INTERACTIVE CONTROL ENGINEERING

Control in Upstream and Midstream Oil and Gas

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 10

Control in Upstream and Midstream Oil and Gas

From the wellbore to the terminal

John D. Hedengren
Editor

Advanced Process Solutions, LLC

Interactive Control Engineering: Principles and Applications
Volume 10: Control in Upstream and Midstream Oil and Gas
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/v10-oil-gas/

How to cite this volume. John D. Hedengren, Ed., *Control in Upstream and Midstream Oil and Gas*, vol. 10 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 follows oil and gas from the wellbore to the terminal. Chapters 2 to 5 cover drilling and production: pressure held between two limits while drilling, a rod pump read through its downhole card, a smart well, and a deepwater flowline kept clear of hydrates. Chapters 6 to 10 cover separation and gas processing. Chapters 11 to 15 cover compression, transmission, product pipelines and liquefied natural gas.

The plants run far from their operators and close to their pressure limits, and many of the quantities that matter are inferred. Each chapter names the technology volume that teaches its main technique.

John D. Hedengren
2026

Contents

Preface v

Notation x

1 Introduction: Pressure, Inventory and Limits From Well to Market 1
- 1.1 The Question This Volume Answers . . . 1
- 1.2 A Short History . . . 2
- 1.3 Notation and Conventions . . . 4
- 1.4 From Reservoir to Market on One Well . . . 4
- 1.5 Instruments and Final Elements in This Volume . . . 12
- 1.6 The Simulator Framework . . . 15
- 1.7 Working With AI Agents While Studying . . . 18
- 1.8 Chapter Roadmap . . . 19
- 1.9 Exercises . . . 22

2 A Managed Pressure Drilling Simulator to Teach Optimization on a Constraint Boundary 24
- 2.1 Introduction . . . 25
- 2.2 Background . . . 26
- 2.3 Process Model . . . 27
- 2.4 Feasible Region Analysis . . . 34
- 2.5 Controller Design . . . 37
- 2.6 Results . . . 40
- 2.7 Technology in Practice . . . 46
- 2.8 Limitations . . . 48
- 2.9 Conclusion . . . 49
- 2.10 Exercises . . . 51

3 Hydraulic Rod Pump Simulator for Downhole Card Inference and Constrained Rate Control 55
- 3.1 Introduction . . . 56
- 3.2 Background . . . 57
- 3.3 Process Model . . . 58
- 3.4 Card Reconstruction and Diagnosis . . . 63
- 3.5 Rate Control . . . 66
- 3.6 Scenarios and Results . . . 68
- 3.7 Technology in Practice . . . 72
- 3.8 Limitations . . . 74
- 3.9 Conclusion . . . 75
- 3.10 Exercises . . . 76

4 Fast Choke Overrides and Slow Waterflood Planning in a Smart Well Simulator 80
- 4.1 Introduction . . . 81
- 4.2 Waterflood Sweep, Mobility Control and Smart Wells . . . 82
- 4.3 A Producer, an Injector and Two Zones . . . 83
- 4.4 Time Scales of the Well and the Flood . . . 89
- 4.5 Control Layers as Implemented . . . 90
- 4.6 Scenarios and Measured Results . . . 93

4.7 Technology in Practice . . . 98
4.8 Limitations . . . 100
4.9 Conclusion . . . 100
4.10 Exercises . . . 102

5 A Deepwater Flowline Simulator for Hydrate Control at the Weakest Point 106
5.1 Introduction . . . 107
5.2 Hydrates, Flowlines and Inhibitors . . . 108
5.3 Model of the Flowline and Riser . . . 109
5.4 The Weakest Point and the Liquid Transit . . . 114
5.5 Instruments, Estimator and Control . . . 116
5.6 Measured Cases . . . 120
5.7 Technology in Practice . . . 127
5.8 Limitations . . . 129
5.9 Conclusion . . . 130
5.10 Exercises . . . 132

6 Offshore Three-Phase Separator Simulator for Slug Absorption and Delayed Quality Measurement 135
6.1 Introduction . . . 136
6.2 Separator Sizing, Riser Slugs and Surge Control . . . 137
6.3 Separator Model . . . 138
6.4 One Inventory, Three Outlets . . . 143
6.5 Implemented Controller . . . 144
6.6 Measured Cases . . . 148
6.7 Technology in Practice . . . 155
6.8 Limitations . . . 158
6.9 Conclusion . . . 158
6.10 Exercises . . . 160

7 Glycol Dehydration Simulator With an Equilibrium Floor and a Reboiler Ceiling 164
7.1 Introduction . . . 165
7.2 Dry Gas for the Pipeline . . . 166
7.3 A Contactor and Its Regenerator as Modeled . . . 168
7.4 A Fast Handle, a Slow Handle and Two Limits . . . 174
7.5 The Controllers in the Code . . . 175
7.6 Nine Cases on the Running Engine . . . 177
7.7 Technology in Practice . . . 181
7.8 Limitations . . . 183
7.9 Conclusion . . . 183
7.10 Exercises . . . 185

8 Selective Amine Sweetening With Delayed Analysis and Loading Protection 189
8.1 Introduction . . . 190
8.2 Selective Treating and Its Modeling History . . . 191
8.3 Unit, States and Constitutive Relations . . . 192
8.4 Measurements and Control . . . 196
8.5 Nine Cases on the Engine . . . 198
8.6 Technology in Practice . . . 202
8.7 Limitations . . . 204
8.8 Conclusion . . . 204
8.9 Exercises . . . 206

9 Ethane Recovery and Rejection Through a Sequenced Turboexpander Plant 210
9.1 Introduction . . . 211
9.2 Process Background and Prior Work . . . 212
9.3 Plant Model . . . 212
9.4 Regulation, Sequencing and Economic Supervision . . . 217
9.5 Cases Measured on the Simulation Engine . . . 218
9.6 Technology in Practice . . . 221
9.7 Limitations and Interpretation . . . 224
9.8 Conclusion . . . 225
9.9 Exercises . . . 226

10 An NGL Fractionation Train Simulator for Plantwide Pairing and Shared Constraints 230
10.1 Introduction . . . 231
10.2 Background . . . 232
10.3 Process Model . . . 233
10.4 Controller Design . . . 236
10.5 Results . . . 239
10.6 Technology in Practice . . . 245
10.7 Limitations . . . 247
10.8 Discussion . . . 248
10.9 Conclusion . . . 248
10.10 Exercises . . . 250

11 Load Sharing by Equal Surge Margin Among Three Compressors on One Header 253
11.1 Introduction . . . 254
11.2 Parallel Machines in Gas Transmission . . . 255
11.3 The Station as Modeled . . . 256
11.4 Control Layers as Implemented . . . 260
11.5 Nine Cases on the Running Engine . . . 263
11.6 Technology in Practice . . . 268
11.7 Limitations . . . 270
11.8 Conclusion . . . 271
11.9 Exercises . . . 272

12 A Gas Transmission Line Simulator for Line Pack, Forecast Feedforward and Leak Detection 276
12.1 Introduction . . . 277
12.2 Gas Transmission Operations . . . 278
12.3 Process Model . . . 278
12.4 Control and Monitoring . . . 281
12.5 Results . . . 284
12.6 Technology in Practice . . . 290
12.7 Limitations . . . 292
12.8 Conclusion . . . 293
12.9 Exercises . . . 295

13 Series Pump Stations, Batches and Surge on a Refined Products Pipeline 298
13.1 Introduction . . . 299
13.2 Batches, Surge and Leak Monitoring on Liquid Lines . . . 300
13.3 The Line as Modeled . . . 301
13.4 The Control Layers as Implemented . . . 305
13.5 Nine Cases, Measured . . . 308

13.6 Technology in Practice . . . 312
13.7 Limitations . . . 315
13.8 Conclusion . . . 315
13.9 Exercises . . . 317

14 An LNG Liquefaction Simulator for Loop Pairing and Inferential Constraint Control 321
14.1 Introduction . . . 322
14.2 Background . . . 323
14.3 Liquefaction Train Model . . . 325
14.4 Pairing Analysis . . . 332
14.5 Implemented Control Structure . . . 333
14.6 Results . . . 337
14.7 Technology in Practice . . . 344
14.8 Limitations . . . 346
14.9 Discussion . . . 347
14.10 Conclusion . . . 347
14.11 Exercises . . . 349

15 An LNG Import Terminal Whose Boil-Off Can Only Leave With Its Send-Out 353
15.1 Introduction . . . 354
15.2 Boil-Off, Recondensers and Rollover at Receiving Terminals . . . 355
15.3 The Terminal as Modeled . . . 356
15.4 Control Layers as Implemented . . . 363
15.5 Nine Cases on the Running Engine . . . 365
15.6 Technology in Practice . . . 369
15.7 Limitations . . . 371
15.8 Conclusion . . . 372
15.9 Exercises . . . 374

16 Summary and Outlook 377
16.1 What the Chapters Showed . . . 377
16.2 Choosing a Method . . . 381
16.3 Where Practice Is Heading . . . 382
16.4 Engineering With AI Agents . . . 383
16.5 Gaps in Oil and Gas Control . . . 386
16.6 Exercises . . . 390

Answers to Selected Exercises 391

Glossary 397

Index of Control Topics 400

Index of Applications 403

About the Editor 407

Volumes in This Series 408

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: Pressure, Inventory and Limits From Well to Market

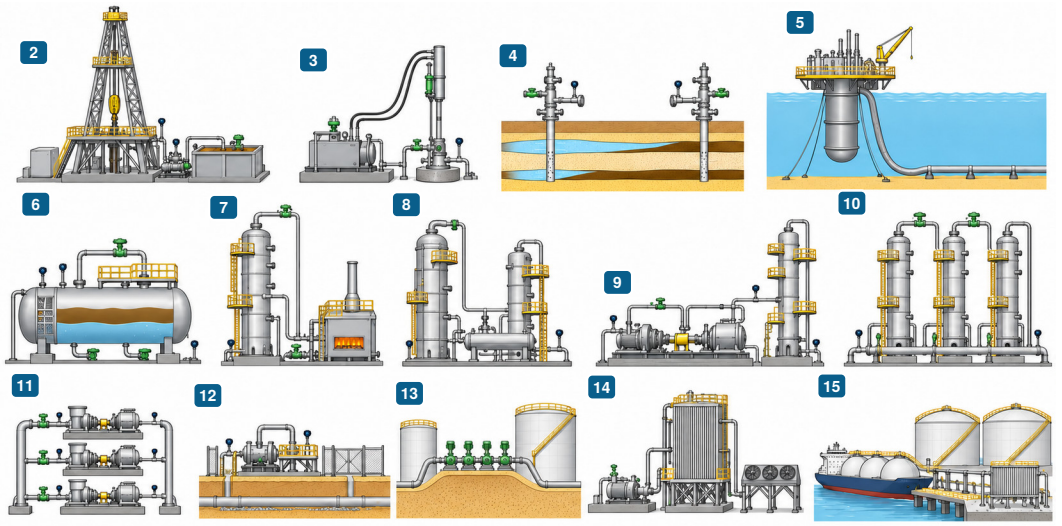

The chapters 2 drilling rig 3 rod-pumped well 4 smart well and flood 5 subsea flowline 6 offshore separator 7 glycol dehydrator 8 amine sweetener 9 turboexpander plant 10 NGL train 11 compressor station 12 gas pipeline 13 products pipeline 14 LNG liquefaction 15 LNG terminal

Notice that the rear row drills, lifts and floods the wells and brings their fluid up a riser, that the middle row separates it, treats the gas and splits its liquids, and that the front row moves gas and refined products by pipeline and gas by ship as LNG. Illustration, not to scale.

IN PLAIN TERMS **Moving oil and gas from the rock to the market**

The main idea. Oil and gas travel from a well through pipes, separators, cleaning plants and long pipelines before they are sold. At most steps a valve, a pump or a machine sets a pressure, and the flow then follows from the rock, the pipe or the buyer.

The challenge. Liquid arrives in surges, and a gas compressor gets into trouble when its flow falls too low. The limit that matters often lies deep in a well or inside a pipe where no instrument can reach, or it is known only minutes late.

Why it matters. Good control keeps people safe and pipes whole, and it decides how much oil and gas reaches the market.

THREE WAYS TO READ THIS CHAPTER

Explore	7 min	This page, Section 1.2 with its timeline and two History boxes, Figure 1.2, the Did you know? boxes in Sections 1.4.2, 1.4.4 and 1.4.5, and Key points in Section 1.8. No equations.
Study	40 min	The whole chapter in order, working each number of the example W-1, then Exercises 1 to 10.
Practice	10 min	This page, Sections 1.4.8 and 1.4.9 with Table 1.2, Section 1.5 with Figures 1.7 and 1.10, and Section 1.8.

1.1 The Question This Volume Answers

Oil and gas reach a market through a chain of plants that were never designed as one. A well is drilled and produced, its fluid crosses a flowline to a separator, the gas is treated, compressed and sent down a pipeline or liquefied, and the liquids leave by their own lines. Each plant in the chain has

a few actuators, many limits and fewer instruments than its limits need [1]. This volume answers one question: which control techniques recur along that chain, and where does what can be measured, rather than what can be actuated, set the limit of what control achieves?

Six control problems define the industry, and each chapter meets at least two of them.

- *Pressure is the handle and the flow follows.* A choke, a pump or a compressor sets a pressure, and the rate follows from a reservoir, a line or a market that the actuator cannot change.
- *Flow arrives in slugs and ramps.* Multiphase lines deliver liquid in surges, and gas demand rises every morning. Separators, surge drums, pipelines and tanks store the difference.
- *Machines work near an instability.* A centrifugal compressor surges when its flow falls too far, and a recycle valve that keeps it out of surge compresses gas for nothing.
- *Products are judged late.* Water, hydrogen sulfide, methane in a liquid product and ethane in propane are read by analyzers whose reports arrive minutes after the gas or liquid they describe.
- *The binding limit sits where no instrument is.* A hydrate margin at the foot of a riser, a fluid level at a pump nine hundred meters down and a temperature approach inside a welded exchanger are computed, not measured.
- *Many limits share one actuator.* A choke protects drawdown, sand, water and pipe pressure at once, and a pump station protects suction and discharge. Selectors and overrides decide which limit owns the actuator.

Fourteen chapters answer the question on fourteen plants, each with an app that the reader runs in a browser. Chapters 2 to 5 are upstream: a well being drilled, a rod-pumped well, a waterflood with a smart producer and a deepwater flowline. Chapters 6 to 10 follow the fluid through separation, gas treating and the recovery and fractionation of natural gas liquids. Chapters 11 to 15 are midstream: a compressor station, a gas transmission line, a refined products line and the two ends of the liquefied natural gas (LNG) chain.

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

- Find the operating point of a well and its choke, and say how much of a choke move reaches the rate.
- Judge from a few numbers whether a riser will slug, and size the level control that absorbs a slug without passing it downstream.
- Compute the line pack of a gas line and say why a liquid line needs overrides that a gas line does not.
- Read a compressor map, place the surge and control lines, and find the throughput below which the recycle opens.
- Explain the floor that the lean solvent sets on an absorber and the cold that an expander makes and a valve does not.
- Build an override on a shared actuator so that it takes over at the limit, and say what each limit is known through.

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, selectors and overrides in general from Volume 4 and predictive control from Volume 9.

1.2 A Short History

The milestones of Figure 1.1 come from three places: the gas plant and its lines, the pumped well and the drill floor. In 1930 a patent described a small group of amines that unite with acidic gases at a low temperature, give them up at a higher one and so circulate in a closed cycle [2]. In 1934 the crystalline plugs found in flowlines above the freezing point of water were shown to be hydrates. An equation for the conditions that form them followed, with inhibitors such as methanol and glycol [3]. Chapter 7 dries gas so that hydrates cannot form, and Chapter 5 doses methanol against them.

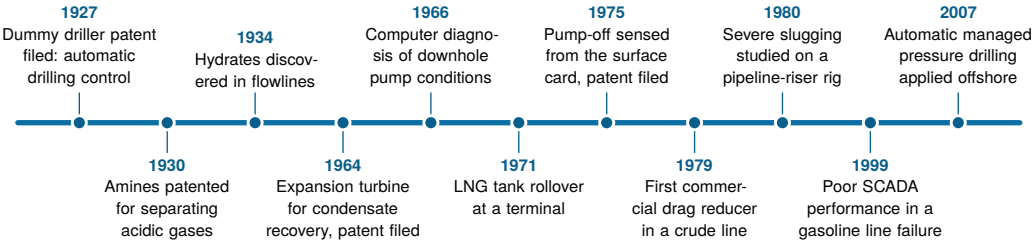

Figure 1.1 Milestones of upstream and midstream oil and gas and of their control. The dates are from [4] (1927), [2] (1930), [3] (1934), [5] (1964), [6] (1966), [7] (1971), [8] (1975), [9] (1979), [10] (1980), [11] (1999) and [12] (2007). A patent is dated by its filing where the card says so. Notice that the drill floor holds the first and the last milestone of the eighty years. Events are evenly spaced, not to a time scale.

About half of the gas transmission network of the United States was installed in the 1950s and 1960s, after the demand for natural gas more than doubled following the Second World War [13]. Gas plants learned to make their cold with a turbine. A patent filed in 1964 expanded cooled feed gas through a turbine that tolerates condensate and used the turbine's work to recompress the residue gas [5], the arrangement of Chapter 9. In 1966 a computer method turned the load and position measured at the top of a rod string into the card at the pump [6]. A patent filed in 1975 sensed that a well had pumped off from a fall in the power input, which the area of the surface card measures, so that the pump could be stopped [8]. Both rest on the two measurements of Chapter 3. The drill floor was automated in steps of its own.

> ⧖ HISTORY **1927 to 2007: Automating the Drill Floor**
>
> A patent filed in 1927 set out to keep a set pressure on a rotary drill hung partly from a cable. A reversible motor paid the cable out or picked it up, and too much pull made it pay out until the strain on the cable fell [4]. A patent filed in 1952 let the tension in the drilling line set how fast the drum unspooled. When the hydraulic pressure that read the tension fell, the bit carried too much weight, the drum slowed and the bit drilled off [14]. In 2007 a choke at the surface held the downhole pressure of a well drilled offshore under automatic control, because narrow drilling margins ask for highly accurate pressure control [12]. Chapter 2 closes all three loops, on weight, on rotation and on pressure.

Pipelines and terminals taught their lessons through events. In 1971 a cargo loaded beneath colder liquefied gas in a terminal tank formed two layers. The layers later mixed suddenly, and about 2000 tonnes of vapor left through the safety valves over a few hours [7]. That rollover is a case of Chapter 15. A polymer that reduces drag first raised the flow of a crude oil line commercially in 1979 [9]. In 1980 a test rig showed that a riser slugs at low gas and liquid rates and that choking removes the slugging [10]. The separator of Chapter 6 takes the slugs as they come, because its wells deliver an imposed stream.

> ⧖ HISTORY **1999: A Control Room That Could Not See**
>
> A pipeline is run from a control room. A control server at a control center gathers field data for the operator, while remote terminal units or programmable controllers at the field sites run the local loops [11]. In June 1999 a 16-inch gasoline line leaked 237 000 gallons, which ignited 1.5 hours later and caused three deaths. During the event the supervisory system performed poorly, which kept the pipeline controllers from seeing an abnormal operation develop. The investigators recommended an offline system on which to test changes to the supervisory database [11]. Chapter 12 and Chapter 13 run below such a room.

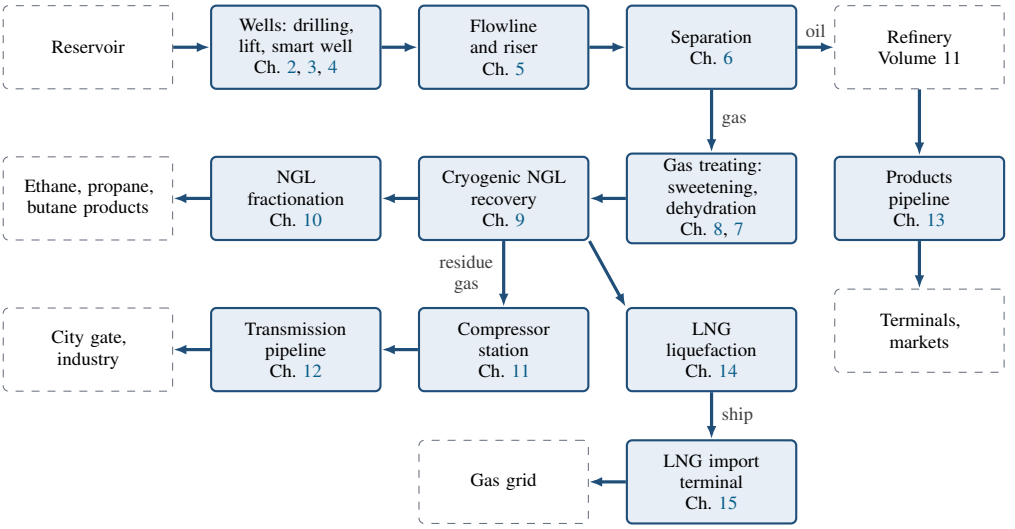

Figure 1.2 The Chain From Reservoir to Market. Filled Blocks Are the Plants of Chapters 2 to 15, Dashed Blocks Lie Outside the Volume. A Liquefaction Plant Usually Treats Its Own Feed Gas.

1.3 Notation and Conventions

Pressures are in bar. A chapter writes barg or bara where gauge and absolute differ, and this chapter uses absolute pressure throughout. Liquid rates are in m^3/h at line conditions, gas in standard cubic meters (Sm^3) or in kg/s, and temperatures in °C. The reservoir pressure is p_R, the flowing bottomhole pressure p_{wf}, the wellhead pressure p_{wh}, the productivity index J and the choke opening x in percent. The liquid holdup, the share of a pipe's volume that liquid fills, is H_L. A separator level is ℓ, because the framework's h is its integration step. A compressor's flow coefficient is ϕ, its head coefficient ψ, its speed n as a fraction of design and its surge margin SM. An absorber's absorption factor is $\mathscr{A}$. The framework's a is the controller action, so a wave speed here is c.

The chapters keep their own symbols, and several collide. Chapter 6 writes level heights as h, and Chapter 5 writes the holdup as H_i. Chapter 11 uses H and W for the shape of its compressor map. Chapter 13 writes the wave speed a, and Chapter 12 writes the isothermal wave speed c, as here. Chapter 3 writes the controller action α and Chapter 7 writes it σ. Several chapters give water content in pounds per million standard cubic feet, the unit of the trade, and this chapter uses mg/Sm^3.

1.4 From Reservoir to Market on One Well

1.4.1 The Path From Reservoir to Market

Figure 1.2 draws the chain and marks where each chapter's plant sits in it. Drilling makes the well, and the well brings oil, gas and water to the surface by its own pressure or with a pump. A flowline carries the mixture to a separator, which splits it into three streams. Oil goes to a refinery, whose control is the subject of Volume 11, and the refined products leave by pipeline. Gas is sweetened and dried, its heavier hydrocarbons are recovered at low temperature and fractionated into products, and the remaining residue gas is compressed into a transmission line or liquefied for a ship. At the other end of the voyage a terminal returns the LNG to gas.

The chapters take the plants one at a time, and the techniques they share are developed below on one running example. The example is a single oil well, W-1, on a small platform, followed from the reservoir to its export lines. The well produces through a choke into a riser and a separator, and its gas is compressed into a short export pipeline. Every number of the example is declared here or computed from these declarations by the script kept with this chapter's record.

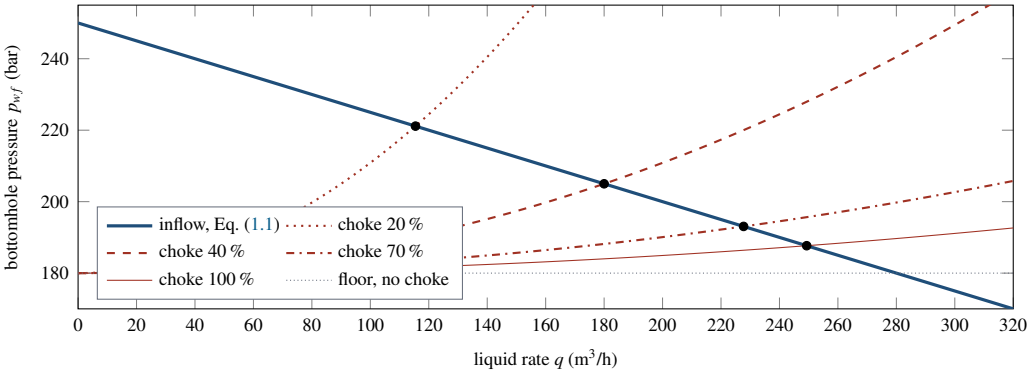

Figure 1.3 Operating Points of W-1. The Inflow Line (Thick Solid) Meets the Outflow Curve of Each Choke Opening (Dotted 20, Dashed 40, Dash-Dotted 70, Thin Solid 100 %) at the Dots. No Opening Takes the Bottomhole Pressure Below the Densely Dotted Floor.

1.4.2 Wellbore, Choke and Flowline Pressure

A reservoir delivers liquid in proportion to the pressure difference that draws it into the well. With the straight-line productivity index, the simplest inflow model, the rate is Equation (1.1).

$$q = J\,(p_R - p_{wf}) \tag{1.1}$$

where $p_R - p_{wf}$ is the drawdown. A well whose oil releases gas below its bubble point delivers less than the straight line at large drawdowns [15], and the example ignores this. W-1 has $p_R = 250\,\text{bar}$ and $J = 4\,\text{m}^3/\text{h}$ per bar.

The fluid then climbs the tubing, crosses the choke and flows to the separator. Each step costs pressure, as shown in Equation (1.2).

$$p_{wf} = p_{wh} + \Delta p_{col}, \qquad q = c\,x\sqrt{p_{wh} - p_{dn}}, \qquad p_{dn} = p_{sep} + \Delta p_{riser} + \Delta p_{fr} \tag{1.2}$$

where Δp_{col} is the weight of the column in the tubing, the choke is a valve with a linear characteristic and coefficient c, p_{dn} is the pressure just downstream of it, and the flowline adds the separator's pressure, the weight of the riser's contents and friction. W-1 has a tubing column of 120 bar, 2000 m of fluid at 0.06 bar/m, a separator at 10 bar and a riser that weighs 50 bar while it flows. Tubing and flowline friction are neglected, so $p_{dn} = 60\,\text{bar}$, and $c = 0.9$ in m^3/h per percent of opening per square root of a bar.

The operating point is where the reservoir delivers what the choke passes. At $x = 40\,\%$ the two meet at $q = 180\,\text{m}^3/\text{h}$, with $p_{wf} = 205\,\text{bar}$, $p_{wh} = 85\,\text{bar}$, 25 bar across the choke and a drawdown of 45 bar. Figure 1.3 draws the inflow line and the outflow curves for four openings. The curves have a floor. With no pressure lost at the choke the bottomhole pressure cannot fall below $\Delta p_{col} + p_{dn} = 180\,\text{bar}$, and the well cannot produce more than $J(p_R - 180) = 280\,\text{m}^3/\text{h}$. The choke wide open still takes 7.7 bar and gives $249.3\,\text{m}^3/\text{h}$.

A choke move therefore reaches the rate only in part. Opening the choke raises the flow through it at once, by $c\sqrt{p_{wh} - p_{dn}} = 4.5\,\text{m}^3/\text{h}$ per percent at the design point. The extra flow drains the wellbore, the bottomhole pressure falls, and the inflow takes back part of the increase. Linearizing Equations (1.1) and (1.2) gives the share that remains at steady state, as shown in Equation (1.3).

$$\frac{\Delta q_\infty}{\Delta q_0} = g = \frac{J}{J + \partial q/\partial p_{wh}}, \qquad \frac{\partial q}{\partial p_{wh}} = \frac{c\,x}{2\sqrt{p_{wh} - p_{dn}}} \tag{1.3}$$

At $x = 40\,\%$ the choke passes $3.6\,\text{m}^3/\text{h}$ more for each bar of wellhead pressure, so $g = 4/7.6 = 0.526$, and the steady gain is $2.37\,\text{m}^3/\text{h}$ per percent. The response is a lead-lag, a jump of 4.5 that relaxes to

2.37 as the wellbore drains. With a wellbore that stores 1.25 m^3 per bar, the relaxation has a time constant of $60 \times 1.25/7.6 = 9.87$ min. Wider open, the choke takes a smaller share of the pressure and the share g falls, so a rate loop tuned at one opening is slower at a wider one.

The choke sets a pressure, and the rate follows from a reservoir pressure the choke cannot raise. A rate loop holding 180 m^3/h opens the choke as the reservoir depletes: 51.6 % at 240 bar, 63.3 % at 235 bar and 89.4 % at 230 bar. At 229 bar it is wide open and the target is lost. Chapter 4 meets the same lead-lag when it reads the wellbore from the choke, and the same loss of travel when its injection plant trips. Chapter 5 balances a well, a choke and a line in the same way. Chapter 2 writes the bottomhole pressure of a well being drilled as the same column plus friction plus a choke's back pressure, held between the pore and fracture pressures, and there closing the choke raises the pressure, a negative process gain. In Chapter 3 the column above the pump turns the pump's intake pressure into a fluid level.

> **+ DID YOU KNOW? Most Oil Wells Need a Pump**
>
> More than seventy percent of the oil wells now producing are lifted artificially [16]. In 1999 sucker rod pumps, the pumps of Chapter 3, were installed in about 90 percent of all oil wells in the United States [17].

1.4.3 Multiphase Flow and Slugging

The riser's weight depends on how much of it is liquid. Gas moves faster than liquid up a pipe, so the liquid holdup H_L exceeds the liquid's share of the volume flow, and the mixture density is Equation (1.4).

$$\rho_m = H_L \rho_L + (1 - H_L)\, \rho_g, \qquad \Delta p_{riser} = \rho_m\, g\, z_r \tag{1.4}$$

Holdup correlations for inclined pipes give H_L from the flow rates [18]. On W-1, oil at 800 kg/m^3 and gas at 60 kg/m^3 give the flowing riser's 0.05 bar/m at a holdup of 0.608 over the riser's 1000 m. A riser full of liquid weighs 78.5 bar instead of 50 bar, and the well behind it feels the 28.5 bar at once. At the same choke opening the rate falls from 180 to 120.9 m^3/h, 32.8 % less. The coupling is a positive feedback. Less flow means slower gas, more holdup, a heavier riser and less flow again.

Severe slugging is the cycle this feedback drives on a flowline that runs down to the foot of a riser [19]. Liquid collects at the riser base and seals it. The gas behind the seal is trapped in the flowline and compresses while the liquid column grows. When the gas pressure exceeds the column's weight, the gas pushes the liquid out in one piece and the cycle starts again. Whether the seal grows faster than the gas can follow decides whether a slug forms. For an isothermal ideal gas trapped in the flowline's gas volume V_g, the two rates compare as in Equation (1.5).

$$\Pi = \frac{\mathrm{d}p_g/\mathrm{d}t}{\mathrm{d}p_{col}/\mathrm{d}t} = \frac{p_s\, Q_{g,s}/V_g}{\rho_L\, g\, q_L/A_r} \tag{1.5}$$

where $p_s\, Q_{g,s}$ is the gas flow at standard pressure, q_L the liquid flow and A_r the riser's bore. A ratio below one means the column outruns the gas and the riser slugs, a criterion of the kind published for flowline and riser systems [20]. W-1's 0.2 m riser fills in 628 s at 180 m^3/h, so its column grows by 0.125 bar/s. Its 18 000 Sm^3/h of gas in 150 m^3 of flowline gas volume raises the trapped pressure by only 0.034 bar/s, so $\Pi = 0.270$ and the riser slugs. It would need 3.70 times the gas, 66 600 Sm^3/h, to keep up. Gas injected at the riser base, a smaller gas volume upstream and feedback that throttles the topside choke on the riser base pressure are the remedies. The last has been tested on flowlines in service [21].

The flowline also delays whatever enters it. Liquid moves as a column, so a change made at the well reaches a point downstream once all the liquid between them has been pushed past. W-1's 5 km flowline of 0.25 m bore holds 147.3 m^3 of liquid at a holdup of 0.6, which takes 49.1 min to displace at 180 m^3/h and 98.2 min at half the rate. The delay is a volume divided by a flow, so it doubles when the flow halves. Chapter 5 doses methanol at the tree and judges it at the riser base hours later,

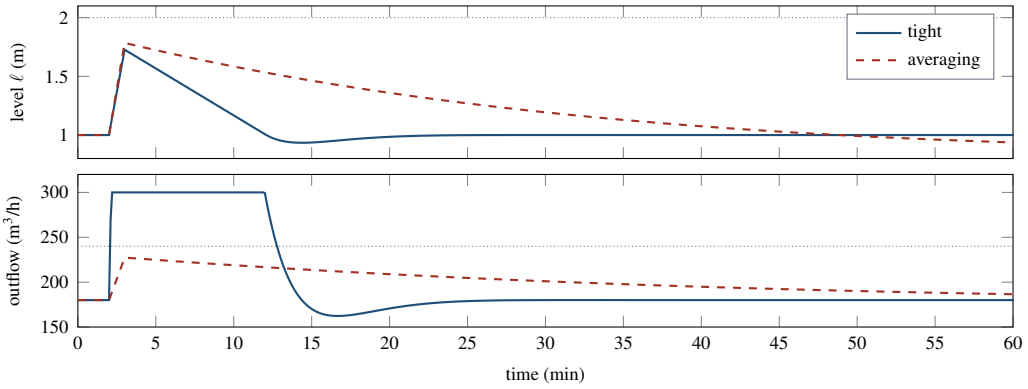

Figure 1.4 A 20 m^3 Slug Into W-1's Separator. Tight Tuning (Solid) and Averaging Tuning (Dashed) Store Nearly the Same Volume. The Dotted Lines Are the High-High Level Trip and the 240 m^3/h Limit of the Equipment Downstream.

through exactly this delay, and its constraint, the hydrate margin, is smallest where the riser's weight keeps the pressure high [22]. Chapter 6 receives the riser's slugs, and Chapter 13 carries its batches through the same displacement.

1.4.4 Separator Inventory and Averaging Level Control

The separator's liquid is an inventory between an inflow that the wells set and an outflow that a valve sets. W-1's separator has a liquid surface of 25 m^2, a normal level of 1.0 m, a high-high trip at 2.0 m, and an oil valve that passes 300 m^3/h wide open at the separator pressure. Its level is an integrator, as shown in Equation (1.6).

$$A_s \frac{\mathrm{d}\ell}{\mathrm{d}t} = q_{in} - q_{out}, \qquad q_{out} = q_{max}\,\frac{u}{100}, \qquad k = \frac{q_{max}}{100\,A_s} \tag{1.6}$$

so one percent of valve moves the level at $k = 0.12\,\mathrm{m/h}$. A PI controller on an integrator gives the characteristic polynomial $\tau_c\tau_I s^2 + \tau_I s + 1$ with $\tau_c = 1/(K_c k)$ and a damping ratio $\zeta = \frac{1}{2}\sqrt{\tau_I/\tau_c}$. Two tunings with the same damping, $\zeta = 0.894$, are compared. The tight one has $K_c = 400\,\%/\mathrm{m}$ and $\tau_I = 4\,\mathrm{min}$, so $\tau_c = 1.25\,\mathrm{min}$. The averaging one has $K_c = 20\,\%/\mathrm{m}$ and $\tau_I = 80\,\mathrm{min}$, so $\tau_c = 25\,\mathrm{min}$, twenty times slower.

A slug of 20 m^3 arrives at 2 min over one minute, an inflow of 1380 m^3/h against a valve that passes at most 300 m^3/h. Whatever the tuning, at least 18 m^3, 0.72 m of level, must be stored during the arrival. Figure 1.4 shows the two runs. The tight loop peaks at 1.727 m, with its valve wide open from 2.15 min, and returns within 5 cm of the setpoint by 11.5 min. It hands the slug to the second stage at the full 300 m^3/h, above a downstream limit of 240 m^3/h for 10.7 min. The averaging loop peaks only 5.8 cm higher, at 1.785 m, never passes more than 227.3 m^3/h and returns within 5 cm by 42.6 min. The largest rate of change of the outflow over six seconds falls from 900 to 47.9 m^3/h per minute. The vessel stores the slug under either tuning. The tuning decides only how fast the stored liquid is passed on.

DID YOU KNOW? Three Barrels of Water for Each Barrel of Oil

The water produced with oil rose from less than 30 million barrels a day in 1990 to about 100 million in 2015 [23]. Wells give up an estimated three barrels of water for every barrel of oil [24], and a separator like that of Chapter 6 must part them.

Averaging control was posed as an optimization: pass an inflow change to the outflow as smoothly as the level band allows [25]. For a step Δq in the inflow, the smoothest response ramps the outflow at a constant rate r. The level then gains the area of a triangle, $\Delta q^2/(2r)$ in volume, which must fit in the band $\Delta\ell$, as shown in Equation (1.7).

$$r \geq \frac{\Delta q^2}{2 A_s \Delta \ell} \tag{1.7}$$

A second well of $60\,\mathrm{m^3/h}$ brought on line, with 0.5 m of band, can be passed on at $2.4\,\mathrm{m^3/h}$ per minute, over 25 min. No tuning does better, and a tighter loop uses less of the band and ramps faster. Chapter 6 measures exactly this trade on a riser slug, with the extra cost that a higher level brings foam nearer the mist extractor. Chapter 10 detunes its surge drum's level loop so that the drum, not the three columns behind it, takes a rate change. Chapter 15 has an integrating inventory too, the vapor space of its tanks, whose outlet capacity is set by the send-out.

1.4.5 Pipeline Inventory: Line Pack and Pressure Waves

A pipeline is also an inventory, and whether it is a large one depends on what it carries. Gas is stored in proportion to its pressure. For an isothermal gas with wave speed $c = \sqrt{ZRT}$ the mass in a line of bore area A_p and length L is Equation (1.8).

$$M = \frac{A_p L \bar{p}}{c^2}, \qquad \frac{\mathrm{d}\bar{p}}{\mathrm{d}t} = \frac{c^2}{A_p L}\left(w_{in} - w_{out}\right) \tag{1.8}$$

where $\bar{p}$ is the mean pressure and w a mass flow. W-1's gas leaves through 50 km of 0.4 m line with $c = 350\,\mathrm{m/s}$, the speed Chapter 12 uses. The line holds 5.13 t of gas per bar, and 307.7 t at a mean of 60 bar. A deficit of 5 kg/s, a tenth of the line's 50 kg/s, lowers the mean pressure by 3.5 bar an hour, so a 2 bar margin lasts 34 min. The line rides through a change in demand on its own pack while the compressor catches up. In steady flow the square of the pressure falls linearly along the line, $p_{in}^2 - p_{out}^2 = K w^2$, and the pressure at the far end follows a change at the near end as a slow diffusion rather than a delay [26].

DID YOU KNOW? Three Million Miles of Gas Pipe

About 3 million miles of mainline and other pipelines link the gas fields and storage of the United States with consumers. In 2022 they delivered about 29.2 trillion cubic feet of natural gas to about 78.3 million consumers [13].

A liquid line stores almost nothing. Oil at $840\,\mathrm{kg/m^3}$ with a pressure wave speed of 1130 m/s in its pipe has an effective bulk modulus $\rho c^2 = 1.07\,\mathrm{GPa}$, so W-1's 20 km oil line of 0.3 m bore, $1414\,\mathrm{m^3}$, takes up only $0.132\,\mathrm{m^3}$ per bar. A mismatch of $18\,\mathrm{m^3/h}$, a tenth of the oil, moves the pressure by 2.3 bar a minute. A sudden stop is worse. A column moving at velocity v that is stopped raises the pressure by the water hammer rise of Equation (1.9) [27].

$$\Delta p = \rho\, c\, \Delta v \tag{1.9}$$

which is 19.0 bar for a 2 m/s stop, arriving at the other end of the 20 km line 17.7 s later. The two lines need different control. The gas line's pack gives the controller time and makes a forecast worth acting

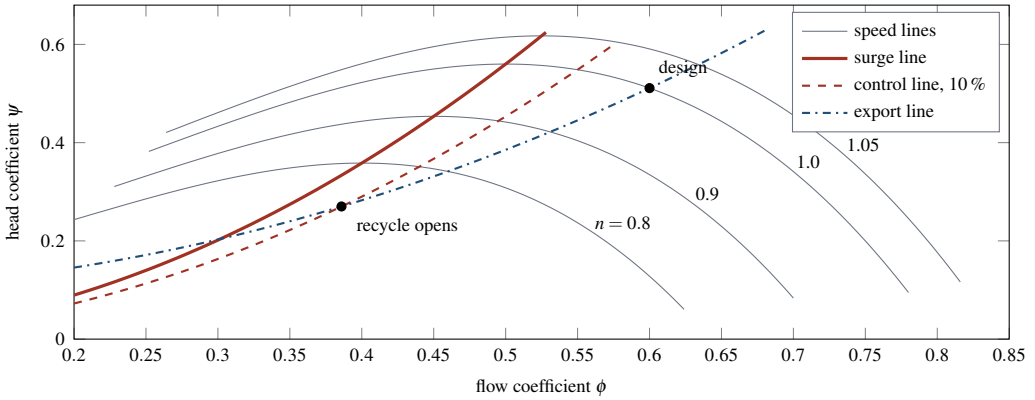

Figure 1.5 W-1's Export Compressor. Thin Speed Lines Scale by the Fan Laws. The Thick Surge Line Passes Through Their Peaks, the Dashed Control Line Lies 10 % to Its Right and the Dash-Dotted Curve Is the Head the Export Line Asks For.

on. The liquid line gives no time, so pumps in series must be held by fast limit loops. Chapter 12 plans its pack before the evening peak, and Chapter 13 joins suction and discharge limits at low selectors because one flow passes every station. Chapter 11 keeps a line pack behind its station, and the tank vapor space of Chapter 15 is an inventory of the same kind.

1.4.6 Compressor Maps and Surge

A centrifugal compressor's head falls as its flow rises, and the map draws that characteristic at each speed. The fan laws scale it, as shown in Equation (1.10).

$$\phi \propto n, \qquad \psi \propto n^2, \qquad P \propto n^3, \qquad \psi(\phi, n) = n^2\,\Psi(\phi/n) \tag{1.10}$$

where Ψ is the characteristic at design speed and P the power. W-1's export compressor uses the cubic characteristic of Chapter 11 [28], which peaks at $\phi/n = 0.5$. Left of the peak the head rises with flow, the compression system loses its steady operating point, and the flow cycles, reversing through the machine [29]. That is surge. Because the peak scales with the fan laws, the surge line is the parabola $\psi = 2.24\,\phi^2$ through the peaks, and the surge margin at the present speed is Equation (1.11).

$$SM = 100\,\frac{\phi - \phi_s(n)}{\phi_s(n)}, \qquad \phi_s(n) = 0.5\,n \tag{1.11}$$

The anti-surge controller holds a control line at 10 % margin, $\phi = 0.55\,n$, by opening a recycle valve that returns discharge gas to suction.

What the compressor must deliver is set by the system it feeds. W-1's export line asks for a head $\psi = 0.1 + 1.142\,\phi^2$: a static part, the pressure ratio the line holds even at low flow, and a part that grows with friction. The design point is $\phi = 0.6$ at full speed, with $\psi = 0.511$ and a margin of 20 %. Figure 1.5 draws the map. As the export flow falls, the speed falls along the system curve and the margin shrinks, 18.3 % at 90 % of design flow and 12.6 % at 70 %. The system curve meets the control line at $\phi = 0.386$ and $n = 0.70$, 64.3 % of design flow. Below it the machine stays on the control line and the recycle carries the difference. At 50 % of design flow the machine runs at $n = 0.608$ and recycles 11.5 % of the forward flow, power spent for nothing. The static head is what pushes the machine toward surge. A system of pure friction would follow a fan-law parabola parallel to the surge line and keep its margin at any speed.

Chapter 11 runs three such machines on one header, where every machine sees the same head and the sharing law decides which nears its surge line first. Chapter 12 places a compressor map at the

Table 1.1 Water Leaving W-1's Two-Stage Contactor (mg/Sm^3) From 1000 mg/Sm^3 at the Inlet, by Absorption Factor and Lean Floor

Absorption factor $\mathscr{A}$	1	1.5	2	3	4	8	∞
Share of the excess removed	0.667	0.789	0.857	0.923	0.952	0.986	1
Floor $y^* = 50$	366.7	250.0	185.7	123.1	95.2	63.0	50.0
Floor $y^* = 25$	350.0	230.3	164.3	100.0	71.4	38.4	25.0

inlet of its line, Chapter 15 stages two boil-off compressors that recycle what they do not send, and the liquefaction train of Chapter 14 is limited by the power of its refrigerant compressor's driver.

1.4.7 Absorption and Cryogenic Gas Processing

Gas leaving the separator carries water and, in a sour field, hydrogen sulfide and carbon dioxide. Both are removed by absorption in a solvent that is regenerated and recycled: glycol for water and an amine for the acid gases [30]. In a counter-current absorber of N ideal stages, the gas leaving the top cannot be cleaner than the gas in equilibrium with the lean solvent entering it, $y^* = mx_{lean}$, where m is the equilibrium ratio. The share of the removable excess that the column removes depends on the absorption factor, the solvent rate relative to the gas rate and the equilibrium, as shown in Equation (1.12).

$$\frac{y_{in} - y_{out}}{y_{in} - y^*} = \frac{\mathscr{A}^{N+1} - \mathscr{A}}{\mathscr{A}^{N+1} - 1}, \qquad \mathscr{A} = \frac{L}{mG} \tag{1.12}$$

which is the Kremser relation, with L and G the molar flows of solvent and gas.

W-1's gas enters a two-stage contactor with 1000 mg/Sm^3 of water, against an example specification of 110 mg/Sm^3. Table 1.1 lists the outlet for two lean purities. With a floor of $y^* = 50\,\text{mg/Sm}^3$, doubling the absorption factor from 2 to 4 lowers the outlet from 185.7 to 95.2 mg/Sm^3, and no circulation reaches below 50. The specification needs $\mathscr{A} = 3.38$. Purer solvent lowers the floor to 25 and the specification then needs $\mathscr{A} = 2.77$, 18 % less circulation. Circulation is the fast handle with a floor, and purity is the slow handle that moves the floor. Purity is made in the regenerator and reaches the contactor only as the solvent inventory turns over, four hours for a 6 m^3 surge drum at 1.5 m^3/h.

Chapter 7 finds the floor on its contactor, with stage efficiencies in place of ideal stages, and a ceiling on its reboiler that caps purity. Chapter 8 has two solutes on one solvent. Hydrogen sulfide approaches equilibrium quickly while carbon dioxide is absorbed at a rate set by contact time, so the solvent rate also sets the selectivity.

Ethane, propane and heavier hydrocarbons are recovered by cooling the gas until they condense. The cold comes from expansion. An expander that does work on a shaft follows, at best, the isentropic path, which for an ideal gas with heat capacity ratio γ ends at the temperature of Equation (1.13).

$$T_{2s} = T_1 \left(\frac{p_2}{p_1} \right)^{(\gamma-1)/\gamma}, \qquad T_2 = T_1 - \eta \left(T_1 - T_{2s} \right) \tag{1.13}$$

with η the isentropic efficiency. W-1's gas at −30 °C and 60 bar expanded to 25 bar with $\gamma = 1.3$ reaches −74.5 °C ideally and −67.8 °C at $\eta = 0.85$. A valve that throttles the same gas does no work and keeps its enthalpy. An ideal gas would not cool at all, and a real gas cools only by its Joule-Thomson effect. On the real-gas model of Chapter 9, expanding the same separator vapor to 23 bar, the expander reaches −70.9 °C where the valve reaches −49.2 °C. That chapter's plant trades recovery against recompression power through the column pressure and switches between recovering and rejecting ethane as prices move, and Chapter 10 splits the recovered liquid into products. Chapter 14 cools the whole stream to liquid, where the closest approach between the warming refrigerant and the cooling gas inside the exchanger decides the power the plant burns.

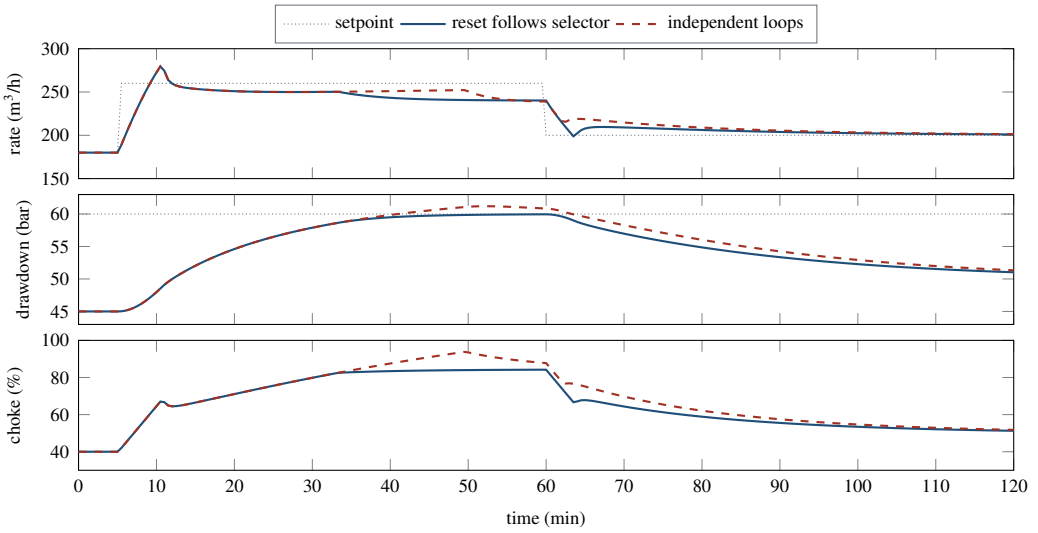

Figure 1.6 A Rate Loop and a Drawdown Override on W-1's Choke. With Both Reset Terms Following the Selected Output (Solid) the Override Holds the 60 bar Limit (Dotted). With Independent Loops (Dashed) the Waiting Override Has Wound Up and the Limit Is Broken.

1.4.8 Override and Constraint Control as Practiced

One actuator often has several jobs. W-1's choke holds the rate, but the rate must not draw the well down by more than 60 bar, a limit set by sand or water coning, so no more than 240 m^3/h may be produced. A rate loop (FIC) and a drawdown limit loop (PDIC) both compute a choke opening, and a low selector applies the smaller, as shown in Equation (1.14).

$$x = \min(u_F, u_D), \qquad u_j = K_{c,j}\, e_j + F_j, \qquad F_{j,k} = F_{j,k-1} + \left(1 - \mathrm{e}^{-\Delta t/\tau_{I,j}}\right)\left(x_k - F_{j,k-1}\right) \tag{1.14}$$

written in the external-reset form that Chapter 13 uses. Each loop's reset term F_j follows the opening actually applied, not the loop's own output. The selected loop is then exactly a PI controller. The other rests above the applied opening by its own proportional term, so it takes over the moment its error says it should [31]. The framework's back-calculation of Section 1.6.2 does the same when the applied output is the selected one.

The example's loops run every 0.1 min: FIC with $K_c = 0.5\,\%$ per m^3/h and $\tau_I = 6$ min on a meter with a 2 min lag, PDIC with $K_c = 3\,\%/\text{bar}$ and $\tau_I = 5$ min, and a choke that travels 5 %/min. The rate setpoint is raised from 180 to 260 m^3/h at 5 min, which the limit forbids, and lowered to 200 m^3/h at 60 min. Figure 1.6 compares two versions. In the first, both reset terms follow the applied opening. The rate overshoots to 282 m^3/h as the choke opens ahead of the draining wellbore, the override holds the choke from 33.3 min, and the drawdown peaks at 59.96 bar. In the second, each loop protects only itself from its own output limit, the arrangement of two independent controllers wired to a selector. The override's integral sits near 100 % while it waits, so it takes the choke only at 49.6 min. The drawdown passes 60 bar at 41.0 min, peaks at 61.2 bar and stays above the limit for 21.8 min. The same two loops with the same tunings protect the well in one wiring and not in the other.

Overrides recur in nearly every chapter, in three shapes. A low selector lets the most restrictive limit close a shared actuator: four loops on the choke in Chapter 4, rate and pressure on the choke in Chapter 5, suction and discharge limits on each pump in Chapter 13, the driver's power on the feed in Chapter 14, the recondenser on the compressor load in Chapter 15 and a high-limit override on the reboiler in Chapter 7. A high selector lets the most demanding protection open an actuator: the methanol dose in Chapter 5, the loading floor on the amine circulation in Chapter 8 and every

anti-surge valve in Chapter 11. A supervisor moves setpoints toward the active limit and backs away from it: the hill climber that backs off the most violated of six constraints in Chapter 2, the grid search above a protection layer that can only cut the pump in Chapter 3, and the rule-based supervisor of Chapter 6. Many chapters also offer a predictive controller, described in Section 1.6.3, as an option measured against the selectors.

1.4.9 Where Measurement Is the Limit

An override can protect only a limit it can see. In this industry the quantity that binds is often not the quantity a transmitter reads, and three cases recur. A product quality is measured late by an analyzer, so a loop on it acts on the past. A constraint sits where no instrument can be placed, so it is inferred from what can be measured, and the loop holds the inference exactly while the true value drifts [32]. An inventory is computed from a few transmitters, so its rate of change, which a leak detector needs, is uncertain during exactly the transients when it matters [33].

The example has one of each. The rate meter reads late, which is why the rate overshoots in Figure 1.6. The drawdown is measured only because the example puts a gauge at the bottom of the well. Without it the drawdown would be computed from the wellhead pressure and a column weight that changes with the holdup, and an error of a few bar in that weight would move the override's point of action by as much. The line pack of the export line is inferred from a few pressures. A leak detector compares inflow, outflow and the change in pack over a window, as shown in Equation (1.15).

$$I = \frac{1}{T_w}\int_{t-T_w}^{t} (w_{in} - w_{out})\,\mathrm{d}t' - \frac{\hat{M}(t) - \hat{M}(t-T_w)}{T_w} \tag{1.15}$$

An error of only 0.1 bar in the change of the estimated mean pressure over a 30 min window is 0.285 kg/s of apparent imbalance, 0.57 % of the line's flow, before any meter error is counted. A smaller leak cannot be told from the estimate's own error.

Table 1.2 lists, for each chapter, the limit that binds, how the controller knows it and the structures the chapter uses. Chapter 16 adds the measured results and asks which techniques recur from well to terminal and where measurement is the limit.

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 flow is metered in thirteen chapters and a pressure read at the surface in ten, and eleven chapters act on a quantity that is inferred or estimated rather than measured. A valve or a choke is moved in eleven chapters, a pump in seven and a compressor or its driver in five.

Each device sets a limit that no tuning removes. Most pressures are read by a diaphragm that bends under the process pressure, with strain gauges bonded to it whose resistance changes as it deforms [34], as drawn in Figure 1.8. Two chapters read the pressure at the bottom of a well. Without such a gauge it is computed from the surface, as Section 1.4.9 explains for W-1. Three chapters read a multiphase meter, which gives the oil, water and gas of a well stream without separating it. A test separator is reliable and accurate but too slow for monitoring in real time, and the water and gas that a meter must read grow as a reservoir matures [35].

The analyzers bring the longest delays. A chromatograph separates the components of a sample in time, as shown in Figure 1.9, so cycles of minutes are not uncommon and some dead time is inevitable [34]. The moisture analyzer of Chapter 7 reports 10 min after its sample left the gas, and the product chromatograph of Chapter 9 samples every 10 min and reports 6 min later.

The final element is most often a choke or a control valve, and on lines and injection wells a variable-speed drive, which sets the speed of a motor by changing the frequency of the power that feeds it [34]. A drive turns the compressor of Chapter 12 and the four pumps of Chapter 13.

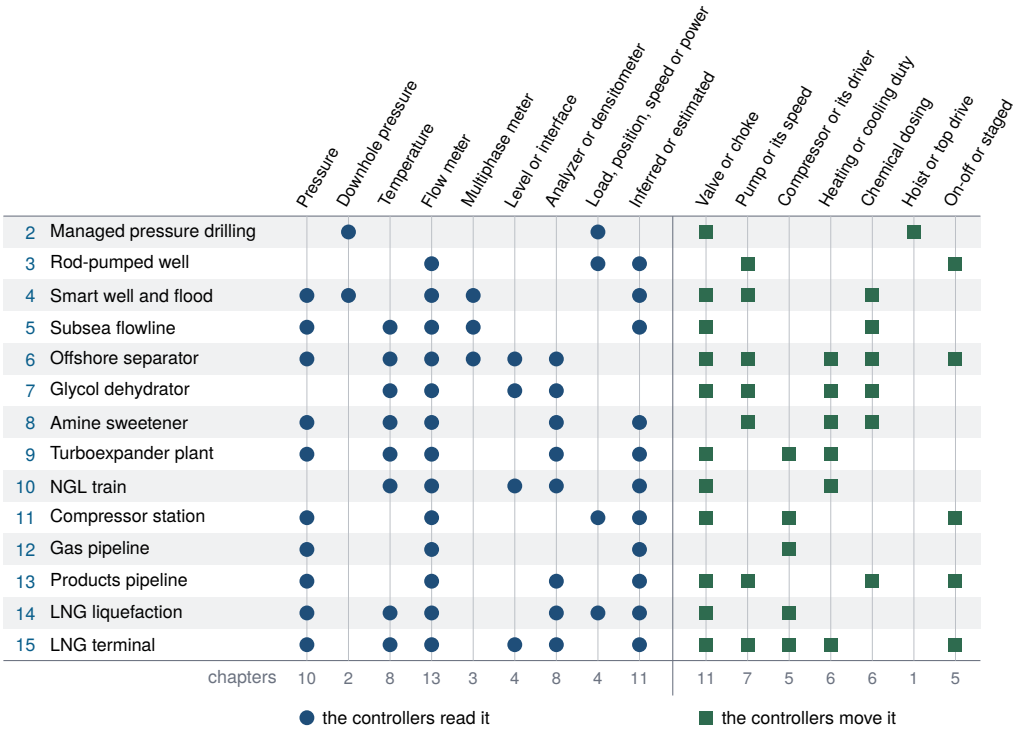

Figure 1.7 Sensors and final elements by chapter. A round mark means that the chapter's controllers or operator read the device, a square mark that they move it. A mark is set only where the chapter's variables table, its instrument text or its section on technology in practice names the device or its reading, so a blank says that the model leaves the device out and not that a real plant lacks it. Notice that eleven of the fourteen chapters act on an inferred or estimated quantity.

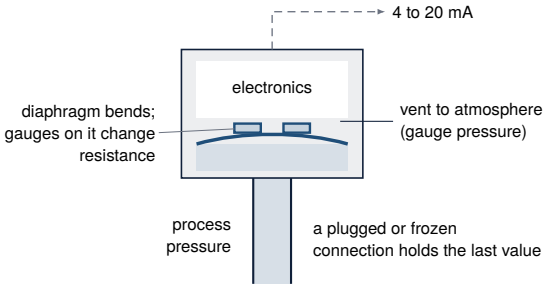

Figure 1.8 Principle of a gauge pressure transmitter. Notice that the reading comes from the bending of one diaphragm, so a metal element that has been overpressured can lose accuracy without any change in the process [34].

Table 1.2 For Each Chapter, the Limit That Binds, How the Controller Knows It and the Structures Used

Ch.	Binding limit	How it is known	Structures
2	Vibration, cuttings, torque, pore and fracture pressure	Six constraint ratios, two of them vibration indicators over 20 s windows	PI loops, supervisor backing off the worst limit, choke MPC
3	Pump fillage, fluid level, rod load	Inferred from the surface load cell through the downhole card	Timer, PI on the inferred level, grid search over a protection layer
4	Drawdown, sand, water, the flood fronts	Drawdown from a model, a delayed multiphase meter, fronts never measured	Four PI loops at a low selector, injection override, planner
5	Hydrate margin at the weakest point	Profile estimator from a few pressures and temperatures	Rate loop with pressure override, methanol feedforward with trim, sequence
6	Levels between trips, oil and water quality	Level transmitters, analyzers behind sample delays	Averaging level, inlet feedforward, rule-based supervisor, MPC
7	Water in the gas, reboiler ceiling	Moisture analyzer reporting late	Ratio feedforward, analyzer trim, high-limit override, economic MPC
8	Hydrogen sulfide in the gas, rich loading	Sampled analyzer, loading inferred from an acid-gas meter	Ratio, loading floor at a high select, analyzer trim, MPC
9	Methane in the liquid product, flooding, power	Chromatograph after a delay, tray temperature as a proxy	Two PI loops, analyzer trim, mode sequence, supervisor
10	Ethane in propane, shared steam	Chromatograph one column downstream, inferential estimate	Local, analyzer and cascade loops, averaging drum level, MPC
11	Surge margin of each machine	Computed from measured flow and head on the commissioned map	Pressure loop, sharing law, anti-surge overrides, sequencer
12	Delivery pressure floor, line pack	Pack inferred from three transmitters, a demand forecast	Cascade, forecast feedforward, pack window, leak detector, MPC
13	Suction and discharge limits, a full hill, batch cuts	Station pressures, densitometer at the terminal, volume balance	Low selectors in external-reset form, back pressure loop, MPC
14	Temperature approach in the exchanger, driver power	Computed from flows, pressures and a sampled refrigerant analysis	Cascade, inferential approach loop, low selector on feed, MPC
15	Tank pressure below the flare, recondenser capacity	Tank pressure, capacity from subcooling, boil-off observer	PI with recondenser override, staging, nomination preview, MPC

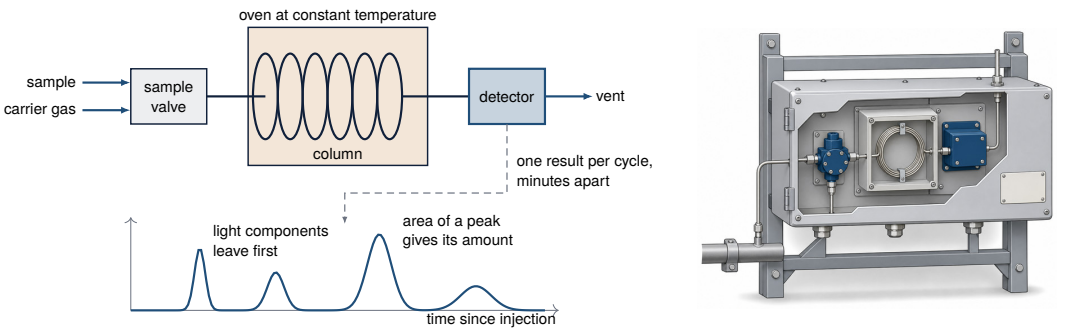

Figure 1.9 Principle of a process gas chromatograph, with a view of the device as installed. Notice that one sample gives one result at the end of the cycle, so a loop on this instrument acts on a value that is always old.

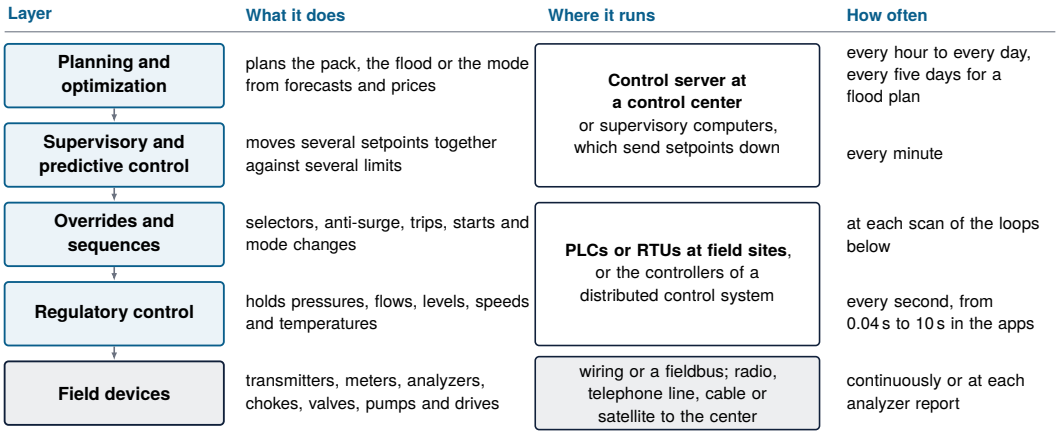

Figure 1.10 The control layers of this volume, the platform that hosts each and how often each acts. The periods of the regulatory, predictive and optimization layers are those of [36, 37], the platforms those of [11], and the ranges in the apps those of the chapters. Notice that each layer sends setpoints 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.

The platform sets the last limit. A pipeline is run through a supervisory system. A control server at a control center reaches the field sites by radio, telephone line, cable or satellite, and remote terminal units or programmable controllers at those sites control the local process. Cycle times run from milliseconds to minutes [11]. Figure 1.10 places the layers on that hardware. Regulatory loops act about every second, predictive control every minute and optimization every hour to every day [36, 37]. The apps run their loops every 0.04 s in Chapter 3, every 1 s in Chapter 13, whose predictive controller acts every 2 min, and every 10 s in Chapter 12.

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.

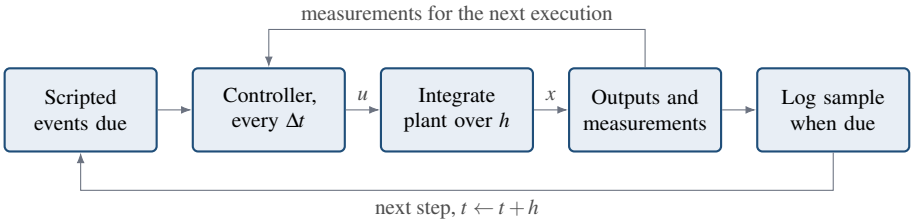

Figure 1.11 Work of the Engine in One Integration Step

1.6.1 The Simulation Engine

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

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

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

Noise comes from a deterministic pseudorandom generator. Its seed is declared by the app, or by a scenario that needs its own noise record, and it is restored whenever a scenario is loaded or Reset is pressed. The framework gives the reader no seed setting, so *default seed* in this book means the seed the app ships with. A run repeats bit for bit when the same scenario is loaded and the same settings are made before Run. A change made by hand during a run lands on a different step each time, which is one reason numeric answers carry a tolerance. The standalone simulators draw noise from the browser's own generator, so a noisy run there repeats in its statistics and not sample for sample.

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

1.6.2 The Velocity-Form PID Algorithm

One PID implementation serves the PID loops of every framework app. It is the velocity (incremental) form: each execution adds the change in the proportional and derivative terms to a running sum, and the integral is a separate term, as shown in Equations (1.16) to (1.19) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.16}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.17}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.18}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.19}$$

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 (1.19) drives I toward the value that reproduces the output applied, which is back-calculation [38]. With no integral action, Equation (1.19) 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 (1.20).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.20}$$

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 (1.21) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q\,(y_k - r_k) + \sum_{j} \Delta u_j^{\top} R\,\Delta u_j + \sum_{i} \rho_i\, s_i^2 \tag{1.21}$$

where r is the setpoint, Q weights the tracking error, R is the move suppression on each input change Δu_j and ρ_i is the price of a slack $s_i \geq 0$ that lets an output band be broken when it cannot be held. Actuator limits are hard constraints. The quadratic program is solved in the browser by a dual active-set method [39]. The first move is applied, and the plan is made again at the next decision.

The plan starts from an estimate of the model state. About half of these apps use a steady-state Kalman filter on the model with an added disturbance state, which absorbs what the model did not predict and removes steady offset without an integral term [40]. Others correct the model output with a filtered bias to the same end. About a dozen apps also offer moving horizon estimation (MHE), which solves the estimation problem over a window of recent samples and can hold bounds that a Kalman filter cannot, such as a concentration that is never negative [41].

With the MPC selected, the tuning fields of the loops it replaces leave the Control panel. In most of these apps the Advanced level shows up to three settings, each a multiple of the shipped value: the prediction horizon, the move suppression and the band price. The trend draws the accepted plan

ahead of the present time, and a list of constraint prices shows which limits bind. A plan that is not optimal, not finite or outside its limits is rejected, and a fallback, in most apps the shipped controller, holds the actuator until a plan succeeds. The Model Fitting and Tuning Studio builds the same kind of controller from the reader's own data in its tab, MPC from the fit.

1.6.4 Scenarios, Progress and Data

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

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

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

Show Model opens the equations, live values of the variables, the editable parameters and the assumptions. The trends plot SP and PV above the MV for each loop, with actuator limits and marks at setpoint moves, disturbances, mode changes and scripted events. The scorecard reports for each loop the integral of the absolute error (IAE), overshoot, settling time, MV travel and share of time at a limit, with the app's own measures, and it does not combine them into one score. IAE and travel accumulate from Reset, so two runs are compared at the same clock reading.

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

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

1.7 Working With AI Agents While Studying

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

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

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

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

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

How the exercises train the habit. One exercise in each chapter is marked *Review the agent*. It presents a tuning, a model or a claim as an AI assistant might propose it, with a flaw, and asks the reader to find the flaw in about five minutes. The flaws are ordinary ones: a unit, a sign, a steady state that does not close, a result measured on the wrong scenario or at a different run length, or a limit that was ignored. Each is found with one of the tests above. The purpose is not distrust. An engineer who can say why an answer is right can hand an assistant much larger tasks and still stand behind the result. That larger picture, in which agents do most of the routine work and the engineer supervises and checks, is the subject of Chapter 16.

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 6, 7, 12 and 15 form a first course: an inventory, an absorber, a line and a tank, each with one main loop and one limit. They rest on Sections 1.4.4 to 1.4.7. Chapters 4, 11 and 13 add the selectors of Section 1.4.8 and are best read together. Chapters 3, 8 and 10 turn on what is inferred or measured late, Section 1.4.9. Chapters 2, 5, 9 and 14 combine several of these and are best left until last. A reader who follows the fluid can read the volume in order, from the well to the terminal.

KEY POINTS

- From well to market a choke, a pump or a compressor sets a pressure, and the flow follows from a reservoir, a line or a market.
- A separator stores a slug under any tuning. The tuning decides only how fast the liquid is passed on.
- A gas line rides through a change in demand on its own pack. A liquid line stores almost nothing and needs fast limit loops.
- Two loops at a selector protect a limit only when the waiting loop's reset follows the output actually applied.
- An override can protect only a limit it can see, and the quantity that binds is often inferred or measured late.

Where else this applies. The same layers and limits return in every industry volume of the series. Selectors and overrides are developed in Volume 4, predictive control in Volume 9, and the refinery that takes the oil in Volume 11.

Table 1.4 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Managed pressure drilling rig	The best drilling rate lies on a constraint boundary that moves as the bit wears.	advanced	3 h
3	Rod-pumped oil well	A loop on an inferred level holds the inference, not the level.	intermediate	2.5 h
4	Smart producer and waterflood	Four limits share a choke, and the most oil now is not the most oil recovered.	intermediate	2 h
5	Deepwater flowline and riser	A minimum that moves along the line, dosed through a transport delay.	advanced	2.5 h
6	Offshore three-phase separator	No tuning keeps a slug out. Tuning decides how fast it leaves.	introductory	2.5 h
7	Glycol dehydrator	A fast handle with a floor, a slow handle with a ceiling.	introductory	2 h
8	Amine sweetener	A loading limit the product analyzer cannot see.	introductory	2 h
9	Turboexpander NGL plant	The order of two setpoint moves decides the specification.	advanced	2 h
10	NGL fractionation train	A specification made one column upstream of where it is measured.	intermediate	2.5 h
11	Compressor station	Equal surge margin shares the load where equal speed does not.	intermediate	2 h
12	Gas transmission line	The line is a tank, and a forecast is worth more than retuning.	introductory	2.5 h
13	Refined products line	One flow through four pumps, held by low selectors.	intermediate	2 h
14	LNG liquefaction train	Pairing on a valve with two effects, and a limit computed from a late analysis.	advanced	3 h
15	LNG import terminal	An integrating pressure whose outlet the market sets.	introductory	2 h

References

[1] B. Foss, "Process control in conventional oil and gas fields—challenges and opportunities," *Control Engineering Practice*, vol. 20, no. 10, pp. 1058–1064, 2012.

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

[3] E. D. Sloan, "Hydrocarbon hydrate flow assurance history as a guide to a conceptual model," *Molecules*, vol. 26, no. 15, p. 4476, 2021.

[4] W. A. Loomis, "Dummy driller," U.S. Patent 1 779 480, Oct. 28, 1930, filed December 14, 1927.

[5] R. W. Bucklin, "Method and equipment for treating hydrocarbon gases for pressure reduction and condensate recovery," U.S. Patent 3,292,380, filed April 28, 1964, granted December 20, 1966, 1966. [Online]. Available: https://patents.google.com/patent/US3292380A/en

[6] S. G. Gibbs and A. B. Neely, "Computer diagnosis of down-hole conditions in sucker rod pumping wells," *Journal of Petroleum Technology*, vol. 18, no. 1, pp. 91–98, 1966.

[7] P. Arjomandnia, "Simulation of LNG rollover in storage tanks," Ph.D. dissertation, Curtin University, Perth, Australia, 2015, open access, hdl.handle.net/20.500.11937/1544.

[8] S. G. Gibbs, "Method for determining the pump-off of a well," U.S. Patent 3 951 209, Apr. 20, 1976, filed June 9, 1975.

[9] E. D. Burger, W. R. Munk, and H. A. Wahl, "Flow increase in the trans alaska pipeline through use of a polymeric drag-reducing additive," *Journal of Petroleum Technology*, vol. 34, no. 2, pp. 377–386, 1982.

[10] Z. Schmidt, J. P. Brill, and H. D. Beggs, "Experimental study of severe slugging in a two-phase-flow pipeline - riser pipe system," *Society of Petroleum Engineers Journal*, vol. 20, no. 5, pp. 407–414, 1980.

[11] 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.

[12] J.-M. Godhavn, "Control requirements for automatic managed pressure drilling system," *SPE Drilling & Completion*, vol. 25, no. 3, pp. 336–345, 2010.

[13] U.S. Energy Information Administration, “Natural gas pipelines,” Energy Explained, last updated March 19, 2024, 2024, accessed November 18, 2026. [Online]. Available: https://www.eia.gov/energyexplained/natural-gas/natural-gas-pipelines.php

[14] J. G. Abraham, “Drilling control,” U.S. Patent 2 759 702, Aug. 21, 1956, filed March 25, 1952.

[15] J. V. Vogel, “Inflow performance relationships for solution-gas drive wells,” *Journal of Petroleum Technology*, vol. 20, no. 1, pp. 83–92, 1968.

[16] S. Fakher, A. Khlaifat, M. E. Hossain, and H. Nameer, “A comprehensive review of sucker rod pumps’ components, diagnostics, mathematical models, and common failures and mitigations,” *Journal of Petroleum Exploration and Production Technology*, vol. 11, no. 10, pp. 3815–3839, 2021.

[17] R. P. Cutler and A. J. Mansure, “Fluid dynamics in sucker rod pumps,” Sandia National Laboratories, Albuquerque, NM, Report SAND99-0093C (Southwest Petroleum Short Course, Lubbock, TX), OSTI 3261, 1999.

[18] D. H. Beggs and J. P. Brill, “A study of two-phase flow in inclined pipes,” *Journal of Petroleum Technology*, vol. 25, no. 5, pp. 607–617, 1973.

[19] Y. Taitel, “Stability of severe slugging,” *International Journal of Multiphase Flow*, vol. 12, no. 2, pp. 203–217, 1986.

[20] B. F. M. Pots, I. G. Bromilow, and M. J. W. F. Konijn, “Severe slug flow in offshore flowline/riser systems,” *SPE Production Engineering*, vol. 2, no. 4, pp. 319–324, 1987.

[21] J.-M. Godhavn, M. P. Fard, and P. H. Fuchs, “New slug control strategies, tuning rules and experimental results,” *Journal of Process Control*, vol. 15, no. 5, pp. 547–557, 2005.

[22] E. D. Sloan, “Fundamental principles and applications of natural gas hydrates,” *Nature*, vol. 426, no. 6964, pp. 353–359, 2003.

[23] H. Gul Zaman, L. Baloo, R. Pendyala, P. Singa, S. Ilyas, and S. Kutty, “Produced water treatment with conventional adsorbents and MOF as an alternative: A review,” *Materials*, vol. 14, no. 24, p. 7607, 2021.

[24] D. Asante-Sackey, S. Rathilal, E. K. Tetteh, and E. K. Armah, “Membrane bioreactors for produced water treatment: A mini-review,” *Membranes*, vol. 12, no. 3, p. 275, 2022.

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

[26] A. R. D. Thorley and C. H. Tiley, “Unsteady and transient flow of compressible fluids in pipelines—a review of theoretical and some experimental studies,” *International Journal of Heat and Fluid Flow*, vol. 8, no. 1, pp. 3–15, 1987.

[27] M. S. Ghidaoui, M. Zhao, D. A. McInnis, and D. H. Axworthy, “A review of water hammer theory and practice,” *Applied Mechanics Reviews*, vol. 58, no. 1, pp. 49–76, 2005.

[28] F. K. Moore and E. M. Greitzer, “A theory of post-stall transients in axial compression systems: Part I—development of equations,” *Journal of Engineering for Gas Turbines and Power*, vol. 108, no. 1, pp. 68–76, 1986.

[29] E. M. Greitzer, “Surge and rotating stall in axial flow compressors—part I: Theoretical compression system model,” *Journal of Engineering for Power*, vol. 98, no. 2, pp. 190–198, 1976.

[30] A. J. Kidnay, W. R. Parrish, and D. G. McCartney, *Fundamentals of Natural Gas Processing*, 3rd ed. CRC Press, 2019.

[31] A. H. Glattfelder and W. Schaufelberger, *Control Systems with Input and Output Constraints*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2003.

[32] 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.

[33] L. Billmann and R. Isermann, “Leak detection methods for pipelines,” *Automatica*, vol. 23, no. 3, pp. 381–385, 1987.

[34] 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.

[35] L. S. Hansen, S. Pedersen, and P. Durdevic, “Multi-phase flow metering in offshore oil and gas transportation pipelines: Trends and perspectives,” *Sensors*, vol. 19, no. 9, p. 2184, 2019.

[36] 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.

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

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

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

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

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

1.9 Exercises

1.1 Theory and calculation · 5 min

W-1's reservoir has fallen to 240 bar. All else is as in Section 1.4.2, with the choke held at 40 %.

(a) Find the rate. *Checkpoint:* 160.5 m^3/h.

(b) Find the drawdown and the pressure drop across the choke. *Checkpoint:* 40.1 bar and 19.9 bar.

1.2 Concept check · 2 min

W-1's rate loop holds 180 m^3/h on the choke while the reservoir pressure falls from 250 bar. Which statement is right?

(a) the choke stays at 40 %, because the rate it passes depends on the choke alone

(b) the choke opens further as the reservoir falls, and at 229 bar it is wide open and the rate starts to fall

(c) the rate can be held at any reservoir pressure above 180 bar, the floor under the bottomhole pressure

(d) the rate is held until the drawdown reaches its 60 bar limit, which the falling reservoir reaches first

1.3 Theory and calculation · 4 min

W-1's choke is opened to 70 %, where the well produces 227.7 m^3/h with 13.07 bar across the choke.

(a) Find the share g of an immediate choke response that remains at steady state, from Equation (1.3). *Checkpoint:* 0.315.

(b) Find the immediate and the steady gain of rate to choke opening. *Checkpoint:* 3.25 and 1.02 m^3/h per percent.

1.4 Theory and calculation · 4 min

W-1's sealed riser fills at 0.125 bar/s. Use Equation (1.5) with a standard pressure of 1.013 bar.

(a) The gas rate is doubled to 36 000 Sm^3/h by gas lift with the flowline gas volume still 150 m^3. Find Π. *Checkpoint:* 0.54.

(b) The gas volume upstream of the riser is also cut to 100 m^3. Find Π and say whether the riser still slugs. *Checkpoint:* 0.81, it still slugs.

1.5 Theory and calculation · 5 min

W-1's separator has a liquid surface of 25 m^2 and an oil valve that passes 300 m^3/h wide open, so $k = 0.12\,\mathrm{m/h}$ per percent.

(a) A PI level loop has $K_c = 100\,\%/\mathrm{m}$ and $\tau_I = 20\,\mathrm{min}$. Find τ_c and ζ. *Checkpoint:* 5 min and 1.0.

(b) A third well adds a step of 90 m^3/h to the inflow and the level may use 0.6 m of band. Find the slowest outflow ramp that keeps the level in the band, and how long the ramp lasts. *Checkpoint:* 4.5 m^3/h per minute for 20 min.

1.6 Concept check · 2 min

Which statement matches the two runs of Figure 1.4?

(a) the tight loop keeps most of the slug out of the vessel, because it opens the valve at once

(b) the averaging loop stores about 18 m^3 more than the tight loop, because it opens the valve slowly

(c) the averaging loop takes the level over the high-high trip, which is the price of a smooth outflow

(d) both store nearly the same volume, and the tuning decides only how fast the stored liquid leaves

1.7 Theory and calculation · 4 min

A gas transmission line is 80 km long with a bore of 0.5 m and an isothermal wave speed of 350 m/s.

(a) Find the gas it stores per bar of mean pressure. *Checkpoint:* 12.82 t per bar.

(b) Demand exceeds supply by 8 kg/s for two hours. Find the fall in mean pressure. *Checkpoint:* 4.49 bar.

1.8 Theory and calculation · 4 min

A two-stage glycol contactor sees 800 mg/Sm^3 of water at the inlet, and the lean glycol sets a floor of 40 mg/Sm^3.

(a) Find the outlet water at an absorption factor of 2, from Equation (1.12). *Checkpoint:* 148.6 mg/Sm^3.

(b) Find the outlet at an absorption factor of 4, and the least water any circulation can reach. *Checkpoint:* 76.2 mg/Sm^3, and 40 mg/Sm^3.

1.9 Concept check · 2 min

The export line of W-1 is replaced by one whose head is pure friction, $\psi = k\phi^2$ with no static part. What happens to the compressor's surge margin as the export flow falls?

(a) it stays the same at every speed, because the operating point follows a fan-law parabola

(b) it shrinks faster than before, because friction adds head at low flow

(c) it shrinks because the surge line moves to the right with the friction

(d) it grows as the speed falls, because the head falls with the cube of speed

1.10 Theory and calculation · 3 min

Natural gas at −20 °C and 70 bar is expanded to 30 bar. Treat it as an ideal gas with $\gamma = 1.3$.

(a) Find the outlet temperature of an isentropic expander. *Checkpoint:* −65.0 °C.

(b) Find the outlet temperature at an isentropic efficiency of 0.8. *Checkpoint:* −56.0 °C.

CHAPTER 2

A Managed Pressure Drilling Simulator to Teach Optimization on a Constraint Boundary

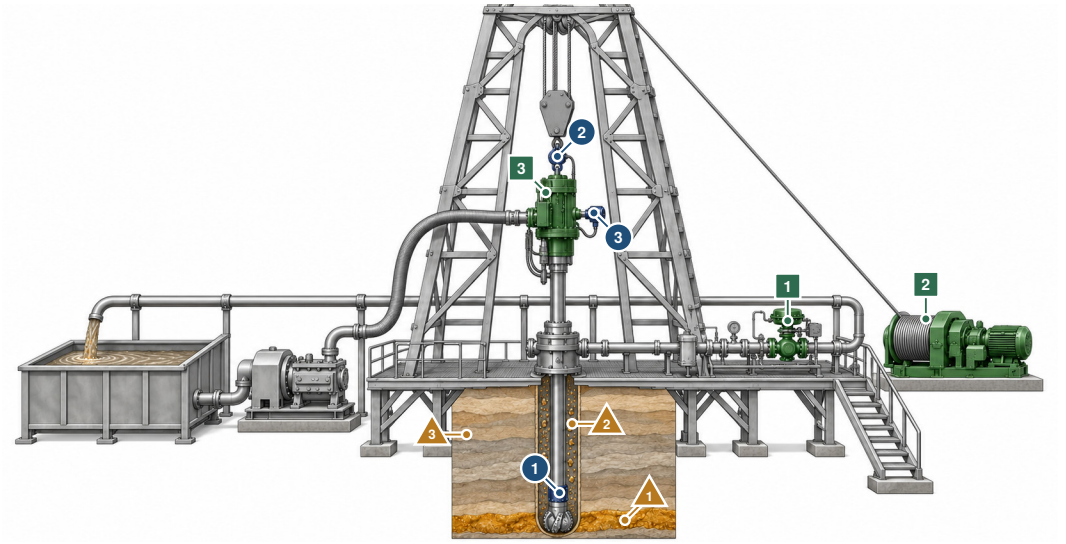

The controller moves 1 choke on the returning mud 2 drawworks that pays out the drilling line 3 top drive that turns the pipe

The controller reads 1 pressure at the bottom of the hole 2 hook load, which gives the weight on the bit 3 speed of the top drive

What disturbs it 1 a harder, more abrasive rock layer 2 rock cuttings carried up in the mud 3 pore and fracture pressure of the rock

The mud that lifts the cuttings up the annulus leaves the wellhead through the choke on the returns skid, so the cuttings and the choke both add to the pressure at the bit. Illustration, not to scale.

IN PLAIN TERMS **Drilling Fast Right at the Edge**

The main idea. A drill bit cuts rock faster when it is pushed harder, turned faster and pressed on less by the mud. A valve on the returning mud lets the crew lower that mud pressure.

The challenge. Every push toward speed brings the well closer to a limit: the pipe twists and jerks, rock chips pile up, or the mud pressure leaves the narrow range the rock allows. The best speed sits right on the edge of these limits.

Why it matters. In the chapter's simulator, a controller working along that edge drills 73 percent more hole than a cautious fixed setting.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? box in the Introduction, the two History boxes in Section 2.2, the Try it boxes and the second Did you know? box in Section 2.6, and Key points in Section 2.9. No equations.
Study	70 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	25 min	This page, Section 2.6 with Figure 2.3, Section 2.7 and Section 2.8 with its Open challenge boxes.

The control question Weight on bit, rotary speed and a lower bottomhole pressure all raise the drilling rate, and each also pushes the well toward a vibration, cuttings or pressure limit. Where does the best feasible operating point sit, what holds it there, and what does operating away from it cost?

Learning objectives

- Explain why the best drilling rate lies on the boundary of the feasible region.
- Derive the cuttings ceiling and use it to pick the handle that raises the achievable rate.
- Compute the choke gain and authority and the ramp offset of the pressure loop.
- Locate the stick-slip corner and the bit-bounce band from G and the axial resonance.
- Separate a conservative limit from a detuned loop in the graded challenge.

Prerequisites Sections 1.4.2 (pressure column and choke), 1.4.8 (constraint control), 1.6.2 (velocity-form PID) and 1.6.4 (the apps). Section 1.6.3 is needed only for the choke MPC.

The app https://apmonitor.com/apps/drilling-mpd/

See also Volume 9, Chapter 8 and Volume 9, Chapter 10 for supervisory setpoints, Volume 9, Chapter 4 for an offset-free MPC, Volume 4, Chapter 7 for rate limits on a utility drill and Volume 4, Chapter 2 for a protective override.

2.1 Introduction

The rate of penetration in rotary drilling is limited by several mechanisms at once. Weight on bit raises the depth of cut until the cutters can no longer clear their own debris [1]. Rotary speed raises the rate sub-linearly, and weight and speed together determine whether the drillstring turns smoothly or enters a torsional limit cycle [2, 3]. The mud column that keeps the formation in the hole also presses each freshly sheared cutting back onto the rock, so overbalance reduces the drilling rate [4]. Every cubic meter of rock removed must be lifted up the annulus, and the mixture of rock and mud is heavier and rougher than mud alone, so drilling faster raises the pressure at the bottom of the hole and slows the drilling down again [5].

The objective and every constraint on it increase together with the two variables the operator manipulates, so the best operating point is not an interior maximum. It is a point on the boundary of a feasible region. At the default settings of the app of this chapter the objective peaks at 111.2 ft/hr with 35.0 klbf of weight and 250 RPM, and that peak is infeasible, because it loads the annulus to 2.29 % cuttings by volume against a limit of 2.2 %. The best point that satisfies every limit produces 106.7 ft/hr and the operating point supplied by the first scenario produces 67.6 ft/hr. The constraint costs 4 % of drilling rate and the conservative operating point costs 37 %, a factor of nine between the two losses.

The feasible region also changes during operation. Bit wear accumulates over hours and lowers the whole objective surface, so a limit that binds at the start of a bit run no longer binds at the end of it, and cuttings loading has a time constant near 340 s, two orders of magnitude slower than the torsional mode that governs stick-slip. An operating point selected from a static evaluation becomes suboptimal within the hour.

> ✦ DID YOU KNOW? **A Deep Well Takes Months**
>
> A gas exploration well finished in the United States in 1974 reached 9583 m. Drilling it took 504 days and 103 bits [6]. In 2016 a well of 6000 m took about three to six months to drill, and one of 8000 m more than a year [6].

The model rests on standard conservation statements, a momentum balance on two torsional inertias, a mass balance on the cuttings inventory in the annulus, and hydrostatic and frictional contributions to bottomhole pressure [7, 8]. The experiments built on the model show integrating dynamics, negative process gain, time-scale separation, anti-windup and supervisory constraint

control, the principles a process dynamics and control course develops with simpler plants [7, 9]. The behavior matters industrially because non-productive time is the largest addressable cost in a well and constraint-aware operation is the mechanism that reduces it [10].

This chapter works with a browser simulator of drilling and managed pressure, set out completely enough to be rebuilt from the text, and locates where its constrained optimum sits and what holds it there. The app gives the cost of a conservative limit as a measured number rather than an assertion. Section 2.2 reviews how the loops the app implements developed. Section 2.3 gives the thirteen-state model, whose steady-state evaluator and dynamic integration are required to agree. Section 2.4 maps the feasible region and derives a closed-form ceiling on the objective imposed by the cuttings constraint, $ROP \leq C_{max} Q \eta(\omega) / A_{bit}$. The ceiling depends on rotary speed alone and is reached at 106.7 ft/hr for every choke position above 52.1 % open, so two of the three manipulated variables contribute nothing beyond that point, while a 3.3 % increase in pump rate adds 4.2 % of feasible drilling rate. Section 2.5 describes the regulatory loops, the constraint-aware search transcribed from the executing code and the optional choke MPC. Section 2.6 reports the scenarios, the graded challenge and the comparison of pressure controllers. Section 2.7 sets the hardware of a real rig beside the model, Section 2.8 states the limitations, and Section 2.9 draws the results together.

2.2 Background

The first drilling automation was mechanical and it controlled exactly one variable. That device, the automatic driller, held weight on bit by paying out drilling line, and the History box below traces it to its first patents. Its structure survives unchanged in the app of this chapter. Weight is not something the rig can set directly, because the string is an elastic member and the bit is destroying the rock the weight rests on, so the manipulated variable is the rate at which line is paid out and the plant between it and the controlled variable is an integrator.

> ⧖ HISTORY **1927 and 1952: Machines That Hold the Weight on the Bit**
>
> A patent application of 1927 described a dummy driller, a device to keep a set force on a rotary drill hung in part from a cable [11]. A reversible motor paid the cable out or picked it up, and too much pull on the cable made the motor pay out until the strain fell. The hydraulic apparatus that already showed the strain on such cables was to work the automatic control as well [11]. A 1952 application, granted in 1956, let the tension in the cable regulate how fast the drawworks drum unspooled [12]. Its hydraulic pressure also read on a gauge calibrated as a weight indicator for the driller. A fall in that pressure meant too much weight on the bit, so the brake slowed the drum, the line paid out more slowly and the bit drilled off [12].

Everything above that loop arrived later as separate developments. Rate-of-penetration modeling became quantitative when the drilling variables were fitted as a multiple regression on offset-well data, which produced the multiplicative form still used today [4], and a mechanistic account for roller-cone bits, in which the rate is limited alternately by rock strength and by the removal of cuttings from beneath the bit, gave the founder behavior a physical basis [1]. Specific energy, the work spent per unit volume of rock destroyed, was proposed as the efficiency measure for rock drilling in 1965 [13] and became the standard real-time surveillance quantity on the rig floor forty years later, where it detects the moment at which extra weight stops making hole and starts grinding and decides when to trip a dull bit [14]. Volume 4, Chapter 7 uses specific energy in the same way, to read the soil a utility drill is cutting.

Drillstring vibration was studied on its own track. Torsional oscillation with the bit stopping completely and then releasing is a self-excited limit cycle driven by a bit torque that falls with speed, and active damping of that mode by feeding measured top-drive speed back through a compliant torque command was demonstrated on full-scale strings [2]. Later work reduced the mechanism to a bit-rock interaction law with a delay set by the depth of cut per revolution [3], and top-drive controllers that absorb torsional waves rather than reflecting them are now a standard option [15].

Hole cleaning followed a third track, with cuttings-transport models built from flow-loop experiments and organized around the annular velocity needed to keep a bed from forming [5].

Managed pressure drilling is the newest of the four developments and the one that made annular pressure a controlled variable. The second History box below follows it from mud density to an automatic choke. The vocabulary and the systems view that holds these pieces together came from a deliberate effort to import automation practice from aviation and process manufacturing into drilling [16], and the case for doing so rests on non-productive time, which is the fraction of rig days spent not making hole and is the largest addressable cost in a well [10]. The application implements the state that arrangement has reached, which is regulatory loops on weight, speed and pressure with a slow supervisory layer above them, rather than a single optimizer replacing them.

> ⧖ HISTORY **2007: The Choke Takes Over Well Pressure**
>
> Pressure at the bottom of a well was long set by the density of the mud, a method that is slow and inaccurate and cannot answer the brief swings of a pipe connection [17]. Managed pressure drilling grew out of underbalanced drilling. An industry committee defines it as an adaptive drilling process that precisely controls the annular pressure profile throughout the wellbore [18]. A rotating control device seals the top of the annulus and sends the returning mud through a choke, whose opening sets the back-pressure [17]. Narrow margins, especially in depleted reservoirs, call for accurate control, and in 2007 automatic control of downhole pressure by a topside choke was applied offshore in the North Sea [19]. Later work put the rate of penetration and the pressure into one controller that moves pump flow, choke, rotation and weight on bit together [20].

2.3 Process Model

The application simulates one vertical hole section drilled by one bit through one formation. Mud is pumped down the drill pipe, passes through the bit, and returns up the annulus between the pipe and the hole wall, through a choke that adds back-pressure and on to the pits. The layout, the measurement points and the symbol for every located quantity are given in Figure 2.1. Everything in the model is SI. Newtons, pascals, radians per second and meters per second are what the equations use, and the field units quoted in this chapter (klbf, RPM, psi, pounds per gallon, feet per hour) are display conversions that change no computed value.

2.3.1 Variable Classification

Three controlled variables, three manipulated variables, the disturbances, the thirteen states and the derived quantities the constraints act on are collected in Table 2.1. The rate of penetration is the objective and is deliberately not a controlled variable, because its achievable value is a property of the whole feasible region rather than of any one loop.

2.3.2 The Objective

The rate of penetration is a product of five factors, one for each mechanism that can slow the bit down, in the reduced form of the multiple-regression law of [4] with the foundering behavior of [1], as shown in Equation 2.1.

$$ROP = K\, f_W(W)\, f_N(\omega)\, f_h(h)\, f_{\Delta P}(P_b - P_{pore})\, f_C(C) \tag{2.1}$$

where K is the drillability coefficient of the rock, divided by the formation strength multiplier $1 + \Delta_{hard}/100$, and each factor is dimensionless and equals unity at its own best condition. The product form matters more than any individual factor. A gain in one factor multiplies the whole objective, which is why the back-pressure choke changes the drilling rate although no rock property changes.

The weight factor rises from a threshold and then founders, as shown in Equation 2.2.

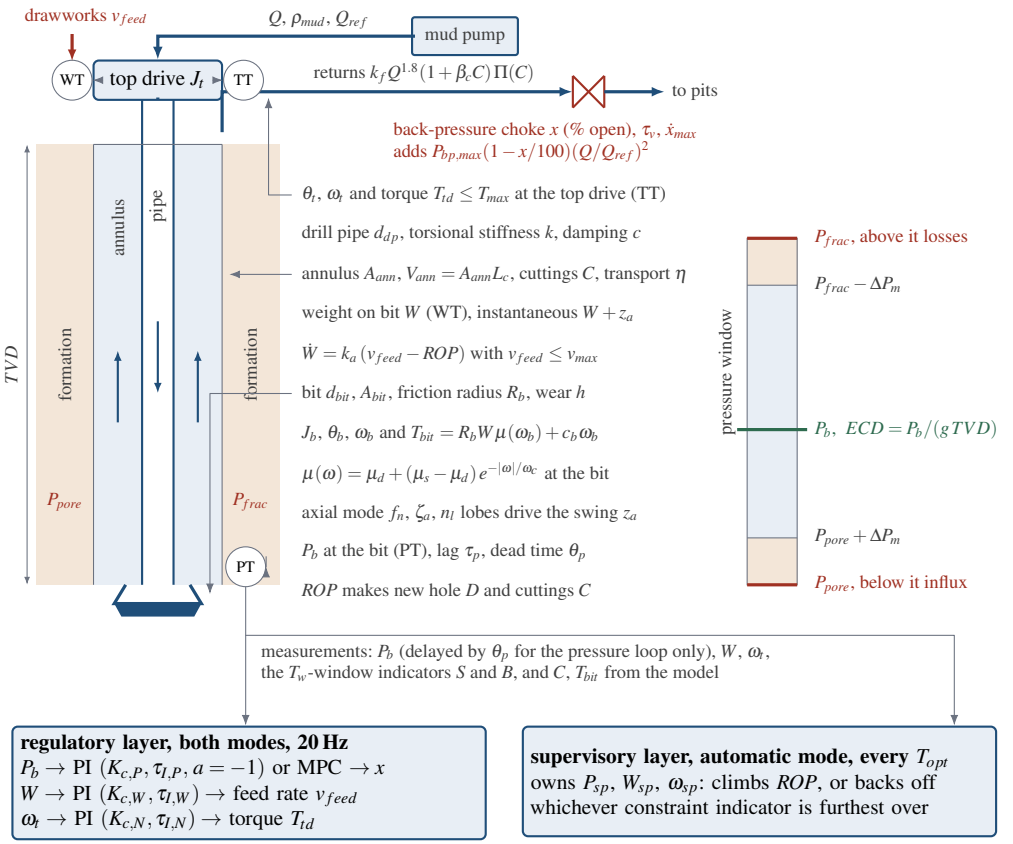

Figure 2.1 Well, surface circulating system and the two control layers. The weight and rotary loops run in both modes. In manual mode the operator sets the choke position and the three setpoints, and the pressure loop tracks the choke.

$$f_W = \chi^m \, \mathrm{e}^{m(1-\chi)}, \qquad \chi = \frac{W - W_t}{W_f - W_t} \tag{2.2}$$

where W_t is the weight below which the cutters do not engage and f_W is zero, W_f is the founder point and m sets how sharply the curve turns over. Differentiating gives $f'_W = m\chi^{m-1}\mathrm{e}^{m(1-\chi)}(1-\chi)$, which vanishes only at $\chi = 1$, so f_W peaks at exactly $W = W_f$ with the value 1 and the founder point is a declared parameter rather than an emergent property. The threshold also creates a region with no usable gradient, because f_W is identically zero below W_t and a gradient search that enters that region receives no direction, a consequence quantified in Section 2.8.

The remaining three factors are given by Equation 2.3.

$$f_N = (\omega/\omega_{ref})^{\gamma} \tag{2.3a}$$

$$f_h = \mathrm{e}^{-h/h_c} \tag{2.3b}$$

$$f_{\Delta P} = \exp\left[-\lambda \, \max(P_b - P_{pore}, 0)\right] \tag{2.3c}$$

$$f_C = \left[1 + (C/C_{50})^2\right]^{-1} \tag{2.3d}$$

Table 2.1 Every Signal in the Model, With Its Role and Range

Symbol	Role	Range (default)	Unit
P_b	CV, state	—	Pa
W	CV, state	≥ 0	N
ω_t	CV, state	—	rad/s
P_{sp}	setpoint	4200–6200 (5352)	psi
W_{sp}	setpoint	0–60 (25.0)	klbf
ω_{sp}	setpoint	0–250 (120)	RPM
x	MV, state	0–100 (65)	% open
v_{feed}	MV	0–0.030	m/s
T_{td}	MV	0–45	kN m
θ_t, θ_b	state, angles	—	rad
ω_b	state, bit speed	—	rad/s
$\bar{\omega}_b$	state, averaged bit speed	—	rad/s
z_a, $\dot{z}_a$	state, axial weight swing	—	N, N/s
C	state, cuttings fraction	0–0.30	—
h	state, bit wear	0–8	—
D	state, new hole	≥ 0	m
Δ_{hard}	DV, formation strength	-40 to $+80$ (0)	%
a_{abr}	DV, abrasiveness	0.4–2.5 (1)	$\times$
ΔP_{pore}	DV, pore shift	± 580 (0)	psi
ΔP_{frac}	DV, fracture shift	-870 to $+580$ (0)	psi
ROP	output, objective	—	m/s
S	output, stick-slip severity	—	—
B	output, bounce index	—	—
T_{bit}	output, bit torque	—	N m
ECD	output, circulating density	—	kg/m^3
MSE	output, specific energy	$\leq 5 \times 10^9$	Pa

where $\gamma < 1$ makes rotary speed sub-linear, h_c is the wear decay constant, λ is the chip hold-down coefficient and C_{50} is the cuttings concentration that halves the drilling rate. Equation 2.3c is the coupling that makes managed pressure drilling effective. Overbalance, the amount by which bottomhole pressure exceeds pore pressure, presses each cutting back onto the rock it was just sheared from, and the bit has to break it loose a second time. At $\lambda = 0.12/\text{MPa}$ the penalty is 7.9 % of drilling rate per 100 psi, which is the economic case for back-pressure control.

The rate of penetration is driven by an averaged bit speed rather than the instantaneous one, as shown in Equation 2.4.

$$\tau_R \frac{\mathrm{d}\bar{\omega}_b}{\mathrm{d}t} = \omega_b - \bar{\omega}_b \tag{2.4}$$

where $\tau_R = 2.0\,\text{s}$ is the depth-of-cut averaging time. Inside a stick-slip cycle the instantaneous bit speed runs from zero to several times surface speed within a second while rock removal integrates over many revolutions, so $\bar{\omega}_b$ is what appears in f_N, in the cuttings generation term and in the wear law, and it is what makes the dynamic model agree with the steady-state evaluator at an oscillating point.

2.3.3 Annular Hydraulics and the Choke

Bottomhole pressure is a hydrostatic head, a friction loss and a back-pressure, with a first-order lag between that sum and the pressure the bit actually sees, as shown in Equation 2.5.

$$\tau_p \frac{\mathrm{d}P_b}{\mathrm{d}t} = P_b^{ss}(C,x) - P_b \tag{2.5a}$$

$$\begin{aligned} P_b^{ss} = {} & \left[\rho_{mud} + C(\rho_c - \rho_{mud})\right] g\,TVD \\ & + k_f Q^{1.8}\,(1+\beta_c C)\,\Pi(C) \\ & + P_{bp,max}\left(1 - \frac{x}{100}\right)\left(\frac{Q}{Q_{ref}}\right)^2 \end{aligned} \tag{2.5b}$$

$$\Pi(C) = 1 + k_{pack}\,\max(C - C_{pack}, 0)/C_{pack} \tag{2.5c}$$

where ρ_c is the density of the cuttings, k_f is the annular friction coefficient, β_c is the extra friction each unit of cuttings loading contributes, Π is the pack-off multiplier and $P_{bp,max}$ is the back-pressure a fully shut choke produces at the sizing flow rate Q_{ref}. Three separate consequences follow from Equation 2.5b and each of them appears in a scenario. The choke position x is percent *open*, so closing it raises the pressure and the process gain of the pressure loop is $\partial P_b/\partial x = -40\,\mathrm{kPa}/\% = -5.80\,\mathrm{psi}/\%$ at the sizing flow. The back-pressure term carries $(Q/Q_{ref})^2$, so at the 174 gpm of the pack-off scenario the choke has only 78 psi of authority against 580 psi at full flow, and a reduced pump rate removes most of the choke authority. The cuttings appear twice, once as added weight and once as added friction, and above C_{pack} the multiplier Π climbs steeply, which is pack-off. Equivalent circulating density, $ECD = P_b/(g\,TVD)$, is reported because it compares directly against a pore or fracture gradient. Section 1.4.2 sets out the same column, friction and back-pressure for a producing well.

The choke itself is first order on its command with a stroking rate limit applied inside the controller, as shown in Equation 2.6.

$$\tau_v \frac{\mathrm{d}x}{\mathrm{d}t} = x_{cmd} - x, \qquad x \in [0, 100] \tag{2.6}$$

where τ_v is the stroke time constant and x is clamped to its travel after every integration step.

2.3.4 Cuttings in the Annulus

The annulus is treated as one well-mixed volume of length L_c carrying a single concentration, with generation set by the bit and removal by the flow, as shown in Equation 2.7.

$$V_{ann} \frac{\mathrm{d}C}{\mathrm{d}t} = ROP \cdot A_{bit} - C\,Q\,\eta(\bar{\omega}_b) \tag{2.7a}$$

$$\eta = \eta_{max} \frac{v_{ann}}{v_{ann} + v_0} \cdot \frac{1 + a_r\,\bar{\omega}_b/(\bar{\omega}_b + \omega_r)}{1 + a_r} \tag{2.7b}$$

where $A_{bit} = \pi d_{bit}^2/4$, $A_{ann} = \pi(d_{bit}^2 - d_{dp}^2)/4$, $V_{ann} = A_{ann}L_c$, $v_{ann} = Q/A_{ann}$ is the annular velocity, η is the transport efficiency, v_0 is the velocity at which half the ceiling efficiency is reached and a_r is the benefit a turning pipe gives by stirring the low side of the hole [5]. Both sides of Equation 2.7a are volumetric flow rates, so the equation is an inventory balance and $V_{ann}/(Q\eta)$ is its time constant, which is 340 s at the default flow and 120 RPM.

Equation 2.7a closes a loop back onto Equation 2.1 and that loop is what keeps the concentration finite. Drilling faster makes more cuttings, cuttings raise both the density and the friction in Equation 2.5b, higher bottomhole pressure means more overbalance, and chip hold-down in Equation 2.3c slows the drilling back down. In the pack-off scenario that negative feedback settles the concentration at 5.63 % rather than letting it run away, which is just above the pack-off threshold of 5.5 %.

2.3.5 Weight on Bit Dynamics

The drawworks pays line out at v_{feed} and the bit drills the hole away at ROP, and the string carries the difference as elastic strain, as shown in Equation 2.8.

$$\frac{\mathrm{d}W}{\mathrm{d}t} = k_a\,(v_{feed} - ROP), \qquad W \geq 0 \tag{2.8}$$

where k_a is the axial stiffness from the drawworks to the bit. Both terms are velocities and k_a is a force per unit length, so the right-hand side is a force per unit time. Equation 2.8 has no self-regulation and is the integrating process of the control literature [7]. A feed rate above the penetration rate raises the weight without bound and a feed rate below it drains the weight to zero, so the weight loop needs integral action to hold a setpoint and an autodriller manipulates a rate rather than a position.

2.3.6 Torsion and the Stick-Slip Indicator

The drillstring is two inertias joined by a twisting pipe, as shown in Equation 2.9.

$$J_t\,\dot{\omega}_t = T_{td} - k(\theta_t - \theta_b) - c(\omega_t - \omega_b) \tag{2.9a}$$

$$J_b\,\dot{\omega}_b = k(\theta_t - \theta_b) + c(\omega_t - \omega_b) - T_{bit} \tag{2.9b}$$

$$T_{bit} = R_b\,W\,\mu(\omega_b) + c_b\,\omega_b \tag{2.9c}$$

$$\mu(\omega) = \mu_d + (\mu_s - \mu_d)\,\mathrm{e}^{-|\omega|/\omega_c} \tag{2.9d}$$

with $\dot{\theta}_t = \omega_t$ and $\dot{\theta}_b = \omega_b$, where J_t and J_b are the top-drive and bit inertias, k and c are the torsional stiffness and damping of the pipe, R_b is the effective friction radius of the bit, c_b is viscous damping at the bit and ω_c is the velocity-weakening scale. Equation 2.9d contains the destabilizing mechanism. A bit that slows down meets a higher friction factor, which slows it further, so the bit torque has a negative slope with respect to bit speed wherever the exponential is still active. Setting that negative damping against the positive damping supplied by the pipe and the mud gives the dimensionless group of Equation 2.10.

$$G = \frac{R_b\,W\,(\mu_s - \mu_d)\,\mathrm{e}^{-\omega/\omega_c}/\omega_c}{c + c_b} \tag{2.10}$$

Both terms in Equation 2.10 are torques per unit of angular velocity, so G is dimensionless and $G = 1$ is the naive onset, the exact one if the top drive held its speed perfectly. The application uses $G_0 = 1.092$ as the boundary, the value at which the linearized string, two inertias, the twisting pipe and the speed loop of Section 2.5 at its default tuning, acquires a growing oscillatory mode. The onset is above unity because the speed loop lets the heavy top drive move and absorb part of the torsional energy, and it rises with the loop gain to 1.16 at four times the default. Near the onset the cycle grows slowly, so a run started at a steady state just above G_0 stays quiet for many minutes, and the instability is subcritical, so a cycle already running survives far below it, as Section 2.4 shows. The exponential is what gives the unstable region its shape. At 60 RPM it is 0.106 and at 250 RPM only 8.7×10^{-5}, so G grows with weight and collapses with rotary speed, and the unstable set is a corner at high weight and low rotary speed.

The severity itself is measured on the running simulation as the peak-to-peak bit speed over twice the mean surface speed within a moving window T_w. The map and the optimizer evaluate hundreds of candidate points per execution and cannot integrate the dynamic model at each one, so they use a correlation fitted to the same limit cycle, as shown in Equation 2.11.

$$\Lambda = \frac{R_b\,W\,(\mu_s - \mu_d)}{2\,\bar{\omega}_t\sqrt{k J_b}} \tag{2.11a}$$

$$S = 1.343\,\Lambda^{0.712}\left[1 - \mathrm{e}^{-12(G - G_0)}\right] \quad \text{for } G > G_0 \tag{2.11b}$$

where Λ is the energy group obtained by winding the pipe up through the torque drop $R_b W(\mu_s - \mu_d)$ and releasing the stored energy into the bit inertia, and $S = 0$ below the onset. The map displays S as written, while feasibility is decided on the amplitude without the ramp factor, $1.343\,\Lambda^{0.712}$ for $G > G_0$. Equation 2.11 is a correlation and not a solution of Equations 2.9, and its accuracy is quantified in Section 2.4 rather than assumed.

2.3.7 The Axial Mode and the Bounce Index

The bottomhole surface develops a pattern of n_l lobes, so a bit turning at ω_b is excited at $n_l \omega_b$. The dynamic part of the weight on bit is the response of a damped second-order axial mode to that excitation, written in newtons so that the equation is the unit-mass normal form, as shown in Equation 2.12.

$$\ddot{z}_a = \omega_n^2 \left[F_0 \sin(n_l \theta_b) - z_a \right] - 2\zeta_a \omega_n \dot{z}_a \tag{2.12a}$$

$$F_0 = \varepsilon_0 \sqrt{W / W_{ref}}\, W \tag{2.12b}$$

where $\omega_n = 2\pi f_n$, ζ_a is the axial damping ratio, ε_0 is the lobe excitation fraction at the reference weight and W_{ref} is that reference. The instantaneous weight on bit is $W + z_a$. The steady amplitude of Equation 2.12 is $F_0 M(r)$ with the textbook magnification factor, which gives the closed-form bounce index of Equation 2.13.

$$B = \varepsilon_0 \sqrt{W / W_{ref}}\, M(r), \quad M(r) = \left[(1 - r^2)^2 + (2\zeta_a r)^2 \right]^{-1/2} \tag{2.13}$$

where $r = n_l \omega / \omega_n$ is the excitation frequency ratio. At $B = 1$ the instantaneous weight reaches zero and the bit leaves bottom. The resonance sits at $\omega = \omega_n / n_l$, which with $f_n = 5.00\,\text{Hz}$ and $n_l = 3$ is exactly 100 RPM, and the index there is 0.833 at 25 klbf against 0.216 at 120 RPM, 0.079 at 150 RPM and 0.019 at 250 RPM. The two vibration constraints differ in shape. Stick-slip occupies a corner because Equation 2.10 couples weight and speed through one exponential, while bit bounce occupies a horizontal band because a resonance depends on frequency and only weakly on weight. The response to stick-slip is to raise rotary speed, and the response to bounce is to move the speed out of the band in whichever direction is nearer.

Pause. Try Exercise 2.3 before reading on.

2.3.8 Bit Wear and Specific Energy

Wear accumulates as a power law in weight and speed, as shown in Equation 2.14.

$$\frac{\mathrm{d}h}{\mathrm{d}t} = \dot{h}_0\, a_{abr} \left(\frac{W}{W_{ref}} \right)^{1.2} \left(\frac{\bar{\omega}_b}{\omega_{ref}} \right)^{0.8}, \qquad h \in [0, 8] \tag{2.14}$$

where $\dot{h}_0$ is the wear rate at the reference weight and speed and a_{abr} is the abrasiveness of the formation. The derivative is forced to zero once h reaches 8. Mechanical specific energy, the quantity a driller uses to decide whether the bit is cutting or rubbing [13], is the axial work plus the rotary work per unit volume removed, as shown in Equation 2.15.

$$MSE = \frac{W}{A_{bit}} + \frac{T_{bit}\, \bar{\omega}_b}{A_{bit}\, ROP} \tag{2.15}$$

The second term is a power divided by a volumetric removal rate and both terms are pressures. It is capped at 5×10^9 Pa so that a bit turning and making no hole logs a large finite number rather than an infinity. MSE does not track the drilling rate, because at steady state stepping the weight from 25 klbf to the founder point at 35 klbf raises the rate of penetration from 67.6 ft/hr to 79.9 ft/hr and the specific energy from 28.5 kpsi to 32.8 kpsi at the same time.

Table 2.2 Model Constants, Their Defaults and Their Adjustable Ranges

Symbol	Meaning	Value	Range	Unit	Symbol	Meaning	Value	Range	Unit
Well and geometry					*Cuttings transport*				
TVD	true vertical depth	3000	1500–5000	m	η_{max}	transport efficiency ceiling	0.90	0.30–1.00	—
d_{bit}	bit diameter	0.2159	0.152–0.311	m	v_0	half-transport velocity	0.80	0.10–2.50	m/s
$d_{dp}^{\dagger}$	pipe outside diameter	0.1270	0.089–0.168	m	a_r	rotation cleaning benefit	0.60	0–2.0	—
L_c	cuttings transport length	200	60–900	m	ω_r	half-benefit rotary speed	8.0	1–25	rad/s
$g^{\dagger}$	gravity	9.81	9.7–9.9	m/s^2	C_{pack}	pack-off concentration	0.055	0.02–0.15	—
Mud and hydraulics					k_{pack}	pack-off friction gain	4.0	0–15	—
ρ_{mud}	mud density	1140	950–1800	kg/m^3	*Drillstring*				
$\rho_c^{\dagger}$	cuttings density	2500	2100–2900	kg/m^3	J_t	top drive inertia	2000	300–6000	kg m^2
Q	pump flow rate	0.030	0.008–0.055	m^3/s	J_b	bit and BHA inertia	120	30–500	kg m^2
$Q_{ref}^{\dagger}$	choke sizing flow	0.030	0.010–0.055	m^3/s	k	torsional stiffness	500	100–2500	N m/rad
k_f	annular friction coefficient	6.5	2–14	10^8 SI	c	string torsional damping	30	0–200	N m s/rad
β_c	cuttings friction loading	12	0–40	—	c_b	bit viscous damping	30	0–200	N m s/rad
$P_{bp,max}$	back-pressure, choke shut	4.0	1–9	MPa	R_b	bit friction radius	0.10	0.03–0.20	m
τ_p	bottomhole pressure lag	1.5	0.2–8.0	s	μ_s	static friction factor	0.80	0.30–1.40	—
θ_p	pressure dead time	0.0	0.0–4.0	s	μ_d	dynamic friction factor	0.25	0.05–0.70	—
τ_v	choke stroke constant	2.0	0.3–8.0	s	ω_c	velocity-weakening scale	2.80	0.50–10.0	rad/s
$\dot{x}_{max}$	choke rate limit	20	2–60	%/s	k_a	string axial stiffness	227	50–700	kN/m
Formation					$\tau_R^{\dagger}$	depth-of-cut averaging	2.0	0.2–10	s
P_{pore}	pore pressure	34.0	24–44	MPa	*Axial mode*				
P_{frac}	fracture pressure	41.5	28–55	MPa	f_n	axial natural frequency	5.00	2.0–9.0	Hz
Rate of penetration					ζ_a	axial damping ratio	0.060	0.015–0.30	—
K	drillability coefficient	0.011	0.002–0.030	m/s	ε_0	lobe excitation fraction	0.100	0.02–0.30	—
W_t	cutting threshold weight	22.2	0–60	kN	$W_{ref}^{\dagger}$	reference weight	111.2	40–220	kN
W_f	founder point	155.7	60–250	kN	n_l	bottomhole pattern lobes	3	2–5	—
m	founder sharpness	3	1.5–6	—	*Rig, detector and optimizer*				
$\omega_{ref}^{\dagger}$	reference rotary speed	12.566	5–25	rad/s	v_{max}	maximum feed rate	0.030	0.005–0.060	m/s
γ	rotary speed exponent	0.60	0.30–1.00	—	T_{max}	top drive torque limit	45	15–80	kN m
h_c	wear decay constant	6.0	2.0–20.0	—	$T_w^{\dagger}$	vibration detector window	20.0	4–60	s
λ	chip hold-down	0.12	0–0.40	1/MPa	T_{opt}	optimizer interval	2.0	0.5–10	s
C_{50}	half-rate cuttings loading	0.050	0.010–0.150	—	$\dot{W}_{sp}$	weight setpoint rate limit	1.5	0.1–8.0	kN/s
$\dot{h}_0$	wear rate at reference	0.50	0–3.0	10^{-3}/s	$\dot{\omega}_{sp}$	rotary setpoint rate limit	0.35	0.02–2.0	rad/s^2
					$\dot{P}_{sp}$	pressure setpoint rate limit	10.0	0.5–40	kPa/s

† not operator adjustable. k_f is in Pa·(s/m^3)$^{1.8}$.

2.3.9 Parameters, Initialization and Numerics

Table 2.2 collects every constant with its default and its adjustable range. Eight are not operator adjustable, because they are universal constants, internal reference values that only non-dimensionalize a ratio, or instrument configuration rather than plant properties.

Five further quantities are exposed as operator-adjustable limits rather than as parameters, because they are engineering choices and not properties of the well. They are the stick-slip limit $S_{max} = 0.40$ on 0.20 to 3.00, the bounce limit $B_{max} = 0.50$ on 0.10 to 1.50, the cuttings limit $C_{max} = 2.2\,\%$ by volume on 0.5 to 8, the bit torque limit $T_{lim} = 20\,\mathrm{kN\,m}$ on 4 to 55, and the pressure window margin $\Delta P_m = 0.7\,\mathrm{MPa}$ on 0.2 to 2.4 MPa, applied at both ends of the window. Only pack-off at C_{pack}, losses above P_{frac}, influx below P_{pore} and the rig limits v_{max} and T_{max} are physics. The distinction is the point of the graded challenge of Section 2.6.7.

Every run begins at a genuine steady state. The initializer takes the weight and rotary setpoints, takes the choke from the operator-set position in manual mode or from the position that puts P_b on setpoint in automatic mode, and then solves the implicit cuttings balance. Because ROP is strictly decreasing in C, both directly through Equation 2.3d and again through the head and friction the cuttings add to Equation 2.5b, the residual $ROP(C)A_{bit} - CQ\eta$ is strictly decreasing and has exactly one root, found by twenty bisections on $[0, 0.30]$. The pipe is initialized already wound up at $\theta_t = T_{bit}/k$, and the axial mode is placed on its steady orbit with $z_a = -F_0 M \sin\phi$ and $\dot{z}_a = F_0 M n_l \omega_b \cos\phi$ where $\phi = \mathrm{atan2}(2\zeta_a r,\ 1 - r^2)$, because at the default 120 RPM a state starting

at rest would ring at 1.5 times the eventual amplitude and the bounce indicator would report it.

The engine and its three clocks are those of Section 1.6.1. Integration is fixed-step fourth-order Runge-Kutta at $\Delta t = 0.005$ s, sixteen steps per cycle of the fastest thing in the model, which is the axial excitation at 12.5 Hz when the rotary speed is at the top of its range. Regulatory execution comes every 0.05 s, exactly ten physics steps, and supervisory execution every $T_{opt} = 2.0$ s, exactly forty regulatory executions. Logging runs at 20 samples per second, and the seed is 20260824, although the model carries no noise at all, so every run repeats exactly.

2.3.10 Assumptions

Each simplification below is stated with the place it fails.

One vertical section. There is no directional work, no sliding, no torque and drag along a curved wellbore and no drag from pipe lying on the low side of the hole. Those terms dominate the torsional problem in a deviated well, and a stick-slip severity computed here would not transfer to one.

Five factors in the drilling rate. A field rate-of-penetration law is regressed per field from offset wells and carries eight or more factors, including hydraulic horsepower at the bit and compaction with depth [4]. The values here are plausible for a soft-to-medium 8.5 in hole section and are not portable.

Two torsional degrees of freedom and one axial mode. A real bottomhole assembly has dozens of coupled modes. Lateral vibration, forward whirl and backward whirl are absent entirely, as is the axial-torsional coupling that makes real bit bounce and real stick-slip feed each other.

Bit bounce is detected and never fed back. The instantaneous weight fluctuation is computed, flagged and used as a constraint, but the bit is never taken off bottom. A real bouncing bit loses contact, so its average depth of cut and its own torque both fall, and it destroys cutters while doing so. Everything the application reports about bit bounce is therefore an indicator and not a consequence.

One well-mixed cuttings concentration. Real hole cleaning is a bed that forms on the low side and whose behavior depends on inclination more than on anything modeled here [5]. The lumped volume uses the residence time of the lower 200 m of annulus but applies the concentration to the whole column when computing *ECD*, which compresses a real bottoms-up time of half an hour into a transient short enough to observe in one session.

A linear choke. Back-pressure goes as $(1 - x/100)$ rather than as an orifice law $\Delta P \propto Q^2/x^2$, so the loop has one process gain to teach rather than three, and the very high gain a real choke shows near the shut end is absent.

Quasi-steady hydraulics. There is no surge and swab, no gas migration, no compressibility wave along the annulus, no cuttings slip velocity and no temperature effect on mud density. One first-order lag τ_p is the only dynamics between the choke position and the bottomhole pressure, and the optional dead time θ_p delays only the measurement the pressure loop reads.

Blind wear. Wear follows Equation 2.14 and nothing else. It does not accelerate with vibration, the bit never fails suddenly, balls up or loses a cone, and tripping for a new bit is free here while in a well it costs hours.

2.4 Feasible Region Analysis

Everything about a candidate operating point is available without integrating anything. Given a weight, a rotary speed and a choke position, the bisection described above returns the steady cuttings concentration, and from it the bottomhole pressure, the drilling rate, the bit torque, the two vibration indicators and the specific energy. That evaluator paints the map, the supervisory optimizer climbs it, and the verification suite compares the integrated model against it. Marching the dynamic model for 1300 s from a deliberately wrong start, no cuttings in the annulus and 60 % of the target weight, and then averaging over a further 60 s, reproduces the evaluator's drilling rate, pressure and concentration at four operating points to within 2.4 %.

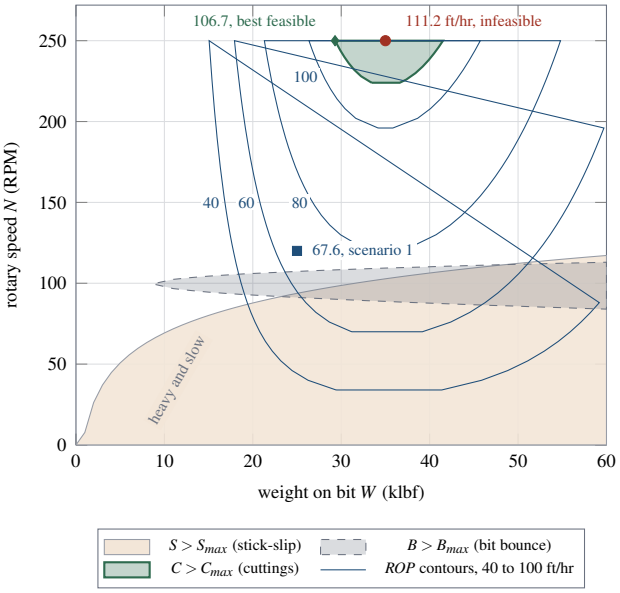

Figure 2.2 Objective contours and the three binding constraints in the weight and rotary speed plane, computed from Equations 2.1 to 2.13 at the scenario-1 choke position of 65 % open with a new bit. Stick-slip forbids the shaded corner with a thin solid outline, bit bounce the band with a dashed outline and the cuttings limit the lens with a thick outline. The circle marks the unconstrained peak, the diamond the best feasible point and the square the scenario-1 point. The unconstrained peak sits inside the cuttings lens, so the constrained optimum lies on the lower edge of the lens.

2.4.1 Shape of the Feasible Region

Figure 2.2 draws the objective and the three constraints that carve into it at the choke position of the first scenario. The objective contours are closed loops, because f_W rises and then founders while f_N increases monotonically, so every level set has a light branch and a heavy branch. The three forbidden sets have three different shapes. Stick-slip occupies the region below a boundary running from 92 RPM at 23 klbf to 117 RPM at 60 klbf. Bit bounce occupies a horizontal stripe from 84 RPM to 113 RPM at the heaviest weight, narrowing to nothing below 9 klbf where even the full resonant magnification cannot swing the weight by half of its mean. The cuttings limit occupies a lens in the top right that contains the unconstrained peak, and that lens is the one that binds.

2.4.2 A Closed-Form Ceiling on the Objective

The cuttings constraint is not an ordinary boundary in the plane. Setting $\frac{\mathrm{d}C}{\mathrm{d}t} = 0$ in Equation 2.7a gives $C = ROP\,A_{bit}/(Q\eta)$, so requiring $C \le C_{max}$ is exactly a bound on the objective itself, as shown in Equation 2.16.

$$ROP \le \frac{C_{max}\,Q\,\eta(\omega)}{A_{bit}} \tag{2.16}$$

Weight on bit does not appear in Equation 2.16 and neither does the choke. The ceiling depends on rotary speed alone, through the transport efficiency, and η saturates, so the ceiling rises from 92.4 ft/hr at 60 RPM to 99.9 ft/hr at 120 RPM and only 106.7 ft/hr at 250 RPM. Two consequences follow and both are visible on the application.

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

The first is that the choke stops contributing beyond a crossover position. The unconstrained peak grows steadily as the choke opens, from 89.4 ft/hr with the choke shut to 111.2 ft/hr at 65 % open and

124.1 ft/hr wide open, because Equation 2.3c multiplies the whole surface. The best feasible point follows it only until the peak crosses into the lens at $x = 52.1\,\%$ open, and past that it is pinned at 106.7 ft/hr. Managed pressure drilling adds 19 % of drilling rate between a shut choke and 52 % open and nothing after that, while the operating point that delivers the ceiling slides along the boundary from 35.0 klbf to 24.7 klbf as the choke opens fully.

The second is that the binding constraint identifies its own remedy, and the remedy is not one of the three manipulated variables. Raising the pump rate from 476 gpm to 491 gpm, an increase of 3.3 %, lifts the ceiling to 111.7 ft/hr, above the unconstrained peak, so the whole plane becomes feasible and the achievable rate rises to 111.2 ft/hr. A further increase in pump rate reduces the drilling rate, because annular friction in Equation 2.5b raises bottomhole pressure and chip hold-down removes the gain, and at 602 gpm the achievable rate has fallen to 105.9 ft/hr. Flow rate is a fourth degree of freedom with an interior optimum of its own, and it is held fixed so that the three degrees of freedom without one can be studied.

2.4.3 Dependence of the Binding Constraint on the Formation

Scaling the whole objective by a constant cannot move an unconstrained optimum, and it does move a constrained one. Firing the hard-stringer event multiplies K by $1/1.45$, and the achievable rate at the best point falls from 106.7 ft/hr to 85.5 ft/hr while the concentration there falls to 1.76 %, under the limit, so the cuttings constraint releases and the answer moves back to the founder point at 35.0 klbf and 250 RPM. Wear does the same thing continuously, and on a bit worn to $h = 6$ of 8 the best point makes 51.4 ft/hr at the founder point with 1.06 % cuttings and is unconstrained. Nothing about the plant changed in either case.

2.4.4 Accuracy of the Vibration Indicators

The bounce index of Equation 2.13 is the exact steady amplitude of Equation 2.12 and is treated as such by the verification suite, which checks the dynamic model against it wherever the bit speed is steady enough for the excitation frequency to mean anything. The stick-slip severity of Equation 2.11 is a correlation, and the difference matters enough to quantify.

The onset G_0 comes from the linearization rather than from a fit. The two-inertia string of Equation 2.9 with the proportional-integral speed loop at its default tuning has a pair of eigenvalues that crosses into the right half-plane at $G = 1.092$, with a period of 2.99 s, at every rotary speed. The integrated model agrees: a settled run kicked by 10^{-6} rad/s of bit speed decays at $G = 1.08$ and grows at $G = 1.10$ at 60 RPM, 80 RPM and 100 RPM. The growth is slow near the onset, with an e-folding time of about 100 s at $G = 1.13$ and 33 s at $G = 1.22$, so a run started at a steady state can take a long time to show the cycle. With the weight and speed chosen and the run started, cells at 70 RPM to 85 RPM with G from 1.18 to 1.23 pass a severity of 0.40 after 487 s to 713 s, cells with G from 1.13 to 1.17 after 863 s to 1659 s, and 15.0 klbf at 80 RPM ($G = 1.096$) not within 2400 s. Every one of these cells lies above G_0 and the map marks it.

Driving the dynamic model at 96 points spanning 9 klbf to 58.5 klbf and 40 RPM to 110 RPM, each started on weight with an empty annulus and frozen wear and run for 600 s, produces a limit cycle at 74 of them. The windowed severity of a developed cycle wanders by tens of percent from one window to the next, so it is averaged over the last 300 s of each run. Comparing that average with the displayed value of Equation 2.11 gives a median relative error of 9.2 %, a ninetieth percentile of 17.6 % and a worst case of 21.1 % over the 72 points at which G exceeds the onset by more than 15 %. That is the accuracy a published field severity chart has. The application's own boundary test, twelve runs of 160 s at points at least 7 % from G_0 across the whole adjustable range, puts every simulated point on the predicted side of G_0.

Inside the 15 % band above the onset the ramp factor in Equation 2.11 lowers the displayed value while the dynamic model, once its cycle has grown, runs at full amplitude. At 36 klbf and 100 RPM, where $G = 1.245$, the simulation settles at a severity of 1.81 against a displayed 1.66, and at 49.5 klbf and 110 RPM ($G = 1.178$) at 1.98 against 1.49. Feasibility is therefore decided on the unramped

amplitude, 1.97 and 2.31 at those points. With that rule the map and the simulation agree about the 0.40 limit at 95 of the 96 points. The exception is 22.5 klbf at 90 RPM, where the map marks stick-slip and the 600 s run has not yet grown its cycle. A run started at the steady state there develops it at 1629 s, so the map is conservative for short runs near the onset and right for long ones.

A cycle once started also survives far below the onset. Moving from the scenario-1 point to 9 klbf and 70 RPM ($G = 0.96$) ten seconds into the run starts a cycle within 6 s that is still at a severity of 1.18 after a further 900 s, while the same point chosen before the run starts never leaves steady rotation. The boundary on the map therefore marks where steady rotation loses stability and not where a running cycle stops.

2.5 Controller Design

2.5.1 The Regulatory Layer

The default three proportional-integral loops run at 20 Hz on the shared velocity-form algorithm of Section 1.6.2, with derivative action set to zero throughout. The weight and rotary loops run in both modes, and the pressure loop runs in automatic mode and tracks the operator's choke in manual mode. Written in the loops' own symbols, with executions indexed by k and separated by $\Delta t = 0.05$ s, the algorithm is Equation 2.17.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k) \tag{2.17a}$$

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

$$\Delta u_k^I = \frac{K_c}{\tau_I}\,e_k\,\Delta t$$

$$q_k = q_{k-1} + \Delta u_k^P$$

$$u_k^* = q_k + I_{k-1} \tag{2.17b}$$

$$u_k = \mathrm{slew}\big(\mathrm{sat}_{[u_{min},u_{max}]}(u_k^*),\ u_{k-1},\ \dot{u}_{max},\ \Delta t\big) \tag{2.17c}$$

$$I_k = I_{k-1} + \Delta u_k^I + \frac{\Delta t}{T_t}\,(u_k - u_k^*) \tag{2.17d}$$

where $a = \pm 1$ is the loop action, q_k is the accumulator that carries the proportional increments without limit (the running sum S_k of Equation (1.17)), u^* is the unclamped output the algorithm demands (the v_k of Equation (1.18)) and $T_t = \tau_I$ is the tracking time of the back-calculation. The output bias is zero and the initial actuator value is carried in I. Anti-windup is the back-calculation term of (2.17d), which corrects the integral alone and never truncates q_k, as Section 1.6.2 explains [21]. At the default $\tau_{I,P} = 4$ s each execution removes $1/80$ of the excess. The tracking term acts in the reported runs mainly in the challenge as delivered, where the drawworks feed sits at zero, its lower clamp, for 68 % of the run. In manual mode the pressure loop sets its integral to $I_k = u_{man} - q_k$, and in automatic mode the operator's choke value follows the controller output after every execution, so transfer is bumpless in both directions with no dedicated transfer logic.

The pressure loop is the loop with negative process gain. Its manipulated variable is the choke position in percent open and opening the choke lowers the pressure, so the loop is declared with $a = -1$. That single declaration is the only place a minus sign appears, and the loop is otherwise identical to the other two. Its output is slew limited at $\dot{x}_{max} = 20\,\%/\mathrm{s}$ before it reaches the stroke lag of Equation 2.6.

Table 2.3 lists the six tuning constants, and three derived quantities follow from them. The pressure loop has $K_c|K_p| = 2.0$ with $\tau_I = 4$ s, giving a velocity constant $K_v = K_c|K_p|/\tau_I = 0.5/\mathrm{s}$ against a plant lag of $\tau_p + \tau_v = 3.5$ s. The weight loop closes on an integrating plant, so its proportional gain sets a first-order closed-loop time constant of $1/(k_a K_{c,W}) = 2.5$ s with $\tau_I = 15$ s well outside it. The rotary loop sees the top-drive inertia through $J_t/K_{c,N} = 4.0$ s, slower than the 3.1 s natural period $2\pi\sqrt{J_b/k}$

Table 2.3 Controller Settings and Their Adjustable Ranges

Symbol	Meaning	Default	Range
$K_{c,P}$	choke per pressure error	5.0×10^{-5} %/Pa	0–2 %/psi
$\tau_{I,P}$	pressure integral time	4.0 s	0.5–60
$K_{c,W}$	feed rate per weight error	1.762×10^{-6}	0–400 (ft/hr)/klbf
$\tau_{I,W}$	weight integral time	15 s	2–300
$K_{c,N}$	torque per speed error	500 N m s/rad	0–3000
$\tau_{I,N}$	rotary integral time	3.0 s	0.2–30

of the torsional mode, so it regulates the mean surface speed and leaves the torsional mode to the string. Stiffening it does not cure stick-slip. At 40 klbf and 60 RPM, $K_{c,N} = 1000\,\mathrm{N\,m\,s/rad}$ lowers the peak severity from 4.80 to 2.90 and starts the cycle sooner (flagged at 14.3 s against 20.2 s), 2000 gives 2.94 and 250 raises it to 6.47, and the cycle persists at all four gains. Field soft-torque systems tune the top drive to absorb the torsional wave, which a plain PI speed loop does not do [2, 15].

The instrument chain is short. Bottomhole pressure passes through a delay line of adjustable dead time θ_p, which is the only transport delay in the model, defaults to zero and is read only by the pressure loop, and weight and surface speed are read exactly with no noise and no bias anywhere. The two vibration indicators come from windowed peak-to-peak detectors of length $T_w = 20\,\mathrm{s}$, which collect the extremes of their signal at every physics step and write the envelope into a ring buffer every 0.05 s, because storing the instantaneous value would alias badly against an axial excitation that reaches 12.5 Hz. Both buffers are pre-filled with the starting value so that a run beginning at a settled point reads zero swing rather than climbing out of a fictitious transient. Those two indicators are slow by construction, so an override on either acts on information up to 20 s old.

2.5.2 The Supervisory Layer

Above the three loops, and only in automatic mode, is a constraint-aware hill climber that owns the three setpoints. It is the third shape of constraint control in Section 1.4.8, a supervisor that moves setpoints toward the active limit and backs away from it. It executes once every T_{opt} seconds of simulated time and takes exactly one of three actions per execution. Every execution begins by measuring how far past its own limit each of the six constraints currently is, as shown in Equation 2.18.

$$[v_1, v_2, v_3, v_4] = \left[\frac{S}{S_{max}}, \frac{B}{B_{max}}, \frac{C}{C_{max}}, \frac{T_{bit}(W, \bar{\omega}_t)}{T_{lim}}\right] - 1 \tag{2.18a}$$

$$v_5 = (P_b - P_{high})/s, \qquad v_6 = (P_{low} - P_b)/s \tag{2.18b}$$

where P_b is the undelayed bottomhole pressure, $P_{low} = P_{pore} + \Delta P_m$, $P_{high} = P_{frac} - \Delta P_m$, $s = \max(\Delta P_m, 0.1\,\mathrm{MPa})$ normalizes the two pressure entries to the same scale as the four ratios, and $\bar{\omega}_t$ is the window mean of surface speed rather than the instantaneous value, because inside a stick-slip cycle the torque swings from nothing to the breakaway value several times a second. Every entry of v is negative when its constraint has room and positive when it does not, and all six are comparable. The most-violated constraint is $j = \arg\max_i v_i$.

Three states follow from v_j, as shown in Equation 2.19.

$$\begin{cases} \text{back off} & v_j > 0 \\ \text{hold} & -0.05 < v_j \le 0 \\ \text{climb} & v_j \le -0.05 \end{cases} \tag{2.19}$$

The hold band in the middle is what makes an override usable on a slow constraint. The cuttings concentration has a time constant of 340 s, so a back-off that ran until the indicator read clear

would keep moving for minutes after the move that fixed it, which is integral windup against a lagging measurement. Backing off is also rate-scaled by how far over the limit the indicator is, through an urgency factor $\kappa = \mathrm{clamp}(v_j/0.15,\ 0.15,\ 1)$, so the move sizes are $\Delta W = \dot{W}_{sp} T_{opt}\,\kappa$ and $\Delta\omega = \dot{\omega}_{sp} T_{opt}\,\kappa$. The setpoint that moves depends on the selected constraint. Stick-slip, cuttings and torque all reduce weight, with the weight setpoint floored at the cutting threshold W_t, below which f_W is identically zero and the climb would find no gradient. Bit bounce moves the rotary speed, and the direction is decided by evaluating the closed form 0.4 rad/s either side and taking the lower, because below the resonance the escape is to slow down and above it to speed up. The two pressure entries move the pressure setpoint by $3\dot{P}_{sp} T_{opt}$.

When nothing is close, the optimizer climbs. It estimates the gradient of the objective by central differences on the steady-state evaluator at the live choke position, with perturbations of 2 % of each adjustable range, as shown in Equation 2.20.

$$g_W = ROP(W_{sp}{+}h_W,\ \omega_{sp}) - ROP(W_{sp}{-}h_W,\ \omega_{sp}) \tag{2.20a}$$

$$g_\omega = ROP(W_{sp},\ \omega_{sp}{+}h_\omega) - ROP(W_{sp},\ \omega_{sp}{-}h_\omega) \tag{2.20b}$$

$$W_{sp} \leftarrow W_{sp} + \dot{W}_{sp} T_{opt}\, g_W / \sqrt{g_W^2 + g_\omega^2} \tag{2.20c}$$

$$\omega_{sp} \leftarrow \omega_{sp} + \dot{\omega}_{sp} T_{opt}\, g_\omega / \sqrt{g_W^2 + g_\omega^2} \tag{2.20d}$$

where $h_\omega = 0.524\,\mathrm{rad/s}$ and h_W starts at 5.34 kN (1.2 klbf) and doubles until g_W is nonzero, so that a probe straddling the dead zone below W_t widens until it reaches weights that cut. Normalizing by the gradient magnitude means the move is always at the rate limit and the direction alone carries the information, so the setpoint trends move at the declared rate limits and the limits can be read from the recorded data.

Separately, and on every climbing execution, the pressure setpoint walks down toward a floor, as shown in Equation 2.21.

$$P_1 = \min\left(P_{low} + 0.15\,\mathrm{MPa},\ P_{high}\right) \tag{2.21a}$$

$$P_2 = P_b\big|_{x=100} + 0.10\,\mathrm{MPa} \tag{2.21b}$$

$$P_{floor} = \max(P_1,\ P_2) \tag{2.21c}$$

Equation 2.21b is the term that binds in practice. A wide-open choke adds no back-pressure at all, so the lowest bottomhole pressure the well can reach is whatever the mud column and the annular friction produce on their own, and that minimum rises as the annulus loads up. Enforcing the floor on every execution rather than only while climbing keeps the pressure loop off its limit when cuttings push the achievable minimum upward underneath it. On a well with a light mud column the choke travel, and not the formation, limits what managed pressure drilling can add.

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

Two further pieces complete the layer. On every execution, in both modes, a 26×22 sweep of the plane at the live choke position finds the highest feasible drilling rate, and a second pass over the points within 2 % of it takes the one with the lowest specific energy. The second pass exists because Equation 2.16 makes the cuttings boundary a contour of equally fast points with a light branch and a heavy branch, and without a tie-break the selected point alternates between them as the grid shifts. On every manual to automatic transition the three supervisory targets are re-seeded from the operator's setpoints, so the supervisory layer is bumpless in the same way the regulatory loops are, and a test that drives the transition measures a setpoint change of zero against a tolerance of 1 %.

2.5.3 Choke Model Predictive Control

Model predictive control (MPC) can replace the pressure PI loop as an option, while the weight and rotary loops, the supervisory layer and its protections stay as they are. The arrangement is the one of Section 1.6.3. The controller plans with a frozen nominal model in the pressure p and a baseline b, both in MPa, an estimated opening v and a command u, both in percentage points, and the current pump flow Q in m^3/s, which it takes as known exactly. With $a = 4(Q/0.030)^2$ MPa, the model is Equation 2.22.

$$\dot{v} = (u - v)/2, \quad \dot{p} = [b + a(1 - v/100) - p]/1.5, \quad \dot{b} = 0. \tag{2.22}$$

Its nominal gain is -40000 Pa per percentage point, and an independent step propagation agrees with it within 1.57×10^{-7} Pa. Every $h = 0.05$ s an exact discretization predicts (p^-, v^-, b^-) from the emitted command. The measured pressure p_m then sets $p = p_m$, $v = v^-$ and $b = b^- + g(p_m - p^-)$, with $g = (1 - e^{-h/5})/(1 - e^{-h/1.5})$. At initialization v is the known command and $b = p_m - a(1 - v/100)$. A change of pump flow adds $650(Q_{new}^{1.8} - Q_{old}^{1.8})$ MPa to b. This is a disturbance observer [22] with a filtered bias in place of a Kalman gain, and it reads no cuttings loading, wear, actual valve state or live model coefficient. The offset-free controller of Volume 9, Chapter 4 carries its disturbance in a Kalman estimator instead.

A decision is made every 1 s. The prediction holds each command over its block and spans 60 s in blocks of $(1, 1, 2, 3, 4, 6, 8, 10, 12, 13)$ seconds. With the currently issued target r in MPa, the block changes Δu_j and one global slack $\varepsilon \geq 0$ in MPa, the shared receding-horizon program [23] minimizes Equation 2.23, a case of Equation (1.21).

$$\sum_{k=1}^{60} (p_k - r)^2 + 0.001 \sum_{j=1}^{10} (\Delta u_j)^2 + 1000\varepsilon^2 + 10\varepsilon. \tag{2.23}$$

The move weight is in MPa^2 per squared percentage point, and the linear slack price is 10 MPa. The hard bounds are $0 \leq u \leq 100$ and $|\Delta u_j| \leq \dot{x}_{max}(1\,\text{s})$. The current limits, taken as known exactly, impose $P_{low}/10^6 - \varepsilon \leq p_k \leq P_{high}/10^6 + \varepsilon$ across the horizon. The emitted commands are also slewed every 0.05 s at 20 percentage points/s by default, and the predictions leave this extra ramp out. No future target or event enters them. A plan is accepted only if it is optimal, of the right dimensions, finite and within every planned bound. On a rejection the PI loop, which has tracked the applied command throughout, resumes until the solver recovers. In manual mode the controller follows the requested command. A forced failure of 12 s produces no jump at the first handover.

2.6 Results

The app follows the pattern of Section 1.6.4. Manual and automatic modes are provided, the drilling rate and each constraint indicator are displayed over the weight and speed plane with the feasible region marked, and the measured and best feasible drilling rates are displayed together. Both SI and field units are supported, and the exported file holds the 20 Hz samples, including the constraint indicators, the supervisory setpoints and the binding constraint at every sample.

The eleven original scenarios appear in Table 2.4 together with the measured performance of each, and Figure 2.3 sets the hole drilled and the rate against the best feasible rate side by side. A scenario fixes only the quantities its own lesson requires, so the change it makes is visible against the baseline.

Performance values were recorded after the elapsed time shown and appear in the right-hand half of the same table. For these original cases, the three standard loop metrics are omitted and quoted below only where they carry information. The integral of absolute error on the pressure loop accumulates in pascal-seconds against a setpoint the supervisory layer is deliberately moving, so it is quoted below as a mean absolute error in psi with its run length attached, and overshoot and settling time are produced by step-response calculations with no step to work with in most of the runs. Every run starts from a steady state at the scenario's own settings, so the first seconds carry no start-up transient.

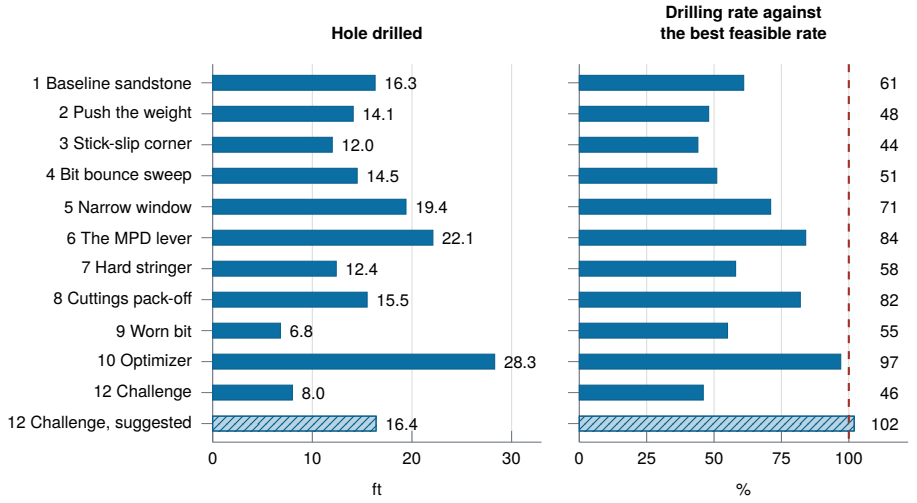

Figure 2.3 The scenarios of Table 2.4 side by side. Solid bars are runs of 900 s. The hatched bar is the suggested challenge, which stops itself at 473 s once 5.0 m is drilled. The dashed line marks the best feasible rate. Notice that only the optimizer and the suggested challenge come near that line, while the delivered challenge, automatic but held back by its conservative limits, stays far below it, and that the narrow window and the MPD lever drill more than the baseline because they carry less overbalance.

Table 2.4 The Original Scenarios, the Setting Each One Changes, and the Measured Performance of Each Run

#	Name	Setting	t (s)	Hole (ft)	ROP vs best (%)	MSE (kpsi)	Min margin (psi)	Peak S
1	Baseline sandstone	manual, 25 klbf, 120 RPM, choke 65 %, nothing scheduled	900	16.3	61	30.4	421	0.00
2	Push the weight	W_{sp} stepped 15, 25, 35, 45, 55 klbf every 90 s	900	14.1	48	82.5	354	0.08
3	Stick-slip corner	40 klbf and 60 RPM, heavy and slow	900	12.0	44	34.1	402	4.80
4	Bit bounce sweep	ω_{sp} stepped 160, 130, 110, 100, 90, 85 RPM	900	14.5	51	26.2	409	1.70
5	Narrow window	$P_{pore} = 31.0$ MPa, $P_{frac} = 33.2$ MPa, mud 1000 kg/m^3	900	19.4	71	25.6	139	0.00
6	The MPD lever	choke stepped 10, 35, 60, 85, 100 % at fixed W and ω	900	22.1	84	32.1	241	0.00
7	Hard stringer	$\Delta_{hard} \to 45$ %, $a_{abr} \to 1.6$ at $t = 150$ s	900	12.4	58	42.9	392	0.00
8	Cuttings pack-off	$Q = 0.011$ m^3/s, restored at $t = 420$ s	900	15.5	82	54.9	332	0.00
9	Worn bit	identical to 1 with $h = 6$ of 8	900	6.8	55	72.7	355	0.00
10	Optimizer	automatic from the settings of scenario 1	900	28.3	97	47.4	263	0.16
12	Challenge	automatic with every limit set conservatively and the pressure loop detuned	900	8.0	46	102.8	309	0.07
12*	Challenge, suggested	$C_{max} = 2.4$ %, $B_{max} = 0.50$, $S_{max} = 0.40$, $\Delta P_m = 0.7$ MPa, $K_{c,P} = 5 \times 10^{-5}$, $\tau_{I,P} = 4$ s	473	16.4	102	49.7	263	0.16

* the run stops itself the moment the 5.0 m interval is drilled, which is why its elapsed time is the score

2.6.1 Weight Steps and the Founder Point

The second scenario steps the weight through five values 90 s apart. At steady state the evaluator gives 26.2 ft/hr, 67.6 ft/hr, 79.9 ft/hr, 71.9 ft/hr and 55.1 ft/hr at 15 klbf, 25 klbf, 35 klbf, 45 klbf and 55 klbf, and the run reads 26.1 ft/hr, 74.0 ft/hr, 87.6 ft/hr, 73.8 ft/hr and 52.1 ft/hr at the end of each step, higher at the middle weights because 90 s is a quarter of the cuttings time constant and the annulus has not yet loaded. The peak is at the founder point in both, where $\chi = 1$ in Equation 2.2, and

past it the extra weight costs drilling rate outright. Two other quantities move during the same sequence and neither moves with the objective. Bit torque climbs monotonically from 2.09 kN m to 6.64 kN m, because Equation 2.9c is proportional to weight and carries no founder term. Specific energy goes 47.1 kpsi, 28.5 kpsi, 32.8 kpsi, 45.8 kpsi and 71.9 kpsi at steady state (47.3 kpsi, 26.1 kpsi, 29.9 kpsi, 44.7 kpsi and 76.0 kpsi in the run), cheapest at 25 klbf rather than at the founder point, so between 25 klbf and 35 klbf the drilling rate rises by 18 % and the specific energy rises by 15 % with it, in both. Whether that is a good trade depends on whether rig time or bit life is the binding cost, and specific energy is the number that distinguishes them [14].

2.6.2 Stick-Slip and Bit Bounce

The third scenario sits at 40 klbf and 60 RPM, and the bit stops and then breaks free at up to 418 RPM, seven times the speed the top drive is holding. The cycle is irregular, with a median period of 5.1 s and a tenth-to-ninetieth percentile range of 3.9 s to 13.5 s, longer than the 3.1 s natural period of the torsional mode. Severe stick-slip is flagged at 20.2 s, and the severity averages 3.39 over the last 600 s and peaks at 4.80 against a limit of 0.40. The cost shows up in two places at once. The run makes 12.0 ft of hole in 900 s against 16.3 ft for the baseline, and the top drive spends 1.06×10^6 N m of total variation on its torque command against 0.12 N m in the baseline, a factor of 8.9×10^6. A torsional limit cycle is a fatigue problem for the string and a duty problem for the drive, and the drilling rate is not the largest cost. Speed cures it reliably only before it starts. Raised a minute into the run, to 150 RPM the severity settles near 1.5, at 200 RPM near 1.3, and only 250 RPM quenches it, although 200 RPM applied at the first flag, 20 s in, still does. A run started at 150 RPM never stick-slips.

> **TRY IT IN THE APP** **3 · Stick-slip corner (heavy and slow)**
>
> Load the case and press Run in Manual. Watch the Rotary speed strip, where the surface trace and the bit trace share one plot. The top drive holds its 60 RPM while the bit stops and then breaks free at up to 418 RPM, over and over at irregular intervals. Severe stick-slip is flagged at 20.2 s.

The fourth scenario sweeps the rotary speed down through the axial resonance at 160 RPM, 130 RPM, 110 RPM, 100 RPM, 90 RPM and 85 RPM. The bounce index peaks at 100 RPM exactly, where Equation 2.13 gives 0.833 against 0.216 at 120 RPM, and the application reports 0.83 there. The sweep also shows that the two vibration limits are neighbors, because at 90 RPM and 25 klbf the run leaves the bounce stripe and enters the stick-slip corner. Severe stick-slip is flagged at 398.9 s, the severity reads 1.69 at the end of the run with a peak of 1.70, and the reported constraint is stick-slip rather than bounce. Escaping a resonance by slowing down works only when the stick-slip region is elsewhere.

2.6.3 Managed Pressure Scenarios

The sixth scenario holds the weight at 30 klbf and the speed at 150 RPM and moves nothing but the choke. Opening it from 10 % to 100 % takes bottomhole pressure from 5678 psi at 90 s to 5177 psi at 430 s, equivalent circulating density from 11.10 ppg to 10.12 ppg, and the drilling rate from 69.3 ft/hr to 97.3 ft/hr. The choke alone adds 40 % of drilling rate with no other change, and the scenario separates a manipulated variable from a disturbance, because the same overbalance would conventionally have been inherited from a mud density chosen days earlier.

The fifth scenario is the case that motivated the commercial adoption of managed pressure drilling. Pore pressure in the depleted sand is 4496 psi and the rock fractures at 4815 psi, a window 319 psi wide of which the two margins take 203 psi. Cutting the mud to 8.35 ppg means the column alone reaches only 4268 psi of hydrostatic head, so the well is held inside its window by the choke rather than by the mud. The feasible band is choke positions of 77.8 % to 99.9 % open, about a fifth of the travel, with the minimum achievable pressure at a wide-open choke sitting 101 psi above pore pressure. The run holds 139 psi of minimum window margin over 900 s and makes 19.4 ft, more than

the baseline, because the reduced overbalance of the light mud adds more drilling rate through chip hold-down than the narrow window removes.

2.6.4 Cuttings Pack-Off

The eighth scenario drills at 35 klbf and 200 RPM with the pumps at 174 gpm, an annular velocity of 90 ft/min. Cuttings are made faster than they are carried away, and the case opens at its steady loading of 5.63 %, past the pack-off threshold of 5.5 %, where the friction multiplier of Equation 2.5c has begun to climb steeply. The concentration is finite rather than running away, and the mechanism that holds it is the chip hold-down coupling of Equation 2.3c, because higher pressure reduces the rock removal rate. Restoring the pumps to 476 gpm at $t = 420$ s first raises bottomhole pressure by 378 psi on the added friction of the loaded annulus and takes the concentration below the pack-off threshold within 6 s. The annulus then cleans out slowly, to 2.43 % at 900 s with the drilling rate at 81.8 ft/hr, and the concentration falls below the 2.2 % limit only at 1030 s. The reduced flow also removes choke authority, 78 psi against 580 psi at full flow, so for the first seven minutes of the run the choke can do little about the pressure.

> **TRY IT IN THE APP** **8 · Cuttings pack-off (fast hole, slow pumps)**
>
> Load the case and press Run in Manual. Read the cuttings trace on the Constraint indicators strip: the annulus opens at 5.63 % by volume, past the pack-off threshold of 5.5 %. At 420 s the pumps come back to 476 gpm. Watch the Bottomhole pressure strip rise by 378 psi at once, then the annulus clean out slowly, to 2.43 % at 900 s.

2.6.5 Bit Wear Effects

The ninth scenario is scenario 1 with $h = 6$ of 8. It makes 6.8 ft in 900 s against 16.3 ft for a sharp bit, at a specific energy of 72.7 kpsi against 30.4 kpsi. The bit spends more than twice as much energy per unit volume of rock to make hole at 42 % of the rate, which is the signature that separates a dull bit from hard rock and the criterion that triggers a trip for a new bit [14].

2.6.6 Supervisory Optimization

The tenth scenario takes the same well and the same opening settings as the first and switches the supervisory layer on, and Figure 2.4 is that run frozen near its end. Over 900 s it makes 28.3 ft of hole against 16.3 ft, an improvement of 73 %, and finishes at 97 % of the best feasible rate against 61 %. What it does to get there is visible in three separate movements. The rotary setpoint runs up to the top of its range at 250 RPM by 52 s, the weight setpoint reaches the founder point at 35.0 klbf by 60 s and is later pulled back by the cuttings override, ending at 28.3 klbf, and the pressure setpoint walks down from 5352 psi to 5193 psi by 110 s, where the choke reaches 97.5 % open and Equation 2.21 stops it. The floor then rises with the cuttings loading and carries the setpoint back to 5225 psi. The binding constraint at the end is cuttings and the concentration reads 2.19 % against a limit of 2.2 %. The scenario demonstrates that a constrained optimum is a point on a boundary and that a supervisory controller stopping short of the boundary gives up drilling rate. The supervisory layer of Volume 9, Chapter 8 meets the same structure, with a lowest-cost point on a constraint that moves with the weather.

Two features of the trends require interpretation. From 30 s to 427 s the measured drilling rate sits above the best-feasible rate, peaking at 140 ft/hr, which is not an inconsistency, because the best-feasible rate is a steady-state quantity and the annulus takes about five minutes to load up (the concentration first passes its limit at 297 s), so the well genuinely drills faster than any point that will be feasible once it has. The second feature is hunting along the boundary. The binding indicator is the cuttings concentration, which answers a weight move only through the cuttings inventory, so the override acts on a measurement that lags its own action by minutes and the result is a slow cycle along the boundary rather than a direct arrival at it. After 300 s the weight setpoint moves between

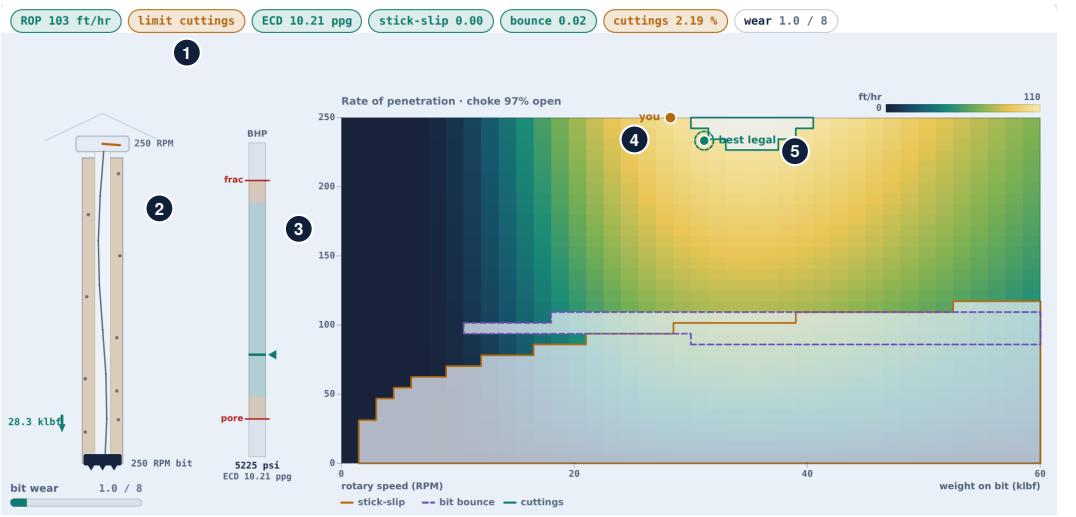

Figure 2.4 The supervisory optimization run of scenario 10 at $t = 891$ s, cropped to the plant view. (1) The chip that names the binding constraint, cuttings. (2) The drill string, with the weight and the surface and bit speeds. (3) The bottomhole pressure gauge between its fracture and pore marks. (4) The operating point on the drilling-rate map, which is evaluated at the choke position the optimizer has reached, 97 % open. (5) The best feasible point on the cuttings boundary. The operating point and the best feasible point have almost converged at 102.7 ft/hr against 106.2 ft/hr with cuttings as the binding constraint.

22.6 klbf and 34.8 klbf and the concentration between 2.08 % and 2.27 %, above its limit 42 % of the time. Every constraint controller acting on a lagging measurement behaves this way, and the alternative, backing off far enough that the lag never matters, is the conservatism the run quantifies.

2.6.7 The Graded Challenge

The challenge asks for 5.0 m of new hole within 900 s while the pressure loop holds bottomhole pressure within 17 psi of a setpoint the optimizer is moving and no constraint flag ever sets. Both conditions are checked at every physics step until the run stops itself, so a single excursion fails the attempt even after it clears. The challenge is delivered with two independent problems. Every constraint limit has been set well inside where it needs to be, with $C_{max} = 0.8\,\%$, $B_{max} = 0.20$, $S_{max} = 0.30$ and $\Delta P_m = 2.0$ MPa on a window only 1088 psi wide, and the pressure loop has been detuned to $K_{c,P} = 8 \times 10^{-6}\,\%/\text{Pa}$ and $\tau_{I,P} = 30$ s.

Only one of those is a control problem. The limits are an engineering choice, and the cuttings limit in particular has been set below the 1.49 % the default operating point already carries, so the cuttings flag is set from the first step and the override takes the weight setpoint down to its floor at $W_t = 5.0$ klbf within 60 s. The annulus then drains, the climb resumes near 300 s, and the weight setpoint cycles between 5.0 klbf and 26.9 klbf for the rest of the run. As delivered the run makes 8.0 ft (2.44 m) in 900 s with the cuttings and bounce flags set, ends at a specific energy of 102.8 kpsi and reaches a worst pressure error of 49.1 psi, so it fails. The detuned loop is the tuning problem, and its cost appears against a moving setpoint. A loop tracking a ramp settles with a standing offset of $\dot{P}_{sp}/K_v$, and the delivered tuning gives $K_v = 0.0107/\text{s}$ against 0.5/s for the default, so the predicted offset falls from 136 psi to 2.9 psi when the gain is multiplied by 6.25 and the integral time divided by 7.5. The two problems cost different things. Loosening only the limits drills the interval in 484 s with a worst pressure error of 94.4 psi, and retuning only the loop leaves 2.44 m at 900 s. Applying the suggested settings, which set the cuttings limit to 2.4 % because the concentration reaches 2.32 %

Table 2.5 Full 900 s Choke Comparison. Entries Within Each Cell Are Shipped PI / MPC / Fixed SIMC PI

Case	Pressure IAE (MPa s)	Outside known band (s)	Command travel (percentage points)
Pore rise	9.036 / 13.947 / 8.899	4.825 / 6.280 / 4.840	184.876 / 189.716 / 159.353
Fracture fall	5.762 / 32.289 / 5.091	44.080 / 42.945 / 43.830	159.903 / 193.732 / 122.204
Rate push	3.021 / 7.374 / 2.679	0.000 / 0.000 / 0.000	118.966 / 42.331 / 88.477
Challenge, shipped	85.065 / 6.959 / 2.883	0.000 / 0.000 / 0.000	26.084 / 42.416 / 63.738
Challenge, suggested	2.663 / 7.163 / 2.353	0.000 / 0.000 / 0.000	105.777 / 36.999 / 78.082
Shifting window	6.984 / 24.721 / 6.754	25.600 / 63.135 / 25.345	230.928 / 224.508 / 181.592
Gain/lag mismatch	13.153 / 27.834 / 17.801	34.255 / 46.500 / 36.990	355.727 / 219.467 / 378.930
Delay mismatch	83.200 / 33.084 / 82.733	78.895 / 62.530 / 75.830	2286.034 / 275.660 / 1935.749

before the interval is drilled, completes the interval in 473 s with no flag set, a worst pressure error of 4.9 psi, 12.1 psi inside the band, and a mean absolute error of 0.80 psi, against 49.1 psi and 13.7 psi for the delivered tuning, a peak stick-slip of 0.16 and 263 psi of minimum window margin. With the default 2.2 % cuttings limit the same settings set the cuttings flag before the interval is complete. The target requires operating at a limit rather than away from it, which is reasonable for a pressure limit measured continuously with a lag of a few seconds and not for a stick-slip limit measured over a 20 s window.

2.6.8 Pressure-Controller Comparison

Table 2.5 compares three pressure controllers on eight cases. Every run lasts 900 s at seed 20260824, with the same instruments and settings for the three controllers in each case and no MPC tuning per case. The third controller is a fixed PI loop derived from a model by the SIMC half rule [24]: $\tau = 2.75$ s, $\theta = 0.75$ s and a chosen $\lambda = 1$ s give $K_c = \tau/[40000(\lambda + \theta)] = 3.92857 \times 10^{-5}$ %/Pa and $\tau_I = 2.75$ s. No tuning search was used.

Scenario 11, the shifting window, raises the pore pressure by 0.9 MPa from 60 to 120 s and lowers the fracture pressure by 1 MPa from 180 to 240 s, with no notice to either controller, and the narrowest band is 0.4 MPa wide. The MPC loses on IAE and on time outside the band (Table 2.5). The supervisor feeds back on the pressure each controller achieves, so the two runs receive different targets (Figure 2.5). At 235 s the PI pressure and target are 32.072 and 32.070 MPa, against 32.110 and 32.310 MPa under the MPC. The upper margin there is 32.100 MPa, which is not the actual fracture pressure of 32.8 MPa, and neither run crosses an actual formation edge.

On the rate push and the suggested challenge the MPC trades a higher IAE for less choke travel. It loses under gain and lag mismatch, with an authority of 7 MPa, a pressure lag of 4 s and a valve lag of 5 s. With a transmitter delay of 2 s it wins on IAE, 33.084 against 83.200 MPa s. The fixed SIMC loop has a lower IAE than the MPC in all six nominal cases. The penalties, the slew limit and the supervisor feedback interact, so these data do not identify their separate effects.

> ✦ **DID YOU KNOW? A Two-Second Delay Keeps the Choke Moving**
>
> From the simulator (Table 2.5). With a transmitter delay of 2 s, the shipped PI loop moves the choke command through 2286.034 percentage points in 900 s and the fixed SIMC loop through 1935.749, while the MPC moves it through 275.660. Without the delay, on the rate push, the same PI loop travels 118.966.

The pore-rise case sets the fracture pressure to 34.5 MPa and raises the pore pressure by 1.6 MPa at 60 s. The pressure then stays below the actual pore pressure for 3.280 s under PI and 3.255 s under MPC, while 0.04922 and 0.04885 m of hole are drilled. The fracture-fall case uses a mud of 950 kg/m^3 and a pore pressure of 29.8 MPa and lowers the fracture pressure by 2 MPa at 15 s. The MPC shortens the time above the actual fracture pressure from 20.700 s to 3.040 s and the hole drilled meanwhile from 0.16336 m to 0.02318 m, while its IAE is worse and the peak excess rises from

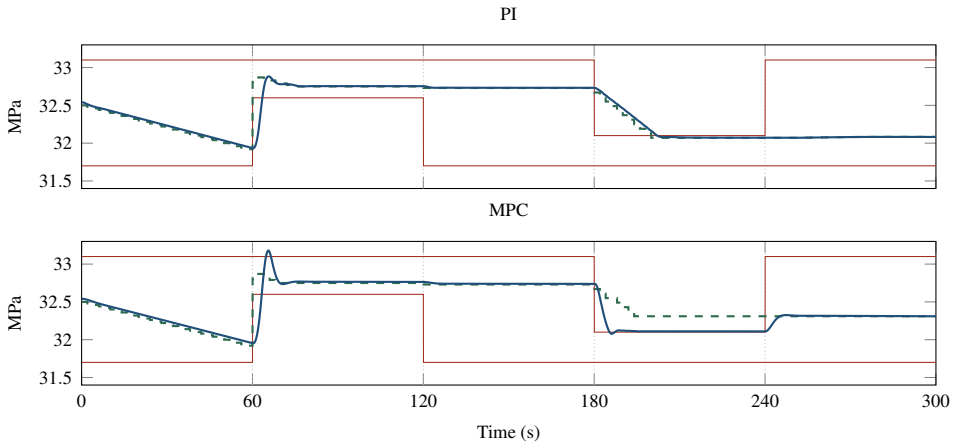

Figure 2.5 Bottomhole pressure in the shifting-window scenario over 300 s, under PI (top) and under MPC (bottom). The thick solid line is the pressure, the thick dashed line the current issued target and the thin solid lines the known band. Dotted vertical lines mark the four window events, and both panels share their axes. Under MPC the tracking error and the band violation remain.

579.111 to 588.859 kPa. In scenario 11 the cuttings limit, which the comparison does not score, is broken for 410.17 s under PI and 410.59 s under MPC, over 3.84795 m in both. In the shipped challenge the bounce limit is broken for 33.30 and 33.15 s, over 0.14840 and 0.14814 m. With the suggested settings all three controllers pass the challenge, reaching 5 m at 473.25, 473.60 and 473.22 s, and with the shipped limits all three fail. The grade concerns the run up to 5 m, while the metrics of the table continue to 900 s.

IAE and travel are computed from the 20 Hz samples, and the limit times and depths on the 5 ms physics grid. Halving that step for all 24 full runs leaves the rankings and grades unchanged. Every MPC run makes 900 solves with no rejection. In a browser the mean times of the solver, of a complete decision and of a complete control scan are 227, 373 and 782 μs at speed 1 and 107, 184 and 570 μs at speed 20. All control work, including the loops kept and the ROP search, takes 0.121 and 1.251% of a core at those two speeds.

2.7 Technology in Practice

Figure 2.6 draws the three regulatory loops and the supervisory layer of Section 2.5 with the hardware a rig installs.

What Measures It

A rotary rig hoists the drill string, turns it, presses the bit down on the rock and circulates mud down the inside of the string, through the bit and up the annulus with the cuttings [6]. The weight on bit is inferred from the pull on the drilling line. A hydraulic load cell turns a force into a liquid pressure that an ordinary pressure transmitter reports as weight [25], and the 1952 feed control read the line tension the same way, on a gauge calibrated as a weight indicator [12]. The return flow is usually metered with a Coriolis mass flowmeter, which needs a continuous flow and a closed system that can take its pressure drop of 3 to 5 psi [18]. Bottomhole pressure comes from a sensor near the bit. The chapter reads weight and surface speed exactly, with no noise and no bias.

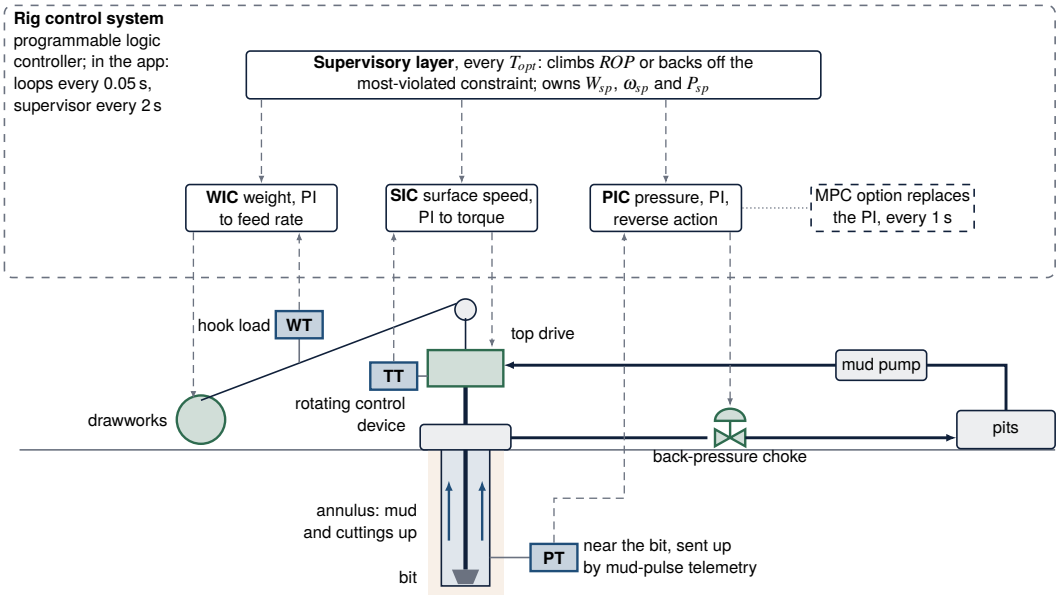

Figure 2.6 The loops of the drilling simulator with their hardware. Lettered boxes are sensing devices: the downhole pressure sensor (PT), the hook load cell on the drilling line (WT) and the top-drive speed sensor (TT). The drawworks drum, the top drive and the back-pressure choke are the final elements, and dashed arrows are signals. Notice that the pressure reading travels up through the mud before any controller sees it, and that the supervisory layer moves only setpoints. The vibration and cuttings indicators the supervisor reads and the setpoint rate limits are left out. Not to scale, and not a piping diagram.

> **TECHNOLOGY** **Sensor: Bottomhole Pressure Sent Up the Mud**
>
> A pressure sensor near the bit sends its reading to the rig through the mud itself. A valve downhole briefly restricts the flow, and the pressure pulses travel up the drill string to the standpipe pressure sensors [6]. Positive- and negative-pulse tools send 0.5 to 3.0 bits per second and continuous-wave tools up to 18, and the rate falls with depth to 1.5 to 3.0 bits per second at 10 000 m. Aerated mud is more compressible and carries the pulses worse. Most such tools work below 150 °C, and hotter holes make them unstable [6]. The chapter reads bottomhole pressure exactly, through a dead time θ_p that is zero by default.

What Moves It

The drawworks pays out drilling line and so sets the weight, the top drive turns the string, and the choke on the returns sets the back-pressure. In the chapter the feed is clamped at $v_{max} = 0.030\,\mathrm{m/s}$ and the top-drive torque at $T_{max} = 45\,\mathrm{kN\,m}$, and the choke is the only final element with a lag of its own in the model.

> **TECHNOLOGY** **Final Element: The Choke Manifold and Its Seal**
>
> A rotating control device seals the annulus around the turning pipe and sends the returns to the managed pressure choke manifold, or to the main flow line for conventional drilling [18]. The manifold controls the well pressure, not the flow rate, and when an influx is detected the choke is moved automatically to raise the surface back-pressure [18]. A back-pressure pump often accompanies the choke. A test model has to reproduce choke plugging and choke swapping [17], and a control system has to handle a blocked choke [19]. The chapter's choke has a 2 s stroke lag, a slew limit of 20 %/s and a linear characteristic.

Where the Controller Runs

Programmable logic controllers serve as local controllers within a supervisory scheme, with programmed functions that include logic, timing and three-mode PID control [26]. Supervisory and lower-level loops of such systems run with cycle times from milliseconds to minutes [26]. The app runs its three PI loops every 0.05 s, the choke MPC every 1 s and the supervisory layer every $T_{opt} = 2.0\,\text{s}$, all inside that range. One field stick-slip prevention system is a PI-type speed controller tuned to damp the torsional oscillation, with all of its software in a standard programmable logic controller [15]. The chapter's structure maps onto three PID blocks, one with reverse action for the choke, under a supervisory sequence that owns their setpoints.

> **IN PRACTICE** **The Safe Drilling Margin Is a Rule**
>
> In United States federal offshore waters, 30 CFR 250.414 asks each drilling prognosis for a planned safe drilling margin between the estimated pore pressure and the lesser of the estimated fracture gradient and the casing shoe pressure integrity test. The equivalent downhole mud weight must exceed the pore pressure and, unless a justified alternative is approved, sit at least 0.5 pound per gallon below the lower of those two limits [27]. Under 30 CFR 250.427 the margin must be kept while drilling. When it cannot be, the operator suspends drilling and submits remedial actions, or notifies the regulator and acts under API Bulletin 92L [28]. The chapter's window margin ΔP_m, 0.7 MPa at each end by default, plays this part.

What Is Modernizing It

The first development is high-speed telemetry. Wired drill pipe carries data at 57,000 bits per second and can send commands down to the tools for closed-loop control, although it is still expensive and a single failed connection stops the link [6]. With it, rate of penetration and bottomhole pressure have been combined in one controller for a managed pressure well [20]. The second is open standards. A 2013 review expected open standards that enable automation to reshape the industry's business model, as they did for industrial automation in the 1990s, and found collaboration on an open environment for drilling data [10]. The third is testing in simulation. A control system for automated managed pressure drilling should pass representative test scenarios in simulation before it is used in the field [17], which is the role the app of this chapter plays on a smaller scale.

2.8 Limitations

The predictor of the choke MPC omits the evolution of the cuttings loading, the wear and future supervisory actions. Unknown formation limits and mismatch in the parameters or in the sensor delay remain limitations. A priced operating band cannot guarantee well control, and the plant models no influx, gas migration or fracture flow.

> **OPEN CHALLENGE** **Seeing a Kick in Time**
>
> A kick is formation fluid that enters the well when the bottomhole pressure falls below the formation pressure, and a late detection can lead to an uncontrolled blowout [18]. Managed pressure drilling usually watches the return flow with a Coriolis mass flowmeter, but the meter reads poorly until its line is full or when the returning mud carries gas, which a 2020 review names as a need for further research. The same review finds no kick circulation method that suits every drilling operation and no single rule for the response [18]. The chapter's plant has no influx at all. Progress would be a detector and a response that hold across muds, depths and rigs.

Section 2.3 lists the modeling assumptions and they are not restated here. Four further shortcomings are properties of the implementation.

The stick-slip boundary on the map is the linear stability limit G_0 of the model at the default speed-loop tuning, and it describes steady rotation only. Just above it a steady start can stay quiet for half an hour or more before the cycle appears, so the map is conservative for short runs there, and a

cycle already running survives far below it, as quantified in Section 2.4. The map keeps G_0 when the speed-loop gain is changed, although the onset moves with it. Inside the band above the onset the displayed severity understates the developed cycle, and feasibility there is decided on the unramped amplitude. Marking the region where a running cycle persists would show the hysteresis on the map as well.

The hill climber of Equation 2.20 has no memory of why it backed off. The back-off is floored at the cutting threshold $W_t = 5.0\,\text{klbf}$ and the weight probe widens until it finds a gradient, so the climber cannot stall in the region where f_W is identically zero. A limit set below the loading of the operating point it starts from still produces a cycle rather than a steady compromise. In the challenge as delivered the weight setpoint is at its floor within 60 s, the climb resumes near 300 s, and from then on the setpoint swings between 5.0 klbf and 26.9 klbf while the cuttings loading swings between 0.70 % and 0.92 % around its 0.8 % limit. The cycle comes from climbing at a fixed rate against a measurement that lags by minutes, the mechanism behind the hunting in the optimizer scenario.

Bit bounce is detected and does not act. The application does not show what a bouncing bit does to a bottomhole assembly, only that the operating point is one at which bounce would happen, and the constraint is therefore an assertion rather than a demonstration.

Flow rate is a parameter rather than a manipulated variable. Section 2.4 shows that it has an interior optimum near 491 gpm and that moving it 3.3 % does more for the achievable drilling rate than the entire upper half of the choke travel, so the most valuable extension is to promote it to a fourth degree of freedom and let the optimizer find that optimum too.

The documented model also supports control methods beyond the implemented PI loops, choke MPC and gradient search. Model predictive control over Equations 2.1 to 2.14 would coordinate weight, rotary speed and choke position in one constrained optimization, treat the pressure window and the cuttings ceiling as explicit constraints rather than override triggers, and act ahead of the 340 s cuttings lag instead of waiting inside the hold band of Equation 2.19. Moving horizon estimation or Kalman filtering would reconstruct the cuttings concentration and the downhole pressure from surface measurements, the soft-sensing role that simplified hydraulics models already fill on managed pressure wells [29] and that the downhole card fills for the pump of Chapter 3. Combined rate-of-penetration and pressure control of this form has been demonstrated on a test well with high-speed telemetry [20]. The APMonitor framework implements the nonlinear estimation and predictive control formulations these extensions require [30], and the Gekko Optimization Suite provides the same capabilities in Python [31].

> ⚠ **OPEN CHALLENGE** **Pressure Waves the Controller Does Not See**
>
> Pressure controllers for managed pressure drilling usually rest on simplified models that ignore pressure waves in part or in full [17]. Those waves act over tens of seconds, or minutes when gas has entered the well, and if the design leaves them out they can degrade the control and even destabilize it. In field data a pump pressure transient reached the choke about 4 s later, at the speed of sound in the mud [17]. The chapter's hydraulics are quasi-steady, with one lag. Progress would be controllers designed on simple models and proven on models that keep the waves before they reach a rig.

2.9 Conclusion

This chapter set out a thirteen-state drilling and managed pressure simulator, with its two control layers and its verification, closely enough for the whole of it to be built again from the text. The rate of penetration is a five-factor product whose weight term founders at a declared point, and every constraint on it increases with the same two factors that increase the objective, so the optimum is a point on a boundary. At the default settings that peak is 111.2 ft/hr and infeasible, the best feasible point makes 106.7 ft/hr, and the operating point the first scenario supplies makes 67.6 ft/hr. The cuttings constraint turns out to be an explicit ceiling on the objective, $C_{max} Q \eta(\omega)/A_{bit}$, which is set by rotary speed and pump rate and carries no dependence at all on weight on bit, the variable adjusted

first in practice. It pins the achievable rate at 106.7 ft/hr for every choke position above 52.1 % open, while a 3.3 % increase in pump rate lifts it above the unconstrained peak. The implemented supervisory layer climbs the objective gradient above three proportional-integral loops and hands the offending setpoint to the most-violated constraint, and it drills 28.3 ft in 900 s against 16.3 ft for the fixed manual point, finishing at 97 % of the feasible ceiling with cuttings as the binding constraint. The stick-slip severity on the map is a correlation with a median error of 9.2 % away from the onset, feasibility is decided on its unramped amplitude, and the onset itself is the linear stability limit of the string at $G = 1.092$, where a disturbance grows slowly enough that a steady start can stay quiet for many minutes. The optional choke MPC reduces the choke travel on the rate push from 118.966 to 42.331 percentage points while raising the IAE from 3.021 to 7.374 MPa s.

A fuller app would feed bit bounce back into the depth of cut, promote flow rate to a manipulated variable, add lateral vibration modes, and replace the hill climber with a model predictive formulation for comparison. As it stands, the app locates an economic optimum on a constraint boundary and measures the cost of operating away from one.

KEY POINTS

- Pushing harder and turning faster speed up the drilling, but every limit grows with them, so the best point sits on the edge of the safe region.
- At the default settings the unconstrained peak of 111.2 ft/hr is infeasible and the best feasible point makes 106.7 ft/hr.
- The cuttings constraint is a ceiling on the drilling rate set by rotary speed and pump rate, not by weight on bit.
- The supervisory layer drills 28.3 ft in 900 s against 16.3 ft for the fixed manual point.

Where else this applies. Volume 9, Chapter 8 moves a supervisory setpoint along a constraint that shifts with the weather.

References

[1] T. M. Warren, “Penetration-rate performance of roller-cone bits,” *SPE Drilling Engineering*, vol. 2, no. 1, pp. 9–18, 1987.

[2] J. D. Jansen and L. van den Steen, “Active damping of self-excited torsional vibrations in oil well drillstrings,” *Journal of Sound and Vibration*, vol. 179, no. 4, pp. 647–668, 1995.

[3] T. Richard, C. Germay, and E. Detournay, “A simplified model to explore the root cause of stick-slip vibrations in drilling systems with drag bits,” *Journal of Sound and Vibration*, vol. 305, no. 3, pp. 432–456, 2007.

[4] A. T. Bourgoyne, Jr. and F. S. Young, Jr., “A multiple regression approach to optimal drilling and abnormal pressure detection,” *Society of Petroleum Engineers Journal*, vol. 14, no. 4, pp. 371–384, 1974.

[5] T. I. Larsen, A. A. Pilehvari, and J. J. Azar, “Development of a new cuttings-transport model for high-angle wellbores including horizontal wells,” *SPE Drilling & Completion*, vol. 12, no. 2, pp. 129–135, 1997.

[6] T. Ma, P. Chen, and J. Zhao, “Overview on vertical and directional drilling technologies for the exploration and exploitation of deep petroleum resources,” *Geomechanics and Geophysics for Geo-Energy and Geo-Resources*, vol. 2, no. 4, pp. 365–395, 2016.

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

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

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

[10] J. D. Macpherson, J. P. de Wardt, F. Florence, C. D. Chapman, M. Zamora, M. L. Laing, and F. P. Iversen, “Drilling-systems automation: Current state, initiatives, and potential impact,” *SPE Drilling & Completion*, vol. 28, no. 4, pp. 296–308, 2013.

[11] W. A. Loomis, “Dummy driller,” U.S. Patent 1 779 480, Oct. 28, 1930.

[12] J. G. Abraham, “Drilling control,” U.S. Patent 2 759 702, Aug. 21, 1956.

[13] R. Teale, “The concept of specific energy in rock drilling,” *International Journal of Rock Mechanics and Mining Sciences & Geomechanics Abstracts*, vol. 2, no. 1, pp. 57–73, 1965.

[14] F. E. Dupriest and W. L. Koederitz, “Maximizing drill rates with real-time surveillance of mechanical specific energy,” in *SPE/IADC Drilling Conference, Amsterdam, The Netherlands*. SPE, 2005.

[15] Å. Kyllingstad and P. J. Nessjøen, “A new stick-slip prevention system,” in *SPE/IADC Drilling Conference and Exhibition, Amsterdam, The Netherlands*. SPE, 2009.

[16] J. Thorogood, W. Aldred, F. Florence, and F. Iversen, “Drilling automation: Technologies, terminology, and parallels with other industries,” *SPE Drilling & Completion*, vol. 25, no. 4, pp. 419–425, 2010.

[17] S. Naderi Lordejani, M. H. Abbasi, N. Velmurugan, C. Berg, J. Å. Stakvik, B. Besselink, L. Iapichino, F. Di Meglio, W. H. A. Schilders, and N. van de Wouw, “Modeling and numerical implementation of managed-pressure-drilling systems for the assessment of pressure-control systems,” *SPE Drilling & Completion*, vol. 35, no. 4, pp. 598–619, 2020.

[18] M. M. Huque, S. Imtiaz, A. Rahman, and M. Hossain, “Kick detection and remedial action in managed pressure drilling: a review,” *SN Applied Sciences*, vol. 2, no. 7, p. 1178, 2020.

[19] J.-M. Godhavn, “Control requirements for automatic managed pressure drilling system,” *SPE Drilling & Completion*, vol. 25, no. 3, pp. 336–345, 2010.

[20] R. Asgharzadeh Shishavan, C. Hubbell, H. Perez, J. Hedengren, and D. Pixton, “Combined rate of penetration and pressure regulation for drilling optimization by use of high-speed telemetry,” *SPE Drilling & Completion*, vol. 30, no. 1, pp. 17–26, 2015.

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

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

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

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

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

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

[27] U.S. Bureau of Safety and Environmental Enforcement, “30 CFR 250.414, what must my drilling prognosis include?” Code of Federal Regulations, Title 30, Part 250, Subpart D, edition of July 1, 2025, govinfo.gov, 2025.

[28] ——, “30 CFR 250.427, what are the requirements for pressure integrity tests?” Code of Federal Regulations, Title 30, Part 250, Subpart D, edition of July 1, 2025, govinfo.gov, 2025.

[29] G.-O. Kaasa, Ø. N. Stamnes, O. M. Aamo, and L. S. Imsland, “Simplified hydraulics model used for intelligent estimation of downhole pressure for a managed-pressure-drilling control system,” *SPE Drilling & Completion*, vol. 27, no. 1, pp. 127–138, 2012.

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

[31] 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.

2.10 Exercises

2.1 Concept check · 2 min

On a climbing execution in automatic mode, the supervisory layer also moves the bottomhole pressure setpoint. Which way does it move it, and what stops it?

(a) Up toward the fracture margin, until the high-pressure entry of v becomes the largest
(b) Nowhere, because the pressure setpoint belongs to the pressure loop and not to the supervisor
(c) Down toward the pore pressure itself, until an influx is flagged at the bit
(d) Down toward a floor set by the pore margin or by the wide-open choke, whichever is higher

2.2 Concept check · 2 min

Equation (2.7a) is a balance on the cuttings carried in the annulus. At steady state, why can no setting of the weight on bit raise the drilling rate above the cuttings ceiling at a given rotary speed and pump rate?

(a) The steady balance gives $C = ROP A_{bit}/(Q\eta)$, and η holds the speed and the flow but no weight
(b) The weight factor founders at W_f, so any weight above the founder point lowers the rate
(c) Added weight raises bottomhole pressure, and chip hold-down removes the gain it makes
(d) The weight loop saturates at its maximum feed rate before the ceiling can be reached

2.3 Concept check · 1 min

Which description of the two vibration constraints on the weight and rotary speed map is right?

(a) Both forbid a corner at high weight and low speed, so a higher rotary speed escapes both
(b) Stick-slip forbids a corner at high weight and low speed, bit bounce a band of speeds around a resonance
(c) Bit bounce forbids a corner at high weight, and stick-slip a band of speeds around 100 RPM
(d) Both forbid bands of rotary speed, whose width grows in proportion to the weight

2.4 Theory and calculation · 7 min

The default well has a 0.2159 m bit, 0.1270 m pipe, a cuttings transport length $L_c = 200\,\mathrm{m}$ and a pump rate $Q = 0.030\,\mathrm{m^3/s}$. Transport uses $\eta_{max} = 0.90$, $v_0 = 0.80\,\mathrm{m/s}$, $a_r = 0.60$ and $\omega_r = 8.0\,\mathrm{rad/s}$, and the cuttings limit is $C_{max} = 2.2\,\%$. Take $\bar{\omega}_b$ equal to the surface speed.

(a) Find the annular area and the annular velocity. *Checkpoint:* $0.023\,94\,\mathrm{m^2}$ and 1.253 m/s.

(b) Find the transport efficiency and the cuttings ceiling on the drilling rate at 120 RPM, in ft/hr. *Checkpoint:* 0.469 and 99.9 ft/hr.

(c) Find the time constant of the cuttings inventory at 120 RPM, and the ceiling at 250 RPM. *Checkpoint:* 340 s and 106.7 ft/hr.

2.5 Theory and calculation · 6 min

The choke adds $P_{bp,max}(1 - x/100)(Q/Q_{ref})^2$ with $P_{bp,max} = 4.0\,\mathrm{MPa}$ and $Q_{ref} = 0.030\,\mathrm{m^3/s}$. The supervisor walks the pressure setpoint at $\dot{P}_{sp} = 10\,\mathrm{kPa/s}$.

(a) Find the process gain $\partial P_b/\partial x$ in psi per percent at the sizing flow, and the whole authority of the choke at the $0.011\,\mathrm{m^3/s}$ of the pack-off scenario. *Checkpoint:* −5.80 psi per percent; 78 psi.

(b) Find the velocity constant $K_v = K_c|K_p|/\tau_I$ and the standing offset $\dot{P}_{sp}/K_v$ behind the ramp for the default tuning ($K_c = 5 \times 10^{-5}\,\%/\mathrm{Pa}$, $\tau_I = 4\,\mathrm{s}$) and for the challenge as delivered ($K_c = 8 \times 10^{-6}\,\%/\mathrm{Pa}$, $\tau_I = 30\,\mathrm{s}$). Compare both with the 17 psi band of the challenge. *Checkpoint:* 0.5/s and 2.9 psi; 0.0107/s and 136 psi.

2.6 Theory and calculation · 7 min

The string uses $R_b = 0.10\,\mathrm{m}$, $\mu_s = 0.80$, $\mu_d = 0.25$, $\omega_c = 2.80\,\mathrm{rad/s}$ and $c = c_b = 30\,\mathrm{N\,m\,s/rad}$, with the onset $G_0 = 1.092$. The axial mode has $f_n = 5.00\,\mathrm{Hz}$, $\zeta_a = 0.060$, $n_l = 3$, $\varepsilon_0 = 0.100$ and $W_{ref} = 111.2\,\mathrm{kN}$. One klbf is 4.448 kN.

(a) Evaluate G of Equation (2.10) at the scenario-3 point, 40 klbf and 60 RPM. *Checkpoint:* $G = 6.18$.

(b) Find the rotary speed above which steady rotation is stable at 40 klbf. *Checkpoint:* 106.3 RPM.

(c) Find the rotary speed of the axial resonance and the bounce index there at 25 klbf, and compare it with the limit $B_{max} = 0.50$. *Checkpoint:* 100 RPM and 0.833.

2.7 App challenge · 8 min

Load scenario 6, The MPD lever. Weight and rotary speed are held at 30 klbf and 150 RPM while the choke opens in steps from 10 % to wide open. Change no setting, use the default seed, run for at least 430 s at speed 20 and export the CSV file.

(a) Read `bhp_psi` and `rop_ftphr` in the rows at 90 and 430 s, and find the ratio of the two drilling rates. *Checkpoint:* 5678 psi and 69.3 ft/hr; 5177 psi and 97.3 ft/hr; ratio 1.404.

(b) Predict the ratio from chip hold-down alone, with $\lambda = 0.12/\text{MPa}$. *Checkpoint:* 1.513.

(c) Read `cuttings_frac` and `bit_wear` in the same rows and account for the difference with Equations (2.3b) and (2.3d) ($h_c = 6$, $C_{50} = 0.050$). *Checkpoint:* Factors 0.968 and 0.959; product 1.404.

Accept the pressures within 1 psi and the rates within 0.1 ft/hr.

2.8 App challenge · 8 min

Load scenario 10, Optimizer against a fixed manual point, which starts in Auto from the settings of scenario 1. Change no setting, use the default seed, run 900 s at speed 20 and export the CSV file. Then load scenario 1, Baseline sandstone, and run it 900 s unchanged.

(a) Read the Hole drilled and ROP vs best legal tiles of both runs, and the binding constraint of scenario 10 at the end (the limit badge, or the column `opt_binding`). *Checkpoint:* 28.3 ft and 97 % against 16.3 ft and 61 %; cuttings.

(b) In the file of scenario 10, find when the rotary speed setpoint (`opt_omega_sp_radps`) reaches 250 RPM, when the weight setpoint (`opt_wob_sp_n`) first reads 35.0 klbf to one decimal, and the lowest pressure setpoint (`opt_bhp_sp_pa`) with its time. *Checkpoint:* 52 s; about 58 s; 5193 psi at 110 s.

(c) Find the interval in which `rop_mps` first exceeds `rop_best_mps`, and explain why that is not an error. *Checkpoint:* From 30 to 427 s.

2.9 App challenge · 9 min

Load the challenge, Drill 5 m in under 15 minutes. Use the default seed and speed 20, and export the CSV file of each run. The worst pressure error is the largest $|\texttt{bhp_pa} - \texttt{bhp_sp_pa}|$ in the file.

(a) Run the case as delivered until it stops itself. Read the Hole drilled tile, the worst pressure error and the flag columns that read 1 at the start. *Checkpoint:* 8.0 ft at 900 s; 49.1 psi; `flag_packoff` from 0 s and `flag_bounce` from 0.15 s.

(b) Press the reload button, press Reveal suggested tuning before Run, and run until the case stops itself. Read the elapsed time, the worst pressure error, the Peak stick-slip and Min window margin tiles, and whether any flag column reads 1. *Checkpoint:* 473 s; 4.9 psi; 0.16 and 263 psi; no flag.

(c) Press the reload button, press Reveal suggested tuning, set Cuttings loading limit back to 2.2 %, and run. Find when the first flag sets, and say whether the attempt passes. *Checkpoint:* `flag_packoff` at 297 s; it fails.

Accept the pressure errors within 0.1 psi.

2.10 Application area · 4 min

Scenario 5, Narrow pressure window, drills a depleted sand at a true vertical depth of 3000 m with a pore pressure of 31.0 MPa, a fracture pressure of 33.2 MPa and a mud of 1000 kg/m^3. The margin is 0.7 MPa at each end of the window.

(a) Find the hydrostatic head of the mud column, in psi, and compare it with the pore pressure. *Checkpoint:* 4268 psi against 4496 psi, 228 psi short.

(b) Find the width of the window and the part left after the two margins. *Checkpoint:* 319 psi, of which 116 psi is left.

(c) The pumps stop for a pipe connection. Say what happens to the bottomhole pressure, what a rig does about it, and whether the app can show the consequence.

2.11 Application area · 4 min

Scenario 9, Worn bit, is scenario 1 with the bit worn to $h = 6$ of 8. In 900 s it makes 6.8 ft against 16.3 ft for a sharp bit, and the MSE tile ends at 72.7 kpsi against 30.4 kpsi.

(a) Compare the ratio of the holes with the wear factor $f_h = \mathrm{e}^{-h/h_c}$, $h_c = 6$, and explain the difference. *Checkpoint:* 0.417 against 0.368.

(b) Compare the ratio of the specific energies with the inverse ratio of the holes, and explain the match from Equation (2.15). *Checkpoint:* 2.39 against 2.40.

(c) A driller sees the specific energy more than double at constant weight and speed. Say what else separates a dull bit from a hard stringer.

2.12 Review the agent · 5 min

An AI assistant is asked how to drill faster in scenario 1 and replies as follows. “Chip hold-down multiplies the whole drilling-rate surface. Opening the choke from 65 % to wide open lifts the unconstrained peak from 111.2 to 124.1 ft/hr, a gain of 11.6 %. The best feasible point, 106.7 ft/hr, sits on the same surface, so it rises by the same share to about 119 ft/hr. Open the choke fully and let the optimizer climb to it.” Find the flaw and say how you would check it.

CHAPTER 3

Hydraulic Rod Pump Simulator for Downhole Card Inference and Constrained Rate Control

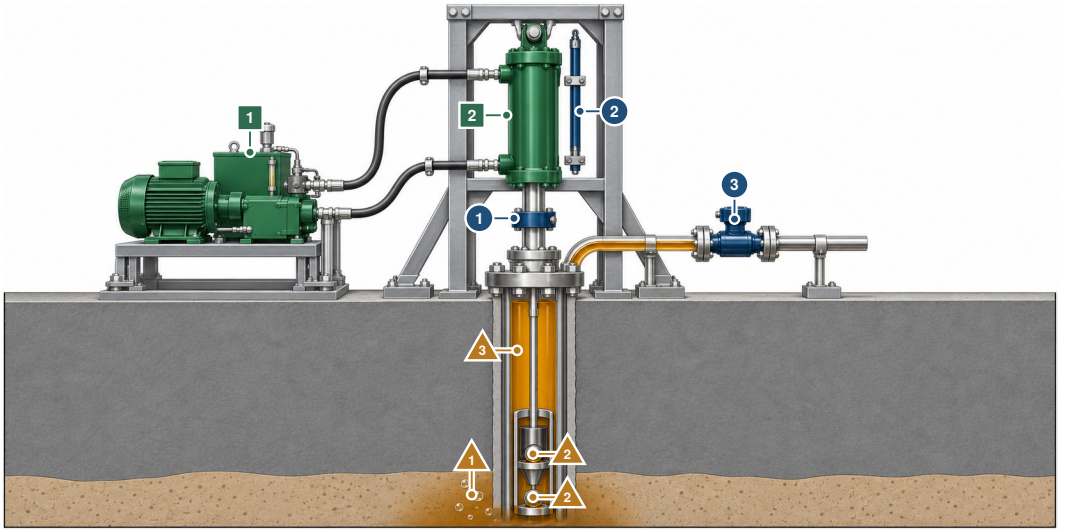

The controller moves 1 hydraulic power unit, which sets the stroke rate 2 cylinder, which sets the stroke length

The controller reads 1 load on the polished rod 2 position of the polished rod 3 liquid rate at the wellhead

What disturbs it 1 free gas at the pump intake 2 leaking valves or a worn plunger 3 thicker oil in a cold spell

The load cell and the position sensor sit at the top of the rod, far above the pump and its two ball valves, so everything the controller knows about the bottom of the well is worked out from the top. Illustration, not to scale.

IN PLAIN TERMS Steering a Pump Nobody Can See

The main idea. A long steel rod works a pump at the bottom of an oil well. The only gauge, at the top, weighs the pull on the rod, and a computer turns that pull into how full the pump is and how much oil stands over it.

The challenge. That answer is worked out, not measured, so a small error at the top grows below. Pumping too fast leaves the pump partly empty, so it hammers on each stroke without lifting much more oil.

Why it matters. In the chapter's simulator, the fastest hand setting lifts 1.6 percent more oil than the best controller but pounds the pump.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? box in Section 3.1, the two History boxes in Section 3.2, the Try it boxes in Section 3.4, the second Did you know? box in Section 3.6, and Key points in Section 3.9. No equations.
Study	50 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 3.6 with Figure 3.4, Section 3.7 and Section 3.8 with its Open challenge boxes.

The control question The only routine instrument on a rod-pumped well is a load cell nine hundred meters above the pump. How far can a controller trust the fillage and fluid level it infers from that card, and which rate and stroke lift the most oil without overloading or fatiguing the rods?

Learning objectives

- Compute the rod stretch, the wave travel time and the fluid load of the well from its geometry and pressures.
- Explain how the downhole card is reconstructed from the surface card, and what a wrong assumed damping does to it.
- Derive the pump-off rate from the reservoir inflow and the net plunger displacement.
- Tell fluid pound from gas interference on the card, and name the pairs the card cannot separate.
- Compare the timer, PI on the inferred level and the constrained optimizer on oil, energy, rod load and pound damage.

Prerequisites Sections 1.4.2 (inflow and the tubing column), 1.4.9 (inferred constraints), 1.4.8 (constraint control), 1.6.2 (velocity-form PID) and 1.6.4 (the apps).

The app https://apmonitor.com/apps/hydraulic-rod-pump/

See also Volume 5, Chapter 11 and Volume 5, Chapter 10 for control on an inferred variable, Volume 9, Chapter 8 for a supervisory grid search, Volume 9, Chapter 4 for an offset-free MPC, and Volume 4, Chapter 2 for a protective override.

3.1 Introduction

Interactive simulations are used across process control teaching for plants that a teaching laboratory cannot hold, such as distillation columns, reactors, grid dispatch, drilling and artificial lift, and others [1–3]. Most of those plants share a convenience that is easy to overlook. The controlled variable is measured. A temperature loop reads a thermocouple, a level loop reads a transmitter, and the tuning work that follows is about dynamics rather than about whether the number on the screen is the number in the vessel.

A rod-pumped oil well does not offer that convenience. The pump sits at the bottom of the tubing, the surface unit sits at the wellhead, and the two are connected by several hundred meters of steel that stretches under load and transmits longitudinal waves at about 5135 m/s. Nothing is installed at the pump on a routine well. Pump fillage, annular liquid level and pump intake pressure are all computed from a polished-rod load cell and a position sensor, through a model, and every control decision made on the well rests on those computed numbers. The two controlled variables sit three inference steps away from the only instrument on the well, as the unmeasured core temperature of Volume 5, Chapter 11 sits behind its soft sensor.

✦ DID YOU KNOW? Most Oil Wells Need a Pump

More than 70 percent of the oil wells producing in 2021 were produced by artificial lift, according to a review of that year [4]. In 1999 a national laboratory report put sucker rod pumps in approximately 90 percent of all oil wells in the United States [5].

The method that produces them is old and it is still standard. The wave-equation description of a rod string and its use to predict surface behavior were published in 1963 [6], and the inverse calculation that turns a measured surface card into a downhole card followed in 1966 [7]. Improvements to the numerical treatment and to the design procedures built on the same model [8–10], and card analysis remains the primary diagnostic on rod-pumped wells [11]. Recent work applies pattern recognition and deep learning to the classification of the resulting cards [12, 13], which is a statement about how central the card is rather than a replacement for it.

However, the literature that a student meets treats the card as a diagnostic artifact and treats the control problem separately, so the connection between the estimator and the loop closed on its output is rarely exercised. The connection is where the interesting failures are. A load cell bias of 450 N on this well is under 1 % of the peak polished-rod load and moves the inferred fluid level by 44 m. A proportional-integral controller closed on that inference drives the estimate back toward its 25 m setpoint while the true level climbs to 87 m within an hour and the oil lifted in that hour falls by 56 %.

This chapter works with a rod pump simulator in which the measurement chain, the card reconstruction, the fault ranking and the rate controller are all present at once, so that a control decision can be traced back through the estimator that produced its input. Section 3.2 gives the history of the measurement and of its automation. Section 3.3 gives the model in enough detail to rebuild it, including an algebraic pump-barrel pressure that makes a check valve integrable at a browser time step. Section 3.4 reconstructs the downhole card, infers the fluid level from it and ranks the faults, all between two strokes. Section 3.5 sets out the two loops, the four strategies and the protection layer. Section 3.6 measures the thirteen scenarios, which separate conditions sharing a single card signature, compares the strategies and reports the verification. Section 3.7 sets the hardware of a real well beside the model, Section 3.8 states the limitations and Section 3.9 draws the results together.

3.2 Background

Rod pumping is the oldest form of artificial lift still in wide use and it lifts the large majority of the world's producing oil wells by count [11]. The machine is a positive displacement pump with a very long connecting rod. A surface unit reciprocates a polished rod, a rod string transmits that motion to a plunger inside a barrel at the bottom of the tubing, and two ball check valves rectify the reciprocation into flow. On the upstroke the traveling valve in the plunger is held shut by the fluid column above it and the rods carry the fluid load, while the standing valve at the bottom of the barrel opens and the well fills the volume the plunger vacates. On the downstroke the barrel contents are compressed until they reach tubing pressure, the traveling valve opens, and the load transfers from the rods to the tubing.

Measurement of what happens at the bottom began with the surface dynamometer, which records polished-rod load against polished-rod position for one stroke. It was read by eye from its shape, and the first History box below follows how a computer came to read it instead. The shapes are ambiguous because the rod string is between the pump and the instrument. A rod string 900 m long stretches 151 mm under the fluid load on the well modeled here, which is 5 % of a three meter stroke, and it rings at its own natural frequencies while the load transfers.

The wave equation resolved the ambiguity. Treating the rod string as a damped elastic bar and solving the resulting boundary value problem downward from the surface produces the load and position at the pump from measurements taken at the wellhead [6, 7]. The calculation is linear, which is what makes it invertible, and the inversion was practical on the computers of the period because each harmonic propagates in closed form. Finite-difference treatments followed for cases where the

boundary conditions are less convenient [9], and a two-phase treatment of the fluid column addressed wells where the gas is not negligible [10].

> ⧖ HISTORY **1963 to 1967: Reading the Pump Through the Rods**
>
> For many years the polished rod dynamometer, which records rod load against position, was the principal tool for judging a rod pumped well [7]. Its card was read by eye. The reading depended on the skill of the analyst and gave only qualitative answers, and the alternative was to run an instrument to the bottom of the well, which was slow and costly [14]. In 1963 a wave-equation model of the rod string, made practical by digital computers, began to replace semi-empirical design formulas [6]. The 1966 diagnosis treated the rod string as a transmission line. Strain waves travel up it at about 16,000 feet per second and bring the pump's condition to the surface in code, which the computer decodes [7]. A patent granted in 1967 claimed downhole cards at every depth, including the pump, from data measured at the polished rod [14].

Automation of the rate followed the measurement. The second History box below follows pump-off control from fixed duty cycles to a controller that reads the downhole card. Continuous rate control through a variable speed drive is more recent and is now common, and the associated energy savings have been measured on real installations [15, 16]. Long-stroke units, in which a hydraulic cylinder or a belt replaces the crank and beam, were developed in part to reach stroke lengths a beam unit cannot and in part to make the stroke length itself adjustable [17].

> ⧖ HISTORY **1975: A Controller That Stops the Pump**
>
> A well is pumped off when the pump can take more than the wellbore holds. The barrel then fills partly, and on the downstroke the load comes off the rods suddenly when the plunger meets liquid, a pounding that can damage equipment downhole and at the surface [18]. A patent filed in 1975 named the simplest cure, a fixed duty cycle of on and off periods set from an occasional test, which fails when well conditions change. Sensing the motor's electric power was another, but that power is not the power reaching the pump. The patent took the area of the surface card as the power put into the rods, stopped the unit when it fell and restarted it after a time learned from the previous cycle [18]. By 1995 controllers sensed pump-off from the downhole card and paired its fillage with a variable-frequency drive to vary the speed instead of stopping [19].

The measurements that a controller can use remain few. Annular liquid level is determined acoustically by firing a pressure pulse down the casing and timing the echo, which is accurate and is not a continuous measurement [20]. Pump leakage can be quantified from dynamometer data with a valve test, which requires stopping the unit [21]. Separating a pumped-off well from a gassy one is still an active subject, and recent field work uses the plunger velocity computed from the same card to do it [22]. Everything a continuous controller has comes from the load cell and the position sensor.

3.3 Process Model

Well HR-118 produces a 14 °API crude at 30 % water cut from a formation whose pressure at the pump datum is 5.5 MPa. The pump is set at 900 m on 25.4 mm rods inside 73 mm tubing, and the tubing hangs inside 159.4 mm casing. A 57.15 mm plunger works in a barrel at the bottom of the tubing between a traveling valve and a standing valve. The surface unit is either a hydraulic long-stroke cylinder or a conventional crank and beam unit, and the two share every component below the polished rod. Figure 3.1 shows the arrangement and each symbol used in the equations that follow. Section 1.4.2 describes the same reservoir inflow and tubing column for a flowing well.

Table 3.1 classifies the variables. The controlled variables are both inferred from the card rather than measured, and that distinction is carried through the rest of the chapter.

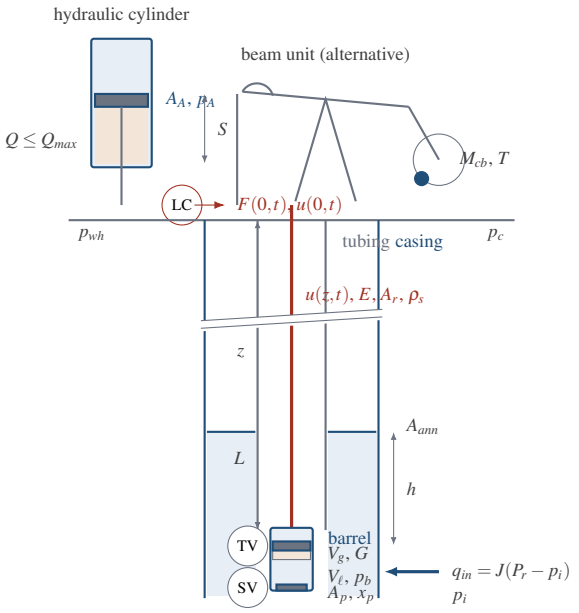

Figure 3.1 Rod pumped well HR-118 with the two alternative surface units. The load cell LC supplies the only routine measurement, at the top of a rod string whose displacement field $u(z,t)$ separates it from the pump.

3.3.1 Rod String

The rod string is a damped elastic bar in one dimension, with z measured downward from the polished rod and u the displacement field, as shown in Equation (3.1).

$$\frac{\partial^2 u}{\partial t^2} = a^2 \frac{\partial^2 u}{\partial z^2} - \nu \frac{\partial u}{\partial t}, \qquad a = \sqrt{E/\rho_s} \tag{3.1}$$

where a is the speed of sound in steel, 5135 m/s for $E = 2.07 \times 10^{11}$ Pa and $\rho_s = 7850\,\mathrm{kg/m^3}$, and ν is a damping rate in reciprocal seconds. The damping has two parts. Structural damping contributes $\nu_0 = 1.0/\mathrm{s}$, and viscous drag on the rod moving through produced fluid in the tubing annulus contributes the rest, as shown in Equation (3.2).

$$f_{drag} = \frac{2\pi\mu\, v}{\ln(r_t/r_r)} \ \text{per meter}, \qquad \nu = \nu_0 + \frac{2\pi\mu}{\rho_s A_r \ln(r_t/r_r)} \tag{3.2}$$

where r_t and r_r are the tubing and rod radii. At 800 cP the drag coefficient is 5.63 N s/m per meter of string, which over 900 m is 5.1 kN at a meter per second, and the total damping rate is 2.42/s. Viscosity therefore supplies 59 % of the damping at the design condition and more than 80 % in a cold well, which is the reason a heavy-oil well is a different control problem from a light-oil well of the same geometry.

Equation (3.1) is discretized into ten elements between eleven lumped masses. Displacements are perturbations from the statically stretched shape, so gravity is balanced by the static internal forces and the states carry only the dynamics. The interior node equations are Equation (3.3).

$$m_i \ddot{u}_i = k\,(u_{i+1} - 2u_i + u_{i-1}) - c_i \dot{u}_i, \qquad k = \frac{EA_r}{\Delta z} \tag{3.3}$$

where $\Delta z = L/10 = 90\,\mathrm{m}$, $m_i = \rho_s A_r \Delta z$ for the interior nodes with half elements plus the polished rod mass at $i = 0$ and the plunger mass at $i = 10$, and $c_i = m_i \nu_0 + \Delta z\, f_{drag}/v$. The top node carries the

Table 3.1 Variable classification for the rod pump application

Symbol	Meaning	Units	Range
Controlled variables (both inferred, not measured)			
$\hat{h}$	fluid over the pump from the card	m	0–900
F	pump fillage from the card	–	0–1
Manipulated variables			
N	stroke rate	SPM	0–12
S	stroke length (hydraulic only)	m	0.6–4.2
Measured variables			
$F(0,t)$	polished rod load	N	$0–9 \times 10^4$
$u(0,t)$	polished rod position	m	0–4.2
q_ℓ	wellhead liquid rate	m^3/s	$0–1 \times 10^{-3}$
Disturbances			
μ	oil viscosity at the pump	Pa s	0.02–5
f_g	free gas fraction at intake	–	0–0.6
ΔP_r	reservoir pressure shift	Pa	-3×10^6–1×10^6
ℓ_{TV}, ℓ_{SV}	valve leak severity	%	0–100
w	plunger and barrel wear	×	1–25
ℓ_T	tubing leak severity	%	0–100
σ_F, b_F	load cell noise and bias	N	≤2700, ±9000
States (28 in total)			
θ, ω	cycle phase and rate	rad, rad/s	–
u_i, v_i	rod node displacement and velocity, $i = 0 \ldots 10$	m, m/s	–
V_ℓ	barrel liquid inventory	m^3	–
G	barrel gas pressure-volume product	Pa m^3	–
h	true fluid over the pump	m	0–900
Q_p	cumulative liquid produced	m^3	–

drive force and the bottom node carries the pump force. The polished-rod load reported to the load cell is $F(0,t) = W_{rf} + \Delta F_0$, where $W_{rf} = 31.51\,\text{kN}$ is the buoyant weight of everything the polished rod holds up and ΔF_0 is the force the drive applies in excess of it.

3.3.2 Rod Stress

Rod fatigue is evaluated with the modified Goodman relation of API RP 11L, which makes the allowable peak stress a function of the minimum stress in the same cycle, as shown in Equation (3.4).

$$S_a = \left(\frac{T_u}{4} + 0.5625\, S_{min} \right) SF, \qquad U_G = \frac{S_{max}}{S_a} \tag{3.4}$$

where $S_{max} = PPRL/A_r$ and $S_{min} = MPRL/A_r$ are the peak and minimum polished-rod stresses over the stroke, $PPRL$ and $MPRL$ are the peak and minimum polished-rod loads, $T_u = 7.93 \times 10^8\,\text{Pa}$ is the minimum tensile strength of an API Grade D rod and SF is the service factor, unity for the sweet service modeled here. A short fast stroke and a long slow stroke of the same displacement do not carry the same U_G, because the inertial contribution to S_{max} goes as the stroke times the square of the rate while the displacement goes as their product. That asymmetry is why U_G enters the supervisory objective in Section 3.5 as a continuous cost rather than only as a limit.

3.3.3 Surface Units

Both units drive the polished rod through a finite stiffness, as shown in Equation (3.5).

$$\Delta F_0 = K_s\,(x_{cmd}(\theta) - u_0) + C_s\,(\dot{x}_{cmd} - v_0) \tag{3.5}$$

where $K_s = 2.0 \times 10^7\,\mathrm{N/m}$ and $C_s = 1.0 \times 10^5\,\mathrm{N\,s/m}$ represent the carrier bar, the wireline and the bridle. The two units differ in x_{cmd} and in what limits the cycle rate.

The hydraulic unit follows two half cosines with an adjustable split, as shown in Equation (3.6).

$$x_{cmd} = \begin{cases} \dfrac{S}{2}\left(1 - \cos\dfrac{\pi\varphi}{f}\right), & \varphi < f \\[2ex] \dfrac{S}{2}\left(1 + \cos\dfrac{\pi(\varphi - f)}{1 - f}\right), & \varphi \geq f \end{cases} \tag{3.6}$$

where $\varphi = \theta/2\pi$ is the cycle fraction and f is the upstroke share of the cycle, adjustable between 0.3 and 0.7. The commanded cycle rate reaches its setpoint through a first-order lag of 0.8 s and is capped by the power unit flow, $\omega \leq (Q_{max}/A_A)/\max|dx/d\theta|$, with $Q_{max} = 7.0 \times 10^{-3}\,\mathrm{m^3/s}$ and $A_A = 6.207 \times 10^{-3}\,\mathrm{m^2}$. Cylinder pressure follows from the rod load, $p_A = (F(0,t) + F_h + p_B A_B)/A_A$, where $F_h = 2500\,\mathrm{N}$ is cylinder and stuffing box friction, $p_B = 6 \times 10^5\,\mathrm{Pa}$ is the return-side back pressure and $A_B = 1.6 \times 10^{-3}\,\mathrm{m^2}$ is the rod-side area. Drive power is $p_A Q/\eta_h$ with $\eta_h = 0.80$ while the cylinder lifts. On the downstroke the rods drive the cylinder, and the power unit is credited $\eta_h r$ of that mechanical power through its accumulator, with a recovery fraction $r = 0.55$, so 44 % of the downstroke work returns.

The beam unit uses the second-order slider-crank approximation, as shown in Equation (3.7).

$$x_{cmd} = \frac{S_u}{2}\,(1 - \cos\theta) + \frac{S_u \lambda}{4}\sin^2\theta \tag{3.7}$$

where $S_u = 2.03\,\mathrm{m}$ is fixed by the pitman hole and $\lambda = 0.28$ is the ratio of crank radius to pitman length. The asymmetry that λ introduces is what gives a beam unit card its characteristic shape. Crank speed is a state, driven by a variable-speed drive against the rod torque and the counterbalance, as shown in Equation (3.8).

$$J_c\,\dot{\omega} = T_m - \Delta F_0 \frac{dx_{cmd}}{d\theta} + M_{cb}\sin(\theta + \tau) - c_g\,\omega \tag{3.8}$$

where the drive torque is $T_m = \mathrm{clamp}(K_d(\omega_{sp} - \omega), -0.7T_{max}, T_{max})$ with the drive gain $K_d = 1.5 \times 10^6\,\mathrm{N\,m\,s/rad}$, $J_c = 12\,000\,\mathrm{kg\,m^2}$, $c_g = 2000\,\mathrm{N\,m\,s/rad}$ is gearbox and belt friction, $\tau = 0.15\,\mathrm{rad}$ is the counterbalance phase angle, $M_{cb} = 40\,\mathrm{kN\,m}$ and $T_{max} = 60\,\mathrm{kN\,m}$. The gearbox torque reported to the operator is $\Delta F_0\,dx_{cmd}/d\theta - M_{cb}\sin(\theta + \tau)$, and the counterbalance term is out of phase with the rod term by design. Drive power on the beam unit is that net torque times ω, charged at $1/\eta_h$ when positive and credited at $\eta_h r$ when negative, as on the hydraulic unit.

3.3.4 Pump Barrel and Valves

The barrel is the part of the model that decides what the card looks like, and it is formulated so that a check valve is integrable at a browser time step. The chamber holds a slightly compressible liquid inventory V_ℓ quoted at a reference pressure and a trapped gas inventory carried as the pressure-volume product G, which is conserved through an isothermal expansion. The chamber volume is set by the plunger position, and the two contents must fill it. Eliminating the volumes gives a quadratic in the pressure, as shown in Equation (3.9).

$$V_\ell\left(1 - \frac{p_b}{B}\right) + \frac{G}{p_b} = V_b \implies \frac{V_\ell}{B}p_b^2 - (V_\ell - V_b)p_b - G = 0 \tag{3.9}$$

where $V_b = A_p(\ell_c - z_{\mathrm{off}} + u_{10})$ is the chamber volume, $\ell_c = 0.12\,\mathrm{m}$ is the pump spacing off bottom and $B = 800\,\mathrm{MPa}$ is an effective bulk modulus for the barrel contents. The positive root is taken in whichever of its two algebraically identical forms avoids cancellation. Solving Equation (3.9) rather

than integrating dp_b/dt is what keeps the model free of stiffness, because the compliance of the chamber changes by four orders of magnitude between a gas-filled barrel and a liquid-full one and no fixed step resolves both as a differential equation.

Valve flow has a feed-forward part and a relaxation part, as shown in Equation (3.10).

$$\begin{aligned} q_{TV} &= \sigma\left(\frac{\Delta_{TV}}{\Delta p}\right)\left[\max(0,-\dot{V}_b) + \kappa C_{ch}\,\Delta_{TV}\right] \\ \Delta_{TV} &= p_b - p_t - p_c \end{aligned} \tag{3.10}$$

where σ is a smooth step over the valve opening band $\Delta p = 0.3\,\mathrm{MPa}$, $p_c = 35\,\mathrm{kPa}$ is the cracking pressure, $\kappa = 60/\mathrm{s}$ is a relaxation rate, and $C_{ch} = V_\ell/B + V_g/p_b$ is the chamber compliance. The feed-forward term is the volume the plunger sweeps, which is what an ideal check valve passes. Scaling the relaxation term by C_{ch} gives it the same time constant whether the barrel holds gas or solid liquid. The standing valve uses the same expression with the intake pressure in place of the tubing pressure and the sign of the sweep reversed.

Liquid available to the standing valve is limited by the annulus, as shown in Equation (3.11).

$$q_{SV,\ell} = \min\left(q_{SV},\, q_{SV}\,s(h) + q_{in}\,(1 - s(h))\right)(1 - f_g) \tag{3.11}$$

where $s(h) = \mathrm{clamp}(h/h_b, 0, 1)$ with $h_b = 1.5\,\mathrm{m}$. With liquid standing over the pump the annulus behaves as a tank and the demand is met in full. Once the level reaches the intake only the reservoir inflow is liquid and the balance of the demand is gas drawn off the casing, which is the mechanism that produces pump-off in this model rather than a correlation imposed on top of one. The force the pump applies to the bottom node is Equation (3.12).

$$F_{pump} = (p_b - p_t)(A_p - A_r) - F_f \tanh\left(\frac{v_{10}}{v_\varepsilon}\right) - c_p v_{10} + F_{tag} \tag{3.12}$$

where $F_f = 700\,\mathrm{N}$ and $c_p = 2000\,\mathrm{N\,s/m}$ are plunger friction terms, $v_\varepsilon = 0.02\,\mathrm{m/s}$ regularizes the Coulomb term at reversal, and F_{tag} is a contact force, capped at 400 kN, that appears when the plunger reaches the standing valve cage.

3.3.5 Wellbore and Reservoir

The annulus is an inventory whose level sets the pump intake pressure, as shown in Equation (3.13).

$$A_{ann}\frac{dh}{dt} = q_{in} - q_{SV,\ell} + q_{SV,leak} + q_T, \qquad p_i = p_{c,h} + \rho_f g h \tag{3.13}$$

where $A_{ann} = 0.015\,77\,\mathrm{m}^2$, $p_{c,h}$ is the casing head pressure, $q_{SV,leak}$ is the flow a leaking standing valve returns to the wellbore, q_T is a tubing leak, and $\rho_f = 986.8\,\mathrm{kg/m}^3$ is the produced fluid density for a 14 °API oil at 30 % water cut. Inflow is a straight-line productivity index, $q_{in} = J(P_r + \Delta P_r - p_i)$ with $J = 1.10 \times 10^{-10}\,\mathrm{m}^3/(\mathrm{s\,Pa})$, which is 0.41 bbl per day per psi. The time constant of the level is Equation (3.14).

$$\tau_h = \frac{A_{ann}}{J\rho_f g} = 4.12\,\mathrm{h} \tag{3.14}$$

A stroke lasts about ten seconds and the fluid level answers in hours, so the two ends of this plant are separated by more than three orders of magnitude in time. Every rate decision made on the well is made against a consequence that will not be visible until the next shift.

3.3.6 Parameters, Initialization and Numerics

Table 3.2 lists the parameters. The rod string is initialized at rest at the bottom of the stroke with the traveling valve open, so the barrel sits at tubing pressure and the polished-rod load equals W_{rf} exactly, while the cycle rate starts at the commanded rate. The engine of Section 1.6.1 integrates the states with fixed-step fourth-order Runge-Kutta at 2 ms. The two fastest modes of the discretized string are the plunger on a closed, liquid-full barrel at 42 Hz and the polished rod on the carrier bar at 47 Hz, which the step resolves at 0.6 rad per step, and the valve transfer is resolved at the same time. The controller and the card sampler execute every 40 ms, which is a realistic scan rate for a wellsite controller, and the historian logs five samples per second.

3.3.7 Assumptions

The rod string is untapered and of a single diameter. Real strings are tapered in two or three sections so that the stress is roughly uniform down the string, and both the forward solution and the inversion then have to be run section by section with matching conditions at each taper.

The barrel contents are given an effective bulk modulus of 800 MPa rather than the 1.5 GPa of clean produced liquid. That spreads the load transfer over about 2 % of the plunger stroke instead of making it instantaneous, which is what allows a 2 ms step to resolve a check valve. The application therefore cannot make a claim about the sharpness of a real load transfer beyond what its own reconstruction can resolve.

Free gas is one volume fraction at the intake, with no bubble point, no solution gas, no gas separator and no casing gas flow. Gas interference and gas lock in this model are consequences of compressing that fraction, while in a real well they are also consequences of where the perforations are relative to the pump.

Inflow is a straight-line productivity index with no skin, no transient term and no two-phase relationship below the bubble point. Leakage is linear in the differential across it, and the cube-of-clearance dependence that makes plunger slip so sensitive to wear is folded into a single wear multiplier.

The well is vertical, so there is no rod-on-tubing contact, no side loading and no deviation-dependent damping. Rod fatigue is reported as the modified Goodman utilization at the present cycle and is not accumulated.

3.4 Card Reconstruction and Diagnosis

3.4.1 Stroke Segmentation and the Gibbs Solution

Surface position and load are sampled at 25 Hz into a ring buffer. Each time the cycle phase wraps, the samples belonging to the completed stroke are resampled onto 128 points uniform in time and the reconstruction runs. Substituting $u = U(z)\exp(\mathrm{i}n\omega t)$ into Equation (3.1) gives an ordinary differential equation in z for each harmonic whose wavenumber is Equation (3.15).

$$k_n = \beta_n + \mathrm{i}\gamma_n, \qquad \begin{aligned} \beta_n &= \frac{n\omega}{a\sqrt{2}}\sqrt{\sqrt{1+(\nu/n\omega)^2}-1} \\ \gamma_n &= \frac{n\omega}{a\sqrt{2}}\sqrt{\sqrt{1+(\nu/n\omega)^2}+1} \end{aligned} \tag{3.15}$$

The displacement and load at depth then follow from the surface coefficients in closed form, as shown in Equation (3.16).

$$\begin{aligned} U_n(z) &= U_n(0)\cosh k_n z + \frac{F_n(0)}{EA_r k_n}\sinh k_n z \\ F_n(z) &= EA_r k_n U_n(0)\sinh k_n z + F_n(0)\cosh k_n z \end{aligned} \tag{3.16}$$

Table 3.2 Model parameters. Hidden parameters are not exposed in the interface

Symbol	Meaning	Value	Units
L	pump setting depth	900	m
d_r	rod diameter	0.0254	m
d_p	plunger diameter	0.05715	m
$d_{t,i}$	tubing inside diameter	0.062	m
$d_{t,o}$	tubing outside diameter	0.073	m
d_c	casing inside diameter	0.1594	m
E	rod Young modulus (hidden)	2.07×10^{11}	Pa
ρ_s	rod density (hidden)	7850	kg/m^3
T_u	rod minimum tensile	7.93×10^{8}	Pa
ν_0	structural damping rate	1.0	1/s
m_{pr}	polished rod mass	60	kg
m_{pl}	plunger mass	35	kg
ℓ_c	pump spacing off bottom	0.12	m
S_{sp}	stroke the pump was spaced for	3.0	m
B	effective barrel modulus	8×10^{8}	Pa
p_c	valve cracking pressure	3.5×10^{4}	Pa
Δp	valve opening band (hidden)	3×10^{5}	Pa
κ	valve relaxation rate (hidden)	60	1/s
k_{slip}	plunger slip conductance	3×10^{-12}	m^3/(s Pa)
F_f	plunger Coulomb friction	700	N
c_p	plunger viscous friction	2000	N s/m
p_{wh}	wellhead tubing pressure	3×10^{5}	Pa
$p_{c,h}$	casing head pressure	2×10^{5}	Pa
API	oil gravity	14	°API
w_c	water cut	0.30	–
P_r	reservoir pressure	5.5×10^{6}	Pa
J	productivity index	1.10×10^{-10}	m^3/(s Pa)
S_u	beam unit stroke	2.03	m
λ	crank radius / pitman	0.28	–
M_{cb}	counterbalance moment	4×10^{4}	N m
J_c	crank inertia	1.2×10^{4}	kg m^2
T_{max}	gearbox torque rating	6×10^{4}	N m
A_A	cylinder bore area	6.207×10^{-3}	m^2
Q_{max}	power unit flow limit	7×10^{-3}	m^3/s
$p_{A,max}$	hydraulic pressure limit	2.1×10^{7}	Pa
P_{max}	motor power limit	4.5×10^{4}	W
p_B	return-side back pressure	6×10^{5}	Pa
A_B	cylinder rod-side area	1.6×10^{-3}	m^2
F_h	cylinder and box friction	2500	N
η_h	power unit efficiency	0.80	–
r	downstroke energy recovery	0.55	–
K_d	drive stiffness	1.5×10^{6}	N m s/rad
c_g	gearbox and belt friction	2000	N m s/rad
τ	counterbalance phase angle	0.15	rad
SF	rod service factor	1.0	–

The mean term is exact and carries no damping, because the cycle average of both the inertia and the damping term vanishes for a periodic displacement, so the mean load is the same at every depth and the mean displacement varies linearly with it. Twenty-four harmonics are retained and evaluated at $z = L - \Delta z/2$, which is where the discretized plant reports its own bottom element force. The whole calculation costs about twenty-five thousand multiplications and twelve thousand sine and cosine evaluations per stroke.

Two properties of the inversion matter to the reader. A truncated Fourier series cannot represent a step, so a load transfer is smeared over a fixed fraction of the cycle however sharp the real one is. In mid-stroke, where the plunger moves fastest, that fraction is a long stretch of travel, and the pumped-off card of Scenario 2 transfers in 7.6 % of the plunger stroke as reconstructed against 2.9 % on the true card. And ν in Equation (3.15) is not measured. It is a setting, and the true value on this well moves from 2.42/s to 6.3/s when the oil goes from 800 cP to 3000 cP.

> **TRY IT IN THE APP 3 · Break the reconstruction on purpose (MANUAL)**
>
> Run in Manual on the Cards tab and read the Card error badge: 1.4 % with the assumed damping at its correct 2.4/s. In Advanced, set the assumed damping to 8/s and the error grows to 45.6 %, then to 0.5/s and it reads 15.5 %. Back at 2.4/s, fire the cold spell to 3000 cP: the error climbs to 30.1 % with no setting changed.

3.4.2 Features and the Fluid Level Inference

Pump fillage is the fraction of the downstroke remaining when the load transfers, measured at the midpoint between the two card plateaus. Transfer sharpness is reported as an effective width, the fluid load divided by the steepest slope the transfer reaches, in units of plunger stroke. Taking it from the steepest slope rather than from two threshold crossings matters because any level taken as a fraction of the peak-to-trough load range is contaminated by rod dynamics at the other end of the card.

The fluid level inference takes the fluid load from the cycle mean of the downhole load rather than from the two card plateaus, as shown in Equation (3.17).

$$F_o = \frac{k_F\,(\bar{F}_{cyc} - W_z)}{\phi_{on}}, \qquad \hat{h} = \frac{p_t - F_o/(A_p - A_r) - p_{c,h}}{\rho_f g} \tag{3.17}$$

where W_z is the static rod tension at the pump, ϕ_{on} is the fraction of the cycle during which the load is above the midpoint of the card, and k_F is a calibration constant. Inertia and Coulomb friction both integrate to zero over a closed cycle whatever shape the trajectory takes, so the cycle mean cancels them exactly. A plateau difference carries the friction twice, and a plateau sum cancels inertia and friction only when the two halves of the stroke are mirror images, which a crank never provides. Computed from the same reconstructed cards at the design submergence, a plateau sum reads 17 m high on the hydraulic unit and 20 m high on the beam unit.

Pause. Try Exercise 3.1 before reading on.

The remaining bias at the design point is a 2.3 % underestimate of the fluid load, and $k_F = 1.023$ removes it. In the field the same constant comes from a single acoustic fluid-level shot [20], which is why it is exposed as a setting rather than buried as a number. An occasional direct measurement that resets the bias of a continuous inference is also how the soft sensor of Volume 5, Chapter 10 is kept honest by its laboratory. The inference is sensitive by construction, because the annular column contributes 0.24 MPa out of the 9.01 MPa of tubing pressure at 25 m of submergence, so 1 % of the fluid load is 8.9 m of level and the estimate is filtered over several strokes before the level controller sees it. The load-on fraction ϕ_{on} is interpolated between the two midpoint crossings rather than counted in samples, because one sample of the 128-point card is 0.8 % of the cycle and would move the estimate by about 12 m in a step that no filter removes. With 90 m over the pump in Scenario 3 and no load cell noise, the calibrated estimate reads between 4.0 m low and 0.1 m high from 3.5 SPM to 6 SPM on the hydraulic unit (1.5 m low at the scenario's 4.9 SPM), and between 0.2 m low and 3.7 m high on the beam unit.

3.4.3 Fault Ranking

Nine conditions are scored against the card features and the corroborating measurements each stroke. Normal operation is scored as the residue of the others rather than as a pattern of its own, because a confident normal indication sitting above a fault the same pipeline had half detected is the worst

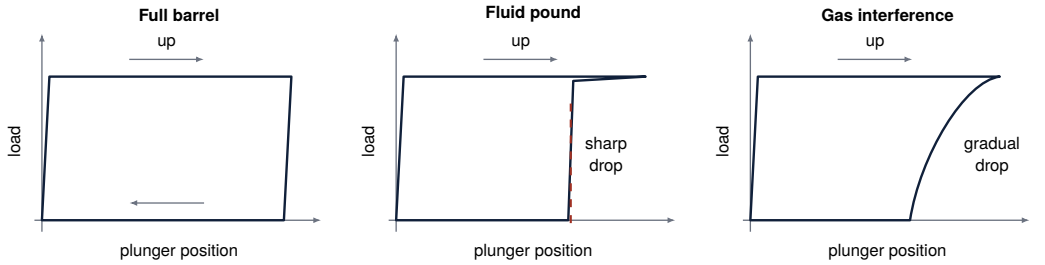

Figure 3.2 Idealized downhole cards, sketched. A full barrel transfers its load at the ends of the stroke, in 2.0 % to 2.7 % of the plunger stroke as reconstructed. The pumped-off barrel of Scenario 2, at 69 % fillage, drops the load sharply in mid-stroke, in 7.6 % as reconstructed, while the gassy barrel of Scenario 10, at 65 % fillage, sheds it gradually over 13.2 %, because the gas has to be compressed before the traveling valve opens. Notice that fillage alone cannot tell the two apart. The shapes are drawn, not computed, and not to scale.

failure mode a diagnostic display has. Each indication carries its evidence, the confirmation it still needs and a recommended response.

Three pairs are deliberately not separated. A leaking traveling valve and a worn plunger are the same leak path from the tubing side and produce the same card, so they are reported as one indication whose confirmation line names the valve test that distinguishes them [21]. A leaking standing valve and a tubing leak both remove production without changing the card, so they are ranked within a few points of each other (0.85 and 0.80 with the standing valve leaking in Scenario 12, 0.29 and 0.28 with the tubing leaking) and the confirmation line says the card cannot separate them. Incomplete fillage and fluid pound are the same condition at different severity, and both are reported when both apply.

The pair that is separated is the one that matters most for the rate decision. Fluid pound and gas interference both give low fillage and require opposite responses, and the effective transfer width separates them. In the shipped scenarios a full barrel transfers its load in 2.0 % to 2.7 % of the plunger stroke as reconstructed. The pumped-off barrel of Scenario 2, at 0.21 MPa intake pressure and 69 % fillage, takes 7.6 %. In Scenario 10 a similar fillage of 65 % produced by 35 % free gas takes 13.2 %, and 60 % free gas at 47 % fillage takes 17.9 %. The diagnostic layer scores pound below a width of 10 % and gas above 7.5 %, so the pumped-off card ranks incomplete fillage first and fluid pound second, and the gassy card ranks gas interference first. The inferred fluid level corroborates the split, because the pounding well reads 0 m and the gassy one 61 m. Field practice reaches the same separation through the plunger velocity computed from the same card [22]. Figure 3.2 sketches the three card shapes.

Pause. Try Exercise 3.2 before reading on.

TRY IT IN THE APP **10 · Gas interference, and then gas lock**

Load the case and press Run, in Auto with PI as loaded, and watch the Diagnostics tab. When the free gas at the intake rises to 35 %, fillage falls to 65 % and gas interference, not fluid pound, becomes the top indication. On the Cards tab the load transfer spreads over 13.2 % of the plunger stroke, against 7.6 % for the pumped-off card of Scenario 2.

3.5 Rate Control

Three distinct models appear in this application and they are kept apart throughout. The simulation model is the wave equation, the pump barrel and the annular inventory of Section 3.3, and it is the plant. The estimation model is the Gibbs propagation of Section 3.4, which maps surface measurements to a downhole card and never predicts forward in time. The controller's internal model is the reduced

steady-state relation of Equation (3.20), which is used only by the supervisory search and knows nothing about the dynamics of either of the other two.

3.5.1 The Two Loops

Both loops use the velocity-form PID of Section 1.6.2, executed every $\Delta t_c = 40\,\text{ms}$, written in this chapter's symbols in Equation (3.18). The running sum s_k is S_k of Equation (1.17), the demand c_k is v_k of Equation (1.18) with no output bias, and α is the action a.

$$\begin{aligned} e_k &= \alpha\,(\text{SP} - \text{PV}_k) \\ s_k &= s_{k-1} + K_c\,(e_k - e_{k-1}) \\ c_k &= s_k + I_{k-1} \\ u_k &= \text{clamp}(c_k,\ u_{min},\ u_{max}) \\ I_k &= I_{k-1} + \frac{K_c}{\tau_I}\,e_k\,\Delta t_c + \frac{\Delta t_c}{T_t}\,(u_k - c_k) \end{aligned} \tag{3.18}$$

where I is the integral state, seeded with the stroke rate or stroke length in force when the controller is built, and $T_t = \tau_I$ is the tracking time of the back-calculation. The derivative term is zero in both loops. The limits are the minimum and maximum automatic stroke rate, 1 and 9 SPM, for the level loop and the stroke length limits, 2.2 and 3.2 m, for the fillage loop. The protection layer, the stroke-length slew limit and the power-unit flow limit all act on the command after the controller, and the integral does not see them, so it keeps accumulating against its own output while any of them holds the unit below that output.

Both loops set $\alpha = -1$. Raising the stroke rate lowers the fluid level and lengthening the stroke lowers the fillage, so both process gains are negative and both controllers are direct acting in the ISA sense. The level loop manipulates stroke rate with $K_c = 0.030\,\text{SPM/m}$ and $\tau_I = 400\,\text{s}$ against a process whose time constant is 4.12 h. The fillage loop manipulates stroke length with $K_c = 3.0\,\text{m}$ per unit of fillage and $\tau_I = 120\,\text{s}$, and it exists only on the hydraulic unit. In manual operation, under the timer and under the optimizer both controllers track the command in force by setting I so that c_k equals it, so a transfer to PI is bumpless (Section 1.6.2), and a transfer to manual during a run copies the commands into the operator stations. Under PI the fillage loop also tracks the stroke station until the first reconstructed card exists, because there is no measured fillage before it.

Stroke length is slew limited to 0.10 m/s because it is a mechanical setting even on a hydraulic unit. A beam unit holds $S = S_u$ at all times.

3.5.2 Four Strategies

The pump-off timer runs at a fixed rate and counts strokes whose fillage falls below 75 %. Two consecutive low strokes stop the unit, the brake holds for a dwell of 1200 s, and the unit restarts. The dwell is the only tuning parameter and nothing measured at the surface computes it.

Proportional-integral control on the inferred fluid level runs the unit continuously against a submergence setpoint. The measurement is Equation (3.17), so anything that changes the fluid load without changing the level moves the quantity the loop is really holding.

The constrained optimizer evaluates a reduced steady-state model on a grid of 429 candidate operating points (33 rates by 13 strokes), drops every candidate that violates a predicted constraint, and takes the best objective, as shown in Equation (3.19). The grid is searched after every stroke in every mode, so the best feasible point is always displayed, and the optimizer acts on the result every five strokes.

$$\max_{N,S}\ [q_{oil} - \lambda_E P - \lambda_W D] \tag{3.19}$$

subject to $F \geq F_{min}$, $h \geq h_{min}$, $PPRL \leq PPRL_{max}$, $U_G \leq U_{G,max}$, $N_{min} \leq N \leq N_{max}$, the power unit flow, pressure and motor limits on the hydraulic unit, the gearbox torque limit on the beam unit,

and a two-sided pump spacing limit. The weights are $\lambda_E = 4 \times 10^{-10}\,\mathrm{m^3/J}$ and $\lambda_W = 2 \times 10^{-5}$, and the damage index is $D = (1-F)^2 F_o / W_{rf} + 2\max(0, U_G - 0.7) + 0.6 U_G^2$. The search takes under a millisecond of computation (0.7 ms on a current desktop processor). The supervisor of Volume 9, Chapter 8 picks its setpoints by the same kind of constrained grid search. The result is rate limited to 0.5 SPM and 0.25 m of stroke per update, and an increase in rate requires three consecutive healthy strokes, while a decrease does not wait.

The reduced model used inside the search resolves the operating point from one comparison. If the pump can remove liquid from the annulus faster than the reservoir can supply it, the well is drawn down and fillage is the ratio of the two. Otherwise the level settles where inflow matches the pump and fillage is complete, as shown in Equation (3.20).

$$F = \min\left(1 - f_g\left(1 - \frac{p_i}{p_t}\right),\ \frac{q_{in}(p_i) + q_{leak,out}}{A_p S_{net} N (1 - f_g)}\right) \tag{3.20}$$

where $S_{net} = S - F_o L/(E A_r)$ is the net plunger stroke. Rod stretch is 151 mm at 25 m of submergence, which is 5 % of a three meter stroke, so omitting the correction overstates displacement by the same amount.

3.5.3 Protection Layer

A protection layer runs underneath the three automatic strategies, moves the commands only downward, and responds within one stroke. Peak polished-rod load above the limit reduces the rate command by 10 %. A fluid pound index above 0.35 reduces it by 15 %. A tagging indication shortens the stroke by 0.15 m on the hydraulic unit or reduces the rate by 15 % on the beam unit. Each test reads the values latched for the last completed stroke. The timer and the PI strategy recompute their commands at every 40 ms scan, so for them each cut is a derate held until a stroke completes inside the limit. The optimizer holds its command between updates, so under the optimizer each cut is applied once, on the scan that processes the new stroke. In the challenge with the optimizer selected and the shipped limits, the start-up stroke peaks at 68.4 kN against a 68 kN limit, the scan at $t = 8.16$ s cuts the command from 9.0 SPM to 8.1 SPM, and the search then walks it down at 0.5 SPM per update to its 5.0 SPM target by $t = 337.4$ s. The layer is separate from the supervisory search on purpose, because a controller that thinks in minutes cannot answer a fluid pound. It is a one-way override in the sense of Section 1.4.8, like the protective loop of Volume 4, Chapter 2 that rests at its clamp until a limit comes near.

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

3.6 Scenarios and Results

Table 3.3 lists the thirteen scenarios in the order they are presented, with what completes each. They progress from reading a healthy card, through the pump-off relation and the estimator that depends on it, to the strategies and then to the faults. Each case is credited when the learner performs the listed actions and the simulation reaches the listed state. The challenge is judged on the case as given, so changing a model parameter or a disturbance during the attempt voids it, and the credit accumulates toward certificates of progress and of completion for the module, as Section 1.6.4 describes. Figure 3.3 shows the application on the card view.

3.6.1 Operating Points and the Pump-Off Boundary

The reservoir delivers at most 50.4 m³/d of liquid, which is 317 bbl/d, when the level is drawn to the pump. A three meter stroke reaches that displacement at 5.15 SPM once rod stretch and the 5 % free gas are counted. Table 3.4 gives the reduced model at five rates and shows what the extra strokes buy. Oil production is flat at 222 bbl/d above the pump-off rate while, from 5.5 SPM to 9.0 SPM, peak polished-rod load rises from 13.71 klbf to 15.28 klbf, rod fatigue utilization rises from 54 % to 62 %, and drive power rises from 16.8 kW to 28.3 kW.

Table 3.3 Scenarios and Completion Requirements

#	Scenario	Required to complete
1	Healthy well by hand	In Manual at 3.5 SPM with 120 m over the pump, run until a full stroke is drawn and compare the surface card with the reconstruction for 10 s, then raise the rate to 4.4 to 4.6 SPM and run 30 s more with the unit stroking at that rate.
2	Pump-off at 8 SPM (demonstration)	In Manual, run until the level is within 2 m of the pump with true fillage below 80 %, and hold that for 300 s. Then apply the balance rate (5.1 SPM, any setting up to 5.2 SPM) and run until the barrel fills to at least 90 % with a pound index below 0.1 for 60 s.
3	Wrong assumed damping	In Manual, set the assumed damping to 7.5/s or more until the card error exceeds 30 %, then to 0.6/s or less, then back to 2.4/s with the 3000 cP cold spell fired, and run 25 s with the card error above 20 %.
4	Pump-off timer	Run the timer for two hours with the 1200 s dwell and, after Reset, two hours with a dwell of 300 s or less, then with a dwell of 3600 s or more until the unit has stood braked for 50 min.
5	PI on the inferred level	Run one hour on PI with no load cell bias. After Reset, set a bias of 400 to 500 N, run until the true level exceeds the estimate by 25 m, and complete a full hour with the bias.
6	Stroke length on fillage	Run one hour on PI with the hydraulic unit (level setpoint 3 m, fillage target 88 %), then after Reset switch to the beam unit and run another hour on PI.
7	Beam unit	Run five strokes in Auto on the beam unit, then change the counterbalance moment from 40 kN m and run 30 s more with it changed.
8	Constrained optimizer	Run two hours on the optimizer with the minimum fillage at 85 %, then after Reset lower it to 62 % or less and run two hours more.
9	Cold spell	After the viscosity step to 3000 cP at 600 s, hold 60 s with the card error above 15 %. Before the oil warms at 3000 s, set the assumed damping to 5.5 to 6.5/s and hold 30 s with the card error below 8 %.
10	Gas interference and gas lock	Hold 60 s with at least 30 % free gas and gas interference as the top indication, then run 300 s on PI in Auto with at least 55 % free gas after the step at 2400 s.
11	Traveling valve washout	After the valve cut at 600 s, hold 60 s with volumetric efficiency below 75 % while true fillage stays above 95 %, then run until the level loop commands at least 6.5 SPM.
12	Two leaks	After 1900 s, with the standing valve leaking, have the standing-valve leak and the tubing leak as the top two indications. After 3900 s, with the tubing leaking, have both among the top three.
13	Challenge	Select PI or the optimizer before the first disturbance at 400 s. At 2400 s, fillage within 9 points of 90 % from 400 s on, peak polished-rod load under 68 kN and rod fatigue under 85 % over the whole run, and pound damage under 1.0, with the model parameters and disturbances as given.[a]

[a] As loaded, the pump-off timer at 9 SPM fails, because a timer is not an accepted strategy and its peak polished-rod load reaches 73.1 kN (worst fillage error 4.9 points, rod fatigue 72 %, no pound damage). Selecting the optimizer with the shipped limits still fails on peak load, at 68.4 kN on its start-up stroke and 68.7 kN after the cold spell, and PI fails on fillage (11.4 points), load (72.6 kN) and damage (6.74). The suggested settings (optimizer between 1.5 and 6.5 SPM and 2.4 and 3.2 m of stroke, 0.4 SPM per update every four strokes) pass with a worst fillage error of 5.2 points, 3.8 to spare, a peak load of 67.1 kN, 0.9 kN to spare, 62 % rod fatigue and no pound damage, whether run directly or after Reset. Four other noise seeds change no verdict, put the PI damage between 6.04 and 6.56, and move no other reading by more than 0.1 in the units shown.

Table 3.4 Reduced model against stroke rate at a three meter stroke

N SPM	F %	h m	q_{oil} bbl/d	PPRL klbf	U %	P kW
4.0	96	124	172	12.51	48	10.9
4.9	95	27	211	13.34	52	14.5
5.5	94	0	222	13.71	54	16.8
6.0	86	0	222	13.92	55	18.2
9.0	57	0	222	15.28	62	28.3

The constrained optimum on the hydraulic unit is 5.25 SPM at a 2.95 m stroke, which produces 222 bbl/d with the fillage constraint binding. It sits on the constraint rather than at the peak of the



Figure 3.3 Card view of the rod pump application in Scenario 2 at $t = 4500$ s, after the well has pumped off with eight strokes per minute requested and 7.18 SPM delivered, 0.7 m of fluid over the pump and 69 % fillage. (1) The surface card at the left leans with rod stretch. (2) The reconstructed downhole card at the right, solid, is drawn over the simulator's hidden true card, dashed, and (3) it transfers its load a third of the way down the stroke. (4) The top-ranked indication is incomplete pump fillage, and (5) the card error is 2.5 %. The trends and the control panel are cropped away.

objective, because the objective is flat above the pump-off rate and the energy and wear terms break the tie downward. On the beam unit the same search returns 8.00 SPM with the stroke rate limit binding, because a 2.03 m stroke has to be run half again as fast for the same displacement.

3.6.2 Strategy Comparison

Table 3.5 compares the strategies over six hours of simulated time from the initial condition of Scenario 5, with the scorecard read at the end of the run. Each row selects that case and changes only the mode, the strategy, the surface unit or the manual rate before the run starts, and Figure 3.4 draws its oil, energy and damage columns. The manual rows keep the 3 m stroke of the operator station, on which the power-unit flow limit delivers 7.18 SPM whether 8 SPM or 9.0 SPM is requested, so the timer row and the 9.0 SPM manual row stroke at the same rate while the unit runs.

Asking the hydraulic unit for 9.0 SPM produces 56.3 bbl against the optimizer's 55.4 bbl, an advantage of 1.6 %, and pays for it with 392.5 units of accumulated pound damage against zero, a peak polished-rod load of 15.41 klbf against 13.68 klbf, and 2.51 kW h per barrel against 1.75 kW h, an increase of 43 %. The pump-off timer produces 51.7 bbl while the unit is stroking for only 67 % of the run, and it accumulates 7.6 units of damage on the strokes between each restart and the shutdown that follows it. Proportional integral control on the inferred level produces 52.3 bbl because it is holding a submergence setpoint of 25 m rather than pushing to the fillage constraint, which costs 5.6 % of the production the optimizer finds, at a peak load of 13.57 klbf against 13.68 klbf. Held by hand at the 4.9 SPM design point, the well settles at 22.7 m of true level and lifts 52.6 bbl.

The beam unit reaches the same production as the hydraulic unit under the optimizer, 55.6 bbl against 55.4 bbl, at 1.49 kW h per barrel against 1.75 kW h, because the counterweights return the rod weight to the gearbox on the downstroke while the hydraulic power unit returns only 44 % of the downstroke work through its accumulator. It pays with a peak polished-rod load of 14.78 klbf against

Table 3.5 Six hours of simulated operation from the same initial condition. Fill is the worst fillage and PPRL the peak polished-rod load of the whole run. Damage is the accumulated fluid-pound severity, one unit per stroke landing at full fluid load on an empty barrel. Three other noise seeds move no entry by more than 0.1 bbl, 0.05 klbf or 0.5 damage units

Strategy	Oil bbl	Rate bbl/d	Energy kWh/bbl	Fill %	PPRL klbf	Dam.	Run %
Hydraulic long-stroke unit							
Manual, 4.9 SPM	52.6	210.4	1.73	95	13.55	0.0	100
Manual, 9.0 SPM	56.3	225.1	2.51	69	15.41	392.5	100
Pump-off timer	51.7	207.0	1.86	75	15.10	7.6	67
PI on fluid level	52.3	209.3	1.72	95	13.57	0.0	100
Constrained optimizer	55.4	221.6	1.75	94	13.68	0.0	100
Conventional beam unit							
PI on fluid level	51.5	205.8	1.45	95	14.60	0.0	100
Constrained optimizer	55.6	222.3	1.49	93	14.78	0.0	100

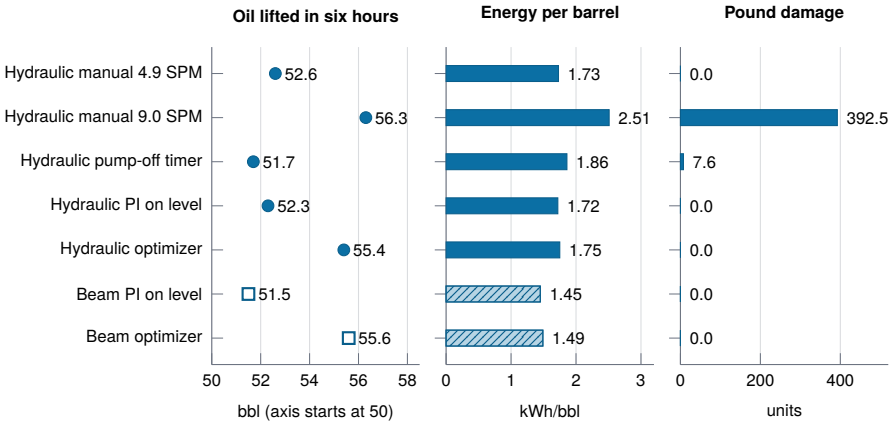

Figure 3.4 Six hours of each strategy from Table 3.5: oil lifted (dots on an axis that starts at 50 bbl), energy per barrel and accumulated pound damage. Filled circles and solid bars are the hydraulic unit, open squares and hatched bars the beam unit. Notice that the 9.0 SPM hand setting buys the most oil with by far the most damage and energy, while the optimizer on either unit comes close to it with none.

13.68 klbf and a fatigue utilization of 60 % against 55 %, because 2.03 m of stroke has to be run at 8.00 SPM to displace about what 2.95 m displaces at 5.25 SPM.

> ✦ **DID YOU KNOW? Two Automatic Strategies Lift Less Than a Hand Setting**
>
> From the simulator (Table 3.5). Over six hours on the hydraulic unit the pump-off timer lifts 51.7 bbl and PI on the inferred level 52.3 bbl, both less than a fixed hand setting of 4.9 SPM at 52.6 bbl. Only the constrained optimizer, at 55.4 bbl, beats that hand setting with no pound damage.

3.6.3 Verification

The application ships 32 automated checks that run headless with every build. Three of them carry most of the weight.

The reduced model used by the optimizer and by the operating map is checked against the wave-equation plant at six operating points spanning both surface units. Each point places the fluid level where the reduced model says it belongs and marches the plant for 700 s. The test requires fillage

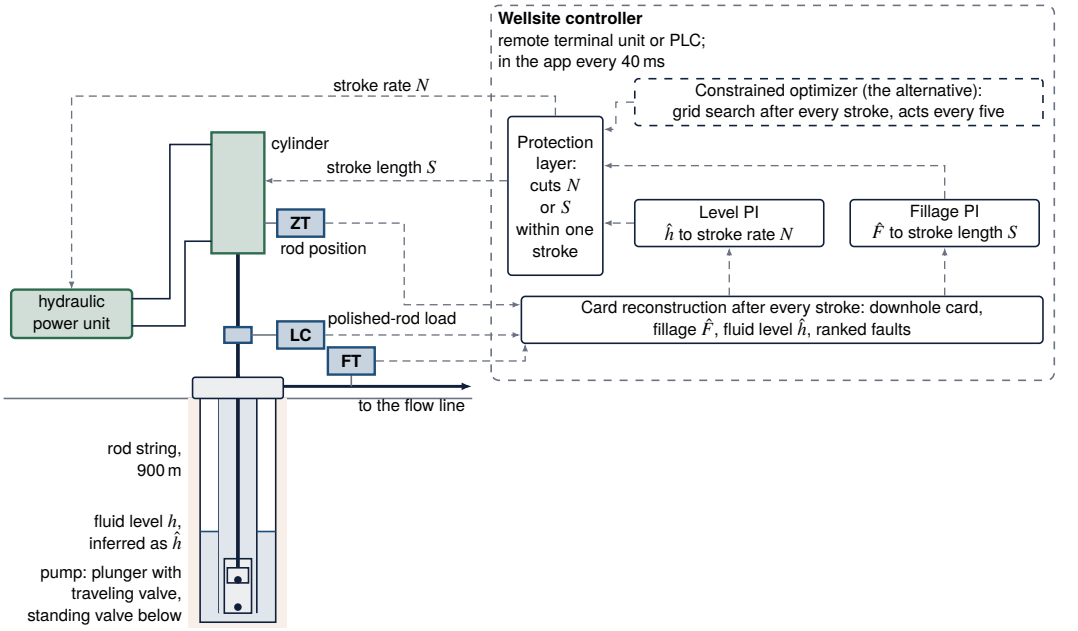

Figure 3.5 The loops of the rod pump simulator with their hardware. Lettered boxes are sensing devices: the polished-rod load cell (LC), the rod position transducer (ZT) and the wellhead liquid rate (FT). The hydraulic power unit and the cylinder are the final elements, and dashed arrows are signals. Notice that both controlled variables come out of the card reconstruction, not out of an instrument, and that every command passes the protection layer on its way down. Not to scale, and not a piping diagram.

within 7 percentage points, oil rate within 5 % and peak polished-rod load within 5 %, and the worst case across the six points is 4.8 percentage points of fillage, 1.4 % of oil rate and 2.0 % of peak load.

The reconstruction is checked against the simulator's hidden true card at four conditions covering a full pump, a pounding pump, a gassy pump and a short stroke. The test requires a root-mean-square error below 6 % of the card load range, and the four conditions return 1.18 %, 2.34 %, 1.25 % and 1.39 %. Raising the assumed damping from 2.4/s to 9/s must increase the error by more than a factor of three, and it raises it from 1.33 % to 50.1 %. Scenario 3 asks the reader to produce the same failure, and in it the card error is 1.4 % at 2.4/s, 45.6 % at 8/s, 15.5 % at 0.5/s, and 30.1 % at 2.4/s once the oil has been taken to 3000 cP.

A kick applied to the plunger reaches the load cell in 0.154 s against the 0.175 s that L/a predicts, a 12 % difference that is the discretization of ten elements. The remaining checks cover equilibrium with the unit stopped, the sign of every actuator, controller direction, the optimizer against a grid five times finer in rate and four times finer in stroke, the protection layer's response inside one stroke, closure of the annular inventory at steady state, the challenge's start-up rate envelope, the balance-rate helper on both units at four gas fractions, and the framework's own checks on the scenarios, the historian and bumpless transfer.

3.7 Technology in Practice

Figure 3.5 draws the two loops, the protection layer and the supervisory search of Section 3.5 with the hardware of a rod-pumped well.

What Measures It

A routine well is measured at the polished rod. A load sensor gives the load and an angle sensor on a beam unit gives the displacement, and the controller turns the two into a pump card through the wave equation [15]. On the hydraulic unit the position comes from the cylinder instead, drawn in Figure 3.5 as a transducer beside it. The wellhead liquid rate of Table 3.1 is the third measurement. In the model the load cell is the only instrument with noise and bias settings.

> **TECHNOLOGY** **Sensor: The Polished-Rod Load Cell**
>
> A load cell is one or more strain gauges bonded to a metal element of precisely known elastic properties, which stretches under force like an extremely stiff spring [23]. On a pumping unit a load sensor measures the polished-rod load and an angle sensor its displacement, and a wave equation turns the two into the pump card [15]. A load cell also reads vibration as force, and it reads stray current, such as from arc welding, as a change in load and can be damaged by it [23]. The chapter's load cell carries noise σ_F up to 2700 N and a bias b_F up to 9000 N either way, and a bias moves the inferred fluid level although nothing downhole has changed.

What Moves It

The controller moves two things, the stroke rate and, on the hydraulic unit, the stroke length. The power unit's flow sets how fast the cylinder can move, which is why the cap Q_{max} of Section 3.3 delivers 7.18 SPM on a 3 m stroke whether 8 SPM or 9.0 SPM is requested. On the beam unit a variable-speed drive moves the crank against the gearbox torque rating T_{max}.

> **TECHNOLOGY** **Final Element: The Hydraulic Long-Stroke Unit**
>
> Hydraulically actuated rod pumping units have been in use since the 1940s. A hydraulic cylinder at the surface reciprocates the polished rod and the rod string [4]. The unit is not subject to damaging torque limits and can work at depths up to 14 000 ft. Its two main drawbacks are higher maintenance cost and more frequent failures than a beam unit [4]. On a beam unit the stroke rate is set instead by a frequency converter that adjusts the motor speed [15]. In the chapter the cylinder's rate follows its command through a 0.8 s lag under the power-unit flow limit, and the stroke length is slew limited to 0.10 m/s.

Where the Controller Runs

Oil wells are run as field sites of a supervisory system. Remote terminal units or programmable logic controllers at each site control the local process, while a control center polls them at defined intervals, such as 5 s or 60 s, and can send new setpoints. Field sites often have remote access for diagnostics and repairs [24]. Loops in such systems run with cycle times from milliseconds to minutes [24]. The app scans its two loops and the card sampler every 40 ms, reconstructs the card once per stroke of about ten seconds, and lets the optimizer act every five strokes. The protection layer has to stay in the wellsite controller, because a center that polls once a minute cannot answer a fluid pound within one stroke. The chapter's structure maps onto two PID blocks in velocity form with output clamps, a card-processing routine that feeds them, a supervisory routine that writes the same two commands as the alternative, and a one-way override beneath all of them.

> **IN PRACTICE** **Reading the Card Is the Daily Job**
>
> Most operators rely on the downhole card rather than the surface card, because it depicts the conditions at the pump more accurately. Its known weaknesses are data delay, noise and inaccurate readings [4]. A 2021 review lists 26 common card shapes, from insufficient liquid supply and gas influence to a leaking traveling or standing valve, a tubing leak and a broken rod [4]. Fluid pound shortens the life of the whole downhole assembly and can end in a broken string [4]. The chapter's diagnostic layer scores nine conditions and leaves three pairs unseparated.

What Is Modernizing It

The first development is variable-speed control with remote supervision. One field system measured the polished-rod load and angle, computed the pump card, and set the stroke speed through a frequency converter to make the card's area as large as possible. It could be watched from a browser on a computer or a smartphone through a cellular data link, and its authors count automatic stroke-speed adjustment among the most effective ways to save energy in the oil field [15]. The second is learning-based diagnosis. Deep learning and machine learning methods now classify downhole cards and electrical power curves [4], and one method diagnoses faults from motor power alone, with its training labels taken from cards [13]. The chapter shows why those labels need care: with a wrong assumed damping, the computed card misses the true card by tens of percent.

3.8 Limitations

The single largest simplification is the untapered rod string. Field strings are tapered and the inversion then propagates section by section, which changes the numerical conditioning of the reconstruction and not its principle.

The effective barrel bulk modulus of 800 MPa is a numerical accommodation. It sets a floor on how sharp a load transfer the plant can produce, and the reconstruction sets a second and larger floor, which in mid-stroke reports a transfer of 2.9 % of the plunger stroke as 7.6 %. Every threshold in the diagnostic layer is calibrated against the reconstruction rather than against the truth for that reason, and the application cannot support a claim about the sharpness of a real transfer.

> **OPEN CHALLENGE** **A Diagnosis That Does Not Rest on One Load Cell**
>
> Dynamometer cards do not always give accurate results and are usually not available in real time [4]. Diagnosis from cards is also limited by the upkeep of the load sensor: its maintenance and calibration, battery replacement and the hazards to the staff who service it. One line of work learns the state of the well from motor power instead, with training data labeled from cards [13], and a 2021 review expects such methods could rival the card in the future [4]. In this chapter a load cell bias under 1 % of the peak load moves the inferred level by 44 m. Progress would be a diagnosis that checks the load cell against a second, independent signal on every well.

Gas is one volume fraction at the intake. The application therefore reproduces gas interference and gas lock as consequences of compression, but it has nothing to say about gas separator design, pump setting depth relative to the perforations, or the casing pressure management that a real gassy well requires.

The reservoir is a straight-line productivity index, so the application cannot be used to study inflow below the bubble point or any transient in the formation. Rod fatigue is reported as a utilization at the present cycle with no cumulative damage model, so the application ranks operating points by fatigue but does not predict a failure.

> **OPEN CHALLENGE** **Rods in a Crooked Hole**
>
> In a deviated well the rods rub on the tubing. The friction wears both and can end in severe leaks and breakage, and centralizers are added to keep them apart [4]. A 2021 review cites work showing that analyzing a deviated well with the models built for vertical wells can lead to serious errors, and it lists crooked holes among the limitations of rod pumping [4]. The chapter's well is vertical, with no rod-on-tubing contact, no side loading and no damping that depends on deviation. Progress would be a card reconstruction that carries the side loads and friction of a deviated string as reliably as the vertical one carries damping.

The four strategies are the ones a rod-pumped well commonly carries. Model predictive control on the inferred fluid level and moving-horizon estimation of the annular height are both natural extensions of the model in Section 3.3 and neither is implemented here, because the supervisory

search was chosen for transparency and for an execution cost of under a millisecond that fits between two strokes. The app therefore does not carry the predictive controller and estimator option of Section 1.6.3. Either extension needs a dynamic optimization package rather than the closed-form search used here [25].

3.9 Conclusion

This chapter studied a heavy-oil rod pumped well as a laboratory for inferential measurement and constrained rate control. Twenty-eight states carry a ten-element discretization of the damped rod-string wave equation, a pump barrel whose pressure is solved algebraically from a liquid inventory and a trapped gas inventory, an annular inventory with a 4.12 h time constant and a productivity index inflow. The downhole card is reconstructed each stroke by propagating twenty-four harmonics of the surface card to depth, and agrees with the simulator's hidden card to between 1.18 % and 2.34 % of the card load range while the assumed damping is correct. Oil production is flat at 222 bbl/d above the 5.15 SPM pump-off rate while peak polished-rod load rises 11 % and drive power rises 68 % from 5.5 SPM to 9.0 SPM, which is the trade the four rate strategies are graded on.

The constrained optimizer sits on the fillage constraint at 5.25 SPM and a 2.95 m stroke rather than at the unconstrained maximum of its own objective. The diagnostic layer reports conditions in ranked groups and declines to separate the three pairs that share a card signature. A fuller model would add a tapered string with section-by-section inversion, a two-phase fluid column with a bubble point, cumulative rod fatigue, moving-horizon estimation of the annular height, and a model predictive rate controller for comparison against the grid search. Even without them, the well shows what control looks like on a plant whose controlled variables are computed rather than measured.

KEY POINTS

- The pump is out of sight, so the fillage and fluid level a controller holds are computed from a load cell and a position sensor at the surface, not measured.
- Oil production is flat at 222 bbl/d above the 5.15 SPM pump-off rate while peak polished-rod load and drive power keep rising.
- The constrained optimizer sits on the fillage constraint at 5.25 SPM and a 2.95 m stroke rather than at the unconstrained maximum of its own objective.
- Fluid pound and gas interference both give low fillage and need opposite responses, and the effective transfer width separates them.

Where else this applies. Volume 5, Chapter 11 also closes its loop on a variable no instrument reaches, through a soft sensor.

References

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

[2] J. Park, R. A. Martin, J. D. Kelly, and J. D. Hedengren, "Benchmark temperature microcontroller for process dynamics and control," *Computers & Chemical Engineering*, vol. 135, p. 106736, 2020.

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

[4] S. Fakher, A. Khlaifat, M. E. Hossain, and H. Nameer, "A comprehensive review of sucker rod pumps' components, diagnostics, mathematical models, and common failures and mitigations," *Journal of Petroleum Exploration and Production Technology*, vol. 11, no. 10, pp. 3815–3839, 2021.

[5] R. P. Cutler and A. J. Mansure, "Fluid dynamics in sucker rod pumps," Sandia National Laboratories, Albuquerque, NM, Report SAND99-0093C (Southwest Petroleum Short Course, Lubbock, TX), OSTI 3261, 1999.

[6] S. G. Gibbs, "Predicting the behavior of sucker-rod pumping systems," *Journal of Petroleum Technology*, vol. 15, no. 7, pp. 769–778, 1963.

[7] S. G. Gibbs and A. B. Neely, "Computer diagnosis of down-hole conditions in sucker rod pumping wells," *Journal of Petroleum Technology*, vol. 18, no. 1, pp. 91–98, 1966.

[8] S. G. Gibbs, "A review of methods for design and analysis of rod pumping installations," *Journal of Petroleum Technology*, vol. 34, no. 12, pp. 2931–2940, 1982.

[9] T. A. Everitt and J. W. Jennings, "An improved finite-difference calculation of downhole dynamometer cards for sucker-rod pumps," *SPE Production Engineering*, vol. 7, no. 1, pp. 121–127, 1992.

[10] D. R. Doty and Z. Schmidt, "An improved model for sucker rod pumping," *Society of Petroleum Engineers Journal*, vol. 23, no. 1, pp. 33–41, 1983.

[11] G. Takacs, "The analysis of sucker-rod pumping installations," in *Sucker-Rod Pumping Handbook*. Elsevier, 2015, pp. 423–503.

[12] B. Zheng, X. Gao, and X. Li, "Diagnosis of sucker rod pump based on generating dynamometer cards," *Journal of Process Control*, vol. 77, pp. 76–88, 2019.

[13] J. Wei and X. Gao, "Fault diagnosis of sucker rod pump based on deep-broad learning using motor data," *IEEE Access*, vol. 8, pp. 222 562–222 571, 2020.

[14] S. G. Gibbs, "Method of determining sucker rod pump performance," U.S. Patent 3 343 409, Sep. 26, 1967.

[15] H. Lv, J. Liu, J. Han, and A. Jiang, "An energy saving system for a beam pumping unit," *Sensors*, vol. 16, no. 5, p. 685, 2016.

[16] J. J. Kilgore, H. A. Tripp, and C. L. Hunt, "Walking beam pumping unit system efficiency measurements," in *SPE Annual Technical Conference and Exhibition*, 1991, paper SPE-22788-MS.

[17] G. Takacs, "Long-stroke sucker-rod pumping," in *Sucker-Rod Pumping Handbook*. Elsevier, 2015, pp. 505–551.

[18] S. G. Gibbs, "Method for determining the pump-off of a well," U.S. Patent 3 951 209, Apr. 20, 1976.

[19] S. G. Gibbs, K. B. Nolen, F. E. Morrow, and W. E. Lynch, "Monitoring and controlling rod pumped wells using downhole pump cards," in *Annual Technical Meeting*, 1995, paper PETSOC-95-89.

[20] O. L. Rowlan, J. N. McCoy, and A. L. Podio, "Acoustic liquid-level determination of liquid loading in gas wells," in *SPE Western Regional/AAPG Pacific Section/GSA Cordilleran Section Joint Meeting*, 2006, paper SPE-100663-MS.

[21] K. B. Nolen and S. G. Gibbs, "Quantitative determination of rod-pump leakage with dynamometer techniques," *SPE Production Engineering*, vol. 5, no. 3, pp. 225–230, 1990.

[22] R. Messer, "Differentiate between pump off and gas interference with real time plunger velocity," in *SPE Artificial Lift Conference and Exhibition - Americas*, 2024, paper D021S006R005.

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

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

[25] L. D. R. Beal, D. C. Hill, R. A. Martin, and J. D. Hedengren, "GEKKO optimization suite," *Processes*, vol. 6, no. 8, p. 106, 2018.

3.10 Exercises

3.1 ⌕ Concept check · 2 min

The fluid level inference of Equation (3.17) takes the fluid load from the cycle mean of the reconstructed downhole load, not from the two plateaus of the card. Why?

(a) The plateaus are smeared by the truncated Fourier series, while the mean is not

(b) The cycle mean is unaffected by a load cell bias, so the level estimate needs no calibration

(c) A plateau difference cancels friction, but the cycle mean averages out more noise

(d) Inertia and Coulomb friction integrate to zero over a closed cycle, whatever the stroke's shape

3.2 ⌕ Concept check · 2 min

A pumped-off well and a gassy well both show a fillage near 65 % to 70 %. Which card feature lets the diagnostic layer tell them apart, and what does each need?

(a) The effective transfer width, narrow for a pound and wide for gas; slow down only the pounding well

(b) The fillage value, always lower for a pound than for gas; slow down only the gassy well

(c) The peak polished-rod load, raised by a pound and not by gas; lengthen the stroke on both wells

(d) The card error, which rises when gas is present; raise the assumed damping on the gassy well

3.3 Concept check · 2 min

In the challenge with the optimizer selected and the shipped limits, the start-up stroke peaks at 68.4 kN against a 68 kN limit. What happens next, and why is that job not left to the optimizer?

(a) The limit is raised to the stroke's peak until the next search confirms the overload

(b) The protection layer cuts the rate command once by 10 %, since the search acts only every five strokes

(c) The unit is braked for the timer dwell, since an overload counts as a pump-off

(d) The stroke is shortened by 0.15 m, since stroke length lowers peak load more than rate

3.4 Theory and calculation · 7 min

The rods are steel with $E = 207\,\text{GPa}$ and $\rho_s = 7850\,\text{kg/m}^3$, 25.4 mm in diameter and 900 m long. The plunger is 57.15 mm, the tubing pressure at the pump is 9.01 MPa, the casing head pressure 0.2 MPa and the produced fluid density 986.8 kg/m^3. Take $g = 9.81\,\text{m/s}^2$.

(a) Find the speed of sound in the rods and the time a disturbance at the plunger takes to reach the load cell. Compare it with the plant's 0.154 s. *Checkpoint:* 5135 m/s, 0.175 s; the plant is 12 % early.

(b) With 25 m of fluid over the pump, find the intake pressure, the fluid load $F_o = (p_t - p_i)(A_p - A_r)$ and the rod stretch $F_o L/(EA_r)$ as a share of a 3 m stroke. *Checkpoint:* 0.442 MPa, 17.6 kN, 151 mm, 5.0 %.

(c) How far does the inferred level move when the fluid load is misread by 1 %? *Checkpoint:* 8.9 m.

3.5 Theory and calculation · 7 min

Use the well of Table 3.2: casing 159.4 mm inside, tubing 73 mm outside, $J = 1.10 \times 10^{-10}\,\text{m}^3/(\text{s Pa})$, $P_r = 5.5\,\text{MPa}$, $p_{c,h} = 0.2\,\text{MPa}$, $\rho_f = 986.8\,\text{kg/m}^3$, 5 % free gas, a 3 m stroke with 151 mm of rod stretch and the plunger area of the previous exercise. Take $g = 9.81\,\text{m/s}^2$.

(a) Find the annular area and the time constant τ_h of the fluid level. *Checkpoint:* 0.015 77 m^2; about 4.1 h.

(b) Find the most liquid the reservoir can deliver, with the level drawn to the pump. *Checkpoint:* 5.83×10^{-4} m^3/s, 50.4 m^3/d.

(c) Find the stroke rate at which the pump's liquid capacity $A_p S_{net} N(1 - f_g)$ equals that inflow. Then add plunger slip, $k_{slip} = 3 \times 10^{-12}\,\text{m}^3/(\text{s Pa})$ acting over half the cycle across $p_t - p_{c,h}$ with $p_t = 9.01\,\text{MPa}$, and find the rate again. *Checkpoint:* 5.04 SPM without slip, 5.15 SPM with it.

3.6 Theory and calculation · 5 min

Viscous drag on the rods is $f_{drag} = 2\pi\mu v/\ln(r_t/r_r)$ per meter (Equation (3.2)), with tubing inside radius $r_t = 31\,\text{mm}$, rod radius $r_r = 12.7\,\text{mm}$, $\rho_s A_r = 3.978\,\text{kg/m}$ and structural damping $\nu_0 = 1.0/\text{s}$.

(a) At 800 cP, find the drag coefficient per meter, the drag force on 900 m of rods moving at 1 m/s, the damping rate ν and the share of it that viscosity supplies. *Checkpoint:* 5.63 N s/m per meter, 5.1 kN, 2.42/s, 59 %.

(b) Repeat for 3000 cP, and say what happens to the reconstructed card if the assumed damping stays at 2.4/s. *Checkpoint:* 6.31/s, 84 % viscous.

3.7 App challenge · 8 min

Load scenario 2, Push the rate until the well pumps off. Change no setting, use the default seed, run at speed 100 to $t = 4500$ s (about 45 s of wall time) and pause. Export the CSV file.

(a) Read the delivered stroke rate on the stroke rate card, the fillage on the stroke length card, the Card error tile, and `level_true_m` and `oil_rate_bpd` in the last row. *Checkpoint:* 7.18 SPM, 69 %, 2.5 %, 0.74 m, 219 bbl/d.

(b) Predict the fillage as the reservoir's 50.4 m^3/d over the liquid the pump sweeps, $A_p S_{net} N(1 - f_g)$, with $S_{net} = 2.849$ m and the delivered rate. *Checkpoint:* 71.8 m^3/d; 70 %.

(c) Press Set the rate that just fills the pump and run on to $t = 5400$ s. Read the stroke rate, the fillage, `oil_rate_bpd` and the Pound damage tile. *Checkpoint:* 5.1 SPM, 95 %, 218 bbl/d; the damage stops climbing.

Any pause from 3500 to 5000 s gives the same readings in (a) within 0.1 m, 0.4 % of card error and 1 bbl/d.

3.8 App challenge · 8 min

Load scenario 5, PI on the inferred fluid level, and switch to SI units. Change no other setting, use the default seed, run at speed 100 to $t = 3600$ s and export the CSV file. Then press Reset, set Load cell bias in Advanced to 450 N, and run to $t = 3600$ s again.

(a) In both runs read the Oil produced tile, and `level_true_m`, `level_est_m` and `spm_cmd` in the last row. *Checkpoint:* 1.41 m^3, 22.9 m, 21.1 m, 4.68 SPM; 0.62 m^3, 87.3 m, 43.2 m, 2.86 SPM.

(b) Find the drop in oil and the gap between the true and the inferred level in the biased run. *Checkpoint:* 56 %; 44 m.

(c) Estimate the gap from the bias, taking $\phi_{on} \approx 0.5$, $A_p - A_r = 2.058 \times 10^{-3}$ m^2 and $\rho_f = 986.8\,\mathrm{kg/m^3}$, and explain why the stroke rate fell. *Checkpoint:* About 45 m.

Readings within 60 s of the end agree within 1 m on the true level, 1.5 m on the estimate and 0.03 m^3 on the oil.

3.9 App challenge · 9 min

Load the challenge, Lift the most oil without breaking the rods, and switch to SI units. It runs on its own seed, which you leave as it is. Run as loaded at speed 100 to $t = 2400$ s and export the CSV file. Then press Reset, press Reveal suggested tuning and run to 2400 s again.

(a) For the run as loaded, read the Peak rod load and Pound damage tiles, the largest value of `goodman_ratio` in the file, and the worst $|\texttt{fill_pct} - 90|$ from 400 s on. Which requirements fail? *Checkpoint:* 73 kN, 0.00, 0.72, 4.9 points; the strategy and the peak load.

(b) For the suggested run, read the challenge result line and find the margin on fillage and on peak load. *Checkpoint:* 5.2 points, 67.1 kN, 62 %, 0.00; 3.8 points and 0.9 kN to spare.

(c) With the optimizer and the shipped limits, the start-up stroke peaks at 68.4 kN. Which suggested setting keeps the start-up under 68 kN, and why? *Checkpoint:* The 6.5 SPM maximum automatic rate.

3.10 Application area · 4 min

The fault ranking deliberately reports three pairs of conditions without separating them.

(a) Name the two pairs that are different faults, and say why the card cannot tell each pair apart. *Checkpoint:* Traveling valve leak and worn plunger; standing valve leak and tubing leak.

(b) For each pair, name the field action that separates them. *Checkpoint:* Valve tests on the stopped unit; pressuring up the tubing.

3.11 Application area · 4 min

Table 3.5 gives six hours of operation from the same start.

(a) From the oil and energy-per-barrel columns, find the drive energy of the 9.0 SPM manual row and of the hydraulic optimizer row, and the energy spent per extra barrel. *Checkpoint:* 141.3 kW h and 97.0 kW h; about 49 kW h per extra barrel.

(b) An operator argues that the 1.6 % of extra oil pays for the faster rate. What else does the row cost, and is the extra oil sustained? *Checkpoint:* 392.5 damage units and 1.73 klbf more peak load; no.

3.12 Review the agent · 5 min

An AI assistant reads the optimizer's display in scenario 8: the target is 5.25 SPM on a 2.95 m stroke and the binding constraint is fillage. It writes: "Fillage is the binding constraint, so it is what holds the rate down. Lower Minimum pump fillage from 85 % to 60 % and the optimizer will move to a faster rate and lift more oil." Find the flaw and say how you would check it in the app.

CHAPTER 4

Fast Choke Overrides and Slow Waterflood Planning in a Smart Well Simulator

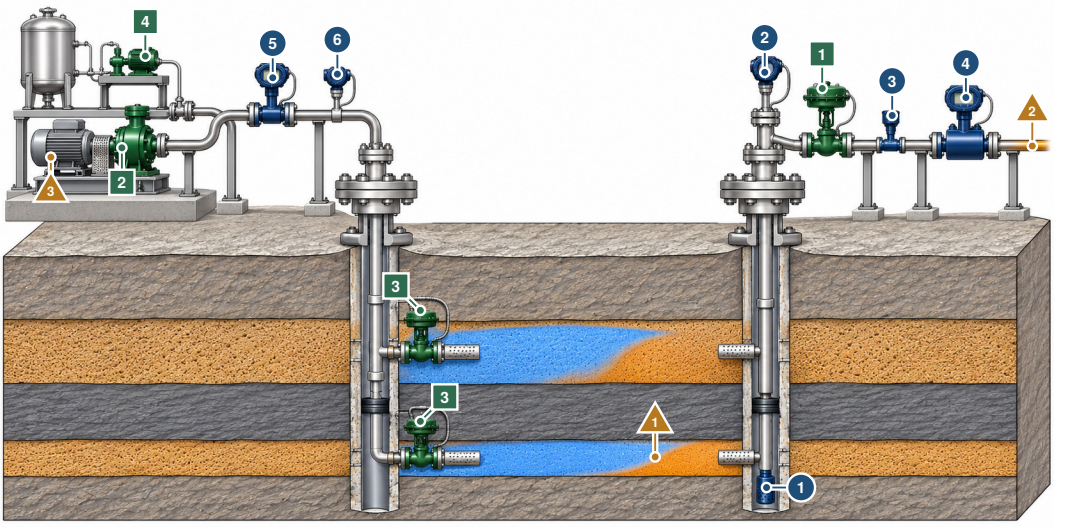

The controller moves 1 production choke 2 injection pump speed
3 two downhole valves that split the water between the layers 4 polymer dosing

The controller reads 1 bottomhole pressure 2 wellhead pressure 3 sand detector
4 multiphase meter for the oil and water rates 5 injection flow 6 injection wellhead pressure

What disturbs it 1 a thin layer that takes most of the water 2 flowline pressure 3 injection pumps out of service

The water in the thin lower layer has run ahead of the water in the thick upper layer, far below the choke that sets the rate. Illustration, not to scale.

IN PLAIN TERMS **Two Clocks on One Oil Well**

The main idea. Water pumped down one well pushes oil through rock to a second well. A valve at the top of that well sets its flow within minutes and holds back sand and extra water, while the water takes weeks to cross the rock.

The challenge. One thin layer of rock lets water through easily, so the water races along it and leaves oil behind. Settings that make the most oil this month can recover less by the end of six months.

Why it matters. In the chapter's simulator, a plan that looks six months ahead makes 2.33 times the oil of a fixed valve setting.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? box in the Introduction, the two History boxes in Section 4.2, the first Try it box in Section 4.4, the second Try it box and the second Did you know? box in Section 4.6, and Key points in Section 4.9. No equations.
Study	55 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 4.6 with Figure 4.5, Section 4.7 and Section 4.8 with its Open challenge boxes.

The control question A production choke changes the rate of an oil well within minutes, but the water it produces was injected weeks earlier, and the layer that water came through decides how much oil the flood leaves behind. How do fast choke overrides and a slow plan for the sweep share one well, and why do the settings that make the most oil this month recover less by the end of a campaign?

Learning objectives

- Compute the lead-lag of the wellbore seen from the choke and the time constants of the two zone pressures from the model's parameters.
- Explain how four PI loops in external-reset form share one choke through a low selector without windup, and predict which limit takes it.
- Estimate a thief zone's breakthrough from the fractional-flow curve and the Welge tangent, and explain how a zone split or polymer delays it.
- Show that a choke can only spend reservoir pressure that injection provides, and that the injector's fracture limit bounds that injection.
- Compare planning horizons over a 180-day campaign and explain why the plan that leads on day 30 does not lead on day 180.

Prerequisites Sections 1.4.2 (inflow, choke and the share of a choke move that reaches the rate), 1.4.8 (selectors and external reset), 1.6.2 (the framework's PID) and 1.6.4 (scenarios, scorecard and CSV).

The app https://apmonitor.com/apps/wellhead-eor/

See also Volume 9, Chapter 7 for a slow store planned days ahead on forecasts with errors. Volume 4, Chapter 14 for a fast actuator that drains a finite store, and Volume 4, Chapter 4 for a feed loop under a low selector with overrides measured late. Volume 5, Chapter 11 for a state that no instrument reaches, inferred from a model.

4.1 Introduction

A waterflooded oil field runs on two clocks. At the wellhead a production choke sets the rate within minutes, against limits on drawdown, sand and produced water that act just as fast. In the reservoir, the water injected to replace the produced volume advances as a front whose speed follows the fractional-flow theory of immiscible displacement [1, 2], and a permeable layer carries water to the producer long before the tighter rock is swept [3]. Downhole inflow control valves decide where the injected water goes [4], polymer thickens the water so that it displaces more oil [5, 6], and optimal control of the valves of smart wells has been used to improve a heterogeneous waterflood [7].

+ DID YOU KNOW? Hundreds of Polymer Floods

Polymer flooding has been evaluated at field scale for more than five decades. A survey published in 1983 already identified 273 polymer projects, abandoned, current, planned or contemplated. A 2017 review still counted large commercial floods as relatively few outside Canada and China [8].

The control problem has a general shape. A fast actuator draws on a store that refills slowly, and the move that maximizes today's output is not the one that maximizes the output over the life of the store. Volume 4, Chapter 14 meets the same shape in a vault of reserves with a floor. Reservoir management has been posed as closed-loop model-based control of a flood over its life [9], long-term recovery and short-term production have been optimized as two layers on two time scales [10], and process systems engineering has asked when production optimization may be treated as static and when it must be dynamic [11]. The arrangement mirrors the layered practice of process control, in which fast regulatory loops run beneath slower model-based layers [12], and the limits that share one actuator are handled by selectors and overrides [13].

The misconception the simulator targets is that the choke setting producing the most oil now also recovers the most. In the campaign of Section 4.6, the shipped configuration, a 1300 m^3/d target with both injection valves open and no polymer, is 48 % ahead of a fixed-choke waterflood on day 30 and only 7.1 % ahead on day 180, while a planner that looks 180 days ahead produces 2.33 times

the fixed-choke oil. A second misconception is that a choke controls the rate. The choke sets the bottomhole pressure, and the rate follows from a reservoir pressure that the choke cannot raise, so when the injection plant trips the rate loop holds its target for 49.6 h and then runs out of travel.

However, the fast and the slow half of this problem are usually studied apart. Waterflood optimization takes well rates or bottomhole pressures as its controls [7, 10], and a tuning study of a choke treats the reservoir as a fixed pressure behind a resistance, so neither shows the drawdown that ties the choke to the front. Good practice in control education favors laboratories, physical or virtual, whose plants behave like the real ones [14].

This chapter joins a producer with its choke overrides and an injector with its valves, pumps and polymer through one two-zone reservoir on one simulated clock, so that a choke move and a sweep decision can be traced to the same barrel of oil. Waterflood theory, override control and receding-horizon planning are all established practice. Section 4.2 traces the domain. Section 4.3 builds the model on two clocks, an hourly transport step that moves exactly the volumes integrated on a one-minute clock, so that the phase balances close to 3×10^{-14}. Section 4.4 reads its time scales from the wellbore to the flood. Section 4.5 gives the two control layers, the four choke loops at a low selector and a planner whose quasi-steady forecast stays within 0.31 % of the dynamic plant over 180 days, and it splits the planner's forecast error into model error and grid coarsening. Section 4.6 runs eight cases and a campaign that separate what the fast loops protect from what the slow decisions recover. Section 4.7 sets the hardware of a real producer and injector beside the model. Sections 4.8 and 4.9 state the limits of the model and what the chapter has shown.

4.2 Waterflood Sweep, Mobility Control and Smart Wells

The engineering of waterflooding rests on a theory of displacement published in 1942. It treats water and oil as flowing together with a water fraction set by their relative mobilities, and it predicts a sharp front whose speed follows from the slope of that fraction [1]. A tangent construction ten years later turned the calculation into a graphical one that gives the front saturation, the volume injected at breakthrough and the average saturation behind the front [2]. The simulator uses the first result as its transport equation and the second as a check that its numerics must pass.

Reservoirs are layered, and a layered flood is limited by its most permeable layer. A 1949 method for waterflood performance used the distribution of permeability to flood the layers in order, the most permeable first, so that the producer sees water from the fastest layer while the slow layers still hold most of their oil [3]. The thief zone of this chapter is that fastest layer. It holds 29 % of the pore volume and takes 77 % of the injected water at the start of a plain flood.

Mobility control followed. A water-soluble polymer lowers the mobility of the injected water so that it pushes oil ahead of it rather than fingering past. The first History box below follows the idea from its first report to field scale.

> ⧖ HISTORY **1964: Thickening the Flood Water**
>
> Engineers had long wanted a better water-oil mobility ratio in waterfloods, above all for viscous crudes. Thicker water would do it, but glycerin, sugar or glycols cost far too much to use [5]. In 1964 a paper reported that a few synthetic polymers, the earliest of them containing acrylamide, cut the mobility of water in rock 5 to 20 times more than their viscosity predicted. Field pilots then recovered 80 and 100 percent more oil than comparable waterflood pilots, in fields with oils of 16 and 130 centipoise [5]. A 1991 monograph gathered the properties of polymer solutions, their retention on the rock and their transport through it [6]. A 2010 review of enhanced recovery counted chemical flooding among the growing trends of recent years, beside carbon dioxide and high-pressure air injection [15].

The well has its own short time scales. Wellbore storage and skin were given an analytical treatment in 1970 that explains why a bottomhole pressure answers a rate change only after the wellbore volume has been drawn down or filled [16]. Multiphase flow through a wellhead choke

Table 4.1 Roles, Defaults and Ranges of the Model Variables

Variable	Role	Default	Range
q_L	CV of FIC-101, MPFM	900 m³/d SP	300–1600
$\Delta\hat{P}$	CV of PDIC-102, model	50 bar limit	5–90
$\dot{m}_s$	CV of SIC-103, SD-103	0.30 g/s SP	0.05–2
q_w	CV of FIC-104, MPFM	760 m³/d SP	100–2000
q_{inj}	CV of FIC-201, FT-201	VRR 1.00	0.6–1.4
p_{iwh}	CV of PIC-202, PT-202	145 bar limit	60–200
x	MV, production choke	56.49 %	0–100
n	MV, injection pump speed	65.75 %	0–100
s, C_{inj}	MVs, thief share, polymer	open, 0 ppm	20–90, 0–2000
p_{fl}	DV, flowline pressure	20 bara	10–60
a	DV, pumps in service	100 %	0–100
$\dot{s}_{pl}$	DV, injector plugging	0 per day	0–10
b_s	DV, sand detector bias	0 g/s	±0.3
p_1, p_2	states, zone pressures	250 bar	—
p_{wf}, x, n	states, well and actuators	213.2 bar	—
s, s_{pl}	states, valve share, skin	0.769, 0	—
V, M, m_s, E	states, 12 cumulatives	0	—
S_j, C_j, Γ_j	slow states, 48 cells	S_{wi}, 0, 0	—

becomes critical, and insensitive to the downstream pressure, once the pressure ratio across it falls low enough [17]. Weak sandstones fail when too much drawdown is taken across the sandface, and the ways in which arriving water triggers sand production have been examined [18].

Automation reached the reservoir through smart wells, completions with downhole valves and gauges that let each zone be produced or injected on its own, reviewed for reservoir management in 2005 [4]. The second History box below follows these wells into closed-loop control, and the application implements the hierarchy that ends it in reduced form, with a receding-horizon enumeration above a regulatory layer of overrides.

> ⧖ HISTORY **2004 to 2011: Closing the Loop Around a Reservoir**
>
> In 2004 optimal control theory set the valves of smart wells in a waterflood of a heterogeneous reservoir. For wells held to a rate, it found large scope to speed up production and raise recovery while the water produced fell sharply [7]. A review of 2005 held that a well with intelligent components is smart only when it maximizes its value over the life of the project [4]. The same year an ensemble Kalman filter updated a reservoir model continuously from measured production, permeability included, and the forecasts improved as data accumulated [19]. In 2011 a hierarchy put the value over the field's life first and short-term production second, and in one study it doubled the first year's production for a loss of value of only 0.3 percent [10].

4.3 A Producer, an Injector and Two Zones

The simulated sector is a strip 300 m long and 300 m wide between an injector and a producer 2000 m deep. It extends the single-zone well W-1 of Section 1.4.2, whose inflow, tubing column and choke it shares, with a second zone, an injector and the fronts between them. The notation follows Section 1.3. Two zones fill it, a tight zone 12 m thick at 700 mD and porosity 0.20 and a thief zone 4 m thick at 7000 mD and porosity 0.25, with no communication between them except through the two wellbores. Their pore volumes are 216 000 and 90 000 m³, and the sector holds 204 000 Sm³ of oil at the connate water saturation. The wells, the zones, the instruments and the symbols of the equations are shown in Figure 4.1. Inside the model time is in hours, pressures in bar and reservoir rates in m³/h, and the daily rates quoted here are conversions, with liquid and water in m³/d and oil in Sm³/d at standard conditions.

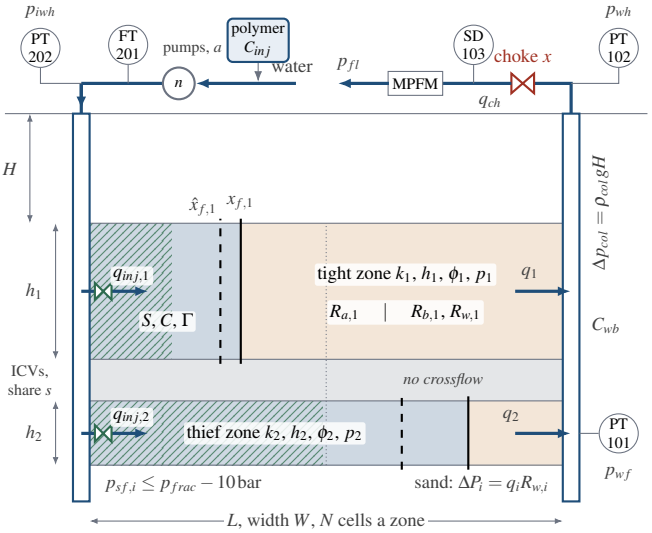

Figure 4.1 Producer, injector and the two-zone sector as modeled. Water (blue, the darker shading) has displaced oil (brown, the lighter shading) behind the true fronts $x_{f,i}$ (solid), the operator's model places them at $\hat{x}_{f,i}$ (dashed), and polymer is hatched. Each zone keeps its pressure p_i at its midpoint (dotted), with the resistance $R_{a,i}$ toward the injector and $R_{b,i} + R_{w,i}$ toward the producer. The inflow control valves (ICVs) split the injection, PT-101 reads the bottomhole pressure p_{wf} behind the wellbore storage C_{wb}, the tubing loses Δp_{col}, and the choke passes q_{ch} past the sand detector SD-103 and the multiphase meter (MPFM) to the flowline at p_{fl}. FT-201 and PT-202 meter the injection. The drawing is not to scale.

Table 4.1 classifies the variables. Nineteen fast states are integrated every minute: the two zone pressures, the bottomhole pressure, the choke and pump positions, the injector skin, the valves' thief share and twelve cumulative quantities that cross the well terminals (injection, production and polymer per zone, choke throughput, surface oil and water, injected water, sand and erosion). The saturation, dissolved polymer and adsorbed polymer of the 48 cells are slow states, advanced once an hour.

4.3.1 Transport in Each Zone

Within each zone the water saturation S and the polymer concentration C in the water move along x from the injector at $x = 0$ to the producer at $x = L$. Relative permeabilities follow the Corey form in the normalized saturation, as shown in Equation (4.1).

$$S_n = \frac{S - S_{wi}}{1 - S_{wi} - S_{or}}, \qquad f_w = \frac{\lambda_w}{\lambda_w + \lambda_o}$$
$$\lambda_w = \frac{k^0_{rw} S_n^{n_w}}{\mu_w R(C)}, \qquad \lambda_o = \frac{k^0_{ro}(1 - S_n)^{n_o}}{\mu_o} \tag{4.1}$$

where λ_w and λ_o are the phase mobilities, f_w is the fractional flow of water and $R(C) = 1 + aC + bC^2$ is the viscosity multiplier of the polymer solution, 7.0 at 1000 ppm. At the end points water is 9.41 times as mobile as the 10 cP oil, and 1200 ppm of polymer ($R = 8.92$) brings the ratio to 1.06.

Water and polymer are transported by the Darcy flux u of the zone, as shown in Equation (4.2).

$$\begin{aligned} \phi \frac{\partial S}{\partial t} + \frac{\partial (u f_w)}{\partial x} &= S \frac{\partial u}{\partial x} \\ \phi \frac{\partial (SC + \Gamma)}{\partial t} + \frac{\partial (u f_w C)}{\partial x} &= SC \frac{\partial u}{\partial x} \\ \Gamma &\leftarrow \max\left(\Gamma,\ \Gamma_{max} \frac{C}{C + C_{1/2}}\right) \end{aligned} \tag{4.2}$$

where u falls linearly from the injection rate at $x = 0$ to the production rate at $x = L$, both divided by the cross-section $A_i = W h_i$, and Γ is the adsorbed polymer per unit pore volume. The right-hand sides remove the share of a voidage imbalance that each phase stores elastically, in proportion to its saturation, so that the saturation bookkeeping stays that of an incompressible displacement while the zone pressure takes up the imbalance. Adsorption is irreversible, so Γ grows toward the Langmuir isotherm and never falls, and the polymer front trails the water that carries it.

Each zone has $N = 24$ cells, and the update is explicit and upwind. Once every simulated hour a zone is advanced by the injected volume, the produced volume and the polymer mass that crossed its two ends in that hour, which are states of the minute-scale integration. The face volumes fall linearly from the injected to the produced volume, each cell's stored share is moved to an elastic ledger, and the step would be sub-cycled if its Courant number passed 0.5. In every shipped case the largest Courant number is 0.065, so no sub-cycling occurs. Injected minus produced equals the change in place plus the elastic ledger to 1.9×10^{-14} in relative terms for oil, water and polymer over 20 days of polymer injection at a voidage ratio of 0.9.

4.3.2 Zone Pressures and Resistances

Each zone keeps one pressure p_i at its midpoint, with the storage $c_t PV_i$, and the resistances between that midpoint and each well come from the mobilities of the cells they cross, as shown in Equation (4.3).

$$\begin{aligned} c_t PV_i \frac{\mathrm{d}p_i}{\mathrm{d}t} &= q_{inj,i} - q_i, \quad q_i = \max\left(0,\ \frac{p_i - p_{wf}}{R_{b,i} + R_{w,i}}\right) \\ R_{a,i} &= \sum_{j<N/2} \frac{\Delta x}{k_i A_i \lambda_{t,j}}, \qquad R_{b,i} = \sum_{j \geq N/2} \frac{\Delta x}{k_i A_i \lambda_{t,j}} \\ R_{w,i} &= \frac{F}{2\pi k_i h_i \lambda_{t,N}}, \qquad R_{wi,i} = \frac{F + s_{pl}}{2\pi k_i h_i \lambda_{t,1}} \end{aligned} \tag{4.3}$$

where $\lambda_t = \lambda_w + \lambda_o$, $\Delta x = L/N$, $F = \ln(r_e/r_w) + s_0 = 2.5$ lumps the radial near-well geometry and a skin s_0, $R_{wi,i}$ is the injector's near-well resistance and s_{pl} is a plugging skin that grows at the rate set by a disturbance. The resistances are recomputed every hour, so a zone whose cells fill with water, the more mobile phase, flows more easily. At the connate saturation the tight zone delivers 6.78 and the thief zone 22.60 $\mathrm{m^3/d}$ per bar of drawdown, a productivity index J of 29.38 $\mathrm{m^3/d}$ per bar in the sense of Equation (1.1), and the thief zone's share of 76.9 % is its share of the permeability-thickness product.

4.3.3 The Producer and Its Choke

The wellbore is a capacitance C_{wb} between the zones and the choke, the tubing carries a constant flowing gradient $\Delta p_{col} = \rho_{col} g H = 127.5$ bar to the wellhead, and the choke passes the flow of

Equation (4.4).

$$\begin{aligned} C_{wb}\frac{\mathrm{d}p_{wf}}{\mathrm{d}t} &= q_1 + q_2 - q_{ch}, \qquad q_{ch} = K_{ch}\,C_v(x)\sqrt{\Delta p_{ch}} \\ \Delta p_{ch} &= \min\big(p_{wh} - p_{fl},\ (1-r_c)\,p_{wh}\big) \\ p_{wh} &= p_{wf} - \Delta p_{col} - k_f q_{ch}^2 \\ C_v(x) &= \frac{R^{x-1} - 1/R}{1 - 1/R} \end{aligned} \tag{4.4}$$

where $x \in [0,1]$ is the opening, $R = 20$ the rangeability of the modified equal-percentage trim and $r_c = 0.5$ the critical pressure ratio. The second branch of Δp_{ch} is critical flow, in which the choke no longer feels the flowline [17], and the friction term is solved in closed form with it. At the design point of 900 m³/d the wellhead is at 82.7 bar and the drop across the choke is 41.3 bar, exactly half, so the 20 bar flowline does not reach the well. The choke carries the inflow's composition at once, and the surface oil and water are $q_{ch}(1-w)/B_o$ and $q_{ch}w/B_w$, where $w = (q_1 f_{w,1} + q_2 f_{w,2})/(q_1 + q_2)$ is the water cut on a reservoir basis evaluated at each zone's producer cell. The actuators follow Equation (4.5).

$$\begin{aligned} \tau_{ch}\frac{\mathrm{d}x}{\mathrm{d}t} &= x_{cmd} - x, \qquad \tau_n\frac{\mathrm{d}n}{\mathrm{d}t} = n_{cmd} - n \\ \frac{\mathrm{d}s}{\mathrm{d}t} &= \mathrm{sat}_{\pm r_s}\left(\frac{s^* - s}{\tau_s}\right) \end{aligned} \tag{4.5}$$

where x_{cmd} and n_{cmd} are the controller's commands of Section 4.5, $\tau_{ch} = \tau_n = 1$ min, s is the thief zone's share of the injection, s^* the share requested, or the natural share when both valves are open, $\tau_s = 5$ min, and $\mathrm{sat}_{\pm r_s}$ limits the travel to r_s, 1 point of share per minute.

4.3.4 The Injector and Its Valves

Two variable-speed pumps in parallel deliver a discharge pressure $P_d = P_{sh}n^2 - k_p(q/a)^2$, where a is the share of pumping capacity in service, and the water-filled tubing adds $\rho_w g H = 201.1$ bar less its friction $k_t q^2$. The two inflow control valves realize a requested split by throttling the zone that needs less pressure, so the injection rate is the intersection of the pump curve with the system curve of Equation (4.6).

$$\begin{aligned} P_{sh}n^2 - k_p\left(\frac{q}{a}\right)^2 + \rho_w g H - k_t q^2 \\ = \max_i \big[p_i + s_i q\,(R_{a,i} + R_{wi,i})\big] \end{aligned} \tag{4.6}$$

where $s_2 = s$ and $s_1 = 1 - s$. With both valves open the zones share one bottomhole pressure p_{bh} and each takes $(p_{bh} - p_i)/(R_{a,i} + R_{wi,i})$. The sandface pressure of zone i is $p_{sf,i} = p_i + s_i q(R_{a,i} + R_{wi,i})$, and the higher one must stay 10 bar below the fracture pressure $p_{frac} = 360$ bar. An injection wellhead pressure of 145 bar therefore puts the higher sandface near 344.5 bar at the design rate, 5.5 bar short of the 350 bar limit. Chapter 2 holds a bottomhole pressure below a fracture pressure of the same kind while a well is drilled.

4.3.5 Sand and Erosion

Sand is an explicitly illustrative proxy. A zone sheds sand once its near-well drawdown $\Delta P_i = q_i R_{w,i}$ passes a critical drawdown that falls as water reaches its sandface, and the sand through the choke accumulates an erosion indicator, as shown in Equation (4.7).

$$\begin{aligned} \dot{m}_s &= \sum_i \kappa_i \left[\frac{\Delta P_i}{\Delta P_{c,i}} - 1\right]_+^2 q_i \\ \Delta P_{c,i} &= \Delta P_{c0,i}\,(1 - \beta_w S_{n,i}), \qquad \frac{\mathrm{d}E}{\mathrm{d}t} = \dot{m}_s\,\frac{2C_d^2\,\Delta p_{ch}}{\rho_{ch}} \end{aligned} \tag{4.7}$$

where $[\cdot]_+$ is the positive part, $S_{n,i}$ the normalized saturation of the producer cell, $\beta_w = 0.45$ the water weakening, $C_d = 0.85$ a discharge coefficient and ρ_{ch} the mixture density in the choke throat. The thief zone fails at 20 bar of drawdown when dry and at 11 bar when fully watered, and the tight zone at 120 bar. The erosion indicator is sand mass times the throat velocity squared, so choking back reduces the sand and speeds up whatever sand remains.

4.3.6 Numbers, Initialization and the Two Clocks

The parameters are listed in Table 4.2. A case in automatic mode starts at the operating point the fast loops would reach. The zone rates that meet the liquid target at the inflow's own water cut are found with every active limit acting as a floor under the bottomhole pressure, the choke opening that passes exactly that inflow is found by Newton iteration, and the pump speed is found by bisection for the injection target, so the default well starts at 56.49 % open with the pump at 65.75 %. The zones start at 250 bar and at the connate saturation with no polymer, except cases 6 to 8, which start from a waterflood of 30 or 40 days computed at reset with the wells solved quasi-steadily. At the reset of a plain flood the largest time derivative of a fast state is 1.4×10^{-12} in model units.

The 19 fast states are integrated with classical fourth-order Runge-Kutta at a one-minute step, the controllers scan once per step before the integration, and the historian logs six samples per simulated hour, the three clocks of Section 1.6.1. A fourth, slow clock is this model's own. It runs inside the post-step correction, once every 60 fast steps, on the volumes the fast clock recorded. Halving the fast step changes the bottomhole pressure 1.5 h into case 1 by 0.0024 bar and the liquid produced by 0.0048 %, and halving the hourly transport step moves a breakthrough by 0.04 d. These are among the application's 43 automated checks, which also cover the phase balances, the Welge front, the quasi-steady reference and bit-for-bit replay, and all 43 pass. The application runs from 0.05 to 120 times one simulated hour per second, one by default, so at the fastest setting a 180-day campaign plays in a nominal 36 s of wall time, which computation and drawing can extend. For the flowline and separator applications the export also carries a produced stream, with mass rates from oil and gas densities of 860 and 0.85 kg/Sm^3 and the gas-oil ratio, and a wellhead temperature $T_{wh} = T_{amb} + (T_{res} - T_{amb})\, q_{ch}/(q_{ch} + q_T)$ with $T_{res} = 90\,^\circ$C, $T_{amb} = 30\,^\circ$C and $q_T = 20\,\mathrm{m}^3/\mathrm{h}$. Nothing in the model reads them.

4.3.7 Assumptions

Two one-dimensional zones. There is no crossflow between the zones, no gravity, no capillary pressure and no areal sweep, so the damage a thief zone does appears entirely as early breakthrough. Real layers cross-flow where they touch, and a real pattern sweeps an area rather than a strip.

Lumped storage in a closed sector. Each zone's elastic storage sits at its midpoint, and no aquifer supports the sector, so injection is its only pressure support. A zone's pressure is therefore a volume average, and the pressure profile along it is not resolved.

A dead oil. The formation volume factor and the gas-oil ratio are constant and no gas leaves solution, so the model is invalid if the bottomhole pressure approaches a bubble point. The fluid is also given constant viscosities and Corey relative permeabilities.

An illustrative polymer. The viscosity multiplier and the irreversible Langmuir adsorption are the whole polymer model, with no shear thinning, no degradation, no permeability reduction and no inaccessible pore volume. A real solution thins in the fast flow near an injector, so the model overstates the injectivity a polymer costs.

Illustrative sand, erosion and fracture limits. None is calibrated to a formation, a choke trim or a completion, and the sand law is chosen to make sand a steep consequence of drawdown.

A lifted tubing. The flowing gradient is constant, so a well at a high water cut does not load up and die, and the tubing transit is folded into the meter delay. The injector hydraulics are quasi-steady and the valves realize any requested split.

Numerical smearing. Explicit upwind transport in 24 cells smears the fronts and brings the thief zone's breakthrough 1.25 days earlier than 48 cells do, which Section 4.4 quantifies.

Table 4.2 Parameters of the Well, the Sector and the Instruments. Visible Ones Can Be Changed in the Parameter Panel and Hidden Ones Cannot

Symbol	Meaning	Value	Unit	Class
Sector				
L, W	spacing, width	300, 300	m	hidden
N	cells per zone	24	—	hidden
h_1, ϕ_1	tight zone	12, 0.20	m, —	hidden
k_1	tight permeability	700	mD	visible
h_2, ϕ_2	thief zone	4, 0.25	m, —	hidden
k_2	thief permeability	7000	mD	visible
p_{init}	initial pressure	250	bar	visible
c_t	total compressibility	2×10^{-4}	1/bar	visible
F	$\ln(r_e/r_w) + s_0$	2.5	—	hidden
Fluids				
μ_o	oil viscosity	10	cP	visible
μ_w	water viscosity	0.5	cP	hidden
S_{wi}, S_{or}	connate water, residual oil	0.20, 0.25	—	hidden
n_w, n_o	Corey exponents	2.5, 2.0	—	hidden
k^0_{rw}, k^0_{ro}	end-point permeabilities	0.40, 0.85	—	hidden
B_o, B_w	formation volume factors	1.20, 1.02	—	hidden
—	gas-oil ratio, stream only	100	Sm^3/Sm^3	hidden
ρ_w	produced water density	1025	kg/m^3	hidden
Polymer and sand				
a, b	viscosity coefficients	$3 \times 10^{-3}, 3 \times 10^{-6}$	1/ppm, $1/ppm^2$	hidden
Γ_{max}	adsorption capacity	80	g/m^3 PV	visible
$C_{1/2}$	half-saturation	100	ppm	hidden
$\Delta P_{c0,1}$	tight critical drawdown	120	bar	hidden
$\Delta P_{c0,2}$	thief critical drawdown	20	bar	visible
β_w	water weakening	0.45	—	visible
κ_1, κ_2	sand coefficients	2, 3	kg/m^3	hidden
Producer				
H	reservoir depth	2000	m	hidden
ρ_{col}	flowing tubing density	650	kg/m^3	hidden
k_f	tubing friction	0.0015	$bar\,h^2/m^6$	hidden
K_{ch}	choke capacity, open	30	$m^3/h/\sqrt{bar}$	visible
R, r_c	rangeability, critical ratio	20, 0.5	—	hidden
C_{wb}	wellbore storage	0.25	m^3/bar	visible
τ_{ch}, r_x	choke lag, travel	1, 5	min, %/min	hid., vis.
C_d, ρ_{ch}	throat coefficient, density	0.85, 500	—, kg/m^3	fixed, hid.
Injector and facility				
P_{sh}, k_p	pump shut-off, droop	230, 0.006	bar, $bar\,h^2/m^6$	hidden
k_t	injection tubing friction	0.0008	$bar\,h^2/m^6$	hidden
τ_n	pump speed lag	1	min	hidden
$P_{d,max}$	discharge rating	200	bar	visible
r_s, τ_s	valve travel, lag	1, 5	%/min, min	vis., fixed
p_{frac}	fracture pressure	360	bar	visible
—	required fracture margin	10	bar	visible
—	water capacity, hourly	800	m^3/d	visible
—	polymer budget	60	t	visible
Operator's model and instruments				
—	thief pore volume, μ_o, c_t	1.35, 0.85, 1.25	×	hidden
—	MPFM delay, averaging	3, 2	min	vis., hid.
—	MPFM noise	2	%	visible
—	SD-103 averaging, noise	1, 0.03	min, g/s	hid., vis.
—	pressure gauge noise	0.2	bar	hidden
—	FT-201 noise	1	%	fixed
—	sand limit, hard	0.5	g/s	visible

Table 4.3 Time Scales of the Default Well and Sector

Process	Time scale
choke and pump actuator lags	1 min
multiphase meter delay and averaging	3 and 2 min
wellbore, choke fixed, Equation (4.8)	10.1 min
wellbore, choke flow held, C_{wb}/J	12.25 min
full choke stroke at 5 %/min	20 min
thief zone pressure, $c_t PV_2(R_{b,2}+R_{w,2})$	19.1 h
tight zone pressure, $c_t PV_1(R_{b,1}+R_{w,1})$	153 h
thief breakthrough in a plain flood (case 4)	27.5 d
polymer benefit at the producer (case 6)	41 d
campaign (challenge)	180 d

4.4 Time Scales of the Well and the Flood

The plant spans more than four orders of magnitude in time, as listed in Table 4.3, and each scale sets a different part of the control problem.

4.4.1 The Wellbore Seen From the Choke

A choke move changes the flow through the choke at once and the inflow from the zones only as the wellbore drains. Linearizing Equation (4.4) about a steady point gives the response of the choke flow to a step δx in the opening as the lead-lag of Equation (4.8).

$$\begin{aligned} \delta q_{ch}(t) &= \frac{\partial q_{ch}}{\partial x}\,\delta x\left[1-(1-g)\left(1-\mathrm{e}^{-t/\tau_{wb}}\right)\right] \\ g &= \frac{J}{J+\partial q_{ch}/\partial p_{wf}}, \qquad \tau_{wb} = \frac{C_{wb}}{J+\partial q_{ch}/\partial p_{wf}} \end{aligned} \tag{4.8}$$

At the design point $\partial q_{ch}/\partial p_{wf} = 0.263\,\mathrm{m^3/h}$ per bar against $J = 1.224\,\mathrm{m^3/h}$ per bar, so the rate jumps with the opening and relaxes to $g = 0.82$ of the jump with $\tau_{wb} = 10.1$ min. The share g is that of Equation (1.3), with the choke's sensitivity taken to the bottomhole pressure so that it includes the tubing friction. A manual step from 56.5 to 60 % takes the rate from 900 $\mathrm{m^3/d}$ to a peak of 1009 $\mathrm{m^3/d}$ 4 min later, behind the actuator lag, and settles it at 992 $\mathrm{m^3/d}$, against the 1017 $\mathrm{m^3/d}$ that the new opening passes at the old bottomhole pressure. When the rate loop holds the choke flow instead, the wellbore relaxes with $C_{wb}/J = 12.25$ min, which is the time constant the drawdown shows in closed loop. The relative gain of the trim, $(\mathrm{d}C_v/\mathrm{d}x)/C_v$, is 3.7 per unit of opening at the design point and 3.2 wide open, so the loop gain in $\mathrm{m^3/d}$ per percent grows roughly with the rate.

4.4.2 The Zone Pressures

With the bottomhole pressure fixed, a zone's pressure relaxes with $c_t PV_i(R_{b,i}+R_{w,i})$, 19.1 h for the thief zone and 153 h for the tight zone. The thief zone holds a small pore volume and delivers most of the flow, so its pressure moves eight times faster and is the first to fall when the voidage is not replaced. On the minute scale both are fixed pressures behind resistances, and in the first case a choke step moves them by at most 0.39 bar in 6 h.

4.4.3 The Fronts

The fractional-flow curves of the default fluids are shown in Figure 4.2. The Welge tangent from the connate saturation touches the water curve at a front saturation of 0.425 with $f_w = 0.742$, so a homogeneous layer breaks through after 0.303 pore volumes, with an average saturation of 0.503 behind the front. For the thief zone that is 27 300 $\mathrm{m^3}$ of injection. In the plain flood of case 4 the thief zone breaks through after 24 191 $\mathrm{m^3}$, 0.269 pore volumes, because upwind transport in 24 cells

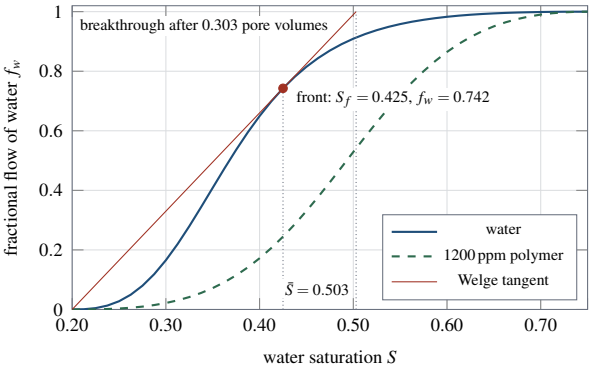

Figure 4.2 Fractional flow of water for the default fluids, without polymer (thick solid) and with 1200 ppm of polymer (thick dashed). The Welge tangent (thin straight line) from the connate saturation gives the front saturation of a plain waterflood and the volume injected at breakthrough.

smears the front ahead of the shock. The same flood with quasi-steady wells breaks through on day 28.25 with 24 cells, 29.5 with 48 and 30.5 with 96, later at each refinement by a smaller step, and with the thief zone at the tight zone's permeability it breaks through on day 85.0, three times later. Polymer moves the curve to the right, and 1200 ppm lowers the fractional flow at the waterflood front saturation from 0.742 to 0.244.

> **TRY IT IN THE APP** **4 · Thief-zone breakthrough**
>
> Load the case, set the speed to 24 hours per second on the Recovery view and press Run. Watch the solid fronts. The thief zone's water reaches the producer on day 27.5, when the front in the tight zone has traveled only 13 % of the way, and the thief zone's share of the injection keeps rising, from 76.9 % to 89.0 % by day 60.

4.4.4 Floors Under the Bottomhole Pressure

Every production limit acts as a floor under the bottomhole pressure, and the choke sets the last one, as on W-1 in Section 1.4.2. With the choke wide open the bottomhole pressure cannot fall below $\Delta p_{col} + p_{fl}$ plus the tubing friction and the small drop left across the open choke, $127.5 + 20 + 2.5 + 1.9 = 151.9$ bar on day 90 of the campaign. A thief zone at 170 bar can then supply only 7.0 bar of near-well drawdown, which is why the best campaigns of Section 4.6 run with the choke wide open for most of their length.

Pause. Try Exercise 4.1 before reading on.

4.5 Control Layers as Implemented

The control structure is shown in Figure 4.3. Its fast layer scans every minute and moves the choke and the pump, and its slow layer re-plans every five days and sets the rate target, the valve split and the polymer concentration.

4.5.1 Instruments

No controller reads a front, a saturation or a zone rate. The multiphase meter reads oil and water after a 3 min transport and processing delay and a 2 min first-order average, with multiplicative colored noise of 2 % and a 3 min correlation time. The acoustic sand detector averages over 1 min and adds 0.03 g/s of noise and a bias that a disturbance can set. The bottomhole, wellhead and injection wellhead gauges carry 0.2 bar of noise, and the injection meter 1 %.

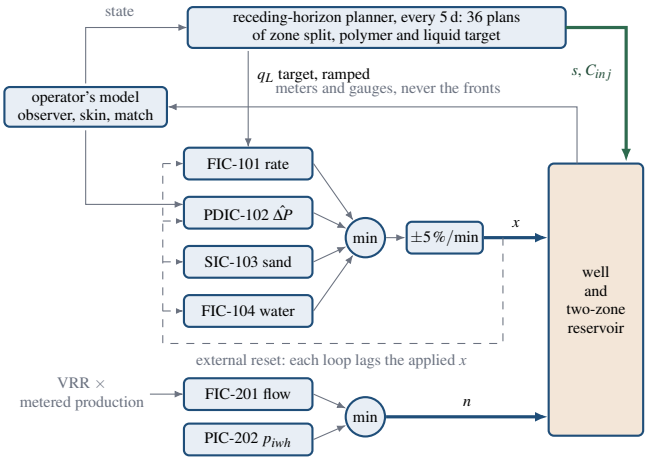

Figure 4.3 Control structure. Four PI loops share the choke through a low selector and a travel limit, and every loop's reset term follows the applied opening. FIC-201 and PIC-202 share the pump the same way. The operator's model supplies the drawdown estimate and the planner's state, and the planner sets the rate target, the zone split and the polymer.

4.5.2 The Choke Selector

Four PI loops share the production choke. FIC-101 holds the metered liquid at its target, PDIC-102 holds the larger of the two estimated drawdowns below the completion limit, SIC-103 holds the detected sand below its setpoint and FIC-104 holds the metered water below its setpoint. Each loop uses the framework's velocity-form PID of Section 1.6.2 with its integral switched off, and supplies its integral action as an external reset, as shown in Equation (4.9).

$$\begin{aligned} e_{j,k} &= \mathrm{SP}_{j,k} - \mathrm{PV}_{j,k}, \qquad P_{j,k} = K_{c,j}\, e_{j,k} \\ S_{j,k} &= S_{j,k-1} + (P_{j,k} - P_{j,k-1}) \\ u_{j,k} &= \mathrm{clamp}\big(S_{j,k} + I_{j,k-1},\ 0,\ 100\big) \\ x_k &= \mathrm{clamp}\big(\min_j u_{j,k},\ x_{k-1} - r_x \Delta t,\ x_{k-1} + r_x \Delta t\big) \\ I_{j,k} &= I_{j,k-1} + \big(1 - \mathrm{e}^{-\Delta t/\tau_j}\big)\big(x_k - I_{j,k-1}\big) \end{aligned} \tag{4.9}$$

where $\Delta t = 1\,\mathrm{min}$, $r_x = 5\,\%/\mathrm{min}$ is the choke's travel rate, and τ_j is the integral time $\tau_{I,j}$ of the selected loop and $\min(\tau_{I,j}, 2\,\mathrm{min})$ for the others. The accumulator $S_{j,k}$ equals $P_{j,k}$ for the gain in force, the clamp acts on the emitted output, and with no integral term the framework's back-calculation never runs, so the reset term I_j is each loop's only memory. For the selected loop away from its limits the last line becomes $I_{j,k} = I_{j,k-1} + (1 - \mathrm{e}^{-\Delta t/\tau_I}) K_c e_{j,k}$, a PI controller with an effective integral time of $\Delta t/(1 - \mathrm{e}^{-\Delta t/\tau_I})$, 8.51 min for the default 8 min. A loop that is not selected follows the applied opening with a lag of 2 min. At a fixed applied opening its output tends toward that opening plus its proportional term, the external-reset form of Equation (1.14) that Chapter 13 also uses on its pumps. During movement the finite tracking lag also affects which loop wins the selector. This feedback keeps the reset terms bounded while the loops wait [13], and Section 1.4.8 shows on W-1 how a waiting override without it breaks its limit. None of the six loops uses derivative action.

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

All six loops act in the same sense. With $e = \mathrm{SP} - \mathrm{PV}$ and positive gains they are reverse acting in ISA terms, since a wider choke raises the rate, the drawdown, the sand and the water together, and a faster pump raises both the injection and its wellhead pressure. The rate loop's setpoint ramps

Table 4.4 Loops, Setpoints and Default Tunings

Loop	PV	SP	K_c	τ_I
FIC-101	liquid, MPFM	900 m^3/d, ramped	0.020 % per m^3/d	8 min
PDIC-102	drawdown, model	50 bar	1.0 %/bar	10 min
SIC-103	sand, SD-103	0.30 g/s	20 % per g/s	10 min
FIC-104	water, MPFM	760 m^3/d	0.030 % per m^3/d	10 min
FIC-201	injection, FT-201	VRR 1.00	0.010 % per m^3/d	3 min
PIC-202	wellhead, PT-202	145 bar	0.30 %/bar	3 min

toward a new target at 15 m^3/d per minute, the way a well is beaned up, and in manual mode it tracks the metered rate so that the change to automatic starts with no error. In manual mode, and for any override switched off, each loop tracks by setting $I_j = x - S_j$, so its output equals the applied opening and both transfers are bumpless. The defaults are listed in Table 4.4. Against the choke step of Section 4.4, 33.4 and 26.2 m^3/d per point at once and at steady state, the default rate gain gives a loop gain of 0.67 at once and 0.52 at steady state, with an effective integral time a little shorter than the wellbore's 10.1 min.

4.5.3 The Injection Pump

FIC-201 holds the metered injection at its target by moving the pump speed, and PIC-202 takes the pump through a second low selector before the injection wellhead passes 145 bar, with the same external reset and no travel limit beyond the pump's lag. The target is either fixed or the metered production in reservoir volume, averaged over 1 h, times the voidage replacement ratio (VRR), and it is always the ratio while the planner runs. When PIC-202 holds the pump, or the pump is flat out, for most of a quarter hour while the injection runs more than 3 % short over the last hour, the application flags the target as unreachable, and it clears the flag when the limit lets go or the shortfall falls below 1 %. No hidden actuator makes up the difference. A dosing skid adds polymer at the commanded concentration and stops at the scan whose dose, reckoned on the metered injection with a 10 % allowance, would pass the 60 t budget.

4.5.4 The Operator's Model

The drawdown estimate and the planner's state come from an operator's model that runs the equations of Section 4.3 with three built-in errors, a thief pore volume 35 % too large, an oil viscosity 15 % too low and a compressibility 25 % too high. Every scan it allocates the metered production between its zones by its own resistances at the measured bottomhole pressure, which yields the drawdown estimate, and every hour it moves its fronts by the metered injection, split as commanded, and by its allocated production. A pressure observer then moves its zone pressures toward the rate the producer shows at its measured bottomhole pressure, a fifth of the correction and at most 2 bar per hour, and an injectivity monitor reads any gap between the injection gauge and the model as injector skin, a fifth at a time. Once a day, if the model's water cut differs from the meter's by more than two points while the model's tight front is short of 80 % of the way, the thief pore-volume multiplier moves toward the measurement by at most 5 %, within 0.5 to 2, and the thief profile is remapped to it. Nothing else is adjusted, and the model never sees how the production divides between the zones. It starts from the plant's initial state, or from its own run of the preflood of cases 6 to 8 with its own parameters. This is a bounded, one-parameter version of the continuous model updating that closed-loop reservoir management performs with an ensemble of models [19].

4.5.5 The Receding-Horizon Planner

Every five days, and at reset, the planner coarsens the operator's model to 12 cells per zone and forecasts 36 candidate plans, four splits (both valves open or a thief share of 65, 50 or 40 %) by three polymer concentrations (0, 600 and 1200 ppm) by three liquid targets (900, 1100 and 1300 m^3/d),

Table 4.5 Forecasts of One Fixed Policy Over 180 Days, a Thief Share of 40 % With 1200 ppm and a 1300 m^3/d Target

Model	Cells, step	Oil (Sm^3)
dynamic plant	24, 1 min	103 120
quasi-steady, true reservoir	24, 6 h	103 093
quasi-steady, true reservoir	12, 1 d	99 290
operator's model, as the planner uses it	12, 1 d	113 083
true reservoir, thief pore volume error only	12, 1 d	106 716
true reservoir, oil viscosity error only	12, 1 d	103 837
true reservoir, compressibility error only	12, 1 d	99 637

each held over a horizon H with the wells solved quasi-steadily in daily steps. In each step the choke delivers the target unless a drawdown, sand, water or wide-open floor under the bottomhole pressure binds first, the injector replaces the voidage unless the injection wellhead limit or the pump binds, the zone pressures advance semi-implicitly with the production following the new pressure, and the dosing stops at the budget. The final forecast interval uses a reduced concentration if its full dose would exceed the remaining polymer. The injection constraint applies PIC-202 directly on wellhead pressure, including the rate-dependent tubing friction term. The objective is the forecast oil less 0.01 Sm^3 per kilogram of polymer, the plan in force is kept unless another beats it by 0.5 %, and only the targets of the chosen plan are applied until the next re-plan. The horizon is the lesser of its setting and the time left in a campaign, and never less than one day. The planner is an enumeration over held policies, not the quadratic program of Section 1.6.3. Volume 9, Chapter 7 plans the releases of another slow store days ahead on forecasts whose error grows with lead time.

4.5.6 Four Models Kept Apart

The plant is the 24-cell reservoir with its wellbore on the minute clock. The operator's model has the same structure with three errors and its own state, the planner's forecast is a 12-cell copy of the operator's model with quasi-steady wells in daily steps, and the fixed-choke reference of Section 4.6 is a quasi-steady run of the true reservoir in 6 h steps. None of the controllers uses an identified model, and the PI loops carry no model at all. Forecasting one fixed policy for 180 days, a thief share of 40 % with 1200 ppm and a 1300 m^3/d target, separates their errors, as listed in Table 4.5. The quasi-steady shortcut is within 0.03 % of the dynamic plant, the planner's coarse grid makes the forecast 3.7 % pessimistic, and the operator's model makes it 13.9 % optimistic. Taken one at a time, the thief pore-volume error accounts for 7.5 %, the oil viscosity error for 4.6 % and the compressibility error for 0.3 %. On its own model the planner ranks a 50 % thief share with 1200 ppm and a 1300 m^3/d target first on day 0, at 113 511 Sm^3, and the same plan at a 40 % share third, at 113 083 Sm^3. On the true reservoir the 40 % plan ranks first, at 99 290 Sm^3, and the planner's choice fourth, at 98 382 Sm^3. The model's errors raise every forecast by far more than they separate these two plans, which differ by 0.38 % on the operator's model and by 0.9 % on the true reservoir.

4.6 Scenarios and Measured Results

Table 4.6 names each case, its event and its lesson, and Table 4.7 collects the measured results of the eight cases. Each begins from the defaults of Table 4.4, with both valves open, no polymer, the planner off and the history match and every override on, apart from the changes named, and the numbers come from headless runs of the engine with its default random seed (Section 1.6.1). Scheduled events act at their nominal minute boundary, and response delays below are measured from that boundary. Case clocks start at reset, so days reported for cases 6 to 8 exclude their initialization waterflood. Reservoir water cuts are used for transport results, while the CSV and water-cut trends use surface volumes. Quantities that the meter noise moves are averaged over the stated window.

Table 4.6 Cases, Events and Lessons

#	Case and event	Lesson
1	target 900 to 1100 m^3/d and back	a choke sets drawdown
2	injection plant down 6 to 96 h	a choke cannot make pressure
3	target 1400 m^3/d	sand comes with drawdown
4	plain waterflood, valves open	the thief zone takes the water
5	thief share 50 % at 12 h	sweep costs injection pressure
6	1000 ppm after a 30 d flood	the benefit arrives late
7	injector skin +5 per day	an unreachable target
8	target 1200 m^3/d at 69 % water	more liquid is mostly water
9	180-day campaign, fixed-choke score	graded challenge

Table 4.7 Measured Results of the Eight Cases

Case	Result
1, 2–3.5 h means	choke 63.6 %, p_{wf} 205.1 bar, ΔP_2 19.7 bar of 20.0
1, after the step	true 1000 m^3/d at 11 min, meter at 16 min
2, no injection	choke open at 55.6 h, 427 m^3/d at 96 h
2, 48 h	p_2 184.3 bar, p_1 234.1 bar
3, SIC-103 on, 3–6 h	1205 m^3/d, 0.30 g/s, peak 0.45, 52 kJ
3, SIC-103 off	1408 m^3/d, 3.24 g/s, 416 kJ
4, breakthrough	day 27.5, tight front at 0.13
4, day 60	water cut 77.8 %, 33 544 Sm^3, model PV $\times 0.96$
5, PIC-202	holds day 1.6 to 3.4, margin $\geq$ 15.0 bar
5, day 60	breakthrough day 44.2, 38.3 %, 49 212 Sm^3
6, water cut	72.2 % on day 41.2, 50.1 % on day 60
6, day 80	24 130 Sm^3 against 13 934 Sm^3 without
7, PIC-202	takes the pump day 14.5, flag day 15.3
7, day 30	757 m^3/d of 943, $\bar{p}$ 216.6 bar, margin $\geq$ 14.3
8, FIC-104 on, 24–48 h	1028 m^3/d, oil 270, water 758
8, FIC-104 off, 24–48 h	1197 m^3/d, oil 314, water 883, 46 of 48 h over

4.6.1 Minutes: The Choke, the Wellbore and the Sand

Case 1 raises the target from 900 to 1100 m^3/d. The setpoint ramps for 13.3 min, the true rate passes 1000 m^3/d 11 min after the change and the meter 16 min after it, and the 10 min mean of the rate is within 2 % of the new target after 40 min. Averaged from 2 to 3.5 h, the choke sits at 63.6 % against 56.5 % before, only 7 points further open for 22 % more rate, and the bottomhole pressure has fallen from 213.2 to 205.1 bar. The thief zone's drawdown has risen from 16.3 to 19.7 bar, against a critical drawdown of 20.0 bar, while the zone pressures have barely moved. The meter noise makes the rate loop wander, so the choke reads 63.11 % at 3 h and 63.9 % in the same run without noise.

Case 3 asks for 1400 m^3/d. The true sand passes 0.01 g/s 32 min after the change, at a thief drawdown of 20.5 bar, and SIC-103 takes the choke 9 min later. From 3 to 6 h the well makes 1205 m^3/d with 0.30 g/s of sand, peaking at 0.45 g/s while the loops hand over, against the 0.5 g/s limit. With SIC-103 off the choke delivers 1408 m^3/d with 3.24 g/s, 17 % more liquid for eleven times the sand, 50.2 kg of it in 6 h. The erosion indicator reaches 416 kJ against 52 kJ, although the throttled choke of the first run passes its smaller sand at a higher throat velocity, with v^2 of 9758 against 8132 m^2/s^2 at 6 h.

Case 8 starts 40 days into a flood at a reservoir water cut of 69.3 %, making 245.8 Sm^3/d of oil in 900 m^3/d of liquid, and raises the target to 1200 m^3/d. FIC-104 takes the choke 17 min after the change, when the metered water reaches 760 m^3/d, and over the second day the well makes 1028 m^3/d with 270 Sm^3/d of oil and 758 m^3/d of water, with the hourly water peaking at 773 m^3/d against the 800 m^3/d capacity. The oil has risen by 24 Sm^3/d for 104 m^3/d more water. With the override off the well makes 1197 m^3/d with 314 Sm^3/d of oil and 883 m^3/d of water over the second day, above the

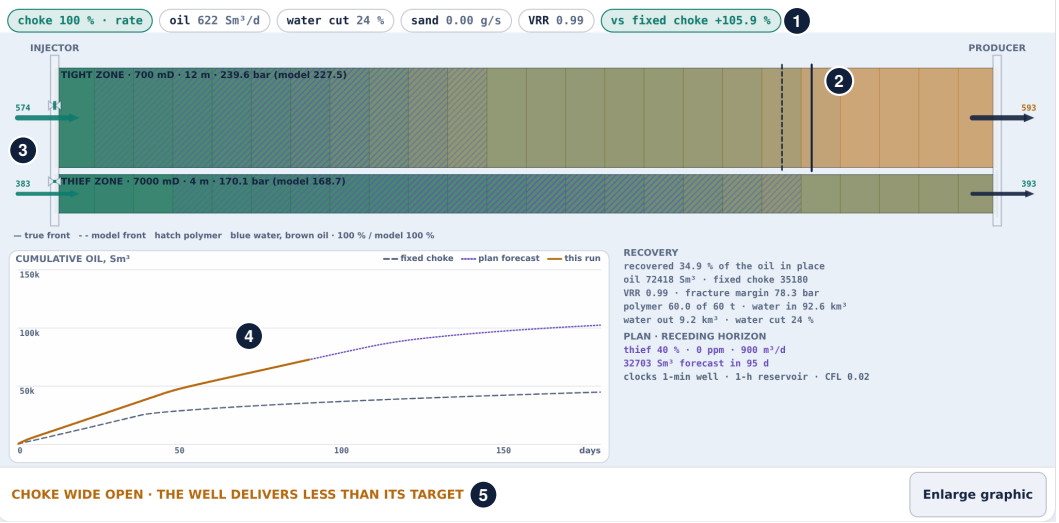

Figure 4.4 Campaign challenge, listed in the scenario menu as Challenge · 180-day campaign, with the planner on and a 180-day horizon, on the Recovery view with the clock at 90d 00:00 (2160 h), cropped to the plant view. (1) The badge reads +105.9 % against the fixed choke, and the first chip shows the choke wide open. (2) The tight zone's true front (solid) near 0.8 of the way, with the operator's model front (dashed) just behind it. Below it the thief zone is watered out from the injector to the producer, polymer hatching shows in both zones, and the labels give the thief zone near 170 bar against 240 bar in the tight zone. (3) 574 m^3/d of injection enter the tight zone and 383 m^3/d the thief zone. (4) The cumulative oil plot shows 72 418 Sm^3 against 35 180 Sm^3 for the fixed choke, with the plan's forecast dotted to day 180. (5) The status line reports the choke wide open with the well short of its target. The ledger gives a recovery of 34.9 % of the oil in place, 60.0 t of the 60 t of polymer used and the plan in force, a 40 % thief share with no polymer and a 900 m^3/d target.

capacity in 46 of the 48 hours.

4.6.2 Days: Pressure Support

Case 2 trips the injection plant at 6 h. The thief zone falls from 249.9 to 218.7 bar in 18 h and to 184.3 bar at 48 h, while the tight zone falls only to 234.1 bar, as their time constants in Table 4.3 predict. The rate loop opens the choke to hold 900 m^3/d until it is wide open at 55.6 h, and the rate then falls with the pressure, to 427 m^3/d at the restart at 96 h. The returning water goes mostly into the depleted thief zone, and the rate is back within 2 % of its target at 146.8 h. With a voidage ratio of 1 the lost pressure is never rebuilt, and at 240 h the tight zone is at 191.9 bar and the thief zone at 199.2 bar. Over 10 days of a plain flood, a ratio of 1 holds the average pressure within 0.03 bar while 0.8 lets it fall 35.1 bar.

Case 7 grows the injector's skin by 5 units a day after a 30-day flood. The sandface climbs from 266.0 bar to the 345 bar that PIC-202 holds from day 14.5, the target is flagged unreachable on day 15.3, and by day 30 the injection is 757 m^3/d against a target of 943 m^3/d, a VRR of 0.80, while the average reservoir pressure has fallen from 249.6 to 216.6 bar. The fracture margin never falls below 14.3 bar. The injectivity monitor infers a skin of 153.6 against 150.0 true on day 30, and the model's zone pressures are within 1.8 bar of the truth.

4.6.3 Weeks: Where the Water Goes

Case 4 is a plain flood with both valves open. The thief zone's share of the injection rises from 76.9 % to 89.0 % by day 60, because a watered-out zone flows more easily, and its water reaches the

producer on day 27.5, when the tight zone's front has traveled 13 % of the way. The reservoir water cut is 46.8 % on day 30 and 77.8 % on day 60, the oil falls from 444 to 173 Sm^3/d, and 33 544 Sm^3 are produced by day 60. The operator's model, with its thief pore volume 35 % too large, places the thief front at 66 % of the way on day 25 when it is at 89 %, and its own water cut first passes 2 % on day 31.6. The history match brings the multiplier within 10 % of the truth on day 32.0 and to 0.96 by day 60. The arriving water also weakens the sandface, the thief zone's critical drawdown falls from 20.0 to 15.5 bar, by about day 58.3. SIC-103 holds an earlier sand episode on day 28.1 to 0.40 g/s.

Case 5 sets the valves to a 50 % thief share at 12 h, and the tight zone's injection doubles from 260 to 531 m^3/d within two hours. The sandface rises from 290 to 339 bar in the first day, PIC-202 first takes the pump on day 1.4, and from day 1.6 to 3.4 it mostly holds the pump with the injection 2.0 % short on average, which the application flags as unreachable from day 2.1 to 3.2. The pressure then eases to 298 bar by day 60 as the near-injector rock fills with more mobile water. The tight zone's pressure rises to 267.3 bar and the thief zone's falls to 206.5 bar, so the producer draws more from the tight zone, whose drawdown rises from 16.3 to about 35 bar. The operator's model reads that drawdown as 30.2 bar on day 20 when it is 35.3 bar, so an override acting on it would act late. The thief zone's water arrives on day 44.2 instead of 27.5, and on day 60 the water cut is 38.3 % against 77.8 % and the oil produced 49 212 Sm^3 against 33 544 Sm^3, 47 % more.

Case 6 adds 1000 ppm of polymer at 12 h, 30 days into a flood that has just broken through. The sandface rises from 266.0 to 283.5 bar by day 60 and the thief zone's share falls from 83.7 to 78.1 %, but for 41 days the water cut keeps rising, to 72.2 % on day 41.2 against 80.6 % in the same flood without polymer. Then the oil the polymer has displaced reaches the producer ahead of it, the water cut falls to 50.1 % on day 60, and the hourly oil peaks at 419 Sm^3/d near day 61 against 128 Sm^3/d without polymer on day 60. The polymer front is still at 88 % of the thief zone on day 60, held back by adsorption. By day 80 the case has made 24 130 Sm^3 against 13 934 Sm^3, for the whole 60 t budget.

> **TRY IT IN THE APP** **6 · Polymer arrival delay**
>
> Load the case, set 24 hours per second on the Recovery view and press Run. The hatched polymer enters both zones at 12 h, yet the water cut keeps rising for 41 days, to 72.2 % on day 41.2. Then the oil the polymer displaced reaches the producer ahead of it, and the water cut falls to 50.1 % on day 60.

4.6.4 Months: The Campaign

The challenge asks for 98 000 Sm^3 in 180 days, 2.2 times fixed-choke waterflooding, with the true sand never above 0.5 g/s, the hourly water never above 800 m^3/d, the higher sandface at least 10 bar below the fracture pressure and no more than 60 t of polymer. The reference holds the choke at 45 %, the largest opening in 5 % steps that keeps a fixed-choke run inside the sand and water limits, with both valves open, no polymer and the pump at the speed that replaced the day-0 voidage. It makes 44 547 Sm^3, and at 50 % its water would reach 921 m^3/d. The full dynamic model run in manual mode with the same choke and pump makes 44 684 Sm^3, within 0.31 %. The configurations are compared in Table 4.8, and Figure 4.5 sets their oil on day 30 and on day 180 side by side.

As shipped, with a 1300 m^3/d target, both valves open, no polymer and the planner off, the campaign is 48 % ahead of the reference on day 30 and only 7.1 % ahead on day 180, having recovered 23.4 % of the oil in place. The rate loop holds its choke for only 0.008 % of the campaign. FIC-104 holds it for 85.5 % of the campaign and SIC-103 for the remaining 14.5 %, at a mean drop of 51.4 bar across the choke against 4.0 bar for the 180-day planner, and the well passes 688 kg of sand against 17 kg, so its erosion indicator reaches 7854 kJ against 146 kJ. With the planner on and a 180-day horizon, the plan chosen on day 0 sends half the water into the tight zone with 1200 ppm of polymer and a 1300 m^3/d target. The budget is spent on day 44.9, the day-45 re-plan keeps 60 % of the water in the tight zone with no polymer and a 900 m^3/d target, and a re-plan on day 120 raises the target to 1300 m^3/d. The campaign makes 103 912 Sm^3, 2.33 times the reference, with a sand peak of

Table 4.8 The Campaign Over 180 Days. Oil in Sm^3, Worst Sand in g/s, Hourly Water in m^3/d and Worst Fracture Margin in bar

Configuration	Day 30	Day 180	RF	Sand	Water	Margin	Pass
fixed choke	19 375	44 547	22.2	0	746	—	—
shipped	28 698	47 731	23.4	0.43	775	55.2	no
planner, 30 d	37 620	89 270	43.7	0.39	774	16.4	no
planner, 90 d	31 762	101 753	49.4	0.37	773	15.1	yes
planner, 180 d	29 183	103 912	50.3	0.36	773	14.9	yes
by hand[a]	25 054	103 120	49.9	0.37	770	14.4	yes

RF, recovery factor in percent. The fixed-choke row is the quasi-steady reference, with no margin record. [a] thief 40 %, 1200 ppm from day 0

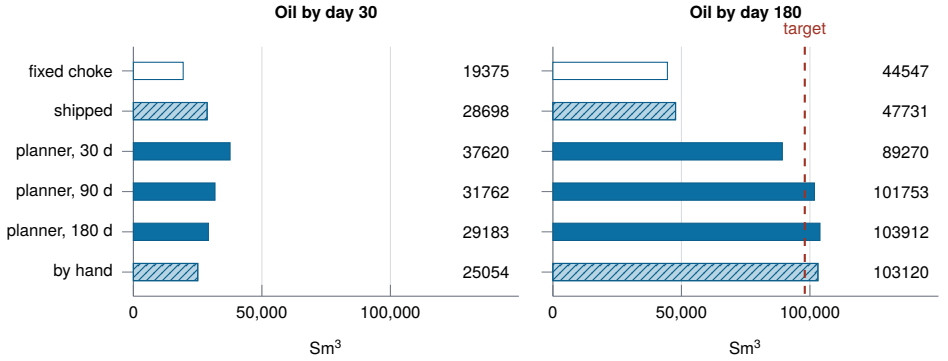

Figure 4.5 The campaign of Table 4.8 on one scale. White is the fixed-choke reference, hatched bars run with the planner off and solid bars with it on, and the dashed line is the 98 000 Sm^3 target. Notice that the order on day 30 is not the order on day 180: the 30-day planner leads early and misses the target, while the hand-set policy starts slowest of the three that pass it.

0.36 g/s, a water peak of 773 m^3/d, a worst margin of 14.9 bar and 60.0 t of polymer, and it recovers 50.3 % of the oil in place.

The horizon decides the plan. A 30-day horizon starts with a 65 % thief share and no polymer, moves to 50 and 40 % on days 10 and 30, and never doses polymer, because inside 30 days its best polymer plan forecasts only 233 Sm^3 more oil than the plan it chose, less than the 562 Sm^3 charged for its 56 t of polymer. It is 29 % ahead of the 180-day planner on day 30 (37 620 Sm^3 against 29 183 Sm^3) and finishes at 89 270 Sm^3, short of the target. A 90-day horizon makes 101 753 Sm^3. Among the three horizons the ranking on day 30 is exactly the reverse of the ranking on day 180. A fixed setting of a 40 % thief share with 1200 ppm, chosen by hand at reset, makes 103 120 Sm^3, within 0.8 % of the planner, so the planner's value here is that it finds such a policy from a model with errors in it, not that it beats the best fixed policy. In the three passing runs the choke is wide open for most of the campaign, from 55.5 % of the time with the 90-day planner and 76.0 % with the 180-day planner to 92.7 % by hand, because moving water into the tight zone draws the thief zone down to about 170 bar and leaves the rate to the reservoir.

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

The run of Figure 4.4 is the 180-day planner on day 90, when it has made 72 418 Sm^3 against 35 180 Sm^3 for the reference, 105.9 % more. The thief zone is watered out, its water having broken through on day 47.5 against 27.5 in the plain flood, and the tight zone's true front is at 0.80 of the way with the model's at 0.77. The history match began on day 48 and took the thief multiplier from 1.35 to 0.94 in eight daily steps, and the model's own water cut first passed 2 % on day 53.0. The liquid

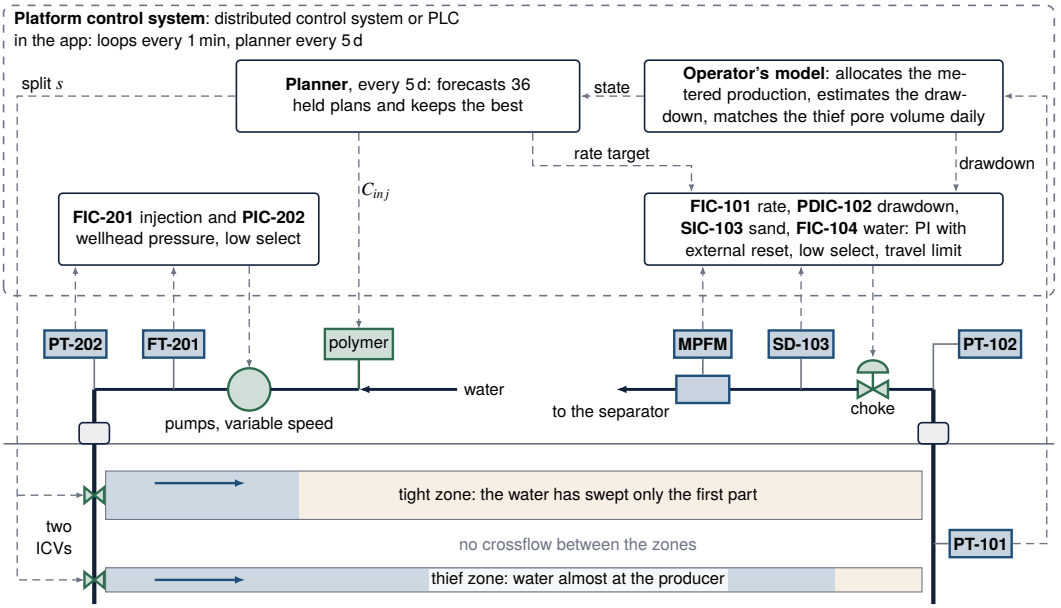

Figure 4.6 The loops of the smart well simulator with their hardware. Lettered boxes are sensing devices with the chapter's tags: the bottomhole gauge PT-101, the wellhead gauge PT-102, the clamp-on sand detector SD-103, the multiphase meter (MPFM), the injection meter FT-201 and the injection gauge PT-202. The production choke, the injection pumps, the polymer skid and the two inflow control valves (ICVs) are the final elements, and dashed arrows are signals. Notice that four loops reach the choke through one selector, and that the planner writes only targets: the rate target, the valve split and the polymer concentration. Not to scale, and not a piping diagram.

is 857 m^3/d against the plan's 900 m^3/d, with the choke wide open and the liquid rate loop saturated 96.7 % of the time so far. The badge and ledger water cut, 24 %, is on a reservoir basis, while the water-cut trend shows all three series on the surface basis: true 27.5, metered 26.8 and the model's 28.1 % at the latest historian sample.

> ✦ **DID YOU KNOW? The Slow Starter Wins the Campaign**
>
> From the simulator (Table 4.8). On day 30 the shipped settings have made 28 698 Sm^3 of oil, ahead of a fixed policy set by hand at reset, a 40 % thief share with 1200 ppm of polymer, at 25 054 Sm^3. On day 180 the hand-set policy has made 103 120 Sm^3 and the shipped settings 47 731 Sm^3.

4.7 Technology in Practice

Figure 4.6 draws the two selectors, the operator's model and the planner of Section 4.5 with the hardware of a producer and an injector.

What Measures It

The rate loop and the water override both read the multiphase meter, so its delay and noise sit inside two of the four choke loops. The sand override reads an acoustic detector clamped to the outside of the flowline. Sand grains striking the inside of a pipe at a bend give off acoustic energy that the detector picks up, and an intrusive variant instead measures the erosion of a probe placed in the stream [20]. In one field case the sand signals rose with the water cut, and a step-rate test found the liquid rate above which they began, so the rate was cut back below it. The same review warns that

such signals must be backed by other data [20]. The model's SD-103 averages over 1 min and carries noise and a bias that a disturbance can set.

> **TECHNOLOGY Sensor: The Multiphase Flow Meter**
>
> The conventional way to measure a well is a test separator, which splits the oil, gas and water and meters each phase alone. It is reliable and accurate but too slow to feed a control loop [21]. A multiphase meter measures the mixture in the line. In one common design gamma rays cross the pipe, and their attenuation gives the density and so the fraction of each phase, while a venturi or an orifice beside it gives the velocity. A meter with a radioactive source needs specially trained staff, and its pressure sensors drift from their zero over time and need recalibration [21]. The chapter's meter reads after a 3 min delay and a 2 min average, with 2 % colored noise.

What Moves It

The choke is the only final element of the fast production layer. On the injector, variable-speed pumps set the rate, the dosing skid sets the polymer concentration, and two downhole inflow control valves split the water between the zones, the smart completion of Section 4.2. The model lets the valves realize any requested split, moving one point of share per minute.

> **TECHNOLOGY Final Element: The Production Choke**
>
> Almost all flowing wells carry a wellhead choke to regulate the rate and to prevent water and gas coning and sand problems [22]. In critical flow the rate through it reaches a maximum that a change of pressure downstream no longer moves. The long-used correlation of 1954 holds only when the upstream pressure is at least 70 percent above the downstream one [22]. Produced sand erodes chokes, completions and surface equipment [20]. Equal-percentage trim counters the droop of a falling pressure drop, but its textbook law never reaches zero flow [23]. The chapter's modified law closes at zero, with a rangeability of 20, a 1 min lag and 5 %/min of travel.

Where the Controller Runs

Wells gathered at one site can be run from a distributed control system, which controls production within one geographic location and whose supervisory level oversees the controllers of each local process [24]. Loops in such systems run with cycle times from milliseconds to minutes [24]. The app's one-minute scan sits inside that range and well below the wellbore's 10.1 min time constant. The chapter's structure maps onto four PID blocks whose reset is fed back from the applied opening, a low-select block and a rate limiter on the choke, and the same pattern on the pump. The operator's model and the planner sit above them in a supervisory layer that writes setpoints once every five days.

> **IN PRACTICE The Injection Pressure Limit Is a Rule**
>
> Federal rules for Class II injection wells in the United States, which cover enhanced recovery, require a maximum wellhead injection pressure. It is calculated so that injection does not start new fractures or extend existing ones in the confining zone [25]. The operator must observe the injection pressure, flow rate and cumulative volume at least monthly for enhanced recovery, report a summary each year and demonstrate mechanical integrity at least every five years [25]. PIC-202 plays a similar part in the chapter. It takes the pump before the injection wellhead passes 145 bar, so the higher sandface stays at least 10 bar below the 360 bar fracture pressure.

What Is Modernizing It

Three developments touch this chapter. The first is the virtual flow meter, which infers the phase rates from ordinary field instruments and empirical algorithms. It costs far less than a multiphase meter but needs much calibration before start-up, and a 2019 review expects it to progress [21]. The second is fiber-optic distributed acoustic sensing along the whole well, which gives sand data in real time and

locates where the sand enters, something a surface detector cannot do [20]. The third is closed-loop reservoir management. Gradient-free methods such as the ensemble Kalman filter connect to any reservoir simulator for history matching and optimization [26]. The chapter's daily match of one pore-volume multiplier is the smallest version of that loop, and its 13.9 % forecast error shows what the full version has to remove.

4.8 Limitations

The reservoir is two strips without crossflow or gravity, so the thief zone costs recovery only through early breakthrough, and the fronts are smeared by 24 upwind cells, which bring the thief breakthrough 1.25 days earlier than 48 cells.

The sand law, the erosion indicator, the fracture pressure and the polymer are illustrative, and none has been calibrated to a formation, a choke or a field fluid. The targets of the challenge are the application's own design choices.

> ⚠ **OPEN CHALLENGE** **When Will the Sand Come?**
>
> Sand control remains an ongoing challenge in sand-rich reservoirs [20]. A 2002 study of deep high-pressure wells found that criteria based on critical drawdown, bottomhole pressure, depletion or fluid velocity missed the outcome by a large margin. All but one of the wells made no sand until water arrived, which the study traced to the loss of a capillary tension of about 1 psi that had held the loose grains together [18]. The chapter's sand law is a critical drawdown that falls as water arrives, made steep on purpose and calibrated to no formation. Progress would be a sand onset predicted for a real well and confirmed by measurement before it happens.

The operator's model corrects one parameter, the thief pore volume, and only while its tight front is short of 80 % of the way, so it cannot repair its oil viscosity or its compressibility and it cannot see the production split. In case 5 its drawdown estimate reads 5.1 bar low. The planner enumerates up to 36 plans held constant over the horizon, so it cannot shape a polymer slug in time. Its daily forecast caps the last dose at the remaining budget and uses that same dose in transport and cost. A partial final dose is spread over the daily interval, so the exact within-day shutoff time is unresolved. Model predictive control over the same equations, with moves that vary along the horizon and an estimator that updates all three uncertain parameters, is the natural extension [9, 19].

> ⚠ **OPEN CHALLENGE** **A Reservoir Model Nobody Can Fully Check**
>
> Closed-loop reservoir management couples automatic history matching with production optimization, and both are large optimization problems [26]. A 2015 review found most studies academic and far from field application. Field engineers care about processing time, the veracity of decisions and ease of use, and one simulation of a large reservoir can take several days. History matching is underdetermined: many parameter sets fit the data, and some of them forecast the future badly [26]. The chapter's operator's model matches its water cut yet forecasts 13.9 % high. Progress would be a loop that states the uncertainty of each forecast beside each plan, at a cost a field team can run.

4.9 Conclusion

This chapter studied a producer and an injector on a two-zone reservoir on one simulated clock, with fast loops that protect the well and a slow planner that decides the sweep. The wellbore answers a choke move with a lead-lag of 10.1 min, the thief zone's pressure moves in 19.1 h and the tight zone's in 153 h, and a plain flood breaks through the thief zone on day 27.5 after 0.269 pore volumes against the Welge value of 0.303. Four PI loops in external-reset form share the choke without windup, so sand can be held at 0.30 g/s and water inside an 800 m^3/d capacity, and an injector at its fracture limit lets the pressure fall to 216.6 bar rather than relax the limit. Holding the thief zone to half the water makes 47 % more oil by day 60, and the 180-day planner makes 2.33 times the fixed-choke oil where

the shipped settings make 1.07 times and a 30-day horizon 2.00 times. The planner's forecasts are 13.9 % optimistic from the operator's model errors and 3.7 % pessimistic from its coarse grid, and a fixed policy chosen by hand comes within 0.8 % of it. The sand law, the polymer and the fracture limit are illustrative, and the targets are the application's own design choices. A fuller model would add crossflow, a shear-thinning polymer, tubing lift and an estimator of every uncertain parameter. The simulator is an accessible place to see that the oil a well makes this month and the oil a flood recovers are different quantities.

KEY POINTS

- A choke changes a well's rate within minutes, but the water that limits its oil was injected weeks earlier, and a permeable layer decides where it goes.
- The choke sets the bottomhole pressure, not the rate, so when the injection plant trips the rate loop holds its target for 49.6 h and then runs out of travel.
- Four PI loops in external-reset form share the choke through a low selector without windup.
- The 180-day planner makes 2.33 times the fixed-choke oil where the shipped settings make 1.07 times, and a fixed policy chosen by hand comes within 0.8 % of it.

Where else this applies. Volume 4, Chapter 14 meets the same shape, a fast actuator drawing on a store that refills slowly, and Chapter 13 uses the same external-reset selectors on its pumps.

References

[1] S. E. Buckley and M. C. Leverett, "Mechanism of fluid displacement in sands," *Transactions of the AIME*, vol. 146, no. 1, pp. 107–116, 1942.

[2] H. J. Welge, "A simplified method for computing oil recovery by gas or water drive," *Journal of Petroleum Technology*, vol. 4, no. 4, pp. 91–98, 1952.

[3] W. E. Stiles, "Use of permeability distribution in water flood calculations," *Journal of Petroleum Technology*, vol. 1, no. 1, pp. 9–13, 1949.

[4] C. A. Glandt, "Reservoir management employing smart wells: A review," *SPE Drilling & Completion*, vol. 20, no. 4, pp. 281–288, 2005.

[5] D. J. Pye, "Improved secondary recovery by control of water mobility," *Journal of Petroleum Technology*, vol. 16, no. 8, pp. 911–916, 1964.

[6] K. S. Sorbie, *Polymer-Improved Oil Recovery*. Springer Netherlands, 1991.

[7] D. R. Brouwer and J.-D. Jansen, "Dynamic optimization of waterflooding with smart wells using optimal control theory," *SPE Journal*, vol. 9, no. 4, pp. 391–402, 2004.

[8] E. Manrique, M. Ahmadi, and S. Samani, "Historical and recent observations in polymer floods: an update review," *CT&F Ciencia, Tecnología y Futuro*, vol. 6, no. 5, pp. 17–48, 2017.

[9] J.-D. Jansen, O. H. Bosgra, and P. M. J. Van den Hof, "Model-based control of multiphase flow in subsurface oil reservoirs," *Journal of Process Control*, vol. 18, no. 9, pp. 846–855, 2008.

[10] G. M. van Essen, P. M. J. Van den Hof, and J. D. Jansen, "Hierarchical long-term and short-term production optimization," *SPE Journal*, vol. 16, no. 1, pp. 191–199, 2011.

[11] B. Foss, B. R. Knudsen, and B. Grimstad, "Petroleum production optimization – a static or dynamic problem?" *Computers & Chemical Engineering*, vol. 114, pp. 245–253, 2018.

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

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

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

[15] V. Alvarado and E. Manrique, "Enhanced oil recovery: An update review," *Energies*, vol. 3, no. 9, pp. 1529–1575, 2010.

[16] R. G. Agarwal, R. Al-Hussainy, and H. J. Ramey, "An investigation of wellbore storage and skin effect in unsteady liquid flow: I. Analytical treatment," *Society of Petroleum Engineers Journal*, vol. 10, no. 3, pp. 279–290, 1970.

[17] F. E. Ashford, "An evaluation of critical multiphase flow performance through wellhead chokes," *Journal of Petroleum Technology*, vol. 26, no. 8, pp. 843–850, 1974.

[18] H. Vaziri, B. Barree, Y. Xiao, I. Palmer, and M. Kutas, "What is the magic of water in producing sand?" in *SPE Annual Technical Conference and Exhibition*, 2002, paper SPE-77683-MS.

[19] G. Nævdal, L. M. Johnsen, S. I. Aanonsen, and E. H. Vefring, "Reservoir monitoring and continuous model updating using ensemble Kalman filter," *SPE Journal*, vol. 10, no. 1, pp. 66–74, 2005.

[20] D. T. Asfha, A. H. A. Latiff, D. A. Otchere, B. N. Tackie-Otoo, I. Babikir, M. Rafi, Z. A. Riyadi, A. D. Putra, and B. A. Adeniyi, "Mechanisms of sand production, prediction: a review and the potential for fiber optic technology and machine learning in monitoring," *Journal of Petroleum Exploration and Production Technology*, vol. 14, no. 10, pp. 2577–2616, 2024.

[21] L. S. Hansen, S. Pedersen, and P. Durdevic, "Multi-phase flow metering in offshore oil and gas transportation pipelines: Trends and perspectives," *Sensors*, vol. 19, no. 9, p. 2184, 2019.

[22] M. Safar Beiranvand, P. Mohammadmoradi, B. Aminshahidy, B. Fazelabdolabadi, and S. Aghahoseini, "New multiphase choke correlations for a high flow rate Iranian oil field," *Mechanical Sciences*, vol. 3, no. 1, pp. 43–47, 2012.

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

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

[25] U.S. Environmental Protection Agency, "40 CFR 146.23, operating, monitoring, and reporting requirements (class II wells)," Code of Federal Regulations, Title 40, Part 146, Subpart C, edition of July 1, 2025, govinfo.gov, 2025.

[26] J. Hou, K. Zhou, X.-S. Zhang, X.-D. Kang, and H. Xie, "A review of closed-loop reservoir management," *Petroleum Science*, vol. 12, no. 1, pp. 114–128, 2015.

4.10 Exercises

4.1 Concept check · 2 min

In scenario 2 the injection plant trips at 6 h. The rate loop holds 900 m^3/d until about 55.6 h, and the rate then falls to 427 m^3/d by the restart at 96 h. Why can the rate loop not hold its target longer?

(a) The rate loop's integral winds up once the sand override has taken the choke

(b) The meter's delay lets the loop overshoot until the drawdown override cuts the rate

(c) The choke sets the bottomhole pressure, and once it is wide open the zone pressures set the rate

(d) The injector's pressure override closes the choke to protect the fracture margin

4.2 Concept check · 1 min

In scenario 3 the rate loop holds the choke while SIC-103 waits, and SIC-103 takes the choke within minutes once the sand passes its setpoint. What lets it take over without a windup delay?

(a) Its integral term is clamped at the top of its range while it waits

(b) The selector picks the loop with the largest error, not the smallest output

(c) Its integral time is the shortest of the four, so it acts first on any error

(d) Its reset term follows the applied opening, so its output waits just above it

4.3 Concept check · 2 min

In the campaign challenge the planner with a 30-day horizon never doses polymer, while the planner with a 180-day horizon spends the whole budget. What explains the difference?

(a) Within 30 days the forecast gain from polymer is smaller than the charge for its polymer

(b) The operator's model makes every polymer plan look worse than a plan of plain water

(c) The fracture limit forbids polymer, because the thicker water raises the sandface pressure

(d) The planner's grid offers polymer only with the lowest of its three liquid targets

4.4 Theory and calculation · 6 min

At the design point the zones deliver $J = 1.224\,\mathrm{m^3/h}$ per bar of drawdown, the choke passes $\partial q_{ch}/\partial p_{wf} = 0.263\,\mathrm{m^3/h}$ more per bar of bottomhole pressure, and the wellbore stores $C_{wb} = 0.25\,\mathrm{m^3/bar}$.

(a) Find the share g of an immediate choke response that remains at steady state, and the time constant τ_{wb} of Equation (4.8) with the choke fixed. *Checkpoint:* 0.82 and 10.1 min.

(b) Find the time constant of the wellbore when the rate loop holds the choke flow. *Checkpoint:* 12.25 min.

(c) The rate loop runs every $\Delta t = 1\,\mathrm{min}$ with $\tau_I = 8\,\mathrm{min}$ and $K_c = 0.020\,\%$ per $\mathrm{m^3/d}$. A one-point choke step raises the rate by 33.4 $\mathrm{m^3/d}$ at once and by 26.2 $\mathrm{m^3/d}$ at steady state. Find the effective integral time and the loop gain at once and at steady state. *Checkpoint:* 8.51 min, 0.67 and 0.52.

4.5 Theory and calculation · 7 min

The sector is 300 m by 300 m. The tight zone is 12 m thick at porosity 0.20 and 700 mD, and the thief zone 4 m thick at porosity 0.25 and 7000 mD. The connate water saturation is 0.20 and $B_o = 1.20$. The end-point relative permeabilities are 0.40 for water and 0.85 for oil, the viscosities 0.5 and 10 cP, and the polymer multiplies the water viscosity by $R(C) = 1 + aC + bC^2$ with $a = 3 \times 10^{-3}/\mathrm{ppm}$ and $b = 3 \times 10^{-6}/\mathrm{ppm^2}$.

(a) Find the pore volume of each zone and the oil in place at standard conditions. *Checkpoint:* 216 000 and 90 000 $\mathrm{m^3}$, 204 000 $\mathrm{Sm^3}$.

(b) Find the thief zone's share of the permeability-thickness product. *Checkpoint:* 76.9 %.

(c) Find the end-point mobility ratio of water to oil with no polymer and with 1200 ppm. *Checkpoint:* 9.41 and 1.06.

(d) The Welge tangent gives breakthrough after 0.303 pore volumes. Find the water the thief zone must take before its water reaches the producer. *Checkpoint:* 27 300 $\mathrm{m^3}$.

4.6 Theory and calculation · 7 min

Take $g = 9.81\,\mathrm{m/s^2}$, a depth of 2000 m, a flowing tubing density of 650 $\mathrm{kg/m^3}$ in the producer and water at 1025 $\mathrm{kg/m^3}$ in the injector. At the design point both wells pass 45 $\mathrm{m^3/h}$ of reservoir fluid, the bottomhole pressure is 213.24 bar, the producer's tubing friction is $k_f = 0.0015$ and the injector's $k_t = 0.0008$, both in $\mathrm{bar\,h^2/m^6}$, the flowline is at 20 bar and the critical pressure ratio is $r_c = 0.5$.

(a) Find the weight of each tubing column. *Checkpoint:* 127.5 bar and 201.1 bar.

(b) Find the wellhead pressure and the drop across the choke, and the flowline pressure below which the choke flow does not depend on it. *Checkpoint:* 82.7 bar, 41.3 bar, 41.3 bar.

(c) PIC-202 holds the injection wellhead at 145 bar. Find the injector sandface pressure and compare it with the fracture pressure of 360 bar less the required margin of 10 bar. *Checkpoint:* 344.5 bar, 5.5 bar inside the 350 bar limit.

4.7 App challenge · 8 min

Load scenario 3, Sand onset. The liquid target rises from 900 to 1400 $\mathrm{m^3/d}$ at 30 min. Use the default seed and a speed of 1, one simulated hour per second, and run each case to 6 h.

(a) Run the scenario as shipped. Read the Erosion indicator and Peak sand tiles, and find in the exported file the mean of the columns `q_liq_m3d` and `sand_gps` over the rows from 3 to 6 h. *Checkpoint:* 52 kJ, 0.45 g/s, 1206 $\mathrm{m^3/d}$ and 0.29 g/s.

(b) Press Reset, set Sand override to Off and run again. Read the same four values. *Checkpoint:* 416 kJ, 3.75 g/s, 1409 $\mathrm{m^3/d}$ and 3.26 g/s.

Accept the means within 5 m³/d and 0.03 g/s. State how much more liquid the second run makes and for how much more sand.

4.8 App challenge · 8 min

Load scenario 2, Declining pressure. The injection plant trips at 6 h and restarts at 96 h. Use the default seed and a speed of 24, and run each case to 240 h.

(a) Run the scenario as shipped. In the exported file find the first row at which `choke_saturated` is 1, the column `q_liq_m3d` at 96 h, and `p_thief_bar` and `p_tight_bar` at 48 h. How long did the choke hold the target after the trip? *Checkpoint:* 55.67 h, 427 m³/d, 184.3 and 234.1 bar, about 49.7 h.

(b) Press Reset, set Voidage replacement ratio to 1.20 on the FIC-201 card before Run, and run again. Read `p_avg_bar` at 240 h in both runs. *Checkpoint:* 194.0 bar as shipped and 214.4 bar at 1.20.

Accept times within 0.2 h, rates within 3 m³/d and pressures within 0.2 bar.

4.9 App challenge · 10 min

Load the challenge, listed as Challenge · 180-day campaign. Use the default seed and a speed of 120, at which each campaign takes a nominal 36 s, and let each run stop itself at day 180. Read the Oil produced, Day 30 vs fixed choke and vs fixed choke tiles after each run.

(a) Run the scenario as shipped. *Checkpoint:* 47 731 Sm³, +48.1 %, +7.1 %.

(b) Press Reset, set Recovery planner to On and Planning horizon to 30 d, and run again. *Checkpoint:* 89 270 Sm³, +94.2 %, +100.4 %.

(c) Set Planning horizon to 180 d, press Reset and run again. *Checkpoint:* 103 912 Sm³, +50.6 %, +133.3 %.

Which runs pass the challenge's 98 000 Sm³, and which plan leads on day 30?

4.10 Application area · 3 min

Forty days into a flood the well makes 245.8 Sm³/d of oil in 900 m³/d of liquid, and the platform treats at most 800 m³/d of water. Raising the target to 1200 m³/d gives 1028 m³/d with 270 Sm³/d of oil and 758 m³/d of water while FIC-104 holds the water, and 1197 m³/d with 314 Sm³/d of oil and 883 m³/d of water with the override off.

(a) Find the extra water per extra unit of oil for the step with the override on, and for the further step with it off. *Checkpoint:* About 4.3 and 2.8.

(b) Name the lever that buys oil on this well, and why the choke is not it. *Checkpoint:* Keeping water out of the well, through the injection split.

4.11 Application area · 4 min

In case 7 the injector's skin grows by 5 units a day. PIC-202 takes the pump on day 14.5, and by day 30 the injection is 757 m³/d against a target of 943 m³/d, the average reservoir pressure has fallen from 249.6 to 216.6 bar, and the fracture margin has never fallen below 14.3 bar, against 10 bar required.

(a) Find the voidage replacement ratio on day 30 and the average rate of pressure loss over the 30 days. *Checkpoint:* 0.80 and about 1.1 bar a day.

(b) A production engineer proposes raising the injection wellhead limit by 4 bar to restore the target. Give two reasons against it, and the remedy the case offers. *Checkpoint:* The skin keeps growing, and the margin would fall toward the fracture limit. Remove the damage.

4.12 Review the agent · 5 min

An AI assistant is asked how to set up the recovery planner for the campaign and replies as follows. “A 30-day horizon is best. In the app it is 29 % ahead of the 180-day horizon on day 30, and because the planner re-plans every five days it can always correct a plan later. A long horizon only adds forecast error, since the operator’s model is 13.9 % optimistic over 180 days.” Find the flaw in the advice and say how it could be checked.

CHAPTER 5

A Deepwater Flowline Simulator for Hydrate Control at the Weakest Point

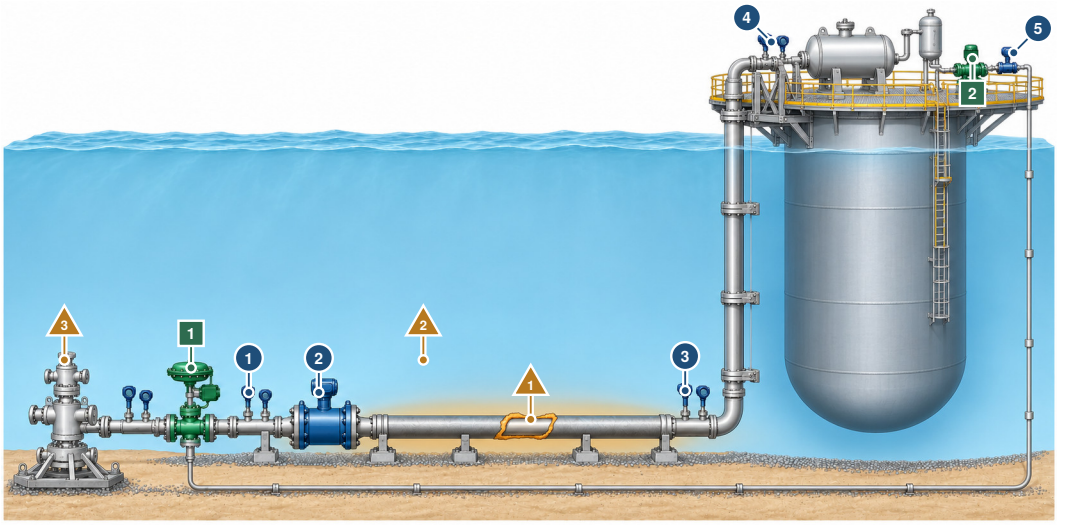

The controller moves 1 subsea choke at the tree 2 methanol pumps on the platform

The controller reads 1 pressure and temperature after the choke
2 multiphase meter for the liquid flow and its water share 3 pressure and temperature at the riser base
4 pressure and temperature where the fluid arrives 5 methanol pump flow

What disturbs it 1 flooded insulation 2 colder seabed water 3 more water from the well

The weakest point is the foot of the riser, where the cooled fluid still carries the seabed pressure, and the methanol reaches it only after the long trip down the umbilical and along the whole line. Illustration, not to scale.

IN PLAIN TERMS **The Coldest Pipe Is Not the Riskiest**

The main idea. Oil, gas and water flow through a long, cold seabed pipe and up to a floating platform. Methanol added at the well keeps the gas and water from locking into solid crystals that could plug the pipe.

The challenge. Pressure lets the crystals form in warmer fluid, so the riskiest spot is the deep foot of the rising pipe, not its colder top. Methanol rides with the liquid and needs hours to get there, so a gauge there asks for it too late.

Why it matters. A dose sent ahead of the cold protects the line, while a dose sent after it comes too late.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the two History boxes and the first Did you know? box in Section 5.2, the first Try it box in Section 5.4, the second Try it box and the second Did you know? box in Section 5.6, and Key points in Section 5.9. No equations.
Study	60 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 5.6 with Figure 5.5, Section 5.7 and Section 5.8 with its Open challenge boxes.

The control question Gas hydrates form wherever the liquid in a subsea line is colder than an equilibrium that climbs with pressure and falls with the methanol in its water, so the constraint is a minimum over kilometers of pipe whose location moves. How is methanol dosed at the tree for a margin that is judged hours downstream, and what can an estimate built from a few transmitters tell the dose?

Learning objectives

- Locate the weakest point of a flowline from its temperature and pressure profiles, and explain why it is not where the fluid is coldest.
- Compute the liquid transit to the riser base from the line's liquid volume and the flow, and the methanol rate for a target concentration.
- Explain why the dose for a turndown or a water breakthrough must come from a feedforward on the predicted demand and the water rate, not from feedback.
- Estimate a stopped line's no-touch time from its lumped cooldown, and judge what two thermometers can and cannot say about lost heat.
- Plan a cold restart through the inhibition lead, the choke ramp and the flow held to the pumps, and grade it on exposure, methanol and oil.

Prerequisites Sections 1.4.3 (holdup and the transport delay of a liquid column), 1.4.8 (selectors and external reset), 1.6.2 (the framework's PID), 1.6.3 (for case M) and 1.6.4 (the apps), and the rate loop under an override on one choke in Chapter 4.

The app https://apmonitor.com/apps/subsea-flow-assurance/

See also Volume 5, Chapter 2 for control through a fuel delay, Volume 5, Chapter 7 for material judged after it has moved on, and Volume 5, Chapter 11 for a state that no instrument reaches, inferred from a model. Volume 4, Chapter 2 for a protective override, and Volume 4, Chapter 5 for a trim behind overrides.

5.1 Introduction

Many process constraints apply at the worst point of a distributed plant rather than at a transmitter. The hottest and the slowest vials of the freeze-dryer in Volume 8, Chapter 5 are a batch example. Transport-reaction processes such as tubular reactors are distributed parameter systems whose control is a subject of its own [1], a drinking-water network is judged by the disinfectant residual that reaches its consumers some time after a booster station has dosed it [2], and a subsea flowline by the smallest difference, anywhere along it, between the liquid temperature and the local hydrate equilibrium [3, 4]. The point that decides such a constraint moves as the operating state moves, and the actuator that protects it acts somewhere else.

A deepwater tieback carries a well's fluid from a subsea tree through an insulated flowline on the seabed and up a riser to a floating platform. Gas hydrates, crystalline solids in which water cages light gas molecules, form wherever free water is colder than an equilibrium temperature set by the pressure, and in deep water both the cold and the pressure are permanent [3]. Insulation keeps the stream warm while it flows and slows its cooling when it stops, and a thermodynamic inhibitor such as methanol lowers the equilibrium in the water it dissolves in [5]. The methanol is pumped on the platform, travels down an umbilical to the tree and then moves with the liquid.

The simulator addresses two beliefs, that the coldest fluid marks the greatest hydrate risk and that an inhibitor can be dosed by feedback on the margin it protects. On the line modeled here the coldest fluid is the fluid arriving at the platform, at 19.4 °C with a margin of +12.2 °C, while the smallest margin, +6.41 °C, sits at the foot of the riser, where the fluid is 3.9 °C warmer but still under 86.5 bar. A step of the methanol pump reaches half strength at that point 3.49 h later at the design rate. A loop that acts only once the estimated margin has fallen therefore doses liquid that has already gone past. A low-order model with one constant delay captures neither the moving minimum nor a delay that stretches as the flow falls.

The class is general. A dose applied at one node and judged at another after a transport delay is the structure of chlorine boosters in water networks [2], and control of a spatial profile rather

than of a point value is the subject of transport-reaction process control [1]. The canal of Volume 9, Chapter 5 has the same structure, a release at the head judged at the far end long after it was made. Three lessons carry across the class. The controlled variable is a minimum whose location moves, the dose belongs to a prediction of demand because feedback at the point of risk arrives a transit late, and an estimate of the profile built from a few sensors can explain the right effect by the wrong cause. However, control studies of pipeline-riser systems have treated the flow itself, such as the riser slugging that feedback on a choke can stabilize [6], while hydrate prevention has been treated as a design calculation of insulation and inhibitor volume [5, 7].

This chapter works with a deepwater flowline simulator in which the hydrate margin is computed slice by slice on the simulated truth, the methanol travels with the liquid, and the methanol loops see only an estimate built from six transmitters, the multiphase meter and the pump's flow meter, so that a turndown, a water breakthrough, an unmeasured heat loss, a trip or a cold restart can each be read as exposure below the hydrate curve. The hydrate correlations, the thermal model, ratio feedforward and override control are established practice. The model of the line is reduced in one particular way. It moves water cut and methanol exactly, as slices of liquid, while the temperatures ride on fixed segments, and it keeps a record of the truth that the instruments never see. The chapter gives it completely enough for the model to be written again from the text. The nine nonchallenge cases and the challenge separate what heat, chemistry, estimation and sequencing can and cannot guarantee.

Section 5.2 traces how hydrates became a control problem on the seabed. Section 5.3 builds the model of the flowline and riser: the thermal segments, the well and choke, the pressure and holdup, the slices of liquid and the hydrate margin. Section 5.4 draws its consequences for the weakest point, the liquid transit and the cooldown of a stopped line. Section 5.5 sets out the instruments, the profile estimator, the choke loops, the methanol feedforward with its trim and override, the shutdown and restart sequence and the optional predictive dosing. Section 5.6 reports the measured cases, the challenge and the comparison with predictive dosing. Section 5.7 sets the hardware of a real tieback beside the model, Section 5.8 states the limitations, Section 5.9 concludes, and the exercises of Section 5.10 run the cases.

5.2 Hydrates, Flowlines and Inhibitors

Hydrates became an engineering problem when natural gas began to travel long distances by pipe. The first History box below tells how the plugs in those lines were identified. By 1945 the formation of hydrates could be predicted from vapor–solid distribution coefficients [8] or from a chart of formation conditions against pressure and gas gravity [9]. The defenses that followed are still the ones in use, removing the water, keeping the gas warm, or adding an alcohol or a glycol that lowers the equilibrium in the free water, and the shortcut relation that carries Hammerschmidt's name is still used to estimate that depression [5]. The curve in the simulator descends from the gas-gravity route through a later correlation of the same kind.

> ⧖ HISTORY **1934: The Plugs Were Not Ice**
>
> Hydrates were first made in a laboratory in 1811, and for a long time they were a matter of scientific curiosity [10]. Before 1934 the crystals that blocked gas lines were blamed on various causes, with much apprehension and little definition. In 1934 Hammerschmidt showed that the plugs, which formed in cold weather above the freezing point of water, were hydrates and not ice [10–12]. He also gave a simple equation for the pressure and temperature at which they form, and he identified inhibitors, among them methanol and ethylene glycol, which are still in common use. The industry could then move from apprehension to avoidance, keeping its lines outside the conditions that such a calculation predicts [10].

Offshore production moved the problem onto the seabed and lengthened it, as the second History box below tells, and partial protection turned out to be a hazard of its own. Insulated lines protected

by methanol during shutdowns and restarts made the time a stopped line takes to cool into the hydrate region a design quantity [5], and the dynamic two-fluid model OLGA [13] brought transient multiphase simulation of such lines into engineering use.

> ⧖ HISTORY **1994 to 1996: Longer Lines in Deeper Water**
>
> In 1994 the fluids of a subsea well could be carried about 10 to 15 miles (16 to 24 km), while leases in about 4000 ft (1200 m) of water could need 50 to 60 miles (80 to 97 km). Chemical injection was regarded as the primary defense, and knowledge of insulation for deepwater lines was described as very limited [7]. Two years later methanol and ethylene glycol were the inhibitors most used offshore. Blockages had followed failures of the chemical injection system, unscheduled shutdowns at the host platform and too little inhibitor at well startup, the three events of cases 6, 7 and 8. Field experiments had shown under-inhibited lines blocking more readily than lines with no inhibitor, and laboratory tests at 1 to 5 percent by weight found both inhibitors speeding up hydrate formation [12].

The emphasis then moved from avoiding hydrates toward managing their risk [14]. Low dosage inhibitors, kinetic inhibitors that delay formation and anti-agglomerants that keep the crystals dispersed, were developed as alternatives to large volumes of methanol [15], and transient formation models now follow the growth of a hydrate phase rather than only its equilibrium [4, 16]. On the same hardware, the severe slugging of a pipeline-riser system has been analyzed as a problem of stabilization by feedback on a choke valve [6]. The simulator keeps to an insulated line protected by methanol from the platform, the arrangement in which the transport delay and the weakest point are plainest.

> ✦ DID YOU KNOW? **A Plug That Is Mostly Water**
>
> A 2021 review puts flowline hydrates at 85 mole percent water, with gas molecules smaller than pentane held in the water cages. Below about 1200 m of water the sea around a flowline is typically near 277 K, close to the 4 °C seabed of this chapter's line [10].

5.3 Model of the Flowline and Riser

The simulated line runs from a subsea tree at 1500 m through 5 km of seabed flowline and a vertical riser of 1.5 km to a spar, with a 0.25 m bore throughout. At the design rate of 1000 Sm^3/d of liquid with a 20 % water cut the choke is open 63.2 % and drops 25.1 bar, the fluid leaves it at 58.7 °C, and the platform holds the arrival at 20 bara. The layout, with its segments, umbilical and transmitters, appears in Figure 5.1.

The controlled, manipulated and disturbance variables and the states are sorted in Table 5.1. The integrator carries 73 states, the fluid and wall temperatures of 32 segments, the flow, the temperature at the wellhead, the two stages of the umbilical, and five running totals of liquid, methanol and energy that close the ledgers. The water share, methanol and exposure of the liquid live instead on 346 slices at the design rate, a discrete state moved after each integration step. Besides the sliders of Table 5.1, discrete events trip the well, request a planned shutdown or a restart, tag a parcel, drop the pump availability to 40 %, raise the water cut by 25 points, cool the seabed water by 2.5 °C, raise the arrival pressure by 25 bar, and fail or repair TT-103.

5.3.1 Segments, Heat and the Sea

The path is divided into 20 seabed segments of 250 m and 12 riser segments of 125 m. Each segment i holds a fluid node and a wall node that stands for the steel and the inner insulation, and the fluid

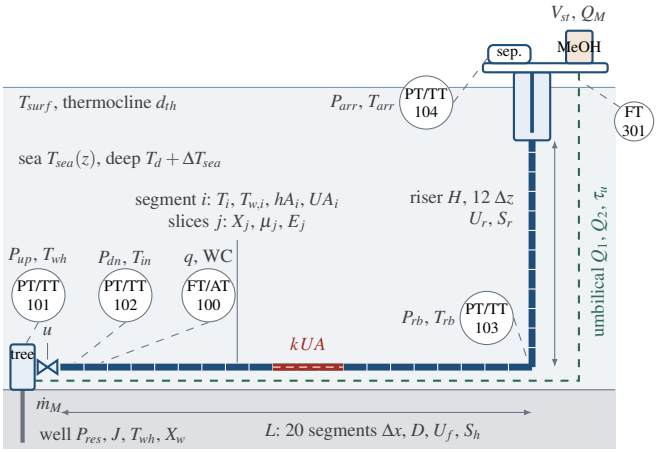

Figure 5.1 The modeled line, not to scale. The well feeds the tree, where the choke at opening u passes the liquid rate q with water share X_w into 20 seabed segments of length Δx and 12 riser segments of height Δz, each with fluid and wall temperatures T_i and $T_{w,i}$, a film hA_i and an insulation UA_i to the sea at $T_{sea}(z)$. The liquid carries its water share X_j, methanol μ_j and exposure E_j as slices j. The section marked kUA (red, with a dashed white center line) can lose its insulation by the factor k. The dashed umbilical takes the pump flow Q_M from the platform tank to the tree through the stages Q_1 and Q_2. PT/TT-101 to 104 measure pressure and temperature upstream and downstream of the choke, at the riser base and at the arrival, FT/AT-100 is the subsea multiphase meter and FT-301 the pump's flow meter.

Table 5.1 Flowline Variables by Role, With Design Values and Ranges

Variable	Role	Design value	Range
q_m	CV of FIC-100 (FT-100)	1000 Sm^3/d	SP 0–2000
$\hat{m}_{min}$	CV of the methanol trim	+6.41 °C, SP 3	SP 0–10
$P_{dn,m}$	CV of PIC-102 (PT-102)	90.4 bar	limit 80–160
u	MV, choke opening	63.2 %	0–100
Q_M	MV, methanol pump	0 m^3/h	0–4
X_w	DV, well water cut	20 %	0–80
P_{arr}	DV, arrival pressure	20 bara	10–60
ΔT_{sea}	DV, seabed water change	0 °C	±4
k	DV, insulation damage	1	1–10
a_p	DV, pump availability	100 %	0–100
b_{WC}, TT-103	DV, meter bias, failure	0, working	±15, 0/1
q, T_{wh}	states, flow, wellhead	1000 Sm^3/d, 60 °C	—
Q_1, Q_2	states, umbilical stages	0 m^3/h	—
T_i, $T_{w,i}$	states, 32 segments	56.0 to 19.4 °C	—
X_j, μ_j, E_j	slice records, 346 slices	0.2, 0, 0	—

carries its heat downstream by first-order upwind transport, as shown in Equation (5.1).

$$\begin{aligned} C_{f,i}\frac{\mathrm{d}T_i}{\mathrm{d}t} &= W_{i-1}\left(T_{i-1}-T_i\right)-hA_i\left(T_i-T_{w,i}\right) \\ C_{w,i}\frac{\mathrm{d}T_{w,i}}{\mathrm{d}t} &= hA_i\left(T_i-T_{w,i}\right)-UA_i\left(T_{w,i}-T_{sea,i}\right) \end{aligned} \tag{5.1}$$

Here $W_{i-1} = q\,c_v(X_{i-1})$ is the heat-capacity flow entering segment i, with $T_{-1} = T_{in}$ and $X_{-1} = X_w$ at the choke, and the other coefficients follow Equation (5.2).

$$\begin{aligned} c_v(X) &= X\rho_w c_w + (1-X)(\rho_o c_o + \mathrm{GOR}\,\rho_{gs} c_g) \\ C_{f,i} &= V_i H_i \left[X_i \rho_w c_w + (1-X_i)\rho_o c_o\right] \\ &\quad + V_i (1-H_i)\,\rho_{g,i} c_g \\ hA_i &= \pi D\,\Delta x_i \left[h_s + (h_f - h_s)\min(q/q_d,\,1.5)^{0.8}\right] \\ UA_i &= \pi D\,\Delta x_i\, U\left[1 + (k-1)\,\phi_i\right] \\ C_{w,i} &= C'_w\,\Delta x_i \end{aligned} \tag{5.2}$$

where c_v counts the gas that travels with the oil, $V_i = A\,\Delta x_i$ is the segment volume, H_i the liquid holdup, $\rho_{g,i}$ the local gas density, U the flowline or riser coefficient on the inner area and ϕ_i the share of segment i inside the damageable section. The film rises from natural convection at rest toward its design value with the Dittus–Boelter exponent. The sea is $T_{sea}(z) = T'_d + (T_{surf} - T'_d)\,\mathrm{e}^{-z/d_{th}}$ at the depth z of a segment center, with $T'_d = T_d + \Delta T_{sea}$, which gives 4.04 °C on the seabed and 18.0 °C in the top riser segment.

At the design rate the heat-capacity flow is 27.7 kW/K, of which the gas carries 1.7 kW/K, and the film and insulation in series conduct 5.78 W/(m·K), so the steady temperature approaches the sea over an e-folding length of 4.79 km. The discrete steady solution replaces the exponential by a factor $W/(W + ua_i)$ per segment, ua_i being the film and insulation of segment i in series, and runs 0.51 °C warm at the riser base against the analytic profile, a first-order error that falls to 0.26 °C with 40 seabed segments and 0.13 °C with 80.

5.3.2 Well, Choke and Flow

The well is a shut-in wellhead pressure $P_{si} = P_{res} - \Delta P_{tub} = 150$ bar behind a productivity index. One lumped liquid rate relaxes toward the rate q^* at which the well, the choke and the line balance, as shown in Equation (5.3), with pressures in bar and rates in Sm³/d.

$$\begin{aligned} P_{si} - \frac{q^*}{J} - \left(\frac{q^*}{K(u)}\right)^2 &= P_{dn,s} + K_f\,q^{*2} \\ \tau_h \frac{\mathrm{d}q}{\mathrm{d}t} &= q^* - q \\ K(u) &= K_{ch}\,\frac{R^{u-1} - R^{-1}}{1 - R^{-1}} \end{aligned} \tag{5.3}$$

Here $K(u)$ is an equal-percentage choke of rangeability R, offset so that zero opening passes no flow, $u \in [0,1]$ its opening, $P_{dn,s}$ the static part of the pressure just downstream of the choke and K_f the summed friction coefficient of the line. The choke actuator travels at 20 %/min. At the design rate the well draws down 34.5 bar, so the choke passes from $P_{up} = 115.5$ bar to $P_{dn} = 90.4$ bar, and a fully open choke would pass 1577 Sm³/d. The fluid entering the flowline follows the wellhead temperature less the Joule–Thomson cooling of the choke, as shown in Equation (5.4).

$$\begin{aligned} \frac{\mathrm{d}T_{wh}}{\mathrm{d}t} &= \phi_q\,\frac{T_{wh,f} - T_{wh}}{\tau'_{up}} \\ &\quad + (1-\phi_q)\,\frac{T_d + \Delta T_{sh} - T_{wh}}{\tau_{dn}} \\ T_{in} &= T_{wh} - \mu_{JT}\,(P_{up} - P_{dn}) \\ P_{up} &= P_{si} - q/J \end{aligned} \tag{5.4}$$

where $\phi_q = \min\!\left(1,\, q/(0.05\,q_d)\right)$ blends warm-up and cool-down and $\tau'_{up} = \tau_{up}\,q_d/\max(q,\,0.25 q_d)$ makes a slow restart warm the well more slowly. The well warms toward 60 °C over about an hour at

the design rate, cools toward 10 °C with a 10 h time constant when shut in, and loses 1.26 °C across the choke at design.

5.3.3 Pressure and Holdup

After each step the pressures are recomputed segment by segment, from the arrival pressure down. In each segment from the top, the gas density follows from its pressure with a constant compressibility, the in-situ gas volume per unit of liquid from the fixed gas-oil ratio, and the holdup from a constant slip ratio between the gas and liquid velocities, as shown in Equation (5.5).

$$\begin{aligned} \rho_{g,i} &= \frac{P_i M_g}{ZR\,(T_i + 273.15)}, \qquad r_i = \frac{\mathrm{GOR}\,\rho_{gs}(1 - X_i)}{\rho_{g,i}} \\ H_i &= \frac{1}{1 + r_i/S_i}, \qquad \rho_{m,i} = H_i \rho_{L,i} + (1 - H_i)\rho_{g,i} \\ \Delta P_i &= \rho_{m,i}\, g\, \Delta z_i + \frac{f\,\Delta x_i}{2D}\,\rho_{ns,i}\left(\frac{(1 + r_i)\,q}{A}\right)^2 \\ P_i &= P_{arr} + \textstyle\sum_{k>i} \Delta P_k + \frac{1}{2}\Delta P_i \\ P_{dn} &= P_{arr} + \textstyle\sum_k \Delta P_k \end{aligned} \tag{5.5}$$

Here $\rho_L = X\rho_w + (1 - X)\rho_o$, $\rho_{ns} = (\rho_L + r\rho_g)/(1 + r)$ is the no-slip density of the homogeneous friction term, Δz_i is zero on the seabed, and two fixed-point passes settle the gas density. At the design rate the riser column puts PT-103 at 86.5 bar, 66.5 bar above the platform, and the 5 km of seabed contribute 0.19 bar of friction (0.28 bar with the riser). Along the seabed the holdup rises from 0.53 at the choke to 0.56 as the cooling gas contracts, and the expanding gas thins the liquid up the riser from a holdup of 0.66 at its base to 0.33 at its top, so the line contains 172.6 m^3 of liquid. When the flow stops, the holdup and the fluid heat capacities keep their last values.

5.3.4 Water and Methanol as Slices of Liquid

The water share and the methanol belong to the liquid and travel with it. Upwind segments would smear a methanol front over several 250 m segments and dilute an under-dosed slug into the dosed liquid around it, so the model carries the liquid as a queue of slices of $\Delta V = 0.5\,\mathrm{m}^3$, 19 m of seabed pipe. Each slice j has a water share X_j, methanol μ_j in kilograms per cubic meter of liquid and its own exposure. After every step the liquid that entered is pushed in at the choke with the well's water share, the methanol that reached the water during the step is shared among the new slices by volume, and whatever exceeds the line's liquid volume $\sum_i V_i H_i$ is delivered to the platform. A slice is placed by its cumulative liquid volume from the choke, so a change made at the tree arrives at any point once all the liquid between the tree and that point has been pushed past it. The methanol reaches the tree through two first-order stages and loses a fixed share to the oil and gas, as shown in Equation (5.6).

$$\begin{aligned} \tfrac{1}{2}\tau_u \frac{\mathrm{d}Q_1}{\mathrm{d}t} &= Q_M - Q_1, \qquad \tfrac{1}{2}\tau_u \frac{\mathrm{d}Q_2}{\mathrm{d}t} = Q_1 - Q_2 \\ \dot{m}_M &= \rho_M\, Q_2\,(1 - f_{loss}) \\ W_j &= \frac{100\, r_j}{1 + r_j}, \qquad r_j = \frac{\mu_j}{\rho_w X_j} \end{aligned} \tag{5.6}$$

where $\dot{m}_M$ is the methanol that reaches the water and W_j the aqueous concentration of slice j in wt %, the quantity the equilibrium sees. The two stages of 7.5 min deliver 59 % of a step after 15 min and 91 % after 30 min. The pump draws from a 120 m^3 tank and stops when it is empty, and the well's water cut approaches its setting through a first-order change with a 1 h time constant.

5.3.5 Equilibrium, Margin and Exposure

The uninhibited equilibrium belongs to one representative natural gas of specific gravity 0.65. It is written $T_{eq,0} = a + b \ln P$ with P in bar and fitted by least squares to seven points of the Motiee gas-

gravity correlation between 20 and 200 bar, which it reproduces within 0.41 °C. Methanol in the water depresses it by the Hammerschmidt relation, used up to about 25 wt %, or by the Nielsen–Bucklin relation, used up to about 50 wt %, as shown in Equation (5.7).

$$\begin{aligned} T_{eq,0}(P) &= a + b \ln P \\ \Delta T_H &= \frac{1297\,W}{32.04\,(100 - W)}, \qquad \Delta T_{NB} = -72 \ln x_w \\ m_{ij} &= T_i - T_{eq,0}(P_i) + \Delta T(W_j) \end{aligned} \tag{5.7}$$

Here $a = -14.63\,^{\circ}\mathrm{C}$, $b = 7.073\,^{\circ}\mathrm{C}$, x_w is the mole fraction of water in the aqueous phase and m_{ij} the margin of slice j in segment i, both depressions in °C. The curve gives 6.56 °C at 20 bar and 17.94 °C at 100 bar, and 20 wt % of methanol depresses it by 10.12 °C by Hammerschmidt and by 9.47 °C by Nielsen–Bucklin. The controlling margin m_{min} is the lowest m_{ij} over every slice carrying free water, a water share above 0.005, so one segment may contain dosed and under-dosed liquid together, and the under-dosed part sets its margin. A negative margin is exposure, not a plug. Each segment accumulates the subcooling of its wet liquid weighted by the share of its liquid volume below the curve, as shown in Equation (5.8).

$$\begin{aligned} \frac{\mathrm{d}E_i}{\mathrm{d}t} &= \sum_{j \in i} \frac{v_{ij}}{V_{L,i}} \max(0, -m_{ij}) \\ \mathrm{PRI} &= 1 - \mathrm{e}^{-\max_i E_i / E_{ref}} \end{aligned} \tag{5.8}$$

The sum runs over the wet slices in segment i, v_{ij} is the part of slice j inside it, $V_{L,i} = V_i H_i$, and E_i is counted in °C·h. The plug-risk index PRI is an explicitly illustrative mapping of the worst segment, with $E_{ref} = 12\,^{\circ}\mathrm{C{\cdot}h}$. The scorecard reads exposure and the worst margin from this record of the truth, never from the instruments.

5.3.6 Grid, Step and Starting State

Table 5.2 gives every parameter value. The 73 states advance by classical fourth-order Runge–Kutta steps of 10 s, the holdup, the heat capacities and the pressures are held fixed through each step, and the slices move after it. The fastest thermal mode, fluid exchanging heat with its own wall, decays in 90 to 143 s during flow, and at the 1577 $\mathrm{Sm^3/d}$ of a fully open choke the top riser segment is swept in 112 s, so the upwind transport runs at a Courant number of 0.089. Halving the step moves the riser-base temperature of case 2 at 12 h by 0.000 01 °C and its worst margin by 0.000 06 °C. Refining the seabed to 40 segments lowers the design margin from +6.41 °C to +6.17 °C and brings the first exposure of case 2 14.7 min earlier, and 80 segments give +6.04 °C. Over 12 h of flow the energy held in fluid and wall matches the inflow less the outflow and the loss to the sea within 10^{-9} of the inflow, once a tracked revaluation of the heat capacities, under 10^{-4} of the inflow, is included, and every kilogram of methanol and cubic meter of liquid is either in the line or delivered, two of the 57 checks the application passes.

A flowing case starts from the discrete steady solution of Equation (5.1) at its flow, water cut and dose, iterated with the hydraulics and the Joule–Thomson cooling of the choke until the inlet temperature settles, with the choke at the opening that holds that flow, so the reset state has zero derivatives. A cold case starts from the same state at 800 $\mathrm{Sm^3/d}$ dosed at 30 wt %, closes the choke and integrates 24 h of cooldown before the run begins, which leaves the fluid of the well at 14.5 °C.

5.3.7 Assumptions

The hydraulics are reduced. A single liquid rate follows the balance of Equation (5.3) with a declared 120 s response, the pressures carry no dynamics of their own, and a constant slip ratio fixes the holdup, which therefore does not rise at low flow and admits neither terrain nor severe riser slugging.

Table 5.2 Flowline Parameters, Marked Visible Where the Panel Adjusts Them, Hidden Where It Does Not and Fixed Where the Code Sets Them

Symbol	Meaning	Value	Class
L, H	flowline, riser length	5.0 km, 1500 m	visible
D	inside diameter	0.25 m	visible
N_h, N_r	seabed, riser segments	20, 12	hidden, fixed
—	damageable section	2.25–3.0 km	visible
U_f, U_r	overall U, inner area	7.5 W/(m^2·K)	visible
C'_w	wall per meter	110 kJ/(m·K)	visible
h_f, h_s	film, design flow, rest	400, 60 W/(m^2·K)	hidden
T_d, T_{surf}	seabed, surface water	4, 22 °C	visible
d_{th}	thermocline scale	250 m	hidden
GOR	gas-oil ratio, all free	100 Sm^3/Sm^3	visible
S_r, S_h	slip, riser, seabed	1.6, 1.0	vis., hid.
ρ_o, c_o	oil	850 kg/m^3, 2100 J/(kg·K)	hidden
ρ_w, c_w	produced water	1025 kg/m^3, 4000 J/(kg·K)	hidden
c_g, Z	gas	2300 J/(kg·K), 0.85	hidden
f	Darcy friction	0.02	hidden
γ, ρ_{gs}	gas gravity, standard	0.65, 0.796 kg/Sm^3	fixed
P_{res}, ΔP_{tub}	reservoir, well column	250, 100 bar	visible
J	productivity index	29 Sm^3/(d·bar)	visible
$T_{wh,f}$, ΔT_{sh}	flowing, shut-in excess	60, 6 °C	vis., hid.
τ_{up}, τ_{dn}	well warm-up, cool-down	1, 10 h	visible
μ_{JT}	Joule–Thomson, choke	0.05 °C/bar	visible
K_{ch}, R	choke capacity, range	720 Sm^3/d/$\sqrt{\text{bar}}$, 25	vis., hid.
—	choke travel rate	20 %/min	visible
τ_h, τ_{WC}	flow response, water cut	120 s, 1 h	visible
q_d, q_{plat}	design, platform limit	1000, 2500 Sm^3/d	hid., vis.
Q_{max}, τ_u	pumps, umbilical	4 m^3/h, 15 min	visible
V_{st}, f_{loss}	storage, loss to oil	120 m^3, 10 %	visible
ρ_M	methanol density	791 kg/m^3	fixed
P_{MAOP}, E_{ref}	MAOP, index scale	130 bar, 12 °C·h	visible
ΔV	slice volume	0.5 m^3	fixed

A stopped line keeps its phases in place, so the liquid in a shut-in riser never falls back to load its base. A two-fluid transient model resolves all three effects [13].

The gas is a fixed 100 Sm^3 per Sm^3 of oil, all of it free and of constant compressibility, so no gas returns to solution or comes out of it, and the choke is the only place where the stream cools on expansion. Oil and water have constant properties and neither compresses nor shrinks, which makes the liquid rate equal to the standard rate.

One fluid node and one wall node represent each segment, heat does not conduct along the pipe, and the energy moves by first-order upwind transport, the source of the 0.51 °C warm error at the riser base. The thermocline keeps a fixed exponential shape, and a cold current lowers only its deep value.

The hydrate curve belongs to one representative gas, fitted between 20 and 200 bar, and is illustrative rather than a design curve for any field. The inhibition relations are shortcut methods, the produced water carries no salt, and a fixed 10 % of the methanol leaves with the oil and gas at every pressure, temperature and gas rate. Exposure is thermodynamic, so an under-dosed slug counts as partly protected, although under-inhibited systems have been seen to block more readily than uninhibited ones [12].

5.4 The Weakest Point and the Liquid Transit

The margin at a point is a race between two profiles, as shown in Figure 5.2. Along the seabed the pressure falls only from 90.4 to 90.2 bar, so the uninhibited equilibrium holds near 17.2 °C while the fluid cools from 58.7 to 23.8 °C toward the 4.04 °C sea. Up the riser the fluid keeps cooling, from 23.3 °C in the first riser segment to 19.4 °C at the top, but the pressure falls by 64.6 bar over the same

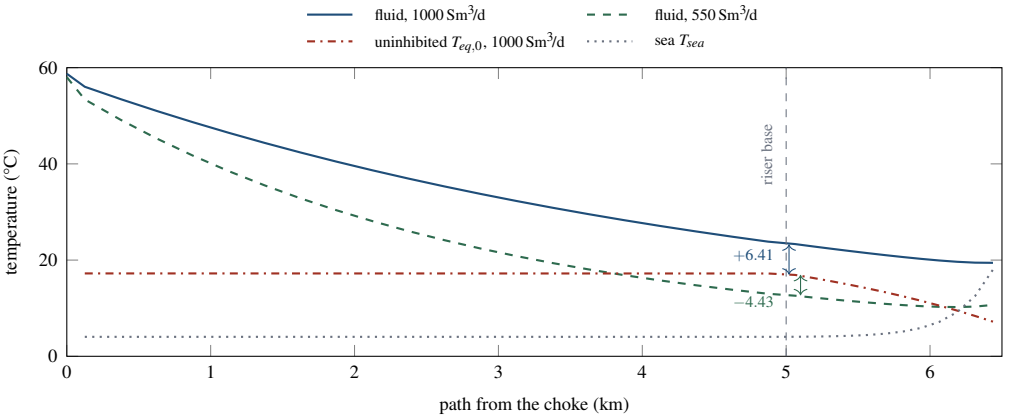

Figure 5.2 Steady fluid temperature, uninhibited hydrate equilibrium and sea temperature along the path at 1000 Sm³/d and 550 Sm³/d, from the model's own steady initialization. The fluid is drawn solid at 1000 Sm³/d and dashed at 550 Sm³/d, the equilibrium dash-dotted and the sea dotted. The equilibrium at 550 Sm³/d lies within 0.12 °C of the one drawn. The arrows give the uninhibited margin in the first riser segment, where both flows reach their minimum.

climb and takes the equilibrium from 16.9 down to 7.2 °C. The minimum of $T - T_{eq}$ therefore falls where the fluid has spent longest cooling but still carries the full seabed pressure, in segment 21 at 5.06 km and 1437.5 m depth, at +6.41 °C. The coldest fluid, at the platform, keeps +12.2 °C.

The equilibrium barely depends on the flow, because the pressure along the seabed is the riser's head plus the platform's pressure, and at the riser base it moves only from 16.91 to 17.02 °C between 1000 Sm³/d and 550 Sm³/d. The fluid temperature carries the whole flow dependence. The uninhibited riser base settles at 23.3 °C at 1000 Sm³/d, 16.6 °C at 700 Sm³/d and 12.6 °C at 550 Sm³/d, so the steady uninhibited margin falls from +6.41 °C to −0.43 °C and −4.43 °C. Without methanol the line cannot hold its margin setpoint at any of the lower rates.

🕮 **Pause.** Try Exercise 5.1 before reading on.

Transport sets the timing. The seabed holds 134.7 m³ of the line's 172.6 m³ at the design rate, so a change made at the tree reaches the riser base 3.23 h later and the platform 4.14 h later, and at 550 Sm³/d the riser base about 5.9 h later. A step of the methanol pump in manual reaches half strength at the riser base after 3.49 h, the umbilical adding about a quarter of an hour. A parcel tagged at reset arrives after 4.15 h at 19.42 °C, its coldest, and its own worst margin, +6.39 °C, was recorded at 5.13 km as it left the first riser segment.

Table 5.3 Transmitters, Lags and Noise

Tag	Point	Lag	Noise, 1σ
PT/TT-101	upstream of the choke	5, 30 s	0.05 bar, 0.05 °C
PT/TT-102	downstream of the choke	5, 30 s	0.05 bar, 0.05 °C
PT/TT-103	riser base	5, 30 s	0.05 bar, 0.05 °C
PT/TT-104	arrival at the spar	5, 30 s	0.02 bar, 0.05 °C
FT-100	multiphase meter, liquid	30 s	1.5 % of reading
AT-100	same meter, water cut	60 s	2 % absolute, bias
FT-301	pump discharge, topsides	5 s	0.5 % of reading

TRY IT IN THE APP 1 · Warm steady operation

Load the case, press Run and open the Profiles tab. Compare the fluid temperature with the hydrate equilibrium along the whole path. The gap is narrowest in the first riser segment, at +6.41 °C, while the fluid reaching the spar, the coldest of the line, keeps +12.2 °C.

A stopped line has a time scale of its own. With the film at natural convection, the fluid first passes its heat to the wall with a time constant of 14 to 16 min, and fluid and wall then cool together toward the sea with the lumped time constant $(C_f + C_w)/UA$, 8.2 h on the seabed and 8.8 h at the riser base. The design line, if it stopped, would have 3.28 h before its first segment reached the curve by the estimator's prediction.

5.5 Instruments, Estimator and Control

The engine is that of Section 1.6.1. A control period of 20 s, two physics steps, serves the loops, the estimator and the sequence alike, and the historian writes one row for each simulated minute. The hydraulics, the slices and the exposure record are brought up to date after each integration step. Speeds range from 60 to 7200 times real time, and at the default of 900 an hour of plant time takes 4 s, so a 24 h case runs in about a minute and a half.

5.5.1 Instruments

The loops never see a segment of the plant. Each transmitter reads the true value at its own point through an exact first-order lag and adds colored noise, correlated over 30 s on the pressures and the pump meter, 60 s on the thermometers, 120 s on the liquid rate and 300 s on the water cut, with the lags and amplitudes of Table 5.3. A failed TT-103 reports nothing, and the estimator is told so rather than fed a guess. The measured water cut passes a 300 s filter before the dose uses it.

5.5.2 The Profile Estimator

The methanol loops and the sequence act on a second copy of the segment model that knows only the declared design, its insulation, geometry and fluids, and the design seawater temperatures, whatever the Parameters panel says. It is driven by FT-100, the filtered water cut, TT-102 as its inlet temperature, the FT-301 pump flow through its own copy of the umbilical stages and PT-104, it carries its own slices, and it advances by forward Euler in two 10 s substeps per control period. Three corrections pull it toward the plant. Its riser pressures are scaled to pass through PT-104 and PT-103, and its seabed pressures are shifted to PT-103 with the mismatch at PT-102 spread linearly back to the choke. The residual at TT-103 is distributed along the seabed in proportion to the distance from the choke, and the residual at TT-104 is distributed the same way up the riser, both with a 300 s time constant. While the line flows, factors θ_f and θ_r multiply the design U of the flowline and the riser, adapted at 4×10^{-5} per °C per second of the riser-base and riser residuals and bounded between 0.3 and 6, with θ_r adapting only while the riser-base residual is under 0.1 °C. Two thermometers fix the total heat the line loses and say nothing about where it is lost. The kiln of Volume 5, Chapter 11 meets

the same gap in a core temperature that no instrument reaches. At reset the estimator starts from its own steady profile at the measured flow, inlet temperature and pump rate, or, after a shut-in, from a line assumed to have cooled everywhere by as much as the riser-base reading shows and to carry the pre-dose throughout.

The estimator also predicts. It marches the steady profile toward which the present flow is heading, with the adapted factors and the design sea, which gives the steady uninhibited margin $\hat{m}_{ss}$, and it predicts the no-touch time, the time left before the first wet segment reaches its inhibited equilibrium if the flow stops, from a lumped cooldown of each segment, as shown in Equation (5.9).

$$t_{nt} = \min_i \frac{C_{f,i} + C_{w,i}}{\theta\, UA_i} \ln \frac{\hat{T}_{m,i} - T_{sea,i}}{T_{eq,i} - T_{sea,i}} \tag{5.9}$$

Here $\hat{T}_{m,i}$ is the heat-capacity-weighted mean of the estimated fluid and wall temperatures and $T_{eq,i}$ the segment's inhibited equilibrium, with $t_{nt} = 0$ once $\hat{T}_{m,i} \le T_{eq,i}$ and no limit while $T_{eq,i}$ lies below the sea.

5.5.3 The Choke Loops

FIC-100 holds the liquid rate on FT-100 by moving the choke, with $K_c = 0.03\,\%$ per $\mathrm{Sm^3/d}$ and $\tau_I = 240\,\mathrm{s}$. PIC-102 watches PT-102 against a 120 bar limit below the 130 bar MAOP, with $K_c = 5\,\%$ per bar and $\tau_I = 120\,\mathrm{s}$. A low select hands the choke to the smaller of the two requests and of a ceiling that the restart sequence ramps, $u = \min(u_Q, u_P, u_{ceil})$. The loop that loses follows what was applied. The integral of FIC-100 is set so that its output equals the applied opening, and the integral of PIC-102 is set to the applied opening itself, so that its output sits above it by its own proportional term and takes the choke exactly when the pressure crosses its limit. At reset both choke loops start at the opening that holds the initial state and the trim at whatever concentration the pump already delivers beyond the feedforward, so no case begins with a bump. This is the low selector with external reset of Section 1.4.8. The three loops of the application are instances of the framework's velocity-form PID of Section 1.6.2, Equations (1.16) to (1.19), with the control period $\Delta t_c = 20\,\mathrm{s}$, a zero bias and the back-calculation tracking time equal to τ_I. The derivative is switched off in all three loops ($\tau_D = 0$), and no output carries a rate limit of its own, although the choke actuator travels at 20 %/min. Every loop uses $e = \mathrm{SP} - \mathrm{PV}$ with a positive gain, so a low rate opens the choke, a high pressure closes it and a low margin adds methanol. In manual the operator's choke opening and pump rate go straight to the actuators, and every loop tracks them by setting its integral term to whatever makes its own output equal the applied value, so switching back to automatic moves nothing.

5.5.4 Methanol Feedforward, Trim and Override

The pump command is a high select of three requests, limited by the pumps still running and by the methanol left in storage, as shown in Equation (5.10).

$$\begin{aligned} W_{ff} &= W^{-1}\big(m_{SP} - \min(\hat{m}_{ss}, \hat{m}_{now})\big) \\ Q_L &= \frac{r(W_{ff} + \Delta W)\,\rho_w\, q_{w,m}}{\rho_M (1 - f_{loss})} \\ r(W) &= W/(100 - W) \\ Q_M &= \min\big(\max(Q_L, Q_{ovr}, Q_{seq}),\ a_p Q_{max}\big) \end{aligned} \tag{5.10}$$

Here W^{-1} inverts the depression relation in use and $\hat{m}_{now}$ is the uninhibited margin of the estimated line as it stands, so the feedforward doses for the worse of the line as it will become and as it is. The dose is proportioned to the measured water rate $q_{w,m}$, the FT-100 rate times the filtered water cut, or, as an option, to the liquid rate with the water cut taken as its 20 % design value. The trim ΔW is a PI on the estimated minimum margin with $K_c = 0.5\,\mathrm{wt\,\%}$ per °C and $\tau_I = 4\,\mathrm{h}$ against a 3 °C setpoint, limited between zero and the concentration the available pumps can still add, so it may add methanol

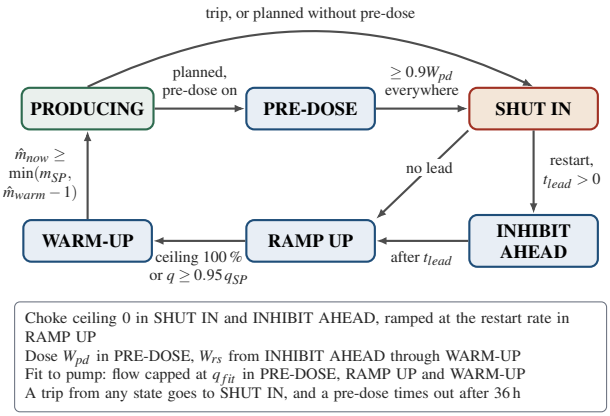

Figure 5.3 Shutdown and restart sequence. The conditions are evaluated on the estimate every 20 s, $\hat{m}_{warm}$ is the steady uninhibited margin at the target flow with the well at its flowing temperature, and q_{SP} stands for the smaller of the setpoint and q_{fit}.

but never remove the feedforward's. The low-margin override Q_{ovr} runs the available pumps at full capacity whenever the estimated margin drops under 1 °C, and lets go 0.5 °C above it, and Q_{seq} is the dose of the sequence. The loop is live only in automatic with the well producing and flowing above 5 % of design. Whenever the override or the sequence requests more methanol than the loop would, the trim freezes at its present value, the conditional integration that keeps a loop behind a selector from winding up [17].

5.5.5 Shutdown and Restart

The sequence is the state machine of Figure 5.3. A trip commands the choke shut from any state. A planned shutdown in automatic, with the pre-dose on, first runs the pumps at the pre-dose concentration $W_{pd} = 30\,\text{wt}\,\%$ until the estimate shows at least nine tenths of it in every segment, then closes the choke. A restart in automatic runs the pumps for the lead t_{lead} with the choke still shut, ramps a ceiling on the choke at the restart rate, and doses the restart concentration $W_{rs} = 30\,\text{wt}\,\%$, first for the target flow and, once the ramp is complete, for the measured flow, until the estimated uninhibited margin reaches the smaller of m_{SP} and the warm steady margin at the target flow less 1 °C. With the fit-to-pump option on, the flow setpoint during pre-dose, ramp and warm-up is capped at the rate at which the available pumps deliver the target concentration, as shown in Equation (5.11).

$$q_{fit} = \frac{a_p Q_{max}\, \rho_M (1 - f_{loss})}{r(W)\, \rho_w\, \text{WC}_m / 100} \tag{5.11}$$

With 4 m³/h, 30 wt % and a 20 % water cut this is 778 Sm³/d, and it moves with the measured water cut WC_m. In manual a restart hands the choke and the pump to the operator at once.

5.5.6 Plant, Estimator and Loop Models

The chapter keeps three models distinct. The plant is the model of Section 5.3, 73 states and the slices, run on the true environment. The estimator's model is the same segment equations on the design environment and the declared default parameters, whatever the learner sets, corrected through six transmitters, and it is the model consulted by the original methanol loops and sequence when the original structure is selected. The feedforward's model is the estimator's steady prediction, and the no-touch prediction of Equation (5.9) is a lumped cooldown that ignores the fast exchange between fluid and wall. No model is identified inside the application. The flow loop faces the 120 s flow lag behind a rate-limited choke, and after the case 2 cut the liquid rate comes within 5 % of its new

setpoint in 17.8 min and stays within 3 % of it from 20 min, the remaining wander being the FT-100 noise that the loop follows. The methanol loop faces a transport delay of hours between its dose and the margin it acts on, which no PI tuning can shorten.

5.5.7 Predictive Methanol Dosing

An additional economic MPC selects only the methanol rate, following the receding-horizon formulation used in industrial control [18] and set up in general in Section 1.6.3. Its sixteen states are fourteen pump-equivalent rates $z_1,\dots,z_{14}$ in m³/h, an uninhibited margin h and a bias b_m in °C. The nominal reduction is

$$\begin{aligned} 450\dot{z}_1 &= v - z_1, & 450\dot{z}_2 &= z_1 - z_2, \\ \dot{z}_i &= \frac{12 q_m}{V_R}(z_{i-1} - z_i), & i &= 3,\dots,14, \\ 7200\dot{h} &= \hat{m}_{ss} - h, & \dot{b}_m &= 0, \\ y &= h + g z_{14} + a + b_m. \end{aligned} \tag{5.12}$$

Here v is the candidate pump rate and q_m the measured liquid rate in m³/s. The nominal observer supplies $\hat{m}_{ss}$, the steady uninhibited minimum at current flow, using the existing known bounded flow target when stopped. If unavailable, h supplies that target. Declared default geometry and steady holdup give $V_R = 138.750\,234\,\mathrm{m}^3$ through the riser-base cell and $172.633\,921\,\mathrm{m}^3$ for the full line. The twelve residence cells replace sharp slice transport. With measured water rate $q_{w,m}$, let $\alpha = \rho_M(1 - f_{loss})/(3600\rho_w q_{w,m})$ and $D(v) = \Delta T[100\alpha v/(1+\alpha v)]$. The slope $g = D'(z_{14})$ and intercept $a = D(z_{14}) - g z_{14}$ linearize the selected inhibition relation. The implemented forward difference uses a ratio increment 10^{-5}. Water cut is floored at 0.1 % and water rate at 10^{-7} m³/s for this map.

A separate profile observer uses the original instrument packet and the inherited declared predose prior for a cold line, with base physical coefficients frozen at their defaults, as the original observer's are. Its inherited thermometer nudging and estimated heat-loss multipliers remain active when enabled. It receives no sea measurement, true profile or future event. Both profile observers and the reduced states advance in Manual and under classical control, using actual pump history. Initial z_i equal the measured pump rate, h starts at the observer's uninhibited margin and $b_m = 0$. Each 20 s scan propagates the reduction in substeps no longer than 5 s. A bounded residual update follows the nominal observer's inferred minimum $\hat{m}_o$, rather than the true minimum,

$$b_m^{+} = \mathrm{clip}_{[-30,30]}\{b_m + \lambda_b[\hat{m}_o - h - D(z_{14}) - b_m]\}. \tag{5.13}$$

Here $\lambda_b = 1 - e^{-20/1800}$, with gain one at priming. This is a residual-bias correction, related to disturbance augmentation [19], without a Kalman gain or an offset-free performance guarantee under mismatch.

Every 120 s, six scans, the model is discretized by an exact zero-order hold on a 300 s grid. The affine intercept is incorporated by shifting h by a. The 6 h horizon has 72 samples and twelve held moves of lengths $[1,1,1,2,3,4,5,7,9,11,13,15]$. The program is

$$\begin{aligned} \min_{v,s} \quad & 0.15\sum_{j=1}^{12}(v_j - v_{j-1})^2 \\ & + 0.04\sum_{k=1}^{72} v_{\beta(k)} + 400 s^2 + s \\ \text{subject to} \quad & 0 \le v_j \le v_{avail}, \quad |v_j - v_{j-1}| \le 0.4, \\ & y_k + s \ge r - (r - \hat{m}_o)_+ e^{-300k/7200}, \\ & s \ge 0. \end{aligned} \tag{5.14}$$

Table 5.4 Cases and What Each Teaches

#	Scenario and event	Teaching point
1	Warm, 1000 Sm^3/d	where the weakest point is
2	Cut to 550 Sm^3/d at 1 h, no feedforward	a transit of delay
3	500 Sm^3/d, water 20 to 45 %, liquid basis	dose the water
4	800 Sm^3/d, sea 3.5 °C colder at 1 h	what the estimate blames
5	Insulation ×6, 2.25–3.0 km, at 1 h	where no sensor looks
6	350 Sm^3/d, pumps to 25 % at 1 h	rate and dose limits
7	Trip at 1 h, no methanol	the no-touch time
8	Cold restart at 0.5 h, no lead, 20 %/min	inhibit before flow
M	Turndown and planned stop, MPC	reduced-model comparison
9	36 h cold restart challenge	graded restart

The index $\beta(k)$ selects the block, v_0 is the previous applied rate clipped into $[0, v_{avail}]$, and $(\cdot)_+$ denotes its positive part. Numerical weights apply in the stated rate and temperature units. There is one global slack s for the entire horizon. Pump availability and storage set v_{avail}, with zero at shut-in or empty storage. The 0.4 m^3/h bound applies per decision or future block boundary. The recovery funnel and margin band are soft. Pressure is handled by the original choke override, not by a QP pressure constraint.

Normally $r = m_{SP}$. Preview can change it only after a planned shutdown request enters pre-dose, to the larger of m_{SP} and the nominal uninhibited margin plus $\Delta T(W_{pd})$ and a declared forecast error. No future restart is available. The sequence owns non-producing doses, and the low-margin override, pump and storage caps retain priority. Acceptance checks optimal status, finite arrays, every held move, slew, slack and recomputed future predictions. Rejection holds the first applied rate before protections, then returns the classical dosing law. A valid recovery plan also holds its first move. Exact trim tracking permits a negative concentration correction below feedforward. Subsequent fallback scans can rise to the classical feedforward floor, and a zero-water inverse cannot encode a positive pump. While MPC remains selected, its nominal observer feeds the sequence, override, fallback and reported estimates even after failure. The original observer carries the same declared defaults, and over 120 runs of 36 h, 100 of them with a physical parameter changed, the two estimates agree at every step, so fallback applies the original dosing law to the same estimate. Retained forecasts describe the last valid plan, separately from actual ownership and the nonacting classical candidate terms.

5.6 Measured Cases

The cases and the lesson each carries are in Table 5.4, and the measured outcome of each is in Table 5.5. Unless a case names a change, the feedforward doses on the water rate, the trim, the low-margin override, the pressure override and the pre-dose are on, the margin setpoint is 3 °C, the lead 30 min, both doses 30 wt % and the restart ramp 5 %/min, and the flow is held to the pumps. Every value comes from a headless run of the engine with its fixed seed, 20261103, and times count hours of simulated time from reset. A setting named for a case is changed after the case is loaded and before it runs, and carrying the same setting through a reset gives the same run. Where the transmitter noise decides a result, the spread over three other seeds is given beside it. The application lists case M as its ninth scenario, "MPC versus feedforward at turndown", and the challenge after it without a number. This chapter keeps the numbering of Table 5.4, in which case 9 is the challenge.

5.6.1 The Weakest Point and a Turndown

Case 1 is the design state of Section 5.4 and needs no methanol. Its +6.41 °C at the riser base is the whole budget of the line, and each later case removes part of it by a different route. In case 2 the platform asks for 550 Sm^3/d at 1 h, and the riser base starts a slow slide from 23.3 °C toward

Table 5.5 Measured Outcome of Each Case

Case	Measured outcome
1	design margin +6.41 °C, segment 21, transit 4.14 h
1, planned shutdown	pre-dose 5.75 h, 22.7 m^3, +2.99 °C at 30 h
2, trim alone	exposed 8.55 h, −4.43 °C at 24 h, 132 °C·h, 47.4 m^3
2, feedforward on	worst +1.03 °C at 7.15 h, 28.3 m^3
3, liquid basis	9.0 wt % at 12 h, worst −1.73 °C, 36.9 °C·h
3, water basis	worst +1.54 °C, pump to 3.49 m^3/h, 63.7 m^3
4, corrections on	worst +1.68 °C, θ_f 1.12 at 24 h
4, corrections off	at 12 h estimate +3.16 °C, true +0.63 °C
5, flowing	exposed 5.59 h, worst −2.19 °C, 19.4 °C·h
5, trip at 12 h	predicted 6.68 h, curve at 2.875 km after 1.36 h
6, as shipped	245 °C·h, worst −3.21 °C
6, setpoint 1000 at 2 h	398 Sm^3/d, 227 °C·h
6, setpoint 100	after 1 min 9.45 °C·h, at 2 h 49.4 °C·h
7, trip	predicted 3.28 h, plant 3.64 h, 3066 °C·h by 24 h
8, no lead	exposed after 12.5 min, worst −4.14 °C, 12.1 °C·h
8, lead, slow ramp	worst +1.17 °C, +2.19 °C
8, lead, Nielsen–Bucklin	worst +0.003 °C, no exposure

the 12.6 °C of the new steady state. The feedforward is off, so methanol can come only from the trim and the override. The trim begins to dose at 4.71 h, once the estimate has sunk below the 3 °C setpoint, and from 7.19 h the override holds the pumps at full capacity, yet each dose enters liquid that needs about 5.9 h more to reach the riser base. The first exposure comes at 8.55 h, 7.55 h after the cut. The heavy dosing then lands around 16 h on liquid that no longer needs it (+20.5 °C), the pumps shut off as the estimate recovers, and the liquid that entered undosed in that pause arrives last, so the day closes at −4.43 °C with 132 °C·h of exposure and 47.4 m^3 of methanol spent. Switched on, the feedforward starts the pump 2.3 min after the cut, on the steady demand the estimator foresees, 14.7 wt % and 1.29 m^3/h at 2 h, the methanol rides in the same liquid as the cold, and the margin bottoms at +1.03 °C (7.15 h) on 28.3 m^3. Over three other noise seeds the trim alone gathers 128 to 130 °C·h and the feedforward bottoms between +1.04 and +1.06 °C, so the comparison does not rest on the noise.

🕮 **Pause.** Try Exercise 5.2 before reading on.

5.6.2 More Water and a Colder Sea

Case 3 holds 500 Sm^3/d at 17.9 wt % (1.31 m^3/h), enough for the 3 °C margin, with the ratio fixed to the liquid rate. From 1 h the well's water cut climbs from 20 toward 45 % with a 1 h time constant. The water rate rises by more than a factor of two while the pump barely moves from 1.3 m^3/h, so the aqueous concentration at the riser base thins from 17.9 to 9.0 wt % by 12 h, while the denser column lifts the pressure at the riser base from 88.0 to 106.1 bar by 24 h, and the equilibrium with it. The extra water brings heat as well, and the riser base warms from 11.2 °C before the breakthrough to 14.2 °C by 24 h, but the margin still reaches −1.73 °C at 12.57 h and the line gathers 36.9 °C·h. A ratio on the water rate raises the pump with the water, to 3.49 m^3/h at most, keeps the margin at +1.54 °C or above (+1.21 to +1.76 °C over three other seeds), and spends 63.7 m^3 against 55.8 m^3. The price of protection scales with the water and not with the oil.

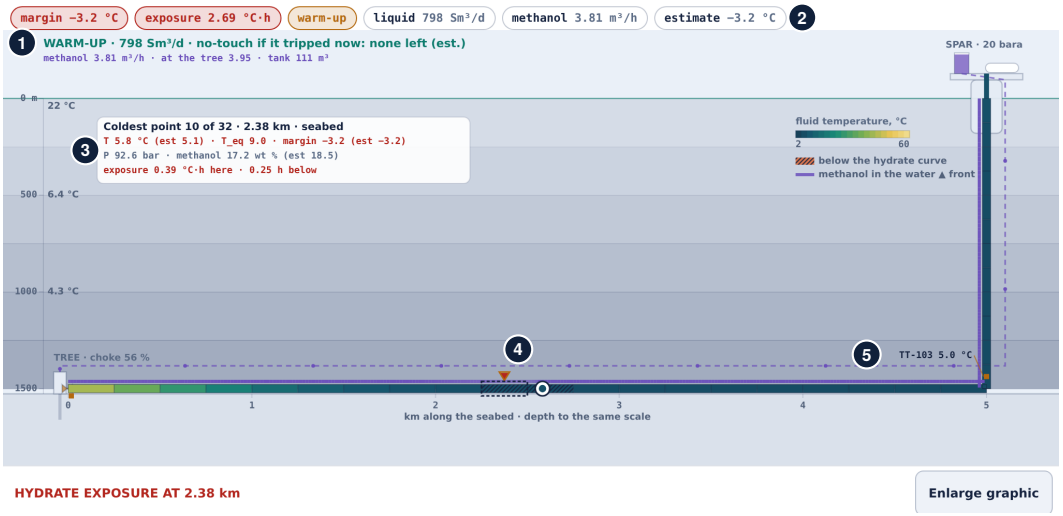

Figure 5.4 Case 8, delayed-injection restart, 3.00 h into the run and 2.5 h after the choke began to open with no inhibitor lead, cropped to the plant view. (1) The header: the sequence is in warm-up at 798 Sm^3/d, by the estimate no no-touch time would be left if the well tripped now, and the pump runs at 3.81 m^3/h with 3.95 m^3/h reaching the tree. (2) The badges: a margin of −3.2 °C and 2.69 °C·h of exposure. (3) The readout of segment 10: 5.8 °C against an equilibrium of 9.0 °C, a margin of −3.2 °C (estimate −3.2 °C), 92.6 bar and 17.2 wt % of methanol in its water (estimate 18.5). (4) Segments 10 and 11, from 2.25 to 2.75 km, are hatched below the hydrate curve, the true and estimated coldest points are both marked in segment 10, and the ringed tagged parcel at 2.58 km marks the front of the liquid that entered after the restart. The new fluid warms the first segments from 42.6 °C while the line beyond 2 km lies between 4.8 and 6.1 °C. (5) TT-103 at the riser base reads 5.0 °C. The status line at the bottom reports hydrate exposure at 2.38 km.

> **TRY IT IN THE APP** **3 · Higher water cut**
>
> Load the case and press Run with the ratio on the liquid rate. The pump stays near 1.3 m^3/h while the water rate more than doubles, and the margin reaches −1.73 °C. Press Reset, set Methanol ratio taken against to Water rate and run again. The pump follows the water, to 3.49 m^3/h at most, and the margin stays at +1.54 °C or above.

Case 4 runs 800 Sm^3/d on 2.3 wt % (0.23 m^3/h) with the riser base at 19.0 °C, and at 1 h a cold current drops the seabed water by 3.5 °C. Nothing tells the estimator. With its corrections on, the profile is pulled through TT-103 and TT-104, the colder line shows up in the present-state margin of the feedforward, the pump climbs to 0.92 m^3/h by 24 h, the riser base settles near 16.5 °C, and the true margin touches +1.68 °C at 5.29 h. The flowline factor rises to $\theta_f = 1.12$ over the day, since a colder sea and 12 % more heat loss through the insulation look the same to two thermometers. With the corrections off, the estimate still reads +3.16 °C at 12 h against a true +0.63 °C, the true margin reaches +0.52 °C at 14.98 h, and the pump stays at or below 0.42 m^3/h.

Pause. Try Exercise 5.3 before reading on.

5.6.3 A Flooded Section and a Pump Shortfall

In case 5 seawater floods the insulation between 2.25 and 3.0 km at 1 h and multiplies its conductance by six. The wall downstream of the damage releases its stored heat for some hours, so at 4 h the truly coldest point is segment 14 at 3.375 km, just past the flood, where the estimate still places it

in segment 21. Methanol starts at 4.71 h, when the estimate first calls for it, and the override runs the pumps at full capacity from 5.55 h, too late for the liquid already on its way. Exposure begins at 5.59 h in segment 18 at 4.375 km, reaching the last seabed segment by 5.70 h and the riser base at 5.06 km by 5.92 h, for 19.4 °C·h in all and a worst margin of −2.19 °C at 8.08 h. The estimator attributes the loss to the whole flowline, $\theta_f = 1.44$ by 24 h. A trip at 12 h then finds an estimated line carrying 10.5 to 27 wt % of methanol from the trim and the override, and the estimator gives it 6.68 h of no-touch time. The flooded section reaches its own curve 1.36 h after the trip, at 2.875 km, far from any thermometer, whereas the undamaged line of case 1 tripped at the same time first reaches the curve in the last seabed segment after 3.63 h, against a prediction of 3.34 h.

Case 6 models a late-life well at 350 Sm^3/d with 30 % water and a platform limited to 400 Sm^3/d of liquid, dosed at 23.6 wt % (1.95 m^3/h). A pump failure at 1 h leaves 1.00 m^3/h. The ratio now calls for 1.97 m^3/h on average, from 1.70 to 2.29 m^3/h as the measured water rate varies, and the loop saturates and stays there. Exposure starts 8.4 h after the failure at 3.875 km, reaches the last seabed segment at 11.07 h and the riser base at 5.06 km at 11.59 h, and totals 245 °C·h by 24 h with a worst margin of −3.21 °C. A setpoint of 1000 Sm^3/d at 2 h runs into the platform limit at 398 Sm^3/d and still collects 227 °C·h. For the lower-rate comparison, the scenario is reloaded at 350 Sm^3/d, run for one minute, then cut to 100 Sm^3/d. That step shrinks the water needing dose but lets the line cool further. Another 24 h gives 9.45 °C·h, with exposure beginning near 15.2 h and a worst margin of −0.67 °C in the last seabed segment. Made at 2 h, after an hour of under-dosing, the same cut leaves 49.4 °C·h with −5.51 °C at 3.125 km. Choosing 100 Sm^3/d before Run instead initializes a line already settled and dosed at that rate, which remains protected at +2.90 °C, and Reset gives the identical startup. A high startup target is capped at the platform capacity, so a 1000 Sm^3/d target in this case starts at 400 Sm^3/d. These are different starting histories from a live rate change. The pumps set the methanol per kilogram of water and the platform sets the flow that keeps the line warm, so no setting of these loops protects the line as it stands at the failure. The application reports the deficit as exposure instead of lending the pumps a capacity they lack.

5.6.4 The No-Touch Time

Case 7 trips the well at 1 h with no methanol in the line. At the trip the estimator gives the line 3.28 h, and the last seabed segment meets its curve 3.64 h after the trip, 0.36 h later than the lumped estimate, because the fluid in the plant hands its heat to the wall before the two cool together. From then on the uninhibited line is simply exposed. At 24 h the riser base is at 5.6 °C under an equilibrium of 16.9 °C, the last seabed segment holds the lowest margin, −11.85 °C, and the exposure has reached 3066 °C·h. The scripted trip of case 7 fires again after every reset and overrides a pre-dose, so the planned shutdown is run on the same design state in case 1. Fired there at 1 h, a planned shutdown pre-doses for 5.75 h, throttling the flow to as little as 703 Sm^3/d to let the pumps reach the dose, and spends 22.7 m^3 before the estimate reports 27 wt % or more in every segment. Only then does the choke close, and the stopped line still has +2.99 °C at 30 h, 23.25 h later. The length of the pre-dose depends on the noise, 5.8 to 9.4 h over three other seeds, because a slice dosed just under 27 wt % while the throttled flow wanders has to reach the platform before the choke may close.

> ✦ DID YOU KNOW? **A Clock That Can Run Fast or Slow**
>
> From the simulator (Table 5.5). On the intact line of case 7 the estimator gives a no-touch time of 3.28 h, and the plant takes 3.64 h. After the flood of case 5 it gives 6.68 h, and the flooded section reaches its curve after 1.36 h.

5.6.5 A Cold Restart

Case 8 opens on a line that was pre-dosed to 30 wt % at 800 Sm^3/d and then left shut in for 24 h, so the riser base sits at 5.1 °C with +5.05 °C of margin and the fluid in the well has cooled to 14.5 °C. At 0.5 h the choke begins to open at 20 %/min, and the pumps start at that instant with no lead. The

Table 5.6 Challenge Runs of 36 h Against the Target

Configuration	Exp. (°C·h)	Worst (°C)	Oil (Sm^3)	MeOH (m^3)	Met
as shipped	155	−4.82	548	103.1	no
suggested[a]	0	+1.38	706	81.9	yes
suggested, flow held to pumps	0	+1.30	529	119.6	no
suggested, no feedforward	1.13	−0.30	706	86.7	no
suggested at 1000 Sm^3/d	0	+0.46	753	74.2	yes

[a] 30 min lead, feedforward on, flow not held to the pumps, 900 Sm^3/d

ceiling reaches fully open within 5 min and the sequence enters warm-up at 0.58 h, after which the dose follows the measured flow, still small, and the pump drops from 4.00 to 0.46 m^3/h before rising again with the flow. Through the umbilical only 59 % of a step has arrived after 15 min, so the first new liquid meets seabed-cold pipe under-dosed. The estimate goes negative 11.0 min after the restart and the truth 12.5 min after it, in the first segment.

From there the under-dosed slug carries its exposure along the line. At 3.00 h it fills segments 10 and 11, from 2.25 to 2.75 km, with its leanest slice at 17.2 wt % and the fluid at 5.8 °C under an equilibrium of 9.0 °C, as shown in Figure 5.4. Behind it comes liquid with the restart dose, 28 to 32 wt %, and ahead of it the pre-dosed contents of the line at 30 wt %. As the slug reaches the foot of the riser the margin falls to −4.14 °C at 4.96 h in the last seabed segment, and the case ends with 12.1 °C·h. A 30 min lead lets the umbilical deliver before the choke opens and keeps the margin at +1.17 °C or better. A 5 %/min ramp without a lead keeps +2.19 °C or better by another route, since the flow grows only as fast as the methanol can follow. Whichever variant is run, the warm-up ends 15.9 h after the restart and 62.1 to 62.6 m^3 have been pumped, about half the tank. The 30 wt % dose also lies past the range of the Hammerschmidt relation. Nielsen–Bucklin gives 15.5 °C of depression there instead of 17.3 °C, and with it the 30 min lead restart keeps only +0.003 °C. The size of the under-dosed slug depends on the noise in the first minutes of flow. Over three other seeds the restart without a lead gathers 3.9 to 6.7 °C·h, the 30 min lead keeps +0.03 to +1.08 °C, and the Nielsen–Bucklin run dips to between −0.07 and −0.96 °C, an uncertainty that only a design margin on the dose can absorb.

5.6.6 The Challenge

The challenge repeats the cold restart at 0.5 h and adds a water breakthrough from 20 to 40 % at 6 h and a platform cut to 600 Sm^3/d at 18 h, over 36 h. Meeting it requires that no wet liquid be exposed, that the flowline stay under its 120 bar limit, that no more than 90 m^3 of methanol be pumped and that at least 650 Sm^3 of oil reach the platform, and Table 5.6 compares five configurations against that target. The shipped settings, which omit the lead and the feedforward and hold the flow to the pumps, expose the line 12.5 min after the restart and end at 155 °C·h, 548 Sm^3 of oil and 103 m^3 of methanol. The suggested settings, a 30 min lead, the feedforward on, the flow no longer held to the pumps and a 900 Sm^3/d setpoint, finish the warm-up at 14.5 h, keep the margin at +1.38 °C or more, and deliver 706 Sm^3 for 81.9 m^3 with the flowline pressure never above 106.1 bar. Restoring the hold on the flow leaves the line cold 5.1 h longer and costs oil and methanol together, since the warm-up lasts until the line is warm and the pumps run at capacity throughout it. Dropping the feedforward leaves the 18 h turndown to the trim, which answers it late enough to expose the riser base from 30.2 h. A 1000 Sm^3/d setpoint passes as well, with more oil and less methanol, but its worst margin is only +0.46 °C, which the target does not penalize. That margin is within the noise. Over three other seeds the 1000 Sm^3/d setpoint fails once, with 0.1 °C·h of exposure, while the suggested settings pass every time with worst margins of +0.06 to +0.91 °C. Figure 5.5 sets the five configurations side by side against the limits of the target.

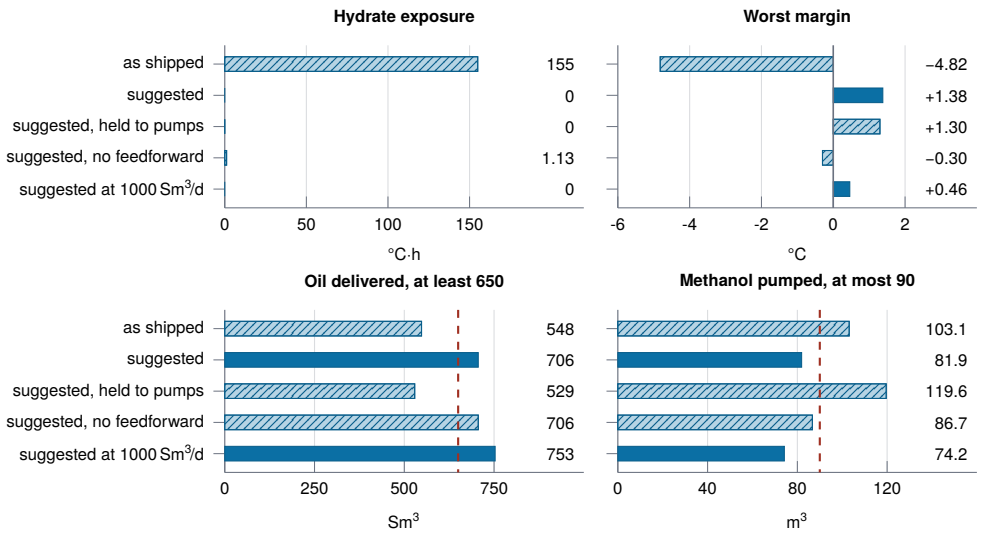

Figure 5.5 The challenge runs of Table 5.6 as bars: exposure, worst margin, oil delivered and methanol pumped over 36 h. Solid bars meet the target, hatched bars miss it, dashed red lines are the oil floor and the methanol cap, and the value of each bar stands at the right. Notice that holding the flow to what the pumps can dose, a cautious rule, costs oil and methanol together, and that of the two runs that pass, the faster setpoint keeps the thinner margin.

The predictive controller made 1080 solves over 36 h in Chrome 157 on the measurement host. Mean times were 56.481 µs for DenseQP alone, 370.648 µs for solve and acceptance, 352.130 µs for model construction and 723.519 µs per decision. Complete controller work, including both observers, took 970.500 ms. Normalized by simulated duration this is 0.000 748 843 % of one core at 1× and 5.391 67 % at 7200×. A short actual 1× window advanced no decisions beyond priming. Actual top-speed playback ran near 60 frames/s without natural frame-budget stops. Guard Off, On and forced exhaustion produced matching 720-step, 60-solve records, with 359 forced stops. The 25 ms guard bounds work within a frame without dropping reached steps or decisions. These measurements show no natural pacing benefit.

5.6.7 Comparison With Predictive Dosing

Table 5.7 adds paired 36 h runs at seed 20261103, with identical instruments and fixed MPC settings. F denotes the feedforward-and-trim structure with the feedforward on and its ratio on the water rate, rather than the feedforward-off setup of case 2 or the liquid-basis setup of case 3. I uses that same law with one PI derived before the cross-case runs. A nominal pump step from 1 to 1.5 m³/h gives gain 0.493 370 °C per wt %, delay 12 575 s and time constant 915 s. The IMC rule $K_c = \tau/[K(\lambda + \theta)]$, $\lambda = \theta$, $\tau_I = \tau$ gives 0.073 741 2 wt % per °C and 915 s for every case [20]. Tables 5.5 and 5.6 keep the windows and settings of the earlier sections.

Case M cuts production at 1 h, requests shutdown at 12 h and requests restart at 30 h. F, I and MPC use 77.818922, 77.809399 and 93.316 154 m³, respectively. F and I have no exposure and a minimum of +1.025 663 °C, whereas MPC reaches −1.739 696 °C and 8.300 350 °C·h. The latter sums segment/slice depths as in Equation (5.8). Integrating only the negative minimum gives the distinct 3.023 019 °C·h. The reduced prediction smears a lean front and the economic objective spends its soft margin. The worst MPC margin occurs at 40 620 s, before shutdown, as shown in Figure 5.6. Both structures close at 124 660 s (34.63 h). The 30 h restart request is ignored during pre-dose, so this case does not demonstrate a completed restart. Separate planned and cold-restart runs test accepted restart commands.

Table 5.7 Additional 36-Hour Comparisons With Water Feedforward, Frozen PI and MPC

Case and law	Minimum (°C)	Exposure (°C·h)	Methanol (m^3)	Margin IAE (°C·h)	Dose travel (m^3/h)	At cap (%)	Mean solve[a] (µs)
Turndown							
F	+1.026	0.000	44.101	34.83	55.91	0.00	—
MPC	−1.740	8.300	105.310	138.85	73.40	36.90	303.0
I	+1.026	0.000	44.812	35.91	56.08	0.00	—
Water cut							
F	+1.544	0.000	92.166	15.96	115.12	0.00	—
MPC	−4.993	62.760	106.136	93.99	47.75	38.52	526.0
I	+1.540	0.000	93.102	18.62	116.05	0.00	—
Cold sea							
F	+1.685	0.000	30.186	12.77	50.25	0.00	—
MPC	−0.134	0.050	79.592	108.56	83.93	22.95	499.2
I	+1.675	0.000	31.053	16.47	51.21	0.00	—
Planned stop/restart							
F	+3.306	0.000	69.981	184.96	55.16	26.91	—
MPC	+3.306	0.000	69.981	184.96	55.16	26.91	523.3
I	+3.306	0.000	69.981	184.96	55.16	26.91	—
Cold restart							
F	−4.141	12.121	63.731	189.19	59.88	21.45	—
MPC	−4.141	12.121	67.535	188.37	72.58	22.84	508.3
I	−4.141	12.121	63.267	187.72	58.90	21.45	—
Pump shortfall							
F	−3.248	486.835	36.925	172.90	2.41	97.22	—
MPC	−3.248	486.548	37.469	172.86	6.14	97.22	520.2
I	−3.248	486.836	36.925	172.90	2.41	97.22	—
Insulation damage							
F	−2.189	19.430	59.713	55.89	70.00	10.69	—
MPC	−2.237	18.955	88.077	104.68	79.51	43.18	508.3
I	−2.191	19.926	60.142	56.36	69.87	10.80	—
TT fault							
F	+0.826	0.000	36.228	51.62	55.79	12.44	—
MPC	−0.510	1.158	68.204	90.57	84.03	19.72	503.9
I	+0.815	0.000	36.485	52.91	55.71	12.48	—
M							
F	+1.026	0.000	77.819	164.76	115.37	0.00	—
MPC	−1.740	8.300	93.316	198.25	117.23	16.99	581.7
I	+1.026	0.000	77.809	164.74	115.34	0.00	—
Challenge, main setup							
F	−4.141	12.121	116.858	124.37	111.51	27.95	—
MPC	−4.696	57.719	120.000	142.64	70.91	47.98	519.3
I	−4.141	12.121	116.858	124.37	111.51	27.95	—

F uses the shipped trim tuning with the feedforward on and its ratio on the water rate. I applies the single predetermined IMC tuning. Preview Off, On, and On with a +3 °C error reproduce every physical row. Every MPC row made 1080 solves and none was rejected, while F and I make no solves. Margin IAE is the shared scorecard integral of setpoint minus the selected observer's estimate in absolute value, divided by 3600. Dose travel is the sum of absolute applied-rate increments. At cap is physical pump-cap occupancy, distinct from the original command-shortfall saturation flag. [a]Fresh Node solve-plus-acceptance wall means vary with host load and are not DenseQP-only or Chrome playback timings.

Preview Off, On and On with a +3 °C error give identical physical results throughout the primary matrix. Candidate targets can change in pre-dose, but sequence ownership leaves zero applied preview

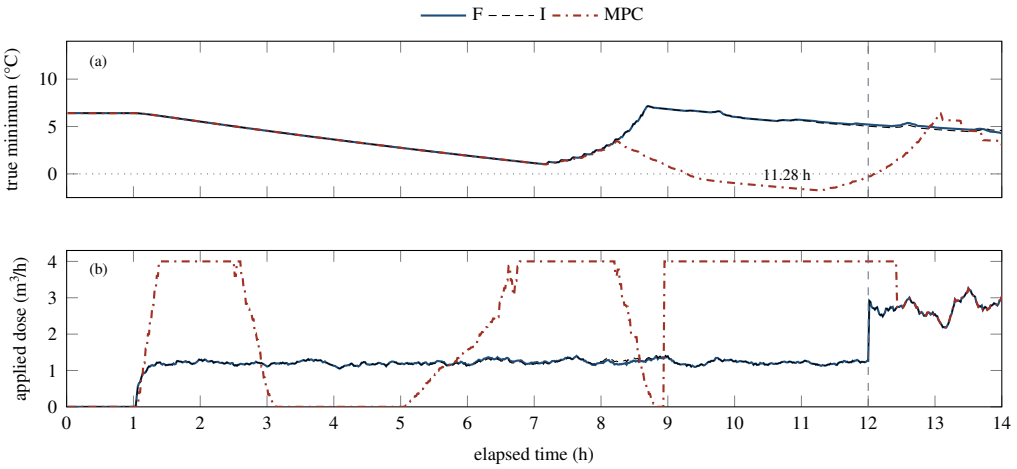

Figure 5.6 Case M through 14 h on common time axes, F solid, I densely dashed and MPC dash-dotted. F and I nearly coincide in true minimum (a), while MPC exposes the lean front despite its larger applied dose (b). The dashed vertical line is the 12 h shutdown request. All curves include the original overrides and sequence, which owns pre-dose actuation. Values are sampled each minute, including the 40 620 s decisive minimum.

benefit. On the unchanged suggested challenge, F and I pass with 81.945284 and 81.669 589 m^3. MPC also avoids exposure and delivers the same 706.130 866 Sm^3 of oil at the same 106.127 bar maximum, but its 102.055 818 m^3 fails the 90 m^3 cap. On the exact original turndown setup, MPC reduces exposure from 276.379 to 8.300 °C·h over 36 h while methanol rises from 75.341 to 105.310 m^3. On the exact liquid-basis water-cut setup it increases exposure from 51.727 to 62.760 °C·h and methanol from 96.416 to 106.136 m^3. The comparison reports these losses as measured.

The 180 primary, 42 diagnostic, 12 exact-setup and 32 forced-fallback rows cover 4821120 finite, physically bounded steps. Halving the physics step changes primary worst margins by at most 0.019 46 °C, exposure by 0.114 68 °C·h and methanol by 1.8443 %, without reversing paired primary rankings. With noise removed, original-versus-fixed PI methanol ordering reverses for turndown and insulation damage. Changed-input diagnostics establish neither numerical convergence nor universal rankings. Twenty corrupted actual solver returns are rejected. Replay of 900 actual packets from each warm and cold history reproduces the nominal observer exactly, whereas a conditionally advanced observer fails that check.

5.7 Technology in Practice

Figure 5.7 draws the choke loops, the profile estimator, the methanol dosing and the sequence of Section 5.5 with the hardware of a deepwater tieback.

What Measures It

The choke loops read the multiphase meter and PT-102. The hydrate margin that the methanol loops protect lives only in the estimator, which the thermometers at the riser base and the arrival pull toward the plant. The dose also rests on the meter's water cut, which matters more as a field matures and the water cut of its wells rises [21]. In the model AT-100 lags 60 s, with 2 % absolute noise and a settable bias.

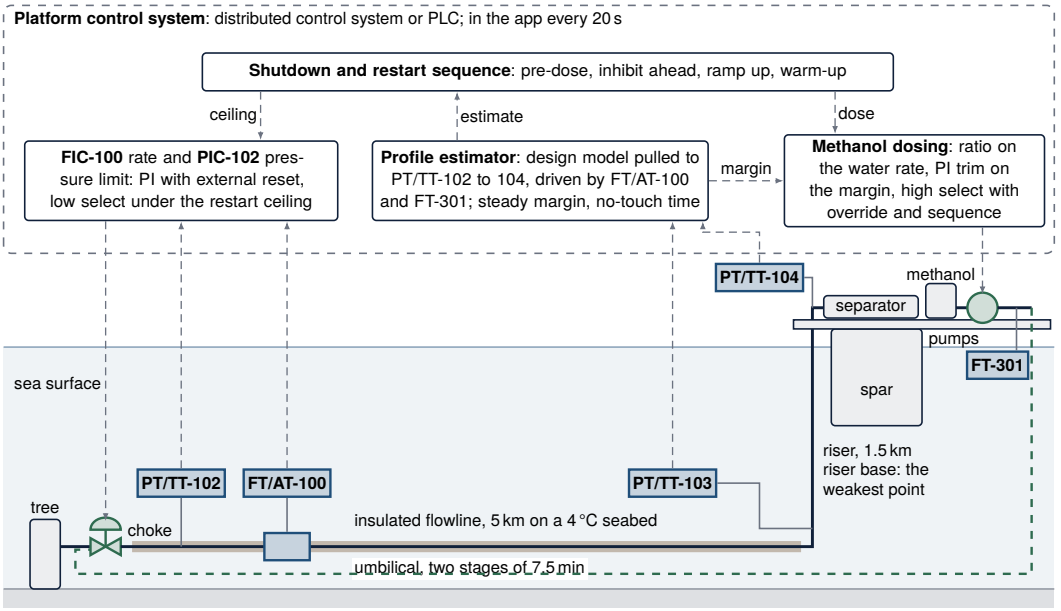

Figure 5.7 The loops of the flowline simulator with their hardware. Lettered boxes are sensing devices with the chapter's tags: the pressure and temperature transmitters PT/TT-102 after the choke, PT/TT-103 at the riser base and PT/TT-104 at the arrival, the subsea multiphase meter FT/AT-100 and the pump's flow meter FT-301. The subsea choke and the methanol pumps are the final elements, the dashed green line is the umbilical and dashed arrows are signals. Notice that no transmitter reads the hydrate margin: every methanol decision passes through the estimator, and the dose leaves the platform hours before it reaches the riser base. Not to scale, and not a piping diagram.

> **TECHNOLOGY** **Sensor: The Temperature Transmitter at the Riser Base**
>
> TT-103 at the riser base is the main witness of the margin estimate. A resistance thermometer is a winding of pure metal wire, usually platinum, whose resistance always rises with temperature [22]. Where it must be replaceable without depressurizing the pipe it sits in a thermowell, whose mass adds a first-order lag. A 4 to 20 mA signal reveals a broken wire by the loss of current, and transmitters that follow the NAMUR levels report a failed sensor below 3.6 mA or above 21 mA [22]. The chapter gives TT-103 a 30 s lag and 0.05 °C of colored noise, and a failed TT-103 reports nothing, so the estimator knows that it is blind there.

What Moves It

Two final elements share the line. The subsea choke at the tree sets the flow for FIC-100 and gives way to PIC-102 near the pressure limit, and in the model its actuator travels at 20 %/min. The methanol pumps on the platform set the dose, and the umbilical carries it down to the tree, where it enters the liquid it has to protect.

TECHNOLOGY Final Element: The Methanol Metering Pump

Metering pumps inject small quantities of a chemical into a process stream. When the dose must follow the treated flow, a 4 to 20 mA signal sets the output between zero and full flow, and many such pumps use a plunger that delivers the chemical in discrete pulses rather than as a steady flow [22]. The chapter's ratio on the water rate is that proportioning. Its pumps deliver 4 m^3/h at most into an umbilical that passes 59 % of a step to the tree after 15 min. When a pump fails in case 6, no setting of the loops protects the line.

Where the Controller Runs

The loops would run in the platform's control system. A distributed control system controls production within one geographic location, usually through feedback and feedforward loops, and loops in such systems run with cycle times from milliseconds to minutes [23]. The app's 20 s control period sits inside that range and is far shorter than the hours the methanol needs to reach the riser base, so no scan rate can make feedback on the margin timely. The chapter's structure maps onto two PID blocks with external reset feeding a low select on the choke, a ratio block with a PI trim and a high select on the pump, a sequence for shutdown and restart, and the profile estimator in a supervisory layer above them.

IN PRACTICE The No-Touch Clock

When a subsea line stops, a 2021 review describes a no-touch time of about half a day, in which the platform tries to resume production while the pressurized line cools toward hydrate conditions at the seabed. For a stop longer than that, operators act to keep hydrates out, for example by displacing the line with dead oil, called bullheading, or by depressurizing it below the hydrate pressure at seabed temperature [10]. The chapter's design line has 3.28 h by its estimator, and the application offers no depressurization.

What Is Modernizing It

Three developments touch this chapter. Low dosage inhibitors, kinetic inhibitors and anti-agglomerants, are an alternative to thermodynamic inhibitors and are used in the field [15]. A 2021 review places them beside alcohols, glycols and salt among the chemical methods of hydrate management, and finds industry trending toward safe operation inside the hydrate region [10]. Virtual flow meters infer the phase rates from ordinary field instruments and empirical algorithms. They cost far less than other multiphase meters but need much calibration before start-up, and a 2019 review expects them to progress [21]. Fiber-optic sensing along a pipeline reads the temperature around it all the way along the fiber [24], the profile that the two thermometers of case 5 cannot give.

5.8 Limitations

The hydraulics are one lumped rate over quasi-steady pressures with a constant slip ratio. Terrain and severe riser slugging, which on a real line are control problems of their own [6] and reach the separator of Chapter 6 as slugs, are absent, and the liquid in a shut-in riser stays where it stopped.

The first-order energy transport on 250 m segments puts the design margin 0.38 °C above its 80-segment value.

The hydrate curve is one representative gas, illustrative and fitted between 20 and 200 bar, the shortcut inhibition relations are used beyond their range at the 30 wt % restart dose, and the share of methanol that partitions into the hydrocarbons is fixed at 10 %.

Exposure is thermodynamic. Nothing forms or obstructs the flow, the plug-risk index is an illustrative mapping whose kinetic uncertainty is an order of magnitude, and a transient formation model would be needed to say whether a given exposure plugs the line [16].

OPEN CHALLENGE When Does Exposure Become a Plug?

Thermodynamics now predicts hydrate formation to within about 10 % in pressure and 1 K in temperature, while kinetic studies give up at least an order of magnitude of that accuracy. Practice has moved from avoiding the hydrate region toward managing operation inside it, and a 2021 conceptual model of how a plug forms on a restart still needs scientific verification [10]. A 2022 review calls for kinetic models of hydrate formation in real crude oils, tested in flow loops [25]. Progress would be a validated prediction of when a given exposure blocks a line.

The estimator knows only the design insulation and sea temperature, and two thermometers are all it has to correct them. Distributed temperature sensing, low dosage inhibitors [15] and depressurization of a shut-in line are not offered.

OPEN CHALLENGE Finding Heat Loss Along Kilometers of Pipe

An optical fiber laid with a pipeline turns the line into a thermometer. A laser pulse scatters inside the fiber, the ratio of two Raman components of the light gives the local temperature, and the travel time gives the place [24]. Over long distances the spatial resolution is limited to the order of meters and worsens with distance, and a cleaner signal costs a longer averaging time [24]. In case 5 two thermometers spread a local flood over the whole flowline. Progress would be a profile fine and fast enough along a subsea line to locate the damage before a shutdown finds it.

The predictive reduction is not an exact forecast of the slice plant. Four nominal pump steps have transient residuals up to 1.5522 °C and RMS 0.3690 °C, although steady gains agree to roundoff. Its bias follows an inferred margin and cannot identify unmeasured local insulation damage. Under forced fallback on changed geometry or physical properties, the nominal and original observers differ by 0 °C, since both carry the declared defaults and meet the changed plant as model mismatch. Neither a positive predicted band nor successful QP solution guarantees the true margin or the challenge's cumulative dose budget.

5.9 Conclusion

The hydrate margin of a deepwater flowline is a minimum over 6.5 km of pipe. The simulator computes it on the truth, slice by slice, while its methanol loops work from an estimate built on six transmitters and two meters. At the design rate the smallest margin, +6.41 °C, sits at the riser base while the coldest fluid keeps +12.2 °C at the platform, and a change made at the tree reaches the riser base 3.23 h later. A turndown dosed by feedback on the estimated margin ends at −4.43 °C with 132 °C·h of exposure, while the feedforward from the predicted demand holds +1.03 °C with 28.3 m^3 of methanol instead of 47.4 m^3. A ratio on the liquid rate lets a water breakthrough thin the riser base to 9.0 wt %, a flooded section of insulation meets its curve 1.36 h after a trip for which the estimate allows 6.68 h, and the first liquid of a restart with no inhibitor lead is exposed 12.5 min after the choke opens, where a 30 min lead keeps +1.17 °C. The model keeps quasi-steady hydraulics, one representative hydrate curve and no kinetics. A fuller model would add a two-fluid hydraulic model with slugging, hydrate formation kinetics with the effect of under-inhibition, salt in the water and distributed temperature sensing. The application is an accessible place to practice sending a dose ahead of its transit and finding a constraint that no instrument measures.

The predictive comparison exposes a failure of model reduction. Case M spends 93.316 m^3 with 8.300 °C·h exposure, while water feedforward spends 77.819 m^3 with none.

KEY POINTS

- The coldest fluid is not the greatest risk: at the design rate the smallest margin sits at the foot of the riser, where cooled fluid still carries the full seabed pressure.
- Methanol travels with the liquid, so a change made at the tree reaches the riser base 3.23 h later, and feedback on the margin there doses liquid that has already gone past.
- A feedforward on the predicted demand holds +1.03 °C through the case 2 turndown, where feedback alone ends at −4.43 °C.
- Two thermometers fix how much heat the line loses and say nothing about where it is lost.

Where else this applies. Volume 9, Chapter 5 judges a release at the head long after it was made, and Volume 5, Chapter 11 infers a core temperature that no instrument reaches.

References

[1] P. D. Christofides, *Nonlinear and Robust Control of PDE Systems*, ser. Systems & Control: Foundations & Applications. Boston, MA: Birkhäuser Boston, 2001.

[2] M. Polycarpou, J. Uber, Z. Wang, F. Shang, and M. Brdys, "Feedback control of water quality," *IEEE Control Systems*, vol. 22, no. 3, pp. 68–87, 2002.

[3] E. D. Sloan, "Fundamental principles and applications of natural gas hydrates," *Nature*, vol. 426, no. 6964, pp. 353–359, 2003.

[4] C. A. Koh, E. D. Sloan, A. K. Sum, and D. T. Wu, "Fundamentals and applications of gas hydrates," *Annual Review of Chemical and Biomolecular Engineering*, vol. 2, no. 1, pp. 237–257, 2011.

[5] E. D. Sloan, Jr. and C. A. Koh, *Clathrate Hydrates of Natural Gases*. CRC Press, 2007.

[6] E. Storkaas and S. Skogestad, "Controllability analysis of two-phase pipeline-riser systems at riser slugging conditions," *Control Engineering Practice*, vol. 15, no. 5, pp. 567–581, 2007.

[7] M. T. Rubel and D. H. Broussard, "Flowline insulation thermal requirements for deepwater subsea pipelines," in *SPE Annual Technical Conference and Exhibition*, 1994.

[8] W. I. Wilcox, D. B. Carson, and D. L. Katz, "Natural gas hydrates," *Industrial & Engineering Chemistry*, vol. 33, no. 5, pp. 662–665, 1941.

[9] D. L. Katz, "Prediction of conditions for hydrate formation in natural gases," *Transactions of the AIME*, vol. 160, no. 1, pp. 140–149, 1945.

[10] E. D. Sloan, "Hydrocarbon hydrate flow assurance history as a guide to a conceptual model," *Molecules*, vol. 26, no. 15, p. 4476, 2021.

[11] E. G. Hammerschmidt, "Formation of gas hydrates in natural gas transmission lines," *Industrial & Engineering Chemistry*, vol. 26, no. 8, pp. 851–855, 1934.

[12] M. H. Yousif, "Effect of under-inhibition with methanol and ethylene glycol on the hydrate control process," in *Offshore Technology Conference*, 1996.

[13] K. H. Bendiksen, D. Maines, R. Moe, and S. Nuland, "The dynamic two-fluid model OLGA: Theory and application," *SPE Production Engineering*, vol. 6, no. 2, pp. 171–180, 1991.

[14] E. D. Sloan, "A changing hydrate paradigm—from apprehension to avoidance to risk management," *Fluid Phase Equilibria*, vol. 228–229, pp. 67–74, 2005.

[15] M. A. Kelland, "History of the development of low dosage hydrate inhibitors," *Energy & Fuels*, vol. 20, no. 3, pp. 825–847, 2006.

[16] L. E. Zerpa, E. D. Sloan, A. K. Sum, and C. A. Koh, "Overview of CSMHyK: A transient hydrate formation model," *Journal of Petroleum Science and Engineering*, vol. 98–99, pp. 122–129, 2012.

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

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

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

[20] D. E. Rivera, M. Morari, and S. Skogestad, "Internal model control: PID controller design," *Industrial & Engineering Chemistry Process Design and Development*, vol. 25, no. 1, pp. 252–265, 1986.

[21] L. S. Hansen, S. Pedersen, and P. Durdevic, "Multi-phase flow metering in offshore oil and gas transportation pipelines: Trends and perspectives," *Sensors*, vol. 19, no. 9, p. 2184, 2019.

[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.

[24] J. Li and M. Zhang, "Physics and applications of Raman distributed optical fiber sensing," *Light: Science & Applications*, vol. 11, p. 128, 2022.

[25] A. A. Almashwali, C. B. Bavoh, B. Lal, S. F. Khor, Q. C. Jin, and D. Zaini, "Gas hydrate in oil-dominant systems: A review," *ACS Omega*, vol. 7, no. 31, pp. 27 021–27 037, 2022.

5.10 Exercises

5.1 Concept check · 1 min

At the design rate of 1000 Sm^3/d, where along the line does the smallest hydrate margin sit, and why there?

(a) At the platform arrival, because the fluid is coldest there at 19.4 °C.

(b) At the foot of the riser, where the fluid has cooled longest but still carries the seabed pressure.

(c) Just downstream of the choke, where the Joule–Thomson cooling of the choke acts.

(d) At the top of the riser, where the pressure falls toward the 20 bara of the platform.

5.2 Concept check · 2 min

In case 2 the platform cuts the flow to 550 Sm^3/d and only the trim on the estimated margin and the low-margin override can add methanol. Why does the riser base still fall below the hydrate curve?

(a) The trim gain of 0.5 wt % per °C is too low, and a larger gain would hold the margin.

(b) The estimator cannot see the riser base, so the trim acts on a margin that is wrong.

(c) Methanol added once the estimated margin falls enters liquid that reaches the riser base about 5.9 h later.

(d) The pumps lack the capacity to protect the line at the lower flow.

5.3 Concept check · 2 min

In case 4 the seabed water turns 3.5 °C colder and nothing tells the estimator. With its corrections on, the flowline heat-loss factor θ_f climbs to 1.12 over the day although the insulation is intact. Why?

(a) The insulation conducts more heat as the sea around it cools.

(b) The film coefficient falls because the flow drops at the colder sea.

(c) TT-103 drifts with the colder water and pulls the estimate down.

(d) Two thermometers fix how much heat the line loses but not why, so a colder sea is read as more heat loss through the insulation.

5.4 Theory and calculation · 6 min

Transport sets the timing of every dose in this line. Use the chapter's model values.

(a) The seabed flowline holds 134.7 m^3 of liquid at the design rate. How long does a change made at the tree take to reach the riser base at 1000 Sm^3/d, and at 550 Sm^3/d with the same liquid volume? *Checkpoint:* 3.23 h and about 5.9 h.

(b) The umbilical is two first-order stages of 7.5 min each. What fraction of a step of the pump rate reaches the tree after 15 and after 30 min, and after how many minutes has half of it arrived? *Checkpoint:* 59 %, 91 % and about 12.6 min.

(c) The chapter measures a methanol step in manual at half strength at the riser base 3.49 h after the pump moves. Add your results of (a) and (b) and account for the difference.

5.5 Theory and calculation · 6 min

Case 6 runs at 350 Sm^3/d with 30 % water, dosed at 23.6 wt %, and at 1 h the pumps fall to 1.00 m^3/h. Use Equation (5.10) with $\rho_w = 1025\,kg/m^3$, $\rho_M = 791\,kg/m^3$ and $f_{loss} = 10\,\%$.

(a) What pump rate delivers 23.6 wt % to the water before the failure? *Checkpoint:* about 1.95 m^3/h.

(b) What aqueous concentration can 1.00 m^3/h deliver at the same water rate? *Checkpoint:* about 13.7 wt %.

(c) By how much does the Hammerschmidt depression fall between the two concentrations, and what does that say about the exposure that follows?

5.6 Theory and calculation · 5 min

The estimator predicts the no-touch time of a stopped line with Equation (5.9), a lumped cooldown of fluid and wall together.

(a) Take the last seabed segment at the design rate: fluid at 23.8 °C, uninhibited equilibrium 17.2 °C, sea 4.04 °C and the lumped time constant 8.2 h. Ignore the difference between fluid and wall. How long after a trip does this segment reach its curve? *Checkpoint:* about 3.3 h.

(b) In case 7 the estimator gives 3.28 h at the trip and the plant's last seabed segment meets its curve after 3.64 h. Is the estimate safe here, and why does the plant take longer?

(c) After the insulation damage of case 5 a trip at 12 h is given 6.68 h, yet the flooded section reaches its curve after 1.36 h. What in the equation was wrong for that section?

5.7 App challenge · 8 min

Load scenario 2, "Lower production rate", default seed, and use speed 900 (a run of 24 h takes about a minute and a half).

(a) Press Run as shipped, with Feedforward from the predicted demand Off, and pause after 24 h. Read the Worst margin tile. The exposure and methanol tiles keep counting after 24 h, so export the CSV and read `exposure_total_Ch` and `metric_meoh_used_m3` in the row at `time_s` 86 400. *Checkpoint:* about –4.43 °C, 132 °C·h, 47.4 m^3.

(b) Press Reset, set Feedforward from the predicted demand to On before Run, run past 24 h again and read the same three values.

(c) In the CSV of the first run, find the first row in which `margin_true_C` is negative and its location in `margin_true_at_km`. Explain both runs with the transit at 550 Sm^3/d.

Give the values to 0.05 °C, 1 °C·h and 0.2 m^3.

5.8 App challenge · 9 min

Load scenario 8, "Delayed-injection restart", default seed, and use speed 3600 (each 8 h run takes about eight seconds). The choke opens at 0.5 h. Make every change before pressing Run, and pause at or after 8 h. The two tiles no longer change after about 6 h.

(a) Run 8 h as shipped (lead 0 min, ramp 20 %/min). Read the Hydrate exposure and Worst margin tiles. *Checkpoint:* about 12.1 °C·h and –4.14 °C.

(b) Press Reset, set Restart inhibition lead to 30 min, and run 8 h. Read the same tiles.

(c) Reload the scenario with its reload button, which restores the lead of 0 min, set Choke ramp on restart to 5 %/min, and run 8 h.

(d) Reload again, set the lead to 30 min and Methanol depression to Nielsen–Bucklin, and run 8 h. What does the Worst margin tile say about the 30 wt % restart dose?

Give the margins to 0.02 °C.

5.9 App challenge · 10 min

Load the Challenge, "cold restart and water breakthrough", default seed, and use speed 7200 (each 36 h run takes about 18 seconds). The target is no exposure, the flowline below 120 bar, at most 90 m^3 of methanol and at least 650 Sm^3 of oil.

(a) Run past 36 h as shipped and pause. Oil and methanol keep counting, so export the CSV and read, in the row at `time_s` 129 600, `exposure_total_Ch`, `metric_worst_margin_C`, `metric_oil_Sm3`, `metric_meoh_used_m3` and `metric_worst_P_bar`. *Checkpoint:* about 155 °C·h, 548 Sm^3, 103.1 m^3.

(b) Press Reset and, before Run, set Restart inhibition lead to 30 min, Feedforward from the predicted demand On, Hold the flow to what the pump can dose Off and Production setpoint, liquid, to 900 Sm^3/d. Run past 36 h and read the same five columns.

(c) Press Reset, set Hold the flow to what the pump can dose back On, run past 36 h, and read oil and methanol in the same row. Use Equation (5.11) to explain them for a water cut of 20 and of 40 %.

Give oil to 2 Sm^3, methanol to 0.2 m^3 and the margin to 0.02 °C.

5.10 Application area · 4 min

The model counts exposure thermodynamically, so an under-dosed slug counts as partly protected. Field experiments have seen hydrate blockages form more readily in under-inhibited lines than in lines with no inhibitor.

(a) In case 8 without a lead, the leanest slice of the slug carries 17.2 wt % against a restart dose of 30 wt %. How should an engineer read the case's 12.1 °C·h of exposure in the light of that field experience?

(b) The restart dose of 30 wt % lies past the range of the Hammerschmidt relation. What would you do before writing it into a procedure?

5.11 Application area · 4 min

The line carries thermometers only at the riser base and the arrival.

(a) In case 5 the insulation floods between 2.25 and 3.0 km. Why does neither thermometer locate the damage, and which measurement would?

(b) Until such a measurement exists, what operating rule would you attach to the estimator's no-touch time after a trip?

5.12 Review the agent · 5 min

An AI assistant reviews case M, the turndown with a planned shutdown, and writes: "Under MPC the line used 93.3 m^3 of methanol against 77.8 m^3 for feedforward and trim. The MPC carries more inhibitor through the turndown, so it gives the larger margin. Switch the line to MPC for safety." Find the flaw and say how you would check the claim.

CHAPTER 6

Offshore Three-Phase Separator Simulator for Slug Absorption and Delayed Quality Measurement

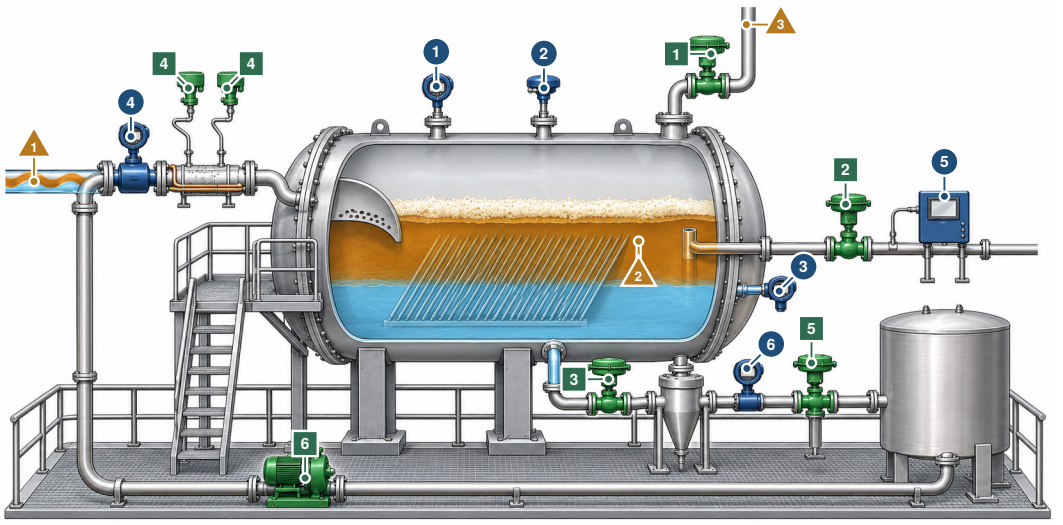

The controller moves 1 gas valve to the compressors 2 oil valve to the second stage 3 water valve to the hydrocyclone 4 demulsifier and antifoam dosing 5 routing valve, sea or hold tank 6 hold-tank recycle pump

The controller reads 1 separator pressure 2 total liquid level 3 oil–water interface 4 inlet multiphase meter 5 water in the export oil 6 oil in the overboard water, four minutes late

What disturbs it 1 a slug from the riser 2 a tighter emulsion 3 gas header pressure

Notice that the slug arrives at the left far faster than the three outlet valves on the right can pass it, and that the oil-in-water reading beside the routing valve describes water sampled minutes earlier. Illustration, not to scale.

IN PLAIN TERMS **A Tank That Takes the Surge**

The main idea. Oil, gas and water arrive mixed from the seabed and flow into a large tank on the platform. After a few quiet minutes the water sinks below the oil, and three valves draw off gas, oil and water.

The challenge. Now and then a huge gulp of liquid arrives far faster than the valves can pass, so the tank must hold it. The meter that checks the cleaned water reports minutes late, after dirty water has gone to sea.

Why it matters. How fast the stored liquid is let out decides whether the next equipment copes and the sea stays clean.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the two History boxes in Section 6.2, the first Did you know? box in the Introduction, the Try it boxes and the second Did you know? box in Section 6.6, and Key points at the end of the Conclusion. No equations.
Study	60 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	25 min	This page, Section 6.6 with Figure 6.5, Section 6.7 and Section 6.8 with its Open challenge boxes.

The control question A riser slug arrives far faster than any outlet valve can pass it, and the analyzers on the oil and water outlets report minutes late. If the separator must store the slug whatever the tuning, what does the level tuning decide, and what keeps the export oil and the overboard water in spec when the measurement arrives after the water has gone?

Learning objectives

- Derive the level gains of a horizontal separator from its geometry and valve capacities, and explain why the gain matrix is triangular.
- Compute the closed-loop time constant and damping of tight and averaging level tuning on an integrating level.
- Explain why no tuning keeps a slug out of the vessel, and measure how the tuning shares it between the vessel, the second stage and the hydrocyclone.
- Estimate what a diversion on a delayed analyzer lets overboard, and compare a faster analyzer, a second hydrocyclone bank and an averaged interface loop.
- Decide which element of a configuration (averaging, feedforward, supervisor or MPC) each clause of a campaign target needs.

Prerequisites Sections 1.4.3 (slugging), 1.4.4 (the level as an integrator, averaging level control), 1.6.2 (the framework's PID), 1.6.3 (for case 9) and 1.6.4 (the apps).

The app https://apmonitor.com/apps/offshore-separation/

See also Volume 6, Chapter 2 for a triangular gain matrix whose relative gain array shows no interaction, and Volume 6, Chapter 7 for inventory control structures. Volume 5, Chapter 7 for diversion on a late analyzer, and Volume 5, Chapter 3 for control through a slow analyzer. Volume 9, Chapter 4 for an offset-free MPC whose estimator carries an unmeasured disturbance, and Volume 9, Chapter 7 for storage that absorbs an inflow under a downstream limit.

6.1 Introduction

An offshore platform sells oil and gas and discharges water, and its first-stage separator decides the quality of all three. The riser delivers the phases together, and the vessel sends oil to a second stage, gas to the compressors and produced water through a hydrocyclone to the sea. The export oil is judged by its water content, the water by its dispersed oil and the gas by the liquid it carries, and each of those depends on droplets that need time and a quiet layer to leave the other liquid [1]. The vessel and the loops that hold its levels set that time, and separator design is organized around retention times, droplet settling and the gas capacity of the vessel [2, 3].

Three properties separate the separator from a set of independent level loops. Every height and the gas pressure follow from volumes in one horizontal cylinder, so the water valve moves the interface and the total level together, the oil valve moves only the total level, and a rising level compresses the gas cap. The separation is a residence-time process, so a disturbance at the inlet reaches the outlet qualities over minutes rather than seconds. The analyzers that report those qualities are late as well, and the one on the overboard line reads through a four-minute sample delay. The class is broad, since any inventory whose level loop exists to smooth or condition an outflow has this shape, from surge drums and flash drums to decanters and settlers, and averaging level control was posed for exactly that class [4].

The wellstream supplies the disturbances. Deepwater risers slug, as liquid collects at the riser base until the gas behind it pushes the column out in one piece [5, 6], water cuts rise as a field ages, and the stability of the emulsion changes with the crude that the wells deliver [7]. Feedback on the topside choke can suppress severe slugging [8, 9], and whatever it leaves reaches the separator. Section 1.4.3 explains how a riser slugs and Section 1.4.4 why a separator level is an integrator.

✦ DID YOU KNOW? More Water Than Oil

A 2021 review reports that produced water rose from less than 30 million barrels a day in 1990 to about 100 million barrels a day in 2015 [10], and a 2022 review estimates about three barrels of it for every barrel of oil [11]. The design wellstream of this chapter is 30 % water.

The simulator addresses the belief that tight level control keeps a slug out of the vessel and out of the equipment downstream at the same time. A slug arrives far faster than any outlet valve can pass it, so the vessel stores most of it under any tuning, and the tuning decides only how fast the stored liquid is handed to the second stage and the hydrocyclone. A second belief is that a diversion triggered by the oil-in-water analyzer protects the sea, when the analyzer reports on water that has already gone overboard. Control-oriented separator models have been published for loop design and estimation [12, 13]. However, a transfer function from valve opening to level carries neither the capacity downstream of the valve nor the delay of the analyzer, and those two quantities decide what a slug costs.

This chapter follows a slug, a water cut or a fault from the riser to each outlet of a simulated offshore three-phase separator, and on to what its analyzers report later, with the geometry, the droplet populations, the water treatment and the instruments modeled together. Gravity separation, averaging level control, feedforward and rule-based supervision are all in routine industrial use, and the optional model predictive controller [14] is measured against the loops on the same slug and inflow step. Every height of the simulated vessel comes from one volume function, its four cells carry three droplet classes per phase with balances that close to round-off, and its record of what left each outlet is kept apart from the instruments. The application offers nine cases and a campaign in which the cost of a tuning choice, an instrument delay or an equipment limit is measured in the units of the plant. Section 6.2 reviews how separators came to be sized and controlled, Section 6.3 gives the model, Section 6.4 measures how its loops are coupled, Section 6.5 states the controller as implemented and Section 6.6 reports the cases. Section 6.7 sets the hardware of a platform separator beside the model, Section 6.8 lists what the model leaves out, and the exercises in Section 6.10 put the cases to work.

6.2 Separator Sizing, Riser Slugs and Surge Control

Separator design began as a sizing problem. The allowable gas velocity through a vessel or a mesh pad is still stated through the capacity factor that Souders and Brown introduced in 1934 for the entrainment limit of fractionating columns [15], and the simulator states the gas load on its mist extractor in that form. Liquid sections were sized by two competing arguments, one requiring a droplet of a design size to settle before its liquid leaves and one requiring the liquid to be held for a design time, and the two theories were compared for oil and water separators in 1990 [2]. Handbook practice for three-phase vessels sizes the oil and water sections with both kinds of argument [3].

Beneath those rules lie droplet settling, coalescence and emulsion stability. The sedimentation of water and oil dispersions, from the coalescence of single droplets to the internals that promote it in gravity separators, has been reviewed [1], as have the stability of crude-oil emulsions and their treatment [7]. Foam received a measure of its own in 1938, when Bikerman defined a unit of foaminess as the ratio of the steady foam volume to the gas flow that sustains it [16], which makes the foam height proportional to the gas velocity through the liquid surface, the form the simulator uses. The produced water then passes treatment whose technologies have been reviewed [17], and the separation achieved by a deoiling hydrocyclone has been measured as its flow varies [18].

Riser slugging became a control problem of its own, as the first History box below tells, and a stability criterion for severe slugging was published in 1986 [6].

> ⧖ HISTORY **1980 to 2002: Slugs From the Riser**
>
> In 1980 an experimental pipeline and riser, built to simulate an offshore line, showed two kinds of slug flow. Severe slugging appeared at low gas and liquid rates where the pipeline sloped down toward the riser, and it produced slugs from one to several riser heights long. A mathematical model agreed well with the measurements, and choking the flow was found to remove the severe slugging [5]. By 2002 field tests and dynamic simulations had shown feedback control holding such lines stable, instead of leaving the receiving facilities to cope with the slugs. In many cases the control also let the pipeline inlet run at a lower pressure, which means a higher production rate [8]. The separator of this chapter is such a receiving facility, and its wells deliver their slugs with no control upstream.

The separator's own loops were treated as a surge problem. A level controller on a surge vessel is detuned on purpose so that the vessel absorbs flow variations [4], and the second History box below follows averaging level control from its optimal form into model predictive control. Three-phase separators with their pressure and level loops have been modeled for control design [12], and a first-principles gravity separator model has been built for control and estimation [13]. The simulator belongs to that line and puts the instruments, the water treatment and a record of what each outlet actually carried into the same model.

> ⧖ HISTORY **1986 and 1989: A Level Allowed to Move**
>
> In 1986 an optimal averaging level control problem was defined and solved for a surge tank, and two controllers were built on its solution, a ramp controller and an optimal predictive controller. Against averaging level controllers reported before, for step and sinusoidal changes in the inlet flow, the optimal predictive controller filtered most of the disturbances better and was very simple to implement [4]. In 1989 Campo and Morari posed the problem with the rate of change of the outlet flow as its objective and then as model predictive control. Their controller minimized the largest rate of change of the outlet flow, kept integral action and respected limits on the tank level and the outlet flow, and a single adjustable parameter traded flow filtering against fast settling [19]. The move price of case 9 makes the same trade.

6.3 Separator Model

The simulated vessel is a horizontal cylinder of diameter $D = 3.0\,\mathrm{m}$ and effective length $L = 10.0\,\mathrm{m}$ between the inlet diverter and the outlets, held at 10.0 barg with the total level at 1200 mm and the interface at 500 mm. At design it receives 149.8 $\mathrm{m^3/h}$ of oil, 64.2 $\mathrm{m^3/h}$ of water and 22 080 $\mathrm{Sm^3/h}$ of gas, a water cut of 30 % and a gas-oil ratio of 150 $\mathrm{Sm^3/Sm^3}$. The riser, the inlet heater, the four cells, the plate pack, the mist extractor, the three outlet valves, the hydrocyclone, the routing valve, the hold tank and every instrument appear in Figure 6.1.

Table 6.1 sorts the quantities into controlled, manipulated and disturbance variables, outputs and states. Thirteen lumped states and 24 dispersed inventories are integrated together with 15 ledgers that accumulate every flow across the vessel boundary, and the slug schedule, the trip latch, the stuck valves and the routing valve are discrete plant state updated after each step.

6.3.1 Heights and the Gas Cap

The liquid volume below a height h follows from the circular segment, as shown in Equation (6.1).

$$\begin{aligned} V(h) &= L\left[R^2\cos^{-1}\frac{R-h}{R} - (R-h)\sqrt{2Rh-h^2}\right] \\ h_{int} &= V^{-1}(V_w + Y) \\ h_{tot} &= V^{-1}(V_w + Y + V_o + X) \end{aligned} \tag{6.1}$$

Here $R = D/2$, the continuous water and oil volumes are V_w and V_o, and X and Y are the total dispersed water in the pad and dispersed oil in the water layer. The inverse is found by table interpolation and Newton steps, exact to 10^{-9} m over the whole diameter, and the scene draws its layers at the same

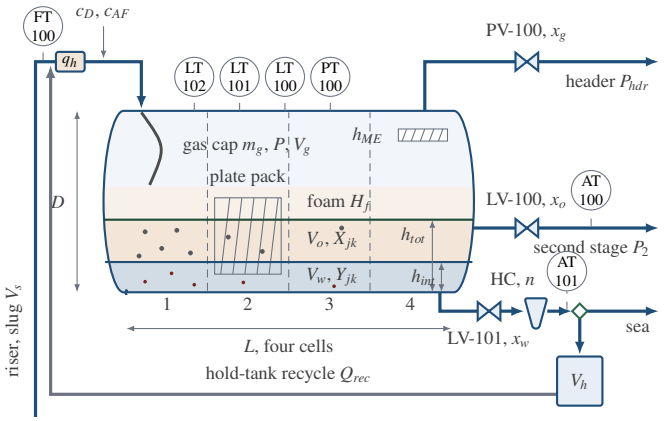

Figure 6.1 Offshore three-phase separator as modeled. The wellstream and its slugs pass FT-100, take in the hold-tank recycle and pass the heater q_h and the chemical injection onto the inlet diverter. The oil pad (V_o, X_{jk}) and the water layer (V_w, Y_{jk}) cross four cells under the foam H_f and the gas cap m_g. Gas leaves through PV-100, oil through LV-100 and AT-100, and water through LV-101, the hydrocyclone HC, AT-101 and the routing valve to the sea or to the hold tank V_h, which recycles to the inlet.

Table 6.1 Classification of the Separator Variables, Design Point and Ranges

Variable	Role	Design value	Range
P	CV, PIC-100 on PT-100	10.00 barg SP	8–11.5
h_{tot}	CV, LIC-100 on LT-100	1200 mm SP	900–1800
h_{int}	CV, LIC-101 on LT-101	500 mm SP	300–850
u_g, u_o, u_w	MV, valve commands	35.3, 50.2, 47.7 %	0–100
c_D, c_{AF}	MV, treating chemicals	8, 2 ppm	0–40, 0–15
q_h, n	MV, heater, banks	1.5 MW, 1	0–3, 1–2
Q_{rec}, r	MV, recycle, throughput	0, 100 %	0–20, 50–110
Q_l, wc	DV, liquid, water cut	214 m^3/h, 30 %	100–300, 0–70
GOR, E	DV, gas-oil ratio, emulsion	150, 1.0	50–450, 0.5–3
P_{hdr}, T_{arr}	DV, header, arrival	8.0 barg, 40 °C	6–12, 15–70
b_l, V_s	DV, LT-101 bias, slug	0 mm, 15 or 40 m^3	±300, events
BS&W, OiW	outputs, AT-100, AT-101	0.104 %, 12.4 mg/L	limits 0.5, 30
m_g, V_o, V_w	states	375 kg, 18.6, 7.7 m^3	—
H_f, T, q_h	states	543 mm, 49.5 °C, 1.5 MW	—
x_g, x_o, x_w	states, valve positions	0.353, 0.502, 0.477	0–1
x_{in}, r	states, inlet valve, wells	1, 1	0–1, 0–1.2
V_h, O_h	states, hold tank	10, 0.02 m^3	0–100
X_{jk}, Y_{jk}	states, 24 dispersions	4 cells, 3 classes	—

heights. A volume change moves a level by $\mathrm{d}V/[L\,w(h)]$ with the surface width $w(h) = 2\sqrt{2Rh - h^2}$, so the vessel holds 29.4 m^3 per meter of level at 1200 mm and 22.4 m^3 per meter at the 500 mm interface. The gas fills whatever the liquid leaves, as shown in Equation (6.2).

$$
\begin{aligned}
P &= \frac{m_g Z \bar{R} T}{M V_g} \\
V_g &= \max\left(\pi R^2 L - V_{liq},\, 0.02\,\pi R^2 L\right) \\
\frac{\mathrm{d}m_g}{\mathrm{d}t} &= w_{g,in} - K_g\, x_g \sqrt{\rho_g\, \Delta P_g} \\
\Delta P_g &= \mathrm{clamp}\left(P - P_{hdr},\, 0,\, P/2\right)
\end{aligned}
\tag{6.2}
$$

Here m_g is the gas mass, $Z = 0.97$, $M = 20.0\,\mathrm{kg/kmol}$, T the liquid temperature in kelvin, $\rho_g = PM/(Z\bar{R}T)$, K_g the effective flow area of the open gas valve and P_{hdr} the pressure of the compressor header, and the cap at $P/2$ stands in for choked flow. At design the cap holds 374.8 kg of gas in 44.3 m^3, 34.0 kg per bar, and the arriving gas replaces it every 72 s. A rising level compresses the gas, by 0.25 bar per cubic meter of liquid at a fixed gas mass.

6.3.2 Outlet Valves, Wells and the Inlet

The two liquid valves are linear, with rated capacities K_v in cubic meters per hour at a differential of 1 bar, as shown in Equation (6.3).

$$\begin{aligned} Q_o &= \phi\, K_{v,o}\, x_o \sqrt{\frac{P - P_2}{SG_o}} \\ Q_w &= \phi\, K_{v,w}\, x_w \sqrt{\frac{P + g(\rho_o h_{pad} + \rho_w h_{int}) - P_{hc}}{SG_w}} \end{aligned} \tag{6.3}$$

where the differentials are expressed in bar, $\phi = \mathrm{clamp}(h_{tot}/0.1\,\mathrm{m}, 0, 1)$ stops both flows as the vessel empties, $h_{pad} = h_{tot} - h_{int}$, $P_2 = 3.0$ barg is the second stage, $P_{hc} = 4.0$ barg the hydrocyclone inlet and SG the specific gravities 0.85 and 1.025. Fully open at design the oil valve passes 298.4 m^3/h and the water valve 134.3 m^3/h. Each position follows its command u in percent through $\tau_v\,\mathrm{d}x/\mathrm{d}t = u/100 - x$, with 3 s for the liquid valves and 2 s for the gas valve, and a stuck stem sets $\mathrm{d}x/\mathrm{d}t = 0$.

The wells deliver an imposed stream that the separator pressure does not affect. At full throughput the oil is the liquid rate Q_l times $1 - wc$, the water the rest, and the gas $w_g = \mathrm{GOR}\, Q_o(\rho_o/\rho_{os})\rho_{gs}$ in kilograms per second, with the standard densities $\rho_{os} = 865\,\mathrm{kg/Sm}^3$ and $\rho_{gs} = 0.846\,\mathrm{kg/Sm}^3$. What arrives is each of these times the throughput r and the inlet valve position x_{in}, plus any slug, and the recycle from the hold tank joins it at the inlet. The wells follow a throughput request through $\tau_w\,\mathrm{d}r/\mathrm{d}t = r_{req} - r$ with $\tau_w = 180$ s, and an inlet shutdown valve closes with a 4 s time constant on a trip. A slug of volume V_s launched at the riser base climbs for 60 s and arrives as a half-sine at the stream's water cut over $t_a = 60$ s, while the gas behind it is held back, as shown in Equation (6.4).

$$q_s = \frac{\pi V_s}{2 t_a} \sin\frac{\pi t'}{t_a}, \qquad w_g = w_{g,0}\left[1 - 0.8 \sin\frac{\pi t'}{t_a}\right] \tag{6.4}$$

Here t' is the time since arrival began. The gas then blows down over 120 s as a half-sine whose peak adds $1.2\,w_{g,0}$. A 15 m^3 slug peaks at 1414 m^3/h on top of the normal 214 m^3/h, and the 40 m^3 slug of the event button arrives over 90 s. The inlet heater duty follows its command through a 60 s lag up to 3 MW, the stream enters at $T_{in} = T_{arr} + q_h/\sum \dot{m} c_p$, 49.5 °C at the design 40 °C arrival and 1.5 MW, and the liquid follows $C_{liq}\,\mathrm{d}T/\mathrm{d}t = \sum \dot{m} c_p\,(T_{in} - T)$ with C_{liq} its heat capacity in the vessel.

6.3.3 Droplets, Cells and Coalescence

The inlet diverter disperses $e_w = 4\,\%$ of the incoming water into the oil pad and $e_o = 0.07\,\%$ of the oil into the water layer. Each dispersion is split into three size classes, 60, 180 and 540 µm for water drops and 15, 45 and 135 µm for oil drops, by a log-normal distribution of spread 0.8 whose median follows Equation (6.5).

$$d_{50} = d_{50,ref}\left(\frac{\rho v^2}{2.5\,\mathrm{kPa}}\right)^{-0.3} \frac{1 + 0.8\,c_D/(c_D + 20)}{E^n} \tag{6.5}$$

Here ρv^2 is the momentum of the whole stream at the 0.40 m nozzle, with its ratio to 2.5 kPa floored at 0.05, c_D the demulsifier dose in ppm, E the emulsion tightness, $n = 1$ for water drops and 0.3 for oil drops, $d_{50,ref}$ is 275 µm and 55 µm, and the class boundaries are the geometric means of neighboring sizes. At design, with 8 ppm of demulsifier and 2.51 kPa of momentum, the water median is 338 µm.

Each class settles or rises at its Stokes velocity, corrected once for the particle Reynolds number in Equation (6.6).

$$v_k = \frac{g\,\Delta\rho\,d_k^2}{18\,\mu_c}\,\frac{1}{1+0.15\,Re_k^{0.687}} \qquad \mu_o = \mu_{o,40}\,\exp\!\left[B_o\left(1/T - 1/313.15\,\mathrm{K}\right)\right] \tag{6.6}$$

where $\Delta\rho = 175\,\mathrm{kg/m^3}$, μ_c is the viscosity of the continuous phase, $\mu_{o,40} = 9.0\,\mathrm{cP}$ and $B_o = 3000\,\mathrm{K}$, and the water follows the same law with 0.65 cP and 1800 K. At the design 49.5 °C the oil is at 6.8 cP, the water classes sink through it at 0.05, 0.45 and 3.86 mm/s against an overflow rate, oil flow over interface area, of 1.86 mm/s, and the oil classes rise through the water at 0.04, 0.35 and 2.81 mm/s.

Along the vessel the pad and the water layer are each four well-mixed cells, the inlet zone, the plate pack and two settling cells, under a level taken to be the same in all four. The dispersed volume X_{jk} of water class k in pad cell j obeys Equation (6.7).

$$\frac{\mathrm{d}X_{jk}}{\mathrm{d}t} = B_j c_{j-1,k} - B_{j+1} c_{jk} - s_k c_{jk} - \kappa_j X_{jk} + \kappa_j X_{j,k-1} \qquad c_{jk} = \frac{N X_{jk}}{V_{pad}}, \qquad s_k = \min\!\left(\frac{v_k\, w(h_{int})\, L}{N}, \frac{V_{pad}}{N\,2\,\mathrm{s}}\right) \tag{6.7}$$

Here $N = 4$, B_j is the pad flow into cell j, taken upwind when it reverses, the first cell receives $e_w Q_{w,in}$ times the class fraction, the last cell's outflow is the pad draw, and the cap on s_k keeps a vanishing layer from stiffening the step. Coalescence moves each class into the next larger one at the rate $\kappa = [1 + 3c_D/(c_D + 20)]\sqrt{\mu_{o,40}/\mu_o}/(\tau_{c,o} E)$ with $\tau_{c,o} = 320\,\mathrm{s}$, four times faster in the plate pack, and the largest class coalesces into drops that fall straight onto the interface, so its coalescence joins the settling flux. The oil drops Y_{jk} follow the same equations with their rise velocities, $\tau_{c,w} = 900\,\mathrm{s}$, $E^{0.3}$ in place of E, no viscosity factor and a plate pack factor of 2.

The cell flows follow from continuity under the common level. The pad inflow is $B_0 = Q_{o,in}(1 - e_o) + e_w Q_{w,in}$, each $B_{j+1} = B_j + \rho_j - \sigma_j - \dot{V}_{pad}/N$ with σ_j and ρ_j the volumes settling out of and rising into pad cell j, and the water layer mirrors it. The continuous phases gain what leaves the dispersions, which Equation (6.8) states.

$$\frac{\mathrm{d}V_o}{\mathrm{d}t} = Q_{o,in}(1 - e_o) + \textstyle\sum_j \rho_j - D_p(1 - C_p) \qquad \frac{\mathrm{d}V_w}{\mathrm{d}t} = Q_{w,in}(1 - e_w) + \textstyle\sum_j \sigma_j - D_w(1 - C_w) \tag{6.8}$$

where D_p and D_w are the draws from the last pad and water cells and C_p and C_w the dispersed fractions there. A layer thinner than its outlet's draw zone, $\delta_o = 250\,\mathrm{mm}$ for the pad and $\delta_w = 200\,\mathrm{mm}$ for the water, makes that outlet take the share $f = [(\delta - h)/\delta]^2$ of its flow from the other phase. The pad draw is then $D_p = (1 - f_{dw})Q_o + f_{do} Q_w + q_{co}$, with q_{co} the liquid carried into the gas, the export oil carries BS&W $= (1 - f_{dw})C_p + f_{dw}(1 - C_w)$, and the separator's water outlet carries $[(1 - f_{do})C_w + f_{do}(1 - C_p)]\rho_o$ of oil, expressed in milligrams per liter. At design 93.9 % of the dispersed water and 74.6 % of the dispersed oil are removed, leaving 0.104 % BS&W and 353 mg/L of oil in the water.

6.3.4 Foam, Mist Extractor and Carryover

Foam grows toward a height proportional to the superficial velocity j_g of the arriving gas through the liquid surface $w(h_{tot})L$, and the gas load on the mist extractor is its capacity factor, as shown in Equation (6.9).

$$\tau_F \frac{\mathrm{d}H_f}{\mathrm{d}t} = \Sigma\, j_g - H_f, \qquad \Sigma = \frac{\Sigma_0\sqrt{\mu_o/\mu_{o,40}}}{1 + c_{AF}/K_{AF}} \qquad K = \frac{Q_g}{A_{ME}}\sqrt{\frac{\rho_g}{\rho_o - \rho_g}} \tag{6.9}$$

Here $\Sigma_0 = 50\,\mathrm{s}$, the antifoam dose c_{AF} halves Σ at $K_{AF} = 3\,\mathrm{ppm}$, $\tau_F = 60\,\mathrm{s}$, Q_g is the actual gas flow through the valve and $A_{ME} = 1.2\,\mathrm{m}^2$. Liquid leaves with the gas at $5\,\mathrm{L/MSm^3}$ through a working mesh, multiplied by $1 + 50\varphi_K + 400\varphi_F + 5000\varphi_L$. The factor φ_K rises smoothly from 0 to 1 as K goes from 0.8 to 1.3 times $K_{max} = 0.10\,\mathrm{m/s}$, $\varphi_F = [1 - g_F/0.3\,\mathrm{m}]^2$ clamped to $[0,1]$ with g_F the clearance between the foam top and the mesh at $h_{ME} = 2.50\,\mathrm{m}$, and φ_L is the same for the liquid surface within 0.1 m. At design the foam is 543 mm, K is 0.051 m/s and the carryover is $5.0\,\mathrm{L/MSm^3}$.

6.3.5 Hydrocyclone, Routing and Hold Tank

Each hydrocyclone bank takes $Q_{hc} = 80\,\mathrm{m^3/h}$, and its grade efficiency depends on its load $\ell = Q_w/(nQ_{hc})$ as shown in Equation (6.10).

$$\begin{aligned} \eta(d) &= 1 - \mathrm{e}^{-(d/d_c)^2} \\ d_c &= d_{c0}\left[1 + 3\max(0, \ell - 1) + 5\max(0, 0.4 - \ell)\right] \end{aligned} \tag{6.10}$$

with $d_{c0} = 12\,\mu\mathrm{m}$, and above 130 % load the share $(\ell - 1.3)/\ell$ bypasses untreated. The reject, 2 % of the treated water with the oil separated from it, goes to a $100\,\mathrm{m^3}$ hold tank, and so does the polished water while the routing valve points to hold, otherwise the water goes overboard. A pump returns the tank contents to the inlet at a commanded rate of up to $20\,\mathrm{m^3/h}$, tapering linearly to zero as the tank falls below $0.5\,\mathrm{m^3}$, and an overflow spills untreated water and is recorded. At design one bank runs at 80 % load, removes 79 % of the 15 µm class and all of the larger ones, and passes 12.4 mg/L.

6.3.6 Trips and the Truth Record

Separate safety transmitters read the total level, the interface and the pressure as the truth plus white noise of 3 mm on each level and 0.005 bar on the pressure, through a 1 s filter sampled on the control scan. A reading beyond its setting for 2 s trips the separator on high-high level at 2100 mm, low-low level at 750 mm, high-high interface at 950 mm, low-low interface at 200 mm or high-high pressure at 12.0 barg, closes the inlet valve and latches. Fifteen ledgers integrate every boundary flow with the same step as the inventories. The scorecard counts the oil exported while its true BS&W was within 0.5 %, the water discharged while its true oil in water exceeded 30 mg/L and the worst carryover against $13\,\mathrm{L/MSm^3}$, all from these records, which no loop and no rule of the supervisor can see.

6.3.7 Numbers, Initialization and Integration

The parameters are listed in Table 6.2, and constants fixed in the code are given where they enter the equations above. Every scenario starts from the steady state at its setpoints, found by fixed-point iteration on the cell equations, which are linear in the concentrations once the cell flows are known, with the hold tank at 10 %. At that state every separator derivative is below 10^{-15}. All 52 states advance together under fourth-order Runge-Kutta integration at a fixed 0.25 s step, in the engine of Section 1.6.1. After the 1200 s slug case the ledgers close the oil and water balances to $1.6 \times 10^{-12}\,\mathrm{m^3}$ and $1.7 \times 10^{-12}\,\mathrm{m^3}$ and the gas balance to 3.5×10^{-10} kg of the 6544 kg that arrived. Halving the step moves the peak level of that case by 9×10^{-8} mm and its peak oil outflow by $2 \times 10^{-10}\,\mathrm{m^3/h}$, because the loops keep their fixed 0.5 s scan and only the integration error changes. Both checks are among the application's 50 passing tests.

6.3.8 Assumptions

The wells deliver an imposed stream, so the separator pressure does not feed back on the riser or the reservoir. A real riser responds to the pressure it discharges against, which is how topside choke control acts on slugging [8].

The level is flat along the vessel and the heads are not counted. Waves appear only as noise on the level transmitter, and a real vessel's heads add volume that the $26.4\,\mathrm{m^3}$ surge figure omits.

Liquids are incompressible at constant density, the gas is ideal at a fixed compressibility, and all the gas arrives free, so flashing, solution gas, the shrinkage of the oil and the inversion of an emulsion at high water cut are absent.

Table 6.2 Parameters of the Separator, Marked Visible (v) Where Show Model Offers a Slider and Hidden (h) Where It Does Not

Symbol	Meaning	Value	
D, L	vessel diameter, length	3.0, 10.0 m	v
h_{ME}, A_{ME}	mist extractor	2.50 m, 1.2 m^2	v
D_{noz}	inlet nozzle	0.40 m	h
δ_o, δ_w	outlet draw zones	250, 200 mm	v
—	level trips, high, low	2100, 750 mm	v
—	interface trips	950, 200 mm	v
—	pressure trip	12.0 barg	v
$K_{v,o}$, $K_{v,w}$	liquid valves	104, 55 m^3/h	v
K_g	gas valve area	113 cm^2	v
τ_v, τ_{vg}	actuator lags	3, 2 s	v, h
P_2, P_{hc}	outlet pressures	3.0, 4.0 barg	v
Q_{exp}	second stage capacity	250 m^3/h	v
ρ_o, ρ_w	liquid densities	850, 1025 kg/m^3	v
ρ_{os}, ρ_{gs}	standard densities	865, 0.846 kg/Sm^3	h
M, Z	gas	20.0 kg/kmol, 0.97	v
$\mu_{o,40}$, B_o	oil viscosity	9.0 cP, 3000 K	v, h
$\mu_{w,40}$, B_w	water viscosity	0.65 cP, 1800 K	h
$c_{p,o}$, $c_{p,w}$, $c_{p,g}$	heat capacities	2.0, 4.1, 2.3 kJ/(kg K)	h
τ_w, τ_{ESD}	wells, inlet valve	180, 4 s	v, h
$q_{h,max}$, τ_H	heater	3 MW, 60 s	v, h
N	cells along the vessel	4	fixed
$d_{50,w}$, $d_{50,o}$	inlet medians	275, 55 µm	v
—	log-normal spread	0.8	h
d_k	water drop classes	60, 180, 540 µm	fixed
d_k	oil drop classes	15, 45, 135 µm	fixed
e_w, e_o	inlet dispersion	4, 0.07 %	v
$\tau_{c,o}$, $\tau_{c,w}$	coalescence times	320, 900 s	v
—	plate pack factor	4	v
—	momentum exponent, reference	0.3, 2.5 kPa	h
K_D	demulsifier half dose	20 ppm	h
—	demulsifier on size, rate	0.8, 3.0	h
—	tightness exponent, oil	0.3	h
Σ_0, τ_F	foaminess, foam lag	50, 60 s	v
K_{AF}, c_0	antifoam, base carryover	3 ppm, 5 L/MSm^3	h
K_{max}	mist extractor capacity	0.10 m/s	v
Q_{hc}, d_{c0}	hydrocyclone bank	80 m^3/h, 12 µm	v
—	reject share	2 %	h
V_{hold}, $Q_{rec,max}$	hold tank, pump	100 m^3, 20 m^3/h	v, h
τ_{PT}, τ_{LT}, τ_{IT}	transmitter lags	1, 2, 3 s	h
τ_{MP}, σ_{MP}	inlet meter	5 s, 4 %	v
τ_{BSW}, θ_{OiW}, τ_{OiW}	analyzers	20, 240, 30 s	v, v, h
—	limits, BS&W, OiW, carryover	0.5 %, 30, 13	v

Settling is computed in four vertically well-mixed cells. A mixed cell removes large droplets less completely than a quiescent layer would, so the model errs toward carrying water forward, and the four cells are a residence-time model of the vessel rather than a grid to refine. Every quality correlation, from the droplet median to the coalescence times, the foam, the mist extractor and the hydrocyclone, is illustrative and fitted to no named crude or device.

The limits of 0.5 % BS&W, 30 mg/L of oil in water and 13 L/MSm^3 of carryover are typical figures rather than a regulation for a named site. The oil in water is scored instant by instant, while many discharge limits are monthly averages. Each transmitter is a lag with correlated noise, only LT-101 can carry a bias, and the valves are linear with no stiction band.

6.4 One Inventory, Three Outlets

The three loops share the vessel through volumes, and manual mode measures the sharing directly. With both level loops in manual at their balance positions, 50.24 % for oil and 47.72 % for water, a

Table 6.3 The Two Level Presets on the Identified Integrators, With $\tau_c = 1/(K_c k)$ and $\zeta = \frac{1}{2}\sqrt{\tau_I/\tau_c}$

Loop and preset	K_c	τ_I	τ_c	ζ
total level, tight	5.0 %/cm	240 s	71 s	0.92
total level, averaging	0.4 %/cm	1500 s	890 s	0.65
interface, tight	4.0 %/cm	400 s	150 s	0.82
interface, averaging	1.5 %/cm	600 s	399 s	0.61

5 % step of either valve changes the level rates by the amounts of Equation (6.11), measured against an unstepped run once the valve has settled, at the end of the minute after the step.

$$\begin{bmatrix} \dot{h}_{tot} \\ \dot{h}_{int} \end{bmatrix} = \begin{bmatrix} -0.0281 & -0.0126 \\ 0.0001 & -0.0167 \end{bmatrix} \begin{bmatrix} \Delta u_o \\ \Delta u_w \end{bmatrix} \quad \frac{\text{mm/s}}{\%} \tag{6.11}$$

The entries follow from the geometry. One percent of the oil valve passes 2.98 m^3/h over the 29.4 m^2 surface at 1200 mm, 0.0282 mm/s, and one percent of the water valve passes 1.34 m^3/h, which lowers the 22.4 m^2 interface by 0.0167 mm/s and the whole liquid by 0.0127 mm/s. The matrix is triangular. The oil valve reaches the interface only through the gas pressure, so the pairing of oil with total level and water with interface is forced, the relative gain array is the identity to within 0.3 %, and the interaction runs one way, from the water valve into the total level. The reverse osmosis train of Volume 6, Chapter 2 has the same shape, a triangular gain matrix whose relative gain array says there is no interaction while one loop still disturbs the other. When the water valve cannot act, the level loop can hold the total level only by taking out oil, which thins the pad from below. Both levels are integrating, because the outlet flows depend on the gas pressure and hardly at all on the liquid heights.

The pressure is not integrating. A 1 % step of the gas valve in manual lowers the pressure by 0.091 bar, with 63 % of the change in 23 s, because the gas flow grows with both the gas density and the differential and so regulates the 34.0 kg per bar cap by itself. The pressure is coupled to both levels through the gas volume, 0.25 bar per cubic meter of liquid at a fixed gas mass, and it drives both liquid valves through their differentials, 10.7 m^3/h of oil and 5.2 m^3/h of water per bar.

The surge volume is set by the trips. Between the 1200 mm normal level and the 2100 mm high-high trip the vessel holds 26.4 m^3, and 12.6 m^3 lie between the normal level and the 750 mm low-low trip, while the water layer can gain 11.5 m^3 before the high-high interface trip and lose 5.7 m^3 before the low-low one.

A proportional-integral loop on an integrator k/s has the characteristic polynomial $\tau_c \tau_I s^2 + \tau_I s + 1$ (Section 1.4.4), and with the identified gains k of Equation (6.11) the two level presets give the time constants and damping ratios of Table 6.3. The averaging preset lengthens the response of the total level twelvefold, from 71 s to 890 s, which is fifteen times the 60 s a slug takes to arrive, and all four settings are underdamped, with ζ from 0.61 to 0.92. The pressure loop, 15 % per bar and 30 s on the 0.091 bar per percent first-order response, has a loop gain of 1.4.

6.5 Implemented Controller

The loops execute every 0.5 s, once per two 0.25 s physics steps, the supervisor every 10 s, and one sample is logged every second. Each step executes the control scan when it is due, integrates the plant, and only then evaluates the trips and adds to the truth record. The page offers speeds of 1, 2, 5, 10, 30 and 60 times real time, with 10 as the default.

6.5.1 Instruments

No loop reads the plant's truth. Each transmitter in Table 6.4 passes the true value through its own first-order lag, solved exactly over the scan, and adds Gaussian noise filtered to the stated correlation time and scaled to the stated standard deviation. The radar level adds waves with a standard deviation

Table 6.4 Instruments. Noise Is One Standard Deviation With Its Correlation Time

Tag	Measures	Lag	Noise
PT-100	separator pressure	1 s	0.01 bar, 1 s
LT-100	total level, radar	2 s	2 mm, 2 s, waves
LT-101	interface, displacer	3 s	3 mm, 2 s, bias b_I
LT-102	interface, profiler	10 s	6 mm, 5 s
LT-102	foam top, profiler	10 s	20 mm, 5 s
TT-100	liquid temperature	10 s	0.05 °C, 10 s
PT-150	gas header	1 s	none
FT-100	inlet oil, water, gas	5 s	4 % each, 8 s
FT-101	produced water flow	2 s	1 %, 2 s
AT-100	BS&W, oil outlet	20 s	0.01 %, 10 s
AT-101	oil in water, overboard	240 s + 30 s	1 mg/L, 20 s
LT-200	hold tank level	2 s	none
ZT	valve positions	1 s	0.03 % on gas

of 4 mm for each unit by which the momentum ratio $\rho v^2/2.5\,\mathrm{kPa}$ exceeds 0.8, which is 0.8 mm at design and 21 mm at the peak of a 15 m^3 slug. The oil-in-water analyzer reads the polished water through a pure sample delay of 240 s and then a 30 s lag.

6.5.2 The PI Algorithm

All three loops use the velocity-form PID routine that every framework app shares, described in Section 1.6.2 and written for this separator as Equation (6.12).

$$\begin{aligned} e_k &= a(\mathrm{SP}_k - \mathrm{PV}_k), \quad P_k = K_c e_k, \quad D_k = -K_c \tau_D a \dot{\mathrm{PV}}^f_k \\ S_k &= S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ c_k &= b_k + S_k + I_{k-1}, \qquad u_k = \mathrm{clamp}(c_k, 0, 100) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t}\right]\Delta t_c \end{aligned} \tag{6.12}$$

Here $a = -1$, b_k is the feedforward bias, $\Delta t_c = 0.5\,\mathrm{s}$ and the tracking time T_t equals τ_I. The derivative is switched off by $\tau_D = 0$ in all three loops, and no output has a rate limit. In manual the loop sets $I_k = u_{op} - b_k - S_k$, so its output equals the operator's valve and the return to automatic is bumpless. Every process gain here is negative, since opening a valve lowers its variable, so $a = -1$ and each loop is direct acting under the ISA convention, opening its valve as its variable rises. The same algorithm shows why a pinned valve does not wind the integral without bound. With u_k at a limit u_{lim} and $T_t = \tau_I$, the last line of Equation (6.12) reduces to $\tau_I\,\mathrm{d}I/\mathrm{d}t = u_{lim} - b - I$, so the integral term settles at $u_{lim} - b$ whatever the error [20].

6.5.3 Loops, Tunings and Feedforward

PIC-100 moves the gas valve on PT-100 with $K_c = 15\,\%$ per bar and $\tau_I = 30\,\mathrm{s}$ against a 10.00 barg setpoint. LIC-100 moves the oil valve on LT-100 and LIC-101 the water valve on LT-101, or on the LT-102 profiler when the operator selects it, with the presets of Table 6.3 written to both loops by one selector and a custom setting once either gain is edited. Any loop can be put in manual on its own, and the mode switch in the header moves all three together. A loop that starts a run in manual holds the balance opening the start-up computed.

With feedforward on, each bias is the opening that would pass the metered inflow of the loop's own phase at the measured differential, from the valve's rated capacity, as shown in Equation (6.13).

$$b = 100\,\frac{\tilde{Q}_{in}}{Q_{max}(P_m)}, \qquad \tau_{FF}\,\frac{\mathrm{d}\tilde{Q}_{in}}{\mathrm{d}t} = Q_{in,m} - \tilde{Q}_{in} \tag{6.13}$$

Here Q_{max} is the flow of the fully open valve by Equation (6.3) or (6.2) at the measured pressures and, for water, the measured head, $\tau_{FF} = 900\,\text{s}$ for the two liquids, so that a slug is not passed straight through, and 2 s for the gas. The recycle flow, which the meter does not see, is added to the metered liquids as 1 % oil and 99 % water. Switching the feedforward moves each integral by the old bias minus the new one, so no valve jumps.

6.5.4 Supervisor, Routing and Diagnostics

A rule-based supervisor, off by default, scans every 10 s. It trims the demulsifier on AT-100 at 0.6 ppm/s per point of BS&W above a 0.20 % target, up to 40 ppm, and then the heater at 30 kW/s per point up to 3 MW, giving back heat first and chemical second at a tenth of those rates once the reading is 0.05 % below the target. It adds 0.5 ppm of antifoam per scan, up to 15 ppm, while the profiler's foam top is within 0.45 m of the mist extractor, and removes 0.1 ppm per scan down to the operator's setting once the gap exceeds 0.8 m. It brings in the second hydrocyclone bank after the measured water flow has stayed above 85 % of one bank for 20 s and withdraws it after 120 s below 70 %. It runs the hold-tank recycle at 12 m^3/h while the tank is above 15 % and the separator has room. Finally it cuts the throughput request in 10 % steps down to 50 % for a reason it cannot otherwise remove, a hold tank above 70 % while diverting, BS&W above its limit for 300 s with both trims at their limits, a level or interface heading for its trip with its valve at its limit or stuck, a pressure within 0.3 bar of its trip with the gas valve open, or an estimated mist extractor load above 85 % of K_{max}. A reason must persist for 30 s and the cuts are 120 s apart, or 180 s and 360 s for the gas load so that the wells can answer, and the request is restored by 5 % after 300 s without a reason.

Independently of the supervisor, the routing valve sends the polished water to the hold tank while AT-101 reads above 30 mg/L and back overboard after 60 s below 25.5 mg/L. Two diagnostics watch the instruments. A valve whose position feedback differs from its command by more than 10 % for 20 s raises a not-following alarm, and LT-101 and LT-102 differing by more than 60 mm for 60 s raise a disagreement alarm.

6.5.5 The Model Predictive Controller

A control structure selector offers one model predictive controller (MPC) in place of the three loops, which stay the default of every case but case 9. Section 1.6.3 describes how the apps set up such an option, and the paragraphs below give this controller's model and objective. It moves all three outlet valves every $T_s = 2\,\text{s}$ over $N_p = 150$ steps (5 min) in 12 move blocks, from 1 step at the start to 105 at the end, and the supervisor and the routing valve act as under the loops. Its model is the balance of Section 6.3 at the declared parameter values, never those of the run, reduced to three inventories, three valve lags and three unmeasured inflows, as in Equation (6.14).

$$\begin{aligned} \frac{\mathrm{d}V_w}{\mathrm{d}t} &= \bar{Q}_{w,in} - Q_w + d_w, \qquad \frac{\mathrm{d}d}{\mathrm{d}t} = 0 \\ \frac{\mathrm{d}V_l}{\mathrm{d}t} &= \bar{Q}_{o,in} + \bar{Q}_{w,in} - Q_o - Q_w + d_l \\ \frac{\mathrm{d}m_g}{\mathrm{d}t} &= \bar{w}_{g,in} - w_g + d_g \end{aligned} \tag{6.14}$$

Here V_l is the liquid volume, the heights, pressure and flows follow Equations (6.1), (6.2) and (6.3) without dispersed volumes and with $\phi = 1$, and d collects what the model cannot explain. The inflows $\bar{Q}_{in}$ are the design wellstream, or with the optional preview the feedforward's filtered meter signals plus the recycle, held over the horizon. At each decision the model is linearized by central differences about the estimate and the applied openings and discretized with an exact hold, so an edited plant parameter is model mismatch, as on a real vessel. Against ±5-point valve steps in manual from the design point, 120 s open loop, it reproduces the plant's −16.2 mm of level, −9.7 mm of interface and −398.4 mbar of pressure within 0.4 mbar and 0.1 mm, because it is the simulator's own balance.

Each plan minimizes Equation (6.15) over the changes of the three openings, taken as offsets from the openings applied now.

$$\begin{aligned} \mathscr{J} = & \sum_{k=1}^{N_p} \left[q_h (e_{int,k}^2 + e_{tot,k}^2) + q_P e_{P,k}^2 \right] + \sum_j \left(a_j \varepsilon_j^2 + b_j \varepsilon_j \right) \\ & + \sum_k \left[\lambda \left(s_o^2 \Delta u_{o,k}^2 + s_w^2 \Delta u_{w,k}^2 \right) + \lambda_G s_g^2 \Delta u_{g,k}^2 \right] \\ & + 10\lambda \left[(Q_{w,N_p} - \hat{Q}_w)^2 + (Q_{l,N_p} - \hat{Q}_l)^2 \right] \end{aligned} \tag{6.15}$$

The errors e are the predicted heights in millimeters and pressure in millibar less their setpoints, with $q_h = 10^{-3}$ and $q_P = 0.01$. Each move is priced by the outflow it changes, s_o, s_w and s_g being the flow one percent of opening passes at the present state (2.98 m³/h of oil and 1.34 m³/h of water at design), with $\lambda = 10$ per (m³/h)² and $\lambda_G = 100$ per (kSm³/h)². The terminal term asks the water and total liquid outflows Q_w and Q_l to pass the expected inflow $\hat{Q} = \bar{Q}_{in} + \hat{d}$ at the horizon's end, so that no plan parks an outflow for the next one to undo. The openings are bounded hard by the PI loops' 0 to 100 %, and a loop in manual pins its valve. Every other limit is soft, with one slack ε_j per output priced at a_j and b_j, the level within 500 mm and the interface within 200 mm of their setpoints at 100λ per mm² and 1000λ per mm, the pressure within 300 mbar at 100 and 1000, and the oil and water outflows under the second stage's Q_{exp} and the nQ_{hc} of the banks in service at 10λ and 100λ per m³/h, judged every third step. The level and interface bands are funnels, the upper bound at step k being $y_{max} + \max(0, y_0 - y_{max})\, \mathrm{e}^{-kT_s/\tau_f}$ with $\tau_f = 120$ s and the lower bound its mirror. As a design note, a slack that covers the whole horizon would charge a violation already present, which no move can change, and leave later violations up to that size free, while the funnel asks for the return at the rate of τ_f. Each decision is one receding-horizon quadratic program [21] of 36 moves and five slacks, solved with Goldfarb and Idnani's dual active-set method [22].

A steady-state Kalman filter for the model at the declared design point corrects its nonlinear prediction every 0.5 s scan, under either structure, with the interface transmitter in use (3 mm for LT-101, 6 mm for LT-102), LT-100 at 3 mm and PT-100 at 10 mbar. The unmeasured inflows follow a random walk of variance $\gamma(2 \times 10^{-4}\ \mathrm{m^3/s})^2$ per scan with $\gamma = 0.01$, which removes offset at steady state [23], as the offset-free controller of Volume 9, Chapter 4 does with a fuel bias. The values of λ and γ and the preview off came from one pass on case 9 that ranked by oil outflow standard deviation the settings holding both bands with no more liquid valve travel than averaging PI.

The PI loops track the applied openings every scan with no preload. A plan the solver did not take to its optimum, or one with an opening outside the valve's travel or not finite, is discarded. The loops then take the valves, the historian is marked MPC FALLBACK, and the next valid plan takes them back. With the noise off, two forced rejections in case 9 at 600 s move a valve at most 0.195 % per scan, against 0.019 % before. Switching the other way, from averaging PI at 600 s with 86.1 m³/h of water against one bank's 80 m³/h, the first plan closes the water valve 9.4 points at once, from 63.9 to 54.6 %, as the soft ceiling acts. In Chrome a plan takes a mean of 1.92 to 1.95 ms in case 9, 0.1 % of one core at real time.

6.5.6 Simulation, Control and Internal Models

Three models should be kept apart. Section 6.3 is the simulation model. The identified control model is that of Section 6.4, two integrators of Equation (6.11) and a first-order pressure response of −0.091 bar per percent with a time constant near 22 s, measured in manual. Under the PI loops the controller's internal model is only the valve equation inside the feedforward, with data-sheet capacities and measured pressures, and the supervisor's estimate of the mist extractor load from the metered gas, the measured pressure and the measured temperature. Neither contains droplets, cells or the hydrocyclone. Under the MPC the model of Equation (6.14) replaces the feedforward and contains none either.

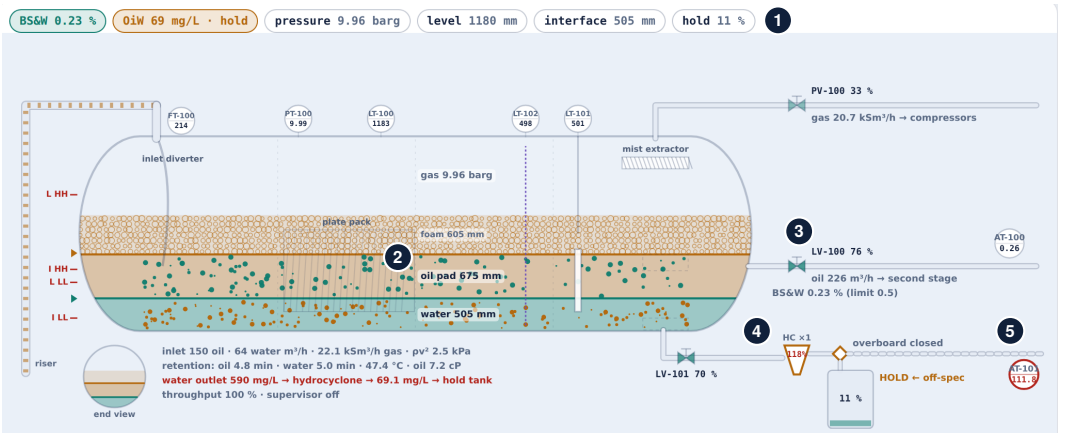

Figure 6.2 Separator application in case 2, liquid slug, at $t = 7.00$ min (420 s) with the shipped tight level tuning, 240 s after the 15 m^3 slug finished arriving, cropped to the plant view. The riser is empty. (1) The badges: the level has fallen back to 1180 mm, the interface is at 505 mm and the hold tank is 11 % full. (2) The vessel, with a foam band 605 mm deep. (3) The oil valve LV-100 at 76 % passes 226 m^3/h of oil at 0.23 % BS&W to the second stage. (4) The water valve LV-101 is at 70 %, and the single hydrocyclone bank runs at 118 %. (5) The routing valve has closed the overboard line and sends the water to the hold tank, while AT-101 reads 111.8 mg/L in alarm on water sampled four minutes earlier and the polished water now leaving the hydrocyclone carries 69.1 mg/L.

Table 6.5 Principal Measured Outcomes, Default Seed

Case	Measured outcome
1, design	retention 7.5 and 7.3 min, 0.104 %, 12.4 mg/L
1, interface 700, 350 mm	0.151 and 0.085 %, 7.5 and 18.0 mg/L
2, tight	317 m^3/h, 268 s over 250 m^3/h, 8.4 m^3 off spec
2, averaging	219 m^3/h, 6.9 m^3 off spec, 1772 L/MSm^3
2, 40 m^3 slug	level trip at 181.0 and 179.0 s
3, one and two banks	134 %, 7.7 m^3 off spec, or worst 14.4 mg/L
4	K 0.113 m/s, 1150 L/MSm^3, antifoam 185 L/MSm^3
4, supervisor	cuts to 90, 80, 70 %, clear from 1225 s
5, alone, supervisor	53.6 m^3 of 75.0 m^3 off spec, or none
6	interface trip at 1404 s
6, supervisor	cuts from 420.5 s to 50 %, 15.7 m^3 deferred
7	11.52 barg with the valve open, undershoot 8.58 barg
8	6.8 m^3 off spec, 6.1 m^3 in spec diverted

6.6 Measured Cases

Table 6.6 sets out the sequence of nine cases and the campaign. Each case starts from the steady state of Section 6.3 with tight level tuning, no feedforward and the supervisor off, as a cautious operator would commission the plant, and differs only as stated, with each changed setting in force from $t = 0$. Every reported number was measured on the engine without its user interface, at the default seed, with plant events resolved to the 0.25 s step and controller actions to the 0.5 s scan, and Table 6.5 gathers the principal outcomes.

6.6.1 The Design Point and the Interface

In case 1 the oil spends 7.5 min in its pad and the water 7.3 min in its layer, and the dispersed water falls from 1.00 % of the pad in the inlet cell to 0.32, 0.18 and 0.104 % along the vessel, because only the 540 µm class outruns the 1.86 mm/s overflow rate without coalescing first. The interface setpoint

Table 6.6 The Nine Cases and the Campaign

#	Scenario and events	Teaching point
1	Balanced feed	retention and the interface
2	15 m^3 slug launched at 60 s	the vessel stores it
3	Water cut 30 to 50 %, 60–300 s	polishing runs out
4	Gas-oil ratio 150 to 330, 60–180 s	foam and the mesh
5	Emulsion ×2.5 at 60 s	a retention time of delay
6	Water valve sticks, water cut 45 %	a trip or a cut
7	Gas header 8.0 to 11.3 barg, 60–240 s	no valve authority
8	LT-101 reads 150 mm high at 60 s	the wrong interface
9	Case 2's slug, wells to 250 m^3/h at 1200 s	an MPC prices outflow
10	2400 s wellstream campaign	graded challenge

trades one outlet against the other at an unchanged throughput. At steady state, 700 mm gives the water 11.8 min and the oil 5.5 min, which lowers the oil overboard to 7.5 mg/L and raises the BS&W to 0.151 %, while 350 mm gives 0.085 % and 18.0 mg/L. Both qualities stay within their limits over that range, so the choice spends margin on one outlet or the other.

6.6.2 A Riser Slug, Stored Either Way

Case 2 launches 15 m^3 at the riser base at 60 s, and it arrives between 120 and 180 s at up to 1414 m^3/h, with its gas held back. The nozzle momentum rises from 2.5 to 15.0 kPa and the water droplet median at the inlet falls from 338 to 198 µm. Neither tuning keeps the slug out of the vessel, since both valves fully open pass only 433 m^3/h at design pressure. When the level peaks, at 177 and 179 s, 14.9 and 15.0 m^3 of the slug have arrived, and the vessel holds 12.8 m^3 of it under the shipped tight tuning and 14.4 m^3 under averaging, at peaks of 1629.6 and 1682.1 mm. The tight loop has its oil valve command at the limit from 138.5 s and its water valve command from 162.0 s, so the 2.1 m^3 it passes on during the arrival is set by the valves rather than by the gain. Across a fifty-fold range of level gain at a 1500 s integral time, from 0.1 to 5 %/cm, the peak level moves only from 1692 to 1640 mm, while the peak oil outflow moves from 175 to 318 m^3/h. The plant at 420 s under tight tuning, with the slug stored and on its way out, is shown in Figure 6.2.

> **TRY IT IN THE APP** **2 · Liquid slug**
>
> Load the case, press Run and watch the level in the vessel. The slug pours in faster than both open valves can pass, so the level climbs to 1629.6 mm at 177 s and the vessel holds 12.8 m^3 of the 15 m^3 slug before the valves draw it back down.

The tuning decides what happens after the arrival, as shown in Figure 6.3. The tight loop returns the stored liquid as fast as the valves pass it. With the gas blowdown lifting the pressure to 10.92 barg at 224 s, the wide-open oil valve delivers 317 m^3/h, more than its 298 m^3/h at design pressure, and the oil outflow exceeds the 250 m^3/h second stage for 268 s. The level first returns within 50 mm of its setpoint at 372 s, undershoots to 1142 mm and stays within the band from 586 s. Averaging spreads the return over a quarter of an hour, with the oil outflow never above 219 m^3/h and the level back within 50 mm at 947 s. The largest rate of change of the oil outflow over 5 s is 15.1 m^3/h per second with tight tuning and 2.09 with averaging, the quantity that optimal averaging control minimizes [4]. At a 1500 s integral time the second stage is respected up to 0.6 %/cm, which peaks at 246 m^3/h, while 0.8 %/cm reaches 272 m^3/h. The pressure transient comes from the gas rather than the levels, with a dip to 9.69 barg at 154 s while the gas is held back, the rise during the blowdown and an undershoot to 9.26 barg at 309 s, and it agrees under both tunings to within 0.09 bar.

Pause. Try Exercise 6.1 before reading on.

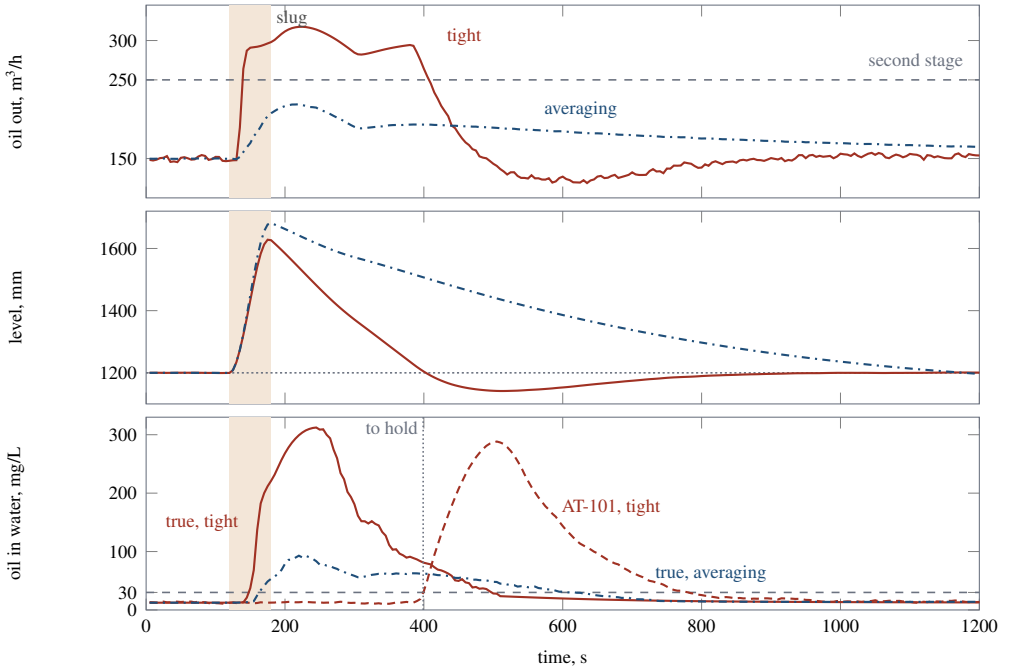

Figure 6.3 The 15 m^3 slug of case 2 under tight (solid) and averaging (dash-dotted) level tuning, from headless runs at the default seed. The shaded band is the arrival of the slug. The oil outflow is shown against the 250 m^3/h second stage (thin dashed), the total level against its 1200 mm setpoint, and the oil in the polished water against the 30 mg/L limit (thin dashed), with the AT-101 reading for the tight case (densely dashed). The vertical dotted line marks the diversion to the hold tank at 399 s.

The price of averaging is freeboard. With the level higher during the blowdown the foam top comes within 21 mm of the mist extractor instead of 193 mm, and the carryover peaks at 1772 L/MSm^3 against 335 L/MSm^3.

6.6.3 Where the Slug's Water Goes

The produced water fares badly under either tuning. The tight interface loop drives the water outflow to 144 m^3/h and the single hydrocyclone bank to 180 % of its capacity, and the polished water peaks at 312 mg/L at 246 s. Averaging lowers the peak load to 134 % but keeps the bank overloaded for 585 s instead of 366 s, so the off-spec water discharged falls only from 8.4 to 6.9 m^3, as listed in Table 6.7. The analyzer sees the excursion late. The true value passes 30 mg/L at 148 s, AT-101 diverts the water to the hold tank at 399 s, and by then the whole 8.4 m^3 has gone overboard. Shortening the analyzer's delay from 240 to 30 s cuts that volume to 1.39 m^3 without removing it. A diversion on a late analyzer always lets through what passed during the delay, the same problem that the continuous drug line of Volume 5, Chapter 7 meets when it routes product on a result that describes material already downstream. The second hydrocyclone bank removes the loss under both tunings, with the tight case peaking at 29.97 mg/L, just under the limit. The supervisor brings that bank in on the measured water flow, at 151 s under tight tuning and 161 s under averaging, and leaves 0.07 m^3 and none off spec. The interface loop can protect the water itself. Averaged to 0.5 %/cm and 1500 s, it holds the water outflow to 82 m^3/h and the single bank to 103 %, the polished water peaks at 25.0 mg/L and nothing goes off spec, while the interface rises only to 680 mm against 673 mm with the preset. Under tight level tuning the same interface loop still protects the water, but the level loop then passes the held-back volume as oil, and the second stage stays overloaded for 351 s instead of 268 s. Feedforward undoes

Table 6.7 Case 2 Variants Over 1200 s. Off-Spec Is Water Discharged Above 30 mg/L. Interface Rows Use $\tau_I = 1500$ s

Configuration	Oil peak m^3/h	Over 250 s	Off spec m^3	Carryover L/MSm^3
tight (shipped)	317.5	268	8.40	335
averaging	219.1	0	6.92	1772
tight, two banks	317.5	268	0	335
averaging, two banks	219.1	0	0	1772
averaging, feedforward 900 s	246.8	0	7.51	2089
averaging, feedforward 300 s	304.6	178	8.23	1203
averaging, supervisor	219.1	0	0	273
averaging, interface 0.5 %/cm	220.8	0	0	2078
tight, interface 0.5 %/cm	317.6	351	0	742
tight, AT-101 delay 30 s	317.5	268	1.39	335
tight, AT-101 delay 120 s	317.5	268	4.84	335
tight, AT-101 delay 480 s	317.5	268	10.77	335

part of the averaging, up to 305 m^3/h with a 300 s filter (Table 6.7).

Pause. Try Exercise 6.3 before reading on.

The slug also shows where tuning stops mattering. Launched from the design state in model runs with slug sizes the application does not offer, the largest slug that arrives over 60 s without a high-high trip is 29.0 m^3 under tight tuning and 27.7 m^3 under averaging, against the 26.4 m^3 surge volume, and the 40 m^3 slug of the event button, fired at 60 s in case 1, trips the level at 181.0 and 179.0 s. Tight tuning buys 1.35 m^3 of slug tolerance and pays for it at the second stage and the hydrocyclone.

DID YOU KNOW? A Faster Analyzer Only Shrinks the Loss

From the simulator (Table 6.7). Under tight tuning in case 2 the off-spec water discharged before AT-101 diverts is 1.39 m^3 with a 30 s delay, 4.84 with 120 s, 8.40 with the shipped 240 s and 10.77 m^3 with 480 s. No delay short of zero keeps it all on board.

6.6.4 A Rising Water Cut and the Hydrocyclone

Case 3 raises the water cut from 30 to 50 % in five steps between 60 and 300 s at an unchanged 214 m^3/h of liquid. The interface loop holds the interface within 52 mm while the water outflow rises from 64 to 107 m^3/h and the water valve settles near 80 %. The separator copes, and its water outlet improves from 353 to 222 mg/L because less oil enters to be dispersed, even though the water retention falls from 7.3 to 4.3 min. The single hydrocyclone bank does not. It runs at 126 % of its capacity by 400 s and 134 % at 1800 s, the polished water passes 30 mg/L at 377 s and peaks at 55.8 mg/L, and AT-101 diverts only at 637 s, after 7.7 m^3 of off-spec water has gone overboard. The hold tank then fills from 10 to 46 % by 1800 s. With both banks in service the worst is 14.4 mg/L, and the supervisor brings the second bank in at 171 s, once the measured water flow has stayed above 85 % of one bank for 20 s, holding the worst at 12.6 mg/L.

6.6.5 More Gas, Foam and the Mist Extractor

Case 4 raises the gas-oil ratio from 150 to 330 in three steps between 60 and 180 s, 2.20 times the design gas. The pressure loop copes, opening the gas valve from 35.3 to 77.6 % after a peak of 10.55 barg and returning within 0.05 bar of setpoint by 345 s. The mist extractor does not. Its capacity factor reaches 0.113 m/s against 0.10 m/s, the foam grows to 1206 mm with its top 92 mm below the mesh, and the carryover passes 13 L/MSm^3 at 129 s and peaks at 1150 L/MSm^3. Antifoam at 8 ppm collapses the foam to 548 mm but leaves 185 L/MSm^3, the overloaded mesh alone. A throughput

request of 75 % with the antifoam lowers the capacity factor to 0.085 m/s and the carryover to 11.6 L/MSm^3 by 1800 s. The supervisor reaches the same answer by itself, raising the antifoam to 6.0 ppm, which it relaxes to 4.7 ppm once the foam has fallen, and cutting the throughput to 90, 80 and 70 % at 350, 710 and 1090 s, and the carryover stays under its limit from 1225 s.

6.6.6 A Tighter Emulsion and a Retention Time of Delay

Case 5 makes the emulsion 2.5 times tighter at 60 s, and the water median at the inlet falls from 338 to 135 µm at once. The outlet does not follow at once. Ten seconds later the BS&W has moved from 0.104 to 0.1065 %, and it takes 340 s to cover 63 % of its rise against a 448 s oil retention time, crosses 0.5 % at 513 s and settles at 0.59 %, so that 53.6 m^3 of the 75.0 m^3 exported by 1800 s is off spec. With 40 ppm of demulsifier from the start the BS&W never passes 0.328 %. The supervisor's trim reaches 40 ppm at 571 s and then raises the heater to 3 MW by 1031 s, warming the liquid to 58.3 °C, so the BS&W peaks at 0.445 % at 516 s and falls back to 0.27 % with nothing off spec. The trim starts only at 240.5 s, after AT-100 has passed its 0.2 % target at 236.5 s, so the rise already held in the pad reaches the outlet before the extra demulsifier and heat can act on it.

6.6.7 Limits No Tuning Moves

In case 6 the water valve stem seizes at 47.6 % at 60 s, and at 120 s the water cut rises to 45 %, so 96 m^3/h of water arrive and 64 m^3/h leave. The not-following alarm is raised at 206 s, and the interface loop's command reaches 100 % at 377 s without effect. The level loop holds the total level by the only route the triangular structure leaves it, taking out 150 m^3/h of oil while 118 m^3/h arrives, so by 1200 s the interface is at 884 mm, the pad has thinned from 700 to 316 mm and the BS&W has risen to 0.314 %. At 1404 s the interface reaches its 950 mm trip and the wells are shut in. The stem freezes wherever valve noise has left it, and over eight noise seeds the trip falls between 1350 and 1448 s. With the supervisor on, the not-following alarm and the rising interface cut the throughput from 420.5 s in 10 % steps to 50 %, the interface peaks at 720 mm at 941 s, nothing trips under any of the eight seeds, and 15.7 m^3 of oil is deferred by 1800 s.

🕮 **Pause.** Try Exercise 6.2 before reading on.

In case 7 the gas header rises from 8.0 to 11.3 barg in four steps between 60 and 240 s. The separator pressure passes 11.0 barg at 187 s, the gas valve command reaches its limit at 231 s, and the pressure settles at 11.52 barg, 0.22 bar above the header and 0.48 bar below the trip. Back-calculation then holds the integral term at 100.0 %, as the algorithm predicts, while the proportional term adds 23.0 %, so the unclamped demand of 123.0 % is the proportional term and nothing more. The higher pressure pushes more liquid out, and the level loops close the oil valve from a mean of 50.0 % over the first minute to 45.3 % over 400 to 600 s and the water valve from 48.0 to 42.1 %, while the denser gas lowers the foam from 543 to 478 mm. A throughput cut would not help, since 80 and 50 % still give 11.44 and 11.36 barg, so the supervisor does not cut. When the header falls back to 8.0 barg at 600 s, the gas valve is still wide open with its integral at the limit, the cap drains and the pressure undershoots to 8.58 barg at 622 s, returning within 0.05 bar only at 807 s. Raised to 11.8 barg at 1200 s instead, the header trips the separator on high pressure 15.5 s later. No tuning creates a differential that the header has removed.

6.6.8 A Biased Interface Transmitter

In case 8 LT-101 starts reading 150 mm high at 60 s, and the interface loop obeys it. The loop drains the water layer at up to 134 m^3/h, the hydrocyclone runs at 168 %, and the polished water peaks at 170.5 mg/L at 111 s. The true value is back under 30 mg/L at 280 s, but AT-101 diverts only at 308.5 s, after 6.8 m^3 of off-spec water has gone overboard, and the diversion then sends 6.1 m^3 of water already in spec to the hold tank. The disagreement between LT-101 and the LT-102 profiler raises its alarm at 119.5 s. By 1800 s the true interface sits at 350 mm, the water keeps 4.3 min instead of 7.3 min, and the separator outlet carries 470 mg/L instead of 353 mg/L, still 18.0 mg/L after the

Table 6.8 Loops and MPC on the Slug and the Inflow Step, Default Seed. IAE and Travel Are the Scorecard's. Interface Averaged Sets the Interface Loop to 0.5 %/cm and 1500 s

Structure	Oil out, m^3/h peak	SD	Water peak m^3/h	Off spec m^3	Level max mm	Level IAE mm·s	Pressure IAE bar·s	Carry-over L/MSm^3	Travel, % oil+ water	gas
Case 9, 2400 s										
tight PI (shipped)	317.5	45.73	144.3	8.40	1630	76 951	138	335	1256	186
averaging PI	219.1	12.25	107.4	6.92	1682	252 904	135	1772	310	184
averaging, interface averaged	220.8	13.24	82.2	0.00	1687	268 133	135	2078	152	184
MPC (as loaded)	202.8	9.38	83.0	0.00	1697	414 246	280	2098	57	151
MPC with preview	238.1	20.04	84.3	0.00	1688	193 206	133	2149	105	459
Case 2, 1200 s										
tight PI (shipped)	317.5	64.35	144.3	8.40	1630	73 102	124	335	604	123
averaging PI	219.1	16.52	107.4	6.92	1682	192 038	120	1772	174	120
averaging, interface averaged	220.8	17.02	82.2	0.00	1687	233 574	120	2078	87	120
MPC	202.8	13.03	83.0	0.00	1697	283 112	124	2098	48	125
Inlet step, 2400 s										
tight PI (shipped)	187.4	6.05	79.4	0.00	1216	5288	17	5	1275	123
averaging PI	179.9	8.45	78.5	0.00	1327	240 338	17	5	269	123
MPC	175.5	7.06	77.9	0.00	1360	305 920	268	5	26	50

MPC runs: 1200 plans in 2400 s and 600 in 1200 s, none rejected.

hydrocyclone, while the thicker pad lowers the BS&W to 0.085 %. With LT-102 as the measurement the loop does not move and nothing goes off spec.

> **TRY IT IN THE APP 8 · Biased interface sensor**
>
> Press Run with Interface measurement on LT-101 displacer. The disagreement alarm comes at 119.5 s, and AT-101 diverts at 308.5 s, after 6.8 m^3 of off-spec water has gone overboard. Set Interface measurement to LT-102 profiler, press Reset and run again: the loop does not move and nothing goes off spec.

6.6.9 Averaging by a Predictive Controller

Case 9 repeats the 15 m^3 slug of case 2 at 60 s, and at 1200 s the wells' liquid rate steps from 214 to 250 m^3/h and stays there. It loads under the MPC, and its comparison is the PI loops with the averaging preset. Table 6.8 measures case 9, case 2 and an inlet step, case 1 with the same rise at 60 s, with the same seed, instruments and noise and the MPC always at the settings case 9 ships, and Figure 6.4 shows case 9. No run in the table trips, and none of its 4200 plans was rejected.

The MPC spends the band on the oil. In case 9 it lets the level climb to 1697 mm, against 1682 mm under averaging, and raises the oil only to 202.8 m^3/h, against 219.1 m^3/h, with a standard deviation of 9.38 against 12.25 m^3/h and 57 % of liquid valve travel against 310 %. The soft ceiling of the one bank holds the water to 83.0 m^3/h, so none goes off spec, against 6.92 m^3 under the preset. That is not the MPC's alone. With the interface loop averaged as well, at 0.5 %/cm and 1500 s as in Table 6.7, the PI loops hold the water to 82.2 m^3/h with 0 m^3 off spec, and the MPC's lead narrows to the oil, a peak of 202.8 against 220.8 m^3/h and a standard deviation of 9.38 against 13.24 m^3/h, and to liquid valve travel, 57 against 152 %.

The MPC loses on the level and the pressure. Its level IAE is 414 246 against 252 904 mm·s under the preset and 268 133 with the interface averaged, and at 1200 s the level is at 1288 mm, where averaging has been within 50 mm since 947 s. The higher level raises the worst carryover to 2098 L/MSm^3 against 1772 L/MSm^3, and the pressure IAE is 280 against 135 bar·s, although its range, 9.75 to 10.80 barg, is narrower than the loops' 9.26 to 10.92 barg. After the step the level peaks at 1316 mm at 2121 s, against 1263 mm under averaging at 2400 s, and run on to 21 600 s it is back within 50 mm from 6289 s and within 10 mm from 13 334 s, with no offset. The estimate of the unseen inflow is slow on purpose, so that a slug does not look like a lasting flood. It still credited

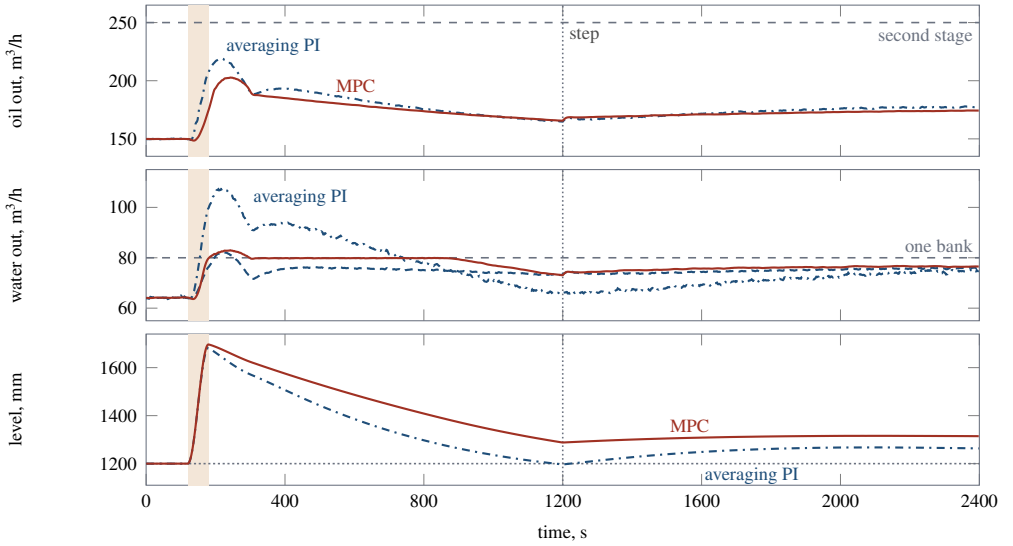

Figure 6.4 Case 9 over the 40 min window as at $t = 2400$ s, default seed, under the MPC as loaded (solid) and averaging PI (dash-dotted), with the slug shaded and the step of the wells dotted. Oil peaks at 202.8 m³/h at 243 s and 219.1 m³/h at 219 s, water at 83.0 and 107.4 m³/h (82.2 m³/h with the interface loop averaged, densely dashed), and the level ends at 1314 and 1263 mm. Thin dashed lines mark the second stage and one hydrocyclone bank.

17.8 m³/h to the slug at 1200 s, when the true excess was zero, and only 27.5 of the step's 36 m³/h at 2400 s, so every plan expected less liquid than arrived. Told the true step, a diagnosis the application does not offer, the MPC returns within 50 mm by 1762 s. The return is a trade of the tuning, not a defect of the estimator.

The move price λ sets how the slug is shared between the vessel and the second stage. At $\lambda = 3$ the oil peaks at 242.8 m³/h and the level is back within 50 mm by 910 s. At the shipped 10 the peak is 202.8 m³/h, and at 30 it is 190.2 m³/h with the level still at 1422 mm at 2400 s. Each step of the price buys a lower oil peak with a slower return of the level.

The preview makes the slug worse, presenting it through the held meter reading as a lasting rise, so the oil peaks at 238.1 m³/h. On the inlet step, where nothing needs absorbing, the shipped tight loops give the lowest oil standard deviation, 6.05 against 7.06 m³/h for the MPC and 8.45 m³/h for averaging, since passing a step straight on is the calmest answer to it, while the MPC moves the liquid valves least, 26 against 1275 and 269 %, at a pressure IAE of 268 against 17 bar·s.

6.6.10 The Wellstream Campaign

The 2400 s challenge brings three 12 m³ slugs, launched at 180, 1500 and 2040 s, a water cut rising to 35 and 40 % at 540 and 720 s, and an emulsion tightening in four steps to 1.8 times its usual stability between 900 and 1260 s. It asks for no trip, 9.0 to 11.5 barg, at most 250 m³/h of oil, BS&W and oil in water within their limits for 85 and 95 % of the time, no hold-tank overflow and at least 90 m³ of in-spec oil. Table 6.9 compares nine configurations. In the shipped configuration the plant exports 69.8 m³ in spec of 114.2 m³, keeps the BS&W within its limit 66.4 % of the time with a worst of 0.98 %, discharges 16.9 m³ of off-spec water with a worst of 264 mg/L, and passes 318 m³/h to the second stage, 564 s above its capacity. The suggested configuration, averaging, feedforward and the supervisor, exports 98.7 m³ in spec of 112.3 m³, 8.7 m³ more than the 90 m³ clause asks, brings the second bank in at 270.5 s, reaches 40 ppm of demulsifier at 1580.5 s, holds the worst BS&W

Table 6.9 Challenge Configurations Over 2400 s. Shares Are the Time Each Quality Stayed Within Its Limit

Configuration	In spec m^3	BS&W %	OiW %	Oil peak m^3/h	Pass
as shipped	69.8	66.4	63.8	318.5	no
suggested[a]	98.7	89.2	100.0	226.0	yes
averaging, supervisor	98.5	89.0	100.0	207.6	yes
supervisor only	87.3	84.4	100.0	318.5	no
averaging only	70.5	63.8	58.0	207.6	no
averaging, feedforward	70.0	63.1	54.9	226.0	no
tight, feedforward, supervisor	87.8	84.8	99.9	312.0	no
MPC, else as shipped	71.2	65.9	86.3	193.8	no
MPC, supervisor	101.2	93.6	100.0	193.4	yes

[a] averaging, feedforward and the supervisor

to 0.613 % and the worst overboard water to 22.7 mg/L, and keeps the pressure between 9.40 and 10.74 barg.

Each element does its own job. The supervisor alone keeps the overboard water in spec and lifts the BS&W share from 66.4 to 84.4 %, just short of the 85 % target, and the oil in spec to 87.3 m^3, short of 90 m^3, but leaves 318 m^3/h at the second stage, which among the PI settings only averaging protects. Averaging alone protects the second stage and fixes neither quality. Without the feedforward the averaging and the supervisor still pass, with 98.5 m^3 in spec, and the feedforward's contribution is a narrower pressure band, 9.40 to 10.74 barg against 9.26 to 10.97 barg, bought with 723 % of gas valve travel against 303 %. Figure 6.5 sets the nine configurations of Table 6.9 against the four clauses of the target that carry a number.

The MPC does not change the campaign's verdict. With everything else as shipped the MPC fails with 71.2 m^3 in spec, since the treating chemicals and the second bank are the supervisor's job, and with the supervisor on it passes with 101.2 m^3 and the pressure between 9.58 and 10.77 barg.

6.7 Technology in Practice

Figure 6.6 draws the three loops, the feedforward, the supervisor, the routing and the optional MPC of Section 6.5 with the hardware of a platform separator.

What Measures It

The level loop reads LT-100, a radar on the vessel top. Echo instruments can be fooled by a layer of foam on the liquid and have difficulty with interfaces that are not distinct, such as emulsions, and guided-wave radar, whose probe carries the signal into the liquid, is the only practical radar for a liquid-liquid interface [24]. In the model LT-100 adds waves that grow with the momentum of the inlet stream. The qualities depend on two analyzers. AT-100 reads the water in the export oil through a 20 s lag, and AT-101 reads the oil in the polished water through a 240 s sample delay. An online oil-in-water monitor is usually mounted on a side stream, and the length of its sample line adds a delay of its own [25].

> **TECHNOLOGY** **Sensor: The Interface Displacer**
>
> LT-101 is a displacer: it weighs an element too heavy to float that hangs in the liquid and reads the buoyant force on it as level. That force is proportional to the density of the liquid around the element, so a change of density reads as a change of level. For an interface the element must stay fully submerged, or the instrument cannot tell a low interface from a low total level [24]. An element in a cage beside the vessel can be isolated for service, one hung inside cannot [24], and an emulsion layer disturbs the reading [26]. The chapter gives LT-101 a 3 s lag, 3 mm of noise and, in case 8, a bias.

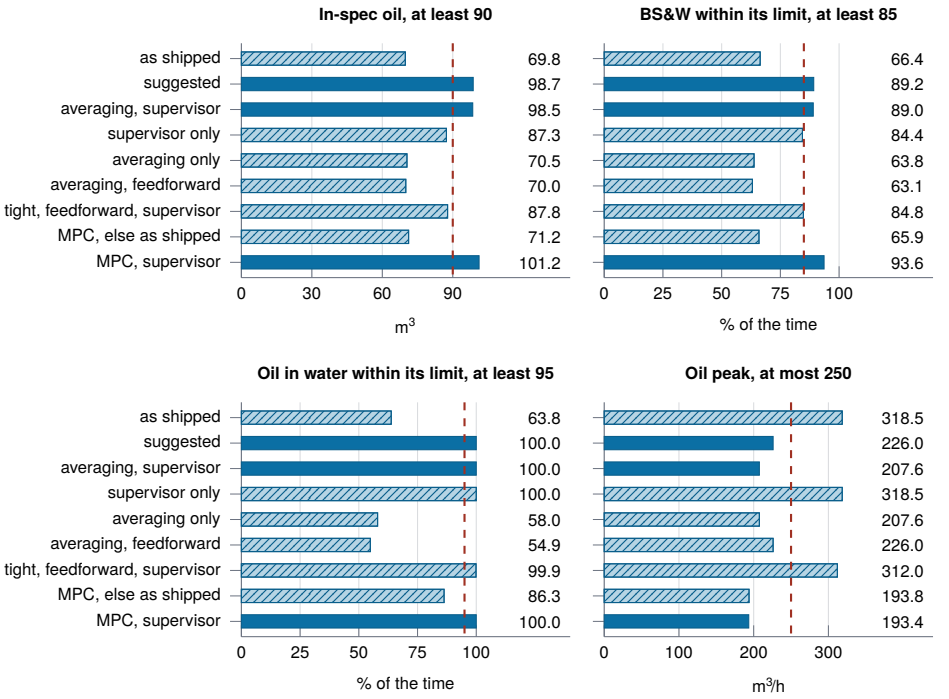

Figure 6.5 The challenge runs of Table 6.9 as bars: in-spec oil, the shares of time the BS&W and the oil in water stayed within their limits, and the oil peak, over 2400 s. Solid bars are the configurations that pass, hatched bars those that fail, dashed red lines are the clauses of the target, and the value of each bar stands at the right. Notice that every passing configuration pairs the supervisor with averaging or with the MPC, and that the supervisor alone misses the oil and BS&W clauses narrowly and the second stage widely.

What Moves It

Three control valves set the outlets: PV-100 on the gas, LV-100 on the oil and LV-101 on the water to the hydrocyclone. Fully open at design the oil valve passes 298.4 m^3/h and the water valve 134.3 m^3/h, while a 15 m^3 slug peaks at 1414 m^3/h. The dosing pumps, the heater, the second hydrocyclone bank, the recycle pump and the routing valve belong to the supervisor and the routing rule.

> **TECHNOLOGY** **Final Element: The Water Valve and Its Positioner**
>
> Packing friction and the force of the process on the plug keep a valve stem from following its actuator exactly, so a positioner compares the stem position with the command and drives the actuator until they agree [24]. With too much friction the stem moves in steps, called stiction, and ignores small reversals of the signal, worst in poorly maintained valves and in valves without a positioner [24]. Electronic positioners report position errors and total stem travel to maintenance and to the control system [24]. The chapter's valves follow their commands through a 3 s lag with no stiction band. Its not-following alarm makes the positioner's comparison of position and command: in case 6 the water valve seizes at 60 s and the alarm comes at 206 s.

Where the Controller Runs

The loops would run in the platform's control system. A distributed control system controls production within one geographic location through feedback and feedforward loops under a supervisory level, and its loops run with cycle times from milliseconds to minutes [27]. The app's loops execute every

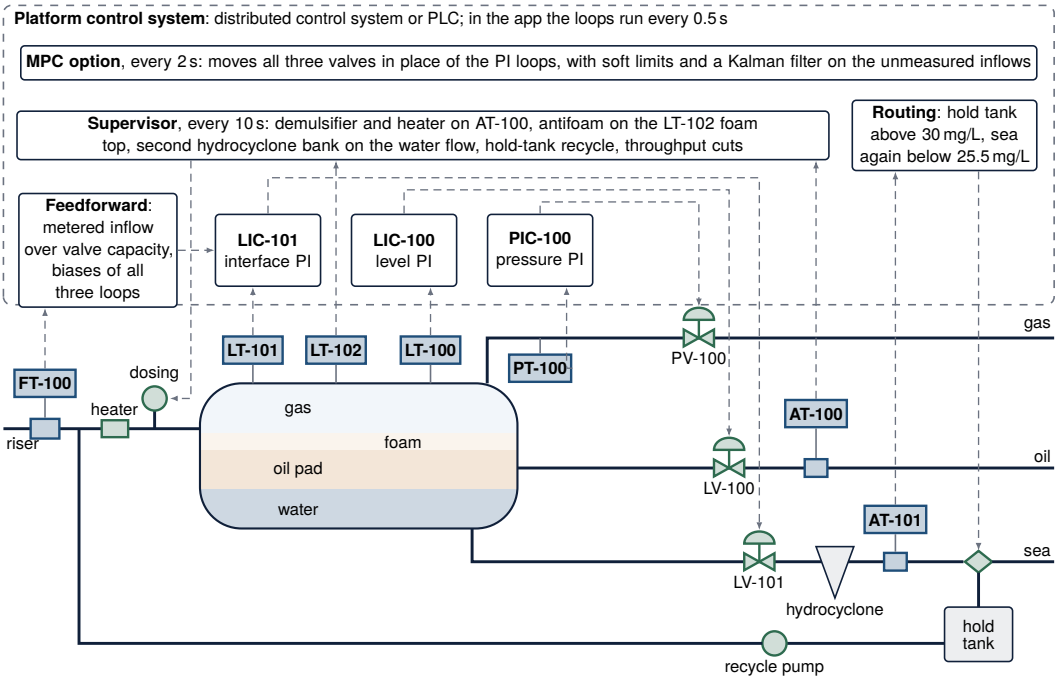

Figure 6.6 The loops of the separator simulator with their hardware. Lettered boxes are sensing devices with the chapter's tags: the level transmitters LT-100, LT-101 and LT-102 on the vessel, PT-100 on the gas line, the inlet meter FT-100 and the analyzers AT-100 on the oil and AT-101 on the polished water. Green symbols are final elements: the gas, oil and water valves, the inlet heater and chemical dosing, the routing valve to the sea or the hold tank, and the recycle pump. Gas goes to the compressor header and oil to the second stage. Dashed arrows are signals. Notice that only the two analyzers read quality, and that AT-101 commands a routing valve that the water it measured has already passed. Not to scale, and not a piping diagram.

0.5 s, its supervisor every 10 s and its MPC every 2 s, all inside that range and all far shorter than the 240 s by which AT-101 reports, so no scan rate makes the diversion timely. The chapter's structure maps onto three PI blocks with a feedforward bias computed from the inlet meter, a rule set in the logic layer for the chemicals, the heater, the second bank, the recycle and the throughput, a routing rule on AT-101 with a timed return, and, as an option, the MPC in a supervisory layer that moves the three valves in place of the PI blocks.

> **IN PRACTICE** **Compliance Is Judged in the Laboratory**
>
> In United States offshore waters the federal rule for produced water limits oil and grease to 42 mg/L on any one day and to 29 mg/L as the average of the daily values over 30 consecutive days [28]. On North Sea installations the oil in water is measured offline by the reference method of ISO 9377-2. The method requires at least two samples a day, samples are taken about three times a day [25]. The chapter's scorecard works differently from both: it judges the true oil in water at every instant against 30 mg/L.

What Is Modernizing It

Three developments touch this chapter. Electronic positioners count total stem travel and can predict when the stem packing will need replacing, flagging a maintenance alarm [24]. Online oil-in-water monitors, most of them fluorescence instruments, have served for monitoring, and a 2017 pilot plant

study used them to follow a hydrocyclone's efficiency as its inflow stepped [25], the dynamic reading that AT-101 lacks. A control-oriented separator model of 2018 adds virtual monitoring of the gas and liquid inflows and of the split of oil and water entering the water layer, quantities whose measurement is expensive, unreliable or impossible [13]. The Kalman filter of Section 6.5.5, which estimates the unmeasured inflows, is an estimator of the same kind.

6.8 Limitations

The wells deliver an imposed stream, so the application cannot represent the coupling between the riser and a topside choke controller on which active slug suppression relies [8].

The vessel is four well-mixed cells under a flat level, and its quality correlations are illustrative. It reproduces the direction and the time scale of each quality effect, but it cannot predict the BS&W of a particular crude or the efficiency of a particular liner [7, 18].

> ⚠ **OPEN CHALLENGE** **Seeing Every Layer in the Vessel**
>
> A crude oil separator holds gas, oil, water, emulsion and sludge, and accurate measurement of its several interfaces is regarded as a challenge. Displacers, differential pressure transmitters and ultrasonic instruments cannot measure the emulsion or the foam, and gamma instruments, which do see the vertical distribution of the phases, carry hazardous radiation. Separator control stays simple for lack of such measurements, and a suitable instrument would give more efficient and compact separators and less chemical use [26]. Progress would be a measurement of the emulsion band and the foam reliable enough for a loop to act on.

There is no flare or relief path for the gas, so a header above the trip pressure ends the run, and the supervisor is a rule set rather than an optimizer or a vendor's logic.

The MPC's model leaves out droplets, foam and carryover, thin-layer draws, transmitter lags and waves, the wells and the supervisor, and a stuck stem. It holds the inflow over the horizon, so it cannot tell a slug from a lasting rise until its slow estimate moves, and its close validation reflects a model built from the simulator's own balance rather than one identified on a real vessel. The pressure keeps a fixed band with one slack for the whole horizon, so a predicted violation leaves any smaller later one unpriced.

> ⚠ **OPEN CHALLENGE** **An Oil-in-Water Reading Fit for Feedback**
>
> Online oil-in-water monitors, most of them fluorescence instruments, serve only for monitoring. Their reading is only approximately linear and is calibrated to one range, one oil may need a different calibration from a slightly different oil, and a fouled view cell drifts. Their precision under changing conditions still needs verification, and none is installed as a feedback transmitter on North Sea installations [25]. Progress would be a reading trusted enough to adjust the water treatment before the discharge goes off spec, not after.

6.9 Conclusion

An offshore three-phase separator is an inventory whose levels are held for the sake of its outflows, and the simulator described here measures that distinction on one vessel. The oil and water valves act on the two levels through a triangular gain matrix, -0.0281 and -0.0126 millimeters per second per percent on the total level against -0.0167 on the interface, so the pairing is forced and a water valve that cannot act makes the level loop drain oil. A 15 m^3 slug is stored under any tuning, 12.8 m^3 of it at the peak under tight tuning and 14.4 m^3 under averaging, and the tuning decides only how fast it returns, at 317 m^3/h and 268 s over the second stage or at 219 m^3/h. Under both presets an overloaded hydrocyclone bank discharges 8.4 or 6.9 m^3 of off-spec water before the analyzer, four minutes late, diverts it, and a second bank or an interface loop averaged to 0.5 %/cm removes the loss. The campaign moves from 69.8 m^3 to 98.7 m^3 of in-spec oil with averaging, feedforward and the supervisor. An optional MPC lowers the slug's oil peak to 202.8 m^3/h against 219.1 and 220.8 m^3/h for the averaging loops, at a level IAE of 414 246 against 252 904 mm·s. The model imposes the

wellstream, mixes each cell and uses illustrative separation correlations. A fuller model would couple the riser to the separator pressure, add a relief path, replace the cells by a settling-zone model and average the discharge limit over time. The case that runs through the chapter is that a level controller chooses where a disturbance goes, not whether it arrives.

KEY POINTS

- A level controller chooses where a slug goes, not whether it arrives: the 15 m^3 slug of case 2 is stored in the vessel under any tuning.
- The tuning decides only how fast the stored liquid returns, at 317 m^3/h and 268 s over the second stage under tight tuning or at 219 m^3/h under averaging.
- The gain matrix is triangular, so the pairing is forced, and a water valve that cannot act makes the level loop drain oil.
- An overloaded hydrocyclone discharges 8.4 or 6.9 m^3 of off-spec water before the analyzer, four minutes late, diverts it, and a second bank or an averaged interface loop removes the loss.

Where else this applies. Volume 9, Chapter 7 spends a reservoir's storage to hold a flood under a downstream limit, and Volume 5, Chapter 7 diverts product on an analyzer that reports after the material has passed.

References

[1] T. Frising, C. Noïk, and C. Dalmazzone, "The liquid/liquid sedimentation process: From droplet coalescence to technologically enhanced water/oil emulsion gravity separators: A review," *Journal of Dispersion Science and Technology*, vol. 27, no. 7, pp. 1035–1057, 2006.

[2] K. E. Arnold and P. J. Koszela, "Droplet-settling vs. retention-time theories for sizing oil/water separator," *SPE Production Engineering*, vol. 5, no. 1, pp. 59–64, 1990.

[3] M. Stewart, "Three-phase oil and water separators," in *Gas-Liquid and Liquid-Liquid Separators*. Elsevier, 2008, pp. 131–174.

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

[5] Z. Schmidt, J. P. Brill, and H. D. Beggs, "Experimental study of severe slugging in a two-phase-flow pipeline - riser pipe system," *Society of Petroleum Engineers Journal*, vol. 20, no. 5, pp. 407–414, 1980.

[6] Y. Taitel, "Stability of severe slugging," *International Journal of Multiphase Flow*, vol. 12, no. 2, pp. 203–217, 1986.

[7] S. Kokal, "Crude-oil emulsions: A state-of-the-art review," *SPE Production & Facilities*, vol. 20, no. 1, pp. 5–13, 2005.

[8] K. Havre and M. Dalsmo, "Active feedback control as a solution to severe slugging," *SPE Production & Facilities*, vol. 17, no. 3, pp. 138–148, 2002.

[9] J.-M. Godhavn, M. P. Fard, and P. H. Fuchs, "New slug control strategies, tuning rules and experimental results," *Journal of Process Control*, vol. 15, no. 5, pp. 547–557, 2005.

[10] H. Gul Zaman, L. Baloo, R. Pendyala, P. Singa, S. Ilyas, and S. Kutty, "Produced water treatment with conventional adsorbents and MOF as an alternative: A review," *Materials*, vol. 14, no. 24, p. 7607, 2021.

[11] D. Asante-Sackey, S. Rathilal, E. K. Tetteh, and E. K. Armah, "Membrane bioreactors for produced water treatment: A mini-review," *Membranes*, vol. 12, no. 3, p. 275, 2022.

[12] A. F. Sayda and J. H. Taylor, "Modeling and control of three-phase gravilty separators in oil production facilities," in *2007 American Control Conference*. IEEE, 2007, pp. 4847–4853.

[13] C. J. Backi, B. A. Grimes, and S. Skogestad, "A control- and estimation-oriented gravity separator model for oil and gas applications based upon first-principles," *Industrial & Engineering Chemistry Research*, vol. 57, no. 21, pp. 7201–7217, 2018.

[14] 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.

[15] M. Souders and G. G. Brown, "Design of fractionating columns I. Entrainment and capacity," *Industrial & Engineering Chemistry*, vol. 26, no. 1, pp. 98–103, 1934.

[16] J. J. Bikerman, "The unit of foaminess," *Transactions of the Faraday Society*, vol. 34, p. 634, 1938.

[17] A. Fakhru'l-Razi, A. Pendashteh, L. C. Abdullah, D. R. A. Biak, S. S. Madaeni, and Z. Z. Abidin, "Review of technologies for oil and gas produced water treatment," *Journal of Hazardous Materials*, vol. 170, no. 2–3, pp. 530–551, 2009.

[18] T. Husveg, O. Rambeau, T. Drengstig, and T. Bilstad, "Performance of a deoiling hydrocyclone during variable flow rates," *Minerals Engineering*, vol. 20, no. 4, pp. 368–379, 2007.

[19] P. J. Campo and M. Morari, "Model predictive optimal averaging level control," *AIChE Journal*, vol. 35, no. 4, pp. 579–591, 1989.

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

[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] 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] 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] 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] P. Durdevic, C. Raju, M. Bram, D. Hansen, and Z. Yang, "Dynamic oil-in-water concentration acquisition on a pilot-scaled offshore water-oil separation facility," *Sensors*, vol. 17, no. 1, p. 124, 2017.

[26] S. Bukhari and W. Yang, "Multi-interface level sensors and new development in monitoring and control of oil separators," *Sensors*, vol. 6, no. 4, pp. 380–389, 2006.

[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 435.13, effluent limitations guidelines representing the degree of effluent reduction attainable by the application of the best available technology economically achievable (BAT), offshore subcategory," Code of Federal Regulations, Title 40, Part 435, Subpart A, edition of July 1, 2025, govinfo.gov, 2025.

6.10 Exercises

6.1 Concept check · 1 min

In case 2 a 15 m^3 slug reaches the separator over one minute. Why does neither tight nor averaging level tuning keep it out of the vessel?

(a) The slug arrives at up to 1414 m^3/h, while both outlet valves fully open pass only about 433 m^3/h at design pressure.

(b) LT-100 lags the level by 2 s, so the loop sees the slug too late to open the oil valve.

(c) The gas held back behind the slug lowers the pressure until the liquid valves pass almost nothing.

(d) The tight preset's 240 s integral time is four times the 60 s the slug takes to arrive.

6.2 Concept check · 2 min

In case 6 the water valve sticks at 60 s and the water cut then rises to 45 %. With the supervisor off, the total level stays on its setpoint while the interface climbs to its trip. Why?

(a) The interface loop winds up and its integral is passed to the oil valve, which then holds the interface high.

(b) The level loop holds the total level the only way left to it, by taking out more oil than arrives, so the water layer grows under a thinning pad.

(c) The supervisor is off, so the rising water cut fills the vessel until the high-high level trips.

(d) The relative gain array changes sign once the water valve sticks, so the level loop drives the interface up.

6.3 Concept check · 2 min

In case 2 under tight tuning, AT-101 sends the polished water to the hold tank at 399 s. Why does 8.4 m^3 of off-spec water still reach the sea?

(a) The routing valve waits 60 s after AT-101 passes 30 mg/L before it closes the overboard line.

(b) The hold tank is full, so the diverted water overflows to the sea.

(c) AT-101 reads through a 240 s sample delay. The true value passed 30 mg/L at 148 s, so the off-spec water had left before the diversion.

(d) The hydrocyclone reject is discharged untreated while the bank is overloaded.

6.4 Theory and calculation · 6 min

The level gains and the two presets follow from the vessel geometry. Use the chapter's values: the oil valve passes 298.4 m^3/h and the water valve 134.3 m^3/h fully open at design, the liquid surface is 29.4 m^2 at 1200 mm and the interface 22.4 m^2 at 500 mm.

(a) Compute the rate of change of the total level per percent of oil valve, and of the interface and the total level per percent of water valve, in mm/s. Why does the oil valve not move the interface? *Checkpoint:* 0.0282 mm/s, 0.0167 mm/s and 0.0127 mm/s.

(b) With $k = 0.0281$ mm/s per percent for the total level, compute $\tau_c = 1/(K_c k)$ and $\zeta = \frac{1}{2}\sqrt{\tau_I/\tau_c}$ for the tight preset (5.0 %/cm, 240 s) and the averaging preset (0.4 %/cm, 1500 s). *Checkpoint:* 71 s and 0.92, 890 s and 0.65.

(c) Compare both closed-loop time constants with the 60 s a slug takes to arrive. What does the comparison predict, and why does the tight loop not behave as its τ_c suggests during the arrival?

6.5 Theory and calculation · 6 min

Water drops dispersed in the oil pad must settle out before the oil leaves. At the design 49.5 °C the oil viscosity is 6.79 cP, its density 850 kg/m^3, and the density difference 175 kg/m^3.

(a) Compute the Stokes velocity $g\,\Delta\rho\,d^2/(18\mu)$ of the 540 and 60 µm water drops. *Checkpoint:* about 4.1 mm/s and 0.051 mm/s.

(b) For the 540 µm drop compute the particle Reynolds number $\rho_o v d/\mu$ and apply the correction of Equation (6.6). *Checkpoint:* $Re \approx 0.28$, about 3.86 mm/s.

(c) The oil flows at 149.8 m^3/h over the 22.4 m^2 interface. Compute this overflow rate and say which of the three classes (60, 180 and 540 µm) settle without first coalescing.

6.6 Theory and calculation · 8 min

A slug's arrival is a volume balance that no controller can change. Use the design values: normal inflow 214 m^3/h, both liquid valves fully open 433 m^3/h, and a surge volume of 26.4 m^3 between the 1200 mm normal level and the 2100 mm high-high trip.

(a) The 15 m^3 slug of case 2 arrives as a half-sine over 60 s. What is its peak flow? *Checkpoint:* about 1414 m^3/h.

(b) If both valves were wide open for the whole arrival, how much of the slug would they pass on, and what is the least volume the vessel must store? *Checkpoint:* 3.65 m^3 and about 11.4 m^3.

(c) The tight run stores 12.8 m^3 at its peak. Estimate the peak level from the surface widths, 29.4 m^2 at 1200 mm and about 29.9 m^2 at 1650 mm, and compare with the measured 1629.6 mm. *Checkpoint:* about 1630 mm.

(d) Repeat (b) for the 40 m^3 slug of the event button, which arrives over 90 s, and say what any tuning can do about it.

6.7 App challenge · 8 min

Run case 2, *Liquid slug*, at the default seed and a speed of 30, so a 1200 s run takes 40 s.

(a) Run to 1200 s as loaded (Level tuning Tight). Read the *Off-spec overboard* and *Carryover* tiles, export the CSV and find the largest value in the `oil_out_m3h` column. *Checkpoint:* 8.4 m^3, 335 L/MSm^3, 317 m^3/h.

(b) Reload the case with the reload button, set Level tuning to Averaging before Run, and repeat. *Checkpoint:* 6.9 m^3, 1772 L/MSm^3, 219 m^3/h.

(c) Keep the averaging level loop, set the Interface loop gain to 0.5 %/cm and its integral time to 1500 s (Level tuning then shows Custom), press Reset and run to 1200 s. Read the same two tiles, the worst value on the *OiW in spec* tile and the largest `water_out_m3h`. What did the interface loop change, and at what cost?

6.8 App challenge · 7 min

Run case 7, *Gas backpressure rise*, at the default seed and a speed of 10.

(a) Pause at any time between 400 and 900 s. Read the separator pressure on the *Pressure* tile, export the CSV and, in its last row, read the gas valve command `gas_valve_cmd_pct`, the integral term `ctrl_I_P` and the proportional term `ctrl_P_P`. *Checkpoint:* 11.52 barg, 100 %, integral 100.0 %, proportional about 23 %.

(b) Move the *Gas export header pressure* slider to 8.0 barg and run on for 300 s. Read the lowest pressure on the *Pressure* tile and, from the `P_barg` column, how long after your move the pressure comes back within 0.05 bar of 10.00 barg to stay.

6.9 App challenge · 9 min

Run the *Challenge, wellstream campaign* at the default seed and a speed of 30, so the 2400 s campaign takes 80 seconds. The verdict is taken at 2400 s, and the tiles keep counting after it, so read the result line the challenge records at 2400 s or the CSV row at `time_s` = 2400.

(a) Turn the Quality supervisor on, leave everything else as loaded, and run to 2400 s. Read the in-spec oil (`metric_oil_inspec_m3`), the BS&W share (or `metric_bsw_offspec_s`, the seconds off spec out of 2400 s) and the peak oil outflow. Which clauses of the target fail? *Checkpoint:* 87.3 m^3, 84.4 %, 318 m^3/h.

(b) Set Level tuning to Averaging as well, press Reset and run to 2400 s again. Read the same three values and the verdict.

(c) Explain what each element contributes, and why the feedforward of the suggested configuration is not needed to pass.

6.10 Application area · 4 min

In case 7 the compressor suction header rises to 11.3 barg, and the separator settles at 11.52 barg with its gas valve wide open, 0.48 bar below the high-high trip. A shift supervisor asks for a faster pressure loop. What is your answer, and what protects a real separator in this situation?

6.11 Application area · 4 min

After the case 2 slug the produced-water engineer has three proposals to stop off-spec water going overboard: a faster AT-101 with a 30 s sample delay, a second hydrocyclone bank in service, or the interface loop retuned to 0.5 %/cm and 1500 s. Rank them with the chapter's results and name a cost of each.

6.12 Review the agent · 5 min

An AI assistant reviewing the separator's control design writes: "The relative gain array of the level and interface loops is the identity to within 0.3 %, so the two loops do not interact. Each can be tuned,

alarmed and maintained on its own, and a fault in the water valve stays inside the interface loop: the level loop and the export oil quality are unaffected.” Find the flaw and say how you would check it in the app.

CHAPTER 7

Glycol Dehydration Simulator With an Equilibrium Floor and a Reboiler Ceiling

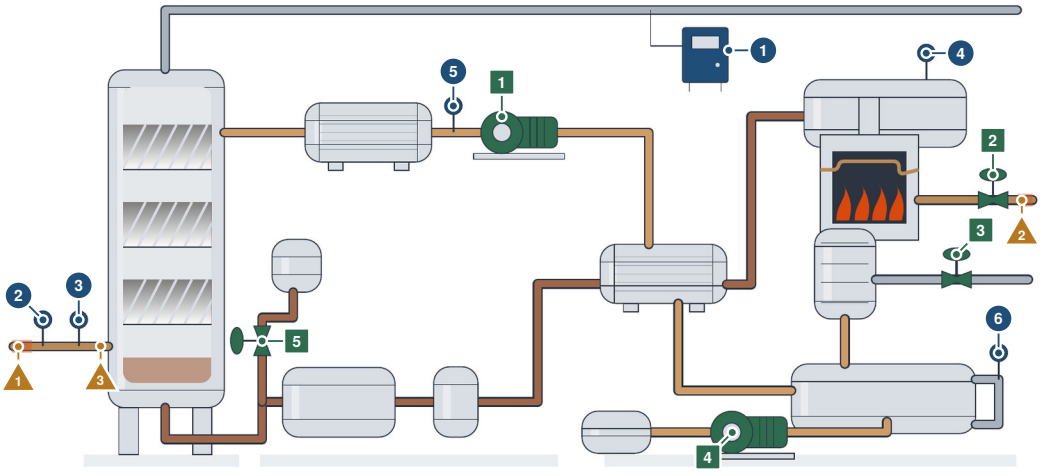

The controller moves 1 glycol pump, the fast handle 2 reboiler fuel valve 3 stripping gas valve 4 glycol makeup pump 5 antifoam injection

The controller reads 1 water left in the dry gas, ten minutes late 2 wet gas flow 3 wet gas temperature 4 reboiler temperature 5 glycol circulation flow 6 surge drum level

What disturbs it 1 more gas, or warmer gas 2 a limit on the reboiler's fuel 3 a slug of liquid hydrocarbon

Notice that the analyzer cabinet samples the dry gas far along its outlet line, and that glycol made purer in the reboiler must pass the stripping column, the surge drum and the exchanger before it reaches the contactor. Illustration, not to scale.

IN PLAIN TERMS **Drying Gas With a Liquid That Drinks Water**

The main idea. Gas from wells carries water that must come out before it enters the pipeline. In a tall tower the gas bubbles up through glycol, a liquid that soaks up water. A heater then boils the water out so the glycol can be used again.

The challenge. More glycol dries the gas within minutes, but only down to a limit set by how dry the glycol is. Drying the glycol further reaches the tower only after about half an hour, burns fuel, and stops before the glycol breaks down.

Why it matters. Wet gas forms icy solids that block pipes, so nearly all gas that is sold passes through such a unit.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the first Did you know? box in the Introduction, the two History boxes in Section 7.2, the Try it boxes and the second Did you know? box in Section 7.6, and Key points at the end of the Conclusion. No equations.
Study	55 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 7.6 with Figure 7.4, Section 7.7 and Section 7.8 with its Open challenge box.

The control question One water specification on the dry gas has two handles: glycol circulation, which acts in minutes but only down to a floor set by equilibrium with the lean glycol, and glycol purity, which moves that floor but reaches the contactor only as the surge drum turns over and is capped by the reboiler ceiling. When is each handle worth its energy, and what goes wrong when a loop reaches for the wrong one?

Learning objectives

- Derive the floor of a two-stage contactor with finite stage efficiency and decide whether any circulation can meet a specification.
- Compute the purity that stripping gas buys below the reboiler from the Kremser relation, and its energy cost.
- Explain why a ratio feedforward on the water load holds a gas step but not a warmer inlet, and what an analyzer trim adds.
- Explain why more circulation makes the gas wetter when the reboiler is short of fuel, and check it on the graded day.
- Weigh the energy of circulation, heat and stripping gas against the water specification, as the economic MPC does.

Prerequisites Sections 1.4.7 (absorption, the Kremser relation and the lean floor), 1.4.8 (selectors and external reset), 1.4.9 (late analyzers), 1.6.2 (the framework's PID) and 1.6.4 (the apps). Section 1.6.3 is needed only for the MPC.

The app https://apmonitor.com/apps/glycol-dehydration/

See also Volume 4, Chapter 3 for limiters at a selector whose integrators track the applied fuel, and Volume 4, Chapter 2 for a protective override that waits at its clamp. Volume 9, Chapter 8 and Volume 9, Chapter 6 for economic predictive control that prices its inputs. Volume 5, Chapter 3 for composition control through a slow analyzer, and Volume 5, Chapter 10 for a product quality inferred and corrected by a late laboratory.

7.1 Introduction

Gas from the production separators arrives carrying all the water it can hold. In the generic unit simulated here, 60 t/h of natural gas, about 60 million standard cubic feet per day, saturated at 35 °C and 60 bar brings 58.2 lb/MMscf of water, 66.2 kg/h in all, into a packed contactor. Against 24.3 L/min of triethylene glycol (TEG) at 99.04 wt%, 22 liters for every kilogram of water arriving, the gas leaves with 6.73 lb/MMscf, a water dew point of −1.6 °C at the contactor pressure, just under a 7 lb/MMscf specification. Doubling the circulation does not halve what is left. From 15 L/min of glycol regenerated at 190 °C, 30 L/min lowers the water from 9.67 to 6.83 lb/MMscf, and 50 L/min still leaves 5.79 lb/MMscf.

The water has to come out because, together with the gas, it forms solids in the pipe well above its own freezing point. Natural gas hydrates are crystalline compounds of water and gas that block valves and lines under pressure, and by 1942 the remedy in practice was to dry the gas upstream of the plant or the line [1]. In 1950 most gas companies took gas only below 7 lb/MMscf, so that virtually all gas sold had to be dehydrated [2], and by 1973 glycols, triethylene glycol above all, were the preferred absorbents [3]. Dehydration is therefore a step that virtually every sales gas passes through, and the specification it serves is the one the simulator keeps as its teaching value.

> **DID YOU KNOW? Far Drier Than the Contract**
>
> In a 2015 test on the transmission system of Spain the gas carried only 4 to 8 parts per million of water by volume, much drier than the European specification, a water dew point below −8 °C at 7 MPa [4]. The 7 lb/MMscf of this chapter is a dew point of −1.0 °C at 60 bar.

The simulator is built around three beliefs that the unit contradicts. The first is that more solvent always removes more water. The gas can be dried only toward the water it would carry in equilibrium

with the lean glycol at the top of the contactor, so with the 45 °C inlet of case 5 and glycol regenerated at 190 °C no circulation brings the water below 8.41 lb/MMscf, and beyond about 570 L/min the glycol, kept 5 K warmer than the gas, heats the contactor and the water rises again. The second belief is that a fixed ratio of glycol to water load holds the dry gas steady. A ratio holds the share of the water removed, so when the inlet warms 2 K and brings 11 % more water, the ratio adds 11 % more glycol and the water left climbs from 6.68 to 7.53 lb/MMscf. The third is that a hotter reboiler always buys purer glycol. The reboiler sits at its bubble point, 99.13 wt% at the 204 °C ceiling in this model, and the last 0.8 wt% has to come from stripping gas, which took the energy of the unit from 154 to 598 kW in case 6.

These features belong to classes of problems that reach well beyond gas processing. Every absorption and regeneration loop, the amine units that sweeten the same gas among them (Chapter 8), delivers a product bounded by the purity of its recycled solvent, and the purity limit of the conventional glycol loop has produced a family of enhanced regeneration schemes of its own [5]. A quality held by a fast inexpensive input and a slow expensive one with a hard limit is a constrained multivariable problem of the kind commercial model predictive control packages are built for [6], and pricing the inputs directly in the objective is the subject of economic predictive control [7]. A ratio of solvent to load used as a feedforward, trimmed by an analyzer, and a high-limit override selected against a temperature loop on a fired vessel are standard structures of process control [8].

However, the material a student meets treats these pieces apart. The equilibrium literature of the TEG, water and gas system is written for contactor design [3, 9], parametric and economic studies of the dehydrator are steady-state simulations [10, 11], and the predictive controller reported for a dehydration unit regulates a feed pressure and a heater temperature on identified first-order models [12]. This chapter quantifies, with a browser simulator of a TEG dehydrator in which the contactor, the still, the surge drum and a delayed moisture analyzer are integrated together, how fixed circulation, a ratio feedforward on the water load, an analyzer trim, a reboiler loop under a high-limit override and an economic model predictive controller (MPC) meet a pipeline water specification through gas and temperature disturbances, foam, a pump trip and a fuel limit. It states the model fully enough for it to be rebuilt.

The correlations are teaching values, and ratio control, override selection and linear MPC are established methods. Every result is measured on a single generic unit at the default random seed. The floor of the contactor is the equilibrium over the lean glycol plus the share of the inlet water that two 90 % stages let through. Under a reboiler fuel limit, extra circulation lowers the purity it was meant to make up for, so an analyzer trim takes the worst sample from 7.04 to 10.47 lb/MMscf. The high-limit override lets the reboiler pass 204 °C by 0.21 K on a sudden loss of load. The model has seventeen states and nine ledger states whose water and glycol balances close to 6.4×10^{-13} of what entered, and a measurement boundary keeps the stage compositions, the lean purity and the foam away from every controller. Each of the nine cases is graded by a completion rule.

Section 7.2 follows gas dehydration from hydrates to enhanced regeneration. Section 7.3 states the model, whose contactor floor is the two-stage, finite-efficiency form of the Kremser relation of Section 1.4.7. Section 7.4 sets the two handles against their two limits, and Section 7.5 states the controllers as coded. Section 7.6 measures the nine cases, Section 7.7 sets the hardware of a real dehydrator beside the model, Section 7.8 lists what the model leaves out, and the exercises in Section 7.10 put the cases to work.

7.2 Dry Gas for the Pipeline

Dehydration began as the cure for the gas hydrate, a solid that blocks gas lines, as the first History box below tells. A subsea flowline that cannot be dried is protected by inhibitor instead (Chapter 5).

HISTORY 1934 and 1942: Ice That Was Not Ice

Before 1934 the plugs that blocked gas lines were blamed on various causes, with a great deal of apprehension. That year Hammerschmidt [13] found that the crystalline plugs forming above the freezing point of water were hydrates of water and gas [14], and with that he brought the subject to the gas industry [1]. In 1942 Carson and Katz described hydrates as white crystalline compounds that exist under pressure well above the freezing point and plug valves and pipes in part or entirely. The cure in practice was to dehydrate the gas before it entered the plant or the pipeline. No method yet predicted when a wet gas would form hydrates, and their equilibrium constants let those conditions be estimated from the composition of the gas [1].

The drying itself was done in several ways. A field study of 1951 named four classes of process then in use, adsorptive, absorptive, chemical reaction and refrigeration, and tested a wellhead unit that dried the gas by expansion through an orifice, noting that the indirect processes spent part of the gas on regeneration and that most of them resisted automation [15]. Calcium chloride was the chemical route. A wellhead unit of 1961 passed the gas through a bed of anhydrous pellets whose brine then absorbed more water on trays below and was dumped into a pit as expendable, and in the San Juan basin and the Big Piney field such units still worked beside glycol dehydrators, although field crews disliked recharging them and could check their performance only by inspecting the bed or running a dew-point test [16].

Absorption into a liquid glycol won on cost, and the second History box below follows the doubts of 1950 to 1973 about how dry it could make the gas.

HISTORY 1950 to 1973: How Dry Can Glycol Make It?

In 1950 adsorption on activated bauxite, alumina or silica gel could make gas virtually bone dry, but a small unit of that kind cost substantially more than an absorption unit that met the pipeline specification, so the search for better absorbent liquids intensified [2]. By 1961 drying with aqueous triethylene glycol had become prominent, yet many doubted that it could dry gas from 120 °F wellheads to the dew point near 32 °F that a contract of 7 lb/MMscf required, a doubt dating from the days when suppliers would not guarantee depressions beyond 65 to 75 °F [9]. In 1973 the published dew points over the glycol still disagreed, and the data above 99.8 wt%, the purity modern contactors needed for depressions of more than 100 °F, were scant and contradictory. Designers answered the doubt with circulation rates higher than necessary, a practice ill suited to offshore platforms where space and power are scarce [3].

Low dew points then moved the problem to the regenerator. At atmospheric pressure a reboiler at 400 °F (204 °C) gives at most about 98.8 wt% while low dew points need up to 99.9 wt%. A hotter reboiler is ruled out, because TEG decomposes thermally at 404 °F (207 °C) and even a 400 °F reboiler can decompose glycol at its hotter film temperatures, so the extra purity comes from stripping gas, best introduced in a short column after the glycol leaves the reboiler [10]. Stripping with nitrogen, part of the dry product gas or a volatile hydrocarbon in the DRIZO process, the Stahl column and the Coldfinger process are among the schemes reviewed for getting past the purity limit of the conventional loop [5], and the trade of circulation, trays and stripping gas against processing cost has been posed as an optimization of the simulated unit [11]. The same loop has an environmental side. Aromatic hydrocarbons are highly soluble in TEG and leave with the regenerator's vent, which is usually open to the atmosphere [10], and condensing the vapor and burning the noncondensable gas in the reboiler was found to be the most cost-effective control [17]. A multivariable predictive controller built on first-order models has been compared with proportional-integral loops on a dehydration unit's feed pressure and heater temperature [12].

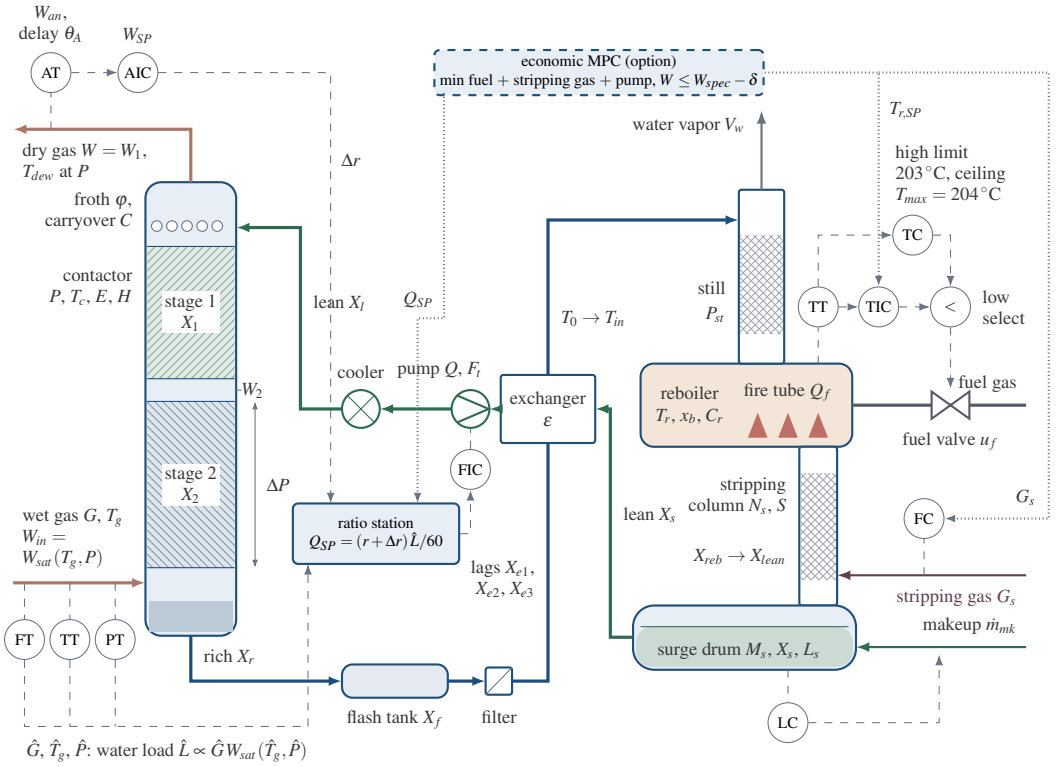

Figure 7.1 TEG dehydration unit as modeled, after the application's wide Unit view. Wet gas G at T_g, saturated with W_{in}, rises through two stages holding X_1 and X_2 at contactor temperature T_c and efficiency E, and leaves with W past the analyzer AT, which reports W_{an} after the delay θ_A. The rich glycol X_r passes the flash tank X_f, the filter, three lags X_{e1} to X_{e3} and the cold side of the lean–rich exchanger of effectiveness ε to the still on the reboiler (T_r, duty Q_f from the fuel valve u_f), which boils off V_w. The stripping column with N_s stages and stripping gas G_s takes X_{reb} to X_{lean}, the surge drum holds M_s at X_s and level L_s, and the pump returns Q through the lean line X_l. The ratio station sets Q_{SP} from the water load $\hat{L}$ and the analyzer trim Δr, the TIC and the high-limit override meet at a low select on u_f, the FC holds G_s, the LC switches the makeup $\dot{m}_{mk}$, and the optional MPC (dotted) sets Q_{SP}, $T_{r,SP}$ and G_s. The contactor level and flash tank pressure loops are not modeled. The drawing shows the process model and is neither an instrument diagram nor to scale.

7.3 A Contactor and Its Regenerator as Modeled

The simulated unit is a generic TEG dehydrator for 60 t/h of a gas of 0.846 kg/Sm^3 at 60 bar, 1182 Sm^3/min or 60.2 million standard cubic feet per day, as shown in Figure 7.1. Wet gas saturated at the inlet temperature enters the bottom of a packed contactor of two theoretical stages and leaves the top past a moisture analyzer. The rich glycol passes a flash tank, a filter and the cold side of the lean–rich exchanger into a still on a fired reboiler at atmospheric pressure, flows down a stripping column into the surge drum, and returns across the exchanger through a pump and a cooler to the top of the contactor. Water contents are carried in milligrams per standard cubic meter at 15 °C and shown in pounds per million standard cubic feet at 60 °F, with 1 lb/MMscf = 16.05 mg/Sm^3. The clock is the minute.

Each variable in Table 7.1 has a role in the loops. Seventeen physical states are integrated, the water per kilogram of TEG on the two stages X_1 and X_2, in the flash tank X_f and in three transport lags X_{e1} to X_{e3}, the reboiler temperature T_r, the TEG in the surge drum M_s and its water ratio X_s, the

Table 7.1 Dehydrator Variables, Their Roles and Design Values (Case 3 at $t = 0$)

Variable	Role	Design	Range
W	CV, water in dry gas (true)	6.73 lb/MMscf	spec 7 lb/MMscf
W_{an}	measured W, 10 min late	6.73 lb/MMscf	2-min samples
T_r	CV, reboiler	200.0 °C	SP 150–204 °C
L_s	CV, surge drum level	60 %	low 30, trip 10
Q	MV, circulation	24.28 L/min	0–50 (standby 16)
r, Δr	MV, ratio and trim	22 L/kg, 0	5–60 L/kg
u_f	MV, fuel valve	44.2 %	0–100 %
G_s	MV, stripping gas	0	0–50 Sm^3/h
$\dot{m}_{mk}$, antifoam	MV, on or off	Auto, off	60 kg/h, on
G	DV, gas flow	60 t/h	30–90 t/h
T_g	DV, inlet temperature	35 °C	20–50 °C
a_f, a_s	DV, fuel, stripping supply	100 %	20–100, lost
pump, h	DV, standby pump, slug	main, 0	16 L/min, 1.5
X_1, X_2	states, stages	0.0177, 0.0466	kg/kg
X_f, $X_{e1..3}$	states, rich transit	0.0466	kg/kg
M_s, X_s	states, surge drum	784.4 kg, 0.00971	—
Q_f	state, absorbed duty	110.6 kW	0–250 kW
φ, a	states, foam, antifoam	0.45, 0	carryover above 1
T_c, E	outputs, contactor	36.09 °C, 0.90	—
T_{dew}	output, dew point at P	−1.61 °C	—
w_{lean}, w_{rich}	outputs, purities	99.04, 95.55 wt%	—
ΔP, losses	outputs, foam signs	25 mbar, 1.12 kg/h	—
energy	output, fuel + strip + pump	151.5 kW	budget 250 kW

lean line X_l, the delivered circulation Q, the absorbed burner duty Q_f, the stripping gas flow G_s, the foam index φ, the liquid hydrocarbon in the glycol h, the antifoam at the packing a and the makeup flow $\dot{m}_{mk}$. Nine more states are ledgers that integrate every boundary flow of water and TEG, so both balances can be checked to round-off. The gas side, the reboiler composition and the stripping column are algebraic at each instant.

7.3.1 Water in the Gas and Its Dew Point

The inlet gas carries the saturated water content of the Bukacek form, a two-term algebraic fit of the McKetta–Wehe chart, and the dew point is its inverse at the contactor pressure, as shown in Equation (7.1).

$$W_{sat}(T,P) = 16.05\left[47\,484\,\frac{p_v(T)}{P} + 10^{6.69449-3083.87/T_R}\right]$$
$$W_{in} = W_{sat}(T_g,P), \qquad W_{sat}(T_{dew},P) = W \tag{7.1}$$

where W_{sat} is in mg/Sm^3, P and the vapor pressure of water p_v are in psia and T_R is the temperature in degrees Rankine. Above 0 °C p_v is the region 4 saturation line of IAPWS-IF97 [18], and below it the Magnus form over supercooled water, so a dew point under 0 °C is that of metastable liquid water. The inlet at 35 °C carries 58.18 lb/MMscf and at 45 °C 96.32 lb/MMscf, 1.66 times as much. Against a published fit of the McKetta–Wehe chart for sweet gas [19] the form gives 0.0, −0.2 and −0.4 % at 35, 38 and 45 °C and 60 bar. A water content of 7 lb/MMscf is a dew point of −1.0 °C at 60 bar.

7.3.2 Water Held by the Glycol

The gas in equilibrium with a TEG solution carries the water activity of the solution times the saturated content at the contactor temperature, with a one-parameter Margules coefficient, as shown in Equation (7.2).

$$W_{eq}(X) = \gamma_w x_w W_{sat}(T_c,P), \qquad \ln\gamma_w = A(T)(1-x_w)^2$$
$$A(T) = 0.256 - \frac{228.7\,\mathrm{K}}{T}, \qquad x_w = \frac{X/M_w}{X/M_w + 1/M_T} \tag{7.2}$$

where X is the water per kilogram of TEG, x_w its mole fraction and $M_w = 18.015$ and $M_T = 150.17$ the molar masses. The coefficient is a teaching correlation, not a fit, with $\gamma_\infty = 0.615$ at 35 °C and 0.800 at 204 °C. Because the activity is nearly proportional to the water left in the glycol, each tenfold cut in that water buys about the same step in dew point, so the depression rises steeply with purity. At a 35 °C contactor the equilibrium dew point at 60 bar is −7.7, −13.4, −22.8, −42.6 and −67.1 °C over 98.5, 99.0, 99.5, 99.9 and 99.99 wt%, a depression from 42.7 to 102.1 K. Against published readings and a published fit of the TEG–water dew-point chart of the Gas Processors Suppliers Association, drawn from the measurements of Parrish, Won and Baltatu, these dew points lie within about 1 K from 98.5 to 99.9 wt% and about 3 K drier at 99.99 wt%. The largest difference among the readings is at 99.9 wt%, 100 °F and 1000 psia, where the model gives −41.5 °C and the chart −40 °C. The measurements above 99.8 wt% are the ones the literature found least consistent [3].

7.3.3 Two Stages and the Floor They Leave

Each stage holds H of TEG and passes a share E of the way to equilibrium on the gas side, and its water ratio follows from a balance on the glycol, as shown in Equation (7.3).

$$\begin{aligned} W_2 &= (1-E)\,W_{in} + E\,W_{eq}(X_2) \\ W &= (1-E)\,W_2 + E\,W_{eq}(X_1) \\ H\frac{\mathrm{d}X_1}{\mathrm{d}t} &= F_t X_l - (F_t - v)X_1 + G(W_2 - W) \\ H\frac{\mathrm{d}X_2}{\mathrm{d}t} &= F_1(X_1 - X_2) + G(W_{in} - W_2) \\ T_c &= T_g + \frac{Lc_{pL}\,\Delta T_\ell + h_{abs}N}{G_m c_{pG} + Lc_{pL}} \end{aligned} \tag{7.3}$$

where G is the gas in $\mathrm{Sm^3/min}$ (with the water contents in $\mathrm{kg/Sm^3}$ in these balances), $F_t = \rho Q/(1+X_s)$ the TEG pumped, v the TEG vaporized into the gas at its partial pressure over the top stage, $F_1 = F_t - C_T - v$ the TEG leaving the top stage after the carryover C_T, L the solution flow, $\Delta T_\ell = 5\,\mathrm{K}$ the lean glycol's approach above the gas, N the water absorbed and G_m the gas mass flow. The contactor temperature is found in two passes, the second with the absorption of the first. At design $T_c = 36.09\,^\circ\mathrm{C}$, the stages hold 0.0177 and 0.0466 kg/kg, the gas leaves stage 2 with 17.89 lb/MMscf and stage 1 with 6.73 lb/MMscf, 88.4 % of the inlet water is absorbed, and the gas in equilibrium with the lean glycol would carry 3.05 lb/MMscf.

With unlimited circulation both stages sit at the lean composition and the outlet tends to $W^* + (1-E)^2(W_{in} - W^*)$, with W^* the equilibrium over the lean glycol, which is the floor no circulation can buy. Two stages at $E = 0.9$ pass 1 % of the inlet water's excess over W^*. This is the floor of Equation (1.12) at an unlimited absorption factor, with each ideal stage replaced by one of efficiency E.

Pause. Try Exercise 7.1 before reading on.

7.3.4 The Still at Its Bubble Point

The reboiler boils the rich glycol down to the composition whose bubble point at the still pressure is its temperature, as shown in Equation (7.4).

$$\begin{aligned} P_{st} &= x_b\,\gamma_w(x_b, T_r)\,p_{sat}(T_r) + (1-x_b)\,p_{TEG}(T_r) \\ p_{TEG}(T) &= p_0 \exp\Big[-\beta\Big(\frac{1}{T} - \frac{1}{T_b}\Big)\Big] \\ X_{reb} &= \min[X_r,\ X(x_b)], \qquad V_w = F_2\,(X_r - X_{reb}) \end{aligned} \tag{7.4}$$

where p_{sat} is the IAPWS-IF97 saturation pressure, $P_{st} = 1.013\,25\,\mathrm{bar}$, p_{TEG} a Clausius–Clapeyron teaching fit through the normal boiling point of TEG, $p_0 = 101.325\,\mathrm{kPa}$ at $T_b = 558.15\,\mathrm{K}$ (285 °C)

with $\beta = 8800\,\mathrm{K}$, $X_r = X_{e3}$ the rich glycol arriving and $F_2 = F_1$ its TEG flow. The bubble point is solved by six Newton steps from the infinite-dilution guess, a fixed count. The reboiler makes 98.76 wt% at 190 °C, 99.04 at 200 °C, 99.11 at 203 °C and 99.13 at 204 °C, with $\gamma_w = 0.823$ at design. These purities are 0.48, 0.36 and 0.34 wt% above the 98.0, 98.5 and 98.8 wt% that a published process simulation gives at 360, 380 and 400 °F (182.2, 193.3 and 204.4 °C) at one atmosphere [10], so this reboiler leaves 0.72 to 0.75 times the water that one does. The fit was kept because it meets the targets of the plan for the application, about 98.5 wt% at 190 °C and 99.1 wt% at 204 °C, and because a refit to the published values would move every case. With 98.68 wt% at 200 °C, the published value interpolated, the design ratio of 22 L/kg would leave 7.71 lb/MMscf and the specification would need about 28 L/kg. The vapor pressure of TEG lies 22 % under Stull's data at 200 °C and moves the purity by 0.016 wt%.

7.3.5 Stripping Below the Reboiler

Dry gas blown up the stripping column removes a further share of the water, from the Kremser relation for N_s theoretical stages, as shown in Equation (7.5).

$$X_{lean} = X_{reb}\,\frac{S-1}{S^{N_s+1}-1}, \qquad S = \frac{\gamma_w\, p_{sat}(T_r)}{P_{st}} \cdot \frac{g/V_m}{\rho_h/M_T} \tag{7.5}$$

where $g = G_s/(0.06\,Q)$ is the stripping gas in $\mathrm{Sm^3}$ per cubic meter of glycol, $V_m = 23.645\,\mathrm{Sm^3/kmol}$, $\rho_h = 1010\,\mathrm{kg/m^3}$ the hot glycol density and $N_s = 2$. At 203 °C and 24.3 L/min the equilibrium ratio is 13.45, and 5, 10, 20 and 45 $\mathrm{Sm^3/h}$ give S = 0.29, 0.58, 1.16 and 2.61 and 99.35 wt%, 99.53, 99.74 and 99.91 leaving the column. Stripping gas also takes from the reboiler 0.9 of the heat that warms it from ambient to T_r (q_s below), and it is counted at its heating value, $36\,\mathrm{MJ/Sm^3}$ or 10 kW per $\mathrm{Sm^3/h}$.

7.3.6 The Reboiler's Heat Balance

The reboiler temperature is a state whose energy balance carries the duty, the sensible heat of the rich glycol after the exchanger, the heat of boiling with reflux and the losses, as shown in Equation (7.6).

$$\begin{aligned} C_r\frac{\mathrm{d}T_r}{\mathrm{d}t} &= Q_f - F_2(1+X_r)\,c_{pL}(T_r - T_{in}) - V_w\lambda\,(1+R) \\ &\quad - UA\,(T_r - T_{amb}) - q_s \\ T_{in} &= T_0 + \varepsilon\,(T_r - T_0), \qquad T_0 = T_c + 30\,\mathrm{K} \end{aligned} \tag{7.6}$$

where every term is in kJ/min, so a duty in kW enters multiplied by 60, T_0 is the rich glycol after the reflux coil, q_s the heat taken by the stripping gas, and the other symbols are listed in Table 7.2. At design the burner delivers 110.6 kW, of which 46.3 kW heats the rich glycol arriving at 159.8 °C, 46.8 kW boils off 58.6 kg/h of water and 17.5 kW is lost, 44.2 % of a 250 kW burner and 147.5 kW of fuel burned at a firing efficiency of 0.75. At a fixed duty the reboiler's own time constant is 112 min. More circulation means more glycol to heat, so at a fixed duty it lowers T_r and the purity with it, which is the coupling that Manual operation and the fuel limit expose.

7.3.7 The Glycol Loop, Losses and Foam

The rich glycol passes the flash tank and three equal lags for the filter, exchanger and lines, $H_f\dot{X}_f = F_2(X_2 - X_f)$ and $(H_{tr}/3)\dot{X}_{ei} = F_2(X_{e,i-1} - X_{ei})$, together 16.4 min of residence at design. The surge

drum and the loop's losses close the circuit, as shown in Equation (7.7).

$$\begin{aligned} \frac{\mathrm{d}M_s}{\mathrm{d}t} &= F_{ret} + \dot{m}_{mk} - F_t, \qquad F_{ret} = F_2 - v_s - r_d \\ M_s \frac{\mathrm{d}X_s}{\mathrm{d}t} &= F_2 X_{lean} - X_s (F_{ret} + \dot{m}_{mk}) \\ H_l \frac{\mathrm{d}X_l}{\mathrm{d}t} &= F_t (X_s - X_l) \\ C &= C_0 \frac{G}{G_0} + C_f \max(0, \varphi - 1)^{1.5} \\ \tau_\varphi \frac{\mathrm{d}\varphi}{\mathrm{d}t} &= \varphi_0 \, \tau_f \left(\frac{G}{G_0} \right)^2 - \varphi, \qquad \tau_f = \tau_{f0} \frac{1 + k_h h}{1 + k_a a} \end{aligned} \tag{7.7}$$

where v_s is the TEG leaving the still overhead after a 98 % reflux recovery, $r_d = k_d[\exp((T_r - T_{max})/4\,\mathrm{K}) - 1]$ the decomposition above the ceiling T_{max}, C the solution carried over with the gas, h the liquid hydrocarbon from a slug, cleared with a time constant of 60 min, and a the antifoam reaching the packing over 5 min. The surge drum level is $L_s = 100 M_s (1 + X_s)/(\rho V_s)$ in percent. The foam lowers the stage efficiency, $E = E_0/[1 + k_E \max(0, \varphi - 1)]$, and raises the pressure drop, $\Delta P = \Delta P_0 (G/G_0)^2 (1 + 0.8\varphi^2)/(1 + 0.8\varphi_0^2)$. At design the unit loses 0.49 kg/h of TEG by carryover, 0.02 kg/h into the gas and 0.61 kg/h from the still, 1.12 kg/h in all, and the surge drum holds 784.4 kg of TEG at 60 %, a turnover of 29.7 min.

7.3.8 Actuators, Instruments and Energy

The pump, burner, stripping gas valve and makeup pump follow their commands through first-order lags of 0.2, 1.0, 0.3 and 0.5 min. The pump delivers at most 50 L/min, or 16 L/min on the standby pump, the stripping gas at most 50 $\mathrm{Sm^3}$/h, and the burner at most 250 kW times the fuel available. The moisture analyzer draws a sample every 2 min and reports it 10 min after it left the gas, 8 min of sample line and 2 min of analysis, with 0.05 lb/MMscf of noise. The gas meter has a 6 s lag and 0.2 % of noise, the inlet temperature a 30 s lag and 0.05 K, the reboiler thermocouple an 18 s lag and 0.1 K, the circulation meter 0.3 % of noise and the level 0.2 percentage points. The noise is seeded as described in Section 1.6.1, so two runs that differ in one setting see identical noise. The energy the challenge judges is the fuel burned, the stripping gas at its heating value and the pump's shaft work, 0.164 kW per L/min against 59 bar, 151.5 kW at design. A burner trip at 210 °C and a pump trip at a 10 % surge level latch and stop the run.

7.3.9 Parameters, Start and Integration

Table 7.2 gives every parameter. None was measured on a plant. They were chosen as typical of a TEG unit of this size and so that the design point closes. A case starts lined out at its own settings, with the circulation its architecture asks for and the reboiler at its setpoint, the contactor solved by Newton's method on the two stage balances, the still and the surge drum iterated to a consistent lean composition, and the stations of Manual lined up with the automatic operating point. At design the derivatives of the compositions and of the reboiler temperature are below 2×10^{-16} of their states, and the surge drum falls only by the losses until the makeup pump runs. Integration uses fourth-order Runge–Kutta steps of 2 s, the controllers execute every 6 s and a row is logged every minute, on the engine of Section 1.6.1. Halving the step moves no lesson number by more than 1.05×10^{-6} of its value and no crossing time or verdict. At the end of every case, as shipped and with its learner moves, the water ledger closes to 6.4×10^{-13} and the TEG ledger to 2.1×10^{-14} of what entered. A simulated hour costs well under 0.1 s of computing in a JavaScript engine, with or without the MPC.

7.3.10 Assumptions and Where They Fail

The contactor is two well-mixed stages with a gas-side efficiency, not a packed bed with mass-transfer coefficients, and the gas holds no water, so the outlet answers a stage at once. That is enough for the

Table 7.2 Model Parameters (A Adjustable in the Application, H Hidden From the Learner, F Constant in the Code)

Symbol	Meaning	Value	Unit	
P	contactor pressure (absolute)	60	bar	A
G_0, ρ_G	design gas, density at 15 °C	60, 0.846	t/h, kg/Sm^3	H
c_{pG}, c_{pL}	gas, glycol heat capacity	2.4, 2.5	kJ/(kg K)	H
E_0	stage efficiency (Murphree, gas)	0.90	—	A
H	TEG holdup per stage	55	kg	A
ΔT_ℓ	lean glycol above inlet gas	5	K	A
h_{abs}	heat of absorption	2400	kJ/kg	H
ρ, ρ_h	lean, hot glycol density	1100, 1010	kg/m^3	H
H_f, H_{tr}	flash tank, transit holdup (TEG)	270, 165	kg	A
H_l	lean line and cooler holdup	55	kg	H
V_s	surge drum volume	1.2	m^3	A
Q_{max}, Q_{sb}	main, standby pump	50, 16	L/min	A
τ_P, η_P	pump lag, efficiency	0.2, 0.6	min, —	H
$\dot{m}_{mk}$, τ_{mk}	makeup rate, lag	60, 0.5	kg/h, min	A, H
L_{lo}, L_{trip}	low alarm, pump trip level	30, 10	%	A
—	Margules constants of $A(T)$	0.256, −228.7	—, K	F
β, T_b	TEG vapor pressure slope, boiling point	8800, 285	K, °C	F
P_{st}	still pressure	1.01325	bar	A
$Q_{f,max}$	burner capacity (absorbed)	250	kW	A
η_F, τ_F	firing efficiency, burner lag	0.75, 1	—, min	H
C_r	reboiler heat capacity	3000	kJ/K	A
ε	lean–rich exchanger effectiveness	0.70	—	A
—	reflux coil heating ($T_0 - T_c$)	30	K	H
λ, R	heat to boil water, reflux	2300, 0.25	kJ/kg, —	H
UA, T_{amb}	reboiler loss, ambient	0.1, 25	kW/K, °C	H
N_s	stripping column stages	2	—	A
$G_{s,max}$, τ_G	stripping gas valve, lag	50, 0.3	Sm^3/h, min	A, H
—	still TEG recovery	0.98	—	H
T_{max}, T_{trip}	ceiling, burner trip	204, 210	°C	A
k_d	decomposition scale	2	kg/h	H
LHV	fuel and stripping gas	36	MJ/Sm^3	H
τ_{f0}, φ_0	foaming tendency, design foam	1.0, 0.45	—	A, H
k_h, k_a, k_E	slug, antifoam, efficiency gains	2.5, 3, 0.6	—	H
τ_h, τ_a, τ_φ	slug, antifoam, foam times	60, 5, 3	min	A, H, H
C_0, C_f, ΔP_0	mist, foam carryover, clean drop	0.5, 113, 25	kg/h, kg/h, mbar	H, A, H

floor, the diminishing return on circulation and plausible sizes, though not for sizing a real contactor. The same temperature serves both stages. The equilibrium scales the saturated content by the water activity, the activity coefficient, the vapor pressure of TEG and the two theoretical stripping stages are teaching correlations, and the reboiler is purer than published simulations by the amounts stated above. The saturated content is that of sweet gas, without corrections for gravity, salt, carbon dioxide or hydrogen sulfide, and below 15 °C the dew point extrapolates the Bukacek form below its stated range. The still is quasi-steady at its bubble point and holds no water, the flash tank and the lines carry composition at constant holdups with no flashed hydrocarbon, and the contactor level and the flash tank pressure are taken as held by their own fast loops, which are not modeled and have no faceplates. Foam, carryover, vaporization and decomposition are synthetic laws chosen for about 1 kg/h of losses at design and about 220 kg/h at the foam peak of case 7. The 204 °C ceiling, the 210 °C burner trip, the 10 % level trip and the 7 lb/MMscf specification are teaching values. The ceiling sits just under the 207 °C decomposition temperature a published design study gives for TEG [10], while real limits come from the glycol supplier's data, the unit's design standard and its burner management system, and the water specification from the transport contract [2].

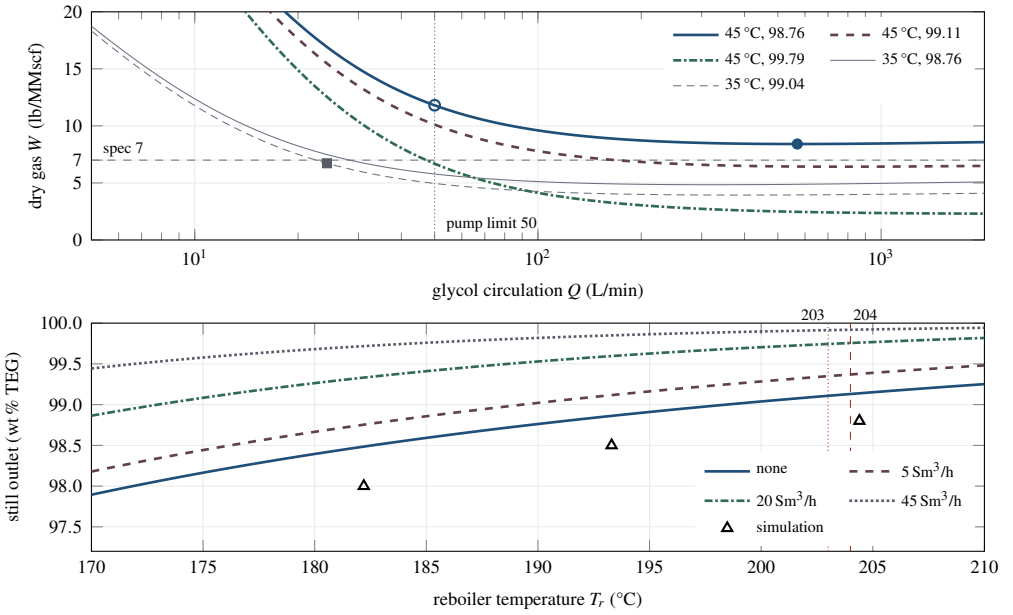

Figure 7.2 The two handles from the application's own functions. Top, steady water in the dry gas against circulation for the 45 °C inlet of case 5 and the 35 °C design inlet, each at the lean purity in wt % the legend names. The 45 °C curves are thick (98.76 solid, 99.11 dashed, 99.79 dash-dotted), and the 35 °C curves are thin and gray (98.76 solid, 99.04 densely dashed). The horizontal dashed line is the specification and the vertical dotted line the pump limit. The filled dot is the 45 °C floor at 98.76 wt%, 8.41 lb/MMscf near 570 L/min, the ring is the state of Figure 7.3, and the square is the design point. Bottom, glycol leaving the still against reboiler temperature at 24.3 L/min with no stripping gas (solid) and with 5 (dashed), 20 (dash-dotted) and 45 Sm^3/h (dotted), beside the 203 °C override (dotted vertical) and the 204 °C ceiling (dashed vertical). Triangles are a published process simulation at one atmosphere [10].

7.4 A Fast Handle, a Slow Handle and Two Limits

The two handles act through different holdups. Circulation reaches the outlet through the stages, 2.08 min of glycol each at design, so in case 1 a step from 15 to 30 L/min at 10 min takes the water under 7 lb/MMscf at 17 min. Purity reaches the outlet only as the surge drum turns over. In case 4 the glycol leaving the still is at 99.11 wt% by 30 min while the glycol entering the contactor is still at 98.88 wt%, and the water crosses under the specification 32 min after the reboiler step.

The steady water in the dry gas against circulation and the purity the still delivers against reboiler temperature, computed with the application's own functions, are shown in Figure 7.2. At the design inlet, 99.04 wt% glycol at 24.3 L/min gives 6.73 lb/MMscf and the least water any circulation reaches is 3.94 lb/MMscf, near 390 L/min. On the 45 °C inlet with 98.76 wt% glycol the least is 8.41 lb/MMscf near 570 L/min, above the specification, and with unlimited glycol the water returns to 8.87 lb/MMscf because the glycol, 5 K warmer than the gas, warms the contactor. At 50 L/min the contactor is at 46.8 °C, where the gas in equilibrium with the lean glycol already carries 6.83 lb/MMscf, and even with perfect stages the least would be about 7.2 lb/MMscf. The floor is a property of the equilibrium and the stage efficiency, so of the operator's handles only purity moves it. At 99.11 wt%, the most the reboiler gives under its override, the least water is 6.43 lb/MMscf, and at the 99.79 wt% that 50 Sm^3/h of stripping gas makes, 50 L/min reaches 6.69 lb/MMscf.

The lower panel shows the ceiling. Heat alone moves the still outlet from 98.76 wt% at 190 °C to 99.11 at 203 °C, a gain of 0.35 wt%, and past that point each further Sm^3/h of stripping gas buys less

purity than the last. At 203 °C the first 5 Sm^3/h raise the purity by 0.24 wt% for 47 kW, while the step from 20 to 45 Sm^3/h raises it by 0.17 wt% for 251 kW. The two limits therefore bound a cheap region and an expensive one. A specification below the floor cannot be met by circulation at all, and above the ceiling every tenth of a percent of purity is paid for in stripping gas.

At a fuel limit the handles stop being independent. When the burner is wide open, every extra liter of glycol is a liter the reboiler cannot heat to its setpoint, so the purity falls and the floor rises. A loop that answers a wet analyzer with more circulation then makes the gas wetter, the opposite of its intent, which is what the challenge of Section 7.6 tests.

🕮 **Pause.** Try Exercise 7.3 before reading on.

7.5 The Controllers in the Code

The controllers execute every $\Delta t_c = 0.1$ min, three physics steps, and read nothing of the plant but a packet holding the analyzer's last report, the gas meter, the inlet temperature and pressure, the reboiler thermocouple, the flows, the level, the contactor pressure drop and the pump's motor contacts. An application test fails if any other key appears in it. The header's Manual and Auto switch the pump and the fuel valve together, and in Manual both sit at their stations while every loop tracks them, so a transfer back to Auto is bumpless (Section 1.6.2).

7.5.1 The Common PID

The reboiler loop and the trim use the framework's velocity-form PID of Section 1.6.2, written in this chapter's symbols in Equation (7.8), with σ for the action a of Equation (1.16).

$$
\begin{aligned}
e_k &= \sigma\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c e_k \\
s_k &= s_{k-1} + P_k - P_{k-1}, \qquad \hat{u}_k = u_0 + s_k + I_{k-1} \\
u_k &= \min\!\big[\max(\hat{u}_k, u_{lo}), u_{hi}\big] \\
I_k &= I_{k-1} + \Big[\frac{K_c}{\tau_I} e_k + \frac{u_k - \hat{u}_k}{T_t}\Big]\Delta t
\end{aligned}
\tag{7.8}
$$

where $\sigma = \pm 1$ sets the direction, $T_t = \tau_I$, and integral times of 600 min and longer turn the integral off. Neither loop uses derivative action. The reboiler loop (TIC) acts with $\sigma = +1$ on the fuel valve, so a reboiler 1 K under its setpoint opens the valve 3 points, with $K_c = 3\,\%/\mathrm{K}$, $\tau_I = 45$ min, limits 0 and 100 % and $\Delta t = \Delta t_c$.

7.5.2 The High-Limit Override

A proportional override with external reset feedback acts 1 K under the ceiling and meets the TIC at a low select, as shown in Equation (7.9).

$$
\begin{aligned}
u_{HL} &= \mathrm{clip}\big[K_{HL}(T_{HL} - \tilde{T}_r) + \tilde{u}_f,\ 0,\ 100\big] \\
u_f &= \min(u_{TIC}, u_{HL}), \qquad \tau_{fb}\frac{\mathrm{d}\tilde{u}_f}{\mathrm{d}t} = u_f - \tilde{u}_f
\end{aligned}
\tag{7.9}
$$

where $K_{HL} = 10\,\%/\mathrm{K}$, $T_{HL} = 203\,^\circ\mathrm{C}$, $\tilde{T}_r$ is the thermocouple through a 0.5 min filter and $\tilde{u}_f$ the applied fuel command through the lag $\tau_{fb} = 2$ min. When the override is selected the TIC tracks its output. Because the feedback is the valve actually applied, the override holds the reboiler just under the limit when it is in command and stays out of the way below it, the external-reset form of Section 1.4.8. A jet engine's fuel control selects among several limiters on the same principle (Volume 4, Chapter 3). After a setpoint step to 204 °C in case 4 the thermocouple peaks at 203.7 °C and settles at 202.7 °C. The override does not hold the ceiling on a sudden loss of load. With the setpoint at 204 °C before the pump trip of case 8, the true reboiler reaches 204.21 °C and stays above 204 °C for 2.3 min, decomposing 2.6 g of TEG, and the gas returning to 60 t/h in case 3

gives 204.12 °C for 1.8 min. The override acts 1 K below the ceiling through a filter on an 18 s thermocouple, while the burner answers in a minute and the load vanishes in seconds. In Manual nothing holds the ceiling, and a fuel valve wide open from 5 min trips the burner at 13.4 min.

7.5.3 Ratio Feedforward and Analyzer Trim

Circulation is fixed at a setpoint, set in ratio to the water load the instruments infer, trimmed by the analyzer, or left to the MPC, as shown in Equation (7.10).

$$\begin{aligned} \hat{L} &= \frac{10^{-3}\,\hat{G}\,W_{sat}(\hat{T}_g,\hat{P})}{\rho_G}, \qquad Q_{SP} = \frac{(r+\Delta r)\,\hat{L}}{60} \\ \Delta r &= \mathrm{PI}(W_{an} - W_{SP}) \\ r_{lo} - r \le \Delta r &\le \min(r_{hi},\ 60\,Q_{cap}/\hat{L}) - r \end{aligned} \tag{7.10}$$

where $\hat{L}$ is the water load in kg/h from the gas meter $\hat{G}$ in t/h and the measured inlet temperature and pressure, $r = 22\,\mathrm{L/kg}$ by default, $r_{lo} = 5$ and $r_{hi} = 60\,\mathrm{L/kg}$, and Q_{cap} the capacity of the running pump. The ratio is a feedforward on both disturbances. The trim is the PID of Equation (7.8) with $\sigma = -1$, $K_c = 3\,\mathrm{L/kg}$ per lb/MMscf, $\tau_I = 20\,\mathrm{min}$ and a target W_{SP} of 6.5 lb/MMscf, below the specification for the analyzer's 10 min delay. A slow analyzer limits composition loops in the same way on a distillation column (Volume 5, Chapter 3). It executes once per fresh report over the time since its last execution, and its upper bound is the ratio the pump can deliver, so it cannot wind up against a saturated pump. Switched off, it tracks no correction, so switching it on is bumpless.

7.5.4 Economic Predictive Control

With MPC selected, a receding-horizon quadratic program sets the circulation, the reboiler setpoint and the stripping gas every 2 min over a 2 h horizon, as shown in Equation (7.11). It specializes the framework's predictive controller of Section 1.6.3 with linear energy prices on the inputs.

$$\begin{aligned} \min_{u}\ & \sum_{k=1}^{60} c^{\top} u_k + \sum_{j} \Delta u_j^{\top} S\,\Delta u_j + 10^{3}\varepsilon^{2} + 10^{4}\varepsilon \\ \text{s.t.}\quad & \hat{W}_k = W_{map}(\bar{z}) + K(z_k - \bar{z}) + b \le W_{spec} - \delta + \varepsilon \\ & z_{k+1} = A_d z_k + B_d u_k + c_d, \qquad u_{lo} \le u_k \le u_{hi} \end{aligned} \tag{7.11}$$

where $u = [Q,\ T_{r,SP},\ G_s]$ and the states z are those inputs lagged through what each passes on its way to the outlet, the circulation through $\tau_Q = 4\,\mathrm{min}$, the reboiler setpoint through $\tau_T = 10\,\mathrm{min}$ and then the surge drum's turnover, and the stripping gas through the surge drum. The map W_{map} is the controller's own copy of the steady contactor and still of Section 7.3, fed by the measured gas, inlet temperature and pressure with the parameters in force when the case started, a Wiener model linearized by central differences at every execution, so the gain K falls toward zero as the circulation approaches the floor. The linear costs c are each input's marginal energy in kW of fuel, about 2.8 kW per L/min of circulation for reheating and pumping, the fuel to hold a hotter reboiler, and 10 kW per $\mathrm{Sm^3/h}$ of stripping gas. The water bound is soft, with a margin $\delta = 1\,\mathrm{lb/MMscf}$ and one slack ε. Move suppression on the block increments Δu_j is $S = \mathrm{diag}(2, 2, 10)$, a weight of 10^{-6} on the inputs themselves keeps the program strictly convex, the bounds are 5 to Q_{cap}, 170 to 203 °C and 0 to 50 $\mathrm{Sm^3/h}$, the moves are limited to 6 L/min, 1 K and 6 $\mathrm{Sm^3/h}$ per execution, and the 60 steps are blocked as 1, 1, 2, 4, 8, 14 and 30. The bias b is updated from each analyzer report against the model's value when the sample left the tap, $b \leftarrow b + 0.2(e - b)$, so the comparison is made at the instant the sample describes and the delay does not appear as a model error.

When the fuel valve has stood at 99.5 % or more for 30 s, the MPC holds the reboiler setpoint where it is, adds the thermal path of circulation through the limited duty to its model, and caps the circulation at what the burner can heat to within 1 K of the setpoint, approached at 1.5 L/min per

Table 7.3 The Nine Cases

#	Case, preset and events	Teaching point
1	Fixed 15 L/min, 190 °C, map view	the operating map
2	Fixed 24.3 L/min, gas +30 % at 30 min	a disturbance, a late analyzer
3	Ratio, gas +30 % at 30 and back at 120 min, inlet +2 K from 150 min	feedforward and trim
4	Ratio, 190 °C, regeneration view	the reboiler ceiling
5	Ratio, 190 °C, inlet 35 to 45 °C over 30–90 min	the floor
6	Ratio, 203 °C	stripping gas and its price
7	70 t/h, makeup off, drum 40 %, slug at 30 min	foam and inventory
8	Ratio, main pump trips at 30 min (standby 16 L/min)	purity without circulation
9	Challenge, 24 h, fixed 24.3 L/min, 200 °C	the whole problem

execution. If the program fails, the last inputs are held. The program is condensed and solved by the framework's shared dense quadratic program solver (`shared/src/framework/core/linmpc.js` and `qp.js`). In the challenge it solves 720 times with no fallback.

7.5.5 Three Models Kept Apart

The plant is the whole of Section 7.3. The identified model is whatever a learner fits to a step, such as the 2.08 min stage holdups behind a circulation step, the 29.7 min surge drum behind a purity step and the 112 min reboiler at a fixed duty. The controllers' model is the MPC's copy of the steady map with lagged inputs and commissioned parameters, corrected only by the analyzer bias, which sees neither the foam nor any parameter edited during the run. The burner and pump trips read the simulator's truth and are teaching guardrails, not certified protection. The exported CSV file pairs each instrument column with the true quantity behind it, and every true column carries the prefix `truth_`, so a controller replayed from the file is kept from reading it.

7.6 Nine Cases on the Running Engine

The nine cases appear in Table 7.3. Each begins from the defaults of Section 7.5, Ratio at 22 L/kg, the reboiler at 200 °C, no stripping gas, the makeup on Auto and the design gas, unless its row states otherwise. Every number below comes from the engine stepped headlessly at the default seed exactly as the browser steps it, with a move named with a time made at that minute of the run (Section 1.6.4 describes how a case is named and run). Opening a case first, selecting it after another case and restarting it with Reset all produce one trajectory. The measured outcomes are collected in Table 7.4.

7.6.1 The Map, the Gas Rate and the Ratio

Case 1 opens at 15 L/min, 13.6 L per kilogram of water, and 98.76 wt% glycol from a 190 °C reboiler, with the gas at 9.67 lb/MMscf. Doubling the circulation takes off 2.84 of those lb/MMscf, and the next 20 L/min only 1.04 more. Raising the reboiler to 204 °C at 15 L/min reaches 8.83 lb/MMscf, because too little glycol carries the water away whatever its purity. The case runs in Auto with the circulation on Fixed for a reason the learner meets in Manual. There the fuel valve is held too, and 30 L/min on the pump cools the reboiler from 190 to 154 °C and leaves 12.2 lb/MMscf, worse than the start.

In case 2 the gas rises 30 % at 30 min and the water load from 66.2 to 86.1 kg/h, so the fixed 24.3 L/min falls from 22 to 16.9 L/kg. The true water passes 7 lb/MMscf one minute later and settles near 7.82 lb/MMscf, while the analyzer, reading gas that left 10 min earlier, reports OFF SPEC at 43 min. About 29.5 L/min is the least circulation that restores the specification.

Table 7.4 Measured Outcomes by Case, Default Seed (Water in lb/MMscf)

#	Setting	Outcome
1	as shipped	9.67, dew point 4.0 °C
	30 or 50 L/min at 10 min	under 7 at 17 min, 6.83 or 5.79 at 240 min
	204 °C at 15 or 25 L/min	8.83 or 6.43
	Manual, pump 30 L/min	reboiler 154 °C, 12.2
2	as shipped	over 7 at 31 min, peak 7.87, OFF SPEC at 43 min
	28, 29.5, 30 L/min at 60 min	7.20, 6.99, 6.92 at 240 min
3	Ratio through the gas step	pump 24.4 to 31.4 L/min, peak 6.86
	warm-up of 2 K	6.68 to 7.53, analyzer over 7 at 175 min
	Ratio + trim at 190 min	analyzer under 7 at 243 min, 6.61 at 360 min
4	204 °C at 10 min	202.7 °C, water under 7 at 42 min, 6.54, 99.10 wt%
5	as shipped	13.4 at 120 min
	ratio 30 L/kg at 100 min	pump at 50 L/min, 11.95
	then 204 °C, then 50 Sm^3/h	10.12, then 6.69 at 99.79 wt%, 741 kW
6	0, 5, 10, 20, 45 Sm^3/h	6.54, 5.87, 5.35, 4.75, 4.27 at 154, 201, 249, 347, 598 kW
7	as shipped	foam index 2.57, ΔP 38 to 185 mbar, 22.3, 219 kg/h carried over
	left alone	133 kg lost by 120 min, level under 30 % at 110 min
	antifoam, makeup at 45 min	under 7 at 53 min, 52 kg lost, drum back at 90 min
8	as shipped	pump 24.4 to 15.9 L/min, 8.62 at 240 min
	204 °C at 45 min	8.45, 128 kW
	204 °C, 15 Sm^3/h at 45 min	under 7 at 139 min, 6.73, 275 kW

TRY IT IN THE APP **2 · Gas flow +30 % at fixed circulation**

Load the case and press Run. Watch the two lines of the Water in dry gas trend after the gas steps up at 30 min. The true water passes 7 lb/MMscf one minute later, but the analyzer, reading gas that left the contactor 10 min earlier, reports OFF SPEC only at 43 min.

Case 3 meets the same step with the ratio, which moves the pump from 24.4 to 31.4 L/min as soon as the meter sees the gas, and the water peaks at 6.86 lb/MMscf. The warm-up is different. At 37 °C the gas brings 11 % more water and the ratio adds 11 % more glycol, which keeps the share removed but lets the water left grow in proportion, and the warmer contactor raises the equilibrium as well. The trim switched on at 190 min raises the ratio to 29.4 L/kg and brings the analyzer back under 7 lb/MMscf at 243 min.

Pause. Try Exercise 7.2 before reading on.

7.6.2 Purity, the Floor and the Ceiling

In case 4 a reboiler step to 204 °C reaches 202.5 °C by 23 min and is held at 202.7 °C by the override. The water crosses under the specification at 42 min and settles at 6.54 lb/MMscf with 99.10 wt% glycol for about 10 kW more fuel. Case 5 warms the inlet to 45 °C, which raises the load from 66.2 to 109.6 kg/h and the water to 13.4 lb/MMscf by 120 min. A ratio of 30 L/kg takes the pump to its 50 L/min and the water only to 11.95 lb/MMscf, against the floor of Section 7.4. The state at 360 min is shown in Figure 7.3. Purity moves the floor, 204 °C to 10.12 lb/MMscf and 50 Sm^3/h of stripping gas to 6.69 lb/MMscf, at 741 kW. In case 6 the first 5 Sm^3/h of stripping gas buys 0.67 lb/MMscf of water for 47 kW, the last 25 Sm^3/h buy 0.48 lb/MMscf for 251 kW, and at 45 Sm^3/h the stripping gas alone costs 450 kW, 3.1 times the reboiler's fuel. The dew point at 60 bar falls from −2.0 to −8.3 °C.

7.6.3 Foam and a Lost Pump

Case 7 runs at 70 t/h with the makeup pump off. The hydrocarbon slug at 30 min takes the foam index to 2.57 within 9 min, the contactor pressure drop from 38 to 185 mbar and the stage efficiency to 0.46, so the water jumps to 22.3 lb/MMscf while 219 kg/h of glycol goes over the top. Left alone, the foam lasts until 140 min, 133 kg of TEG is lost by 120 min and the drum passes its 30 % alarm at 110 min. Antifoam and makeup at 45 min clear the foam and the specification by 53 min, after 52 kg has gone.

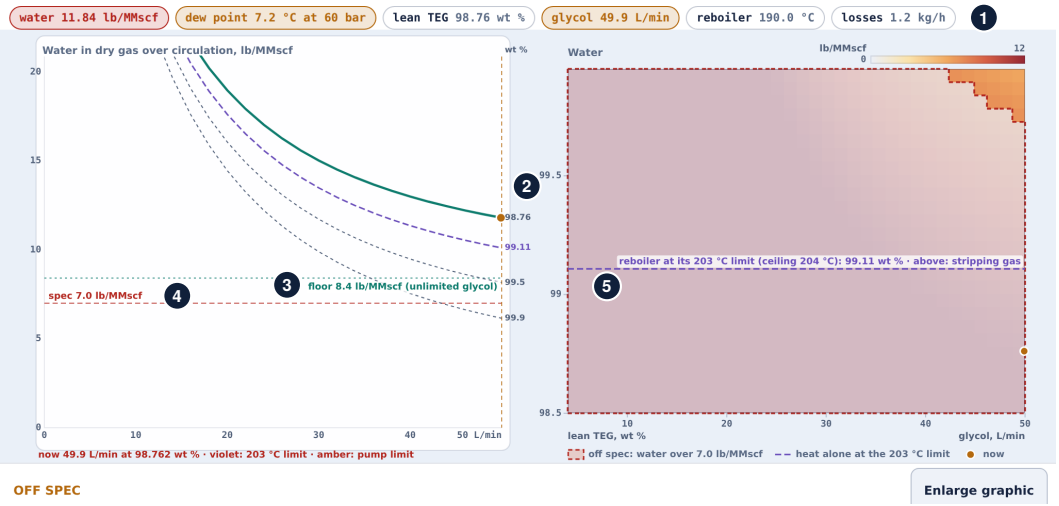

Figure 7.3 The equilibrium floor in case 5 at $t = 360$ min, Dew-point map view cropped to the map, with the Circulation ratio raised to 30 L/kg before the run (the case ships with 22 L/kg). The inlet gas has warmed to 45 °C and carries 109.2 kg/h of water, and the pump is at its limit, 49.9 L/min or 27.5 L/kg. (1) The status chips: the reboiler is back at 190.0 °C, the lean glycol is at 98.76 wt%, and the dry gas carries 11.81 lb/MMscf (analyzer 11.84 lb/MMscf, dew point 7.2 °C at 60 bar). The status strip under the map reads OFF SPEC. (2) The bold solid curve ends at the pump line (amber, dashed). (3) The floor at 8.4 lb/MMscf (teal, dotted) lies above (4) the 7.0 lb/MMscf specification (red, dashed). The thinner dashed curves for 99.11 (violet, long dashes), 99.5 and 99.9 wt% lie below the bold one, and only the last crosses the specification within the pump's range. (5) On the right panel the map over circulation and purity is off specification except at high circulation with purities well above the long-dashed violet line, the 99.11 wt% that heat alone reaches at the 203 °C limit.

> **TRY IT IN THE APP** **7 · Foaming, carryover and the glycol inventory**
>
> Press Run and watch the contactor's pressure drop after the hydrocarbon slug at 30 min. It climbs from 38 to 185 mbar as the foam builds, and the water in the dry gas jumps to 22.3 lb/MMscf. At 45 min turn Antifoam injection On and set the Makeup pump to Auto. The water is back under 7 lb/MMscf at 53 min, after 52 kg of glycol has gone.

The main pump trips at 30 min in case 8 and the standby delivers at most 16 L/min, two thirds of design, 14.5 L/kg. The plan for the application called for half circulation, but at 12 L/min no purity holds the specification, the least being 8.20 lb/MMscf at 204 °C and 50 Sm^3/h, so the lesson of holding the water by purity would fail. On 16 L/min the reboiler at 204 °C alone settles near 8.45 lb/MMscf. With 15 Sm^3/h of stripping gas as well, at 99.78 wt%, the water returns under 7 lb/MMscf at 139 min, 94 min after the moves, and reaches 6.73 lb/MMscf at 275 kW against 128 kW.

7.6.4 A Day Under an Energy Budget

The challenge runs 24 h. The gas ramps from 60 to 80 t/h between 6 and 10 h, the inlet warms to 38 °C between 12 and 13.5 h and cools to 35 °C by 16 h, and then the reboiler's fuel is limited to 55 % of the burner. The target is the true water of every analyzer sample at or under 7 lb/MMscf, the reboiler never above 204 °C, no trip and at most 250 kW of energy on average. Two features depart from the plan for the application, each for a measured reason. The plan asked for a hot afternoon 10 K above the inlet, which no setting can meet, since with every handle at its limit all day the worst

Table 7.5 The Challenge Over 24 h. Worst Analyzer Sample in lb/MMscf, Samples Over 7, Reboiler Peak and Average Energy

Settings	Worst	Over	°C	kW	Verdict
shipped (fixed 24.3 L/min, 200 °C)	10.37	506	200.4	165	fails
Ratio	8.04	321	201.0	180	fails
Ratio + trim	10.48	242	202.0	188	fails
Ratio + trim, 204 °C	10.47	195	203.2	187	fails
and 10 Sm^3/h all day	6.94	0	203.2	263	fails (energy)
and 15 Sm^3/h from 16 h	7.04	5	203.2	230	fails
and 8 from 12, 12 Sm^3/h from 16 h	6.81	0	203.2	229	passes
every handle at its limit	4.56	0	203.1	711	fails (energy)
suggested (MPC)	6.73	0	203.2	232	passes

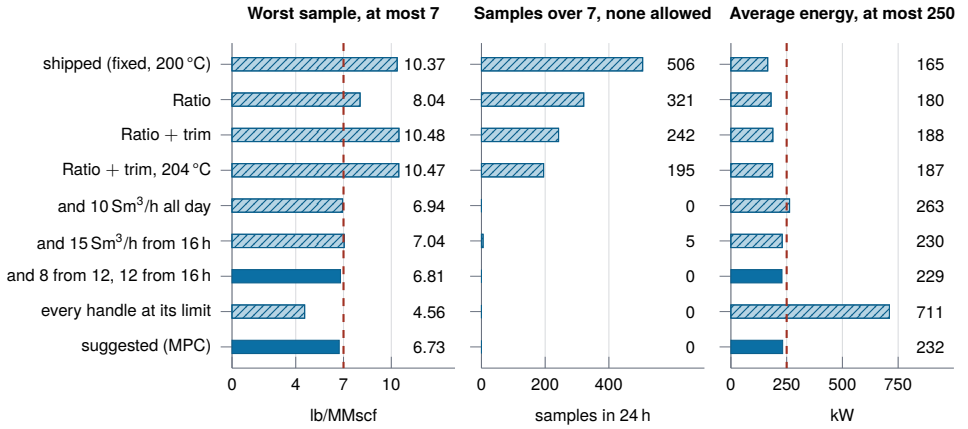

Figure 7.4 The challenge runs of Table 7.5 as bars: the worst analyzer sample, the samples over 7 lb/MMscf and the average energy over 24 h. Solid bars are the settings that pass, hatched bars those that fail, dashed red lines are the limits of the target, and the value of each bar stands at the right. Notice that the two settings that pass add stripping gas only where it is needed, while every handle at its limit holds the water best and fails on energy.

sample is 8.57 lb/MMscf with 49 samples over. It also asked for minimum energy without a number, and without a budget brute force passes, since every handle at its limit holds every sample at or under 4.56 lb/MMscf at 711 kW. The variants appear in Table 7.5. At the top playback speed of 300 min per second the 24 h challenge plays in about 5 s.

Figure 7.4 sets the nine settings of Table 7.5 against the two limits of the target that carry a number and the count of samples over the specification.

Fixed circulation fails as the gas ramps, from 428 min, and the ratio alone holds the morning but not the warm afternoon, from 750 min, because it keeps the share removed rather than the water left. Ratio and trim at 204 °C miss the afternoon by 0.04 lb/MMscf with the pump at its limit and then fail badly under the fuel limit, where the worst sample reaches 10.47 lb/MMscf. The trim answers the wet analyzer with more glycol that the starved reboiler cannot heat, the purity falls and the trim asks for more. Stripping gas added where nothing cheaper works, 8 Sm^3/h from 12 h and 12 from 16 h, passes at 229 kW, while 10 Sm^3/h left open all day meets the water target but spends 263 kW, over the budget. The MPC finds a plan of that kind by itself. It holds the reboiler setpoint at its 203 °C limit and raises the circulation from a mean of 28.9 L/min in the first six hours to 38.6 after the gas ramp. Before 16 h it uses stripping gas only in the hot afternoon, with the pump at or near its limit, apart from up to 6 Sm^3/h in the first 10 min while it raises the reboiler setpoint. Under the fuel limit it trades circulation down to a mean of 31.3 L/min for up to 27.1 Sm^3/h of stripping gas while the

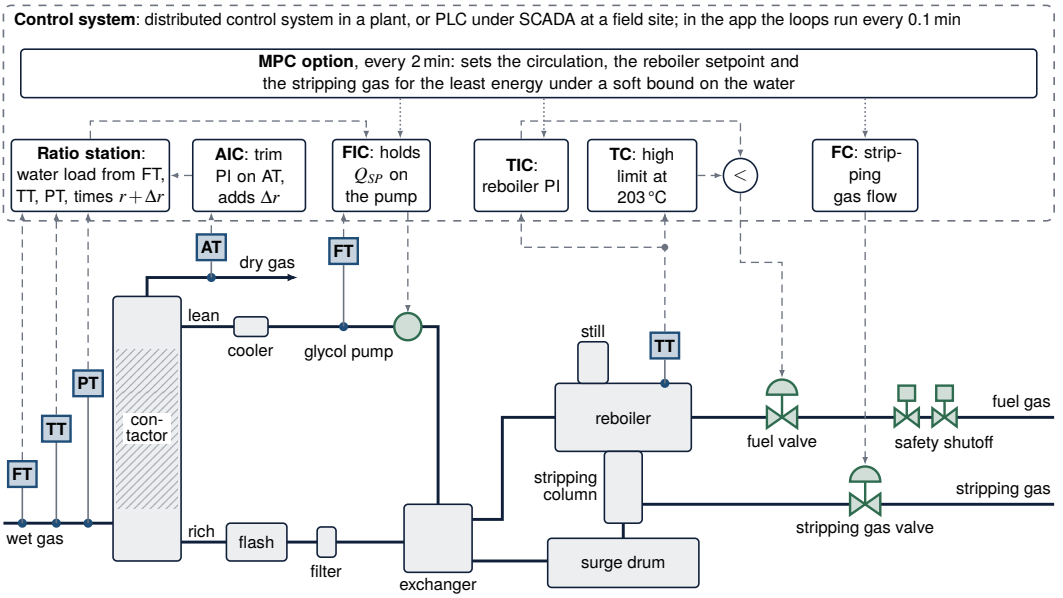

Figure 7.5 The loops of the dehydrator simulator with their hardware. Lettered boxes are sensing devices with the chapter's tags: FT, TT and PT on the wet gas, AT on the dry gas, FT on the lean glycol and TT on the reboiler. Green symbols are final elements: the glycol pump, the fuel valve with the two safety shutoff valves of the fuel train, and the stripping gas valve. Dashed arrows are signals, dotted arrows the setpoints the optional MPC writes. Notice that only AT reads the water the specification is about, and that the reboiler's two controllers meet at a low select before the fuel valve. The makeup pump, its level switch and the antifoam are left out. Not to scale, and not a piping diagram.

reboiler falls to 190.0 °C and takes 3.75 h to recover. Every sample stays at or under 6.73 lb/MMscf, with 232 kW on average, 188 kW of it fuel. At seeds 1 and 2 the MPC passes the same way, at 231.6 and 231.9 kW, and the shipped settings fail from 428 min. The budget makes the challenge an economic problem, and the hand-timed plan that passes uses 2.5 kW less than the MPC because its moves were placed with the day's schedule in view. A data hall whose cheapest operating point sits on a constraint that moves with the weather poses the same economic problem (Volume 9, Chapter 8).

> ✦ DID YOU KNOW? **Feedback That Raised the Worst Case**
>
> From the simulator (Table 7.5). On the graded day, adding the analyzer trim to the ratio cut the samples over the specification from 321 to 242 but raised the worst sample from 8.04 to 10.48 lb/MMscf. The trim lowered the count and deepened the worst miss.

7.7 Technology in Practice

Figure 7.5 draws the loops of Section 7.5 with the hardware of a dehydrator in place of the schematic's symbols.

What Measures It

One instrument reads the quantity the specification is about: the moisture analyzer AT on the dry gas. Many analyzers stand some distance from the process and receive a sample through a system of tubes, heaters, filters, pumps and regulators, and problems with process analyzers quite often stem from sample systems that are badly built or maintained. A plugged sample line leaves an analyzer

reading stale sample with no way of knowing it [20]. The model gives AT a sample every 2 min reported 10 min after it left the gas, 8 min of that in the sample line. The feedforward reads ordinary transmitters: the gas meter with a 6 s lag, the inlet temperature with 30 s and the reboiler thermocouple with 18 s.

> **TECHNOLOGY Sensor: The Moisture Analyzer**
>
> Hygrometers for natural gas work on several principles: a chilled mirror that finds the dew point, an electrolytic cell, a laser spectrometer, and polymer or aluminum oxide sensors. Condensation and spectroscopic instruments work only at low pressure, so they are fed through pressure-reducing sample lines. In a 2015 comparison in natural gas at up to 6 MPa, the instruments saw a step in humidity almost at once, the electrolytic analyzer needed up to two hours to settle, and every instrument was slower when the gas dried than when it wetted. Most aluminum oxide probes drifted downward, reading drier as time went on [4], which on a trend looks like a unit slowly improving by itself. The chapter's AT carries 0.05 lb/MMscf of noise and no drift.

What Moves It

Three final elements carry the two handles. The glycol pump sets the circulation, the fast handle, through a 0.2 min lag up to 50 L/min, or 16 L/min on the standby pump. The fuel valve sets the reboiler duty, the slow handle, and the stripping gas valve adds the expensive end of the purity handle up to 50 Sm^3/h.

> **TECHNOLOGY Final Element: The Reboiler's Fuel Train**
>
> The fuel valve u_f is the modulating valve of a fuel gas train. Upstream of it, two safety shutoff valves with a vent between them, a double block and bleed, isolate the fuel when a burner management system sees low or high fuel pressure or a lost flame. The shutoff valves are electrically actuated and close by spring when power is lost, and an ultraviolet flame sensor is used because it cannot be fooled by the glow of a hot firebox. The same system purges the firebox and sequences the ignition at start-up [20]. In the chapter the fuel valve acts through a 1 min burner lag, the override reaches it through the low select, and the burner trip at 210 °C stands in for the burner management system as a teaching guardrail.

Where the Controller Runs

Inside a gas plant the loops would run on a distributed control system, which controls production within one location through feedback and feedforward loops under a supervisory level. At a remote field site a programmable logic controller (PLC) or a remote terminal unit controls the local process and reports to a supervisory control and data acquisition (SCADA) center. Such loops run with cycle times from milliseconds to minutes [21]. The app's loops execute every 0.1 min and its MPC every 2 min, both far shorter than the analyzer's 10 min, so no scan rate repairs the delay. The chapter's structure maps onto standard blocks: a PID for the reboiler and a proportional high limit with external reset meeting at a low selector, a ratio station that multiplies the computed water load, a PID trim on the analyzer that adds to the ratio, and, as an option, the MPC in a supervisory layer that writes the circulation, the reboiler setpoint and the stripping gas.

> **IN PRACTICE A Rule of Thumb Written Into Law**
>
> Under the United States emission standards for oil and gas production, a glycol dehydrator at an area source outside urban areas must compute its optimum circulation from the gas flow, the water it removes and the industry rule of thumb of 3.0 gallons of glycol per pound of water, times 1.15 for a margin of safety, and must not circulate more. A unit that cannot meet its sales gas moisture specification at that rate calculates and documents an alternate rate [22]. The rule belongs to the standards for hazardous air pollutants: the glycol absorbs other constituents of the gas along with the water, and the unit's process vents are the reboiler vent and the flash tank vent [23].

What Is Modernizing It

Two developments touch this chapter. Laser absorption analyzers shine a tunable laser through the sample and read the gas from the light its molecules absorb [20]. In the 2015 comparison of hygrometers in natural gas the spectroscopic and electrolytic analyzers showed no appreciable drift, where most aluminum oxide probes drifted [4]. Regeneration is being reworked for energy: in a simulated plant, routing the flash tank's gas to the regenerator as stripping gas instead of to the flare allowed less circulation and cut the reboiler duty by 36.2 % [24]. The economic MPC of Section 7.5 prices the same handle in kilowatts of fuel and opens it only where nothing cheaper is left.

7.8 Limitations

The unit is generic, and none of its sizes or correlations was fitted to a plant. The contactor is two well-mixed stages, so the circulation that meets a specification here is a teaching number, not a design. For 7 lb/MMscf from gas at 90 °F and 500 psia over 98.8 wt% glycol the model needs 27.9 L/kg with perfect stages and 79 L/kg with its efficiency of 0.9, where a published simulation of two equilibrium stages with the glycol at the gas temperature needs about 3 gal per pound of water, or 25 L/kg [10]. The reboiler is 0.34 to 0.48 wt% purer than the same simulation, a fit kept because it meets the targets of the plan for the application while a refit would move every case, and the equilibrium is about 3 K drier than the chart at 99.99 wt%. The saturated water content ignores acid gas and gravity, the dew point below 0 °C is that of supercooled water, and the hydrate point of the same gas lies higher. Hydrocarbon absorption, flash gas and the aromatic emissions of the still are absent [17].

The contactor level and the flash tank pressure have no loops, the high-limit override passes the ceiling by 0.21 K on a sudden loss of load, and nothing protects the ceiling in Manual. The challenge departs from the plan for the application in two measured ways, a milder afternoon and an energy budget, because the original afternoon is infeasible and the original objective is passed by brute force, and the pump trip leaves two thirds of the circulation for the same kind of reason. The MPC's model is the plant's own steady map with first-order dynamics, so it meets model error only through foam, a changed parameter and the analyzer, and its bias learns slowly through a 10 min delay. Disturbances arrive as scheduled steps, ramps and pulses.

> **OPEN CHALLENGE** **A Moisture Reading Fit for the Loop**
>
> Moisture instruments for natural gas are usually calibrated at atmospheric pressure in nitrogen or air, far from the conditions of the process. Turning a water content into a dew point depends on the composition, temperature and pressure of the gas, and the extrapolations can move the dew point by up to 10 °C where 1 to 3 °C is wanted [4]. A 2020 study still saw a gap between the analyzers available and what industry asks for: a response fast enough for online use in natural gas, valid across gas mixtures with a single calibration [25]. Progress would be a reading at line pressure fast and trusted enough to close the dew point loop on it.

7.9 Conclusion

This chapter studied how a glycol dehydrator meets one water specification with a handle that is fast but bounded by equilibrium and a handle that is slow, capped and expensive. The simulator computes both from a seventeen-state model whose water and glycol ledgers close to round-off. Circulation lowers the water from 9.67 to 6.83 lb/MMscf when it doubles, but on a 45 °C inlet with 98.76 wt% glycol no circulation reaches 7 lb/MMscf, the least being 8.41 lb/MMscf, so only purity can meet it. Heat buys purity up to 99.11 wt% under the override, and stripping gas the rest at 10 kW per Sm^3/h. A ratio on the water load holds a gas step at a peak of 6.86 lb/MMscf but lets a warmer inlet through, a trim recovers it slowly through a 10 min analyzer, and under a fuel limit the trim makes the gas wetter. An economic MPC holds every sample of a graded day at or under 6.73 lb/MMscf at 232 kW on average. The correlations are teaching values, the reboiler is about 0.4 wt% purer than published simulations and two plant loops are left out. A fuller model would add a packed-bed contactor, the contactor level and flash pressure loops, aromatic emissions and a reboiler refitted to measured

purities. The simulator shows that a specification can lie below what any amount of the fast handle can reach.

KEY POINTS

- Pumping more glycol dries the gas only toward a floor set by how dry the glycol is, so on a 45 °C inlet with 98.76 wt% glycol no circulation reaches 7 lb/MMscf.
- Heat buys purity up to 99.11 wt% under the override, and stripping gas the rest at 10 kW per Sm^3/h.
- A ratio on the water load holds a gas step at a peak of 6.86 lb/MMscf but lets a warmer inlet through, and a trim recovers it slowly through a 10 min analyzer.
- Under a fuel limit the trim makes the gas wetter, while an economic MPC holds every sample of the graded day at or under 6.73 lb/MMscf.

Where else this applies. Chapter 8 sweetens the same gas with a solvent whose regenerated purity bounds the product, and Volume 9, Chapter 8 prices its inputs in an economic predictive controller on a constraint that moves with the weather.

References

[1] D. B. Carson and D. L. Katz, “Natural gas hydrates,” *Transactions of the AIME*, vol. 146, no. 1, pp. 150–158, 1942.

[2] J. A. Porter and L. S. Reid, “Vapor-liquid equilibrium data on the system natural gas-water-triethylene glycol at various temperatures and pressures,” *Journal of Petroleum Technology*, vol. 2, no. 8, pp. 235–240, 1950.

[3] A. Rosman, “Water equilibrium in the dehydration of natural gas with triethylene glycol,” *Society of Petroleum Engineers Journal*, vol. 13, no. 5, pp. 297–306, 1973.

[4] J. G. Gallegos, R. Benyon, S. Avila, A. Benito, R. M. Gavioso, H. Mitter, S. Bell, M. Stevens, N. Böse, V. Ebert, M. Heinonen, H. Sairanen, A. Peruzzi, R. Bosma, and M. Val’ková, “An investigation of the comparative performance of diverse humidity sensing techniques in natural gas,” *Journal of Natural Gas Science and Engineering*, vol. 23, pp. 407–416, 2015.

[5] Z. Y. Kong, A. Mahmoud, S. Liu, and J. Sunarso, “Revamping existing glycol technologies in natural gas dehydration to improve the purity and absorption efficiency: Available methods and recent developments,” *Journal of Natural Gas Science and Engineering*, vol. 56, pp. 486–503, 2018.

[6] 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.

[7] 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.

[8] 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.

[9] F. R. Scauzillo, “Equilibrium ratios of water in the water-triethylene glycol-natural gas system,” *Journal of Petroleum Technology*, vol. 13, no. 7, pp. 697–702, 1961.

[10] V. N. Hernandez-Valencia, M. W. Hlavinka, and J. A. Bullin, “Design glycol units for maximum efficiency,” in *Proceedings of the Seventy-First GPA Annual Convention*. Tulsa, OK: Gas Processors Association, 1992, pp. 310–317.

[11] R. Chebbi, M. Qasim, and N. Abdel Jabbar, “Optimization of triethylene glycol dehydration of natural gas,” *Energy Reports*, vol. 5, pp. 723–732, 2019.

[12] A. Wahid, R. Mauricio, and N. S. Maro, “Control of gas dehydration unit using multivariable model predictive control (MMPC) to obtain more optimal control performance,” *E3S Web of Conferences*, vol. 67, p. 03013, 2018.

[13] E. G. Hammerschmidt, “Formation of gas hydrates in natural gas transmission lines,” *Industrial & Engineering Chemistry*, vol. 26, no. 8, pp. 851–855, 1934.

[14] E. D. Sloan, “Hydrocarbon hydrate flow assurance history as a guide to a conceptual model,” *Molecules*, vol. 26, no. 15, p. 4476, 2021.

[15] L. R. Records and D. H. Seely, “Low temperature dehydration of natural gas,” *Journal of Petroleum Technology*, vol. 3, no. 2, pp. 61–66, 1951.

[16] O. W. Fowler, “Calcium-chloride wellhead gas dehydrators,” in *Gas Technology Symposium*. Society of Petroleum Engineers, 1961, paper SPE-22-MS.

[17] P. L. Grizzle, “Hydrocarbon emission estimates and controls for natural gas glycol dehydration units,” in *SPE/EPA Exploration and Production Environmental Conference*. Society of Petroleum Engineers, 1993, paper SPE-25950-MS.

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

[19] A. Bahadori, H. B. Vuthaluru, and S. Mokhatab, “Rapid estimation of water content of sour natural gases,” *Journal of the Japan Petroleum Institute*, vol. 52, no. 5, pp. 270–274, 2009.

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

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

[22] U.S. Environmental Protection Agency, “40 CFR 63.764, general standards, national emission standards for hazardous air pollutants from oil and natural gas production facilities,” Code of Federal Regulations, Title 40, Part 63, Subpart HH, edition of July 1, 2025, govinfo.gov, 2025.

[23] ——, “40 CFR 63.761, definitions, national emission standards for hazardous air pollutants from oil and natural gas production facilities,” Code of Federal Regulations, Title 40, Part 63, Subpart HH, edition of July 1, 2025, govinfo.gov, 2025.

[24] S. A. Affandy, A. Kurniawan, R. Handogo, J. P. Sutikno, and I.-L. Chien, “Technical and economic evaluation of triethylene glycol regeneration process using flash gas as stripping gas in a domestic natural gas dehydration unit,” *Engineering Reports*, vol. 2, no. 5, p. e12153, 2020.

[25] N. Takeda, P. Carroll, Y. Tsukahara, S. Beardmore, S. Bell, K. Yamanaka, and S. Akao, “Trace moisture measurement in natural gas mixtures with a single calibration for nitrogen background gas,” *Measurement Science and Technology*, vol. 31, no. 10, p. 104007, 2020.

7.10 Exercises

7.1 Concept check · 2 min

In case 1 the reboiler holds 190 °C, so the lean glycol stays at 98.76 wt%. Doubling the circulation from 15 to 30 L/min takes the water in the dry gas from 9.67 to 6.83 lb/MMscf, far from half. Which explanation is right?

(a) The outlet approaches a floor set by equilibrium with the lean glycol and by what two stages of efficiency 0.9 let through

(b) The stage holdups of about two minutes delay the response, so the water has not settled when it is read

(c) The extra glycol cools the reboiler, so the lean glycol gets wetter and cancels part of the gain

(d) The analyzer reports gas that left the contactor ten minutes earlier, so it hides part of the fall

7.2 Concept check · 1 min

Case 3 runs on Ratio. From 150 min the inlet gas warms from 35 to 37 °C and brings about 11 % more water, and the ratio raises the circulation by the same 11 %. What does the water in the dry gas do?

(a) It stays near 6.68 lb/MMscf, because the glycol per kilogram of water is unchanged

(b) It rises, because a fixed ratio keeps the share removed and the warmer contactor raises the equilibrium

(c) It falls, because the extra glycol cools the contactor below the temperature of the incoming gas

(d) It rises on the analyzer only, because the inlet temperature transmitter lags the true inlet

7.3 Concept check · 2 min

The challenge is run on Ratio + trim with the reboiler setpoint at 204 °C. At 16 h the reboiler fuel is limited to 55 % of the burner, the analyzer reads wet and the trim raises the circulation ratio. What follows?

(a) The water returns under the specification within one analyzer delay, because circulation is the fast handle

(b) The high-limit override takes the fuel valve and holds the reboiler at 203 °C until the limit clears

(c) The gas gets wetter, because the burner cannot heat the extra glycol and the purity falls, so the trim asks for more

(d) Nothing changes, because the trim stops at the ratio that the fuel-limited burner is able to heat

7.4 Theory and calculation · 7 min

The reboiler sits at its 203 °C override and makes 99.11 wt% glycol. At 24.3 L/min the equilibrium ratio $\gamma_w p_{sat}(T_r)/P_{st}$ in Equation (7.5) is 13.45. Use $V_m = 23.645\,\mathrm{Sm^3/kmol}$, $\rho_h = 1010\,\mathrm{kg/m^3}$, $M_T = 150.17$ and $N_s = 2$.

(a) Find the stripping gas per cubic meter of glycol g and the stripping factor S for 20 $\mathrm{Sm^3/h}$. *Checkpoint:* $g = 13.72\,\mathrm{Sm^3/m^3}$ and $S = 1.16$.

(b) Find the share of the reboiler's water that stays in the glycol leaving the column, the water ratio X_{lean} and the purity leaving the column. *Checkpoint:* 0.285, 0.002 56 kg/kg and 99.74 wt%.

(c) Repeat for 5 $\mathrm{Sm^3/h}$ and 45 $\mathrm{Sm^3/h}$, and compare the purity bought per $\mathrm{Sm^3/h}$ by the first 5 $\mathrm{Sm^3/h}$ with that bought by the step from 20 to 45 $\mathrm{Sm^3/h}$. *Checkpoint:* 99.35 and 99.91 wt%, about 7 times less per $\mathrm{Sm^3/h}$ for the last step.

7.5 Theory and calculation · 7 min

At design the inlet gas carries $W_{in} = 58.18\,\mathrm{lb/MMscf}$, the gas leaves stage 2 with 17.89 lb/MMscf and stage 1 with 6.73 lb/MMscf, the stage efficiency is $E = 0.9$, and the gas in equilibrium with the lean glycol would carry $W^* = 3.05\,\mathrm{lb/MMscf}$.

(a) Use the gas-side relations of Equation (7.3) to find the water that gas in equilibrium with each stage's glycol would carry, and the share of the inlet water absorbed. *Checkpoint:* $W_{eq}(X_2) = 13.41$ and $W_{eq}(X_1) = 5.49\,\mathrm{lb/MMscf}$, 88.4 % absorbed.

(b) Estimate the floor with unlimited circulation at the design contactor temperature, and explain why the least water any circulation reaches in the model, 3.94 lb/MMscf near 390 L/min, is higher. *Checkpoint:* 3.60 lb/MMscf.

(c) In case 5 at 50 L/min the inlet is at 45 °C ($W_{in} = 96.32\,\mathrm{lb/MMscf}$) and $W^* = 6.83\,\mathrm{lb/MMscf}$. Estimate the floor the same way and say whether a bigger pump can meet 7 lb/MMscf. *Checkpoint:* 7.72 lb/MMscf, above the specification.

7.6 Theory and calculation · 7 min

Use the reboiler balance of Equation (7.6) at the design point: $T_r = 200\,^\circ\mathrm{C}$, $T_c = 36.09\,^\circ\mathrm{C}$, $\varepsilon = 0.70$, 58.6 kg/h of water boiled off with $\lambda = 2300\,\mathrm{kJ/kg}$ and $R = 0.25$, $UA = 0.1\,\mathrm{kW/K}$, $T_{amb} = 25\,^\circ\mathrm{C}$, a firing efficiency of 0.75 and 46.3 kW to heat the rich glycol at 24.28 L/min.

(a) Find the rich glycol's temperature entering the reboiler, the heat to boil the water and the loss, and from them the absorbed duty and the fuel burned. *Checkpoint:* 159.8 °C, 46.8 kW and 17.5 kW, 110.6 kW absorbed and 147.5 kW burned.

(b) Find the fuel that heating the rich glycol costs per L/min of circulation (the sensible heat only) and add the pump's 0.164 kW per L/min. Compare the sum with the cost the MPC places on circulation. *Checkpoint:* About 2.7 kW per L/min.

(c) In the challenge the burner is limited to 55 % of 250 kW while the gas runs at 80 t/h. Assume the boil-off grows in proportion to the gas, the loss is unchanged and each L/min still takes the design sensible heat. Estimate the most circulation the burner can hold at 200 °C. *Checkpoint:* About 30 L/min.

7.7 App challenge · 8 min

Run case 1, *The operating map: set circulation and purity yourself*, at the default seed. It ships with the Circulation setpoint at 15 L/min and the Reboiler temperature setpoint at 190 °C. For each setting below, reload the case, make the setting before Run, run 60 min at speed 20 (3 s), and read the column `truth_water_lbMMscf` of the exported CSV at 60 min. The Water (analyzer) tile carries the analyzer's noise and may differ by up to 0.06 lb/MMscf.

(a) Circulation setpoint 20, 30 and 50 L/min. Give the water for each, and the water removed per extra L/min over each interval, starting from 9.67 lb/MMscf at 15. *Checkpoint:* 8.24, 6.83 and 5.79 lb/MMscf (each within 0.03).

(b) Reboiler temperature setpoint 204 °C with the circulation left at 15 L/min. Give the water, the Lean TEG tile and the Reboiler tile, and say why purity alone does not meet 7 lb/MMscf here. *Checkpoint:* 8.81 lb/MMscf, 99.107 wt%, 202.9 °C.

(c) Circulation setpoint 25 L/min and reboiler 204 °C together. Give the water. *Checkpoint:* 6.43 lb/MMscf (within 0.03).

7.8 App challenge · 8 min

Run case 3, *The ratio feedforward on the water load*, at the default seed. Dew point control ships on Ratio at 22 L/kg.

(a) Run 360 min as shipped at speed 60 (6 s) and export the CSV. Read the circulation, column `circulation_Lmin`, at 29 and 60 min. Then read the column `truth_water_lbMMscf`: its largest value between 30 and 120 min, and its values at 150 and 360 min. *Checkpoint:* About 24.2 and 31.4 L/min (within 0.3), a peak of 6.86 lb/MMscf, then 6.68 and 7.52 lb/MMscf (within 0.02).

(b) Reload the case, select Ratio + trim in Dew point control before Run, and run 360 min again. Read `truth_water_lbMMscf` at 150 min, its largest value between 150 and 360 min, and `circulation_ratio_Lkg` at 360 min. Explain why this run holds the warm-up when the run in which the trim is switched on at 190 min lets the analyzer stay above 7 lb/MMscf until 243 min. *Checkpoint:* 6.40 lb/MMscf, a peak of 6.97 lb/MMscf near 180 min, and 29.6 L/kg (within 0.1).

7.9 App challenge · 9 min

Run case 9, the Challenge, *a day on the dehydrator*, at the default seed. Each run is 24 h at speed 300 (about 5 s). The run stops at 1440 min and its verdict appears, and the So far tile shows the worst sample and the average energy.

(a) Run as shipped. Give the worst sample, the number of samples over 7 lb/MMscf and the average energy. *Checkpoint:* 10.37 lb/MMscf, 506 samples, 165 kW.

(b) Reload, select Ratio + trim and set the Reboiler temperature setpoint to 204 °C before Run. Give the same three values, and read `circulation_Lmin`, `reboiler_C` and `truth_lean_teg_wtpct` at 950 and 1440 min in the CSV. Explain what happened after the fuel limit at 960 min. *Checkpoint:* 10.47 lb/MMscf, 195 samples, 187 kW. From 39.6 to 49.9 L/min, from 202.8 to 161.0 °C, from 99.11 to 97.28 wt%.

(c) Reload and select MPC before Run. Give the verdict, and read the largest `stripping_gas_Sm3h` and the lowest `reboiler_C` after 960 min. *Checkpoint:* 6.73 lb/MMscf and 232 kW, passed. About 27.1 Sm^3/h and 190.0 °C (within 0.3).

7.10 Application area · 4 min

Case 7 runs at 70 t/h with the makeup pump off and the surge drum at 40 %. The drum holds 784.4 kg of TEG at 60 %, and its low alarm is at 30 %.

(a) How much TEG can the unit lose before the low alarm? *Checkpoint:* About 131 kg.

(b) At the peak carryover of 219 kg/h, how long would that take? The case passes the alarm about 80 min after the slug. Why longer? *Checkpoint:* About 36 min.

(c) What does an operator see first, and which two actions put it right?

7.11 Application area · 3 min

A sales contract asks for at most 7 lb/MMscf of water. The export line runs at 60 bar and its coldest point is at 4 °C.

(a) What water dew point at 60 bar does the model give for 7 lb/MMscf, and does gas on the specification drop free water in this line?

(b) What does the chapter say about hydrates, and what should be checked instead of the dew point?

(c) Why does the reboiler stop at 204 °C, and where does purity beyond the ceiling come from?

7.12 Review the agent · 5 min

An AI assistant reads case 5 with the inlet at 45 °C. As shipped, the pump runs at about 40 L/min and the water reaches 13.4 lb/MMscf. With the ratio raised to 30 L/kg the pump sits at its 50 L/min limit and the water falls to 11.95 lb/MMscf. The assistant writes: “The last 10 L/min removed 1.45 lb/MMscf. A second pump in parallel gives 100 L/min, five more such steps, so the water would fall by about 7.3 lb/MMscf to about 4.7 lb/MMscf, well under the specification. Buy the pump.” Find the flaw, and say how you would check the claim in the app and in the chapter.

CHAPTER 8

Selective Amine Sweetening With Delayed Analysis and Loading Protection

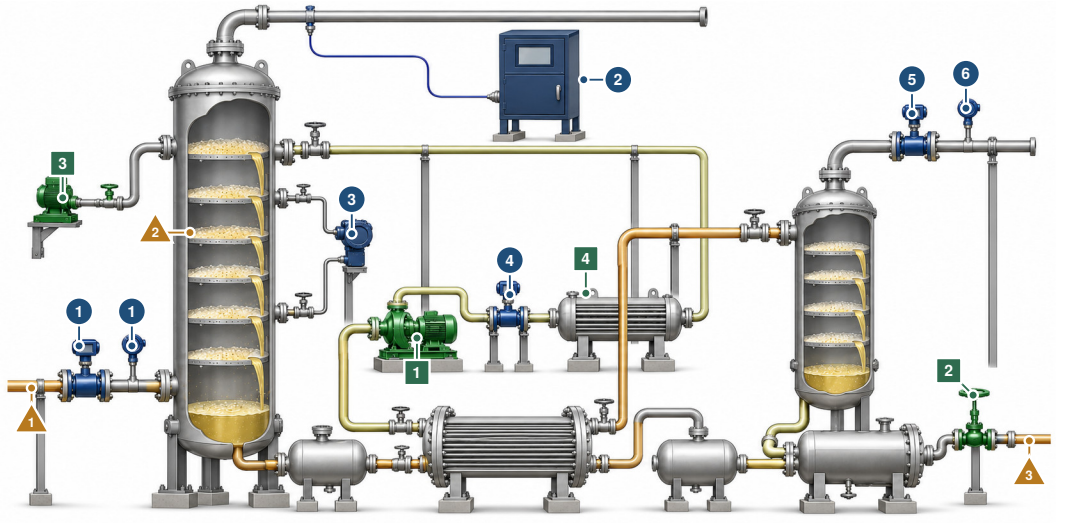

The controller moves 1 lean amine pump, the circulation 2 reboiler steam valve 3 antifoam injection 4 lean amine cooler

The controller reads 1 sour gas flow and temperature 2 hydrogen sulfide left in the sweet gas, minutes late 3 pressure drop across the absorber 4 lean amine flow 5 acid gas flow, which gives the loading 6 regenerator overhead temperature

What disturbs it 1 more gas, warmer gas or a sourer well 2 a hydrocarbon slug that makes foam 3 a dip in the steam supply

Notice the long sample line from the sweet gas to the analyzer cabinet, and the lean amine's road from the reboiler through the surge drum, the exchanger, the pump and the cooler back to the absorber's top. Illustration, not to scale.

IN PLAIN TERMS Washing a Sour Gas Out of Natural Gas

The main idea. Some natural gas carries hydrogen sulfide, a sour gas that buyers allow only in tiny traces. In a tall tower the gas rises through a falling liquid called amine that soaks it up. A column heated by steam then boils the sour gas out, so the liquid can be used again.

The challenge. The instrument that checks the cleaned gas reports each sample minutes after that gas has left. And the liquid can be loaded so full that it eats the plant's steel while the cleaned gas still passes its test.

Why it matters. A plant has to protect both the gas it sells and its own equipment, and each needs a different instrument.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? box at the end of the Introduction, the two History boxes in Section 8.2, the Try it boxes and the second Did you know? box in Section 8.5, and Key points at the end of the Conclusion. No equations.
Study	40 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 8.5 with Figure 8.3, Section 8.6 and Section 8.7 with its Open challenge box.

The control question A sweetener must keep hydrogen sulfide in the sales gas under 4 ppm, but its analyzer reports each sample minutes after the gas has gone, and a cap on the rich amine loading protects the plant where that analyzer cannot see. Which measurement can protect which limit, and what does each defense cost in steam?

Learning objectives

- Explain why hydrogen sulfide approaches stage equilibrium in MDEA while carbon dioxide uptake follows the liquid holdup, and predict what a holdup change does to each.
- Compute the ratio circulation and the loading floor from the gas meter, the laboratory composition and the acid-gas meter.
- Predict when a sampled, delayed analyzer first reports an exceedance, and bound that delay.
- Show from runs that the rich loading can pass its cap while the sweet gas stays on specification, and choose the instrument that protects each limit.
- Compare the defenses of the twelve-hour challenge by the limit each holds and the steam each costs.

Prerequisites Sections 1.4.7 (absorption), 1.4.8 (the high selector), 1.4.9 (late and inferred measurements), 1.6.2 (the framework's PI) and 1.6.4 (the apps). Section 1.6.3 is needed only for the predictive option. Chapter 7 treats the same structures on one solute.

The app https://apmonitor.com/apps/amine-sweetening/

See also Volume 5, Chapter 3 for composition control through a slow analyzer, and Volume 5, Chapter 10 for a quality and a constraint that are both inferred. Volume 4, Chapter 2 for a protective override that rests at its clamp in healthy operation, and Volume 4, Chapter 6 for a safety limit that bounds throughput. Volume 9, Chapter 10 for a slow quality trim on a setpoint loop, priced in energy.

8.1 Introduction

A gas sweetener can meet its product specification while damaging its solvent circuit. In the teaching unit considered here, an increase in feed hydrogen sulfide at 30 min takes the rich amine over 0.5 mol/mol at 38.5 min, yet the sweet gas stays below 3.43 ppm of H_2S, inside a 4 ppm limit. Conversely, cutting the pump from 64 to 50 m^3/h sends the true gas over 4 ppm at 13.8 min, while the analyzer first reports it at 24.7 min. A feedback loop cannot respond to a measurement before the sample exists. A loading limit therefore needs a different instrument and a different action from a sales-gas limit.

The distinction is especially useful for a tertiary amine. Reaction kinetics reported for tertiary amines put carbon dioxide on a slower pathway than hydrogen sulfide [1, 2], and selective absorption measurements with aqueous methyldiethanolamine (MDEA) show why residence and interfacial transfer matter [3, 4]. Equilibrium models of acid gas in aqueous amines provide a more detailed basis for real design [5, 6]. The simulator of this chapter compresses that chemistry to six liquid stages and a teaching correlation. It lets the reader test which change buys carbon dioxide pickup, which preserves the H_2S specification, and which merely consumes steam.

The control problem is one of unequal information and timescales. Gas flow is metered in seconds, the sweet-gas analyzer samples every 10 min, the rich loading inference filters an acid-gas meter for 12 min, and the solvent takes tens of minutes to return from the regenerator. Ratio feedforward, analyzer feedback and a loading override consequently protect different parts of the same unit. Model predictive control (MPC), a recognized industrial approach [7], extends that comparison to constrained operation. The reaction chemistry, the equilibrium form, the PI loops and the optimizer are all established methods, applied here to one teaching unit. Section 1.4.9 places this unit among the plants of the volume whose binding limit is measured late or only inferred.

This chapter measures how selective absorption, delayed analysis, loading protection and steam availability interact in a reproducible MDEA sweetener, through nine cases that make those interactions visible. Section 8.2 places the simplified chemistry beside published equilibrium and rate work. Section 8.3 specifies the absorber stages, the regenerator, the solvent return path and the foam

state. Section 8.4 gives the measurements and the control laws: the PI loops, the gas ratio, the loading floor, the analyzer trim and the predictive option. Section 8.5 reports the nine cases, from a manual circulation cut to the twelve-hour graded shift. Section 8.6 sets the hardware of a real sweetener beside the model, Section 8.7 states what the synthetic model cannot predict, and Section 8.8 summarizes which measurement protects which limit.

> ✦ DID YOU KNOW? **Where the Hydrogen Sulfide Goes**
>
> In 2024 the United States produced an estimated 8.2 million metric tons of elemental sulfur, recovered at petroleum refineries, natural gas processing plants and coking plants, 81 plants in 25 states [8]. The acid gas a sweetener's regenerator sends toward a Claus unit is the gas plants' share, and Volume 11, Chapter 13 follows it there.

8.2 Selective Treating and Its Modeling History

An aqueous amine absorbs acid gases from natural gas and releases them when heated at lower pressure. Section 1.4.7 gives the single-solute absorber and its Kremser relation. Here two solutes share one solvent. Hydrogen sulfide, carbon dioxide and water interact with the amine differently, so a prediction based on one fixed removal fraction does not transfer from a lean, cool contactor to a loaded, warm one. The 1981 equilibrium model of Deshmukh and Mather represented both H_2S and CO_2 solubility in aqueous alkanolamines [5]. The later electrolyte-NRTL treatment of Austgen and coauthors gave another equilibrium framework for acid-gas amine systems [6]. Neither result makes a six-stage teaching absorber a plant-design model, but both establish that the two solutes share a solvent and influence its equilibrium.

> ⧖ HISTORY **1930: A Solvent That Gives the Gas Back**
>
> On 7 October 1930 Robert Roger Bottoms applied for a patent on separating acidic gases, chiefly carbon dioxide, sulfur dioxide and hydrogen sulfide, from industrial gas mixtures. The usual alkaline carbonates held little gas for their volume. He had found that a small group of amines unite with acidic gases at a low temperature and give them up as gas at a higher temperature, regenerated, while their own vapor pressure stays low. The absorbent could therefore circulate continuously in a closed cycle between a cool absorption stage and a hot regeneration stage, with heat interchangers to save heat, and he preferred triethanolamine for its higher boiling point [9]. Amine scrubbing has separated carbon dioxide from natural gas and hydrogen ever since [10].

The rate distinction was investigated separately, first for carbon dioxide and tertiary amines, as the second History box below tells. A rate-based 2004 model treated selective H_2S absorption and desorption in aqueous MDEA [4], while a separate 2004 study measured selective absorption of H_2S from mixed H_2S-CO_2 gas into MDEA and another amine [3]. Those works motivate keeping an explicit carbon-dioxide rate in the absorber rather than multiplying both gases by the same removal efficiency.

> ⧖ HISTORY **1988 and 1990: The Slow Path of Carbon Dioxide**
>
> In 1988 Versteeg and van Swaaij absorbed carbon dioxide into aqueous MDEA and two other tertiary alkanolamines through a liquid surface that looked completely smooth, in a stirred vessel, and inferred the reaction rate from the enhanced mass transfer. The reaction fitted a base-catalysis mechanism proposed earlier, and a Brønsted relation tied the logarithm of its rate constant to the strength of the amine as a base [1]. In 1990 Littel, van Swaaij and Versteeg measured three tertiary amines from 293 to 333 K, found the reaction first order in the amine, and traced the mechanism to Donaldson and Nguyen. They argued that most earlier work had overstated the part played by hydroxide ions, which are depleted near the gas-liquid interface [2].

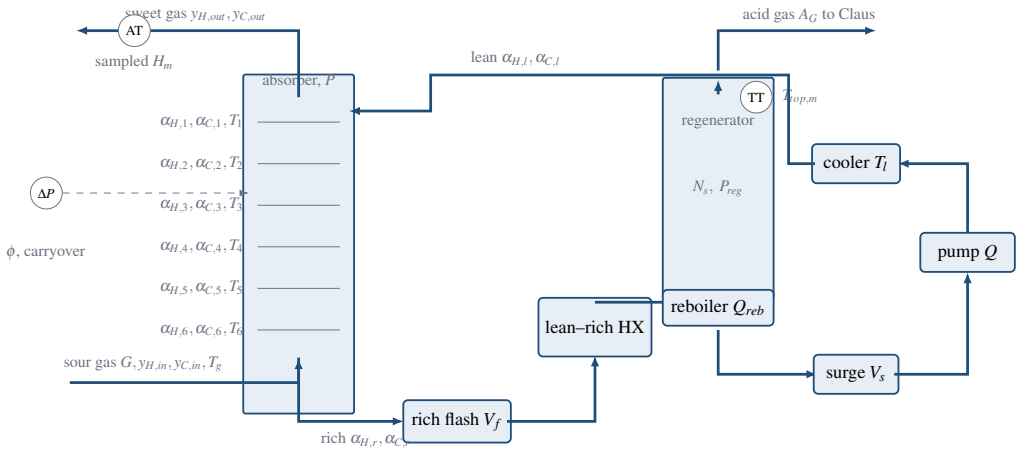

Figure 8.1 The simulated absorber and amine loop. Six absorber liquid stages carry H_2S and CO_2 loadings $\alpha_{H,n}$ and $\alpha_{C,n}$ and temperature T_n. Gas flow G, composition $y_{j,in}$ and temperature T_g enter below. Circulation Q returns lean solvent through the cooler. The rich flash, exchanger, regenerator and surge tank impose the return path seen by analyzer trim. AT reports a sampled and delayed H_2S value H_m, while the differential-pressure measurement catches foam. The diagram describes a teaching process model, not a piping design.

Published equilibrium data can also test the shape of a compact correlation without pretending to validate a complete plant. A 2020 patent tabulates H_2S partial pressures of 10, 30 and 100 kPa at loadings 0.21, 0.40 and 0.69 mol/mol for 50 wt% MDEA at 40 °C [11]. After the simulator's concentration factor is scaled from 45 to 50 wt%, its three values are 9.78, 31.34 and 98.17 kPa, errors of -2.23, $+4.46$ and -1.83%. This comparison tests three points on one solute's equilibrium curve. It does not check mixed-gas kinetics, hydraulic behavior, metal corrosion or the size of any dynamic upset.

The controller structures follow a common industrial progression. A fixed solvent rate is easy to operate but cannot follow a flow step. A measured gas-flow ratio acts before product analysis, while a product analyzer supplies a slower correction. A loading floor and a differential-pressure antifoam trigger add protection for limits the analyzer does not measure. PID tuning may be related to model reduction [12] and internal-model-control design [13], but the three loops here are shipped as declared teaching tunings. Chapter 7 applies the same progression of ratio feedforward, analyzer trim and economic MPC to a single solute. An economic MPC uses the same measurement packet to trade steam and pump work against predicted quality and loading constraints [7, 14, 15].

8.3 Unit, States and Constitutive Relations

Figure 8.1 follows the Unit view. Sour gas at 60 t/h, 60 bar absolute, 35 °C, 2 mol% H_2S and 4 mol% CO_2 rises through six lumped stages. The lean stream contains 45 wt% MDEA at a design flow of 64 m^3/h. Rich solvent leaves the absorber bottom, enters a flash drum and a lean–rich exchanger, then feeds a four-theoretical-stage stripping section over a reboiler. The acid gas overhead goes toward a Claus unit. A 40 m^3 surge tank, pump and cooler return lean solvent to the absorber top.

The model clock is minutes. Gas flow is converted from tonnes per hour to kmol per minute using the inlet mixture molecular weight, and the solvent's amine concentration is $c_A = \rho w_A / M_A = 3.92749\,\mathrm{kmol/m^3}$ with $\rho = 1040\,\mathrm{kg/m^3}$, $w_A = 0.45$ and $M_A = 119.16\,\mathrm{kg/kmol}$. Table 8.1 separates variables a learner may set from states and observations. Table 8.2 lists the full plant defaults and which are visible to the learner.

Table 8.1 Principal Variables at the Design Point

Quantity	Role	Initial	Range or test
G	disturbance, metered	60 t/h	30–90 t/h
$y_{H,in}$, $y_{C,in}$	disturbances	2, 4 mol%	0.5–5, 0–10 mol%
Q	manipulated, measured	64 m^3/h	0–110 m^3/h
$T_{top,SP}$	manipulated	106 °C	98–112 °C
H_m	controlled, delayed	2.64 ppm	4 ppm ceiling
α_r	state, displayed	0.42965	0.5 mol/mol cap
Q_{reb}	state, metered	4654 kW	6000 kW capacity
ϕ	state, foam index	0.45	1 carry threshold
V_s	state, surge volume	40 m^3	decreases on carry, pump starved below 0.5 m^3
ΔP	measured	120 mbar	1.3× clean alarm

8.3.1 Equilibrium, Transfer and the Six Stage Balances

The teaching equilibrium pressure of solute $j \in \{H, C\}$ over a stage is

$$
\begin{aligned}
p_j^{*} &= K_j \frac{\alpha_j \alpha_t}{1-\alpha_t} \exp(-b\alpha_t) \exp(\theta_j), \\
\theta_j &= \frac{\Delta H_j}{R}\left(\frac{1}{313.15} - \frac{1}{T+273.15}\right), \qquad \alpha_t = \alpha_H + \alpha_C .
\end{aligned}
\tag{8.1}
$$

The implementation floors $1-\alpha_t$ at 0.02 during excursions and nonnegative loadings at zero. At 40 °C, $K_H = 245$ kPa, $K_C = 100$ kPa and $b = 2.1$. The heats are 40 MJ/kmol for H_2S and 55 MJ/kmol for CO_2. The 40 °C curve has an approximately quadratic low-loading slope, then rises as free amine is exhausted. All four constants and the nonideality term are fitted teaching choices, rather than a regression supplied by the equilibrium papers.

For stage n, counted from the top, the incoming gas is from $n+1$. The model takes H_2S a fraction $E_H = 0.95$ toward equilibrium at each contacted stage, while CO_2 depends on an effective rate k_C and the liquid amine inventory:

$$
\begin{aligned}
N_{H,n} &= G_{n+1} E_H^{eff} \left(y_{H,n+1} - p_{H,n}^{*}/P\right), \\
K_{C,n} &= k_C c_A V_L (1-\alpha_{t,n}) f_c \\
&\quad \times \exp\left[-\frac{E_a}{R}\left(\frac{1}{T_n+273.15} - \frac{1}{313.15}\right)\right], \\
E_{C,n} &= \frac{K_{C,n} P}{G_{n+1} + K_{C,n} P}, \\
N_{C,n} &= G_{n+1} E_{C,n} \left(y_{C,n+1} - p_{C,n}^{*}/P\right), \\
G_n &= G_{n+1} - N_{H,n} - N_{C,n}, \\
y_{j,n} &= \frac{G_{n+1} y_{j,n+1} - N_{j,n}}{G_n}.
\end{aligned}
\tag{8.2}
$$

Here P is in kPa, N and G in kmol/min, $k_C = 10^{-4}\,(\text{min kPa})^{-1}$, $E_a = 30$ MJ/kmol, $V_L = 1.2$ m^3 per lumped stage, and $f_c = 1/(1+0.8\max(\phi-1,0))$ is foam's contact factor. Foam scales E_H by the same f_c. The top $6 - N_{contact}$ stages are dry if the learner lowers the feed point, so both rates there are zero. The gas contracts as it loses solute and has no separate holdup.

The stage loadings integrate the liquid counterflow and the transfers, while one thermal state per

stage forms the bulge:

$$\begin{aligned} c_A V_L \dot{\alpha}_{j,n} &= F_n(\alpha_{j,n-1} - \alpha_{j,n}) + N_{j,n}, \\ C_n \dot{T}_n &= \frac{Q_n \rho c_{p,L}}{60}(T_{n-1} - T_n) + G_{kg} c_{p,G}(T_{n+1} - T_n) \\ &\quad + \Delta H_H N_{H,n} + \Delta H_C N_{C,n}. \end{aligned} \tag{8.3}$$

F_n is kmol/min of circulating amine, Q_n is $\mathrm{m^3/h}$, $C_n = V_L \rho c_{p,L} C_{stage}$, $c_{p,L} = 3.6\,\mathrm{kJ/(kg\,K)}$, $c_{p,G} = 2.5\,\mathrm{kJ/(kg\,K)}$ and $C_{stage} = 1.3$. The top incoming liquid is lean solvent and the bottom incoming gas is the sour feed. At design $T_1 = 43.93\,^\circ\mathrm{C}$, $T_5 = 59.22\,^\circ\mathrm{C}$ and $T_6 = 55.01\,^\circ\mathrm{C}$. The gas and cold lean solvent cool the ends while absorption heats the lower interior.

8.3.2 Regenerator, Return Path and Foam

The rich flash drum of $10\,\mathrm{m^3}$ and exchanger line of $3\,\mathrm{m^3}$ each mix the loadings toward the stream arriving from the absorber at rate $F_r/(c_A V)$, and flash-drum temperature follows rich temperature at $Q_r/(60 V_f)$. At 1.9 bar absolute the reboiler temperature is the IAPWS-IF97 saturation temperature [16] plus a declared 3 K boiling-point elevation. A lean–rich exchanger of effectiveness 0.8 preheats the rich stream. The stripping section uses a four-stage Kremser fraction for each solute, the stripping counterpart of Equation (1.12):

$$\begin{aligned} V_b &= \frac{60 \max(0, Q_{reb} - Q_{loss})}{\lambda}, \quad S_j = \beta_j \frac{V_b}{F_r}, \\ f_j &= \frac{S_j - 1}{S_j^{N_s+1} - 1}, \quad \alpha_{j,out} = f_j \alpha_{j,feed}. \end{aligned} \tag{8.4}$$

The $S_j = 1$ limit is $1/(N_s + 1)$. In the implementation $Q_{loss} = 100\,\mathrm{kW}$, $\lambda = 39\,700\,\mathrm{kJ/kmol}$, $N_s = 4$, $\beta_H = 1.9$ and $\beta_C = 2.3$. A sensible-heat demand on the rich liquid and desorption heats reduce the steam left for the overhead. The overhead temperature target is the water dew point at water partial pressure $y_W P_{reg}$, with a teaching floor 15 K below the hot rich feed. Its state follows this target with a 2 min lag. The two rich-side loadings in the $8\,\mathrm{m^3}$ reboiler sump, two loadings in the $40\,\mathrm{m^3}$ surge tank and two in the $2.5\,\mathrm{m^3}$ lean line each obey a well-mixed flow balance. The lean cooler target is $T_{hot} - 0.9 u_c (T_{hot} - T_{cw})$, with 30 °C cooling water, filtered over 2 min. Pump flow and duty follow their commands with 0.2 and 1 min lags.

The foam state ϕ moves toward a load- and contamination-dependent target in 0.25 min. Hydrocarbon in the solvent grows if the lean amine is below the gas dew point, taken here as T_g, and drains with a 45 min time constant. Antifoam reaches the trays over 1 min and decays in 90 min. The clean pressure drop is $120(G/60)^2$ mbar. The foam correction raises it by $(1 + 0.8\phi^2)/(1 + 0.8(0.45)^2)$, while solvent carryover is $12 \max(\phi - 1, 0)^{1.5}\ \mathrm{m^3/h}$. The sweet-gas scrubber drains at $1.5\,\mathrm{m^3/h}$ and stops the run when its $1\,\mathrm{m^3}$ capacity fills. A separate stop follows 60 min continuously above rich loading 0.5. Neither stop is a certified equipment-protection prediction.

8.3.3 Initialization, Numerics and Scope

Each case starts from a steady solution for its own disturbances, circulation, overhead setpoint and lean temperature. The solver closes the lean solvent around the absorber, flash, exchanger, regenerator, sump, surge tank and cooler, then initializes every first-order state at its target. The integrated state has 18 absorber quantities, four foam and antifoam quantities, seven rich-side quantities, six lean-side quantities, pump flow, duty, overhead temperature and scrubber volume, plus ten boundary ledgers. The levers are clamped to their declared ranges after each step. The engine of Section 1.6.1 integrates with a classical fourth-order Runge–Kutta method in 1 s steps, while control runs every 5 s. The full disturbed-run ledger in the application tests closes H_2S and CO_2 to relative error below 10^{-10} and solvent volume below $10^{-9}\,\mathrm{m^3}$. Repeating the 31-seed challenge at 0.5 s retains every verdict.

Table 8.2 All Plant Parameter Defaults. V is visible on the parameter panel and H is hidden from the learner. Values are generic teaching choices, and the software uses the variable names shown.

Key	Value	Unit		Key	Value	Unit	
P	60	bar	V	tauQ	1	min	H
G0	60	t/h	H	tauTop	2	min	H
MWhc	17	kg/kmol	H	Qloss	100	kW	H
cpG	2.5	kJ/(kg K)	H	Vsurge0	40	m^3	V
VL	1.2	m^3	V	Vlean	2.5	m^3	H
EH	0.95		V	Tcw	30	°C	H
kC	10^{-4}	1/(min kPa)	V	epsC	0.9		H
EaC	30000	kJ/kmol	H	tauCool	2	min	H
Cstage	1.3		H	Q0	64	m^3/h	H
wMDEA	0.45		H	QpMax	110	m^3/h	V
rho	1040	kg/m^3	H	tauP	0.2	min	H
cpL	3.6	kJ/(kg K)	H	foamTend	1		V
KH	245	kPa	V	phi0	0.45		H
bAct	2.1		H	kHC	3		H
KC	100	kPa	H	tauHC	45	min	V
dHH	40000	kJ/kmol	H	kCond	0.006	per K min	H
dHC	55000	kJ/kmol	H	kAF	4		V
Vflash	10	m^3	V	tauAFin	1	min	H
Vexch	3	m^3	H	tauAF	90	min	H
Vreg	8	m^3	V	tauFoam	0.25	min	H
eps	0.8		H	Cf	12	m^3/h	V
Preg	1.9	bar	H	kEf	0.8		H
bpe	3	K	H	dp0	120	mbar	H
lamW	39700	kJ/kmol	H	Vscrub	1	m^3	V
Ns	4		V	scrubDrain	1.5	m^3/h	H
betaH	1.9		V	capLoad	0.5	mol/mol	V
betaC	2.3		H	capHold	60	min	V
QrebMax	6000	kW	V	dTapp	5	K	H
Vsuct	0.5	m^3	H				

The gas has no dynamic holdup. Six liquid cells stand for a many-tray absorber, and the gas-film efficiency and CO_2 uptake are teaching closures. Regeneration is quasi-steady in its stripping section and uses a simplified Kremser expression. Foam, corrosion and hydrocarbon condensation are phenomenological limits chosen to make a process-control lesson. There is no validated hydraulic prediction, corrosion-rate law, electrolyte species balance, absorber flooding map or Claus chemistry. The purpose of the states is to expose delay, selectivity and competing limits while keeping the mass transfers internally consistent.

8.3.4 State Transport and Boundary Ledgers

The generic solvent holdup equation makes the lags in Figure 8.1 explicit. For either gas j, a well-mixed rich-side vessel of volume V and flow Q_r obeys $V\dot{\alpha}_j = (Q_r/60)(\alpha_{j,in} - \alpha_j)$. The flash drum and exchanger use $V = 10$ and $3\,m^3$ in series. The regenerator sump uses $8\,m^3$ and the stripped outlet of Equation (8.4) as its incoming loading. The surge tank receives the sump outlet, and the lean line receives the surge outlet. Its loadings therefore satisfy

$$\begin{aligned} V_s\dot{\alpha}_{j,s} &= (Q_r/60)(\alpha_{j,r} - \alpha_{j,s}),\\ V_{lean}\dot{\alpha}_{j,l} &= (Q/60)(\alpha_{j,s} - \alpha_{j,l}),\\ \dot{V}_s &= (Q_r - Q)/60. \end{aligned} \tag{8.5}$$

Carryover makes $Q_r < Q$, so the surge volume falls. The pump draws on the tank, and over its last $V_{suct} = 0.5\,m^3$ it loses suction: the delivered flow is $Q = Q_p \min(1, V_s/V_{suct})$ for a pump flow Q_p, and the carryover cannot exceed Q. A drained tank therefore starves the circulation until the foam no longer carries solvent over, and its outflow then equals its inflow. The stage loop in Equation (8.3) receives $\alpha_{j,l}$ and closes the acid-gas balance. The rich flash temperature approaches the absorber-bottom temperature at rate $Q_r/(60V_f)$. The reboiler sump and surge loadings react to a duty change separately, so an overhead temperature move can precede an improvement in the gas by tens of minutes.

The regenerator heat closure accounts for water boil-up, rich-liquid sensible heating and desorption:

$$\begin{aligned} D_j &= F_r(\alpha_{j,feed} - \alpha_{j,out}), \quad A_G = D_H + D_C, \\ q_{sens} &= (Q_r/60)\rho c_{p,L}(T_{reb} - T_{rich,hot}), \\ q_{des} &= \Delta H_H D_H + \Delta H_C D_C, \\ S_{ov} &= \max[0, V_b - (q_{sens} + q_{des})/\lambda], \\ y_W &= S_{ov}/(S_{ov} + A_G). \end{aligned} \tag{8.6}$$

The last ratio is zero when its denominator vanishes. The overhead target is the greater of the water saturation temperature at $y_W P_{reg}$ and $T_{rich,hot} - 15\,\mathrm{K}$. A first-order 2 min state follows it. Gas absorption heat appears in each stage balance and solvent sensible heat in the regenerator duty. Heating supplied by the steam header is accumulated in an energy ledger. The ledgers integrate inlet H_2S and CO_2, gas outlet, acid-gas outlet and carryover with the same numerical step as their stage states. Thus their residuals reveal lost material if a transport or closure term is edited. They are diagnostic states, not additional equipment inventories.

A compact foam state connects the approach violation to the pressure measurement. With hydrocarbon contamination h_c, antifoam at the trays a_f and inlet gas at its assumed hydrocarbon dew point, the implementation uses

$$\begin{aligned} \dot{h}_c &= 0.006\max(T_g - T_l, 0)G/G_0 - h_c/45, \\ \phi^* &= 0.45\frac{(1+3h_c)}{(1+4a_f)}(G/G_0)^2\sqrt{Q/Q_0}, \\ 0.25\dot{\phi} &= \phi^* - \phi. \end{aligned} \tag{8.7}$$

The hydrocarbon source is zero if lean amine is warmer than sour gas. A manual or automatic antifoam shot first enters a one-minute delivery state, then the tray state decays in 90 min. These time constants and coefficients are teaching settings. They express a fast pressure response and a slow contamination inventory without representing a measured emulsion chemistry.

8.4 Measurements and Control

The simulated plant contains the true concentrations and temperatures, but a controller receives only a measurement packet. At each 5 s execution, gas flow has a 6 s lag and 0.2% multiplicative noise. Sour gas temperature, lean temperature, regenerator overhead and absorber differential pressure have 0.5, 0.3, 0.3 and 0.03 min lags. The acid-gas meter has a 0.5 min lag and 1% noise, while circulation has 0.3% noise. The sweet-gas analyzer samples every 10 min, spends 4.5 min in analysis and 0.5 min in its sample line, then reports with 0.05 ppm noise. A value displayed between reports is the last reported sample, not the current true gas. The true sample is retained separately to grade the challenge.

8.4.1 Fixed Flow, Gas Ratio and Loading Floor

In Auto, the circulation base command is either a fixed setpoint, a feedforward ratio or the MPC's command. The ratio uses measured gas flow G_m and the last laboratory feed composition $y_{H,lab}, y_{C,lab}$, not the current feed assay:

$$Q_{base} = \begin{cases} Q_{SP}, & \text{Fixed}, \\ 1000 r G_m(y_{H,lab} + y_{C,lab})/(M_{mix}c_A), & \text{Ratio}, \\ Q_{MPC}, & \text{MPC}. \end{cases} \tag{8.8}$$

The implemented load calculation converts G_m from t/h to kmol/h before multiplying by composition, and $r = 1.29$ mol MDEA per mol acid gas at design. The rich loading inference filters the measured

acid-gas molar flow A_m and circulation separately with $\tau = 12$ min. It adds an assumed lean loading of 0.004 mol/mol. For Fixed and Ratio only, the inferred loading produces a max-selected flow floor, the high-selector shape of Section 1.4.8:

$$\begin{aligned} \hat{\alpha}_r &= 0.004 + \frac{\bar{A}_m}{c_A \max(1, \bar{Q}_m)}, \\ Q_{floor} &= \frac{\bar{A}_m}{c_A \max(0.05, \alpha_{lim} - 0.004)}, \\ Q^* &= \min(110, \max(Q_{base}, Q_{floor})). \end{aligned} \tag{8.9}$$

The floor is active only when its switch is On. The pump follows Q^* with a 0.2 min time constant. A sourer well whose new composition has not reached the laboratory changes A_m after the absorber, so the floor supplies information the flow-only ratio lacks. The loading cap of 0.5 is a teaching corrosion guard. The learner may set the floor below it, at a default 0.45 and 0.42 in the suggested challenge.

8.4.2 Temperature, Analyzer and Antifoam Loops

The regenerator overhead temperature is held by the steam valve with the framework's velocity-form PI controller of Section 1.6.2 [12, 13]. Its proportional gain is 4% of valve per kelvin, integral time 6 min, and valve bounds 0–100%. A second PI holds lean amine at $T_g + 5$ K or at a fixed temperature by increasing cooler opening when the amine is too warm. Its gain is 3% per kelvin and integral time 4 min. Both use back-calculation with a tracking interval equal to each integral setting. Derivative time is zero. In Manual the three stations follow the operator's pump, steam and cooler commands and the inactive integrals track those commands, so the subsequent Auto transfer has no command jump.

The optional analyzer PI changes the regenerator temperature setpoint only on a new report, a cascade onto the temperature loop. It is reverse acting in the sense that measured H_2S above the 3 ppm target raises the temperature setpoint, with 1.5 K/ppm gain and 40 min integral time. The trim itself is bounded from -4 to $+6$ K and the resulting setpoint from 98 to 112 °C. It can therefore lower steam use during periods of ample quality margin. A differential-pressure ratio above 1.3 times clean triggers one antifoam shot in Auto, with at least 5 min between shots. These comparisons use filtered measured pressure, not the true foam state.

8.4.3 Two-Input Predictive Option

The MPC executes every 5 min. Section 1.6.3 describes the general setup, and this paragraph gives what is specific to this unit. It uses the controller's own copy of the steady absorber and regenerator map $m(Q, T_{SP}) = [\ln H, \alpha_r, Q_{reb}]$, updated with measured gas flow and the laboratory composition. The map is linearized about the current internal states by a 2 m^3/h circulation perturbation and a 0.5 K temperature perturbation. Three lags represent circulation to the sweet gas (5 min), overhead temperature (4 min) and the sump–surge–lean return, whose time constant is $\max[5, 60(V_r + V_s + V_{lean})/Q_m]$ min. The analyzer bias is corrected against the model state at the time the reported sample left the absorber, with a 0.3 update fraction. The loading bias takes the same fraction of the measured inference residual.

The quadratic program predicts 36 five-minute steps over three hours in move blocks spanning 1, 1, 2, 4, 8 and 20 prediction steps. It minimizes the linearized steam and pump duty plus moves while penalizing upper-limit slack on $\ln H$ and α_r. Pump work is estimated from 60 bar and 70% efficiency. The H_2S prediction limit is the learner's 4 ppm setting minus the declared 1 ppm margin, and rich loading is bounded by the lesser of the loading setpoint and the 0.5 cap. Circulation is bounded 30–110 m^3/h, setpoint 98–112 °C, and each move by 10 m^3/h and 1 K. When the steam valve has been wide open for 2 min, the temperature move is held and the map uses measured available duty. If the quadratic program fails or proposes a nonfinite or out-of-bounds move, the last bounded pair is retained until the next execution. This is a local control model, not the plant's six-stage dynamic

Table 8.3 Case Sequence and Main Student Action

No.	Event	Action or comparison
1	pump 64 to 50 m^3/h at 10 min	restore manually before stop
2	gas +20% at 30 min	Ratio then Fixed and Reset
3	gas warms 35 to 40 °C	enable analyzer trim
4	liquid holdup at 50–200%	compare H_2S and CO_2 profiles
5	feed H_2S 2 to 2.6%	engage loading floor
6	hydrocarbon slug at 30 min	enable antifoam on ΔP
7	lean setpoint 40 to 30 °C	restore gas-plus-approach
8	steam availability to 60%	trace delayed lean loading
9	sour well, slug, steam dip	meet 12 h scorecard

Table 8.4 Measured Landmarks in the Shipped Cases. The time column names the first relevant threshold or recovery, not an estimated settling time.

Case	Intervention or comparison	Time (min)	Engine observation
1	cut circulation to 50 m^3/h at 10	13.8 true, 24.7 report	true H_2S crosses 4 ppm, then the analyzer reports it
2	20% more gas on Ratio / Fixed	41.2 true on Fixed	2.66 / 5.04 ppm true peak
3	turn analyzer trim On at 75	154.7 analyzer recovery	report returns below 4 ppm after 79.7 min
4	double tray holdup at 10	120 readout	53.0% feed CO_2 absorbed, 2.25 ppm H_2S
5	sourer well at 30, loading floor at 45	38.5 cap, 51.3 return	rich loading peaks at 0.510 mol/mol
6	hydrocarbon slug at 30, antifoam Auto	30.167 alarm	206 mbar peak, no scrubber carryover
7	restore gas-plus-approach at 45	45.0 ΔP alarm	159 mbar peak, no carryover stop
8	steam availability 60% at 30	59.3 true, 64.7 report	6.62 ppm true peak at 129.5
9	suggested four-action line-up	720 completion	3.78 ppm worst sampled truth, 5.25 MW average

model. The app also keeps the learner's identified ratio or loading impression separate from the MPC's map.

8.5 Nine Cases on the Engine

Each case loads its own steady state and scripted events. Results in Tables 8.3, 8.4 and 8.5 are with the simulation engine at the shipped random seed and 1 s physics step, with operator actions entered at the times named. Rich loading is mol of H_2S plus CO_2 per mol MDEA. The analyzer is a delayed, noisy report, whereas the true gas is an immediate simulation quantity used for comparisons. The same case selected first, selected after another and rerun through Reset gives the same CSV in all nine load-path checks.

8.5.1 A Limit the Product Analyzer Cannot See

In case 1 the manual circulation cut takes true sweet-gas H_2S over 4 ppm at 13.8 min, rich loading past 0.5 at 14.7 min, and the analyzer over 4 at 24.7 min. Holding the cut gives 0.561 rich loading and approximately 23 ppm H_2S at one hour, followed by a corrosion stop at 74.5 min. Restoring 64 m^3/h at 50 min clears the loading cap at 53.8 min and the analyzer limit at 64.7 min. Holding smaller cuts finds the least circulation near the design point. About 57 m^3/h holds 4.00 ppm and 0.492 rich loading, whereas 55 m^3/h gives 5.85 ppm and 0.512 and ends in the corrosion stop. Both limits therefore bind at nearly the same circulation, and the result is no more precise than the synthetic equilibrium shape and the one modeled temperature profile that fix it. The delayed staircase explains why an operator who waits for the product report has less time to protect the rich line than the gas trend suggests.

The gas-rate step of case 2 separates fixed circulation from feedforward. Ratio raises circulation from 63.9 to 77.1 m^3/h and keeps the true H_2S peak at 2.66 ppm. Fixed circulation reaches 5.04 ppm true gas, crossing 4 at 41.2 min, while the analyzer does not cross until 54.7 min. The ratio's rich loading changes from 0.430 to 0.434, compared with 0.431 to 0.497 on Fixed. However, the ratio reads the last laboratory feed composition. In case 5 a sourer well raises feed H_2S from 2 to 2.6%, so rich loading crosses its 0.5 cap at 38.5 min while true gas stays below 3.43 ppm. Left alone, the

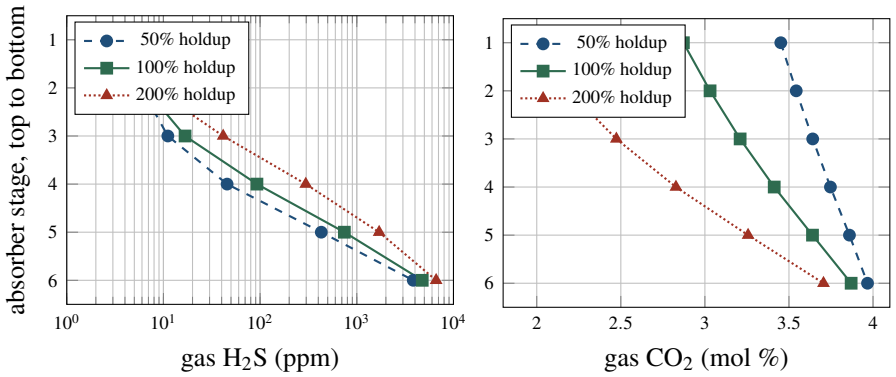

Figure 8.2 Gas profiles calculated from case 4 at 120 min, with 80 m^3/h circulation and three liquid holdups. Stage 1 is the top. H_2S uses a logarithmic ppm axis, while CO_2 uses mol%. The app's own Selectivity view also draws hollow equilibrium rings at each stage. Circles with dashed lines mark 50% holdup, squares with solid lines 100% and triangles with dotted lines 200%. The plotted points are simulation truth computed from the case, not analyzer samples.

unit stops at 98.4 min. Switching on the loading floor at 45 min brings the rich loading under 0.5 by 51.3 min and to 0.447 by 240 min, with 2.28 ppm sweet-gas H_2S. A product analyzer cannot be a substitute for a rich-loading estimate when these two limits separate.

> **TRY IT IN THE APP** **5 · The loading cap on a sourer gas**
>
> Load the case and press Run. Watch the Rich loading trend and the H_2S in sweet gas trend after the feed turns sourer at 30 min. The rich loading crosses its 0.5 cap at 38.5 min, yet the true H_2S stays below 3.43 ppm and the analyzer never reports a problem. Left alone, the corrosion stop ends the run at 98.4 min.

8.5.2 The Contact-Time Experiment

At 80 m^3/h circulation, case 4 varies tray liquid holdup without changing the number of stages. Halving it takes only 16.0% of feed CO_2, compared with 30.5% at design and 53.0% at double holdup. The corresponding sweet-gas H_2S values at 120 min are 3.26, 2.85 and 2.25 ppm. The latter moves slightly because the regenerator's steam duty changes with coabsorbed CO_2, from 4.84 MW at design to 5.81 MW at double holdup. The independent point is the nearly twofold difference in CO_2 pickup, not an assertion of strictly constant H_2S. Four stages in contact give 6.76 ppm H_2S because a complete approach-to-equilibrium stage has been removed. Figure 8.2 shows each gas concentration leaving the six stages at half, design and double holdup.

Pause. Try Exercise 8.2 before reading on.

8.5.3 A Slow Correction and Two Fast Alarms

Gas warming in case 3 makes true H_2S pass 4 ppm at 62.3 min. The analyzer reports the exceedance at 74.7 min. Enabling its trim at 75 min lifts the overhead setpoint toward 109 °C, moves reboiler duty from about 4.68 to 5.82 MW, and gets the analyzer below 4 ppm at 154.7 min. The sample delay is only part of that 80-minute recovery. Leaner solvent must also traverse the sump, 40 m^3 surge tank and cooler line.

A hydrocarbon slug is much faster. In case 6 the absorber pressure drop reaches its alarm ten seconds after the 30 min slug, at 30.167 min, peaks at 360 mbar untreated and fills the sweet-gas scrubber to a stop at 44.98 min. The delayed analyzer does not report more than 4 ppm until 44.7 min, so the pressure alarm leads it by 14.5 min. Automatic antifoam gives one shot, limits the pressure

Table 8.5 Challenge Results at the Default Seed

Configuration	Time (min)	Worst (ppm)	Peak rich (mol/mol)	Mean duty (MW)
Fixed, shipped	214.5	5.27	0.565	4.86
Ratio, floor, antifoam, trim	720	3.78	0.458	5.25
Same, no trim	720	3.63	0.458	5.54
Same, no antifoam	367.0	3.13	0.458	5.04

drop peak to 206 mbar and avoids carryover. Instantaneous true H_2S still peaks at 4.40 ppm between samples, while the analyzer peaks at 2.70 ppm. This outcome is allowed by a challenge that grades sampled gas rather than continuous truth.

> **TRY IT IN THE APP** **6 · Foam: a hydrocarbon slug, ΔP and antifoam**
>
> Press Run and watch the Absorber ΔP trend after the hydrocarbon slug at 30 min. The pressure drop passes its alarm ten seconds later and climbs to 360 mbar, and the scrubber fills until the run stops at 44.98 min. Set Antifoam on high ΔP to Auto on the Foam card, press Reset and run again: one shot holds the peak to 206 mbar with no carryover.

In case 7 lowering lean amine from 40 to 30 °C briefly improves absorption but condenses hydrocarbon, and leaving it there ends at 84.6 min. Restoring gas-plus-approach at 45 min lets the pressure drop briefly reach 159 mbar, just above its 156 mbar alarm, then avoids carryover.

Steam supply is a third constraint. In case 8 a dip to 60% availability at 30 min makes true H_2S pass 4 at 59.3 min and the analyzer pass it at 64.7 min. The true peak is 6.62 ppm at 129.5 min. A steady sweep of circulation from 30 through 110 m^3/h at that duty finds no flow holding 4 ppm, because low flow overloads the solvent while high flow receives too little stripping per mol. The case asks the learner to recognize an infeasible loop rather than keep moving the pump.

Pause. Try Exercise 8.1 before reading on.

8.5.4 The Twelve-Hour Shift

The final case introduces feed H_2S from 2 to 3% over hours 2–3, a hydrocarbon slug at hour 6 and a one-hour steam availability dip to 85% at hour 9. Its verdict requires every analyzer sample's true H_2S to be at most 4 ppm, continuous true rich loading at most 0.5, no stop and an average duty at most 5.35 MW. A learner who changes a disturbance, model parameter or tray holdup receives no grade. The shipped Fixed setting stops on corrosion at 214.5 min. Table 8.5 shows the main controlled comparisons.

Pause. Try Exercise 8.3 before reading on.

Figure 8.3 sets the four configurations of Table 8.5 against the four clauses of the target.

The suggested configuration selects Ratio, sets the loading floor at 0.42, enables automatic antifoam and analyzer trim. Its exact worst sample is 3.779 645 ppm, peak rich loading 0.458294 and mean duty 5247.400 kW over 720 min. Removing trim keeps the sampled gas below 4 ppm but raises mean duty to 5539.637 kW, above the 5350 kW ceiling. Removing antifoam ends at 367.0 min on carryover. The simulator tests repeat the shipped and suggested verdicts at the default seed and seeds 0–29, each at 1 s and 0.5 s: all 62 shipped runs fail and all 62 suggested runs pass. This is numerical robustness for the declared model and scripted shift, not a probabilistic reliability claim for a real sweetener.

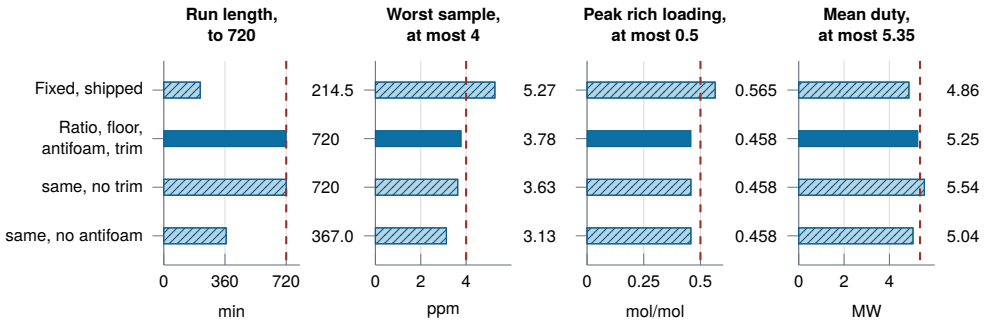

Figure 8.3 The challenge runs of Table 8.5 as bars: how long each run lasted, the worst sampled H_2S, the peak rich loading and the mean reboiler duty. The solid bars are the one configuration that passes, hatched bars those that fail, dashed red lines are the targets, and the value of each bar stands at the right. Notice that each configuration that fails breaks a different clause, and that the cleanest gas belongs to a run that ended early.

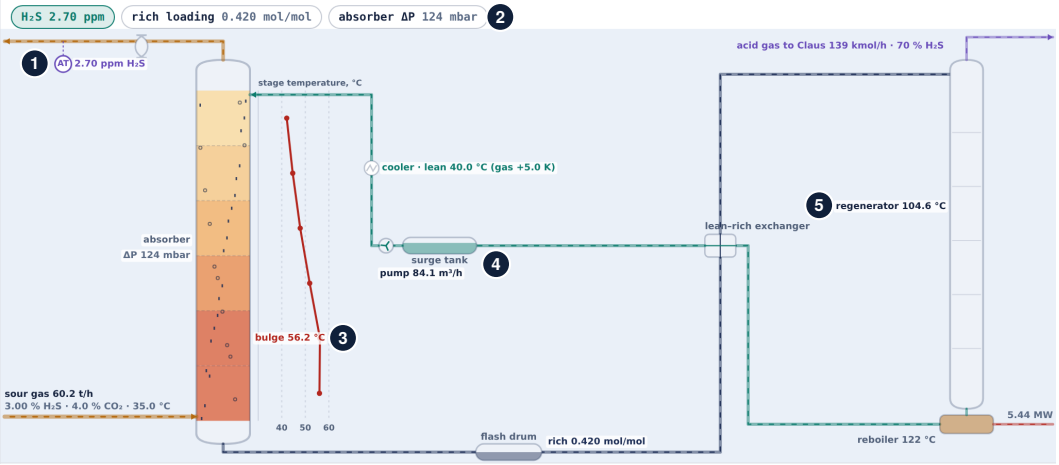

Figure 8.4 Suggested configuration at the end of the twelve-hour challenge, Unit view cropped to the plant drawing. (1) The analyzer AT on the sweet gas line with its current report, 2.70 ppm. (2) The status chips: the same report, the rich loading and the absorber pressure drop. (3) The absorber with the temperature of each stage, warmest in the lower interior where absorption heats the liquid. (4) The surge tank and the pump, which circulates about 84 m^3/h of lean amine. (5) The regenerator, its overhead at 104.6 °C, with the acid gas leaving above it toward the Claus unit and the reboiler below. The current report is not the grade: the challenge grades the complete run.

> ✦ **DID YOU KNOW?** **Feedback That Saved Steam**
>
> From the simulator (Table 8.5). Over the twelve-hour shift the analyzer trim left a slightly worse worst sample, 3.78 ppm against 3.63 ppm without it, and yet the unit used less steam, 5.25 MW against 5.54 MW. A loop that can only add heat when the gas is sour also takes it away when the gas is clean.

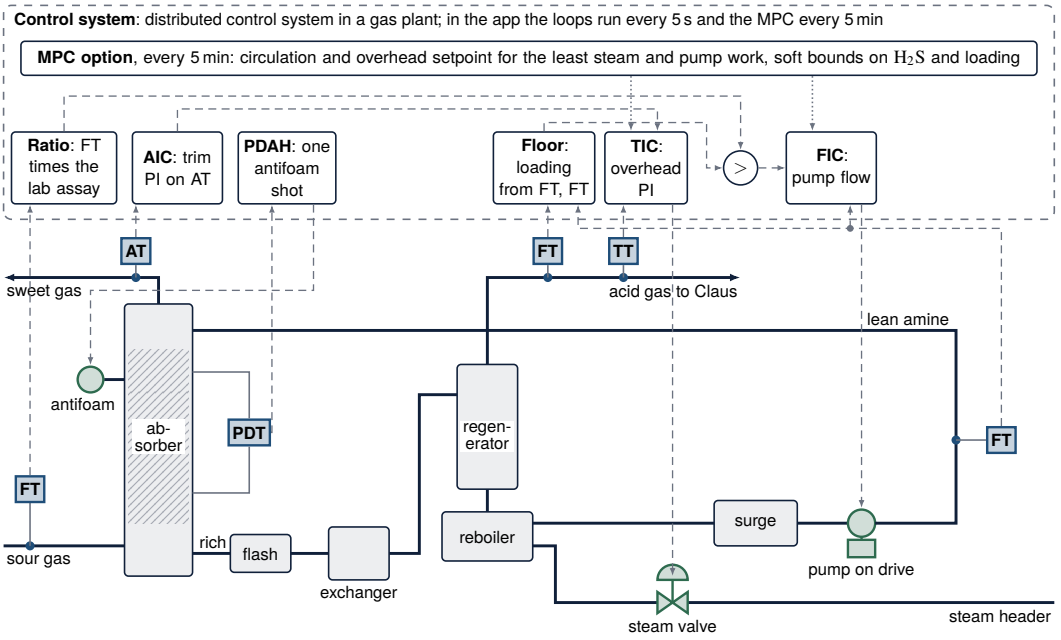

Figure 8.5 The loops of the sweetener simulator with their hardware. Lettered boxes are sensing devices with the chapter's tags: FT on the sour gas, AT on the sweet gas, PDT across the absorber, FT and TT on the acid gas overhead and FT on the lean amine. Green symbols are final elements: the lean amine pump on its drive, the reboiler steam valve and the antifoam pump. Dashed arrows are signals, dotted arrows the setpoints the optional MPC writes. Notice that AT reads the specification but not the loading cap, which only the acid gas meter and the lean flow meter infer, and that the floor meets the ratio at a high select. The lean cooler and its approach loop on the sour gas temperature are left out. Not to scale, and not a piping diagram.

8.6 Technology in Practice

Figure 8.5 draws the loops of Section 8.4 with the hardware of a sweetener in place of the schematic's symbols.

What Measures It

One instrument reads what the sales gas specification is about: the analyzer AT on the sweet gas. A process gas chromatograph, one common kind, separates a sample by its travel down a column, and cycle times of minutes are not uncommon even for online chromatographs in control loops [17]. Figure 8.6 shows the principle. The model's AT samples every 10 min with 0.05 ppm of noise. The loading floor reads two ordinary meters instead: the acid gas meter, with a 0.5 min lag, and the circulation.

> **TECHNOLOGY** **Sensor: The Sweet-Gas Analyzer**
>
> A 2018 review found the sulfur methods in use for natural gas accurate but slow, with few analyses, a long cycle and a delayed result [18]. An analyzer that separates its sample in time carries a dead time by principle, and enough dead time makes a feedback loop oscillate. An analyzer for hydrogen sulfide is checked against a span gas of known concentration, often 25 or 50 ppm, because analyzers drift more than pressure, level or flow transmitters [17]. A plugged sample line leaves it measuring stale sample with no way of knowing, so the trend goes quiet while the plant moves.

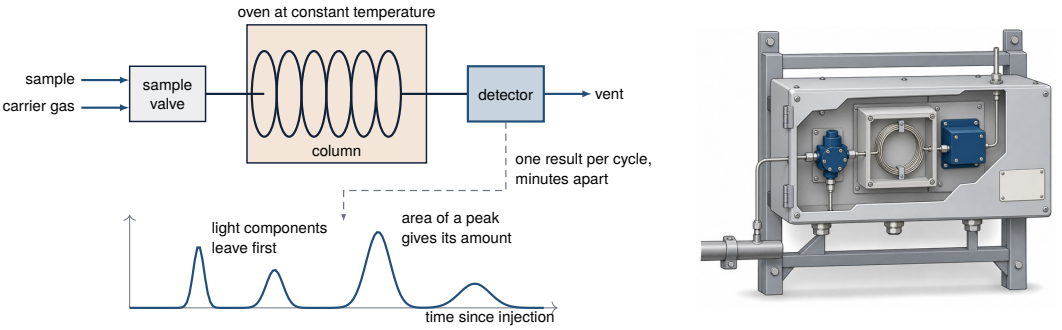

Figure 8.6 Principle of a process gas chromatograph, one kind of online analyzer for the sweet gas. A valve injects a sample into carrier gas, the column holds back each species for its own time, and the detector sees them leave one by one. Notice that a result exists only at the end of each cycle, which is the staircase the chapter's AT draws.

What Moves It

Four final elements carry the chapter's handles. The lean amine pump sets the circulation, the fast handle, from 0 to 110 m^3/h. The reboiler steam valve sets the duty that strips the solvent, the slow handle, and the steam header can limit it, as in case 8. The cooler holds the lean amine warmer than the gas, and the antifoam pump gives one shot on an alarm.

> **TECHNOLOGY** **Final Element: The Lean Amine Pump**
>
> The lean amine pump sets the circulation, the chapter's fast handle. A pump's flow can be set by throttling its discharge, by recycling to its suction, or by changing its speed with a drive. Changing the speed uses less energy than running at full speed and throttling, and it takes a control valve out of the line. A drive limits speed, current and the rate of acceleration and deceleration, and shuts down on overcurrent. A centrifugal pump, the most common kind in process plants, delivers less pressure as its flow rises [17]. In the chapter the pump follows its command through a 0.2 min lag and loses suction over the last 0.5 m^3 of the surge tank, so a drained tank shows on the trend as delivered flow falling below its command.

Where the Controller Runs

Inside a gas plant the loops would run on a distributed control system, which controls production within one location through feedback and feedforward loops under a supervisory level. Such loops run with cycle times from milliseconds to minutes [19]. The app's loops execute every 5 s and its MPC every 5 min, while the analyzer reports every 10 min, so no scan rate repairs the delay. The chapter's structure maps onto standard blocks: a ratio station that multiplies the gas meter by the laboratory composition, a high selector that lets the loading floor override it ahead of the FIC on the pump, a PID on the regenerator overhead with its setpoint trimmed by a PI on the analyzer, and a pressure drop alarm that fires the antifoam. The optional MPC sits in a supervisory layer above them.

IN PRACTICE **The Acid Gas Meter Is Required Anyway**

Under the United States performance standard for sulfur dioxide from onshore natural gas processing, written for plants built or modified from 1984 to 2011, a covered sweetening unit must measure its acid gas flow continuously and record it at least once an hour. At least one sample a day gives the hydrogen sulfide in that acid gas, and the sulfur produced each day is measured within 2 percent. Together they fix the sulfur recovery the plant must reach [20, 21]. The acid gas meter that the chapter's loading floor divides by the circulation is one such a plant already runs.

What Is Modernizing It

Three developments touch this chapter. Loading that the chapter infers from two meters is most often measured by wet titration in a laboratory, which can take two hours while the sample flashes carbon dioxide and the amine degrades. Near-infrared spectroscopy has been calibrated on a gas plant's solvent to read the amine, the carbon dioxide loading and the water at once, and its authors offer it to replace those methods [22]. Laser photoacoustic detectors read hydrogen sulfide in natural gas from a few ppm to several percent, despite the overlap of carbon dioxide's absorption, and are calibrated in place [23]. Model predictive control is established in industry [7], and the chapter's MPC adds a forecast of the loading and the sweet gas that does not wait for the next report.

8.7 Limitations

All vessel sizes, tunings, disturbances, foam laws and stop thresholds are generic teaching values. The three-point patent comparison supports only the rough shape of one H_2S equilibrium curve at 40 °C. The absorber's six cells do not resolve real tray hydraulics, mass-transfer area or gas-side holdup. The rate law compresses carbonate chemistry into one effective free-amine term, and the regenerator's quasi-steady stripping section cannot predict a moving temperature profile. Foam is one scalar state driven by hydrocarbon and antifoam, so the 10 s pressure alarm and 45 min scrubber stop are properties of this scripted model. A real unit needs measured fluid properties, a fitted hydraulic and kinetic model, and operating procedures for contamination and relief.

The challenge's quality requirement checks the true gas at analyzer sample instants, not continuously. Case 6 shows why that distinction matters: automatic antifoam allows a 4.40 ppm instantaneous peak while the analyzer sees at most 2.70 ppm. The feedback receives only reported samples and noisy process measurements, but the teaching overlay and grading can see the synthetic truth. The MPC uses a commissioned parameter copy and local map, so unmodeled foaming, solvent degradation and property drift can invalidate its forecast. No conclusion about industrial controller superiority is drawn from its local optimization or from the 62 deterministic seeded runs.

OPEN CHALLENGE **What Makes Amine Foam?**

Foaming is a stubborn problem of sour gas sweetening: in 2019 two incidents at gas plants cost solvent, put sweet gas off specification and cut the gas supply. Foam comes from contaminants in the feed gas or from amine degradation and corrosion, and identifying and quantifying the organic contaminants in the solvent remains a challenge. In a 2022 study with plant solvents, unsaturated fatty acids and alcohol ethoxylates promoted foam at 10 ppm, while linear hydrocarbons had little effect even at 2000 ppm [24]. The chapter's foam state, driven by hydrocarbon alone, is a teaching stand-in. Progress would be an online reading of the solvent's tendency to foam, so that antifoam goes in before the pressure drop rises.

8.8 Conclusion

The modeled sweetener couples a fast H_2S specification to a slower solvent inventory and to an analyzer that reports after each 10 min sample. At a circulation cut the true gas crosses 4 ppm at 13.8 min and the analyzer does so at 24.7 min. A sourer feed shows the other side of the problem: rich loading exceeds 0.5 while sweet gas remains below 3.43 ppm. Ratio circulation protects a metered

gas-rate change, an acid-gas-meter loading floor protects the rich circuit when the feed assay is stale, and differential pressure gives a response to foam long before product analysis can. Analyzer trim then trades additional margin for steam. In the declared twelve-hour challenge those four actions keep the worst sampled H_2S at 3.78 ppm, peak rich loading at 0.458 and average duty at 5.25 MW, while the shipped Fixed case stops on corrosion. The simulation's value is a reproducible comparison of which measurement can protect which limit, with its synthetic boundaries stated openly.

KEY POINTS

- A sweetener can keep its sales gas clean while overloading its own solvent: a sourer feed takes the rich loading over 0.5 while the gas stays below 3.43 ppm.
- At a circulation cut the true gas crosses 4 ppm at 13.8 min and the analyzer does so only at 24.7 min.
- Ratio circulation protects a metered gas-rate change, an acid-gas-meter loading floor protects the rich circuit when the feed assay is stale, and differential pressure answers foam first.
- Ratio, loading floor, antifoam and analyzer trim together keep the worst sampled H_2S at 3.78 ppm, peak rich loading at 0.458 and average duty at 5.25 MW, where Fixed stops on corrosion.

Where else this applies. Chapter 7 applies the same ratio feedforward and analyzer trim to a single solute, and Volume 4, Chapter 2 holds a protective override that rests at its clamp.

References

[1] G. Versteeg and W. van Swaaij, "On the kinetics between CO2 and alkanolamines both in aqueous and non-aqueous solutionsii. tertiary amines," *Chemical Engineering Science*, vol. 43, no. 3, pp. 587–591, 1988.

[2] R. J. Littel, W. P. M. Van Swaaij, and G. F. Versteeg, "Kinetics of carbon dioxide with tertiary amines in aqueous solution," *AIChE Journal*, vol. 36, no. 11, pp. 1633–1640, 1990.

[3] B. Mandal, A. Biswas, and S. Bandyopadhyay, "Selective absorption of H2S from gas streams containing H2S and CO2 into aqueous solutions of n-methyldiethanolamine and 2-amino-2-methyl-1-propanol," *Separation and Purification Technology*, vol. 35, no. 3, pp. 191–202, 2004.

[4] M. Bolhar-Nordenkampf, A. Friedl, U. Koss, and T. Tork, "Modelling selective H2S absorption and desorption in an aqueous MDEA-solution using a rate-based non-equilibrium approach," *Chemical Engineering and Processing: Process Intensification*, vol. 43, no. 6, pp. 701–715, 2004.

[5] R. Deshmukh and A. Mather, "A mathematical model for equilibrium solubility of hydrogen sulfide and carbon dioxide in aqueous alkanolamine solutions," *Chemical Engineering Science*, vol. 36, no. 2, pp. 355–362, 1981.

[6] D. M. Austgen, G. T. Rochelle, X. Peng, and C. C. Chen, "Model of vapor-liquid equilibria for aqueous acid gas-alkanolamine systems using the electrolyte-NRTL equation," *Industrial & Engineering Chemistry Research*, vol. 28, no. 7, pp. 1060–1073, 1989.

[7] 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.

[8] U.S. Geological Survey, "Sulfur," in *Mineral Commodity Summaries 2025*. Reston, VA: U.S. Geological Survey, 2025, pp. 172–173.

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

[10] G. T. Rochelle, "Amine scrubbing for CO2 capture," *Science*, vol. 325, no. 5948, pp. 1652–1654, 2009.

[11] B. Delfort, A. Nicola, T. Huard, A. Wender, C. Lefebvre, N. Laloue, and K. Gaillard, "Polyamines, synthesis process thereof and use of same for the selective elimination of H2S from a gaseous effluent comprising CO2," U.S. Patent 10,710,021 B2, 2020. [Online]. Available: https://patents.google.com/patent/US10710021B2/en

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

[13] D. E. Rivera, M. Morari, and S. Skogestad, "Internal model control: PID controller design," *Industrial & Engineering Chemistry Process Design and Development*, vol. 25, no. 1, pp. 252–265, 1986.

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

[15] 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.

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

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

[18] X. Li, L. Shen, Q. Luo, L. Zhou, F. Yang, Z. He, and W. Yang, "Standard methods and related techniques for on-line detection of the total sulfur content in natural gas," *Natural Gas Industry B*, vol. 5, no. 2, pp. 99–104, 2018.

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

[20] U.S. Environmental Protection Agency, "40 CFR 60.646, monitoring of emissions and operations, standards of performance for SO2 emissions from onshore natural gas processing," Code of Federal Regulations, Title 40, Part 60, Subpart LLL, edition of July 1, 2025, govinfo.gov, 2025.

[21] ——, "40 CFR 60.642, standards for sulfur dioxide, standards of performance for SO2 emissions from onshore natural gas processing," Code of Federal Regulations, Title 40, Part 60, Subpart LLL, edition of July 1, 2025, govinfo.gov, 2025.

[22] M. Yusop Nurida, D. Norfadilah, M. R. Siti Aishah, C. Zhe Phak, and S. M. Saleh, "Monitoring of CO2 absorption solvent in natural gas process using Fourier transform near-infrared spectrometry," *International Journal of Analytical Chemistry*, vol. 2020, p. 9830685, 2020.

[23] A. Szabó, Á. Mohácsi, G. Gulyás, Z. Bozóki, and G. Szabó, "In situ and wide range quantification of hydrogen sulfide in industrial gases by means of photoacoustic spectroscopy," *Measurement Science and Technology*, vol. 24, no. 6, p. 065501, 2013.

[24] Q. Chen, X. Peng, J. Xue, J. Li, and C. Pan, "Identification and quantification of organic contaminants and evaluation of their effects on amine foaming in the natural gas sweetening industry," *ACS Omega*, vol. 7, no. 50, pp. 46 421–46 427, 2022.

8.9 Exercises

8.1 ⌕ Concept check · 2 min

In case 8 the steam available to the reboiler falls to 60 % at 30 min, and a steady sweep at that duty finds no circulation from 30 to 110 m^3/h that holds 4 ppm of H_2S. Why does neither more nor less circulation restore the specification?

(a) The analyzer reports each sample 5 min late, so every circulation move reaches the absorber after the gas has left.

(b) The amine pump is already at its 110 m^3/h capacity, so the circulation cannot be raised any further.

(c) The regenerator loop winds up against the open steam valve and holds its temperature setpoint at the upper bound.

(d) Less circulation loads each mol of amine closer to its capacity, while more circulation gets less stripping steam per mol and returns richer lean amine.

8.2 ⌕ Concept check · 2 min

In case 4 at 80 m^3/h, doubling the tray liquid holdup raises the share of feed CO_2 absorbed from 30.5 to 53.0 %, while the sweet-gas H_2S moves only from 2.85 to 2.25 ppm. Which statement explains the difference?

(a) Each stage takes H_2S a fixed 95 % of the way to equilibrium whatever its holdup, while CO_2 uptake grows with the amine inventory on the stage.

(b) More holdup cools the stages, and the larger heat of absorption of CO_2 makes its equilibrium pressure fall faster than that of H_2S.

(c) The extra liquid is drawn from the surge tank, so the circulation falls and the lean amine has less capacity left for H_2S.

(d) All the H_2S is absorbed on the bottom stage, so no change above it can move the H_2S leaving the top.

8.3 Concept check · 1 min

With automatic antifoam in case 6, the instantaneous true H_2S peaks at 4.40 ppm between two analyzer samples, while the analyzer never reports more than 2.70 ppm. Why can a run with such an excursion still meet the H_2S clause of the twelve-hour challenge?

(a) The challenge grades the analyzer's reports, and the trim holds them 1 ppm under the limit.

(b) The challenge grades the true H_2S only at the analyzer's sample instants.

(c) The challenge grades the average H_2S over the twelve hours, not its peak.

(d) The challenge forgives one sample over 4 ppm for each antifoam shot.

8.4 Theory and calculation · 7 min

The Ratio station sets the circulation from the gas meter and the laboratory's feed analysis, Equation (8.8), with $r = 1.29\,\text{mol/mol}$. Take the hydrocarbon gas at 17 kg/kmol (Table 8.2), H_2S at 34.08 kg/kmol and CO_2 at 44.01 kg/kmol.

(a) Compute the amine concentration c_A, the molar mass of the sour gas at 2 % H_2S and 4 % CO_2, and its molar flow at 60 t/h. *Checkpoint:* 3.927 kmol/m^3, 18.42 kg/kmol and 3257 kmol/h.

(b) Compute the acid-gas load and the Ratio circulation at design and after the 20 % gas step of case 2. Compare with the 63.9 and 77.1 m^3/h the chapter reports. *Checkpoint:* 195.4 kmol/h, 64.2 m^3/h and 77.0 m^3/h.

(c) In case 5 the feed H_2S rises to 2.6 % while the laboratory still reports 2 %. Compute the true acid-gas load, the mol ratio the unchanged circulation now gives, and the circulation that would restore 1.29 mol/mol.

8.5 Theory and calculation · 6 min

The equilibrium of Equation (8.1) can be checked against the three patent points of Section 8.2 for H_2S alone ($\alpha_C = 0$, so $\alpha_t = \alpha_H$).

(a) Compute p_H^* at 40 °C for loadings 0.21, 0.40 and 0.69 with $K_H = 245\,\text{kPa}$ and $b = 2.1$. *Checkpoint:* 8.80, 28.2 and 88.4 kPa.

(b) The amine concentration scales the equilibrium pressure, so multiply by 50/45 to compare with the patent's 10, 30 and 100 kPa for 50 % MDEA. Give the three errors. *Checkpoint:* 9.78, 31.34 and 98.17 kPa.

(c) In case 3 the lean amine warms from 40 to 45 °C with the gas. By what factor does p_H^* rise at a fixed loading, and what does that do to the sweet gas?

8.6 Theory and calculation · 6 min

In every case the analyzer draws a sample at 9.5, 19.5, 29.5 min and every 10 min after, and reports it 5 min later (0.5 min of sample line and 4.5 min of analysis), at 14.5, 24.5, 34.5 min and so on. The displayed value holds until the next report.

(a) For an exceedance that begins at an arbitrary time and persists, give the shortest and the longest delay until a report shows it. *Checkpoint:* 5 min and 15 min.

(b) The true gas passes 4 ppm at 13.8 min in case 1, 62.3 min in case 3 and 59.3 min in case 8. Predict when the analyzer first reports it in each case, and the delay.

(c) In case 6 the slug arrives at 30 min, the pressure-drop alarm sounds at 30.167 min and, without antifoam, the scrubber stops the run at 44.98 min. When does the first analyzer report of the upset arrive, and what is it worth to the operator?

8.7 App challenge · 8 min

Load scenario 1, "Manual: cut the amine circulation", default seed. Before Run, set the Amine pump (FIC card) to 56 m^3/h. The plant starts from the design steady state and the pump moves at $t = 0$. Run at speed 20 (each simulated minute takes 0.05 s) to 300 min or until the run stops, then export the CSV.

(a) Does the run end in a stop? Read the `exercise_stop` column, the time of the last row, and the first row where `truth_rich_loading` exceeds 0.5. *Checkpoint:* a corrosion stop at 104.1 min, cap first passed at 44.2 min (both ± 0.2 min).

(b) Press Reset, which keeps your settings, set the pump to 57 m^3/h and run to 300 min. Read `truth_h2s_ppm` and `truth_rich_loading` in the row at 300 min, and the largest `h2s_analyzer_ppm` of the run.

(c) Explain why the corrosion stop in (a) comes exactly when it does, and why 57 m^3/h is the least whole-number setting with no margin on either limit.

8.8 App challenge · 9 min

Load scenario 5, "The loading cap on a sourer gas", default seed. Before Run, turn the Loading limit On (FIC card) and leave its setpoint at 0.45. Run at speed 20 to 240 min and export the CSV.

(a) Read the first row with `loading_limit_active` equal to 1, the largest `truth_rich_loading` and the first later row where it is back under 0.5. *Checkpoint:* 45.8 min, 0.506 and 51.3 min.

(b) The true loading passed 0.5 at 38.5 min. Read `rich_loading_inferred` in that row and explain why the floor engaged seven minutes later.

(c) Press Reset, set the Loading limit setpoint to 0.42 and run again to 240 min. Read the largest `truth_rich_loading` and `circulation_m3h` at 240 min, and compare the circulation with the first run's.

8.9 App challenge · 9 min

Load scenario 9, "Challenge · a twelve-hour shift", default seed. Use speed 300, so that 720 min takes about 2.4 s. The run stops at 12 h or at a stop, and its verdict appears in What changed?

(a) Run as shipped. What ends the run, at what time, and which clauses of the target fail? *Checkpoint:* a corrosion stop at 214.5 min.

(b) Reload the scenario. Before Run, set Circulation control to Ratio, the Loading limit On with setpoint 0.42, Antifoam on high ΔP to Auto and the Analyzer trim On. Run to 12 h and read the verdict's worst sample, peak rich loading and average duty.

(c) Press Reset, turn the Analyzer trim Off and run again. Read the same three values and explain how a loop that raises the regenerator temperature when the gas is too sour saves steam.

8.10 Application area · 3 min

In case 7, on a cold night, someone lowers the lean amine setpoint from 40 to 30 °C because colder amine absorbs H_2S better.

(a) What happens over the next hour, and how does it end if nobody acts?

(b) Why does a sweetener hold its lean amine a few kelvin above the sour gas temperature, and what does that margin cost in absorption?

8.11 Application area · 4 min

In case 4 at 80 m^3/h, doubling the tray holdup raises the feed CO_2 absorbed from 30.5 to 53.0 % and the reboiler duty from 4.84 to 5.81 MW. The feed is 3257 kmol/h of gas with 4 % CO_2.

(a) Compute the extra CO_2 absorbed and the heat to desorb it at 55 MJ/kmol, and compare it with the rise in duty. *Checkpoint:* 29.3 kmol/h and 0.45 MW, about half of 0.97 MW.

(b) Name two costs of coabsorbed CO_2 for the plant, and two ways a selective unit limits it.

8.12 Review the agent · 5 min

An AI assistant reviews case 5 and writes: “The analyzer trim already protects the rich line. When the sourer gas overloads the amine, the H_2S slip rises, the trim raises the regenerator temperature, the leaner amine regains capacity, and the rich loading falls back under 0.5. A separate loading limit only duplicates the trim and costs circulation.” Find the flaw, and say how you would check it in the app.

CHAPTER 9

Ethane Recovery and Rejection Through a Sequenced Turboexpander Plant

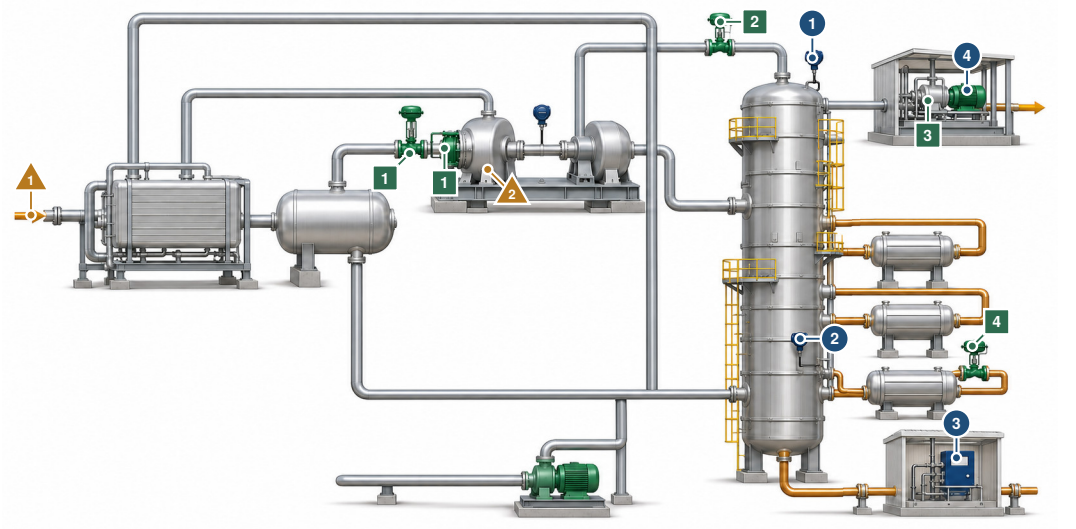

The controller moves 1 expander guide vanes, with the valve beside them that takes the rest 2 reflux split valve 3 residue compressor, which holds the column pressure 4 trim heat to the bottom reboiler

The controller reads 1 column pressure 2 control tray temperature 3 methane in the liquid product, reported minutes late 4 compressor driver power and its limit

What disturbs it 1 a richer or larger inlet gas 2 an expander trip ▲ a fall in the ethane price

Notice the one shaft that joins the expander to the booster, and the analyzer shelter at the far end of the liquid product line, the last place the methane shows. Illustration, not to scale.

IN PLAIN TERMS Pulling Liquids Out of Natural Gas

The main idea. A gas plant chills natural gas until ethane and its heavier parts turn to a liquid it can sell. A spinning machine called an expander makes the deepest cold, and a tall column boils the last methane out of the liquid.

The challenge. When ethane is worth more in the sales gas, the plant warms the column and raises its pressure to let the ethane go. If the pressure rises first, methane spoils the liquid, and the instrument that checks it reports minutes late.

Why it matters. The same two moves can keep or ruin a day's product, depending only on their order.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History boxes in Section 9.1 and Section 9.2, the Did you know? boxes in Sections 9.3 and 9.5, the Try it boxes in Section 9.5, and Key points at the end of the Conclusion. No equations.
Study	40 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 9.5 with Figure 9.2, Section 9.6 and Section 9.7 with its Open challenge box.

The control question When ethane is worth less as a liquid than as gas, a gas plant moves its demethanizer from ethane recovery to rejection and back, while a chromatograph reports the methane in its liquid product minutes late. In what order must the column pressure and the tray temperature move, and what does the move cost?

Learning objectives

- Explain how one expander shaft and the feed-heated reboilers tie refrigeration, recompression and column temperature together, and compare expander and valve cooling.
- Compute the frac spread and the margin of a mode from prices, and read the pressure map as recovery traded against recompression power.
- Explain why a tray temperature held on its setpoint stops meaning a methane specification, and how a delayed chromatograph trims it.
- Predict which setpoint must move first in a mode switch, and show from runs that the order and the return margin are both needed.
- Judge a controller by every limit of a run, not by its setpoint tracking or its margin alone.

Prerequisites Sections 1.4.7 (expansion cooling), 1.4.9 (late and inferred measurements), 1.6.2 (the framework's PI) and 1.6.4 (the apps). Section 1.6.3 is needed only for the predictive comparison. Chapter 8 introduces the analyzer trim.

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

See also Volume 8, Chapter 13 for a supervisor that holds a mode as a latched state, Volume 8, Chapter 9 for late results matched to the product they describe, and Volume 8, Chapter 3 for a composition loop that holds its setpoint while the product misses. Volume 9, Chapter 8 for setpoints chosen from a cost model, and Volume 9, Chapter 10 for a slow quality trim on a setpoint loop, priced in energy.

9.1 Introduction

A single steady recovery number hides the hardest moment in a natural-gas-liquids (NGL) plant. When ethane's value as liquid falls below its value in the residue gas, the desired column pressure and bottoms temperature change together. The NGL must still meet its methane limit while cold metal, column inventory and a shaft continue moving. A static process diagram cannot show whether a pressure move arrives before the reboiler's effect or whether a chromatograph reports an earlier inventory. Dynamic studies of NGL recovery and demethanizer control establish why the process has coupled, delayed responses [1, 2].

The plant of this chapter joins the economic choice to its physical consequence. Its cold box preheats the column's reboilers, so an action that changes the column returns to the feed temperature. The expander's refrigeration simultaneously drives a residue booster. Column pressure trades ethane recovery against recompression power, and a warmer control tray strips methane but also rejects ethane. Pressure, tray temperature and reflux therefore cannot be assigned independent economic values during a transition. Published optimization studies have shown the pressure and heat-integration tradeoffs in turboexpander processes [3–5]. The simulator of this chapter adds a reproducible dynamic experiment in which the order of two ordinary setpoint moves is a measurable constraint, with the economics and the source of each measurement visible. Section 1.4.7 gives the expansion cooling on which the plant rests, and Table 1.2 places its limits beside those of the other plants of the volume.

This chapter explains how the synthetic turboexpander plant's sampled NGL specification depends on the mode-switch sequence. Section 9.2 places the process in its technical history. Sections 9.3 and 9.4 give the model and the controls: the thermodynamics, stages, exchangers and economics, then the two regulatory loops, the analyzer trim, the mode sequencer, the economic supervisor and the predictive comparison. Section 9.5 reports the nine measured cases, ending with the twelve-hour shift. Section 9.6 sets the hardware of a real gas plant beside the model, Section 9.7 states what the model cannot represent, and Section 9.8 summarizes the measured lessons.

> ⧖ HISTORY **1964: Expansion That Pays for Its Own Recompression**
>
> On 28 April 1964 an engineer in Houston applied for a patent on treating high-pressure natural gas that had to leave the plant at a lower pipeline pressure. Feed gas already cooled well below the freezing point of water was expanded through a turbine of a kind that tolerates the liquid it condenses, down to a stripping pressure below the line pressure, chosen for the best liquid recovery. The turbine's own work then drove the compressor that restored the gas to line pressure, and that compression load steadied the turbine. On its way to the compressor, the cold gas from the stripper precooled the incoming feed [6]. The expander, the booster on its shaft and the feed exchanger of this chapter's plant are those three pieces.

9.2 Process Background and Prior Work

Early cryogenic NGL recovery relied on cooling and expansion to condense heavy components, followed by a demethanizer that rejected methane overhead. The gas-subcooled process (GSP) added a subcooled reflux branch to increase ethane capture relative to a single expander arrangement, and the History box below tells of the patent that split the cold gas this way. Honeywell UOP describes GSP as a high-recovery standard solution and labels its ethane recovery band 80–92%, its recovery-mode propane band 92–100%, and its ethane-rejection option as partial [7]. These are licensor categories, not measurements on the present model. The brochure defines partial ethane rejection by the ethane share of the liquid product, so its category is not itself a target of 8.1% recovery.

> ⧖ HISTORY **1977: Splitting the Cold Gas for a Reflux**
>
> In a patent application of 14 July 1977, two engineers named the limit of the expander plants of their day: the demethanizer worked largely as a stripping column, so the ethane lost overhead depended on the column's top temperature at a given pressure. Their answer split the cold gas before expansion. One part was cooled until it had substantially condensed, then expanded and fed above the rest, while the rest went through the expander to a feed point lower in the column. The split also diluted the carbon dioxide in the rising vapor and lowered its risk of freezing, so a colder top, and more ethane, became practical. The patent was granted on 14 July 1981 [8]. The reflux split and subcooler of this chapter's plant follow the same arrangement.

The separation and its power draw remained objects of process optimization. Chebbi and coauthors varied demethanizer pressure and feed in a conventional turboexpander study in 2010 [3]. Later comparisons examined reflux and alternative flowsheets [4, 5], while exergy studies allocated loss among expansion, heat exchange and recompression [9, 10]. This work motivates showing the energy cost alongside recovery instead of asking for maximum recovery alone. The 2013 demethanizer control study and later dynamic NGL study addressed interacting column and plant controls [1, 2]. A 2022 separation-train study compared a cascade with boilup approximation and pressure compensation with conventional temperature control under feed disturbances [11]. Their models and aims differ from this teaching plant, but together they establish that an inferential tray target and a measured product composition need not respond on the same clock.

The simulator's four pure-component reference points use the methane, ethane, propane and normal-butane boiling and critical data compiled in the NIST Chemistry WebBook [12]. The model deliberately does not use a production equation of state or a calibrated binary-interaction matrix. Matching a licensor recovery band is only a scale check for its selected feed and equipment settings, and Section 9.3.2 reports the sensitivity to exchanger-cell count.

9.3 Plant Model

9.3.1 Flow Arrangement, States and Measurements

The feed is 100 MMscfd, represented as 4980 kmol/h at 80 bar absolute and 40 °C. Its mole fractions are $(0.900, 0.055, 0.025, 0.020)$ for methane, ethane, propane and a normal-butane surrogate for heavier hydrocarbons. A fraction $s = 0.27$ of the feed passes E-2, where it gives heat to two side reboilers and the feed-heated bottom reboiler. The rest passes the counter-current E-1 exchanger

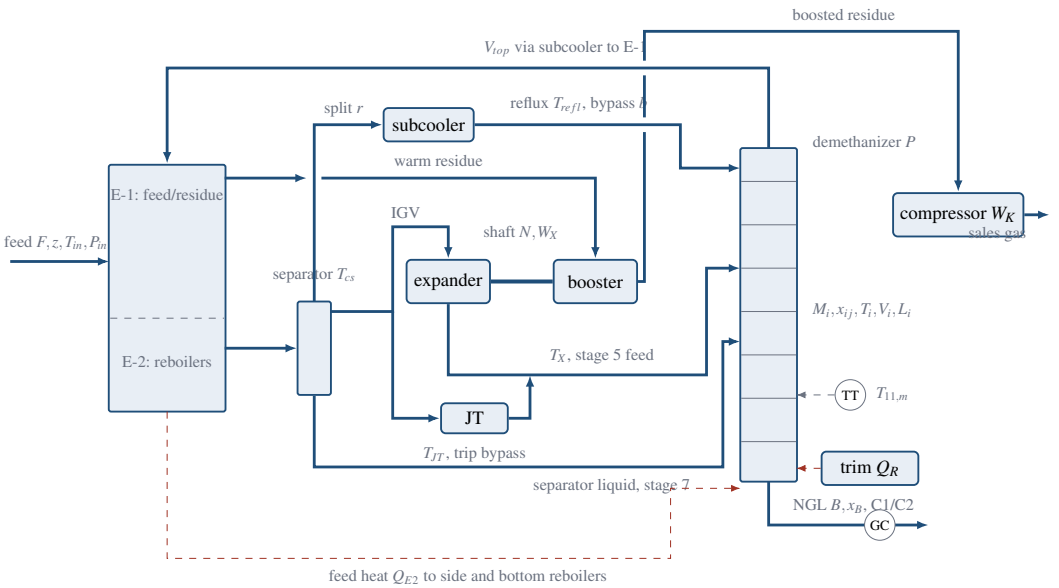

Figure 9.1 Teaching model of the gas-subcooled recovery plant. The expander and residue booster share a shaft. Solid arrows are material paths. Dashed links denote heat or work coupling and chromatograph information. E-2 returns feed heat to the side and bottom reboilers. The column has twelve equilibrium stages. The marked feeds and measured stage 11 are the key control locations. This is a model diagram, not an equipment piping specification.

against cold column overhead. The portions reunite in a cold separator. Separator vapor divides into subcooled reflux and a path through the expander or its parallel Joule–Thomson (JT) valve. Separator liquid enters the lower column. Column overhead returns through the subcooler and E-1 before a shaft-driven booster and a residue compressor deliver gas to the sales line. NGL leaves the sump, and Chapter 10 splits such a liquid into products. Figure 9.1 follows the application's Plant view.

The integrated state has 127 quantities: each of twelve stages stores total liquid and three independent mole fractions, the sump stores four quantities, E-1 has 48 metal temperatures, E-2 has three, and the reflux subcooler has 16. Shaft speed, column pressure, delivered trim duty and a flooding lag add four states, with four integrated component ledgers completing the vector. The engine of Section 1.6.1 integrates it with a two-second fourth-order Runge–Kutta physics step, a ten-second control and history interval, and ten-minute chromatograph cycles. It has no gas-side dynamic holdup beyond the one column-pressure inventory. The screens distinguish measured values from optional simulation truth.

9.3.2 Declared Thermodynamics and Stage Balances

For component j, a two-constant vapor-pressure line is fixed by its normal boiling point $T_{b,j}$ and critical point $(T_{c,j}, P_{c,j})$. Ideal-solution K values approach unity as pressure approaches the declared convergence pressure P_k:

$$B_j = \frac{\ln(P_{c,j}/P_b)}{1/T_{b,j} - 1/T_{c,j}}, \quad A_j = \ln P_{c,j} + \frac{B_j}{T_{c,j}},$$
$$\ln P_j^{sat} = A_j - \frac{B_j}{T}, \qquad \ln K_j = \left(1 - \frac{P}{P_k}\right) \ln \frac{P_j^{sat}}{P}. \tag{9.1}$$

Here $P_b = 1.013\,25\,\text{bar}$ and $P_k = 140\,\text{bar}$. A stage bubble point solves $\sum_j x_j K_j = 1$ in ten Newton iterations. The methane line above its critical temperature is a hypothetical volatility, not a physical

saturation curve. The four reference-point pairs reproduce P_b at each T_b and the corresponding P_c at each T_c to machine precision. This does not validate a mixture near its critical region.

The gas stream uses a common $c_{p0} = 0.036\,\mathrm{MJ/(kmol\,K)}$ and a declared nonideal volume $v = RT/P - A/T^2$. Its consistent enthalpy and entropy are

$$\begin{aligned} h_V &= c_{p0}(T - T_0) - 3A(P - P_0)/T^2, \\ s_V &= c_{p0}\ln T - R\ln P - 2A(P - P_0)/T^3. \end{aligned} \tag{9.2}$$

The enthalpy reference is $T_0 = 273.15\,\mathrm{K}$ at $P_0 = 1\,\mathrm{bar}$ in Eq. (9.2), distinct from the normal-boiling reference pressure in Eq. (9.1). The constant is $A = 420$ in the application's MJ, kmol, K and bar units. The liquid enthalpy subtracts $\sum_j x_j \lambda_j(T)$ from vapor enthalpy, and the two-phase enthalpy subtracts $(1 - \psi)$ times that sum. The latent heat at each normal boiling point is RB_j and follows a Watson exponent 0.38, with reduced temperature capped at 0.97. An isothermal Rachford–Rice flash supplies vapor fraction ψ, liquid and vapor compositions. These closures are small enough to inspect but do not include heats of mixing or composition-dependent gas heat capacities.

The expander's isentropic reference outlet satisfies $s_V(T_{2s}, P_2) = s_V(T_1, P_1)$. With efficiency $\eta_X = 0.87$, the actual specific work is $\eta_X[h_V(T_1, P_1) - h_V(T_{2s}, P_2)]$. Equation (1.13) is its ideal-gas counterpart. The JT outlet instead solves $h(T_{JT}, P_2) = h_1$. For the same separated vapor and 23 bar absolute destination, the model gives $-70.9\,^\circ\mathrm{C}$ after expansion and $-49.2\,^\circ\mathrm{C}$ after throttling. The measured flash enthalpy difference on the JT path is below 10^{-6} MJ/kmol. Shaft speed N changes as expander power less booster work and bearing loss is stored in $J\omega_0^2 N\dot{N}$. Booster head follows $h_{B0}N^2$.

The twelve equilibrium stages use the shared stage-stack component balance, with a stage-specific energy correction to the vapor flow. The subcooled reflux enters stage 1, expanded gas stage 5, separator liquid stage 7, and the control tray is stage 11. For component j in stage i, the stored liquid balance is

$$\begin{aligned} \frac{d(M_i x_{ij})}{dt} &= L_{i-1}x_{i-1,j} + V_{i+1}y_{i+1,j} \\ &\quad + \sum_{k \in i} F_k z_{k,j} - L_i x_{ij} - V_i y_{ij}. \end{aligned} \tag{9.3}$$

where flows in kmol/h are divided by 60 on the minute clock. The stage energy relation adjusts V_i for incoming vapor and liquid sensible heat, feed quality and side heat. Setting that correction to zero would remove an important path by which expander refrigeration and reflux subcooling affect the column. The upper and lower side reboilers warm liquid leaving stages 8 and 10, while the bottom reboiler receives feed heat and a separately manipulated trim duty Q_R. Constant stage liquid holdup is 15 kmol and sump target 80 kmol. The dynamic sump level regulator changes NGL withdrawal, so instantaneous component recovery can exceed 100% while inventory drains.

For clarity, the stage heat correction used in Eq. (9.3) is

$$\begin{aligned} \varepsilon_i &= \frac{1}{\lambda_i}\big[V_{i+1}\{h_V(T_{i+1}) - h_V(T_i)\} \\ &\qquad + L_{i-1}(h_{L,i-1} - h_{L,i}) + Q_i\big], \\ V_i &= V_{i+1} + \sum_{k \in i}(1 - q_k)F_k + \varepsilon_i, \\ q_k &= \{h_V(T_i) - h_{F,k}\}/\lambda_i. \end{aligned} \tag{9.4}$$

The liquid and vapor cascades are recalculated in four passes. Side-reboiler duties are evaluated in the second pass, and the subsequent passes update sensible-heat correction and sump boilup. A floor of 5 kmol/h prevents a stage from condensing its entire upward vapor stream, while a nonnegative liquid-flow bound prevents evaporating more liquid than arrives. The sump boilup also has a dry-inventory bound proportional to its liquid holdup. These numerical and physical bounds are included in the declared parameter defaults.

Table 9.1 Plant Parameter Defaults. V is shown in the parameter panel. H is hidden. All values are teaching choices. The key identifies the corresponding source parameter.

Key	Meaning	Value	Unit	Panel
`Tin`	Inlet gas temperature	40	°C	V
`Pin`	Inlet and cold separator pressure	80	bar a	V
`split`	Feed share through the reboilers (E-2)	0.27		V
`UAh1`	E-1 conductance, feed side	5	MW/K	H
`UAc1`	E-1 conductance, residue side	5	MW/K	H
`Cm1`	E-1 metal heat capacity	60	MJ/K	H
`UAh2`	E-2 conductance, feed side	1.2	MW/K	H
`UAc2`	Side reboiler conductance, column side	0.3	MW/K	H
`UAcB`	Feed-heated bottom reboiler conductance	0.4	MW/K	H
`Cm2`	E-2 metal heat capacity per reboiler	10	MJ/K	H
`UAsc`	Reflux subcooler conductance	1	MW/K	H
`Csc`	Reflux subcooler metal heat capacity	12	MJ/K	H
`etaExp`	Expander isentropic efficiency	0.87		V
`hB0`	Booster isentropic head at design speed	0.833	MJ/kmol	H
`etaB`	Booster efficiency	0.78		H
`Jshaft`	Shaft inertia (J ω_0^2)	15	MJ	H
`fShaft`	Bearing and seal losses (share of design power)	0.02		H
`Wexp0`	Design expander power	1.5	MW	H
`Mtray`	Liquid held on each stage	15	kmol	H
`Msump`	Sump holdup at its level setpoint	80	kmol	H
`kLevel`	Sump level loop gain	60	1/h	H
`kDry`	Boilup limit per kmol in the sump	400	1/h	H
`Vmin`	Least vapor a stage sends up	5	kmol/h	H
`Kflood`	Tray capacity at flooding	14650	kmol/h sqrt (K/bar)	V
`kE`	Efficiency lost per unit of flooding	8		H
`Emin`	Least stage efficiency when flooded	0.4		H
`tauFl`	Flooding onset lag	1	min	H
`tauReb`	Trim reboiler response	2	min	H
`Qmax`	Trim reboiler capacity	3	MW	V
`specC1`	Bottoms C1/C2 limit	2	mol %	V
`CP`	Gas held per bar between the column and the compressor	15	kmol/bar	H
`dPres`	Pressure drop through the cold box	1	bar	H
`dPac`	Pressure drop through the booster aftercooler	0.5	bar	H
`dTac`	Aftercooler approach to the cooling water	10	K	H
`etaC`	Residue compressor efficiency	0.78		H
`Fc0`	Compressor flow at full capacity and reference suction	6000	kmol/h	H
`Psuc0`	Reference suction pressure	34	bar a	H
`Tsuc0`	Reference suction temperature	40	°C	H
`Wmax`	Residue compressor driver limit	4.2	MW	V
`Psales`	Sales line pressure	70	bar a	H
`cp0`	Ideal-gas heat capacity	0.036	MJ/(kmol K)	H
`Agas`	Real-gas constant A (v = RT/P − A/T^2)	420	MJ K^2/(kmol bar)	H
`Pk`	Convergence pressure of the K values	140	bar a	H

For a hot-stream cell, the implicit outlet temperature solves

$$F_h(h_{in} - h_{out}) = UA_h(T_{out} - T_m), \qquad C_m\dot{T}_m = Q_h - Q_c. \tag{9.5}$$

Flows are converted from hourly units before calculating MW, and the metal derivative is converted to the minute clock. The cold-side enthalpy calculation runs in the opposite direction through the metal nodes. E-2 uses a flash of each column-side liquid, so boiling changes its temperature rather than being represented by a constant liquid heat capacity. The feed enthalpy is interpolated on a flash-based temperature table that is regenerated when composition or pressure changes. Table 9.2 identifies the physical variables separately from the source keys in Table 9.1.

Each exchanger cell moves enthalpy between a stream and a metal state. E-1 has 48 cells, the reflux subcooler 16, and E-2 has three metal states. The conductances and heat capacities in Table 9.1 are generic. The cell count matters: the same design calculation gives 84.5% ethane recovery with 24/8 E-1/subcooler cells, 88.8% with 48/16 and 90.8% with 96/32. Thus the reported 88.8% is a result at a declared resolution, not a converged property prediction.

Table 9.2 Variables Used in the Plant and Control Model

Symbol	Unit	Meaning
F, z_j	kmol/h, –	Inlet flow and composition
M_i, x_{ij}, y_{ij}	kmol, –	Stage liquid and phase compositions
L_i, V_i	kmol/h	Stage liquid and vapor flows
T_i, T_m	K	Fluid and exchanger metal temperatures
P	bar absolute	Column top pressure
$B, x_{B,j}$	kmol/h, –	NGL flow and composition
N	–	Shaft speed relative to design
W_X, W_K	MW	Expander and residue-compressor power
Q_R, Q_{E2}	MW	Trim and feed-supplied reboiler heat
r, b	–	Reflux split and subcooler bypass
c_m	mol %	Reported NGL methane/ethane ratio
m	\$/h	Incremental processing margin

Column pressure accumulates top vapor minus compressed residue with capacity $C_P = 15\,\text{kmol/bar}$. The residue compressor flow scales with its commanded capacity and suction density until the calculated power reaches its 4.2 MW driver limit. Vapor load is the maximum of $V_i\sqrt{T_i/P}/K_f$ across stages, with $K_f = 14650$ in the code's units. Exceeding one initiates a one-minute flooding lag and reduces effective stage efficiency. The 2% bottoms C1/C2 limit, power limit and flooding threshold are intentionally declared teaching values. Real thresholds come from a column, compressor and product specification for a particular installation.

9.3.3 Economics and Numerical Checks

For NGL component flows B_j, the processing margin is the incremental value of liquids relative to leaving those molecules in the gas:

$$m = \sum_{j=2}^{4} B_j(p_{L,j} - p_G H_j) - p_W W_K - p_Q Q_R. \tag{9.6}$$

Here $p_{L,j}$ is the liquid price converted from cents per gallon to dollars per kmol, p_G is dollars per MMBtu, H_j is gross heating value, and power and heat are charged in dollars per MWh. Methane in the NGL earns no separate liquid credit. Base prices are 25, 75 and 100 cents per gallon for ethane, propane and the butane lump, gas is 3 dollars per MMBtu, power 60 dollars per MWh and trim heat 12 dollars per MWh. These are scenario inputs. The ethane frac spread, liquid price less heating-value opportunity cost, is $+5.12$ cents per gallon at the base prices and -9.88 when ethane falls to 10 cents per gallon.

> **DID YOU KNOW? Ethane Leads the Liquids**
>
> In 2025 the natural gas processing plants of the United States produced an average of 7.5 million barrels a day of natural gas liquids, and 3.1 million of them were ethane, more than any other liquid [13]. Ethane is also the one liquid a plant like this chapter's can choose to leave in its sales gas.

The initial condition is a nonlinear steady solve at the selected mode's pressure and tray setpoints. Changing a starting control before Run solves a different initial condition, while changing it during a run applies a step. Reset and the reload button behave as Section 1.6.1 describes, and Reset solves the start again with the reader's choices. The load-path test confirms that each of the nine cases gives the same CSV loaded first, selected after another case, or reset with its original settings.

The disturbed-run component ledgers compare net inlet minus products with the accumulated column and sump component inventories. The simulator's regression checks close every component to below 10^{-6} kmol for a trip, restart and feed change. A separate design energy check includes feed enthalpy, warm residue, NGL, expander work and trim heat and closes within 0.02 MW. This is an internal conservation check, not evidence for the selected constitutive laws. A two-to-one-second refinement after a reboiler step changes the reported temperature and recovery by less than 0.02 in their displayed units. The larger cell-count sensitivity reported above is a separate spatial approximation.

9.4 Regulation, Sequencing and Economic Supervision

9.4.1 Two Regulatory Loops and a Delayed Composition Trim

The pressure PI increases residue-compressor capacity when pressure exceeds its target. Its gain is 6% of rated capacity per bar, integral time 8 min, and command bounds 0–120%. The tray-temperature PI increases trim duty when stage 11 is colder than its target, using 0.02 MW/K, 15 min integral time and 0–3 MW bounds. Both use the framework's velocity-form controller of Section 1.6.2, with back-calculation and a tracking time equal to the integral setting. The header's Manual mode hands these two commands to the reader, and the inactive controllers track them for a bumpless return to Auto. Guide vanes, reflux split and subcooler bypass remain operator stations in either header mode. Ordinary PI tuning ideas apply locally [14], but the large operating-mode move also changes the heat integration and the appropriate temperature target.

The product chromatograph samples every 10 min and reports after a 6 min analysis delay at the default cycle. At each fresh recovery-mode report, the analyzer changes the tray target by $K_A(c_m - c_{SP})$, clipped to ± 3 K per report and a total ± 25 K. The default K_A is 1.5 K per percentage point per report and $c_{SP} = 1.5\%$, below the 2% limit. The trim is held during a switch, an expander trip or a compressor-limit condition. A high C1/C2 therefore warms the recovery tray only after the report arrives. The tray sensor is an inferential measurement in the sense of Section 1.4.9, not an instantaneous assay of the bottoms. Chapter 8 trims a setpoint from a sampled, delayed analyzer in the same way.

9.4.2 A Mode Is a State and a Procedure

Recovery owns pressure and stage-11 targets of 23 bar absolute and $-1.3\,^{\circ}$C. Rejection owns 27 bar absolute and $45.5\,^{\circ}$C. The interface names RECOVERY, REJECTION or SWITCHING with a step count. A mode is reached when the effective setpoints finish moving, even though the column inventory may still be settling. The flow-diversion supervisor of Volume 8, Chapter 13 is another mode that is held as a state with its own procedure.

The Loops architecture steps both setpoints when the requested mode changes. The Sequencer instead ramps one target over a default 20 min, then ramps the other over 20 min. Its Declared order moves tray temperature first in both directions. Reversed moves pressure first. A return to recovery sets a positive analyzer-trim margin before the ramp, so the first temperature move stops warmer than

the ordinary recovery target. The delayed analyzer subsequently works that margin down after the mode is reached. The default margin is 12 K, increased to 20 K in the suggested challenge.

This procedure is specific to the implemented dynamics. On the way to rejection, preheating the lower column before raising pressure prevents the methane excursion observed with the reverse path. On the return, neither the order nor the larger margin alone is sufficient in the tested shift. The successful path retains the warm margin while pressure and the richer-feed inventory change. The result supports that tested procedure, not a universal rule for real GSP plants.

The Supervisor evaluates a stored steady-state table every 15 min. It interpolates the lean and rich feed cases using inlet-chromatograph ethane, scales product and energy flows with metered feed, and prices each mode using Eq. (9.6). It requests the other mode only if its estimated advantage exceeds 50 dollars per hour, at least 120 min have passed since the last switch, and no switch is in progress. It then uses the same sequencer. Its estimate prices the destination steady states and has no prediction of transition loss. The supervisor of Volume 9, Chapter 8 likewise chooses setpoints for a regulatory layer from a cost model, there by constrained search.

9.4.3 Predictive Control as a Comparison

An optional linear MPC replaces the pressure and temperature PIs. It is a receding-horizon quadratic program, consistent with the industrial MPC class surveyed by Qin and Badgwell [15], and Section 1.6.3 gives the general setup. Its four-state model was identified near recovery at the design feed. Compressor capacity has pressure gain -0.326 bar per percent with a six-minute lag, tray gain -0.156 K per percent with a forty-minute lag, and power gain 0.0587 MW per percent with a two-minute lag. Trim duty has tray gain 57 K per MW and a twenty-two-minute lag.

The MPC executes each minute with a ninety-minute prediction horizon. It corrects its predicted outputs by the measured current residual, holds that bias over the horizon and carries the internal state forward with the commands actually applied. Default pressure and temperature weights are 10 and 1, while compressor-move suppression is 0.01 and duty-move suppression is 10 in their respective units. Input bounds match the two actuators. A soft power ceiling is 0.05 MW below the driver limit, with quadratic slack penalty 10^4. The model has no flooding or methane state. A failed or nonfinite solution, or a first move outside the box, returns control to the tracked PIs. This small local model fails the mode-switch challenge and is presented as a comparison, not its solution. The software context of the apps follows APMonitor and GEKKO [16, 17], while this plant executes its own local model and controller in the browser.

9.5 Cases Measured on the Simulation Engine

Table 9.3 lists the reader's sequence, and Section 1.6.4 describes how a case is loaded, run and exported. All values below were regenerated on the built engine at the default seed, with operator actions applied during the run at their stated times. Measured recovery is a ten-minute filtered ratio of measured NGL ethane to measured inlet ethane. True recovery uses instantaneous component flows. The difference matters while liquid inventory changes. C1/C2 is 100 times bottoms methane mole fraction divided by bottoms ethane mole fraction, not methane percent of total NGL.

9.5.1 Refrigeration and Pressure Tradeoffs

In case 1, the header is Manual. Closing guide vanes to 50% at 10 min reduces shaft work and leaves more separator vapor on the JT path. By 70 min, pressure is 28.07 bar absolute, shaft speed about 68% and cold-separator temperature $-18.4\,^{\circ}$C rather than the initial $-27.1\,^{\circ}$C. Measured ethane recovery is 41.3%, compared with instantaneous true recovery 32.6%. The measured value first falls below 60% at 39.7 min. Restoring full guide vanes at 70 min brings it back above 80% at 153.5 min. The warm separator is evidence of the reboiler/feed thermal loop, while the delayed recovery is evidence of both inventory and filtering.

Case 2 separates pressure from the expander station. At fixed NGL C1/C2 of 1.5%, the stored steady map has the values in Table 9.4. The maximum base-price margin within the driver capacity is

Table 9.3 Nine Cases and Their Main Operator Action

No.	Case	Action or comparison
1	Expander	Close guide vanes, then restore
2	Pressure map	Lower and raise column pressure
3	Specification	Enable delayed analyzer trim
4	Reflux	Change reflux split or bypass
5	Trip	Restart the expander after the trip
6	Flooding	Raise rejection pressure at high feed
7	Switching	Compare the two setpoint orders
8	Economics	Let the supervisor respond to prices
9	Challenge	Meet the twelve-hour target

Table 9.4 Steady Pressure Map at Base Prices and 1.5% C1/C2

P (bar absolute)	C2 recovery (%)	Power (MW)	Margin ($/h)
21.0	97.9	4.15	3147
21.5	96.5	4.03	3150
23.0	88.8	3.74	3138
25.0	76.4	3.39	3113

near 21.5 bar absolute, not at maximum ethane recovery. At 21 bar the compressor uses 4.15 MW, close to its 4.2 MW ceiling. A running pressure reduction to 21.5 bar at 2 min yields 96.1% measured recovery at 92 min. Raising pressure to 25 bar at that instant gives a chromatograph report of 3.17% C1/C2 at 122 min until analyzer trim compensates. A pressure map with a fixed composition target is therefore not the transient path of an untuned column.

9.5.2 Inferential Quality, Reflux and Equipment Limits

Case 3 holds tray temperature while the feed becomes richer at 20 min. With analyzer trim Off, the chromatograph exceeds 2% at 46 min, and true C1/C2 reaches 3.33% at 150 min. Switching trim On then brings the reading to 1.8% or less from 196 min. The warmer setpoint costs recovery: the measured value changes from about 88.9% to 66%. A controlled tray temperature has not guaranteed a controlled product composition.

Pause. Try Exercise 9.3 before reading on.

> **TRY IT IN THE APP** **3 · The methane specification: tray temperature and analyzer**
>
> Load the case and press Run. Watch the Tray 11 temperature and Bottoms C1/C2 trends on the Column view after the feed turns richer at 20 min. The tray holds its setpoint, yet the chromatograph reads above 2 % from 46 min. Near 150 min set Analyzer trim to On on the TIC + AIC card: the reading is 1.8 % or less from 196 min.

Case 4 changes the reflux branch after 10 min. At 150 min, reducing split to 15% gives 77.3% recovery, while increasing it to 32% gives 91.5%. A 50% subcooler bypass gives 65.7% recovery and 4.13 MW compression, with the top at $-88.5\,^{\circ}$C. The expander is still running in all three trials. The comparison makes the additional top reflux mechanism visible rather than assigning every recovery change to expander efficiency.

The expander trips at 15 min in case 5. Its gas moves to the isenthalpic bypass and its booster loses shaft drive. The residue compressor reaches its power limit at 15.87 min, column pressure is 29.53 bar

at 30 min and 30.15 bar at 60 min. The chromatograph exceeds 2% at 36 min, and measured recovery falls below 40% at 47.7 min. A restart at 60 min opens guide vanes at 20% per minute and restores measured recovery above 80% at 167 min. During the inventory recovery, instantaneous propane recovery exceeds 100%, reaching about 137% at 90 min. The numerator includes previously stored liquid, so that value does not violate component conservation.

Pause. Try Exercise 9.2 before reading on.

Case 6 starts in rejection. A 15% feed increase at 20 min causes flooding at 29.3 min, with maximum vapor load 101.3% of the declared capacity. Raising pressure to 29 bar at 60 min reduces load to 99.2% by 240 min, at a propane-recovery cost from about 92.2% to 87.0%. The case uses greater feed rate because making the feed richer alone reduced the top-stage load in the tested model. This is a physical distinction between vapor throughput and a verbal expectation that every heavier feed must flood.

9.5.3 Order, Prices and the Twelve-Hour Shift

Case 7 requests rejection at 20 min with twenty-minute ramps. The pressure-first path reaches true C1/C2 of 3.40% at 43.1 min, with the worst reported chromatograph value 3.17%. The tray-first path stays near 1.51%. Both finish moving their setpoints at 60.2 min. Slowing the reversed ramps from five through 120 minutes does not meet the sampled limit: the tested worst readings span 2.10–4.66%. The infeasibility claim is limited to that ramp sweep with the same reversed order and other settings.

Pause. Try Exercise 9.1 before reading on.

Case 8 changes ethane to ten cents per gallon at 60 min. The supervisor requests rejection at that evaluation and finishes its setpoint sequence at 100.4 min. The measured ethane recovery is still 39.6% at that instant, demonstrating why the mode label does not mean a settled inventory. Actual margin bottoms near 97 dollars per hour at 108.5 min during the move, then reaches 2984 dollars per hour at 150 min, compared with a contemporaneous recovery-mode table estimate of 2322. A profitable destination can have an expensive path.

> **TRY IT IN THE APP 8 · The frac spread flips**
>
> Press Run and open the Economics view. At 60 min the ethane price falls to 10 cents per gallon and the supervisor asks for rejection. Watch the Margin trend: the setpoints reach rejection at 100.4 min while the measured recovery still reads 39.6 %, the margin falls to near 97 dollars per hour at 108.5 min, and it reaches 2984 dollars per hour by 150 min.

The challenge schedules rejection at 180 min, a C2+ feed increase from 10 to 12 mol% at 360 min, and recovery at 540 min. Its whole-run target is C1/C2 at or below 2% at each chromatograph sample, no flooding, no time at the compressor limit, and integrated margin at least 95% of the better mode's lined-out margin at the current feed and prices. The quality score uses the true composition when each sample was drawn, not its noisy report. The controller sees only the delayed report. Changing model parameters, disturbance sliders or event buttons makes the challenge ungraded.

Table 9.5 separates the two requirements of the successful return. Declared order with the original 12 K margin fails at a sample near 600 min, and the 20 K margin with reversed order fails on the outward switch. The suggested combination gives a 1.511% worst sample at the default seed, peak vapor load 94.43%, no driver-limit time and 98.15% of the better-mode integral. A supervisor using the same sequence also passed a four-seed comparison when the app was built. The optional local MPC follows its pressure and tray targets but reaches 3.69% at a sample and 101.7% vapor load, quantities that its local prediction model does not constrain. Figure 9.2 sets the six configurations against the three limits of the target.

The default seed and seeds 0–29 were repeated at two-second and one-second physics steps with the control and instrument clocks held fixed. All 62 shipped trials fail and all 62 suggested trials pass. Suggested worst samples range from 1.5082 to 1.5369%, with peak vapor load no greater than

Table 9.5 Default-Seed Twelve-Hour Scorecard. Each row differs only in the settings named.

Configuration	Worst sample C1/C2 (%)	Peak vapor load (%)	Driver-limit time (min)	Margin share (%)	Verdict
Shipped: reversed, 12 K margin	3.65	95.2	0	97.7	Fail
Declared order only	3.20	94.4	0	98.2	Fail
20 K margin only	3.14	95.2	0	97.6	Fail
Suggested: declared, 20 K margin	1.51	94.4	0	98.2	Pass
Both setpoints stepped	11.36	93.6	0	98.8	Fail
Suggested sequence with MPC	3.69	101.7	0	98.7	Fail

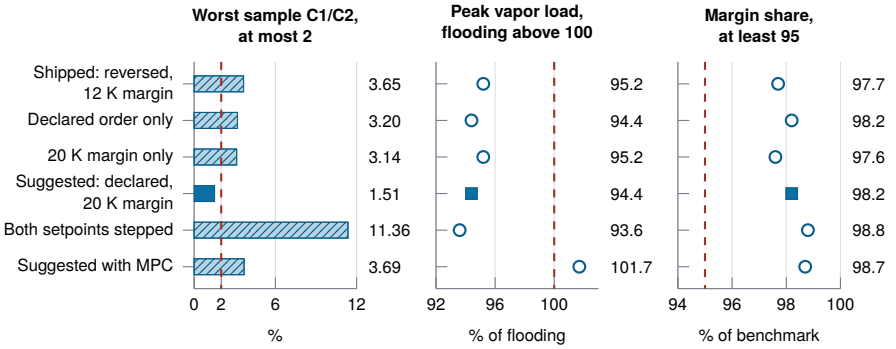

Figure 9.2 The twelve-hour scorecard of Table 9.5 against the target. The solid bar and the filled squares are the one configuration that passes, hatched bars and open circles the five that fail, and dashed lines the limits. Notice that the configuration with the largest margin share has by far the worst sample, and that only the predictive controller crosses into flooding. No row spends time at the driver limit.

94.437% and margin from 98.147 to 98.155% of the benchmark. This tests the declared noise and numerical step, rather than the reliability of a real plant.

> ✦ DID YOU KNOW? **The Richest Run Failed Worst**
>
> From the simulator (Table 9.5). Stepping both setpoints at once earned the largest margin share of the six configurations, 98.8 % of the better-mode benchmark, and also the worst sample by far, 11.36 % C1/C2 against the 2 % limit. The one configuration that passed earned 98.2 %.

9.6 Technology in Practice

Figure 9.5 draws the loops of Section 9.4 with the hardware of a gas plant in place of the schematic's symbols.

What Measures It

The NGL specification is read by one instrument, the slowest of the plant: the product chromatograph AT on the column's bottoms. The two regulatory loops read fast instruments instead: a pressure transmitter at the column top and a temperature element on stage 11, which infers the bottoms composition only as long as pressure and feed hold still. The supervisor reads a third pair, the inlet gas meter and an inlet chromatograph that gives the feed's ethane.

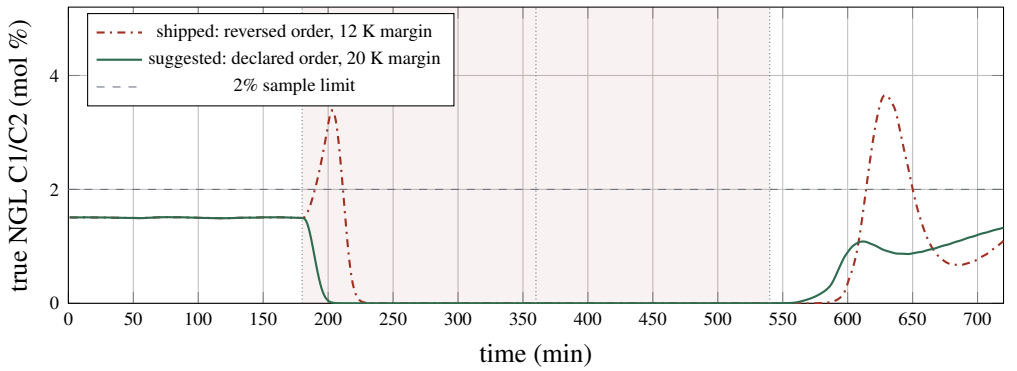

Figure 9.3 True NGL C1/C2 through the challenge, regenerated from the engine at one-minute spacing. The shaded interval has low ethane price. Vertical dotted lines mark the rejection request, richer feed and recovery request at 180, 360 and 540 min. The dash-dotted path combines reversed order with the original 12 K return margin. The solid path uses the declared order and 20 K margin. The dashed line is the 2% limit. The scorecard checks sample instants on these trajectories rather than every plotted instant.

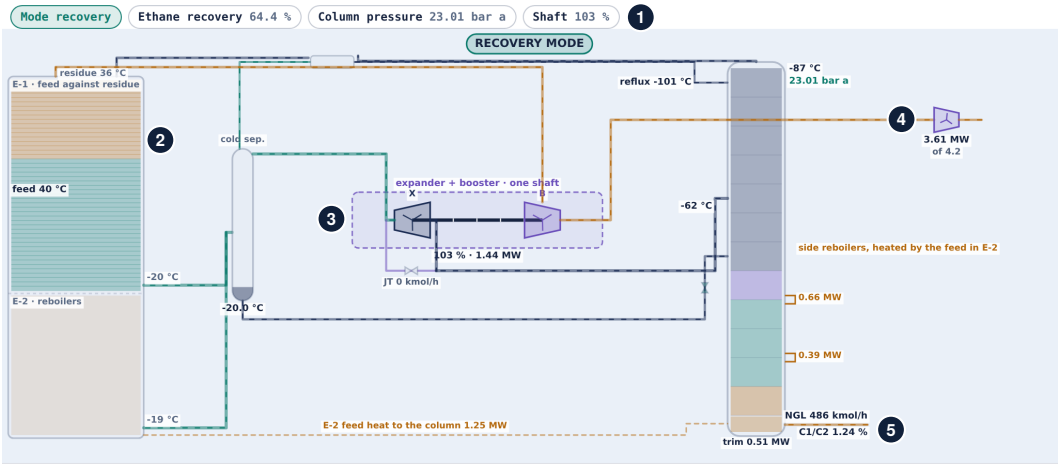

Figure 9.4 Suggested challenge at 720 min in the Plant view. Callout 1 marks the status chips: mode, measured ethane recovery, column pressure and shaft speed. Callout 2 is the cold box, E-1 above E-2, and callout 3 the expander and booster on one shaft with the JT valve below. Callout 4 is the residue compressor and callout 5 the NGL leaving the sump with its last chromatograph report. The final instruments read 23.01 bar absolute, 64.4% ethane recovery and 1.24% C1/C2, with 3.61 MW residue compression. These are the current values, not the whole-run worst sample or margin integral. The mode is again Recovery while analyzer trim is still reducing its warm return target.

TECHNOLOGY **Sensor: The NGL Chromatograph**

A heater at the inlet of the product chromatograph vaporizes its liquid sample, and the valve and column are kept hot so that it does not condense again, since liquid in a gas analyzer can cause anything from a delay to a gross error. Peak areas are converted with response factors measured on a calibration gas, often injected on a schedule. Separating species in time makes cycle times of minutes and a dead time inherent [18]. The chapter's samples every 10 min and reports 6 min later.

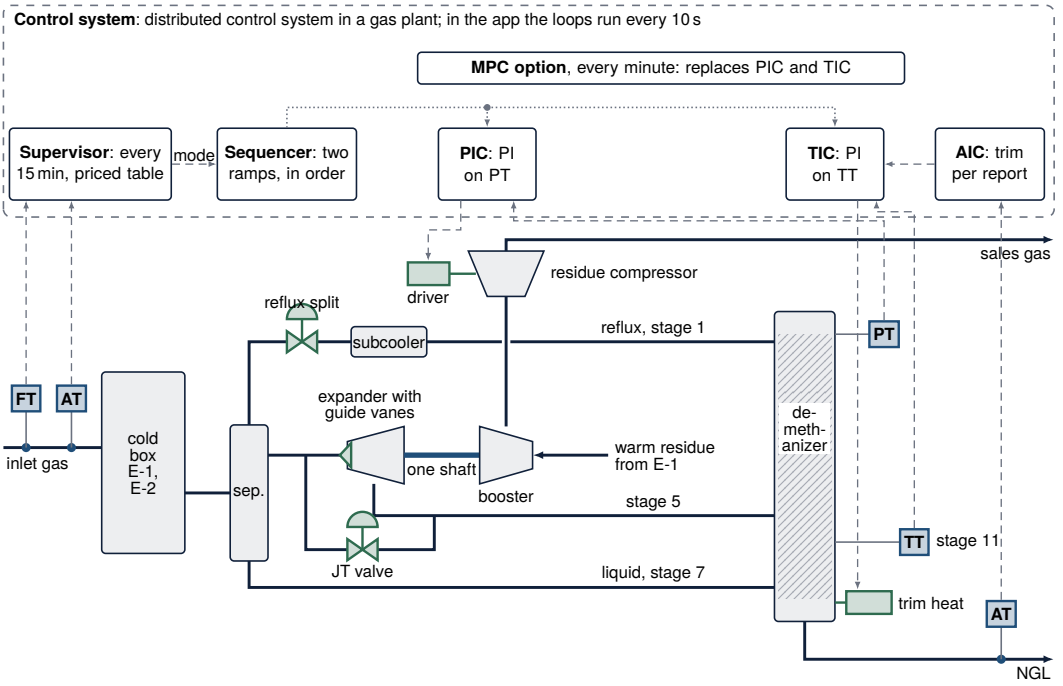

Figure 9.5 The loops of the turboexpander plant with their hardware. Lettered boxes are sensing devices with the chapter's tags: PT at the column top, TT on stage 11, AT on the NGL and FT with AT on the inlet gas. Green symbols are final elements: the expander's guide vanes with the Joule–Thomson valve beside them, the reflux split valve, the residue compressor and the trim reboiler's heat. Dashed arrows are signals, dotted arrows the setpoints the sequencer writes. Notice that the analyzer reaches the reboiler only through the tray setpoint, and that the supervisor reads the inlet meters, not the product. The subcooler bypass, the side reboilers and E-2 are left out. Not to scale, and not a piping diagram.

What Moves It

Five final elements carry the chapter's handles. The residue compressor's capacity holds the column pressure, between 0 and 120 % of its rating and below its 4.2 MW driver. The trim reboiler's heat, from 0 to 3 MW, holds the tray. The expander's guide vanes, the reflux split and the subcooler bypass are operator stations in either header mode, and of these only the guide vanes move refrigeration and recompression together.

> **TECHNOLOGY** **Final Element: The Turboexpander and Its Guide Vanes**
>
> A turboexpander is a turbine wheel driven by the expanding gas, here with the booster's wheel on the same shaft. In gas plants it often gives the coldest refrigeration on the site, and much liquid can condense inside it. The seventh edition of API 617, of 2002, covered cryogenic turboexpanders in its Chapter 4, and magnetic bearings are making inroads in gas processing [19]. The inlet guide vanes set how much gas the expander takes, and the Joule–Thomson valve beside it takes the rest. On a trip the vanes shut and the booster loses its drive, so the column pressure climbs and, in case 5, the residue compressor reaches its limit at 15.87 min.

Where the Controller Runs

Inside a gas plant these loops would run on a distributed control system, which controls production within one location under a supervisory level and nests loops so that one loop's setpoint comes from

another's measurement. Such loops run with cycle times from milliseconds to minutes [20]. The app's controllers execute every 10 s, its MPC every minute and its supervisor every 15 min, while the chromatograph cycles every 10 min, so no scan rate repairs the analyzer's delay. The chapter's structure maps onto standard blocks: two PI blocks with back-calculation for pressure and tray temperature, a sampled trim that writes the tray setpoint, a sequence of two setpoint ramps for a mode change, and a supervisory layer that chooses the mode from a priced table.

> ⚠ **IN PRACTICE** **The Cold Box Has a Speed Limit**
>
> Brazed aluminum plate-fin exchangers serve natural gas processing down to about −130 °C [21], and their makers' association sets integrity operating windows, with API 668 as a controlled document. Near the nozzles a stream should change by less than 60 °C an hour and 5 °C a minute in a start-up, shutdown or trip, and by less than 1 °C a minute in steady operation. Stream-to-stream and stream-to-metal differences should stay under about 30 °C. Outside a window the operator stops opening or throttles the valve, and steady swings call for tuning the controllers, because the damage is fatigue that ends in leaks [22]. The 48 metal temperatures of the chapter's E-1 are where such a limit would apply.

What Is Modernizing It

Two developments touch this chapter. The first removes the analyzer from the loop: a 2022 study controlled a demethanizer without composition analyzers, with a cascade on an approximation of the boilup and a pressure compensator, and found it gave the best performance and the lowest energy use among the schemes compared [11]. A pressure compensator speaks to the failure of case 7, a tray temperature that stops meaning a composition when the pressure moves. The second is model predictive control, established in industrial practice [15]. The chapter's comparison shows what it still needs here: a model that carries the methane and the flooding.

9.7 Limitations and Interpretation

The simulated gas is dry and has four pseudo-components. Water freezing, hydrate formation, carbon-dioxide solids, contaminants and heavier-end characterization are absent. The ideal-solution volatility with a convergence correction is not an equation-of-state calculation, and methane's extrapolated vapor-pressure line cannot reproduce a mixture critical locus. The shared heat capacity and gas correction omit composition effects. Agreement with the GSP brochure's broad recovery category at one selected design therefore cannot validate the physical correlations over a plant envelope.

The finite exchanger grid has a visible numerical effect on recovery. Doubling cells from the default changes design ethane recovery from 88.8% to 90.8%. The time-step study and challenge half-step sweep test temporal integration, while that spatial dependence remains. A future model could reduce this cell sensitivity or replace the teaching closure with measured mixture-property data, but the present reported results must retain the stated 48/16 resolution.

Liquid stage holdups are fixed, separator level is assumed controlled, and the flooding-efficiency law is prescribed rather than a hydraulic design correlation. There is one pressure inventory instead of distributed vapor holdup. The machine models omit surge, choking, a detailed expander efficiency map, mechanical protection and compressor-control hardware. All limits are educational boundaries. These omissions make a fast interactive lesson possible but exclude equipment sizing and operating-procedure certification.

The switch study uses one feed family, specified prices and a fixed delay structure. The better-mode benchmark assumes instantaneous lined-out operation and therefore intentionally ignores the cost of getting there. Integrated margin can also temporarily exceed both table modes when the column releases inventory. A more complete economic optimization would price inventories and transition costs explicitly. The supervisor's interpolation and the local MPC are separate approximations, each visible in the source and each limited to the information its instruments provide.

Finally, deterministic seed repetition measures sensitivity to the declared instrument noise, not industrial reliability. The physics step, sampling cycle, noise generator and scripted disturbances define those trials. Reported pass counts have meaning only alongside the actual quality, flooding, power and margin limits. The application provides those quantities rather than replacing them with a single unexplained success label.

OPEN CHALLENGE **A Colder Top Against Frozen Carbon Dioxide**

A colder demethanizer top keeps more ethane, but carbon dioxide in the gas can freeze on the column's internals, and its concentration can rise in the colder stages above its level in the feed [8]. Solid carbon dioxide blocks and plugs equipment and can rupture it, and freezing calculations enter the design of the exchangers and turboexpanders that serve demethanizers. Empirical correlations have not proved trustworthy for industrial use, and process simulators have not reliably matched data even for methane with carbon dioxide [23]. Progress would be a freezing margin predicted well enough to serve, beside flooding and the driver limit, as a constraint of the supervisor.

9.8 Conclusion

This chapter's teaching plant connects ethane economics to a timed physical procedure. At its 23 bar absolute recovery design it captures 88.8% of ethane with 3.74 MW residue compression, while the JT comparison loses about 22 K of refrigeration relative to the expander. A guide-vane closure warms the separator through the integrated reboilers, a richer feed defeats a fixed tray inference, and a profitable mode change briefly reduces actual margin to 97 dollars per hour. These outcomes require the cold box, column inventory and measurement delay to be represented together.

The twelve-hour experiment demonstrates why its successful switch needs both the declared tray-first order and the 20 K return margin. Each change alone fails the sampled 2% C1/C2 target. Their combination keeps the default-seed worst sample at 1.511%, peak vapor load at 94.43% and integrated margin at 98.15% of the better-mode benchmark. The optional predictive controller provides a useful counterexample because its local model omits quality and flooding. The plant is a transparent, locally executed process experiment, and its measured lesson is the distinction between a profitable destination, a feasible path and the delayed evidence available to its controller.

KEY POINTS

- Ethane can be sold as a liquid or left in the gas, and a profitable switch can have an expensive path.
- A controlled tray temperature has not guaranteed a controlled product composition.
- The successful switch needs both the tray-first order and the 20 K return margin, and each alone fails.
- The predictive controller follows its targets but reaches 3.69 % at a sample and 101.7 % vapor load, which its local model does not constrain.

Where else this applies. Chapter 8 trims a setpoint from a sampled, delayed analyzer in the same way, and Volume 8, Chapter 13 holds a mode as a state with its own procedure.

References

[1] W. L. Luyben, "NGL Demethanizer Control," *Industrial & Engineering Chemistry Research*, vol. 52, no. 33, pp. 11 626–11 638, 2013.

[2] S. Yoon, J.-S. Oh, and J.-K. Kim, "Dynamic simulation and control of Natural Gas Liquids recovery process," *Journal of Cleaner Production*, vol. 257, p. 120349, 2020.

[3] R. Chebbi, N. Al-Amoodi, N. Abdel Jabbar, G. Husseini, and K. Al Mazroui, "Optimum ethane recovery in conventional turboexpander process," *Chemical Engineering Research and Design*, vol. 88, no. 5-6, pp. 779–787, 2010.

[4] L. Kherbeck and R. Chebbi, "Optimizing ethane recovery in turboexpander processes," *Journal of Industrial and Engineering Chemistry*, vol. 21, pp. 292–297, 2015.

[5] Y. Li, F. Xu, and C. Gong, "System optimization of turbo-expander process for natural gas liquid recovery," *Chemical Engineering Research and Design*, vol. 124, pp. 159–169, 2017.

[6] R. W. Bucklin, "Method and equipment for treating hydrocarbon gases for pressure reduction and condensate recovery," U.S. Patent 3,292,380, filed April 28, 1964, granted December 20, 1966, 1966. [Online]. Available: https://patents.google.com/patent/US3292380A/en

[7] Honeywell UOP, "Ortloff Natural Gas Liquids Recovery," 2019. [Online]. Available: https://uop.honeywell.com/content/dam/uop/en-us/documents/industry-solutions/gas-processing/ngl-recovery/honeywell-uop-ortloff-ngl-recovery-brochure.pdf

[8] R. E. Campbell and J. D. Wilkinson, "Hydrocarbon gas processing," U.S. Patent 4,278,457, filed July 14, 1977, granted July 14, 1981, 1981. [Online]. Available: https://patents.google.com/patent/US4278457A/en

[9] A. E. S. Konukman and V. R. Akkaya, "Minimisation of exergy destruction in an ethane recovery turboexpander process using Aspen Plus process simulator," *International Journal of Exergy*, vol. 8, no. 4, p. 425, 2011.

[10] I. A. Argueta Flores, A. P. Meneguelo, C. Marangoni, Y. N. Nariyoshi, and M. S. Bacelos, "Exergy and Exergoeconomic Analysis of the Gas Subcooled Process for the Brazilian Market: An Approach to Enhance Natural Gas Liquids Recovery," *ACS Omega*, vol. 10, no. 26, pp. 28 092–28 111, 2025.

[11] M. Mandis, J. A. Chebeir, J. A. Romagnoli, R. Baratti, and S. Tronci, "Effect of the demethanizer improved control strategy on the separation train for the NGL separation process," *IFAC-PapersOnLine*, vol. 55, no. 7, pp. 889–894, 2022.

[12] National Institute of Standards and Technology, "NIST Chemistry WebBook, Standard Reference Database 69: Methane, Ethane, Propane and Butane," accessed October 14, 2026. [Online]. Available: https://webbook.nist.gov/chemistry/

[13] U.S. Energy Information Administration, "U.s. natural gas plant field production, annual, thousand barrels per day," Petroleum and Other Liquids data series, release of October 30, 2026, 2026, accessed November 18, 2026. [Online]. Available: https://www.eia.gov/dnav/pet/pet_pnp_gp_dc_nus_mbblpd_a.htm

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

[15] 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.

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

[17] 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.

[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] J. Jummonville, "Tutorial on cryogenic turboexpanders," in *Proceedings of the Thirty-Third Turbomachinery Symposium*. Texas A&M University, Turbomachinery Laboratories, 2004.

[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] Brazed Aluminium Plate-Fin Heat Exchanger Manufacturers' Association, "Applications," accessed November 18, 2026. [Online]. Available: https://alpema.org/applications.html

[22] ——, "Integrity operating windows for brazed aluminium plate-fin heat exchangers," Jun. 2025. [Online]. Available: https://alpema.org/pdf/ALPEMA_Integrity_Operating_Window_Table.pdf

[23] G. De Guido, S. Langè, S. Moioli, and L. A. Pellegrini, "Thermodynamic method for the prediction of solid CO_2 formation from multicomponent mixtures," *Process Safety and Environmental Protection*, vol. 92, no. 1, pp. 70–79, 2014.

9.9 Exercises

9.1 ○ Concept check · 2 min

In case 7 with the Reversed order, the sequencer raises the column pressure from 23 to 27 bar absolute first, while the tray loop still holds stage 11 at its recovery setpoint. The true C1/C2 of the bottoms reaches 3.40 %. Which statement explains the rise?

(a) The higher pressure floods the top stage, and the flooded stages send methane down the column with the liquid.

(b) The residue compressor reaches its driver limit, so the pressure overshoots its new setpoint.

(c) Every bubble point rises with the pressure, so liquid at the same tray temperature holds more methane until the tray setpoint is raised.

(d) The analyzer trim cools the recovery tray setpoint during the switch, which lets methane into the bottoms.

9.2 Concept check · 2 min

In case 5 the expander restarts at 60 min, and at 90 min the simulator's instantaneous propane recovery reads about 137 %. How can more propane leave in the NGL than enters with the feed?

(a) Chromatograph noise on the NGL composition inflates the recovery while the reading is near its range.

(b) The JT valve condenses more propane than the feed carries while the expander spins up.

(c) The column's component ledger stops closing during the restart, so propane appears from the integration.

(d) The sump level regulator draws down liquid that the column stored earlier, so the NGL carries propane that entered before.

9.3 Concept check · 1 min

In case 3 the feed turns richer at 20 min. The tray loop holds stage 11 on its −1.3 °C setpoint, yet the chromatograph reads above 2 % C1/C2 from 46 min. Why?

(a) A richer liquid boils warmer at the same methane content, so the held temperature now allows more methane.

(b) The trim reboiler reaches its 3 MW capacity and can no longer hold the tray temperature.

(c) The richer feed floods the column, and the flooded stages pass methane to the bottoms.

(d) The chromatograph's 6 min delay makes it report methane that the bottoms no longer contain.

9.4 Theory and calculation · 8 min

The cold separator vapor leaves at −27.1 °C and 80 bar and expands to 23 bar. The chapter's model gives −70.9 °C after the expander ($\eta_X = 0.87$) and −49.2 °C after the JT valve. Take $R = 8.314\,\mathrm{kJ/(kmol\,K)}$.

(a) Treat the vapor as an ideal gas with $c_{p0} = 36\,\mathrm{kJ/(kmol\,K)}$. Find γ and the expander outlet from Equation (1.13), isentropic and at $\eta = 0.87$. *Checkpoint:* $\gamma = 1.300$, −88.65 °C and −80.65 °C.

(b) An ideal gas does not cool through a valve. Use the chapter's vapor enthalpy, Equation (9.2) with $A = 420$ in MJ, kmol, K and bar units and $P_0 = 1$ bar, to find the vapor-only JT outlet. Start from $T = 217\,\mathrm{K}$. *Checkpoint:* about 216.8 K, −56.4 °C.

(c) Both estimates are colder than the model. The expander outlet in the model is 6.6 % liquid. Explain the difference.

9.5 Theory and calculation · 6 min

Ethane in the NGL is worth its liquid price, and ethane left in the residue gas is worth its heating value, Equation (9.6). The app uses 22.31 gal/kmol of liquid ethane and a gross heating value of 1.4786 MMBtu/kmol, with gas at 3 dollars per MMBtu. The design feed carries 273.9 kmol/h of ethane.

(a) Find ethane's value as gas in cents per gallon, and the frac spread at 25 and at 10 cents per gallon. *Checkpoint:* 19.88 cents/gal, +5.12 and −9.88 cents per gallon.

(b) Recovery keeps 88.8 % of the ethane and rejection 8.1 %. Considering the ethane term alone, what does rejection gain per hour at 10 cents per gallon?

(c) The supervisor requests the other mode only for an advantage above 50 dollars per hour. On the ethane term alone, how far below the breakeven price must ethane fall before rejection clears that hysteresis?

9.6 Theory and calculation · 6 min

The product chromatograph draws a sample every 10 min and reports it 6 min later. In recovery each fresh report moves the tray setpoint by $K_A(c_m - c_{SP})$ with $K_A = 1.5$ K per percentage point and $c_{SP} = 1.5$ %, clipped to 3 K per report.

(a) A methane excursion starts at the bottoms. What are the shortest and the longest times before a report can show it? *Checkpoint:* 6 min and 16 min.

(b) Compute the setpoint move for reports of 1.80, 3.17 and 3.65 %.

(c) After a return to recovery the tray setpoint starts the return margin above its recovery value. If every report then read 0.50 %, how many reports would remove a 12 K and a 20 K margin?

9.7 App challenge · 9 min

Load scenario 5, "The expander trips", default seed. Run at speed 5 (five simulated minutes per second). The expander trips at 15 min.

(a) Without pressing Restart, run to 60 min and pause (anywhere from 60 to 65 min is fine). Read from the exported CSV the first row with `comp_at_limit` equal to 1, `p_meas_bara` at 30 and 60 min, the first row with `c1c2_meas_pct` above 2 and the first row with `recovery_meas_pct` below 40. *Checkpoint:* 16.0 min, 29.52 and 30.17 bar, 36.2 min, 47.8 min.

(b) The Compressor power tile reads 4.20 MW. Explain why the pressure loop cannot hold 23 bar after the trip.

(c) Press Restart the expander (Manual actuator card), run to 240 min and read when the Ethane recovery tile is back above 80 %.

9.8 App challenge · 8 min

Load scenario 7, "The switch to rejection: in what order", default seed. It ships with the Sequencer order Reversed. Run at speed 5 and, between 15 and 25 min, set Mode to Rejection on the Mode card. Run to 120 min and export the CSV.

(a) Read the largest `truth_c1c2_pct` after the request and its time, and the largest chromatograph report, `c1c2_meas_pct`.

(b) Reload the case (↻), set Sequencer order to Declared before Run, and repeat the request at the same time. Read the same two maxima.

(c) Readers who requested at different times report true peaks within 0.02 % of each other but chromatograph peaks that differ by up to 0.5 %. Why?

9.9 App challenge · 9 min

Load the challenge, "a day with two switches", default seed. Run four times to 720 min, where it stops, at speed 100: as shipped, with only the Sequencer order set to Declared, with only the Return margin set to 20 K, and with both. Reload the case (↻) before each run and make the named settings on the Mode card before Run. Read the closing note, or the So far tile's tooltip, after each run.

(a) Tabulate the worst sample C1/C2, the most flooding, the minutes at the compressor limit and the margin share for each run. *Checkpoint:* 3.65, 3.20, 3.14 and 1.51 % worst samples.

(b) From the Bottoms C1/C2 trend, say in which switch each failing run breaks the limit.

(c) Explain why the return to recovery needs both changes.

9.10 Application area · 4 min

In case 6 a 15 % larger feed floods the top stage of the rejecting column at 27 bar, with a peak vapor load of 101.3 %. Raising the rejection pressure to 29 bar at 60 min brings the load to 99.2 % by 240 min, while propane recovery falls from about 92.2 to 87.0 %.

(a) The load is $V\sqrt{T/P}/K_f$. What would the pressure change alone do to the load? Why is the measured relief smaller?

(b) The feed carries 2.5 % propane. How much propane per hour moves to the residue gas, and what other handle could clear the flooding?

9.11 Application area · 3 min

A manager reads Table 9.4 and asks to run the column at 21.5 bar permanently, since the margin there is the highest. Use the table to advise.

9.12 Review the agent · 5 min

An AI assistant compares the rows of Table 9.5 and writes: “With the suggested sequence, the MPC earns 98.7 % of the better mode’s margin against 98.2 % for the PI loops, and it follows its pressure and tray-temperature targets. It is the better controller for the twelve-hour day, so the two PI loops should be replaced by it.” Find the flaw, and say how you would check it in the app.

CHAPTER 10

An NGL Fractionation Train Simulator for Plantwide Pairing and Shared Constraints

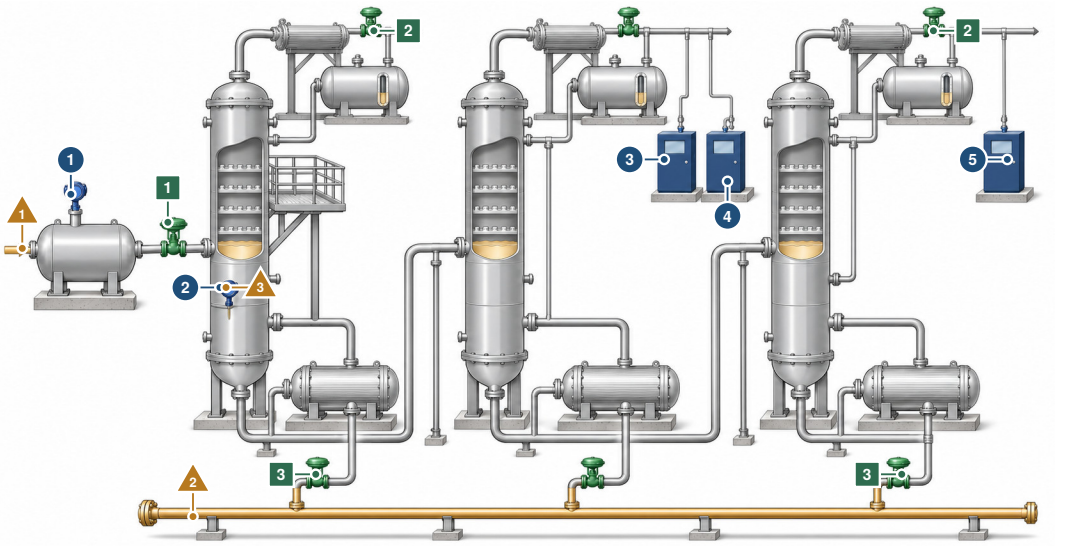

The controller moves 1 feed valve, which holds the surge drum level
2 three overhead product valves, which set the cuts
3 three reboiler steam valves, which hold the boilups the operator sets

The controller reads 1 surge drum level 2 first column's temperature, an estimate of the ethane in its bottoms
3 ethane in the propane 4 butane in the propane 5 pentane in the butane

What disturbs it 1 a richer, heavier or larger raw feed 2 a steam header that runs short 3 an estimate that drifts

The analyzer that judges the ethane in the propane sits on the second column's product line, while the valve that sets it is on the first column, and all three reboilers draw on one steam pipe. Illustration, not to scale.

IN PLAIN TERMS Which Column Owns Each Product

The main idea. A gas plant sends a mixed liquid through three tall columns in a row. Each column boils the lightest of what is left off its top. Every product is sold under a limit on the unwanted liquid it may carry.

The challenge. The ethane in the propane is checked at the second column, yet only the first column can take it out. A quick estimate of that ethane can drift without any alarm. The three columns also share one steam supply that can run short.

Why it matters. When nobody decides which product gives way, the one with no controller watching it is the one that fails.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History boxes in Section 10.1 and Section 10.2, the Did you know? boxes in Sections 10.3 and 10.5, the Try it boxes in Section 10.5, and Key points at the end of the Conclusion. No equations.
Study	50 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 10.5 with Figure 10.4, Section 10.6 and Section 10.7 with its Open challenge box.

The control question Ethane in the propane product is measured on the depropanizer, set by the deethanizer cut one column upstream, and three reboilers share one steam header. Which valve owns each specification, which measurement can see a drifting soft sensor, and which product loses when the header runs short?

Learning objectives

- Show from the shortcut split why the depropanizer cut cannot control ethane in the propane, and compute the distillate that dilution would need.
- Distinguish the cut, which trades one key impurity against the other, from the boilup, which improves both at the price of steam.
- Compare the local, product analyzer and cascade structures on a drifted inferential, and explain why the failure appears as giveaway.
- Predict which specification a proportional cut of the steam header leaves broken, and say what layer would decide instead.
- Judge a predictive controller against a cascade by its depth over a limit as well as its minutes, and trace its deeper surges to its disturbance model.

Prerequisites Sections 1.4.4 (averaging level), 1.4.8 (constraint control), 1.4.9 (late and inferred measurements), 1.6.2 (the framework's PID) and 1.6.4 (the apps). Section 1.6.3 is needed only for the MPC. Chapter 9 makes the liquid this train splits.

The app https://apmonitor.com/apps/ngl-train/

See also Volume 6, Chapter 2 for a pairing decided before any loop is tuned, Volume 6, Chapter 7 for the recycle this train does not have, and Volume 6, Chapter 8 for plantwide structure on a plant with sampled analyzers. Volume 5, Chapter 3 for composition control through a slow analyzer, and Volume 5, Chapter 10 for a soft sensor whose bias a late measurement corrects.

10.1 Introduction

A control course teaches loop pairing on a single unit. Two inputs, two outputs, a relative gain array, and a choice. On a plant the question arrives in a different form and the form is harder: a product specification is written on a stream that leaves one column, and the valve that sets it belongs to a column upstream. On the train modeled here, ethane in the propane product is measured on the depropanizer overhead, and moving the depropanizer cut from 290 kmol/h to 345 kmol/h leaves the ethane flow in that distillate at 48.5 kmol/h at both ends. It does not move at all. The percentage falls from 16.7 to 14.0 mol% purely by dilution, and reaching the 2 mol% contract limit that way would need a distillate of about 2430 kmol/h against a column feed of 918.

The misconception this corrects is that a controlled variable can be paired with any manipulated variable that moves it. Ethane is lighter than the depropanizer's light key, so its recovery to that distillate is one for every cut and every boilup, and the only thing a local loop can do is change the denominator. A relative gain array computed on the depropanizer alone would not say so, because the gain is not zero. It is a dilution gain, and dilution is not control. Volume 6, Chapter 2 shows a second way the array misleads, on a pairing it rates as free of interaction. Tracing a specification back to the unit that makes it, before drawing a single loop, is the first step of every plantwide control design procedure [1, 2], and it is the step a unit-level course has no example for.

The structure generalizes well beyond gas processing. Units in series with specifications made upstream of where they are measured, plus a shared utility header that couples them all, describes refinery crude and vacuum trains, aromatics complexes and most of what an operator means by the plant rather than the column. Plantwide design procedures exist for exactly this class [1, 3, 4] and the constraint arbitration that a real train delegates to a multivariable controller [5] is deliberately left out of the loop structures so that its absence can be observed. An optional model predictive controller plans the three cuts against their contract limits but leaves the steam header to its proportional rule, so the header arbitration stays absent under it as well. This chapter studies an interactive three column

NGL train whose separation factors, fractionation coefficients and analyzer dynamics are hidden, so that the pairing must be discovered from experiment and a shared constraint resolves itself by arithmetic rather than by judgment. The shortcut split, ratio control and plantwide design procedures are established methods. The train is a shortcut split model with its instrument chain, three quality architectures and a steam header allocation rule, written out in the detail needed to rebuild it. Its scenarios build to a boiler trip whose casualty is the specification nobody chose to sacrifice, and every run exports its complete time history. An optional model predictive controller of a standard linear form is measured against the cascade on the same instruments, and the comparison is reported as a negative result with its mechanism.

Section 10.2 follows NGL fractionation and its control from the shortcut design methods to the plantwide procedures. Section 10.3 states the train model, its shortcut split and the shared header, and Section 10.4 the loops, the three ethane architectures and the predictive controller. Section 10.5 measures the scenarios, from the transport of one disturbance to the MPC against the cascade. Section 10.6 sets the hardware of a real train beside the model, Section 10.7 lists what the model leaves out, Section 10.8 draws three general lessons, Section 10.9 sums up, and the exercises in Section 10.10 put the scenarios to work.

> ⧖ HISTORY **1964 to 1973: Control Structure Becomes the Question**
>
> A textbook of 1964, *Techniques of Process Control*, gave the first comprehensive discussion of plantwide control, in a chapter on overall process control that treated material balance control, production rate control and buffer tanks used as low-pass filters. In 1973 a critique of process control theory said that the central issue for new theories was the control structure: which variables to measure, which inputs to manipulate and which links to make between the two. A review of 2000 found industrial practice still close to the 1964 chapter, and noted that even a plant with no recycle needs the plantwide view, because in a string of units in series each one disturbs the next [6]. The train of this chapter is such a string.

10.2 Background

Raw natural gas liquids arrive from a gas plant, such as the turboexpander plant of Chapter 9, as a mixture of ethane, propane, butanes and a C5-plus cut called natural gasoline. Three columns in series take it apart, each one's bottoms being the next one's feed, and each product is sold against a contract written in mole percent. Nothing recycles, which makes the train tractable: one pass down it is the steady state.

The industry that produces these streams grew with the gas transmission network and with petrochemical demand for light feedstock. Ethane recovered from natural gas became the dominant steam cracker feed in North America, and the economics of recovering it rather than leaving it in the sales gas have driven gas plant design for decades [7]. Growth in gas processing elsewhere has followed the same pattern more recently [8]. The fractionation train itself is old technology, and its design methods are older still: the Fenske minimum stage count [9], the Underwood minimum reflux [10] and the Gilliland correlation joining them [11] are the shortcut sequence a designer still uses to size these columns before any rigorous simulation is run, and shortcut methods remain a subject of active development [12].

Control developed on a different track. The individual columns were run first on tray temperatures and manual cut adjustments, with laboratory samples confirming the contract. Process chromatographs made continuous product measurement possible and made the modern quality loop practical, but a chromatograph on the propane product measures a specification whose valve is on another column, and the retrofit that follows is a plantwide cascade rather than a local loop. The plantwide control literature then formalized that step, as the History box at the end of this section tells. Heuristic rules for the same decisions were collected in parallel [4]. Soft sensors entered the same trains as local inferentials, usually pressure-compensated tray temperatures, and the review literature on their maintenance and drift is directly relevant to what scenario 6 of this application shows [13], and

Table 10.1 Classification of Variables at the Design Operating Point

Symbol	Role	Nominal	Range
x_{C2} in propane	CV, analyzer	1.80 mol%	0.2–3.0 SP
x_{C4} in propane	CV, analyzer	2.30 mol%	0.5–4.0 SP
x_{C5} in butane	CV, analyzer	1.90 mol%	0.4–4.0 SP
$\hat{x}_{C2,B}$	CV, inferential	0.50 mol%	0.05–4.0 SP
h (V-101)	CV, level	50 %	20–80 %
D_1/F_1	MV	8.16 %	2–25 %
D_2/F_2	MV	35.14 %	20–55 %
D_3/F_3	MV	49.80 %	30–72 %
$(V/F)_{1,2}$	MV, economic	1.80, 1.90	0.8–3.2
$(V/F)_3$	MV, economic	1.40	0.6–2.6
F_0	DV, measured	1000 kmol/h	600–1400 kmol/h
z_{C2}	DV, unmeasured	0.08	0.02–0.20
V_{max}	DV, measured	5200 kmol/h	3000–7000 kmol/h
inferential bias	DV, unmeasured	1×	0.5–2×

Section 1.4.9 places them among the other inferred and late measurements of this volume. Constraint arbitration under a shared utility header arrived last, and arrived as multivariable control rather than as a regulatory structure [5], which is why a train without it behaves the way Section 10.5 reports.

> HISTORY **1997 and 2000: Procedures for the Whole Plant**
>
> In 1997 a heuristic design procedure set out to give a control structure for an entire flowsheet rather than for single units. Its nine steps centered on energy management, production rate, product quality, constraints, inventories, component balances and optimization [1]. The first steps establish the control objectives and count the control degrees of freedom by the independent valves, and the quality step recommends the usual rule of pairing close [6], a rule the ethane in this chapter's propane puts to the test. In 2000 self-optimizing control asked which controlled variables, held at constant setpoints, keep operation near the optimum with an acceptable loss [2].

10.3 Process Model

The train takes 1000 kmol/h of raw NGL at 8 mol percent ethane, 32 percent propane, 30 percent butanes and the balance natural gasoline. A surge drum feeds deethanizer T-101, whose bottoms feed depropanizer T-102, whose bottoms feed debutanizer T-103. Four products leave under five contract limits, two of them on the propane. Figure 10.1 shows the arrangement and every symbol used in the equations.

> DID YOU KNOW? **The Products of One Train, by the Million Barrels**
>
> In 2025 the natural gas processing plants of the United States produced an average of 2.3 million barrels a day of propane, 0.7 million of normal butane, 0.5 million of isobutane and 0.9 million of natural gasoline [14]. A train like this chapter's makes all of them, with the two butanes lumped into one product.

Table 10.1 classifies the variables. Flows are in kmol/h, compositions in mole percent and time in minutes. Nineteen states are integrated: four component holdups in the surge drum, four distillate component flows per column, each a first-order lag on its shortcut split, and three reboiler thermal lags. The bottoms of each column follow by difference from its feed. The analyzer samplers, their averaging filters and the local inferential lag are discrete states updated at each controller scan. A case starts at the steady state of the settings in force, with the drum at its level setpoint, the boilups at what the header passes and every instrument settled on its start-up reading.

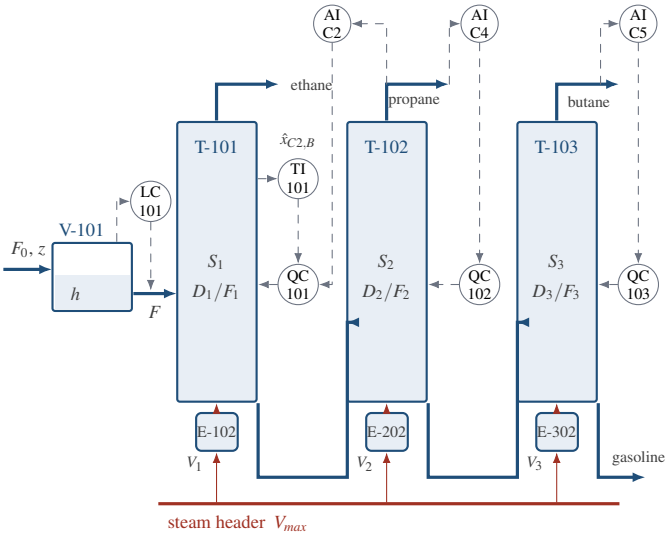

Figure 10.1 NGL fractionation train. Raw liquids F_0 of composition z enter surge drum V-101 at level h and leave as train feed F. Deethanizer T-101, depropanizer T-102 and debutanizer T-103 each carry a separation factor S_i set by the boilup ratio and a cut ratio D_i/F_i that places the material balance. Reboilers E-102, E-202 and E-302 draw boilup V_i from one steam header of capacity V_{max}. LC-101 holds the drum level with the train feed. AI-C2 and AI-C4 are the propane product chromatographs and AI-C5 the butane product chromatograph, TI-101 is the deethanizer local inferential $\hat{x}_{C2,B}$, and QC-101 to QC-103 manipulate the three cut ratios, QC-101 from TI-101, AI-C2 or both in cascade. The ethane specification is measured on the T-102 overhead and set by the T-101 cut.

10.3.1 The Shortcut Split

Each column is represented by two numbers, a separation factor S that says how sharply it fractionates and a cut rate D that says where the material balance is placed. Everything lighter than the light key goes entirely overhead, everything heavier than the heavy key goes entirely to the bottoms, and the two keys share according to Equation (10.1).

$$\frac{r_L}{1-r_L} = S\,\frac{r_H}{1-r_H} \tag{10.1}$$

where r_L and r_H are the recoveries of the light and heavy keys to the distillate. The equation is solved by bisection on r_H subject to the overall material balance $D = \sum_i r_i F z_i$.

The first sentence of that paragraph contains the whole plantwide lesson. Ethane is lighter than the depropanizer's light key, so $r_{C2} = 1$ for that column, always, at every cut and every boilup.

10.3.2 Fractionation and Energy

The separation factor is built from a stage count and an efficiency that saturates in boilup, Equation (10.2).

$$S = \alpha^{N\eta}, \qquad \eta = 1 - \exp\left[-k\left(\frac{V}{F} - \left(\frac{V}{F}\right)_0\right)\right] \tag{10.2}$$

The coefficients k_i are calibrated so that each column reproduces its design split at its design boilup ratio, and $(V/F)_0 = 0.50$ is the boilup below which nothing separates. Because η saturates, the first increment of steam buys a great deal of separation and the fourth buys very little, which is the economic content of the boilup setpoint.

Table 10.2 Model Parameters

Symbol	Meaning	Value	Class
α_1	C2/C3 volatility	2.10	adjustable
α_2	C3/C4 volatility	2.40	adjustable
α_3	C4/C5 volatility	2.60	adjustable
N_i	theoretical stages, each	25	adjustable
k_1, k_2, k_3	fractionation coefficients	0.4703, 0.3384, 0.4587	hidden
$(V/F)_0$	boilup with no separation	0.50	adjustable
τ_1	deethanizer composition	25 min	adjustable
τ_2	depropanizer composition	30 min	adjustable
τ_3	debutanizer composition	45 min	adjustable
θ	analyzer cycle	10 min	adjustable
x_{C2}^{max}	ethane in propane	2.0 mol%	visible
x_{C4}^{max}	butane in propane	2.5 mol%	visible
x_{C5}^{max}	pentane in butane	2.0 mol%	visible
$x_{C4,gas}^{max}$	butane in gasoline (RVP)	2.0 mol%	visible
x_{C3}^{max}	propane in ethane	5.0 mol%	visible

10.3.3 Cut Against Boilup

The two handles are not interchangeable and the difference is measurable. On the deethanizer, six more kmol/h of cut takes the ethane left in the bottoms from 1.60 kmol/h to 0.62 kmol/h and pushes the propane lost overhead from 3.20 kmol/h to 8.22 kmol/h, which is a trade along a fixed curve. Raising the boilup ratio from 1.6 to 2.2 at a fixed cut takes the ethane from 2.97 kmol/h to 0.45 kmol/h and the propane from 4.57 kmol/h to 2.05 kmol/h, so both improve, which no cut change can do. That asymmetry is why the quality loop goes on the cut and the boilup stays an economic setpoint.

🕮 **Pause.** Try Exercise 10.2 before reading on.

10.3.4 Why the Manipulated Variable Is a Ratio

The quality loops manipulate D/F rather than D. An 18 percent rate increase moves the correct cut flow of every column by 18 percent and leaves the correct cut ratio unchanged at a fixed feed composition, because the split depends only on the cut and boilup ratios, so a ratio station removes the disturbance before any controller sees it. This is feedforward written into the pairing rather than added beside it, and it is the same argument the reflux ratio makes on a superfractionator at a different scale.

10.3.5 The Shared Header

Three reboilers draw on one steam header of capacity V_{max}. When the sum of the three demands exceeds it, every boilup is multiplied by the same factor, Equation (10.3).

$$\phi = \min\left(1,\ \frac{V_{max}}{V_1 + V_2 + V_3}\right), \qquad V_i^{app} = \phi\, V_i \tag{10.3}$$

Nothing in this rule decides which column should give way. That is deliberate, and Section 10.5 reports what the plant does about it.

10.3.6 Assumptions

Each column is a shortcut split with a single separation factor rather than a stage-by-stage model, so tray hydraulics, flooding and efficiency loss at turndown are absent. The column responds as a first-order lag on composition, which represents the dominant time constant and not the shape of the approach.

The train has no recycle. Real gas plants return some streams and a depropanizer overhead is sometimes recycled for reprocessing, which would change the plantwide analysis considerably.

Component lists are lumped into four pseudo-components, so the isobutane and normal butane split that a real debutanizer has to manage is invisible.

The steam header allocation is a proportional cut with no priority. Real plants carry a load shedding scheme or a multivariable controller that decides which column to starve, and the absence of one here is what scenario 8 exists to show.

The natural gasoline Reid vapor pressure specification is represented by a butane mole fraction limit rather than by a vapor pressure calculation.

Analyzers are exact samplers with a cycle time and an averaging filter, and the inferential carries an explicit bias fault. Where a case adds analyzer repeatability noise, the noise enters after the sampler, so a published result wanders between samples (0.030 mol% RMS about the held sample in the challenge) where a real chromatograph reports one value per cycle. A real chromatograph has more ways to be wrong than a gain error.

10.4 Controller Design

10.4.1 Velocity Form PID

Every loop uses the framework's velocity-form PID of Section 1.6.2, Equations (1.16) to (1.19), with the derivative taken on the measurement, executed every $\Delta t_c = 0.2$ min. The integral is seeded with the output in force when the case starts, and the tracking time of the back-calculation is $T_t = \tau_I$, so each scan at a limit removes $\Delta t_c / \tau_I$ of the clipped excess from the integral alone. Clamping the running output instead would rectify symmetric analyzer noise into a one-way drift away from whichever limit the output rests against. No cut ratio and no setpoint written by the trim reaches its limit in any of the ten original cases as shipped or in the cascade runs of Section 10.5.7, and under the predictive controller the loops only track the cuts it applies, so the tracking term does not act in any result reported here. In manual the integral is set so that the output equals the operator's cut ratio exactly, which makes the transfer to automatic bumpless. Derivative action is not used on any loop here ($\tau_D = 0$), because the measurements are sampled and a derivative on a staircase is a sequence of impulses [15]. Integration is fixed-step fourth-order Runge-Kutta at 0.05 min on the engine of Section 1.6.1, so every run reproduces exactly.

10.4.2 The Three Ethane Architectures

The ethane specification selector offers four closed-loop configurations besides off. Comparing the three below is the subject of Section 10.5, and the fourth, the model predictive controller of Section 10.4.3, is compared with the cascade in Section 10.5.7.

Local closes a loop from the deethanizer inferential, a lagged estimate of ethane left in the bottoms, to the deethanizer cut ratio, with action $a = -1$ in Equation (1.16) because more cut is less ethane left behind, and $K_c = 1.20$ % of feed per mol% and $\tau_I = 28$ min.

Analyzer closes a loop from the propane product chromatograph, one column downstream, directly to the deethanizer cut ratio, also with $a = -1$, $K_c = 0.60$ % of feed per mol% and $\tau_I = 120$ min. It is slow because it waits for a composition to travel a column and an analyzer to sample it. Volume 5, Chapter 3 tunes a composition loop against the same kind of slow analyzer on a single column.

Cascade puts the analyzer on the setpoint of the local loop, with $a = +1$ on the setpoint, $K_c = 0.40$ and $\tau_I = 240$ min inside 0.05 mol% to 4.0 mol%, which keeps the fast local response and removes the inferential's offset. The trim is not told when the inner loop's cut reaches its limit.

The butane in propane loop manipulates the depropanizer cut, with $a = +1$, $K_c = 0.50$ and $\tau_I = 40$ min, and the pentane in butane loop manipulates the debutanizer cut, with $a = +1$, $K_c = 0.60$ and $\tau_I = 45$ min. These are the tunings of scenarios 4 to 9. The surge drum level loop, $K_c = 25$ kmol/h per percent and $\tau_I = 90$ min, is detuned so that the drum absorbs rate changes rather than passing them to three columns at once, the averaging level control of Section 1.4.4.

10.4.3 Model Predictive Control

A fifth setting of the ethane selector, never the default and loaded only by scenario 10, hands the cut ratios to one model predictive controller (MPC), a plant-specific form of the framework's option of Section 1.6.3. Once a minute ($T_s = 1$ min, every fifth controller scan) a quadratic program plans the deethanizer cut ratio $u_1 = D_1/F_1$, and $u_2 = D_2/F_2$ and $u_3 = D_3/F_3$ whenever the butane and pentane loops are in automatic, and the cuts take the plan's first move and keep it for that minute, as in the receding-horizon practice of industry [5]. A cut whose loop is in manual stays where the operator set it, and its impurity leaves the objective. The ratio stations, the cut ranges and the header rule of (10.3) act on its moves as on the loops' outputs, and it moves no boilup. Predictive control has been evaluated on an industrial depropanizer [16], and the control of a whole NGL train has been studied [17].

The controller's model is identified on the application. Each cut ratio was stepped by ± 0.1 percent of feed on the deethanizer and ± 0.2 on the other two, in manual with the noise off, where the products sit on their targets of 1.80, 2.30 and 1.90 mol% (cut ratios of 7.489, 35.62 and 49.79 %), and each 480 min response of the logarithm of an impurity was fitted by least squares on first-order lags at the declared time constants of Table 10.2. Logarithms are used because a cut changes an impurity by a factor, as the inferential bias does. The deethanizer cut reaches the downstream products only through its bottoms as the inferential reads them, because the same cut leaves different bottoms on a different feed. The result is Equation (10.4).

$$
\begin{aligned}
\tau_m \frac{\mathrm{d}a_{jm}}{\mathrm{d}t} &= u_j - u_j^0 - a_{jm}, \qquad \tilde{y}_B = \textstyle\sum_{m=1}^{3} \gamma_m a_{1m} + d \\
\tau_2 \frac{\mathrm{d}p}{\mathrm{d}t} &= \tilde{y}_B - p, \quad \tau_3 \frac{\mathrm{d}q}{\mathrm{d}t} = \tilde{y}_B - q, \quad \tau_3 \frac{\mathrm{d}r}{\mathrm{d}t} = p - r \\
\tilde{y}_{C2} &= 1.0137\,p - 0.0281\,a_{22} \\
\tilde{y}_{C4} &= -0.4214\,p + 0.8195\,a_{22} + 0.0081\,a_{23} \\
\tilde{y}_{C5} &= -0.0827\,p - 0.0786\,q - 0.0592\,r \\
&\quad - 0.9079\,a_{22} + 1.3635\,a_{23} + 0.5566\,a_{33}
\end{aligned}
\tag{10.4}
$$

where a_{jm} is cut j through the lag τ_m, d is a disturbance at the deethanizer bottoms held constant over a plan, $\gamma = (-1.3673,\ -0.0855,\ 0.0137)$, p, q and r carry the bottoms downstream, and $\tilde{y}_B$, $\tilde{y}_{C2}$, $\tilde{y}_{C4}$ and $\tilde{y}_{C5}$ are the model's changes in the logarithms of $\hat{x}_{C2,B}$, x_{C2}, x_{C4} and x_{C5} of Table 10.1. The coefficients on a_{jm} are in natural logarithm per percent of feed and those on p, q and r are dimensionless, a term of $-6 \times 10^{-7}\,a_{23}$ in $\tilde{y}_{C2}$ is omitted, and the model is anchored at the first readings of a run, which it takes as a steady state, so u_j^0 is the cut ratio at $t = 0$. Recomputed by central differences on the declared train, the steady log gains differ from the model's by at most 0.2 %. Free fits to one-sided steps of 0.1 % of feed return 25.8 min from the deethanizer cut to the inferential, 29.3 min from the depropanizer cut to the butane and 44.2 min from the debutanizer cut to the pentane, against the declared 25, 30 and 45, and the model's own responses depart from these steps by 0.1 to 6.1 % RMS of their span. At the design cut ratios, with 0.50 mol% ethane in the propane, the deethanizer cut moves the ethane 1.41 times as hard as this linear model says.

The model is stepped at every scan with the cut ratios applied. The inferential updates d at every scan, and each analyzer result updates an output bias β_i on its own product, as shown in Equation (10.5).

$$
\begin{aligned}
d &\leftarrow d + \lambda_I \left(\ln \hat{x}_{C2,B} - \ln \hat{x}^0_{C2,B} - \tilde{y}_B \right) \\
\beta_i &\leftarrow \beta_i + \lambda_A \left(\ln x_i^{an} - \ln x_i^0 - \tilde{y}_i(t_s) - \beta_i \right)
\end{aligned}
\tag{10.5}
$$

where $\lambda_I = 0.2$ per scan, $\lambda_A = 0.5$ per result, a superscript 0 marks a first reading, and $\tilde{y}_i(t_s)$ is the model's prediction recorded when the sampler drew the sample whose result x_i^{an} is read 7.5 min later, once the 1.5 min averaging has settled. Compared at the moment of reading instead, the ethane in

scenario 10 would peak at 3.60 rather than 3.28 mol%. The analyzer biases find a drifted inferential, taking up half of the 25 percent drift of scenario 10 in 47.8 min and 90 percent in 97.9 min. The estimate of product i is $\hat{y}_i = \ln x_i^0 + \tilde{y}_i + \beta_i$, with one integrating disturbance per measured signal, the number shown sufficient for offset-free control [18], so the analyzer dead time lives in the estimator and the plan predicts the true product that the contract judges.

Each decision minimizes Equation (10.6) over $N = 120$ steps, a horizon of 120 min, with the moves held in eight blocks of 1, 2, 3, 5, 8, 13, 20 and 68 steps, a program of 24 moves and three slacks.

$$\begin{aligned} \min_{u,\varepsilon} \; & \sum_{k=1}^{N} \sum_{i} w_i \left(\hat{y}_{i,k} - \ln x_i^*\right)^2 + \sum_{b=1}^{8} \sum_{j=1}^{3} S_j \, g_j^2 \, \Delta u_{j,b}^2 \\ & + \sum_{i} \left(\rho_2 \, \varepsilon_i^2 + \rho_1 \, \varepsilon_i\right) \\ \text{s.t. } & u_j^{lo} \le u_{j,b} \le u_j^{hi}, \qquad \varepsilon_i \ge 0 \\ & \hat{y}_{i,k} \le \ln x_i^{max} + \phi_{i,k} + \varepsilon_i, \quad k = 1, 6, 11, \ldots \end{aligned} \tag{10.6}$$

where $\hat{y}_{i,k}$ is the predicted logarithm of product i at step k, x_i^* is the learner's target clipped to the contract limit x_i^{max} of Table 10.2, w_i is 1 when the plan moves the impurity's own cut and 0 otherwise (its band is then dropped), $\Delta u_{j,b}$ is the move of cut j into block b, the first measured from the cut applied now, $g = (1.459,\ 0.828,\ 0.557)$ is each cut's steady log gain on its own impurity, $S_1 = 20000$, $S_2 = S_3 = 1$, $\rho_1 = 10^4$ and $\rho_2 = 10^5$. The bounds are the ratio stations' ranges of Table 10.1 and are hard, and an optional deethanizer rate limit is off as shipped. The contract limits are soft constraints, and the shared solver carries a single slack for each output over the whole horizon. The funnel $\phi_{i,k} = \max(0,\ \hat{y}_{i,k}^h - \ln x_i^{max})(1 - s_{i,k})$, with $\hat{y}^h$ the prediction with the cuts held and $s_{i,k}$ the normalized step response of impurity i to its own cut, lets an excess already on its way decay as fast as that cut could remove it. Without the funnel the plan spends 750 min over the ethane limit in scenario 10 instead of 153.

Tracking the targets replaced a linear price on giveaway, which with one slack per impurity either paid the plan to ride just over a limit or, priced to forbid that, left no move price that both approached a target and rode out a surge. Giveaway is therefore measured, not priced. The slack is still bought against the move price, in 130 of the 1440 plans of scenario 10, all between 210 and 565 min, none above 0.031 in the logarithm and none with a cut at a bound, and the first of them re-solved with the deethanizer move priced a hundredth as much buys none. The prices were chosen once, on scenario 10, as the least summed loop IAE with no more cut travel and no more minutes over any limit than the shipped cascade, then the largest prices within 10 percent of that IAE.

The controller keeps a plan only if the solver reports success with an optimal status, every number it returns and every prediction converted back to mole percent is finite, the slacks are non-negative, and every planned move over the horizon, not only the first, lies inside its range. Otherwise the cascade and the two downstream loops, which track the applied cuts at every scan with no preload, take over until a plan is accepted again. With the noise off, forced rejections at five instants of scenario 10 hand over within 0.017 % of feed, and none of the measured runs rejected a plan. A switch to the MPC in mid-run is the plan's first decision rather than a bump, but it can move the debutanizer cut by 0.72 % of feed at once, about twice the largest move between consecutive plans of the day. The first plan is solved at $t = 0$. In Chrome a plan takes 2.1 ms, which occupies 0.21 % of a processor core when the run plays one simulated minute per second and 25 % at the top speed of 120 simulated minutes per second. The plant on the run's parameters, the instruments and the controller's identified model are three separate models, and replaying the 7201 measurement packets of scenario 10 in front of four different hidden plants reproduces all 1439 plans bit for bit.

Table 10.3 Depropanizer Cut Swept on a Rich Feed. The Ethane Flow in the Distillate Does Not Move.

Quantity	Cut 290 kmol/h	Cut 345 kmol/h
Ethane in distillate (kmol/h)	48.5	48.5
Ethane in distillate (mol%)	16.7	14.0
Butane dragged overhead (kmol/h)	0.24	0.95

10.5 Results

Eleven scenarios run from an uncontrolled disturbance to a graded challenge. Scenario 1 follows one disturbance down the train in manual, 2 shows that the unit that measures the ethane cannot fix it, and 3 sets the cut against the boilup. Scenarios 4 to 6 build the local, product analyzer and cascade architectures, 7 puts a number on giveaway, 8 trips the steam header, 9 lands a heavier feed and then a rate step, 10 sets the MPC against the cascade and 11 is the graded challenge. Each is credited by a completion rule, as Section 1.6.4 describes, and the exercises of Section 10.10 are built on them. Every run below is the case as it loads, with any setting named in the text changed before the run starts. Reset restarts a case from $t = 0$ with the learner's settings kept (Section 1.6.1), so scenarios 5, 6 and 10 reach their comparison runs by Reset, with the architecture changed before the scripted disturbance arrives.

10.5.1 Following a Disturbance Down the Train

Stepping the raw NGL from 8 to 13 mol% ethane at 30 min with nothing controlling gives the transport structure directly. On the run clock the deethanizer inferential is 30 percent of the way to its new value at 60 min and half way at 90. The propane product is 10 percent at 60 min, 28 percent at 90 and half way at 129. With no loop closed the ethane in the propane finishes at 15.0 mol% against a limit of 2.0. The first hour belongs to the surge drum rather than to the deethanizer. The bottoms follow the feed by difference, so the inferential moves as soon as the drum composition does, and 1500 kmol of drum holdup at 1000 kmol/h is a 90 min mixing lag. With the drum capacity at its 500 kmol minimum, set before Run, the inferential is half way 8.9 min after the step instead of 60. The further 39 min the product takes is mostly the depropanizer's 30 min composition lag.

> **TRY IT IN THE APP** **1 · Follow one disturbance down the train (MANUAL)**
>
> Load the case and press Run. At 30 min the raw liquid turns richer in ethane and nothing on the plant moves. Watch the ethane trend: the deethanizer inferential moves first and is half way to its new value at 90 min, the propane product follows and is half way at 129, and the ethane in the propane finishes at 15.0 mol% against a limit of 2.0.

10.5.2 Where the Specification Is Made

Table 10.3 reports the depropanizer cut ratio stepped from 31.6 to 37.6 % of its feed on an already rich feed of 13 mol% ethane, as scenario 2 asks.

The values are the steady state at each cut. A cut chosen before Run starts the case at its steady state, so the scenario's first 90 min hold reads 16.7, and the second, after the step to 37.6 %, reads 14.2 mol%. The ethane flow reads 48.5 kmol/h at every step of the run. The percentage falls only by dilution, and reaching 2 mol% that way would require a distillate of about 2430 kmol/h against a column feed of 918. The valve that sets this specification is the deethanizer cut, one column and about an hour upstream.

Pause. Try Exercise 10.1 before reading on.

Table 10.4 Ethane in Propane With a Drifted Inferential, Target 1.80 mol%

Architecture	15 h (mol%)	30 h (mol%)
Local inferential alone	0.705	0.705
Cascade	1.595	1.751
Analyzer alone	1.856	1.800

10.5.3 Architecture Audit

Scenario 6 drifts the deethanizer inferential 30 percent low at 30 min and steps the feed to 12 mol% ethane a minute later. Table 10.4 reports the ethane in the propane product against a 1.80 mol% target for each architecture, selected before the run. The scenario's own route to the local run, Reset and then Local, gives the same 0.705.

The direction matters more than the magnitude. A soft sensor reading low makes the local loop let more ethane through than its setpoint intends, but the 0.50 setpoint sits so far inside the 1.80 mol% target that the product still settles below it. The failure appears as giveaway rather than as an alarm, and a plant can run that way for months. The cascade recovers most of the error because the analyzer above it corrects the inferential setpoint, moving it from the shipped 0.50 to 1.24 mol% by 30 h. The trim is deliberately slow, and the product reaches 1.80 ± 0.01 mol% only 47 h into the run. The analyzer alone is exact because nothing biased stands between it and the product. What the local loop buys in exchange is speed, and what it costs is a measurement nobody is checking, which is the standard soft sensor maintenance problem [13].

Pause. Try Exercise 10.3 before reading on.

> **DID YOU KNOW? The Slow Loop Was Closer**
>
> From the simulator (Table 10.4). At 15 h into the drifted-inferential case, the product analyzer alone, the slowest of the three structures, held the propane at 1.856 mol% against the 1.80 mol% target, while the cascade read 1.595 and the local loop 0.705. At 30 h the analyzer alone read exactly 1.800.

10.5.4 Giveaway as a Quantity

Running at 0.60 mol% ethane and 1.00 mol% pentane against limits of 2.0, then stepping both targets up at 240 and 245 min, adds 3.39 kmol/h of propane and 6.30 kmol/h of butane four hours later, which is about half a tonne an hour. By then the slow trim has taken ethane in the propane only to 1.11 mol%, and the propane gain grows to 6.46 kmol/h as ethane approaches its target over the next three days. The steam consumption moves from 4380 kmol/h to 4387 kmol/h, which is to say it does not change. Giveaway on this train is not paid for in energy. It is paid for in product that left in the wrong stream, and it is invisible on every trend that has an alarm on it. The low pentane target also held butane in the natural gasoline at 2.74 mol% at 240 min, above its 2.0 limit, so the careful-looking operation was already off specification on the stream whose specification has no loop.

> **TRY IT IN THE APP 7 · What running below the limit costs**
>
> Press Run. The cascade holds ethane in the propane at 0.60 and pentane in the butane at 1.00 mol%, and at 240 and 245 min both targets step up. Four hours later the propane rate is up 3.39 kmol/h and the butane 6.30 kmol/h, while the Steam badge moves only from 4380 to 4387. Butane in the natural gasoline already read 2.74 mol% at 240 min, over its 2.0 limit.

10.5.5 The Shared Header

Tripping a boiler at 60 min multiplies every boilup ratio by 0.912, and the header stays short for the rest of the run. Table 10.5 reports the state 15 h into the run.

Table 10.5 Steam Header Trip From 5200 kmol/h to 4000 kmol/h. The Specification Without a Controller Is the One That Fails.

Product quality	Peak	15 h	Limit	Loop
Ethane in propane (mol%)	1.49	1.49	2.0	yes
Butane in propane (mol%)	2.56	2.29	2.5	yes
Pentane in butane (mol%)	2.36	1.90	2.0	yes
Butane in gasoline (mol%)	3.81	3.81	2.0	no
Propane recovery (%)	—	97.1	—	—

Propane recovery falls from 98.3 to 97.1 %. The three controlled impurities are inside their contracts at 15 h, but two of them leave on the way while the cut controllers move the material balances to make up for the lost separation. Butane in the propane spends 65 min above its 2.5 mol% limit and pentane in the butane 194 min above 2.0. Ethane is still climbing toward its target under the slow trim. Butane in the natural gasoline crosses 2.0 mol% 45 min after the trip, climbs to 3.81 mol% and stays there. The debutanizer defends the specification that has a controller on it, using the only cut it has, and the specification on the other side of the same cut is the casualty. Nothing in the control system chose that outcome. The header cut everybody by the same factor and each loop then did what it was told, and the result is the specification nobody would have picked to sacrifice. Resolving that properly requires a layer that knows the relative value of the four products, which on a real train is a multivariable controller [5] and on this one is deliberately absent.

10.5.6 The Graded Challenge

The challenge runs a 1400 min shift with four disturbances, with noise on the analyzers and the inferential. Raw NGL ethane steps from 8 to 12.5 mol% at 60 min, the deethanizer inferential drifts to 0.75 of its reading at 240 min, butanes step from 30 to 33.5 mol% at 460 min, and the rate rises from 1000 to 1150 kmol/h at 700 min. The case loads on the local inferential alone with both downstream loops in manual and a cautious tuning. It is met when ethane in the propane stays within 0.30 mol% of target and butane in the propane within 0.20 mol% from 924 min to the end, with all five product limits met at 1400 min. It is judged on the case as given, and changing a model parameter or a disturbance voids the attempt.

With the shipped line-up (Local, both downstream loops in Manual, cautious tuning) the worst errors over the scored window are 0.74 and 1.79 mol%, and butane in the natural gasoline ends at 22.9 mol%, so the case fails on all three counts. With the suggested line-up (Cascade, both loops in Auto, the tunings of Section 10.4, boilup ratios 1.95, 2.00 and 1.80) they are 0.22 and 0.09 mol%, inside the bands by 0.08 and 0.11, with every product inside its limit, whether the line-up is chosen before Run or carried through a Reset. Over eight other seeds of the analyzer and inferential noise the suggested ethane error ranges from 0.21 to 0.25 mol% and every run passes. The suggested boilups ask the header for about 5406 kmol/h after the rate step, so every reboiler is trimmed by 4 % and the case still passes.

10.5.7 Three Cuts Planned Together

Scenario 10 starts at the design cut ratios, 0.50 mol% ethane in the propane against a 1.80 mol% target, with analyzer and inferential noise of 0.04 and 0.03 mol%. Raw NGL ethane steps from 8 to 12 mol% at 180 min, the inferential drifts 25 percent low at 600 min and the butanes step from 30 to 33.5 mol% at 1020 min. The MPC runs as the case loads, and the cascade on the tunings of Section 10.4 with its inferential setpoint at 1.82 mol%, selected before the run as the scenario's Reset route does. Table 10.6 compares the two with three cascade variants and the MPC restricted to the deethanizer cut, and then the MPC with the cascade, or the challenge's shipped local loop, on scenarios 5, 6, 7 and 9 and the challenge, always on the settings of scenario 10 without retuning.

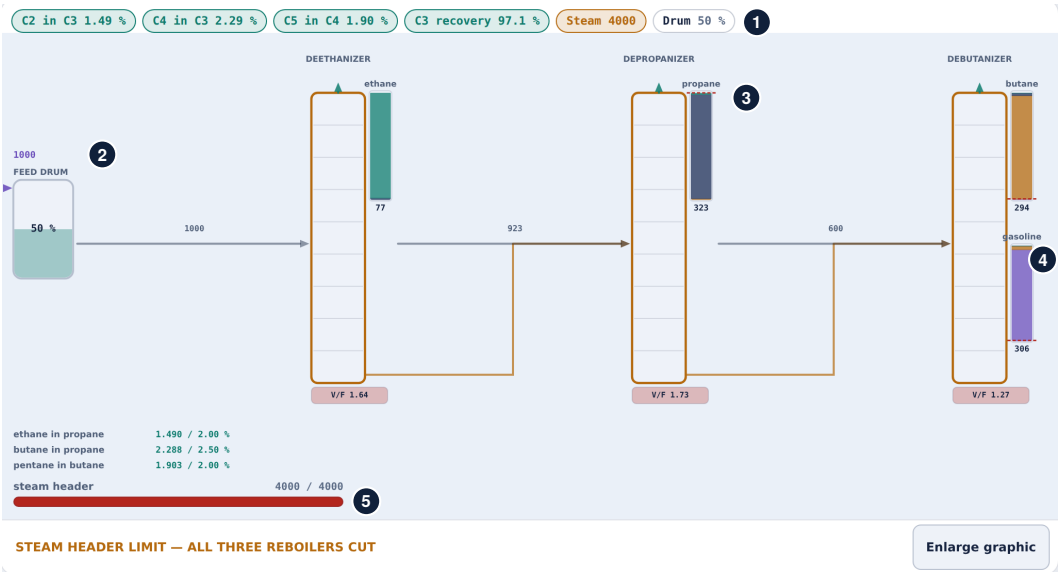

Figure 10.2 The NGL fractionation train application running scenario 8, the steam header trip, at $t = 15.00$ h with the ethane cascade and both downstream quality loops in automatic, cropped to the plant view. (1) The status chips for the three controlled impurities, the propane recovery, the steam header and the drum. (2) The feed surge drum, drawn with the three columns on one line in the order the material flows. (3) Beside each column is a stacked bar of its overhead composition, ethane at the top and then propane, butanes and gasoline, with the contract limit drawn as a red line on the band it applies to, short-dashed. (4) The debutanizer also shows its bottoms, because the gasoline vapor pressure specification lives there and has no loop of its own. (5) The bar under the scene is the steam header, here passing 4000 of 4000 kmol/h, and every column outline has turned from its dark default to amber because the reboilers ask for more than it can pass. The boilup ratios read 1.64, 1.73 and 1.27 against setpoints of 1.80, 1.90 and 1.40. Ethane, butane and pentane in their products read 1.490, 2.288 and 1.903 mol%.

Giveaway is the room left below the ethane, pentane and gasoline butane limits, in kmol as the application reports it and in tonnes at each component's molar mass.

Over the day the MPC spends 153 against 513 min over the ethane limit, but its excess is 86.7 against 65.8 mol% min and its peak 3.28 against 2.54 mol%, as shown in Figure 10.3(a). The two structures lose in different windows. In the ethane step the MPC is over for 153 against 103 min with 86.7 against 34.1 mol% min, and in the drift window the cascade's slow trim holds the propane over the limit for 394 min while the MPC, on this seed, stays under it. Both rankings hold at half the physics step and on six other noise seeds, where the MPC spends 124 to 374 min over with 84.6 to 92.5 mol% min against 498 to 544 min with 62.1 to 66.4 mol% min.

In moles the two give away the same, 63.4 against 63.5 kmol, and by mass the MPC gives away 3.59 against 4.42 t, because the room it leaves is lighter, 23.1 against 7.7 kmol below the ethane limit in the propane and 25.2 against 43.1 kmol below the pentane limit in the butane. The molar tie lies within the noise. On the six other seeds the MPC gives away 61.4 to 65.6 kmol against 63.5 to 63.8, while its tonnage stays lower, 3.52 to 3.66 against 4.42 to 4.44 t.

Three cascade variants locate the difference. With the MPC's mean purities as targets and the matching inferential setpoint, the cascade gives away more, 73.7 kmol, so the gain is not where the targets sit. Restricted to the deethanizer cut, the MPC gives away 78.5 kmol, so the pentane gain comes from moving the three cuts together. Tuned once on this case by the rule that tuned the MPC,

Table 10.6 MPC Against the Loop Structures, Each Case at Its Own Seed, Noise and Duration

Case	Structure	Over limit (min) C2 / C4 / C5	Excess C2 (mol% min)	Peak (mol%) C2 / C4 / C5	Giveaway kmol (t)	Plans, outcome
10, 1440 min	MPC	153 / 0 / 0	86.7	3.28 / 2.50 / 1.93	63.4 (3.59)	1440
three cuts	cascade	513 / 0 / 0	65.8	2.54 / 2.34 / 1.98	63.5 (4.42)	—
planned	cascade, giveaway by setpoint[a]	373 / 0 / 0	46.3	2.51 / 2.32 / 1.90	73.7 (5.27)	—
together	cascade, hindsight[b]	111 / 0 / 0	52.1	2.77 / 2.35 / 1.98	49.3 (2.81)	—
	cascade, tuned once by the MPC's rule[c]	256 / 0 / 0	80.5	2.90 / 2.34 / 1.97	54.5 (3.40)	—
	MPC, deethanizer cut only[d]	154 / 0 / 0	89.1	3.30 / 2.34 / 1.91	78.5 (4.88)	1440
5, 1200 min[e]	cascade	107 / 206 / 324	40.3	2.59 / 5.56 / 2.60	92.9 (4.84)	—
13 mol% ethane	MPC	122 / 0 / 0	111.3	3.55 / 2.41 / 1.90	50.0 (2.75)	1200
6, 1800 min	cascade	58 / 38 / 0	8.1	2.21 / 2.57 / 1.91	93.7 (4.73)	—
drift and step	MPC	109 / 0 / 0	86.7	3.34 / 2.41 / 1.90	64.8 (3.61)	1800
7, 960 min	cascade	0 / 0 / 0	0.0	1.42 / 2.32 / 1.90	80.5 (3.57)	—
target steps	MPC	0 / 0 / 16	0.0	1.68 / 2.32 / 2.08	77.1 (3.46)	960
9, 1440 min	cascade	0 / 124 / 50	0.0	1.68 / 3.07 / 2.02	112.9 (6.12)	—
heavier feed	MPC	52 / 140 / 197	7.7	2.23 / 2.73 / 2.43	101.4 (4.47)	1440
11, 1400 min	local, loops in manual, as shipped	783 / 0 / 0	1369.5	5.22 / 2.32 / 1.90	152.3 (10.72)	fails
challenge	MPC, loops in manual	198 / 0 / 0	118.8	3.38 / 2.32 / 1.90	174.6 (11.32)	1400, fails
	cascade, suggested	76 / 231 / 143	15.4	2.31 / 4.47 / 2.10	160.1 (9.31)	passes
	MPC, otherwise suggested	219 / 108 / 136	90.3	3.19 / 2.72 / 3.56	152.0 (7.99)	1400, fails

Minutes over and peaks are read from the true products at every physics step, and the excess integrates the ethane above 2.0 mol%. No plan was rejected. [a]Targets at the MPC's mean purities, 1.764, 2.247 and 1.656 mol%, and the inferential setpoint at 1.765. [b]The fewest minutes over among 216 fixed cascade tunings, $K_c = 1$ and $\tau_I = 120$ min on the trim, 0.6 on the local loop, 1 and 2.4 on the butane and pentane loops. [c]The same grid tuned once on scenario 10 by the rule of Section 10.4.3, 0.4, 120 min, 0.6, 0.5 and 1.2. [d]A copy of the controller that plans the deethanizer cut alone. [e]Scenario 5's ethane step run on the cascade instead of its product analyzer.

the cascade gives away 54.5 kmol with 256 min over, and the best of the 216 fixed tunings, chosen in hindsight, beats the MPC on every count, with 49.3 kmol, 111 min, 52.1 mol% min and a peak of 2.77 mol%. A firmer and faster trim, a gentler local loop and faster downstream loops buy more than the plan does. Figure 10.4 sets the six structures of scenario 10 side by side on minutes over the limit, peak and giveaway.

The MPC wins where its estimator or its coordination matters. After the drift its analyzer biases recalibrate the inferential within about 100 min, a correction the slow trim spreads over most of a shift, although on two of the six other seeds the MPC is also over the limit there, for 185 and 142 min at peaks of 2.05 and 2.06 mol%, against 386 to 396 min for the cascade. After the heavier feed it moves the debutanizer cut with the other two and gives away 16.2 against 25.9 kmol. Through scenario 5's ethane step the cascade's local loop drives the deethanizer cut to 21.85 % of feed and holds butane in the propane over its limit for 206 min, where the MPC keeps it inside, as do the hindsight and the once-tuned cascades of that case, whose local loop gain is 0.6 instead of 1.20.

Every ethane surge goes deeper under the MPC, with an excess of 86.7 against 8.1 mol% min in scenario 6 and 111.3 against 40.3 through scenario 5's step. The plan holds the bottoms disturbance it sees constant while the 1500 kmol surge drum keeps delivering the richer feed for about 90 min. In the ethane step of scenario 10 its present estimate of the ethane is within 0.004 in the logarithm on average, but its predictions 30 to 90 min ahead run low by 0.11 to 0.26 on average and by up to 1.1. No change of the move price, the band price, either estimator gain, the analyzer alignment or the coupling cures the depth, because across them the peak in scenario 10 stays between 3.03 and 3.60 mol%. A deethanizer move priced at 300 instead of 20000 still peaks at 3.08 mol% and puts butane over its limit for 124 min. Told the drum's eventual disturbance and its pace, which no

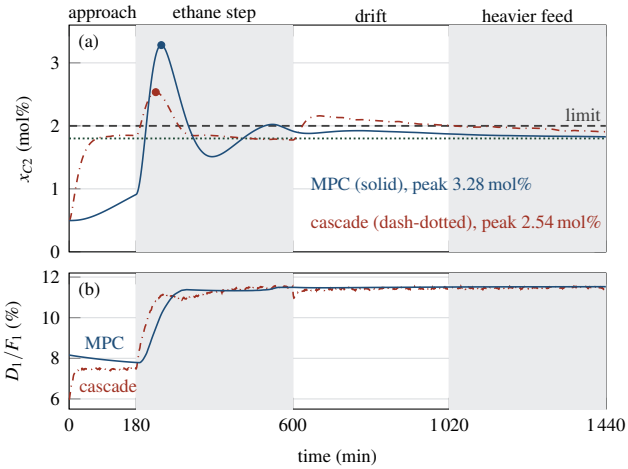

Figure 10.3 Scenario 10 under the MPC as it loads and under the cascade selected before the run, at the scenario's own noise seed, with the ethane step and the heavier feed windows shaded. The MPC is the solid line and the cascade the dash-dotted line in both panels. (a) True ethane in the propane product against the 2.0 mol% limit (densely dashed) and the 1.80 mol% target (densely dotted). The MPC peaks at 3.28 mol% at 247 min and the cascade at 2.54 mol% at 232 min. (b) The deethanizer cut ratio each structure applies.

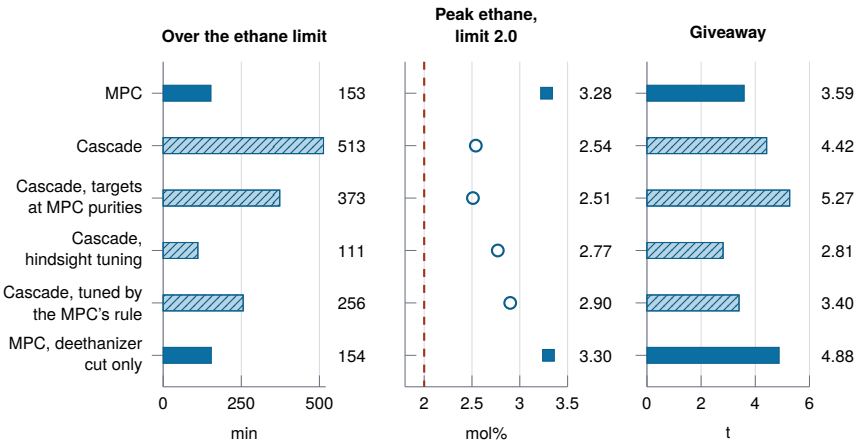

Figure 10.4 The scenario 10 rows of Table 10.6 as a chart: minutes over the 2.0 mol% ethane limit, the peak ethane in the propane against that limit (dashed line), and the giveaway in tonnes. MPC rows are solid bars and filled squares, cascade rows hatched bars and open circles. Notice that the two MPC rows spend far fewer minutes over the limit than the shipped cascade but reach the highest peaks, and that the cascade tuned in hindsight is the lowest on minutes and on giveaway.

instrument here measures, the same plan peaks at 2.24, 2.25 and 2.52 mol% in scenarios 10, 6 and 5, a bound and not a design, and with that information revealed 10 or 30 min late it peaks at 2.45 and 3.02 mol% in scenario 10.

The approach is slow as well, as shown in Figure 10.3(b). The MPC lowers the deethanizer cut from 8.16 to 7.79 % of feed over the first 180 min, and the ethane reaches 0.56, 0.72 and 0.91 mol% at 60, 120 and 180 min, where the cascade, whose local loop drops the cut to 5.96 % at once, is at 1.74, 1.84 and 1.84. The move price sets that pace. At 10000, 5000 and 2000 the plan reaches 1.24,

1.57 and 1.74 mol% by 180 min, but at 5000 and below butane goes over its limit in the surge, for 42 and 81 min.

Pentane is the other loss, over its limit for 16 min in scenario 7 where the cascade has none and for 197 against 50 min in scenario 9, with an excess of 52.7 against 0.5 mol% min, although the MPC gives away less in both. In scenario 7 the cause is the downstream move price, because priced at 10 or more instead of 1 the same plan keeps pentane inside. In scenario 9 it is the estimate, because priced at 1000 the plan still spends 190 min over, and while the pentane is over its limit the plan's estimate of it is low by 0.12 in the logarithm on average.

Judged by its own rule, the challenge is not met under the MPC. With the suggested line-up and the MPC selected, the worst errors from 924 min are 1.36 and 0.30 mol% against bands of 0.30 and 0.20, with every product inside its limit at the end, and it breaks limits that the rule leaves unscored, with 219, 108 and 136 min over the ethane, butane and pentane limits, a pentane peak of 3.56 mol% and butane in the natural gasoline over its limit for 369 min. As shipped, with the downstream loops in manual so that the plan moves the deethanizer cut alone, it fails at 1.53 and 1.77 mol% with a product outside its limit and gives away 174.6 against 152.3 kmol, the only case of Table 10.6 in which it gives away more than the structure it is set against.

10.6 Technology in Practice

Figure 10.5 draws the loops of Section 10.4 with the hardware of a fractionation train in place of the schematic's symbols.

What Measures It

Four products leave under five limits, and three of the limits have an instrument a loop can use. Ethane and butane in the propane are read by the chromatographs AI-C2 and AI-C4 on the depropanizer overhead and pentane in the butane by AI-C5 on the debutanizer overhead, product analyzers of the kind Chapter 9 describes, each on a 10 min cycle. Butane in the natural gasoline and propane in the ethane have limits and no loop. The one fast composition measurement is the deethanizer's local inferential, TI-101, a lagged estimate of the ethane left in its bottoms.

> **TECHNOLOGY** **Sensor: The Tray Temperature Behind the Inferential**
>
> Product compositions in a distillation column are rarely measured directly for control, because analyzers such as gas chromatographs bring long delays and high investment and maintenance costs. Tray temperature control has widely taken their place, on the assumption that the product meets its specification while a well-chosen tray holds its temperature. The assumption weakens when the feed composition or rate changes in a multicomponent column, and pressure changes move the temperature too. Remedies include placing the measurement where the nonkey components hardly change and using a temperature difference to compensate for pressure [19]. TI-101 of this chapter is such an inference, with its drift written as an explicit bias.

What Moves It

Seven final elements carry the chapter's handles. Three distillate valves set the cuts, each under a ratio station that multiplies its column's feed, so the quality controllers move ratios rather than flows. Three steam valves set the boilups, which stay economic setpoints of the operator, and the train feed valve holds the surge drum level under the slow averaging loop LC-101. When the header passes less than the reboilers ask, every boilup falls by the same factor, as Equation (10.3) writes.

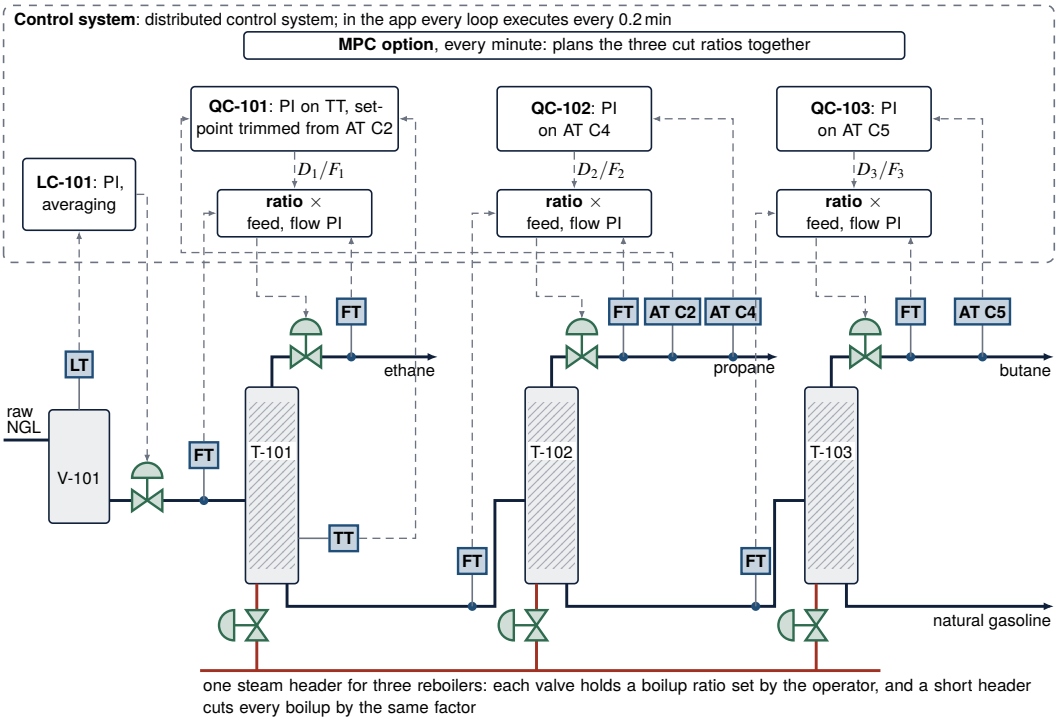

Figure 10.5 The loops of the NGL train with their hardware. Lettered boxes are sensing devices: LT on the surge drum, TT for the deethanizer inferential, AT C2 and AT C4 on the propane, AT C5 on the butane, and FT for the flow meters that the ratio stations need. Green symbols are final elements: the train feed valve, the three distillate valves and the three reboiler steam valves. Dashed arrows are signals. Notice the long arrow from AT C2 on the depropanizer overhead back to QC-101 on the deethanizer, and that every cut reaches its valve through a ratio station that reads the column's feed. Condensers, reflux drums and reboilers are left out, and a steam valve stands for each reboiler. Not to scale, and not a piping diagram.

> **TECHNOLOGY** **Final Element: The Distillate Valve Under a Ratio Station**
>
> A plant builds a cut ratio from a flow loop. The flow transmitter on the column's feed, the wild flow, is multiplied by the ratio in a ratio station, and the product becomes the setpoint of a flow controller that moves the distillate valve, the captive flow. The ratio holds only while the two transmitters' calibrations agree. Pressure-based flowmeters read well over a turndown often no more than three to one, perhaps ten to one with better transmitters, and a dulled orifice edge or uneven liquid in the impulse lines gives large errors at low flow [20]. The chapter's stations multiply exactly, so at a fixed feed composition a rate change is removed before any controller sees it.

Where the Controller Runs

A train like this runs on a distributed control system, which controls production within one location, and its supervisory and lower loops run with cycle times from milliseconds to minutes [21]. The control structure itself is layered by time scale: scheduling over weeks, site-wide optimization over a day, local optimization over an hour, supervisory and predictive control over minutes and regulatory control over seconds [2]. The app's loops execute every 0.2 min and its MPC every minute, while the chromatographs cycle every 10 min. The chapter's structure maps onto standard blocks: PI blocks for the quality and level loops, three ratio stations, a cascade in which the analyzer trim writes the

remote setpoint of the inferential loop and, in a real plant, a load shedding scheme or a multivariable controller where the app has its proportional header rule.

> ⚠ **IN PRACTICE** **Commission the Inner Loop First**
>
> A cascade is commissioned from the inside out. The secondary loop is tuned first, because it does not depend on the primary, while the primary can do nothing at all if the secondary is in manual. The secondary has three modes: manual, automatic on a local setpoint, and cascade on the remote setpoint the primary writes, and leaving it in automatic quietly defeats the cascade. The secondary must also respond faster than the primary [20]. Here the inner loop is the deethanizer's inferential loop with an integral time of 28 min, and the analyzer trim above it has 240 min.

What Is Modernizing It

Two developments touch this chapter. The first is the inference itself. Soft sensors are used to stabilize quality, to run closer to specifications and to validate online analyzers by comparison, and recursive models and model bias updates keep them current as a process changes [19]. The analyzer bias that the MPC of Section 10.4.3 updates at every result is of this kind. The second is predictive control, which sits in the supervisory layer above the regulatory loops [2]. The chapter's comparison shows what it still lacks on this train: a model of the surge drum's mixing, or a feed analyzer, so that it does not hold a disturbance constant while the drum is still delivering it.

10.7 Limitations

The default debutanizer cut of 49.80 % starts both products inside specification: pentane in butane is 1.899 mol% and butane in gasoline is 1.559 mol%.

The columns are shortcut splits, so nothing hydraulic can be studied. Flooding, weeping and efficiency loss at turndown are all absent, and a real train is often limited by one of them rather than by steam.

The composition dynamics are first-order lags fitted to the dominant time constant. The shape of the approach on a real column is not first order, and a student fitting a model from this application would find it easier than the same exercise on a rigorous simulation.

The four pseudo-components remove the isobutane and normal butane separation that a real debutanizer manages, and the natural gasoline vapor pressure specification is a mole fraction limit rather than a computed vapor pressure.

The steam header allocation has no priority scheme. That is the point of scenario 8, but it also means the application cannot demonstrate a well designed load shedding scheme working correctly, only the consequence of not having one.

The optional MPC has limits of its own. Its model is linear in the logarithms at one operating point, and at the design point the deethanizer cut moves the ethane 1.41 times as hard as it believes. It holds the disturbance at the deethanizer bottoms constant over the plan, with no model of the surge drum's mixing and no feed analyzer, which is why its surges go deeper. After large deethanizer moves its estimates of the downstream products run high, by up to 0.24 in the logarithm for butane and 0.37 for pentane in the ethane step of scenario 10, until analyzer results correct them, and its biases correct an offset, not a gain. It leaves the boilups, the steam header and the natural gasoline to the operator and the plant, so the arbitration missing in scenario 8 is still missing under it. Its limits are soft, with one slack per impurity for the whole horizon, and its prices were chosen once, on scenario 10.

Finally, there is no recycle anywhere in the train. One pass down it is the steady state, which is what makes the plantwide reasoning tractable and also what removes the snowball effect that plantwide control literature spends most of its attention on [1]. Volume 6, Chapter 7 studies that effect on a reactor and column joined by a recycle.

OPEN CHALLENGE **Keeping a Soft Sensor Honest**

Engineers surveyed in industry named model maintenance the most critical issue for soft sensors, and the problem cited most often was accuracy lost as process characteristics change. An estimate deteriorates as equipment changes and as feed properties vary, and rebuilding the models by hand each time is too heavy a workload. Recursive methods and bias updates follow slow change but cannot cope with abrupt change [19]. Progress would be the high-performance, maintenance-free virtual sensing that the review calls for, an inference whose drift shows before the giveaway does.

10.8 Discussion

Three results from this train generalize.

The first is that a nonzero process gain does not establish a pairing. The depropanizer cut moves ethane in the propane product by 2.7 mol% across the swept range, which is a large gain and a useless one, because the mechanism is dilution and the mechanism runs out. Before pairing a loop it is worth asking not whether the manipulated variable moves the controlled variable but by what physical route, and whether that route has an end.

The second is that a soft sensor fails in the direction that does not alarm. The drifted inferential in Table 10.4 settles the plant at 0.705 mol% against a 1.80 mol% target, which is inside every contract and costs money continuously. A failure that produces giveaway rather than an off-specification product is a failure a plant can run with for a long time, and the only defenses are a periodic comparison against a reference measurement and a supervisory layer above it.

The third is that a shared constraint resolved by arithmetic will be resolved against the specification that is least defended. The header cut all three reboilers by 8.8 % and each controller defended what it was given, and the product that left its contract and stayed out was the one with no controller on it. That is not a failure of any loop. It is a statement that the decision of which product to sacrifice was never made, and a control system that does not make it will make it implicitly.

10.9 Conclusion

A fractionation train is where loop pairing stops being a two by two exercise. The simulated train of this chapter measures ethane in the propane product on the depropanizer overhead and cannot change it there: the ethane flow in that distillate is 48.5 kmol/h across the entire cut range and the percentage falls only by dilution, which would require a distillate of 2430 kmol/h against a column feed of 918 to reach the contract. The valve is the deethanizer cut, one column upstream. Measured on the same disturbance, a deethanizer inferential drifted 30 percent low settles the propane at 0.705 mol% against a 1.80 mol% target while, at 30 h, the cascade above it reaches 1.751 and the analyzer alone 1.800, so the soft sensor failure costs giveaway rather than an alarm. A steam header trip from 5200 kmol/h to 4000 kmol/h multiplies every boilup by 0.912, leaves the three controlled impurities at 1.49, 2.29 and 1.90 mol% inside their limits at 15 h after excursions of up to 194 min, and takes butane in the natural gasoline to 3.81 mol% against a limit of 2.0, because that is the specification with no loop on it. The predictive controller offered as an option, planning the three cuts together, spends 153 against 513 minutes over the ethane limit in scenario 10 but peaks at 3.28 against 2.54 mol%, because it holds constant a disturbance that the surge drum is still delivering, and a cascade retuned in hindsight beats it on minutes, depth and giveaway. Further work would add a recycle so that the snowball effect becomes reachable, a load shedding scheme with an explicit priority, and a rigorous stage model on one column so that the shortcut split can be checked against it. On the train as it stands, a reader can find by experiment where each specification is made.

KEY POINTS

- The ethane in the propane is measured on the second column but made on the first, so the valve that sets it is one column upstream.
- A nonzero process gain does not establish a pairing: the depropanizer cut moves the ethane only by dilution, and the mechanism runs out.
- A soft sensor fails in the direction that does not alarm, and the drifted inferential settles the propane at 0.705 mol% against a 1.80 mol% target.
- A shared constraint resolved by arithmetic is resolved against the specification that is least defended, here butane in the natural gasoline.

Where else this applies. Volume 5, Chapter 3 closes a composition loop through a slow analyzer on one column, and Volume 5, Chapter 10 corrects a soft sensor's bias with a late measurement.

References

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

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

[3] ——, "Economic plantwide control," in *Plantwide Control*. John Wiley & Sons, 2012, pp. 229–251.

[4] W. L. Luyben, "Heuristics for plantwide control," in *Plantwide Control*. John Wiley & Sons, 2012, pp. 97–119.

[5] 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.

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

[7] B. F. Greek, "Natural gas liquids remain strong petrochemical feedstock," *Chemical & Engineering News Archive*, vol. 62, no. 11, pp. 17–38, 1984.

[8] L. Li, "Development of natural gas industry in China: Review and prospect," *Natural Gas Industry B*, vol. 9, no. 2, pp. 187–196, 2022.

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

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

[11] 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.

[12] F. Duanmu, D. N. Chia, and E. Sorensen, "A shortcut design method for complex distillation structures," *Chemical Engineering Research and Design*, vol. 180, pp. 346–368, 2022.

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

[14] U.S. Energy Information Administration, "U.s. natural gas plant field production, annual, thousand barrels per day," Petroleum and Other Liquids data series, release of October 30, 2026, 2026, accessed November 18, 2026. [Online]. Available: https://www.eia.gov/dnav/pet/pet_pnp_gp_dc_nus_mbblpd_a.htm

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

[16] C. A. Larsson, C. R. Rojas, X. Bombois, and H. Hjalmarsson, "Experimental evaluation of model predictive control with excitation (MPC-X) on an industrial depropanizer," *Journal of Process Control*, vol. 31, pp. 1–16, 2015.

[17] W. L. Luyben, "Control of a train of distillation columns for the separation of natural gas liquid," *Industrial & Engineering Chemistry Research*, vol. 52, no. 31, pp. 10 741–10 753, 2013.

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

[19] 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.

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

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

10.10 Exercises

10.1 Concept check · 2 min

In scenario 2 the raw NGL carries 13 mol% ethane. Taking the depropanizer cut ratio from 31.6 to 37.6 % of its feed lowers the ethane in the propane product from 16.7 to 14.0 mol%. Why is the depropanizer cut not a usable handle for the ethane specification?

(a) The gain is small, and a cut range a little wider than the slider's would reach the limit

(b) Ethane is lighter than the column's light key, so all of it goes overhead and the cut only dilutes it

(c) The analyzer cycle of 10 min is too slow for a loop closed on the depropanizer cut

(d) The cut also drags butane overhead, so an ethane loop would fight the butane loop

10.2 Concept check · 1 min

The deethanizer cut is held fixed and its boilup ratio is raised from 1.6 to 2.2. What happens to the ethane left in its bottoms and the propane lost overhead?

(a) The ethane falls and the propane rises, along the curve a larger cut also moves along

(b) Both rise, because more vapor carries more of each key component up the column

(c) Both fall, because a larger separation factor sharpens the split at the same cut

(d) Neither moves, because the cut alone fixes how the material balance is placed

10.3 Concept check · 2 min

In scenario 6 the local loop alone holds the deethanizer inferential at its 0.50 mol% setpoint when the inferential drifts to read 30 % low. The product target is 1.80 mol%. Where does the ethane in the propane settle, and what does the control room see?

(a) Above the 2.0 mol% limit, and the propane analyzer alarm sounds within the hour

(b) On the 1.80 mol% target, because the loop's integral action removes the bias

(c) Near 0.35 mol%, because a low reading makes the loop cut harder than needed

(d) Near 0.7 mol%, above what the setpoint intends but below target, with no alarm

10.4 Theory and calculation · 5 min

In scenario 2 the deethanizer bottoms carry 48.51 kmol/h of ethane into the depropanizer, whose feed is 918.4 kmol/h. The cut ratio is the distillate as a percentage of that feed.

(a) Find the distillate and the ethane in the propane at cut ratios of 31.6 and 37.6 %. *Checkpoint:* 290.2 and 345.3 kmol/h, 16.72 and 14.05 mol%.

(b) Find the distillate that would bring the ethane to 2.0 mol% by dilution alone, as a multiple of the column feed. *Checkpoint:* About 2425 kmol/h, 2.6 times the feed.

(c) Find the slope of the ethane percentage against the distillate at each cut, and explain why a nonzero gain is not enough to pair the loop. *Checkpoint:* -0.058 and -0.041 mol% per kmol/h.

10.5 Theory and calculation · 7 min

The deethanizer has $\alpha_1 = 2.10$, $N = 25$, $k_1 = 0.4703$ and $(V/F)_0 = 0.50$, and at design it is fed 80 kmol/h of ethane and 320 kmol/h of propane.

(a) Use Equation (10.2) to find η and S at boilup ratios 1.6, 1.8, 2.0 and 2.2, and the factor by which each step of 0.2 multiplies S. *Checkpoint:* $S \approx$ 1790, 4840, 11900 and 27200, factors 2.70, 2.47 and 2.28.

(b) At a boilup ratio of 1.8 the chapter reports 1.60 kmol/h of ethane left in the bottoms and 3.20 kmol/h of propane overhead, and 0.62 and 8.22 kmol/h with six more kmol/h of cut. Find the key recoveries and S from Equation (10.1) in each case. *Checkpoint:* About 4850 and 4860.

(c) At a boilup ratio of 2.2 and the design cut the chapter reports 0.45 and 2.05 kmol/h. Find S the same way and compare it with (a). *Checkpoint:* About 27400.

10.6 Theory and calculation · 7 min

In scenario 8 the three reboilers ask for 4387 kmol/h of boilup when the header capacity drops to 4000 kmol/h. The boilup setpoints are 1.80, 1.90 and 1.40.

(a) Find ϕ of Equation (10.3) and the boilup ratios the three columns receive. *Checkpoint:* $\phi = 0.912$, ratios 1.64, 1.73 and 1.28.

(b) The debutanizer has $\alpha_3 = 2.60$, $N = 25$, $k_3 = 0.4587$ and $(V/F)_0 = 0.50$. Find its S before the trip and after it. *Checkpoint:* About 3230 and 1290.

(c) The pentane loop holds pentane in the butane on target with the debutanizer cut. Use Equation (10.1) to say which way the cut must move once S has fallen, and what that does to the butane in the natural gasoline.

10.7 App challenge · 8 min

Run scenario 2, *The unit that measures it cannot fix it*, at the default seed. It loads in Manual with the raw NGL at 13 mol% ethane and the Depropanizer cut ratio at 35.14 % of feed.

(a) Set the Depropanizer cut ratio to 31.6 % before Run and run 90 min at speed 20 (20 min per second, about 5 s). Read the C2 in C3 and C4 in C3 badges, and in the exported CSV the columns `d2_kmolh` and `b1_kmolh` at 90 min. Compute the ethane flow in the distillate. *Checkpoint:* 16.72 % and 0.08 %, 290.2 and 918.4 kmol/h, 48.5 kmol/h of ethane.

(b) Set the cut ratio to 37.6 % and run another 90 min. Read the same four values and compute the ethane flow again. *Checkpoint:* 14.16 % (within 0.02) and 0.27 %, about 342.6 and 918.4 kmol/h, 48.5 kmol/h of ethane.

(c) From your readings, find the distillate that would put the ethane on 2.0 mol%, and say which valve sets this specification.

10.8 App challenge · 8 min

Run scenario 6, *The cascade, and an inferential that has drifted*, at the default seed. At 30 min the deethanizer inferential drifts to 0.70 of its reading and at 31 min the raw NGL steps to 12 mol% ethane. Run each case 1800 min at speed 120 (120 min per second, 15 s) and export the CSV.

(a) Run as shipped, on Cascade. Read `c2_in_c3_pct` at 900 and 1800 min, `c2_local_sp_pct` at 1800 min, and the Giveaway and Off-spec propane tiles at the end. *Checkpoint:* 1.595 and 1.751 %, 1.244 %, 93.7 kmol and 96 min.

(b) Reload, select Local in Ethane-in-propane control before Run, and read the same values. *Checkpoint:* 0.705 and 0.705 %, 0.50 %, 168.9 kmol and 39 min.

(c) Reload, select Product analyzer before Run, and read the same values. Then rank the three structures for a plant that sells propane against the 2.0 mol% limit. *Checkpoint:* 1.856 and 1.800 %, 0.50 % (unused), 62.5 kmol and 622 min.

10.9 App challenge · 8 min

Run scenario 8, *One header, three reboilers*, at the default seed. It loads on Cascade with both downstream loops in Auto, and at 60 min the header capacity falls from 5200 to 4000 kmol/h. Run 900 min at speed 60 (60 min per second, 15 s) and export the CSV.

(a) Read `steam_cut` in the first rows after 60 min, and `vf1`, `vf2` and `vf3` at 900 min. *Checkpoint:* 0.912, then 1.639, 1.730 and 1.275 (within 0.005).

(b) At 900 min, read the C2 in C3, C4 in C3, C5 in C4 and C3 recovery badges and the gasoline C4 on the debutanizer card. *Checkpoint:* 1.49, 2.29 and 1.90 %, 97.1 %, 3.81 %.

(c) Read the three Off-spec tiles at 900 min, and the time at which `c4_in_c5_pct` first exceeds 2.0. Say which product the trip sacrificed and why. *Checkpoint:* Propane 64 (within 1), butane 194 and gasoline 795 min, and a crossing at 105 min.

10.10 Application area · 3 min

After the boiler trip of scenario 8 the header shares 4000 kmol/h of boilup in proportion to what each reboiler asks for.

(a) Name the two kinds of layer the chapter says a real train uses instead, and what such a layer must know that the proportional rule does not.

(b) At 15 h ethane in the propane reads 1.49 % against 2.0, butane in the propane 2.29 against 2.5 and pentane in the butane 1.90 against 2.0. Which column would you ask to give up steam first to protect the gasoline, and what else would you check?

(c) Which simplifications of the model would make you distrust the 3.81 mol% figure on a real plant?

10.11 Application area · 4 min

A gas plant runs its deethanizer on a local loop that holds a pressure-compensated tray temperature inferential.

(a) Why does a 30 % low drift of that inferential raise no alarm in scenario 6?

(b) Name the two defenses the chapter gives, and say how the trim output of a cascade can itself serve as a drift monitor.

(c) In scenario 10 the predictive controller meets a 25 % drift. How quickly do its analyzer biases take it up, and how quickly does the cascade's trim?

10.12 Review the agent · 5 min

An AI assistant summarizes scenario 10: "Over the day the MPC spends 153 min over the ethane limit against the cascade's 513, and the Giveaway tiles read 63.4 against 63.5 kmol. The MPC is better on ethane at no cost in giveaway, so it should replace the cascade." Find the flaw, and say how you would check the claim in the app and in the chapter.

CHAPTER 11

Load Sharing by Equal Surge Margin Among Three Compressors on One Header

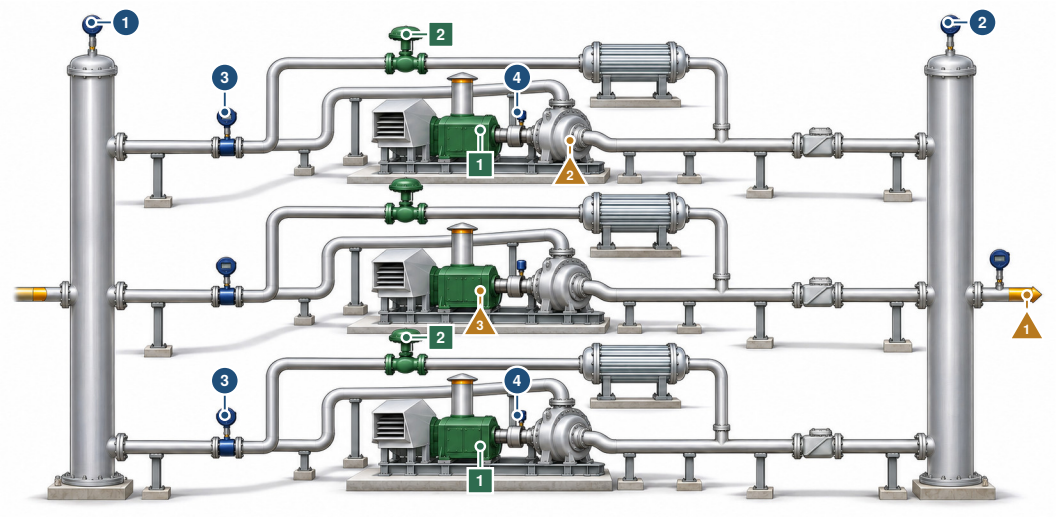

The controller moves 1 gas turbine on each unit, sped up or slowed by the station
2 recycle valve on each unit, which sends gas back to the inlet

The controller reads 1 suction header pressure 2 discharge header pressure 3 flow through each machine
4 speed of each shaft

What disturbs it 1 what the customers take 2 fouling of unit 1 3 a turbine trip on unit 2

Notice that each unit's recycle line leaves its discharge ahead of the check valve, passes the cooler and returns to the suction side, so gas it recycles is compressed for nothing. Illustration, not to scale.

IN PLAIN TERMS Machines That Share a Push, Not a Load

The main idea. A compressor station pushes natural gas down a pipeline with several spinning machines side by side. They all push into one pipe, so each must push as hard as the others, and the gas splits among them by itself.

The challenge. A machine that runs a little slower or is a little dirty carries less gas and drifts toward surge, where its flow stalls and swings backward. To stay safe it sends some gas back around to its own inlet, which burns fuel for nothing.

Why it matters. A good sharing rule keeps every machine out of surge without sending gas in circles.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History boxes in Section 11.2, the Did you know? boxes in Sections 11.3 and 11.5.7, the Try it boxes in Section 11.5, and Key points in Section 11.8. No equations.
Study	50 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 11.5 with Figure 11.4, Section 11.6 and Section 11.7 with its Open challenge box.

The control question Compressors on one header share a head, not a load, so a machine a little slower or a little fouled drifts toward surge while its neighbors take the gas. Which sharing law, start sequence and trip action keep every machine out of surge and off its recycle while the station holds its pressure?

Learning objectives

- Predict where the load goes when one of two machines on common headers changes speed.
- Compute a surge margin from the speed-scaled map, and the speed that gives a fouled machine a clean one's margin.
- Compare equal speed, equal flow and equal surge margin by the gas recycled and the fuel burned in a turndown.
- Explain why a trip needs its own valve action and why a fast station loop oscillates against anti-surge loops on recycle.
- Justify each layer of the challenge by the failure measured without it.

Prerequisites Sections 1.4.6 (map, fan laws and surge margin), 1.4.8, 1.4.5, 1.6.2 and 1.6.4, with Section 1.6.3 only for the MPC.

The app https://apmonitor.com/apps/compressor-station/

See also Volume 4, Chapter 2 for one machine with a Greitzer surge cycle and a protective override, Volume 4, Chapter 3 for a surge margin that is never measured, and Volume 4, Chapter 12 for one total allocated among unlike actuators. Volume 6, Chapter 2 for interacting loops and a static decoupler. Chapter 12 for the line the station feeds.

11.1 Introduction

Two identical compressors that draw from one suction header and deliver into one discharge header must make one head, because the header pressures fix it. In the generic station simulated here each machine carries 150 kg/s at 90.5 % of rated speed, 28.0 % to the right of its surge flow. Holding unit 1's speed and moving unit 2's setpoint to 83 % takes unit 2 to a margin of 12.4 % within 6 s. Five minutes later unit 1 passes 153.5 kg/s without having been asked to and unit 2 passes 127.4 kg/s. The load moved because a speed line crossed a fixed head further to the left, toward the peak of the characteristic, and not because any controller divided it.

A common belief holds that parallel machines share a load by themselves, and the simulator is designed to overturn it. Machines on a common header share a head, and their flows follow from their speeds and their maps. A machine that is a little slower, a little fouled or still coming up to speed drifts toward its surge line while its neighbor takes the gas, and the anti-surge loop that then opens its recycle valve compresses gas only to send it back. With one machine fouled by 6 % of its head and the demand cut to 245 kg/s, sharing by equal speed recycles 15.0 t of gas in an hour and sharing by equal surge margin recycles none. The class of problem is wider than gas transmission. Parallel pumps, fans, chillers and refrigeration compressors on a common header raise the same question of who carries the load, and the stability of pumping systems with two or more elements in parallel was surveyed with the other pumping instabilities in 1981 [1]. In a pipeline station the same question is posed as the distribution of a varying demand among the running units, together with when to start and stop them [2–4], a layer that divides a total among parallel actuators with unequal limits above local loops that protect each one. The car of Volume 4, Chapter 12 allocates one force among three actuators in the same way.

However, the literature a student is likely to meet takes the pieces one at a time. Surge is modeled on a single machine with lumped models of the Greitzer type [5, 6], station dynamics are studied by detailed transient simulation [7, 8], and load sharing is posed as an optimization over identified performance maps [3, 4]. This chapter places a station pressure loop, three sharing laws, three anti-surge overrides, start and trip sequences and a predictive controller around one simulated three-unit station on one clock, and measures what each does to the gas recycled, the fuel burned and the surge count. The speed-scaled map, the proportional-integral (PI) loops, the override and the model

predictive controller (MPC) are established methods. The results are measurements on one generic station at one seed. Among them, a station loop tuned fast enough to ring and settle on clean machines holds a limit cycle of 0.33 bar once both units recycle, changed only in its fifth significant figure when the integration step is halved, and the predictive controller matches equal-margin sharing rather than improving on it.

Section 11.2 follows the domain, from the optimization of pipeline operation to the load-sharing problem. Section 11.3 sets out the station, a twenty-state model whose map and surge margin equal those of the single-compressor plant of Volume 4, Chapter 2 term for term, given to its last default, with controllers that read only instruments. Section 11.4 describes the control layers. Section 11.5 measures nine cases, each with a completion rule that requires both a change the reader makes and a state the station then reaches, and the six-hour challenge. Section 11.6 sets the hardware of a real station beside the model, Section 11.7 states the limitations, and the exercises in Section 11.9 work through the cases on the app.

11.2 Parallel Machines in Gas Transmission

Gas pipelines are powered by compressor stations [2], and the History box below tells how their operation became an optimization subject and their surge protection a technology of its own.

> ⧖ HISTORY **1968 and 1994: Optimizing the Line, Guarding the Machine**
>
> The expense of operating natural gas pipelines made their operation an optimization subject early. In 1968 dynamic programming was applied both to the transient distribution of gas at a terminal and to the steady transmission problem, because the inherent nonlinearities and numerous constraints of a pipeline system suited the method [9]. Protecting the machines took a path of its own. A technology assessment of 1994 described the surge schemes in use on transmission compressors as detection devices, which anticipate surge or catch its inception, joined to control devices that act to prevent it. Its patent search over the previous twenty years found most of the effort in surge control rather than in detection, and a trend likely to continue toward incremental improvement of the conventional scheme [10].

How the units of one station should be controlled together was argued with field operation in view. A 1986 analysis of parallel centrifugal gas compressors compared the common approaches to their control by energy efficiency, stability of control, equipment protection and level of automation, and described an improved system with an example from field operation [11]. Detailed simulation of station dynamics followed. A 1991 model of a station combined pipe flow, the quasi-steady flow through valves, constrictions and compressors and the control system [7]. Whether two units should run in series or in parallel was examined by transient simulation of a two-unit station with variable-speed drivers in 2002 [12]. The emergency shutdown emerged as the hard case for anti-surge systems, as the second History box tells.

> ⧖ HISTORY **1994 and 2004: When a Station Shuts Down**
>
> Models of whole stations made the shutdown a subject of its own. A model of 1994 joined the gas dynamics, the control system and the inertias of the compressor and turbine shafts. It showed that a cooler in the recycle path changes how fast a compressor shaft decelerates on shutdown and, depending on the shaft's inertia against the cooler's volume, can spin the shaft backward, which harms its dry seals [8]. Ten years later the data of a compressor that surged during an emergency shutdown were analyzed. Most anti-surge systems avoid surge in normal operation, the study found, but unplanned emergency shutdowns are a significant challenge, and avoiding surge in them depends largely on the station's layout. Its model gave simpler rules for sizing the piping on either side of the machine and the control elements [13].

The load distribution itself became an optimization problem. A 2013 formulation optimized both the load distribution and the on and off switching of parallel units in pipeline stations, with the

Table 11.1 Variables by Role, Design Values With Two Units on Line

Symbol	Role	Design	Range
P_d	CV, discharge header (bar a)	70.00	SP 60–76
SM_i	CV, surge margins (%)	28.05	line 3–25
$n_{SP,i}$	MV, speed setpoints (% rated)	90.47	70–105
$v_{cmd,i}$	MV, recycle valve commands (%)	0	0–100
F^*	MV of the PIC, flow demand (kg/s)	300	0–250N
w_{del}	DV, far-end demand (kg/s)	300	100–500
P_{up}	DV, upstream line (bar a)	47.5	40–55
a_i	DV, head lost to fouling (%)	0	0–12
n_i, g_i	states, speed, turbine (%)	90.47, 74.73	—
v_i, θ_i	states, valve, surge phase	0, 0	—
P_s, P_1, P_2	states (bar a)	45.00, 67.85, 63.34	—
$w_{c,i}$, $w_{f,i}$, $w_{r,i}$	outputs, flows (kg/s)	150, 150, 0	—
h_p, Q	outputs (kJ/kg, m^3/s)	60.13, 4.202	—
M_{pack}	output, gas in the line (t)	1324.05	—

effect of each start and stop on the machines' expected life among the costs [2]. Online tracking of each unit's performance map was then joined to the load sharing optimization [3], and real-time optimization of the sharing under uncertainty followed in 2020 [4]. Model predictive control has meanwhile been applied to the process and anti-surge loops of a single centrifugal machine [14]. A degraded machine changes all of this, because the gas turbine and its driven compressor degrade as one system and change how they match each other and the pipeline [15].

11.3 The Station as Modeled

No station stands behind the simulated one, and every size, rating and threshold is a teaching value. An upstream line at 47.5 bar a feeds a 300 m^3 suction header. Three identical 15 MW gas turbines drive three centrifugal compressors that lift the gas into a 300 m^3 discharge header held at 70.0 bar a. Each machine has a discharge line with a check valve and a recycle line that returns gas from ahead of the check valve through a cooler and an anti-surge valve to the suction header. A 40 km, 0.9 m line in two cells of 20 km carries the gas to customers at the far end. The gas is 17.4 kg/kmol with $Z = 0.90$ at 20 °C. Pressures are in bar absolute, flows in kg/s, heads in kJ/kg, powers in MW and time in seconds. Figure 11.1 shows the layout and every symbol of the equations below, with the values of case 4 at 16 min, and Figure 11.2 shows the three maps at the same instant.

> ✦ **DID YOU KNOW? Up to Sixteen Machines in One Station**
>
> Mainline transmission compressors are usually rated at 1,000 horsepower or more, centrifugal machines driven by turbines or reciprocating ones. The larger stations may hold 10 to 16 units, 50,000 to 80,000 horsepower in all, and pass more than three billion cubic feet of gas a day [16]. Most burn gas drawn from the line they drive.

The model integrates twenty states. Per unit there are the speed n_i as a fraction of rated speed, the power turbine's output g_i as a fraction of its rating, the recycle valve position v_i and a deep-surge phase θ_i. The four holdups are the suction header P_s, the discharge header P_d and the two line cells P_1 and P_2. Four ledgers integrate the gas supplied and delivered, the fuel and the gas recycled, so the mass balance and the scorecard's totals close exactly. The variables are listed by role in Table 11.1.

11.3.1 One Map, Scaled by the Fan Laws

Every machine uses the characteristic of the single-compressor plant of Volume 4, Chapter 2, the Moore–Greitzer cubic [17] with its published shape, scaled to a pipeline machine and to any speed

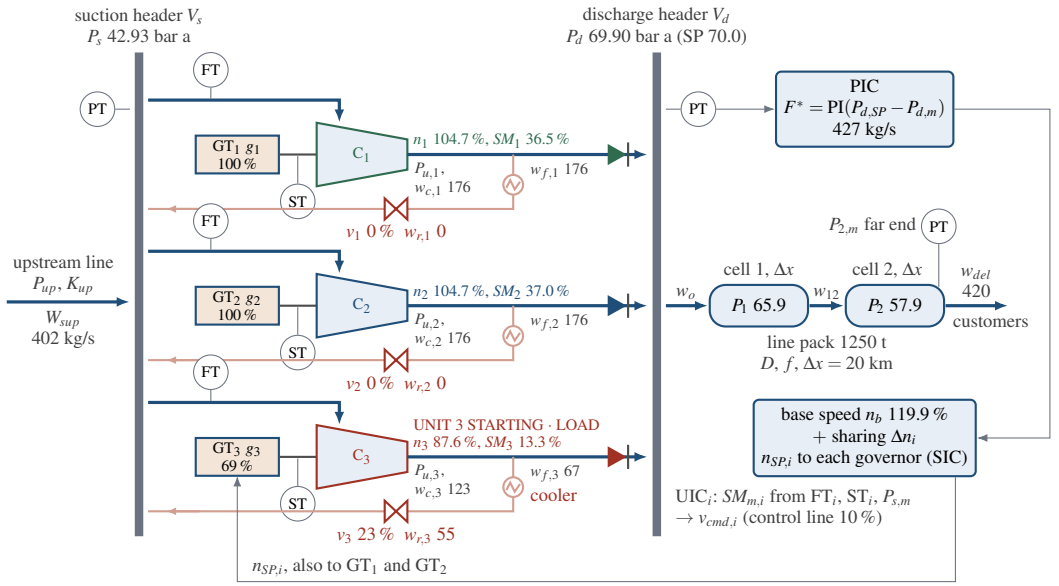

Figure 11.1 The station as modeled, drawn after the Station view of case 4 at $t = 16$ min, with *Start a unit* pressed at 13.7 min: units 1 and 2 on line at 104.7 % speed with their recycle valves shut, unit 3 loading at 87.6 % with 123 kg/s through the machine, 67 kg/s forward through its check valve and 55 kg/s back through its cooler and anti-surge valve, 23 % open. Flow elements (FT), speed pickups (ST) and pressure transmitters (PT) feed the station pressure loop (PIC), the governors and the anti-surge loops (UIC). Flows in kg/s. Process model only, not to scale.

by the fan laws of Equation (1.10), as shown in Equation (11.1).

$$\begin{aligned} \psi_c(\phi, n) &= (1-a)\, n^2\, \Psi(\phi/n) \\ \Psi(t) &= \psi_0 + H\left[1 + \tfrac{3}{2}\left(\frac{t}{W} - 1\right) - \tfrac{1}{2}\left(\frac{t}{W} - 1\right)^3\right] \end{aligned} \tag{11.1}$$

where $\phi = Q/Q_{ref}$ is the inlet volume flow, $\psi = h_p/H_{ref}$ the polytropic head, $\psi_0 = 0.2$, $H = 0.18$, $W = 0.25$, $Q_{ref} = 7.255\,\mathrm{m^3/s}$, $H_{ref} = 160\,\mathrm{kJ/kg}$ and a the fraction of head lost to fouling. The peak of Ψ lies at $\phi/n = 2W$, so the surge line is the parabola through the peaks and the margin is the definition of Section 1.4.6 taken at the present speed, as shown in Equation (11.2).

$$\begin{aligned} \phi_s(n) &= 2Wn, \qquad \psi_s(n) = (1-a)\, n^2 (\psi_0 + 2H) \\ SM &= 100\,\frac{\phi - \phi_s(n)}{\phi_s(n)} = 50\,(x-1), \qquad x = \frac{\phi}{nW} - 1 \end{aligned} \tag{11.2}$$

At $n = 1$ and $a = 0$ both functions equal those of the single-compressor plant term for term. Its own model, imported read only by the technical review's script, gives a largest difference of exactly zero in ψ_c over ϕ from 0.30 to 0.85 and in the margin, with the surge point at $\phi = 0.5$, $\psi = 0.56$ in both. The fan laws hold as written. A similar point ($x = 1.5$) at 70, 80, 90, 100 and 105 % speed keeps $Q/n = 4.5344\,\mathrm{m^3/s}$, $h_p/n^2 = 77.000\,\mathrm{kJ/kg}$ and a margin of 25.0000 %. An alternative convention measures the distance at the same head, in reduced coordinates, and at the design point it reads 41.4 % where Equation (11.2) reads 28.0 %. The choke line at $\phi/n = 0.70$ is drawn as a guide, and the cubic is not cut there. It falls to zero head at $\phi/n = 0.8037$.

Fouling multiplies the head at every flow coefficient by $1 - a$, removes $a/2$ of the efficiency and leaves the surge flow at each speed where it was, so the anti-surge loops' commissioned surge

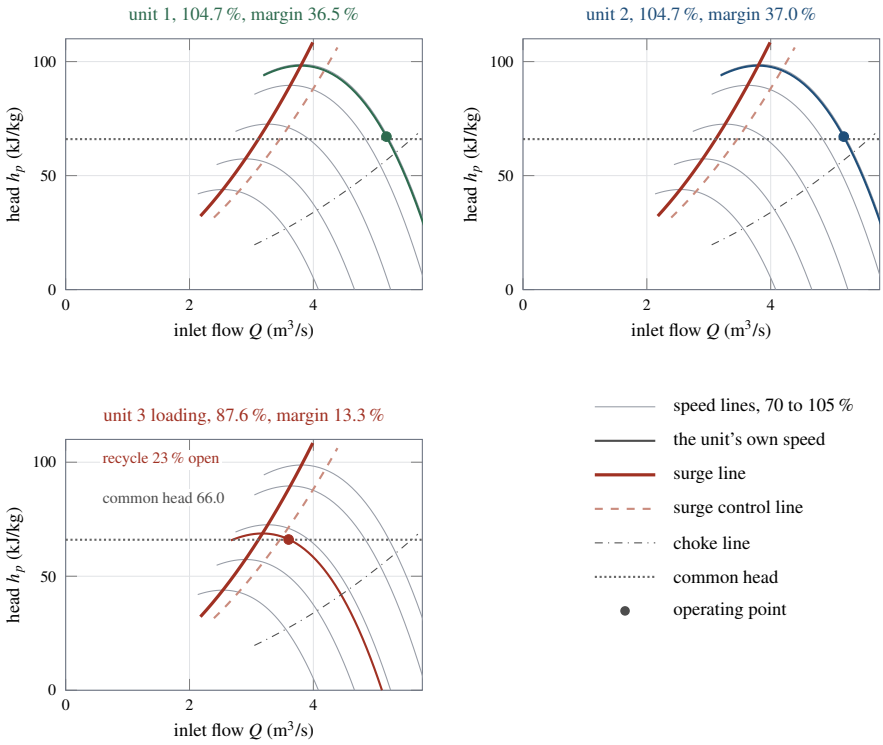

Figure 11.2 The Maps view of case 4 at the instant of Figure 11.1: speed lines from 70 to 105 % (thin), each unit's own speed line (thick), the surge line (very thick), the surge control line 10 % to its right (dashed), the choke line (dash-dotted), the common head of 66.0 kJ/kg between the headers (dotted) and each operating point (dot). The key in the lower right names the line styles.

line stays true. This is a teaching simplification. Deposits degrade the gas turbine and its driven compressor together [15], and whether they also move a centrifugal stage's surge line depends on where they sit. A machine whose surge line had moved toward higher flow would surge while its controller still showed margin, and the app keeps that out so that case 5 teaches one thing.

11.3.2 One Machine Between Its Headers

The gas in an impeller and its piping settles in far less than a physics step, so each machine sits on its characteristic at the head that its own discharge pressure P_u asks for. The head, the flow the map passes and the balance that fixes P_u are given by Equation (11.3).

$$\begin{aligned} h_p &= \frac{ZRT/M}{\sigma}\left[\left(\frac{P_u}{P_s}\right)^{\sigma} - 1\right], \qquad w_c = \rho_s Q_{ref}\, nW(x+1) \\ 0 &= w_c(P_u) - w_r(P_u) - w_f(P_u) \\ w_r &= K_r v\sqrt{P_u^2 - P_s^2}, \qquad w_f = K_d\sqrt{P_u^2 - P_d^2}\ \ (P_u > P_d) \end{aligned} \tag{11.3}$$

where $\sigma = (m-1)/m = 0.2814$, with m the polytropic index, is held at its design value, $ZRT/M = 126.07\,\mathrm{kJ/kg}$, ρ_s the suction density, $x = \phi/(nW) - 1$ the value on the falling branch at which Ψ equals $h_p/(H_{ref} n^2(1-a))$ (a cubic in x solved in closed form), $K_r = 4.44\,\mathrm{kg/s}$ per bar the recycle valve wide open and $K_d = 20\,\mathrm{kg/s}$ per bar the discharge line and check valve, which passes no gas backward. The root is found by an Illinois regula falsi between P_s and the pressure at the curve's peak, each term being monotone in P_u. The machine's flow is then set to $w_r + w_f$, so the unit holds no

gas and the header balances close whatever the solver's tolerance. A stopped, tripped or purging unit has its block valve shut and no forward flow. At design the machine makes 60.13 kJ/kg at 70.40 bar a, while the head between the headers, the dotted line on the maps, is 59.31 kJ/kg, because the discharge line takes 0.40 bar at 150 kg/s.

11.3.3 Surge as a Prescribed Cycle

When no root exists below the peak, the machine's outlets cannot take the peak flow at the peak pressure and the machine is in surge. The flag latches after the step and the flow then follows the cycle of Equation (11.4).

$$w_c = \rho_s Q_{ref}\, 2Wn\left(0.4 + 0.8\sin 2\pi\theta\right), \qquad \frac{\mathrm{d}\theta}{\mathrm{d}t} = \frac{1}{T_{surge}} \tag{11.4}$$

with $T_{surge} = 2\,\mathrm{s}$, a teaching value, so the flow falls to –40 % of the surge flow once a cycle. Surge clears when the machine has a steady point at least 2 % right of its surge line. The single-compressor plant resolves this cycle with the Greitzer model [5], while a station stepping at 0.05 s cannot carry each machine's Helmholtz dynamics, so here the cycle is imposed. The onset is therefore the static one, the loss of a steady point right of the peak. With that plant's cubic and a throttle parabola, Greitzer's dynamic onset lies 10.5, 4.6, 1.1 and 0.10 % of flow left of the peak at B of 1, 1.5, 3 and 10, and a 300 m^3 discharge header makes B large. Unit 2 slowed by hand with its recycle held shut latches surge at 80.68 % speed with its last steady margin at 1.55 %, because the margin falls as the square root of the distance to the fold and one step crosses it.

11.3.4 Shaft, Gas Turbine and Fuel

Each shaft follows its torque balance, written in power as Equation (11.5).

$$\begin{aligned} I\,n\frac{\mathrm{d}n}{\mathrm{d}t} &= gP_r - \frac{w_c h_p}{\eta} - k_m P_r n^2, \qquad I = 2H_{in}P_r \\ \tau_{gt}\frac{\mathrm{d}g}{\mathrm{d}t} &= g_{cmd} - g, \qquad \tau_v\frac{\mathrm{d}v}{\mathrm{d}t} = v_{cmd} - v \\ \eta &= \mathrm{clip}\left(\eta_{max} - \kappa_\eta(\phi/n - \phi_\eta)^2,\ 0.3,\ 0.95\right)(1 - a/2) \\ Q_{fuel} &= (P_r/\eta_{gt})\left(a_0 + (1 - a_0)\,g\right) \end{aligned} \tag{11.5}$$

where $P_r = 15\,\mathrm{MW}$, $H_{in} = 4\,\mathrm{s}$ ($I = 120\,\mathrm{MJ}$), $k_m = 0.01$, $\tau_{gt} = 1.5\,\mathrm{s}$, $\tau_v = 1\,\mathrm{s}$, $\eta_{max} = 0.82$ at $\phi_\eta = 0.60$ with $\kappa_\eta = 4$, the turbine efficiency at rating $\eta_{gt} = 0.35$ and $a_0 = 0.30$ of rated fuel at zero output, a Willans line. In surge the efficiency is halved. At design each machine runs at $\eta = 0.8135$ with 11.09 MW of gas power, its turbine at 74.7 % of rating and burning 35.28 MW of fuel. An overspeed trip at 110 % is a teaching value.

11.3.5 Headers, Supply and Line Pack

The headers and the two cells are isothermal holdups, $C = V/c^2$ with $c^2 = ZRT/M$, and the flows between them follow the steady friction law of the gas transmission line of Chapter 12, as shown in Equation (11.6).

$$\begin{aligned} C_s\frac{\mathrm{d}P_s}{\mathrm{d}t} &= W_{sup} - \textstyle\sum_i w_{f,i}, \qquad C_d\frac{\mathrm{d}P_d}{\mathrm{d}t} = \textstyle\sum_i w_{f,i} - w_o \\ C_c\frac{\mathrm{d}P_1}{\mathrm{d}t} &= w_o - w_{12}, \qquad C_c\frac{\mathrm{d}P_2}{\mathrm{d}t} = w_{12} - w_{del} \\ P_a^2 - P_b^2 &= Kw|w|, \qquad K = fLc^2/(DA^2) \end{aligned} \tag{11.6}$$

where W_{sup} comes from the upstream line with $P_{up}^2 - P_s^2 = K_{up}W_{sup}|W_{sup}|$ and $K_{up} = 2.569\times10^{-3}$ bar^2 per $(\mathrm{kg/s})^2$, w_o crosses half a cell and w_{12} a whole one, $f = 0.0095$, $D = 0.9\,\mathrm{m}$, A the pipe's cross section, and L in meters, K being divided by 10^{10} to work in bar. With $c = 355.07\,\mathrm{m/s}$ each

header holds 237.96 kg per bar and each cell 10 092 kg per bar. The customers take their demand w_{del} from the far cell unless it falls below 40 bar a, below which delivery is curtailed linearly to zero at 35 bar a. Recycled gas returns to the suction header, so only the forward flows enter the balances. At design the line holds 1324.05 t, and over the shipped challenge, with a surge, a trip, a restart and the discharge below the suction, the inventory closes to 9.1×10^{-8} kg on 7009 t moved.

11.3.6 Instruments

The controllers read one packet. The three pressures carry 0.01 bar of colored noise (time constant 0.5 s), each machine's flow element 0.4 % of reading (0.3 s), and each speed pickup 0.05 % of rated speed (0.2 s), all from the seeded generator, and the turbine output is read directly. The anti-surge loops compute the margin from the suction pressure, the flow element and the speed pickup alone, as shown in Equation (11.7).

$$SM_{m,i} = 100\left(\frac{w_{m,i}}{\rho(P_{s,m})\,Q_{ref}\,2W\,n_{m,i}} - 1\right) \tag{11.7}$$

One scan's measured margin carries 0.54 point of noise (one standard deviation). The packet carries no fouling, demand or supply value, and of the line's pressures only the far end's, and one of the app's tests fails if a controller touches any other key. In the exported file the simulator's own values carry the prefix `truth_`, beside the measured ones.

11.3.7 Parameters, Initialization and Numerics

Table 11.2 gives every parameter with its default. Every case begins at a steady state set by its own demand and settings. The line's pressures follow from the throughput by the friction law, each loaded machine's speed comes from the case's sharing law by bisection, and when clean machines would sit below the control line their recycle valves are placed where the anti-surge loops would hold them. Fourth-order Runge–Kutta steps of $\Delta t = 0.05$ s advance the states (Section 1.6.1), the controllers execute every 0.1 s and the historian records at 2 Hz, so a 2 s surge cycle shows on the trend. Halving the step changes the challenge's recycled gas by less than 0.001 t and its worst pressure deviation by less than 0.0001 bar, and leaves the time of case 7's surge and case 4's count of surges and cycles unchanged.

11.3.8 Assumptions

Each machine is quasi-steady on its map, which a station model with transient pipe flow also assumes for its compressors [7], and fails within a surge cycle, which is prescribed. Rotating stall is absent. The headers and the line are isothermal at 20 °C, so the recycle coolers are perfect and the heat of compression is not tracked, and the line is two friction-dominated cells with no momentum states and hence no pressure waves. The three machines share one map and the polytropic exponent is held at its design value. Fouling is a single head factor with the surge line fixed. The fuel each turbine burns is metered and not drawn from the suction header. There is no station bypass. If the discharge header falls below the suction header, gas passes straight through the machines and their check valves with no head made, a regime that only the shipped challenge reaches (Section 11.5.7). Every safety number, the 10 % control line, the speed limits, the overspeed trip, the surge period and the curtailment pressure, is a teaching value. API Standard 617 covers the design of centrifugal compressors and API Standard 670 their machinery protection, and a station's own surge tests and protection studies set its real numbers.

11.4 Control Layers as Implemented

Each physics step begins with the controllers, then the states advance, as in Section 1.6.1, and the surge latch, the trips and the challenge's scoring run afterward. The controllers execute every $\Delta t_c = 0.1$ s on the packet of Equation (11.7). All the PI loops use the velocity-form PID of Section 1.6.2, without derivative action. Every loop acts on $e = r - y$, with r the setpoint and y the measurement, and its

Table 11.2 Parameters and Defaults. A, Adjustable in the Advanced Panel. C, a Control on a Card. F, Fixed in the Source

Symbol	Meaning	Value	Unit	
M, Z, T	gas	17.4, 0.90, 293.15	kg/kmol, –, K	A
σ	polytropic exponent	0.2814	—	A
ψ_0, H, W	map shape	0.20, 0.18, 0.25	—	A
H_{ref}, Q_{ref}	map scale	160, 7.255	kJ/kg, m^3/s	A
ϕ_{ch}	choke line, ϕ/n	0.70	—	A
$\eta_{max}, \kappa_\eta, \phi_\eta$	efficiency curve	0.82, 4, 0.60	—	A
P_r, H_{in}	turbine rating, inertia	15, 4	MW, s	A
$\tau_{gt}, \eta_{gt}, a_0$	lag, fuel line	1.5, 0.35, 0.30	s, –, –	A
k_m	mechanical loss	0.01	—	A
—	overspeed trip	110	%	A
K_r, τ_v	recycle valve, lag	4.44, 1.0	kg/s/bar, s	A
K_d	discharge line	20	kg/s/bar	A
T_{surge}	deep-surge period	2.0	s	A
—	surge clears at	2	%	A
V_s, V_d	header volumes	300, 300	m^3	A
P_{up}, K_{up}	upstream line	47.5, 2.569e-3	bar a, see (11.6)	A
$\Delta x, D, f$	line cells	20, 0.9, 0.0095	km, m, –	A
—	curtailment below	40	bar a	A
$\dot{v}_{max}$	recycle stroking	50	%/s	C
$t_{purge}, t_{acc}, t_{load}$	start	20, 60, 120	s	C
t_{unload}	stop sequence	120	s	C
t_{up}, t_{down}	sequencer	60, 300	s	C
—	minimum run, rest	600, 900	s	C
—	noise on P, w, n	0.01, 0.4, 0.05	bar, %, %	F
—	speed range, base clamp	70–105, 55–120	%	F

output is clipped to its range and then rate-limited where the loop declares a stroking rate. The tracking time of Equation (1.19) equals τ_I, so when the output saturates, back-calculation adjusts the integral and nothing else. In Manual the PI solves its integral for the operator's value, so every transfer is bumpless.

11.4.1 Station Loop and Base Speed

The station PIC turns the discharge pressure error into a station flow demand F^* in kg/s, with $K_c = 20\,\text{kg/s}$ per bar, $\tau_I = 60\,\text{s}$ and an output from 0 to 250 kg/s per loaded unit. The demand becomes one base speed for every loaded unit, the speed at which a clean machine passes F^*/N on the commissioned map at the measured pressures, so the loop's gain is nearly constant from turndown to full speed. The base speed is clamped to 55 to 120 %, past the governors' 105 % limit, so that a unit trimmed below the base can still rise. In case 4 it reaches 119.9 % while the discharge sags. When a unit joins or leaves the shared load, the base speed is reset to the loaded units' mean setpoint, their trims to their offsets from it, and the PIC's output to the flow that mean passes, so no setpoint jumps. With the header on Manual the operator holds the three speed setpoints and the PIC tracks what they amount to.

11.4.2 Three Sharing Laws

Each loaded unit's speed setpoint is the base plus a trim, as shown in Equation (11.8).

$$\begin{aligned} n_{SP,i} &= \text{clip}(n_b + \Delta n_i,\ 0.70,\ 1.05) \\ \frac{\mathrm{d}\Delta n_i}{\mathrm{d}t} &= k_b(\bar{e} - e_i), \qquad \textstyle\sum_i \Delta n_i = 0 \end{aligned} \tag{11.8}$$

where $\bar{e}$ is the mean of the loaded units' e_i and $k_b = K_b/(100 \cdot 60)$, with the sharing rate $K_b = 5\,\%$ of speed per point of difference per minute. Equal surge margin sets $e_i = SM_{m,i}$ and equal flow (sharing by capacity, the machines being identical) sets e_i to each unit's measured flow as a percent of the

mean. A trim never pushes a unit further into a speed limit, each is held within $\pm 15\,\%$, and after each scan the trims are re-centered to sum to zero, so the sharing never fights the pressure loop for the total. Equal speed lets the trims bleed away with a time constant of $600/K_b = 120\,\mathrm{s}$. Each governor is a PI from speed setpoint to turbine output, $K_c = 4\,\%$ per percent and $\tau_I = 4\,\mathrm{s}$.

11.4.3 Anti-Surge Overrides

Each unit's anti-surge loop is the single-compressor plant's override on the measured margin, the protective override of Section 1.4.8 that rests at its clamp in healthy operation, a PI from $SM_{SP} - SM_{m,i}$ to the recycle valve with $K_c = 3\,\%$ of valve per point of margin, $\tau_I = 5\,\mathrm{s}$, an output of 0 to 100 % and a stroking limit of 50 %/s, with the surge control line $SM_{SP} = 10\,\%$. Above the line the output rests at zero, and because the tracking time equals τ_I the error terms cancel there and the integral decays toward zero instead of winding down, so the valve opens at the first scan below the line. The loops have their own Auto and Manual switch. A start forces the valve open through the purge and acceleration, a stop through the run-down, and *Open at once* drives a tripped unit's valve open in the trip's own scan.

11.4.4 Starts, Stops, Trips and the Sequencer

A start purges for 20 s with the block valve shut and the recycle open, lights the turbine and ramps the speed setpoint to 70 % in 60 s on full recycle, then loads the unit by ramping its setpoint to the base speed over the load time (120 s by default) while its anti-surge loop closes the recycle from wide open, and joins it to the sharing at the end of the ramp. Setting *Recycle while loading* to *Shut at once* instead closes the valve when loading begins. A stop takes the unit out of the sharing, unloads it to 70 % over 120 s, then cuts its fuel and opens its recycle. A trip cuts the fuel at once. The *Stop a unit* helper stops the running unit with the lowest measured margin. The optional sequencer starts the next stopped unit when every loaded setpoint has sat at 105 % with the measured discharge more than 0.3 bar low for 60 s, and stops the highest-numbered running unit when one fewer could carry the measured forward flow at rated speed for 300 s, with the framework's shared stage sequencer holding a minimum run of 600 s, a minimum rest of 900 s and a minute between changes.

11.4.5 Model Predictive Control on the Speeds

The MPC replaces the PI and the sharing law, in the general form of Section 1.6.3. Every 1 s it linearizes the discharge header about the measured point, from the clean map at the measured pressures and speeds, as shown in Equation (11.9).

$$\begin{aligned} C_d \dot{P}_d &= \textstyle\sum_i k_i (n_i - \bar{n}_i) + C_d b \\ &\quad - \left(k_o + \textstyle\sum_i k_{P,i}\right)(P_d - \bar{P}_d) \\ \tau_n \dot{n}_i &= n_{SP,i} - n_i \\ SM_i &= \bar{SM}_i + a_i (n_i - \bar{n}_i) - c_i (P_d - \bar{P}_d) \end{aligned} \tag{11.9}$$

where k_i, $k_{P,i}$, a_i and c_i are central differences of the clean map, k_o is the outlet line's conductance at the measured flow, $\tau_n = 3\,\mathrm{s}$ stands for the governors, and b is the measured trend of P_d filtered over 5 s, the drift by which the line pack it cannot see enters the model. The exact zero-order-hold discretization at 1 s enters the quadratic program of Equation (11.10) over 60 steps in 12 graduated move blocks.

$$\begin{aligned} \min\ & \textstyle\sum_k w_P (P_{d,k} - P_{d,SP})^2 + w_{\Delta u} \sum_j \|\Delta n_{SP,j}\|^2 \\ & + w_{fuel} \textstyle\sum_{k,i} \gamma_i n_{SP,i,k} + \rho\, \varepsilon^2 \\ \text{s.t. } & 70 \le n_{SP,i} \le 105, \quad |\Delta n_{SP,i}| \le 1\,\% \text{ per step} \\ & SM_{i,k} \ge SM_{min} - \varepsilon, \quad \varepsilon \ge 0 \end{aligned} \tag{11.10}$$

where $w_P = 30$ per bar^2, $w_{\Delta u} = 1$ per percent^2, $w_{fuel} = 1$, γ_i the unit's measured turbine output per kg/s through it times its flow per percent of speed and the fuel per shaft power, $SM_{min} = 10\,\%$ on

Table 11.3 The Nine Cases

#	Case and events	Teaching point
1	Manual, two units, reader slows unit 2	one head, two loads
2	demand 300 to 340 kg/s at 60 s	the station loop, line pack
3	unit 1 fouled 6 %, demand to 245 kg/s at 60 s	equal speed against equal margin
4	demand 360 to 420 kg/s over 1 to 11 min	start the third unit
5	Manual, unit 1 fouls to 8 % over 1 to 6 min	the fouled map
6	unit 1 fouled, demand to 180 kg/s at 60 s	a turndown no law serves
7	three units at 360 kg/s, unit 2 trips at 120 s	a trip
8	200 kg/s, PIC 200 kg/s per bar and 5 s, SP +0.5 bar at 60 s	station against anti-surge
9	Challenge: 6 h profile, unit 1 fouled, trip at 2, restart at 3 h	the whole station

the control line and $\rho = 10$ per percent^2. The framework's shared solver uses the dual method of Goldfarb and Idnani [18], only the first move is applied, and a solve that fails or is not finite hands the speeds back to the PI and the sharing law, which track the MPC's speeds without a bump. The margin floor sits on the control line and its price is soft for a measured reason. On the challenge with the suggested settings, a floor of 12 % priced at 1000 held the margins by letting the discharge rise 12.55 bar above 70 bar a, and a floor of 10 % priced at 1000 recycled 48.09 t, while the shipped price of 10 recycles 31.97 t and passes.

11.4.6 Three Models

The plant is Section 11.3. What a reader identifies from case 1 or case 5 is a static relation, how far one machine's margin moves for a point of speed on a common head. The controllers' models are the commissioned map inside the base speed and the MPC's linearization, which knows neither the fouling nor the line pack beyond the measured trend, and the measured margin of Equation (11.7), which is true because fouling leaves the surge line where it was.

11.5 Nine Cases on the Running Engine

The nine cases appear in Table 11.3. They share the defaults above, Auto with the PI and equal speed, two units on line at 300 kg/s, and a row lists only what differs, and Section 1.6.4 describes how a case is named and run. The figures quoted below were produced by stepping the engine in Node.js at seed 20261022 in the same increments the page uses. The technical review's measuring script, re-run on the app as shipped, reproduced its record line for line. Opening a case first, switching to it from another one and pressing Reset all reproduce the same run bit for bit.

11.5.1 One Head, Two Loads, and the Line Between

In case 1 a single move of unit 2's setpoint to 83 % at 20 s takes its margin from 28.05 to 12.38 % within 6 s, below 18 % 6.05 s after the move. Five minutes later the two units pass 153.5 and 127.4 kg/s, together 280.9 kg/s against the customers' 300, so the line unpacks, the discharge falls to 69.35 bar a and the lower common head lets unit 2 back out to 17.7 %. Brought down a point every 2, 5 or 10 s its margin bottoms at 15.6, 16.1 and 16.0 % instead, because the discharge is falling as the speed comes down. At 82 % the margin dips to 8.1 % and its anti-surge valve opens to 4.0 %.

🕮 **Pause.** Try Exercise 11.1 before reading on.

In case 2 the customers take 40 kg/s more at 60 s. The extra gas comes first out of the line pack, the header dips only 0.14 bar, and the station delivers 322 kg/s 10 min after the step, 330 after 14.2 min and 338 after 26.2 min, and at 1 h both machines run at 99.1 % speed, 34.7 % from surge. Raised to 71.0 bar a, the discharge is within 0.2 bar of the new setpoint 3.8 min after the change.

> **TRY IT IN THE APP 2 · The station pressure loop**
>
> Load the case and press Run. Touch nothing. At 60 s the customers take 40 kg/s more. Watch the Discharge pressure trend: the header dips only 0.14 bar, because the extra gas comes first out of the line. Watch the delivered flow beside the station flow demand: the station delivers 322 kg/s 10 min after the step and 338 kg/s after 26.2 min.

11.5.2 Who Recycles First

In case 3 unit 1 has lost 6 % of its head. At one speed it meets the common head further left on its map, nearer surge and with less gas through it. As the station turns down to 245 kg/s at equal speed, unit 1 reaches its control line at 15.6 min, recycles 15.01 t in the hour and ends 9.9 % from surge against unit 2's 18.8 %, both at 83.2 %. Switched to equal surge margin at 90 s, the trims speed unit 1 to 83.8 % and slow unit 2 to 81.3 %, both sit 12.5 to 12.6 % from surge, nothing is recycled and the hour's fuel falls from 63.33 to 62.76 MWh. Equal flow lands between them at 11.2 and 13.9 %. The reason is a floor. Two clean machines cannot pass less than 234.0 kg/s right of their control lines at 70 bar a. With unit 1 fouled that floor is 237.8 kg/s when both sit on their lines, as equal margin puts them, and 251.6 kg/s at equal speed, where unit 1 reaches its line first, so 245 kg/s lies between the two. Figure 11.3a sweeps the demand on the same plant and confirms the order. After two hours equal speed recycles 6.2 kg/s at 245, while equal flow and equal margin recycle nothing measurable there and only 1.86 and 0.09 kg/s at 240, and at 200 kg/s the three laws recycle 49.2, 40.3 and 36.9 kg/s.

Case 5 isolates the map. Both units hold 90.5 % by hand while unit 1 fouls to 8 %. At 6 min the headers ask 58.8 kJ/kg of both (59.57 and 59.64 kJ/kg at the machines), and unit 1 meets it at 22.9 % and 144.2 kg/s against unit 2's 28.5 % and 150.8 kg/s. At 94.3 % speed the margins match at 27.56 and 27.60 %, unit 1 then passing 155.5 kg/s and unit 2 149.2 kg/s. Equal margin asks more speed of the fouled machine and gives it more of the load.

11.5.3 A Start That Pushes Its Neighbors

In case 4 the demand climbs from 360 to 420 kg/s, more than two units can deliver. Units 1 and 2 reach 105 % at 7.9 min and the discharge passes below 69.5 bar a at 13.4 min (67.44 bar a at 30 min with no start). Started at 13.7 min, unit 3 purges for 20 s, accelerates on full recycle for 55.8 s behind its shut check valve, and 50.2 s into the load ramp its check valve opens. As it delivers, the station loop takes gas back from units 1 and 2, whose margins fall from 38.4 to 17.6 % before all three settle at 20.1 % and 90.3 % speed, while the header overshoots to 71.71 bar a. The load time hardly matters. Ramps of 10, 30, 120 and 300 s leave the neighbors' least margin at 17.20, 17.46, 17.57 and 17.63 %, because the station loop and not the ramp sets their speeds. Shut at once when loading begins, unit 3 is dead-headed and surges eight separate times, 18 prescribed cycles in all, before its head reaches the header's. The instant at 16 min is drawn in Figures 11.1 and 11.2.

Pause. Try Exercise 11.2 before reading on.

11.5.4 A Turndown No Law Can Serve

Case 6 is infeasible on purpose. With unit 1 fouled and the demand cut to 180 kg/s, far below the 237.8 kg/s two machines need, unit 1 recycles from 7.8 min while unit 2 never does, and by 1 h the station recycles 66.8 kg/s with unit 1's valve 28.8 % open, 199.4 t in the hour. Across the three laws and station tunings of 10, 20 and 80 kg/s per bar the recycle at 40 min is 54.8 to 69.9 kg/s, least with equal margin, and the MPC recycles 56.9 kg/s. Stopping unit 1 at 8.8 min leaves unit 2 half an hour later at 95.8 % speed and 36.9 % from surge with no recycle, and the station's fuel falls from 61.5 to 37.7 MW.

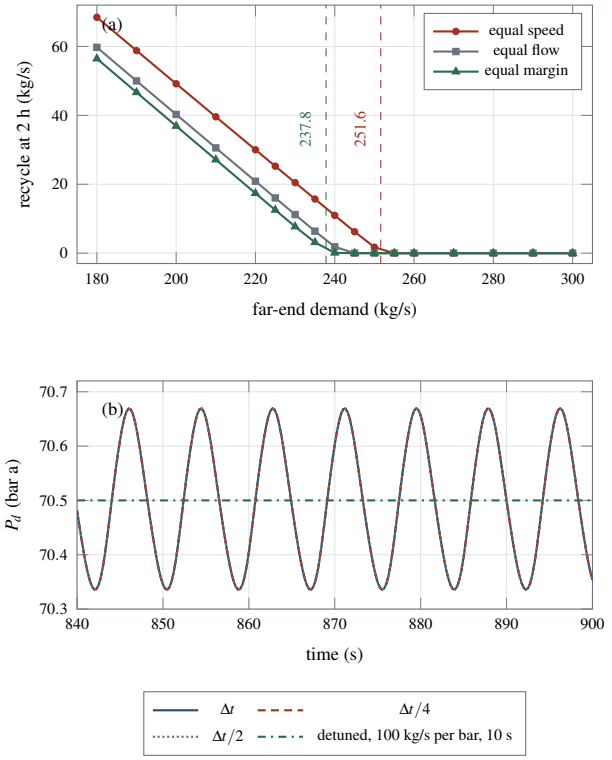

Figure 11.3 Two measurements from the app. (a) Case 3's plant, unit 1 fouled 6 %, with the demand stepped to each value at 61 s: the recycle averaged over the last 10 min of 2 h for each sharing law, equal speed with circles, equal flow with squares and equal margin with triangles. Dashed lines mark the computed floors with unit 1 fouled, both units on their control lines (237.8 kg/s) and equal speed (251.6 kg/s). (b) Case 8 with noise off, the true discharge pressure over 840 to 900 s at Δt (solid), $\Delta t/2$ (dotted) and $\Delta t/4$ (dashed), whose traces coincide within 3.4×10^{-6} bar, and with the station loop detuned (dash-dotted).

TRY IT IN THE APP **6 · A demand too small for two machines**

Load the case and press Run. Unit 1 is fouled, and at 60 s the demand falls to 180 kg/s. Watch the Recycle valves trend: unit 1's valve opens from 7.8 min and unit 2's never does. Press Stop a unit at 8.8 min. Half an hour later unit 2 runs alone at 95.8 % speed, 36.9 % from surge, with no recycle.

11.5.5 A Trip

In case 7 three units share 360 kg/s at 85.4 % speed, 11.4 % from surge, when unit 2's turbine trips at 120 s. Its shaft loses 3.0 % of its speed in the first second, its head falls below the header's, its check valve shuts, and 0.75 s after the trip, at 84.0 % speed, it surges before its PI has opened the valve far enough, entering surge twice for one full cycle. With *Open at once* it coasts down with its margin never below 10.8 % while above 30 % speed. Either way units 1 and 3 run to their 105 % limit, their margins rise to 38.1 %, and the discharge dips to 68.59 bar a 33 s after the trip and holds 69.75 bar a at 15 min. This is the emergency-shutdown problem of field stations [13], in which feedback on the margin is too slow for a shaft that loses head within a second.

Pause. Try Exercise 11.3 before reading on.

Table 11.4 Measured Cases at the Default Seed

Case	Configuration	Result
1	unit 2 to 83 % in one move	margin 28.05 to 12.38 % in 6 s
2	demand +40 kg/s	dip 0.14 bar, 322 kg/s after 10 min
3	equal speed / margin / flow	15.01 / 0 / 0.01 t, 63.33 / 62.76 / 62.76 MWh
3	MPC	0.84 t, 62.72 MWh, within 0.046 bar
4	load 10 to 300 s	neighbors' least margin 17.20–17.63 %
4	recycle shut at once	8 surges, 18 cycles
5	fouled 8 %, both 90.5 %	22.9 / 28.5 %, match at 94.3 %
6	3 laws × 3 tunings	54.8–69.9 kg/s recycled at 40 min
7	trip to the PI / open at once	surge after 0.75 s / none, $\geq$ 10.8 %
8	Δt, $\Delta t/2$, $\Delta t/4$	0.333869 / 0.333882 / 0.333885 bar
8	detuned 100, 10 s	4.4×10^{-5} bar

Table 11.5 The Six-Hour Challenge at the Default Seed

Settings	Surges	$\lvert P_d - 70\rvert$ (bar)	Recycled (t)	Fuel (MWh)	Verdict
shipped	1	24.97	6.90	380.54	fails
suggested	0	1.048	30.74	499.28	passes
at equal speed	0	1.237	97.75	501.31	fails
at equal flow	0	1.056	38.68	498.86	fails
sequencer by hand	0	24.95	0.00	380.02	fails
trip to the PI	2	1.049	30.74	499.28	fails
station 20, 60 s	0	1.714	37.52	499.65	fails
station 120, 15 s	0	1.020	30.48	499.21	passes
on MPC	0	1.281	31.97	499.78	passes

11.5.6 Two Loops on One Flow

In case 8 two clean units at 200 kg/s sit on their control lines at 79.4 % speed with their valves 7.1 % open, recycling 33.2 kg/s, and the station loop is tuned at 200 kg/s per bar and 5 s. On recycle a unit's forward flow answers both its speed and its valve. When the station loop raises a speed the margin rises, the override closes the valve, and more gas goes forward than the speed alone would send, so the anti-surge loop adds gain and lag to the station loop. With the noise off the discharge then swings 0.333 869 bar peak to peak over the last 300 s of 900 s, with an 8.353 s period, the valves swinging 12.2 % and the speeds 1.73 %. The swing is 0.333 882 bar at $\Delta t/2$ and 0.333 885 bar at $\Delta t/4$, and it is 0.3309 over 300 to 600 s and 0.3351 bar over 1200 to 1500 s, a limit cycle of the tuning and not of the integration, as shown in Figure 11.3b. Detuned to 100 kg/s per bar and 10 s it falls to 4.4×10^{-5} bar. The same fast tuning at 300 kg/s, with no recycle, rings 0.74 bar between 60 and 180 s and then settles to 0.0003 bar. At 180 and 230 kg/s, recycling 39.2 and 35.6 kg/s, it holds 0.303 and 0.286 bar, and at 250 kg/s, where nothing recycles, 2×10^{-5} bar.

11.5.7 The Six-Hour Challenge

The challenge's demand runs at 300 kg/s, climbs to 400 for the peak from 1.4 to 1.9 h, falls to 340 for the trip of unit 2 at 2 h, rises again and runs down to 250 for the night, with unit 1 fouled 6 % and unit 2 started again at 3 h. Its target, scored from the simulator's truth at every physics step over the whole run, is no surge on any unit, the discharge within ± 1.5 bar of 70.0 bar a from 10 min on, at most 35 t recycled by units on line, and the surge control line kept at 10 % or above, on the plant as shipped. It ships with equal speed, starts and stops by hand, a tripped unit's valve left to its PI and the station at 20 kg/s per bar and 60 s. The outcomes are in Table 11.5.

As shipped nobody starts unit 3 for the peak, unit 2 surges on its trip, and with one unit left the discharge falls 24.97 bar by 3.01 h. In that run the discharge header falls below the suction header from 2.87 h, by up to 1.10 bar. With no bypass in the model, gas then passes straight through the machines and their check valves with no head made. Unit 1 runs out to $\phi/n = 0.8027$, past the drawn

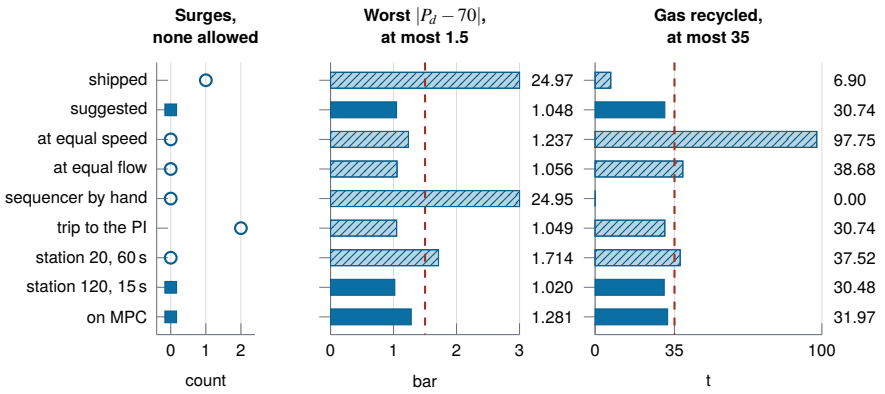

Figure 11.4 The six-hour challenge of Table 11.5 against its target. Solid bars and filled squares are the three settings that pass, hatched bars and open circles the six that fail, and dashed lines the limits. The rows below the suggested one each change one setting of it. Two pressure bars are cut at the axis end, and the column beside each panel gives the table's value. Notice that each failing row misses a different limit, and that the trip left to the PI misses only by its two surges. The fuel, not graded, is in the table.

choke line at 0.70 and next to the cubic's zero-head end at 0.8037, off the map from 2.28 to 3.18 h, and unit 2, restarting at 3.0 h, reads a measured margin of 821 % at one scan and 517.5 % logged, from a flow that is not on its curve, while its margin on the curve never exceeds 60.7 %. This windmill regime is documented rather than modeled, because only the failing shipped run reaches it.

The suggested settings change four layers, equal surge margin, the automatic sequencer, *Open at once* and the station at 80 kg/s per bar and 20 s, and pass with no surge, 1.048 bar at the trip and 30.74 t recycled. Dropping any one fails. Equal speed recycles 97.75 t through the night with unit 1 fouled, the trip left to the PI surges unit 2 twice, the slow station loop lets the discharge stray 1.714 bar at 3.06 h, and without the sequencer the peak and the trip are never covered. The predictive controller passes too, at 1.281 bar, 31.97 t and 499.78 MWh. In case 3 it recycled 0.84 t on 62.72 MWh against equal margin's 0 and 62.76 MWh, so it matches equal-margin sharing within 0.07 % of fuel and does not beat it. Figure 11.4 sets the nine settings of Table 11.5 against the three limits of the target.

The target grades the gas recycled rather than the fuel, because across the three laws the six hours' fuel differs by less than half a percent (498.86, 499.28 and 501.31 MWh) while the gas recycled differs threefold (38.68, 30.74 and 97.75 t). The 35 t budget was first checked at eleven seeds. The suggested settings recycle 28.34 to 32.80 t (mean 30.17) and pass throughout that initial set, equal flow recycles 37.11 to 41.51 t (mean 39.24) and fails at every seed, so the budget sits 2.2 t above the worst passing run and 2.1 t below the best failing one, and halving the step moves the recycle by less than 0.001 t. A wider check of the default seed and seeds 0–29 gives 29 of 31 passing runs at each step. Seeds 12 and 21 recycle 35.21 t and 37.17 t, respectively, while retaining zero surges and worst pressure deviations below 1.06 bar. Both misses persist at half the step. The default certificate is reproducible, but the suggestion does not satisfy the recycle budget for every tested noise realization. The band is 1.5 bar because over eleven station tunings the trip's dip never fell under 1.001 bar, and from 160 kg/s per bar on the units surge through case 8's interaction (3 to 6 surges). The verdict was tightened in the technical review. Judged against the setpoint in force with nothing else checked, equal speed with a 66 bar a setpoint, equal flow with an 8 % control line and equal speed with unit 1's fouling removed had each passed (the first 5.07 bar off 70 bar a) without the lesson. The band is now judged against 70.0 bar a, a control line under 10 % is a miss, and a changed disturbance or model parameter makes the run read "not graded", so all three now fail or go ungraded while the shipped and suggested verdicts and every run's log are unchanged. At the top speed of 600 s per second,

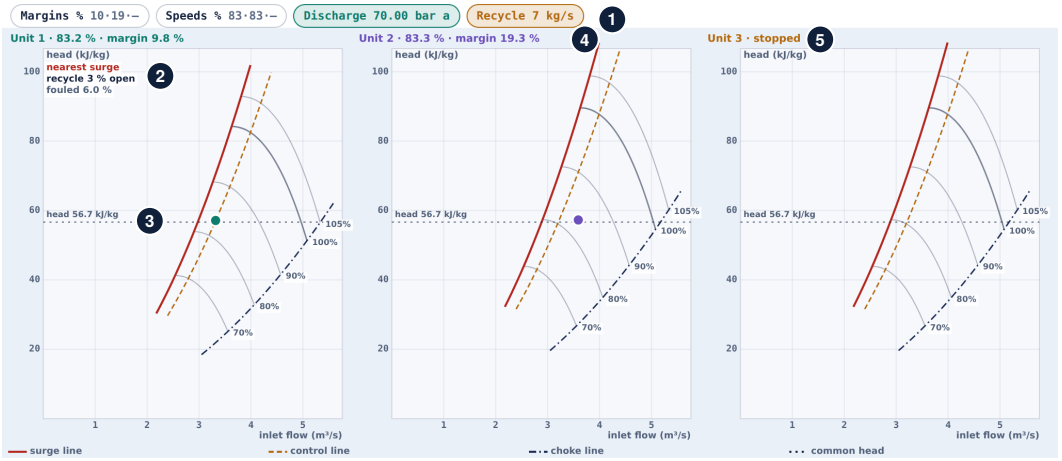

Figure 11.5 Case 3 at equal speed, $t = 40$ min, on the Maps view, light theme, cropped to the maps. Both units run at 83.2 % and 83.3 % on a common head of 56.7 kJ/kg, the dotted line at callout 3, and fouled unit 1 (left) sits on its control line with measured margin 9.8 %, flagged nearest surge with its recycle valve 2.9 % open (callout 2; the flag rounds to 3 %), while clean unit 2 stands 19.3 % from surge (callout 4) and unit 3 is stopped (callout 5). The badges at callout 1 read the discharge at 70.00 bar a and 7 kg/s recycled.

sustained in Chrome at 1366×768 (599.8 s/s) and 390×844 (599.6 s/s), the six hours play in 36 s. The lesson of case 3 at 40 min is shown in Figure 11.5.

> **DID YOU KNOW? Two Surges That the Totals Do Not Show**
>
> From the simulator (Table 11.5). Over the six hours the suggested settings, and the same settings with the tripped unit's valve left to its loop, both recycle 30.74 t and burn 499.28 MWh, and their worst pressure deviations are 1.048 and 1.049 bar. Only the surge count tells them apart, 0 against 2.

11.6 Technology in Practice

Figure 11.6 draws one unit's loops and the station loop of Section 11.4 with the hardware of a transmission station in place of the schematic's symbols.

What Measures It

The station loop reads one pressure transmitter, on the discharge header, and the base speed adds the one on the suction header. Each anti-surge loop reads three instruments, the suction pressure, its machine's flow element and its speed pickup, and computes the margin of Equation (11.7) from them, so one scan's margin carries the noise of all three, 0.54 point in the model.

> **TECHNOLOGY Sensor: The Flow Element Behind Every Margin**
>
> A pressure-based flow element such as an orifice plate or a venturi accelerates the gas, and the flow follows the square root of the pressure drop across it divided by the density [19]. The square root makes the reading less certain as the flow falls, the direction of surge. The practical turndown of such a meter is often no more than 3 to 1, perhaps 10 to 1 with a modern transmitter, and a worn element or uneven liquid in the impulse tubes then gives large flow errors at the low end of the range [19]. The model's element carries 0.4 % of reading as noise behind a 0.3 s lag.

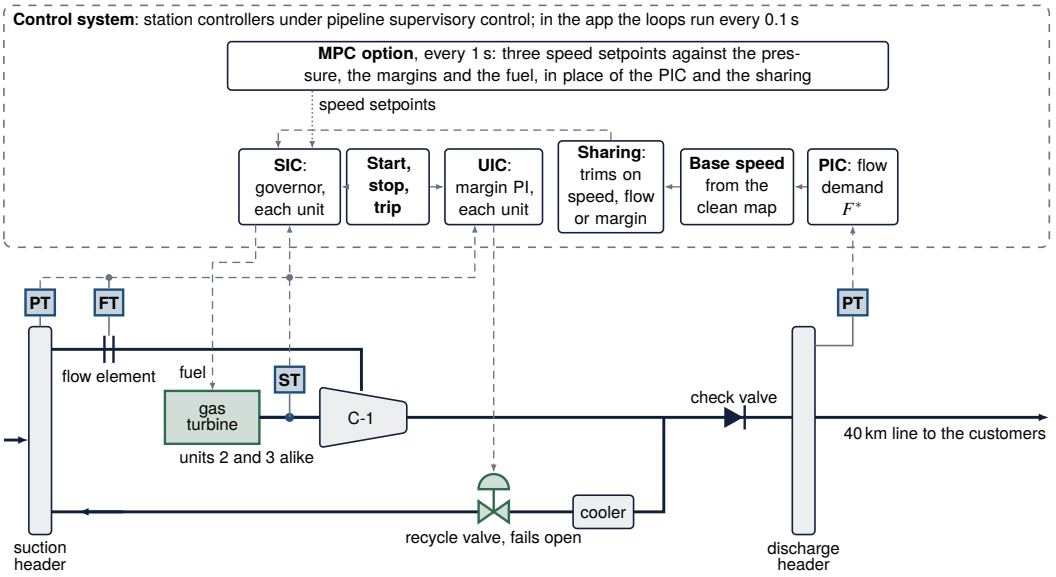

Figure 11.6 The loops of one unit and of the station with their hardware. Lettered boxes are sensing devices with the chapter's tags: PT on each header, FT on the machine's inlet and ST on its shaft. Green symbols are final elements: the gas turbine's fuel, set by the governor (SIC), and the anti-surge recycle valve, which fails open. Dashed arrows are signals, dotted arrows the speed setpoints that the MPC option writes in place of the sharing law. Notice that the anti-surge loop reads three instruments and no controller reads the fouling, and that the trip reaches the recycle valve without passing through the margin. Units 2 and 3 repeat unit 1. Not to scale, and not a piping diagram.

What Moves It

Two final elements per unit carry the chapter's handles. The gas turbine's fuel moves the shaft speed, through a governor whose setpoint comes from the sharing law, and the recycle valve moves the flow through the machine without moving the flow delivered. Speed is the efficient handle and recycle the protective one, because gas sent around the recycle has been compressed for nothing. A start, a stop and a trip also drive the valve open.

> **TECHNOLOGY** **Final Element: The Anti-Surge Recycle Valve**
>
> Surge is a violent and potentially self-destructive action of a centrifugal compressor when the pressure across it is too high and the flow through it too low. A recycle valve from discharge to suction prevents it, keeping enough flow through the machine while unloading the pressure across it [19]. A textbook loop diagram for surge control makes the valve air-to-close, so that it opens when it loses its air, and the transmitter of the pressure across the machine reverse-acting, so that a broken wire reads as a high pressure and the controller opens the valve. The safest failure decides: a process safer with the valve open needs an air-to-close valve [19]. The model's valve strokes at 50 %/s behind a 1 s lag, and case 7 shows when that is too slow.

Where the Controller Runs

A gas pipeline is run through supervisory control and data acquisition. A control server at a control center gathers the field data, while programmable logic controllers or remote terminal units at the field sites control each local process, and control loops run with cycle times from milliseconds to minutes [20]. Compressor surge control is a standard example of a safety function that must respond within a guaranteed time [19]. The app's controllers execute every 0.1 s and its MPC every 1 s. The

chapter's structure maps onto standard blocks: a pressure PI whose output is a flow demand, a characterizer that turns it into a base speed, sharing trims that re-center to zero, a PI governor per unit, an anti-surge PI per unit that rests at its clamp above the control line, and a sequence for the start, the stop and the trip. A predictive controller for one centrifugal machine has been timed on an industrial programmable logic controller [14].

> ⚠ **IN PRACTICE** **What the Federal Rules Ask of a Station**
>
> United States pipeline safety rules set a station's last layer of protection. Every station, except an unattended one of 1,000 horsepower or less, must have an emergency shutdown system that blocks gas out, blows down the piping and shuts down the compressors, operable from two places or more outside the gas area [21]. Every prime mover but an electric motor needs an automatic shutdown before it or its compressor passes a maximum safe speed, and every unit a shutdown or alarm for inadequate cooling or lubrication [22]. Each remote shutdown device is tested at least once each calendar year and within 15 months of the last test [23]. The chapter's overspeed trip at 110 % is a teaching value.

What Is Modernizing It

Three developments touch this chapter. The first is the drive: most compressor units burn gas taken from the line, but electric drives have been growing, mainly for environmental reasons [16], and converting rotating machines to variable speed is an ongoing trend [24]. The second is the sharing layer. A 2022 study ran online feedback optimization of a station's load sharing, with Gaussian process regression learning the mismatch between the assumed and the real maps: imperfect maps had cost 5 % more power than perfect ones, and the learning cut that to 0.8 %, by a method meant for programmable controllers and edge devices [24]. The third is security. In May 2021 more than 5,500 miles of refined products pipeline shut down after ransomware encrypted the operator's business systems, with no evidence then that it had reached the control systems [20]. The pipeline of Chapter 13 carries such products.

11.7 Limitations

Every size, rating and threshold is a teaching value picked so that the design point is consistent, two 15 MW units carrying 300 kg/s at 70 bar a, and none was fitted to a station. Each machine is quasi-steady, surge is a prescribed 2 s cycle whose onset is the static surge line and not Greitzer's dynamic boundary, rotating stall is absent, and the three machines share one map. Fouling is one head factor that leaves the surge line fixed, a simplification that keeps the controllers' margins true where real deposits may not. The headers and line are isothermal, the recycle coolers perfect and the line two cells without pressure waves. There is no station bypass, so the shipped challenge's windmill regime is outside the model's intended range. The margin is the same-speed distance of Section 1.4.6, which reads lower than the same-head distance many industrial controllers use.

The controllers carry the commissioned map, the declared defaults of every model parameter, so the base speed and the MPC meet fouling and any change of the plant's parameters only through the measurements, while the anti-surge loops' surge line, a protection limit, follows the plant's map. The MPC is linearized every second on the clean map with a drift for the line pack and carries no guarantee beyond its horizon. There is no decoupling between the station loop and the anti-surge loops, so case 8's remedy is detuning. A static feedforward between two interacting loops, the remedy of Volume 6, Chapter 2, is not modeled. The challenge grades recycle in place of fuel because the fuel differs by less than half a percent. Every number comes from one model at seed 20261022. The challenge was also run at ten other seeds initially, then over 31 seeds at both steps, with the two recycle-budget misses reported above. Case 8 uses two further refinements, $\Delta t/2$ and $\Delta t/4$.

OPEN CHALLENGE Sharing the Load With Maps That Drift

Every sharing law of this chapter trusts a map, and a real map drifts. Predicting how wear and tear change an engine's performance is still a matter of discussion, and a gas turbine and the compressor it drives degrade as one system [15]. Rotating machines degrade over time, accurate models of them are generally not available, and the optimizers that track each machine's map fit a polynomial of fixed degree online, which may give low model fidelity and so a poorer share [24]. Progress would be a sharing layer that learns each machine's map as it fouls, runs on the station's own controller and still keeps every unit right of its control line.

11.8 Conclusion

A compressor station is a set of machines that share a head and not a load, and this simulator logs each machine's speed, flow, margin and recycle beside the station's pressure, fuel and line pack. Slowing one of two machines to 83 % moves its margin from 28.05 to 12.38 % in 6 s while its neighbor takes the gas. With one machine fouled, equal speed recycles 15.01 t in an hour of turndown where equal surge margin recycles none, because the two laws reach the control lines at 251.6 and 237.8 kg/s. A start pushes the running units to about 17.6 % whatever the load time, a trip handled by feedback alone surges the tripped unit in 0.75 s, and a fast station loop on recycling units holds a 0.33 bar limit cycle that halving the step does not move. The predictive controller matches equal-margin sharing, and the six-hour challenge needs all four suggested layers. The model is lumped and its surge cycle imposed. A fuller model would add a station bypass, distinct maps with a surge line that fouling moves, a decoupler between the station and anti-surge loops, and a Greitzer duct for each machine at a finer step. The simulator shows that a station's surge count and its recycle are decided by the layers above its loops.

KEY POINTS

- Compressors on one header share a head, not a load, so a machine a little slower or fouled drifts toward surge while its neighbors take the gas.
- With one machine fouled, equal speed recycles 15.01 t in an hour where equal surge margin recycles none.
- A trip handled by feedback alone surges the tripped unit in 0.75 s, and a valve opened by the trip signal prevents it.
- A fast station loop on recycling units holds a 0.33 bar limit cycle that halving the step does not move.
- The predictive controller matches equal-margin sharing and does not beat it.

Where else this applies. Three producers on one shared header return in Volume 18, Chapter 9, one total divided among unlike actuators in Volume 4, Chapter 12, and parallel pumps started and stopped under minimum run and rest times in Volume 9, Chapter 11.

References

[1] E. M. Greitzer, "The stability of pumping systems—the 1980 Freeman Scholar Lecture," *Journal of Fluids Engineering*, vol. 103, no. 2, pp. 193–242, 1981.

[2] F. Paparella, L. Domínguez, A. Cortinovis, M. Mercangöz, D. Pareschi, and S. Bittanti, "Load sharing optimization of parallel compressors," in *2013 European Control Conference (ECC)*. IEEE, 2013, pp. 4059–4064.

[3] A. Cortinovis, M. Mercangöz, M. Zovadelli, D. Pareschi, A. De Marco, and S. Bittanti, "Online performance tracking and load sharing optimization for parallel operation of gas compressors," *Computers & Chemical Engineering*, vol. 88, pp. 145–156, 2016.

[4] P. Milosavljevic, A. G. Marchetti, A. Cortinovis, T. Faulwasser, M. Mercangöz, and D. Bonvin, "Real-time optimization of load sharing for gas compressors in the presence of uncertainty," *Applied Energy*, vol. 272, p. 114883, 2020.

[5] E. M. Greitzer, "Surge and rotating stall in axial flow compressors—part I: Theoretical compression system model," *Journal of Engineering for Power*, vol. 98, no. 2, pp. 190–198, 1976.

[6] J. T. Gravdahl and O. Egeland, "Centrifugal compressor surge and speed control," *IEEE Transactions on Control Systems Technology*, vol. 7, no. 5, pp. 567–579, 1999.

[7] K. K. Botros, P. J. Campbell, and D. B. Mah, "Dynamic simulation of compressor station operation including centrifugal compressor and gas turbine," *Journal of Engineering for Gas Turbines and Power*, vol. 113, no. 2, pp. 300–311, 1991.

[8] K. K. Botros, "Transient phenomena in compressor stations during surge," *Journal of Engineering for Gas Turbines and Power*, vol. 116, no. 1, pp. 133–142, 1994.

[9] P. J. Wong and R. E. Larson, "Optimization of natural-gas pipeline systems via dynamic programming," *IEEE Transactions on Automatic Control*, vol. 13, no. 5, pp. 475–481, 1968.

[10] K. K. Botros and J. F. Henderson, "Developments in centrifugal compressor surge control—a technology assessment," *Journal of Turbomachinery*, vol. 116, no. 2, pp. 240–249, 1994.

[11] N. Staroselsky and L. Ladin, "More effective control for centrifugal gas compressors operating in parallel," in *ASME 1986 International Gas Turbine Conference and Exhibit, Volume 5*. Düsseldorf, West Germany: American Society of Mechanical Engineers, 1986, paper V005T14A001.

[12] S. Ohanian and R. Kurz, "Series of parallel arrangement in a two-unit compressor station," *Journal of Engineering for Gas Turbines and Power*, vol. 124, no. 4, pp. 936–941, 2002.

[13] R. Kurz and R. C. White, "Surge avoidance in gas compression systems," *Journal of Turbomachinery*, vol. 126, no. 4, pp. 501–506, 2004.

[14] G. Torrisi, S. Grammatico, A. Cortinovis, M. Mercangöz, M. Morari, and R. S. Smith, "Model predictive approaches for active surge control in centrifugal compressors," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 6, pp. 1947–1960, 2017.

[15] R. Kurz and K. Brun, "Degradation in gas turbine systems," *Journal of Engineering for Gas Turbines and Power*, vol. 123, no. 1, pp. 70–77, 2001.

[16] U.S. Energy Information Administration, "Natural gas pipelines: Transportation process and flow," About U.S. Natural Gas Pipelines, archived analysis, based on data through 2007/2008 with selected updates, 2008, accessed November 18, 2026. [Online]. Available: https://www.eia.gov/naturalgas/archive/analysis_publications/ngpipeline/process.html

[17] F. K. Moore and E. M. Greitzer, "A theory of post-stall transients in axial compression systems: Part I—development of equations," *Journal of Engineering for Gas Turbines and Power*, vol. 108, no. 1, pp. 68–76, 1986.

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

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

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

[21] Pipeline and Hazardous Materials Safety Administration, "49 CFR 192.167, compressor stations: Emergency shutdown," Code of Federal Regulations, Title 49, Part 192, edition of October 1, 2025, govinfo.gov, 2025.

[22] ——, "49 CFR 192.171, compressor stations: Additional safety equipment," Code of Federal Regulations, Title 49, Part 192, edition of October 1, 2025, govinfo.gov, 2025.

[23] ——, "49 CFR 192.731, compressor stations: Inspection and testing of relief devices," Code of Federal Regulations, Title 49, Part 192, edition of October 1, 2025, govinfo.gov, 2025.

[24] M. Zagorowska, M. Degner, L. Ortmann, A. Ahmed, S. Bolognani, E. A. del Rio Chanona, and M. Mercangöz, "Online feedback optimization of compressor stations with model adaptation using Gaussian process regression," *Journal of Process Control*, vol. 121, pp. 119–133, 2023.

11.9 Exercises

11.1 Concept check · 2 min

Case 1 starts with both units on Manual at 90.5 % speed, each passing 150 kg/s at 28.0 % from surge. The reader lowers unit 2's speed setpoint to 83 % and leaves unit 1 alone. What has happened five minutes later?

(a) Unit 1 still passes 150 kg/s, because its speed setpoint and its map have not changed

(b) Unit 1 passes more gas, because the discharge has fallen and unit 1 meets the lower head at a larger flow

(c) Both units pass about 140 kg/s, because the anti-surge loops split the station's flow between them

(d) Unit 2 is in deep surge, because its 83 % speed line no longer reaches the header's head

11.2 Concept check · 1 min

In case 4 the reader starts unit 3 while units 1 and 2 run at their 105 % limit. Load ramps of 10, 30, 120 and 300 s leave the two running units' least margins at 17.20, 17.46, 17.57 and 17.63 %. Why does the length of the ramp hardly matter?

(a) The running units' anti-surge loops open as unit 3 starts and pin their margins at one value

(b) Unit 3's check valve stays shut through the ramp, so units 1 and 2 never see its gas

(c) The station loop sets the running units' speeds from the discharge pressure and takes back what unit 3 delivers

(d) The sequencer's minimum run time overrides the ramp and loads unit 3 on a schedule of its own

11.3 Concept check · 2 min

In case 7 three units share 360 kg/s at 85.4 % speed, 11.4 % from surge. Unit 2's turbine trips at 120 s with *Recycle on a trip* on *Anti-surge loop*, and unit 2 surges 0.75 s later. Which single change, made before Run, keeps unit 2 out of surge?

(a) Raise the anti-surge gain to its maximum of 20, so the PI opens the valve faster

(b) Raise the surge control line to 15 %, so the valves start the trip partly open

(c) Choose equal surge margin sharing, so units 1 and 3 take unit 2's gas more quickly

(d) Set *Recycle on a trip* to *Open at once*, so the trip itself opens the valve

11.4 Theory and calculation · 6 min

At the design point a clean unit runs at $n = 0.9047$ with an inlet flow of $Q = 4.202\,\mathrm{m^3/s}$, and the suction header is at 45.00 bar a. Use Equations (11.1) and (11.2) with $\psi_0 = 0.2$, $H = 0.18$, $W = 0.25$, $Q_{ref} = 7.255\,\mathrm{m^3/s}$ and $H_{ref} = 160\,\mathrm{kJ/kg}$, and $ZRT/M = 126.07\,\mathrm{kJ/kg}$.

(a) Find ϕ, ϕ/n and the surge margin, once from the definition and once as $50(x-1)$. *Checkpoint:* $\phi = 0.5792$, $\phi/n = 0.6402$, $SM = 28.0\,\%$.

(b) Find $\Psi(\phi/n)$ and the polytropic head the machine makes. *Checkpoint:* $\Psi = 0.4592$, $h_p = 60.1\,\mathrm{kJ/kg}$.

(c) Find the suction density, the mass flow and the surge flow in kg/s at this speed, and check the margin from the two mass flows. *Checkpoint:* 35.69 kg/m^3, 150.0 kg/s and 117.1 kg/s.

11.5 Theory and calculation · 6 min

Fouling multiplies a machine's head by $1-a$ at every flow coefficient and leaves its surge flow where it was.

(a) Two machines make the same head. Show from Equations (11.1) and (11.2) that they have equal margins when $n_1\sqrt{1-a_1} = n_2\sqrt{1-a_2}$. State what the result neglects. *Checkpoint:* Equal margin means equal ϕ/n, so $(1-a)n^2$ must match.

(b) In case 5 unit 2 is clean at 90.47 % and unit 1 has lost 8 % of its head. Find the speed that gives unit 1 unit 2's margin. *Checkpoint:* 94.32 %.

(c) In case 3 with equal surge margin, unit 1 is fouled 6 % and clean unit 2 runs at 81.25 %. Predict unit 1's speed. *Checkpoint:* 83.80 %.

11.6 Theory and calculation · 7 min

The gas is $M = 17.4\,\text{kg/kmol}$ with $Z = 0.90$ at 293.15 K. The line is two cells of 20 km with $D = 0.9\,\text{m}$ and $f = 0.0095$, the headers hold 300 m^3 each, and the outlet flow w_o crosses half a cell. Use Equation (11.6).

(a) Find c, and the gas held per bar by one header and by one cell. *Checkpoint:* 355.07 m/s, 237.96 kg and 10 092 kg per bar.

(b) Find K for one cell in bar^2 per $(\text{kg/s})^2$, and the steady pressures P_1 and P_2 for 300 kg/s with the discharge at 70.00 bar a. *Checkpoint:* $K = 6.576 \times 10^{-3}$, $P_1 = 67.85\,\text{bar}\,a$, $P_2 = 63.34\,\text{bar}\,a$.

(c) Find the gas held in the two cells. *Checkpoint:* 1324.05 t.

11.7 App challenge · 9 min

Load scenario 3, "Equal speed against equal margin, one machine fouled", at the default seed. Play at speed 60, so an hour takes a minute, and pause just after 60 min. Read the row at 3600 s of the exported file.

(a) Run the case as shipped, at equal speed. Read the gas recycled (column `recycled_online_t`) and the fuel (`fuel_MWh`). *Checkpoint:* 15.0 t and 63.33 MWh, within 0.1.

(b) Press the reload button, set *Sharing law* to *Equal surge margin* before Run, and repeat. Read the same two columns, both units' speeds in their titles on the *Maps* view, and the columns `truth_u1_margin_pct` and `truth_u2_margin_pct`. *Checkpoint:* 0.0 t, 62.76 MWh, 83.8 % and 81.3 %, margins 12.5 and 12.6 %, within 0.1.

(c) Explain why the fouled unit runs faster under equal margin, and why the fuel falls by less than one percent while the recycle falls to zero. *Checkpoint:* The fan-law result of the second theory exercise, and gas compressed for nothing costs only its compression work.

11.8 App challenge · 8 min

Load scenario 7, "A trip: the others take the load", at the default seed, and play at speed 10.

(a) Run to 300 s as shipped, with *Recycle on a trip* on *Anti-surge loop*. Read the *Surges* tile. *Checkpoint:* 2.

(b) Set *Recycle on a trip* to *Open at once*, press Reset and run to 900 s. Read the *Surges* tile. In the exported file, find the lowest value of the column `truth_pd_bara` and its time. At 900 s, read that column and two more, `truth_u1_margin_pct` and `truth_u3_margin_pct`. *Checkpoint:* 0 surges, 68.59 bar a at 153 s, then 69.75 bar a and 38.1 % for both units.

(c) The discharge dip is the same 68.59 bar a with either trip action. Explain why, and why the discharge settles short of 70.0 bar a. *Checkpoint:* The dip comes from the loss of unit 2's forward gas and the limit of the other two.

11.9 App challenge · 8 min

Load scenario 8, "Station loop against anti-surge loops", at the default seed. Set *Instrument noise* to *Off* in the Advanced disturbances before Run, play at speed 10 and run to 900 s.

(a) In the exported file, find the peak-to-peak swing over 600 to 900 s of two columns, `truth_pd_bara` and `u1_recycle_pct`, and the period of the pressure swing. Read the *Lowest margin* tile. *Checkpoint:* 0.334 bar within 0.005, 12.2 % within 0.3, a period of 8.4 s within 0.2, and 5.9 %.

(b) Press Reset, set *Station gain* to 100 kg/s per bar and *Station integral time* to 10 s before Run, and repeat. *Checkpoint:* Below 1×10^{-4} bar, and 9.9 %.

(c) Explain why the same fast tuning rings and then settles when the units carry 300 kg/s with no recycle. *Checkpoint:* Without recycle the anti-surge loops are idle and add no gain.

11.10 Application area · 4 min

A station manager asks why the six-hour challenge grades the gas recycled and not the fuel, since fuel is what the station pays for. Use Table 11.5.

(a) For the suggested settings at equal speed, equal flow and equal margin, find the spread of the six hours' fuel as a percent, and the ratio of the largest recycle to the smallest. *Checkpoint:* 2.45 MWh, just under half a percent, and 3.2 times.

(b) Give two reasons, from the table and the chapter, why fuel is a poor grade for this lesson. *Checkpoint:* It barely separates the laws, and the failing shipped run burns the least.

11.11 Application area · 5 min

The chapter measures the surge margin at the present speed. Many industrial anti-surge controllers measure it at the same head, on a map in reduced coordinates. On this map the surge line is $\psi = 0.56\, n_s^2$ at $\phi = 0.5\, n_s$.

(a) At the design point, $\phi = 0.5792$ at $n = 0.9047$ with $\Psi = 0.4592$, find the speed n_s whose surge point has the same head, and the same-head margin. *Checkpoint:* $n_s = 0.8193$, 41.4 %.

(b) On the 10 % control line, $\phi/n = 0.55$ and $\Psi = 0.5485$. Find the same-head margin there. *Checkpoint:* 11.2 %.

(c) A vendor states a control line of 10 % in the same-head convention. What does this mean for comparing the vendor's numbers with this chapter's, and where do a real station's numbers come from? *Checkpoint:* About a point apart at the line, thirteen at design, set by surge tests.

11.12 Review the agent · 5 min

An AI assistant reviews the six-hour challenge. It writes: "The suggested settings pass with a worst pressure deviation of 1.048 bar, set by the trip at 2 h, and that is the binding target. Keep equal surge margin, the automatic sequencer and *Open at once*, and stiffen the station loop to 200 kg/s per bar with a 20 s integral time. A faster loop recovers the trip's dip sooner, so it can only improve the result." Find the flaw, and say how to check it on the app.

CHAPTER 12

A Gas Transmission Line Simulator for Line Pack, Forecast Feedforward and Leak Detection

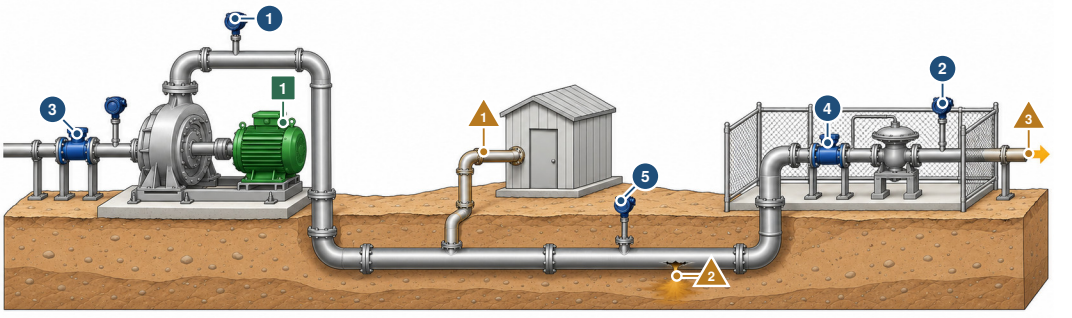

The controller moves 1 compressor speed, through the motor's drive

The controller reads 1 pressure where the gas leaves the station 2 pressure at the city gate
3 gas metered in, for the leak detector 4 gas metered out, for the leak detector
5 pressure at the middle of the line, for the line pack

What disturbs it 1 gas taken by a side branch 2 a leak 3 what the city takes

Notice that the controller moves the compressor on the left but must hold the pressure at the city gate on the right, with a long buried pipe of stored gas between them that only the midpoint transmitter looks into. Illustration, not to scale.

IN PLAIN TERMS **A Pipe That Is Also a Tank**

The main idea. A long gas pipeline also stores gas, and the higher its pressure, the more it holds. A machine called a compressor pushes gas in at one end, and a city takes it out at the other.

The challenge. A push at the start takes most of an hour to reach the city, so waiting for the city's pressure to fall is too late. No meter weighs the stored gas. So a leak is found by comparing gas in with gas out, and a meter that reads a little low looks like a leak.

Why it matters. The city must get its gas above a promised pressure, without wasted power or false alarms.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History box in Section 12.2, the Did you know? boxes in Sections 12.3 and 12.5, the Try it boxes in Section 12.5, and Key points in Section 12.8. No equations.
Study	45 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 12.5 with Figure 12.4, Section 12.6 and Section 12.7 with its Open challenge box.

The control question A transmission line stores its own gas and takes most of an hour to pass a station change to the city gate, which must stay above a contract floor. What must the station's control see, and how small a leak can a mass balance on three pressure transmitters find?

Learning objectives

- Compute the steady pressure profile, the line pack and the three time scales of a gas transmission line from its friction law.
- Explain why a station loop on its discharge leaves the delivery short, and why a cascade cannot avoid the dip.
- Use a demand forecast as feedforward and trade the pressure margin it saves against compression cost and the lead time.
- Predict the detection time and the smallest detectable leak of a windowed mass balance, and explain its false alarms.
- Judge what a predictive controller that carries the floor as a constraint buys over a cascade with a higher setpoint.

Prerequisites Sections 1.4.5 (line pack), 1.4.9 (inferred inventories and leak detection), 1.6.2 (the framework's PI) and 1.6.4 (the apps). Section 1.6.3 is needed only for the predictive comparison. Chapter 11 describes the station at the head of such a line.

The app https://apmonitor.com/apps/gas-pipeline/

See also Volume 3, Chapter 2 for loop structure matched to where a disturbance enters, Volume 3, Chapter 6 for feedforward hours ahead of the feedback, Volume 6, Chapter 11 for an inventory in transit ordered from a forecast, and Volume 5, Chapter 11 for a quantity inferred from the measurements around it.

12.1 Introduction

Gas pipelines are usually drawn as a pipe with a pressure drop, which is correct only at steady state. Between steady states a transmission line is a long, pressurized storage vessel. Every bar of pressure along its length holds gas, the inventory rises and falls through the day, and that stored gas carries the city through the morning ramp while the compressor catches up [1]. The transient behavior of such lines has been modeled for decades, from implicit and explicit schemes for networks to reviews of compressible flow in long pipelines [2–4].

The control problem follows from the storage. The delivery pressure at the city gate is the quantity the contract names, but the compressor acts at the other end of a line whose response is a slow diffusion of pressure rather than a transport delay or a first-order lag. A station that holds its own discharge pressure keeps its own number and loses the customer's. A delivery loop above it has to act through tens of minutes of line response, and demand for gas is largely predictable, which makes the forecast worth more than any retuning [5]. Scheduling of compressors against forecast demand, electricity price and pressure limits is the central optimization problem of gas transmission and has a long literature of its own [6–8].

The misconception the simulator targets is that a pipeline is a pipe. The broader class of problem is a distributed inventory controlled from its boundaries and measured at a few points: water mains, oil lines, district heating networks and irrigation canals all share it. Volume 9, Chapter 5 releases water ahead of demand along a canal, Volume 9, Chapter 6 sends heat ahead of a morning peak, and Chapter 13 runs a liquid line with a leak detector of its own. Three consequences carry across the class. A far-end quantity needs its own loop above the local one. A forecastable load belongs in feedforward. An inventory that is not measured has to be inferred, and the inference is wrong during exactly the transients in which a leak detector most needs it [9, 10].

This chapter studies a transmission line simulator in which the station loop, the delivery loop, the forecast, the pack window and the leak detector can each be changed while a winter day runs, and the result read as time under the contract floor, compression cost and alarms. The flow model, cascade control, forecast feedforward and computational pipeline monitoring are all established methods. The

simulator joins them in a method-of-lines model with a compressor map and a leak detector that infers the line pack from three transmitters, as a simple real system does, and its scenarios run from a manual speed step through a full day to the detection limits of a mass balance. The application also offers a model predictive controller (MPC) that sets the discharge setpoint in place of the delivery loop, a method with its own literature on gas networks [6, 11], and its measured comparison with the cascade is reported here.

Section 12.2 traces transmission operations from telephone dispatch to transient simulation and leak detection. Section 12.3 states the twenty-cell line model, the compressor, the demand and the leak. Section 12.4 describes the cascade, the forecast feedforward, the pack window, the mass-balance leak detector and the predictive controller. Section 12.5 measures the scenarios, from the line as a tank to the floor carried as a constraint. Section 12.6 sets the hardware of a real line beside the model, Section 12.7 lists what the model leaves out, Section 12.8 sums up, and the exercises in Section 12.9 put the scenarios to work.

12.2 Gas Transmission Operations

Long-distance gas transmission grew with welded large-diameter pipe and centrifugal compressors in the middle of the twentieth century. Early lines were run by station operators who held discharge pressures by hand, with dispatchers setting those pressures by telephone from a few readings at the delivery end. The flow was steady enough, and the pipe forgiving enough, that the transient behavior of the line could be left to experience.

Numerical transient simulation changed that, and the optimization of compressor operation grew up beside it [3, 6], as the History box tells.

> ⧖ HISTORY **1971 to 2015: Optimizing a Line That Is Never Steady**
>
> Computing how a gas line moves between steady states came first. A technique of 1984 calculated the transient flow in high-pressure transmission systems, with and without compressors, from a partial differential equation and a numerical scheme for it [2]. Optimizing that motion proved harder. Hierarchical control of transient flow in gas networks was proposed in 1971, and an algorithm of that kind from 1998, tested on part of the British transmission network over 24 hours in steps of 2 hours, could not solve the day exactly in reasonable time [7]. For a long time the operators could plan a day as if consumption stayed constant, because cities and industry drew gas predictably on day-ahead contracts, and early studies optimized steady flows. By 2015 that assumption no longer held, and the transient problem was posed as optimal control of the whole network [8].

Leak detection developed alongside. The mass-balance method compares metered inflow and outflow with the change in computed inventory, and model-based methods for pipelines were formalized in the 1980s [9]. Its limits are the meter accuracy and the inventory estimate, and a survey of gas leak detection places it among hardware-based and acoustic alternatives with different sensitivities and detection times [10]. Section 1.4.9 places the computed inventory among the other inferred and late measurements of this volume. Line pack itself has more recently been treated as a quantity of its own, with dynamic depletion models for networks under stress [1].

12.3 Process Model

A single compressor station takes gas at 40 bar and pushes it through 100 km of 0.9 m pipe to a city gate that must never fall below 45 bar. At a steady 210 kg/s with the delivery at 48 bar, the station discharges at 60.35 bar, the compressor runs at 73.2 % speed and 17.3 MW, the line holds 2825 t of gas, and the gas itself takes about 3.7 h to cross. The arrangement is shown in Figure 12.1.

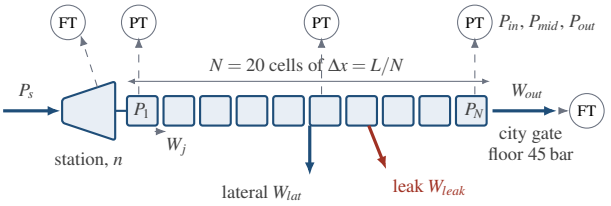

Figure 12.1 Transmission line model. A compressor at speed n takes gas at suction pressure P_s and discharges into the first of N pressure cells of length Δx, joined by faces carrying mass flows W_j. A lateral at 50 km takes W_{lat}, a leak at 60 km loses W_{leak}, and the city gate delivers W_{out} as long as P_N allows. Three pressure transmitters PT at the inlet, midpoint and outlet and two flow meters FT feed the leak detector.

Table 12.1 Variables at the Design Operating Point

Symbol	Role	Nominal	Range
P_{out}	CV, delivery pressure	48 bar	SP 45–60
P_{in}	CV (inner), discharge	60.35 bar	$\leq$ MAOP 72
$P_{SP,dis}$	MV of outer loop	60.35 bar	20–70
n	MV, compressor speed	73.2 %	50–100
$P_1 \ldots P_{20}$	states, cell pressures	—	≥ 1 bar
$W_1 \ldots W_{19}$	states, face flows	210 kg/s	—
W_{dem}	DV, city demand	140–285 kg/s	daily profile
W_{lat}	DV, lateral takeoff	0 kg/s	0–60
W_{leak}	DV, leak at 60 km	0 %	0–10
b_{in}, b_{out}	DV, meter biases	0 %	± 3

> ✦ **DID YOU KNOW? Three Million Miles of Pipe**
>
> In 2022 about 3 million miles of mainline and other pipelines in the United States delivered about 29.2 trillion cubic feet of natural gas to about 78.3 million consumers. About half of the mainline transmission network was laid in the 1950s and 1960s, after demand more than doubled following World War II [12].

The variables are classified in Table 12.1. Forty states are integrated: twenty cell pressures, nineteen internal face flows and the compressor speed.

12.3.1 Mass and Momentum on Twenty Cells

Isothermal one-dimensional flow is discretized by the method of lines. Each cell stores gas in proportion to its pressure, and each internal face carries a flow driven by the pressure difference and opposed by friction, Equations (12.1) and (12.2).

$$\frac{A\,\Delta x}{c^2}\frac{\mathrm{d}P_i}{\mathrm{d}t} = W_i - W_{i+1} - W_{lat}\,\delta_{i,i_{lat}} - W_{leak}\,\delta_{i,i_{leak}} \tag{12.1}$$

$$\frac{\Delta x}{A}\frac{\mathrm{d}W_j}{\mathrm{d}t} = P_{j-1} - P_j - \frac{f\,c^2\,\Delta x\,W_j|W_j|}{2DA^2\,\bar{P}_j} \tag{12.2}$$

where A is the pipe area, $c = \sqrt{ZRT} = 350\,\mathrm{m/s}$ the isothermal wave speed, f the Darcy friction factor, D the diameter and $\bar{P}_j$ the mean of the two neighboring cell pressures. The inlet face carries the compressor flow and the outlet face the delivered flow. In steady flow the square of the pressure falls linearly along the line, as shown in Equation (12.3).

$$P_{in}^2 - P_{out}^2 = K\,W^2, \qquad K = \frac{f\,L\,c^2}{DA^2} \tag{12.3}$$

Table 12.2 Line and Station Parameters, Marking Those the Interface Hides

Symbol	Meaning	Value	Class
L, D	length, diameter	100 km, 0.90 m	visible
f	Darcy friction factor	0.0095	visible
c	isothermal wave speed	350 m/s	visible
MAOP	pressure limit	72 bar	visible
—	contract floor	45 bar	visible
—	curtailment onset	38 bar	visible
P_s	suction pressure	40 bar	visible
W_r, r_{max}	rated flow, ratio	300 kg/s, 1.90	visible
α	map curvature	0.40	hidden
τ_N	speed response	20 s	visible
—	speed ramp limit	1.0 %/s	visible
η	compressor efficiency	0.80	visible
x_{lat}, x_{leak}	lateral, leak location	50, 60 % of L	visible

Between the first and last cell centers, 95 km apart, the discretized line gives $(N-1)KW^2/N$, which sets the design discharge pressure. The line pack is $M = \sum_i A\,\Delta x\,P_i/c^2$, the cell-by-cell form of Equation (1.8). Three time scales follow from Equations (12.1) and (12.2). The acoustic wave crosses in under five minutes, friction damps it within about half a minute of travel, and what reaches the outlet is a diffusion of pressure with a time scale of order $L^2 f|W|/(2DA\bar{P})$, about 53 min at design with the mean line pressure of 54.4 bar, and shorter at low flow. The gas velocity is a few meters per second, so the gas crosses in about 3.7 h.

Pause. Try Exercise 12.2 before reading on.

12.3.2 Compressor, Delivery and Leak

The compressor map gives the pressure ratio $r = P_1/P_s$ as a function of speed fraction $s = n/100$ and flow, and inverting it gives the flow the station delivers against the ratio the line imposes, Equation (12.4).

$$W_{in} = W_r \sqrt{\frac{s^2 - (r-1)(1-\alpha)/(r_{max}-1)}{\alpha}} \tag{12.4}$$

with rated flow $W_r = 300\,\mathrm{kg/s}$, ratio $r_{max} = 1.9$ at full speed and rated flow, and curvature $\alpha = 0.4$. The flow falls as the ratio rises, which is the only self-regulation the line has. Surge is flagged below $0.5\,W_r s$. The shaft power is $W_{in} c_p T_s (r^{0.3/1.3} - 1)/\eta$ with $c_p = 2.3\,\mathrm{kJ/(kg\,K)}$, $T_s = 288\,\mathrm{K}$ and $\eta = 0.80$, and the compression cost integrates it against a daily electricity price of 35 to 140 \$/MWh. The speed follows its command as a 20 s lag limited to 1 %/s.

The city takes its demand as long as the pressure allows. Below 38 bar the delivery is curtailed linearly to zero at 33 bar. The demand follows a winter weekday profile with a morning peak near 270 kg/s at 08:00 and an evening peak near 285 kg/s at 19:00, scaled by a demand factor and multiplied by a slow random wander of 2 % with a 600 s time constant. The leak is a fixed opening, $W_{leak} = (\ell/100)\,(210\,\mathrm{kg/s})\sqrt{P/60\,\mathrm{bar}}$, with ℓ in percent of design flow.

Pressure transmitters carry 0.02 bar of colored noise and the flow meters 0.3 % with their biases. The inlet pressure is flagged above the 72 bar maximum allowable operating pressure (MAOP) and the delivery below the 45 bar floor.

Table 12.2 gives the parameter values, and every scenario opens on the steady profile of Equation (12.3) for its initial demand.

12.3.3 Assumptions

Flow is isothermal with constant wave speed and friction factor. Temperature and compressibility changes shift the line pack and the profile by a few percent on a real line [3].

Twenty cells of 5 km reproduce the diffusion behavior and smear a wave front. The model is not suitable for surge or rupture transients, which need a much finer grid and the full momentum equation [4].

One compressor with a quadratic map and fixed suction pressure represents the station. Parallel units, recycle, driver limits and fuel gas consumption are omitted. Chapter 11 places parallel machines with recycle valves on shared headers ahead of such a line.

Demand follows a known profile with a slow wander, so the forecast is nearly perfect. Real forecast errors are larger and weather-driven [5].

The city network behind the gate is not modeled, and delivery is curtailed by a linear rule below 38 bar.

12.4 Control and Monitoring

The controllers run every 10 s and the physics advances by fourth-order Runge-Kutta steps of 2 s, on the engine of Section 1.6.1. Both loops are reverse acting in the ISA sense, because both process gains are positive, so a measurement below its setpoint raises the output ($a = +1$). Both use the velocity-form PI of Section 1.6.2, with no derivative term in either loop and the forecast of Section 12.4.2 as PIC-D's output bias u_b in Equation (1.18). Its running sum is never clipped, and at a limit only the integral is corrected, by back-calculation with a tracking time $T_t = \tau_I$, so each scan of Δt_c removes $\Delta t_c / \tau_I$ of the clipped excess, 1/240 for the delivery loop and 1/12 for the station loop. The speed lag and ramp limit act after the station loop's output and are not seen by its integral, but the loop's own 50 to 100 % range is the tighter limit.

12.4.1 The Cascade

The inner loop PIC-S holds the station discharge pressure at its setpoint by adjusting the compressor speed, with $K_c = 8$ % of speed per bar, $\tau_I = 120$ s and a speed range of 50 to 100 %. The outer loop PIC-D holds the delivery pressure by moving that setpoint, with $K_c = 1.0$ bar per bar, $\tau_I = 2400$ s and an output range from 20 bar to 2 bar below MAOP. In manual the operator sets the speed and both loops track, each setting its integral so that the output before limiting equals the value applied, the operator's speed for PIC-S and the measured discharge pressure for PIC-D. In automatic the manual speed is set to the station loop's output after every execution, so transfer is bumpless in both directions. A manual start holds the steady speed unless the reader sets another.

12.4.2 Forecast Feedforward and the Pack Window

The feedforward is the steady inlet pressure that would deliver the demand forecast a lead time T_{lead} ahead at the delivery setpoint, the static inverse of Equation (12.3), and it enters as the output bias u_b of PIC-D, Equation (12.5).

$$P_{ff}^2 = P_{SP}(t + T_{lead})^2 + K\left[W_{dem}(t + T_{lead}) + W_{lat}\right]^2 \tag{12.5}$$

The bias uses the full-length K, so at design flow it sits 0.58 bar above the 60.35 bar the twenty-cell line needs, and the integral carries the difference. The default lead is 30 min, a little over half the line's own response time, and switching the feedforward moves the integral by the change in bias so the setpoint does not bump. The pack window raises the delivery setpoint by a boost ΔP_{pack} between two hours of the day, storing gas in the line ahead of a peak.

12.4.3 The Mass-Balance Leak Detector

Every simulated minute the detector computes the line pack from the three transmitters by assuming the square of the pressure varies linearly between them, which is exact only in steady flow, Equation (12.6), where $P^2(x)$ is the piecewise linear interpolant through P_{in}^2, P_{mid}^2 and P_{out}^2 evaluated at the cell centers x_i.

$$\hat{M} = \frac{A\,\Delta x}{c^2} \sum_{i=1}^{N} \sqrt{P^2(x_i)} \tag{12.6}$$

The imbalance over a window T_w is the metered inflow less the metered outflow and the lateral, less the change in computed pack, expressed in percent of 210 kg/s, Equation (12.7), the form of Equation (1.15) with the lateral metered.

$$I = \frac{100}{210\,T_w}\Big[\int_{t-T_w}^{t} (\hat{W}_{in} - \hat{W}_{out} - W_{lat})\,dt \\ -\hat{M}(t) + \hat{M}(t-T_w)\Big] \tag{12.7}$$

An alarm is raised when $|I|$ exceeds the threshold for the persistence time, and cleared when it falls below 70 % of the threshold. The defaults are a 30 min window, a 2 % threshold and a 10 min persistence. Until a full window of samples exists the balance uses the shorter record it has, so the first half hour of a run is the noisiest. A meter bias enters I one for one, and the error in $\hat{M}$ during a ramp enters as a false imbalance.

12.4.4 Model Predictive Control

The delivery control has a second setting, never the default and pinned only by scenario 10, in which a model predictive controller (MPC) writes the discharge setpoint in place of PIC-D. Every $T_s = 10$ min it plans the setpoint over $N_p = 72$ steps (12 h), and only the first move of each plan reaches the line, the receding-horizon scheme of industrial predictive control [13]. The move is ramped over its period, the applied setpoint on scan $j = 0,\ldots,59$ being $u_a + (u_1 - u_a)(j+1)/60$, with u_a the setpoint when the plan was made and u_1 its first value, and the plan models the ramp exactly. PIC-S still turns the setpoint into speed. The setpoint is clamped at every scan to PIC-D's range, 20 bar to MAOP less its 2 bar margin as the run sets them, so MAOP stays a protection above the plan, and PIC-D tracks the applied setpoint with no preload. Section 1.6.3 describes the MPC option of the series in general. The controller here is written for this line, with a nonlinear model and its own estimator.

The controller's model is the line's own mass and friction balance reduced to five nodes and written from the declared parameter defaults of Table 12.2, never from the run's values, as shown in Equation (12.8).

$$\begin{aligned} \tau_S \frac{\mathrm{d}p_0}{\mathrm{d}t} &= u + d_{in} - p_0 \\ C_j \frac{\mathrm{d}p_j}{\mathrm{d}t} &= Q_j - Q_{j+1} - \delta_{j,2}\,(W_{lat} + d_{mid}), \quad j = 1,\ldots,4 \\ Q_j &= \frac{\Delta_j}{\sqrt{k_j\left(|\Delta_j| + k_j W_\varepsilon^2\right)}}, \qquad \Delta_j = p_{j-1}^2 - p_j^2 \\ Q_5 &= \hat{W}_{dem} + d_{out} + d_{pass} \end{aligned} \tag{12.8}$$

where p_0 to p_4 are pressures in bar at the centers of cells 1, 6, 11, 16 and 20 (2.5, 27.5, 52.5, 77.5 and 97.5 km), so that p_0, p_2 and p_4 sit at the three transmitters, u is the discharge setpoint, $\tau_S = 66$ s is the closed-loop lag of the station loop identified from setpoint steps on the plant at design flow, Q_j is the flow through segment j, $k_j = Kn_j/N$ with n_j equal to 5, 5, 5 and 4 cells and $K = 0.031\,95$ bar^2 per (kg/s)^2 from Equation (12.3), $W_\varepsilon = 5$ kg/s keeps the flow law smooth through zero flow and changes a 100 kg/s flow by 0.13 %, C_j is the cell capacity $A\,\Delta x/c^2 = 2597$ kg/bar times the node's share of 3, 5, 5, 4.5 and 2.5 cells, $\delta_{j,2}$ places the lateral at the midpoint node, $\hat{W}_{dem}$ is the city's nomination, and d_{in}, d_{mid}, d_{out} and d_{pass} are the disturbances of the estimator below. In steady flow the model's squared pressures fall linearly as in Equation (12.3), and at design it needs 60.36 bar at the inlet for 48 bar delivered, against 60.35 bar on the twenty-cell line.

Stepped by 1 bar up and down with the delivery loop frozen, the plant and the model have steady gains of 1.2572 and 1.2573 at 210 kg/s and 1.4236 and 1.4233 at 280 kg/s, 63 % times of 54 and 55 min and of 68 and 70 min, and RMS differences of 0.66 and 0.68 % of the response, 6.9 % at 120 kg/s, where the plant's falling step is not linear. One identified lag would not serve, because the

gain rises 1.3 times from 120 to 280 kg/s (case 4). Run open loop from its hourly estimate, the model predicts the delivery 1 to 12 h ahead within 0.02 to 0.06 bar RMS with the wander and the noise removed and within 0.14 to 0.51 bar RMS as shipped.

The estimator is a steady-state Kalman filter on the three pressure transmitters, and no flow meter enters it. The model is augmented with a station bias d_{in} and unexplained offtakes d_{mid} at the midpoint node and d_{out} at the city gate, all three random walks, and with a passing offtake d_{pass} at the city gate that the plan lets decay with $\tau_W = 1$ h, so that a swing of the wander is not held for the whole horizon. As many integrating disturbances as transmitters absorb a persistent mismatch at each measured point, which makes the tracking offset-free [14]. Each scan predicts with the nonlinear model and corrects, as shown in Equation (12.9).

$$\begin{bmatrix} \hat{\mathbf{p}}_k \\ \hat{\mathbf{d}}_k \end{bmatrix} = \begin{bmatrix} \mathbf{F}\left(\hat{\mathbf{p}}_{k-1}, u_{k-1}, \mathbf{w}_{k-1}, \hat{\mathbf{d}}_{k-1}\right) \\ \mathbf{E}\,\hat{\mathbf{d}}_{k-1} \end{bmatrix} + \mathbf{K}_f\left(\mathbf{y}_k - \mathbf{H}\mathbf{F}_k\right) \tag{12.9}$$

where $\hat{\mathbf{p}} = (p_0, \ldots, p_4)$ and $\hat{\mathbf{d}} = (d_{in}, d_{mid}, d_{out}, d_{pass})$ are the estimates, $\mathbf{F}$ integrates Equation (12.8) over one scan Δt_c with the setpoint and the nominations $\mathbf{w} = (\hat{W}_{dem}, W_{lat})$ that held over it, $\mathbf{F}_k$ is that prediction, $\mathbf{H}$ selects p_0, p_2 and p_4, $\mathbf{y}_k = (P_{in}, P_{mid}, P_{out})$ is the transmitter reading, $\mathbf{E} = \mathrm{diag}(1, 1, 1, e^{-\Delta t_c/\tau_W})$, and $\mathbf{K}_f$ is the 9×3 steady-state gain of the model linearized at design flow and sampled at Δt_c, with per-scan variances of 10^{-6} bar^2 on each node, 10^{-4} bar^2 on d_{in}, 10^{-3} (kg/s)^2 on d_{mid} and d_{out} and 0.8 (kg/s)^2 on d_{pass}, and 4×10^{-4} bar^2 on each transmitter, the square of its 0.02 bar noise. The gain, from 110,318 Riccati iterations, is stored, and the filter runs under either structure so that the MPC, once selected, starts from a current estimate.

The preview $\mathbf{w}$ is the dispatcher's nominations, the day's demand profile on the demand scale (or the flat demand) without its wander and the lateral, from the function that gives the feedforward of Equation (12.5) its forecast, read at the middle of each step, and with the preview off the present nomination is held. The realized demand, the wander, a leak and the meters are never read, the challenge's lateral at 10:00 is absent from the 09:50 plan, and a change to a nomination reaches the preview when it is made, as it reaches the feedforward.

A decision is the quadratic program of Equation (12.10), with the setpoint held in twelve blocks of $1, 1, 1, 1, 2, 2, 3, 4, 6, 9, 15, 27$ steps.

$$\begin{aligned} \min_{\mathbf{u}, s} \; & \sum_{k=1}^{N_p} \left(\hat{p}_{4,k} - r_k\right)^2 + \lambda \sum_{b=1}^{12} \left(u_b - u_{b-1}\right)^2 \\ & + w_E \sum_{k=1}^{N_p} e_k\, u_{b(k)} + \rho\, s^2 + \sigma s \\ \text{s.t. } & 20\,\text{bar} \le u_b \le P_{MAOP} - 2\,\text{bar}, \quad s \ge 0 \\ & \hat{p}_{4,k} \ge P_{fl} + m_k - s \\ & m_k = 0.3 + 0.5\left(1 - e^{-kT_s/\tau_m}\right) \end{aligned} \tag{12.10}$$

where $\mathbf{u} = (u_1, \ldots, u_{12})$ are the block values of the setpoint, $\hat{p}_{4,k}$ is the predicted delivery pressure at the end of step k, u_0 is the setpoint applied now, $u_{b(k)}$ is the block that holds step k, r_k is the dispatcher's delivery setpoint at that time with the pack window included, raised to $P_{fl} + m_k$ where it lies below so that no target pulls against the floor, $\lambda = 8$ weighs a squared bar of setpoint move against a squared bar of delivery error at one step, e_k is the compression energy in MWh that one more bar of discharge pressure costs over step k at the nominated flow, from the power law of Section 12.3 with the declared suction pressure and efficiency, $w_E = 0.05$ bar^2 per MWh makes that energy a tie-breaker, $\rho = \sigma = 1000$ price the slack s in bar, $P_{MAOP} = 72$ bar, $P_{fl} = 45$ bar is the contract floor, and m_k is a margin in bar with $\tau_m = 1$ h, worth 0.38 at the first step, 0.62 at 1 h, 0.78 at 3 h and 0.80 at 12 h. The margin grows with the look-ahead because the forecast's error, the wander, grows with

it. Where the delivery is already below $P_{fl} + 0.3$ bar the floor row is lowered by the shortfall times $e^{-kT_s/45\,\mathrm{min}}$, a funnel back to the band. The bounds on u are hard and the floor is soft, with one slack over the whole horizon because the shared solver carries one per output, not one per step.

The prediction is the nonlinear model run from the estimate under a nominal plan, with the passing offtake decaying, plus the departure from that plan through the model linearized at the estimate and sampled exactly over a ramped step. Two passes are made, around the last plan moved on by one decision and then around the first pass's result. Over scenario 10's 144 plans the program's prediction under the accepted plan departs from the nonlinear model's by a median of 0.0095 bar and at most 0.105 bar after the priming plan, which departs by 1.43 bar, and three to six passes move that plan's first move by under 0.001 bar and leave the day unchanged.

The margin's limit of 0.8 bar, λ and τ_W were fixed once, before the comparison, on case 7's day over seeds 1 to 10, which the comparison does not use, as the least mean cost ($40,443) of 60 combinations (limits of 0.6 to 1.5 bar, λ of 2 to 16, τ_W of 10 to 60 min) with no minute under the floor and no false alarm on any seed.

A plan is kept only if the solver reports an optimal status, checked even on a successful return, the arrays have the problem's shapes, the first value equals the first move, every block lies inside the clamp before rounding, and every value returned and every prediction published is finite. A specification test injects three such faults and rejects each. Otherwise PIC-D, which has tracked the setpoint, takes it over and the historian records a fallback, until a valid plan resumes on a ramp from where PIC-D left the setpoint. With the noise off, a rejection forced at 17:10 of scenario 10, the return, and switches to and from the cascade and from manual move the setpoint 0.009 to 0.014 bar, about one scan of either structure's own motion. No measured run rejected a plan, and the first plan is made at $t = 0$. In Chrome a plan takes 1.3 ms, under 0.8 % of a core at the top speed of 3600 simulated seconds per second.

The plant on the run's parameters, the noisy transmitters and the five-node model on the declared defaults are kept apart, so a reader's change to the pipe is mismatch the estimator carries, and replaying the same transmitter readings over another hidden plant, leak, meter bias and noise gives the same plans bit for bit.

12.5 Results

The scenarios are listed in Table 12.3 and the measured outcomes in Table 12.4. Section 1.6.4 describes how a scenario, its briefing and its completion rule work, and the exercises in Section 12.9 are built on these cases. Every figure quoted was measured on the application's own engine at its 2 s step with the default seed, with each setting changed before the run starts unless stated otherwise. Results that depend on a small margin were repeated on five other seeds of the demand wander and the instrument noise, and the spread is given beside them. Cases 1 to 4 use a flat demand with the wander removed so that the step is the only disturbance.

12.5.1 The Line as a Tank

A speed step in manual from the steady 73.16 % to 74.2 % reaches the delivery pressure only after about 12 min (a rise of 0.01 bar), and an hour later the delivery is 0.30 bar up and still rising at 0.35 bar/h while the inlet is 0.47 bar up. Most of the extra gas has gone into line pack. Holding the discharge constant with the station loop alone, a 30 kg/s demand increase takes the city gate under the floor 47 min later and leaves it near 43.5 bar, with the station on its setpoint at 60.35 bar within 0.02 bar. The delivery loop above it removes the offset, with a dip of about 1.8 bar to 46.2 bar 28 min after the step and a return to within 0.2 bar of setpoint after 124 min. Doubling K_c and halving τ_I raises the bottom of the dip only to 46.47 bar and overshoots to 48.59 bar, because the plant between the two loops is a diffusion that tuning cannot speed up.

🕮 **Pause.** Try Exercise 12.1 before reading on.

Table 12.3 Scenario Sequence

#	Scenario	Teaching point
1	Step the compressor	the line as a tank
2	The station loop alone	the wrong quantity held
3	The delivery loop	cascade
4	Night flow	flow-dependent dynamics
5	A winter day, feedback	margin carried by feedback
6	The same day, forecast	feedforward
7	A peak the line cannot carry	storage against a constraint
8	A meter that reads low	false alarm
9	A real leak	detection limits
10	The floor as a constraint	MPC against cascade
11	Challenge	graded day

Table 12.4 Measured Results

Case and condition	Result
1, speed to 74.2 %	outlet moves after 12 min, +0.30 bar at 1 h
2, station loop, +30 kg/s	delivery settles near 43.5 bar
3, cascade, same step	dip 1.8 bar to 46.2, recovery 2 h
4, 63% time at 120 / 210 / 280 kg/s[a]	34 / 54 / 68 min
4, gain at 120 / 210 / 280 kg/s[a]	1.09 / 1.25 / 1.42
4, cascade to 90%, night / design	83 / 58 min
5, feedback, 48 bar	low 44.3 bar, 52 min under floor, $37,300
6, forecast 30 min, 48 bar	low 47.0 bar
6, forecast 30 min, 46 bar	low 45.0 bar, $34,400
6, forecast 90 min, 46 bar	low 42.2 bar
7, demand +8 %, forecast, 46 bar	55 min under floor, discharge at limit
7, with 4 bar pack 12:00–21:00	low 45.0 bar, +$1,400, 1 false alarm
8, meter −1.5 %, threshold 1%	false alarm after 34 min
8, same, threshold 2%	no alarm in 24 h
9, leak 3% at 3 h, threshold 2 / 1%	found in 34 / 23 min
9, leak 1.5% from start, threshold 2 / 1%	not in 3 h / 32 min
9, leak 1% from start	not in 3 h at either

[a] Model test with the delivery loop off and the discharge setpoint stepped 1 bar directly, which the interface does not offer.

The line changes with flow. With the station holding a fixed discharge, the delivery answers a 1 bar change with a 63 % time of 34, 54 and 68 min at 120, 210 and 280 kg/s, and with gains of 1.09, 1.25 and 1.42, measured by stepping the discharge setpoint directly in the model. The cascade is nevertheless slower at night, 83 min to 90 % of a 1 bar setpoint rise against 57 min rising and 58 min falling at design, because the lower gain matters more than the faster response.

12.5.2 A Winter Day

Over a winter day at a 48 bar setpoint, feedback alone lets the delivery fall to about 44.3 bar on the morning ramp and spend 52 min under the floor (40 to 62 min on the other seeds), for about $37,300 of compression. With the 30 min forecast the same setpoint holds the delivery above 47.0 bar (46.7 to 47.2 bar on the other seeds), as shown in Figure 12.2, and the setpoint can be lowered to 46 bar with the lowest delivery at 45.01 bar, for about $34,400, 8 % less. That margin is 0.01 bar, and on two of the five other seeds the same day spends 25 and 38 min under the floor. The forecast reduces the margin feedback must carry, and the margin is paid for in compression all day. At 46 bar a lead of 15 min gives a lowest delivery of 45.15 bar, and 45 min already falls to 44.41 bar. A lead of 90 min is too long, because the setpoint then moves ahead of the line's own response and the delivery sags to about 42.2 bar.



Figure 12.2 The transmission line running scenario 6, a winter day with the demand forecast as feedforward at a 48 bar delivery setpoint, at 23:59, cropped to the plant view. Callout 1 marks the pressure along the line, with MAOP and the contract floor dashed and the three leak detector transmitters as dots. Callout 2 is the compressor with its speed, and the pipe runs from it to the city gate at callout 3 with tracers moving at the gas velocity. Callout 4 sets the true line pack against the pack computed from three transmitters, and callout 5 the leak detector imbalance against its threshold, beside the clock with the electricity price. The run's scorecard, cut from this view, read a lowest delivery of 47.05 bar, no time below the floor, a highest inlet of 70.03 bar, no leak alarms and $36,412 of compression.

> **TRY IT IN THE APP** **5 · A winter day, feedback only**
>
> Load the case and press Run at a high speed. Touch nothing. Watch the Delivery pressure trend against its floor line through the morning and evening ramps. On the morning ramp the delivery falls to about 44.3 bar, and over the day it spends 52 min below the floor. Read the Time below the floor and Compression cost tiles at the end of the day.

12.5.3 A Peak the Line Cannot Carry

With demand raised 8 %, the evening peak of 308 kg/s equals the line's steady capacity with 70 bar in and 45 bar out, and the demand wander takes it past that. With the forecast and a 46 bar setpoint, the discharge setpoint sits at its 70 bar limit for 76 min and the delivery still spends 55 min under the floor, with a low of 44.27 bar. No tuning changes that outcome, because the constraint and not the controller sets it. A 4 bar pack from 12:00 to 21:00 stores gas in the line during the afternoon and raises the evening low to 45.58 bar for about $1,400 more compression, and the pressure ramp it causes sets off one false leak alarm at 12:40. The lowest delivery of the packed day, 45.00 bar, then comes at 21:55, after the window closes and the setpoint falls back 4 bar. The capacity-limit comparison is within the spread of the wander: on the five other seeds the unpacked day spends 0 to 60 min under the floor and the packed day 8 to 32 min, on the morning ramp or after the window closes.

Pause. Try Exercise 12.3 before reading on.

Table 12.5 Challenge Outcomes Over One Day

Configuration	Low (bar)	Under floor	False	Cost ($)
as loaded	44.33	52 min	7	42,758
threshold 2.5%	44.33	52 min	0	42,758
forecast, 48 bar, 2.5%	46.05	0	0	41,917
forecast, 47 bar, 2.5%	45.62	0	0	40,934

12.5.4 What the Mass Balance Can Find

The leak detector is as good as its meters and its inventory estimate. An outlet meter reading 1.5 % low with a 1 % threshold produces an alarm 34 min later with no leak at all (30 to 34 min on the other seeds), and a 2 % threshold does not alarm in a day. A 3 % leak at 60 km is found 34 min after it starts at a 2 % threshold and 23 min after at 1 % (31 to 37 and 18 to 23 min on the other seeds). The leak sizes are set at 60 bar, so at the 52.6 bar of the leak site the 3, 1.5 and 1 % leaks lose 2.81, 1.41 and 0.94 % of design flow. Because the scripted leak at 3 h replaces whatever the slider holds, the smaller leaks are started at $t = 0$. A 1.5 % leak is not found in the three hours before the scripted leak at 2 %, and at 1 % the alarm comes 32 min after it starts (11 to 47 min on the other seeds). A 1 % leak is not found in three hours at either. Each of these times includes the 10 min persistence, and by its alarm the 3 % leak has released 12.1 t of gas at the 2 % threshold and 8.2 t at 1 %. The threshold that avoids the false alarm from the biased meter is the threshold that misses the 1.5 % leak, which is the trade the method is known for [9, 10].

> **TRY IT IN THE APP** **9 · A real leak, and how long the balance takes**
>
> Load the case and press Run. At three hours a leak of 3 % of design flow opens at 60 km. Watch the Leak detector imbalance trend climb toward its threshold, then read the Leak found after tile: about 34 min. Press Reset, set Leak alarm threshold to 1 % before Run, and run again. The alarm now comes after about 23 min.

12.5.5 The Graded Challenge

The challenge runs a winter day with a gas turbine plant starting on the lateral at 10:00 and the outlet meter reading 1 % low, starting with feedback only at 48 bar and a 1 % leak threshold. It is met with the delivery never under the floor, the inlet never over MAOP, no false alarms and a compression cost of $41,500 or less, on the case as given: changing a model parameter or a disturbance, or firing an extra disturbance event, voids the attempt. Table 12.5 reports the measured outcomes. As loaded, the day fails on three counts, with 52 min below the floor, seven false alarms and $42,758, $1,258 over budget.

A higher threshold alone removes the seven false alarms and nothing else. The forecast at 48 bar removes the floor violations but costs $41,917, over the target, and only the combination of the forecast with the setpoint lowered to 47 bar meets every condition, at $40,934, $566 under budget, with a lowest delivery of 45.62 bar and a highest inlet of 70.08 bar. At a 2 % threshold the same day raises two false alarms, at 07:48 on the morning ramp and at 10:40 after the lateral starts. On the five other seeds the suggested configuration passes four times and fails once, on one false alarm at 10:40 at 2.5 %.

12.5.6 The Floor as a Constraint

Scenario 10 runs case 7's day (demand 8 % above its profile, a 46 bar setpoint, the forecast on) under the MPC. Table 12.6 compares it with the cascade on the same forecast and with variants that locate the difference, and sets the two structures side by side on cases 5 and 6 and the challenge, at the default seed and on seeds 11 to 30, each a full day with the same transmitters, noise and nominations for both, and with the MPC on scenario 10's settings throughout.

Table 12.6 Model Predictive Control Against the Cascade Over Full Days, at the Default Seed and on Seeds 11 to 30

	Default seed							Seeds 11 to 30			
Structure	Under (min)	Lowest (bar)	Cost ($)	False alarms	IAE (bar h)	Travel (%)	Plans	Under, min mean (worst)	Days under	Cost ($)	False alarms
Case 10, case 7's day, demand 108 %, 46 bar											
cascade + forecast (case 7)	54.6	44.27	40,213	0	7.65	306	—	20.2 (75.6)	16	39,938	0
cascade + forecast at 46.4 bar	21.9	44.44	40,593	0	7.93	303	—	1.5 (25.8)	3	40,337	0
cascade + forecast at 47.2 bar	9.6	44.76	41,347	0	8.61	297	—	—	—	—	—
+ pack 4 bar 12:00 to 21:00 (case 7)	0.0	45.00	41,633	1	14.71	417	—	22.4 (53.4)	17	41,441	19
+ pack 2 bar 16:00 to 19:00, hindsight[a]	0.0	45.11	40,294	0	8.24	338	—	31.8 (98.5)	18	40,043	0
cascade, feedback only	333.6	41.57	41,315	0	31.79	306	—	343.8 (365.1)	20	41,033	0
MPC (scenario 10)	10.4	44.73	40,623	0	9.54	217	144	1.4 (25.2)	2	40,381	0
MPC, preview off	301.6	42.79	41,690	0	30.64	247	144	269.6 (326.1)	20	41,449	1
MPC, 1 h horizon	24.7	44.40	41,372	0	16.90	232	144	5.1 (29.6)	5	41,102	0
MPC without the floor row[b]	96.5	44.30	40,103	0	10.58	187	144	63.7 (162.4)	19	39,820	0
Case 5, 48 bar, no forecast											
cascade, feedback only	51.6	44.33	37,330	0	26.56	298	—	54.7 (86.3)	20	37,001	0
MPC, preview off	48.0	44.09	37,417	0	31.06	213	144	61.7 (112.4)	20	37,138	0
Case 6, 48 bar											
cascade + forecast	0.0	47.05	36,418	0	5.78	294	—	0.0 (0.0)	0	36,093	0
MPC	0.0	46.64	36,245	0	9.17	178	144	0.0 (0.0)	0	35,979	0
Case 6, 46 bar											
cascade + forecast	0.0	45.01	34,396	0	5.89	299	—	9.5 (32.4)	11	34,069	0
MPC	0.0	45.26	34,762	0	8.64	218	144	0.4 (7.3)	1	34,431	2
Challenge (case 11), as loaded											
cascade, feedback only	51.6	44.33	42,758	7	26.82	285	—	54.7 (86.3)	20	42,524	128
MPC	0.0	46.13	41,734	7	10.69	180	144	0.0 (0.0)	0	41,507	133
Challenge (case 11), suggested settings[c]											
cascade + forecast	0.0	45.62	40,934	0	7.85	324	—	0.0 (0.0)	0	40,659	7
MPC	0.0	45.62	40,782	0	10.01	184	144	0.0 (0.0)	0	40,517	1

Under is the time with the true delivery below the 45 bar floor, counted at every 2 s step, and false alarms are leak alarms with no leak. IAE is the integrated absolute delivery error and travel the total travel of the compressor speed, as the application's loop metrics report them. Each MPC row makes 144 plans a day and rejects none, at a mean of 1.3 ms a plan in Chrome. [a]The cheapest of 288 pack windows on this day, chosen knowing the outcome. [b]A copy of the controller with the floor constraint removed, tracking the 46 bar setpoint with the same forecast. [c]Forecast with a 30 min lead, 47 bar and a 2.5 % threshold, with the MPC selected in its row. The challenge is met in both suggested rows and in neither as-loaded row, and on seeds 11 to 30 with the suggested settings on 13 days under the cascade and 19 under the MPC.

On the shipped day the MPC spends 10.4 min under the floor against 54.6 min for the cascade, for $40,623 against $40,213, and on the further days it goes under on 2 against 16. The floor in its problem raises the setpoint sooner, as shown in Figure 12.3. At 17:40 its setpoint is at the 70 bar clamp and the cascade's at 68.1 bar, with 3029 t against 2983 t in the line, but the extra gas is modest, 30 t at 17:00 and 43 t at 18:00. The wander decides the evening, taking the demand to 327.6 kg/s at 19:06 against a nominated 307.8 kg/s, and both deliveries are lowest at 19:08, 44.73 bar under the MPC and 44.27 bar under the cascade. With the wander removed the two days stay at or above 45.71 and 45.32 bar.

The ablations credit the gain to the floor carried as a constraint, not to packing as a plan. Without the floor row the same controller tracks 46 bar and spends 96.5 min under the floor, more than the cascade, and without its preview it spends 301.6 min, against 333.6 min for feedback alone (48.0 against 51.6 min on case 5's day). A 1 h horizon gives 24.7 min for $749 more, and 3 and 6 h give 13.8 and 9.9 min. A fixed margin buys most of the same. Raised 0.4 bar, the cascade spends 21.9 min

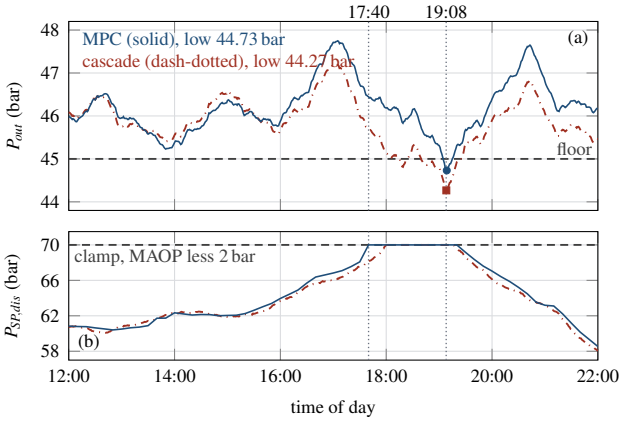

Figure 12.3 Scenario 10's afternoon and evening under the MPC (solid) and under the cascade with the same forecast (dash-dotted), at the default seed. (a) True delivery pressure against the 45 bar floor (dashed), with each day's lowest value at 19:08 marked, a circle for the MPC and a square for the cascade. (b) The discharge setpoint applied, against its 70 bar clamp (dashed), which the MPC reaches at 17:40 and the cascade at 18:00. Dotted verticals mark 17:40 and 19:08.

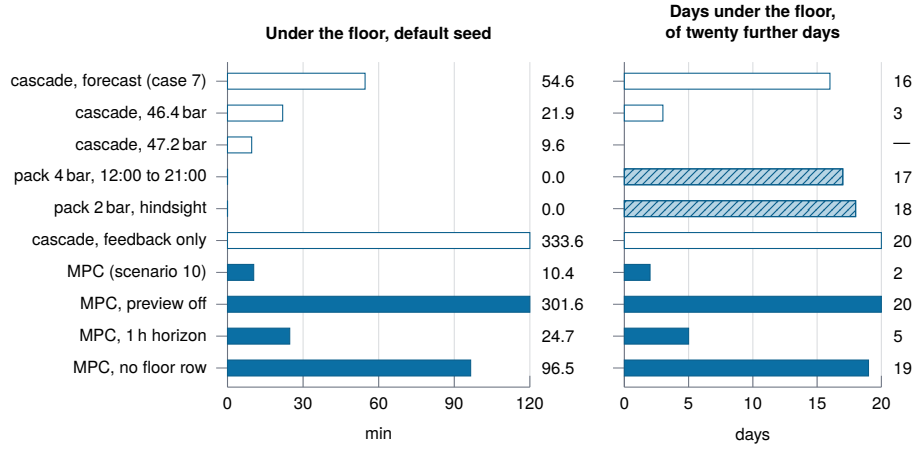

Figure 12.4 Case 10's day under each structure of Table 12.6: minutes under the 45 bar floor at the default seed (left) and days under it among the twenty further days of seeds 11 to 30 (right). Open bars are the cascade, hatched bars the two pack schedules and solid bars the MPC, and every row but feedback only carries the forecast. Two bars of the left panel are cut at the axis end, the column beside each panel gives the table's value, and the row at 47.2 bar was run at the default seed only. Notice that the pack schedules keep the floor on the shipped day and fail on most further days, while the MPC and the cascade raised 0.4 bar go under on 2 and 3.

under the floor on the shipped day, but on the further days 1.5 min on average against 1.4 and 3 days against 2, for \$40,337 against \$40,381. Only on the shipped day, a hard one, is the plan clearly ahead, because the cascade needs 47.2 bar and \$724 more to reach 9.6 min. Figure 12.4 sets the shipped day beside the further days for each structure.

Storage scheduled by hand does better on the day it is placed for. Case 7's 4 bar pack from 12:00 to 21:00 keeps the floor for \$41,633, and a 2 bar pack from 16:00 to 19:00, the cheapest of 288 windows in hindsight, for \$40,294, \$329 less than the MPC. On the further days the two leave the

delivery under the floor on 17 and 18 days, because their lost minutes fall outside their windows, 166 min in sum on the morning ramp for both, 282 min after 21:00 for the 4 bar pack and 434 min between 16:00 and 21:00 for the 2 bar pack, where the MPC loses 27 min. The 4 bar pack also raises 19 false alarms on those days.

The MPC loses where the forecast alone suffices. On case 6's day at 48 bar it tracks the setpoint less closely than the cascade with the forecast, an IAE of 9.17 against 5.78 bar h, with 178 against 294 % of speed travel. The move price is the cause, since with $\lambda = 2$ the same plan reaches 5.28 bar h for 210 % of travel and $275 more. At 46 bar both keep the floor on the shipped day and the MPC costs $366 more. On the further days it goes under on 1 against 11, but on one of them it raises two false leak alarms, at 08:23 and 19:04, where the cascade raises none. Its setpoint had swept 4.4 and 2.7 bar in the 40 minutes before each, and the pack computed from three transmitters trailed the true pack by 12 to 14 t, the mechanism of case 7's false alarm.

The MPC reads no meter and leaves the leak detector as it was. Through case 9's 3 % leak of 5.91 kg/s its persistent unexplained offtake lies between 4.00 and 6.17 kg/s from 03:30 on, and the balance finds the leak 34.0 min after it starts under either structure (23.0 min at 1 %). With case 8's meter 1.5 % low its offtake estimate is −0.19 kg/s at 06:00, and the false alarm comes 34.0 min after the fault under both.

As loaded, the challenge fails under the MPC as it does under the cascade. The floor holds, against 51.6 min under it for the cascade, but the 1 % threshold on a 1 % meter error raises seven false alarms under both, and the MPC's compression, $41,734, is $234 over budget. With the suggested settings both meet it at the default seed, the MPC at $40,782 against $40,934, and on the further days the MPC meets it on 19 and the cascade on 13, the cascade failing on false alarms at 10:40 after the lateral starts and the MPC once, at 08:43. Of the limits the rule leaves unscored, the inlet passes the 70 bar discharge clamp by at most 0.01 bar for 3.1 min under the MPC and 0.08 bar for 44.0 min under the cascade. MAOP stays a protection above the plan. No run at the default seed took the applied setpoint over the clamp, and the highest inlet of any measured day, 70.13 bar on scenario 10 under the MPC (70.10 bar under the cascade), is the station loop's overshoot of its 70 bar setpoint, 1.87 bar under MAOP.

> ✦ DID YOU KNOW? **More Than a Hundred False Alarms**
>
> From the simulator (Table 12.6). Run as loaded, with a 1 % leak threshold on an outlet meter reading 1 % low, the challenge day raises 128 false leak alarms over the twenty further days under the cascade and 133 under the MPC. With the suggested settings the counts fall to 7 and 1.

12.6 Technology in Practice

Figure 12.5 draws the station loop, the delivery loop and the leak detector of Section 12.4 with the hardware of a transmission line in place of the schematic's symbols.

What Measures It

The station loop reads the pressure transmitter at the discharge and the delivery loop the one at the city gate, so the contract's quantity is measured where no controller sits. The leak detector adds the midpoint transmitter and both flow meters. A gas meter at a city gate can be a transit-time ultrasonic meter: sound pulses sent with and against the flow arrive at different times, the difference gives the velocity, several paths average the flow profile, and the speed of sound the meter infers can be checked against the gas composition to catch a faulty sensor [15]. The app's case 8 names this meter, losing a path and reading low.

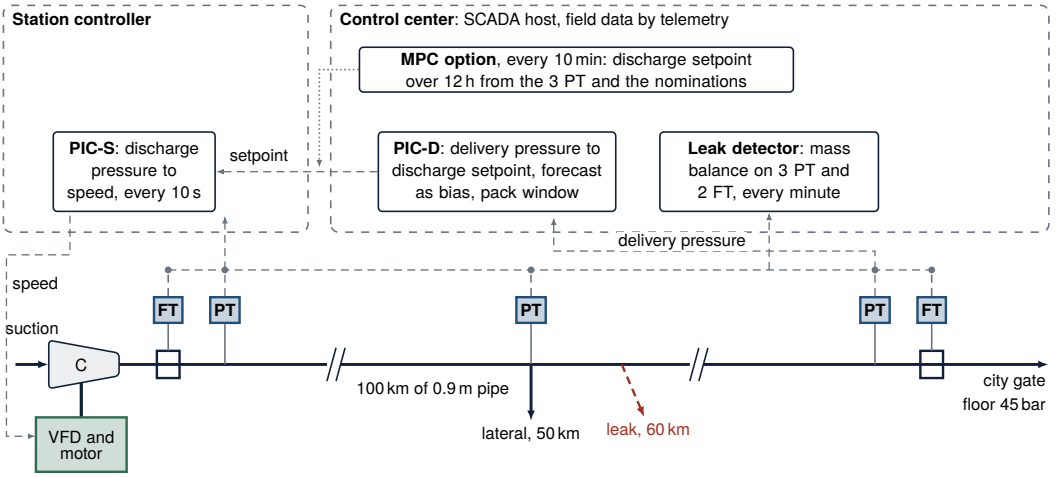

Figure 12.5 The loops of the line with their hardware. Lettered boxes are sensing devices with the chapter's tags: a PT at the station, at the midpoint and at the city gate, and an FT at each end. The green box is the final element, the variable-frequency drive and motor that turn the compressor. Dashed arrows are signals, the dotted arrow the discharge setpoint that the MPC option writes in place of PIC-D. Notice that the delivery loop reads one transmitter 100 km from the machine it moves, through telemetry, while the leak detector reads all five instruments and moves nothing. Not to scale, and not a piping diagram.

> **TECHNOLOGY** **Sensor: The Pressure Transmitter at the City Gate**
>
> An electronic pressure transmitter turns pressure into an electrical signal through a strain gauge, a capacitance cell or a resonating element. In the capacitance cell a taut diaphragm between two fixed plates flexes with the pressure, and the transmitter sends the reading on a 4 to 20 mA loop or as a digital signal [15]. The live zero of the loop shows a broken wire as no current. A transmitter that fails drives its current below 3.6 mA or above 21 mA, so its trend jumps off scale, and the control system can put the loop in manual [15]. The model's three transmitters carry 0.02 bar of colored noise, and the delivery loop reads the one 100 km from the station.

What Moves It

One handle moves the whole line, the compressor's speed. The delivery loop writes only the discharge setpoint, clamped 2 bar below MAOP, and that clamp stands for a legal limit. No segment of a United States gas line may run above its maximum allowable operating pressure, the lowest of several limits that include the design pressure of its weakest element and its test pressure after construction divided by a factor of 1.1 to 1.5 [16].

TECHNOLOGY Final Element: The Compressor's Variable-Frequency Drive

Most pipeline compressors burn gas taken from the line, but electric drives have been growing, mainly for environmental reasons [17]. An induction motor turns at a speed set by the frequency of its supply. A variable-frequency drive rectifies the line power to direct current, smooths it and inverts it back to alternating current at the frequency the speed calls for. Such drives limit the rate of acceleration and deceleration, and they shut themselves down on overcurrent or overvoltage [15]. The model's 20 s speed lag and 1 %/s ramp limit stand for such a drive, and the station loop's output is its speed command.

Where the Controller Runs

A transmission line is run through supervisory control and data acquisition, which lets the control room direct flows and hold pressures [17]. A control server at a control center gathers the field data, remote terminal units or programmable logic controllers at the field sites control each local process, and loops run with cycle times from milliseconds to minutes [18]. The app's loops execute every 10 s, the detector every simulated minute and the MPC every 10 min. The chapter's structure maps onto standard blocks: a PI in the station's controller from discharge pressure to speed, a PI in the control center from delivery pressure to that setpoint with the forecast on its bias input and the pack window on its setpoint, and the mass balance as a pipeline monitoring application beside them.

IN PRACTICE What the Control Room Rules Ask

In the United States an operator whose controllers run a transmission line with a compressor station through supervisory control and data acquisition must follow written control room procedures. Its alarm plan identifies, at least once each calendar month, the safety points that generated false alarms, had alarms inhibited, were taken off scan or held forced values. Alarm set points and descriptions are verified each calendar year, within 15 months, and the controllers' workload is reviewed so that they have time to react to alarms [19]. Case 8's false alarm from a meter reading low is the kind of point the monthly review lists.

What Is Modernizing It

The first of three developments is transient planning: a 2015 formulation posed a day of compressor operation as optimal control of the network's transient flow, minimizing compression cost or the deliveries left unmet, with gains in security and efficiency over methods that assume steady flow [8]. The MPC option of Section 12.4.4 is one line's version of that plan. The second is leak detection beyond the mass balance. Optical fiber laid with a pipeline picks up the temperature around it, and the travel time of the light locates a leak, with a spatial resolution of the order of meters on long-distance systems [20]. The third is security. A penetration test of a natural gas utility's business network strayed into a part connected to its control system, locked that system up, and left the utility unable to send gas through its pipelines for four hours [18].

12.7 Limitations

The flow is isothermal and the friction factor fixed, so the line pack carries errors of a few percent against a model with an energy balance and a real gas equation of state [3].

The grid is coarse. Wave fronts are smeared over several cells and the model says nothing about surge or rupture.

The forecast is nearly perfect because the demand is a known profile with a small wander. With realistic forecast errors, part of the benefit shown here returns to the feedback loop [5].

The station is one compressor with no fuel use and no recycle, and the electricity price is a fixed daily shape.

The leak detector is the simplest computational pipeline monitor: a windowed balance with a steady-state inventory estimate. Transient-model detectors and acoustic methods have different limits [10].

The predictive controller's model leaves out what decides the floor on the peak day. Its preview carries no wander, which on scenario 10's day takes the evening demand 19.8 kg/s above the nominated peak and both structures under the floor, while without the wander neither goes under. A real forecast's own error, larger and driven by the weather [5], would fall on the same margin. The model also omits the inertia of the gas, the curvature of the profile between nodes and the compressor map, its last block spans the final 4.5 h of a plan, the solver carries one slack over the horizon rather than one per step, and it has no model of the leak detector, whose pack estimate from three transmitters its ramps can outrun.

> ⚠ **OPEN CHALLENGE** **Gas-Fired Power Plants on the Line**
>
> The chapter's day follows a known profile, and its challenge adds one power plant on the lateral. Real lines increasingly feed gas-fired generators, which change their output quickly and are dispatched to balance wind and solar power. Their intermittent dispatch creates rapidly changing withdrawals and stresses that ad hoc methods find harder and harder to keep within design limits, and gas interrupted to generators can force electric load shedding [8]. Optimization of transient flow remains one of the most challenging research areas, still in a developing phase [7]. Progress would be a line that plans its pack from the power system's schedule and still holds every delivery floor when a plant starts unannounced.

12.8 Conclusion

A transmission line is a storage vessel with pressure transmitters at three points, and the simulator of this chapter makes that consequential. A 1 % speed step takes 12 min to reach the city gate and an hour later is still arriving. A station loop that holds its own discharge leaves the delivery at 43.5 bar after a demand step, and a delivery loop above it recovers with a 1.8 bar dip over two hours. Over a winter day, feedback alone at 48 bar spends 52 min under the 45 bar floor, while a 30 min forecast lets the setpoint fall to 46 bar with no violation on the default seed and 8 % lower compression cost. When the peak exceeds the line's capacity at its pressure limit, only storage helps, and a 4 bar afternoon pack keeps the floor for about $1,400, with a margin the demand wander can erase. The mass balance finds a 3 % leak in about 34 min and cannot tell a 1.5 % leak from a 1.5 % meter error. The optional predictive controller, holding the floor as a soft constraint over a twelve-hour plan, spends 10.4 min under it on scenario 10's day against 54.6 min for the cascade with the same forecast and goes under on 2 of twenty further days against 16, although a setpoint 0.4 bar higher matches it on those days and a pack chosen in hindsight keeps the floor on the shipped day for less. The model is isothermal, coarse and forecast-friendly. A fuller model would add an energy balance, several compressor units, forecast error and a transient-model leak detector. As it stands, the line is a plant on which to study the control of an inventory that cannot be measured.

KEY POINTS

- A gas line stores its own gas, so a speed step at the station takes 12 min to reach the city gate.
- A station loop on its discharge leaves the delivery short after a demand step, and a delivery loop above it recovers only after a 1.8 bar dip.
- A 30 min forecast lets the setpoint fall from 48 to 46 bar at 8 % lower compression cost.
- The mass balance cannot tell a 1.5 % leak from a 1.5 % meter error.
- The predictive controller goes under the floor on 2 of twenty further days against 16, but a setpoint 0.4 bar higher matches it on those days.

Where else this applies. Water released ahead of demand along Volume 9, Chapter 5, heat sent ahead of a morning peak in Volume 9, Chapter 6, and a liquid line with a leak detector of its own in Chapter 13 share the same distributed inventory.

References

[1] S. Chevalier and D. Wu, "Dynamic linepack depletion models for natural gas pipeline networks," *Applied Mathematical Modelling*, vol. 94, pp. 169–186, 2021.

[2] A. Osiadacz, "Simulation of transient gas flows in networks," *International Journal for Numerical Methods in Fluids*, vol. 4, no. 1, pp. 13–24, 1984.

[3] A. Thorley and C. Tiley, "Unsteady and transient flow of compressible fluids in pipelines—a review of theoretical and some experimental studies," *International Journal of Heat and Fluid Flow*, vol. 8, no. 1, pp. 3–15, 1987.

[4] T. Kiuchi, "An implicit method for transient gas flows in pipe networks," *International Journal of Heat and Fluid Flow*, vol. 15, no. 5, pp. 378–383, 1994.

[5] B. Soldo, "Forecasting natural gas consumption," *Applied Energy*, vol. 92, pp. 26–37, 2012.

[6] D. Marqués and M. Morari, "On-line optimization of gas pipeline networks," *Automatica*, vol. 24, no. 4, pp. 455–469, 1988.

[7] R. Z. Ríos-Mercado and C. Borraz-Sánchez, "Optimization problems in natural gas transportation systems: A state-of-the-art review," *Applied Energy*, vol. 147, pp. 536–555, 2015.

[8] A. Zlotnik, M. Chertkov, and S. Backhaus, "Optimal control of transient flow in natural gas networks," in *2015 54th IEEE Conference on Decision and Control (CDC)*, 2015, pp. 4563–4570.

[9] L. Billmann and R. Isermann, "Leak detection methods for pipelines," *Automatica*, vol. 23, no. 3, pp. 381–385, 1987.

[10] P.-S. Murvay and I. Silea, "A survey on gas leak detection and localization techniques," *Journal of Loss Prevention in the Process Industries*, vol. 25, no. 6, pp. 966–973, 2012.

[11] A. Gopalakrishnan and L. T. Biegler, "Economic nonlinear model predictive control for periodic optimal operation of gas pipeline networks," *Computers & Chemical Engineering*, vol. 52, pp. 90–99, 2013.

[12] U.S. Energy Information Administration, "Natural gas pipelines," Energy Explained, last updated March 19, 2024, 2024, accessed November 18, 2026. [Online]. Available: https://www.eia.gov/energyexplained/natural-gas/natural-gas-pipelines.php

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

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

[15] 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.

[16] Pipeline and Hazardous Materials Safety Administration, "49 CFR 192.619, maximum allowable operating pressure: Steel or plastic pipelines," Code of Federal Regulations, Title 49, Part 192, edition of October 1, 2025, govinfo.gov, 2025.

[17] U.S. Energy Information Administration, "Natural gas pipelines: Transportation process and flow," About U.S. Natural Gas Pipelines, archived analysis, based on data through 2007/2008 with selected updates, 2008, accessed November 18, 2026. [Online]. Available: https://www.eia.gov/naturalgas/archive/analysis_publications/ngpipeline/process.html

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

[19] Pipeline and Hazardous Materials Safety Administration, "49 CFR 192.631, control room management," Code of Federal Regulations, Title 49, Part 192, edition of October 1, 2025, govinfo.gov, 2025.

[20] J. Li and M. Zhang, "Physics and applications of Raman distributed optical fiber sensing," *Light: Science & Applications*, vol. 11, p. 128, 2022.

12.9 Exercises

12.1 Concept check · 2 min

In scenario 2 the station loop holds the discharge at 60.35 bar while the city demand steps from 210 to 240 kg/s. The delivery pressure settles near 43.5 bar and stays there. Why does it not recover?

(a) With the inlet held, $P_{out}^2 = P_{in}^2 - KW^2$, so a larger flow leaves a lower delivery for good, and no loop measures it.

(b) The extra flow drains the line pack at a steady rate, so the delivery keeps sinking slowly toward the curtailment pressure.

(c) The compressor reaches its 100 % speed limit, so the station cannot supply the extra 30 kg/s.

(d) The city gate curtails its offtake below the floor, which pins the delivery at the curtailment pressure.

12.2 Concept check · 1 min

A pressure wave crosses the 100 km line in under five minutes, yet in scenario 1 a 1 % speed step first moves the delivery pressure after about 12 min, and an hour later it is still rising. Which statement explains this?

(a) Pressure travels with the gas, which takes about 3.7 h to cross, so the delivery moves only as the new gas arrives.

(b) Friction damps the wave within about half a minute of travel, and what reaches the outlet is a slow diffusion of pressure as the line pack fills.

(c) The speed follows its command through a 20 s lag and a 1 %/s ramp, which holds back the extra flow for several minutes.

(d) The colored noise of the delivery transmitter hides the first movement until the change grows larger than the noise.

12.3 Concept check · 2 min

In scenario 7 with a 4 bar pack from 12:00 to 21:00, the leak detector raises one alarm at 12:40, although nothing leaks and both meters read true. What produces the imbalance?

(a) The 0.3 % noise of the two flow meters adds up over the 30 min window until it crosses the 2 % threshold.

(b) The higher delivery setpoint raises the flow the city takes, and the outlet meter does not register the extra delivery.

(c) The pack computed from three transmitters assumes a steady profile, so it trails the true pack during the ramp, and gas stored in the line is counted as lost.

(d) The compressor nears surge as the window opens, and the recycled gas passes the inlet meter twice.

12.4 Theory and calculation · 7 min

Use the line of Table 12.2: $L = 100\,\text{km}$, $D = 0.90\,\text{m}$, $f = 0.0095$, $c = 350\,\text{m/s}$, $N = 20$ cells.

(a) Compute K of Equation (12.3) in bar^2 per $(\text{kg/s})^2$. *Checkpoint:* $K = 0.031\,95$.

(b) For 210 kg/s delivered at 48 bar, find the inlet pressure with the full-length K and with the twenty-cell line's $(N-1)K/N$. Which does the feedforward of Equation (12.5) use, and what carries the difference? *Checkpoint:* 60.93 bar and 60.35 bar.

(c) Take the mean pressure of a steady line as $\bar{P} = \frac{2}{3}(P_{in}^3 - P_{out}^3)/(P_{in}^2 - P_{out}^2)$. Find the line pack from Equation (1.8) and the time the gas takes to cross, M/W.

12.5 Theory and calculation · 7 min

Case 4 measures the delivery's answer to a 1 bar step of a held discharge: 63 % times of 34, 54 and 68 min and gains of 1.09, 1.25 and 1.42 at 120, 210 and 280 kg/s. Take the delivery at 48 bar, the twenty-cell $0.95K$ with $K = 0.031\,95$, and $A = 0.6362\,\mathrm{m}^2$.

(a) Differentiate Equation (12.3) at fixed W to find the steady gain $\mathrm{d}P_{out}/\mathrm{d}P_{in}$, and evaluate it at the three flows. *Checkpoint:* 1.09, 1.26 and 1.43.

(b) Evaluate the diffusion time $L^2 fW/(2DA\bar{P})$ of Section 12.3 at the three flows, with $\bar{P}$ from the formula of the previous exercise. *Checkpoint:* 33, 53 and 66 min.

(c) The line answers faster at night, yet the cascade tuned at design reaches 90 % of a 1 bar setpoint step in 83 min at night against 58 min at design. Explain.

12.6 Theory and calculation · 6 min

The leak detector of Section 12.4.3 uses a 30 min window, a 10 min persistence and the leak law of Section 12.3. Case 9 opens a 3 % leak at 60 km, where the line is at 52.6 bar.

(a) Find the leak flow in percent of design flow and in kg/s. *Checkpoint:* 2.81 %, 5.90 kg/s.

(b) Model the imbalance as rising linearly from zero to the leak over one window. When does it cross a 2 % and a 1 % threshold, and when would the alarm come? Compare with the measured 34 and 23 min. *Checkpoint:* 21.4 and 10.7 min, alarms at 31.4 and 20.7 min.

(c) How much gas has the leak released by the measured alarms? What is the smallest leak setting that a 2 % threshold can ever see here?

12.7 App challenge · 7 min

Load scenario 1, "Step the compressor, time the outlet (MANUAL)", default seed, at speed 60 (one simulated minute per second). Press Run. When the Clock tile reads between 00:05 and 00:15, set Compressor speed to 74.2 % on the station card and note the time. Run one more hour of simulated time, pause and export the CSV.

(a) After the step, find the first row in which `p_out_bar` is at least 0.01 bar above its value at the step, and the changes in `p_out_bar` and `p_in_bar` one hour after the step. *Checkpoint:* 12 to 13 min, then +0.30 and +0.47 bar.

(b) Read `pack_t` at the step and an hour later. Use Equation (1.8) to turn the gain into a rise of the mean pressure, and compare it with (a).

(c) Over this hour, does the delivery behave more like a first-order lag or like an integrator? What will finally stop the rise?

12.8 App challenge · 9 min

Load scenario 6, "The same day with a forecast", default seed. Set Delivery pressure setpoint to 46 bar before Run. Run the full day at speed 3600 (one simulated hour per second) and export the CSV. In the row at `time_s` = 86400 read `metric_min_pout_bar`, `metric_below_floor_s` and `metric_cost_usd`. The Lowest delivery pressure, Time below the floor and Compression cost tiles show the same values at 24:00, but the cost keeps counting after midnight. Repeat with Forecast lead set to 15, 45 and 90 min, each after a Reset and before Run.

(a) Tabulate the lowest delivery, the minutes below the floor and the cost for the leads of 15, 30, 45 and 90 min. *Checkpoint:* 30 min lead: 45.01 bar, 0, $34,396.

(b) Case 5 at 48 bar without the forecast costs $37,330. What does the 30 min forecast at 46 bar save? Is its margin one to operate on?

(c) Explain the trend of the lowest delivery and the cost with the lead.

12.9 App challenge · 9 min

Load scenario 10, "The floor as a constraint: MPC against the cascade with the same forecast", default seed. It opens with Delivery control on MPC. Run one day at speed 3600, export the CSV and read `metric_below_floor_s`, `metric_min_pout_bar` and `metric_cost_usd` in the row at `time_s` = 86400. Then Reset, set Delivery control to Cascade before Run, and repeat. Then Reset, set Delivery pressure setpoint to 46.4 bar before Run, and repeat.

(a) Tabulate the three runs. *Checkpoint:* MPC: 10.4 min, 44.73 bar, $40,623.

(b) In the first two files, read `discharge_sp_bar` at 17:40 (`time_s` = 63600) and `pack_t` at 17:00 and 18:00. What does the MPC do differently?

(c) Reload the scenario, set MPC preview to Off (an Advanced control on the delivery card) before Run and run the day again. Where does the MPC's advantage come from, and what would you tell a manager who asks whether this line needs it?

12.10 Application area · 4 min

In scenario 7 the demand is 8 % above its profile, so the evening peak is $1.08 \times 285\,\mathrm{kg/s}$.

(a) With the inlet at the 70 bar clamp and the delivery at the 45 bar floor, what steady flow can the twenty-cell line carry, from Equation (12.3) with $0.95K$? *Checkpoint:* 307.8 kg/s.

(b) Packing is the control room's own answer. From the chapter, what does case 7's 4 bar pack cost, what does it do to the leak detector, and how does it fare on other days?

12.11 Application area · 3 min

A control room's detector runs at a 1 % threshold and calls the field crew out every time a meter drifts. The supervisor proposes 2 %. Using cases 8 and 9, state what the change gains and loses, and name two ways to keep the sensitivity without the false alarms.

12.12 Review the agent · 5 min

An AI assistant writes about the challenge: "Set Delivery control to MPC and raise the Leak alarm threshold to 2.5 %. As loaded, the MPC keeps the delivery above the floor all day, so its only failure is the seven false alarms from the 1 % meter error, and the higher threshold removes them. Keep the 48 bar setpoint for margin." Find the flaw, and say how you would check it in the app.

CHAPTER 13

Series Pump Stations, Batches and Surge on a Refined Products Pipeline

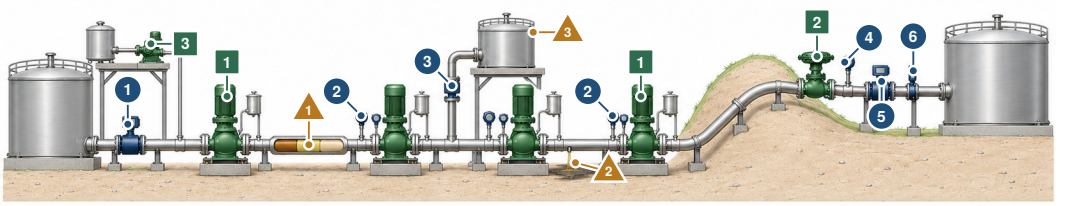

The controller moves 1 pump speed at the four stations 2 the far terminal's valve, holding back pressure 3 drag reducer dose

The controller reads 1 flow leaving the origin 2 pressures before and after each station 3 flow taken by the side terminal 4 pressure at the far terminal's valve 5 flow into the far terminal 6 fuel density at the far terminal

What disturbs it 1 a change of fuel 2 a leak 3 what the side terminal takes

The controller sets the flow at the first station on the left, the other three pass it on, and the valve at the far terminal on the right holds the pressure that keeps the hill full. Illustration, not to scale.

IN PLAIN TERMS One Pipe, Four Pumps, Several Fuels

The main idea. A long pipe carries gasoline, jet fuel and diesel one batch after another. Four pumping stations along the way push them. The same flow passes every station, so only the first may set it, and the others must pass on what arrives.

The challenge. If a valve at the far end shuts, a jump in pressure races back up the pipe in seconds. The pumps keep squeezing fuel in behind it. The leak alarm compares the fuel sent with the fuel received, and a change of fuel can fool it.

Why it matters. The line must stay below its pressure limit, keep each fuel clean and find a real leak fast.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History boxes in Section 13.2, the Did you know? boxes in Sections 13.3 and 13.5, the Try it boxes in Section 13.5, and Key points in Section 13.8. No equations.
Study	50 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 13.5 with Figure 13.4, Section 13.6 and Section 13.7 with its Open challenge boxes.

The control question One flow passes every pump station of a refined products line, while its batches move at walking pace and its pressure news at the speed of sound. Which controller may set that flow, what must the other stations hold, and what protects the line when a valve closes, a station trips, a batch changes or a hill near the terminal runs slack?

Learning objectives

- Explain why only one station of a series line can set the flow, and how low selects of suction and discharge loops protect the rest.
- Compute the storage of a liquid line per bar, the drift of a suction under a flow mismatch and the Joukowsky rise of a closure.
- Tell the pressure front of a valve closure from the packing that follows it, and say which of the two the stroke time controls.
- Compute the friction gradient with and without a drag reducer, and explain why its gain in flow arrives over hours.
- Judge the settings of a volume-balance leak detector against false alarms at a batch change and the time to alarm a leak.

Prerequisites Sections 1.4.5 (line pack and water hammer), 1.4.8 (selectors in the external-reset form) and 1.4.9 (balance leak detection), with Section 1.6.2 for the terminal loop. Section 1.6.3 is needed only for the back-pressure MPC.

The app https://apmonitor.com/apps/products-pipeline/

See also Volume 4, Chapter 3 for min–max limiter selection on one fuel flow and Volume 4, Chapter 2 for a protective override. Volume 6, Chapter 7 for the choice of the loop that sets the throughput. Volume 8, Chapter 13 for a diversion valve that acts on what its instruments report.

13.1 Introduction

A products pipeline is a long column of nearly incompressible liquid moved by pumps that cannot see one another. In the generic line simulated here, 1000 m^3/h of diesel crosses 300 km of 16-inch pipe at 2.3 m/s, so a batch takes about 40 h to reach the far terminal while a pressure change arrives in 4.4 min. Shutting the terminal valve in 5 s raises the pressure there 18.5 bar in ten seconds, and the stations behind it go on packing the line until it sits at about 87 bar g, close to its maximum allowable operating pressure (MAOP). The hydraulics of such lines, from the friction gradient to the stations' heads, are textbook material [1], and the pressure wave of a sudden closure has been analyzed for more than a century [2].

Two common beliefs about such a line are wrong, and the simulator shows why. The first is that each pump station on a line can hold its own discharge pressure. One flow passes all four stations, so a station that holds its discharge pumps whatever its discharge demands, and the pipe upstream of it fills or empties. In case 3, with the override selectors off, a 10 % cut in the lead flow drains station 2's suction to the low-suction alarm in 4.6 min and trips all four stations within 10.2 min. The second belief is that a slow valve makes a closure safe. A stroke of 120 s instead of 5 s lowers the steepest ten-second rise from 17.0 to 8.3 bar, while the blocked-in line still ends near 87 bar g. Both corrections belong to a wider class, units in series that share one throughput and are protected by constraint overrides, which includes booster stations on water mains, compressor trains, cascades of reactors and canal pools. Override and selector control of constrained loops is treated in general form in [3], and the station, SCADA and monitoring layers of pipelines in particular in [4].

The published material divides this problem into pieces. Batch scheduling is posed as mixed-integer optimization of volumes and timing [5], interface mixing is estimated without control [6, 7], leak detection is reviewed method by method [8, 9], and slack flow is an operating topic of its own [10]. This chapter puts the pieces on one simulated line, on which the station overrides, the terminal valve, the batch cut, the drag reducer, the back pressure over a hill and a balance leak detector can be changed while the line runs and read against a scorecard. The lumped transient model, the

Darcy–Weisbach friction, the Austin–Palfrey mixing form, the override structure and the linear MPC are all established methods. The model has 70 states, and its volume balance closes to 2.3×10^{-13} of the throughput through batch changes, leaks and slack. A measurement boundary keeps every controller, the tracker and the detector away from the simulator's truth, so the detector's false alarm at a batch change can be split into a meter shift and an unmeasured change of line pack.

Section 13.2 follows the domain, from batch interfaces and water hammer to drag reducers and leak monitoring. Section 13.3 sets out the line as modeled: thirty lumped segments, friction and the drag reducer, the pump stations and terminals, the batches, the declared slack treatment and the instruments. Section 13.4 describes the control layers: the station overrides, the terminal back-pressure loop, the MPC that sets only that back pressure, the tracker, the cut and the leak detector. Section 13.5 reports the nine cases, each graded by a completion rule, from one flow through four pumps to a full day on the line. Section 13.6 sets the hardware of a real line beside the model, Section 13.7 states the limitations, Section 13.8 draws the conclusions, and the exercises of Section 13.9 put the cases to work.

13.2 Batches, Surge and Leak Monitoring on Liquid Lines

Refined products have long shared pipe. Instead of one line per fuel, a products line carries a sequence of batches, and every interface between two products leaves a mixed zone that must be cut out at the receiving terminal and sent to a transmix tank. How the length of that mixed zone came to be predicted is told in the History box below.

> **HISTORY 1954 to 1963: How Long Is a Mixed Zone?**
>
> Once several products shared one pipe, the length of the mixed zone between two batches became a design question. In 1954 G. I. Taylor treated dispersion in turbulent pipe flow with a virtual coefficient of diffusion and tested it on brine injected into a pipe of three-eighths of an inch. When one fluid is forced into a pipe full of another with which it can mix, his theory gave the length over which the interface spreads, and it agreed with tests made on long pipelines in America [11]. In 1963 Austin and Palfrey drew new formulae from pipes of one-eighth of an inch to 40 inches and from a few feet to hundreds of miles long. They found two regions of mixing in turbulent flow, with mixing growing fast below a critical Reynolds number, and offered the work for the design of products pipelines [6].

The pressure side of the problem has an older history. Joukowsky's water hammer relation, a rise of $\rho a \Delta V$ when a column moving at ΔV is stopped by a wave of speed a (Equation (1.9)), and the method-of-characteristics calculations that followed it are reviewed in [2]. Liquid lines over high ground meet the opposite extreme as well. Column separation, the vapor cavity that opens where a transient pulls the pressure down to the vapor pressure, has flow regimes of its own [12]. Its steady counterpart, slack flow downhill of a summit at low rates, is avoided in operation by holding back pressure [10].

Capacity came to be bought chemically, as the second History box tells.

Table 13.1 Variables and Their Roles, Valued at the Design Point

Variable	Role	Design	Range
Q_{in}	CV, lead flow (origin meter)	1000 m³/h	SP 200–1400
$p_{s,k}$	CV, suctions S2–S4	8 bar g	SP 2–40
$p_{d,k}$	CV (override), discharges	76.9–60.1 bar g	limit 85, 75
p_B	CV, Terminal B back pressure	5 bar g	SP 1–40
s_k	MV, pump speeds	99.8–83.7 %	0–105
x_B	MV, Terminal B valve	47.3 %	0–100
$p_{B,SP}$	MV of the MPC	5 bar g	3–40
c	MV, drag reducer dose	0 ppm	0–30
q_A	DV, Terminal A delivery	150 m³/h	0–300
ℓ	DV, leak at km 180	0	0–3 %
—	DV, product entering	diesel	3 products
—	DV, station 3 trip, B close	—	events
p_i, Q_j	states, 30 + 30	72.0–6.5 bar g	—
x_m	carried, interfaces	—	0–300 km
V_{slack}	carried, slack volume	0	—
r	output, leak residual	0	—

HISTORY 1979: Buying Flow With a Polymer

Dilute solutions of long, randomly coiling molecules reduce friction in turbulent pipe flow, and a 1975 review gathered the experimental evidence for the effect [13]. Drag reduction had been seen in the laboratory for almost 35 years, and its mechanism was still not well understood, when the first commercial use of a polymer additive to raise the flow of a crude oil pipeline began in 1979 on the trunk line across Alaska. Two experiments on that line led to the decision for full-scale use. The gain in pumpability was spent on more flow without exceeding the line's safe pressure limits, and until then drag reducers had not been used routinely in crude oil pipeline operations [14].

Automation moved the stations from local operators to a control center reached through SCADA, with station pressure control and leak detection among its functions [4]. Computational pipeline monitoring grew up in the same control rooms. Its simplest form is the compensated volume balance, metered inflow less metered outflow less the change of line pack over a window [8] (Section 1.4.9), while transient-model methods compare the measurements with a running hydraulic simulation [9, 15]. API Recommended Practice 1130 frames how such a detector's sensitivity, reliability and response time are stated together. Model predictive control entered process plants in the same decades [16], and on the simulated line of this chapter it sits above the station overrides rather than in place of them.

13.3 The Line as Modeled

The simulated line is generic and has no proprietary source. It runs 300 km from an origin tank farm to Terminal B, in NPS 16 steel of 392.1 mm inside diameter and 45 µm roughness, with pump stations S1 to S4 at km 0, 70, 140 and 210, as shown in Figure 13.1. Terminal A takes a delivery of 150 m³/h from station 3's suction header at km 140. The ground climbs from 20 m at the origin to 220 m at station 4, rises to a hill of 440 m at km 250, and falls to the terminal at 160 m. At 1000 m³/h of diesel with every intermediate suction at 8 bar g, the stations discharge at 76.9, 75.7, 59.6 and 60.1 bar g at speeds of 99.8, 96.2, 83.7 and 94.0 %, draw 7.98 MW, and deliver 851.8 m³/h to Terminal B at 5 bar g.

Each variable carries the role listed in Table 13.1. Seventy states are integrated: thirty segment pressures, thirty face flows, four pump speeds and six ledgers (the volumes through the origin, the two terminals, the leak and the relief valves, and the pumping energy) that let conservation be checked exactly. Batches, the drag reducer and any slack volume are carried beside the states and set the segment properties after each step. The clock is in minutes.

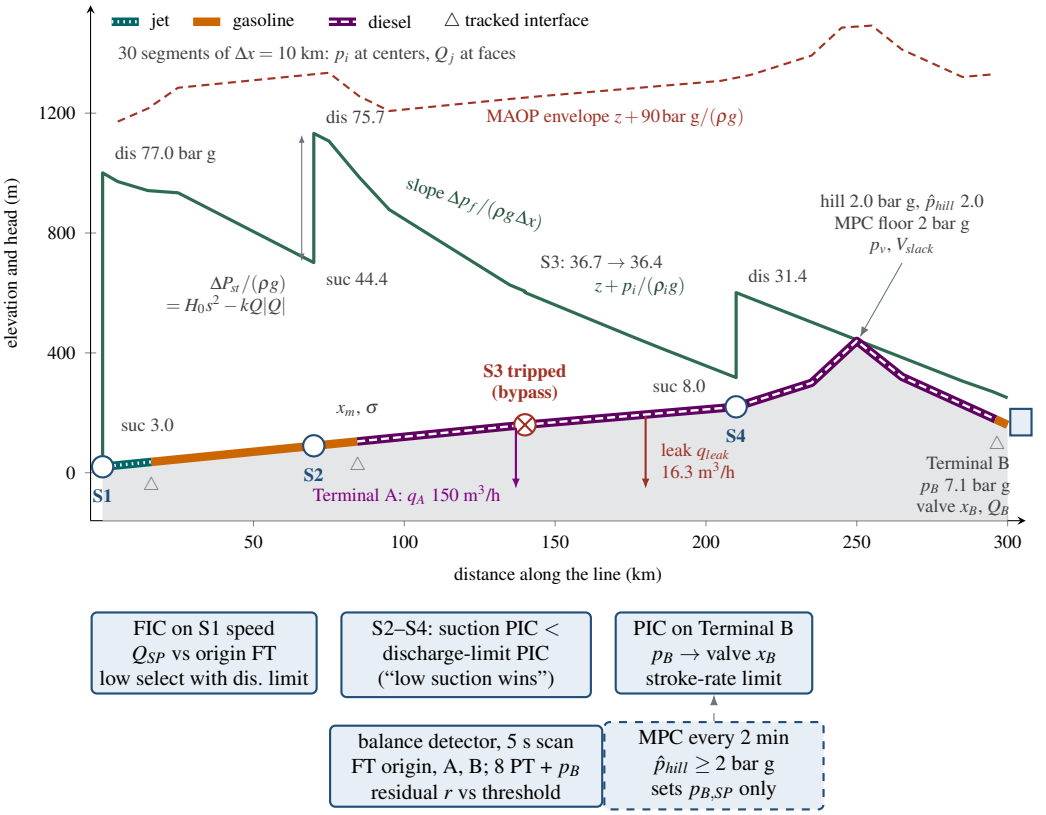

Figure 13.1 The line as modeled, redrawn from the Profile view of the challenge with the suggested settings at $t = 1000$ min. The ground $z(x)$ carries the pipe painted by batch, jet fuel (dotted band) entering behind gasoline (plain band), diesel (dashed band) and gasoline at Terminal B, with the tracked interfaces x_m marked by triangles. The very thick solid green hydraulic gradient $z + p_i/(\rho_i g)$ falls at the friction slope and jumps by the pump head at each running station. Station 3 has tripped and passes the flow through its bypass. The dashed MAOP envelope is the ground plus 90 bar g of head. Terminal A delivers q_A, the 2 % leak at km 180 loses q_{leak}, Terminal B holds the back pressure p_B with its valve x_B, and the hill holds 2.0 bar g against the MPC's 2 bar g floor. Station pressures are true values in bar g. The two rows below show the control layer, with the dashed MPC block setting only the setpoint of Terminal B's PIC. Not to scale across the pipe.

13.3.1 Thirty Lumped Segments

The line is divided into thirty segments of $\Delta x = 10$ km in the lumped inertia and compressibility form. Each segment holds a pressure at its center and each face between two segments carries a flow, as shown in Equation (13.1).

$$\begin{aligned} \frac{\mathrm{d}p_i}{\mathrm{d}t} &= \frac{\rho_i a_i^2}{A\,\Delta x}\left(Q_i - Q_{i+1} - q_{ext,i}\right) \\ \frac{\mathrm{d}Q_j}{\mathrm{d}t} &= \frac{A}{\rho_j\,\Delta x}\left(p_{j-1} - p_j - \Delta p_{f,j} - \rho_j g\,\Delta z_j + \Delta P_{st,j}\right) \end{aligned} \tag{13.1}$$

where A is the bore area, $\rho a^2 = K$ is the bulk modulus that the steel pipe presents to each product, q_{ext} collects Terminal A, the leak and the relief valves, Δz_j is the rise between neighboring centers, and ΔP_{st} is a pump's pressure rise on the face where a station sits. Face 0 is station 1, fed by the origin booster at 3 bar g through half a segment. Pressures are gauge throughout. A segment holding two

products averages ρ and $1/K$ by volume, and a face takes the geometric mean of the two viscosities. Each 10 km of diesel stores 0.113 m^3 per bar, so a 100 m^3/h mismatch between what station 1 sends and station 2 takes moves station 2's suction by 2.1 bar a minute. A liquid line stores almost nothing, as Section 1.4.5 shows for the oil line of the running example.

The scheme was timed on the app's own derivatives. On the line at rest a 0.1 bar step crosses from km 45 to km 245 at 1110.5 m/s in diesel against the declared 1130 m/s, 1093.5 m/s in jet fuel against 1110 and 1014.5 m/s in gasoline against 1030, the same at half and a quarter of the step. Friction is strong on this line. The front of a closure loses most of its height as it travels, so in case 2 only about 7.9 bar of the 18.6 bar front remain 75 km upstream.

13.3.2 Friction and the Drag Reducer

Friction follows Darcy–Weisbach with the Swamee–Jain explicit form of the Colebrook factor, laminar below a Reynolds number of 2300 and blended up to 4000. A polymer additive scales it down in turbulent flow, as shown in Equation (13.2).

$$\begin{aligned} \Delta p_f &= f(1-DR)\frac{\Delta x}{D}\frac{\rho V|V|}{2}, \quad DR = DR_{max}\frac{c}{c+c_{1/2}} \\ \frac{1}{\sqrt{f}} &= -2\log_{10}\!\left(\frac{\varepsilon}{3.7D}+\frac{5.74}{Re^{0.9}}\right) \end{aligned} \tag{13.2}$$

where $V = Q/A$ and c is the local additive concentration. Diesel at 1000 m^3/h runs at $Re = 257\,717$ with $f = 0.0159$ and loses 0.902 bar per km, about 11 m of head, against 0.596 bar/km at 800 m^3/h. The additive is injected at every running station at the dose c, carried with the flow at the liquid's speed, and sheared to nothing in the next station's pumps. At $DR_{max} = 60\,\%$ and $c_{1/2} = 8$ ppm, 10 ppm removes a third of the friction. Both constants are a declared teaching curve rather than a fit to an additive [13, 14].

13.3.3 Pump Stations, Terminals, the Leak and Relief

Each station is one variable-speed pump on a quadratic curve, with a check-valved bypass that carries the flow when the pump is stopped, as shown in Equation (13.3).

$$\begin{aligned} \Delta P_{st} &= \max\!\left(\rho g\,(H_0 s^2 - k\,Q|Q|),\; -K_{byp}Q^2\right) \\ \frac{\mathrm{d}s}{\mathrm{d}t} &= \frac{s_{cmd}-s}{\tau_s}, \qquad P = \sum_k \frac{Q_k\,\Delta P_{st,k}}{\eta} \end{aligned} \tag{13.3}$$

where s is the speed fraction, $H_0 = 1050$ m at stations 1 to 3 and 780 m at station 4, $\tau_s = 3$ s (20 s for a coast-down after a trip), the command ramps at 2 % per second, and $\eta = 0.8$. The bypass loses 1 bar at 1000 m^3/h, and check valves forbid reverse flow through any station. Terminal B's delivery valve is an orifice that passes 1200 m^3/h wide open across 2 bar into a tank at 0.5 bar g, so $Q_B = C_B x_B \sqrt{p_B - 0.5\,\mathrm{bar\,g}}$, with p_B the pressure at the valve inlet, extrapolated from the last segment center by half a segment of friction and the 23 m drop to the terminal. Terminal A delivers its nomination through its own flow controller, ramping 10 m^3/h per minute, and closes as its header falls from 2 to 1 bar g. A leak is a fixed opening split between the two nearest segment centers, $q_{leak} = \ell\,(1000\,\mathrm{m^3/h})\sqrt{p/30\,\mathrm{bar\,g}}$, so a 1 % leak at km 180 passes 10.0 m^3/h at the design pressure of 30.1 bar g there. Relief valves at each station discharge and at Terminal B open at the MAOP and pass 900 m^3/h at 4 % above it.

13.3.4 Batches and Their Mixed Zones

An interface moves with the local flow, and its mixed zone grows with the square root of the distance traveled, as shown in Equation (13.4).

$$\frac{\mathrm{d}x_m}{\mathrm{d}t} = \frac{Q(x_m)}{A}, \qquad c_{behind}(x) = \tfrac{1}{2}\,\mathrm{erfc}\!\left(\frac{x - x_m}{\sqrt{2}\,\sigma}\right)$$
$$4.65\,\sigma = C_m\sqrt{D\,(x_m - x_{m,0})}\,Re^{-0.1} \tag{13.4}$$

where $Q(x_m)$ is interpolated between faces, $C_m = 11.75$ is the turbulent-flow constant of the Austin–Palfrey form [6], and 4.65σ is the length from 1 to 99 %. For gasoline behind diesel at 850 m³/h the zone reaches 1079 m after 300 km, 130 m³ of line. The products are generic teaching values at 15 °C, given as density, viscosity and wave speed: gasoline 740 kg/m³, 0.6 cSt and 1030 m/s, jet fuel 800, 1.5 and 1110, and diesel 840, 3.5 and 1130. A segment whose product changes keeps its liquid volume referred to zero gauge, $A\Delta x(1 + p/K)$, so its pressure is scaled by K_{new}/K_{old} and the flows absorb the difference. Without that rule the line gained liquid wherever an interface moved under pressure, a defect that a technical review of the app found and removed.

13.3.5 Slack Line, a Declared Teaching Treatment

The hill at km 250 stands 280 m above Terminal B, 23.1 bar of diesel. At low flow the friction no longer holds its summit above the vapor pressure, and the treatment of Equation (13.5) applies after each step.

$$p_i < p_v \Rightarrow V_{slack,i} \mathrel{+}= \frac{(p_v - p_i)A\,\Delta x}{K_i}, \quad p_i = p_v$$
$$V_{slack,i} > 0 \Rightarrow Q_{i+1} \le Q_i \tag{13.5}$$

where p_v is the vapor pressure, −0.46 bar g for gasoline and −1.0 bar g for jet and diesel. Liquid arriving later fills the slack volume before the pressure rises. This is an outflow clamp, not a model of a partly full pipe [12], and its slack volume is a time-step artifact. In case 6 the volume grows by what the downhill pipe draws within one step before the clamp acts, 13.5 m³/h at the 1 s step, 6.8 at 0.5 s and 3.4 at 0.25 s, so the volume at 90 min falls from 15.3 to 7.7 to 3.9 m³ as the step halves. What does not depend on the step is when the hill reaches the vapor pressure (22.63, 22.62 and 22.60 min) and which back pressure refills it (15 bar g does at every step, 10 bar g does not). The app and this chapter therefore read the slack volume as a sign of slack, never as a prediction of how much vapor a real line would hold.

> **+ DID YOU KNOW? Slack Flow Over a Mountain**
>
> One liquid line of about 800 km, with six pump stations, crossed more than 1000 m of height between a mountain top and a river valley and ran slack there, which made its leak detection inaccurate. A pressure regulating station that kept the line full also made its leak detection more reliable [10].

13.3.6 Instruments

Pressure transmitters at each station's suction and discharge and at Terminal B have a 1 s lag and 0.03 bar of noise. The three flow meters, at the origin and at each terminal, have a 2 s lag, 0.1 % noise and a meter factor proved on diesel, so each reads gasoline high by the declared error of 0.25 % and jet fuel by half of it. A densitometer at Terminal B reads the composition at the valve through a 30 s sample delay with 0.3 kg/m³ of noise. The noise draws per execution are fixed in number, so two runs that differ in one setting see identical noise sequences.

13.3.7 Start-Up State and Integration

Table 13.2 collects the parameters. Each one is a generic teaching number for a 16-inch line, chosen to be self-consistent and taken from no pipeline. Every case opens on a steady line at its own settings,

Table 13.2 Parameters. A: Adjustable in the Parameters Panel. V: Fixed and Stated in the App. H: Fixed and Not Displayed

Symbol	Meaning	Value	Unit	
L, Δx, D	length, segment, bore	300, 10, 392.1	km, km, mm	V
ε	wall roughness	45	µm	A
a	wave speed scale	100	% of declared	A
MAOP	max. allowable pressure	90	bar g	A
—	high-high discharge trip	MAOP − 2	bar g, 2 s	A
—	low suction alarm, low-low trip	3, 1 (5 s)	bar g	A
—	Terminal B relief set	90	bar g	A
p_{tank}	origin booster	3	bar g	H
H_0	pump shutoff head	1050, 780	m	V
k	pump curve coefficient	1944, 1037	m s^2/m^6	H
K_{byp}	stopped-station bypass	1	bar at 1000 m^3/h	H
η	pump efficiency	0.8	—	V
τ_s	drive lag, coast-down	3, 20	s	V
C_B	Terminal B valve, wide open	1200	m^3/h at 2 bar	H
DR_{max}, $c_{1/2}$	drag reducer curve	60, 8	%, ppm	A
C_m	Austin–Palfrey constant	11.75	—	A
—	meter error on gasoline	0.25	%	A
—	densitometer delay	0.5	min	A
x_{leak}	leak location	180	km	A
ρ, ν, a	products, Section 13.3.4	—	—	V

found by marching from Terminal B's pressure back to the origin with every intermediate suction on its setpoint, and if the stations cannot carry the requested flow within their speed range and discharge limits a bisection lowers the flow until they can. Integration uses fourth-order Runge–Kutta steps of 1 s, inside the engine of Section 1.6.1. Halving the step moves the largest ten-second rise of case 2 from 16.89 to 16.53 bar (−2.1 %), because the valve's 5 s stroke from its 47 % opening is resolved in only a few steps. At half and a quarter of the step the shipped challenge still fails and the suggested one still passes. The truth volume balance closes to 2.3×10^{-13} of the volume pumped over the challenge's day.

13.3.8 Assumptions

The line is isothermal at 15 °C, with no exchange with the ground and no thermal expansion, which removes the largest source of false alarms on real lines. A wave front is spread over about one segment, so the model reproduces the wave speed and the Joukowsky rise but not the shape of a sharp front, and it says nothing about column separation during the front [12]. Each station is one pump at constant efficiency, with no multi-pump sequencing, start-up transient or cavitation below the suction trip. Products are constant in density, viscosity, bulk modulus and vapor pressure. The slack line is the treatment above. The safety numbers, the MAOP of 90 bar g, discharge limits of 85 and 75 bar g, the trips and the 3 bar g suction alarm, are teaching values. Real ones come from the pipe's design under ASME B31.4 and from the pumps' net positive suction head requirements, and real detector settings from API RP 1130 and the operator's own testing.

13.4 The Control Layers as Implemented

The controllers execute every 1 s of simulated time, the SCADA scan and the leak detector every 5 s, and the MPC every 2 min. Each physics step runs the controllers, the integrator, the post-step corrections (check valves, slack, interfaces, additive transport, trips, alarms) and the outputs. Everything the controllers, the tracker and the detector read passes through a measurement packet built from the instruments above, and the simulator's truth reaches only the scene's teaching overlay, the scorecard and the exported columns prefixed `truth_`.

13.4.1 Station Overrides

Station 1's flow loop (FIC) holds the origin meter at the line setpoint. At stations 2 to 4 a suction loop holds 8 bar g, and every station has a discharge-limit loop. With the selectors on, each pump takes the lower of its two speed requests, which is the arrangement known as "low suction wins" [3, 4]. With them off, station 1 runs on flow alone and stations 2 to 4 on discharge alone. The override loops are written in the external-reset form of Equation (13.6).

$$\begin{aligned} u_k &= K_c\, e_k + F_{k-1}, \qquad s_{cmd} = \min(u_{main},\, u_{dis}) \\ F_k &= F_{k-1} + \left(1 - \mathrm{e}^{-\Delta t/\tau_I}\right)\left(s_{applied} - F_{k-1}\right) \end{aligned} \tag{13.6}$$

where e is the error in the loop's own direction, u_{main} is the flow or suction request, $s_{applied}$ is the clamped speed the drive was given, and each loop's reset filter F follows that applied speed. This is the external-reset form of Equation (1.14), and Section 1.4.8 explains why neither loop winds up and why the transfer between them, or from Manual to Auto, is bumpless. In Manual each filter is reset to make its output match the operator's speed. The defaults are $K_c = 0.02$ % per m³/h with $\tau_I = 30$ s for the flow loop, 1.0 % per bar with 60 s for the suction loops, and 4.0 % per bar with 15 s for the discharge loops, with speeds limited to 0 to 105 %. Two tracking PID objects in a low select were tried first and ratcheted the speed downward on measurement noise, which is why this form was chosen. The same selection of the most restrictive request, over a speed governor and its limiters on one fuel flow, protects the jet engine of Volume 4, Chapter 3. Choosing the one loop that sets the throughput is the plantwide question that Volume 6, Chapter 7 studies on a reactor and its recycle.

Pause. Try Exercise 13.1 before reading on.

13.4.2 Terminal B Back Pressure

The back-pressure loop (PIC) moves Terminal B's valve with the velocity-form PID of Section 1.6.2, with no derivative term and back-calculation on the integral with a tracking time equal to τ_I, as shown in Equation (13.7).

$$\begin{aligned} e_k &= -(p_{B,SP} - p_{B,k}), \qquad S_k = S_{k-1} + K_c(e_k - e_{k-1}) \\ \hat{x}_k &= x_0 + S_k + I_{k-1}, \qquad x_k^* = \min(100, \max(0, \hat{x}_k)) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{x_k^* - \hat{x}_k}{\tau_I}\right]\Delta t \\ x_{B,k} &= x_{B,k-1} + \mathrm{clip}\left(x_k^* - x_{B,k-1},\, \pm\tfrac{100\,\Delta t}{t_s}\right) \end{aligned} \tag{13.7}$$

where the minus sign makes a pressure above setpoint open the valve, S is the unclamped running sum of Equation (1.17), x_0 the output at the start, $K_c = 2$ % per bar and $\tau_I = 60$ s. The stroke time t_s, 60 s by default, limits every move of the valve, an emergency close included.

13.4.3 Model Predictive Control of the Back Pressure

The MPC chooses one setpoint, Terminal B's back pressure, every 2 min. Its model is the control room's steady hydraulic profile on the tracked batches, rebuilt at each execution with any tripped station as a bypass. The profile uses the declared line of Table 13.2, with its 45 µm wall roughness, and never the plant's parameters, so a roughness changed in the app is model mismatch. In case 6 at 200 µm the MPC reads the hill 0.66 bar low at 120 min and holds 0.6 bar more back pressure than an MPC given the plant's roughness, and at 10 µm it reads the hill 0.07 bar high, which leaves the hill at 1.86 bar g, under its 2 bar g floor. From the measured terminal pressure it infers the lowest pressure $\hat{p}_{hill}$ over the five segment centers from km 235 to 275, and a finite difference of 1 bar gives the gain g of the hill on the back pressure. The profile is built at the smaller of the origin meter's flow and Terminal B's flow plus Terminal A's delivery, so a leak upstream of the hill cannot make the estimate

optimistic. The prediction model and the quadratic program, with a horizon of thirty 2-min steps grouped in graduated blocks, are given in Equation (13.8).

$$\begin{aligned} &\tau_h \dot{x} = g\,(u - u_0) - x, \qquad \hat{p}_{hill} = \hat{p}_{hill,0} + x \\ &\min \textstyle\sum_k w_E \frac{Q_{in}}{1000} u_k + w_{\Delta u} \Delta u_k^2 + 100 w_s s_k^2 + 10 w_s s_k \\ &\text{s.t. } \hat{p}_{hill,k} \geq 2 \text{ bar g} - s_k, \quad 3 \leq u_k \leq 40, \quad |\Delta u_k| \leq 4 \end{aligned} \tag{13.8}$$

where u is the back-pressure setpoint in bar g, $\tau_h = 3$ min, and $w_E = w_{\Delta u} = w_s = 1$. The linear cost prices back pressure in proportion to the flow that must be pumped against it, about 0.03 MW per bar at 1000 m³/h, paid by station 4 on the 850 m³/h past Terminal A. The shared linear MPC module of Section 1.6.3 solves it with the dual active-set algorithm of Goldfarb and Idnani [17] and sends only the first move to the loop below. A solve that is not optimal and finite is rejected and the setpoint stays where it was, with the station overrides running underneath in every case. Over the challenge's day the MPC solves 720 times with no rejection.

13.4.4 Tracker, Cut and Leak Detector

The tracker moves its own copy of each interface at the metered flow over the bore area, the origin meter upstream of km 140 and the origin meter less Terminal A's beyond it, with the line taken as rigid. It confirms and corrects an interface when the densitometer passes the midpoint of the density change. The cut sends the delivery from the ahead product's tank to a transmix tank when the density has moved a set share of the way, 10 % by default, and on to the next product's tank at a second share, 90 %. Like the flow diversion valve of Volume 8, Chapter 13, the cut acts on what its instrument reports, here a densitometer read through a sample delay. The leak detector is a compensated volume balance [8], as shown in Equation (13.9), the form of Equation (1.15) written in volumes.

$$\begin{aligned} r &= \frac{1}{W}\Big[\int_{t-W}^{t} (\hat{Q}_{in} - \hat{Q}_A - \hat{Q}_B)\,dt - \Delta \hat{V}_{pack}\Big] \\ \hat{V}_{pack} &= \textstyle\sum_i A\,\Delta x\,\big(1 + \hat{p}_i / \hat{K}_i\big) \end{aligned} \tag{13.9}$$

where $\hat{p}_i$ is interpolated along each leg in piezometric head from the eight station pressures and Terminal B's, and $\hat{K}_i$ comes from the tracked batches. An alarm is raised when r has stayed above the threshold for the persistence time, judged only once a full window has elapsed, and cleared below 80 % of the threshold. The defaults are 6 m³/h over a 15 min window with 2 min of persistence. Chapter 12 watches a gas transmission line with a detector of the same kind, where the pack is large and no instrument measures it.

13.4.5 Trips, Flags and Three Separate Models

Protection sits outside the controllers and reads the truth. It trips a station whose discharge stays within 2 bar of the MAOP for 2 s, or whose suction stays below 1 bar g for 5 s, and the trip latches until the operator restarts it. Status flags mark MAOP exceeded, low suction, slack line, an open relief and surge, defined as any segment rising more than 10 bar within 10 s. Three models are kept apart. The plant is Section 13.3. The identified model is what a learner reads from the trends, such as the friction gradient or the storage of a leg. The controller's model is the steady profile on the tracked batches, which knows neither the leak nor a meter's product shift nor the transient that follows a trip.

13.4.6 Design Choices

Six design choices shape the line and its control. First, the hill sits at km 250, in the last leg, and station 4 at km 210. An earlier layout placed the hill at km 190. A station cannot raise the pressure behind it, so a hill upstream of station 4 could not be held full from it, while in the last leg Terminal B's back pressure is a real handle. The leak stays at km 180. Second, the surge lesson states that a slower stroke removes the front but not the packing, which the measurements of Section 13.5 bear out. Third,

Table 13.3 The Nine Cases

#	Case and events	Teaching point
1	Manual at 800 m^3/h, diesel	the series line by hand
2	Terminal B emergency close at 3 min, 5 s stroke	Joukowsky front, packing
3	Selectors off; flow SP 1000 to 900 m^3/h at 10 min	low suction wins
4	Gasoline at B, diesel 12 km behind	tracking and the cut
5	Flow SP 1250 m^3/h, station 1 at top speed	drag reducer
6	Flow SP 1000 to 500 m^3/h at 10 min	slack line, back pressure
7	Gasoline behind diesel at 10, 1 % leak at 180 min; 3 m^3/h, 5 min, no persistence	detector tuning
8	Station 3 trips at 10 min	a trip under overrides
9	Challenge, 24 h: three batches, trip at 8 h, 2 % leak at 16 h	the whole line

Table 13.4 Case Outcomes Measured on the Engine

Case and condition	Result
1, S4 alone to 94 %	low-low suction trip at 2.32 min
1, origin downstream, 8 min apart	992 m^3/h at 60 min, suctions ≥ 8
2, stroke 5 s	+18.5 bar in 10 s, SURGE at 6 s
2, stroke 120 / 600 s	steepest 10 s rise 8.3 / 4.1 bar
2, end state, any stroke	87.0–87.4 bar g at B
3, selectors off	LOW SUCTION 14.6 min, S2 trip 15.9
3, selectors on	S2 lowest suction 7.05, 899 m^3/h
4, cut 10/90	transmix 73.4, off 2.01 / 0.74 m^3
4, cut 30/70	transmix 29.0, off 7.82 / 3.56 m^3
4, cut 5/95	transmix 92.6, off 1.05 / 0.30 m^3
5, no additive	1063 m^3/h, S1 at 105 %
5, 10 / 5 / 20 ppm at 2 min	1250 at 8 h / 1189 / 1251 at 4 h
6, 5 bar g	SLACK LINE at 25.05 min
6, 10 / 15 / 20 bar g at 40 min	not refilled / 44.2 / 42.0 min
6, MPC from Reset	never slack, 15.8 bar g at 90 min
7, as shipped	first alarm 11.4 min, 247 before the leak
7, persistence 5 min only	quiet, leak in 5.5 min
7, 4 / 6 m^3/h, 15 min, 2 min	quiet, leak in 3.7 / 7.0 min
8, after the trip	781 m^3/h, S1 and S2 at 85 bar g
8, restart at 70 min	above 990 m^3/h at 74.7 min

the MPC sets only the back pressure. In a series line of variable-speed stations the total head at a given flow is fixed by friction, elevation and the terminal pressure, and the suction setpoints only redistribute it, so the energy handles are the flow, the drag reducer, with a transport delay of 8.45 h a leg, and the back pressure the hill needs. Fourth, energy is reported but not judged in the challenge. Fifth, the stations use speed control only, without discharge throttle valves. Sixth, the slack line is the declared outflow clamp of Equation (13.5).

13.5 Nine Cases, Measured

The cases appear in Table 13.3 and the principal measurements in Table 13.4. All numbers come from the engine of Section 1.6.1 run without a browser on the default seed, advanced with the browser's own step, and any move the case's text asks the learner to make is made after the run begins. Seeds 1 to 3 were used to check the margins of cases 7 and 9. Every case loads identically when selected first, chosen after another case, or rerun with Reset.

13.5.1 One Flow Through Four Pumps

Case 1 starts lined out at 800 m^3/h in Manual, with speeds of 83.6, 79.3, 66.8 and 69.9 %. Raised alone to 94 %, station 4 takes its extra flow from the pipe just upstream before anything upstream can send more, and trips on low-low suction at 2.32 min. Raised alone to 96 %, station 2 pulls its own suction to −1.7 bar g and trips at 2.3 min while station 3's suction rises to 20.8 bar g. Raised from the

origin downstream, 8 min apart, to 99.8, 96.2, 83.7 and 94.0 %, the line reaches 992 m³/h by 60 min with no suction below 8 bar g, and the first leg's gradient steepens from 0.596 to 0.889 bar/km.

Case 3 ships the stations on discharge control at their design values and cuts the lead flow to 900 m³/h at 10 min. Station 2 keeps pumping what its discharge needs, the first leg empties at the storage rate above, LOW SUCTION appears at 14.6 min, and station 2 trips on low-low suction at 15.9 min. Stations 3 and 4 follow at 17.3 and 18.6 min as each leg empties, and station 1, pushing into a line that no longer moves, trips on high-high discharge at 88 bar g at 20.2 min. With the selectors on, each suction loop takes its pump as the suction falls below 8 bar g, station 2's lowest suction is 7.05 bar g and the line settles at 899 m³/h. In case 8 the same structure carries a trip of station 3. Stations 2 and then 1 reach their 85 bar g limits, their discharge loops take the pumps, and the line settles at 781 m³/h with station 4 still on its suction and no other trip. Restarted at 70 min, station 3 brings the flow back above 990 m³/h at 74.7 min.

> **TRY IT IN THE APP** **8 · Station 3 trips**
>
> Load the case and press Run. At 10 min station 3's drive trips. Watch the Flows trend: the line settles at about 781 m³/h with no other trip, while stations 1 and 2 hold their discharges at 85 bar g. Once it has settled, press Restart tripped stations, below the cards. Restarted at 70 min, station 3 brings the flow back above 990 m³/h at 74.7 min.

13.5.2 A Front That Wears Down, and Packing That Does Not

Case 2 closes Terminal B at 3 min with a 5 s stroke, as shown in Figure 13.2. Diesel at 840 kg/m³ stopping 1.96 m/s at 1130 m/s gives a Joukowsky rise of 18.6 bar, and the terminal rises 18.5 bar in the first 10 s. SURGE is flagged 6 s after the close begins. On the 5 s trends, Terminal B and station 4's discharge first stand 2 bar above their earlier values 5 s and 90 s after the close, an apparent 1060 m/s over the 90 km between them, slower than the declared speed because friction wears the front down as it travels [2]. Behind the front the stations go on pumping into a line that no longer delivers, and friction's head turns into pressure. The terminal is 38.4 bar above its start at 1 min and 80.3 bar at 5 min, peaks at 89.7 bar g about 9 min after the close, and sits near 87.4 bar g at 30 min once the discharge loops hold every station at its limit. A 120 s stroke lowers the steepest ten-second rise to 8.3 bar and a 600 s stroke to 4.1 bar, neither raising SURGE, and both end within 0.4 bar of the fast closure. These steepest rises are taken over every segment center, the measure behind the SURGE flag. The stroke time controls the front and not the packing, which is why a terminal closure on a real line is coordinated with the stations rather than left to the valve.

Pause. Try Exercise 13.2 before reading on.

13.5.3 Interfaces, Additive and the Hill

In case 4 the tracker predicts the interface at Terminal B at 102.3 min and the true interface arrives at 102.4 min. When the densitometer confirms it, the tracked position is 0.071 km, or 8.5 m³, off. The cut trades transmix against contamination, as Table 13.4 lists. Cutting at 10 and 90 % of the density change sends 73.4 m³ to transmix, with 2.01 m³ of diesel in the gasoline tank and 0.74 m³ of gasoline in the diesel tank, and 2 and 98 % raise the transmix to 115.7 m³ with 0.33 and 0.12 m³ off. The gasoline tank always receives more, because the 30 s sample delay, about 7.1 m³ of line at 850 m³/h, makes every switch late.

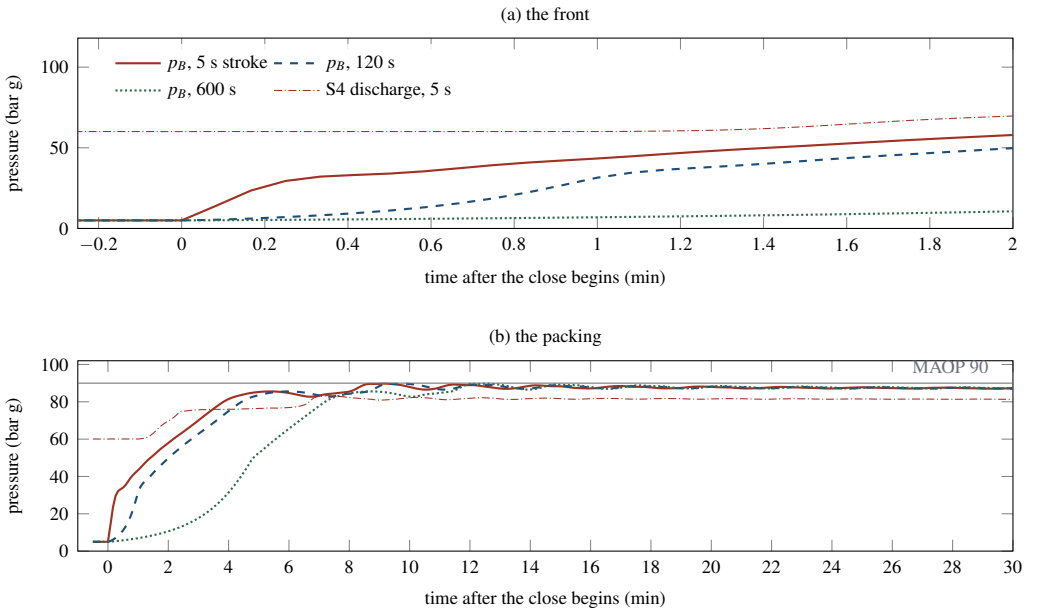

Figure 13.2 Terminal B's valve-inlet pressure after the emergency close of case 2 for three stroke times, 5 s (solid), 120 s (dashed) and 600 s (dotted), with station 4's discharge for the 5 s stroke (dash-dotted), logged every 5 s. The thin gray line is the MAOP. (a) The 5 s stroke raises 18.5 bar in 10 s, and the front reaches station 4 worn down. (b) All three strokes pack the line to the same end, near 87 bar g.

DID YOU KNOW? A Narrow Cut Is Not a Clean Cut

From the simulator (Table 13.4). Cutting case 4's interface at 30 and 70 % of the density change instead of 10 and 90 % sends only 29.0 m^3 to transmix instead of 73.4 m^3, but puts 7.82 m^3 of diesel into the gasoline tank instead of 2.01 m^3.

In case 5 the line delivers 1063 m^3/h with station 1 at its 105 % top speed. Dosed at 10 ppm at 2 min, the flow climbs leg by leg as treated diesel fills each one, to 1103 m^3/h after an hour, 1192 after four and the 1250 m^3/h setpoint after eight. Five ppm levels off near 1189 m^3/h, and 20 ppm reaches the setpoint after four hours. Used to save power at 1000 m^3/h instead, 10 ppm lowers the pumping power from 7.98 to 5.65 MW ten hours after the dose and toward 5.48 MW once every leg is treated.

TRY IT IN THE APP 5 · Drag reducer buys flow

Load the case and press Run. After the run starts, set Drag reducer dose on the Drag reducer card to 10 ppm, then run at 60 or 120 min/s. The line flow climbs leg by leg as treated diesel fills each one: about 1103 m^3/h after an hour, 1192 after four and the 1250 m^3/h setpoint after eight. Run again at 5 ppm: the flow levels off near 1189 m^3/h.

In case 6 the nomination falls to 500 m^3/h at 10 min. The hill reaches the vapor pressure at 22.6 min and SLACK LINE is flagged at 25.05 min. Raising the back pressure at 40 min to 10 bar g does not refill it by 120 min, 13 bar g refills it at 88.7 min, 15 bar g at 44.2 min and 20 bar g at 42.0 min. The MPC, from a Reset, holds the true hill at or above 1.74 bar g after 40 min with 15.8 bar g of back pressure and never lets it go slack, and at 1000 m^3/h it lowers the back pressure to its 3 bar g bound because friction keeps the hill full there. With 5 bar g at the terminal, a cut to 700 m^3/h leaves 1.7 m^3

of slack after two hours and a cut to 750 m³/h none. Those slack volumes carry the step artifact described in Section 13.3, while the onset does not.

13.5.4 What the Balance Can and Cannot Tell

Case 7 ships the detector at 3 m³/h over 5 min with no persistence, sends gasoline in behind the diesel at 10 min and opens a 1 % leak, 9.96 m³/h, at km 180 at 180 min. The first alarm comes at 11.4 min with no leak open, and the alarm comes and goes 247 times before the leak. Two causes add up. Once gasoline passes the origin meter, which is proved on diesel and reads gasoline 0.25 % high, the meters show 2.5 m³/h more entering than leaving. The line pack worked out from eight transmitters also misses part of the change as the lighter, more compressible gasoline fills the first leg, so from 30 min to 180 min the residual averages 3.08 m³/h with a standard deviation of 0.67, the remaining 0.6 m³/h of the mean coming from the unmeasured line pack. The residual sits on the threshold and the noise does the rest. With the meter error set to zero the same detector alarms twice, and without the batch change it stays quiet at all four seeds. Five minutes of persistence alone keeps it quiet at all four seeds and alarms the leak 5.5 min after it opens, and 6 m³/h over 15 min with 2 min of persistence alarms in 7.0 min. Two minutes of persistence at the shipped threshold is quiet at the default seed but not at seeds 1 and 3.

Pause. Try Exercise 13.3 before reading on.

The detector is more sensitive than a real one, because the model has no temperature and temperature is the largest source of false alarms on real lines. Its other false alarms come from transients nobody tells it about. With the default 6 m³/h, 15 min and 2 min, LEAK ALARM with no leak is raised from 20.1 to 28.7 min after case 2's closure, from 22.3 to 44.2 min while case 3's legs empty, from 17.0 to 25.1 min after case 6's rate cut and from 78.8 to 81.1 min after a restart at 60 min in case 8. A leak smaller than a meter's product shift can hide under it. On case 4's line, with gasoline arriving at Terminal B and diesel entering, the residual sits at −1.49 m³/h, and a 0.2 % leak of about 2 m³/h only brings it to +0.43 m³/h. Of 160 detector settings, with windows of 5 to 60 min, persistence of 0 to 5 min and thresholds of 1 to 6 m³/h, the 60 that stay quiet through case 7's batch change all miss that leak for an hour, the trade the monitoring literature describes [8, 9].

13.5.5 A Day on the Line

The challenge runs a day with jet fuel, gasoline and diesel in the line, gasoline entering at 4 h and jet fuel at 14 h, a trip of station 3 at 8 h and a 2 % leak at km 180 at 16 h. It is met with no MAOP exceedance, no suction below 3 bar g, no protective trip, every cut within 100 m³ of transmix and 2 m³ of the other product, and the leak alarmed within 10 min with no false alarm before it. The state of the suggested run at 1000 min is shown in Figure 13.3, and the configurations are compared in Table 13.5. As shipped, without selectors, the suctions run below their floor from 81.9 min, stations 4, 1 and 2 trip on protection after the drive fault at station 3, and the tight detector false-alarms 17 times. The selectors alone turn the trip into a lower flow and stop the protective trips, but at that flow the hill runs slack, up to 12.5 m³ by the app's measure, and with the shipped detector and cut the day carries 24 false alarms and 7.59 m³ off-product. The suggested configuration, selectors on, the MPC, 6 m³/h over 15 min with 2 min of persistence and a cut at 5 and 90 %, passes at the default seed and at seeds 1 to 3, alarming the leak in 8.6 to 8.75 min, and the line carries 727 to 729 m³/h between 600 and 900 min. The MPC's hill estimate is a steady inference and lags a fast transient. At 484 min, four minutes after the trip, it reads 4.51 bar g against a true −0.34 bar g, and the true hill touches −0.46 bar g, the gasoline vapor pressure, and holds at most 0.25 m³ of slack, under the 0.5 m³ that raises SLACK LINE. After the leak the estimate and the truth agree within 0.1 bar, 2.04 against 1.97 bar g at 1000 min. At the fastest playback speed of 300 min/s the day plays in about 5 s.

A fixed back pressure of 10 bar g also passes. Its day uses 102.3 MWh against the MPC's 107.9 MWh, but it pumps 19 324 m³ against 19 749 m³, because with the higher back pressure the line carries 880 to 955 m³/h in the hours before the trip against 942 to 989 m³/h under the MPC. The

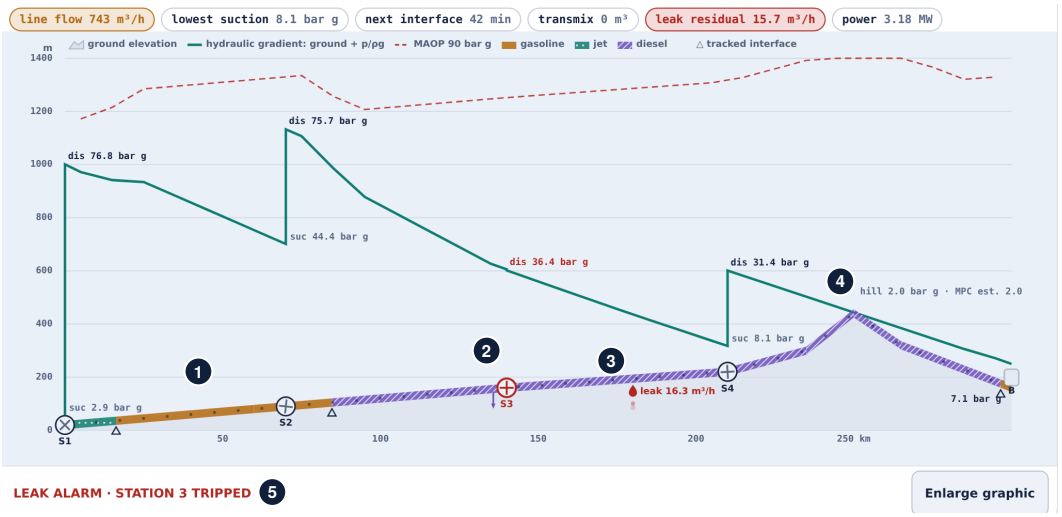

Figure 13.3 The challenge with the suggested settings at $t = 1000$ min (16.67 h) on the Profile view, 40 min after the 2 % leak opened and eight hours after station 3 tripped, cropped to the plant view. The badges above the profile read 743 m^3/h on the origin meter, a lowest suction of 8.1 bar g, the next interface 42 min away by the tracker (50 min in truth, because the leak takes flow the tracker counts), a residual of 15.7 m^3/h and 3.18 MW. (1) Jet fuel entering behind gasoline, diesel and the gasoline reaching Terminal B, with the tracked interfaces marked by triangles. (2) Station 3 in red and with no jump in the gradient, which runs straight through its bypass. (3) The leak at km 180, losing 16.3 m^3/h. (4) The hill at 2.0 bar g with the MPC's estimate beside it (1.97 true and 2.04 inferred). (5) The status strip with LEAK ALARM, raised 8.75 min after the leak.

Table 13.5 Challenge Configurations Scored at 24 h

Configuration	False alarms	Leak (min)	Transmix (m^3)	Off (m^3)	Verdict
shipped	17	—	28	7.48	fails[a]
selectors on only	24	0.9	31	7.59	fails
suggested, shipped detector	13	0.7	84	1.03	fails
suggested, cut 10/90	0	8.75	74	2.08	fails
suggested, PID at 5 bar g[b]	1	8.4	86	0.98	fails
suggested, PID at 10 bar g	0	8.4	85	1.00	passes
suggested (MPC)	0	8.75	84	1.03	passes

[a] Suctions below the floor for 598 min and stations 1, 2 and 4 tripped by protection. [b] The hill's slack refill after the trip raises the false alarm. This row rests on the slack treatment and holds at the 1 and 0.5 s steps but not at 0.25 s, so it is not a converged result.

challenge does not hold the delivered volume fixed, so the energy totals of different configurations describe different days, which is why energy is reported on its tile and not judged. Figure 13.4 sets each configuration against the challenge's limits on alarms and off-product.

13.6 Technology in Practice

Figure 13.5 draws the station overrides, the back-pressure loop, the cut and the leak detector of Section 13.4 with the hardware of a liquid line in place of symbols.

What Measures It

The lead flow loop reads the meter at the origin, and every suction and discharge loop reads a pressure transmitter beside its own pump. The leak detector adds the meters at both terminals and the station

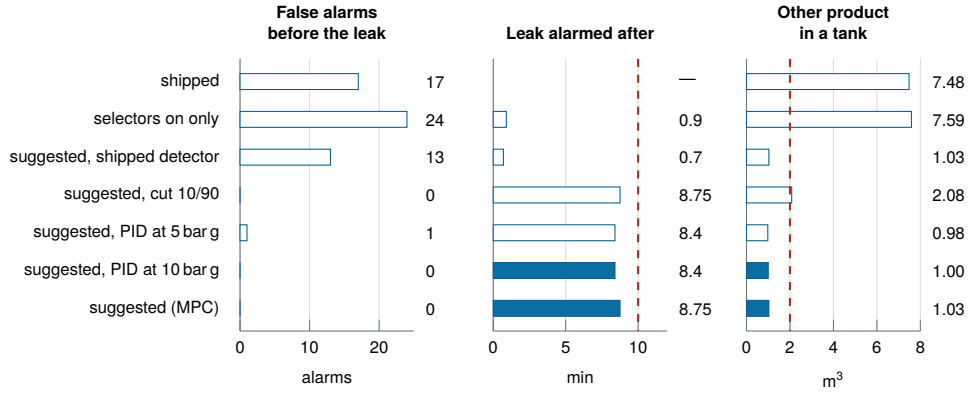

Figure 13.4 The configurations of Table 13.5: false alarms before the leak (left), minutes from the leak to its alarm (middle) and the other product put into a tank by the worst cut (right). Solid bars pass the challenge and open bars fail it. Dashed lines are the limits of 10 min and 2 m^3, a false alarm is not allowed, and the shipped row has no leak time. Notice that the suggested settings with the shipped detector alarm the leak fastest but fail on 13 false alarms, and that the cut at 10 and 90 % fails with 2.08 m^3 against the limit of 2.

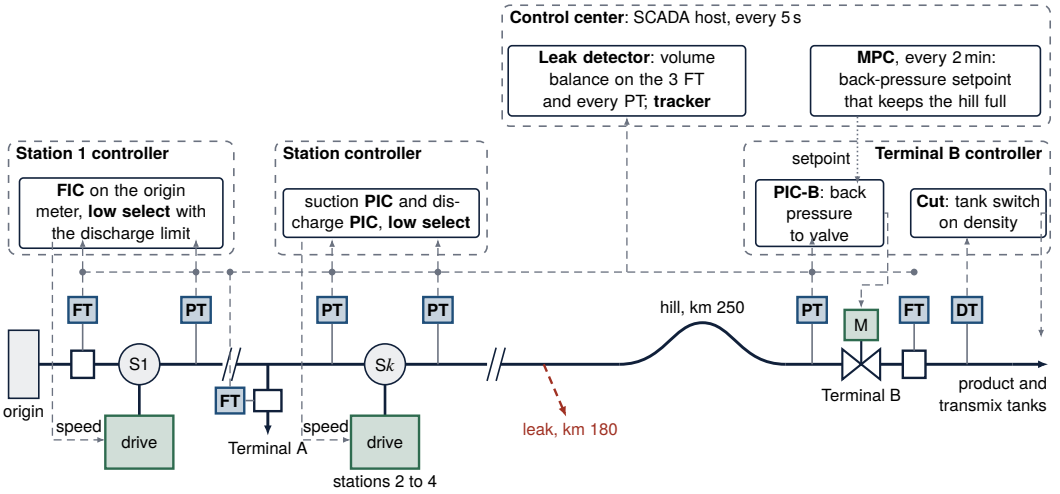

Figure 13.5 The loops of the line with their hardware. Blue boxes are sensing devices with the chapter's tags: the turbine meters (FT) at the origin and at both terminals, a suction and a discharge transmitter (PT) at each station, the transmitter at Terminal B's valve and the densitometer (DT). Green boxes are final elements: the pump drives and Terminal B's motor-operated valve. Dashed arrows are signals, the dotted arrow the back-pressure setpoint that the MPC writes. One station stands for stations 2 to 4. Notice that each pump takes the lower of two speed requests in its own station, while the leak detector in the control center reads the meters and transmitters of the whole line and moves nothing. Not to scale, and not a piping diagram.

and terminal pressures, and the cut reads the densitometer at Terminal B through its sample delay.

TECHNOLOGY Sensor: The Turbine Meter and Its Prover

A turbine meter measures flow with a free-spinning wheel in the stream, much like a small windmill. A pickup coil senses each passing blade, so the pulse frequency is proportional to the volumetric flow, and the turndown commonly exceeds 10 to 1. About 20 pipe diameters of straight pipe upstream and 5 downstream keep swirl away. For high accuracy the pulses per unit volume are found with a piston prover, best connected on site. Diesel lubricates the bearings, but viscosity slows the wheel, most at low flow, and a wheel that coasts after a sudden stop shows flow that is no longer there [18]. The model's meters are proved on diesel and read gasoline 0.25 % high, a declared error that stands for such a change of product.

What Moves It

Each station's loops move one pump's speed, and the back-pressure loop moves Terminal B's valve. The model's MAOP, its discharge trips and its relief valves stand for legal limits. A United States liquid line may not run above its maximum operating pressure, the lowest of limits that include the internal design pressure of the pipe and 80 % of its test pressure. A surge may not exceed 110 % of that limit, and the operator must provide the controls and protective equipment that hold it [19]. Relief valves and other pressure-limiting devices are inspected and tested every calendar year, at intervals of no more than 15 months [20].

TECHNOLOGY Final Element: The Terminal's Motor-Operated Valve

Electric motors have long driven large valves, first as shutoff valves and now in throttling service as well. A worm gear reduces the motor's fast rotation to the slow turn that moves a large valve. A throttling actuator needs a position sensor beside its limit switches, and a hand crank moves the valve if the motor or its power fails [18]. The model's stroke time, 60 s by default, limits every move of Terminal B's valve, an emergency close included, and case 2 shortens it to 5 s. The chapter's surge results turn on that one setting: it decides the front, while the stations decide the packing.

Where the Controller Runs

A pipeline is run through supervisory control and data acquisition: a control server at a control center gathers the field data, remote terminal units or programmable logic controllers at the field sites control the local process, and loops run with cycle times from milliseconds to minutes [21]. The app's controllers execute every 1 s, its SCADA scan and leak detector every 5 s and its MPC every 2 min. The chapter's structure maps onto standard blocks: two PI blocks in external-reset form and a low select in each station's controller, a velocity-form PI at Terminal B, and the tracker, the cut and the volume balance as applications on the control server, with the MPC above them. The server's performance is a safety matter. In June 1999 a 16-inch gasoline line leaked 237,000 gallons that ignited 1.5 hours later and killed three people, after a poorly performing supervisory system had kept the controllers from seeing the abnormal operation develop. The investigators asked for an offline system to test changes to the supervisory database [21].

IN PRACTICE What the Leak Detection Rule Asks

In the United States every hazardous liquid line that carries liquid in a single phase must have an effective leak detection system. The operator evaluates it against the length and size of the line, the product carried, the swiftness of detection, the location of the nearest response personnel and the leak history. A computational pipeline monitoring system must comply with API RP 1130 in its operation, maintenance, testing, record keeping and dispatcher training [22], and a new one is designed to section 4.2 of that practice [23]. Case 7's trade between a quiet detector and a fast alarm is what that evaluation weighs.

What Is Modernizing It

The first of three developments is the transient-model leak detector that the chapter's Conclusion names as the next step, which a 2019 review calls the most sensitive but also the most complex tech-

nique in use. It compares the measured values with those of a running hydraulic simulation [24]. The second is fiber-optic cable laid along the pipe. Leaking hydrocarbon changes the cable's temperature, small leaks can be found, and the timing of the scattered light locates them [24]. The third is security. In May 2021 a ransomware attack on the business systems of an operator of over 5,500 miles of products pipeline, carrying more than 100 million gallons a day, led it to shut the line down, with no evidence that the attack reached the control system [21].

13.7 Limitations

Every size, product property and limit is a generic teaching value, and none was fitted to an operating pipeline. The line is isothermal, so the detector is more sensitive than a real one, whose false-alarm band is set largely by temperature. Its false alarms here come from the meters' product shift and from the line pack that eight transmitters cannot measure at a batch change, and from transients no one reports to it. The lumped scheme smears a front over a segment and represents neither cavitation nor column separation during the front. The slack line is a declared teaching treatment, an outflow clamp whose slack volume is a time-step artifact of about 13.5 m^3/h of growth at the 1 s step that halves with the step. When the hill reaches the vapor pressure, and which back pressure refills it, do not depend on the step, but the slack volume must not be read as a physical prediction, and one row of Table 13.5 depends on it. A real slack line drains the downhill pipe much faster.

> **OPEN CHALLENGE** **How Friction Wears Down a Front**
>
> Hydraulic transients in pipes have been studied for more than a hundred years, yet a 2005 review found two questions still open: how the turbulence responds to a transient wave, and how the flow loses its symmetry about the pipe axis. Both matter for modeling how a transient dissipates its energy [2]. In case 2 friction takes most of the closure front's height before it reaches station 4. Progress would be a friction model, checked on real lines, that predicts how high a front still stands at the next station upstream.

The controllers carry simplifications of their own. The MPC moves only the terminal back pressure and infers the hill from a steady profile, so it lags a fast transient by several bar for a few minutes. Station setpoints, the flow and the additive remain the operator's. Stations are single pumps on speed control without throttle valves, sequencing or start-up. Pumping energy is reported but not judged, because the challenge does not fix the delivered volume.

> **OPEN CHALLENGE** **A Small Leak Under Temperature and Batch Changes**
>
> A 2019 review of leak detection counts temperature among the uncertainties that affect flow monitoring most, because it changes the density, viscosity and friction of the liquid. Temperature-dependent flow models had been limited to short lines, and the review asks for detectors that find the valid leaks with fewer false alarms. A volume balance, it notes, is sensitive to the dynamics of the line and cannot locate a leak [24]. On this isothermal line every detector setting that stayed quiet through a batch change missed a small leak for an hour. Progress would be a detector that follows temperature and batches along the whole line and still alarms a leak smaller than the meters' product shift.

13.8 Conclusion

A refined products line is a single flow shared by pumps in series, carrying batches that move at walking pace and pressure news that moves at 1110 m/s, and this chapter measured both on one simulated line. Without override selectors a 10 % flow cut trips all four stations within 10.2 min, while with them a station trip becomes a steady 781 m^3/h. A 5 s closure raises the terminal 18.5 bar in ten seconds, and slower strokes remove that front but end at the same packed pressure near 87 bar g. The densitometer's delay puts more of the following product into the leading tank at every cut, a drag reducer buys 187 m^3/h over eight hours, and a hill near the terminal needs about 15 bar g of back pressure at 500 m^3/h, where the MPC holds 15.8 bar g, and only its 3 bar g bound at full flow. A

balance detector tight enough to catch a 1 % leak in minutes also alarms on a meter's product shift at every batch change. The model is isothermal and lumped, and its slack treatment is declared rather than physical. A fuller model would add temperature to the line pack, a partly full pipe model for the hill, multi-pump stations with throttling, and a transient-model leak detector. The line is a compact setting for practicing override control on units in series.

KEY POINTS

- One flow passes every pump station, so only the lead station sets it and the others give way.
- Without override selectors a 10 % flow cut trips all four stations within 10.2 min. With them a station trip becomes a steady 781 m^3/h.
- A 5 s closure raises the terminal 18.5 bar in ten seconds. Slower strokes remove that front but end at the same packed pressure.
- A hill near the terminal needs about 15 bar g of back pressure at 500 m^3/h.
- A balance detector tight enough to catch a 1 % leak in minutes also alarms on a meter's product shift at every batch change.

Where else this applies. The Joukowsky rise of a fast closure also governs the penstock of Volume 1, Chapter 7, and an override that gives up throughput to hold a limit protects Volume 11, Chapter 3 as well.

References

[1] E. S. Menon, *Liquid Pipeline Hydraulics*. CRC Press, 2004.

[2] M. S. Ghidaoui, M. Zhao, D. A. McInnis, and D. H. Axworthy, "A review of water hammer theory and practice," *Applied Mechanics Reviews*, vol. 58, no. 1, pp. 49–76, 2005.

[3] A. H. Glattfelder and W. Schaufelberger, *Control Systems with Input and Output Constraints*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2003.

[4] M. S. Yoon, C. B. Warren, and S. Adam, Eds., *Pipeline System Automation and Control*. ASME Press, 2007.

[5] R. Rejowski and J. M. Pinto, "Scheduling of a multiproduct pipeline system," *Computers & Chemical Engineering*, vol. 27, no. 8–9, pp. 1229–1246, 2003.

[6] J. E. Austin and J. R. Palfrey, "Mixing of miscible but dissimilar liquids in serial flow in a pipeline," *Proceedings of the Institution of Mechanical Engineers*, vol. 178, no. 1, pp. 377–389, 1963.

[7] R. M. Baptista, F. B. de Freitas Rachid, and J. H. Carneiro de Araujo, "Estimating mixing volumes between batches in multiproduct pipelines," in *Proceedings of the 2000 International Pipeline Conference, Volume 2: Integrity and Corrosion; Offshore Issues; Pipeline Automation and Measurement; Rotating Equipment*. American Society of Mechanical Engineers, 2000, p. V002T08A008.

[8] M. Henrie, P. Carpenter, and R. E. Nicholas, "Mass balance leak detection," in *Pipeline Leak Detection Handbook*. Elsevier, 2016, pp. 41–55.

[9] A. F. Colombo, P. Lee, and B. W. Karney, "A selective literature review of transient-based leak detection methods," *Journal of Hydro-environment Research*, vol. 2, no. 4, pp. 212–227, 2009.

[10] T. Ma, O. O. Youzwishen, and M. Hylton, "Avoidance of slack flow in liquid pipelines," in *2004 International Pipeline Conference, Volumes 1, 2, and 3*. ASMEDC, 2004, pp. 491–497.

[11] G. I. Taylor, "The dispersion of matter in turbulent flow through a pipe," *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 223, no. 1155, pp. 446–468, 1954.

[12] A. Bergant and A. R. Simpson, "Pipeline column separation flow regimes," *Journal of Hydraulic Engineering*, vol. 125, no. 8, pp. 835–848, 1999.

[13] P. S. Virk, "Drag reduction fundamentals," *AIChE Journal*, vol. 21, no. 4, pp. 625–656, 1975.

[14] E. D. Burger, W. R. Munk, and H. A. Wahl, "Flow increase in the trans alaska pipeline through use of a polymeric drag-reducing additive," *Journal of Petroleum Technology*, vol. 34, no. 2, pp. 377–386, 1982.

[15] J. C. P. Liou and J. Tian, "Leak detection—transient flow simulation approaches," *Journal of Energy Resources Technology*, vol. 117, no. 3, pp. 243–248, 1995.

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

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

[18] 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] Pipeline and Hazardous Materials Safety Administration, "49 CFR 195.406, maximum operating pressure," Code of Federal Regulations, Title 49, Part 195, edition of October 1, 2025, govinfo.gov, 2025.

[20] ——, "49 CFR 195.428, overpressure safety devices and overfill protection systems," Code of Federal Regulations, Title 49, Part 195, edition of October 1, 2025, govinfo.gov, 2025.

[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] Pipeline and Hazardous Materials Safety Administration, "49 CFR 195.444, leak detection," Code of Federal Regulations, Title 49, Part 195, edition of October 1, 2025, govinfo.gov, 2025.

[23] ——, "49 CFR 195.134, leak detection (design requirements)," Code of Federal Regulations, Title 49, Part 195, edition of October 1, 2025, govinfo.gov, 2025.

[24] M. A. Adegboye, W.-K. Fung, and A. Karnik, "Recent advances in pipeline monitoring and oil leakage detection technologies: Principles and approaches," *Sensors*, vol. 19, no. 11, p. 2548, 2019.

13.9 Exercises

13.1 Concept check · 2 min

The line runs at 1000 m^3/h of diesel with the override selectors On, and the shipper cuts the nomination to 900 m^3/h. Which loop makes the line carry 900 m^3/h, and what do the other stations do?

(a) Every station's discharge loop lowers its discharge in proportion, so the four pumps share the cut

(b) Station 4's suction loop slows its pump until Terminal B's meter reads 900 m^3/h

(c) Terminal B's back-pressure loop closes its valve until the delivery has fallen by the cut

(d) Station 1's flow loop alone sets the flow, and stations 2 to 4 hold their suctions under discharge limits

13.2 Concept check · 1 min

In case 2 the stroke time of Terminal B's valve is raised from 5 s to 600 s and the closure is run again. Which statement matches the measurements of Section 13.5?

(a) The steepest ten-second rise falls to about a quarter, and the terminal still ends near 87 bar g

(b) Both the steepest rise and the final pressure fall, because the stations see the closure coming

(c) The steepest rise is unchanged, because $\rho a \Delta V$ holds no time, while the final pressure falls

(d) Neither changes, because the line stops in either case and the same volume is packed

13.3 Concept check · 2 min

In case 7 the balance detector false-alarms after gasoline follows diesel into the line, with no leak open. What causes the false alarms on this line?

(a) The gasoline is warmer than the diesel, and the detector applies no temperature correction

(b) The origin meter, proved on diesel, reads gasoline high, and the line pack worked out from eight transmitters misses part of the change

(c) The leak at km 180 weeps before it opens, and a tight detector sees it early

(d) Transmitter noise alone, which a longer window would average away

13.4 Theory and calculation · 5 min

Use the diesel values of Section 13.3: inside diameter 392.1 mm, density 840 kg/m^3 and wave speed 1130 m/s.

(a) Compute the bore area, the bulk modulus $K = \rho a^2$, and the volume stored per bar of pressure by one 10 km segment and by the 70 km first leg. *Checkpoint:* 0.1207 m^2, 1.073 GPa, 0.113 m^3 and 0.788 m^3 per bar.

(b) Station 1 sends 100 m^3/h less than station 2 takes. How fast does station 2's suction fall? *Checkpoint:* 2.1 bar per minute.

(c) At that rate, how long does the suction take to fall from its 8 bar g setpoint to the 3 bar g low-suction alarm? Case 3 measures 4.6 min from the cut. Why is the measurement longer? *Checkpoint:* 2.4 min.

13.5 Theory and calculation · 6 min

Terminal B delivers 851.8 m^3/h of diesel at the design point when its valve is shut in 5 s.

(a) Compute the velocity stopped at the valve and the Joukowsky rise of Equation (1.9) with $a = 1130\,\mathrm{m/s}$. *Checkpoint:* 1.96 m/s and 18.6 bar.

(b) How long would a sharp front take to cover the 90 km back to station 4 at the declared 1130 m/s and at the 1110.5 m/s timed on the model? The trends show it about 85 s after the close. *Checkpoint:* 79.6 s and 81.0 s.

(c) Compare the time a pressure change takes to cross the 300 km line with the time a batch takes to reach Terminal B at 1000 m^3/h to km 140 and 850 m^3/h beyond. *Checkpoint:* 4.4 min and 39.6 h, a ratio of about 540.

13.6 Theory and calculation · 6 min

Use Equation (13.2) for diesel at 1000 m^3/h (kinematic viscosity 3.5 cSt, roughness 45 µm).

(a) Compute the velocity, the Reynolds number, the Swamee–Jain friction factor and the friction gradient in bar per km and in meters of diesel per km. *Checkpoint:* 2.30 m/s, 257 700, 0.0159, 0.902 bar/km and 10.9 m/km.

(b) Compute the drag reduction at 5, 10 and 20 ppm, and the gradient at 10 ppm if the whole leg is treated. *Checkpoint:* 23.1, 33.3 and 42.9 %. 0.601 bar/km.

(c) The additive moves with the liquid and is sheared to nothing at the next station. How long does it take to fill a 70 km leg at 1000 m^3/h, and what does that mean for using it on a nomination? *Checkpoint:* 8.45 h.

13.7 App challenge · 8 min

Load scenario 3, Low suction wins, with the default seed. Run at 5 min/s and export the CSV file after each run.

(a) Run as shipped, with Override selectors Off, to 30 min. Find the first row in which `truth_suc_min_barg` falls below the 3 bar g alarm, and list the order in which the stations trip, with the protection that trips each (status strip and station cards). *Checkpoint:* 14.58 min, within 0.1 min. Stations 2, 3 and 4 on low-low suction, then station 1 on high-high discharge.

(b) Set Override selectors to On, press Reset and run to 60 min. Read the Worst suction tile, the lowest value of `suc2_barg`, and `truth_flow_in_m3h` at 60 min. *Checkpoint:* 6.9 bar g. 7.0 bar g, within 0.1 bar. 899 m^3/h, within 2 m^3/h. No LOW SUCTION and no trip.

Explain the difference between the two runs in terms of who sets the flow.

13.8 App challenge · 8 min

Load scenario 2, A valve shut in five seconds, with the default seed. Each run goes to 30 min at 5 min/s. Export the CSV file after each run.

(a) Run as shipped, with the 5 s stroke. Read the largest value of `truth_surge_10s_bar`, `pB_meas_barg` at 30 min, and the Least MAOP margin tile. *Checkpoint:* 16.6 bar. 87.2 bar g, within 0.2 bar. 0.3 bar.

(b) Set Terminal B valve stroke time (PIC card) to 120 s, press Reset and repeat the three readings. Then do the same with 600 s. *Checkpoint:* 8.3 bar, 87.2 bar g, 0.3 bar, no SURGE. 4.1 bar, 87.3 bar g, 0.3 bar, no SURGE.

(c) Which of the three readings does the stroke time control, and why does the margin to the MAOP not improve?

13.9 App challenge · 9 min

Load scenario 6, Slack line over the hill, with the default seed. Run at 20 min/s and export the CSV file after each run.

(a) Run as shipped (back pressure 5 bar g) to 90 min. Find the first row with `truth_slack_m3` above 0.5 m^3, where SLACK LINE is raised, and the slack volume at 90 min. *Checkpoint:* 25.1 min, within 0.1 min. 15.3 m^3.

(b) Press the reload button, set Terminal B back pressure (PIC card) to 10 bar g before Run and run to 120 min. Then set it to 15 bar g, press Reset and run again. Is SLACK LINE raised in each run? *Checkpoint:* At 10 bar g, yes, from 26.7 min, with 11.8 m^3 at 120 min. At 15 bar g, no, the largest `truth_slack_m3` is 0.37 m^3.

(c) Press the reload button, choose MPC for Line control before Run, and run to 120 min. Read `pB_sp_barg` and `mpc_hill_est_barg` at 90 min. *Checkpoint:* 15.8 bar g and 2.1 bar g, each within 0.2 bar. No slack.

Why is the slack volume in (a) a sign of slack and not a prediction?

13.10 Application area · 3 min

The densitometer at Terminal B reads its sample 30 s late.

(a) How much line passes the valve during that delay at 850 m^3/h? *Checkpoint:* 7.1 m^3.

(b) In case 4 a cut at 10 and 90 % put 2.01 m^3 of diesel into the gasoline tank and 0.74 m^3 of gasoline into the diesel tank. Explain the asymmetry, and name two changes that would reduce the larger figure on a real line.

13.11 Application area · 4 min

Station 4 pays for Terminal B's back pressure on the 850 m^3/h that pass Terminal A, at a pump efficiency of 0.8.

(a) Compute the pumping power per bar of back pressure, and the energy a day of 10 bar g costs over the MPC's 3 bar g bound at full flow. *Checkpoint:* 0.0295 MW per bar, and about 5.0 MWh.

(b) Over the challenge's day the suggested settings use 107.9 MWh and pump 19 749 m^3 under the MPC, and 102.3 MWh and 19 324 m^3 with a fixed 10 bar g. Compute the energy per cubic meter of each, and explain why the challenge reports energy and does not judge it. *Checkpoint:* 5.46 kWh/m^3 and 5.29 kWh/m^3.

13.12 Review the agent · 5 min

An AI assistant proposes a tuning for the detector of case 7: "Set the leak alarm threshold to 2 m^3/h, the balance window to 60 min and the persistence to zero. A 60 min window averages away the

transmitter noise that causes the false alarms, so the detector will stay quiet through the batch change and catch leaks as small as 0.2 %.” Find the flaw, and say how you would check the proposal in the app.

CHAPTER 14

An LNG Liquefaction Simulator for Loop Pairing and Inferential Constraint Control

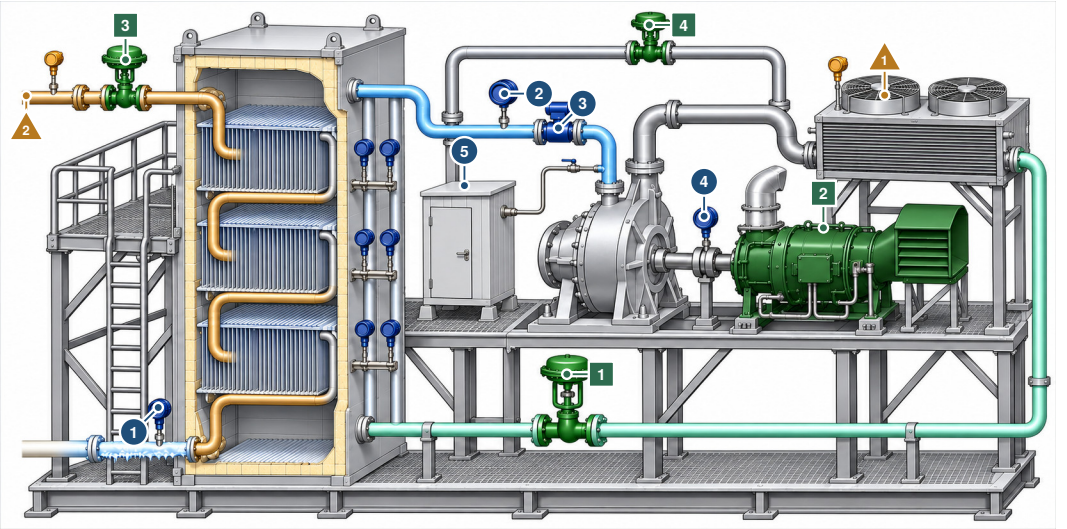

The controller moves 1 Joule-Thomson valve, which sets the circulation 2 compressor driver, which sets the speed 3 feed gas valve, cut back when the driver runs short 4 anti-surge recycle valve

The controller reads 1 LNG outlet temperature 2 suction pressure 3 refrigerant circulation 4 shaft power 5 refrigerant composition, every twelve minutes

What disturbs it ▲1 warm air at the air cooler, which also derates the driver ▲2 a richer feed, or a cut in it ▲ light ends lost from the refrigerant

Notice the throttling valve on the cold return, whose two effects the chapter separates, and that the gauges drawn on the cold box wall are not in the chapter, whose approaches inside the box are computed from the refrigerant analysis in the cabinet rather than measured. Illustration, not to scale.

IN PLAIN TERMS Cooling Gas Until It Turns to Liquid

The main idea. A liquefaction train chills natural gas until it becomes a liquid that ships can carry. A very cold fluid boils in a block of metal passages and draws heat out of the gas.

The challenge. One valve sets both how much cold fluid flows and how warm it boils, and the two pull in opposite directions. The closest gap between gas and fluid lies where no instrument reaches, so the plant computes it from a sample up to twelve minutes old.

Why it matters. The machine that drives the cooling is the largest on the site, so the energy it spends on each tonne matters most.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History boxes in Section 14.2, the Did you know? boxes in Sections 14.1 and 14.6.5, the Try it boxes in Section 14.6, and Key points in Section 14.10. No equations.
Study	70 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	25 min	This page, Section 14.6 with Figure 14.6, Section 14.7 and Section 14.8 with its Open challenge box.

The control question The Joule-Thomson valve of a mixed refrigerant train sets the circulation and the suction pressure together, and the two reach the LNG temperature with opposite signs, while the closest approach inside the exchanger is computed from an analysis twelve minutes old. Which loop structure makes the valve a usable handle, and what does holding an inferred limit cost when its input goes stale?

Learning objectives

- Separate the two effects of the Joule-Thomson valve by experiment and explain why no PI tuning repairs a gain that changes sign.
- Compute the refrigerant boiling temperature, a zone duty and the bundle lag from the model's equations.
- Classify circulation as a capacity variable and suction pressure as the temperature trim, and the blend as a max-min problem.
- Explain how a stale analysis drives the computed cold-end approach the wrong way, and read the gap from the exported file.
- Compare the cascade with the predictive controller by time, depth and integral of the driver excess.

Prerequisites Sections 1.4.7 (cryogenic processing), 1.4.6 (surge margin), 1.4.8 (low selector), 1.4.9 (inferred limits), 1.6.2 and 1.6.4, with Section 1.6.3 only for the MPC. Chapter 11 for anti-surge recycle.

The app https://apmonitor.com/apps/lng-liquefaction/

See also Volume 6, Chapter 2 for a valve whose gain changes sign near design, and Volume 6, Chapter 8 for a plantwide structure on a plant with sampled analyzers. Volume 5, Chapter 3 and Volume 5, Chapter 10 for quality held through a late analysis. Volume 4, Chapter 2 for the surge cycle under a protective override, and Volume 4, Chapter 3 for limiters selected on a margin that is never measured.

14.1 Introduction

Every loop on a liquefaction train holds a quantity that an instrument reports. Suction pressure comes from a transmitter, refrigerant circulation from a flowmeter, and the product temperature from a thermocouple in the outlet nozzle. The quantity that decides how much power the plant burns per tonne is none of those. It is the minimum internal temperature approach, the closest the warming refrigerant and the cooling gas come to each other anywhere inside the exchanger, and in the coil-wound bundle modeled here three of the four places that minimum can sit are inside a welded aluminum block with no instrument in them. What the control room works with instead is an arithmetic result computed from measured flows, measured pressures and a refrigerant composition that arrives from a gas chromatograph once every twelve minutes. Holding a variable that is calculated rather than measured is inferential control, a structure posed and analyzed for chemical processes half a century ago [1] and now the ordinary arrangement wherever the economically interesting property is slow or impossible to measure [2].

The class is wider than cryogenics. A distillation column controls a tray temperature because the composition analyzer is slow, a fired heater controls a tube skin because the coke layer is invisible, and a polymerization reactor controls a jacket temperature because the melt index arrives from a laboratory once a shift. Volume 5, Chapter 3 controls a column through a slow analyzer, and Volume 5, Chapter 10 steers a dryer by its exhaust air while the moisture arrives from a laboratory. Each of those plants runs a loop on a proxy, and each inherits the same failure mode. The loop holds its setpoint exactly while the quantity it stands for walks away, because the model that connects the two has an input that is stale. Which controlled variable to hold, and how much the choice costs when the plant moves, is the self-optimizing control question that the plantwide literature poses in general terms [3] and that the LNG literature has posed for this particular plant [4, 5].

A second difficulty on this train is a pairing one, and it is the reason a transfer function is a poor substitute for the plant. The Joule-Thomson valve does two things at once. It sets how much refrigerant circulates, which is more cooling, and, because whatever it passes has to be taken away again by the compressor, it sets the pressure at which that refrigerant boils, which is warmer refrigerant. At steady state, opening the valve from 54.9 to 62 % puts 4.6 kg/s more refrigerant through the bundle

and raises the suction pressure from 2.90 to 3.17 bar, and the two effects together buy 4.1 °C of product temperature before the gain turns over and becomes negative. A first-order-plus-dead-time model fitted to this loop never shows that sign change, because the sign change is the thing the fit averages away. A poorly behaved loop is easily taken for a tuning problem, and here it is not one. No proportional-integral setting repairs a gain of 3.4 °C per percent in one direction, 0.27 °C per percent in the other and negative beyond 60 % valve. The repair is structural, and interaction measures have been the standard way of choosing among such structures since the relative gain array was proposed [6].

Liquefaction is where an LNG project spends its energy. The refrigerant compressor driver is the largest machine on the site, and specific power is the number a commercial team argues about, which is why optimization of the cycle and of its operation has its own substantial literature [7, 8] and why the fuel is assessed on the efficiency of the chain that delivers it [9]. Published treatments of LNG control select controlled variables and evaluate structures on steady-state or full process-simulator models that cannot be opened, run or perturbed in an hour. This chapter works with a browser-hosted single mixed refrigerant train whose every equation, constant and measured number is given in the text, so that its sixteen states, its three exchanger zones, its compressor map and its five loops can be rebuilt, and uses it to pull apart two effects that one valve delivers together. The pairing is measured: the open-loop product temperature turns over at 60 % valve travel at fixed compressor speed and becomes linear at 9.2 °C per bar once a suction pressure loop is closed, and refrigerant circulation is a capacity variable whose gain changes by a factor of 16 across its working range. The model is an effectiveness-NTU three-zone cold box with bundle metal inertia, a sampled composition analyzer, an inferential approach estimator and a low selector on throughput, written out as implemented rather than as intended. An optional model predictive controller is compared with the cascade on five measured cases, losses included. The simulator reports both the true minimum approach and the control room estimate of it, so the difference between the variable a loop holds and the variable that limits the plant can be watched on the same screen. The cycle, the effectiveness-NTU relation, the cascade and the predictive controller are all established methods.

> ✦ **DID YOU KNOW?** **Part of the Gas Pays for the Cold**
>
> A federal estimate puts the gas used for liquefaction at about 7 to 15 % of the gas delivered to United States LNG export plants, mostly to run the equipment on site. All but one of those plants drive their refrigerant compressors with gas turbines [10]. This chapter's train spends 0.321 kWh of shaft work per kilogram.

Section 14.2 traces the cycle from the classical cascade to the mixed refrigerant and the published control work. Section 14.3 states the train model, its constants and its assumptions. Section 14.4 measures the two effects of the valve, the pressure handle, the circulation as capacity and the blend as a max-min problem. Section 14.5 describes the five loops, the low selector, the instruments and the predictive controller. Section 14.6 measures the scenarios, from two handles by hand to the graded shift and the predictive controller under a driver limit. Section 14.7 sets the hardware of a real train beside the model, Section 14.8 lists what the model leaves out, Section 14.9 weighs the economics and the methods that fit the same equations, Section 14.10 sums up, and the exercises in Section 14.11 put the scenarios to work. The LNG leaving this train is returned to gas at the terminal of Chapter 15.

14.2 Background

Liquefying natural gas is an old problem with a young constraint. How the mixed refrigerant, a blend that boils over a range of temperature, replaced the classical cascade of pure fluids and its three machine trains with one is the subject of the first History box below [8, 11]. What the substitution bought in capital cost it charged for in operability, because the shape of the refrigerant warming curve is set by its composition, and composition is not an actuator that anybody moves in a hurry.

HISTORY 1934 to 1970: From a Staircase of Pure Fluids to a Blend

A patent application of 1934 proposed a refrigerant made of a mixture of gases and vapors, compressed by one machine and condensed in stages, whose composition could be adjusted while the system was still in operation [12]. Natural gas liquefiers first took another road. A 1966 application for a liquefier of gas to be shipped across water described the practice of its day: separate cycles of essentially pure propane, ethylene and methane, each over its own band of temperature, whose curve is a staircase that in its example needed nine exchangers. It proposed one fluid of methane, ethane, propane, butanes and pentanes, with perhaps a little nitrogen, whose composition could be changed to keep the plant near its optimum [13]. An application of 1970 named the gain: the curves of single refrigerants do not closely match the continuous cooling curve of the gas, and horsepower is wasted [14].

The constraint acquired a name and a method in a different discipline. Pinch analysis, developed for heat exchanger networks in the late 1970s and formalized as the pinch design method, established the minimum approach temperature as the single quantity that trades exchanger area against utility consumption, and established that the network can only be understood by plotting both composite curves and reading the gap between them [15]. Liquefaction is that argument compressed into one exchanger over 185 degrees of temperature. Refrigerant synthesis methods that choose a blend by matching composite curves over the whole temperature range followed, and they treat the minimum approach as a constraint on the optimization rather than as a setpoint [16]. The design approach on this train is 5 °C, and the reason it is a constraint and not a target is that the composition which maximizes the smallest approach is an interior maximum. Going lighter opens the cold end and closes the warm end, and going heavier does the reverse.

The published control work asks which variables a liquefaction plant should hold constant when the ambient temperature and the feed move, and answers it with steady-state loss evaluations rather than with a tuning rule, from the first study of optimal operation of a simple mixed refrigerant cycle, told in the second History box below [4], to controlled-variable selection for a full LNG process on a dynamic first-principles model of the TEALARC plant [5] and a control structure designed from a dynamic model with the same question in view [17]. Three results are common to that work and are the ones this application is arranged around. Suction pressure, and not refrigerant flow, is the handle that sets product temperature. The compressor driver is an active constraint for much of the year, so throughput is a manipulated variable rather than a schedule. And the approach temperature has to be inferred, so the estimator and its inputs are part of the control design rather than an instrument detail.

HISTORY 2006: Operating the Plant, Not Only Designing It

A 2006 study of the simplest single mixed refrigerant process began by noting that, for all the work that goes into the design of LNG processes, their operation received surprisingly little attention. It counted nine manipulated inputs and three constraints active at the optimum, one of them the LNG leaving at −155 °C or colder. In design the minimum approach is specified, but in operation it is free to move. With the design areas held, the optimum closed it from 1.2 to 0.537 °C and needed 2.60 % less compressor power. The study then asked which variables, held at constant setpoints, keep the plant close to that optimum as disturbances arrive [4].

The ambient constraint deserves its own sentence, because it is the one a control room meets daily. Warm air raises the temperature at which the refrigerant condenses, which raises the discharge pressure and therefore the work per kilogram, and it derates the gas turbine driving the compressor at the same time. The two effects move production in the same direction and are the subject of a dedicated body of mitigation work in the LNG industry [18]. On the train modeled here an ambient change from 18 to 36 °C raises the shaft power required from 23.1 to 26.9 MW while lowering what the driver can deliver from 26.0 to 22.7 MW, and nothing on the product temperature trend says so until the plant runs out of power.

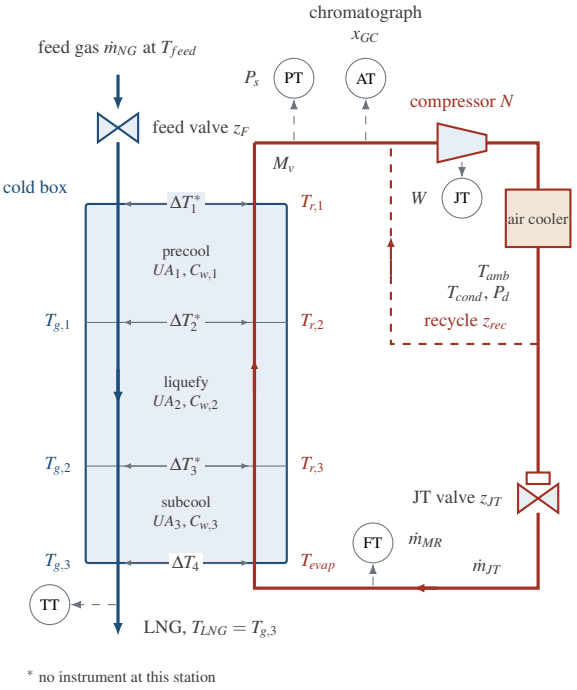

Figure 14.1 Single mixed refrigerant train with the cold box drawn as one block containing three counter-current zones. Feed gas runs down the left from T_{feed} to T_{LNG} and refrigerant runs up the right from T_{evap} to $T_{r,1}$, and the four stations where the two streams can come closest are marked ΔT_1 through ΔT_4. Only the cold end ΔT_4 has a temperature measurement on both sides of it.

14.3 Liquefaction Train Model

The application simulates a single mixed refrigerant train of the kind used at mid scale, taking 20.0 kg/s of natural gas at 30 °C and 55 bar and delivering LNG at −155.0 °C, which is 72.0 t/h and about 0.58 million tonnes a year once availability is allowed for. The refrigerant is compressed to 46.1 bar, condensed against ambient air, throttled through the Joule-Thomson valve to 2.90 bar and boiled back through the same exchanger it is cooling. The arrangement, the four places the two streams can pinch and the symbol for every quantity used below appear in Figure 14.1. The single-refrigerant cycle is used in preference to a propane-precooled arrangement because it carries the whole argument with one loop instead of three.

14.3.1 Variable Classification

Table 14.1 classifies every signal the loops and the scorecard interact with. Two entries need a word of explanation. The refrigerant light fraction x is a state and a setpoint at the same time, because the makeup system moves the inventory toward an operator target at a bounded rate of 0.010/min rather than setting it. The true minimum approach ΔT_{min} is an output that no instrument produces and that the implemented controllers never see, and it appears on the screen only because the simulator can compute what the plant cannot measure.

14.3.2 Three Counter-Current Zones

The cold box is one block carrying three lumped zones in series, precool, liquefaction and subcool, with the gas passing 1 to 3 and the refrigerant 3 to 1. Each zone is solved as a counter-current effectiveness-NTU exchanger rather than on a mean temperature difference. Heat-capacity rates

Table 14.1 Every Signal the Loops and the Scorecard Touch. The Nominal Column Is the Design Line-Up That Opens Each Scenario.

Symbol	Role	Nominal	Range
T_{LNG}	CV, measured	−155.0 °C	−168 to −146 SP
ΔT_4	CV, inferred	7.00 °C	2–16 SP
P_s	CV, measured	2.90 bar	1.9–5.2 bar
$\dot{m}_{MR}$	CV, measured	33.0 kg/s	14–48 kg/s
W	CV, measured	23.10 MW	16–28 MW SP
SM	CV, computed	44.9 %	6–30 % SP
z_{JT}	MV	54.9 %	0–100 %
N	MV	92.0 %	55–105 %
z_F	MV	100 %	40–110 %
z_{rec}	MV	0 %	0–100 %
T_{amb}	DV, measured	18.0 °C	0–45 °C
f_{rich}	DV, measured	1.00×	0.85–1.25×
f_{avail}	DV, measured	1.00×	0.40–1.05×
f_b	DV, unmeasured	1.00×	0.70–1.05×
f_{ac}	DV, unmeasured	1.00×	0.70–1.10×
$T_{g,1..3}$	states, °C	−35.0, −119.9, −155.0	−200 to +120
$T_{r,1..3}$	states, °C	+25.0, −40.0, −124.9	−200 to +120
$Q_{1..3}$	states, kW	3379, 9515, 2384	−4 to 40 MW
P_s	state	2.901 bar	0.6–14 bar
T_{cond}	state	30.16 °C	—
N_a, $z_{\cdot,a}$	actuator states	as MV	as MV
x	state and target	0.420	0.30–0.55
T_{evap}	output	−162.0 °C	—
ΔT_{min}	output, unmeasured	5.00 °C	—
$\Delta\hat{T}_{min}$	output, inferred	5.00 °C	—
w	output	0.321 kWh/kg	—

follow from the flows and the per-zone heat capacities, as shown in Equation (14.1).

$$C_{g,k} = \dot{m}_{NG}\, c_{g,k}, \qquad C_{r,k} = \dot{m}_{MR}\, c_{r,k} \tag{14.1}$$

where $\dot{m}_{NG}$ and $\dot{m}_{MR}$ are the gas and refrigerant mass flows in kg/s and $c_{g,k}$ and $c_{r,k}$ are the zone heat capacities in kJ/kg/°C, so every rate is in kW/°C. The smaller of the two rates is written $C_{\min,k}$. Feed richness raises the gas heat capacities by 20, 60 and 15 % of the excess $f_{rich} - 1$ in zones 1, 2 and 3, which puts most of the extra latent load in the middle zone. The refrigerant heat capacities move with composition, as shown in Equation (14.2).

$$c_{r,k} = \max\left(c_{r,k}^{\min},\; c_{m,k}\left[1 + s_k (x - x_0)\right]\right) \tag{14.2}$$

where x is the light fraction of the blend, x_0 is its design value, $c_{m,k}$ is the design heat capacity of zone k and s_k is the sensitivity of that zone to the blend, with $s_1 = -1.40$, $s_2 = -0.15$ and $s_3 = +1.60$ per unit of light fraction. The floors $c_{r,k}^{\min}$ are 0.35, 0.60 and 0.35 kJ/kg/°C. Those three signs are the physical content of the mixed refrigerant. Heavy components boil at the warm end and light ones at the cold end, so making the blend lighter moves duty from the top of the bundle to the bottom.

Zone effectiveness follows from the number of transfer units and the heat-capacity-rate ratio, as shown in Equation (14.3).

$$\varepsilon_k = \frac{1 - \mathrm{e}^{-\mathrm{NTU}_k (1 - C_{r,k}^{*})}}{1 - C_{r,k}^{*}\, \mathrm{e}^{-\mathrm{NTU}_k (1 - C_{r,k}^{*})}}, \qquad \mathrm{NTU}_k = \frac{U A_k\, f_b}{C_{\min,k}} \tag{14.3}$$

where $C^{*}_{r,k}$ is the ratio of the smaller heat-capacity rate to the larger, UA_k is the design zone conductance in kW/°C and f_b is the cold box condition multiplier. The balanced branch $\varepsilon = \mathrm{NTU}/(1+\mathrm{NTU})$ is used whenever $C^{*}_{r,k} > 0.999$, and on this plant that branch is the normal case rather than an edge case. At the design point $C_{g,1} = C_{r,1} = 52.0\,\mathrm{kW/^{\circ}C}$ and $C_{g,2} = C_{r,2} = 112.0\,\mathrm{kW/^{\circ}C}$, because matching the two rates in every zone is precisely what blending a mixed refrigerant is for. The duty a zone would deliver at its present inlets is then Equation (14.4).

$$Q_{ss,k} = \varepsilon_k C_{\min,k}\left(T_{g,in,k} - T_{r,in,k}\right) \tag{14.4}$$

with the gas inlets $T_{g,in} = (T_{feed}, T_{g,1}, T_{g,2})$ and the refrigerant inlets $T_{r,in} = (T_{r,2}, T_{r,3}, T_{evap})$, the two lists running in opposite directions because the block is counter-current. An effectiveness formulation cannot produce a crossed pair of streams at steady state. As the approach closes the effectiveness delivers less duty and the zone simply stops transferring, which is what the plant does, whereas a mean-driving-force cell will return a refrigerant colder than the gas at one end and warmer at the other and report a negative pinch that nobody can act on. In a transient the cold-end approach can still fall below zero for a while, because the boiling temperature of (14.9) follows the suction pressure and the blend at once while the gas holdup lags. It reaches $-0.52\,^{\circ}\mathrm{C}$ after the feed cut of scenario 10 and $-5.2\,^{\circ}\mathrm{C}$ at the instant of the relief event in the graded challenge run with its suggested settings.

14.3.3 Bundle Metal, Stream Holdup and the Duty Limiter

Five hundred tonnes of aluminum sit between the duty the conditions call for and the duty being delivered. The metal is carried as a first-order lag on each zone duty rather than as a temperature node, as shown in Equation (14.5).

$$\tau_k \frac{\mathrm{d}Q_k}{\mathrm{d}t} = Q_{ss,k} - Q_k, \qquad \tau_k = \frac{C_{w,k}}{60\,UA_k\,f_b} \tag{14.5}$$

where $C_{w,k}$ is the bundle heat capacity of zone k in kJ/°C and the factor 60 converts the seconds implied by kJ/°C over kW/°C into the minutes the clock runs in. At the design conductances the three lags are 4.44, 4.55 and 6.65 min, which is why nothing about a cold box answers quickly.

A lagged duty must never exceed what the two streams can physically exchange at the inlets they have at this instant, so a ceiling is computed at every step, as shown in Equation (14.6).

$$Q_{\lim,k} = C_{\min,k}\left(T_{g,in,k} - T_{r,in,k}\right) \tag{14.6}$$

The duty that reaches the temperature balances, $Q_{eff,k}$, is the state Q_k clipped between zero and $Q_{\lim,k}$ when $Q_{\lim,k}$ is positive and between $Q_{\lim,k}$ and zero when it is negative. Without (14.6) a sudden loss of feed would leave stale duty in the metal that cools the gas straight past the refrigerant and out the other side, and every approach in the application would go negative on a transient that a real bundle rides out. The two stream temperatures then follow from lumped holdups, as shown in Equations (14.7) and (14.8).

$$C_{NG}\frac{\mathrm{d}T_{g,k}}{\mathrm{d}t} = 60\left[C_{g,k}\left(T_{g,in,k} - T_{g,k}\right) - Q_{eff,k}\right] \tag{14.7}$$

$$C_{MR}\frac{\mathrm{d}T_{r,k}}{\mathrm{d}t} = 60\left[C_{r,k}\left(T_{r,in,k} - T_{r,k}\right) + Q_{eff,k}\right] \tag{14.8}$$

where C_{NG} and C_{MR} are the gas-side and refrigerant-side holdups per zone in kJ/°C. All six stream temperatures are clipped to -200 to $120\,^{\circ}\mathrm{C}$ and the three duties to -4 to 40 MW after every step.

14.3.4 Where the Refrigerant Boils

Two quantities set the cold end of the cycle and nothing else does, as shown in Equation (14.9).

$$T_{evap} = T_{evap,0} + k_P\left(P_s - P_{s,0}\right) - k_x\left(x - x_0\right) \tag{14.9}$$

where P_s is the suction pressure in bar, x is the light fraction, $k_P = 9.5\,^{\circ}\mathrm{C}$ per bar and $k_x = 120\,^{\circ}\mathrm{C}$ per unit of light fraction. Eight points of blend is therefore 9.6 °C of boiling temperature, which is more than the entire cold-end margin. Neither the Joule-Thomson valve nor the compressor speed appears in (14.9). They reach the boiling temperature only through P_s, which is the fact the pairing argument rests on.

14.3.5 Suction Inventory, Compressor and Driver

The suction pressure is a vapor inventory rather than a resistance, so its rate of change is set by the flow imbalance, as shown in Equation (14.10).

$$\frac{\mathrm{d}P_s}{\mathrm{d}t} = \frac{60\,P_s}{M_v}\left(\dot{m}_{JT} + \dot{m}_{rec} - \dot{m}_{comp}\right) \tag{14.10}$$

where M_v is the suction system vapor holdup in kg and the three flows are in kg/s. The inventory regulates itself, because a higher pressure feeds the compressor of (14.12) more vapor and leaves the Joule-Thomson valve a smaller pressure difference, so a speed change moves the pressure to a new value rather than ramping it. After a 1 % speed step with every loop open the pressure settles 0.043 bar lower, with 63 % of the change in 0.65 min. The two valves are square-root devices on the pressure difference across them, as shown in Equation (14.11).

$$\begin{aligned} \dot{m}_{JT} &= k_{JT}\,\frac{z_{JT}}{100}\sqrt{\max(0, P_d - P_s)} \\ \dot{m}_{rec} &= k_{rec}\,\frac{z_{rec}}{100}\sqrt{\max(0, P_d - P_s)} \end{aligned} \tag{14.11}$$

The compressor is a single-line map in which flow is proportional to speed and to suction density and is reduced by whatever head is being asked of it, as shown in Equation (14.12).

$$\dot{m}_{comp} = \dot{m}_{c0}\,n\,\frac{P_s}{P_{s,0}}\left[1 - a_M\left(\frac{P_d/P_s}{n^2\,P_{d,0}/P_{s,0}} - 1\right)\right] \tag{14.12}$$

where $n = N/N_0$ is the fractional speed, with the result floored at 0.5 kg/s and the pressure ratio floored at 1.05. The flow at the surge line for the same speed and suction condition is $\dot{m}_{surge} = \phi_s \dot{m}_{c0} n P_s / P_{s,0}$, and the surge margin reported to the operator and to the anti-surge loop is $\mathrm{SM} = \dot{m}_{comp}/\dot{m}_{surge} - 1$, the same fractional margin as Equation (1.11) of Section 1.4.6, with the surge flow in place of a map. Shaft power is polytropic head times the whole flow through the machine, recycle included, as shown in Equation (14.13).

$$W = \frac{\dot{m}_{comp}\,c_{p,g}\,(T_{r,1} + 273.15)}{\eta_c}\left[\left(\frac{P_d}{P_s}\right)^{\kappa} - 1\right] \tag{14.13}$$

where $\kappa = (\gamma - 1)/\gamma$ and the suction temperature is the refrigerant return temperature $T_{r,1}$. Recycled gas appears in $\dot{m}_{comp}$ and does no cooling, which is why anti-surge protection has a price. The air cooler and the driver close the loop on the weather, as shown in Equations (14.14) and (14.15).

$$\tau_{cond}\,\frac{\mathrm{d}T_{cond}}{\mathrm{d}t} = T_{amb} + \Delta T_{app} f_{ac}\left(\frac{\dot{m}_{comp}}{\dot{m}_{c0}}\right)^{0.8} - T_{cond} \tag{14.14}$$

$$\begin{aligned} P_d &= P_{d,0}\,\mathrm{e}^{k_d (T_{cond} - 30)} \\ W_{max} &= W_{driver}\left[1 - \beta\,(T_{amb} - 18)\right] \end{aligned} \tag{14.15}$$

where $\beta = 0.007/^{\circ}\mathrm{C}$ is the gas turbine derate. Both P_d and W_{max} move against the plant as the air warms.

14.3.6 Actuators, Makeup and Feed

Four actuator positions are states, each a first-order lag on its command with time constants of 0.15, 0.30, 0.08 and 0.20 min for the Joule-Thomson valve, the compressor speed, the recycle valve and the feed valve. The refrigerant inventory moves toward the makeup target at a bounded rate, as shown in Equation (14.16).

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \mathrm{clamp}\left(\frac{x_{SP} - x}{\tau_{make}},\ -\dot{x}_{max},\ +\dot{x}_{max}\right) \tag{14.16}$$

with $\tau_{make} = 25\,\mathrm{min}$ and $\dot{x}_{max} = 0.010/\mathrm{min}$, so a blend correction of five points takes about half an hour. Feed is $\dot{m}_{NG} = \dot{m}_{NG,0}\,(z_F/100)\,f_{avail}$.

14.3.7 The Four Approaches and What the Control Room Sees

The four candidate pinch stations are the warm end, the two zone boundaries and the cold end, and the minimum internal temperature approach is the smallest of them, as shown in Equation (14.17).

$$\begin{aligned}\Delta T_{min} = \min\big(&T_{feed} - T_{r,1},\ \ T_{g,1} - T_{r,2},\\ &T_{g,2} - T_{r,3},\ \ T_{g,3} - T_{evap}\big)\end{aligned} \tag{14.17}$$

At the design point the four are 5.01, 5.00, 5.00 and 7.00 °C, so the bundle is balanced rather than pinched at any one station and the application names whichever is smallest at the time. Only the fourth term of (14.17) has a temperature measurement at both ends of it, and even that one requires T_{evap}, which (14.9) obtains from the composition. The chromatograph is a zero-order hold of period $t_{GC} = 12\,\mathrm{min}$, refreshed at the first controller scan at or after each multiple of t_{GC}, so the estimate the control room computes carries an age, as shown in Equation (14.18).

$$\hat{T}_{evap} = T_{evap}\big(\hat{P}_s, x_{GC}\big), \qquad \Delta\hat{T}_4 = \hat{T}_{LNG} - \hat{T}_{evap} \tag{14.18}$$

where $\hat{P}_s$ and $\hat{T}_{LNG}$ are the filtered transmitter readings and x_{GC} is the last analysis. That is the cold-end approach the approach controller holds. The estimated minimum reported beside the true one evaluates the same cold-end term with the plant's own suction pressure and product temperature, so its error is the composition error alone. The two internal boundary estimates are shifted by $0.55\,k_x\,(x_{GC} - x)$ to represent an estimator that runs the same zone model on a stale composition, the warm-end estimate is taken as observed, and the estimated minimum $\Delta\hat{T}_{min}$ is the smallest of the four estimates. The gap between ΔT_{min} and $\Delta\hat{T}_{min}$ is the whole subject of the application, and it is zero only while the analysis is current.

14.3.8 Parameters, Initialization and Numerics

Tables 14.2 and 14.3 list every constant. Seven are hidden and never displayed, and those seven set the inertia and the speed of the suction inventory that the scenarios explore. The rest are visible and adjustable from the parameter panel.

No scenario begins on a transient the operator did not cause. The initializer solves the suction pressure at which the Joule-Thomson valve passes exactly what the compressor takes away, by bisection on the residual $\dot{m}_{JT} + \dot{m}_{rec} - \dot{m}_{comp}$ over 0.6 to 14 bar, which is monotone because raising the pressure feeds the machine and starves the valve. The three duties then follow from (14.4) by direct solution rather than iteration, because marching the gas down and the refrigerant up makes the three duties linear in each other. At $z_{JT} = 54.9\,\%$ and $N = 92.0\,\%$ the initializer returns 2.901 bar, 32.99 kg/s of refrigerant, −154.95 °C of LNG and 23.09 MW of shaft power, and the design operating point specified directly as 2.90 bar and 33.0 kg/s returns −155.00 °C and 23.10 MW. The application's verification suite, which the shared test harness of the apps runs against the shipped code, checks that second point, the absence of crossed streams at steady state over a grid of 75 operating conditions, and eleven other properties.

Table 14.2 Feed, Cold Box and Refrigerant Parameters

Symbol	Meaning	Value	Unit	Visibility
$\dot{m}_{NG,0}$	feed at 100 % valve	20.0	kg/s	adjustable
T_{feed}	feed temperature	30	°C	adjustable
$c_{g,1}$	gas c_p, precool	2.60	kJ/kg/°C	adjustable
$c_{g,2}$	gas c_p, liquefy	5.60	kJ/kg/°C	adjustable
$c_{g,3}$	gas c_p, subcool	3.40	kJ/kg/°C	adjustable
UA_1	precool conductance	676	kW/°C	adjustable
UA_2	liquefy conductance	1904	kW/°C	adjustable
UA_3	subcool conductance	401	kW/°C	adjustable
$C_{w,1}$	precool bundle metal	180000	kJ/°C	hidden
$C_{w,2}$	liquefy bundle metal	520000	kJ/°C	hidden
$C_{w,3}$	subcool bundle metal	160000	kJ/°C	hidden
C_{NG}	gas holdup per zone	6000	kJ/°C	hidden
C_{MR}	refrigerant holdup	8000	kJ/°C	hidden
$c_{m,1}$	refrigerant c_p, warm	1.576	kJ/kg/°C	adjustable
$c_{m,2}$	refrigerant c_p, mid	3.394	kJ/kg/°C	adjustable
$c_{m,3}$	refrigerant c_p, cold	1.949	kJ/kg/°C	adjustable
s_1	warm zone blend slope	−1.40	per unit x	adjustable
s_2	mid zone blend slope	−0.15	per unit x	adjustable
s_3	cold zone blend slope	+1.60	per unit x	adjustable
x_0	design light fraction	0.420	—	adjustable
$T_{evap,0}$	boiling T at design	−162.0	°C	adjustable
k_P	boiling T per bar	9.5	°C/bar	adjustable
k_x	boiling T per unit x	120	°C	adjustable
τ_{make}	makeup response	25	min	adjustable
$\dot{x}_{max}$	fastest makeup rate	0.010	1/min	adjustable

The engine of Section 1.6.1 advances the state vector on a fixed 0.01 min step (0.6 s) under classical fourth-order Runge-Kutta, the six controllers execute on a 0.05 min (3 s) scan, and the historian keeps ten samples for every simulated minute. The clock runs in minutes, and the speed setting counts simulated minutes per second of real time. The noise seed is 20260913, and no scenario switches the noise on, so every result in this chapter is deterministic. A speed of 20, which runs of several hours need, alters only the rate at which simulated time is consumed and leaves the trajectory untouched.

14.3.9 Assumptions

The model is deliberately detailed on the pinch and thin on thermodynamics, and each simplification fails somewhere.

Three lumped zones stand in for a continuous bundle. The pinch can only appear at a zone boundary or at an end, and on a real coil-wound exchanger the tightest point is usually in the middle of a zone. The zone heat capacities are fixed by the design temperature ranges, so a long way off design the capacity assigned to each zone stops being physical.

The refrigerant is two numbers. A light fraction and the heat capacity it implies per zone replace four components and a bubble curve. The signs and the orders of magnitude in (14.2) are right, but a real blend has a phase envelope that a linear coefficient cannot represent, and a real makeup system adds one component at a time rather than moving a single fraction. Equation (14.9) is linear in pressure and in composition, which is a fair fit over the range the sliders allow and not beyond it.

The gas is a sensible-heat stream. The latent heat of condensation is carried by the elevated $c_{g,2}$ of the middle zone, so the application can show how much extra duty a rich feed brings but not what it does to the shape of the cooling curve. There is no process-side pressure drop.

The bundle metal is a lag on duty and not a temperature node. That carries the thermal inertia, which is what matters for control, and loses the internal metal profile, which is what matters for cool-down and for thermal stress.

Table 14.3 Machinery, Valve, Instrument and Limit Parameters

Symbol	Meaning	Value	Unit	Visibility
$\dot{m}_{c0}$	flow at rated speed	39.9	kg/s	adjustable
N_0	design speed	100	%	hidden
$P_{s,0}$	design suction	2.90	bar	adjustable
a_M	flow lost per excess PR	0.55	—	adjustable
ϕ_s	reduced flow at surge	0.62	—	adjustable
$c_{p,g}$	vapor heat capacity	1.90	kJ/kg/°C	adjustable
κ	$(\gamma-1)/\gamma$	0.230	—	adjustable
η_c	compressor efficiency	0.720	—	adjustable
W_{driver}	driver rating at 18 °C	26000	kW	adjustable
β	driver derate	0.0070	1/°C	adjustable
M_v	suction vapor holdup	1000	kg	hidden
$P_{d,0}$	discharge at 30 °C	46.0	bar	adjustable
k_d	discharge rise	0.018	1/°C	adjustable
ΔT_{app}	air cooler approach	14.15	°C	adjustable
τ_{cond}	air cooler lag	1.5	min	adjustable
k_{JT}	JT valve capacity	9.14	kg/s per $\sqrt{\text{bar}}$	adjustable
k_{rec}	recycle valve capacity	3.6	kg/s per $\sqrt{\text{bar}}$	adjustable
τ_{JT}	JT valve travel	0.15	min	adjustable
τ_N	speed response	0.30	min	adjustable
τ_{rec}	recycle valve travel	0.08	min	adjustable
τ_F	feed valve travel	0.20	min	adjustable
T_{spec}	tank bubble point	−150.0	°C	adjustable
ΔT^{lo}_{min}	pinch warning	1.5	°C	adjustable
SM_{alarm}	surge alarm	0.10	—	adjustable
τ_{Tm}	cryogenic TT response	0.35	min	adjustable
τ_{Pm}	pressure transmitter	0.08	min	adjustable
τ_{Wm}	power transducer	0.15	min	adjustable
t_{GC}	chromatograph cycle	12	min	adjustable
p_{LNG}	value of on-spec LNG	0.44	$/kg	adjustable
p_{fuel}	cost of shaft power	0.030	$/kWh	adjustable

The machinery is a single-point map. Equation (14.12) is one line with a surge line and no choke, there is no compressor curve to read, and the driver is a rating with a linear ambient derate rather than a thermodynamic model. The surge margin is reported and defended, and the reversal itself is not simulated.

The chromatograph is a perfect sample and hold. Real analyzers have a sample line with its own lag, a calibration that drifts, and a habit of failing in ways that leave the last good value on the screen. Failure of the analyzer is the most common cause of the inferential problem this application demonstrates, and it is outside the model.

Three disturbances are stylized. A relief event that removes five points of light ends at once is faster than any real release, but it separates the composition upset from the analyzer delay cleanly, which a slow drift would not. An ambient change of eighteen degrees arrives as a step, which the air cooler follows within about five minutes, rather than over the hours a real afternoon takes, because the constraint crossing is the subject and not the ramp rate. A feed cut of 35 % in a single step would trip most plants, and it is used because the over-refrigerated bundle it produces is the clearest illustration of specific power moving while every controlled variable stays on setpoint.

Cool-down from ambient is absent. Every scenario starts with a cold plant. The sixteen to thirty hours a real train takes to get there is the operation that most limits how often anybody is willing to shut one down.

The arithmetic is too reliable. Freezing the step and the scan gives the same trajectory on every

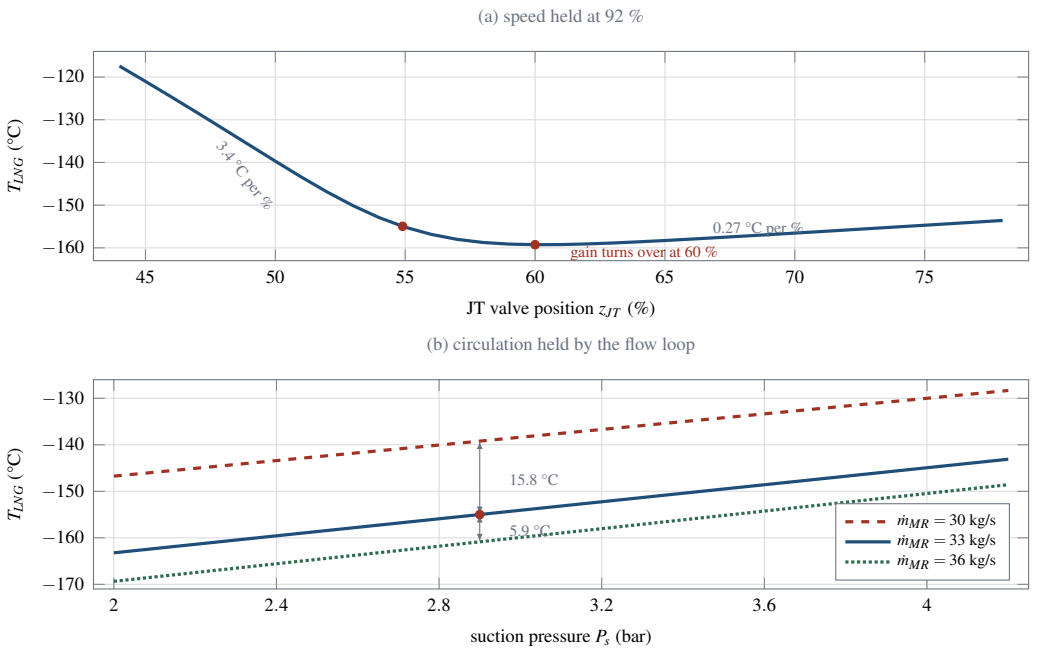

Figure 14.2 Steady-state characteristics computed from the implemented model. Panel (a) sweeps the Joule-Thomson valve with the compressor speed fixed, so that both of the valve's effects act, and the product temperature reaches a minimum of −159.26 °C at 60 % travel. Panel (b) sweeps the suction pressure with the circulation held, so that only one effect acts, and the three curves are straight and nearly parallel: 30 kg/s dashed, 33 kg/s solid and 36 kg/s dotted. The dots mark the design point and, in (a), the turning point.

machine, and a real control system delivers scan jitter, coarse measurement resolution and an execution order that moves.

14.4 Pairing Analysis

14.4.1 Two Effects on One Valve

The Joule-Thomson valve reaches the product temperature by two paths that oppose each other. More valve opening is more refrigerant through the bundle, which is more cooling and arrives over the four to seven minutes of the metal lags of (14.5). More valve opening also delivers more vapor to a compressor that cannot swallow it, so by (14.10) the suction pressure rises, and by (14.9) the refrigerant boils warmer. That second path is 63 % complete in 0.7 min and 95 % complete in 3.2 min after a step to 62 %. Figure 14.2(a) is the resulting open-loop characteristic with the compressor held at 92.0 % speed.

Closing the valve from 54.9 to 44 % costs 37.5 °C of product temperature at steady state, and in scenario 2 the product is 31.1 °C warmer than at 54.9 % an hour after the step and still moving, because the plant is short of refrigerant and the whole bundle has to warm up. Opening it by the same eleven points buys 3.0 °C, which the scenario delivers within eight minutes by putting the product onto the boiling temperature, and past 60 % the product gets warmer again, reaching −157.95 °C at 66 % against −159.08 °C at 62 %. The average gain over those two excursions is 3.4 °C per percent in one direction against 0.27 °C per percent in the other, a factor of 12, with a sign change in between. A relative gain array computed on this pair flags the interaction, because both candidate manipulated variables reach both candidate controlled variables [6], although a linear interaction measure cannot

show the change of sign. No proportional-integral setting repairs it, and the standard remedy is not a better tuning but a structural one, which is to give one of the two effects to another loop. The seawater train of Volume 6, Chapter 2 has a valve whose gain reverses sign a few percent either side of design for a similar reason, two terms of opposite sign on one handle.

14.4.2 What the Pressure Loop Buys

Closing a suction pressure loop on the compressor takes the pressure path away from the valve. What remains is Figure 14.2(b), and it is a straight line. Over 2.6 to 3.2 bar at 33 kg/s the product temperature moves 9.15 °C per bar, and the same sweep at 30 and 36 kg/s gives 8.37 and 9.44 °C per bar, so the gain of the outer loop is within 6 % of constant across the whole working range of circulation. The shaft power moves the other way over the same sweep, from 24.34 to 22.00 MW, so colder product is bought with power and the trade is visible on the same screen.

The vertical spacing of the three curves in Figure 14.2(b) carries the second result. At 2.90 bar the step from 30 to 33 kg/s is worth 15.8 °C while the step from 33 to 36 kg/s is worth 5.9 °C and the next three kilograms are worth 1.0 °C. Below about 31 kg/s there is not enough refrigerant to carry the duty at all, and above about 35 kg/s the cold end is already sitting on the boiling temperature and the surplus is pumped around the loop for nothing at about 1 MW of shaft power for every 3 kg/s. The ratio of the low-flow gain to the high-flow gain is 16.1. Circulation is therefore a capacity variable that decides how much cooling the plant has, and suction pressure is the trim variable that decides what temperature that cooling is delivered at, which is the pairing the published control studies of this plant arrive at from steady-state loss arguments [4, 5].

Pause. Try Exercise 14.3 before reading on.

14.4.3 The Blend Is a Max-Min Problem

Composition is the third handle and it is not a control handle. A heavy blend at $x = 0.360$ raises the boiling temperature to −154.80 °C, so at 2.90 bar no circulation can bring the product to its −155 °C setpoint. At 33 kg/s the product settles at −146.04 °C with the smallest approach at 3.22 °C on the subcool boundary, and 39 kg/s buys −154.66 °C only by closing the cold end to 0.14 °C. A light blend at $x = 0.500$ lowers the boiling temperature to −171.60 °C and gives away the warm end instead, where the approach closes to 1.22 °C because a light mixture has almost nothing left to boil at −40 °C. The design blend is neither. Searching x over 0.32 to 0.54 at 2.90 bar and 33 kg/s, the application returns $x = 0.419$ with a smallest approach of 4.97 °C, against 5.00 °C at the nominal $x_0 = 0.420$. The objective has an interior maximum, the makeup system needs half an hour to move a few points, and the measurement that says where the blend is arrives every twelve minutes, so a proportional-integral controller pointed at composition would not know which way to go. The simulator offers the search as a manual helper instead.

14.5 Implemented Control Structure

14.5.1 Five Loops and a Selector

Five loops are implemented, and the control structure switch selects how many of them are in service. Figure 14.3 shows the arrangement in its full form. The two inner loops hold pressure and flow on their own setpoints. The cascade adds the product temperature above the pressure loop and, when the approach mode is enabled, the inferred cold-end approach above the flow loop. The throughput switch puts the feed valve behind a low selector with the driver power controller. A fourth setting of the structure switch, never the default, hands the suction pressure and circulation setpoints and the makeup target to the model predictive controller of Section 14.5.5.

14.5.2 The Discrete Algorithm

All five loops and the anti-surge loop run the velocity-form PID of Section 1.6.2, Equations (1.16) to (1.19), executed every $\Delta t_c = 0.05$ min. Anti-windup is back-calculation with tracking time constant

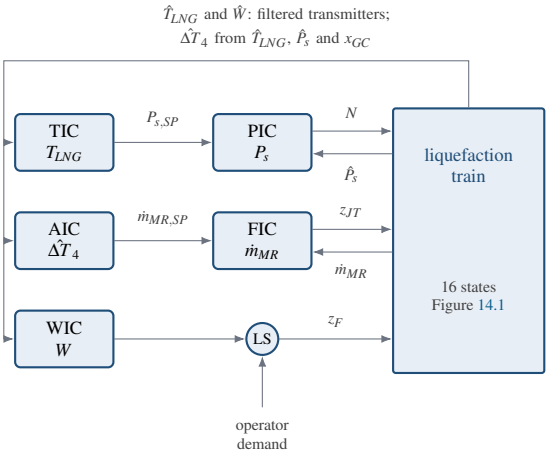

Figure 14.3 Control structure in its full form. The temperature controller drives the suction pressure setpoint, the inferred approach controller drives the circulation setpoint, and a low selector gives the feed valve to whichever of the operator and the power controller asks for less.

Table 14.4 Default Tuning of the Five Adjustable Loops. An Action of +1 Lowers the Output When the Measurement Rises.

CV	MV	K_c	τ_I (min)	a	Output limits
T_{LNG}	$P_{s,SP}$	0.075	14	+1	1.9–5.2 bar
P_s	N	14	1.6	−1	55–105 %
$\dot{m}_{MR}$	z_{JT}	2.0	0.4	+1	0–100 %
$\Delta\hat{T}_4$	$\dot{m}_{MR,SP}$	0.18	60	−1	14–48 kg/s
W	z_F	0.004	30	+1	40–110 %

$T_t = \tau_I$ on every loop [19]. Under a held error the integral comes to rest at the limit itself, so the demand before clipping exceeds the limit only by the proportional term and the output leaves the limit as soon as the error reverses. Away from the limits the recursion is algebraically identical to the positional law with the same gains, and on a limit it is not. Derivative action is zero on all six loops, so the derivative filter never runs, and no loop carries a rate limit, so the clipped output reaches the actuator model directly. In MANUAL the loops keep executing and track the operator's actuator position exactly, and in AUTO the operator's stations are set to the controller outputs after every execution, so transfer in either direction is bumpless with no separate logic. Output bias is zero on every loop and the integral is seeded at reset with the starting actuator value, so a loop that starts on its setpoint starts without a bump.

The action flag is the ISA sign convention. A loop with $a = +1$ is reverse acting, so a rising measurement lowers the output, and a loop with $a = -1$ is direct acting. Table 14.4 lists the five adjustable loops with their pairings, defaults and limits. The anti-surge loop is fixed at $K_c = 220$ and $\tau_I = 0.25$ min on the recycle valve, is reverse acting on surge margin, and is not offered for tuning because it is protection rather than regulation.

Three of the five tunings deserve a comment. The pressure loop carries 14 % of speed per bar with 1.6 min of reset, an order of magnitude more gain and an order of magnitude less reset than the temperature loop above it, because (14.10) is a fast, self-regulating inventory, 63 % complete in 0.65 min after a speed step, against the 4.44 to 6.65 min metal lags the temperature loop works through. The temperature loop carries 0.075 bar per degree with 14 min of reset, which is slow

because everything between the pressure and the thermocouple is aluminum. The approach loop carries 60 min of reset, because its measurement is refreshed by an analyzer every twelve minutes and a loop that reacts inside its own measurement interval is reacting to the hold rather than to the plant.

14.5.3 The Low Selector on Throughput

The feed valve belongs to the operator until the driver runs out of power. The power controller computes a demand against the smaller of the operator power setpoint and the derated driver rating W_{max} of (14.15), and the selector passes the lower of that demand and the operator's own, as shown in Equation (14.19).

$$z_F = \min(z_{F,op},\, u_W)\,, \qquad W_{lim} = \min(W_{SP},\, W_{max}) \tag{14.19}$$

The loser is told what was selected and tracks it through the same MANUAL expression, so the losing controller does not wind up while it is out of service and hands over without a bump when the selection changes, as the reset that follows the selector does in Section 1.4.8. The operator's own demand is never overwritten by the selected value, so the selector hands the valve back as soon as the driver can afford it. MANUAL holds the valve at the selected z_F (88.0 against 100 % at 40 min of scenario 9) and AUTO restores the demand, because the selector acts in AUTO only. That path is slow by construction. The feed valve does not reach the compressor directly. It changes the heat load, which changes the product temperature, which the temperature controller answers with suction pressure, which is what the machine feels, so the whole cold box is inside the power loop and the 30 min reset is a consequence rather than a choice.

14.5.4 Instruments and the Analyzer

The product thermocouple, the suction pressure transmitter and the power transducer are first-order filters of 0.35, 0.08 and 0.15 min, and the refrigerant flow measurement carries a fixed 0.08 min filter. Optional colored noise can be added to the temperature and pressure signals with a standard deviation of up to 2 °C and 0.15 bar. The composition analyzer is different in kind. It is a sample, a column and an integration, and it arrives as a step every 12 min carrying the plant that produced it. Between analyses the estimator of (14.18) runs on a number that is up to twelve minutes old, which is the mechanism the seventh scenario shows.

14.5.5 Model Predictive Control

A fourth setting of the structure switch, never the default and selected by no shipped scenario except scenario 12, replaces the temperature and approach controllers with one model predictive controller (MPC), an instance of the option described in general in Section 1.6.3. Every $T_s = 0.5$ min, it solves a quadratic program for the three setpoints $u = (P_{s,SP},\, \dot{m}_{MR,SP},\, x_{SP})$ and keeps only its first move, held for one decision period, the receding-horizon arrangement of industrial predictive control [20], also proposed in adaptive form for natural gas liquefaction [21]. The pressure and flow loops carry out the first two setpoints, the makeup law of (14.16) carries out the third, and the low selector and the anti-surge loop stay above the controller on the run's own settings.

The controller's model is the cold box of Section 14.3 written again with the declared values of Tables 14.2 and 14.3 and the design disturbances, $T_{amb} = 18\,^{\circ}\mathrm{C}$ and $f_{rich} = f_{avail} = f_b = f_{ac} = 1$, never with the run's, so a plant the reader changes is model mismatch. A tilde marks the controller's copy of a quantity. The inner loops enter as their identified responses to the setpoints, and the gas heat-capacity rates carry an estimated load factor $\tilde{L}$, as shown in Equation (14.20).

$$\begin{aligned} \tilde{P}_s &= \alpha\, \tilde{p}_1 + (1-\alpha)\, \tilde{p}_2, & \tau_i \frac{\mathrm{d}\tilde{p}_i}{\mathrm{d}t} &= P_{s,SP} - \tilde{p}_i \\ \tau_m \frac{\mathrm{d}\tilde{m}}{\mathrm{d}t} &= \dot{m}_{MR,SP} - \tilde{m}, & \tilde{C}_{g,k} &= \tilde{L}\, \dot{m}_{NG,0}\, \frac{\tilde{z}_F}{100}\, c_{g,k} \end{aligned} \tag{14.20}$$

where $\tilde{p}_1$ and $\tilde{p}_2$ are parallel lags, $\tilde{m}$ is the circulation and $\tilde{z}_F$ follows the feed command through its 0.20 min lag. The constants $\alpha = 0.413$, $\tau_1 = 0.70$ min, $\tau_2 = 3.90$ min and $\tau_m = 0.307$ min are fits to setpoint steps of either sign in scenario 3, with residuals of 1.2 and 2.1 % RMS of a unit step. The model carries 15 states, and its shaft power and compressor speed are (14.13) and (14.12) plus biases $\tilde{d}_W$ and $\tilde{d}_N$. Driven open loop on the declared plant through a 150 min sequence of steps, it predicts the LNG temperature within 0.036 °C RMS and every approach within 0.054 °C RMS. One linearization at the design point would put the LNG at −144.95 °C for 30 kg/s of circulation where the model settles at −139.21 °C, so every decision linearizes the model anew around the current estimate and samples it over T_s with an exact zero-order hold.

Each decision minimizes Equation (14.21) over $N = 120$ steps, a horizon of 60 min, with the moves held in nine blocks lasting 1, 1, 2, 2, 4, 6, 9, 15 and 80 steps, a program of 27 moves and 16 slacks.

$$
\begin{aligned}
\min_{u,\varepsilon} \sum_{k=1}^{N} \Big[& \big(\tilde{T}_{LNG,k} - T^*\big)^2 + q_A \big(\Delta\tilde{T}_{4,k} - \Delta T_4^*\big)^2 \\
& + r \sum_{j=1}^{3} g_j^2 \big(u_{j,k} - u_j^*\big)^2 \Big] + \lambda \sum_{b=1}^{9} \sum_{j=1}^{3} g_j^2 \Delta u_{j,b}^2 \\
& + \sum_{\ell=1}^{16} \rho_\ell \big(\varepsilon_\ell^2 + \varepsilon_\ell\big) \\
\text{s.t. } & u_j^{lo} \le u_{j,b} \le u_j^{hi}, \quad |\Delta u_{j,1}| \le \delta_j, \quad \varepsilon_\ell \ge 0 \\
& \tilde{y}_{i,k} \le \bar{y}_i + \phi_{i,k} + \varepsilon_\ell \quad \text{(LNG, power, speed)} \\
& \tilde{y}_{i,k} \ge \bar{y}_i - \phi_{i,k} - \varepsilon_\ell \quad \text{(approaches, surge margin)}
\end{aligned}
\tag{14.21}
$$

where k counts the steps and b the blocks, $\Delta u_{j,b}$ is the move of setpoint j into block b, the first measured from the setpoint in force, T^*, ΔT_4^* and u^* are the targets of (14.22), q_A is 1 when the circulation is the controller's and 0 when it is the operator's, $r = 0.01$, $\lambda = 10$, and $g = (9.2,\ 2.0,\ 145)$ °C per bar, per kg/s and per unit of light fraction prices each setpoint in the degrees of LNG it is worth near design. The bounds are the clamps the cascade and the makeup apply, 1.9 to 5.2 bar, 14 to 48 kg/s and 0.30 to 0.55, and the first-move cap is $\delta = (0.05\,\text{bar},\ 0.5\,\text{kg/s},\ 0.02)$, both hard. Every output limit $\bar{y}_i$ is soft. The LNG ceiling is $T_{spec} - 1\,°\text{C}$, the three internal approaches are held at $\Delta T_{min}^{lo} + 1\,°\text{C}$, which is 2.5 °C as shipped, the cold-end guard is 0.5 °C, the power limit is the selector's W_{lim} of (14.19), or 0.5 MW under it when the feed is the operator's, the speed limit is 104 % and the surge floor is 2 % above the anti-surge setpoint. The slack prices ρ_ℓ are 10^4 for the internal approaches, 10^3 for the LNG ceiling and the power and 10^2 for the rest, so the approach floor outranks the driver tenfold. The power and surge limits are soft because the shared solver has no hard output rows, which leaves the selector and the anti-surge loop as the hard protections. That solver also carries one slack per soft output, so each limit enters twice, over the first 10 min and over the rest, with a funnel $\phi_{i,k} = v_i\, e^{-kT_s/\tau_i}$ that returns an output already past its limit by v_i, with τ_i of 2 to 5 min and 25 min for the internal approaches in the second window. No slack per step is claimed.

A steady-state Kalman filter, designed at the design point for the 3 s scan, corrects 14 states and the random-walk disturbances $\tilde{L}$, $\tilde{d}_W$ and $\tilde{d}_N$ at every scan from the five signals the shipped loops read, which are $\hat{T}_{LNG}$, $\hat{P}_s$, the circulation, the shaft power and the speed command. Such disturbance states keep a predictive controller free of steady offset under model error [22]. The blend is not estimated. It is the last analysis x_{GC} carried forward by (14.16), which is what the control room knows. The approach limit is therefore a limit on a prediction, because the warm end and the two zone boundaries of (14.17) have no instrument, and in scenario 7 the smallest approach the controller predicts misses the true minimum by up to 5.37 °C before the vented light ends reach an analysis. The cold end, with a measurement on both sides, is tracked to its setpoint instead and guarded only against a cross.

The targets are the solution of Equation (14.22) on the model's steady map, found by up to eight steps of sequential quadratic programming inside a trust region of 0.3 bar, 3 kg/s and 0.04.

$$\begin{aligned}\min_{u^*}\ & \left(\tilde{T}_{ss}-T_{LNG,SP}\right)^2+q_A\left(\Delta\tilde{T}_{4,ss}-\Delta T_{SP}\right)^2\\ &+2500\left(x^*-x_0\right)^2+c_W\,\tilde{W}_{ss}+10^{-3}\sum_{j=1}^{3}g_j^2\left(u_j^*-u_j\right)^2\\ &+\sum_i\rho_i\left(\varepsilon_i^2+\varepsilon_i\right)+10^4\left(\varepsilon_T^2+\varepsilon_T\right)\end{aligned}\tag{14.22}$$

where ss marks the model's steady value at u^*, u_j is the setpoint in force, ΔT_{SP} is the cold-end setpoint, ε_i carry the limits of (14.21) other than the cold-end guard at the same prices, and ε_T forbids an LNG target colder than the setpoint. While the selector has the feed, c_W rises from 0 to 10 per MW as the steady power climbs from 1.5 to 0.5 MW under W_{lim}, because a megawatt saved is feed the selector can give back. A squared degree of LNG error then costs as much as a tenth of a megawatt, and the target runs the LNG warm whenever power or a predicted approach is worth more.

A plan is discarded and marked unless the solver reports an optimum with every value finite and every move inside its bounds, and the temperature and approach controllers, which track the applied setpoints at every scan, then hold the setpoints until the next valid plan. Forced rejections hand over within 0.0024 bar of suction setpoint with the noise off, and every plan in the measured runs was accepted. The priming call at $t = 0$ already solves a plan. In Chrome a plan takes 3.43 ms on average in scenario 12, 0.68 % of a core at one simulated minute per second and 13.7 % at the top speed of 20. The plant on the run's settings, the instruments and the controller's copy on the declared values are thus three separate models, and the 8800 measurement packets of scenario 12 replayed over four other hidden plants give the same 800 plans bit for bit.

14.6 Results

The thirteen scenarios of Table 14.5 run as a laboratory sequence. Each one loads a starting mode, a control structure, five tunings, a set of disturbance positions and whatever timed events belong to it, and carries the briefing, completion rule and data export of Section 1.6.4. One hundred twenty-five columns, twenty-seven of them the predictive controller's, leave the historian at ten samples per simulated minute, among them the true minimum approach beside the estimated one, the light fraction beside the last analysis, the age of that analysis, all four approaches and the running specific power. The first two open in MANUAL with deliberately weak tunings loaded, so that a reader who switches to AUTO and closes the loops sees what a badly tuned plant looks like before the tuned settings arrive in scenario 3. No scenario switches the measurement noise on, and the results below are the same whether a case is loaded from the selector, reloaded or restarted.

14.6.1 Manual Operation and the Two Handles

Scenario 1 makes the pairing argument with two open-loop steps and no controller, each made here from the design line-up five minutes after Run, as the scenario asks (a station set before Run starts the plant at its steady state, so it would show no response). Opening the Joule-Thomson valve from 54.9 to 62 % raises circulation from 32.99 to 37.63 kg/s and the suction pressure from 2.901 to 3.167 bar within about five minutes, so the refrigerant that boiled at −161.99 °C now boils at −159.46 °C. Twenty minutes after the move the product is at −159.46 °C, on the boiling temperature itself, and the cold-end approach has closed from 7.04 °C to zero, which raises the temperature-cross alarm 24.3 min after the move. The whole 4.6 kg/s of extra refrigerant bought 4.51 °C. The model's steady state at 62 % is −159.08 °C with 0.38 °C left at the cold end, and the run creeps toward it only over hours (−159.17 °C 240 min after the move). Raising the compressor speed from 92.0 to 96.0 % instead moves the suction to 2.735 bar and the boiling temperature to −163.57 °C, and twenty minutes later the product is at −156.60 °C (−156.73 °C at steady state), with the four approaches

Table 14.5 The Thirteen Scenarios and What Each One Is Arranged to Show

#	Scenario	What it exercises
1	Two handles by hand	the valve's two effects, separated by experiment
2	The wrong pairing	direction-dependent gain and a sign change
3	Pressure and flow on setpoints	the linear handle, and a fast self-regulating inner loop
4	The cascade	temperature onto pressure, and what a degree costs
5	Circulation is capacity	gain varying by a factor of 16
6	The blend	an interior maximum in the smallest approach
7	Nitrogen to the flare	a stale analysis and an inferential error
8	A hot afternoon	two constraints crossing with no loop misbehaving
9	Throughput on the driver	a low selector with the plant inside the loop
10	Upstream cuts the feed	specific power as the diagnostic
11	Turndown into surge	protective override and what recycle costs
12	Three handles under one driver limit	one predictive plan for three setpoints, priced against the driver
13	A shift on the train (challenge)	380 minutes with four upsets and a heavy makeup target

Table 14.6 Measured Closed-Loop Responses, With the Model's Steady State in Parentheses. Times Are Simulated Minutes From the Step.

Step	Measured response
$P_{s,SP}$ −0.30 bar at $t = 0$, scenario 3	9.3 °C per bar, 63 % of the move in 8.3 min and 95 % in 18.0 min
P_s itself, same step	63 % of the move in 2.2 min
$T_{LNG,SP}$ −155 → −158 °C, scenario 4	after 70 min P_s is 0.32 bar lower and shaft power 1.31 MW higher, about 440 kW per degree, with the product at −157.83 °C
$\dot{m}_{MR,SP}$ 33 → 30 kg/s, scenario 5	product 11.4 °C warmer at 55 min and still moving (15.8 °C, to −139.2 °C)
$\dot{m}_{MR,SP}$ 30 → 36 kg/s	product 17.0 °C colder at 60 min, 5.6 °C below the start (21.7 °C, of which the last 3 kg/s contributes 5.9 °C)
$\dot{m}_{MR,SP}$ 36 → 39 kg/s, model only	(product cools 1.0 °C for 1 MW of shaft power)
Blend 0.420 → 0.360, scenario 6	at 105 min T_{evap} is −154.91 °C and the product −146.16 °C (−154.80 °C and −146.04 °C)
Blend 0.360 → 0.500	warm-end approach down to 0.63 °C at 88 min while the cold end keeps 5.4 °C or more (1.22 °C at 0.500)

essentially unchanged at 5.18, 5.09, 4.96 and 6.97 °C. One handle moved the capacity and spoiled the profile, and the other moved the temperature and left the profile alone.

🕮 **Pause.** Try Exercise 14.1 before reading on.

14.6.2 Closed-Loop Steps

Table 14.6 collects the measured responses of the tuned loops. Every entry is a change read at a stated time in the scenario or a time to a stated fraction of the final change, with the model's steady state beside it where the scenario moves on before the plant settles, so none of them is an accumulated integral. The application also reports integral of absolute error, integral of squared error, overshoot, settling time, actuator travel and saturated fraction on its scorecard, and the two integral measures keep accruing after a transient ends. An IAE figure from this plant is only interpretable with the observation window printed beside it, so the comparisons below are made with peak deviations, settling times, actuator travel and power instead.

The inner loop is faster than the outer one by the margin a cascade needs. Suction pressure reaches 63 % of a setpoint move in 2.2 min against 8.3 min for the product temperature, a ratio of 3.7, and everything between the two is aluminum.

Table 14.7 Measured Disturbance Results. Each Entry Compares the Condition Before the Upset With the Stated Later Instant.

Quantity	Before	After	Scenario
Ambient 18 to 36 °C, feed fixed, read at 170 min			
Discharge pressure (bar)	46.1	63.8	8
Shaft power (MW)	23.1	26.9	8
Driver rating (MW)	26.0	22.7	8
Same afternoon with the low selector, read at 200 min			
Feed valve (%)	100	88.9	9
Production (t/h)	72.0	64.0	9
Shaft power (MW)	23.1	22.7	9
Upstream removes 35 % of the feed, read 40 min later			
Product temperature (°C)	−155.0	−156.2	10
Shaft power (MW)	23.1	17.0	10
Production (t/h)	72.0	46.8	10
Specific power (kWh/kg)	0.321	0.362	10
Turndown to 20 kg/s, read at 60 and 120 min			
Surge margin (%)	45	6.5	11
Recycle valve (%)	0	10.6	11
Recycle shaft power (MW)	0	1.8	11
Product temperature (°C)	−155.0	−80.5	11

14.6.3 What an Inferential Measurement Costs

Scenario 7 is the central result of the chapter. The cascade is in service with the approach controller on circulation, the plant is settled at the design point, and at $t = 15$ min a relief event vents five points of light ends from the refrigerant inventory at once. Nine minutes later, just before the next analysis lands, the refrigerant is boiling at −160.3 °C instead of −162.0 °C, the product has walked warm to −151.9 °C, and the cold-end approach the control room computes has gone *up* to 12.7 °C, because it is still subtracting the boiling temperature the last analysis implies. The estimated minimum approach reads 6.5 °C while the true minimum is 4.2 °C.

🕮 **Pause.** Try Exercise 14.2 before reading on.

Both controllers then do the right thing for the wrong reason. The approach controller sees a wide approach and adds refrigerant, which is the correct action, and the temperature controller pulls the suction pressure down to chase the product back. The analysis at 24 min closes the gap between the estimated and true minimum to under 0.5 °C, but the makeup keeps moving the blend between analyses and the gap reopens to 0.9 °C by 36 min and stays above 0.5 °C until 48 min. By the hour mark the analysis reads 0.407 against a true 0.412 and the plant is back at −155.6 °C, on 23.7 MW instead of 23.1 MW for the same product. Figure 14.4 shows the screen at 26 min, two minutes after the first analysis that follows the event. The application's verification suite reproduces the transient directly and requires the gap between the estimated and true minimum approach to exceed 1.5 °C five minutes after the event and to have fallen by more than half by 45 min. Section 1.4.9 places this inferred constraint among the other measurement limits of the volume.

14.6.4 Constraints, Throughput and Turndown

Table 14.7 collects the disturbance scenarios. Each row reports the settled condition before the upset and the condition at the stated later instant, so the figures are comparisons rather than accumulations.

The eighth scenario is the one with no misbehaving loop in it. Warm air raises the condensing temperature, so the discharge pressure of (14.14) climbs from 46.1 to 63.8 bar and every kilogram costs a sixth more to compress, while (14.15) takes the driver from 26.0 to 22.7 MW. The two

numbers cross going in opposite directions, the temperature controller holds its setpoint by spending power the plant no longer has, and nothing on the temperature trend says so. Scenario 9 hands the feed valve to the selector of (14.19) and the plant settles at 88.9 % feed, 64.0 t/h instead of 72.0 t/h, with the shaft on the derated driver and the product on specification through the afternoon. Getting there is slow by construction. The shaft stays above the derated driver for 76 min after the ambient step, by as much as 1.31 MW, before the power controller brings it under. When a cold front arrives at $t = 200$ min the selector returns the valve on its own and the feed is back at 100 % within a minute, faster than the temperature cascade can follow, so the product runs warm to −146.6 °C at 203.7 min and spends 4.4 min off specification. By 400 min the shaft is at 21.8 MW of an available 27.8 MW.

> **TRY IT IN THE APP** **8 · A hot afternoon**
>
> Load the case and press Run. Touch nothing. Watch the Shaft power trend with its derated driver line. When the air warms from 18 to 36 °C, the driver falls from 26.0 to 22.7 MW while the shaft climbs from 23.1 to 26.9 MW. The LNG outlet temperature trend shows nothing wrong.

Scenario 10 is the case where the trends look perfect and the plant is worse off. Removing 35 % of the feed leaves the refrigerant loop circulating the same 33 kg/s into two thirds of the heat load, so the product runs down onto the boiling temperature and the cold-end approach collapses to nothing. The temperature controller answers by raising the suction pressure, which is exactly right, and forty minutes later the plant is at −156.2 °C on 17.0 MW, drifting slowly toward −155.1 °C and 16.5 MW by the time the feed returns at 140 min. Dividing rather than reading gives the diagnosis. That 17.0 MW is buying 46.8 t/h instead of 72.0 t/h, so specific power has gone from 0.321 to 0.362 kWh/kg, 13 % more energy per tonne with every loop on setpoint, and it is still 10 % above design at 140 min.

> **TRY IT IN THE APP** **10 · Upstream cuts the feed**
>
> Load the case and press Run. Upstream takes 35 % of the feed away. Watch the Power and Rate chips above the plant. Forty minutes after the cut the shaft is down from 23.1 to 17.0 MW, and production from 72.0 to 46.8 t/h, so each kilogram costs 0.362 against 0.321 kWh. Every loop stays on its setpoint.

Scenario 11 separates a machine limit from a process limit. Cutting circulation to 20 kg/s with the anti-surge system bypassed takes the surge margin from 45 to 6.5 %, inside the alarm at 10 % and inside the region where a real machine risks reversing. Restoring the recycle loop at 60 min buys the margin back at 10.6 % valve opening and 1.8 MW of shaft power, all of it spent compressing refrigerant that is throttled straight back to the suction. The product is the other half of the answer. At 20 kg/s the outlet is at −80.5 °C by 120 min and at 16 kg/s it is at −58 °C, so nothing on specification is produced for the entire scenario, which is the real answer to deep turndown on a single-refrigerant train.

The graded challenge runs 380 min from a plant whose makeup target is four points too heavy, whose circulation setpoint is two kilograms low and whose bundle has lost a tenth of its conductance, and applies an afternoon warming from 22 to 33 °C at $t = 30$, a rich feed at $t = 80$, a relief event at $t = 130$ and an evening cooling to 14 °C at $t = 180$ min. The stated target has four parts, and the application judges all four exactly as stated. The product must stay within ±2 °C of −155 °C and the true minimum approach above 1.5 °C over the last two hours, from 260 to 380 min, both read from the true values logged ten times a minute rather than from the thermocouple or the estimate. The shaft must not exceed the derated driver at any integration step of the whole run, and the specific power, shaft energy divided by on-specification product since $t = 0$, must finish below 0.36 kWh/kg. The approach is not judged before 260 min, which matters, because even with the suggested settings the cold end touches the boiling temperature twice in the first 80 min and the relief event crosses it by up to 5.2 °C for 2.5 min. The heavy makeup target and the relief event both act on the blend rather than on any loop, and the blend is the part no controller handles.

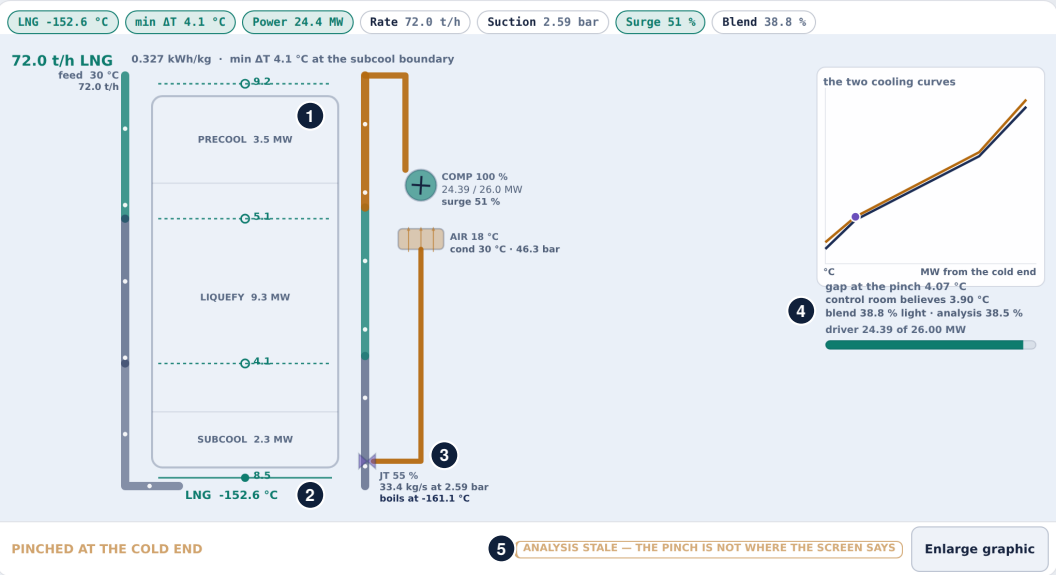

Figure 14.4 The stale-analysis transient of scenario 7 at $t = 26$ min, eleven minutes after the relief event took five points of light ends out of the refrigerant, with the cascade in service and the cold-end approach controller holding the circulation setpoint at 33.34 kg/s. A fresh analysis has since landed, so the readout beside the inset reads 38.8 % light against an analysis of 38.5 % and a true pinch of 4.07 °C against the 3.90 °C the control room computes for it. The product temperature has walked warm to −152.57 °C against a −155.00 °C setpoint and a −150.00 °C tank limit, and the temperature controller has pulled the suction pressure setpoint down to 2.563 bar after it. The cold box animation marks the four approach stations, three of them drawn hollow because no instrument exists there (1), and the one measured station at the cold end (2). The Joule-Thomson valve carries its opening, the circulation and the boiling temperature (3), the inset plots the two cooling curves against cumulative duty with the true and the computed pinch beside it (4), and the status line warns that the analysis is stale (5).

Each scenario is credited when the reader performs the actions its completion rule lists and the simulation reaches the stated state. The challenge is judged on the case as given, so changing a model parameter or a disturbance during the attempt voids it. With the shipped settings (pressure and flow loops only, weak tunings, makeup target 0.380) every part fails. The worst product error over the last two hours is 31.2 °C, the smallest approach 0.47 °C, the shaft spends 149.1 min above the derated driver and the specific power ends at 4.39 kWh/kg. With the suggested settings (the full structure of Figure 14.3 with the default tunings, makeup target 0.420 and a 20.5 MW power limit) every part passes, with a worst error of 0.49 °C, a smallest approach of 3.84 °C, no time over the driver (0.59 MW of margin at the closest) and 0.338 kWh/kg. Revealed before Run or carried through Reset, they give the same run.

14.6.5 Three Handles Under One Driver Limit

Scenario 12 is the afternoon of scenario 9, ambient 18 to 36 °C at 25 min and a cold front at 200 min with the selector in service, run first with the MPC and then, after Reset, with the cascade and its approach controller. Table 14.8 compares the structures on scenarios 12, 8, 10 and 7 and the challenge, the MPC on the settings of scenario 12 without retuning. In scenario 12 the model-information row gives the cascade the MPC's steady-state blend target, and the hindsight row is the best of 20 fixed

Table 14.8 Cascade Against the MPC, Default Seed, No Noise, Full Duration of Each Case

Structure	Over driver (min)	Peak over (MW)	Excess (MW min)	Under 1.5 °C (min)	Smallest ΔT_{min} (°C)	On-spec LNG (t)	Off spec (min)	LNG IAE (°C min)	w (kWh/kg)	Smallest SM (%)
Scenario 12, 400 min, hot afternoon and selector, 800 MPC plans										
MPC	18.6	2.78	20.4	41.5	−0.37	455.1	4.6	274.0	0.320	14.8
cascade	85.1	1.31	28.9	61.5	−0.16	450.5	4.4	275.3	0.329	18.6
cascade, MPC's blend target[a]	49.5	1.30	21.9	45.4	−0.14	455.1	3.7	321.5	0.323	16.8
hindsight cascade[b]	53.3	0.37	2.4	12.3	−0.50	457.5	1.5	140.6	0.312	16.2
Scenario 8, 260 min, feed on the operator, 520 plans per MPC row										
MPC	150.0	2.10	25.5	0.0	2.12	312.0	0.0	602.0	0.315	15.5
cascade	150.8	4.43	620.5	0.0	4.09	312.0	0.0	20.2	0.351	30.3
MPC, no power limit[c]	150.7	4.18	618.2	0.0	4.13	312.0	0.0	4.1	0.351	31.1
Scenario 10, 240 min, feed cut and return, 480 MPC plans										
MPC	0.0	0.00	0.0	36.4	−1.09	215.9	18.1	587.7	0.360	15.4
cascade	0.0	0.00	0.0	152.6	−0.59	230.4	6.0	337.4	0.337	25.9
cascade at 21.0 kg/s[d]	0.0	0.00	0.0	175.5	−0.59	118.7	99.1	1446.4	0.656	12.9
Scenario 7, 180 min, stale analysis, 360 MPC plans										
MPC	0.0	0.00	0.0	2.8	1.03	216.0	0.0	62.0	0.323	44.9
cascade	0.0	0.00	0.0	0.2	1.06	216.0	0.0	88.9	0.326	44.9
Scenario 13, 380 min, challenge, 760 plans per MPC row										
MPC, suggested, passes	0.0	0.00	0.0	77.3	−4.32	401.3	17.2	908.6	0.325	14.2
cascade, suggested, passes	0.0	0.00	0.0	177.2	−5.20	388.8	0.0	239.5	0.338	18.7
MPC, rest as shipped, fails	150.0	3.03	317.9	293.7	0.00	358.6	81.2	1854.9	0.408	23.6

Over driver counts every physics step with the true shaft power above the derated driver, and the approach columns read the true minimum. SM is the smallest surge margin. No MPC plan was rejected. [a]Makeup target set to the MPC's blend target as it moved, no plan. [b]Best by time under 1.5 °C of makeup 0.40 to 0.48, shipped or twice-as-fast temperature and approach loops, approach setpoint 5 or 7 °C, here 0.44, fast, 5 °C. [c]A copy with no power row or price. [d]Circulation forced to 21.0 kg/s over 100 to 139.5 min.

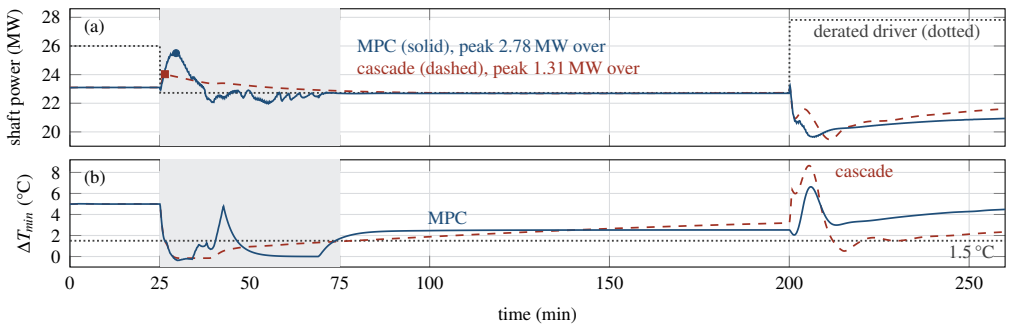

Figure 14.5 Scenario 12 with the MPC that the scenario loads and with the cascade selected after Reset, default seed, with the window from 25 to 75 min shaded. (a) The true shaft power against the derated driver. The MPC is over for 18.6 min with a peak of 2.78 MW and 20.4 MW min, and the cascade for 85.1 min with a peak of 1.31 MW and 28.9 MW min. (b) The true minimum approach against the 1.5 °C warning. In both panels the MPC is solid, the cascade dashed and the limit dotted, and the dot and the square mark the two peaks.

cascade settings chosen after the runs. Figure 14.6 sets the MPC and cascade rows of the five cases side by side.

Under the MPC the shaft is over the derated driver for 18.6 against 85.1 min, all of it before 73.2 min, but it goes further over, with a peak of 2.78 against 1.31 MW and 8.6 against 6.2 min more than 1 MW over, while the excess is 20.4 against 28.9 MW min, as shown in Figure 14.5(a). Most of

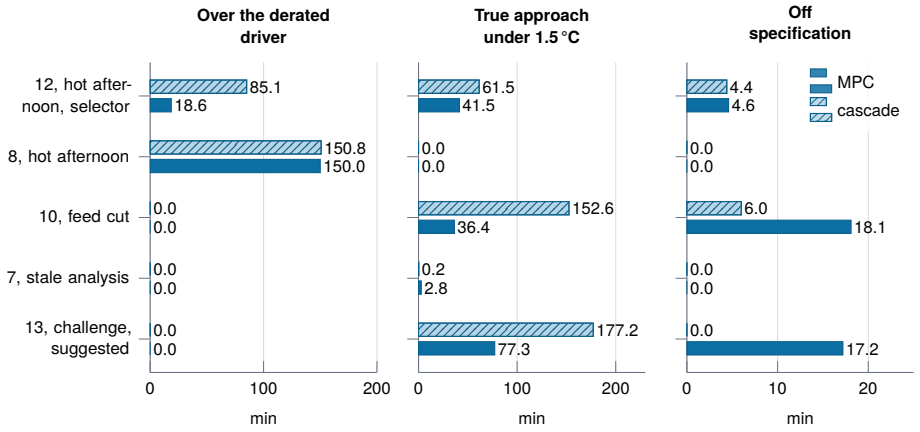

Figure 14.6 The MPC and the cascade of Table 14.8 on the five cases, labeled by scenario number, solid bars for the MPC and hatched bars for the cascade, with the table's value beside each bar. Notice that the MPC cuts the time over the driver markedly only in scenario 12, and that it spends more time off specification than the cascade in scenarios 10 and 13 while spending less time with the approach under 1.5 °C. The other rows of the table are not drawn.

the cascade's time is the selector holding the shaft just above its limit, more than 0.1 MW over for only 55.1 min. The deeper excursion is the price order at work. In the seven minutes after the ambient step the plan raises the circulation setpoint from 32.5 to 36.0 kg/s and the makeup target from 0.42 to 0.55 to hold its predicted internal approaches, accepting up to 1.84 MW of power slack, the selector cuts the feed to 80.2 % against 87.0 %, and the shaft reaches 25.50 MW against 24.03 MW. With the approach floor at 10^2 or the power row at 10^4 the peak falls to 2.11 or 2.09 MW, but neither variant ships, because repricing after measuring would be tuning the case.

The true minimum approach is under 1.5 °C for 41.5 against 61.5 min, crossing at the cold end to −0.37 °C against −0.16 °C near 30 min, and the plant makes 455.1 against 450.5 t of LNG at 0.320 against 0.329 kWh/kg. The LNG IAE is nearly equal, 274.0 against 275.3 °C min, but 223.9 against 121.9 in the afternoon and 48.8 against 120.8 after the cold front, and the MPC does not hold the setpoint. Its target sits more than 0.5 °C warm for 166.5 of the 400 min, as warm as −151.00 °C, because near the selector's limit a tenth of a megawatt or a shortfall of predicted approach is worth more to it than a degree of LNG.

The production gain is the model's blend and not the plan, because the cascade given only the MPC's blend target makes the same 455.1 t, with 49.5 min over the driver at a peak of 1.30 MW and 21.9 MW min. What the plan adds is time, 18.6 against 49.5 min over the driver and 41.5 against 45.4 min under 1.5 °C, bought with the deeper peak. The hindsight cascade beats the MPC on everything but minutes over the driver, with 457.5 t, 12.3 min under 1.5 °C and 53.3 min over at a peak of 0.37 MW and 2.4 MW min. The ranking holds at half the integration step, 17.8 min over at the same 2.78 MW peak and 20.4 MW min, and at three seeds with instrument noise, where the MPC peaks at 2.65 to 2.92 MW over the driver against 1.28 to 1.34 MW.

Scenario 8 is the loss the design accepts. With the feed on the operator only the power row protects the driver, and the MPC holds the LNG at its ceiling, −151.01 °C on average from 40 to 170 min, an IAE of 602.0 against 20.2. Both structures spend about 150 min over the driver, but the MPC is more than 0.1 MW over for only 23.9 of them, at a peak of 2.10 against 4.43 MW and 25.5 against 620.5 MW min, and handed no power limit it tracks the LNG with an IAE of 4.1, so the whole temperature loss is the price of the power row. In scenario 10 the MPC matches the circulation to the reduced feed, 21.0 kg/s against 29.2 kg/s just before the feed returns unannounced at 140 min, so the bundle is then short of refrigerant, the LNG warms to −119.3 °C against −139.3 °C, and the

MPC ends 18.1 against 6.0 min off specification and 14.5 t behind. The cascade forced to the same 21.0 kg/s does worse, 99.1 min off specification, so the loss is the circulation chosen by a controller that sees the feed only through its load estimate, and its surge margin falls to 15.4 against 25.9 %, above the 10 % alarm. In scenario 7 the stale prediction costs 2.8 against 0.2 min under 1.5 °C.

DID YOU KNOW? Less Refrigerant, Nearly Twice the Energy per Tonne

From the simulator (Table 14.8). In scenario 10 a cascade made to circulate only 21.0 kg/s through the feed cut ended at 0.656 kWh/kg, against 0.337 kWh/kg for the cascade left alone. Less refrigerant drew less power, and it made 118.7 t of on-specification LNG against 230.4 t.

With the challenge's suggested settings and the MPC selected the target is met, at a worst error of 1.69 °C against the cascade's 0.49 °C, but limits the rule does not score are broken on the way. The warm target keeps the product off specification for 17.2 min where the cascade has none, the true approach is under 1.5 °C for 77.3 against 177.2 min, and at the instant of the relief event at 130 min the cold end crosses to −4.32 °C under the MPC and −5.20 °C under the cascade. Half a minute later the MPC's model, on the analysis taken before the event, still put the cold end at 1.92 °C against a true −3.19 °C. With every other setting as shipped the MPC fails, at 7.05 °C with 150.0 min over the driver at a peak of 3.03 MW and 317.9 MW min.

14.7 Technology in Practice

Figure 14.7 draws the five loops of Section 14.5 and the anti-surge loop with the hardware of a liquefaction train in place of the schematic's symbols.

What Measures It

The loops read four instruments and one analyzer. A thermocouple in the outlet nozzle gives the product temperature, a transmitter the suction pressure, a flowmeter the circulation and a transducer the shaft power. The gas chromatograph on the suction line reports the blend, and the cold-end approach that the approach controller holds is computed from that analysis, the suction pressure and the product temperature. The other three stations of the bundle have no instrument at all.

TECHNOLOGY Sensor: The Thermocouple in the LNG Outlet

A thermocouple is a junction of two unlike metals whose voltage depends on its temperature. It is read against a second, reference junction, so it measures a difference and the instrument must compensate for the reference [23]. Thermocouples are more rugged than resistance thermometers and span more temperature, though usually less accurate. Copper against constantan, type T, covers −300 to 700 degrees Fahrenheit and resists both oxidizing and reducing gases well at low temperatures, even wet. The most common failure is to fail open, called burnout. An open input picks up noise that reads as a wildly varying temperature, so the instrument is set to drive high or low on burnout, whichever is safer [23]. The chapter gives the product thermocouple a 0.35 min filter.

What Moves It

Four final elements and one slow system carry the handles. The Joule-Thomson valve sets the circulation, the compressor speed sets the suction pressure, the feed valve sets the heat load under the low selector, and the recycle valve keeps the machine away from surge. The makeup system moves the blend, at no more than 0.010 of light fraction per minute. Behind the speed stands the driver, whose rating falls as the air warms, and that rating is the limit the selector defends.

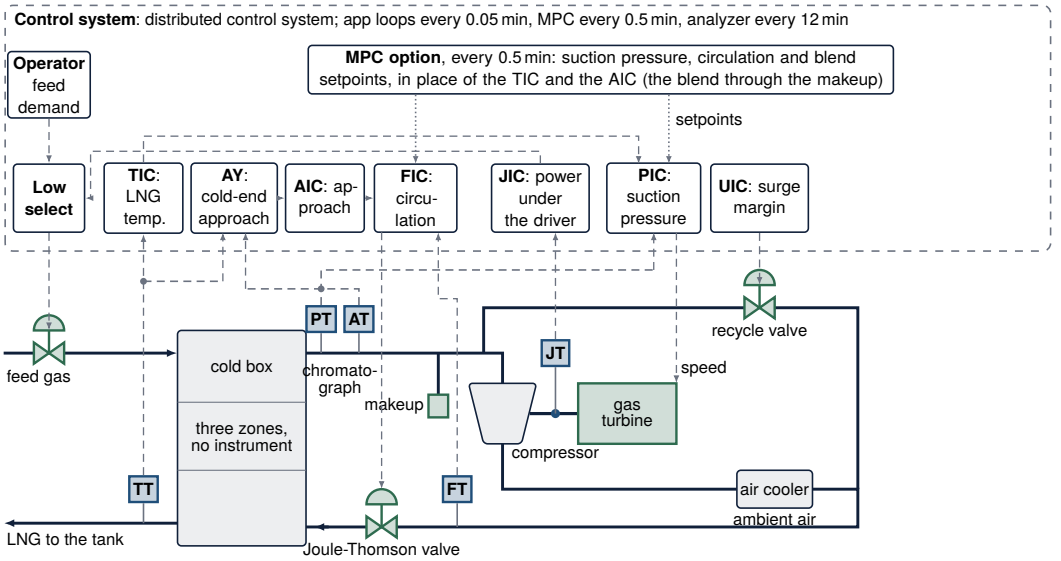

Figure 14.7 The loops of the train with their hardware. Lettered boxes are sensing devices with the chapter's tags: TT on the LNG outlet, PT and AT on the suction line, FT on the circulation and JT on the shaft. Green symbols are final elements: the Joule-Thomson valve, the gas turbine that sets the compressor speed, the feed valve, the recycle valve and the makeup of the blend. Dashed arrows are signals, dotted arrows the three setpoints that the MPC option writes in place of the temperature and approach controllers. Notice that the approach controller reads an analyzer that reports every 12 min, and that no instrument stands inside the cold box. Not to scale, and not a piping diagram.

> **TECHNOLOGY** **Final Element: The Gas Turbine on the Compressor Shaft**
>
> In oil and gas service a gas turbine often turns a compressor directly, with no generator between. Its own compressor squeezes the inlet air, fuel heats it, and the expansion turbine drives both that compressor and the load [24]. Warm air is less dense, so less of it passes, and it takes more work to compress, so power and efficiency both fall as the inlet warms. Ratings are quoted at sea level and 59 degrees Fahrenheit, and inlet-air cooling by refrigeration, a water spray or stored cold recovers part of the loss. At part load the turbine inlet temperature is lowered, which costs efficiency too [24]. The chapter keeps a speed lag of 0.30 min and a rating that falls with the air, from 26.0 to 22.7 MW between 18 and 36 °C.

Where the Controller Runs

A distributed control system controls production within one location, and its supervisory and lower-level loops run with cycle times from milliseconds to minutes [25]. Control of a process plant is usually layered by time scale, with supervisory and predictive control acting in minutes and regulatory control in seconds [3]. The app's loops execute every 0.05 min, its MPC every 0.5 min, and its analyzer reports every 12 min, so the chapter spans all three of those speeds. Its structure maps onto standard blocks: a temperature PI cascaded onto a pressure PI on speed, a flow PI on the valve under an approach PI that reads a calculation block, a power PI and the operator's demand into a low selector with tracking, and an anti-surge PI on the recycle valve.

> **IN PRACTICE What the Federal LNG Rules Ask of the Controls**
>
> United States rules for LNG facilities require an attended control center from which operations and warning devices are watched, and from which every remotely actuated and automatic shutdown control can be worked [26]. Operations are watched for gas, temperature, pressure, vacuum and flow alarms [27]. Each control system must be adjusted to operate within design limits, one out of service for 30 days or more is tested before it returns, and those normally in operation in a base load plant are inspected and tested every calendar year, at intervals of 15 months at most [28]. Cooldown must keep thermal stresses within design limits, and the cold piping is then checked for leaks [29]. The chapter starts every scenario cold.

What Is Modernizing It

Three developments touch this chapter. The first is the drive. Liquefaction uses a share of the feed itself, and the one United States export plant whose refrigerant compressors run on electric motors converts most of its feed gas into LNG [10]. New plants are increasingly driven by electricity under a power management system, and a modern drive can cut the power to the process almost instantaneously [30]. The second is the estimate. Soft sensors are used beside online analyzers, among other things to validate them, and a survey of industrial applications names model maintenance as their most critical issue [31]. The moving horizon estimator of Section 14.9 is the chapter's version of that idea. The third is the predictive layer: receding-horizon control is established in industry [20] and has been proposed in adaptive form for natural gas liquefaction [21], and scenario 12 shows that the trade it makes is set by its prices.

14.8 Limitations

Sizing an exchanger or a machine from this model is not a legitimate use of it, because the model is a control study and not a process simulator. The compositions are not a property package, the refrigerant is two numbers rather than four components with a phase envelope, and the boiling temperature relation of (14.9) is a linear fit that is credible only over the range the sliders allow. Three lumped zones cannot put the pinch anywhere except at a zone boundary or at an end, which is the simplification that most limits what the application can say about a real coil-wound bundle, where the tightest point usually sits inside a zone. The zone heat capacities are constants fixed by the design temperature ranges, so far from design the duty split between zones stops being physical. The compressor is a single line with a surge line and no choke, and the surge reversal itself is not simulated, so scenario 11 reports the approach to the limit and not what happens after it. Volume 4, Chapter 2 simulates the surge cycle itself, with a protective override on a recycle valve. The gas turbine is a rating with a linear derate rather than a thermodynamic model, which is adequate for showing two constraints crossing and inadequate for predicting the crossing point on a particular machine. Cool-down from ambient is absent, and the sixteen to thirty hours it takes is the operation that decides whether a train is shut down at all on low feed. The analyzer is a perfect sample and hold, so the failure modes that cause most real inferential errors, a plugged sample line or a calibration that has drifted, cannot be demonstrated. Perfect repeatability is itself a departure from practice, since one frozen step and one frozen scan hand back the same trajectory everywhere while a distributed control system never does. The features that survived the simplification, the two-effect valve, the fast suction pressure inventory, the metal inertia, the interior maximum in the blend, the sampled composition and the two constraints that move together with the weather, all survived because a control engineer has to design around each of them. Section 14.3.9 states where each individual simplification stops being true.

The predictive controller inherits the inferential problem rather than solving it. Its approach floor is on predictions from the last analysis, wrong by up to 5.37 °C in scenario 7, and the effect of the blend on the boiling temperature of (14.9) is taken from that analysis carried forward by the makeup law, with nothing inferring the blend between analyses. Its model holds the declared parameters, so with the bundle fouled to 0.85 its warm-end approach is off by up to 1.78 °C. It has no feed measurement

and no model of the selector, its power and surge limits are soft with one slack per window, and its prices rather than its setpoints decide what it gives up.

> **OPEN CHALLENGE** **Running On the Power That Is Left**
>
> New LNG plants are increasingly driven by electric motors under a site power management system. A 2022 steady-state study found that a mixed fluid cascade plant could keep running on about 30 % less power, 60 of 200 MW, instead of shutting down. Its model left the transient open: the LNG could warm toward saturation and the refrigerant compressors move toward surge while the plant moves to its new point. Progress, as the authors set it out, means tests on an industrial simulator and in plant experiments, and a study of the dynamics and of the need for more complex control structures [30]. The chapter's low selector meets a slower form of the limit in scenario 9.

14.9 Discussion

The economic argument on a liquefaction train is a ratio and not a rate. At the design point the plant turns 20.0 kg/s of feed into 72.0 t/h of LNG for 23.10 MW of shaft power, which is 0.321 kWh/kg, and that ratio is what a percent of efficiency is measured against. Scenario 10 moves it to 0.362 kWh/kg without a single loop leaving its setpoint, and scenario 7 moves it from 23.1 to 23.7 MW for the same product while the approach display reads better than the truth. Both of those are invisible on the controlled variables and both are visible in the ratio, which is the argument for putting specific power on the wall of a control room rather than production. The same reasoning is why the LNG operation literature evaluates candidate controlled variables by the steady-state loss they cause rather than by their closed-loop response [3–5].

Methods beyond the implemented controllers fit the same equations. Sixteen states with an active driver constraint and a bounded approach make a natural nonlinear model predictive control problem, because the controller can trade the approach against the driver limit directly instead of relying on a selector to notice the collision, and the successively linearized controller of Section 14.5.5 shows in scenario 12 that the trade it makes is set by its prices. Moving horizon estimation on the same equations, using the sampled composition as an infrequent measurement alongside the continuous temperatures, is the standard remedy for the error that scenario 7 puts on the screen, and it is the direction the soft-sensor literature takes for exactly this situation [1, 2]. An economic formulation would place the value of on-specification LNG against the cost of shaft power, both of which the application already carries as adjustable parameters, and would decide the feed rate and the approach setpoint together rather than separately.

14.10 Conclusion

This chapter studied a single mixed refrigerant liquefaction train as an interactive laboratory for inferential constraint control. The simulator integrates sixteen states covering three counter-current effectiveness-NTU zones, the bundle metal as a lag of 4.44, 4.55 and 6.65 min on the three zone duties, a self-regulating suction pressure inventory, four actuator positions and the refrigerant light fraction, and it reports the true minimum internal temperature approach beside the estimate a control room would compute from an analysis that is refreshed every 12 min. Swept open loop with the compressor speed held, the product temperature reaches a minimum of −159.26 °C at 60 % valve travel, giving 3.4 °C per percent in one direction, 0.27 °C per percent in the other, and a sign change in between. Closing a suction pressure loop removes one of the valve's two effects and leaves a handle worth 9.2 °C per bar that is straight across the working range, with the pressure reaching 63 % of a setpoint move in 2.2 min against 8.3 min for the product temperature. Circulation is a capacity variable whose gain falls by a factor of 16 between 30 and 39 kg/s, and the refrigerant blend is an interior maximum at $x = 0.419$ rather than a setpoint. A relief event that removes five points of light ends sends the computed cold-end approach up to 12.7 °C while the true minimum falls to 4.2 °C, and the plant recovers on 23.7 MW instead of 23.1 MW for the same product. An optional model

predictive controller spends 18.6 rather than 85.1 min over the derated driver in scenario 12 but goes further over it, 2.78 MW at its peak against 1.31 MW with 20.4 against 28.9 MW min, owes its 4.6 t of extra LNG to the blend its model chooses, and loses scenario 10 with 18.1 against 6.0 min off specification. A fuller model would add a continuous bundle discretization so the pinch can sit inside a zone, a component-wise makeup system, an analyzer failure mode, and a moving horizon estimator to be compared against the sample-and-hold inference. The liquefaction train puts a plant whose limiting variable is computed rather than measured within reach of anyone with a browser.

KEY POINTS

- One valve sets both how much refrigerant circulates and how warm it boils, and the two effects pull the LNG temperature opposite ways.
- A suction pressure loop takes one effect away and leaves a handle worth 9.2 °C per bar, straight across the working range.
- Circulation is a capacity variable, and the blend is an interior maximum rather than a setpoint.
- A relief event sends the computed cold-end approach up to 12.7 °C while the true minimum falls to 4.2 °C.

Where else this applies. A valve whose gain changes sign returns in Volume 6, Chapter 2, and quality held through a late analysis in Volume 5, Chapter 3.

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] 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] J. B. Jensen and S. Skogestad, “Optimal operation of a simple LNG process,” *IFAC Proceedings Volumes*, vol. 39, no. 2, pp. 241–246, 2006.

[5] F. A. Michelsen, B. F. Lund, and I. J. Halvorsen, “Selection of optimal, controlled variables for the TEALARC LNG process,” *Industrial & Engineering Chemistry Research*, vol. 49, no. 18, pp. 8624–8632, 2010.

[6] 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.

[7] B. Austbø, S. W. Løvseth, and T. Gundersen, “Annotated bibliography—use of optimization in LNG process design and operation,” *Computers & Chemical Engineering*, vol. 71, pp. 391–414, 2014.

[8] W. Lim, K. Choi, and I. Moon, “Current status and perspectives of liquefied natural gas (LNG) plant design,” *Industrial & Engineering Chemistry Research*, vol. 52, no. 9, pp. 3065–3088, 2013.

[9] S. Kumar, H.-T. Kwon, K.-H. Choi, W. Lim, J. H. Cho, K. Tak, and I. Moon, “LNG: An eco-friendly cryogenic fuel for sustainable development,” *Applied Energy*, vol. 88, no. 12, pp. 4264–4273, 2011.

[10] U.S. Energy Information Administration, “Liquefied natural gas,” Natural Gas Explained, last updated June 21, 2024, 2024, accessed November 18, 2026. [Online]. Available: https://www.eia.gov/energyexplained/natural-gas/liquefied-natural-gas.php

[11] G. Venkatarathnam, “Natural gas liquefaction processes,” in *Cryogenic Mixed Refrigerant Processes*, ser. International Cryogenics Monograph Series. New York, NY: Springer, 2008, pp. 149–220.

[12] W. J. Podbielniak, “Art of refrigeration,” U.S. Patent 2,041,725, filed July 14, 1934, granted May 26, 1936, 1936. [Online]. Available: https://patents.google.com/patent/US2041725A/en

[13] J. C. Perret, “Method and apparatus for the cooling and low temperature liquefaction of gaseous mixtures,” U.S. Patent 3,364,685, filed March 16, 1966, granted January 23, 1968, 1968. [Online]. Available: https://patents.google.com/patent/US3364685A/en

[14] L. S. Gaumer, Jr. and C. L. Newton, “Combined cascade and multicomponent refrigeration system and method,” U.S. Patent 3,763,658, filed January 12, 1970, granted October 9, 1973, 1973. [Online]. Available: https://patents.google.com/patent/US3763658A/en

[15] B. Linnhoff and E. Hindmarsh, “The pinch design method for heat exchanger networks,” *Chemical Engineering Science*, vol. 38, no. 5, pp. 745–763, 1983.

[16] G. C. Lee, R. Smith, and X. X. Zhu, “Optimal synthesis of mixed-refrigerant systems for low-temperature processes,” *Industrial & Engineering Chemistry Research*, vol. 41, no. 20, pp. 5016–5028, 2002.

[17] A. Singh and M. Hovd, “Dynamic modeling and control structure design for a liquefied natural gas process,” in *2007 American Control Conference*, 2007, pp. 1347–1352.

[18] S. Huang, M. Chapeaux, and B. Martinez, “Strategies for mitigating impacts on LNG plant capacity by variations in ambient temperature,” in *Proceedings of the 2nd Annual Gas Processing Symposium*. Elsevier, 2010, pp. 455–463.

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

[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] W. Won and K. S. Lee, “Delta-operator-based adaptive model predictive control and online optimization of a natural gas liquefaction process,” *Chemical Engineering Science*, vol. 162, pp. 21–32, 2017.

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

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

[24] U.S. Environmental Protection Agency Combined Heat and Power Partnership, “Catalog of CHP technologies, section 3. technology characterization: Combustion turbines,” U.S. Environmental Protection Agency, Washington, DC, Tech. Rep., 2017, march 2015, revised September 2017. [Online]. Available: https://www.epa.gov/sites/default/files/2015-07/documents/catalog_of_chp_technologies_section_3._technology_characterization_-_combustion_turbines.pdf

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

[26] Pipeline and Hazardous Materials Safety Administration, “49 CFR 193.2441, liquefied natural gas facilities: Control center,” Code of Federal Regulations, Title 49, Part 193, edition of October 1, 2025, govinfo.gov, 2025.

[27] ——, “49 CFR 193.2507, liquefied natural gas facilities: Monitoring operations,” Code of Federal Regulations, Title 49, Part 193, edition of October 1, 2025, govinfo.gov, 2025.

[28] ——, “49 CFR 193.2619, liquefied natural gas facilities: Control systems,” Code of Federal Regulations, Title 49, Part 193, edition of October 1, 2025, govinfo.gov, 2025.

[29] ——, “49 CFR 193.2505, liquefied natural gas facilities: Cooldown,” Code of Federal Regulations, Title 49, Part 193, edition of October 1, 2025, govinfo.gov, 2025.

[30] I. Y. Xuan, C. Skourup, J. B. Jensen, T. Haugen, and N. F. Thornhill, “Flexible operation of a mixed fluid cascade LNG plant for electrical power management,” *Energy*, vol. 250, p. 123794, 2022.

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

14.11 Exercises

14.1 Concept check · 2 min

Scenario 1 starts in Manual at the design line-up. Five minutes after Run the Joule-Thomson valve is opened from 54.9 to 62 % with the compressor speed left at 92.0 %. Twenty minutes later the circulation has risen by 4.6 kg/s, yet the product is only 4.51 °C colder and sits on the refrigerant boiling temperature. Why?

(a) The flow loop holds the circulation at its setpoint, so most of the extra opening is lost

(b) The extra refrigerant reaches the product through the bundle lag, so most of the gain is still to come

(c) The extra vapor raises the suction pressure, so the refrigerant boils warmer and the product rests on it

(d) The air cooler warms with the extra flow, so the discharge pressure falls and the valve passes less

14.2 Concept check · 2 min

In scenario 7 a relief event at 15 min vents five points of light ends. At 24 min, just before the next analysis lands, the cold-end approach the control room computes reads 12.7 °C, wider than its 7 °C setpoint, while the true minimum approach has fallen to 4.2 °C. What makes the computed approach move the wrong way?

(a) The product thermocouple lags by 0.35 min, so the computed approach trails the true one

(b) The approach controller has raised the circulation, and the extra refrigerant has opened the cold end

(c) The suction transmitter reads high after the event, so the estimate puts the boiling point too low

(d) The estimate subtracts a boiling temperature from the last analysis, which predates the vented light ends

14.3 Concept check · 2 min

At a suction pressure held at 2.90 bar, circulation steps of 3 kg/s from 30 to 33, 33 to 36 and 36 to 39 kg/s are worth 15.8, 5.9 and 1.0 °C of product temperature. Which role does this chapter give the circulation?

(a) A capacity setting that carries the duty, with the suction pressure as the temperature trim

(b) The temperature handle, with a gain schedule that spans the factor of 16 in its gain

(c) A ratio to the feed rate with no feedback, since the duty follows the feed

(d) The surge protection handle, held at the flow that gives the widest surge margin

14.4 Theory and calculation · 5 min

Equation (14.9) gives the refrigerant boiling temperature with $T_{evap,0} = -162.0\,°\mathrm{C}$, $k_P = 9.5\,°\mathrm{C}$ per bar, $k_x = 120\,°\mathrm{C}$ per unit of light fraction, $P_{s,0} = 2.90\,\mathrm{bar}$ and $x_0 = 0.420$.

(a) In scenario 1 the valve step raises the suction pressure from 2.901 to 3.167 bar at the design blend. Compute the boiling temperature before and after. *Checkpoint:* −161.99 °C and −159.46 °C.

(b) A relief event takes 0.05 of light fraction out of the blend at 2.90 bar. Compute the new boiling temperature. *Checkpoint:* −156.0 °C, 6.0 °C warmer.

(c) Find the suction pressure that brings the boiling temperature back to −162.0 °C at $x = 0.370$, and say whether the temperature controller's output limits of Table 14.4 allow it. *Checkpoint:* 2.27 bar, inside the 1.9 to 5.2 bar range.

14.5 Theory and calculation · 6 min

Zone 1 (precool) runs at the design point with $\dot{m}_{NG} = 20.0\,\mathrm{kg/s}$, $\dot{m}_{MR} = 33.0\,\mathrm{kg/s}$, $c_{g,1} = 2.60\,\mathrm{kJ/kg/°C}$, $c_{r,1} = c_{m,1} = 1.576\,\mathrm{kJ/kg/°C}$ at $x = x_0$, $UA_1 = 676\,\mathrm{kW/°C}$, $C_{w,1} = 180000\,\mathrm{kJ/°C}$, $T_{feed} = 30\,°\mathrm{C}$ and a refrigerant inlet $T_{r,2} = -40.0\,°\mathrm{C}$.

(a) Compute the two heat-capacity rates and their ratio $C^*_{r,1}$, and say which branch of Equation (14.3) applies. *Checkpoint:* 52.0 and 52.0 kW/°C, $C^* = 0.9998$, the balanced branch.

(b) Compute NTU_1, ε_1 and the duty of Equation (14.4), then the refrigerant leaving the zone and the warm-end approach. *Checkpoint:* 13.0, 0.929, 3380 kW, then 25.0 °C and 5.0 °C.

(c) Compute the metal lag τ_1 of Equation (14.5) at $f_b = 1$ and for the challenge's bundle at $f_b = 0.9$. *Checkpoint:* 4.44 min and 4.93 min.

14.6 Theory and calculation · 7 min

The ambient air steps from 18 to 36 °C, as in scenario 8. Use Equations (14.14), (14.15) and (14.13) with the constants of Table 14.3, $f_{ac} = 1$ and the compressor flow held at 33.0 kg/s.

(a) Compute the derated driver rating W_{max} at 36 °C. *Checkpoint:* 22.72 MW.

(b) Compute the steady condensing temperature and the discharge pressure at 18 and at 36 °C. *Checkpoint:* 30.16 °C and 46.13 bar, then 48.16 °C and 63.78 bar.

(c) With the suction pressure at 2.90 bar and the suction temperature unchanged, find the factor by which the shaft power rises, and the shaft power that follows from the design 23.10 MW. *Checkpoint:* 1.164, 26.9 MW against 22.7 MW available.

14.7 App challenge · 8 min

Load scenario 1, Two handles by hand (Manual), default seed, at the default speed of 2. Press Run, pause at any time from 5 to 7 min, set the JT valve to 62 %, and run on until 20 min after the move. Export the CSV file.

(a) In the row 20 min after the move, read the circulation `mr_kgs`, the suction pressure `psuc_bar`, the boiling temperature `tev_C`, the product `lng_C` and the cold-end approach `app_cold_C`. *Checkpoint:* 37.63 kg/s, 3.167 bar, −159.46 °C for both temperatures and 0.00 °C, each within 0.02.

(b) Press the reload button, run again, and in the same window set Compressor speed to 96.0 % instead, leaving the valve at 54.9 %. Read `psuc_bar` and `lng_C` 20 min after the move, and compute the product change per bar of suction change from the first row (2.901 bar, −154.95 °C). *Checkpoint:* 2.735 bar and −156.60 °C, about 10 °C per bar.

(c) Compare the warm-end approach `app_warm_C` 20 min after each move with its starting 4.99 °C. *Checkpoint:* 17.76 °C after the valve move and 5.18 °C after the speed move, within 0.1 °C.

14.8 App challenge · 8 min

Load scenario 7, Nitrogen to the flare, and a stale analysis, and change no setting. Run to 60 min at the default speed with the default seed and export the CSV file.

(a) In the row at 24.0 min, read `tev_C`, `lng_C`, the computed cold-end approach `appr_meas_C`, `mita_est_C`, `mita_C`, `x_light` and `x_gc`. *Checkpoint:* −160.33 °C, −151.89 °C, 12.67 °C, 6.52 °C, 4.22 °C, 0.385 and 0.420, temperatures within 0.02 °C.

(b) In the same row, compare $k_x(x_{GC} - x)$ with the difference between `appr_meas_C` and the true cold-end approach `app_cold_C`. *Checkpoint:* 4.19 °C against $12.67 - 8.44 = 4.23$ °C.

(c) Read `appr_meas_C` and `x_gc` in the row at 24.1 min, and the shaft power `power_kW` at 15 and 60 min. *Checkpoint:* 8.49 °C and 0.385, then 23.10 MW and 23.72 MW.

14.9 App challenge · 10 min

Load scenario 12, Three handles under one driver limit, which opens with Control structure on MPC and Throughput on the driver. Run 400 min at speed 20 with the default seed and export the CSV file. Press Reset, set Control structure to Cascade before Run, run the same 400 min and export again.

(a) For each file, read `metric_overpower_min` at 400 min. Then take the excess of `power_kW` over `driver_kW` in every row, find its largest value, and sum its positive values times the 0.1 min row spacing. *Checkpoint:* MPC 18.6 min, 2.78 MW, 20.4 MW min, and cascade 85.1 min, 1.31 MW, 28.9 MW min, sums within 0.1 MW min.

(b) Read `metric_pinch_min`, the minutes with the true minimum approach under 1.5 °C, and `metric_lng_t` at 400 min. *Checkpoint:* 41.5 min and 455.1 t under the MPC, 61.5 min and 450.5 t under the cascade.

(c) The three driver measures rank the structures differently. Say why, and which one a plant manager should be shown.

14.10 Application area · 4 min

Table 14.3 prices on-specification LNG at 0.44 $/kg and shaft power at 0.030 $/kWh.

(a) At the design specific power of 0.321 kWh/kg, what does the shaft power cost per kilogram of LNG, and what share of the LNG's value is that? *Checkpoint:* 0.0096 $/kg, 2.2 %.

(b) In scenario 9 the low selector holds the feed on the derated driver. At the design specific power, how much LNG per hour does one megawatt of shaft power saved let the selector give back, what is it worth, and what does the megawatt cost in fuel? *Checkpoint:* 3.12 t/h, about 1370 $/h, against 30 $/h.

14.11 Application area · 3 min

The Limitations section notes that the model's analyzer cannot fail. Suppose the sample line plugs and the chromatograph keeps reporting its last good value, 0.420, while the inventory drifts to 0.400 light over a shift with the suction pressure at 2.90 bar. The approach controller holds the computed cold-end approach at 7 °C.

(a) By how much, and in which direction, is the computed approach wrong, and where does the true cold-end approach settle? *Checkpoint:* 2.4 °C too wide, so the true approach settles near 4.6 °C.

(b) Name one check that would expose the fault.

14.12 Review the agent · 5 min

An AI assistant reviews scenario 1 and writes: *Opening the JT valve from 54.9 to 62 % made the LNG 4.5 °C colder, a process gain of −0.63 °C per percent with a time constant of a few minutes. Pair the LNG temperature directly with the JT valve under a PI loop tuned from that step, leave the compressor at fixed speed, and drop the suction pressure loop as unnecessary.* Find the flaw and say how you would check it in the app.

CHAPTER 15

An LNG Import Terminal Whose Boil-Off Can Only Leave With Its Send-Out

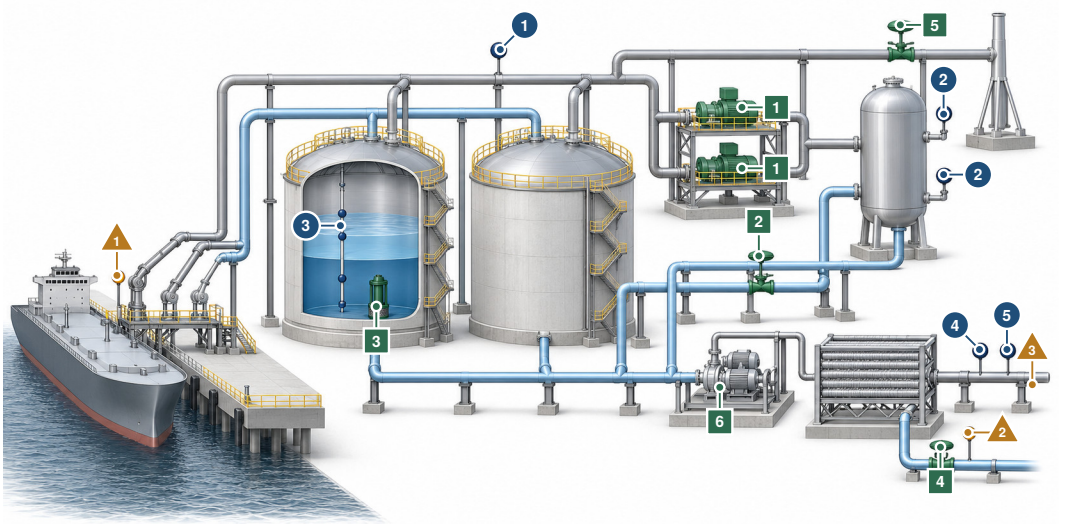

The controller moves 1 two boil-off compressors, loaded and unloaded in stages
2 valve that sets how much LNG goes through the recondenser 3 in-tank pumps 4 seawater supply to the vaporizers
5 flare valve, opened by a high tank pressure 6 high-pressure pumps, which set the send-out

The controller reads 1 tank pressure in the shared vapor space 2 recondenser pressure and level
3 density profiler in a layered tank 4 send-out flow 5 gas temperature after the vaporizers

What disturbs it 1 the ship: displaced vapor and cargo flash 2 seawater temperature
3 the grid's nomination, which sets the send-out

Notice that the compressors' gas has one way out, into the recondenser, where the LNG on its way to the pipeline condenses it, so a low send-out leaves the boil-off in the tanks. Illustration, not to scale.

IN PLAIN TERMS Where the Gas From a Cold Tank Goes

The main idea. An import terminal stores liquefied natural gas in tanks so cold that a little of it boils all the time. Compressors push that gas into a vessel where the cold liquid leaving for the pipeline turns it back into liquid.

The challenge. The outgoing liquid can take back only a small share of its own weight. On a night of low demand the gas has nowhere to go, and more compressor power does not help. The tank pressure climbs until the gas must be burned at a flare.

Why it matters. Knowing the demand hours ahead lets the terminal lower the tank pressure before the night and keep the gas.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? box at the end of the Introduction, the History box in Section 15.2, the Try it boxes in Sections 15.3.8 and 15.5, the second Did you know? box in Section 15.5.8, and Key points in Section 15.8. No equations.
Study	45 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 15.5 with Figure 15.4, Section 15.6 and Section 15.7 with its Open challenge box.

The control question The boil-off of an LNG import terminal can leave only by condensing into the LNG it sends out, so the outlet of an integrating tank pressure shrinks with a flow the market sets. How low can the send-out go before the tanks reach the flare, and how much of a low-demand night can a published nomination prepare the tanks for?

Learning objectives

- Compute the recondenser's capacity per kilogram of send-out from the subcooling and the heat each kilogram of boil-off must give up, and the minimum send-out that follows.
- Predict the rate of rise of the tank pressure from the vapor-space capacitance and the boil-off the recondenser cannot take.
- Explain why the compressors cannot rescue a low-send-out night, and how the recondenser override takes the load away from the pressure loop.
- Compare feedback alone, a preview of the published nomination and an MPC on flaring, compressor energy and boil-off.
- Judge the trade between tank pressure, cargo flash and compressor energy during a ship's unloading.

Prerequisites Sections 1.4.4 and 1.4.5 (inventories), 1.4.8 (selectors in external-reset form), 1.6.2 (the PI loops) and 1.6.4 (cases and data). Section 1.6.3 is needed only for the predictive controller. Chapter 11 for staged machines.

The app https://apmonitor.com/apps/lng-regasification/

See also Volume 6, Chapter 7 for inventory choices that decide where a rate change lands. Volume 4, Chapter 2 and Volume 4, Chapter 3 for protective overrides and limiter selection. Volume 3, Chapter 6 for feedforward that acts before feedback can.

15.1 Introduction

Two compressors that can move 20 t/h between them cannot hold the pressure of tanks that make 6.9 t/h of gas. In the generic terminal simulated here that happens every night the grid nominates 30 t/h. The compressors deliver the boil-off to a recondenser, where it condenses into LNG drawn from the tanks at −159.14 °C on its way to the high-pressure pumps, and the recondenser at 8 bar g saturates at −126.26 °C. Each kilogram of that LNG can absorb only its 32.89 K of subcooling, while each kilogram of compressed boil-off must give up 694.3 kJ, so the recondenser condenses at most 0.1658 kg per kilogram of LNG through it. At 450 t/h of in-tank pump flow that is 73.5 t/h. With the nomination ramped to 30 t/h, the recondenser fills 91.4 min into the case, its pressure override cuts the compressors to 4.5 t/h, and the tank pressure climbs 23.7 mbar an hour to the flare at 4.69 h.

A common belief holds that a gas inventory is limited by the machines that empty it, and the simulator is built to show it wrong. The vapor space of the tanks is an integrator, but the outlet that matters is a capacity proportional to a product flow set by the market, and the binding constraint moves from the compressors to the recondenser without any alarm on the compressors. The class is general. A refinery fuel gas header whose only consumers are its furnaces (Volume 11, Chapter 14), a digester whose biogas goes to an engine that follows the electrical load, a cooling system whose heat sink is a process stream, and a surge tank whose outflow is set downstream all integrate a surplus against a sink that another decision sets. Import terminals meet the problem in its plainest form, and their literature poses boil-off handling as a question of design and operation together [1–3].

DID YOU KNOW? A Tank Boils Away a Fraction of a Percent a Day

LNG storage tanks are typically designed so that less than 0.05 % of their content boils away a day, though the figure ranges from 0.02 to 0.1 %. Unloading a ship typically makes 8 to 10 times the boil-off of a tank at rest [4]. The in-leak of this chapter's tanks alone boils 0.082 % of two full tanks a day.

The published treatments of terminal boil-off are rigorous simulation-based optimizations and dynamic simulations of detailed process models, one of them built in a commercial simulator, which

a student cannot open or perturb [3, 5, 6], and the textbooks describe the equipment [7, 8]. This chapter works instead with a browser simulation of an LNG import terminal whose recondenser recomputes its capacity from the subcooling at every step. It quantifies how feedback, feedforward from a published nomination and a predictive controller manage a boil-off inventory whose sink is the send-out. The balances, the PI loops, the override and the MPC are standard. The measurements belong to one generic terminal at one seed, among them that the capacity per kilogram of send-out is $r/(1+r) = 0.1422$ with a minimum send-out of 49.6 t/h by hand and 50.0 t/h on the running model, and that a constant low tank pressure saves compressor energy while it raises the day's boil-off from 210.0 to 271.5 t through cargo flash. The model has 32 states, its tank and recondenser mass and liquid energy close to 3×10^{-5} kg and 3×10^{-4} kJ over a disturbed day, and its controllers read only instruments and published schedules. Each of its nine cases credits a learner only for a change of their own followed by the state it produces in the terminal.

Section 15.2 follows the domain, from boil-off handling to rollover. Section 15.3 sets out the terminal as modeled, and Section 15.4 its control layers: the pressure loop with its recondenser override, the staging, the boil-off observer, the preview and the MPC. Section 15.5 measures the nine cases and the graded day. Section 15.6 sets the hardware of a real terminal beside the model, Section 15.7 states the limitations, Section 15.8 draws the conclusions, and the exercises of Section 15.9 work through the cases.

15.2 Boil-Off, Recondensers and Rollover at Receiving Terminals

Natural gas crosses oceans as a liquid near −160 °C, and the gas-processing literature covers LNG as part of an international trade [7]. The receiving end of that trade stores the liquid, pumps it to pipeline pressure and vaporizes it. A handbook of the industry gives separate chapters to the operation of regasification terminals and to the process control and automation of LNG plants and import terminals [8]. No tank insulation is perfect, so the stored liquid boils continuously, and the gas that results has three possible destinations: compression into the export pipeline, recondensation into the send-out, or the flare.

Operating that choice has a literature of its own. A 2008 study optimized the operation of the boil-off compression process with a model of the boil-off rate of the storage tanks [9], and boil-off management at the Spanish terminals was surveyed in 2010 [1]. In the same year a simulation-based optimization over a regasification superstructure was posed to recover the boil-off made during ship unloading with the least energy, as a strategy for minimizing flare emissions at receiving terminals [3]. A 2012 dynamic study called the recondensation system one of the most important facilities at a receiving terminal, whose failure can cause boil-off loss or accidents, and improved its control algorithm by tracking recondenser pressure, liquid level and high-pressure pump suction pressure while the LNG and boil-off loads fluctuated [5]. A retrofit design of the boil-off handling process followed [2], and the boil-off generated at the jetty while ships load was simulated dynamically at export terminals, with less flaring as the aim [6].

The composition of stored LNG changes as it boils, because the lightest components leave first. A weathering model validated against industry data showed that heat ingress and boil-off release shift the liquid's composition and density over its storage, and that the nitrogen content governs the early boil-off through nitrogen's preferential evaporation [10]. A stratified tank adds a sharper hazard. Rollover is the sudden mixing of stratified layers, which can release large amounts of boil-off, and the History box at the end of this section tells its best documented incident and the first model to reproduce it [11, 12]. A 1997 study modeled the behavior of stratified LNG in storage tanks [13]. A 2014 review found no reliable model of industrial-scale rollover in the open literature, estimated a critical stability ratio from the 1997 data, and observed that the lumped models that describe La Spezia predict rollover too early at other incidents because they overestimate the mass transfer between the layers [11]. The handling literature adds that heat flowing into stratified layers equalizes their densities, and that a fall in the boil-off rate may be the only warning [14]. The simulator draws

Table 15.1 Variables by Role, Design Values With Tanks Half Full

Symbol	Role	Design	Range
p	CV, tank pressure (mbar g)	150	band 50–220
p_R, L_R	CV, recondenser (bar g, %)	8.0, 50	SP 6–10
F_{SO}	CV, send-out (t/h)	450	0–700
T_g	CV, gas temperature (°C)	5.00	≥ 1
D	MV, compressor load (t/h)	6.57	0–20
x_R	MV, LNG share to recondenser	0.0894	0–1
F_{LP}	MV, in-tank pumps (t/h)	443.4	0–700
F_{sw}	MV, seawater (t/h)	19,622	10–100 %
n	MV, compressors loaded	1	0, 1, 2
F_{nom}	DV, nomination (t/h)	450	0–700
Q_{unl}	DV, unloading rate (m^3/h)	12,000	0–14,000
T_{sw}	DV, seawater (°C)	15	0–30
M_v	state, vapor (t)	283.3	—
M_A, T_A	state, tank A	36.0 kt, −159.44°C	—
M_{vR}, M_R	states, recondenser (kg)	662, 8100	—
$\dot{m}_{bog}$	output, boil-off (t/h)	8.35	—
$\dot{W}_{max}$	output, capacity (t/h)	73.5	—

on this history in two places. Its recondenser carries the capacity that makes low send-out a flaring problem, and its rollover is a declared illustration of the mechanism rather than a model of any incident.

> **HISTORY 1971: A New Cargo Under an Old One**
>
> On 21 August 1971 a carrier discharged a new cargo into the bottom of a storage tank at an Italian terminal. The tank was already half full of LNG colder than the cargo, and the two liquids did not mix. They formed two layers of different density. About 31 hours after the transfer began, the layers mixed suddenly and the tank pressure rose fast. About 2,000 tonnes of vapor left through the tank's safety valves over a few hours and damaged its roof [15]. A 1983 model of the heat and mass transfer between the layers reproduced the documented time history and boil-off of that event [12]. A survey of the years 1970 to 1982 counted 41 rollover incidents in 22 plants and traced over half of them to new LNG loaded into a tank that already held some [15].

15.3 The Terminal as Modeled

The simulated terminal copies no real site. Its sizes, properties and thresholds are all teaching values. Two 160 000 m^3 full-containment tanks share one vapor space and one boil-off header at a design pressure of 150 mbar g. A ship unloads through a liquid arm into both tank roofs and takes part of the displaced vapor back through a vapor arm. Two 10 t/h compressors deliver the boil-off to a recondenser at 8 bar g, submerged in-tank pumps send LNG to it and around it, high-pressure pumps raise everything to 80 bar g, open-rack vaporizers warm the gas on seawater, and a flare valve opens on the tanks' pressure. The arrangement, the symbols of the equations and the values of two lesson instants are shown in Figures 15.1 and 15.2. Pressures in the tanks are gauge in mbar, those of the recondenser and send-out gauge in bar, flows in t/h and time in minutes.

The model integrates 32 states, listed by role with the other variables in Table 15.1. The physical states are the vapor mass M_v, tank A's liquid mass and temperature, tank B's upper and lower layers (mass, temperature, reference density), the recondenser's vapor and liquid, the compressors' net flow, the recondenser's LNG share, the in-tank pumps' flow, the send-out, the flare valve, the seawater flow, the gas temperature and the ship's cargo pumps. Thirteen running totals integrate the mass and energy that cross each boundary, so the balances can be checked exactly.

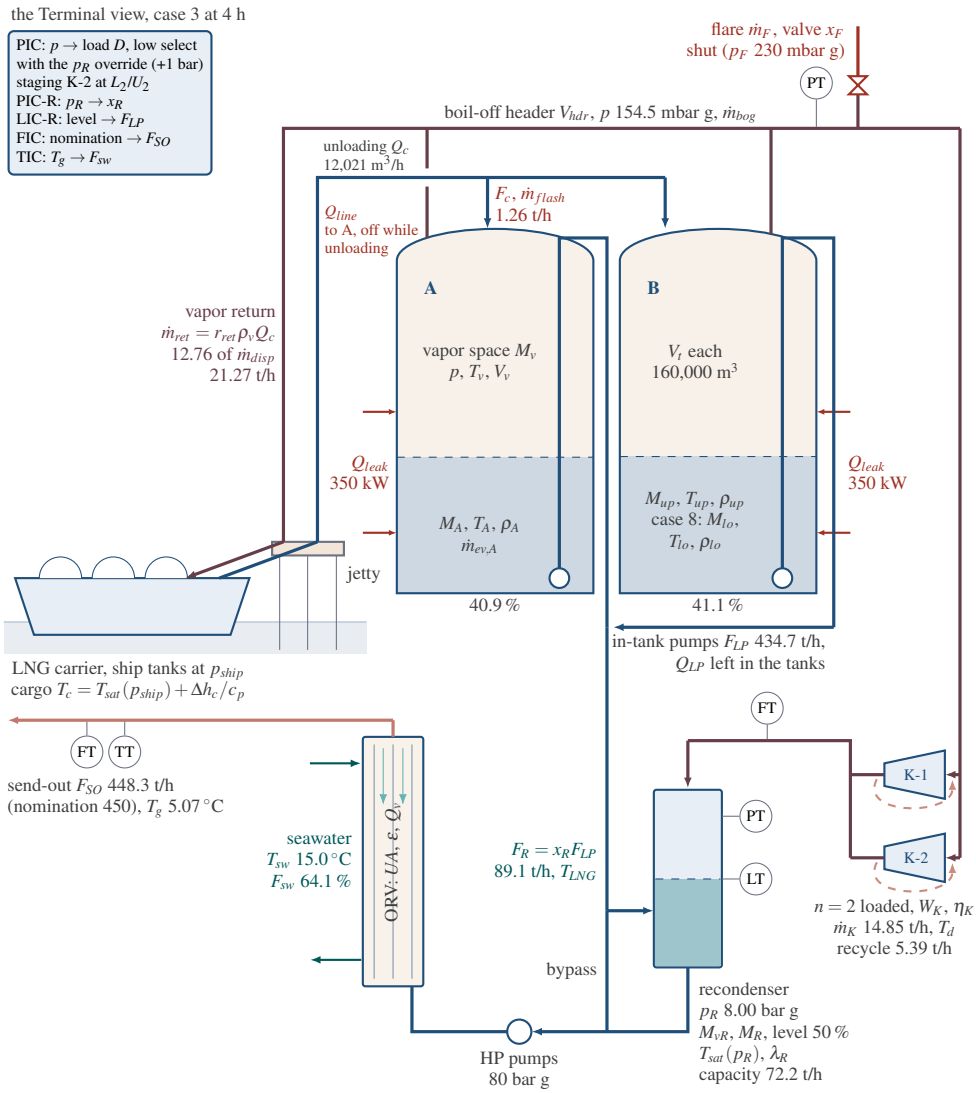

Figure 15.1 The terminal as modeled, drawn after the Terminal view of case 3 at $t = 4$ h, with every stream flowing: the ship unloading 12 021 m³/h into both roofs, the vapor arm taking back 12.76 t/h of the 21.27 t/h displaced, 1.26 t/h of cargo flash, both compressors loaded with 14.85 t/h to the recondenser and 5.39 t/h round their recycles, and 448.3 t/h sent out. The ship, the tanks and the header are drawn on top, and the compressors, the recondenser, the high-pressure pumps and the vaporizers below, with the flow from right to left. Flows in t/h. Process model only, not to scale.

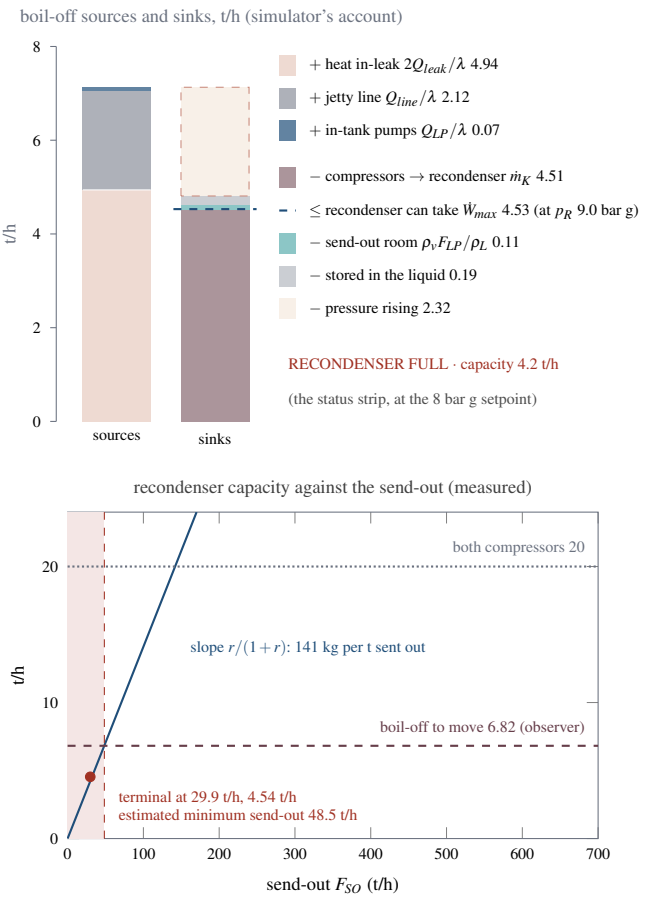

Figure 15.2 The Balance view of case 1 at $t = 4$ h, a night at 30 t/h. Top: the boil-off's sources against its sinks, the recondenser's limit of 4.53 t/h at 9.0 bar g as a dashed line across the sinks bar, and 2.32 t/h left as "pressure rising". Bottom: the measured capacity line (solid) with the terminal's point left of the estimated minimum send-out (thin dashed), under the boil-off to move (dashed) and both compressors' limit (dotted).

15.3.1 The Shared Vapor Space

Both tanks and the header hold one vapor mass at one pressure, so the vapor space is a single integrator, a gas inventory of the kind described in Sections 1.4.4 and 1.4.5, its volume set by the room the liquid leaves, as given in Equation (15.1).

$$\begin{aligned} \frac{\mathrm{d}M_v}{\mathrm{d}t} &= \dot{m}_{bog} - \dot{m}_{ret} - \dot{m}_K - \dot{m}_F - \dot{m}_{rel} \\ p = \frac{M_v R T_v}{\mathscr{M} V_v}, &\qquad V_v = 2V_t - \sum_k \frac{M_k}{\rho_k} + V_{hdr} \end{aligned} \tag{15.1}$$

where $\dot{m}_{bog}$ is the boil-off made by evaporation and flash, $\dot{m}_{ret}$ the vapor returned to a ship, $\dot{m}_K$ the compressors' net flow to the recondenser, $\dot{m}_F$ the flare, $\dot{m}_{rel}$ the relief valves, R the gas constant, $\mathscr{M}$ the molar mass, T_v the vapor temperature, V_t the capacity of each tank, V_{hdr} the header, and M_k, ρ_k the mass and density of each liquid layer. The capacitance $\partial M_v/\partial p$ is 243.6 kg/mbar with both tanks half full and 97.9 kg/mbar at 80 %, so a surplus of 1 t/h raises the pressure 4.1 or 10.2 mbar an hour.

Sharing the space means that a send-out drawn from either tank makes the same room, that unloading into both tanks squeezes one volume, and that a surge from one tank raises the pressure over both.

15.3.2 Saturation, Density and Evaporation

The LNG is one methane-rich lump. Its saturation line is Clausius–Clapeyron from a declared bubble point and its density falls linearly with temperature, as shown in Equation (15.2).

$$\frac{1}{T_{sat}(p)} = \frac{1}{T_b} - \frac{R}{\lambda \mathscr{M}} \ln \frac{p}{p_0}, \qquad \rho_k = \rho_{ref,k} - \beta\,(T_k - T_{ref}) \tag{15.2}$$

where T_b is the bubble point at $p_0 = 1.013\,25\,\mathrm{bar\,a}$, λ the latent heat, $\rho_{ref,k}$ a layer's reference density and $T_{ref} = 113.65\,\mathrm{K}$. At 150 mbar g the line gives −159.444 °C, a liquid of 449.92 kg/m^3 and a vapor of 1.7653 kg/m^3 at T_v. Heat that reaches a surface evaporates liquid there, and a weak exchange between the bulk and the surface lets a rising pressure store heat in the liquid, as shown in Equation (15.3).

$$\begin{aligned} \dot{m}_{ev,j} &= \frac{Q_{s,j} + h_S\,(T_j - T_{sat}(p))}{\lambda} \\ Q_{s,A} &= Q_{leak} + Q_{line}(1 - x_u) + Q_{LP,A} \\ Q_{LP} &= \frac{F_{LP}\,\Delta p_{LP}}{\rho_L}\left(\frac{1}{\eta_{LP}} - 1\right) \end{aligned} \tag{15.3}$$

where j is tank A or tank B's upper layer, Q_{leak} the in-leak per tank, Q_{line} the jetty line's recirculation heat, returned to tank A between ships and stopped while the ship's pumps run at fraction x_u, and Q_{LP} the in-tank pumps' losses, shared by the draw. Tank B's surface receives Q_{leak} less what a lower layer takes. At design the in-leak alone makes 4.941 t/h, 7.06 kg per kWh and 0.082 % of two full tanks a day, the line 2.118 t/h and the pumps' 182.5 kW 1.288 t/h, for 8.347 t/h of boil-off. The send-out's drawdown leaves room for 1.739 t/h, so the compressors move 6.574 t/h. Each liquid layer balances its mass and energy, as shown for tank A in Equation (15.4).

$$\begin{aligned} \frac{\mathrm{d}M_A}{\mathrm{d}t} &= \tfrac{1}{2}F_c' - d_A F_{LP} - \dot{m}_{ev,A} \\ M_A c_p \frac{\mathrm{d}T_A}{\mathrm{d}t} &= \tfrac{1}{2}F_c' c_p (T_c' - T_A) - h_S (T_A - T_{sat}) \\ &\quad - \dot{m}_{ev,A}\, c_p (T_{sat} - T_A) \end{aligned} \tag{15.4}$$

where F_c' is the cargo that stays liquid, $T_c' = \min(T_c, T_{sat})$ its temperature once it has flashed, d_A tank A's share of the draw (one half, or zero when the send-out is drawn from tank B first) and c_p the liquid heat capacity. Tank B's upper layer obeys the same balance plus the liquid an overturn lifts into it (Section 15.3.8).

15.3.3 Unloading: Displacement, Return and Flash

The cargo's volume squeezes the same volume of vapor, the ship takes a share of it back, and cargo warmer than the tanks' saturation flashes as it enters, as shown in Equation (15.5).

$$\begin{aligned} Q_c &= x_u Q_{unl}, \qquad F_c = \rho_L(T_c)\,Q_c \\ T_c &= T_{sat}(p_{ship}) + \Delta h_c / c_p \\ \dot{m}_{disp} &= \rho_v Q_c, \qquad \dot{m}_{ret} = r_{ret}\,\rho_v Q_c \\ \dot{m}_{flash} &= \frac{F_c\, c_p \max\left(0,\, T_c - T_{sat}(p)\right)}{\lambda}, \quad F_c' = F_c - \dot{m}_{flash} \end{aligned} \tag{15.5}$$

where Q_{unl} is the unloading rate, $p_{ship} = 150\,\mathrm{mbar}\,g$ the ship's tank pressure, $\Delta h_c = 0.3\,\mathrm{kJ/kg}$ the heat its pumps and the line add, and $r_{ret} = 0.6$ the share returned (the ship's own boil-off fills the

rest of its emptying tanks). The cargo pumps ramp over 30 min, and they begin to stop when the cargo left equals what the ramp down will still deliver, so a 140 000 m^3 cargo ends at 140 000 m^3. The cargo arrives 0.0857 K above the tanks' saturation at 150 mbar g, so at 12 000 m^3/h (5398 t/h) it flashes 3.18 t/h at 150 mbar g, and the flash falls by $F_c c_p (\mathrm{d}T_{sat}/\mathrm{d}p)/\lambda = 0.3995\,\mathrm{t/h}$ for each mbar the tanks run higher, reaching zero at 157.9 mbar g. This top-fill equilibrium flash is strong on purpose. A bottom fill would suppress part of it with the hydrostatic head. The two ship values are teaching values chosen so that a compressor trip during unloading can reach the flare.

15.3.4 Compressors

Each compressor is an on/off machine whose capacity scales with the suction pressure. What the load demand does not take goes round its recycle, and a loaded machine draws its full power, as shown in Equation (15.6).

$$\begin{aligned} W_K(p) &= W_{K,0}\,\frac{p}{p_{des}}, \qquad w = \frac{c_{p,v} T_v \left[(p_R/p)^{\kappa} - 1\right]}{\eta_K} \\ T_d &= T_v\left(1 + \frac{(p_R/p)^{\kappa} - 1}{\eta_K}\right), \qquad P_K = n\,W_K\,w \\ \tau_K \frac{\mathrm{d}\dot{m}_K}{\mathrm{d}t} &= u\,n\,W_K - \dot{m}_K, \qquad u = \mathrm{clip}\big(D/(nW_K), 0, 1\big) \end{aligned} \tag{15.6}$$

where $W_{K,0} = 10\,\mathrm{t/h}$ at $p_{des} = 150\,\mathrm{mbar}\,g$ in absolute terms, $\kappa = (\gamma - 1)/\gamma = 0.26$, $\eta_K = 0.75$, $c_{p,v}$ the vapor heat capacity, T_d the discharge temperature, n the number loaded and D the load demand. At design $w = 274.54\,\mathrm{kJ/kg}$ and a loaded machine draws 762.6 kW, 18.30 MWh a day, whatever its recycle. The discharge leaves at −15.21 °C.

15.3.5 The Recondenser and Its Capacity

The recondenser is in equilibrium. The LNG sent through it leaves saturated at its pressure, and the compressed boil-off it condenses is exactly what the subcooling of that LNG absorbs, as shown in Equation (15.7).

$$\begin{aligned} \dot{m}_{cond} &= \frac{x_R F_{LP}\, c_p \left(T_{sat}(p_R) - T_{LNG}\right)}{\lambda_R} = r\,x_R F_{LP} \\ \lambda_R &= \lambda_{R,lat} + c_{p,v} \max\left(0,\, T_d - T_{sat}(p_R)\right) \\ \frac{\mathrm{d}M_{vR}}{\mathrm{d}t} &= \dot{m}_K - \dot{m}_{cond}, \qquad \frac{\mathrm{d}M_R}{\mathrm{d}t} = F_{LP} + \dot{m}_{cond} - F_{SO} \\ p_R &= \frac{M_{vR} R T_{vR}}{\mathscr{M}\,(V_{RC} - M_R/\rho_{RC})} \end{aligned} \tag{15.7}$$

where x_R is the share of the in-tank pumps' flow F_{LP} sent through it, T_{LNG} the draw-weighted tank temperature plus 0.3 K of warming in the line, $\lambda_{R,lat}$ the latent heat at the recondenser's pressure, V_{RC} its volume, ρ_{RC} its liquid density and T_{vR} its vapor temperature. By hand from the declared constants, $T_{sat}(p_R) = -126.258\,^{\circ}\mathrm{C}$, $T_{LNG} = -159.144\,^{\circ}\mathrm{C}$, the subcooling is 32.886 K, $\lambda_R = 450 + 2.2 \times 111.05 = 694.31\,\mathrm{kJ/kg}$ and

$$r = \frac{3.5 \times 32.886}{694.31} = 0.16578, \qquad \frac{r}{1+r} = 0.14220. \tag{15.8}$$

With all the LNG through it ($x_R = 1$) the recondenser condenses rF_{LP}, and the condensate joins the send-out, $F_{SO} = (1+r)F_{LP}$, so its capacity per kilogram of send-out is $r/(1+r)$, as shown in Equation (15.8). Written on F_{LNG}, the LNG through the recondenser, the capacity would be r per kilogram. Written on the send-out it is $r/(1+r)$, which is why the plotted line has the slope 0.142 kg/kg and not 0.166 kg/kg. The engine at $t = 0$ of case 2 reproduces every figure above to the printed digit, and its condensed flow equals Equation (15.7), with the recondenser's enthalpy balance closing to machine precision, at every step of case 1. The in-leak and jetty line alone need

7.06 t/h moved, which divided by 0.1422 gives a minimum send-out of 49.6 t/h. A higher recondenser pressure condenses at a higher temperature: $r = 0.1737$ at 9 bar g and 0.1808 at 10 bar g, for 814 kW and 861 kW a machine. The capacity the operator reads on the status strip and the trend is the same formula from measurements at the pressure setpoint, while the Balance view's limit line is the simulator's capacity at the actual pressure. In case 1 at 4 h they read 4.21 and 4.53 t/h, the difference being the override's extra bar and not truth against measurement, because the capacity computed from measurements at the measured pressure reads 4.526 t/h.

15.3.6 High-Pressure Pumps and Vaporizers

The high-pressure pumps raise the mixed stream from p_R to $p_{SO} = 80\,\text{bar}\,g$ and leave their losses in it. Open-rack vaporizers are a lumped counter-current bank whose conductance grows with the seawater film, as shown in Equation (15.9).

$$\begin{aligned} T_{Lv} &= T_{HP} + \left(\frac{1}{\eta_{HP}} - 1\right)\frac{p_{SO} - p_R}{\rho_L c_p} \\ Q_v &= \varepsilon\, C_{min}(T_{sw} - T_{Lv}), \qquad UA = UA_0\left(\frac{F_{sw}}{F_{sw,0}}\right)^{n_{UA}} \\ \tau_V \frac{\mathrm{d}T_g}{\mathrm{d}t} &= T_{Lv} + \frac{Q_v}{F_{SO}\, c_{eff}} - T_g \end{aligned} \tag{15.9}$$

where T_{HP} is the flow-weighted temperature of the saturated recondenser outlet and the bypass, ε the counter-flow effectiveness at NTU $= UA/C_{min}$ and C_{min}/C_{max} with $C_g = F_{SO}c_{eff}$ and $C_{sw} = F_{sw}c_{sw}$, c_{eff} the effective heat capacity of the vaporizing LNG and $F_{sw,0} = 20\,000\,\text{t/h}$. At design the pumps take 2667 kW and warm the LNG 1.524 K, and 19 622 t/h of 15 °C seawater at $\varepsilon = 0.9409$ delivers 105.50 MW, 844 kJ per kilogram of gas, cooling the seawater 4.84 K.

15.3.7 Flare and Relief

The flare valve's flow is $\dot{m}_F = x_F W_F \sqrt{(p - p_0)/230\,\text{mbar}}$ with $W_F = 80\,\text{t/h}$, and it is fully open 40 mbar above its 230 mbar g set pressure, passing 86.7 t/h at 270 mbar g. The relief valves pass $k_{rel}\max(0, p - p_{rel})$ with $k_{rel} = 20\,\text{t/h}$ per mbar above $p_{rel} = 290\,\text{mbar}\,g$.

15.3.8 Tank B's Layers (Illustrative)

Case 8 gives tank B a denser lower layer that receives the floor's share of the in-leak and the walls' share in proportion to its depth, and has no surface. The upper layer weathers as it boils, and once the densities cross the lower layer overturns at a declared rate, as shown in Equation (15.10).

$$\begin{aligned} Q_{lo} &= Q_{leak}\left(f_{fl} + f_w \frac{M_{lo}}{M_{lo} + M_{up}}\right) \\ M_{lo} c_p \frac{\mathrm{d}T_{lo}}{\mathrm{d}t} &= Q_{lo}, \qquad \dot{m}_{lift} = k_{Ro} M_{lo} \\ M_{up} \frac{\mathrm{d}\rho_{ref,up}}{\mathrm{d}t} &= \tfrac{1}{2} F_c'(\rho_{ref} - \rho_{ref,up}) + w_W \dot{m}_{ev,B} \\ &\quad + \dot{m}_{in}(\rho_{ref,lo} - \rho_{ref,up}) \\ \dot{m}_{flash,L} &= \frac{\dot{m}_{lift}\, c_p \max\left(0, T_{lo} - T_{sat}\right)}{\lambda} \end{aligned} \tag{15.10}$$

where $f_{fl} = 0.15$ and $f_w = 0.60$, $w_W = 40\,\text{kg/m}^3$ per unit of the layer boiled off, $k_{Ro} = 1$ per hour applies once $\rho_{lo} \le \rho_{up}$, and $\dot{m}_{in} = \dot{m}_{lift} - \dot{m}_{flash,L}$ is the liquid that joins the upper layer. Tank B's pumps draw the lower layer while it lasts. The layers exchange no heat or mass across their interface, the mechanism that the rollover models of the literature integrate [11, 12], so this model shows the sequence and not the size or timing of a real event.

Table 15.2 Parameters and Defaults. A, Adjustable in the Advanced Panel. C, a Control on a Card. F, Fixed in the Source

Symbol	Meaning	Value Unit	
V_t, V_{hdr}	tank (each), header	160,000, 500 m^3	A
ρ_{ref}, β	LNG density, its slope	450, 1.4 kg/m^3, per K	A
c_p, λ	liquid heat capacity, latent heat	3.5, 510 kJ/(kg K), kJ/kg	A
$\mathcal{M}$, T_b	molar mass, bubble point	16.8, 112.0 kg/kmol, K	A
T_v	vapor-space temperature	133.15 K	A
Q_{leak}, Q_{line}	in-leak per tank, jetty line	350, 300 kW	A
h_S	bulk–surface exchange per tank	20 kW/K	A
Δp_{LP}, η_{LP}	in-tank pump head, efficiency	10, 0.6 bar, –	A
—	LNG warming to the recondenser	0.3 K	A
r_{ret}, p_{ship}	vapor returned, ship pressure	0.6, 150 –, mbar g	A
Δh_c, V_{cargo}	cargo heat, cargo	0.3, 140,000 kJ/kg, m^3	A
$W_{K,0}$, η_K, κ	compressor (each)	10, 0.75, 0.26 t/h, –, –	A
$c_{p,v}$, τ_K	vapor heat capacity, response	2.2, 0.3 kJ/(kg K), min	A
V_{RC}, V_{RL}	recondenser volume, at 100 %	80, 40 m^3	A
ρ_{RC}, T_{vR}	recondenser liquid, vapor	405, 165 kg/m^3, K	A
$\lambda_{R,lat}$	latent heat at p_R	450 kJ/kg	A
$F_{SO,max}$, η_{HP}	HP pumps	700, 0.75 t/h, –	A
UA_0, n_{UA}	vaporizer conductance	1935, 0.6 kW/K, –	A
c_{eff}, c_{sw}	vaporizing LNG, seawater	5.3, 4.0 kJ/(kg K)	A
$F_{sw,max}$, τ_V	seawater pumps, lag	30,000, 2 t/h, min	A
T_{min}	minimum gas temperature	1 °C	A
W_F, p_{rel}	flare valve, relief set	80, 290 t/h, mbar g	A
f_{fl}, f_w	lower layer's heat shares	0.15, 0.60 –	A
w_W, k_{Ro}	weathering, overturn rate	40, 1.0 kg/m^3, per h	A
$\Delta\rho_0$, ΔT_0	case 8 layer gap	0.35, 1.5 kg/m^3, K	A
p_{SP}, p_F	tank setpoint, flare	150, 230 mbar g	C
L_2, U_2	K-2 loads, unloads	95, 60 % of a machine	C
$p_{R,SP}$, $L_{R,SP}$	recondenser setpoints	8.0, 50 bar g, %	C
—	override above $p_{R,SP}$	1.0 bar	F
—	preview floor, ceiling margin	65, 30 mbar g, mbar	F

TRY IT IN THE APP 8 · Rollover in tank B (illustrative)

Load the case and press Run. Touch nothing. Watch the Tank B layer densities trend on the Tanks view as its two lines close. They meet at 21.34 h, the status strip reads ROLLOVER, the boil-off jumps to 97.9 t/h and the tanks reach 234.2 mbar g. The Flared tile ends near 3.84 t.

15.3.9 Instruments, Parameters and Numerics

The controllers read one packet. The tank pressure carries a 5 s lag and 0.3 mbar of noise, the recondenser pressure 2 s and 0.005 bar, its level 2 s and 0.2 %, the compressor discharge meter 5 s and 1 %, the in-tank pump, recondenser LNG and send-out meters 2 s and 0.3 %, the unloading meter 5 s and 0.3 %, the flare's ultrasonic meter 5 s and 3 %, the gas temperature 10 s and 0.1 K, and tank B's density profiler each layer to 0.02 kg/m^3. Every execution draws the same 18 deviates whatever the settings. The parameters are listed in Table 15.2. Every case begins at the steady state its settings imply, found by a fixed point between the pump losses and the condensate, with the vaporizer's seawater found by bisection. At $t = 0$ of case 2 every actuator and vaporizer derivative is under 0.0013 per minute and the vapor drift is zero, and after 15 min the pressure reads 149.92 mbar g. The classical fourth-order Runge-Kutta method advances the states in 2 s steps, the control layers run on a 6 s scan, and a logged row follows every 30 s, as the engine of Section 1.6.1 runs them. Over the suggested challenge with a compressor trip, the tanks' mass closes to 3.0×10^{-5} kg on 43.6 kt, the recondenser's to 6×10^{-8} kg and the liquids' energy to 3×10^{-4} kJ on 3.13 TJ. Halving the step leaves the flare onset of case 1 (281.50 min), the peak of case 4 (186.712 mbar g), the flared mass of case 6 (13.960 t) and the rollover of case 8 (1280.60 min, 97.913 t/h) unchanged, and moves the RECONDENSER FULL instants of cases 1 and 4 by 0.01 min.

15.3.10 Assumptions

The LNG is one lump with constant heat capacity and latent heat, and its boil-off is the same lump, while real boil-off is richer in nitrogen and methane [10]. Only tank B's declared layers weather. Against the NIST Chemistry WebBook's saturation table for pure methane, a fluid with a reference equation of state fitted to its vapor pressure and saturated densities [16], the declared line gives 146.89, 149.13 and 151.21 K at 8, 9 and 10 bar g against 146.90, 149.17 and 151.28 K, within 0.07 K at the recondenser. At 1 atm and 150 mbar g it sits 0.33 K and 0.34 K above pure methane, as an LNG lump should, and its slope at 150 mbar g is 0.010 79 K/mbar against 0.01097. The latent heat matches methane's 510.5 kJ/kg at 1 atm, while $\lambda_{R,lat} = 450$ is 6 % above methane's 423.2 kJ/kg at 8 bar g, and $c_p = 3.5$ against 3.48 kJ/(kg K). The vapor is an ideal gas at one temperature whose energy is not tracked. The recondenser is in equilibrium, so the subcooling left at the high-pressure pumps' suction is logged but their net positive suction head is not modeled, which is what bounds a low recondenser pressure in a real terminal. The vaporizers are one lumped bank without ice. The flare (230 mbar g), relief (290 mbar g), pressure band and 1 °C gas minimum are teaching values. Real ones come from the tank system's design code (EN 14620-1 or API 625, which cover refrigerated liquefied gas tank systems, with EN 14620-1 limiting design pressure to 500 mbar), its venting design (API 2000, normal and emergency venting) and the pipeline's gas quality specification.

15.4 Control Layers as Implemented

Within a physics step the control layers act first, the states advance next, and last the status flags latch and the challenge is scored. The control scan is $\Delta t_c = 6$ s, always on the packet. Every loop is the velocity-form PID of Section 1.6.2 used without derivative action, so Equations (1.16) to (1.19) apply with $\tau_D = 0$, a control period Δt_c, the back-calculation time $T_t = \tau_I$ and an output bias u_0 that carries a feedforward where a loop has one. A loop not driving its actuator tracks the one that is, so Manual to Auto, an architecture change and an MPC fallback are bumpless.

15.4.1 Tank Pressure, Override and Staging

The tank pressure loop (PIC) sets the load demand D in t/h, acting so that a higher pressure asks for more load, with $K_c = 0.5$ t/h per mbar, $\tau_I = 120$ min and an output of 0 to 20 t/h. Only this loop answers to the header's Manual/Auto switch. A recondenser high-pressure override, a PI on p_R at $p_{R,SP} + 1$ bar with $K_c = 5$ t/h per bar and $\tau_I = 2$ min, gives the most load the recondenser allows, and $D = \min(D_{PIC}, D_{OVR})$. The loop not selected tracks the selected output plus its own proportional term, the external-reset form of Section 1.4.8, and the override acts in Manual too, because it is protection. Volume 4, Chapter 2 treats a protective override on a compressor in depth. Two compressors stage on D in percent of one machine's capacity at the measured pressures, where Chapter 11 shares a load among three machines on one header. K-1 loads above 5 % and unloads below 1 %, K-2 loads above L_2 and unloads below U_2, each machine runs and rests at least 10 min and the sequencer allows one change a minute. The loaded machines pass D and recycle the rest. The flare valve opens in proportion above its set pressure, $x_F = \mathrm{clip}\big((p_m - p_F)/40\,\mathrm{mbar}, 0, 1\big)$.

15.4.2 Recondenser, Send-Out and Gas Temperature

The recondenser pressure loop (PIC-R) moves x_R, more LNG through the vessel for a higher pressure, with $K_c = 10$ % per bar and $\tau_I = 2$ min. The level loop (LIC-R) trims the in-tank pumps' flow around the measured send-out, $F_{LP,SP} = F_{SO,m} + \mathrm{PI}$, with $K_c = 2$ t/h per percent and $\tau_I = 20$ min. RECONDENSER FULL is declared when $x_R \geq 99.5$ % and the pressure is more than 0.2 bar above its setpoint. The send-out loop (FIC) moves the high-pressure pumps' speed toward the nomination, or a setpoint by hand, with $K_c = 0.3$ % per t/h and $\tau_I = 0.3$ min. The gas temperature loop (TIC) moves the seawater pumps between 10 % and 100 %, with $K_c = 5$ % per K and $\tau_I = 5$ min, and has its own Auto/Man switch.

15.4.3 Boil-Off Observer

The preview and the MPC need the boil-off, which no instrument measures. A two-state observer on the vapor balance estimates it from the measured pressure, compressor, flare and unloading flows and the room the measured send-out makes, as shown in Equation (15.11).

$$\begin{aligned} \hat{p}_{k+1} &= \hat{p}_k + \Delta t_c \frac{\hat{g}_k + d_m - \rho_v F_{LP,m}/\rho_L - \dot{m}_{K,m} - \dot{m}_{F,m}}{C_t} \\ &\quad + 2\zeta\omega\,\Delta t_c\,(p_m - \hat{p}) \\ \hat{g}_{k+1} &= \hat{g}_k + C_t\,\omega^2\,\Delta t_c\,(p_m - \hat{p}) \end{aligned} \tag{15.11}$$

where C_t is the vapor capacitance computed from the tank gauges and the declared capacity, in t/h per mbar/min, $d_m = (1 - r_{ret})\rho_v Q_{c,m}$ the displaced vapor the ship does not take back, $\omega = 0.1$ per minute and $\zeta = 0.8$, the error being taken after the prediction step.

15.4.4 Preview From the Published Nomination

The preview reads 12 hours of the published nomination $F_{nom}(t)$ and ship schedule in 10 min steps and adds up the boil-off that the recondenser and the compressors will not take, as shown in Equation (15.12).

$$\begin{aligned} e_i &= \hat{g} - \dot{m}_{flash,m} + s_i - \frac{\rho_v}{\rho_L} F_{nom,i} \\ &\quad - \min\left(0.9\frac{r}{1+r}F_{nom,i},\ 0.9\,n_{max}W_K\right) \\ m_{room} &= \max_j \sum_{i \le j} e_i\,\Delta\tau \\ p'_{SP} &= \mathrm{clip}\left(p_F - 30 - \frac{m_{room}}{C_v},\ 65,\ p_{SP}\right) \end{aligned} \tag{15.12}$$

where s_i is the published ship's displaced vapor and flash, r is computed from the measured LNG and discharge temperatures at $p_{R,SP}$, C_v is the capacitance in kg per mbar and the 0.9 derates the capacities. The PIC then follows p'_{SP} moved at most 1 mbar/min, and the measured unloading's boil-off enters as its bias u_0. The same idea, acting on a forecast before the demand arrives, plans the line pack of Chapter 12 before its evening peak.

15.4.5 Model Predictive Control

Every 5 min the MPC solves a quadratic program over 8 hours, 96 steps in 12 graduated move blocks, on the vapor space as the integrator it is, as shown in Equation (15.13).

$$\begin{aligned} p_{k+1} &= p_k + \frac{T_s}{C_t}(w_k - u_k) \\ w_k &= \hat{g} - \dot{m}_{flash,m} + s_k - \frac{\rho_v}{\rho_L}F_{nom,k} \\ \min\ & \textstyle\sum_k 0.01 w_p (p_k - p_{SP})^2 + w_{\Delta u}\sum_j \Delta u_j^2 \\ & \quad + 0.1 w_{pow}\textstyle\sum_k u_k + 10^3\|\varepsilon\|^2 \\ \text{s.t. } & 0 \le u_k \le n_{max}W_K, \quad u_{k-1} \le 0.95\frac{r}{1+r}F_{nom,k} + \varepsilon \\ & 65 - \varepsilon \le p_k \le p_F - 30 + \varepsilon \end{aligned} \tag{15.13}$$

where $T_s = 5$ min, u is the compressor flow and $w_p = w_{\Delta u} = w_{pow} = 1$ by default. A second state carries the last move so that the capacity bound applies step by step. The shared quadratic programming module of Section 1.6.3 solves it with Goldfarb and Idnani's dual algorithm [17]. Only the first move reaches the compressors, as industrial MPC implementations do [18], and any solve that is infeasible, not finite or outside its bounds returns the compressors to the tracked PIC. At the Advanced level the horizon, the move suppression and a multiple of the slack price can be changed.

Table 15.3 The Nine Cases

#	Case and events	Teaching point
1	Manual, tanks 80 %, nomination to 30 t/h over 0.5–1.5 h	low send-out means flaring
2	ship at 5000 m^3/h from 20 min, K-2 unloads at 90 %	hysteresis
3	tanks 30 %, ship at 12 000 m^3/h from 30 min	displacement and flash
4	tanks 75 %, nomination to 40 t/h over 0.5–2.5 h, back at 7 h	minimum send-out
5	seawater 15 to 4 °C over 1 h from 20 min	the vaporizers' limit
6	tanks 55 %, ship from 5 min, K-2 trips at 2 h	flare or not
7	tanks 70 %, nomination to 20 t/h from 5 to 11 h, loops	the nomination as preview
8	tank B stratified, 0.35 kg/m^3 and 1.5 K	rollover (illustrative)
9	Challenge: ship 2–14 h, night at 20 t/h from 18 h	the whole terminal

15.4.6 Three Models

The plant is Section 15.3. What a learner identifies is a capacitance (kg per mbar) and a capacity ratio (kg per kg of send-out). The controllers' models are the observer, the preview's running sum and the MPC's integrator, which know the declared properties exactly but neither the in-leak, the line heat, the bulk exchange nor tank B's layers. A test of the application perturbs those undeclared constants and finds the preview's and the MPC's decisions unchanged.

15.5 Nine Cases on the Running Engine

The cases are listed in Table 15.3. Every number below was measured by stepping the engine headless at the default seed 20261020 in the increments the page uses, along the path its sentence names. A second run of the measurement script printed an identical record, and each case logs an identical run when it is loaded first, picked after some other case, or reset (Section 1.6.1).

15.5.1 A Night the Compressors Cannot Fix

In case 1 the compressor load sits in Manual at the 6.60 t/h the evening needed and the nomination falls to 30 t/h. RECONDENSER FULL comes at 91.4 min, the override cuts the compressors to what a full recondenser condenses at 9.0 bar g, 4.53 t/h at 2 h and 4.51 at 4 h, while the tanks still boil 6.94 t/h at 4 h. The difference stays in a vapor space of 97.9 kg/mbar, the pressure climbs 23.7 mbar/h to 213.7 mbar g at 4 h, and the flare opens at 4.69 h. The balance at that instant is drawn in Figure 15.2. Switched to Auto at 2 h, the PIC asks for all 20 t/h, the override still holds the compressors at 4.49 t/h at 5 h, and the pressure passes 200 mbar g at 3.42 h just the same. A night at 60 t/h, above the minimum send-out, never fills the recondenser, whose capacity reads 8.90 t/h at 5 h. The lesson is that "low send-out means flaring" is a statement about the recondenser's capacity, and the compressors' 20 t/h never enters it.

Pause. Try Exercise 15.1 before reading on.

15.5.2 Staging and Its Hysteresis

Case 2 brings the boil-off to move to 9.02 t/h, just under one machine's 10.01 t/h. With K-2 loading at 95 % and unloading at 90 % it starts 10 times in six hours, each time pulling the pressure down until the demand falls below the unload threshold. Unloading at 60, 70 or 80 % it starts once and stays loaded, recycling a mean of 10.46 t/h between 2 and 6 h against 5.15, and the compressors use 7.52 MWh against 5.91. Each extra loaded hour costs about 0.76 MWh whether its gas is needed or not.

> **TRY IT IN THE APP 2 · Compressor staging and hysteresis**
>
> Load the case and run it to 6 h. Watch K-2 on the Terminal view and the K-2 starts tile: it starts 10 times. Set K-2 unloads below to 60 on the Staging card, press Reset and run again. K-2 now starts once and stays loaded, and the Compressor energy tile reads 7.52 MWh against 5.91 MWh. Fewer starts cost more power.

15.5.3 The Ship: Displacement, Return and Flash

In case 3 at 4 h, with the ship at full rate, the cargo displaces 21.27 t/h of vapor, the ship takes 12.76 t/h back and 8.51 t/h stays, against 1.26 t/h of cargo flash and a surface boil-off of 7.46 t/h. The compressors move 14.70 t/h, K-2 loads at 61 min, and the pressure peaks at 158.9 mbar g. The flash is the term the tank pressure controls. Held by the setpoint, the tanks flash 4.43 t/h at 146.9 mbar g, 1.26 at 154.8 and none at 158.6, 0.40 t/h per mbar between the first two as the hand derivation predicts. A setpoint of 200 mbar g removes the flash and lowers the boil-off at 4 h from 7.46 to 6.16 t/h, while the displaced vapor, a volume, stays. The ship's pumps begin to stop at 12.17 h with the 140 000 m^3 cargo in, and both tanks end 69.1 % full.

🕮 **Pause.** Try Exercise 15.2 before reading on.

15.5.4 The Minimum Send-Out and the Recondenser Pressure

Case 4 ramps the nomination to 40 t/h. RECONDENSER FULL comes at 149.3 min, at 42.4 t/h, and the pressure peaks at 186.7 mbar g at 7.03 h before the ramp back restores the capacity. At 4 h the capacity ratio computed from measurements is 0.1622 kg per kilogram of LNG and the boil-off to move 6.98 t/h, and the true minimum send-out is 50.0 t/h against the 49.6 t/h of Equation (15.8). Taken off the nomination at RECONDENSER FULL, a send-out of 50 or 55 t/h holds while 45 and 48 t/h refill the recondenser. A recondenser at 10 bar g raises the ratio to 0.1774 and cuts the peak to 173.6 mbar g, for 8.96 MWh of compressor energy against 8.59 over 9 h. Case 5 shows a different limit on the send-out. With the seawater cooled to 4 °C, SEND-OUT COLD comes at 76.3 min with the seawater pumps at 100 % and the gas at −0.03 °C, and the most the vaporizers warm to 1 °C is 416.8 t/h, 92.6 % of the nomination (420, 410 and 400 t/h give 0.91, 1.20 and 1.47 °C).

15.5.5 A Trip While the Ship Unloads

In case 6 K-2 trips at 2 h with the tanks at 157.4 mbar g, 72.6 mbar under the flare, and 13.18 t/h to move. K-1 alone takes 10.20 t/h, the pressure climbs 17.8 mbar/h in a vapor space of 219 kg/mbar, and the flare opens at 6.15 h, burning 13.96 t before the unloading ends. Of the boil-off, 8.64 t/h is displaced vapor the ship does not take back, the one source an operator can turn down. Slowing the ship to 7000 m^3/h at 2.2, 5 or 5.2 h keeps the tanks off the flare, with peaks of 159.9, 207.3 and 210.4 mbar g.

15.5.6 The Nomination as a Preview

Case 7 publishes a night at 20 t/h from 5 h. On the loops alone the pressure sits at 150 mbar g until the night, the measured capacity collapses from 73.52 t/h at 5 h to 2.65 at 6 h, and the flare opens at 9.12 h and burns 7.15 t, as shown in Figure 15.3. With the preview the setpoint holds 150 mbar g until 233.5 min, when the room the night needs passes what 150 mbar g leaves under the ceiling, and then falls 1 mbar/min. The pressure is at 98.1 mbar g at 5 h, bottoms at 60.0 mbar g at 5.90 h, and peaks at 196.9 mbar g, for 11.43 MWh against 10.19 over 12 h. The MPC, reading the same schedule, peaks at 199.8 mbar g for 11.30 MWh with no fallback in its solves.

15.5.7 Rollover, Declared and Illustrative

Case 8 opens with tank B's lower layer at 12.61 kt, 0.350 kg/m^3 denser than the upper one and 1.500 K over the surface's saturation, 129.8 t of boil-off stored. Before the rollover tank B boils 2.40 t/h at 6 h against 3.11 as one layer, the warning the handling literature describes [14]. The densities meet at 21.34 h with 7.87 kt below storing 92.4 t. The boil-off jumps to 97.9 t/h logged (99.9 t/h in the step it begins), the pressure reaches 234.2 mbar g and 3.84 t are flared. Only 53.1 t of the 92.4 t flashes, the rest staying in the liquid as the rising pressure lifts its saturation. The observer's estimate reads 7.7 t/h as the surge begins and peaks at 51.5 t/h 26 min later, the lag of a 0.1 min^{-1} observer. Drawn from tank B first, the send-out takes the warm lower liquid with it, the layers meet at 18.07 h with 4.59 kt storing 53.7 t, the surge peaks at 61.4 t/h and 203.4 mbar g, and nothing is flared. Every one

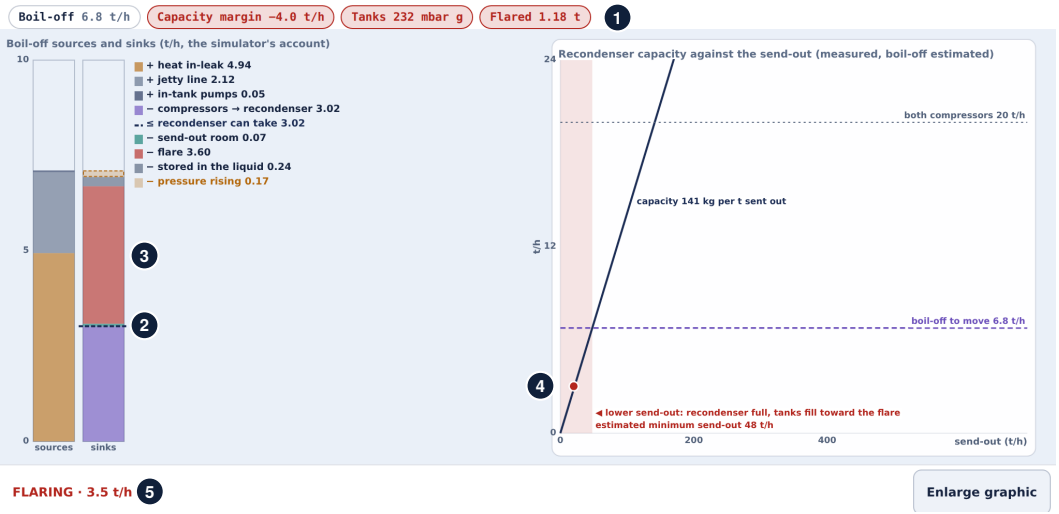

Figure 15.3 Case 7 as shipped on Loops (PI), $t = 9.5$ h, on the Balance view, light theme, cropped to the view. The status strip (5) reads FLARING · 3.5 t/h with the tanks at 231.9 mbar g against a setpoint of 150. The sources bar stacks 4.94 t/h of heat in-leak, 2.12 t/h from the jetty line and 0.05 t/h from the in-tank pumps, and the sinks bar stops the compressors at the recondenser's 3.02 t/h, its dashed line (2), with 3.60 t/h to the flare (3), 0.24 t/h stored in the liquid and 0.17 t/h as pressure rising. The plot puts the terminal (4) at the 20 t/h nomination, far left of the estimated minimum send-out of 48 t/h, under a boil-off to move of 6.8 t/h, and the capacity margin badge among the chips (1) reads −4.0 t/h, beside 1.18 t flared.

Table 15.4 Measured Cases at the Default Seed

Case	Configuration	Result
1	Manual, night 30 t/h	FULL 91.4 min, flare 4.69 h
1	Auto from 2 h	200 mbar g at 3.42 h, $\dot{m}_K$ 4.49 t/h
2	K-2 at 95/90 against 95/60	10 against 1 start, 5.91 / 7.52 MWh
3	setpoint 150 / 200 mbar g	flash 1.26 / 0 t/h, boil-off 7.46 / 6.16
4	minimum send-out	49.6 by hand, 50.0 measured
4	recondenser 8 / 10 bar g	peak 186.7 / 173.6, 8.59 / 8.96 MWh
5	seawater 4 °C	416.8 t/h reach 1 °C (92.6 %)
6	trip, as shipped / slowed	13.96 t / 0 t flared
7	loops / preview / MPC	7.15 / 0 / 0 t, 10.19 / 11.43 / 11.30 MWh
8	both tanks / B first	97.9 / 61.4 t/h, 3.84 / 0 t flared

of these numbers is a declared choice. The layers meet at 21.3, 37.1 and 47.3 h from gaps of 0.35, 0.70 and 1.05 kg/m^3, and from 1.40 kg/m^3 the send-out drains the lower layer before they meet. The overturn rate sets the surge, 53.4 t/h at 0.5 per hour with nothing flared and 183.9 t/h at 2 per hour with 11.45 t flared.

🕮 **Pause.** Try Exercise 15.3 before reading on.

15.5.8 The Graded Day

The challenge runs 24 h with both tanks 30 % full, a ship from 2 h whose pumps begin to stop at 13.67 h, and the nomination falling to 20 t/h from 18 h. Six targets are scored from the simulator's truth on every physics step: the pressure between 50 and 220 mbar g after 5 min, under 0.1 t flared, each hour's send-out within 2 % of that hour's published nomination, the gas at or above 1 °C, at most 30.6 MWh of compressor energy and at most 230 t of boil-off. A changed disturbance or model

Table 15.5 The Challenge Day at the Default Seed

Settings	p range (mbar g)	Flared (t)	Energy (MWh)	Boil-off (t)	Verdict
shipped, loops, K-2 50/40	135.6–231.9	6.61	33.08	190.8	fails
suggested, preview, 95/60	58.9–196.1	0	30.24	210.0	passes
preview, shipped staging	58.9–196.1	0	32.84	210.0	fails
loops, 95/60	135.6–231.9	6.61	28.07	190.8	fails
MPC, 95/60	64.7–196.7	0	29.75	218.4	passes
constant 70, 95/60	69.8–199.6	0	31.05	271.5	fails
constant 70, 95/80	69.8–199.6	0	28.93	271.5	fails
constant 92, 120/100	91.8–219.2	0	27.75	263.3	fails
preview, $p_{R,SP}$ 6 bar g	59.0–203.6	0	25.82	209.9	passes

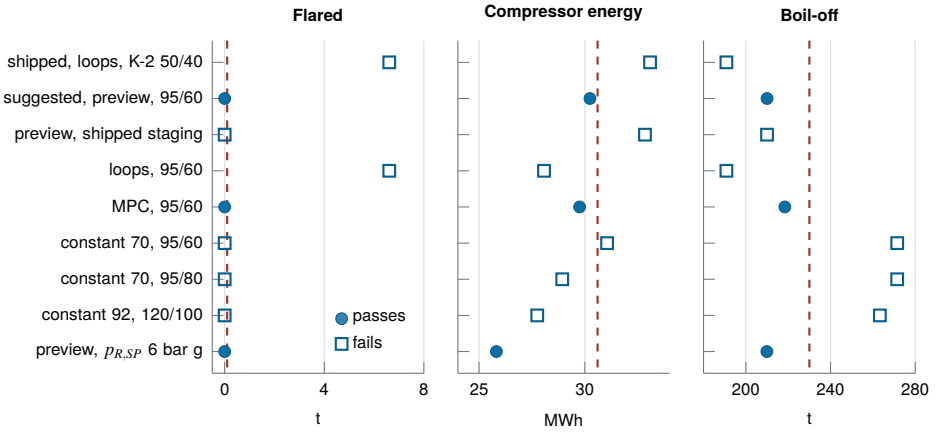

Figure 15.4 The challenge day of Table 15.5 as three dot panels, filled circles for the runs that pass and open squares for those that fail, with the day's limit as a dashed line: under 0.1 t flared, at most 30.6 MWh and at most 230 t of boil-off. Notice that no single panel sorts the runs: the loops fail on the flare, the shipped staging on energy and the constant setpoints on boil-off. The pressure range stays in the table.

parameter, or a press of the Ship button, marks the run "not graded". The outcomes are listed in Table 15.5 and the day is plotted in Figure 15.5. Figure 15.4 sets the table's flared mass, energy and boil-off against their limits.

As shipped, the pressure is still at 147.6 mbar g when the night begins, the night takes it to the flare at 22.16 h and 6.61 t burn, and K-2, loading above half a machine, runs all day for 33.08 MWh. The suggested settings, the preview with K-2 at 95 and 60 %, bank the tanks to 59.2 mbar g by 18 h and pass. Either change alone fails, the preview on power at 32.84 MWh and the staging on the flare. The MPC with the same staging passes at 29.75 MWh, and over seeds 0 to 4 and at half the step the suggested run uses 30.16 to 30.24 MWh and the MPC 29.72 to 29.80 MWh, every verdict unchanged.

The sixth target exists because energy could not separate two strategies. The energy counts loaded machine-hours, and a loaded machine draws slightly less at a lower suction pressure, so holding the tanks low all day costs nothing in energy. With K-2 unloading at 80 %, constant setpoints from 65 to 92 mbar g used 29.13 to 28.25 MWh, inside the budget and cheaper than the preview. Without the boil-off target they would pass. What a low tank pressure does cost is the cargo flash of Equation (15.5). At 70 mbar g the flash peaks at 27.27 t/h during the unloading against 4.38 t/h with the preview, and the day's boil-off is 271.5 t against 210.0 with the preview and 218.4 with the MPC, as Figure 15.5b shows. The 230 t budget sits 11.6 t above the MPC's worst seed and 33 t below the best constant setpoint. The Ship button is ungraded for a similar reason, because stopping the ship at 5 h with only

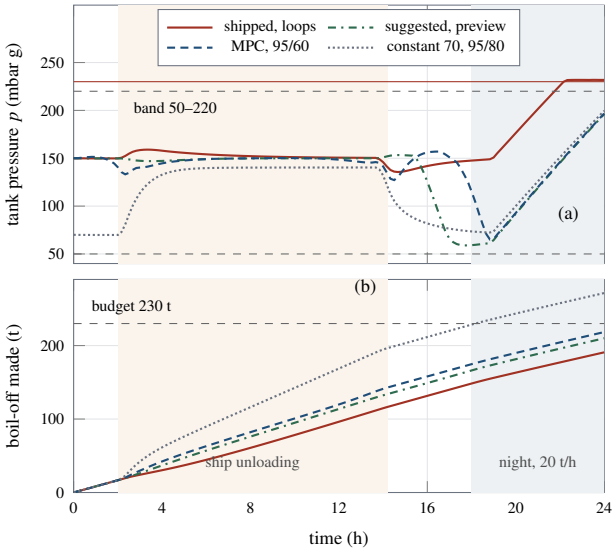

Figure 15.5 The challenge day under four strategies: the shipped loops (solid), the suggested preview (dash-dotted), the MPC (dashed) and the constant setpoint (dotted), with the ship's hours and the night shaded and labeled. (a) True tank pressure, with the judged band and the flare at 230 mbar g. (b) Boil-off made in the tanks since $t = 0$, against the 230 t budget, ending at 190.8 (shipped), 210.0 (preview), 218.4 (MPC) and 271.5 t (constant setpoint). The constant setpoint stays in the band without flaring but flashes cargo through the whole unloading.

36 000 m^3 in let the loops with K-2 at 95 and 80 % pass at 21.35 MWh.

> ✦ **DID YOU KNOW? The Two Cheapest Days Both Fail**
>
> From the simulator (Table 15.5). The two runs of the graded day with the least compressor energy both fail. The loops with K-2 at 95 and 60 % use 28.07 MWh but flare 6.61 t. A constant 92 mbar g with K-2 at 120 and 100 % uses 27.75 MWh but boils off 263.3 t. The suggested settings pass at 30.24 MWh.

One lever the target does not forbid is the recondenser setpoint. With the preview and the recondenser at 6 bar g the day passes at 25.82 MWh, 209.9 t and 59.0 to 203.6 mbar g, because a lower discharge pressure cuts the work per kilogram while the reduced capacity still covers the day. In a real terminal the high-pressure pumps' suction margin bounds that lever, and the model logs only the subcooling left there.

15.6 Technology in Practice

Figure 15.6 draws the loops of Section 15.4 with the hardware of a receiving terminal in place of the schematic's symbols.

What Measures It

The loops read a pressure transmitter on the shared vapor space, a pressure and a level transmitter on the recondenser, flowmeters on the compressors' discharge, the in-tank pumps, the recondenser's LNG, the send-out and the unloading line, a temperature element on the send-out gas, and an ultrasonic meter on the flare line. Tank B carries a density profiler. No instrument measures the boil-off itself, which the observer infers from the vapor balance. Each instrument model of Section 15.3, a lag and a noise level, stands for one of these devices. The flare meter's 5 s lag and 3 % noise, for example, are the ultrasonic meter below.

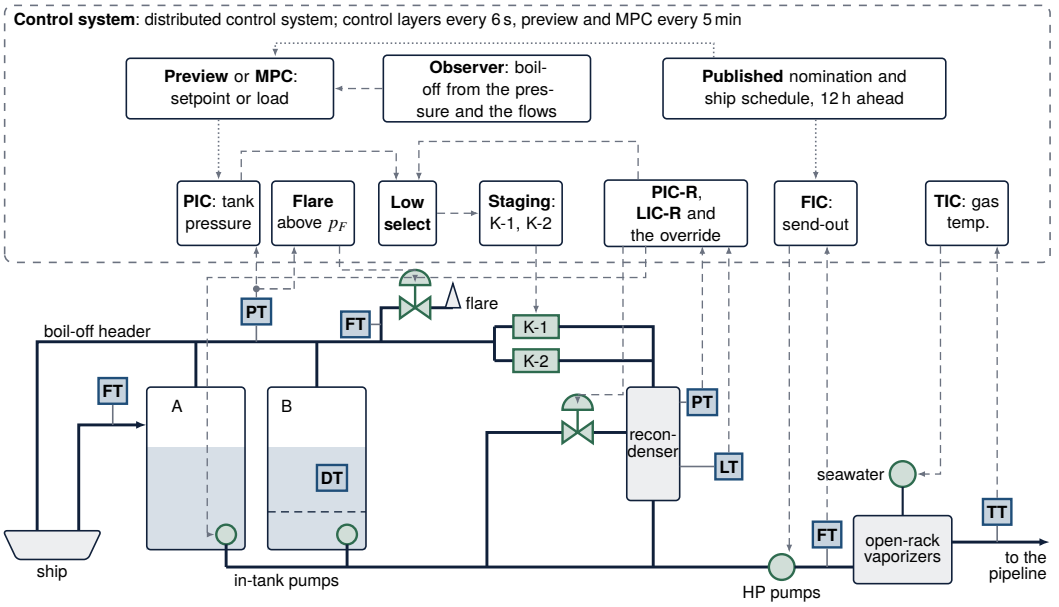

Figure 15.6 The loops of the terminal with their hardware. Lettered boxes are sensing devices with ISA letters: PT on the boil-off header and on the recondenser, LT on the recondenser, FT on the unloading arm, the flare line and the send-out, TT on the gas and DT, the density profiler, in tank B across its two layers. Green symbols are final elements: the in-tank pumps, the compressors K-1 and K-2, the valve that sets the recondenser's share of the LNG, the flare valve, the high-pressure pumps and the seawater pumps. Dashed arrows are signals, dotted arrows the published schedule and what the preview or the MPC writes. Notice that the compressors' load passes a low selector where the recondenser's override can cut it, and that no instrument measures the boil-off. Not to scale, and not a piping diagram.

> **TECHNOLOGY** **Sensor: The Ultrasonic Meter on the Flare Line**
>
> A transit-time ultrasonic flowmeter sends sound pulses between a pair of opposed sensors, with the flow and against it. The pulse carried downstream arrives sooner, and the flow follows from the two travel times in a form that does not depend on the speed of sound, so changes of composition, temperature or pressure have little effect. The meter also reports that speed of sound, a check that can reveal a fault in the sensors or their electronics. Several paths across the pipe are averaged where the velocity profile changes with the flow [19]. A flare is among the hardest services: almost no flow most of the time, a very large flow in an emergency, and a gas of varying composition. Multipath transit-time meters suit it, but their cost on large flare pipes limits them [19].

What Moves It

Five handles carry the control: the load of two on/off compressors, which recycle what the demand does not take, the share of the in-tank pumps' LNG sent through the recondenser, the in-tank pumps' flow, the speed of the high-pressure pumps and the seawater pumps of the vaporizers. Behind them stand the flare valve, which opens on the tank pressure once the handles have run out, as in case 1 at night and in case 8 at the rollover, and the tanks' relief valves above it.

TECHNOLOGY Final Element: The Flare and Relief Valves

Relief and safety valves open on the process pressure itself, with no signal from a controller. A relief valve opens in proportion to the overpressure. A safety valve, common in gas service, snaps fully open at its lift pressure and closes again only well below it, and the difference between the two pressures is its blowdown. In a pilot-operated valve a small pilot senses the pressure and sets the lift, so the setting is widely adjustable. An improper pilot setting played a role in a 1999 pipeline rupture, which shows why such settings need the same care as any instrument [19]. The chapter's flare valve opens in proportion over the 40 mbar above its set pressure, as a relief valve would, and its relief valves pass 20 t/h per mbar above 290 mbar g.

Where the Controller Runs

A distributed control system controls production within one location, and its supervisory and lower-level loops run with cycle times from milliseconds to minutes [20]. Control of a process plant is usually layered by time scale, with supervisory and predictive control acting in minutes and regulatory control in seconds [21]. The app's control layers run on a 6 s scan and its MPC every 5 min, so the chapter spans both of those layers. Its structure maps onto standard blocks: a pressure PI and a recondenser override PI into a low selector with external reset, a staging sequencer with hysteresis and minimum run and rest times, the recondenser's pressure and level PIs, a flow PI on the pumps' speed and a temperature PI on the seawater. The observer, the preview and the MPC sit above them as calculation blocks.

IN PRACTICE Before the Ship Pumps, Compare the Cargoes

United States rules for LNG plants require written transfer procedures. Before a transfer, personnel verify that the receiving tank has room for the cargo and that the incoming liquid, as it warms, will not overfill or overpressure it. For a bulk transfer into a partly filled tank they determine any difference in temperature or specific gravity between the incoming and the stored LNG and, if necessary, provide a means to prevent rollover. During the transfer they watch flow rates, liquid levels and vapor returns, and stop the flow by hand before overfilling or overpressure occurs [22]. Vaporization must keep the gas within the limits of the vaporizer and the downstream piping [23]. Cases 5, 6 and 8 rehearse three of these duties.

What Is Modernizing It

Three developments touch this chapter. The first is estimation. Soft sensors stabilize quality through online estimation and are used to validate online analyzers, and a survey of industrial applications names model maintenance as their most critical issue [24]. The boil-off observer of this chapter is a soft sensor of that kind, reading only pressures and flows. The second is predictive control, established in industry with only the first move applied [18], which the chapter's MPC follows on the published nomination. The third is design and operation studied together. A 2010 optimization recovered the unloading boil-off with the least energy as a strategy against flaring [3], and a 2012 dynamic simulation checked an improved recondenser control algorithm against fluctuating LNG and boil-off loads before its use at other terminals [5].

15.7 Limitations

The sizes, properties and thresholds were picked as teaching values for one self-consistent design, two tanks boiling 8.35 t/h at 150 mbar g beside a 450 t/h send-out, and none was fitted to a terminal. The LNG is one lump on a Clausius–Clapeyron line, within 0.07 K of pure methane at the recondenser but without composition or weathering outside case 8. The vapor space has one temperature, the recondenser is in equilibrium with no suction head modeled at the high-pressure pumps, the vaporizers are one bank without ice, and the cargo flashes at equilibrium as a top fill would. The rollover is illustrative. Its layers exchange nothing across their interface, its crossing time is set by the chosen initial gap and its surge by the declared overturn rate, so the case teaches the sequence and the remedy rather than a prediction.

The controllers know the declared properties exactly, and the observer lags a fast surge by tens of minutes. The compressors are on/off machines without load steps. The preview and the MPC read a perfect published schedule, and a nomination that changes after publication is outside the cases. The challenge grades boil-off because energy cannot separate a preview from a constant low pressure, and it leaves the recondenser setpoint free, which a real terminal's pump suction margin would not. All figures come from a single model at seed 20261020, the challenge alone being repeated over seeds 0 to 4 and at a halved step.

OPEN CHALLENGE **Measuring a Rollover Before Predicting One**

Industrial rollover incidents are rarely documented in the open literature, mainly because operating and commercial data stay confidential, and the report of the 1971 incident remains the most detailed record of a real event. Very little quantitative data exists. Progress, as a 2015 study set it out, means direct measurements of a rollover: the initial composition of each layer, its temperature and density against depth as they evolve, the boil-off rate, and the heat flux into the tank by position. Such measurements are costly, but the scale of the industry probably justifies them [15]. A model would then track the stability of the layers until its critical value is reached [11].

15.8 Conclusion

An LNG import terminal with a recondenser is a gas inventory whose outlet is a fraction of its product flow, and the simulator of this chapter logs the boil-off, the recondenser's capacity and the tank pressure beside the instruments and controllers that see them. The capacity is $r = 0.1658$ per kilogram of LNG and $r/(1+r) = 0.1422$ per kilogram of send-out, which puts the minimum send-out at 49.6 t/h by hand and 50.0 t/h on the running plant. A 30 t/h night fills the recondenser in 91.4 min and reaches the flare at 4.69 h whatever the compressors are asked to do. The ship's displaced vapor and its cargo flash of 0.40 t/h per mbar make the tank pressure a lever on the boil-off itself. Read from the published nomination, the preview passes the graded day at 30.24 MWh and the MPC at 29.72 to 29.80 MWh, while a constant low pressure flashes enough cargo to raise the day's boil-off to 271.5 t. The rollover is declared, with a chosen gap and overturn rate. A fuller model would add a multicomponent liquid with weathering, interfacial transfer between layers, the high-pressure pumps' suction margin, load-stepped compressors and an uncertain nomination. The simulator shows that the boil-off a terminal can take back is decided by the send-out it is asked to make.

KEY POINTS

- The tanks boil all the time, and the gas can leave only by condensing into the LNG the terminal sends out, so a night of low demand drives the tank pressure toward the flare.
- The recondenser takes at most 0.1422 kg per kilogram of send-out, so the minimum send-out is 49.6 t/h by hand and 50.0 t/h on the running plant.
- A 30 t/h night reaches the flare at 4.69 h whatever the compressors are asked to do.
- Read from the published nomination, the preview and the MPC pass the graded day, while a constant low pressure raises the day's boil-off to 271.5 t through cargo flash.

Where else this applies. A surplus integrated against a sink that another decision sets returns in Volume 11, Chapter 14, and a forecast acted on before the demand arrives in Chapter 12.

References

[1] E. Querol, B. Gonzalez-Regueral, J. García-Torrent, and M. J. García-Martínez, "Boil off gas (BOG) management in Spanish liquid natural gas (LNG) terminals," *Applied Energy*, vol. 87, no. 11, pp. 3384–3392, 2010.

[2] C. Park, K. Song, S. Lee, Y. Lim, and C. Han, "Retrofit design of a boil-off gas handling process in liquefied natural gas receiving terminals," *Energy*, vol. 44, no. 1, pp. 69–78, 2012.

[3] C. Liu, J. Zhang, Q. Xu, and J. L. Gossage, "Thermodynamic-analysis-based design and operation for boil-off gas flare minimization at LNG receiving terminals," *Industrial & Engineering Chemistry Research*, vol. 49, no. 16, pp. 7412–7420, 2010.

[4] Đ. Dobrota, B. Lalić, and I. Komar, "Problem of boil-off in LNG supply chain," *Transactions on Maritime Science*, vol. 2, no. 2, pp. 91–100, 2013.

[5] Y. Li and X. Chen, "Dynamic simulation for improving the performance of boil-off gas recondensation system at LNG receiving terminals," *Chemical Engineering Communications*, vol. 199, no. 10, pp. 1251–1262, 2012.

[6] Y. M. Kurle, S. Wang, and Q. Xu, "Dynamic simulation of LNG loading, BOG generation, and BOG recovery at LNG exporting terminals," *Computers & Chemical Engineering*, vol. 97, pp. 47–58, 2017.

[7] A. J. Kidnay, W. R. Parrish, and D. G. McCartney, *Fundamentals of Natural Gas Processing*, 3rd ed. CRC Press, 2019.

[8] S. Mokhatab, J. Y. Mak, J. V. Valappil, and D. A. Wood, *Handbook of Liquefied Natural Gas*. Elsevier, 2014.

[9] M. W. Shin, D. Shin, S. H. Choi, and E. S. Yoon, "Optimal operation of the boil-off gas compression process using a boil-off rate model for LNG storage tanks," *Korean Journal of Chemical Engineering*, vol. 25, no. 1, pp. 7–12, 2008.

[10] C. Migliore, C. Tubilleja, and V. Vesovic, "Weathering prediction model for stored liquefied natural gas (LNG)," *Journal of Natural Gas Science and Engineering*, vol. 26, pp. 570–580, 2015.

[11] P. Arjomandnia, M. O. Tade, V. Pareek, and E. F. May, "Analysis of available data from liquefied natural gas rollover incidents to determine critical stability ratios," *AIChE Journal*, vol. 60, no. 1, pp. 362–374, 2014.

[12] J. Heestand, C. W. Shipman, and J. W. Meader, "A predictive model for rollover in stratified LNG tanks," *AIChE Journal*, vol. 29, no. 2, pp. 199–207, 1983.

[13] S. Bates and D. S. Morrison, "Modelling the behaviour of stratified liquid natural gas in storage tanks: a study of the rollover phenomenon," *International Journal of Heat and Mass Transfer*, vol. 40, no. 8, pp. 1875–1884, 1997.

[14] R. G. Scurlock, *Stratification, Rollover and Handling of LNG, LPG and Other Cryogenic Liquid Mixtures*, ser. Springer-Briefs in Energy. Springer International Publishing, 2016.

[15] P. Arjomandnia, "Simulation of LNG rollover in storage tanks," Ph.D. dissertation, Curtin University, Perth, Australia, 2015, open access, hdl.handle.net/20.500.11937/1544.

[16] U. Setzmann and W. Wagner, "A new equation of state and tables of thermodynamic properties for methane covering the range from the melting line to 625 K at pressures up to 1000 MPa," *Journal of Physical and Chemical Reference Data*, vol. 20, no. 6, pp. 1061–1155, 1991.

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

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

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

[22] Pipeline and Hazardous Materials Safety Administration, "49 CFR 193.2513, liquefied natural gas facilities: Transfer procedures," Code of Federal Regulations, Title 49, Part 193, edition of October 1, 2025, govinfo.gov, 2025.

[23] ——, "49 CFR 193.2503, liquefied natural gas facilities: Operating procedures," Code of Federal Regulations, Title 49, Part 193, edition of October 1, 2025, govinfo.gov, 2025.

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

15.9 Exercises

15.1 Concept check · 2 min

In case 1 the nomination falls to 30 t/h for the night. The two compressors can move 20 t/h, nearly three times the 6.94 t/h the tanks boil, yet the tank pressure climbs to the flare. Why?

(a) The pressure loop is in Manual, so the compressor load stays at the evening value

(b) The compressors' capacity falls with the suction pressure, so they lose their 20 t/h at night

(c) The recondenser condenses only a fraction of the send-out, and its override cuts the compressors to it

(d) The vaporizers cannot warm so small a send-out, so the high-pressure pumps return LNG to the tanks

15.2 Concept check · 1 min

In case 3 a ship unloads 12 000 m^3/h into tanks held at 150 mbar g. The tank pressure setpoint is raised to 200 mbar g before Run. What happens to the cargo flash and to the displaced vapor the ship does not take back?

(a) Both fall, because a higher pressure compresses the vapor the cargo displaces

(b) Both rise, because a higher tank pressure raises the cargo's arrival temperature

(c) The flash rises and the displaced vapor is unchanged, because warmer tanks boil more

(d) The flash stops and the displaced vapor, set by the cargo's volume, stays about the same

15.3 Concept check · 2 min

In case 8 tank B holds a denser, warmer lower layer. Before the layers roll over, which reading is the warning that the handling literature describes?

(a) Tank B boils less than it would as one layer, because the lower layer keeps the heat it receives

(b) Tank B boils more than it would as one layer, because the lower layer is warmer than the surface

(c) The tank pressure creeps up, because the lower layer vents its vapor into the shared space

(d) The observer's boil-off estimate jumps, because the observer reads tank B's density profiler

15.4 Theory and calculation · 6 min

The recondenser runs at 8 bar g, where the declared saturation line gives $T_{sat}(p_R) = -126.258\,^{\circ}\mathrm{C}$. The LNG reaches it at $T_{LNG} = -159.144\,^{\circ}\mathrm{C}$ with $c_p = 3.5\,\mathrm{kJ/(kg\,K)}$. The compressed boil-off arrives at $T_d = -15.21\,^{\circ}\mathrm{C}$ with $c_{p,v} = 2.2\,\mathrm{kJ/(kg\,K)}$, and the latent heat at p_R is $\lambda_{R,lat} = 450\,\mathrm{kJ/kg}$.

(a) Find the heat each kilogram of compressed boil-off must give up to condense, λ_R. *Checkpoint:* 694.3 kJ/kg.

(b) Find the boil-off condensed per kilogram of LNG through the recondenser, r, and per kilogram of send-out with all the LNG through it. *Checkpoint:* 0.1658 and 0.1422.

(c) The in-leak and the jetty line make 7.06 t/h of boil-off. Find the lowest send-out at which the recondenser can take it, and the capacity at a 30 t/h night. *Checkpoint:* 49.6 t/h and about 4.27 t/h.

15.5 Theory and calculation · 6 min

Case 1 opens with both tanks 80 % full. Each tank holds $V_t = 160\,000\,\mathrm{m^3}$, the header $V_{hdr} = 500\,\mathrm{m^3}$, and the vapor has $\mathscr{M} = 16.8\,\mathrm{kg/kmol}$ at $T_v = 133.15\,\mathrm{K}$, with $R = 8314.46\,\mathrm{J/(kmol\,K)}$.

(a) Find the vapor-space capacitance $\partial M_v/\partial p$ from Equation (15.1), in kg per mbar. *Checkpoint:* 97.9 kg/mbar.

(b) At 4 h the Balance view leaves 2.32 t/h of boil-off as "pressure rising". Find the rate of rise. *Checkpoint:* 23.7 mbar/h.

(c) The true pressure is 213.7 mbar g at 4 h. Estimate when the tanks reach the flare at 230 mbar g. *Checkpoint:* 4.69 h.

15.6 Theory and calculation · 5 min

A ship unloads 12 000 m^3/h, which is $F_c = 5398$ t/h of cargo, into tanks at 150 mbar g. The cargo arrives 0.0857 K above the tanks' saturation temperature, with $c_p = 3.5$ kJ/(kg K) and $\lambda = 510$ kJ/kg. Near 150 mbar g the saturation temperature rises 0.010 79 K/mbar.

(a) Find the cargo flash at 150 mbar g. *Checkpoint:* 3.18 t/h.

(b) Find how much the flash falls for each mbar the tanks are held higher. *Checkpoint:* About 0.40 t/h per mbar.

(c) Find the tank pressure above which the cargo no longer flashes. *Checkpoint:* 157.9 mbar g.

15.7 App challenge · 9 min

Load case 1, A night at low send-out (Manual), and change no setting. Run to 300 min at the default speed of 30 with the default seed and export the CSV file.

(a) Find the time of the first row with `rc_full` equal to 1 and of the first row with `flare_meas_tph` above zero. *Checkpoint:* 91.5 min and 281.5 min, within 0.5 min.

(b) In the row at 240 min read `truth_bog_tph`, `truth_comp_tph`, `truth_rc_capacity_tph` and `rc_capacity_tph`. *Checkpoint:* 6.94, 4.51, 4.53 and 4.21 t/h, within 0.02 t/h.

(c) Press the reload button, run at speed 5, pause at any time between 110 and 130 min, switch the header to Auto and run on to 300 min. Read `comp_load_req_tph` and `truth_comp_tph` at 300 min and the first row with `flare_meas_tph` above zero. *Checkpoint:* 20.00 t/h, 4.49 t/h and 281.5 min.

15.8 App challenge · 8 min

Load case 6, A compressor trip during unloading: flare or not. Use the default seed and speed 60, at which 12 h take 12 s.

(a) Run as shipped to 720 min. Read the Flared tile and, in the exported file, the first row with `flare_meas_tph` above zero. *Checkpoint:* 13.96 t and 369 min (6.15 h).

(b) Press the reload button and run at speed 10. Pause at any time between 2.2 and 5 h, set Unloading rate to 7000 m^3/h, and run to 720 min. Read the Flared and Peak pressure tiles. *Checkpoint:* 0.00 t flared, peak between about 161 and 208 mbar g.

15.9 App challenge · 9 min

Load case 7, The nomination as a preview. Run three times to 720 min at speed 60 with the default seed, choosing Tank pressure control before Run: Loops (PI), then Loops + nomination preview, then MPC on the compressors. For each run read the Flared, Compressor energy and Peak pressure tiles.

(a) Report the three runs. *Checkpoint:* Loops 7.15 t, 10.19 MWh, 233.0 mbar g. Preview 0.00 t, 11.43 MWh, 197.4 mbar g. MPC 0.00 t, 11.30 MWh, 200.1 mbar g.

(b) In the preview run, find in the exported file the first row with `p_sp_eff_mbarg` below 150 and the value of `truth_p_mbarg` at 300 min. *Checkpoint:* 233.5 min and 98.1 mbar g.

Accept an energy within 0.02 MWh and a peak within 0.5 mbar. Explain why the preview costs more energy than the loops.

15.10 Application area · 3 min

With the preview and K-2 at 95 and 60 %, the graded day passes at 30.24 MWh with the recondenser at 8 bar g and at 25.82 MWh with it at 6 bar g.

(a) Explain why the lower recondenser pressure saves energy and what it costs in capacity. *Checkpoint:* Less work per kilogram, smaller r.

(b) An engineer proposes to run a real terminal's recondenser at its lowest pressure for the same saving. Name the limit the model leaves out and the logged quantity that hints at it. *Checkpoint:* The high-pressure pumps' suction margin. The subcooling at their suction.

15.11 Application area · 4 min

A terminal engineer wants to use case 8 to decide how long a stratified tank may be held before its lower layer must be sent out.

(a) Which results of the case are set by declared choices rather than by physics the model resolves? *Checkpoint:* The crossing time and the surge.

(b) What can the case support, and which remedy does the app teach? *Checkpoint:* The sequence and the remedy. Tank B first.

15.12 Review the agent · 5 min

An AI assistant is asked how to pass the graded day of case 9 more cheaply than the preview and replies: "Keep Tank pressure control on Loops (PI), set the Tank pressure setpoint to 70 mbar g, K-2 loads above to 95 % and K-2 unloads below to 80 %. The low pressure leaves room for the night, so nothing flares, and a loaded compressor draws slightly less power at a lower suction pressure, so the day uses about 29 MWh, inside the 30.6 MWh budget and below the preview's 30.24 MWh. No forecast is needed." Find the flaw and say how you would check it.

CHAPTER 16

Summary and Outlook

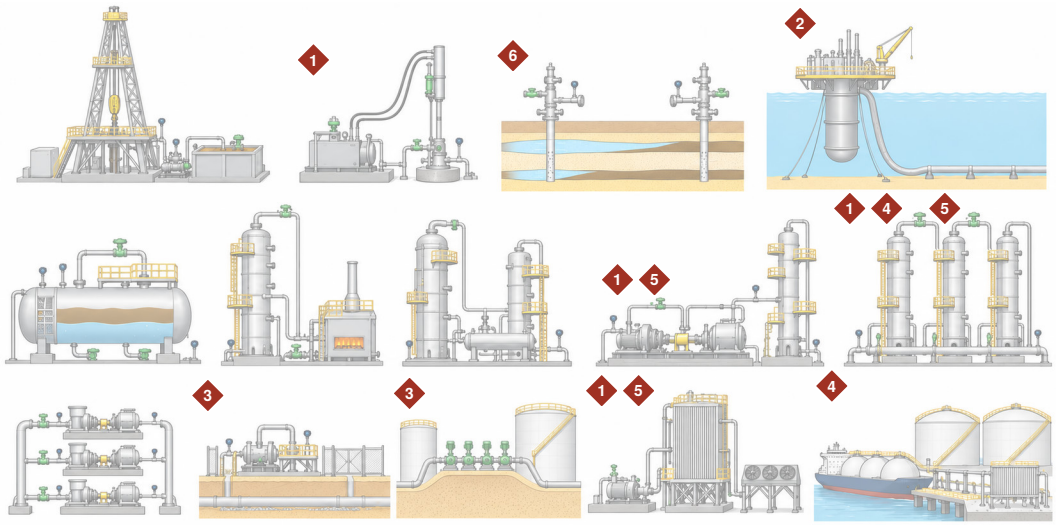

Open gaps 1 drifting inferences 2 bounded minimum 3 leak or meter 4 chain coordination 5 models without the limit 6 payoff in months

Notice that each number stands beside every plant where a chapter met that gap, so the first gap appears four times and the NGL train carries three. Illustration, not to scale.

IN PLAIN TERMS What Fourteen Oil and Gas Plants Taught

The main idea. Fourteen chapters followed oil and gas from a well being drilled to a terminal that sends gas on to its customers. The same few control ideas came back in plant after plant, such as one valve that answers to whichever limit is closest.

The challenge. The thing that must stay inside its limit often cannot be seen directly. It is worked out from other readings or reported late, so a controller can hold a number that looks right while the real value drifts away.

Why it matters. A reading that looks right but is wrong can let a pipeline plug with ice-like solids, hide a leak or send out gas that is too wet.

Fourteen chapters followed oil and gas from a well being drilled to a terminal sending gas into a grid. Table 1.2 listed, for each plant, the limit that binds, how the controller knows it and the structures the chapter uses. This chapter adds what worked, what failed and the measured result, and reads the comparison for the two questions of the volume, which techniques recur from well to terminal and where measurement, not actuation, sets the limit. It then gives a short guide for choosing a design, looks at where practice is heading, reports how the volume was made with AI agents and what they got wrong, and lists the problems that remain open.

16.1 What the Chapters Showed

Tables 16.1 and 16.2 compare the chapters, split where the fluid leaves treating for the liquids recovery plant. The second column condenses the first two columns of Table 1.2. The last column gives the

Table 16.1 What Decided Each Result, Chapters 2 to 8

Chapter	Binding limit, how known	What worked	What failed	Measured result
2, drilling rig	Cuttings, vibration and the pressure window. Cuttings answer a weight move minutes late.	A supervisor that climbs the rate gradient above three PI loops and backs off the most violated limit.	The fixed manual point. Limits set inside the starting state.	28.3 ft against 16.3 ft in 900 s, 97 against 61 % of the best feasible rate (Table 2.4). Cuttings over their limit 42 % of the time while it hunts (Section 2.6).
3, rod-pumped well	Fillage and fluid level, computed through the downhole card from a surface load cell.	An optimizer on the fillage limit, 5.25 SPM at a 2.95 m stroke.	9.0 SPM by hand. A pump-off timer. A card computed with a wrong damping.	55.4 bbl in 6 h with no pound damage, against 56.3 and 392.5 damage units by hand and 51.7 on the timer (Table 3.5). Card error 1.33 to 50.1 % (Section 3.6).
4, smart producer	Drawdown, sand and water from a model and a late meter. The flood fronts are never measured.	Four PI loops in external-reset form at a low selector, under a 180-day planner.	The shipped choke at a 51.4 bar mean drop. A 30-day horizon.	Oil 2.33 times the fixed choke's in 180 days, against 1.07 shipped and 2.00 at 30 days (Table 4.8). Sand 0.30 against 3.24 g/s with the sand loop off (Table 4.7).
5, deepwater flowline	Hydrate margin at the weakest point, 3.23 h downstream, estimated from six transmitters and two meters.	Methanol feedforward on the predicted demand, ratioed to water. A 30 min lead before a restart.	Trim on the estimate alone. A ratio on the liquid rate. The MPC on its reduced model.	Turndown margin –4.43 °C on 47.4 m^3 against 1.03 °C on 28.3 m^3 (Table 5.5). In case M the MPC spends 93.3 m^3 with 8.3 °C h exposed, the feedforward 77.8 and none (Table 5.7).
6, separator	Levels between trips, read directly. Oil and water quality from analyzers behind sample delays.	Averaging level tuning. A second hydrocyclone bank. The supervisor with averaging.	Tight tuning on a slug. The supervisor alone. A longer analyzer delay.	Oil to the second stage 317.5 against 219.1 m^3/h. Off-spec water 1.39, 4.84 and 10.77 m^3 at delays of 30, 120 and 480 s (Table 6.7). In-spec oil 69.8 to 98.7 m^3 (Table 6.9).
7, glycol dehydrator	Water in the gas, read 10 min late. A reboiler ceiling and a pump limit.	A ratio on the water load. Purity by heat, then stripping gas. Economic MPC under the energy budget.	Fixed circulation on a gas step. Ratio and trim under a fuel limit, where the trim makes the gas wetter.	Gas-step peak 6.86 against 7.87 lb/MMscf (Table 7.4). Over the day the worst sample is 6.73 at 232 kW on the MPC and 10.47 on ratio and trim (Table 7.5).
8, amine sweetener	Hydrogen sulfide from a sampled analyzer. Rich loading inferred from an acid-gas meter.	Ratio circulation, a loading floor at a high select, antifoam on pressure drop, analyzer trim.	Fixed circulation. The product analyzer as the only guard of the loading.	True gas over 4 ppm at 13.8 min, analyzer at 24.7 min (Table 8.4). The shift passes at 3.78 ppm, loading 0.458 and 5.25 MW, and Fixed stops at 214.5 min (Table 8.5).

measured result with the table or section of the chapter that reports it.

The columns of the two tables make five points. Figure 16.1 draws one pair of results from each of thirteen chapters.

The selector recurs, and its wiring decides. Six chapters share one actuator among limits through a low selector and three open one through a high selector (Section 1.4.8). On the smart producer

Table 16.2 What Decided Each Result, Chapters 9 to 15

Chapter	Binding limit, how known	What worked	What failed	Measured result
9, turboexpander plant	Methane in the liquid, from a delayed chromatograph. Tray temperature as a proxy. Flooding.	Tray setpoint first, then pressure, with a 20 K return margin.	Either change alone. Both setpoints stepped. An MPC whose model has no quality or flooding.	Worst sample C1/C2 1.51 against 3.65 shipped, 11.36 stepped and 3.69 % on the MPC at 101.7 % vapor load (Table 9.5).
10, NGL train	Ethane in propane, measured one column downstream of the cut that sets it.	The analyzer cascaded onto the deethanizer inferential. A cascade retuned in hindsight.	The inferential alone after it drifts. An MPC holding constant what the surge drum still delivers.	Drifted inferential 0.705 against a 1.80 mol% target, cascade 1.751 (Table 10.4). MPC 153 against 513 min over, peak 3.28 against 2.54 mol% (Table 10.6).
11, compressor station	Surge margin of each machine, computed from flow and head on the commissioned map.	Equal surge margin sharing, a sequencer, the anti-surge valve opened at once on a trip.	Equal speed with one machine fouled. A trip left to the anti-surge PI.	Recycle 15.01 against 0 t in an hour. Surge 0.75 s after a trip (Table 11.4). Over six hours 30.74 against 97.75 t at equal speed, and 31.97 on the MPC (Table 11.5).
12, gas line	Delivery floor. Line pack inferred from three transmitters. A demand forecast.	A 30 min forecast feedforward at a lower setpoint. An afternoon pack for a peak over capacity.	Feedback alone at 48 bar. A 1 % leak threshold against a biased meter.	52 min under the floor on feedback, none with the forecast. A 3 % leak found in 34 min at a 2 % threshold, a 1.5 % meter error alarmed in 34 min at 1 % (Table 12.4).
13, products line	Suction and discharge at four stations, a hill near the terminal, batch cuts at a densitometer.	Low selectors in external-reset form on every pump. Back pressure by MPC on an inferred hill.	Discharge control without selectors. A tight balance across a batch change.	Without selectors a cut at 10 min trips every station by 20.2 min, with them a trip settles at 781 m^3/h (Section 13.5). Over the day 17 false alarms shipped and none suggested (Table 13.5).
14, LNG train	Minimum approach, computed from an analysis refreshed every 12 min. Driver power.	A suction pressure loop under the product temperature. A low selector on the feed.	The computed approach after a relief. The MPC on the feed cut and in peak power.	Estimated minimum 6.5 against 4.2 °C true (Section 14.6). MPC 18.6 against 85.1 min over the driver, peak 2.78 against 1.31 MW, feed cut 18.1 against 6.0 min off spec (Table 14.8).
15, LNG terminal	Tank pressure below the flare. Recondenser capacity from subcooling.	A preview from the published nomination with staging at 95 and 60 %. An MPC on the same schedule.	The loops alone before a low night. A constant low pressure that flashes cargo.	A 30 t/h night fills the recondenser in 91.4 min and flares at 4.69 h (Table 15.4). Over the day the preview uses 30.24 MWh and the MPC 29.75, and a constant pressure boils off 271.5 t (Table 15.5).

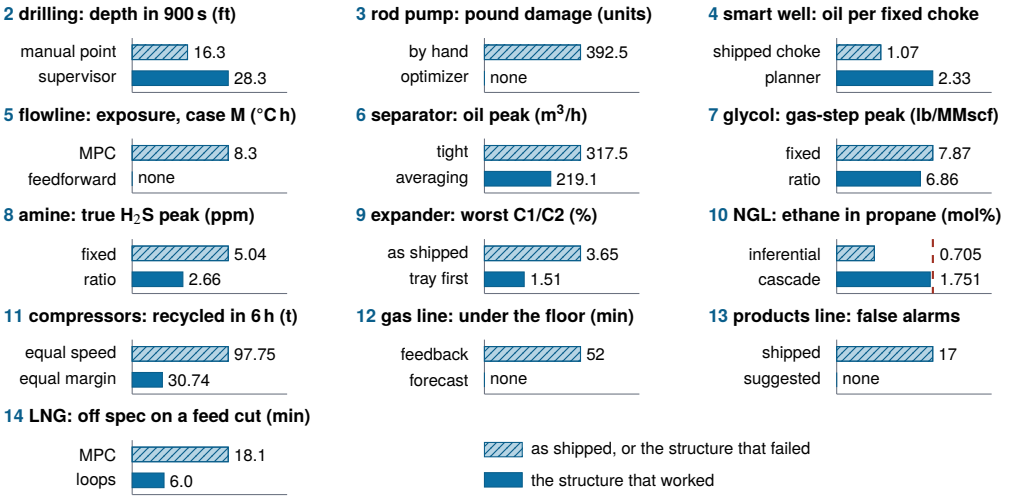

Figure 16.1 One pair of results from each of thirteen chapters, with the numbers as Tables 16.1 and 16.2 print them (the sweetener pair as this section quotes it). Each panel has its own scale from zero. A larger value is better in the panels of Chapters 2 and 4, and the NGL panel marks its 1.80 mol% target with a dashed line. Table 16.2 gives the LNG terminal no like pair. Notice that in four panels the structure that worked removed the harm entirely.

the sand loop held 0.30 g/s where the same rate target without it passed 3.24 g/s. On the products line the selectors turned a station trip into a steady lower flow where discharge control alone tripped every station. On the compressor station the valve of a tripped unit had to open at once, because its anti-surge PI let it surge in 0.75 s. All of this rests on the condition of Figure 1.6. The same two loops with the same tunings held a drawdown limit when both reset terms followed the selected output, and broke it by 1.2 bar for 21.8 min when the waiting loop had wound up.

An inventory decides where a disturbance goes, not whether it arrives. The separator stored a 15 m^3 slug under any tuning, and the tuning decided only whether 317.5 m^3/h or 219.1 m^3/h reached the second stage. The gas line took 12 min to pass a speed step to the city gate, and a peak above its capacity could be met only by packing it beforehand, for about $1,400. The terminal's recondenser filled on a low night whatever the compressors were asked to do. In each case the handle that worked was set before the disturbance: the averaging tuning before the slug, the pack before the peak, a lower tank pressure before the night.

A measured load ahead of a late quality beats a faster loop on the quality. A ratio on the glycol pump held a gas step at 6.86 lb/MMscf where fixed circulation reached 7.87. A ratio on the amine circulation held the true hydrogen sulfide at 2.66 ppm where fixed circulation reached 5.04 ppm. Methanol sent on the predicted demand spent less and kept the margin positive, and the gas line's forecast removed 52 min under the floor. The load must be the one the limit depends on. A methanol ratio on the liquid rate let a water breakthrough thin the inhibitor at the riser base to 9.0 wt% and the margin to −1.73 °C, while a ratio on the water held 1.54 °C (Section 5.6). The glycol ratio keeps the share of water removed, so a warmer inlet passed through it, and the amine ratio reads the last laboratory assay, so a sourer well took the loading past its cap.

Where measurement is the limit, the loop holds the measurement. The three cases of Section 1.4.9 all appeared. A late analyzer turns its delay into product. The amine analyzer reported 10.9 min after the gas, the glycol analyzer 12 min after the water crossed its limit, and on the separator

a longer analyzer delay sent 1.39, 4.84 and 10.77 m^3 of off-spec water overboard. An inferred constraint is held exactly while the truth drifts. The rod pump's card error rose from 1.33 to 50.1 % with a wrong damping, the turboexpander held its tray temperature while the true product reached 3.33 % C1/C2 (Section 9.5), and the subsea estimator with its corrections off read a margin of 3.16 °C where the truth was 0.63 °C, because a colder sea and more heat loss look the same to two thermometers (Section 5.6). The direction of the error decides its cost. The drifted NGL inferential gave away product, while the stale analysis of the LNG train read the approach wide when it was narrow. An inventory computed from a few transmitters cannot see a small leak. The gas line could not tell a 1.5 % leak from a 1.5 % meter error, and on the products line the 60 of 160 detector settings that stayed quiet through a batch change all missed a 0.2 % leak for an hour (Section 13.5).

Prediction paid where it carried a preview or priced several handles. The glycol MPC passed the energy-limited day that ratio and trim failed, and the terminal's MPC passed at 29.75 MWh. On the gas line the MPC spent 10.4 min under the floor against 54.6 min for the cascade with the same forecast, and 301.6 min with its preview off (Table 12.6). Prediction lost where its model left out the limit that binds. The turboexpander's model had no quality or flooding, the NGL plan held constant a disturbance the surge drum was still delivering, the subsea plan smeared a lean front, and the LNG plan cut the minutes over the driver while its peak over it rose from 1.31 to 2.78 MW. On the compressor station it matched equal margin sharing within 0.07 % of fuel and did not beat it. Volume 9 treats that choice in full.

> ⧖ HISTORY **1927 to 2016: The Automatic Driller, Still Unfinished**
>
> The timeline of Chapter 1 opened with a 1927 patent for a dummy driller and closed with automatic control of downhole pressure by a topside choke, applied offshore in 2007 [1]. Mud-pulse telemetry, which sends downhole readings up as pressure pulses at 0.5 to 3.0 bits a second, remained the most common method in 2016. Wired drill pipe, which carries 57,000 bits a second and can send commands down for closed-loop control, was first used commercially in 2006 [2]. A 2016 review found drilling digitized, visual and automated, yet its degree of automation still not enough for real industry [2]. The other chapters stand at the same point: the limit that binds is often what no instrument reads.

16.2 Choosing a Method

The chapters suggest an order of questions for a new plant in the chain. Each names the chapter where its answer decided the result. Figure 16.2 sets the same questions out as a flowchart.

1. *Which limit binds, and is it measured, inferred or late?* Fill a row of Table 1.2 before choosing any structure. The rest of the design follows from that column.
2. *Does one actuator serve several limits?* Use a selector, and make every waiting loop's reset follow the applied output (Chapter 4, Chapter 13). Test the wiring by a disturbance that hands the actuator from one loop to the other.
3. *Is there an inventory between the disturbance and the product?* Decide where the disturbance should go, and tune the level or pack for it (Chapter 6). If the capacity behind the inventory is the limit, only an earlier move helps (Chapter 12, Chapter 15).
4. *Is the load measured before the quality?* Add feedforward on the load the limit depends on, and keep a slow trim from the analyzer (Chapter 7, Chapter 5).
5. *Is the binding variable inferred?* Keep a second source that can disagree with it, such as an analyzer cascade or a second transmitter, and alarm on the disagreement (Chapter 10, Chapter 6). Ask which direction of error is unsafe.

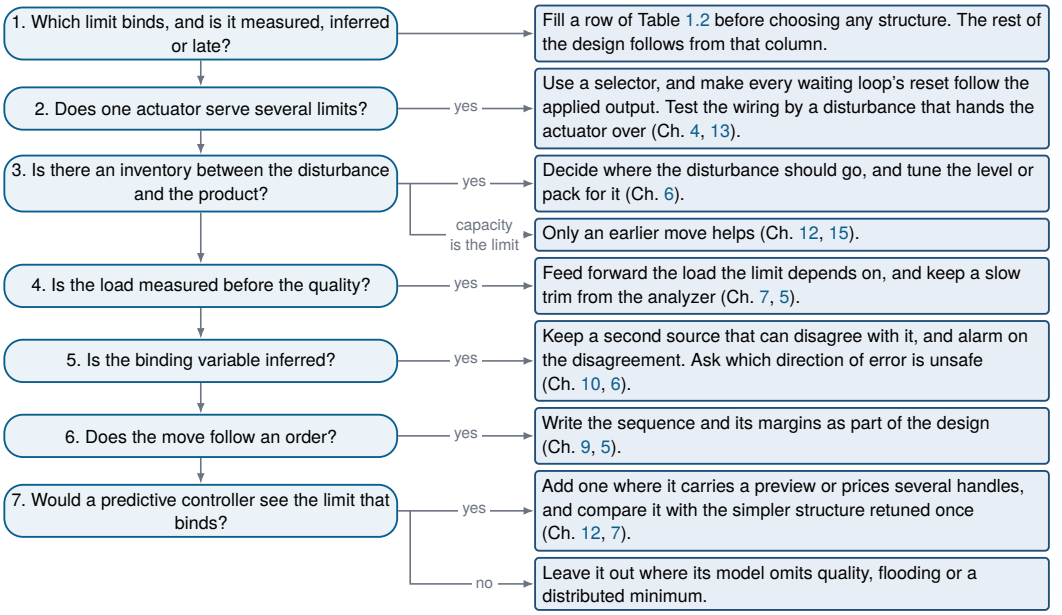

Figure 16.2 The seven questions of this section in order, each with the answer the list gives and the chapters where that answer decided the result. The third and the last question branch. Notice that the first question, which limit binds and how it is known, settles what every later answer acts on.

6. *Does the move follow an order?* Write the sequence and its margins as part of the design. The turboexpander needed both the order and a wider return margin, and the restart needed its methanol lead (Chapter 9, Chapter 5).
7. *Would a predictive controller see the limit that binds?* Add one where it carries a preview or prices several handles, and compare it with the simpler structure retuned once (Chapter 12, Chapter 7). Leave it out where its model omits quality, flooding or a distributed minimum.

16.3 Where Practice Is Heading

Automation reaches the well. Drilling has traditionally been operated by hand. Automatic managed pressure drilling, with downhole pressure held by topside choking, was applied offshore in 2007, and its requirements were stated for normal operations such as connections and for failures such as gas kicks and a blocked choke [1]. A 2013 review found broad interest in drilling-systems automation but no agreement on what it entails, and named well complexity, data overload and the loss of skilled people among its drivers [3]. Real-time production optimization has been surveyed as a chain of data acquisition, model updating and planning that ends in choke, valve, compressor and pump settings [4]. Optimization of gas transmission has its own body of work [5]. In the editor's view the supervisors of Chapter 2, Chapter 3 and Chapter 4 are the form this takes on one well, and the work will be in coordinating many of them.

Measurement moves into models. Virtual flow metering computes multiphase rates from pressures and temperatures instead of a hardware meter. A 2020 review found it able to serve as a standalone meter or a backup, and asked for automatic calibration as fluid properties change, for methods that state the uncertainty of the estimate, and for state estimation that stays accurate in transients [6]. Slug control has been designed with and without pressure sensors upstream of the riser [7]. Digital twins of oil and gas assets were judged in 2020 to be at an early stage, limited to isolated applications [8].

Hydrate practice has been described as moving from avoidance toward the management of risk [9], which is what a computed margin and an exposure integral, as in Chapter 5, make possible. In the editor's view every limit in Table 1.2 that is now inferred will be inferred by a larger model, and the chapters show why each such model needs a second source that can contradict it.

Integrity and emissions are watched more closely. Pipeline leak detection methods have been reviewed for their advantages, limitations and performance, together with the standards that govern them [10]. For fugitive methane, a review of six classes of technology, from handheld instruments to satellites, found minimum detection limits from under 1 g/h to 7.1×10^6 g/h, found the uncertainties poorly constrained, and proposed screening sites with a wide survey before confirming the source up close [11]. In the editor's view the trade measured in Chapter 12 and Chapter 13, a threshold low enough to find a small leak and high enough to ignore a meter, will be met more often as these programs grow.

Fiber turns lines and wells into thermometers. A 2022 review calls Raman distributed fiber sensing of temperature mature, widely applied to finding pipeline leaks and used to check the integrity of oil and gas wells [12]. The same review states its limits: the sensing distance and the spatial resolution cannot both be had, and a cleaner signal costs measurement time. In the editor's view a measured temperature profile is the second source that gaps 2 and 3 of Section 16.5 ask for, on a flowline read through two thermometers and on lines that look for a leak in an inventory computed from a few transmitters.

Electric drives join the power system. New LNG plants are increasingly driven by electricity under a site power management system, which in an electrical contingency typically disconnects loads until the consumption meets the power that remains [13]. A 2022 analysis found that a mixed fluid cascade plant could instead keep running on 30 % less power or more [13]. In the editor's view the driver limit of Chapter 14 and the load sharing of Chapter 11 will more often be set by the power system than by the machine.

Optimizers learn the machines they run. Online feedback optimization steers a plant toward its optimum from measurements and a linear model, without solving a nonlinear optimization problem online. A 2023 study added a model that learns how each compressor departs from its map. On a computed station sharing load among parallel machines, imperfect maps cost 5 % more power and the learning cut the excess to 0.8 %, with a method meant to run at the edge with moderate modeling effort [14]. In the editor's view this answers the drifting maps of Chapter 11, so far only on a model. Figure 16.3 sets the developments of this section on three lanes, by how far each has come into use.

16.4 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 [21]. 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.

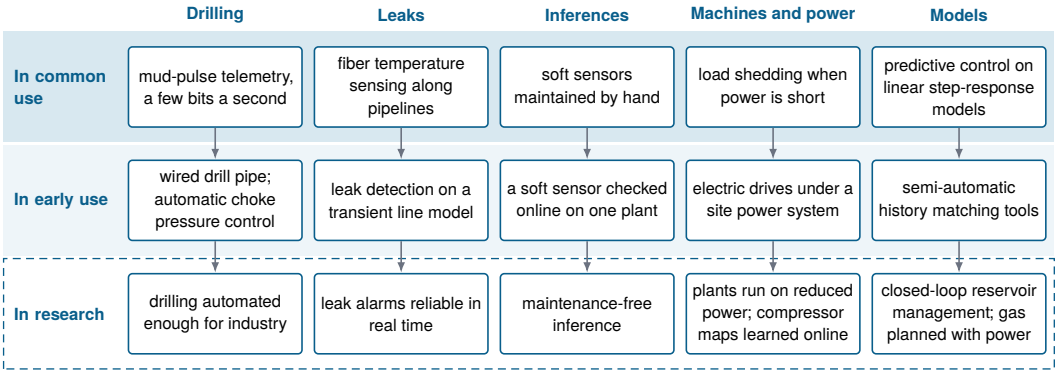

Figure 16.3 The developments of this section in five threads. The top lane holds what the cited source calls the most common, typical or widely applied practice, the middle lane what is in service at some sites or sold as a tool, and the bottom lane what is reported from simulation or as a need. The placing is the editor's reading of the sources, and no box carries a date. Sources by thread: drilling [1, 2], leaks [12, 15], inferences [16, 17], machines and power [13, 14] and models [18–20]. Notice that each bottom box asks that a measurement or a model be trusted to act on in real time, with no person checking it first.

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 [22]. 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.

16.4.1 How This Volume Was Made

Each of the fourteen 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. At least five of the fourteen reports are self-reviews by the agent that wrote the paper, and six papers had a further referee pass on predictive control added later. 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 fourteen logs written beside the papers hold 105 entries with a severity. Of these, 28 are wrong behavior or output of an app, 48 are documentation or content, 20 are cosmetic or display and 9 are other kinds, such as a test harness or a load path. The counts understate the work. Three of the logs hold one entry or none, and eight apps keep a second log of their build in their own folders, which is not counted here. Five cases show the errors that matter for a volume about what a controller can know.

Controllers that knew what no instrument tells. Five apps logged a controller that read values a real plant would not give it. The subsea estimator and its dose conversions read the run's own parameters. At identical measurements an edited heat transfer coefficient moved the estimated minimum margin from 4.575 to 3.660 °C, and handing only the controllers an edited copy moved 49 of 100 six-hour runs. The compressor station's MPC built its model from the learner's parameters, and 16 of 17 edits moved its plan. The products line's MPC read the wall roughness of the plant, so its inferred hill moved by up to 2.63 bar with nothing measured changing. The dehydrator and the sweetener started their controllers from true plant states. No shipped run moved when these were fixed. Each was caught by the same probe, a hidden value changed at identical measurements to see whether the controller moves. In a volume where the hydrate margin and the hill are the limits, a controller that reads them from the plant's parameters looks better than any real one could.

Tests of that boundary that passed bad code. After the first truth tests were written, reviewers wrote deliberately wrong controllers, called mutants, to see whether the tests would fail. On the separator the two original tests caught 1 of 8 mutants. On the NGL train 9 of 12 mutants had passed. On the LNG train a mutant that reused two shipped analyzer statements to hand the MPC the true refrigerant blend on every scan passed the test. Each set of tests was then strengthened. A test that has never failed says nothing about the code it guards.

A limit checked where the instrument was not. The products line's margin tile reported 2.15 bar to the maximum allowable pressure while the terminal valve inlet peaked at 89.7 bar g, 0.26 bar under the limit. The check scanned the segment centers and the station discharges and never the terminal, the lowest point of the last leg. The referee found it by reading the data behind a figure. Once fixed, only the surge case moved, its least margin from 2.1 to 0.3 bar.

Measurement and truth confused in the text. The subsea draft said the loops see only the estimator, yet the flow loop and the pressure override read their transmitters directly. Its screenshot

caption said the flow was held by the pump limit, and a run showed a lower flow under a setpoint computed from a measured water cut of 18.15 % against a true 20 %. Three referee reports, on the separator, the smart producer and the flowline, found numbers inherited from the app's own guide that the engine did not reproduce. One belonged to another instant of the same run. Each was caught by measuring again instead of copying.

The book round. The chapters were written after a repair round that applied the book's earlier defect list to apps and papers, and each chapter agent first confirmed that its paper had not changed since conversion. None of the fourteen had. The fourteen chapter records list 216 rows of numeric verification, each a stated value or group of values set beside its reproduction, and thirteen report that every rerun number reproduced. The NGL record found three printed values that differ from the run, among them a steam demand of 5235 kmol/h where the run asks about 5406, and left them for the editor. The book round's defect list holds thirteen entries for this volume. The one that touches control was found on the dehydrator. A pump capacity edited in the panel never reached the controls, so at 30 L/min the plant delivered 29.8 while the controller still reported 50 and no saturation, and the trim could wind up against a pump at its limit. The agent writing Chapter 1 first chose an averaging tuning whose outflow peaked at 273 m^3/h against a 240 m^3/h limit, and caught it in its calculation script before writing.

Every error above was caught by a check that could fail: a hidden value moved at fixed measurements, a test run against code wrong on purpose, a limit checked at the plant's worst point, a number measured instead of copied. Four working rules follow. Name the instrument behind every value a controller reads. Test each test with code that should fail it. Check a limit where the plant is worst, not where the transmitter is. Let a person decide what changes.

16.5 Gaps in Oil and Gas Control

Six gaps belong to this volume's topic, and the gaps common to agent-era practice follow them in Section 16.5.1. In the editor's view none is closed. Each item gives the evidence of the chapters, the state of the literature, why the gap is hard and what progress would look like. A sentence on difficulty or progress that carries no citation is the editor's view. Figure 16.4 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. *Inferences that report their own drift.* The rod pump's card, the turboexpander's tray temperature, the NGL inferential and the LNG train's computed approach each held a loop on a number that was wrong for a time. What is missing is an inference that carries its own uncertainty and says when the data no longer support it, the request the flow-metering review makes for its field. A review of soft sensors found that most have no automated mechanism for their maintenance [16]. The Open challenge box after this list takes the gap further.

2. *A distributed minimum estimated with bounds.* The hydrate margin of Chapter 5 is a minimum over 6.5 km known through two thermometers that cannot tell a colder sea from worse insulation. A dosing law needs a lower bound on that minimum, not a best estimate. A 2021 history of flow assurance finds industry trending toward safe operation inside the hydrate domain [23]. A fiber can measure the profile, but its resolution, of the order of meters, cannot be reconciled with a long reach [12]. The gap is hard because the weakest point moves while the instruments stay at the ends. Progress would be a dose set from a stated lower bound and checked against a measured profile.

3. *Leak detection that tells a leak from a meter.* On both lines of Chapter 12 and Chapter 13 the threshold that ignored a biased meter or a batch change missed a small leak. Transient models, temperature and product properties in the balance are the known routes, and none removes the

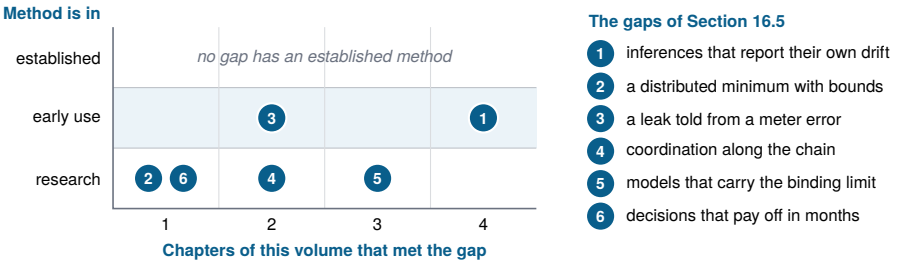

Figure 16.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: four for gap 1 (Chapters 3, 9, 10 and 14), one for gap 2 (5), two for gap 3 (12 and 13), two for gap 4 (10 and 15), three for gap 5 (9, 10 and 14) and one for gap 6 (4). Up is the maturity of the method available, which is the editor's reading of the sources cited in each item and not a measurement. Notice that only the two gaps of measurement have a method in use anywhere.

trade. A 2019 review finds most methods that work from the fluid's own measurements prone to false alarms, and gaps remaining before reliable real-time detection is achieved [15]. The gap is hard because a leak and a meter error move the same balance the same way. Progress would be a threshold set from a stated meter error that still alarms a leak smaller than that error.

4. *Coordination along the chain.* Each plant in this volume was controlled alone, with its neighbors as disturbances. The NGL header trip showed a decision of which product to sacrifice that no layer made, and the terminal's flare was decided by the send-out the market asked for. Methods that coordinate selectors and supervisors across plants owned by different operators are thin. Gas lines now feed generators that balance wind and solar power, and a 2015 study found the coupling invalidates many traditional approaches to operation [20]. The gap is hard because each plant answers to its own operator and contract. Progress would be a layer that sees which plant's limit binds along the chain and moves the cut there.

5. *Prediction models that carry the binding limit.* The predictive controllers of Chapter 9, Chapter 10 and Chapter 14 lost because their models left out quality, flooding, a surge drum or a peak. A routine way to tell, before commissioning, whether a reduced model contains the limit that will bind is not available. The 2003 survey of industrial predictive control listed a lack of methods to verify whether an identified model is adequate for control [24], and in 2023 new applications were still built mainly on linear models from step responses [18]. The gap is hard because the limit that binds in service may never move in the plant test. Progress would be a check that each limit expected to bind is in the model.

6. *Decisions whose payoff arrives in months.* The waterflood planner of Chapter 4 forecast 13.9 % high from its model's errors, and a fixed policy chosen by hand came within 0.8 % of it. Closed-loop reservoir management needs a way to value a plan against a simple policy when the truth is known only years later. A 2015 review found most of its studies academic and far from large-scale field application, and automatic closed-loop management not yet achieved there [19].

OPEN CHALLENGE **An Inference That Knows When It Is Wrong**

Gap 1 is the one most chapters met. A 2013 review warns that product quality deteriorates if a soft sensor's estimates are blindly believed and used in a control system [17]. Comparing the estimate with each analysis judges it only when an analysis exists. A check between analyses meets a harder problem: a model that must follow changes in the process and the feed sees its new data as faults. One scheme that checked its soft sensors online ran on an ethylene plant and passed only reliable estimates to a predictive controller [17]. Progress would be such a check in routine use, one that tells a drifting inference from a changing plant.

16.5.1 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 [25]. 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 [22]. 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 [26]. 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 [27]. 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 [28]. 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 controller holds the reading it is given exactly, even while the true value drifts away.
- A selector shares one actuator among limits, and its wiring decides whether the limit holds.
- An inventory decides where a disturbance goes, not whether it arrives.
- A measured load ahead of a late quality beats a faster loop on the quality.

- Prediction paid where it carried a preview or priced several handles, and lost where its model left out the limit that binds.

References

[1] J.-M. Godhavn, "Control requirements for automatic managed pressure drilling system," *SPE Drilling & Completion*, vol. 25, no. 3, pp. 336–345, 2010.

[2] T. Ma, P. Chen, and J. Zhao, "Overview on vertical and directional drilling technologies for the exploration and exploitation of deep petroleum resources," *Geomechanics and Geophysics for Geo-Energy and Geo-Resources*, vol. 2, no. 4, pp. 365–395, 2016.

[3] J. D. Macpherson, J. P. de Wardt, F. Florence, C. D. Chapman, M. Zamora, M. L. Laing, and F. P. Iversen, "Drilling-systems automation: Current state, initiatives, and potential impact," *SPE Drilling & Completion*, vol. 28, no. 4, pp. 296–308, 2013.

[4] H. P. Bieker, O. Slupphaug, and T. A. Johansen, "Real-time production optimization of oil and gas production systems: A technology survey," *SPE Production & Operations*, vol. 22, no. 4, pp. 382–391, 2007.

[5] R. Z. Ríos-Mercado and C. Borraz-Sánchez, "Optimization problems in natural gas transportation systems: A state-of-the-art review," *Applied Energy*, vol. 147, pp. 536–555, 2015.

[6] T. Bikmukhametov and J. Jäschke, "First principles and machine learning virtual flow metering: A literature review," *Journal of Petroleum Science and Engineering*, vol. 184, p. 106487, 2020.

[7] F. Di Meglio, N. Petit, V. Alstad, and G.-O. Kaasa, "Stabilization of slugging in oil production facilities with or without upstream pressure sensors," *Journal of Process Control*, vol. 22, no. 4, pp. 809–822, 2012.

[8] T. R. Wanasinghe, L. Wroblewski, B. K. Petersen, R. G. Gosine, L. A. James, O. De Silva, G. K. I. Mann, and P. J. Warrian, "Digital twin for the oil and gas industry: Overview, research trends, opportunities, and challenges," *IEEE Access*, vol. 8, pp. 104 175–104 197, 2020.

[9] E. D. Sloan, "A changing hydrate paradigm—from apprehension to avoidance to risk management," *Fluid Phase Equilibria*, vol. 228–229, pp. 67–74, 2005.

[10] H. Lu, T. Iseley, S. Behbahani, and L. Fu, "Leakage detection techniques for oil and gas pipelines: State-of-the-art," *Tunnelling and Underground Space Technology*, vol. 98, p. 103249, 2020.

[11] T. A. Fox, T. E. Barchyn, D. Risk, A. P. Ravikumar, and C. H. Hugenholtz, "A review of close-range and screening technologies for mitigating fugitive methane emissions in upstream oil and gas," *Environmental Research Letters*, vol. 14, no. 5, p. 053002, 2019.

[12] J. Li and M. Zhang, "Physics and applications of Raman distributed optical fiber sensing," *Light: Science & Applications*, vol. 11, p. 128, 2022.

[13] I. Y. Xuan, C. Skourup, J. B. Jensen, T. Haugen, and N. F. Thornhill, "Flexible operation of a mixed fluid cascade LNG plant for electrical power management," *Energy*, vol. 250, p. 123794, 2022.

[14] M. Zagorowska, M. Degner, L. Ortmann, A. Ahmed, S. Bolognani, E. A. del Rio Chanona, and M. Mercangöz, "Online feedback optimization of compressor stations with model adaptation using Gaussian process regression," *Journal of Process Control*, vol. 121, pp. 119–133, 2023.

[15] M. A. Adegboye, W.-K. Fung, and A. Karnik, "Recent advances in pipeline monitoring and oil leakage detection technologies: Principles and approaches," *Sensors*, vol. 19, no. 11, p. 2548, 2019.

[16] 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.

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

[18] S. Skogestad, "Advanced control using decomposition and simple elements," *Annual Reviews in Control*, vol. 56, p. 100903, 2023.

[19] J. Hou, K. Zhou, X.-S. Zhang, X.-D. Kang, and H. Xie, "A review of closed-loop reservoir management," *Petroleum Science*, vol. 12, no. 1, pp. 114–128, 2015.

[20] A. Zlotnik, M. Chertkov, and S. Backhaus, "Optimal control of transient flow in natural gas networks," in *2015 54th IEEE Conference on Decision and Control (CDC)*, 2015, pp. 4563–4570.

[21] 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.

[22] 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.

[23] E. D. Sloan, "Hydrocarbon hydrate flow assurance history as a guide to a conceptual model," *Molecules*, vol. 26, no. 15, p. 4476, 2021.

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

[25] 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.

[26] 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.

[27] 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.

[28] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

16.6 Exercises

16.1 Discussion · 7 min

A new well is to share its choke between a rate loop and a drawdown limit. A colleague proposes two PI controllers into a low selector, each with its own anti-windup against its own output limits. Using Figure 1.6, Chapter 13 and Chapter 11, say what to expect and what to check before the scheme is commissioned.

16.2 Discussion · 8 min

Section 1.4.9 names three ways in which measurement limits control: a late analyzer, an inferred constraint and an inventory computed from a few transmitters. For Chapter 8, Chapter 10 and Chapter 12, say which applies, what it cost in the chapter, and what second source could contradict the measurement the loop uses.

16.3 Discussion · 7 min

A ratio controller is usually tuned at the design point, where it protects the product. Using case 3 of Chapter 5 and the warm-up of case 3 in Chapter 7, explain how a ratio that is right at design can protect the wrong quantity, and how the basis of a ratio should be chosen.

16.4 Discussion · 8 min

A vendor proposes an MPC for the mode switch of the turboexpander plant of Chapter 9, and cites the economic MPC of Chapter 7 as evidence that prediction pays in gas processing. Using both chapters and Section 16.2, say what you would ask before accepting it.

16.5 Discussion · 7 min

An agent delivers a new estimator for the hill near the terminal of Chapter 13 and reports that it tracks the true hill within 0.1 bar on every shipped case. Using Section 16.4.1, say what checks you would make before the estimator goes into the back-pressure controller.

16.6 Discussion · 8 min

A field operator plans to replace the physical multiphase meters on its smart producers with a virtual flow meter computed from pressures and temperatures, and to run the water-capacity loop of Chapter 4 on it. Using Section 16.3 and the chapters on inferred limits, discuss what the operator gains, what it risks and what it should keep.

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) 160.5 m^3/h. (b) A drawdown of 40.1 bar, with 19.9 bar across the choke.

1.3 (a) $g = 0.315$. (b) 3.25 m^3/h per percent at once and 1.02 m^3/h per percent at steady state.

1.5 (a) $\tau_c = 5$ min, $\zeta = 1.0$. (b) 4.5 m^3/h per minute, 270 m^3/h per hour, over 20 min.

1.7 (a) 12.82 t per bar. (b) 4.49 bar.

1.9 (a). On $\psi = k\phi^2$ the ratio ϕ/n stays fixed, and so does the margin.

Chapter 2. Managed Pressure Drilling Optimization

2.1 (d). Overbalance costs drilling rate through chip hold-down, so the setpoint walks down to the floor of Equation (2.21).

2.3 (b). G couples weight and speed through one exponential, while the bounce index is a resonance in the excitation frequency.

2.5 (a) −5.80 psi per percent open; 78 psi at 0.011 m^3/s against 580 psi at full flow. (b) Default 0.5/s and 2.9 psi, inside the band. Delivered 0.0107/s and 136 psi, eight times the band.

2.7 (a) 5678 psi, 69.3 ft/hr; 5177 psi, 97.3 ft/hr; ratio 1.404, a 40 % gain. (b) 1.513. (c) The cuttings factor (0.968) and the wear factor (0.959) bring the prediction to 1.404.

2.9 (a) 8.0 ft, 49.1 psi, cuttings and bounce flags from the start: fail. (b) 473 s, 4.9 psi, 0.16, 263 psi, no flag: pass. (c) The cuttings flag sets at 297 s, and the attempt fails.

2.11 (a) 0.417 against 0.368: a slower bit loads the annulus less. (b) 2.39 against 2.40: at fixed weight and speed the torque is fixed, so MSE goes nearly as $1/ROP$. (c) The time history: a stringer steps MSE at the interface, wear raises it gradually.

Chapter 3. Rod Pump Card Inference and Rate Control

3.1 (d). Over a closed cycle the inertia and friction terms average to zero, so the mean keeps only the fluid load and the rod weight.

3.3 (b). At $t = 8.16$ s the layer cuts 9.0 SPM to 8.1 SPM; a search that acts every five strokes cannot answer within one stroke.

3.5 (a) 0.015 77 m^2, $\tau_h \approx 4.11$ h (the chapter's 4.12 h). (b) 50.4 m^3/d. (c) 5.04 SPM, and 5.15 SPM with the slip of 1.32×10^{-5} m^3/s, the chapter's pump-off rate.

3.7 (a) 7.18 SPM delivered for 8 SPM asked, 69 % fillage, 2.5 % card error, 0.74 m over the pump, 219 bbl/d of oil. (b) 70 %. (c) 5.1 SPM, 95 % fillage and the same oil, about 218 bbl/d, with no further pound damage.

3.9 (a) 73 kN, no pound damage, fatigue 72 %, 4.9 points; it fails because the timer is not an accepted strategy and the load passes 68 kN. (b) 5.2 points, 67.1 kN, 62 %, 0.00; 3.8 points and 0.9 kN to spare. (c) The rate envelope of 6.5 SPM.

3.11 (a) $56.3 \times 2.51 = 141.3$ kW h against $55.4 \times 1.75 = 97.0$ kW h: 44.4 kW h for 0.9 bbl, about 49 kW h per barrel against 1.75 kW h on average. (b) 392.5 units of pound damage and a peak load 1.73 klbf higher; above the pump-off rate the sustained oil rate is flat at 222 bbl/d.

Chapter 4. Smart Well Chokes and Waterflood Planning

4.1 (c). The choke can only lower the bottomhole pressure. With it wide open the rate follows the falling zone pressures.

4.3 (a). The polymer pays back after weeks. A 30-day forecast sees only 233 Sm^3 more oil, less than the 562 Sm^3 charged.

4.5 (a) 216 000 and 90 000 m^3, 204 000 Sm^3. (b) 76.9 %. (c) 9.41, and 1.06 with $R = 8.92$. (d) About 27 300 m^3.

4.7 (a) 52 kJ, 0.45 g/s, 1206 m^3/d, 0.29 g/s. (b) 416 kJ, 3.75 g/s, 1409 m^3/d, 3.26 g/s. About 17 % more liquid for about eleven times the sand and eight times the erosion.

4.9 (a) 47 731 Sm^3, +48.1 % on day 30 and +7.1 % at the end. (b) 89 270 Sm^3, +94.2 and +100.4 %. (c) 103 912 Sm^3, +50.6 and +133.3 %. Only (c) passes, and the ranking on day 30 is the reverse of the ranking on day 180.

4.11 (a) $757/943 = 0.80$, and $33.0/30 = 1.1$ bar a day. (b) The plugging keeps growing, so the extra headroom is spent within days, and the margin would fall to about 10 bar, where a fracture could carry water straight to the producer. Acid stimulation removes the skin.

Chapter 5. Deepwater Flowline Hydrate Control

5.1 (b). The equilibrium follows the pressure, which falls only up the riser, so the margin is smallest where the cooled fluid is still under the full seabed pressure.

5.3 (d). The estimator only knows the design sea, and a colder sea and 12 % more heat loss look the same at TT-103 and TT-104.

5.5 (a) 1.95 m^3/h. (b) 13.7 wt %. (c) From 12.5 °C to 6.4 °C, a loss of about 6.1 °C of depression that no control setting can restore.

5.7 (a) −4.43 °C, 131.65 °C·h, 47.4 m^3. (b) +1.03 °C, 0 °C·h, 28.3 m^3. (c) At 8.57 h (the chapter's 8.55 h on the step), at 5.06 km, the riser base.

5.9 (a) 155.21 °C·h, −4.82 °C, 548 Sm^3, 103.1 m^3, 106.9 bar: fails. (b) 0 °C·h, +1.38 °C, 706 Sm^3, 81.9 m^3, 106.1 bar: passes. (c) 529 Sm^3 and 119.6 m^3: fails. The cap is 778 Sm^3/d at 20 % water and about 389 Sm^3/d at 40 %.

5.11 (a) Each sees only the integrated heat loss upstream of it. Distributed temperature sensing along the line, or an insulation survey, would locate it. (b) Treat the estimate as an upper bound when the heat-loss factors have moved, and inhibit or depressurize early rather than wait for the predicted time.

Chapter 6. Offshore Separator Slug Absorption

6.1 (a). The arrival rate is several times what the open valves can pass, so most of the slug must be stored whatever the controller does.

6.3 (c). The analyzer reports on water that has already gone overboard, so the diversion comes after the loss.

6.5 (a) 4.1 mm/s and 0.051 mm/s. (b) $Re = 0.28$, a factor 1.062, 3.86 mm/s. (c) 1.86 mm/s. Only the 540 µm class outruns it. The 180 µm class at 0.45 mm/s and the 60 µm class settle only in part, unless they coalesce into the next class first.

6.7 (a) 8.4 m^3, 335 L/MSm^3, 317 m^3/h (±1). (b) 6.9 m^3, 1772 L/MSm^3, 219 m^3/h (±1). (c) 0.0 m^3, worst 25 mg/L, 2078 L/MSm^3, water peak about 82 m^3/h. Holding the water back keeps the single hydrocyclone bank near its capacity. The cost is a higher level and foam, so the carryover rises further.

6.9 (a) 87.3 m^3, 84.4 % (about 374 s off spec), 318 m^3/h: three clauses fail, the in-spec oil, the BS&W share and the second stage. (b) 98.5 m^3, 89.0 % (about 264 s), 207.6 m^3/h (207.5 in the CSV): it passes. (c) The supervisor fixes the water and most of the BS&W, averaging protects the second stage, and the feedforward only narrows the pressure band.

6.11 The second bank and the averaged interface loop each remove the loss, while the faster analyzer still lets 1.39 m^3 through. The bank needs equipment and power, the retuned loop moves the interface further and, with a tight level loop, overloads the second stage longer.

Chapter 7. Glycol Dehydration Floor and Ceiling

7.1 (a). More glycol only moves both stages toward the lean composition, and the outlet tends to $W^* + (1 - E)^2 (W_{in} - W^*)$.

7.3 (c). At the fuel limit every extra liter cools the reboiler, the lean glycol gets wetter and the floor rises, so the trim's answer makes the water worse.

7.5 (a) 13.41 and 5.49 lb/MMscf, 88.4 % absorbed. (b) 3.60 lb/MMscf. More glycol, 5 K warmer than the gas, warms the contactor and raises W^*. (c) 7.72 lb/MMscf, so no.

7.7 (a) 8.24, 6.83 and 5.79 lb/MMscf, so 0.29, 0.14 and 0.05 lb/MMscf per L/min. (b) 8.81 lb/MMscf, lean glycol 99.107 wt%, reboiler 202.9 °C. (c) 6.43 lb/MMscf.

7.9 (a) 10.37 lb/MMscf, 506 samples over, 165 kW: fails on water. (b) 10.47 lb/MMscf, 195 over, 187 kW. The circulation rises from 39.6 to 49.9 L/min while the reboiler falls from 202.8 to 161.0 °C and the lean glycol from 99.11 to 97.28 wt%. (c) Passes at 6.73 lb/MMscf and 232 kW. After the fuel limit it opens up to about 27.1 Sm^3/h of stripping gas while the reboiler dips to 190.0 °C.

7.11 (a) −1.0 °C, below the line's 4 °C, so no free water condenses. (b) The dew point below 0 °C is that of supercooled water and the hydrate point of the same gas lies higher, so the hydrate curve must be checked. (c) TEG decomposes near 207 °C. Stripping gas buys the rest.

Chapter 8. Selective Amine Sweetening

8.1 (d). With the duty fixed, the circulation only trades capacity against stripping, and at 60 % steam no point on that trade holds 4 ppm.

8.3 (b). The clause requires the true H_2S of every analyzer sample to be at most 4 ppm, so a peak that falls between samples is not graded.

8.5 (a) 8.80, 28.2 and 88.4 kPa. (b) 9.78, 31.34 and 98.17 kPa, errors −2.2, +4.5 and −1.8 %. (c) A factor 1.27. The lean amine holds back more H_2S at the top, and the true gas passes 4 ppm at 62.3 min with nothing the ratio measures changed.

8.7 (a) A corrosion stop at 104.1 min, the cap first passed at 44.2 min. (b) 4.00 ppm and 0.492 (±0.01 ppm, ±0.002). The analyzer peaks at 4.23 ppm (±0.05). (c) The stop follows 60 min continuously over the cap. At 57 m^3/h the gas sits on its 4 ppm limit and the loading just under the cap.

8.9 (a) A corrosion stop at 214.5 min: 5.27 ppm worst sample with three samples over, peak loading 0.565, 4.86 MW. (b) Target met: 3.78 ppm, 0.458, 5.25 MW. (c) 3.63 ppm, 0.458, 5.54 MW, which fails the 5.35 MW clause.

8.11 (a) 29.3 kmol/h, 0.45 MW, about half of the 0.97 MW rise. (b) Steam, a weaker acid gas for the Claus unit and lost H_2S capacity. Short contact (fewer trays, low weirs) and a lower lean feed point, within the H_2S limit.

Chapter 9. Sequenced Turboexpander Plant

9.1 (c). At a higher pressure the same stage-11 temperature corresponds to a liquid richer in methane, and the tray setpoint moves only in the second step.

9.3 (a). The relation between the tray temperature and the bottoms composition changes with the feed, so a held temperature is no longer a held composition.

9.5 (a) 19.88 cents/gal, +5.12 and −9.88 cents per gallon. (b) 221.0 kmol/h, 4931 gal/h, about 487 dollars per hour. (c) About 1.01 cents/gal, an ethane price near 18.87 cents/gal.

9.7 (a) 16.0 min, 29.52 bar and 30.17 bar (±0.05 bar), 36.2 min and 47.8 min (±0.2 min). (b) The booster on the expander's shaft did part of the recompression. The residue compressor is at its driver limit, so more capacity gives no more flow. (c) 167 min to 170 min for a restart between 60 and 65 min.

9.9 (a) Shipped 3.65, 95.2, 0, 97.7. Declared only 3.20, 94.4, 0, 98.2. Margin only 3.14, 95.2, 0, 97.6. Both 1.51, 94.4, 0, 98.2, the only pass. (b) Shipped: both switches. Declared only: the return. Margin only: both, its worst on the way out. (c) The declared order protects the way out, and the warm margin with the declared order protects the return with the richer feed.

9.11 The gain is 12 dollars per hour, about 0.4 % of the margin, for 0.29 MW more power and only 0.17 MW of headroom to the 4.2 MW driver limit.

Chapter 10. NGL Fractionation Train Plantwide Pairing

10.1 (b). The ethane flow in the distillate is the same at every cut, and reaching 2 mol% by dilution would need a distillate about 2.6 times the column feed.

10.3 (d). It settles at 0.705 mol%, inside every contract, so the error shows only as giveaway.

10.5 (a) $\eta = 0.404$, 0.457, 0.506 and 0.550, $S \approx 1790$, 4840, 11900 and 27200, factors 2.70, 2.47 and 2.28. (b) $r_L = 0.980$, $r_H = 0.0100$, $S \approx 4850$, and $r_L = 0.992$, $r_H = 0.0257$, $S \approx 4860$: the cut moves along one curve. (c) $S \approx 27400$, the value of (a) within the rounding of the flows.

10.7 (a) 16.72 and 0.08 %, 290.2 and 918.4 kmol/h, 48.5 kmol/h. (b) 14.16 and 0.27 %, about 342.6 and 918.4 kmol/h, still 48.5 kmol/h. (c) About 2425 kmol/h, 2.6 times the column feed. The deethanizer cut sets it.

10.9 (a) 0.912, then 1.639, 1.730 and 1.275. (b) 1.49, 2.29 and 1.90 %, recovery 97.1 %, gasoline 3.81 %. (c) 64, 194 and 795 min, the gasoline crossing at 105 min. The natural gasoline, the one product with no loop.

10.11 (a) The product settles at 0.705 mol%, inside every contract, so the cost is giveaway. (b) A periodic comparison against a reference measurement and a supervisory layer above the loop. The trim moved the inferential setpoint from 0.50 to 1.24 mol%, and a trend of it shows the drift. (c) Half in 47.8 min and 90 % in 97.9 min. The trim spreads it over most of a shift.

Chapter 11. Compressor Station Load Sharing

11.1 (b). Both machines must make the head the headers set. With unit 2 slower the line unpacks, the discharge falls to 69.35 bar a, and unit 1 meets the lower head at 153.5 kg/s.

11.3 (d). The trip signal opens the valve in the trip's own scan, before the measured margin has fallen, and unit 2 coasts down with no surge.

11.5 (a) The margin depends on ϕ/n alone, and at one head $(1-a)n^2\Psi(\phi/n)$ is the same for both, so equal ϕ/n needs equal $(1-a)n^2$. It neglects the small difference in head between the machines' discharge lines. (b) $90.47/\sqrt{0.92} = 94.32$ %, against the 94.3 % found by hand in the app. (c) $81.25/\sqrt{0.94} = 83.80$ %, against the app's 83.8 %.

11.7 (a) 15.0 t recycled and 63.33 MWh. (b) No gas recycled, 62.76 MWh, unit 1 at 83.8 % and unit 2 at 81.3 %, both about 12.5 % from surge. (c) The fouled machine needs $1/\sqrt{0.94}$ times the speed for the same margin, and the fuel falls 0.57 MWh in the hour because the recycle costs only the work spent compressing it.

11.9 (a) 0.334 bar peak to peak with a period of about 8.4 s, the valve swinging 12.2 %, and a lowest margin of 5.9 %. (b) A swing of 4.4×10^{-5} bar and a lowest margin of 9.9 %. (c) With no unit on recycle only the speeds move the forward flow, so the station loop sees its design gain and settles.

11.11 (a) $\psi = 0.3759$, $n_s = 0.8193$, $\phi_s = 0.4096$ and a same-head margin of 41.4 %, against 28.0 % at the same speed. (b) 11.2 % against 10 %. (c) The two conventions agree within about a point near surge but differ by thirteen points at design, so a margin must be quoted with its convention. Real control lines come from the station's surge tests and protection studies.

Chapter 12. Gas Pipeline Line Pack and Leak Detection

12.1 (a). A fixed inlet and a higher flow mean a larger friction drop, and only a loop on the delivery pressure would move the inlet.

12.3 (c). In a pressure ramp the estimate $\hat{M}$ lags the true pack, so inflow less outflow less $\Delta\hat{M}$ looks like a leak.

12.5 (a) P_{in}/P_{out}: 1.09, 1.26 and 1.43. (b) 33, 53 and 66 min, within a few percent of the measured times. (c) The loop gain falls by about 13 % at night, and with the slow integral of PIC-D the lower gain outweighs the faster plant.

12.7 (a) 12 to 13 min after the step, then +0.30 and +0.47 bar (±0.01 bar). (b) 2825.1 t to 2843.4 t, +18.3 t, a mean rise of 0.35 bar. (c) Nearly an integrator. Only the compressor map, which delivers less as the ratio rises, slowly brings the flows back into balance.

12.9 (a) MPC 10.4 min, 44.73 bar, \$40,623. Cascade 54.6 min, 44.27 bar, \$40,213. Cascade at 46.4 bar 21.9 min, 44.44 bar, \$40,593. (b) 70.00 bar against 68.09 bar, and about 30 and 43 t more gas at 17:00 and 18:00. (c) Preview off: 301.6 min under, 42.79 bar, \$41,690. The forecast and the floor in its problem make the difference, and a fixed margin on the cascade buys most of it.

12.11 It removes the false alarm from a 1.5 % meter error, but a 3 % leak is found 11 minutes later and releases about 4 t more, and a 1.5 % leak is not found at all. Better meters or redundant metering, and evidence that does not depend on the meters.

Chapter 13. Refined Products Pipeline Control

13.1 (d). One flow passes all four stations, so only the lead station's flow loop sets it, and the others give way through their suction loops.

13.3 (b). A meter shift of about 2.5 m^3/h and an unmeasured change of line pack of about 0.6 m^3/h put the residual on the threshold.

13.5 (a) 1.96 m/s, 18.6 bar, against the 18.5 bar measured in the first ten seconds. (b) 79.6 s and 81.0 s, and the front arrives later because friction wears it down. (c) 4.4 min against 39.6 h.

13.7 (a) 14.58 min. Stations 2, 3, 4 trip on low-low suction and station 1 last on high-high discharge. (b) Worst suction 6.9 bar g, station 2 lowest near 7.0 bar g, 899 m^3/h.

13.9 (a) 25.1 min, 15.3 m^3. (b) Slack at 10 bar g, none at 15 bar g. (c) 15.8 bar g with a hill estimate of 2.1 bar g, never slack. The volume is a time-step artifact of the outflow clamp.

13.11 (a) 0.0295 MW/bar, about 5.0 MWh a day for 7 bar. (b) 5.46 and 5.29 kWh/m^3: the days carry different flows and volumes, so their energy totals are not comparable.

Chapter 14. LNG Liquefaction Pairing and Inferred Limits

14.1 (c). The suction pressure rises from 2.901 to 3.167 bar, and by Equation (14.9) the refrigerant boils at −159.46 °C instead of −161.99 °C.

14.3 (a). Its gain changes by a factor of 16 across the working range, while the suction pressure moves the product by a nearly constant 9.2 °C per bar.

14.5 (a) 52.0 kW/°C each, $C^* = 0.9998$, balanced. (b) NTU = 13.0, $\varepsilon = 0.929$, $Q = 3380$ kW, $T_{r,1} = 25.0$ °C, warm-end approach 5.0 °C. (c) 4.44 and 4.93 min.

14.7 (a) 37.63 kg/s, 3.167 bar, −159.46 °C, −159.46 °C, 0.00 °C. (b) 2.735 bar, −156.60 °C, about 10 °C per bar. (c) 17.76 against 5.18 °C: the valve spoiled the profile and the speed did not.

14.9 (a) MPC 18.6, 2.78, 20.4, and cascade 85.1, 1.31, 28.9. (b) 41.5 against 61.5 min, 455.1 against 450.5 t. (c) The MPC wins on time and on integrated excess and loses on depth. All three belong in front of a manager, the peak first.

14.11 (a) The computed approach reads 2.4 °C too wide, so the true cold end sits near 4.6 °C while the display shows 7. (b) For example, an alarm on an analysis that has not changed over several cycles, or on the sample flow.

Chapter 15. LNG Regasification Terminal Boil-Off

15.1 (c). At 30 t/h of send-out a full recondenser at 9 bar g takes 4.53 t/h, and the override holds the compressors there.

15.3 (a). At 6 h tank B boils 2.40 t/h against 3.11 as one layer.

15.5 (a) 97.9 kg/mbar. (b) 23.7 mbar/h. (c) About 4.69 h, as the run gives.

15.7 (a) 91.5 min and 281.5 min (4.69 h). (b) 6.94, 4.51, 4.53 and 4.21 t/h. (c) The PIC asks for 20.00 t/h, the compressors move 4.49 t/h, and the flare opens at 281.5 min as in Manual.

15.9 (a) Loops: 7.15 t, 10.19 MWh, 233.0 mbar g. Preview: nothing flared, 11.43 MWh, 197.4 mbar g. MPC: nothing flared, 11.30 MWh, 200.1 mbar g. (b) 233.5 min and 98.1 mbar g. The preview compresses the night's boil-off ahead of time, while the recondenser can still take it.

15.11 (a) The time the densities meet, set by the chosen initial gap, and the size of the surge, set by the declared overturn rate. (b) The sequence of a rollover and the remedy of sending out tank B first, not a prediction of when a real tank rolls over.

Chapter 16. Summary and Outlook

16.1 A good answer predicts that the waiting override winds up while the rate loop holds the choke, so it takes over late and the limit is broken. It cites Chapter 1, where the same loops and tunings broke the 60 bar limit by 1.2 bar for 21.8 min when wired independently and held it when both reset terms followed the applied opening. It asks for the external-reset form used on every pump of the products line, and a test that hands the actuator from one loop to the other. It notes that a protection that must act at once, like the anti-surge valve on a trip, should not wait for a PI to integrate.

16.3 A good answer says that a ratio keeps the proportion between two flows and protects the limit only while that proportion is what the limit depends on. On the flowline a ratio on the liquid rate thinned the methanol in the water to 9.0 wt% when the water cut rose, and the margin reached −1.73 °C, while a ratio on the water held 1.54 °C. On the dehydrator the ratio keeps the share of water removed, so a warmer gas with more water left more water in the product. The basis is the flow that moves the limit in the limit's own balance.

16.5 A good answer first checks what the estimator reads. It moves a hidden plant value, such as the wall roughness, with the measurements held, and the estimate must not follow, since the earlier MPC's inferred hill moved by up to 2.63 bar that way while every shipped run stayed the same. It runs the truth tests against deliberately wrong code to see that they can fail. It tests a fast transient, where a steady inference lags, as at a station trip. It checks the limit at the hill's worst point and time, not only at the shipped cases. A person approves the change.

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.

absorption Removing a gas from a gas stream by dissolving it in a liquid solvent that flows the other way through a column. *Chapter 1*.

acid gas Hydrogen sulfide and carbon dioxide, the gases that a sweetener removes from natural gas before it is sold. *Chapter 8*.

analyzer delay The age of an analyzer reading: the time in the sample line plus the time the analysis takes. *Chapter 6*.

anti-surge control A loop that opens a recycle valve to keep a compressor's flow on the safe side of its surge line. *Chapter 1*.

artificial lift Lifting oil to the surface with a pump when the pressure of the reservoir alone cannot bring it up. *Chapter 3*.

averaging level control Level control tuned to let the level swing within a band so that the outflow changes as smoothly as possible. *Chapter 1*.

back-calculation A way to stop integral windup: when the output is limited, the integral is pulled toward the output actually applied. *Chapter 1*.

boil-off gas Vapor formed as stored liquefied natural gas takes up heat through its tank's insulation and boils. *Chapter 15*.

bottomhole pressure The pressure at the bottom of a well, where the reservoir and the equipment at the surface meet. *Chapter 1*.

bumpless transfer Switching a loop between manual and automatic without a jump in the controller's output. *Chapter 1*.

cascade control Two loops in series in which the outer controller sets the setpoint of the inner one. *Chapter 8*.

choke A valve whose opening restricts the flow from a well, or the mud returning from a well being drilled, and so sets a pressure. *Chapter 1*.

chromatograph An analyzer that separates the components of a sample in time and reports one composition at the end of each cycle. *Chapter 1*.

compressor map A chart of a compressor's head against its flow at each speed, with the surge line drawn through the peaks. *Chapter 1*.

compressor station A site of several compressors that raises the pressure of the gas in a transmission pipeline. *Chapter 11*.

constraint control Control that holds a plant at the limit that binds, with selectors, overrides or a supervisor. *Chapter 1*.

dehydration Removing water vapor from natural gas so that it cannot form hydrates or liquid water in a pipeline. *Chapter 7*.

distributed control system (DCS) The control computers of a plant, which run its loops in controllers spread through the plant and show them on operator screens. *Chapter 1*.

drag reduction Lowering the friction of a liquid in a pipe by adding a small dose of a long-chain polymer. *Chapter 13*.

drawdown The difference between the reservoir pressure and the flowing pressure at the bottom of a well, which drives oil into the well. *Chapter 1*.

dynamometer card A plot of the load on a pumping rod against its position over one stroke, read to see what the pump is doing. *Chapter 3*.

equilibrium floor The cleanest gas an absorber can deliver, set by the lean solvent it is fed, however much solvent circulates. *Chapter 1*.

ethane rejection Running a gas plant so that ethane stays in the sales gas, chosen when ethane is worth more as gas than as liquid. *Chapter 9*.

external reset A form of controller whose integral follows the output actually applied, so a loop waiting at a selector is ready to take over. *Chapter 1*.

fast and slow handles Two inputs that move one output, one quickly but only down to a limit and one slowly by moving that limit. *Chapter 1*.

feedforward Acting on a measured disturbance before it has changed the controlled variable. *Chapter 5*.

final control element The device that changes a flow of material or energy at a controller's command, such as a valve, a pump or a heater. *Chapter 1*.

flow assurance Keeping oil and gas moving through a line by preventing hydrates and other blockages. *Chapter 5*.

flowline The pipe that carries a well's mixed oil, gas and water to a separator or a platform. *Chapter 1*.

fractionation train A series of distillation columns that splits recovered gas liquids into products such as propane and butane. *Chapter 10*.

fracture pressure The pressure at which the rock around a well breaks open and drilling fluid is lost into it. *Chapter 2*.

gas hydrate An ice-like solid of water and light gas that forms above the freezing point at high pressure and can plug a line. *Chapter 5*.

giveaway Product purer than its specification requires, paid for by valuable product lost to another stream. *Chapter 10*.

glycol contactor The column in which triethylene glycol absorbs water from natural gas. *Chapter 7*.

hydrate inhibitor A chemical, such as methanol or glycol, dosed into a stream so that hydrates need a lower temperature to form. *Chapter 5*.

inferential control Control of a quantity computed from other measurements because it cannot be measured directly or in time. *Chapter 1*.

inflow control valve A downhole valve that sets how much a zone of a well produces or takes. *Chapter 4*.

integrating process A process whose output keeps changing for as long as its inflow and outflow differ, as a tank level does. *Chapter 1*.

Joule–Thomson valve A valve through which a gas expands without doing work, cooling far less than in an expander. *Chapter 1*.

Kalman filter An estimator that blends a model's prediction with noisy measurements in proportion to how much each is trusted. *Chapter 6*.

kick An unwanted flow of fluid from the rock into a well being drilled, when the well's pressure falls below the rock's. *Chapter 2*.

Kremser relation A formula for the share of a removable gas that an absorber with a given number of stages takes out. *Chapter 1*.

leak detection Finding a leak in a pipeline by comparing the metered flows in and out with the change in the line's inventory. *Chapter 1*.

line pack The gas stored in a pipeline by its pressure, which lets the line ride through a change in demand. *Chapter 1*.

liquefied natural gas (LNG) Natural gas cooled until it is a liquid, so that a ship can carry it across the sea. *Chapter 14*.

load sharing Dividing the total flow of a station among parallel compressors by a chosen rule. *Chapter 11*.

managed pressure drilling (MPD) Drilling with the well sealed at the top and a choke on the returning mud, so the pressure downhole can be changed quickly. *Chapter 2*.

maximum allowable operating pressure (MAOP) The highest pressure at which a pipeline may be operated under the safety rules that govern it. *Chapter 12*.

methyldiethanolamine (MDEA) An amine solvent that takes up hydrogen sulfide faster than carbon dioxide, so its circulation sets which gas it removes. *Chapter 8*.

minimum internal temperature approach The smallest temperature difference between warm and cold streams inside an exchanger, which no instrument reads. *Chapter 14*.

mixed refrigerant A refrigerant blended from several components so that it boils over a range of temperatures and follows the cooling curve of the gas. *Chapter 14*.

model predictive control (MPC) Control that predicts the process with a model and chooses future moves that best meet its targets within limits. *Chapter 1*.

multiphase flow Oil, gas and water flowing together in one pipe, each at its own speed. *Chapter 1*.

multiphase meter A meter that gives the flows of oil, water and gas in a well stream without separating them. *Chapter 1*.

natural gas liquids (NGL) Ethane, propane, butanes and heavier hydrocarbons recovered from natural gas as liquids. *Chapter 1*.

no-touch time The time a stopped line can wait before its first wet section reaches the conditions at which hydrates form. *Chapter 5*.

override control A loop that takes a shared actuator away from the normal loop when a limit is approached. *Chapter 1*.

PID control The proportional, integral and derivative controller that holds most loops in this volume. *Chapter 1*.

polymer flooding Adding polymer to injected water to thicken it, so that it sweeps oil more evenly through the rock. *Chapter 4*.

pore pressure The pressure of the fluid in the rock, below which a drilling well must not let its own pressure fall. *Chapter 2*.

process gain How much a process output changes at steady state for a unit change of an input. *Chapter 1*.

produced water Water that comes up a well with the oil and gas and must be separated and cleaned. *Chapter 4*.

productivity index How much liquid a well delivers per hour for each bar of drawdown. *Chapter 1*.

programmable logic controller (PLC) A rugged computer that runs local control and logic at a plant or a field site. *Chapter 1*.

pump fillage The share of a downhole pump's barrel that fills with liquid on each stroke. *Chapter 3*.

pump-off The state of a rod-pumped well whose

pump can lift more than flows into the well, so the barrel only partly fills. *Chapter 3*.

rate of penetration How fast a drill bit deepens the hole. *Chapter 2*.

ratio control Holding one flow in a set proportion to another, such as an inhibitor to the water it protects. *Chapter 5*.

receding-horizon planning Planning a series of future moves, applying only the first, and planning again later with new data. *Chapter 4*.

recondenser A vessel at an LNG terminal in which cold liquid condenses compressed boil-off gas back into the send-out. *Chapter 15*.

recycle valve A valve that returns compressed gas to a compressor's inlet to keep the flow through the machine above surge. *Chapter 1*.

relative gain array A table of how strongly each input of a plant acts on each output, used to choose which input controls which output. *Chapter 6*.

remote terminal unit (RTU) A field device that reads sensors and moves actuators at a distant site and reports to a control center. *Chapter 1*.

rich loading The amount of acid gas each mole of amine carries when it leaves the absorber. *Chapter 8*.

riser The rising pipe that carries production from the seabed or a low point up to a platform. *Chapter 1*.

rollover The sudden mixing of two layers of stored LNG of different density, which releases a burst of vapor. *Chapter 15*.

selector A function that passes on the lowest or the highest of several controller outputs. *Chapter 1*.

sequence control Logic that steps a plant through a fixed series of actions, such as a start, a trip response or a change of mode. *Chapter 5*.

severe slugging A cycle in which liquid seals the foot of a riser, gas builds up behind it and then blows the liquid out at once. *Chapter 1*.

smart well A well with downhole valves and gauges that let each zone be produced or injected on its own. *Chapter 4*.

soft sensor A model that computes a quantity from measured ones and stands in for an instrument. *Chapter 3*.

stick-slip A vibration in which the drill bit stops and then spins fast while the pipe at the surface turns steadily. *Chapter 2*.

supervisory control A layer that moves the setpoints of the loops below it toward an economic or a safe target. *Chapter 1*.

supervisory control and data acquisition (SCADA) The system that gathers data from distant field sites and sends commands to them from a control center. *Chapter 1*.

surge Unstable flow in a compressor whose flow has fallen too low, with gas reversing through the machine. *Chapter 1*.

surge drum A vessel whose level is allowed to swing so that it, and not the units behind it, takes a change in flow. *Chapter 1*.

surge margin How far a compressor's flow lies above the surge flow at its present speed, in percent. *Chapter 1*.

sweetening Removing hydrogen sulfide and carbon dioxide from natural gas, here by absorbing them in an amine solution. *Chapter 8*.

thief zone A thin, permeable layer of rock that takes most of the injected water and leaves the oil elsewhere behind. *Chapter 4*.

transmix The mixed fuel at the boundary between two batches in a products pipeline, cut out into a tank of its own. *Chapter 13*.

transport delay The time a change takes to travel with the fluid from where it is made to where it is seen. *Chapter 1*.

triethylene glycol (TEG) The glycol that absorbs water from natural gas in the contactor of a dehydrator. *Chapter 7*.

turboexpander A turbine in which gas expands and cools as it does work, here driving a compressor on the same shaft. *Chapter 1*.

variable-speed drive An electronic drive that sets the speed of a motor by changing the frequency of the power that feeds it. *Chapter 1*.

water dew point The temperature at which water begins to condense from a gas at a given pressure. *Chapter 7*.

water hammer The pressure rise when a moving liquid column is stopped suddenly, which travels along the line as a wave. *Chapter 1*.

waterflooding Injecting water into a reservoir to push oil toward the producing wells and hold up the pressure. *Chapter 4*.

weight on bit The downward force on the drill bit, the share of the drill string's weight that the hoist lets rest on it. *Chapter 2*.

Index of Control Topics

absorption, 10, 169
  selective, 190
    history, 191
acid gas
  monitoring, 204
active constraint, 25, 67
AI agents
  engineering practice, 383
  open problems, 388
  reproduction, 386
  truth audit, 385
  verification record, 385
  while studying, 18
analyzer delay, 135, 172, 189, 196, 217, 380
anti-surge control, 253
averaging level control, 7, 135, 236, 380
  history, 138
  optimal, 8

back-calculation, 37, 67, 197, 217, 261, 281, 306
bumpless transfer, 17, 37, 117, 175, 197, 217, 261, 281, 306

capacity variable, 333
cascade control, 197, 217, 236, 276, 321, 333
  commissioning, 247
closed-loop reservoir management, 83
  history, 83
  open problems, 100
compressor map, 9, 256
compressor staging, 363
compressor surge, 9, 259
conditional integration, 118
constrained optimization, 24
constraint control, 11, 189, 321, 335
  design procedure, 381

data export, 18
distributed constraint, 106
distributed parameter system, 109, 277
  lumped segments, 302
distributed temperature sensing, 383
  open problems, 130
disturbance estimation, 147
disturbance model, 336
  integrating, 283
disturbance observer, 40
drilling automation
  history, 3

economic model predictive control, 119, 164, 180, 197
effectiveness-NTU method, 325
electric drive, 383
equilibrium floor, 170
external reset, 306

fast and slow handles, 174, 235
fault diagnosis, 65
feasible region, 35
feedforward, 117, 145, 176, 196, 276
  demand forecast, 281
  from a published schedule, 364
feedforward control
  ratio, 380
final control element, 12
flexible operation, 347
flow scheme, 4

giveaway, 240
glycol circulation
  rule of thumb, 182
grid search, 67

heat exchanger
  thermal stress, 224
hill climbing, 38
history matching, 92
hydrate risk management
  open problems, 130
hysteresis, 31, 363

inferential control, 12, 55, 196, 217, 236, 306, 321, 322
  drift, 380
  stale input, 339
integrating process, 7, 31, 62, 353, 358
interaction, 143, 266
interface level measurement
  open problems, 158
inventory control, 353

Kalman filter, 17, 147, 283
  steady-state, 336
Kremser relation, 171

leak detection, 12, 276
  mass balance, 281, 301
limit cycle, 31, 266
line pack, 8, 259, 276, 303, 380
  packing ahead of a peak, 281
load sharing, 253
  equal flow, 262
  equal speed, 262
  equal surge margin, 261
loading constraint, 190
loop pairing, 230, 321, 323, 332

max-min objective, 333
measurement bias, 152
method of lines, 279
mode switching, 210, 217
model mismatch, 93
model predictive control, 40, 146, 218, 237, 262, 282, 306, 335, 364
  against simpler control, 381
  disturbance model, 238
  in the apps, 17
  move blocking, 40, 176, 197, 238
  soft constraints, 40, 119, 147, 176, 197, 218, 238, 263, 284
model reduction, 125
moisture analyzer
  open problems, 183
moving horizon estimation, 17, 347
multiphase flow, 6
multirate simulation, 80, 87

observer, 364
oil and gas control
  open problems, 386
oil-in-water analyzer
  open problems, 158
online feedback optimization, 383
override control, 11, 38, 80, 164, 197, 262, 298, 363, 378
  external reset, 91, 117, 175
  low suction wins, 306

PID control
  anti-windup, 16, 37, 67, 145, 333
  velocity form, 16, 37, 67, 117, 175, 197, 217, 236, 281, 306, 333
PID tuning
  SIMC rule, 45
pinch analysis, 324
plantwide control, 230
  design procedure, 233
  history, 232
prediction horizon, 97
process gain
  flow-dependent, 285
  negative, 30
  of a well, 5, 89
produced water
  discharge limits, 157
productivity index, 5
protection layer, 68
protective override, 340
pump-off control, 55

ratio control, 117, 164, 196, 235
  ratio station, 246
receding-horizon planning, 80, 92
relative gain array, 144, 231, 323
resonance, 32

sampled measurement, 329
scenarios, 18
  completion rules, 18
scorecard, 18
selector, 11, 91, 117, 175, 197, 306, 335, 363, 378
self-optimizing control, 322
  history, 324
self-regulating process, 328
separation factor, 234
sequence control, 118, 210, 217, 262
shared constraints, 230
shortcut column model, 234
simulation engine, 16, 34
simulator framework, 15
slugging, 6
soft sensor, 65, 232, 329, 382
  maintenance, 248
  online validation, 387
  tray temperature, 245
stability boundary, 31
state estimation, 116
steady-state optimization, 34, 67, 218
steady-state target calculation, 336
step test, 143
stuck valve, 152
supervisory control, 24, 146, 218
  history, 3

surge margin, 257, 328
surge vessel, 138

transport delay, 6, 106, 239
    variable, 112
trip, 142

turboexpander, 10

valve authority, 152
velocity error constant, 37

wave equation, 59

Index of Applications

absorption equilibrium, 169
acoustic fluid-level survey, 58
amine sweetening, 189
    history, 191
annular fluid level, 62
anti-surge recycle, 340
antifoam, 194
API RP 11L, 60
artificial lift, 55
autodriller, 26
    history, 26

back-pressure choke, 29
    hardware, 47
batch interface, 300
    history, 300
batch tracking, 307
beam pumping unit, 61
bit bounce, 32
bit wear, 32
boil-off compressor, 360
boil-off gas, 355
bottomhole pressure, 29
bubble point, 170
bundle metal inertia, 327
burner management system, 182

carbon dioxide
    freezing, 225
cargo flash, 359
centrifugal compressor, 256
centrifugal pump, 303
check valve, 258, 303
chip hold-down, 29
chromatograph, 217
    liquid sample, 222
city gate, 278
Claus unit, 192
coil-wound heat exchanger, 325
cold box, 211, 325
cold restart, 123
column separation, 300
compressor map, 280
compressor station, 253, 276
    federal safety rules, 270
    history, 255
computational pipeline monitoring, 281, 301
contract pressure floor, 278
control room management, 292
control valve
    positioner, 156
corrosion, 194
cuttings transport, 30
    ceiling on the drilling rate, 35

debutanizer, 233
deepwater tieback, 107
deethanizer, 233
demethanizer, 210
densitometer, 304
depropanizer, 233
deviated well
    rod pumping, 74
displacer level transmitter, 155
drag reduction, 300
    history, 301
drilling automation, 24, 382
    history, 381
drillstring
    torsional model, 31
droplet settling, 137
dynamometer card, 57
    history, 58
    in practice, 73

emergency shutdown, 265
    history, 255
emulsion, 137
equivalent circulating density, 30
ethane rejection, 211

fault diagnosis
    open problems, 74
feed gas valve, 335
flare, 361
flooding, 216
flow assurance, 106
flow element
    differential pressure, 268
foam, 141, 171
foaming, 194
    open problems, 204

fouling, 257
founder point, 27
frac spread, 216
fractional flow, 82
fractionation train, 230
fracture pressure, 86
friction factor, 303

gas chromatograph, 322, 335
gas hydrate, 166
  history, 167
gas hydrates, 107
  history, 108
gas processing, 10
gas transmission, 253
gas turbine, 256
  ambient derate, 324
  mechanical drive, 345
gas-electric coordination, 293
gas-subcooled process, 212
  history, 212
glycol contactor, 170
glycol reboiler, 170

Hammerschmidt relation, 113
heavy oil, 55
hole cleaning, 30
hydrate equilibrium curve, 112
hydrate exposure, 113
hydrate inhibitor
  low dosage, 109
  thermodynamic, 107
hydraulic long-stroke unit, 61
  hardware, 73
hydrocyclone, 142
hydrogen sulfide, 189
  analyzer, 202

import terminal, 353
inflow control valve, 83
injection well
  regulation, 99
injectivity, 95
instrument, 12

Joule–Thomson cooling, 111
Joule–Thomson valve, 213
Joule-Thomson valve, 332

kick detection, 48

leak detection, 383
  federal rule, 314
  open problems, 315
liquefied natural gas, 353
LNG carrier, 359
LNG facility
  federal safety rules, 346
LNG liquefaction, 321
LNG transfer procedures, 371
load cell
  polished rod, 73
load sharing
  degraded maps, 271

managed pressure drilling, 24
  history, 27
maximum allowable operating pressure, 280, 299
mechanical specific energy, 32
meter factor, 304
  product shift, 311
metering pump, 129
methane emissions, 383
methanol injection, 107
methyldiethanolamine, 190
minimum internal temperature approach, 322, 329
mist extractor, 141
mixed refrigerant, 321, 323
  history, 324
mobility ratio, 82
moisture analyzer, 172
  technologies, 182
motor-operated valve, 314
mud-pulse telemetry, 47
multiphase flow meter, 99
multiphase meter, 116

natural gas
  sweetening, 189
natural gas dehydration, 164
natural gas liquids, 10, 210, 230
natural gas pipeline, 276
  history, 278
Nielsen–Bucklin relation, 113
no-touch time, 117
  practice, 129
non-productive time, 27

offshore platform, 135
oil and gas production, 1
  history, 2
open-rack vaporizer, 361
overbalance, 29

pack-off, 30
pipeline, 8
pipeline insulation, 109
planned shutdown
  pre-dose, 123
plug-risk index, 113
polymer flooding, 82
  adsorption, 85
  history, 82
  oil bank, 96
power management, 383
pressure transmitter
  city gate, 291
pressure-wave propagation, 49
process economics, 216
produced water, 94, 136
production choke, 5, 85, 111
  critical flow, 82
  hardware, 99
products pipeline, 298
pump
  variable-speed drive, 203
pump fillage, 65
pump station, 301
pump trip, 307
pump-off controller, 58
  history, 58
pump-off timer, 67
pumping energy, 312

rate of penetration, 25
reboiler, 192
recondenser, 360
recycle valve, 262
  fail open, 269
refrigerant blend, 333
refrigerant boiling temperature, 328
refrigerant compressor, 328
refrigerant makeup, 329
refrigeration cycle
  cascade, 323
regenerator, 192
relative permeability, 84
relief valve, 303, 371
residue compressor, 213
rich loading, 190
riser, 107
riser slugging, 136
  history, 138
rod fatigue, 60
rod string, 59
rollover, 355
  history, 356
  open problems, 372
rotary drilling, 25

safe drilling margin, 48
sand production, 83
  erosion, 86
  open problems, 100
SCADA, 301
slack line, 300
smart well, 80
Souders–Brown capacity factor, 137
specific power, 340, 347
standing valve, 57
steam header, 235
stick-slip, 31
Stokes settling velocity, 141
stratification, 355
stripping gas, 171
subsea flowline, 106
  history, 109
sucker-rod pumping, 55
suction pressure, 328
surge control line, 262
surge drum, 171, 233
surge tank, 192

temperature transmitter, 128
thermocouple
  cryogenic, 344
thief zone, 80
three-phase separator, 135
transmix, 300
traveling valve, 57
triethylene glycol, 164
  history, 167
turbine flowmeter, 314
turbine trip, 262
turboexpander, 210
  guide vanes, 223
  history, 212

ultrasonic flowmeter
  flare gas, 370
umbilical, 107

variable speed drive, 58
variable-frequency drive
  pipeline compressor, 292
virtual flow metering, 382
voidage replacement, 92

water dew point, 169
water hammer, 300
    open problems, 315
waterflooding, 80
weathering, 355
weight on bit, 26
Welge tangent, 82
wellbore storage, 82
Willans line, 259

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.