind the loops, and Sections 10.7 and 10.8 state the limitations and the three results that carry over to other plants. The oligomerization chemistry, mid-ranging and bias updating are all established methods, and the chapter uses them as they stand.

10.2 Background

Polyalphaolefins are the base stock in most full synthetic motor oils. They are made by oligomerizing a linear alpha olefin, most commonly 1-decene, then hydrogenating the residual unsaturation and distilling the product into grades. The grades are named by their kinematic viscosity at 100 °C, so PAO 4, PAO 6 and PAO 8 are motor oil base stocks and the heavier cuts go into gear oils and grease [4]. What the customer pays for is not the viscosity alone. Viscosity index, the degree to which viscosity resists temperature, is the reason a synthetic oil exists at all, and it is measured on finished base stocks alongside the pour point [5].

The chemistry determines the control problem. Cationic oligomerization over boron trifluoride was the original route and gives a broad distribution. Metallocene catalysts, which arrived in the 1990s and have been developed for olefin oligomerization specifically [6], give a narrower one and higher viscosity index, and the chain length distribution is set by the competition between propagation and chain transfer. Since transfer carries the larger temperature dependence, a hotter reactor makes a lighter oil, and the whole product is described by one number, the probability that a growing chain adds another monomer.

> ⧖ HISTORY **Synthetic Oils for Jet Engines (1961 to 1973)**
>
> A patent application of 1961 sought synthetic lubricants for aircraft engines flying high, which needed a low pour point and a high viscosity index, and for jet engines, whose oil had to stay stable at about 500 degrees Fahrenheit and above. Alpha olefins polymerized over catalysts such as boron trifluoride already gave a low pour point and a high viscosity index, but hydrogenating them for stability raised the pour point sharply. The answer was to distill the dimer out before hydrogenation [7]. A 1971 application oligomerized the olefins at about 10 °C to 60 °C over boron trifluoride promoted by an alcohol, hydrogenated the product, and kept the low pour point without removing the dimer [8]. The dimer that troubled both is the light cut that the column of this chapter sends overhead.

Control of such a plant follows the polymer industry rather than the refining one. The controlled variable is a property of a distribution, it is measured in a laboratory on a cycle of an hour or more, and the practical answer is an on-line measurement corrected by the laboratory. Inferential estimation of polymer properties has a long literature and is now pursued with data-driven methods [3], and the maintenance and drift problems that come with any soft sensor are surveyed generally [2]. What is specific to a plant with a downstream fractionation is that the property can be moved after the reactor as well as inside it, and the two handles have different speeds and different costs.

The structure that arranges them is mid-ranging, also called valve position control, whose history the box at the end of this section tells. Volume 4, Chapter 10 adds a mid-ranging layer to a power cycle whose three actuators differ in speed. Its defining feature is that the outer loop measures the inner loop's output rather than a process variable, and its characteristic failure is a sign error, because the process gain seen by the outer loop runs through the inner controller. Above it, the economic trade between yield and viscosity index is what a real time optimizer or a multivariable controller with an

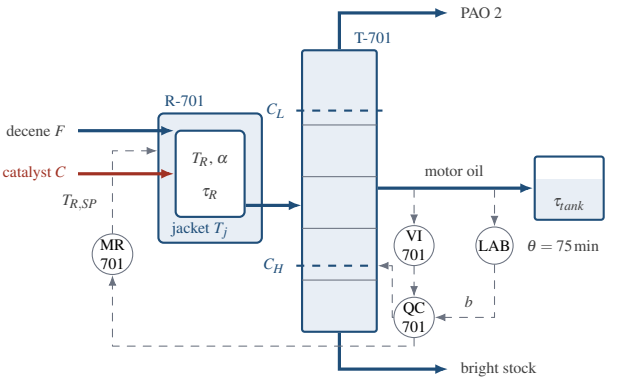

Figure 10.1 Synthetic base oil plant. Decene F and catalyst C react in R-701 at temperature T_R with residence time τ_R to a distribution described by the propagation probability α. Column T-701 cuts at C_L and C_H, sending PAO 2 overhead, motor oil base stock from the middle at viscosity KV, viscosity index VI and pour point PP, and bright stock from the bottom. The product tank has turnover τ_{tank}. VI-701 is the on-line viscometer, which drifts, and the laboratory reports every θ minutes on a sample drawn 12 min earlier. QC-701 holds the meter reading corrected by the laboratory bias b on the bottom cut point, and MR-701 returns that cut point to its target by moving the reactor setpoint $T_{R,SP}$, which a cap driven by the laboratory viscosity index (not drawn) can only lower.

economic objective handles [9], and the plantwide question of which unit owns which specification is the same one that arises on any train of units in series [10].

> ⧖ HISTORY **One Idea Under Many Names (1981 to 2005)**
>
> Mid-ranging has been in the control literature since the 1980s, traced to a publication of 1981, and has gone by names such as valve position control, input resetting, coordinated control and habituating control [11]. The problem is always the same: two inputs and one output, with the faster input in some way costlier to use than the slow one [12]. A 1998 study analyzed it on paper drying with steam and infrared heat [1, 12]. By 2003 most mid-ranging controllers in industrial use were valve position controllers, designed and tuned largely ad hoc, and a systematic method was shown on a pilot plant and on paper mill examples [13]. In 2005 the idea was described as common in the process industry and normally built from PID controllers [12].

10.3 Process Model

Reactor R-701 takes 12.0 t/h of 1-decene with 260 ppm of metallocene at 124.5 °C, giving a propagation probability of 0.633. Column T-701 cuts at 25 and 67 carbons, and the middle cut is 6.000 cSt at 100 °C with a viscosity index of 133.6, a pour point of −59.5 °C and a yield of 57.4 %. Figure 10.1 shows the arrangement.

10.3.1 The Plant Makes a Distribution

If α is the probability that a growing chain adds another decene rather than terminating, the mass fraction of the n-mer is Equation (10.1), and α is set by the ratio of chain transfer to propagation, Equation (10.2).

$$w_n \propto n\,(1-\alpha)\,\alpha^{n-2} \tag{10.1}$$

$$\alpha = \frac{1}{1+r}, \qquad r = r_0\, e^{E/R\,(1/T_{ref}-1/T_R)} \left(\frac{C}{C_0}\right)^m \tag{10.2}$$

The C90+ bin sums $n \geq 9$: its unnormalized mass is $\alpha^7[9+\alpha/(1-\alpha)]$. All bins are normalized by $(2-\alpha)/(1-\alpha)$. Chain transfer carries the larger activation energy, so hotter is lighter. Measured near

Table 10.1 Classification of Variables at the Design Operating Point

Symbol	Role	Nominal	Range
KV100	CV, inferred	6.000 cSt	5.70–6.30 spec
C_H	CV of MR-701	67.0 carbons	52–88
VI	constraint	133.6	≥ 130
PP	constraint	−59.5 °C	≤ -50
bright stock	constraint	30.4 cSt	≥ 33
C_H	MV, fast	67.0 carbons	52–88
T_R	MV, slow	124.5 °C	100–150
C_L	MV	25.0 carbons	18–34
C	MV	260 ppm	80–600
F	DV, measured	12.0 t/h	6–18
catalyst lot	DV, unmeasured	1×	0.5–1.8×
cut sharpness	DV, unmeasured	1×	0.6–3×
meter fouling	DV, unmeasured	off	0–4×

the design point, the viscosity moves −0.0281 cSt per degree of reactor temperature and 0.0016 cSt per part per million of catalyst removed. At $\alpha = 0.633$ the mass is 19.7 percent dimer, 18.7 percent trimer and 15.8 percent tetramer, with 11.7 percent in the lump of C90 and heavier.

10.3.2 The Column Decides Which Tank, Not What

Two cut points and a sharpness give each oligomer a fraction that stays in the product, Equation (10.3).

$$f_i = \sigma\left(\frac{C_i - C_L}{s}\right)\left[1 - \sigma\left(\frac{C_i - C_H}{s}\right)\right] \tag{10.3}$$

Measured over two carbons either side of the design cut, the viscosity moves 0.0522 cSt per carbon of bottom cut, nearly twice the reactor gain, and arrives in 14 min against the reactor's 38 plus 14. Nothing in Equation (10.3) changes a single w_n. It moves mass between three tanks, each with a specification of its own.

10.3.3 Three Functionals of One Shape

Viscosity blends on the Refutas index, Equation (10.4), and the two other specifications are correlations on the same cut and on the reactor temperature, Equations (10.5) and (10.6).

$$\nu = \sum_i x_i \,\mathrm{VBI}(\nu_i), \qquad \mathrm{VBI}(\nu) = a \ln\ln(\nu + 0.8) + b \tag{10.4}$$

$$\mathrm{VI} = c + d\ln\nu - e\,x_{dimer} - g\,[T_R - T_{iso}]^{+} \tag{10.5}$$

$$\mathrm{PP} = h + k\ln\nu + l\,x_{C70+} \tag{10.6}$$

The last term of Equation (10.5) is what stops the plant simply running hotter for yield: a hot metallocene isomerizes the chain as well as growing it, and the cost is 0.55 points of viscosity index per degree above 118 °C. It is also the only constraint on the unit that no continuous measurement can see.

10.3.4 The Tank

What is shipped is the tank, not the line, and the tank is a first-order filter with a 210 min turnover. Measured, a forty minute excursion reaches 6.297 cSt in the line and 6.049 cSt in the tank. The same excursion left for four hundred minutes takes the line to 6.315 cSt, above the band from 102 to 461 min with 41 t booked off specification, while the tank reaches only 6.266 cSt. The tank hides exactly the excursions that are easy to see and shows exactly the ones that are hard, because the things that move slowly are drifting instruments, fouling columns and aging catalyst.

Table 10.2 Model Parameters

Symbol	Meaning	Value	Class
r_0	transfer to propagation at T_{ref}, C_0	0.5152	adjustable
E/R	activation temperature of r	4100 K	adjustable
m	order of r in catalyst	0.38	adjustable
τ_R	reactor residence time	38 min	adjustable
τ_{col}	column response to a cut move	14 min	adjustable
τ_{tank}	product tank turnover	210 min	adjustable
s	cut sharpness	1.6 carbons	adjustable
g	VI lost per °C of reactor	0.55	adjustable
T_{iso}	isomerization onset	118 °C	adjustable
τ_{vis}	on-line viscometer response	3.0 min	adjustable
drift	viscometer fouling rate	0.020 cSt/h	adjustable
θ	laboratory cycle	75 min	adjustable
τ_{lab}	laboratory transport delay	12 min	adjustable
VI minimum	grade specification	130	visible

10.3.5 Assumptions

With the corrected C90+ sum, the nominal bright cut is 30.4 cSt, below its illustrative 33 cSt target; the challenge grades the motor oil, not that side product. The distribution is Schulz-Flory with a single propagation probability. A real metallocene gives a narrower distribution than that and the shape depends on the ligand [6].

Oligomers are lumped by carbon number with one viscosity each and everything from C90 upward in a single lump whose viscosity stands for a long tail.

The column is two logistic cut points with a sharpness parameter. A real column produces that shape from a tray count and a reflux ratio, and the sharpness would be a consequence rather than a knob.

Hydrogenation of the residual unsaturation is assumed complete and is not modeled.

Viscosity index and pour point are fitted correlations in the blend viscosity, the dimer content and the reactor temperature. Real structure-property relationships for these oils are considerably less tidy [5].

The reactor is one well mixed stage whose distribution follows the instantaneous one with a first-order lag, and the catalyst does not deactivate within a run.

10.4 Controller Design

10.4.1 Velocity Form PI

Every loop uses the velocity-form PID of Section 1.6.2, Equations (1.14) to (1.17), with the derivative time zero, so each of the three loops of this plant is a PI controller. The output bias u_b is zero, so the integral state starts at the actuator's initial value, and the tracking time of the back-calculation is $T_t = \tau_I$. The controller scan is $\Delta t = 0.1$ min and integration is fixed-step fourth-order Runge-Kutta at 0.02 min. The third PI loop, not drawn, holds the reactor at $T_{R,SP}$ with the jacket ($K_c = 3.0$, $\tau_I = 9$ min).

10.4.2 Mid-Ranging, and Its Sign

QC-701 holds the inferred viscosity on the bottom cut point. A wider cut keeps heavier molecules, so the viscosity rises with the cut and the loop takes the action $a = +1$, which is reverse acting in the convention of Section 1.6.2. Its gain is $K_c = 7.0$ carbons per cSt, its integral time $\tau_I = 30$ min and its output range 52 to 88 carbons.

MR-701 holds the cut point at its target by moving the reactor setpoint between 100 and 150 °C. The sign must be worked out rather than guessed, because the gain runs through QC-701. A hotter

reactor makes a lighter oil, so the viscosity controller answers a hotter reactor by opening the cut, and along the line of constant 6.00 cSt at the design point $\mathrm{d}C_H/\mathrm{d}T_R = +0.58$ carbons per degree. A cut point above its target therefore means the reactor is too hot and its setpoint must come down, which is a positive process gain and the action $a = +1$, a reverse acting controller in the convention of Section 1.6.2. The app's Mid-ranging action control uses the same convention: Reverse for this setting and Direct for the wrong sign. Its gain is 0.55 °C per carbon and its integral time is 180 min against the inner loop's 30, because two controllers acting on the same viscosity at similar speeds argue rather than cooperate [1]. What has to differ is closed-loop speed rather than integral time, and the reactor's hour of lag already slows MR-701, as Section 10.5 measures.

10.4.3 The Bias Update

The on-line viscometer answers in three minutes and fouls, reading about 0.02 cSt an hour high. The laboratory is exact and reports every seventy-five minutes on a sample drawn twelve minutes before, so the correction is a bias, Equation (10.7).

$$b \leftarrow b + \lambda \left(\nu_{lab} - \hat{\nu}\big|_{t-\tau_{lab}} \right), \qquad \hat{\nu} = \nu_{meter} + b \tag{10.7}$$

Because $\hat{\nu}$ carries the bias, Equation (10.7) is the laboratory update of Equation (1.11) with the raw meter in the place of the model, as long as the laboratory cycle is longer than the transport delay, so that no update falls between a sample and its report. Volume 5, Chapter 10 corrects a psychrometric soft sensor from a delayed laboratory in the same way. Two details decide whether such an update converges. The lab result describes a sample drawn $\tau_{lab} = 12$ min before it is reported, so it must be compared against what the controller believed then rather than against what it believes now, and against the corrected value rather than the raw meter, because comparing against the raw meter makes the corrections add up instead of converging. The gain $\lambda = 0.40$ is a compromise between injecting laboratory scatter into the product at $\lambda = 1$ and letting the drift outrun the correction at $\lambda = 0.1$ [2, 3].

🕮 **Pause.** Try Exercise 10.3 before reading on.

10.4.4 The Viscosity Index Cap

A one-directional override lowers the reactor setpoint when the laboratory viscosity index comes within 2.0 points of its minimum, to the operator's setpoint less 2.5 °C per point of shortfall. It can only lower, and it acts on the slow handle because the viscosity index is a property of what the reactor makes rather than of what the column keeps. The cap is computed before MR-701 runs and becomes MR-701's upper limit, so the back-calculation of Equation (1.17) unwinds the mid-ranging integral against it while it holds.

10.4.5 Model Predictive Control

An optional third structure, never the default, hands both handles to a quadratic program, set up as Section 1.6.3 describes for every app. Every $T_s = 2$ min from $t = 0$ it chooses the bottom cut C_H and the reactor setpoint request $T_{R,SP}$ over $N = 90$ steps (180 min) in twelve move blocks of $\{1,1,1,1,2,2,3,4,6,9,15,45\}$ steps, sends only the first move, which stands until the next plan [9, 14]. The jacket loop, the cap below the request and every alarm stay as they are, and the plan acts only with the temperature and viscosity loops in Auto.

The controller's model is not the plant of Section 10.3. It is that plant's balances, linearized a single time about the design point using the declared parameter values, so a parameter changed in a run is mismatch the controller is not told about, and it sees the plant only through the corrected meter, the filtered reactor temperature and the laboratory index. In deviations from the design point, marked

by a tilde, it is Equation (10.8).

$$\begin{aligned}
\tau_J \dot{\tilde{T}}_j &= -\tilde{T}_j + \tilde{I}_T + K_T\left(\tilde{T}_{R,SP} - \tilde{T}_m\right) \\
\tau_{I,T} \dot{\tilde{I}}_T &= K_T\left(\tilde{T}_{R,SP} - \tilde{T}_m\right) \\
\tau_T \dot{\tilde{T}}_R &= \tilde{T}_j - (1-h)\,\tilde{T}_R, \qquad \tau_m \dot{\tilde{T}}_m = \tilde{T}_R - \tilde{T}_m \\
\tau_R \dot{\tilde{\alpha}} &= a_T\,\tilde{T}_R - \tilde{\alpha} \\
\tau_{col} \dot{\widetilde{\mathrm{KV}}} &= k_\alpha \tilde{\alpha} + k_H \tilde{C}_H - \widetilde{\mathrm{KV}} \\
\tau_{col} \dot{\widetilde{\mathrm{VI}}} &= v_T \tilde{T}_R + v_\alpha \tilde{\alpha} + v_H \tilde{C}_H - \widetilde{\mathrm{VI}}
\end{aligned} \tag{10.8}$$

Here T_j is the jacket temperature with lag $\tau_J = 2.5$ min, I_T the integral of the temperature PI at its commissioned $K_T = 3.0$ and $\tau_{I,T} = 9$ min, $\tau_T = 7$ min the reactor's thermal lag, $h = -0.0319$ the slope of the reactor-to-jacket temperature difference that the reaction heat sustains, and T_m the measured temperature behind its $\tau_m = 1$ min filter. Central differences of Equations (10.1) to (10.5) give $a_T = -0.00602$ per °C, $k_\alpha = 4.660$ cSt, $k_H = 0.0485$ cSt per carbon, $v_T = -0.55$ per °C, $v_\alpha = 21.44$ and $v_H = 0.199$ per carbon. Two constant output biases b_{KV} and b_{VI} complete nine states, the predicted outputs being $\nu = \widetilde{\mathrm{KV}} + b_{KV}$ and $\mathrm{VI} = \widetilde{\mathrm{VI}} + b_{VI}$, and an exact zero-order hold over T_s discretizes the model. A degree of request moves the steady viscosity $k_\alpha a_T = -0.0281$ cSt, the reactor gain of Section 10.3.

In step tests on the plant with the temperature loop closed, QC-701 and MR-701 open and no noise, read 270 min after a step at 60 min, ±2 °C of request gives −0.0281 and –0.0279 cSt per degree against the model's −0.0280, with root-mean-square residuals under 0.0002 cSt. The cut gain is local. Two carbons up moves the viscosity 0.0676 cSt per carbon and two down 0.0368 cSt, against the model's 0.0485, with residuals of 0.037 and 0.023 cSt, and the mean of the two is the 0.0522 cSt secant gain of Section 10.3.

The estimator propagates the jacket and catalyst lags from the controller's own applied commands, the distribution from the measured reactor temperature through Equation (10.2), and the viscosity and index through the declared correlations and the column lag at the applied cuts. The viscosity bias filters, over 12 min, the difference between the corrected meter $\hat{\nu}$ of Equation (10.7) that QC-701 uses and the model's viscosity. The index bias moves by 0.4 of each laboratory residual, taken against the index the model held when the sample was drawn. The two biases make the prediction offset free [15], and no future event or target is shown to the controller.

Each plan minimizes Equation (10.9) subject to Equation (10.10).

$$\begin{aligned}
J = &\sum_{k=1}^{N} q\,(\nu_k - \nu_{SP})^2 + p\,(\nu_N - \nu_{SP})^2 + \rho_\nu \varepsilon_\nu^2 + \rho_{VI} \varepsilon_{VI}^2 \\
&+ \sum_{j=1}^{12} \left[s_C\,\Delta C_{H,j}^2 + s_T\,\Delta T_{R,SP,j}^2 \right] \\
&+ r \sum_{k=1}^{N} \left(C_{H,k} - C_H^{tgt} \right)^2
\end{aligned} \tag{10.9}$$

$$\begin{aligned}
&52 \le C_{H,j} \le 88, \qquad 100 \le T_{R,SP,j} \le 150 \\
&5.70 - \varepsilon_\nu \le \nu_k \le 6.30 + \varepsilon_\nu, \qquad \mathrm{VI}_k \ge 130 + g - \varepsilon_{VI}
\end{aligned} \tag{10.10}$$

Here $q = p = 100$ per cSt^2, $s_C = 0.1$ per carbon squared and $s_T = 0.2$ per degree squared weigh the block moves, the first measured from the command applied now, and $r = 0.02$ per carbon squared prices each step the cut spends away from its target C_H^{tgt}. That term is the whole of the mid-ranging. An industrial valve position controller returns the fast handle with a slow integrating loop [1, 13],

Table 10.3 Scenario Sequence

#	Scenario	Teaching point
1	The plant makes a shape	the CV is a functional
2	Two ways to make the same number	the two routes differ
3	The fast handle	speed and cost
4	Where the handle ends up	integrators keep corrections
5	Mid-ranging	a loop on a loop's output
6	The sign that runs through another controller	no oscillation, wrong place
7	The meter that drifts	believing an instrument
8	Fast and wrong, slow and right	the bias update
9	What the viscosity controller cannot see	perfect control, bad product
10	The constraint, and how far to stand back	what margin is worth
11	The tank remembers	filtering quality
12	Mid-ranging as a weight	a loop or a price
13	Challenge: a day on a drifting plant	graded run

Table 10.4 Two Routes to 6.30 cSt

Route	VI	Yield (%)	PP (°C)
Reactor to 113.2 °C	138.5	51.0	−58.6
Cut out to 70.8 carbons	134.8	60.9	−56.6

and the plan replaces that loop by a price, while the reactor's displacement costs nothing. The input limits are hard and are the clamps the PI loops use. The bands are soft, with nonnegative slacks priced at $\rho_v = 1000$ and $\rho_{VI} = 100$ and the cap's guard g (2.0 points by default). Each output has a single slack shared by every step of the horizon, so a plan with zero slack is a forecast and not a guarantee at every step. The program has 26 variables and is solved by a dual active-set method. With no constraint active its first move agrees with independently formed normal equations to 3×10^{-14}.

A plan is applied only if it is optimal and its moves, predictions, cost and slacks are finite and inside their bounds. Otherwise the applied commands are held for that scan, the run is marked, QC-701 and MR-701 take over, and the next valid plan takes back control. Both PI loops track the applied cut and request while the plan acts. A failure forced from 300 to 306 min in scenario 12 with the noise off moves neither handle at the rejection, and the largest single-step moves during the fallback are 7.5×10^{-5} carbon and 5.3×10^{-5} °C.

10.5 Results

Thirteen scenarios run from the distribution to a graded challenge. Table 10.3 lists them with the names the app shows, where the challenge carries no number, and Section 1.6.4 describes how a scenario is run and scored.

10.5.1 Two Routes to One Number

Table 10.4 reports both ways of reaching 6.30 cSt in scenario 2, read at 180 min for the reactor route and at 480 min for the cut point route. In Manual nothing acts on the reactor setpoint, so the scenario cools the reactor through its jacket, from 93.6 to 81.9 °C.

Ten percentage points of yield against nearly four points of viscosity index, at 6.300 cSt and 6.298 cSt. Neither is wrong. They are different products and the choice belongs to whoever is selling the oil, and what is wrong is not knowing which one the plant just made.

🕮 **Pause.** Try Exercise 10.1 before reading on.

Table 10.5 A More Active Catalyst Lot, at 600 min

Configuration	C_H	T_R (°C)	VI
Viscosity loop only	68.04	124.5	133.6
Mid-ranging, reverse acting	67.52	123.3	134.3
Mid-ranging, direct acting	69.14	127.7	131.8

TRY IT IN THE APP **2 · Two Ways to Make the Same Number**

Load the scenario and press Run. It runs in Manual, first cooling the reactor through its jacket, then returning it and opening the bottom cut instead. Watch the Viscosity index minimum / laboratory trend over about eight hours. Both moves take the oil to 6.30 cSt, and the index settles near 138.5 on the reactor route and 134.8 on the cut route.

10.5.2 Mid-Ranging and Its Sign

Table 10.5 reports a catalyst lot fifteen percent more active, with and without the outer loop, and with the outer loop signed the wrong way, all read 600 min into the run, just before the viscosity index cap first acts in the wrong-sign case. The action is named in the convention of Section 1.6.2, with the app's label of the Mid-ranging action control in parentheses.

Without the outer loop the cut point keeps its correction permanently, because an integrator moves error out of its measurement and into its output. With the outer loop correctly signed the cut is a third of a carbon from its target at 900 min (67.34) and within a sixth of a carbon (67.16) by 1400 min, with the reactor at 122.7 °C carrying the disturbance.

Pause. Try Exercise 10.2 before reading on.

Equal integral times do not make the two loops argue here, because the reactor's lag already slows MR-701. At $\tau_I = 30$ min the cut point undershoots by 0.08 carbon and stays within 0.02 of its target after 745 min, sooner than at 180. Only the 10 min slider minimum produces a cycle, about 380 min long, with the cut between 66.74 and 67.15 carbons and the reactor between 121.6 and 123.1 °C.

With it signed the wrong way the reactor walks 4.5 °C hotter than the correct sign puts it and the viscosity index sits 1.8 points from its minimum instead of 4.3. The viscosity ripples by 0.06 cSt and never leaves its band. What stops the walk is the cap, not the controller. Each time the laboratory index falls inside the cap's guard, on every third result from 600 min, the cap pulls the setpoint to about 124 °C, and the wrong sign walks it back. For the rest of the day the reactor cycles between 123.2 and 127.6 °C, the index in the line stays above 131.7 and nothing is downgraded. A sign error running through another controller does not announce itself, and here the protection does the controller's work.

DID YOU KNOW? **A Wrong Sign That Does Not Oscillate**

From the simulator: a wrong sign is usually pictured as an oscillation, but here it is a quiet walk. At 600 min (Table 10.5) the wrongly signed loop holds the reactor at 127.7 °C against 123.3 °C and the viscosity index at 131.8 against 134.3, while the viscosity never leaves its band.

10.5.3 Perfect Control and a Bad Product

Table 10.6 reports a column that loses cut sharpness at 60 min, with and without the two protective layers, after the 1400 min day the scenario runs.

The viscosity loop does its job perfectly and the plant makes a hundred and fifty tonnes of downgraded oil. The blurred cut takes the dimer in the product from 1.45 to 9.1 %, which pulls the viscosity down, and the controller puts it back by opening the bottom cut six carbons without removing one molecule of dimer. Viscosity index is not made of viscosity.

Table 10.6 A Column That Loses Cut Sharpness

Configuration	KV (cSt)	C_H	VI	Downgraded (t)
Viscosity loop alone	6.000	73.9	129.4	163.5
Mid-ranging and VI cap	6.004	67.7	133.5	5.4

Table 10.7 Yield Is Bought With Viscosity Index at Constant Viscosity

T_R (°C)	Yield (%)	VI
120	53.6	136.2
124.5	57.3	133.6
128	59.7	131.7
130	61.0	130.5

The mid-ranging loop does the protective work in this case. Armed alone it gives the same 5.4 t and 133.5, because returning the cut toward 67 cools the reactor to 115.3 °C, below the isomerization onset. The cap alone leaves 58 t downgraded. It is proportional to the laboratory shortfall and releases as soon as a laboratory result is back above its guard, so the reactor alternates between 124.5 and 117.5 °C on a cycle of about 300 min and the index in the line drops below 130 in every one.

> **TRY IT IN THE APP 9 · What the Viscosity Controller Cannot See**
>
> Run the scenario for a day. The Product viscosity trend stays on 6.000 cSt while the Bottom cut point opens to 73.9 carbons, the index falls to 129.4 and 163.5 t are booked off specification. Reset, switch on Viscosity index cap and Mid-ranging (cut → reactor) before the column fouls at 60 min, and run again: 5.4 t.

A controller in perfect control is not evidence that the plant is well. It is evidence that one measurement is where it was put, and the more tightly it is held the less that measurement can say.

10.5.4 The Drifting Meter

With the bias update off, a fouling viscometer leads the controller to hold the meter within 0.023 cSt of 6.00 cSt for the first 525 min while the real oil falls to 5.865 cSt. By then MR-701 has warmed the reactor to bring the falling cut point back, the laboratory index has dropped inside the cap's guard, and the cap starts pulling the setpoint down. From there the reactor moves between 116.2 and 129.8 °C and the meter stays within 0.073 cSt of its setpoint. At 1400 min the real oil is at 5.540 cSt, well outside the band, the viscosity index is 129.9, and 40 t are downgraded with an indication that never looked wrong.

With the update on, the true viscosity holds at 5.945 cSt and nothing is downgraded. The remaining five hundredths is a ramp-following offset, because the meter drifts continuously and the correction arrives in steps. The scenario has no laboratory scatter, so $\lambda = 1$ does better, at 5.984 cSt, while $\lambda = 0.1$ lets the oil sag to 5.77 cSt by 1350 min, inside the band, with nothing downgraded.

10.5.5 The Back-Off

Table 10.7 reports the line of constant 6.00 cSt viscosity with the cut point moved to match each reactor temperature.

The specification minimum is 130, so there is a cliff at about 131 °C. The arithmetic either side of it is wildly asymmetric. A percentage point of yield moves roughly 0.12 t/h from the bright stock tank into the motor oil tank and is worth about sixty dollars an hour. Crossing the viscosity index specification downgrades the entire motor oil production, about seven tonnes an hour at fifteen hundred dollars of lost value, which is about eleven thousand dollars an hour. The gradient toward the constraint is tens of dollars per degree and the cliff is thousands.



Figure 10.2 The synthetic base oil application at $t = 1400$ min of scenario 9, What the viscosity controller cannot see, with the viscosity loop in automatic, the mid-ranging loop and the viscosity index cap off, and the column cut sharpness tripled at 60 min, cropped to the plant view, its status chips and its alarm strip. (1) The status chips: the viscosity is 6.000 cSt in the line and in the tank, the cut has opened to 73.9 carbons, and 163.5 t are booked off specification. (2) The reactor and what it is making. (3) The chain length distribution drawn as bars by carbon number with the two cut points as dashed lines across it, so that moving a cut changes the color of the bars and not their height. Below it are the three product tanks with their rates and their own viscosities. (4) The meter, the bias, the laboratory result and the viscosity index, 129.37 against its minimum of 130. (5) The inset plots viscosity index against reactor temperature along the line of constant viscosity with the specification minimum marked.

The cliff is also invisible. Viscosity index is a laboratory number that is twelve minutes old when it arrives and eighty-seven minutes old when the next one replaces it, on a plant whose reactor takes thirty-eight minutes to respond, so anything done now is committed for about two hours before the plant can be told it was a mistake. That is what sets the back-off, and the number comes from the measurement rather than from the tuning.

10.5.6 The Graded Challenge

The challenge runs twenty hours with a catalyst lot change at 90 min, a fouling viscometer from 300 min, a rate increase to 13.5 t/h at 520 min and a loss of cut sharpness at 700 min, under viscometer noise and laboratory scatter. It starts with cautious tuning, three of the five structural decisions switched off and the reactor at 128 °C. It is met when the viscosity in the line stays within 0.10 cSt of target from 792 min, the last third of the run, the viscosity index in the line stays above 130 throughout, and less than five tonnes are downgraded. As loaded the run fails all three parts. The worst viscosity error from 792 min is 0.299 cSt, the lowest index 127.9 and 63.7 t are downgraded. With the suggested configuration, which also starts the reactor at 124.5 °C, the worst error is 0.084 cSt, the lowest index 131.2 and nothing is downgraded, so the target is met by 0.016 cSt on the viscosity. The app gives the same run whether the configuration is revealed before Run or carried through Reset. The noise seed is fixed, and 11 other seeds give 0.077 to 0.093 cSt, all passing. No part of the configuration is spare. Leaving out any one of the bias update, mid-ranging, the cap, the 124.5 °C

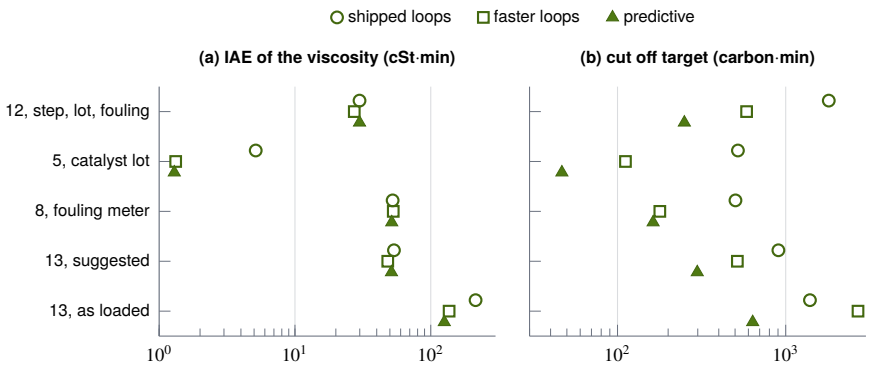

Figure 10.3 Table 10.8 drawn on logarithmic axes, one row per case. (a) The integral of the absolute viscosity error. (b) The time the cut spends away from its target, in carbon·min. Notice that the predictive controller (filled triangles) holds the cut closest to its target in every case, while on viscosity error it trails the faster loops (open squares) in two of the five cases.

start or the faster tuning fails, and the structural changes without the faster tuning leave a 0.196 cSt error.

10.5.7 Comparison With Model Predictive Control

Table 10.8 runs three structures for 1200 min on each case: the shipped loops at their shipped tuning, the same loops faster (QC-701 at $K_c = 14.0$ carbons per cSt and $\tau_I = 14$ min, MR-701 at $\tau_I = 30$ min), and the plan with the settings of scenario 12. Each triple shares its initial state, seed, instruments, noise and events, and neither structure is shown a disturbance or target in advance. The faster tuning is one stated setting, not an optimized bound on classical control. Scenarios 5 and 8 are the catalyst lot and the fouling meter above, run here to 1200 min, and the challenge runs as loaded and with its suggested configuration. Figure 10.3 draws the viscosity error and the cut's distance from its target of every row.

Scenario 12 raises the viscosity target from 6.00 to 6.20 cSt at 60 min, brings a catalyst lot fifteen percent more active at 360 min and starts the meter fouling at 600 min, without noise and with every structural element of the loops armed at the shipped tuning. The plan answers the step with both handles at once. It opens the cut to 73.12 carbons by 64.1 min and takes the request down to 115.18 °C by 84.1 min, so the viscosity reaches 6.18 cSt 7.8 min after the step against 110 min under the loops, and then overshoots to 6.254 cSt at 78 min, 27 percent of the step. From 116 min the cut stays within 0.2 carbon of its target to the catalyst lot. Under the loops the cut is still 2.04 carbons out at 240 min, with the request at 122.31 °C, as shown in Figure 10.4.

Over the whole 1200 min the viscosity error is a tie, 29.889 against 29.886 cSt · min, and the faster loops beat both at 27.341. Divided at the events, the plan wins the step (4.63 against 9.24 from 60 to 360 min, 2.50 faster) and loses the fouling (23.67 against 19.13 from 600 to 1200 min, 24.04 faster). Between 64 and 88 percent of each total accrues after the meter starts fouling, where all three end 0.052 to 0.060 cSt low, the ramp-following offset of the shared bias update. With half the integration step the totals are 29.968, 29.893 and 27.321, and with the challenge's instrument noise on eight other seeds the plan is below the shipped loops on five, with means of 29.71, 30.49 and 29.34, so the faster loops lead and the other two are not separated.

What the price buys is centering. The cut spends 249.7 carbon·min away from its target against 1803.7 and 585.5, and 275 to 312 against 1766 to 1832 under noise. The plant pays with the largest excursion, 6.12 carbons against 2.16 and 3.20, and with three and almost four times the shipped loops' travel, 18.65 carbons of cut and 21.85 °C of request. One likely cause of the excursion and the

Table 10.8 Shipped Loops, Faster Loops and Model Predictive Control Over 1200 min, Seed 20260923

Case	Structure	IAE KV	Cut off target	Peak cut off	Cut travel	Request travel	Down-graded	Lowest VI	Plans, fallbacks
12, target step, catalyst lot, fouling meter	shipped loops	29.886	1803.7	2.16	6.10	5.90	0	133.63	—
	faster loops	27.341	585.5	3.20	12.06	13.32	0	133.63	—
	predictive	29.889	249.7	6.12	18.65	21.85	0	133.62	600, 0
5, catalyst lot	shipped loops	5.141	519.9	0.72	1.24	1.71	0	133.57	—
	faster loops	1.323	111.5	0.70	1.41	2.05	0	133.60	—
	predictive	1.292	46.7	0.50	1.02	2.09	0	133.58	600, 0
8, fouling meter	shipped loops	52.396	502.7	0.76	4.78	4.12	0	132.45	—
	faster loops	52.946	178.5	0.51	10.44	8.15	0	132.27	—
	predictive	51.611	162.8	0.27	12.29	14.67	0	132.22	600, 0
13, suggested	shipped loops	53.576	904.8	1.81	21.07	13.99	0	131.17	—
	faster loops	48.240	515.9	2.96	43.16	26.04	0	131.53	—
	predictive	51.361	298.0	2.60	21.15	26.20	0	131.26	600, 0
13, as loaded	shipped loops	213.852	1393.7	1.83	5.53	0.00	67.28	127.90	—
	faster loops	136.523	2690.9	4.33	39.95	0.00	68.30	127.95	—
	predictive	125.788	636.4	9.04	39.71	31.07	3.80	129.44	600, 0

Shipped and faster are the PI loops at the two tunings of the text, and predictive is the plan. IAE is the integral of the absolute difference between the viscosity target and the true viscosity in the line, in cSt · min, the definition of the application's scorecard. Cut off target integrates $|C_H - C_H^{tgt}|$, in carbon·min, and peak cut off is its largest value, in carbons. Travel sums the absolute changes of the logged cut, in carbons, and request, in °C. Downgraded oil is in tonnes. The challenge's worst viscosity error in its last third is 0.084, 0.067 and 0.074 cSt with the suggested configuration, all passing, and 0.299, 0.291 and 0.251 cSt as loaded, all failing. A plan took 0.31 to 0.38 ms on average in Node.js on the machine that ran the table.

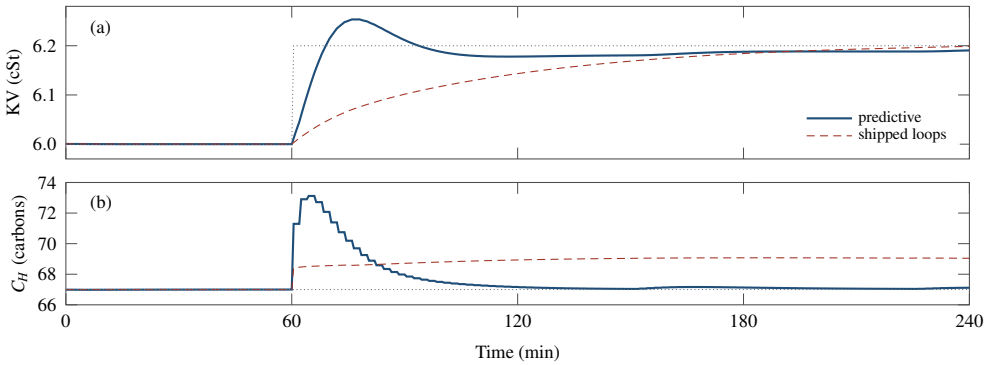

Figure 10.4 Viscosity and bottom cut in the first 240 min of scenario 12, with the plan (thick solid line) and with the shipped loops (dashed) in control. (a) True viscosity in the line against the target (dotted), which steps from 6.00 to 6.20 cSt at 60 min. (b) Bottom cut point against its 67 carbon target (dotted).

overshoot is the local cut gain, since the model understates the plant's answer to two carbons up by 28 percent.

Taking the price away shows what it is for. With $r = 0$ and every other setting of scenario 12 unchanged, the 1200 min run gives the lowest viscosity error of the three, 28.7 against 29.9 with the price and 29.9 under the shipped loops, and the least request travel, 5.3 °C against 21.8 and 5.9 °C. The cut then spends 2564 carbon·min away from its target against 250 with the price and 1804 under the loops. Without the price nothing returns the fast handle, so the plan leaves each correction in the cut, as the viscosity loop alone does in scenario 4, and the range the cut has used is gone when the

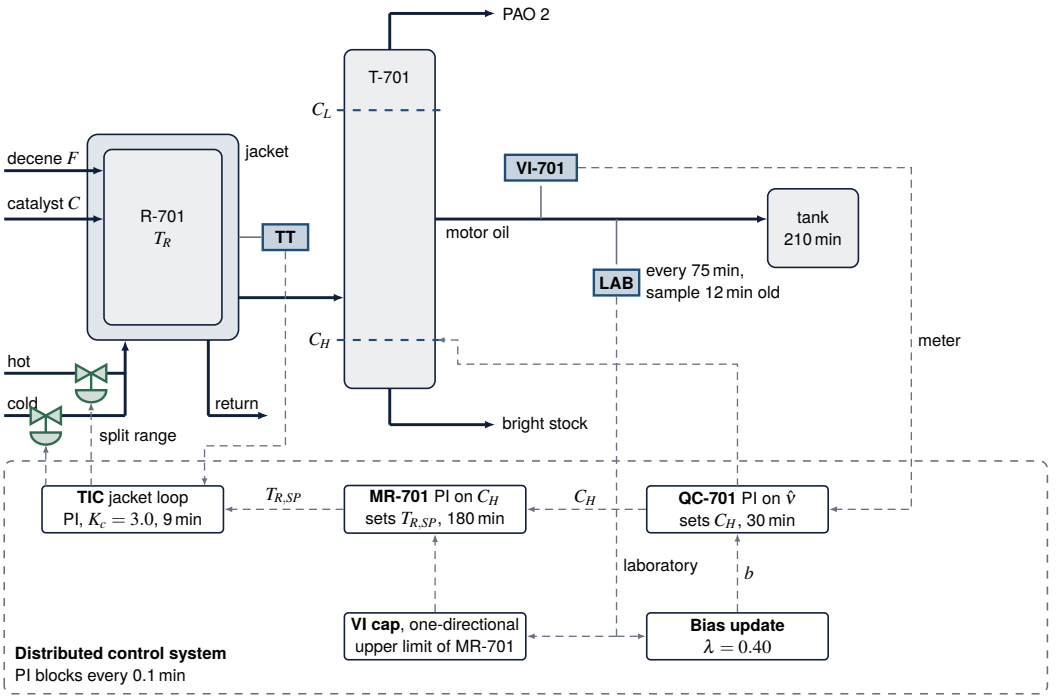

Figure 10.5 The loops of the base oil plant drawn with their hardware. The viscometer VI-701 on the motor oil line and the laboratory feed QC-701 through the bias update, QC-701 sets the bottom cut C_H, MR-701 reads the cut and moves the reactor setpoint under the viscosity index cap, and the jacket loop holds the reactor through split-ranged hot and cold water valves. Notice that MR-701 has no transmitter of its own: its measurement is another block's output. The split-ranged valves are one real form of the jacket the model lumps into a lag. The top cut, the catalyst and the optional predictive controller are left out.

next disturbance asks for it.

On the catalyst lot of scenario 5 the plan has the lowest error, 1.292 against 5.141 and 1.323, and the cut closest to its target, 46.7 against 519.9 and 111.5 carbon·min, although the faster loops alone come within 2.4 percent of its error. On the fouling meter of scenario 8 the errors lie within 2.6 percent of one another and the plan moves most, 12.29 carbons and 14.67 °C against 4.78 and 4.12. With the suggested configuration of the challenge all three pass.

The challenge as loaded is not like-for-like, and the plan fails it. The loaded loops have cautious gains with mid-ranging, the bias update and the cap off, while the plan prices the cut's distance and holds a soft index floor, so it supplies structure they lack. It also reads the uncorrected meter and predicts a temperature loop at 3.0 and 9 min that runs at 1.0 and 25 min. It downgrades 3.80 t against 67.28, but the index falls to 129.44 and the viscosity is 0.251 cSt off in the last third, failing two of three parts with an index slack of 0.50. The faster gains on the loaded structure still downgrade 68.30 t and the loops with the suggested structure none, so the reduction measures the missing structure rather than prediction. No plan was rejected in any run, nor in any of the thirteen scenarios run 1200 min under the plan.

10.6 Technology in Practice

The loops of Section 10.4 run on real instruments, valves and a control system, drawn together in Figure 10.5. This section names them and says what each one does in service.

What Measures It

The viscosity has two measurements of different kinds. The on-line viscometer VI-701 reads the oil continuously in the line. The laboratory measures the kinematic viscosity by ASTM D445 and calculates the viscosity index by ASTM D2270 from the kinematic viscosities at 40 and 100 °C, and the pour point has a test method of its own [16]. That is why the index exists in this chapter only as a laboratory number.

> **TECHNOLOGY** **Sensor: An In-Line Viscometer**
>
> An in-line viscometer measures the oil without drawing a sample. A rotational type turns a spindle and reads the shear stress, a pipe type reads the pressure drop along a straight pipe, and a vibrating type infers viscosity and density from the amplitude and frequency of a vibrating fork and sends both on a 4 mA to 20 mA signal [17]. A vibrating element responds to the product of density and viscosity [18], so the kinematic viscosity of a specification needs the density as well [19]. In field trials a rotational unit plugged with solids and read consistently above the standard bench instrument, and a vibrating unit was checked against manual readings at the start of every shift [17]. No range or accuracy is given here because no open, vendor-neutral source states one.

In the model the meter is a 3.0 min lag that fouls and reads 0.020 cSt/h high, and the laboratory reports every 75 min on a sample drawn 12 min earlier, exact apart from a repeatability term.

What Moves It

The fast handle is the bottom cut point. The model treats it as a number of carbons that the product follows in 14 min and does not name the column hardware that moves it, so this section describes the slow handle's hardware. The reactor temperature loop drives the jacket, which the model holds as a 2.5 min lag between 40 and 150 °C.

> **TECHNOLOGY** **Final Element: Split-Ranged Jacket Valves**
>
> A reactor that must be heated or cooled can take hot or cold water into one jacket through two valves driven by one controller. With exclusive sequencing both valves are shut at half output, the cooling valve opens over the lower half of the signal and the heating valve over the upper half, so water flows through one or the other but never both. Each valve fails in the position safest for the process: in the textbook example the heating valve fails closed and the cooling valve fails open, so a lost signal or lost instrument air gives full cooling [19]. On this plant full cooling makes a heavier oil, and QC-701 would answer by closing the bottom cut.

Where the Controller Runs

A distributed control system holds this structure. Its loops may be nested so that the setpoint of one comes from another, at cycle times from milliseconds to minutes [20]. A valve position controller is normally a PID block and can be built on any such system [12, 13]. In this chapter every PI loop runs every 0.1 min, the bias moves on each laboratory report, and the optional predictive controller plans every 2 min. MR-701 reads QC-701's output rather than a transmitter, and the cap enters as MR-701's upper limit before it runs.

> **IN PRACTICE** **Tuning the Slow Loop of a Mid-Ranging Pair**
>
> Most mid-ranging controllers in plants are valve position controllers, and their tuning has been largely ad hoc [13]. A published rule of thumb keeps the response of the fast input an order of magnitude faster than that of the slow one, to avoid exciting cross-coupling and driving the slow input into its limits [12]. The slow controller is traditionally integral-only and tuned slow. Its target need not be mid-range: on a paper mill feed line, a valve in series with a variable-speed pump is held near 90 percent open, to save pump energy yet keep room to reject disturbances [11].

What Is Modernizing It

Three developments bear on this loop. The first is on-line measurement: the demand to measure and manage viscosity on line is growing [18], and where properties are mainly measured in a laboratory, real-time instruments are sought because the laboratory hinders real-time optimization and decisions [17]. The second is the soft sensor. A data-driven estimate must be maintained and tuned on a regular basis, and most still have no automated mechanism for it [2], which is the work the bias update of Equation (10.7) does on one number. The third is model predictive control, with a solid foundation in refining and petrochemicals [9]. A multivariable controller can produce the mid-ranging effect directly [12], and the plan of Section 10.5.7 does so with a price on the cut's distance from its target.

10.7 Limitations

The distribution is Schulz-Flory with one parameter, which is broader than a real metallocene product [6], and the oligomers are lumped by carbon number with a single tail lump.

The column is two logistic cut points rather than a tray model, so the cut sharpness is a disturbance rather than a consequence of reflux, and the application cannot show what a column operator would actually do about it.

Viscosity index and pour point are fitted correlations. A reader would not learn here how either is measured or why the correlations are difficult, and the absolute values carry less meaning than their changes [5].

Hydrogenation is absent and the catalyst does not deactivate, so a campaign has no history and the slow drifts a real plant manages are represented only by the explicit disturbance sliders.

The laboratory is exact apart from a repeatability term. A real laboratory has systematic differences between analysts and between instruments, which is the harder half of maintaining a bias update [2].

The predictive controller keeps its design-point linearization and the commissioned temperature tuning whatever the plant does. The plant moves 39 percent more than its cut gain for two carbons up and 24 percent less for two down, and a catalyst lot, a loss of cut sharpness or a retuned temperature loop reaches it only through the two output biases. It omits the 3 min viscometer lag, the jacket's 40 to 150 °C limits and the cap, and its objective leaves out the tank, the pour point and the side products. Its index is corrected by laboratory results up to 87 min old. On its own scenario it ties the shipped loops and loses to a faster tuning, 29.889 against 27.341 cSt · min.

Finally, the economic trade between yield and viscosity index is left to the reader rather than to an optimizer [9], and the plantwide question of which stream carries which specification is posed rather than answered [10].

> **OPEN CHALLENGE** **Tuning Two Handles of Different Speed**
>
> Tuning the slow controller of a mid-ranging pair has long rested on trial and error. A 2003 study named the lack of a systematic procedure, a 2013 application found the tuning not straightforward, and a 2025 method that turns the pair into one PI loop covers only inputs with identical dynamics and different gains [11]. Anti-windup for the moment the fast input saturates was said to need further attention in 2005 [12]. The reactor and the cut of this chapter differ in speed, not only in gain. Progress would be a rule that takes the speed, the sign and the limits of both handles and returns settings a plant can commission without hours of field tests.

10.8 Discussion

Three results generalize past base oils.

The first is the pairing rule for one controlled variable with two manipulated variables of different speeds. Control with the fast one, because control needs speed, and position it with the slow one, because the slow one is the one that costs something. The outer loop's measurement is then another controller's output, which is the defining feature of mid-ranging and the source of its characteristic

sign error [1]. Working that sign out on paper is worth doing, because the field test is slow and the primary measurement never moves.

> ✦ DID YOU KNOW? **The Same Loop in a CD Player**
>
> Mid-ranging is not confined to process plants. A 2005 study names it as a key idea in positioning the pickup head of a CD player [12], and a 2025 paper lists uses from paper machines, separators and bioreactors to drug delivery and cryogenic systems [11].

The second is that tight control of one variable can conceal the failure of another. The measured case here is a viscosity held at exactly 6.000 cSt while a hundred and fifty tonnes are downgraded on viscosity index in a day. The defense is not a better viscosity loop. It is to watch the manipulated variable, because the cut point is three carbons from where it lives 29 min after the column fouls, an hour before the laboratory reports the first low viscosity index at 150 min.

The third is that the distance a plant keeps from a constraint is a design variable with a price on it, and the price is set by the age of the measurement. With laboratory results up to eighty-seven minutes old and a thirty-eight minute reactor, this plant is committed for two hours before it can be told it was wrong, and the back-off has to cover that. The value of a faster measurement is exactly the back-off it removes, which is a number an engineer can compute and put in a proposal.

10.9 Conclusion

A base oil plant makes a distribution, and the specification is three functionals of it. In the simulated plant of this chapter both the reactor and the column cut point reach 6.30 cSt from the design point, and the two oils differ by 3.7 points of viscosity index and 9.9 points of yield. The cut point moves viscosity by 0.0522 cSt per carbon in fourteen minutes and the reactor by −0.0281 cSt per degree in fifty-two, so the cut point controls and the reactor positions it. Mid-ranging returns the cut to within a sixth of a carbon of its target, and the same loop signed the wrong way walks the reactor 4.5 °C hotter in ten hours, with the viscosity trend showing nothing, until the viscosity index cap catches it. A column that loses cut sharpness is corrected to exactly 6.000 cSt by opening the cut six carbons and downgrades 164 t on viscosity index in a day, against 5 t with the mid-ranging loop and the viscosity index cap armed. A fouling viscometer takes the true oil to 5.540 cSt without a bias update and holds 5.945 cSt with one. A predictive controller that does the mid-ranging with a price ties the loops on viscosity error (29.889 against 29.886 cSt · min) and keeps the cut 249.7 rather than 1803.7 carbon·min from its target, at three to four times the movement of both handles. Further work would add a tray model so that cut sharpness becomes a consequence rather than a disturbance, a moment model for the distribution breadth, catalyst deactivation over a campaign, a laboratory with systematic as well as random error, and an optimizer for the yield against viscosity index trade. The base oil plant is a browser platform for learning what a second handle is for.

KEY POINTS

- A base oil plant makes a distribution, and its specification is three properties of it.
- The cut acts in fourteen minutes and the reactor in fifty-two, so the cut controls and the reactor positions it.
- Signed the wrong way, mid-ranging walks the reactor 4.5 °C hotter with a flat viscosity trend.
- A viscosity held at exactly 6.000 cSt can hide 164 t downgraded on viscosity index in a day.
- The back-off from a laboratory-measured limit is set by the age of the measurement, not the tuning.

Where else this applies. Mid-ranging in Volume 4, Chapter 10, and a laboratory-corrected soft sensor in Volume 5, Chapter 10.

References

[1] 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.

[2] 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.

[3] C. Zhu and J. Zhang, "Inferential estimation of polymer melt index using deep belief networks," in *2018 24th International Conference on Automation and Computing (ICAC)*. IEEE, 2018, pp. 1–6.

[4] K. Hope, "Polyalphaolefin lubricant applications," in *Friction, Lubrication, and Wear Technology*. ASM International, 2017, pp. 207–212.

[5] M. Warren, A.-F. Asfour, and J. Gao, "Viscometric behaviour of viscosity-index improvers in lubricant base oil over a wide temperature range: Ii. polyalphaolefin synthetic base oil," *Journal of Synthetic Lubrication*, vol. 22, no. 3, pp. 249–258, 2005.

[6] C. Janiak and F. Blank, "Metallocene catalysts for olefin oligomerization," *Macromolecular Symposia*, vol. 236, no. 1, pp. 14–22, 2006.

[7] L. A. Hamilton and F. M. Seger, "Polymerized olefin synthetic lubricants," U.S. Patent 3,149,178, filed 11 July 1961, granted 15 September 1964, 1964.

[8] R. L. Shubkin, "Synthetic lubricants by oligomerization and hydrogenation," U.S. Patent 3,780,128, filed 3 November 1971, granted 18 December 1973, 1973.

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

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

[11] J. D. F. d. O. Luna, D. O. Machado, and J. E. Normey-Rico, "A SISO equivalent mid-ranging control fast tuning method for a class of TISO systems," in *Anais do XVII Simpósio Brasileiro de Automação Inteligente (SBAI 2025)*, 2025.

[12] S. Haugwitz, M. Karlsson, S. Velut, and P. Hagander, "Anti-windup in mid-ranging control," in *Proceedings of the 44th IEEE Conference on Decision and Control and the European Control Conference*, Seville, Spain, 2005, pp. 7570–7575.

[13] B. J. Allison and S. Ogawa, "Design and tuning of valve position controllers with industrial applications," *Transactions of the Institute of Measurement and Control*, vol. 25, no. 1, pp. 3–16, 2003.

[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] Q. Hu, K. Zeng, S. Han, J. Xu, W. Hu, and J. Li, "Synthesis and performance evaluation of metallocene polyalphaolefins (mPAO) base oil with anti-friction and anti-wear properties," *Polymers*, vol. 16, no. 19, p. 2828, 2024.

[17] N. Liu, D. Zhang, H. Gao, Y. Hu, and L. Duan, "Real-time measurement of drilling fluid rheological properties: A review," *Sensors*, vol. 21, no. 11, p. 3592, 2021.

[18] T. Yanaseko, H. Sato, I. Kuboki, K. Mossi, and H. Asanuma, "Vibration viscosity sensor for engine oil monitoring using metal matrix piezoelectric composite," *Materials*, vol. 12, no. 20, p. 3415, 2019.

[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.

10.10 Exercises

10.1 Concept check · 2 min

From the design point, the motor oil cut is moved from 6.00 to 6.30 cSt twice: once by cooling the reactor and once by opening the bottom cut. Which statement matches the two oils the model makes?

(a) The reactor route has the higher viscosity index and the lower yield
(b) The cut route has the higher viscosity index and the higher yield
(c) Both routes give the same viscosity index, because it follows the viscosity
(d) The reactor route has the lower viscosity index and the higher yield

10.2 Concept check · 1 min

In scenario 3, the viscosity loop alone absorbs a catalyst lot fifteen percent more active and returns the viscosity to 6.000 cSt. Mid-ranging is off. Where is the disturbance fifteen hours later?

(a) In the reactor temperature, which the viscosity loop has lowered by about a degree
(b) In the bottom cut point, which stays about a carbon above its 67-carbon target
(c) In the viscosity, which settles a few hundredths of a centistoke below target
(d) Nowhere, because the integral action has removed it from the plant

10.3 Concept check · 2 min

A laboratory viscosity is reported at time t for a sample drawn at $t - 12\,\text{min}$. The on-line meter drifts steadily. Against which value must the bias update compare the laboratory result so that the bias converges to the meter's own error, whatever the plant does in those twelve minutes?

(a) The raw meter reading at t, when the result arrives
(b) The raw meter reading at $t - 12\,\text{min}$, when the sample was drawn
(c) The corrected meter value the controller held at $t - 12\,\text{min}$
(d) The corrected meter value the controller holds at t

10.4 Theory and calculation · 4 min

Equation (10.1) gives the mass fraction of the n-mer. At the design propagation probability $\alpha = 0.633$ the model puts 19.7 % of the mass in dimer.

(a) Show that the ratio of the mass fractions of two successive oligomers is $w_{n+1}/w_n = \alpha\,(n+1)/n$. *Checkpoint:* $w_{n+1}/w_n = \alpha(n+1)/n$

(b) From the dimer fraction, compute the trimer and tetramer mass fractions to one decimal place (the key accepts ± 0.1 point). *Checkpoint:* 18.7 % and 15.8 %

10.5 Theory and calculation · 6 min

At the design point the reactor runs at 124.5 °C with $\alpha = 0.633$ and $E/R = 4100\,\text{K}$ in Equation (10.2). At fixed cuts the steady viscosity moves $k_\alpha = 4.660\,\text{cSt}$ per unit of α, and the local cut gain is $k_H = 0.0485\,\text{cSt}$ per carbon.

(a) Differentiate Equation (10.2) with respect to T_R at fixed catalyst and evaluate $a_T = \mathrm{d}\alpha/\mathrm{d}T_R$ at the design point, to three significant figures. *Checkpoint:* $a_T = -0.00602$ per °C

(b) Compute the steady reactor gain $k_\alpha a_T$ and the reactor cooling it predicts for a rise of 0.30 cSt, to one decimal place. *Checkpoint:* −0.0281 cSt per °C, about 10.7 °C

(c) With QC-701 holding the viscosity, find the slope $\mathrm{d}C_H/\mathrm{d}T_R$ along the line of constant viscosity, to two decimal places. State the action a MR-701 needs and name it in the convention of Section 1.6.2. *Checkpoint:* +0.58 carbons per °C, $a = +1$, reverse acting

10.6 Theory and calculation · 8 min

The viscometer drifts high at $\rho = 0.020\,\text{cSt/h}$. The laboratory is exact, reports every $\theta = 75\,\text{min}$ and each result describes a sample drawn $\tau_{lab} = 12\,\text{min}$ earlier. QC-701 has integral action, so over a laboratory cycle it holds the average of the corrected meter $\hat{v} = v + D + b$ at the setpoint, where v is the true viscosity and D the accumulated drift.

(a) Let $e_k = D(t_{s,k}) + b_{k-1}$ be the error of the corrected meter at the k-th sampling instant. Use Equation (10.7) to show that $e_{k+1} = (1-\lambda)e_k + \rho\theta$, and find the settled e for $\lambda = 0.40$ to three decimal places. *Checkpoint:* $e = \rho\theta/\lambda = 0.0625\,\text{cSt}$

(b) Show that the average of $D + b$ over a cycle is $(1-\lambda)e + \rho(\tau_{lab} + \theta/2)$, and predict the average true viscosity for a setpoint of 6.00 cSt at $\lambda = 0.40$ and at $\lambda = 1$, to three decimal places. *Checkpoint:* 5.946 cSt and 5.984 cSt

10.7 App challenge · 8 min

Load scenario 2, Two ways to make the same number, as shipped (Manual, default seed). Run it to at least 480 min at the default speed of 20 min per second, then export the CSV, whose time column is in hours.

(a) In the row at 3 h (180 min, the reactor route), read `kv_cst` to three decimal places and `vi_true`, `yield_pct` and `tr_C` to one (key ±0.1). *Checkpoint:* 6.300 cSt, VI 138.5, 51.0 %, 113.2 °C

(b) In the row at 8 h (480 min, the cut route), read the same three properties. *Checkpoint:* 6.298 cSt, VI 134.8, 60.9 %

(c) Explain in two sentences which term of Equation (10.5) makes most of the difference in viscosity index, and why the cut route gains yield. *Checkpoint:* the isomerization term, the cut moves mass between tanks

10.8 App challenge · 8 min

Run scenario 5, Mid-ranging, and scenario 6, The sign that runs through another controller, each as shipped at the default seed. Scenario 6 differs from scenario 5 only in the Mid-ranging action control, which it sets to Direct. Run each to 1440 min at 60 min per second and export the CSV (time in hours).

(a) In the row at 10 h (600 min) of each file, read `cut_h` (key ±0.02 carbon), `tr_C` and `vi_true` (key ±0.1). *Checkpoint:* scenario 5: 67.52, 123.3 °C, 134.3, scenario 6: 69.14, 127.7 °C, 131.8

(b) In the scenario 6 file, find the lowest and highest `kv_cst` over the whole day. Does the true viscosity ever leave its 5.70 to 6.30 cSt band? *Checkpoint:* about 5.97 to 6.03 cSt, no

(c) Name the action of MR-701 in each scenario in the convention of Section 1.6.2, and say which layer stops the walk in scenario 6. *Checkpoint:* 5: $a = +1$, reverse acting, 6: $a = -1$, direct acting, the viscosity index cap

10.9 App challenge · 9 min

Load scenario 9, What the viscosity controller cannot see, as shipped (default seed): the viscosity loop is in Auto, mid-ranging and the viscosity index cap are off, and the column loses cut sharpness at 60 min.

(a) Run to 1400 min at 60 min per second and export the CSV. In the row at 1400 min (23.333 h), read `kv_cst`, `cut_h`, `vi_true` and `metric_offspec_t`, the last to one decimal place (key ±0.1). *Checkpoint:* 6.000 cSt, 73.9 carbons, 129.4, 163.5 t

(b) Press Reset, set Viscosity index cap and Mid-ranging (cut → reactor) to On before Run, and run to 1400 min again. Read `metric_offspec_t`, `vi_true` and `tr_C` in the same row (key ±0.1). *Checkpoint:* 5.4 t, 133.5, 115.3 °C

(c) The chapter reports that the cap alone still downgrades 58 t. Explain in two sentences why the cap does so much worse than mid-ranging on this disturbance. *Checkpoint:* the cap waits for the laboratory and releases at once, mid-ranging acts on the cut point

10.10 Application area · 4 min

The app's economics put the motor oil cut $500 per tonne above the bright stock and a downgraded tonne $1500 below on-specification motor oil. The plant takes 12.0 t/h of decene at a yield of 57.4 %.

(a) What is one percentage point of yield worth per hour, and what does an hour of downgraded motor oil cost? Give both to two significant figures. *Checkpoint:* about $60 per hour and about $10,000 per hour

(b) How many hours of one extra point of yield pay for one hour off specification, and what sets how far the reactor should stay from the viscosity index limit? *Checkpoint:* about 170 h, the age of the laboratory result and the reactor lag

10.11 Application area · 3 min

The product tank turns over in 210 min. Treat it as a first-order filter on the viscosity in the line. A bottom cut left four carbons wide lifts the line to about 6.30 cSt from 6.00 cSt.

(a) What fraction of that rise reaches the tank if the excursion lasts 40 min, and if it lasts 400 min? Give two decimal places. *Checkpoint:* 0.17 and 0.85

(b) Which disturbances does the tank hide from the customer, and which does it pass on? *Checkpoint:* it hides short upsets and passes slow drifts

10.12 Review the agent · 5 min

An AI assistant reviewing the mid-ranging loop writes: "A valve position controller must be much slower than the loop it watches, so its integral time should be at least five times the inner loop's. MR-701 at 180 min against QC-701's 30 min is right, and setting it to 30 min would make the two loops fight over the viscosity and cycle." Find the flaw and say how to check it in the app.

CHAPTER 11

Real-Time Optimization of a Williams–Otto Reactor Whose Model Lacks a Reaction

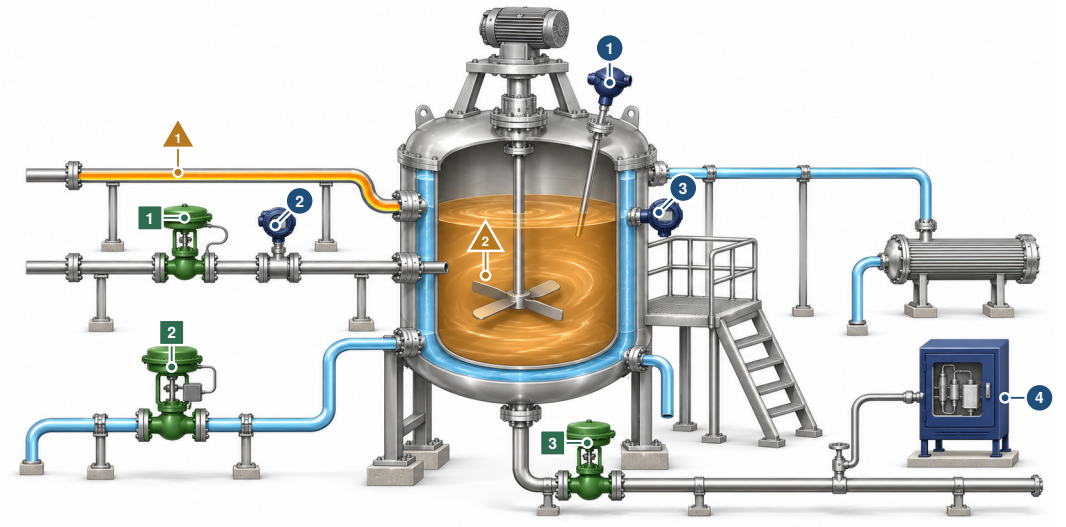

The controller moves 1 B feed valve, whose flow setpoint the optimizer chooses
2 jacket water supply, which sets the reactor temperature the optimizer chooses 3 outlet valve, which holds the level

The controller reads 1 reactor temperature 2 B feed flow 3 level
4 analyzer, the optimizer's only view of what the reactor makes

What disturbs it 1 a change in the A feed, which a meter reports
2 a waste reaction that slows, which the optimizer is not told of

Notice the analyzer cabinet at the right: the optimizer above the three loops sees the product only through its reports, minutes old and noisy. Illustration, not to scale.

IN PLAIN TERMS An Optimizer That Stops Short

The main idea. A computer above the control loops of a chemical reactor picks the temperature and feed rate that should earn the most money. It decides with a simple model of the reactor, which it adjusts after each new set of readings.

The challenge. The model leaves out one of the three reactions in the tank, so it matches what the plant makes but not how it changes when a setting moves. Finding out takes small test moves, a wait for the plant to settle, and noisy readings.

Why it matters. A computer that stops in the wrong place loses money every hour while its own model says all is well.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the history box in Section 11.2, the try-it and fact boxes in Section 11.6, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 11.6, Section 11.7, and Section 11.8 with its open challenges

The control question A real-time optimizer moves the setpoints of a reactor's loops to the best steady state its model can find, and its model lacks one of the plant's three reactions. Refitted at every steady state, the model agrees with the plant where it stands but not on the slopes. Where does such an optimizer stop, and what does it cost to correct the slopes from plant data?

Learning objectives

- Explain why a two-step optimizer with a structurally wrong model stops where the fitted model's slopes vanish, not at the plant's optimum.
- Write the modifier-adaptation corrections and estimate a plant gradient from three steady states, with the noise such an estimate carries.
- Set the period, steady-state check, filter gain and back-off of an optimization layer, and judge the profit each costs or saves.
- Compare an economic MPC with modifier adaptation after a change the layer is told of and after one it can only infer.

Prerequisites Sections 1.4.7 (two-step optimization and modifier adaptation), 1.6.2 (the loops beneath the layer), 1.6.4 (scenarios, scorecard and data) and 1.6.3 (needed only for the economic MPC).

The app https://apmonitor.com/apps/williams-otto-rto/

See also Volume 9, Chapter 8 for an optimum that sits on a moving constraint and Volume 9, Chapter 10 for where a supervisory setpoint belongs. Volume 1, Chapter 11 for a model that fits a record without the right structure, and Volume 1, Chapter 12 for fitting models from plant data. Volume 7, Chapter 2 for a jacketed reactor whose heat release can outrun its cooling.



11.1 Introduction

Feed B to a stirred tank at 4.00 kg/s with 1.8275 kg/s of A, hold it at 75 °C, and the simulated Williams–Otto reactor earns 151.12 $/s, 84.64 % of the 178.540 $/s its best permitted point makes at 4.9747 kg/s and 84.322 °C, where the waste sits on its 8 % limit. Hand the two setpoints to a real-time optimizer whose model is refitted at every steady state, and after six hourly runs it settles at 4.4454 kg/s and 78.915 °C, earning 165.736 $/s, or 92.83 % of the best. The model fitted there predicts 180.92 $/s and a profit slope of zero in both directions, while the plant's own slopes there are –24.44 $/s per kg/s of B and +5.76 $/s per degree.

That run corrects a belief that most engineers carry into optimization, that a model good enough to reproduce the plant's measurements is good enough to optimize it. The model of this optimizer has two adjustable constants, it is refitted to five measured fractions at every run, and it agrees with the plant on the waste to within half a point at its stopping point. What it does not reproduce is how the profit and the waste change when the setpoints move, and an optimizer moves along those changes. The class is general. Every steady-state optimizer above the control system of a refinery unit, an ethylene plant or a utility runs on a model that is wrong somewhere, and a model updated only in its parameters or its output biases can converge to a point that is not the plant's optimum [1–3]. Correcting the model's gradients with plant gradients removes that error at convergence [3, 4], and estimating plant gradients from noisy steady states is where the remaining difficulty lies [3].

However, the benchmark on which these methods are usually compared is a static model, on which a steady state is available at a call, while a plant must be waited for, its steady states must be detected from noisy data [5], and its analyzer reports what it sampled minutes earlier. This chapter places the reactor-only Williams–Otto benchmark under regulatory loops in the app, with an optimization layer that reads only the measurements, and measures what the two-step method and modifier adaptation do there. The plant equations, the two-reaction model and both methods are published [4, 6, 7], and Section 1.4.7 introduces both methods on the running example of this volume. The findings are measurements. The two-step fixed point is reached from 21 starting points and is set by the model's structure. The perturbation rule a running modifier adaptation needs costs about three percent of profit on every third run. The settings first suggested for the graded challenge pass at only 21 of 31

noise seeds, while a lower filter gain with a waste back-off of 0.15 point passes at all 31. The plant is dynamic, its steady-state map is exactly the published one, checked against an independent solver to every printed digit, and a measurement boundary keeps the optimizer away from the simulator's truth. The application also offers an economic model predictive controller (MPC) as a fourth method of the layer, on the optimizer's own model with the reactor's holdup and the same modifiers. The two tie on profit, and the economic controller's one gain, a lower waste peak after a metered feed rise, comes from its dynamic model. Ten cases run from a hand search to a graded 100 h challenge.

Section 11.2 follows the benchmark and the methods from 1960 to economic MPC. Section 11.3 sets out the plant: its three reactions, the energy balance and jacket the app adds, the profit and the measurements. Section 11.4 gives the optimizer's two-reaction model and shows where a model refitted at every steady state stops. Section 11.5 describes the loops, the steady-state detector, two-step, modifier adaptation with its safeguards, and the economic MPC (Section 11.5.6). Section 11.6 measures the ten cases, the economic MPC against the other methods, the worst cases and the challenge. Section 11.7 names the instruments, valves and computers behind the loops and the layer. Section 11.8 states the limitations and Section 11.9 the conclusions. The simulator framework the app runs on is described in Section 1.6.

11.2 From a Benchmark Plant to Modifier Adaptation

Williams and Otto proposed their plant in 1960 because the computer-control schemes then offered to the chemical and petroleum industries could not be compared directly, and the joint studies of computer makers and process companies were proprietary [6]. Their answer was a generalized chemical processing model on which such schemes could be compared in the open. The part that survived into the optimization literature is the reactor alone, a stirred tank fed with A and B in which a desired product P is made through an intermediate C and lost to a waste G, with a byproduct E that sells for little. It has become a common test case of real-time optimization, used for gray-box neural models [8], for evaluations of optimization strategies and online parameter adaptation [9, 10], and for modifier adaptation with Bayesian and derivative-free optimizers [7], in the form whose constants this simulator takes from those sources.

> ⧖ HISTORY **A Computer Above the Loops (1959)**
>
> In March 1959 a digital computer went online on the catalytic polymerization unit of a refinery in Texas, among the first to control a process plant in closed loop [11, 12]. Like the optimization layer of this chapter, it did not drive the valves. The first process control computers were supervisory: they adjusted the setpoints of the controllers below them, so the plant kept working when the computer was off [11, 12]. Product quality from process stream analyzers was its feedback, and every hour it adjusted 19 parameters. Conversion of propylene to gasoline rose from 80 to 91 percent, and the computer paid for itself in less than three years. By 1965 the United States had 400 such computers [11]. The benchmark of this chapter was proposed a year after that first unit came online.

Industrial practice arrived at the layer the benchmark tests from a different direction. The argument that a steady-state optimizer must sit above multivariable control to capture a plant's profit was made in 1983 [13], on-line optimization was described as a two-phase approach in an application study of 1987 [14], and identifying a steady state on line from noisy data became a subject of its own [5]. Industrial practice of the layer was reviewed and assessed in 2011 [15].

The weakness of that two-step structure had been pointed out before the industrial argument was published. In 1979 Roberts showed that a model which is not an accurate representation of the process can still lead to the optimum if the optimization includes a term comparing the model's output derivatives with the process's, a modified two-step approach in his words, and demonstrated the algorithm on the feed rate and temperature setpoints of a simulated exothermic reactor [1]. In 1994 the model accuracy an economic optimizer needs was examined for the case of a bias update [2], and adaptation strategies for real-time optimization were reviewed in 2009 [16]. Modifier adaptation,

Table 11.1 Variables by Role, With Values in Case 3 at 8 h

Variable	Role	Value, range
ϕ	objective, profit (\$/s)	165.75, best 178.54
x_G	constraint, waste by mass	7.19 %, limit 8
T_{SP}, $F_{B,SP}$	MVs of the RTO, decisions u	78.93, 4.447 in a box of 70–100, 3–7
T	CV of TIC (°C)	78.91
F_B	CV of FIC (kg/s)	4.438
L	CV of LIC (%)	50.0, SP 40–60
T_{ws}, v_B, v_o	MVs of the loops	66.9 °C, 44.5, 41.8 %
F_A	DV, A feed (kg/s)	1.8275, 1.5–2.4
a_3	DV, waste reaction activity	100 %, 60–140
p_P, p_E, p_A, p_B	DV, prices (\$/kg)	published
σ_{an}	DV, analyzer noise	0.2 %, 0–2
m_j	states, six component masses	$W = 2105$ kg
T, T_j, T_{ws}	states, temperatures	78.9, 73.7, 66.8 °C
v_B, v_o	states, valve positions	—
θ_1, θ_2	RTO model constants	−2.519, −1.797
ε, λ_ϕ, λ_G	modifiers (MA only)	zero under two-step

published the same year, adds to the model's cost and constraints the differences between the plant's and the model's values and gradients, so that at convergence the modified problem satisfies the plant's conditions of optimality [4], and it needs no parameter estimation [3]. The gradients are the price. Dual formulations of 2010 and 2011 [17, 18], which a 2016 thesis states as constraints keeping each new point off the line through the last two [9], keep the operating points spread enough to estimate them, and a 2016 review gathers the variants, the estimators of gradients from noisy measurements and the applications that followed [3].

The wait for steady state has drawn two answers of its own. Modifier adaptation was reformulated in 2014 to use transient measurements, and on a simulated stirred tank it then converged in about one settling time of the plant where the static scheme needed more than five iterations [19]. Economic MPC takes the other route and puts the profit itself into a receding-horizon controller on a dynamic model, in place of a setpoint handed down from a steady-state layer. Its average performance is never worse than the best steady operation, an analysis of 2012 shows under its assumptions [20], and a tutorial review of 2014 collects its formulations [21].

11.3 The Reactor as Modeled

The simulated plant is a jacketed stirred tank fed with pure A at F_A and pure B at F_B, as shown in Figure 11.1. A level loop (LIC) holds the holdup W through the effluent valve v_o, a flow loop (FIC) holds the B feed through its valve v_B, and a temperature loop (TIC) holds the reactor temperature T by the temperature T_{ws} of the tempered water sent to the jacket. An analyzer (AT) on the effluent reports the six mass fractions x_j every 5 min, 5 min after it drew the sample. Above the loops the real-time optimization (RTO) layer sets the two setpoints T_{SP} and $F_{B,SP}$. Table 11.1 classifies the variables, with their values in case 3 at 8 h, where two-step has settled.

11.3.1 Reactions and Component Balances

The plant runs the three reactions of the benchmark, A + B → C, B + C → P + E and C + P → G, at rates in kilograms per second given by Equation (11.1).

$$\begin{aligned} r_1 &= k_1 x_A x_B W, \quad r_2 = k_2 x_B x_C W, \quad r_3 = a_3 k_3 x_C x_P W \\ k_i &= A_i \exp\big(-E_i/(T+273.15)\big) \end{aligned} \tag{11.1}$$

where $x_j = m_j/W$ is the mass fraction of component j, $W = \sum_j m_j$ the holdup in kilograms, T in °C, and a_3 the activity of the waste reaction that the drift disturbance lowers. The constants, $A_i = 1.6599 \times 10^6$, 7.2117×10^8 and 2.6745×10^{12} s^{-1} with $E_i = 6666.7$, 8333.3 and 11 111 K, are

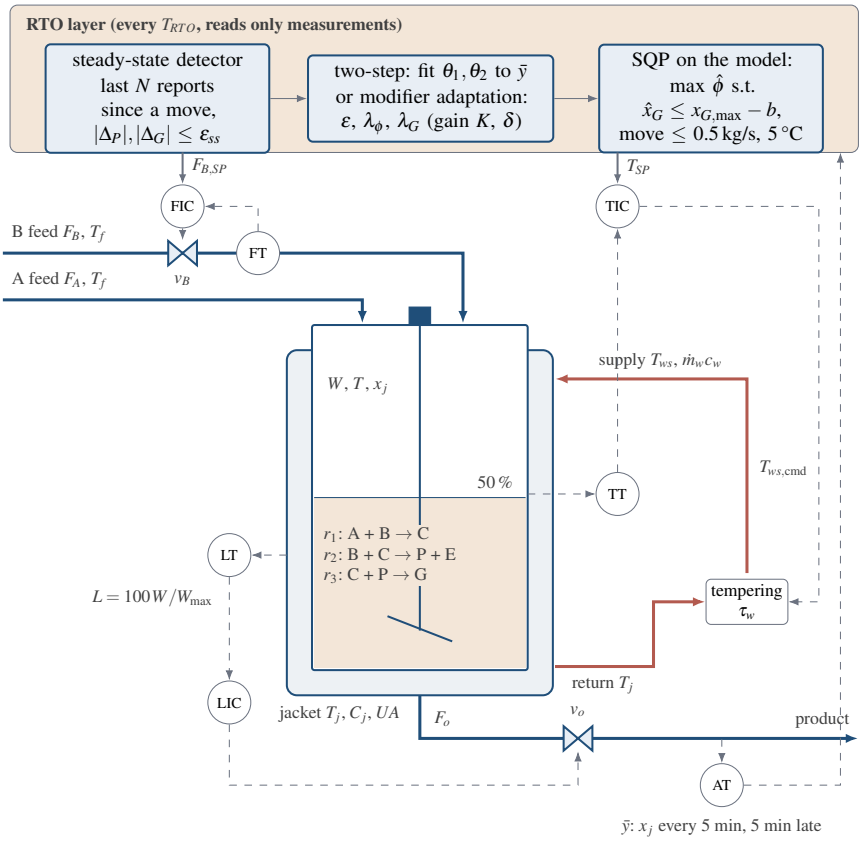

Figure 11.1 The Williams–Otto reactor under its three loops and the RTO layer, after the application's Reactor view of case 3 at 8 h with two-step settled. Feeds A and B at T_f enter the tank of holdup W at temperature T, where the three reactions run at r_1 to r_3. TIC moves the jacket supply T_{ws} through a tempering lag τ_w, the jacket of capacity C_j and conductance UA returns at T_j, FIC moves v_B and LIC moves v_o, and the analyzer reports the fractions x_j that the layer averages into $\bar{y}$. At that instant $T = 78.91\,^\circ$C (setpoint 78.93), $F_B = 4.438$ kg/s (setpoint 4.447), $T_{ws} = 66.8\,^\circ$C, $T_j = 73.7\,^\circ$C, $F_o = 6.260$ kg/s, the level is 50.0 %, the analyzer reads $x_G = 7.18$ %, and the plant earns 165.75 \$/s against the best 178.54 \$/s. Process model only, not to scale.

those of the reactor-only benchmark as the sources give them [8, 9]. The component masses follow Equation (11.2).

$$\frac{\mathrm{d}m_j}{\mathrm{d}t} = F_{j,in} + \sum_{i=1}^{3} \nu_{ij} r_i - F_o x_j$$

$$\nu = \begin{bmatrix} -1 & -1 & 2 & 0 & 0 & 0 \\ 0 & -1 & -2 & 2 & 1 & 0 \\ 0 & 0 & -1 & 0 & -\tfrac{1}{2} & \tfrac{3}{2} \end{bmatrix} \tag{11.2}$$

with columns in the order A, B, C, E, P, G, $F_{A,in} = F_A$, $F_{B,in} = F_B$ and the others zero. Each row of ν sums to zero, because the benchmark gives no molecular weights and the balances are written in mass, so mass is conserved reaction by reaction. Over 3 h with a 10 °C step and the A feed raised to 2.2 kg/s, the holdup differs from the integrated flows by 1.7×10^{-10} kg, and over 100 h by 2.9×10^{-8} kg.

11.3.2 Energy, Jacket and Valves

The benchmark is a steady-state problem with the temperature as a decision. The simulator turns the temperature into a controlled variable with an energy balance and a jacket, as shown in Equation (11.3).

$$\begin{aligned} W c_p \frac{\mathrm{d}T}{\mathrm{d}t} &= (F_A + F_B) c_p (T_f - T) \\ &\quad + \textstyle\sum_i (-\Delta H_i) r_i - UA(T - T_j) \\ C_j \frac{\mathrm{d}T_j}{\mathrm{d}t} &= \dot{m}_w c_w (T_{ws} - T_j) + UA(T - T_j) \\ \tau_w \frac{\mathrm{d}T_{ws}}{\mathrm{d}t} &= T_{ws,\mathrm{cmd}} - T_{ws}, \qquad \tau_v \frac{\mathrm{d}v}{\mathrm{d}t} = v_{\mathrm{cmd}} - v \end{aligned} \tag{11.3}$$

where the last equation holds for both valves, $F_B = F_{B,\max} v_B$ and $F_o = F_{o,\max} v_o$. Every constant in it is a teaching value, listed in Table 11.2. At the best point the reactions release 439.2 kW, heating the feeds from 80 °C absorbs 88.2 kW, and the jacket holds 79.93 °C with a supply of 74.08 °C, so the jacket cools across the operating box. The level loop returns the holdup to 2105.2 kg and the temperature loop puts T on its setpoint, so every steady state of the dynamic plant is the published one. The heats decide only how long the plant takes to get there. The balance has the form of the jacketed reactor of Volume 7, Chapter 2, where heat release can outrun the jacket and the reactor ignites. Here the jacket only sets the pace of each move.

11.3.3 Profit, Measurements and Numerics

The profit pays for P and E in the effluent and charges for both feeds, as shown in Equation (11.4).

$$\begin{aligned} \phi &= (p_P x_P + p_E x_E)(F_A + F_B) - p_A F_A - p_B F_B \\ x_G &\le x_{G,\max} \end{aligned} \tag{11.4}$$

in dollars per second, the benchmark's unit, with $p_P = 1143.38$, $p_E = 25.92$, $p_A = 76.23$ and $p_B = 114.34\,\$/\mathrm{kg}$ [9, 10]. The 8 % waste limit is a declared teaching value. The best permitted point earns 178.540 \$/s, about 643 000 \$/h. The analyzer draws a sample every 5 min and reports it one cycle later, each fraction multiplied by $1 + \sigma_{an} n$ with n a unit normal deviate, so a reading is always 5 to 10 min old. The thermowells have a lag of 0.3 min, the reactor reading 0.02 °C of noise, and the flow meters a lag of 0.1 min and 0.2 % relative noise. The noise comes from the framework's seeded generator (Section 1.6.1), and every execution draws the same count of deviates whatever the noise settings, so changing one noise level leaves the other channels' deviates as they were. The 13 states, the six masses, three temperatures, two valve positions and two flow ledgers that check the mass balance, advance by fourth-order Runge–Kutta steps of 6 s. The spectral radius of the Jacobian is 646 h^{-1} at 7 kg/s and 100 °C, a time constant of 5.6 s, and 279 to 281 h^{-1} at the best point and at three other corners of the box, so $\lambda \Delta t$ is at most 1.08 against the method's real-axis limit of 2.785. With the step halved, the temperature and waste 1.5 h after a 10 °C step differ by 0.0026 °C and 1.2×10^{-5}. Each case starts lined out at its setpoints, the steady state found by Newton's method on the six component balances.

11.3.4 Assumptions

The tank is perfectly mixed, the reactor-only benchmark carries no separator or recycle, and every balance is in mass. The steady state is assumed unique. The plant's Newton solve from one fixed start converged at all 5440 points of a grid over the box (B feed every 0.25 kg/s, temperature every 2 °C, the activity from 60 to 140 % and the A feed from 1.5 to 2.4 kg/s), which shows convergence but does not prove uniqueness, and a plant with several steady states would need a different initialization. The energy balance and the jacket are not part of the benchmark. They make each move take about half an hour, which is what an optimization layer above real loops has to wait for, and they cannot change where the plant settles. The waste limit, the prices' use as disturbances and the noise levels

Table 11.2 Constants. P: Published, Shown in the Model Dialog. A: Adjustable in the Parameters or Disturbance Panel. H: Hidden

Symbol	Meaning	Value	Mark
A_i, E_i	Arrhenius constants	see text	P
W_0	holdup at 50 % level	2105.2 kg	P
$W_{\max}$	holdup at 100 % level	4210.4 kg	H
F_A	A feed	1.8275 kg/s	P, A
$p_{P,E,A,B}$	prices ($/kg)	1143.38, 25.92, 76.23, 114.34	P, A
c_p	heat capacity of the mass	3.0 kJ/(kg K)	A
T_f	feed temperature	80 °C	A
$\Delta H_{1,2,3}$	heats of reaction	−150, −250, −100 kJ/kg	A
UA	jacket conductance	80 kW/K	A
C_j	jacket heat capacity	2000 kJ/K	A
$\dot{m}_w c_w$	jacket water flow capacity	60 kW/K	A
τ_w, τ_v	tempering, valve lags	2, 0.5 min	A
$F_{B,\max}$, $F_{o,\max}$	valve capacities	10, 15 kg/s	A
ψ_1, ψ_2, T_{ref}	model temperature terms	−17, −29, 383.15 K	P
θ^0	model constants at start	(−3, −4)	P
$x_{G,\max}$	waste limit	8 %	A
Δt	integration and control step	6 s	fixed

are teaching values. A real permit and its margin would come from the plant's operating license and its own data.

11.4 The Optimizer's Model and the Mismatch

11.4.1 Two Reactions Where the Plant Has Three

The RTO's model is the published two-reaction approximation, A + 2B → P + E and A + B + P → G, with no C, whose rate laws and reparametrized constants the application takes from the supporting code of a 2021 study [7], as shown in Equation (11.5).

$$\begin{aligned} \hat{r}_1 &= \hat{k}_1 x_A x_B^2 W, \qquad \hat{r}_2 = \hat{k}_2 x_A x_B x_P W \\ \hat{k}_i &= \exp(\theta_i)\exp\big(\psi_i\,(T_{ref}/(T+273.15)-1)\big) \end{aligned} \tag{11.5}$$

with mass coefficients A $-\hat{r}_1-\hat{r}_2$, B $-2\hat{r}_1-\hat{r}_2$, E $2\hat{r}_1$, P $\hat{r}_1-\hat{r}_2$ and G $3\hat{r}_2$, each reaction conserving mass, and the same profit and limit. The model has five fractions, two constants to fit and fixed temperature dependences, so it can match the plant's compositions only approximately. Its predictions are written $\hat{x}_j(u,\theta)$, $\hat{\phi}(u,\theta)$ and $\hat{x}_G(u,\theta)$ with $u=(F_B,T)$. A model with the wrong structure can still fit a record well, as the fitted radiative term of Volume 1, Chapter 11 does in data generated without one.

11.4.2 The Published Point and the Plant Optimum

At 4.787 kg/s and 89.70 °C the plant gives the fractions 0.0875, 0.3896, 0.0153, 0.2906, 0.1095 and 0.1075 for A, B, C, E, P and G, within 0.00075 of the operating point a 2019 thesis tabulates [10], and earns 190.991 $/s. That table labels its fourth and fifth fractions as P and E, against its own state order and the balances, which make them E and P, and it lists the point as an initial condition and prints no profit. That the point is the unconstrained optimum is a computation. The application's SQP finds it at (4.7875 kg/s, 89.703 °C) with 10.75 % waste, and with the 8 % limit the optimum moves to (4.9747 kg/s, 84.322 °C), 178.540 $/s, with the waste on the limit. A solver written independently of the application (steady states by time integration polished by Newton, optima by Nelder–Mead and a boundary search on the limit) reproduces both points and both profits to every printed digit. The residence time $W/(F_A+F_B)$ is 5.16 min at the best point and 5.30 min at the published one.

11.4.3 Where a Refitted Model Stops

Section 1.4.7 shows on the running example S-1 why a refitted model can stop short of the plant's optimum. Two-step RTO from (4.0 kg/s, 75 °C) and $\theta^0 = (-3, -4)$, iterated on the static plant, converges to (4.4454 kg/s, 78.915 °C) with $\theta = (-2.5189, -1.7964)$. The plant earns 165.736 \$/s there, 92.83 % of its best, with 7.187 % waste, while the model, fitted there, predicts 7.645 % and 180.92 \$/s. The same point is reached from 20 other starts spread over the box, to within 4.9×10^{-8}, from $\theta^0 = (-1, -1)$, and by the independent solver from five starts. The model's optimum is interior, so the limit plays no part. The stop is where the fitted model's gradient vanishes, and the plant's gradient there is (−24.44 \$/s per kg/s, +5.763 \$/s per degree), which says that, locally, the plant would earn more with less B and a hotter reactor. Modifier adaptation with exact plant gradients and a filter gain of 0.5, from the same start, reads 84.641, 96.510, 98.648, 99.465 and 99.803 % of the best at its first iterations and 100.000 % from the tenth, on the waste limit, as the independent solver confirms.

Pause. Try Exercise 11.1 before reading on.

11.5 Loops and Optimization Layer as Implemented

11.5.1 Regulatory Loops

The three loops execute every 6 s on the velocity-form algorithm of Section 1.6.2, with no derivative action and with the tracking time of the back-calculation equal to each loop's integral time. TIC acts on $T_{SP} - T$ with $K_c = 2.5$ °C per °C and $\tau_I = 6$ min and commands T_{ws} between 20 and 140 °C. FIC uses 6 % of valve per kg/s and 0.5 min. Both are reverse acting, with action $a = +1$ in Equation (1.14). LIC uses 6 % per percent and 6 min and is direct acting, $a = -1$, so that a high level opens the effluent valve. From 75 °C, a 3 °C setpoint step overshoots by 0.47 °C and stays within 0.1 °C of the setpoint from 13.5 min after the step, and a 5 °C step overshoots by 0.78 °C and settles in 21.9 min. The composition follows with its 5 min residence and the analyzer reports 5 min late, so a move takes about half an hour to read. The header's Manual places all three loops on their stations and pauses the layer, and each transfer is bumpless, as Section 1.6.2 describes.

11.5.2 When the Layer Runs

The layer reads only the measurement packet: the analyzer reports, the meters, and the prices and limit it is configured with. It executes no sooner than the period T_{RTO} after its last setpoint move and, with the steady-state check on, only when the detector of Equation (11.6) passes, over the last N reports drawn since that move.

$$\begin{aligned} \Delta_j &= \hat{\beta}_j\,(t_N - t_1) \\ \text{steady} &\iff |\Delta_P| \le \varepsilon_{ss} \text{ and } |\Delta_G| \le \varepsilon_{ss} \end{aligned} \tag{11.6}$$

where $\hat{\beta}_j$ is the least-squares slope of fraction j against the sample times $t_1 \ldots t_N$, so Δ_j is how far the fitted line moves over the window. The defaults are $T_{RTO} = 1$ h, $N = 6$ and $\varepsilon_{ss} = 0.2$ percentage points. The steady state the run uses, $\bar{y}$, is the mean of the last N reports with the mean measured flows. Every move is bounded to 0.5 kg/s and 5 °C by a move box inside the decision box of 3 to 7 kg/s and 70 to 100 °C, for both methods.

11.5.3 Two-Step

Each two-step run fits the two constants by least squares on five fractions and then moves to the fitted model's best permitted point, as shown in Equation (11.7), the method of Equation (1.12) with the waste limit added.

$$\begin{aligned} \theta_k &= \arg\min_{\theta} \textstyle\sum_{j\in\{A,B,E,P,G\}} \big(\bar{y}_j - \hat{x}_j(u_k, \theta)\big)^2 \\ u_{k+1} &= \arg\max_{u} \hat{\phi}(u, \theta_k) \;\; \text{s.t. } \hat{x}_G(u, \theta_k) \le x_{G,\max} - b \end{aligned} \tag{11.7}$$

where b is the waste back-off, zero unless the learner sets it. The fit is Gauss–Newton from the last θ with a step capped at one unit and halved until the residual falls. The optimization is a sequential quadratic program (SQP) in two variables scaled so that 1 kg/s of B counts as 10 °C, with central-difference gradients, a central-difference Hessian of the Lagrangian shifted to be positive definite, an exact-penalty backtracking line search, and each quadratic subproblem handed to the shared solver of the MPC option (Section 1.6.3), an implementation of the dual algorithm of Goldfarb and Idnani [22]. The linearized waste row is elastic, a slack penalized 10^5 times the profit gradient's scale, which is the application's own device so that a linearization that cannot be met inside the move box yields the step that violates it least instead of no step. Over the 100 h challenge on the suggested settings the layer made 99 runs and 255 SQP solves, none of them failing, at 3.43 iterations each on average.

11.5.4 Modifier Adaptation

Modifier adaptation keeps the model as last fitted (fitting it once if it never was) and corrects it at every run k, as shown in Equation (11.8).

$$\begin{aligned} \phi_m(u) &= \hat{\phi}(u) + \lambda_\phi^\top (u - u_k) \\ G_m(u) &= \hat{x}_G(u) + \varepsilon_k + \lambda_G^\top (u - u_k) \\ \varepsilon_k &= \bar{y}_G - \hat{x}_G(u_k) \\ \lambda &\leftarrow (1 - K)\,\lambda + K\,(\nabla_p - \nabla\hat{}) \end{aligned} \tag{11.8}$$

and moves to the maximum of ϕ_m subject to $G_m \le x_{G,\max} - b$. Against Equation (1.13), the profit carries no value offset, which would not move the maximum, and the predicted waste carries one. The waste offset ε is taken whole at every run, because filtered at K it left the steady waste 0.10 to 0.13 points above the limit with no noise at all, against 0.02 to 0.04 unfiltered, as the application's development record notes. The gradient modifiers λ are filtered with gain K, 0.5 by default. The model's gradients $\nabla\hat{}$ come from central differences, and the plant's from the last three steady states read, as shown in Equation (11.9).

$$\begin{bmatrix} (u_{k-1} - u_k)^\top \\ (u_{k-2} - u_k)^\top \end{bmatrix} \nabla_p \phi = \begin{bmatrix} \phi_{k-1} - \phi_k \\ \phi_{k-2} - \phi_k \end{bmatrix} \tag{11.9}$$

with the same system for the waste. The estimate is used only when the inverse condition number of the scaled difference matrix is at least 0.1. On three exact plant points around (4.6 kg/s, 82 °C) it returns (−19.03, +4.081) against central differences of (−18.81, +4.117), and with $K = 1$ the corrected gradient equals the plant's, which confirms the sign conventions.

11.5.5 Three Safeguards and a Switching Rule

Three points on a line give no slope across it, and a running optimizer that moves straight toward an optimum lines its points up. Without excitation the three-point estimate of the application's first version collapsed to undefined modifiers and then to a corner of the box. The application therefore keeps the next point at least δ (default 1 °C, with 1 kg/s counted as 10 °C) off the line through the last two, solving the corrected problem once on each side of that line with $a^\top (u - u_k) \ge \delta$ as an extra row and taking the better side, unless the uncorrected move already lies that far off the line, and the first run after a fresh start probes 0.1 kg/s per degree of δ in B. The rule resembles the minimum-perturbation constraint of the dual approach of Rodger and Chachuat [18] as a 2016 thesis states it in its Equations 2.16 and 2.17 [9], without that approach's bound on the step length and with three-point differences rather than Broyden updates, so it is presented as the application's design and not as a reproduction of a published algorithm. The move box is the second safeguard. The third is a plant-change test of the application's own design: once three plant gradients have been used, a steady state whose profit change differs from the corrected model's prediction by more than 4 % of the profit, or whose waste differs by more than one point, clears the gradient history. The move limits, the test

and the elastic row were added after the challenge's drift drove the application's first version, as its development record notes, to (3.0 kg/s, 100 °C) with 16 points of excess waste on the default seed.

🕮 **Pause.** Try Exercise 11.2 before reading on.

Switching the method to modifier adaptation also clears its steady-state history, its gradient count, its last prediction and its modifiers. A usability review of the application found why. Two-step's last steady states sit within 0.002 kg/s and 0.03 °C of one another, so slopes estimated from them are noise, and before the rule a switch at 8 h in case 3 sent the plant to 55.8, 5.1, −25.6 and −7.8 percent of the best at the next four runs. With the rule the same switch reads 91.22, 90.60, 97.48 and 98.37 percent, and runs 14 to 24 average 98.87 %.

11.5.6 Economic Model Predictive Control

The economic MPC is the fourth and last method of the layer, and no case except one starts with it. Where the MPC option of Section 1.6.3 tracks setpoints with the objective of Equation (1.19), this controller prices the profit itself and chooses the setpoints. It writes the same two setpoints every $T_s = 5$ min, the analyzer cycle, counted as 50 executions of the 6 s loops, and plans them over 2 h, $N = 24$ steps grouped into nine move blocks, four single steps and then blocks of 2, 2, 3, 4 and 9 steps. The three loops run underneath as they do under the other methods. Its model is the optimizer's two-reaction model of Equation (11.5) made dynamic by the holdup, with the two setpoint loops as first-order lags, as shown in Equation (11.10).

$$\begin{aligned} \frac{\mathrm{d}\hat{x}_j}{\mathrm{d}t} &= \frac{1}{W_0}\left(F_{j,in} + \textstyle\sum_{i=1}^{2} \hat{\nu}_{ij}\hat{r}_i - (F_A + F_B)\,\hat{x}_j\right) \\ \frac{\mathrm{d}T}{\mathrm{d}t} &= \frac{T_{SP} - T}{\tau_T}, \qquad \frac{\mathrm{d}F_B}{\mathrm{d}t} = \frac{F_{B,SP} - F_B}{\tau_B} \end{aligned} \tag{11.10}$$

where j runs over A, B, E, P and G, $\hat{\nu}_{ij}$ are the mass coefficients listed after Equation (11.5), the rates use the fitted θ, and $W_0 = 2105.2$ kg is the declared holdup that the level loop holds, not the run's. The lags come from closed-loop step tests of ± 2 °C and ± 0.2 kg/s at the plant's best point with the layer off and the analyzer noise off, averaged over the two signs and fitted as first order plus dead time. The temperature loop gives a gain of 0.988, a lag of 1.06 min and a dead time of 1.33 min, with a residual of 0.045 of the step (rms) that is mostly the loop's overshoot, and the B feed loop a gain of 1.000, 0.71 min and 0.04 min, with 0.002. Each lag in the model is the sum of the fitted lag and dead time, $\tau_T = 2.39$ min and $\tau_B = 0.74$ min. The model's steady state is exactly the optimizer's, which a 40 h integration confirms to 7×10^{-15} in mass fraction. The model's five fractions come from an observer, not from the analyzer, which integrates them at every execution from the measured temperature, B feed and A feed (the last filtered over one decision period), starting from the model's steady state when the method takes over, and the analyzer reaches the controller only through the offset ε read at steady states.

Every decision begins by linearizing the model, by central differences, at its steady state for the plan's target u_t and discretizing it for a zero-order hold over T_s, with its affine remainder, as $z_{k+1} = A_d z_k + B_d u_k + c_d$ for $z = (\hat{x}_A, \hat{x}_B, \hat{x}_E, \hat{x}_P, \hat{x}_G, T, F_B)$ and $u = (F_{B,SP}, T_{SP})$. An analytic Jacobian with a separate scaling-and-squaring exponential reproduces A_d, B_d and c_d to 1.2×10^{-11}. At the target of case 9's decision at 13.0 h, (5.459 kg/s, 84.95 °C), the residence time is 5.18 min and the steady gains of the waste fraction are −0.02127 per kg/s of B and +0.004277 per degree, equal to the static model's. The plan maximizes the corrected profit of Equation (11.8) along the horizon, as

shown in Equation (11.11).

$$\begin{aligned}
\min_{u} \sum_{k=1}^{N} \Big[& -T_s\, g^{\top} y_k + \tfrac{T_s}{2} (s_k - u_t)^{\top} \tilde{H} (s_k - u_t) \Big] \\
& + \textstyle\sum_{k=2}^{N} \left(\rho_1 \sigma_k + \rho_2 \sigma_k^2 \right) + \sum_k \Delta u_k^{\top} S \Delta u_k \\
\text{s.t. } & \hat{x}_{G,k} + \varepsilon + \lambda_G^{\top} (\zeta_k - u_r) \le x_{G,\max} - b + \sigma_k \\
& \sigma_k \ge 0, \qquad k = 2, \ldots, N \\
& u_{lo} \le u_k \le u_{hi}, \qquad |\Delta u_k| \le \Delta u_{\max} \\
& u_k = u_t \text{ from the first block that } \Delta u_{\max} \text{ allows}
\end{aligned} \tag{11.11}$$

where $y_k = (\hat{x}_P, \hat{x}_E, F_B, T)$ and $\zeta_k = (F_B, T)$ are predicted, s_k are the setpoints in force at step k, u_r are the setpoints of the steady state last read, and g is the gradient of the corrected profit at the target's steady state $\bar{x}$, with $\bar{F} = F_A + F_{B,t}$,

$$g = \left(p_P \bar{F},\ p_E \bar{F},\ p_P \bar{x}_P + p_E \bar{x}_E - p_B + \lambda_{\phi,1},\ \lambda_{\phi,2} \right). \tag{11.12}$$

The matrix $\tilde{H}$ is the model's steady-profit curvature at the target, by central differences and floored as the layer's SQP floors it (in scaled variables, each eigenvalue at least 1 % of the trace), and the modifiers add none because they are linear. At the decision at 13.0 h it is $[[28.377, -2.369], [-2.369, 0.773]]$ in \$/s per $(\mathrm{kg/s})^2$, per kg/s and degree, and per degree squared, and the floor did not act. The move suppression is a tenth of one period's curvature, $S = 0.1\, T_s\, \mathrm{diag}(\tilde{H})$. The rate limit $\Delta u_{\max}$ is 0.125 kg/s and 1.25 °C a decision, so eight moves cross the move box corner to corner and the target can always be reached inside the horizon. The box $[u_{lo}, u_{hi}]$ is modifier adaptation's move box about the last read point intersected with the decision box. The waste ceiling is the corrected waste of Equation (11.8) applied to the predicted transient and is soft, with one slack per step priced at $\rho_1 = 10^8$ and $\rho_2 = 10^{11}$ dollars per unit of mass fraction and per unit squared, so that one point of excess held for one step costs about $\$1.1 \times 10^7$, some 500 times in the linear term alone what a point of waste is worth near the limit (6.69 \$/s, or \$2006 a step). The first step is not judged, because the earliest the waste can answer a move is through the loop and the holdup, and the ceiling is judged from 10 min ahead. Judged only from 30 min, the plan ran hot inside the unjudged half hour and the waste peak after the feed rise of Section 11.6 was 9.05 % on three seeds against 8.63 %. A ridge of 10^{-6} keeps the program strictly convex, and the shared dual active-set solver [22] computes it, so each decision is one SQP for the target and one quadratic program for the path, with no nonlinear solver.

The target u_t is modifier adaptation's own decision. Under this method the layer keeps running the estimator of Equations (11.8) and (11.9) with its period, detector, filter, perturbation rule and change test, but hands its decision to the plan instead of moving. A probe or a perturbation is held at the requested point at once, because planned economically it earned nothing on the model, was put off at every decision, and left the plant at 74.8 to 75.5 % of the best for want of a gradient. A corrected optimum is reached along the plan. When the measured A feed has moved more than 1 %, or a price or the waste limit has changed, since the last steady state was read, the target is re-solved at every decision by the layer's SQP on the corrected model at the conditions then in force, inside the same move box. This is the only information the economic MPC acts on that modifier adaptation must wait an hour to use. A plan whose first move and target both lie within 0.02 kg/s and 0.2 °C of the setpoints holds them, and a move larger than that restarts the detector and the period, so that steady states are read at setpoints that have held.

The arrival rows of the last line of Equation (11.11) pin every block from the first one the rate limits allow to the target, and the setpoint that binds takes its first move as a full step. With the arrival left to the end of the horizon, or to a block counting down toward it, each new plan put the arrival off by one more decision. Without the arrival rows the plan earned 99.6 to 101.1 % of the best by riding

the transient, and held the true waste above 8.1 % at steady state for 5.8 to 7.4 h in 24 h (three seeds each, with no change, a feed rise or a price cut).

The modifiers are validated before the plan takes them. Modifier adaptation arms its change test only after three gradients, and in case 7 the plan's steady states fell so that one slope spanned the drift, giving $\lambda_{\phi,1} = +90\,\$/\mathrm{s}$ per kg/s and walking the first version of the plan to the corner of its move box. The rule adopted applies the same test, with the same thresholds of 4 % of the profit and one point of waste, to every steady state read under the plan from the second on, against the plan's own prediction of it. A new gradient estimate enters the plan's λ, through the filter of Equation (11.8), only when the two newest of the three steady states behind it passed. Otherwise the plan keeps the λ it last accepted and, while its modifiers differ from the layer's, makes modifier adaptation's decision itself with its own. The offset ε is taken at every read, as modifier adaptation takes it. The rule was fixed before it was run, and its results, good and bad, are in Section 11.6.5.

A plan is accepted only when the target is finite, the solver reports an optimal solution, every array has its shape, every planned setpoint is finite and inside its box, the first move is inside the rate limit, the last block sits on the target, every slack is finite and not negative, and the cost and the predicted waste are finite. Anything else is rejected and named, the setpoints hold where they are, and modifier adaptation moves them at its next steady state until a plan is accepted again, which then returns at the rate limit. A test injects each fault and finds each rejected without a step in any setpoint. The solver's plans were accepted in every run this chapter reports. In Chrome on the real frame loop of case 9, with another process loading the host, the decisions took 1.38 ms on average and 1.7 % of a core at the default speed, and 9.4 % at ten times that speed, close to the apps' budget of 10 % of a core for one controller.

11.5.7 Plant, Model and What the Optimizer May Read

Three models are kept apart. The plant of Section 11.3 is the first. The identified model is the two-reaction approximation with its fitted θ, and the corrected model adds ε and λ. The plant's best point, the profit map, the gap to the best and the challenge's yardstick are computed by the simulator from its own steady state, for display and grading, and no controller or optimizer reads them. The comma-separated export marks those columns with a `truth_` prefix. The economic MPC's model is the identified one with the declared holdup and the two identified loop lags, and it reads what the layer reads together with the measured temperature and B feed and the setpoints in force. One of the application's tests raises each of the 13 parameters, each disturbance setting other than the prices and the hidden masses of C and E by 20 % in place, without advancing the plant, in case 9 at 13 h, and the controller's inputs and decision stay the same bit for bit.

11.6 Ten Cases, Measured

The cases are listed in Table 11.3, and the Exercises section turns several of them into runs for the reader. Each number in this section was measured by stepping the engine directly, outside any browser, at the application's seed (the default seed of Section 1.6.1) unless six seeds are named, with settings entered as each case's instructions say, and every case runs identically whether it is opened first, after another case, or again through Reset. A percentage quoted for run k is the plant's true profit, as a share of the best, at the steady state run k read, which is where run $k-1$ left the plant.

11.6.1 A Hand Search and a Model That Points Away

In case 1 at 4.00 kg/s of B, raising the setpoint from 75 to 78.0 °C takes the true profit from 151.1 to about 168.1 \$/s, with the waste reaching 8 % at 78.002 °C. The profit would keep rising, to 182.51 \$/s at 84.70 °C, but with 11.25 % waste, so the limit and not the peak ends the search, as the lowest-cost point of Volume 9, Chapter 8 sits on a constraint. A move past the limit shows quickly. A setpoint of 79 °C entered at 0.1 h raises the waste alarm at 0.33 h, 0.23 s later at the default speed of 1 h per second. Case 2 starts at the plant's best point and lets the layer fit the model there, 35 min in, with a misfit of 0.408 points rms over five fractions. The fitted model predicts 8.35 % waste where the

Table 11.3 The Ten Cases

#	Start and events	Teaching point
1	(4.00, 75), RTO off	hand search, active limit
2	at the best point, RTO off	plant against model map
3	(4.00, 75), two-step from 0.1 h	where two-step stops
4	two-step point, MA	modifier adaptation
5	best point, MA, noise 0.2 to 1 % at 6 h	noise and the gain K
6	(4.79, 89.7), MA, limit 12 % then 8 % at 12 h	a limit turns active
7	best point, MA, a_3 to 80 % at 6 h	unannounced drift
8	(4.00, 75), MA every 0.25 h, check off	period and detector
9	best point, economic MPC, A feed 1.83 to 2.2 kg/s at 12 h, P price −10 % at 24 h	planning through the transient
10	Challenge: 100 h, drift at 30 h, P price −10 % at 60 h, two-step every 0.25 h as shipped	the whole problem

Table 11.4 Measured Outcomes at the Application's Seed

#	Setting	Outcome
1	78.0 °C at 4.00 kg/s	168.1 $/s, waste 8.00 %
2	model fitted at the best	model best (4.52, 80.46), 195.07 $/s
3	two-step, runs 6–12	165.65–165.99 $/s, 92.78–92.97 %
4	MA, runs 10–19	mean 98.96 %, dips 96.66–97.09 %
	$\delta = 0$, 0.5, 2 °C	mean 81.44, 99.33, 97.99 %
5	$K = 0.5$ as shipped, 24 h	Profit so far 97.18 %
	$K = 0.3$, 0.2 from 6.2 h	98.40, 98.68 %
6	reads after the cut	8.10, 7.97 % waste at 99.66, 99.58 %
7	MA against Off, 24 h	98.50 against 95.26 %
8	0.25 h, check off	94 runs, 87.78 %, waste +1.97
	check on	38 runs, 97.68 %, waste +0.21
	period 1 h	23 runs, 96.60 %, waste +0.04

plant makes 8.00 %, so it believes the plant is over its limit, and it places its best point at (4.52 kg/s, 80.46 °C), promising 195.07 $/s. Two-step switched on at 0.7 h goes there in one run, and the next steady state is 96.61 % of the best. The path ends at (4.45 kg/s, 78.9 °C), the model fitted at the right point having pointed away from it.

11.6.2 Two-Step Against Modifier Adaptation

In case 3 two-step switched on at 0.1 h makes its first run at 1 h, to (4.387 kg/s, 77.65 °C), and by the sixth it has stopped at (4.445 kg/s, 78.93 °C), the static fixed point. The steady states it reads from 6 to 12 h earn 165.65 to 165.99 $/s, 92.78 to 92.97 % of the best, while the fitted model promises 180.49 to 182.05 $/s at the same runs. At 8 h the model's best point (4.448 kg/s, 78.93 °C) sits on the setpoints (4.447 kg/s, 78.93 °C), as the application's Reactor view of that instant shows in Figure 11.3. Switched to modifier adaptation at that instant, the same run climbs toward the plant's best point and reads 99.46 % of it at 20 h, at (4.676 kg/s, 82.36 °C), as shown in Figure 11.2. The corrected model there carries $\varepsilon = -0.75$ points and a temperature slope of −1.24 $/s per degree in its uncorrected part, against a plant estimate of +4.36 $/s per degree.

> **TRY IT IN THE APP** **3 · Two-Step RTO Settles Below the Best**
>
> Load the scenario, press Run and set RTO method to Two-step. Watch the numbered path on the profit map: it climbs for a few hourly runs and by the sixth it stops near 4.45 kg/s and 78.9 °C, short of the star. At 8 h Gap to best reads 7.16 %. Switch Map shows to RTO model, and the model's own best point sits where the path stopped.

Case 4 starts at the two-step point with modifier adaptation. The steady states runs 1 to 6 read are 91.21 to 98.20 % of the best while the gradient estimates build up. From run 7 on, two runs in three

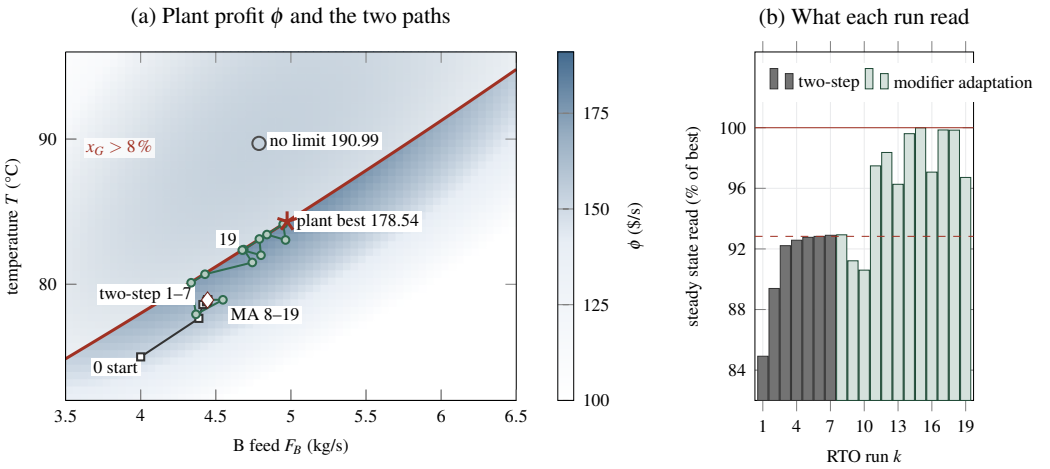

Figure 11.2 The lesson of case 3, computed from the application's own functions. (a) The plant's steady profit over B feed and temperature, with the 8 % waste limit (thick solid line) and the region beyond it faded, the plant's best point (star, 178.54 $/s), the two-step fixed point where the refitted model's best point sits (diamond), and the unconstrained optimum (open circle, 190.99 $/s). Two-step (squares) from the start point 0 stops at run 7, and modifier adaptation (filled circles) from 8 h climbs to run 19 at (4.676 kg/s, 82.36 °C), 99.46 % of the best at 20 h. (b) The steady state each run read, as a share of the best, two-step as dark bars and modifier adaptation as light bars, with the static two-step value of 92.83 % dashed and the best solid. Every third modifier-adaptation run reads lower because the run before it perturbed off the line of its predecessors.

read 98.61 to 100.09 %, with runs 13, 14, 16, 17 and 20 within 0.1 %, and every third, after a move the perturbation rule pushed off the line, reads 96.66 to 97.09 % (runs 9, 12, 15 and 18). Runs 10 to 19 average 98.96 %. Over runs 8 to 19 the model's own temperature slope is -2.56 to -1.08 $/s per degree while the plant estimate is $+3.13$ to $+5.41$, so the model would have cooled the reactor and the correction carried it hotter. The dips are the price of the perturbation and scale with it. With $\delta = 0$ the estimate collapses and runs 10 to 19 average 81.44 %, with $\delta = 0.5$ °C they average 99.33 %, and with 2 °C they average 97.99 % with dips to 93.35 %. The shipped 1 °C earns less here than 0.5 °C, and the application's development record keeps it for the margin it leaves the waste limit under noise in the challenge.

> **+ DID YOU KNOW? The Price of Probing, in Another Plant**
>
> The desalter of Volume 11, Chapter 2 pays the same price: its extremum-seeking controller perturbs the mixing pressure drop, and with its learning switched off the probing alone leaves more salt than a fixed setting. On the Williams–Otto reactor, a 2022 review found such model-free estimators reach the optimum without a model, but significantly more slowly than model-based ones [23].

11.6.3 Noise, a New Limit, Drift and the Period

In case 5 the analyzer noise rises from 0.2 to 1 % of the reading at 6 h. At 1 % one report of P scatters by about 0.0011 in mass fraction and the mean of six by 0.00044, about 3.4 $/s of profit, and a difference of two such means over 1 °C is about 4.85 $/s per degree, as large as the plant's own slope of 3.89 $/s per degree near its best point. As shipped ($K = 0.5$) the plant-change test reads the noise step as a plant change at 7 h, a false alarm (profit +4.0 against +13.6 $/s predicted), runs 14 to 16 read 91.7, 89.2 and 90.9 %, and Profit so far ends at 97.18 % at 24 h. With K lowered to 0.3 at 6.2 h it ends at 98.40 %, and with 0.2 at 98.68 %. The lesson depends on the noise drawn. Over the default

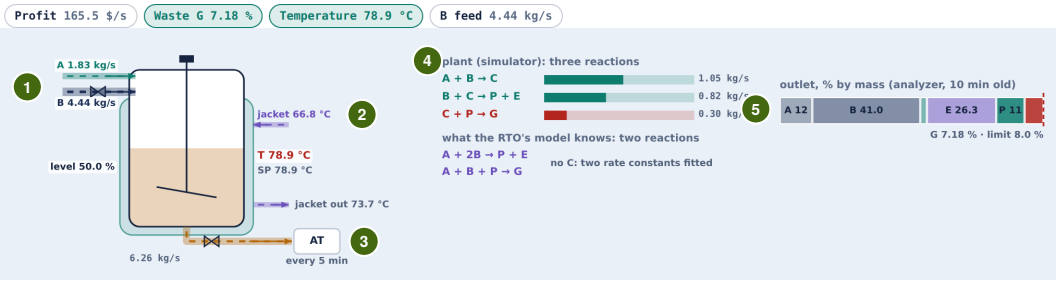

Figure 11.3 Case 3 at $t = 8$ h on the Reactor view, two-step settled after seven hourly runs, cropped to the plant view and its status chips. (1) The measured A and B feeds, 1.83 and 4.44 kg/s. (2) The jacket supply at 66.8 °C and its return at 73.7 °C. (3) The level at 50.0 % and the effluent at 6.26 kg/s to the analyzer. (4) The plant's three reaction rates (1.05, 0.82 and 0.30 kg/s) beside the model's two reactions with no C. (5) The outlet bar, the analyzer's report, 10 min old, with $x_G = 7.18$ % under the 8 % limit.

seed and seeds 1 to 5 the gain from $K = 0.3$ over 0.5 is $+1.22$, $+1.89$, -0.03, $+0.68$, $+0.47$ and $+2.05$ points, so one realization in six shows no benefit, and the case runs at its own seed so that a reader sees the first of these.

In case 6 the waste permit falls from 12 to 8 % at 12 h. The run at 12 h reads 11.53 % waste (11.55 % true), and the offset ε, measured at one steady state and taken whole, puts the corrected model's limit onto the plant's at once. The next two steady states hold 8.10 and 7.97 % waste at 99.66 and 99.58 % of the new best. In case 7 the waste reaction slows to 80 % at 6 h, unannounced, and the best possible rises from 178.54 to 242.88 $/s at (5.261 kg/s, 89.24 °C), 36.0 % more. The plant-change test fires at 7 h, a real change this time, and the runs at 9 and 10 h read 99.37 and 100.07 % of the new best. Profit so far at 24 h is 98.50 % against 95.26 % with the layer off, where the plant stays at 228.6 $/s with 6.96 % waste and its limit no longer binding.

Case 8 runs modifier adaptation every 15 min with the check off, so each run reads a plant still moving from the last. It makes 94 runs in 24 h, ends at 87.78 %, and leaves the waste up to 1.97 points over its limit at steady state. With the check on it makes 38 runs and earns 97.68 % with at most 0.21 point over, and hourly with the check off 23 runs and 96.60 % with at most 0.04. In case 3 the detector says steady only while the true fractions move less than 0.0073 h^{-1}, and under 1 % noise in case 5 while they move up to 0.0102 h^{-1}, so the detector alone admits a little transient that a longer period removes.

11.6.4 Economic MPC Against Modifier Adaptation

The comparison starts every method at the plant's best point, with the method chosen before the run, applies one change at 12 h and judges the true profit from 12 to 24 h as a share of the clairvoyant best. The changes are a rise of the A feed from 1.83 to 2.2 kg/s, which the A meter reports, a 10 % cut of the P price, which the layer is told at once, and no change, which measures each method's own running cost. Every row uses the default seed and five more, with the same analyzers and noise, and the economic MPC runs with the settings case 9 ships, which are the layer's defaults, retuned for no case. The results are listed in Table 11.5. At the hourly period the steady-state check makes no difference to either RTO method, because the detector has always passed by the time the period allows a run, so the rows with the check on and off are identical. Holding the setpoints earns more than the best permitted point only by running over the waste limit, for 11.6 h after the feed rise.

On profit the economic MPC ties modifier adaptation, and it does not earn more. After the feed rise it earns 98.98 % against 99.38 %, after the price cut 98.12 % against 97.81 %, and with no change 98.83 % against 98.88 %, each difference inside the spread of the six seeds. The premise that an RTO

Table 11.5 The Economic MPC Against the Layer's Other Methods Over Six Seeds. Profit Is the True Profit From 12 to 24 h as a Percentage of the Best (Mean, With the Range After the Feed Rise). Peak Is the Largest True Waste in That Window and h Over the Longest Time at Steady State Above 8.1 % in 0–24 h, Over the Six Seeds

	A feed rise (metered)			P price cut		No change	
Method	profit %	peak %	h over	profit %	h over	profit %	h over
Off, setpoints held	102.88	9.36	11.58	97.03	0.00	100.02	0.00
Two-step, 1 h, check on	92.15 (92.13–92.18)	8.14	0.57	90.19	0.00	92.84	0.00
Two-step, 1 h, check off	92.15 (92.13–92.18)	8.14	0.57	90.19	0.00	92.84	0.00
MA, 1 h, check on (shipped)	99.38 (98.77–99.79)	9.38	1.27	97.81	0.00	98.88	0.00
MA, 1 h, check off	99.38 (98.77–99.79)	9.38	1.27	97.81	0.00	98.88	0.00
MA, 15 min, check on	99.05 (98.25–99.47)	9.55	2.12	97.13	1.95	99.09	1.72
MA, 15 min, check off	93.45 (68.78–98.74)	10.98	16.21	39.09	17.74	78.23	17.74
Economic MPC	**98.98** (97.95–99.42)	**8.90**	**1.56**	**98.12**	**0.00**	**98.83**	**0.00**
Economic MPC without ε, λ	95.65 (95.59–95.72)	8.24	0.00	100.08	11.42	96.52	0.00
Economic MPC without the holdup	99.10 (98.22–99.50)	9.42	1.46	98.06	0.01	98.86	0.00
Economic MPC without the slope check	99.08 (98.84–99.37)	8.81	1.58	97.93	0.00	98.87	0.00

loses profit while it waits for steady state is not borne out on this plant, where a move settles in about half an hour and an hourly layer gives up little to the wait. Two counterfactual runs of the same controller attribute what it does. Without the modifiers, optimizing its fitted model alone, it earns 95.65 % after the feed rise and 96.52 % with no change, and after the price cut it earns 100.08 % only by holding the waste above the limit for 11.4 h, because without ε the model's waste is the one two-step trusts. The profit is therefore modifier adaptation's. Without the holdup, planning on the static map, it earns the same within the seed spread but lets the waste peak at 9.42 % after the feed rise, against 8.90 % with the holdup and 9.38 % for modifier adaptation. The lower peak comes from acting at once on the A meter through a model of how the outlet answers a move. After the feed rise the economic MPC still spends longer just over the limit at steady state, at most 1.56 h against 1.27 h, because both methods learn the new waste only from a steady state read after the change. Modifier adaptation every 15 min with the check on, the cheapest way to make the shipped method react sooner, earns 99.05 % after the feed rise with a waste peak of 9.55 % and up to 2.1 h over, and with the check off it acts as early as the economic MPC but reads transients as steady states, earning 39.09 to 93.45 % across the three changes with up to 17.7 h over. With the integration and control step halved, on three seeds, the peaks after the feed rise rank the same way, 8.83 % against 9.45 %. Figure 11.4 draws the profits and the waste peaks of Table 11.5.

🕮 **Pause.** Try Exercise 11.3 before reading on.

In case 9 itself, at the application's seed, the A feed rises at 12 h and the P price falls 10 % at 24 h. The economic MPC moves 5.1 min after the rise, at its next decision, and first lowers the temperature setpoint to 83.19 and then 82.64 °C while it raises the B feed, which holds the waste down while the extra A arrives, before it takes the setpoints to (5.459 kg/s, 84.98 °C) by 12.34 h, as shown in Figure 11.5. Modifier adaptation ran at 12.00 h, from a window the rise had not yet reached, and holds that move for an hour. The true waste peaks at 8.63 % under the economic MPC and 9.33 % under modifier adaptation, and at steady state the two spend 1.53 h and 1.23 h above 8.1 % in the 36 h run. Over 12 to 24 h modifier adaptation earns 99.45 % of the best and the economic MPC 99.22 %, over 24 to 36 h, after the price cut, 96.72 % and 97.25 %, and Profit so far at 30 h reads 98.46 % against 98.45 %. At 24 h it reads 98.66 % under the economic MPC and 98.83 % under modifier adaptation. The same case without the modifiers earns 95.62 % over 12 to 24 h, and without the holdup its waste peaks at 9.39 %.

A change that no instrument reports removes the economic MPC's advantage. Case 7, whose drift

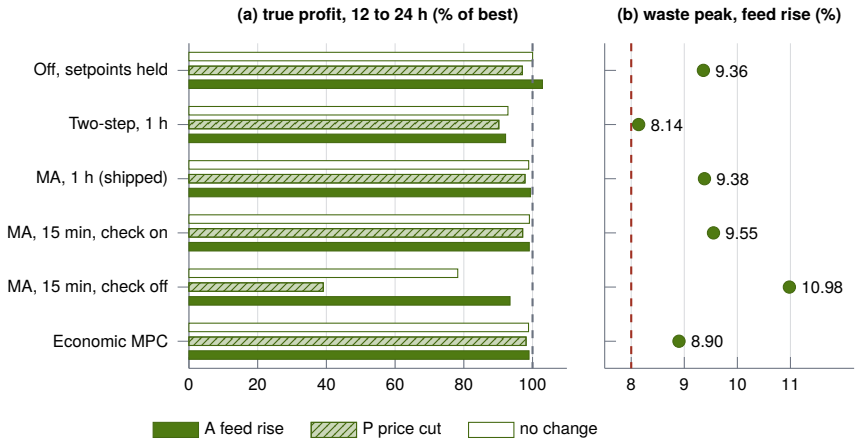

Figure 11.4 Table 11.5 drawn as a chart, mean of six seeds. (a) The true profit from 12 to 24 h as a share of the best after the metered A feed rise (dark bars), the P price cut (hatched) and with no change (white), with the best dashed. Two-step stands for its identical rows with the check on and off. (b) The largest true waste after the feed rise, with the 8 % limit dashed. Notice that the economic MPC's bars match those of hourly modifier adaptation within a point while its waste peak is lower. Two-step keeps the waste lowest but stops short on profit. Only holding the setpoints earns more than the best, and only by staying over the limit.

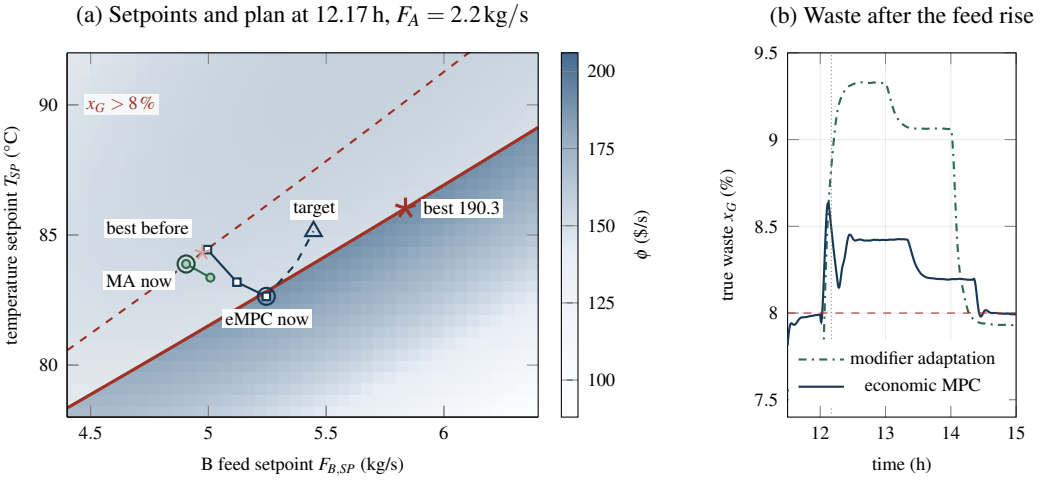

Figure 11.5 Case 9 at $t = 12.17$ h, ten minutes after the A feed rose to 2.2 kg/s, drawn from the engine's data for the Landscape view at the application's seed. (a) The plant's profit map and 8 % waste boundary (thick solid line) at the new feed, the boundary before the rise (dashed), and the plant's best point before (pale star, 178.54 $/s) and after (star, 190.33 $/s). The economic MPC (squares) has moved twice since 12 h, to (5.246 kg/s, 82.64 °C), and plans (dashed) to its target (triangle) at (5.446 kg/s, 85.14 °C) with a predicted corrected waste of 8.00 % at the end of its horizon. Modifier adaptation (filled circles) holds the point it chose at 12.00 h, (4.905 kg/s, 83.89 °C), until 13.00 h. (b) The true waste of both runs, peaking at 8.63 % under the economic MPC (solid) and 9.33 % under modifier adaptation (dash-dotted), with the limit dashed and the instant of (a) dotted.

at 6 h the layer is not told of, run to 24 h under both methods, each chosen before the run, gives a peak of the true waste after 6 h of 8.21 % under the economic MPC against 8.15 % under modifier

Table 11.6 Hours at Steady State With the True Waste Above 8.1 %, Mean (Max) Over Six Seeds, and Mean Profit of the Whole Run. eMPC, First: Without the Slope Check

Case	MA	eMPC, first	eMPC	profit %, MA, eMPC
1 and 3	0.00 (0.00)	0.22 (0.72)	0.00 (0.00)	97.41, 96.54
2	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	98.68, 98.59
4	0.18 (0.58)	0.43 (1.44)	0.24 (1.44)	98.22, 98.04
5	1.39 (3.34)	1.32 (3.18)	0.84 (2.59)	96.02, 93.80
6	0.40 (1.06)	0.73 (1.47)	0.32 (0.67)	98.35, 97.70
7	1.92 (4.18)	1.59 (3.65)	0.00 (0.00)	98.15, 98.51
8	23.99 (32.50)	12.84 (28.44)	24.32 (30.69)	83.09, 81.20
10, shipped	5.71 (6.30)	4.76 (6.49)	4.45 (6.04)	98.40, 98.31
10, suggested	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	97.43, 97.14

adaptation, and Profit so far of 97.72 % against 98.50 %. Both methods must wait for steady states to see what the model gets wrong, so acting at once gains nothing. Cases 9 and 7 together separate a disturbance the controller is told of from a plant change it can only infer.

> **TRY IT IN THE APP** **9 · Economic MPC Against Modifier Adaptation**
>
> Run the scenario to 30 h and watch the Waste G trend after the A feed rises at 12 h. Reset, set RTO method to Modifier adaptation and run again. The economic MPC moves at once and holds the true waste peak to 8.63 % against 9.33 %, yet Profit so far at 30 h reads 98.46 % against 98.45 %. The trend, logged every 0.05 h, peaks slightly lower.

11.6.5 Worst Cases and the Slope Check

Because the economic MPC shares the waste limit with every other method, each existing case was also run under it and under modifier adaptation with that case's own settings, events and duration (36 h, the challenge 100 h), over six seeds, and the time the true waste spends above 8.1 % at steady state is listed in Table 11.6. The first version of the plan took modifier adaptation's slopes whenever the layer formed them. At the application's seed it then spent 1.59 h over in case 7, where modifier adaptation spends none, and 26.84 h against 6.53 h in case 8. The cause in case 7 is the unarmed window of the change test. The plan's steady states at 5.25, 6.67 and 7.84 h fell across the drift at 6 h, the slope formed at 7.84 h gave $\lambda_{\phi,1} = +90.4\,\$/\mathrm{s}$ per kg/s, and the test that fired at the next read cleared the history but not that λ. Modifier adaptation has the same weakness. With the drift moved to 3, 4 or 4.5 h it spends 2.08, 2.06 and 2.05 h over, and over six seeds of case 7 it spends 1.92 h on average and up to 4.18 h, with the true waste reaching 11.12 % at one seed.

With the slope check of Section 11.5.6, case 7 spends no time over at any of the six seeds, case 8 spends 0.09 h at the application's seed, and over six seeds the economic MPC is at or under modifier adaptation's mean in seven of the nine rows of Table 11.6. The two exceptions are case 4, 0.24 against 0.18 h, and case 8, 24.32 against 23.99 h, a configuration that reads transients under every method. Cases 1 and 3 share their start, so with the method chosen before the run they are one row. The check has a price. Under the fivefold analyzer noise of case 5 more than half of the slopes fail it (10.2 of 18 on average), learning slows, and the run earns 93.80 % of the best against 96.02 %, and 89.15 % against 97.35 % at the application's seed. It costs a little in cases 1, 3 and 6 as well (96.54 against 97.41 %, 97.70 against 98.35 %), and gains in case 7 (98.51 against 98.15 %). On the challenge's suggested settings both methods keep the waste at steady state within the tolerance at every seed. The economic MPC meets both of the challenge's targets at the application's seed with 97.18 %, below modifier adaptation's 97.47 %, but earns as little as 96.88 % at another seed, short of the 97 % target.

11.6.6 The Challenge and Its Verdict

The challenge runs 100 h from (4.00 kg/s, 75 °C), with the waste reaction slowed to 80 % at 30 h and the P price cut 10 % at 60 h, so the best point moves from 178.5 \$/s at (4.97, 84.3) to 242.9 at

Table 11.7 The Challenge Over 100 h. Waste Is the Worst Excess Over 8 % at Steady State, in Points

Settings	Default seed	Other seeds
two-step, 0.25 h (shipped)	fails, 92.83 %	seeds 0–5: 92.82–92.85 %
two-step, 19 settings	fail, 86.10–92.89 %	—
MA, 1 h, K 0.2, b 0.15 (suggested)	passes, 97.47 %, +0.000	31/31 pass, 97.18–97.53 %, at $\Delta t/2$ 31/31, 97.10–97.53 %
MA, 1 h, K 0.3, b 0	passes, 98.19 %, +0.080	21/31 pass, at $\Delta t/2$ 16/31
MA, 1 h, K 0.5	passes, 98.24 %, +0.055	seeds 3, 4, 5 fail, +0.103 to +0.163
MA, 0.25 h, K 0.3	fails, 98.62 %, 4.97 h over	seeds 0–5 fail
RTO off	fails, 78.54 %	—

(5.26, 89.2) and then 152.4 at (4.77, 85.7). Its target is a cumulative true profit of at least 97 % of the clairvoyant best's, switched the instant a condition changes, and a true waste that never exceeds 8.1 % at steady state, which is the case's 8 % limit with a 0.1-point tolerance for an optimizer that holds the limit through a noisy analyzer. Steady state means the true fractions of P and G both move less than $0.002\,\mathrm{h}^{-1}$. A lined-out plant moves at most $0.001\,55\,\mathrm{h}^{-1}$ under its own noise, so it is always judged, and judgment resumes 27 min after a 3 °C move. Two verdict rules were added in a technical review of the application. The waste is judged against the case's 8 % and not against the learner's limit control, because with that control raised to 9 % the suggested settings had earned the certificate with the true waste up to 9.36 %. A run whose limit, disturbances or model constants leave the case's values beyond its two scripted events, or in which an event button is pressed, reads "not graded". With the limit at 9 % the run now reads not graded, and with a learner's back-off of 0.3 point, which is allowed, it is graded and fails at 96.31 %. The outcomes are listed in Table 11.7.

As shipped, two-step every 15 min earns 92.83 %. It cannot pass with any setting tried: over periods of 0.25 to 4 h, the check on and off, windows of 6 and 12 reports, a back-off and a tighter threshold, 19 settings in all, it earns 86.10 to 92.89 %, the lowest with a 12-report window every 15 min and the check off, which reads transients. Its fixed point belongs to its model. The application first suggested modifier adaptation hourly with $K = 0.3$, which passes at the default seed with 98.19 %. The technical review ran it at 31 noise seeds and found it passing at only 21, and at 16 of 31 with the step halved, the default seed itself failing at $\Delta t/2$ by 0.135 point. Every failure at Δt began between 5.39 and 8.61 h, after runs made before the gradient estimates were good, and left the waste at up to 8.28 % for 19 to 68 min. The suggestion became $K = 0.2$ with a 0.15-point waste back-off, a margin of the kind Chapter 10 keeps from a limit that its laboratory sees late. It was the only one of fifteen settings in the review's sweep to pass all 62 runs, at 97.18 to 97.53 % with at most 0.052 point over at Δt and 97.10 to 97.53 % with at most 0.059 at $\Delta t/2$, and it passes at $\Delta t/4$. At the default seed it earns 97.47 %, against 98.19 % for the old suggestion, so the margin on waste is paid for in profit. On the suggested settings the 100 h run takes about 0.2 s of computing in Node on the development host, and at 10 h per second the app plays it in ten seconds.

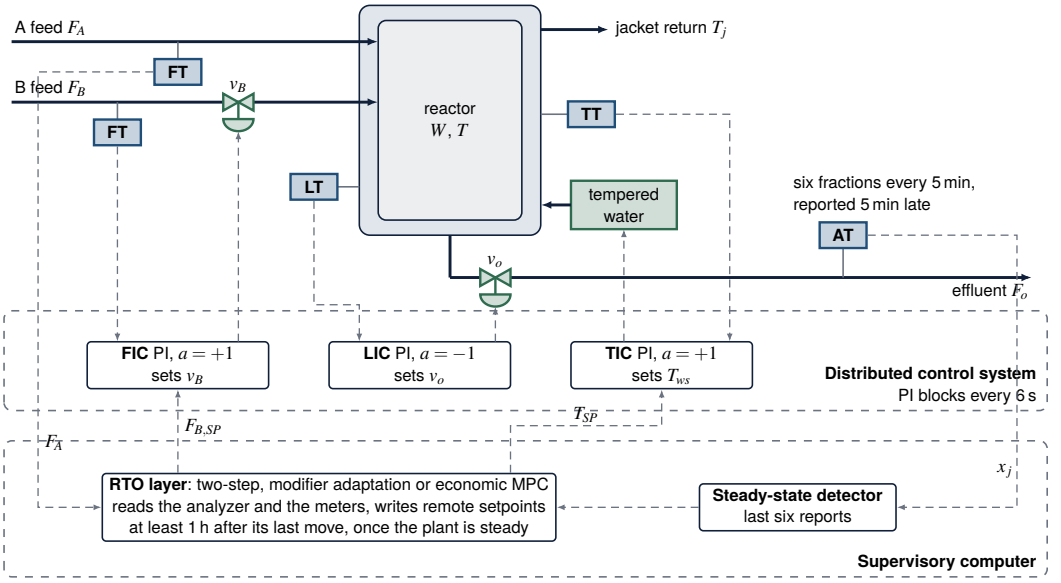

Figure 11.6 The loops and the optimization layer of the Williams–Otto reactor drawn with their hardware. Three PI blocks of the distributed control system hold the level, the B feed and the reactor temperature, the last through a tempered-water unit that sets the jacket supply. A supervisory computer reads the analyzer and the A feed meter and writes two remote setpoints, T_{SP} and $F_{B,SP}$. Notice that the layer reads no temperature or valve: it sees the plant through the analyzer, whose reports are minutes old, and it moves the plant only through the setpoints of loops that do the regulating.

DID YOU KNOW? The Most Profitable Setting Fails

From the simulator: in the 100 h challenge of Table 11.7 the highest profit, 98.62 % of the best, belongs to modifier adaptation run every 15 min with a filter gain of 0.3, and it fails, with the true waste over the tolerance for 4.97 h. With the layer off the plant earns 78.54 %.

11.7 Technology in Practice

The loops and the layer of Section 11.5 run on real instruments, valves and computers, drawn together in Figure 11.6.

What Measures It

The layer's eyes are the analyzer. In the model it reports six mass fractions every 5 min, each one a cycle after its sample was drawn, with a relative noise of 0.2 %. That timing is a process chromatograph's. Chromatographs separate the species in time, so cycle times of minutes are common and some dead time is inevitable [24], as Chapter 2 describes. A spectroscopic analyzer is the faster alternative for a liquid product.

TECHNOLOGY Sensor: A Near-Infrared Analyzer on the Product

A near-infrared analyzer reads the overtone and combination bands of bonds such as C–H and O–H, so nearly any molecule that carries them shows up in the spectrum. It is fast, one minute or less per sample, does not destroy the sample and suits in-line use, and optical fibers let one instrument watch several points of a plant up to 100 m away [25]. Its numbers come from a multivariate model calibrated on samples whose composition a reference method has determined, so it can never be better than that reference [25]. Its weak

point is the sample. When the line that feeds an analyzer plugs, the instrument measures old, stale sample, has no way of knowing it, and reports meaningless numbers [24].

The optimizer averages six reports before it trusts a steady state, and a move takes about half an hour to read, against lags of 0.3 and 0.1 min on the thermowells and flow meters.

What Moves It

The temperature loop does not move a valve in the model. Its output is the temperature of the tempered water sent to the jacket, which follows its command with a lag of 2 min. A higher supply warms the reactor, so the loop is reverse acting. The action follows the final element: a reactor whose loop throttles an air-to-open steam valve to its jacket needs reverse action, and a vessel whose loop opens a chilled glycol valve needs direct action [24]. The B feed and effluent valves follow their commands with a lag of 0.5 min.

Where the Controller Runs

The three PI blocks run in a distributed control system, every 6 s in the model. Its supervisory level sends set points to the field controllers and requests their data, at cycle times from milliseconds to minutes [26].

TECHNOLOGY Platform: A Supervisory Computer Writing Setpoints

An optimization layer runs on a supervisory computer. It sends setpoints over a network to the loop controllers and reads their measurements back, a two-layer system with the PID loops below and the computer and its process models above. Each loop then has three modes: manual, automatic with a local setpoint from the operator, and automatic with a remote setpoint from the computer [24]. In the classical hierarchy a unit's optimizer runs every hour, the predictive controller every minute and the basic loops every second [27]. Because the loops do the regulating, the plant keeps working when the computer stops [11]. On a trend the setpoints simply go flat while every loop holds them.

In this chapter the layer runs no sooner than 1 h after its last move and only once the detector has passed, and the economic MPC decides every 5 min. TIC and FIC take their setpoints from it, LIC keeps a setpoint of its own, and the header's Manual pauses the layer and places all three loops on their stations.

IN PRACTICE Squeeze the Variance, Then Shift the Setpoint

An output that must not cross its limit, even for a moment, is run with a back-off: its setpoint sits a margin inside the limit, and any back-off from an active constraint costs money, the constraint's shadow price times the margin [28]. The rule of thumb that follows is called squeeze and shift: improve control to squeeze the variance, then shift the setpoint closer to the limit [28]. A margin is added for imperfect control and noise, and a constraint whose shadow price is small can be given a large one [23]. The waste back-off of 0.15 point in this chapter's challenge is such a margin, and its price is measured in Section 11.6.

What Is Modernizing It

Three developments bear on this layer. The first, feedback optimizing control, turns the economic objective into controlled variables, needs no numerical optimization online, and avoids the wait for steady state by using measurements as they come [23]. A 2022 review compares such methods on the Williams–Otto reactor with the waste held at 0.08 by the reactor temperature, the limit of this chapter [23]. The second learns the mismatch instead of modeling it. The same review corrects the two-reaction model with a radial basis function network fitted to the gap between measured and predicted cost [23], and Gaussian processes serve as modifiers in schemes that steer their own probing [7]. The third is economic model predictive control, which this chapter tests. One review

advises it only for small problems or where the optimization and control layers cannot be separated in time scale, because for practical implementation and tuning the combined layer is rarely a good solution [28]. On this plant it tied modifier adaptation on profit.

11.8 Limitations

The plant is the reactor-only benchmark, perfectly mixed and in mass, with no separator or recycle, and its energy balance, jacket, heats of reaction and valves are teaching values that set the dynamics and not the steady states. The mismatch is the literature's two-reaction model, a single structural error of a known kind, while an industrial model is wrong in many places at once. Plant gradients come only from past steady states, so the layer has no dedicated plant tests, no Broyden or quadratic surrogate updates and no dual formulation with a step bound [3, 17]. The perturbation rule, the move limits, the plant-change test and the elastic waste row are the application's design and were tuned on this plant. Every result is a run at the application's seed, checked at 31 seeds for the challenge and at six for the noise case, whose lesson is seed dependent. The waste tolerance of 0.1 point and the steady-state definition of the verdict are the application's choices. The analyzer is unbiased and never fails, and the prices and limit reach the optimizer exactly, which an industrial layer cannot assume.

> ⚠ **OPEN CHALLENGE** **Why Plants Seldom Keep Their Optimizers**
>
> The business benefits of real-time optimization are not disputed, yet its use has stayed relatively low [7]. A review names four technical reasons: the cost of building the model, the wait for a steady state before each model update, numerical failures of large optimizations, and conflicts with the planning layer above [23]. An optimizer needs a team of skilled engineers to maintain and monitor it, and without one the application is often turned off [23]. Progress would be an optimizer that learns from routine data while the plant moves, flags its own model when it drifts, and stays switched on without a dedicated team.

The economic MPC's model leaves out the intermediate C and the third reaction, which is the shared mismatch and stays the lesson, the energy balance and jacket, which it replaces by a lag, the temperature loop's overshoot, the analyzer's delay, and any retuning of the loops after the step tests, which becomes mismatch in the lags. Its composition is an open-loop observer corrected only through ε at steady states, so it learns the drift, as modifier adaptation does, only from steady states. It is linearized once per decision at its target and solved as one quadratic program, not as a nonlinear program. The arrival rows, the slope check, the dead band and the waste prices are the application's design, chosen on this plant before the comparison. During a fallback the layer moves with its own modifiers, which may include a slope the plan had rejected, a path that no run reported here exercised.

> ⚠ **OPEN CHALLENGE** **Slopes From Noisy Data**
>
> Inferring the plant's gradients from noisy measurements is hard, and it decides how well and how reliably modifier adaptation works. Newer schemes learn the mismatch between model and plant with Gaussian processes and steer their probing with an acquisition function, but their theory and practical performance are yet to be analyzed in depth, and probing that favors the feasible side of the process constraints is left to future work [7]. The challenge of this chapter shows the stakes: a setting chosen at one seed passed at 21 of 31. Progress would be an estimator that probes only as hard as the noise demands and says how far its slopes can be trusted.

11.9 Conclusion

This chapter ran the Williams–Otto reactor under three loops and an optimization layer that reads only its analyzer and meters. Its steady states are the published ones to every digit an independent solver prints, 178.540 \$/s at (4.9747 kg/s, 84.322 °C) under the 8 % waste limit. A model with two reactions, refitted at every steady state, takes the plant to (4.4454 kg/s, 78.915 °C) and stops at 92.83 % of that from any start, where its own slopes vanish and the plant's are -24.4 and $+5.76$. Modifier adaptation

reaches 100.000 % with exact gradients and averages 98.96 % on the running plant, every third run about three points short because it must perturb.

A filter gain, the detector, the period and a plant-change test decide how much of that survives noise, drift and a short period, and the challenge's suggested settings pass at 31 of 31 seeds only with a waste back-off of 0.15 point, after a setting without it passed at 21. An economic MPC on the same model and modifiers ties modifier adaptation on profit, 98.98 % against 99.38 % after a metered feed rise over six seeds, and owes its lower waste peak, 8.90 % against 9.38 %, to its dynamic model.

The safeguards are the application's own and the mismatch is a single known one. A fuller study would add a dual formulation with a step bound, gradient estimation by dedicated plant tests, an analyzer that drifts or fails, and the separator and recycle of the original plant. The simulator shows that an optimizer's model must be right about slopes, not only about values.

KEY POINTS

- An optimizer whose model is wrong in structure can stop short of the plant's best, here at 92.83 % of it, where its model's slopes vanish.
- Modifier adaptation corrects the model with plant gradients and reaches 100.000 % with exact ones, 98.96 % on average on the running plant.
- Every third run is about three points short, the price of the perturbation the gradient estimate needs.
- The challenge's suggested settings pass at 31 of 31 seeds only with a waste back-off of 0.15 point.
- An economic MPC on the same model ties on profit and owes its lower waste peak to its dynamic model.

Where else this applies. An optimum on a moving constraint in Volume 9, Chapter 8, and a slope learned by probing the plant in Volume 11, Chapter 2.

References

[1] P. D. Roberts, "An algorithm for steady-state system optimization and parameter estimation," *International Journal of Systems Science*, vol. 10, no. 7, pp. 719–734, 1979.

[2] J. F. Forbes and T. E. Marlin, "Model accuracy for economic optimizing controllers: The bias update case," *Industrial & Engineering Chemistry Research*, vol. 33, no. 8, pp. 1919–1929, 1994.

[3] A. Marchetti, G. François, T. Faulwasser, and D. Bonvin, "Modifier adaptation for real-time optimization—methods and applications," *Processes*, vol. 4, no. 4, p. 55, 2016.

[4] A. Marchetti, B. Chachuat, and D. Bonvin, "Modifier-adaptation methodology for real-time optimization," *Industrial & Engineering Chemistry Research*, vol. 48, no. 13, pp. 6022–6033, 2009.

[5] S. Cao and R. R. Rhinehart, "An efficient method for on-line identification of steady state," *Journal of Process Control*, vol. 5, no. 6, pp. 363–374, 1995.

[6] T. J. Williams and R. E. Otto, "A generalized chemical processing model for the investigation of computer control," *Transactions of the American Institute of Electrical Engineers, Part I: Communication and Electronics*, vol. 79, no. 5, pp. 458–473, 1960.

[7] E. A. del Rio Chanona, P. Petsagkourakis, E. Bradford, J. E. Alves Graciano, and B. Chachuat, "Real-time optimization meets Bayesian optimization and derivative-free optimization: A tale of modifier adaptation," *Computers & Chemical Engineering*, vol. 147, p. 107249, 2021.

[8] F. A. Cubillos, G. Acuña, and E. L. Lima, "Real-time process optimization based on grey-box neural models," *Brazilian Journal of Chemical Engineering*, vol. 24, no. 3, pp. 433–443, 2007.

[9] J. E. A. Graciano, "Real time optimization in chemical process: evaluation of strategies, improvements and industrial application," Ph.D. dissertation, Escola Politécnica, Universidade de São Paulo, São Paulo, Brazil, 2016.

[10] J. O. A. Matias, "Real-time optimization with persistent parameter adaptation using online parameter estimation," Ph.D. dissertation, Escola Politécnica, Universidade de São Paulo, São Paulo, Brazil, 2019.

[11] M. Pitt, "Taking a look back at control: Part 2," *The Chemical Engineer*, Oct. 2023, published online 26 October 2023. [Online]. Available: https://www.thechemicalengineer.com/features/taking-a-look-back-at-control-part-2/

[12] Computer History Museum, "Computers in control," CHM Revolution, online exhibition, section Real-Time Computing, accessed 18 November 2026. [Online]. Available: https://www.computerhistory.org/revolution/real-time-computing/6/130

[13] C. R. Cutler and R. T. Perry, "Real time optimization with multivariable control is required to maximize profits," *Computers & Chemical Engineering*, vol. 7, no. 5, pp. 663–667, 1983.

[14] C. Y. Chen and B. Joseph, "On-line optimization using a two-phase approach: an application study," *Industrial & Engineering Chemistry Research*, vol. 26, no. 9, pp. 1924–1930, 1987.

[15] M. L. Darby, M. Nikolaou, J. Jones, and D. Nicholson, "RTO: An overview and assessment of current practice," *Journal of Process Control*, vol. 21, no. 6, pp. 874–884, 2011.

[16] B. Chachuat, B. Srinivasan, and D. Bonvin, "Adaptation strategies for real-time optimization," *Computers & Chemical Engineering*, vol. 33, no. 10, pp. 1557–1567, 2009.

[17] A. Marchetti, B. Chachuat, and D. Bonvin, "A dual modifier-adaptation approach for real-time optimization," *Journal of Process Control*, vol. 20, no. 9, pp. 1027–1037, 2010.

[18] E. A. Rodger and B. Chachuat, "Design methodology of modifier adaptation for on-line optimization of uncertain processes," *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 4113–4118, 2011.

[19] G. François and D. Bonvin, "Use of transient measurements for the optimization of steady-state performance via modifier adaptation," *Industrial & Engineering Chemistry Research*, vol. 53, no. 13, pp. 5148–5159, 2014.

[20] D. Angeli, R. Amrit, and J. B. Rawlings, "On average performance and stability of economic model predictive control," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1615–1626, 2012.

[21] 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.

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

[23] D. Krishnamoorthy and S. Skogestad, "Real-time optimization as a feedback control problem: A review," *Computers & Chemical Engineering*, vol. 161, p. 107723, 2022.

[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] C. Pasquini, "Near infrared spectroscopy: fundamentals, practical aspects and analytical applications," *Journal of the Brazilian Chemical Society*, vol. 14, no. 2, pp. 198–219, 2003.

[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] 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.

[28] S. Skogestad, "Advanced control using decomposition and simple elements," *Annual Reviews in Control*, vol. 56, p. 100903, 2023.

11.10 Exercises

11.1 ○ Concept check · 2 min

In case 3 the two-step layer stops at (4.4454 kg/s, 78.915 °C), where the plant earns 92.83 % of its best with 7.19 % true waste under the 8 % limit, while the model refitted there predicts 180.92 $/s. What stops the layer at that point?

(a) The waste limit, which the fitted model believes is active there

(b) The move box, which caps every run at 0.5 kg/s and 5 °C

(c) Analyzer noise, which hides the remaining slope from the least-squares fit

(d) The fitted model's slopes vanish there, though the plant's slopes do not

11.2 ○ Concept check · 1 min

Why does the application's modifier adaptation keep each new operating point at least δ off the line through its last two points?

(a) Three points on one line give no estimate of the plant's slope across that line
(b) The waste offset ε can be measured only after a move of at least δ
(c) The steady-state detector passes only after a setpoint move of δ or more
(d) The SQP needs a step of at least δ to leave a stationary point

11.3 Concept check · 2 min

After the metered A feed rise of Table 11.5, the true waste peaks at 8.90 % under the economic MPC and at 9.38 % under modifier adaptation. The same controller without the reactor's holdup in its model peaks at 9.42 %. What gives the economic MPC its lower peak?

(a) Its modifiers, which carry the plant's slopes into the plan
(b) Acting at once on the A meter through a model of how the outlet answers a move
(c) Its slope check, which rejects gradients formed across the change
(d) Its soft waste ceiling, priced far above what a point of waste earns

11.4 Theory and calculation · 6 min

Near the plant's best point the outlet holds a P fraction of 0.1079 and $F_A + F_B = 6.80\,\mathrm{kg/s}$, with $p_P = 1143.38\,\$/\mathrm{kg}$. In case 5 the analyzer noise is 1 % of the reading (one standard deviation), and the layer averages $N = 6$ reports drawn independently.

(a) Compute the standard deviation of one report of x_P, of the mean of six, and of the profit that mean implies through the P term of Equation (11.4). *Checkpoint:* 0.00108, 0.00044 and 3.4 \$/s.

(b) A slope is the difference of two such means taken 1 °C apart. Compute its standard deviation in \$/s per degree and compare it with the plant's own slope of 3.89 \$/s per degree near its best point. *Checkpoint:* About 4.84 \$/s per degree, as large as the slope.

(c) The gradient modifiers follow $\lambda \leftarrow (1-K)\lambda + K(\nabla_p - \hat{\nabla})$. How many runs does it take before a step change in the true correction is 90 % taken up, for $K = 0.5$, 0.3 and 0.2? How many hours is that at the default period? *Checkpoint:* 4, 7 and 11 runs; 4, 7 and 11 h.

11.5 Theory and calculation · 6 min

At the best permitted point, 4.9747 kg/s of B and 84.322 °C with $F_A = 1.8275\,\mathrm{kg/s}$, the reactions release 439.2 kW. Use the energy balance of Equation (11.3) at steady state and the constants of Table 11.2.

(a) Compute the heat taken up in warming the feeds from T_f to T, and the residence time $W_0/(F_A + F_B)$. *Checkpoint:* 88.2 kW; 5.16 min.

(b) Compute the heat the jacket removes and the jacket temperature T_j. *Checkpoint:* 351.0 kW; 79.93 °C.

(c) Compute the tempered-water supply T_{ws} the jacket needs, and say why none of these constants moves the plant's best point. *Checkpoint:* 74.08 °C.

11.6 Theory and calculation · 6 min

The economic MPC of Section 11.5.6 decides every $T_s = 5\,\mathrm{min}$ over $N = 24$ steps grouped into move blocks of 1, 1, 1, 1, 2, 2, 3, 4 and 9 steps, with rate limits of 0.125 kg/s and 1.25 °C a decision inside a move box of 0.5 kg/s and 5 °C either side of the last steady state read.

(a) Check that the blocks cover the horizon, give its length, and count the moves each setpoint may make in one plan. *Checkpoint:* 24 steps, 2 h; 9 moves each.

(b) How many decisions does a setpoint need to cross the move box corner to corner, and how long is that in minutes? Can one plan reach such a target before its horizon ends? *Checkpoint:* 8 decisions, 40 min; yes.

(c) The ceiling's slack costs $\rho_1 = 10^8$ per unit of mass fraction and $\rho_2 = 10^{11}$ per unit squared, each step. Compute the cost of one point of excess waste held for one step, and compare its linear part with what a point of waste earns near the limit, 6.69 \$/s, over one step. *Checkpoint:* $\$1.1 \times 10^7$; about \$2000 a step, a ratio near 500.

11.7 App challenge · 8 min

Open scenario 1, *1 · Find the best temperature by hand*, at the default seed and the default speed of 1 h per second. Leave the B feed setpoint at 4.00 kg/s. A value typed before Run starts the case lined out at that setpoint.

(a) Type 78.0 into the *Temperature setpoint* before Run and run to 2 h. Read the *Best possible*, *Gap to best* and *Waste G, by mass* tiles, and compute the true profit. Give the gap to within 0.1 point and the waste to within 0.05 point. *Checkpoint:* 178.5 \$/s, 5.82 %, 7.99 %; about 168.1 \$/s.

(b) Reload the case, type 79.0 before Run and run to 2 h. Read the same tiles. Why does the higher profit not count? *Checkpoint:* 3.51 %, 8.45 %, waste tile in alarm.

(c) The completion rule asks for a temperature at which the true profit is at least 165 \$/s and the true waste at most 8.1 % for half an hour. Interpolate the two waste readings to find the highest setpoint on the control's 0.1 °C grid that the rule accepts. *Checkpoint:* 78.2 °C.

11.8 App challenge · 8 min

Run the two methods of Section 11.5 at the default seed and 1 h per second, each with its RTO method chosen before Run. Export the CSV where asked: its time column, `Time (h)`, is in hours, and its first lines start with #.

(a) Open scenario 3, *3 · Two-step RTO settles below the best*, set *RTO method* to *Two-step* and run to 8 h. Read *Gap to best* and *Profit so far* to within 0.15 point, and compare `FB_sp_kgps` and `T_sp_C` in the last row of the CSV with `rto_model_opt_FB_kgps` and `rto_model_opt_T_C`. *Checkpoint:* About 7.1 % and 91.3 %; the two points agree to the third decimal, near (4.447 kg/s, 78.93 °C).

(b) Explain from part (a) why more runs of two-step will not close the gap.

(c) Open scenario 4, *4 · Modifier adaptation climbs to the best*, set *RTO method* to *Modifier adaptation* and run to 20 h. Read *Profit so far* to within 0.05 point. On the RTO view, which of runs 7 to 19 read lowest, and why? *Checkpoint:* 97.66 %; runs 9, 12, 15 and 18.

11.9 App challenge · 9 min

Compare the economic MPC and modifier adaptation after a change the layer is told of and after one it is not. Use the default seed, 2 h per second, and read the peak of the true waste from the truth line of the *Waste G* trend or the CSV column `truth_xG_pct`. Give every reading to within 0.05 point.

(a) Open scenario 9, *9 · Economic MPC against modifier adaptation: a feed rise and a price cut*, as shipped and run to 24 h. Read *Profit so far*, pausing by 24.5 h, and the peak of the true waste after 12 h. Press Reset, set *RTO method* to *Modifier adaptation* and run to 24 h again. *Checkpoint:* 98.66 % and 8.59 %; 98.83 % and 9.33 %.

(b) In the two exported files, find the rows after 12 h where `FB_sp_kgps` or `T_sp_C` first change. Which of those moves could have used the new A feed? *Checkpoint:* Economic MPC from 12.10 h; modifier adaptation at 12.05 and 13.05 h, of which only the second.

(c) Open scenario 7, *7 · Plant drift: the RTO retracks*, set *RTO method* to *Economic MPC* and run to 24 h. Then reload the case, which runs modifier adaptation, and run to 24 h. Read the same two numbers, with the peak after 6 h, and explain the difference from part (a). *Checkpoint:* 97.72 % and 8.17 %; 98.50 % and 8.15 %.

11.10 Application area · 3 min

At the benchmark's prices the best permitted point earns 178.54 \$/s, and two-step stops at 165.74 \$/s, 92.83 % of it. Modifier adaptation averages 98.96 % of the best on the running plant of case 4.

(a) What is the two-step gap worth per hour and per day? *Checkpoint:* About \$46 100 an hour and \$1.11 million a day.

(b) What does modifier adaptation's perturbation cost per hour against the best, and what does it recover per hour against two-step? Name one reason from Section 11.8 why a plant's layer would recover less. *Checkpoint:* About \$6700 and \$39 400 an hour.

11.11 Application area · 4 min

Section 11.8 assumes an unbiased analyzer. Suppose the G channel reads 0.3 point high under modifier adaptation with no back-off. Near the best point the plant's best profit changes by about 10 \$/s for each point the waste limit moves.

(a) Where does the true waste settle, and what does the bias cost? *Checkpoint:* About 7.7 %; about 3 \$/s.

(b) If the channel read 0.3 point low instead, what would happen, and which setting of the layer and which practice protect the permit?

11.12 Review the agent · 5 min

An AI assistant reviews the challenge and writes: "Modifier adaptation hourly with $K = 0.3$ and no back-off passes at the default seed with 98.19 %, more than the suggested setting's 97.47 %. The 0.15-point back-off only costs profit. Ship $K = 0.3$ with $b = 0$." Find the flaw in the recommendation and say how you would check it.

CHAPTER 12

Summary and Outlook

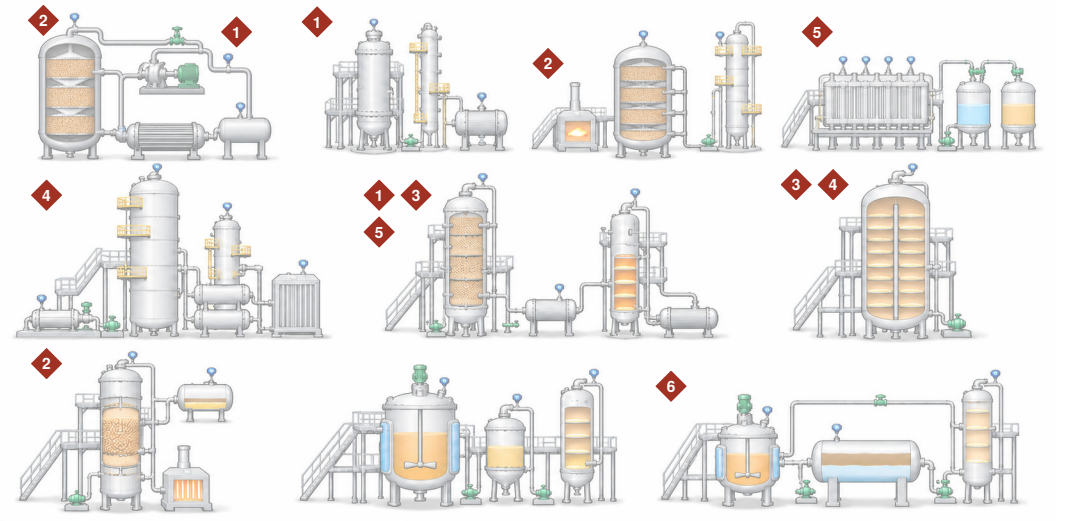

Open gaps 1 inventory structure 2 distance to the cliff 3 broken inferences 4 performance unwatched 5 late obligations 6 noisy gradients

Notice that the capture plant carries three gaps, and that gap 2 marks two converters and a column alike. Illustration, not to scale.

IN PLAIN TERMS What Ten Chemical Plants Taught

The main idea. Each chapter of this book ran one chemical plant in a simulator. In most of them, unused material goes round again, a reaction stops partway, and energy is the largest bill.

The challenge. Material that goes round again turns a small error into a large one that grows for hours. The best way to run a plant often sits next to a cliff, such as a reactor that goes out and will not light again by itself.

Why it matters. A plant can hold every reading on target while it wastes energy or spoils product, so what is controlled matters more than how it is tuned.

Ten chapters put ten chemical plants under control. Table 1.2 listed, for each plant, what recycles, what equilibrium limits, where the energy goes and the structure the chapter implements. This chapter adds what worked, what failed and the measured result, and reads the comparison for the question of the volume: how recycle, equilibrium and energy cost shape each plant's control structure. It then gives a short guide for choosing a structure, 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.

12.1 What the Chapters Showed

Tables 12.1 and 12.2 compare the chapters, five in each. The second column condenses the recycle, equilibrium and energy columns of Table 1.2 into the difficulty each plant posed. The last column gives the measured result with the table, figure or section of the chapter that reports it.

Table 12.1 How Recycle, Equilibrium and Energy Shaped Each Structure, Chapters 2 to 6

Chapter	Difficulty	What worked	What failed	Measured result
2, ammonia loop	Loop gas circulates 5.95 times the makeup. The purge sets the inerts and a 5.1 h time constant. The converter heats its own feed and can go out.	Makeup on pressure, hot spot on quench, and the slow inert and ratio loops in service from the start of the shift.	The fast loops alone: inerts past 16 % at 3.1 h and 21.0 % at the end. A quench opened to 22 % put the converter out.	Challenge passed with a worst pressure error of 0.04 bar and 12.47 % inerts (Section 2.6). Inert IAE 56.6 under prediction against 1843.7 % min supervised (Table 2.7).
3, urea loop	Excess ammonia returns with no exit ($g_R = 1$), so the N/C ratio integrates every feed error. No instrument reads it.	A PI on the N/C inferred from density, with a laboratory bias. A cascade onto the condenser drum through the turndown.	The ratio held on the flows. Wet ammonia drained the loop and the run stopped at 9 h 48 min.	Suggested settings pass 22 of 22 runs, samples 2.965 to 3.033, biuret 0.882 wt% (Table 3.6). The MPC passes with 0.0001 to 0.0054 wt% of margin (Section 3.6).
4, contact acid plant	Four adiabatic beds on an exothermic equilibrium. An ignition floor near 395 °C. No recycle in the gas.	Four bed inlet loops on the bypasses, and the dilution water on ratio with an analyzer trim.	Bypasses held through a turndown to 60 %: bed 1 out at 2 h 17 min, then the others. Either half of the fix alone.	Conversion 99.878 against 95.969 % shipped, stack at most 260 against 48 468 ppmv. The MPC reaches 99.904 % with bed 1 at 409.4 °C (Table 4.6).
5, chlor-alkali cell	No recycle and no equilibrium. The current is the load, power is priced by the hour, and the gas side moves faster than feedback.	Feeds in ratio to the current with trims. Chlorine pressure fed forward to the pressure-difference cascade. A price scheduler.	Feeds in Manual. Feedback alone on the gas side, which trips the rectifier 22 s after the compressor stops.	81.0 t of 86.8 t off specification with the feeds in Manual (Table 5.6). Graded day 72.3 t on specification for $10 908 against $13 853 at constant current (Table 5.7).
6, air separation	One condenser-reboiler couples two columns, 3.11 bar per bar. The compressor runs near its 752 kW rating.	Full cascade, pressure on the guide vanes with a power override. Production given up when the exchanger fouls.	Duty bought back with the pressure setpoint, a 786 kW peak at 5.85 bar. Purity loops blind to a heat leak.	Challenge passed, oxygen within 0.049 %, peak 722 kW (Section 6.6). A heat leak took recovery from 90.8 to 64.9 % while purity rose (Table 6.5).

The columns of the two tables make five points. Figure 12.1 draws one like pair of results from each chapter's row that has one.

A recycle tells the controller which inventory to hold. The three plants with a recycle each needed a loop on the recycled quantity itself. In Chapter 2 a front end 5 % low on ratio settles the loop at 1.80, an amplification of eight, while the pressure never leaves 150.0 bar (Section 2.10). Only the slow inert and ratio loops see it, and they need hours. Put in service at 20 min they finish the challenge at 12.89 % inerts. Waiting two hours gives 13.21 and fails. The urea loop has no exit for its excess ammonia, so a ratio held on two flow meters integrates the meters' error, and only a loop on the inferred N/C held it. In Chapter 7 the return gain is negative, and an inner loop on the lean loading made the outer gain the same at ten minutes and at eight hours. Equation (1.5) predicts all three behaviors from one relation and the value of its return gain.

An equilibrium puts the optimum next to a cliff. Hotter is faster and worse on an exothermic equilibrium. The ammonia converter's own optimum lies near 480 °C and is worth 11 % of production, yet in manual a quench of 15.3 % keeps it lit and 15.4 % puts it out (Section 2.6). The acid plant's best inlet is the coolest one the catalyst can still light, and its MPC held bed 1 within a kelvin of its 410 °C bound. The MTBE column shows the same shape as a fold. A setpoint 1.4 K under tray 12's

Table 12.2 How Recycle, Equilibrium and Energy Shaped Each Structure, Chapters 7 to 11

Chapter	Difficulty	What worked	What failed	Measured result
7, amine capture	Solvent returns in a 70 min lap with $g_R = -3.4$. The reboiler duty is bought at an hourly price.	Capture on circulation over a lean-loading loop on the duty. A supervisor that times regeneration to the price.	A weak inner loop, 3 MW/°C, which leaves the circulation on its stop. A tight level loop that passes the imbalance to the reboiler.	Gain 0.177 at once and 0.040 after 8 h with the duty frozen, 0.177 held with the inner loop (Figure 7.2). Weak loop ends at 93.03 against 94.01 % (Table 7.4).
8, dividing-wall column	No recycle. The spare liquid split moves only the duty, and the minimum moves with the feed.	The four-point structure, a loop on the split, with feedforward from the feed rate.	A fixed split of 0.40. The four-point loop or the feedforward alone. An MPC whose model lacks an input disturbance.	A fixed split costs 17.7 % more steam (Section 8.4). Challenge worst hour 0.3 % above the minimum, 32 of 32 seeds, against 12.4 % shipped (Table 8.6).
9, MTBE column	Etherification on the trays. At 42.5 MW three states convert 98.22, 90.05 and 94.89 %. The duty decides the branch.	The temperature loop on tray 11 at 120 °C, below the low branch's range, so a loop on the low branch cuts the duty.	Tray 12 at 142 °C, 1.4 K under the high branch's peak. On the low branch the same loop raises the duty to 47.76 MW.	Suggested settings keep at least 99.70 mol% MTBE and 97.2 % conversion at seven seeds. Shipped falls to 67.58 mol% (Table 9.7).
10, base oil plant	No recycle and no equilibrium. Yield is traded against a viscosity index the laboratory reports late.	Viscosity on the cut, mid-ranging to the reactor, a bias update and a one-way index cap.	The viscosity loop alone on a blurred cut. Mid-ranging signed the wrong way, which walks the reactor 4.5 °C hotter.	163.5 t downgraded at exactly 6.000 cSt against 5.4 t (Table 10.6). Challenge worst error 0.084 cSt, nothing downgraded, against 63.7 t as loaded (Section 10.5).
11, Williams–Otto RTO	An irreversible reactor whose model lacks a reaction. The optimum lies on the 8 % waste limit.	Modifier adaptation every hour with filter gain 0.2 and a 0.15-point waste back-off.	Two-step, which stops where its own model's slopes vanish. Filter gain 0.3 with no back-off.	Suggested 31 of 31 seeds, 97.47 % of the best. Two-step 86.10 to 92.89 % over 19 settings (Table 11.7).

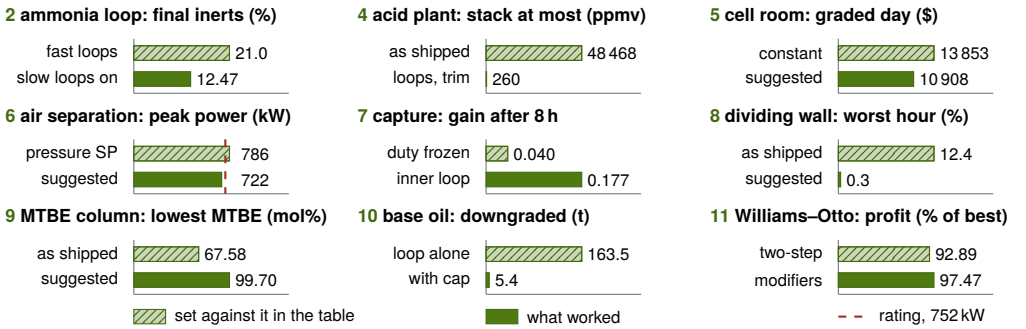

Figure 12.1 One pair of results from each chapter, with the numbers as the "Measured result" and "What failed" columns of Tables 12.1 and 12.2 print them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 7, 9 and 11. Chapter 3 has no like pair in its row and no panel. Notice that each pair sets one structure or measurement against another.

peak lost the column, and the fix moved the measurement to a tray where a loop on the wrong branch acts in the right direction. In each case the structure carried a margin from the cliff or a measurement chosen so that crossing it could be undone.

Energy is spent by the handle no purity loop watches. A dividing-wall column holding all three purities burned 17.7 % more steam at the wrong split. A fouled air separation exchanger cost eight percent of the oxygen while every controlled variable sat on its setpoint and the compressor power fell. A heat leak raised the oxygen purity while it cut the recovery to 64.9 %. The base oil loop held 6.000 cSt exactly and downgraded 163.5 t. The structures that worked closed a loop on the spare handle (the four-point split loop), or guarded the quality the main loop could not see (the index cap). Purity and viscosity are specifications, and recovery and specific energy are performance.

Timing the energy is worth what the storage allows. The chlor-alkali scheduler made a day's caustic for \$9 008 against \$10 617 at constant current, limited mainly by its tank (Table 5.6). The capture supervisor earned a 0.99 % larger margin over a price day at the same captured tonnage, and on a five-hour peak the lean tank reached its floor after 98 min (Section 7.6). The price decides when, and the tank decides how much.

Prediction paid on slow coupled inventories, and a model needs the right slopes. The predictive controller cut the ammonia inert IAE from 1843.7 to 56.6 % min, kept the oxygen on specification through a day the full cascade spent 724.7 min under it (Table 6.6), and earned \$5 105 more than the capture supervisor on one price day, 70 % of which a supervisor given the same prices recovered (Section 7.10). It failed the dividing-wall challenge with the benzene at 98.29 mol%, because its model lacked an input disturbance, and it passed the urea shift by thousandths. The optimizer of Chapter 11 fitted every measurement and stopped at 92.83 % of the best, because its model's slopes vanished where the plant's did not. Modifier adaptation corrected the slopes and averaged 98.96 % (Section 11.9).

12.2 Choosing a Method

The chapters suggest an order of questions for a chemical plant's control structure. Each names the chapter where its answer decided the result, and Figure 12.2 sets the same questions as a flowchart.

1. *What recycles, and what is its return gain?* A gain near one makes a slow composition that a purge or a slow loop must hold (Chapter 2). A gain of one makes an integrator that needs a loop on the recycled ratio itself, measured or inferred (Chapter 3). A negative gain wants an inner loop on the returning quality (Chapter 7).
2. *Which inventory does the makeup hold?* Pairing the makeup with pressure or with production changed where a catalyst loss appeared: 17 t/d of production under one, 2.3 bar of pressure margin under the other (Section 2.10).
3. *Where is the cliff?* Find the ignition floor, the extinction point or the fold before choosing a setpoint, keep a stated margin from it, and measure where the loop's sign holds on every branch (Chapter 4, Chapter 9).
4. *Is the quality known late?* Infer it from fast measurements, correct the inference with each laboratory result (Section 1.4.6), and protect a second quality with a cap or a back-off (Chapter 3, Chapter 10).
5. *Is there a handle that moves only energy?* Close a loop on it or price it, and trend recovery and specific energy beside purity (Chapter 8, Chapter 6).

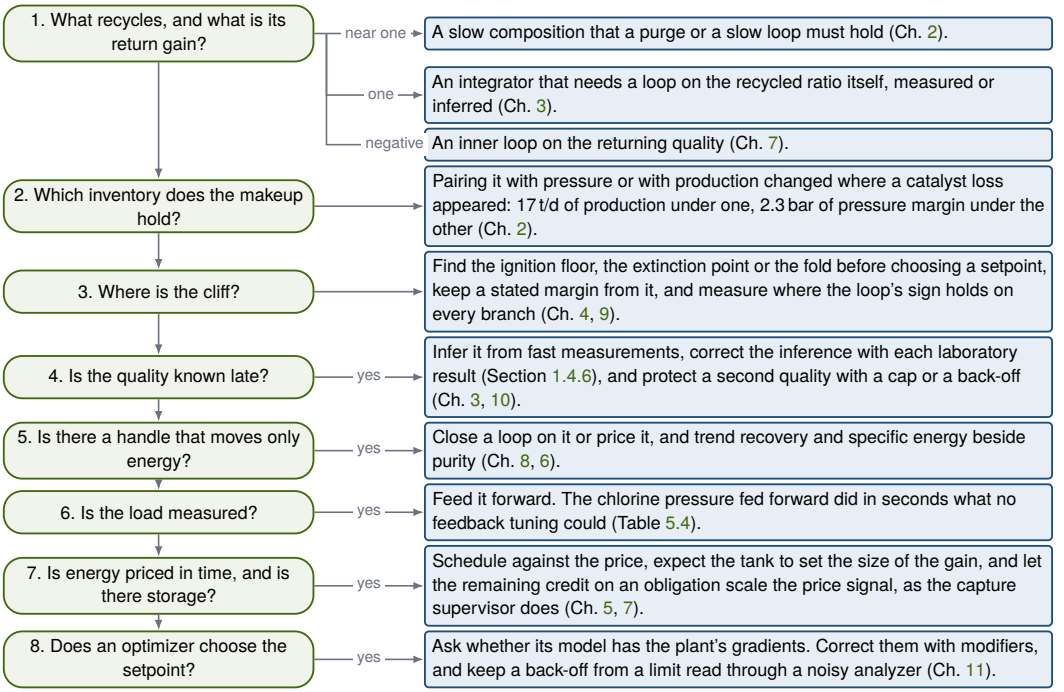

Figure 12.2 The eight questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the first three questions are about the plant and are answered before any loop is tuned.

6. *Is the load measured?* Feed it forward. The chlorine pressure fed forward did in seconds what no feedback tuning could (Table 5.4).
7. *Is energy priced in time, and is there storage?* Schedule against the price, expect the tank to set the size of the gain, and let the remaining credit on an obligation scale the price signal, as the capture supervisor does (Chapter 5, Chapter 7).
8. *Does an optimizer choose the setpoint?* Ask whether its model has the plant's gradients. Correct them with modifiers, and keep a back-off from a limit read through a noisy analyzer (Chapter 11).

A predictive controller earned its place where it coordinated two or more slow inventories or handles that the supervisory loops moved one at a time. It did not replace a missing loop or a wrong structure, and Volume 9 treats the choice in full.

12.3 Where Practice Is Heading

Plants follow the power price. A scheduling study of an air separation unit under time-sensitive electricity prices found significant savings from operating in a transient regime, raising production when energy is cheap and cutting it at peak price while using the storage available, and it built the schedule on low-order dynamic models so that the plan stays feasible [1]. A later framework lets power-intensive chemical processes take part in short-term electricity markets, whose prices move in seconds to minutes, and stresses that frequent changes of production target mean transient operation, so the schedule must respect the process dynamics [2]. The chlor-alkali scheduler and the capture supervisor of this volume are small cases of the same idea. In the editor's view the control structure

of a power-intensive plant will increasingly be judged by how far and how fast it lets the rate move without leaving its specifications.

Synthesis loops and capture plants run flexibly. A study of the Haber–Bosch process for power-to-ammonia concluded that the process can be operated over a wide operating envelope while keeping its hydrogen intake and ammonia production flexible [3]. Dynamic models and the optimization of flexible operation of post-combustion capture plants have been reviewed as a field of their own [4]. In the editor's view this moves the hard part of the ammonia loop from the design point to the turndown, where Chapter 2 found a converter that goes out between two quench positions a tenth of a percent apart, and the hard part of a capture plant to the obligation that outlasts its tanks.

> ⧖ HISTORY **2014 to 2020: A Capture Plant That Must Follow**
>
> The timeline of Chapter 1 closed in October 2014, when a capture plant on a 110 MW coal unit opened as the world's first fully commercial plant of its kind [5]. A second commercial plant reported in 2020 that it was designed for 50 to 100 percent of its carbon dioxide load and functioned best at 100. Its low-load point was set at 70 percent, limited by the turndown of the gas turbine whose waste heat raised most of its steam, and a hot start took 2 to 9.5 hours [6]. That year a review found that a capture plant needs more than two hours to settle after a cut in reboiler duty, far slower than the station upstream, and that most control studies still treated it apart from the station [7]. The plant of 2014 was built to capture, and gap 5 of Section 12.5 asks what it takes to follow as well.

Optimizers correct their gradients. An assessment of real-time optimization practice discusses the layer's place in the plant's hierarchy of decisions and its coordination with MPC [8]. Modifier-adaptation schemes reach the plant's optimum on convergence despite structural mismatch between plant and model, and need no parameter estimation. They depend on gradients estimated from noisy measurements, and a 2016 review surveys the methods for that estimate along with the open issues of the field [9]. Chapter 11 shows the cost of that estimate: every third run about three points short, and a back-off that bought robustness to noise with profit.

Intensified units bring control into design. A review of process systems engineering for intensification lists two key open questions: how to derive intensified designs systematically, and how to ensure the operability and optimality of the structures derived while they deliver their function [10]. Dynamic process intensification treats changes to a process's dynamics, operation and control as a source of efficiency in its own right, and names challenges in its design, operation, control and fault detection [11]. The dividing-wall column and the MTBE column are both intensified units, and both chapters found that the hard decision was which variable to control. In the editor's view the control structure of an intensified unit will be chosen with its design, not after it.

Compositions are read in line. Urea plants read the ratio with vibrating density sensors, Raman spectroscopy, chromatographs or the laboratory, and sensors that touch the carbamate corrode and drift. A gamma-ray gauge, which reads the density without touching the stream, has gained attention for that reason [12]. The loading of an amine solvent is still most often found by titration, which can take two hours, and a study of a gas plant's solvent found near-infrared spectroscopy able to replace it [13]. In the editor's view such a reading would have exposed the capture plant's inference that read 0.238.

Cells and catalysts lower the energy. Oxygen-depolarized cathodes make hydroxide from oxygen instead of making hydrogen, and lower the cell voltage by about 1 V at 4 kA/m^2 and the electricity by about 30 %. A reference document of 2014 lists them as an emerging technique, with one plant at industrial scale [14]. Cesium-promoted catalysts for the acid plant stay active at 380 to 400 °C and

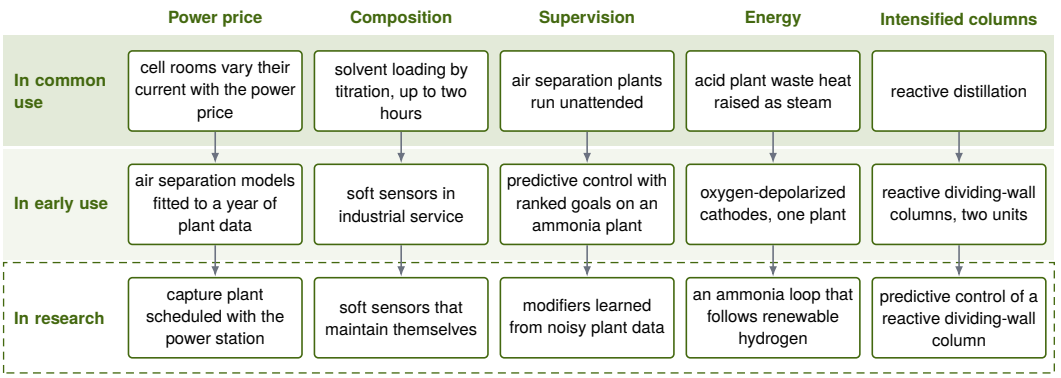

Figure 12.3 Developments in five threads. The top lane holds what the cited source calls standard or common practice, the middle lane what it reports in service at some sites or on plant data, and the bottom lane what it reports from simulation, laboratory work or as a need, in the editor's reading. Sources by thread: power price [7, 14, 17], composition [13, 18], supervision [15, 16, 19], energy [3, 14, 15, 20] and intensified columns [21]. Notice that the cell room is the only plant of the volume whose rate follows the price in common use.

can serve in the first bed to lower its inlet temperature [15]. In the editor's view the ignition floor of Chapter 4 would then sit lower and still be a floor.

Plants run with fewer people. Computer-based control lets cryogenic air separation plants run unattended or with minimal staffing. Responses that an operator would give informally at an attended plant must then be designed into the controls, and the instrumented system must shut the plant down safely in an unplanned upset with no manual action [16]. In the editor's view this is the last working rule of Section 12.4.1 applied to a whole plant. Figure 12.3 sets these developments on three lanes.

12.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [22]. 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 [23]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

These five are the professional form of the habit introduced in Section 1.7: check before trusting, and know enough to check.

12.4.1 How This Volume Was Made

Each of the ten apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace every number to a run, and a referee report was written on each paper. Three of the ten reports call themselves self-reviews. The others do not say who wrote them. 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. The ten logs hold 101 entries with a severity, from two in the urea and acid logs to 22 in the chlor-alkali and capture logs. Of these, 22 are wrong behavior or output of an app, 62 are documentation, 10 are cosmetic or display and 7 are other kinds, such as a check that tested the wrong thing or a completion rule met on the clock alone. Five cases show the errors that matter for a volume about recycle, equilibrium and energy.

Helpers that found the extinguished converter. The ammonia app offered an operator aid that finds the makeup valve opening for a pressure setpoint by bisection. The bisection assumed that the settled pressure rises with the valve. It does not, because below about 45 % the steady-state solver, started from a fixed guess, lands on the extinguished branch, where nothing is consumed and the pressure runs to its ceiling. At the design line-up the aid set the valve to 20.0 % in place of 59.3 and reported a converter out at 260 bar. A second aid, for the purge, returned a 1.0 % purge for every inert setpoint from 14.5 to 18 %, which also puts the converter out. The paper agent found the first while measuring the app and flagged the same exposure in the second, and both aids now search the lit

branch only. Later, the capture app's predictive controller was found linearizing its absorber on a root with the bottom stage loaded past its maximum, from a fixed starting guess, which gave a capture gain of −7884 percent per unit of loading where the design gives −139. An equilibrium with more than one solution traps a solver the way it traps a plant.

A self-test that measured a pump on its stop. The capture app checked its own headline result, that the lean-loading loop holds the capture gain, with a built-in test. Two statements in the test ran in the wrong order, so the circulation command went to 300 and clamped at its stop of 170. The test divided the capture change by a step of 20 that had in fact been 40, and reported 0.271. The true value is 0.177. The 0.271 had spread to the slider hint, the controller comment, the model dialog, a scenario card and the guide. The paper agent found it by instrumenting the test and seeing the circulation on its stop for eight hours. A check that passes is also a claim.

A price that moved nothing. The capture app charged its carbon price per kilogram on a quantity already converted to tonnes, so the carbon cost was a thousand times too small. At design the stack loses carbon dioxide worth about $3 900 an hour, and the scorecard showed a few dollars after an hour. The price slider, visible from 0 to 300 $/t, changed nothing a student could see. In a volume where energy cost shapes the structure, the probe is simple: move the price and ask whether the margin moves.

A converged helper that left out the recycle. After the ethylene oxide finding of Volume 12, the chapter agents of this volume checked whether each app helper behind a printed number had converged. Nine of the ten records report the check, one of them for a helper that is explicit and needs none, and the tenth chapter uses no helper. Every iterative helper converged. One that converged still answered the wrong question. The capture app's "duty for this lean loading" solves the stripper once through, at the present rich loading. In the recycle the leaner solvent returns a lap later as a leaner rich feed. A request of 0.200 settled at 0.1792, and 0.260 at 0.2765. The chapter now says the duty holds the loading only until the changed solvent comes round, and the app owner has the fix.

A step no learner can take, and a setting tried at one seed. Nine of the ten records report checking that every value an exercise asks for lies on its control's grid, and four report walking the completion steps themselves. The air separation agent found that scenario 4 asks for a level loop gain of 0.4 on a control that moves in steps of 0.5, so the typed 0.4 becomes 0.5 and the step can never be met. The app's own test passed, because it set 0.4 through the program interface and not through the control. In Chapter 11 the first suggested setting passed the challenge at the default seed with 98.19 %. A technical review ran it at 31 seeds and found it passing at 21. The suggestion became a smaller filter gain with a waste back-off, which passes all 31 at 97.47 % (Section 11.6).

The book round. Each chapter agent first confirmed that its paper had not changed since conversion, and none of the ten had. The ten records list 154 rows of numeric verification, each a stated value or group of values set beside its reproduction, and all ten report that every rerun number reproduced. The book round's defect list holds ten entries for this volume. In four apps a loop's action was labeled the opposite of Chapter 1's convention: the acid plant's bed 4 and analyzer loops, the dividing-wall split loop, the MTBE temperature loop in one case text and the base oil mid-ranging control. Each chapter now names the action as Section 1.6.2 does, and notes the app's label where it still differs. The base oil app has since been relabeled, and its Mid-ranging action control now reads Reverse for the setting Chapter 10 calls reverse acting. The MTBE paper's flow sheet labeled the overhead vapor with 7651, which is the boilup. The overhead vapor is about 10 932 kmol/h. The chapter agent left the label as printed and reported it, and the orchestrating agent removed the number from the

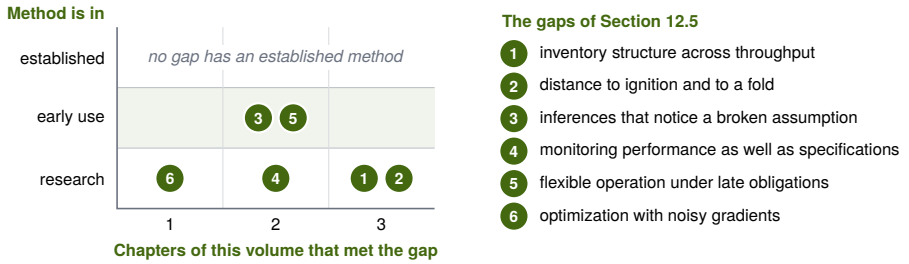

Figure 12.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: three for gap 1 (Chapters 2, 3 and 7), three for gap 2 (2, 4 and 9), two for gap 3 (7 and 8), two for gap 4 (6 and 8), two for gap 5 (5 and 7) and one for gap 6 (11). Up is the maturity of the method available, the editor's reading of the sources cited in each item. Notice that the two gaps most chapters met have the least mature methods.

chapter's copy. The agent writing Chapter 1 first took a converter feed of 5000 and fifty passes from a well-mixed shortcut, and a mixer balance gave 4900 and forty.

Every error above was caught by a check that could fail: a helper's answer run in the plant, a test instrumented step by step, a price moved, a value typed into the control a learner uses, a setting run at many seeds, a balance closed around the mixer. None was caught by rereading. Four working rules follow. Start every solver where the plant is, and reject a root the plant cannot hold. Close the balance around the recycle before trusting a gain or a helper. Move every price and see what it moves. Walk every step through the controls a person will use.

12.5 Gaps in Chemical Plant Control

Six gaps belong to the control of chemical plants, and the gaps common to agent-era practice follow them in Section 12.5.1. In the editor's view none is closed. A sentence on difficulty or progress that carries no citation is the editor's view. Figure 12.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the plants in the picture that opens this chapter.

1. *Inventory structure across throughput.* The ammonia loop needed its slow loops in service within the first hours, and its pairing decided whether a loss appeared as production or as pressure margin (Chapter 2). The urea loop and the capture plant needed a loop on the recycled quantity too (Chapter 3, Chapter 7). A method that derives the inventory pairing and the slow composition loops of a recycle plant from its model, with stated time scales for each, and rechecks them when the rate changes, is not routine. A 2023 review finds a lack of understanding of how such control systems should be designed, and calls choosing the structure from a model a very difficult combinatorial problem even for small systems [24]. The slow loops are hard to judge because they act over hours that a steady-state analysis does not see. Progress would be a procedure that returns each pairing with its time scale.

2. *Distance to ignition and to a fold, measured online.* The ammonia converter goes out between two quench positions a tenth of a percent apart, and the MTBE column's branch monitor reads one tray and can name the wrong branch near a fold (Chapter 2, Chapter 9). The acid plant's best inlet is the coolest one its catalyst can still light (Chapter 4). An estimate of the margin to the nearest extinction or fold, from plant measurements, would let a structure keep a known margin rather than a guessed one. The Open challenge box after this list takes the gap further.

3. *Inferred compositions that notice a broken assumption.* The capture plant's inferred lean loading read 0.238 while the true value was 0.274 after a pressure drift, and the dividing-wall loops read

their setpoints while the true purities had moved (Chapter 7, Chapter 8). An inference that detects when a secondary variable has changed its meaning, before the laboratory does, remains open. A 2009 review finds that telling a drift of the process from a drift of the sensor needs expert knowledge, and that most soft sensors have no automated maintenance [18]. It is hard because a variable that has changed its meaning still moves smoothly with the plant. Progress would be an inference that flags its own broken assumption.

4. *Monitoring performance as well as specifications.* Fouling took eight percent of the oxygen with every loop on setpoint, a heat leak cut the recovery to 64.9 %, and a column on specification burned 17.7 % more steam (Chapter 6, Chapter 8). Monitors that track recovery, specific energy and the position of an energy minimum, and say which loop to change, are thinly developed. A dividing-wall column loses its energy saving away from its best vapor split, and the window of minimum energy may be narrow and move with the feed [25]. It is hard because recovery and specific energy have no setpoint and no alarm. Progress would be a monitor that reports the distance from the energy minimum beside each purity.

5. *Flexible operation under obligations that bind late.* Seven hours of high price put the capture plant's lean tank on its floor within 73 min against a four-hour average obligation, and the chlor-alkali scheduler was limited by its tank (Chapter 7, Chapter 5). Schedules that carry the plant's dynamics, storage and rolling obligations together, and degrade gracefully when a forecast is wrong, are still research. Chlor-alkali plants already vary their current with the electricity price [14]. A 2020 review finds that most control studies treat a capture plant apart from the power station it is coupled to, and that steady-state optimization cannot see the transients [7]. It is hard because an obligation averaged over hours answers for decisions taken hours before.

6. *Optimization with a wrong model and noisy gradients.* Modifier adaptation reached the optimum of Chapter 11 only by perturbing the plant, and a waste back-off traded profit for robustness to noise. Real-time optimization is used less than one would expect [26], and gradients inferred from noisy measurements are key to the reliability of modifier adaptation [19]. It is hard because each estimate costs a perturbation that the profit pays for. Gradient estimates that need less perturbation and come with a bound on the limit violation they allow would make such layers easier to trust.

> ⚠ **OPEN CHALLENGE** **How Far Is the Cliff?**
>
> Gap 2 ties with gap 1 as the gap most chapters met. A 2018 study of an autothermal ammonia converter for power-to-ammonia found that operating outside its envelope shuts down the autothermal operation or overheats the catalyst, and that such converters are run only at their optimum [3]. Multiple steady states appear even in ideal distillation, and the result is highly relevant in practice, because industrial columns set their inputs on a mass or volume basis [27]. The margin to a cliff is hard to measure because it is a property of a branch the plant is not on. Progress would be an online estimate of that margin, with its uncertainty, that a setpoint and an alarm can use.

12.5.1 Gaps Common to Agent-Era Practice

The final chapter of each volume lists gaps specific to its topic. Seven more are common to every topic in the series. In the editor's view none of them is closed.

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [28]. 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 [23]. 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 [29]. 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 [30]. 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 [31]. 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 plant that sends material round again needs a loop on what goes round.
- An equilibrium puts the optimum next to a cliff, so the structure carries a margin from it or a measurement chosen so that crossing it can be undone.
- Energy is spent by the handle no purity loop watches, so recovery and specific energy are trended beside purity.
- Timing the energy is worth what the storage allows. Prediction paid on slow coupled inventories, and an optimizer needs the right slopes.

References

[1] R. C. Pattison, C. R. Touretzky, T. Johansson, I. Harjunkoski, and M. Baldea, "Optimal process operations in fast-changing electricity markets: Framework for scheduling with low-order dynamic models and an air separation application," *Industrial & Engineering Chemistry Research*, vol. 55, no. 16, pp. 4562–4584, 2016.

[2] J. I. Otashu and M. Baldea, "Grid-level "battery" operation of chemical processes and demand-side participation in short-term electricity markets," *Applied Energy*, vol. 220, pp. 562–575, 2018.

[3] I. I. Cheema and U. Krewer, "Operating envelope of Haber–Bosch process design for power-to-ammonia," *RSC Advances*, vol. 8, no. 61, pp. 34 926–34 936, 2018.

[4] M. Bui, I. Gunawan, V. Verheyen, P. Feron, E. Meuleman, and S. Adeloju, "Dynamic modelling and optimisation of flexible operation in post-combustion CO_2 capture plants—A review," *Computers & Chemical Engineering*, vol. 61, pp. 245–265, 2014.

[5] M. Bui, C. S. Adjiman, A. Bardow, E. J. Anthony, A. Boston, S. Brown, P. S. Fennell, S. Fuss, A. Galindo, L. A. Hackett, J. P. Hallett, H. J. Herzog, G. Jackson, J. Kemper, S. Krevor, G. C. Maitland, M. Matuszewski, I. S. Metcalfe, C. Petit, G. Puxty, J. Reimer, D. M. Reiner, E. S. Rubin, S. A. Scott, N. Shah, B. Smit, J. P. M. Trusler, P. Webley, J. Wilcox, and N. Mac Dowell, "Carbon capture and storage (CCS): the way forward," *Energy & Environmental Science*, vol. 11, no. 5, pp. 1062–1176, 2018.

[6] Petra Nova Power Holdings LLC and G. Kennedy, "W.A. Parish post-combustion CO_2 capture and sequestration demonstration project (final technical report)," U.S. Department of Energy, Office of Scientific and Technical Information, Tech. Rep. DOE-PNPH-03311, Mar. 2020.

[7] X. Wu, M. Wang, P. Liao, J. Shen, and Y. Li, "Solvent-based post-combustion CO_2 capture for power plants: A critical review and perspective on dynamic modelling, system identification, process control and flexible operation," *Applied Energy*, vol. 257, p. 113941, 2020.

[8] M. L. Darby, M. Nikolaou, J. Jones, and D. Nicholson, "RTO: An overview and assessment of current practice," *Journal of Process Control*, vol. 21, no. 6, pp. 874–884, 2011.

[9] A. Marchetti, G. François, T. Faulwasser, and D. Bonvin, "Modifier adaptation for real-time optimization—methods and applications," *Processes*, vol. 4, no. 4, p. 55, 2016.

[10] Y. Tian, S. E. Demirel, M. M. F. Hasan, and E. N. Pistikopoulos, "An overview of process systems engineering approaches for process intensification: State of the art," *Chemical Engineering and Processing: Process Intensification*, vol. 133, pp. 160–210, 2018.

[11] M. Baldea and T. F. Edgar, "Dynamic process intensification," *Current Opinion in Chemical Engineering*, vol. 22, pp. 48–53, 2018.

[12] S. Z. Islami Rad and R. Gholipour Peyvandi, "A novel calibration method for online determination of nitrogen-to-carbon ratio in urea synthesis," *Scientific Reports*, vol. 16, no. 1, p. 36571, 2026.

[13] M. Yusop Nurida, D. Norfadilah, M. R. Siti Aishah, C. Zhe Phak, and S. M. Saleh, "Monitoring of CO_2 absorption solvent in natural gas process using Fourier transform near-infrared spectrometry," *International Journal of Analytical Chemistry*, vol. 2020, p. 9830685, 2020.

[14] T. Brinkmann, G. Giner Santonja, F. Schorcht, S. Roudier, and L. Delgado Sancho, "Best available techniques (BAT) reference document for the production of chlor-alkali," European Commission, Joint Research Centre, Luxembourg, JRC Science and Policy Report EUR 26844 EN, 2014.

[15] European Commission, "Reference document on best available techniques for the manufacture of large volume inorganic chemicals: Ammonia, acids and fertilisers," European Commission, Integrated Pollution Prevention and Control, Tech. Rep., Aug. 2007.

[16] European Industrial Gases Association, "Guideline for safe practices for cryogenic air separation plants," European Industrial Gases Association, Brussels, EIGA Doc 147/22, 2022.

[17] C. Tsay, Y. Cao, Y. Wang, J. Flores-Cerrillo, and M. Baldea, "Identification and online updating of dynamic models for demand response of an industrial air separation unit," *IFAC-PapersOnLine*, vol. 54, no. 3, pp. 140–145, 2021.

[18] 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.

[19] E. A. del Rio Chanona, P. Petsagkourakis, E. Bradford, J. E. Alves Graciano, and B. Chachuat, "Real-time optimization meets Bayesian optimization and derivative-free optimization: A tale of modifier adaptation," *Computers & Chemical Engineering*, vol. 147, p. 107249, 2021.

[20] C. Smith, A. K. Hill, and L. Torrente-Murciano, "Current and future role of Haber–Bosch ammonia in a carbon-free energy landscape," *Energy & Environmental Science*, vol. 13, no. 2, pp. 331–344, 2020.

[21] A. A. Kiss and M. Jobson, "Taking reactive distillation to the next level of process intensification," *Chemical Engineering Transactions*, vol. 69, pp. 553–558, 2018.

[22] 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.

[23] 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.

[24] S. Skogestad, "Advanced control using decomposition and simple elements," *Annual Reviews in Control*, vol. 56, p. 100903, 2023.

[25] D. Dwivedi, J. P. Strandberg, I. J. Halvorsen, H. A. Preisig, and S. Skogestad, "Active vapor split control for dividing-wall columns," *Industrial & Engineering Chemistry Research*, vol. 51, no. 46, pp. 15 176–15 183, 2012.

[26] D. Krishnamoorthy and S. Skogestad, "Real-time optimization as a feedback control problem: A review," *Computers & Chemical Engineering*, vol. 161, p. 107723, 2022.

[27] E. W. Jacobsen and S. Skogestad, "Multiple steady states in ideal two-product distillation," *AIChE Journal*, vol. 37, no. 4, pp. 499–511, 1991.

[28] 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.

[29] 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.

[30] 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.

[31] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

12.6 Exercises

12.1 Discussion · 8 min

An agent proposes to hold the N/C ratio of a new urea loop with a ratio station on the ammonia and carbon dioxide flow meters, retuned twice as fast as before, and reports that its simulation holds the ratio of the two flows within 0.1 percent. Using Chapter 3 and Chapter 2, explain why that report does not show the loop is held, and say what structure and what evidence you would ask for.

12.2 Discussion · 8 min

An optimization study shows that the ammonia converter of Chapter 2 would make 11 % more product with its hot spot near 480 °C, and an agent recommends moving the setpoint there. List what you would check before agreeing, and compare the margin the acid plant of Chapter 4 keeps from its ignition floor.

12.3 Discussion · 7 min

A monitoring agent reports that a dividing-wall column and an air separation unit are healthy because every loop has been on its setpoint for a week. Using Chapter 8 and Chapter 6, explain why that report could be true while both plants lose money, and name the measurements you would add to it.

12.4 Discussion · 8 min

Section 12.4.1 describes two operator aids in the ammonia app that returned an extinguished converter, and a capture helper that converged but left out the recycle. You are asked to write the acceptance tests an agent-written steady-state helper must pass before its answer is shown to an operator. List them, and say which defect each would have caught.

12.5 Discussion · 8 min

A plant manager asks whether a new amine capture unit, a chlor-alkali cellroom and an air separation unit on one site should each get a predictive controller. Using the comparisons of Chapter 7, Chapter 5, Chapter 6 and Chapter 8, argue for an order of investment and say what each predictive controller would have to beat.

12.6 Discussion · 9 min

An agent installs a two-step real-time optimizer on a reactor and reports that the fitted model matches the plant's five measured fractions within half a point and that the optimizer has settled. Using Chapter 11 and Section 12.4.1, explain why neither fact shows the plant is at its optimum, and propose the evidence you would require before the layer is left running.

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) 0.423. (b) 0.398: the inert lowers the partial pressures and the conversion.

1.3 (a) 12.0 %, 7.5 h. (b) About 22.5 h, three time constants.

1.5 (a). With no exit for the excess, the return gain is one and the loop ratio integrates the feed error.

1.7 $y_{out} = 0.00668$, 83.3 % removed.

1.9 (d). The two-step method stops where the fitted model is flat, and a model can match the values and miss the slopes.

Chapter 2. Ammonia Loop Recycle and Inventory Pairing

2.1 (c). Pressure is the controlled inventory under this pairing, so the loss appears as production and nothing else.

2.3 (a). An inert leaves only through the purge, so its time constant is $N/\dot{n}_P$, about 5.1 h at design.

2.5 (a) $K_p = 1.291 \times 10^{-2}$ and 8.671×10^{-3}, a factor of 1.49. (b) $\kappa = 0.3295$, $u = 0.7262$, $y_A^{eq} = 17.4\,\%$. (c) $\kappa = 0.1770$, $y_A^{eq} = 11.0\,\%$. A hotter bed reacts faster toward a lower ceiling.

2.7 (a) 1.853, 640 t/d, and the pressure within 0.05 bar of setpoint. (b) 7.6 against 8.0 settled. (c) 3.40, still 2.18 short of 5.58.

2.9 (a) 10.36 % at 81.5 min, purge 56.2 %. (b) 11.98 % at 333 min, 10.73 % at the end. (c) IAE 56.6 against 1843.7, for $924 less margin.

2.11 About $22 300 per day. The fix is at the reformer. In the meantime the inert setpoint can be raised to trade production against purged hydrogen.

Chapter 3. Urea Loop Ratio Inference

3.1 (b). The excess ammonia has no exit, so the N/C is an inventory and a fresh ratio error makes it ramp, not step.

3.3 (d). The plant moves during the analysis, so the result is compared with what the inference said when the sample was drawn.

3.5 (a) 0.135. (b) 2.792, which matches the 2.791 at which case 3's index reaches 1. (c) 1.41, so the effective floor rises to 2.817.

3.7 (a) 3.030 and 3.083; 965.7 kg/m^3 and 956.1 kg/m^3. (b) 0.0156 per minute, about 4 % above 0.01499. (c) 2 h 9 min after the step.

3.9 (a) It stops at 9 h 48 min on the corrosion floor, with samples 2.95 and 2.86. (b) Samples 2.97, 3.02, 3.03, 61.3 %, 0.88 wt%, 142.9 bar a: met. (c) The biuret, 0.92 wt%.

3.11 The laboratory removes calibration errors that the inference cannot see. An 8 h period leaves such an error uncorrected for most of a shift. Keep the 4 h period, or check the gauge directly if samples are cut.

Chapter 4. Sulfuric Acid Plant and Its Ignition Floor

4.1 (b). The equilibrium curve falls with temperature, so a parallel line that starts hotter meets it at a lower conversion.

4.3 (d). With the product removed, the equilibrium of bed 4's own gas lies far above the feed's curve, so bed 4 can convert most of what is left.

4.5 (a) $k_{hi} = 0.1189$ and $k = 0.0594$. (b) 5.05×10^{-3}. (c) 3.93 % with the break, 59.5 % without it.

4.7 (a) 20.4, 40.0, 66.0 and 36.7 %, all opened from 15.0, 29.6, 42.6 and 30.0 %. (b) 214.5 and 132.8 ppmv. (c) It closes, from 36.7 to about 33.7 %. Bed 4's loop is direct acting.

4.9 (a) Fails on all four limits. (b) Fails on the acid alone, at 99.38 %. (c) Passes with 99.88 %, 260 ppm, 417 °C and 98.45 to 98.63 %.

4.11 (a) All four inlets lie in the conventional range, and bed 1's 630.6 °C outlet lies in the upper range. (b) The wrong-way rise can carry the catalyst past the range where activity is lost. (c) From the catalyst supplier, and the MPC's lower bound of 410 °C is the bottom of the conventional range.

Chapter 5. Chlor-Alkali Electrolyzer Ratio Control

5.1 (a). The ratio follows the current at once but carries its own 1.8 % error, and only the trim could remove it.

5.3 (c). The monitor sees one cell drifting from the others, not a stack that moves together.

5.5 (a) $\tau_e = 180$ s, 133.2 mbar, $t_m = 21.5$ s. (b) 53.9 s, and $K_c/\tau_I > 4.44\,\text{min}^{-1}$. (c) 1.295 mbar/s and 18.3 s. The estimate assumes a perfect slave and a continuous master.

5.7 (a) 86.8 t, 5.6 t on specification, 81.0 t off, 177.9 g/L, 93.22 %. (b) 88.4 t, all on specification, 209.7 g/L, 94.90 %. (c) The weak anolyte and strong caustic of (a) cost current efficiency.

5.9 (a) Fails on tonnage, ΔP, the trip and a stockout: \$4880, 9.8 t. (b) Met: \$10 908, 72.3 t, 26.4 to 39.0 mbar. (c) Fails on the trip, the tonnage and a 109 min stockout. The true minimum is just below 20 mbar and the tile rounds it.

5.11 About \$0.49 million over the remaining two and a half years. The monitor will not warn of it. Measure the hardness of the polished brine and switch the ion-exchange columns before breakthrough.

Chapter 6. Air Separation With Thermally Coupled Columns

6.1 (d). The duty and the conductance are unchanged, so the driving force stays at 1.894 K, and the nitrogen must condense 1.55 K hotter, which it does only at a higher pressure.

6.3 (b). It moves the purity by 378 ppm per percent, against −174 ppm for the vanes and −70 ppm for the split.

6.5 (a) $\lambda_{11} = 0.471$, $\lambda_{12} = 0.529$. (b) $K_{11} - K_{12}K_{21}/K_{22} = -0.0149\,\%/\%$, 2.12 times K_{11}. (c) A loop tuned with the argon loop open has about twice the gain once the argon loop closes.

6.7 (a) 5.600, 1.400, 95.36, 93.47, 722.0 kW, then 6.303, 1.626, 96.91, 95.01, 767.5 kW. (b) 3.11 bar per bar at 1.894 K. (c) 722.0 and 767.5 kW.

6.9 (a) 28.66, 1042, 44.56 and 51.79 %. (b) 0, below 1 m³/h from 0.046 to 0.450 h, and 47.37 %. (c) The tripled gain refills the level faster by stopping the oxygen supply for about 24 minutes.

6.11 (a) \$110.44 and \$96.29 an hour, \$14.15 less at 78 %. (b) 0.0536 \$ per cubic meter, 0.975 of the oxygen price. (c) Refrigeration is not modeled.

Chapter 7. Amine Capture Recycle Gain and Energy Timing

7.1 (c). A frozen duty strips the larger flow less deeply, and a lap later the leaner-than-before solvent stops arriving.

7.3 (a). The inference assumes 190 kPa, so holding the old temperature holds a richer solvent, and the capture loop makes it up with circulation.

7.5 (a) $g_R = 1 - 0.177/0.040 = -3.4$. (b) About 153 with the inner loop, 230 m³/min with the duty frozen, which the pump cannot reach. (c) 0.62 and 0.14, a ratio of 4.4; the tuning assumes the inner loop is in service.

7.7 (a) 93.55 and 90.81 %, gains 0.177 and 0.040 points per m³/min; lean loading 0.263. (b) 90.25 and 91.73 %; host output 463.7 falls to 452.3 MW.

7.9 (a) Worst 3.19 points, rolling minimum 86.96 %, amine peak 3.70 mg/Nm^3, rich tank 11.1 %: fails; 17.2 t short. (b) 0.58 points, 89.10 %, 3.74 mg/Nm^3, 12.6 %: passes; \$37,240 per hour. (c) 1.11 points, 88.93 %, 3.80 mg/Nm^3, 4.7 %: passes; \$37,518 per hour, about \$4,400 more over the sixteen hours.

7.11 (a) $2^{(128-115.87)/8} = 2.9$; \$12,960 per hour. (b) The amine emission limit of 4.0 mg/Nm^3, which the plant holds only by giving up capture or cooling the absorber top.

Chapter 8. Dividing-Wall Column and Its Energy Valley

8.1 (c). Three purities and the component balances fix the three product flows, so the fourth handle can only redistribute the reflux and the boilup.

8.3 (a). A tray's bubble point moves with every component on it, and the line was fitted for one feed.

8.5 (a) 32.34 kJ/mol and 0.008 984 MW per kmol/h. The duties are 5.83, 6.86 and 7.08 MW, so the outer points lie 17.7 % and 21.4 % above the minimum. (b) 0.628, 0.56 and 0.50. (c) 0.675.

8.7 (a) 6.86 MW and +17.5 %. (b) At 0.34, 5.98 MW and +2.5 %. At 0.28, 6.86 MW and +17.5 % at 4 h. (c) 7.02 MW and +20.3 % at 12 h, still climbing toward the curve's 7.08 MW.

8.9 (a) Lowest 98.63, 98.15 and 98.13 mol %, worst hour +12.4 %: both halves fail. (b) 99.11, 98.07 and 99.07 mol %, worst hour +0.3 %: passes. (c) The box snaps to 0.340, and the run passes with 99.11, 98.08 and 99.07 mol % and +0.0 %, but only because the day's feed was known in advance.

8.11 (a) 8.0 % of the duty, or $48.2 \times 0.008984 = 0.433$ MW. (b) The operator trades purity margin for throughput: at 0.05 mol % above the specifications the column fits under the steam limit and makes all three products, but a disturbance now has less room before a product goes off specification.

Chapter 9. MTBE Reactive Distillation and Its Branches

9.1 (b). Past a Hopf point the state is lost by a growing oscillation, not by a jump, and at 1 MW per hour the duty moves on before it has grown.

9.3 (d). The monitor reads one tray against one saddle, and the excess methanol has reached the bottoms while the unstripped C4 keeps tray 8 cool.

9.5 (a) $K = 271.2$ at 25 °C and 12.76 at 100 °C. (b) $K_{ads} = 1/\varepsilon = 20$ and $1/\varepsilon^2 = 400$. (c) 0.47 and 0.81 at $\varepsilon = 0.05$, 0.72 and 0.91 at $\varepsilon = 0.02$.

9.7 (a) 3.35 h, 44.82 MW. (b) 10.15 h, 40.88 MW, a loop 3.93 MW wide. (c) 1.32 MW past the Hopf point and 0.87 MW past the fold.

9.9 (a) Low from 6 h 40 min, 67.58 mol % and 73.0 %, fails. (b) Never, 99.70 mol % and 97.2 %, passes. (c) Never, 99.64 mol % and 97.3 %, passes. The same tuning passes with margin under the peak, so the shipped setting fails on its setpoint, not its tuning.

9.11 (a) The catalyst mass, the rate term ε and the latent heat λ, teaching values on which the folds move by several megawatts. (b) The stripping-tray temperature that stops rising with the duty, and a catalytic tray that jumps between the branches.

Chapter 10. Synthetic Base Oil Mid-Ranging Control

10.1 (a). Table 10.4: the reactor route gives 138.5 at 51.0 % yield, the cut route 134.8 at 60.9 %.

10.3 (c). Equation (10.7): the corrected value $\hat{v}$ at the instant the sample was drawn.

10.5 (a) −0.00602 per °C. (b) –0.0281 cSt per °C, about 10.7 °C of cooling. (c) +0.58 carbons per °C, with $a = +1$, reverse acting (labeled Reverse in the app).

10.7 (a) 6.300 cSt, 138.5, 51.0 %, 113.2 °C. (b) 6.298 cSt, 134.8, 60.9 %. (c) The isomerization term: the colder reactor is below its onset and the cut route pays it at 124.5 °C. The cut route only moves bright stock into the motor oil tank.

10.9 (a) 6.000 cSt, 73.9 carbons, 129.4, 163.5 t. (b) 5.4 t, 133.5, 115.3 °C. (c) The cap acts only after a late laboratory index falls and releases as soon as one recovers, while mid-ranging sees the cut open within minutes and cools the reactor for good.

10.11 (a) $1 - e^{-40/210} = 0.17$ and $1 - e^{-400/210} = 0.85$. (b) It hides fast upsets and passes slow drifts such as fouling instruments, columns and aging catalyst.

Chapter 11. Williams–Otto Real-Time Optimization

11.1 (d). Two-step moves to the fitted model's optimum, and where that optimum is the point itself the layer stops, whatever the plant's slopes are.

11.3 (b). It moves within one decision of the rise and plans the setpoints through the transient on a model with the holdup. Planned on the static map, it peaks as high as modifier adaptation.

11.5 (a) 88.2 kW and 5.16 min. (b) 351.0 kW and 79.93 °C. (c) 74.08 °C. The level and temperature loops return the holdup to W_0 and T to its setpoint, so the steady state is the published one whatever the heats.

11.7 (a) 178.5 \$/s, a gap of 5.82 % and 7.99 % waste, so the true profit is about 168.1 \$/s. (b) 3.51 % and 8.45 %, over the limit. (c) 78.2 °C.

11.9 (a) Economic MPC 98.66 % with a peak of 8.59 %, modifier adaptation 98.83 % with 9.33 %. (b) All the economic MPC's moves from 12.10 h, but of modifier adaptation's moves only the one at 13.05 h. (c) 97.72 % and 8.17 % against 98.50 % and 8.15 %: with nothing measured to act on, the economic MPC gains nothing.

11.11 (a) The layer holds the reading at 8 %, so the true waste settles near 7.7 % and about 3 \$/s of profit is given away. (b) The true waste would sit near 8.3 %, over the permit. A waste back-off b at least as large as the bias, and a laboratory check of the analyzer, protect it.

Chapter 12. Summary and Outlook

12.1 A good answer notes that the excess ammonia has no exit, so the return gain is one and the reactor N/C integrates any error between the meters and the true feeds. A ratio of flows is a ratio of meters, not of the process. Two weight percent of water carried the ratio held on the flows to the 2.80 floor in 6 h 49 min, and the shipped challenge stopped at 9 h 48 min. It asks for a loop on the N/C itself, inferred from density and temperature and corrected by each laboratory result, a run with a meter or feed error, and the true N/C, not the flow ratio, in the report. It cites the ammonia loop's eightfold amplification as the same mechanism with a purge.

12.3 A good answer separates specifications from performance. With all three purities held, the dividing-wall column's liquid split moves only the duty, and a fixed split of 0.40 needs 17.7 percent more steam than the minimum. In the air separation unit, fouling held every setpoint while production fell eight percent, and a heat leak raised the oxygen purity while recovery fell from 90.8 to 64.9 percent. Measurements to add: specific energy (duty per product, kWh per cubic meter of oxygen), recovery, the position of the split against the valley's minimum, and the exchanger's temperature difference.

12.5 A good answer puts structure and storage first. In the capture plant the inner lean-loading loop and the price supervisor carry most of the value, and the MPC's \$5 105 extra on a price day was 70 percent recovered by a supervisor given the same prices. The chlor-alkali gains come from ratio trims, feedforward and a scheduler limited by its tank. In the air separation unit the MPC kept oxygen on specification where full cascade spent 724.7 min under it, because two slow handles interact, but it lost when the pressure setpoint lay beyond the motor rating. The dividing-wall MPC failed its challenge (benzene 98.29 mol percent) for lack of an input disturbance model. Each MPC must beat the best supervisory structure, tuned, on the same events.

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.

absorber A column in which a rising gas meets a falling solvent that takes one component out of the gas. *Chapter 1*.

absorption factor The solvent flow divided by the gas flow and by the slope of the equilibrium line; the larger it is, the more an absorber captures. *Chapter 1*.

adiabatic reactor A reactor that exchanges no heat with its surroundings, so the heat of reaction changes the temperature of the gas passing through. *Chapter 1*.

amine scrubbing Removing carbon dioxide or another acid gas with an amine solution, which takes the gas up when cool and gives it up when hot. *Chapter 1*.

ammonium carbamate The salt that ammonia and carbon dioxide form first in a urea reactor, which then loses water slowly to become urea. *Chapter 3*.

analyzer An instrument that takes a sample of a stream and reports its composition or a property, usually some minutes later. *Chapter 1*.

averaging level control Level control tuned loosely, so that a vessel absorbs upsets instead of passing them downstream as changes of flow. *Chapter 6*.

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 disturbances and measurement error do not carry the plant over the limit. *Chapter 10*.

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*.

biuret An unwanted byproduct that forms when urea solution is held hot, and whose content in the product is limited. *Chapter 3*.

branch monitor A check that reads a catalytic tray's temperature to tell on which of its steady states the reactive column is running. *Chapter 9*.

carbon capture Removing carbon dioxide from a flue gas so that it is not released to the air. *Chapter 7*.

cascade control Two loops in series, in which the outer controller sets the setpoint of a faster inner one. *Chapter 3*.

caustic soda Sodium hydroxide, made together with chlorine and hydrogen when brine is electrolyzed. *Chapter 5*.

chemical equilibrium The point at which a reversible reaction stops changing because it runs as fast backward as forward; temperature and pressure move it. *Chapter 1*.

chlor-alkali process Electrolysis of salt water that makes chlorine, caustic soda and hydrogen in fixed proportion to the electric current. *Chapter 5*.

circulator The compressor that keeps the gas of a synthesis loop moving around the loop. *Chapter 2*.

condenser-reboiler One exchanger in which a vapor condensing on one side boils a liquid on the other, joining two columns. *Chapter 6*.

contact process Making sulfuric acid by burning sulfur, oxidizing the sulfur dioxide over a catalyst and absorbing the product in strong acid. *Chapter 1*.

continuous emissions monitor (CEMS) An analyzer that reads a pollutant in the stack gas without pause, for the record a regulator keeps. *Chapter 4*.

control valve A valve moved by an actuator to the position the controller asks for, the most common final element. *Chapter 1*.

conversion–temperature diagram A plot of conversion against temperature on which an adiabatic bed moves along a straight line toward the equilibrium curve. *Chapter 1*.

corrosion index A teaching measure of liner corrosion that rises steeply as the ratio of ammonia to carbon dioxide falls toward its floor. *Chapter 3*.

current efficiency The fraction of the electric current that does the wanted work, here carrying sodium through the membrane. *Chapter 5*.

cut point The chain length, counted in carbons, at which the product is divided from the lighter or heavier oligomers. *Chapter 10*.

dead time The time between a change in the process and the first sign of it in a measurement. *Chapter 1*.

degrees of freedom The number of handles left free to move once the levels and the product specifications are held. *Chapter 1*.

demand response Changing a plant's power use with the hour's electricity price or the needs of the grid. *Chapter 5*.

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.*

dividing-wall column One distillation column split inside by a vertical wall, so that it makes three products with less energy than two columns. *Chapter 1.*

double absorption Absorbing the sulfur trioxide between converter beds and again at the end, so the last bed can convert almost all that is left. *Chapter 4.*

energy valley The curve of a dividing-wall column's heat duty against its liquid split, whose minimum moves with the feed. *Chapter 8.*

equilibrium-limited reaction A reversible reaction whose conversion is capped by its equilibrium, however much catalyst is used. *Chapter 1.*

etherification A reaction that joins an alcohol to an olefin to make an ether, such as methanol and isobutene to MTBE. *Chapter 9.*

feed and effluent exchanger An exchanger that heats a reactor's feed with the hot gas leaving it, so that the reactor supplies its own heat. *Chapter 2.*

feedforward Control that moves a handle when a disturbance is measured, before the disturbance reaches the controlled variable. *Chapter 1.*

final control element The device that changes a flow of material or energy, such as a valve, a pump, a compressor or a rectifier. *Chapter 1.*

flexible capture Running a capture plant harder when power is cheap and lighter when it is dear, within a limit on the rolling average. *Chapter 7.*

guide vanes Adjustable blades at a compressor's inlet that set how much air it draws. *Chapter 6.*

ignition and extinction The jump of a catalyst bed between a lit, reacting state and a cold, dead one as its inlet temperature crosses a threshold. *Chapter 4.*

ignition floor The lowest inlet temperature at which a bed stays lit; just below it the conversion collapses. *Chapter 4.*

inert A gas that enters a loop but takes no part in the reaction, so it builds up until the purge removes it. *Chapter 1.*

inferential control Control of a quality that is not measured fast, through a value computed from other measurements. *Chapter 1.*

input multiplicity One value of an output reached at two different values of the same input, so the sign of the process gain changes between them. *Chapter 1.*

integrating process A process whose output keeps moving as long as its inputs are out of balance, with no level at which it settles. *Chapter 3.*

interpass absorber The absorber partway through the converter that takes sulfur trioxide out of the gas before the last bed. *Chapter 4.*

inventory control Holding the amount of material in a vessel or a loop, through a level, a pressure or a flow. *Chapter 1.*

Kremser relation A formula for what leaves a counter-current absorber or stripper, from its number of stages, its absorption factor and the entering solvent. *Chapter 1.*

lap time The time the solvent takes to go once around its loop, from the absorber through the stripper and back. *Chapter 1.*

lean loading The carbon dioxide still held by the regenerated solvent that returns to the absorber. *Chapter 1.*

liquid split The fraction of the liquid at the top of the wall that is sent to the prefractionator side of a dividing-wall column. *Chapter 8.*

main air compressor The machine that compresses all the air an air separation unit takes in, driven by a motor whose power the controllers cap. *Chapter 6.*

membrane cell An electrolysis cell whose anode and cathode sides are divided by a membrane that lets sodium through. *Chapter 5.*

mid-ranging control Using a fast, costly handle for the moment while a slow, cheap one moves to bring the fast one back to mid-range. *Chapter 10.*

model predictive control (MPC) Control that predicts the plant over a horizon with a model and plans moves against limits, again at every period. *Chapter 1.*

modifier adaptation Correcting an optimizer's model at each run with the measured error and slope, so that its best point becomes the plant's. *Chapter 1.*

MTBE Methyl tert-butyl ether, the ether that a reactive column makes from methanol and the isobutene of a C4 cut. *Chapter 9.*

multiple steady states More than one steady state for the same settings of the handles, so the plant's history decides where it sits. *Chapter 1.*

N/C ratio The molar ratio of ammonia to carbon dioxide in a urea reactor, which sets the conversion and the corrosion. *Chapter 3.*

optimum inlet temperature The inlet temperature at which a bed converts most, balancing a faster catalyst against a lower equilibrium. *Chapter 1.*

override control A selector that lets a second controller take over a handle when a limit comes near. *Chapter 2.*

Petlyuk column A prefractionator and a main column joined by their liquid and vapor flows, the arrangement a dividing-wall column builds in one shell. *Chapter 8.*

PID control Feedback that moves the handle in

proportion to the error, its running sum and its rate of change. *Chapter 1*.

plantwide control Choosing the control structure of a whole plant, with its recycles and inventories, rather than one unit at a time. *Chapter 1*.

polyalphaolefin (PAO) A synthetic base oil made by joining decene molecules into short chains. *Chapter 10*.

prefractionator The first column, or the feed side of the wall, which splits a three-component feed roughly before the main column. *Chapter 8*.

purge A small stream drawn from a loop so that inerts and excess reactants have a way out. *Chapter 1*.

quench converter A reactor of several catalyst beds in which cold feed gas is mixed in between beds to cool the gas. *Chapter 2*.

ratio control Setting one flow in proportion to another measured flow, often with a trim that corrects the ratio from a composition. *Chapter 1*.

reactive distillation A column in which a reaction runs on catalytic trays while distillation removes the product, so the reaction goes further. *Chapter 1*.

real-time optimization (RTO) A layer that waits for a steady state, fits a model to it and moves the setpoints to the model's best point. *Chapter 1*.

rectifier Equipment that turns alternating current into the direct current a cell room uses, at the level the controllers set. *Chapter 5*.

recycle Material sent back to an earlier point of the plant, which multiplies the gains and the time constants it passes through. *Chapter 1*.

regeneration energy The heat a stripper uses to drive carbon dioxide out of the solvent so that the solvent can absorb again. *Chapter 1*.

relative gain array A table that shows how strongly each pairing of an input and an output is changed when the other loops close. *Chapter 6*.

relative volatility How much more readily one component goes into the vapor than another; the larger it is, the easier the separation. *Chapter 1*.

safety interlock Logic that takes the plant to a safe state when a measurement crosses a limit, apart from the control loops. *Chapter 5*.

self-optimizing control Choosing what to hold constant so that holding it keeps the plant near its best operation when disturbances come. *Chapter 8*.

snowball effect A large, slow swing of the flows in a recycle whose return gain is close to one. *Chapter 1*.

stripper A column that drives a dissolved gas out of a solvent, usually with heat, so that the solvent can absorb again. *Chapter 1*.

supervisory control A layer that sets the setpoints of the controllers below it rather than moving valves itself. *Chapter 1*.

synthesis loop A reactor, a separator and a recycle compressor in a ring, so that unconverted gas returns to the reactor. *Chapter 1*.

thermal coupling Joining two columns or column sections by their flows or by a shared exchanger, so that the heat of one serves the other. *Chapter 1*.

thermocouple A temperature sensor made of two different metals joined at a junction, which gives a small voltage that depends on temperature. *Chapter 4*.

throughput manipulator The one flow chosen to set a plant's production rate, around which every other inventory loop is arranged. *Chapter 1*.

tray temperature control Holding the temperature on one tray of a column as a fast stand-in for a product purity. *Chapter 8*.

two-step method The usual real-time optimization: fit the model to the measured steady state, then move to the fitted model's best point. *Chapter 1*.

Underwood method A shortcut that gives the least vapor a column needs for a sharp split, from the feed's composition and volatilities. *Chapter 1*.

urea stripper An exchanger in which heat and fresh carbon dioxide drive unconverted ammonia and carbon dioxide out of the reactor's effluent. *Chapter 3*.

valve positioner A device on a control valve that compares the stem position with the controller's signal and corrects the air to the actuator. *Chapter 1*.

vanadium catalyst The catalyst of the contact process, vanadium pentoxide with alkali sulfates on silica, which is a liquid melt when hot. *Chapter 4*.

viscosity index A number that says how little an oil's viscosity changes with temperature; a higher index is better. *Chapter 10*.

Williams–Otto reactor A benchmark stirred reactor with three reactions, used since 1960 to test optimization methods. *Chapter 11*.

Index of Control Topics

absorption, 8
adiabatic reactor, 4, 76
AI agents
  completion checks, 290
  engineering practice, 288
  open problems, 292
  reproduction, 290
  solver checks, 289
  unit checks, 290
  verification record, 289
  while studying, 18
analyzer delay, 193
averaging level control, 139, 169, 174

back-off, 273, 275
bias update, 10, 61, 197, 239
bumpless transfer, 16, 197
butterfly valve, 93
bypass control, 82

cascade control, 156
  fixing a recycle gain, 167
  pressure difference, 109
challenge scenario, 90
chemical equilibrium, 4, 58, 76
chemical plant control
  open problems, 291
computer process control
  history, 257
constant molar overflow, 191
constraint back-off, 233, 243
constraint control, 140, 200
constraint margin, 41
continuation, 218
continuous emission monitoring, 180
control valve
  steam, 227
controller action, 238
conversion–temperature diagram, 5, 77
Coriolis flowmeter, 227
creeping reaction front, 85

data export, 18
degrees of freedom, 187
demand response, 101, 174, 285
differential pressure transmitter, 117

distillation
  degrees of freedom, 9
distillation control, 124, 187
  one-point control, 227
dividing-wall column, 9

economic model predictive control, 170, 264
economic supervisory control, 156, 169
energy efficiency
  control structure, 285
equilibrium-limited reactor, 283

feedforward, 10, 174, 198
  from a measured disturbance, 112
Fenske equation, 194
final control element, 13
flexible operation, 47
flow scheme, 4
fold bifurcation, 219
four-point control structure, 190

Gilliland correlation, 194
gradient estimation, 263
  open problems, 276

heat integration, 60, 128
Hopf bifurcation, 219
hysteresis, 222

ignition and extinction, 84
ill-conditioning, 197
infeasible setpoint, 66
inferential control, 10, 53, 163, 197, 210, 221, 233
  pressure compensation, 173
infrared gas analyzer, 179
input multiplicity, 10, 210, 220
integrating process, 54, 134
inventory control, 8, 35

Kalman filter, 17, 37
  input disturbance, 141
Kremser relation, 131

lap time, 168
linear programming
  load scheduling, 109

liquid-to-vapor ratio, 196
loop interaction, 138, 197
loop pairing, 138, 197

measurement noise
  gradient estimates, 268
mid-ranging control, 233
  as a price, 245
  history, 236
  open problems, 249
  tuning, 248
model limitations, 94
model predictive control, 63, 140, 198, 239
  against supervisory loops, 285
  disturbance model, 37
  in the apps, 17
  move blocking, 140
  price preview, 177
  setpoint selection, 87
modifier adaptation, 11, 263
moving horizon estimation, 17
multiple steady states, 10, 27, 210, 219
  in distillation, 211
  margin, 292

numerical integration
  implicit reaction step, 217

online analyzer
  composition, 287
optimum inlet temperature, 6
override control, 140, 239
  low selector, 37
oxygen analyzer
  paramagnetic, 149

PID control
  anti-windup, 16, 139, 168
  velocity form, 16, 63, 86, 168, 221
plant–model mismatch, 10, 255
plantwide control, 11, 24, 236
  design procedure, 285
  recycle, 283
power-to-ammonia, 287
process intensification, 287
purge, 7

ratio control, 10, 54, 87, 221
  feeds in ratio to a current, 99
  in a recycle, 7, 62
  with an analyzer trim, 108
reaction rate
  inhibition, 79
reactive dividing-wall column, 228
real-time optimization, 10, 48, 255
  modifier adaptation, 287
  open problems, 276
  two-step method, 262
rectifier, 118
recycle, 7, 63
  gain and time constant, 7
  gain decay, 157
  negative return gain, 8, 156
recycle systems
  amplification, 25
  inert accumulation, 34
regeneration energy, 9
relative gain array, 138, 197
relative volatility, 191
resistance temperature detector, 202

safety interlock, 106
sampled analyzer, 87, 217, 258
scenarios, 17, 241
  completion rules, 17
scorecard, 18
self-optimizing control, 188
  history, 190
sequential quadratic programming, 263
simulation engine, 14
simulator framework, 14
snowball effect, 34, 54
soft sensor, 54, 243
  maintenance, 71
spare degree of freedom, 194
steady state
  convergence, 290
steady-state detection, 262
steady-state gain
  sign change, 220
step-response model, 198
stripping, 8
supervisory control, 221, 275

thermal coupling, 9, 124, 189
thermocouple, 92
three-point control structure, 197
throughput manipulator, 8, 25
transport delay, 156
  volume-keyed, 162
tray temperature control, 204

unattended operation, 288

Underwood method, 9, 194

valve position control, 235
valve positioner, 46

variable-speed drive, 70, 179

wrong-way behavior, 86

Index of Applications

absorber, 159
    temperature bulge, 162
acid strength, 83
adsorption term, 216
air separation, 9, 124
    demand response, 151
    history, 127
amine emission, 164
amine scrubbing, 8, 156
    history, 158
ammonia synthesis, 24
    history, 3, 27
    startup, 47
ammonium carbamate, 55
analyzer, 193, 258
argon sidearm, 126
autothermal reactor, 27
azeotrope, 216

base oil, 233
biuret, 60
    limit, 70
branch monitor, 221
bubble point, 215

C4 cut, 211
carbamate condenser, 56
carbon capture, 156
    flexible operation, 181
    history, 159, 287
catalyst deactivation, 41, 89
caustic soda
    concentration, 119
cell-voltage monitor, 110
chemical plants, 1
    history, 3
chlor-alkali
    history, 3
chlor-alkali electrolysis, 99
chlorine compressor trip, 110
circulator, 29
condenser-reboiler, 124
    explosions, 130
    submergence, 149
constant molar overflow, 215
contact process, 76
    history, 3, 79
continuous emissions monitor, 83
    practice, 93
cooling water temperature, 175
corrosion index, 60
crude argon, 133
current efficiency, 102
cut point, 237

decene oligomerization, 235
density gauge, 69
dilution water, 87
distillation column, 187
dividing-wall column, 187
    history, 189
    operation, 204
double absorption, 78
double column, 126

emission limit, 78
energy penalty, 164
energy valley, 194
equilibrium stage, 215
etherification, 212
evaporator, 60

feed and effluent exchanger, 28
flexible capture, 159

gas chromatograph
    process, 45
gas header, 106
gas treating, 158
Gillespie-Beattie correlation, 30
guide vanes, 130
    hardware, 149

heat of reaction, 215

ignition floor, 84
instrument, 13
interpass absorber, 79
ion exchange
    brine polishing, 101
ion-exchange resin, 212
isobutene, 211

liquid split, 190
liquid splitter, 203

main air compressor, 130
membrane cell, 100
  hardness damage, 104
  history, 102
metallocene catalyst, 235
methanol excess, 224
monoethanolamine, 157
MTBE, 210

near-infrared analyzer, 274

on-line viscometer, 239
  principles, 248
oxygen-depolarized cathode, 287

Petlyuk column, 189
polyalphaolefin, 233
  history, 235
prefractionator, 190
product tank, 237
purge, 28
  economics, 40
  history, 27

quench converter, 28

reactive distillation, 10, 210
  industrial use, 212
reactor jacket
  split-ranged valves, 248
reboiler, 200
Refutas blending index, 237
regeneration energy, 163
rolling emissions obligation, 169

safety interlock
  practice, 118
Schulz–Flory distribution, 236
side draw, 190
smelter gas, 94
solvent degradation, 164
solvent storage, 174
specific power, 134
start-up, 224
start-up emissions, 89
stirred tank reactor, 258
stripper, 159
sulfur burner, 82
sulfur dioxide oxidation, 76
sulfuric acid plant, 76
surge tank, 169
synthesis loop, 7

tray temperature, 197
turndown, 85

urea stripper, 56
urea synthesis, 53
  history, 56
  total recycle, 55

vanadium catalyst, 78
vapor–liquid equilibrium
  amine solution, 160
viscosity index, 235

waste-heat boiler, 79
Williams–Otto reactor, 255
Wilson activity model, 215

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.