INTERACTIVE CONTROL ENGINEERING

Override, Split-Range, and Constraint Control

Interactive Control Engineering: Principles and Applications

Edited by John D. Hedengren

Control technology

1 Process Dynamics, Modeling, and Identification
2 Practical PID Control
3 Feedforward, Cascade, and Ratio Control
4 Override, Split-Range, and Constraint Control
5 Delays, Slow Measurements, and State Estimation
6 Multivariable and Plantwide Control
7 Nonlinear Dynamics, Oscillations, and Stabilization
8 Batch, Run-to-Run, and Hybrid Control
9 Model Predictive Control and Economic Optimization

Industry applications

10 Control in Upstream and Midstream Oil and Gas
11 Control in Petroleum Refining
12 Control in Petrochemicals and Polymers
13 Control in Chemical Plants
14 Control in Semiconductor Manufacturing
15 Control in Mining and Mineral Processing
16 Control in Iron, Steel, and Metals Production
17 Control of Energy Conversion and Storage
18 Control of Nuclear and Fusion Systems
19 Control of Biological Systems and Human Performance
20 Motion Control, Robotics, and Autonomous Systems

INTERACTIVE CONTROL ENGINEERING

Principles and Applications

Volume 4

Override, Split-Range, and Constraint Control

When limits decide what the controller may do

John D. Hedengren

Editor

Advanced Process Solutions, LLC

Interactive Control Engineering: Principles and Applications
Volume 4: Override, Split-Range, and Constraint Control
Edited by John D. Hedengren

First edition, 2026

Copyright © 2026 Advanced Process Solutions, LLC. All rights reserved.

Produced by Advanced Process Solutions, LLC.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations in a review or scholarly work.

The simulators described in this book are teaching models. Their equations, parameters and results are simplified for instruction and are not validated for the design, operation or safety assessment of any real plant, vehicle or medical device. The editor and the publisher accept no liability for use of this material outside that purpose.

Names of products and companies are used for identification only and may be trademarks of their owners.

Every chapter has a companion simulator that runs in a web browser at https://apmonitor.com/apps/. This volume online, with chapter files and citation records:
https://apmonitor.com/apps/books/v04-override-constraint/

How to cite this volume. John D. Hedengren, Ed., *Override, Split-Range, and Constraint Control*, vol. 4 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026.

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

Preface

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

How to read this book with the apps

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

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

How AI agents were used

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

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

About this volume

This volume is about limits. A plant earns the most close to a constraint and is damaged beyond it, and a controller that ignores its limits winds up or fights its neighbors. The chapters show how simple controllers share actuators and respect constraints with selectors, overrides, external reset, split-range and mid-ranging, and how effort is allocated when there are more actuators than objectives. The plants range from a compressor at its surge line and a turbofan fuel valve to a tunnel face, a drilling head, a company short of cash and a currency defended with finite reserves.

Volume 3 covers the cascade and feedforward structures that these schemes build on. Volume 9 treats the same constraints with prediction and optimization.

John D. Hedengren
2026

Contents

Preface v

Notation x

1 Introduction: Selectors, Overrides, and Shared Actuators 1

1.1 The Question This Volume Answers . . . 1
1.2 A Short History . . . 2
1.3 Notation and Conventions . . . 3
1.4 Limits, Selectors, and Shared Actuators . . . 4
1.5 Instruments and Final Elements in This Volume . . . 13
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 Teaching Limit Cycles and Protective Override with a Compressor Surge Simulator 24

2.1 Introduction . . . 25
2.2 Background . . . 26
2.3 Compression System Model . . . 27
2.4 Stability Analysis . . . 31
2.5 Controller Design . . . 33
2.6 Performance Metrics for Override Control . . . 36
2.7 Results . . . 37
2.8 Technology in Practice . . . 43
2.9 Limitations . . . 45
2.10 Conclusion . . . 46
2.11 Exercises . . . 48

3 Limiters That Vote on One Fuel Valve in a Turbofan Simulator 51

3.1 Introduction . . . 52
3.2 Fuel Schedules, Limiters and the Five-Second Rule . . . 53
3.3 The Engine as Modeled . . . 54
3.4 The Controller as Implemented . . . 59
3.5 Ten Cases, Measured . . . 61
3.6 Technology in Practice . . . 68
3.7 Limitations . . . 70
3.8 Conclusion . . . 71
3.9 Exercises . . . 72

4 Feed Rate, Grain Loss and a Rain Deadline in a Combine Harvester Simulator 76

4.1 Introduction . . . 77
4.2 Automation of the Combine Harvester . . . 78
4.3 Machine and Field Model . . . 81
4.4 What Shapes the Control Problem . . . 85
4.5 Implemented Controller . . . 86
4.6 Measured Scenarios . . . 90
4.7 Technology in Practice . . . 96
4.8 Limitations . . . 98

4.9 Conclusion . . . 99
4.10 Exercises . . . 101

5 Shear Heating, Overrides and Delayed Density Feedback in a Cereal Extrusion Simulator 104
5.1 Introduction . . . 105
5.2 Extrusion Cooking, Its Measurement and Its Control . . . 106
5.3 Model of the Line . . . 107
5.4 Where the Heat Comes From . . . 114
5.5 Implemented Controller . . . 115
5.6 The Cases as Measured . . . 120
5.7 Technology in Practice . . . 128
5.8 Limitations . . . 130
5.9 Conclusion . . . 131
5.10 Exercises . . . 133

6 The Drying Front Sets the Line Speed of a Battery Electrode Coater 136
6.1 Introduction . . . 137
6.2 Electrode Lines, Solvent Ovens and Binder Migration . . . 138
6.3 The Coating Line as Modeled . . . 139
6.4 Why the Drying Front Carries the Cap . . . 143
6.5 The Control Layers as Implemented . . . 145
6.6 Code Verification and Agreement With Sources . . . 149
6.7 Nine Cases and a Challenge, Measured . . . 150
6.8 Technology in Practice . . . 156
6.9 Limitations . . . 158
6.10 Conclusion . . . 158
6.11 Exercises . . . 160

7 Horizontal Directional Drilling Where the Soil Chooses the Binding Limit 164
7.1 Introduction . . . 165
7.2 Trenchless Crossings, Mud Pressure and the Pilot Bore . . . 166
7.3 The Crossing as Modeled . . . 167
7.4 Sequencer, Guidance, Autopilot and Optimizer as Coded . . . 173
7.5 What the Nine Runs Show . . . 175
7.6 Code Checks, Correlations and the Limits of the Evidence . . . 179
7.7 Technology in Practice . . . 180
7.8 Limitations . . . 182
7.9 Conclusion . . . 183
7.10 Exercises . . . 184

8 Growth Control Against a Cash Floor in a Simulated Product Startup 188
8.1 Introduction . . . 189
8.2 From Learning Curves and Market Growth to the Venture Round . . . 190
8.3 The Company as Stocks and Flows . . . 191
8.4 Why Growth Costs Cash . . . 196
8.5 The Board, Its Loops and Its Planner . . . 198
8.6 Measured Cases . . . 201
8.7 Technology in Practice . . . 207
8.8 Limitations . . . 209
8.9 Conclusion . . . 210
8.10 Exercises . . . 211

9 Valve-Position Control of Chilled Water in a Gas Turbine Trigeneration Plant 215
9.1 Introduction . . . 216
9.2 From Peaking Turbines to Trigeneration . . . 217
9.3 Plant Model . . . 218
9.4 Controllers as Implemented . . . 224
9.5 Measured Behavior . . . 227
9.6 Technology in Practice . . . 235
9.7 Limitations . . . 237
9.8 Conclusion . . . 238
9.9 Exercises . . . 239

10 A Supercritical CO_2 Brayton Simulator for Three-Layer Mid-Ranging Control 243
10.1 Introduction . . . 244
10.2 Background . . . 245
10.3 Process Model . . . 246
10.4 Actuator Characteristics and Open-Loop Response . . . 250
10.5 Controller Design . . . 252
10.6 Results . . . 256
10.7 Technology in Practice . . . 263
10.8 Limitations . . . 265
10.9 Discussion . . . 266
10.10 Conclusion . . . 266
10.11 Exercises . . . 268

11 Holding a Tunnel Face With Its Own Muck Through Sixty-Six Stops 271
11.1 Introduction . . . 272
11.2 Closed Shields, Face Support and the Trough . . . 273
11.3 The Drive as Modeled . . . 274
11.4 Why the Face Is Hard to Hold . . . 279
11.5 The Control Layers as Implemented . . . 280
11.6 Results . . . 283
11.7 Verification, Agreement and What Is Not Claimed . . . 289
11.8 Technology in Practice . . . 290
11.9 Limitations . . . 291
11.10 Conclusion . . . 292
11.11 Exercises . . . 294

12 Allocating Tires, Suction Fans and Cold-Gas Thrusters on One Nitrogen Tank 297
12.1 Introduction . . . 298
12.2 Suction, Cold Gas and the Allocation of Thrust . . . 299
12.3 The Car as Modeled . . . 300
12.4 Three Prices and One Budget . . . 305
12.5 The Controllers as Implemented . . . 306
12.6 Ten Cases, Measured . . . 309
12.7 Technology in Practice . . . 312
12.8 Limitations . . . 314
12.9 Conclusion . . . 315
12.10 Exercises . . . 317

13 Quadruped Balance Simulator for Sharing Ground Forces Inside Friction Cones 320
13.1 Introduction . . . 321
13.2 Balance on Legs, From Creeping Gaits to Force Planning . . . 322

13.3 Planar Torso, Legs and Contacts . . . 323
13.4 Loops, Allocation, Gait and Planner . . . 326
13.5 What the Feet Can Deliver . . . 331
13.6 Results Across the Cases . . . 333
13.7 Verification of the Engine . . . 338
13.8 Technology in Practice . . . 339
13.9 Limitations . . . 341
13.10 Conclusion . . . 341
13.11 Exercises . . . 343

14 Defending a Currency Peg With a Finite Vault and Too Few Instruments 346
14.1 Introduction . . . 347
14.2 From a Gold Stockpile to a Franc Floor . . . 348
14.3 A Vault, a Band and a Money Market . . . 350
14.4 Why the Vault Drains Faster Than Its Clock . . . 354
14.5 Two Handles, a Calendar and an Override . . . 356
14.6 Nine Defenses, Measured . . . 359
14.7 Technology in Practice . . . 365
14.8 Limitations . . . 367
14.9 Conclusion . . . 368
14.10 Exercises . . . 369

15 Summary and Outlook 373
15.1 What the Chapters Showed . . . 373
15.2 Choosing a Structure . . . 376
15.3 Where Practice Is Heading . . . 377
15.4 Engineering With AI Agents . . . 379
15.5 Gaps in Override and Constraint Control . . . 381
15.6 Exercises . . . 386

Answers to Selected Exercises 387

Glossary 393

Index of Control Topics 396

Index of Applications 398

About the Editor 401

Volumes in This Series 402

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: Selectors, Overrides, and Shared Actuators



The chapters 2 compressor near surge 3 turbofan engine 4 combine harvester 5 cereal extruder 6 electrode dryer 7 directional drill 8 product startup 9 trigeneration plant 10 CO_2 power cycle 11 tunnel boring shield 12 fan car 13 quadruped robot 14 currency peg

Notice that the units from the compressor to the drill each guard one main actuator with several limits, while most from the trigeneration plant to the robot share one task among several actuators. Illustration, not to scale.

IN PLAIN TERMS **Which Goal Gets the Handle**

The main idea. A plant usually runs best close to a limit, such as a temperature it must not pass or a valve that is already wide open. The methods in this book decide which goal gets a shared handle, and which of several handles moves first.

The challenge. A guard that waits unused for hours must still take over smoothly at the right moment. A supply that can be spent quickly, like a cold store or cash in the bank, runs out, so it becomes a limit of its own.

Why it matters. The same few ideas protect jet engines, drills, power plants, robots and a country's money.

THREE WAYS TO READ THIS CHAPTER

Explore	8 min	This page, Section 1.2 with its timeline and two History boxes, the Did you know? boxes in Sections 1.1, 1.4.3 and 1.4.7, Figure 1.10, and Key points in Section 1.8. No equations.
Study	50 min	The whole chapter in order, with each number of the reactor example worked, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.10 with Table 1.4, Section 1.5 with Figures 1.6 and 1.8, and Section 1.8 with Figure 1.10.

1.1 The Question This Volume Answers

A controller is usually taught with one measurement, one actuator and a setpoint well inside the actuator's range. Plants do not stay there for long. A cooling valve opens wide, a motor reaches its rating, a pressure approaches what the ground or the vessel can hold, and the most profitable operating point turns out to lie on one of these limits. This volume answers two questions that arise there. When several variables must each stay inside a limit and only one actuator can move them, which variable gets the actuator? When several actuators can move one variable, which one should move, and for how long?

The first question is answered by selectors. A low or high selector passes one of several controller demands to a single actuator, so that the normal controller runs the plant until a limit comes near and an override controller takes over to hold it. The second question is answered by split range, mid-ranging and control allocation, which share one demand among actuators that differ in range, speed and price. A third kind of actuator joins the two questions. A reserve of chilled water, compressed gas, cash or foreign currency can be spent quickly but not forever, so an actuator that depletes is also a limit that the plant approaches with time.

DID YOU KNOW? Two Selectors Keep a Ship's Boiler From Smoking

A ship's boiler is tuned for efficiency at sea but must run with excess air in harbor, where black smoke brings heavy penalties. A maximum and a minimum selector around its air and oil flow loops keep the air ahead of the fuel as the power demand rises or falls. Selectors commonly implement logic in engines and power systems [1].

Thirteen chapters test these structures, each on a plant with an app that runs in a browser. Chapter 2 and Chapter 3 protect machines that can damage themselves within seconds, a compressor near surge and a jet engine whose single fuel valve answers to six limits. Chapters 4 to 7 push throughput against limits that move: a combine harvester, a cereal extruder, a battery electrode coater and a horizontal drill. Chapter 8 runs a company against a cash floor. Chapters 9 to 11 return a fast actuator to the middle of its range in a trigeneration plant, a supercritical carbon dioxide power cycle and a tunnel boring machine. Chapters 12 and 13 allocate forces among tires, fans, thrusters and feet, and Chapter 14 defends an exchange rate with a vault that empties.

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

- Name the limits that are active at the plant's best operating point, and state what each degree of back-off from a limit costs.
- Build an override with a low or high selector and external reset, and predict when the override takes over.
- Choose the split point of a split-range controller from the actuator gains, and say when mid-ranging should replace it.
- Tune a valve-position controller from the gains of its two actuators, and state what its resting position costs.
- Write an allocation rule for actuators with different limits and prices, and compute how long an actuator that depletes will last.

This chapter develops each skill on one small example and names the chapters that use it. PID tuning is not repeated here. Volume 2 develops it, and every loop in this chapter is a PI tuned from its model by the rules of that volume. Standard texts treat selectors, overrides and split range in fuller detail [2–4].

1.2 A Short History

The milestones of Figure 1.1 follow selectors, overrides and shared actuators through the kinds of plant in this volume. Much of the development of PID control occurred in the power and process industries, where many of the bottom-up architectures around it also appeared, among them selector and midranging control [1]. Pneumatic control became the standard of the oil, gas and chemical industries during the 1940s [5], and an override built from such controllers needed a way to stop the integral of the controller it did not pass.

HISTORY 1931 Onward: A Reset Line From the Valve Stem

An early pneumatic controller left an offset, which the operator removed by resetting the setpoint by hand. A circuit that copied that correction, called automatic reset, is the integral action of today, and a controller with it was marketed in 1931 [6]. Some pneumatic controllers come with an option for external reset: the reset bellows takes its signal from where the final effect of the output is seen, such as a position transmitter on the valve stem. If the stem meets a stop or a jam, that signal freezes and integration ceases instead of winding up [16]. Section 1.4.4 feeds the override's integral from the selector's output in the same way.

Engines carried caps and floors from the start. The first jet engine built in the United States, in 1942, had a governor whose fuel metering valve was given a minimum-flow stop against flameout

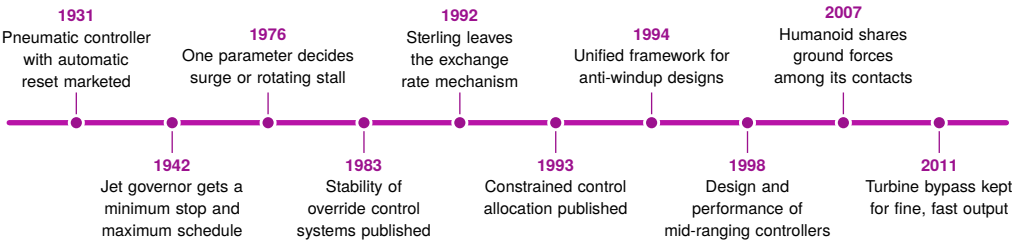

Figure 1.1 Milestones of override, selector and constraint control in the plants of this volume. The dates are from [6] (1931), [7] (1942), [8] (1976), [9] (1983), [10] (1992), [11] (1993), [12] (1994), [13] (1998), [14] (2007) and [15] (2011). Notice that a currency, a robot and a power cycle sit among the engines and the theory: the same structures served them all. Events are evenly spaced, not to a time scale.

and a maximum-flow schedule against overtemperature [7]. For compressors, a study of 1976 found one nondimensional parameter that decides whether a compression system at its stall line falls into surge or into rotating stall [8]. Chapter 2 calls it B.

Papers on the stability of override control systems appeared in 1983 [9], on constrained control allocation in 1993 [11] and on the design and performance of mid-ranging controllers in 1998 [13]. The same structures served far from process plants. On one day in September 1992 a central bank reportedly spent as much as half its foreign exchange reserves defending its band, and its currency left the exchange rate mechanism at the end of the day [10]. A humanoid robot of 2007 shared the ground force it needed among any number of contact points [14]. A 2011 report on supercritical carbon dioxide cycles kept the turbine bypass, the simplest and fastest handle, for fine and fast adjustments of the generator output [15].

> ⧖ HISTORY **1994: Every Anti-Windup Scheme in One Frame**
>
> Windup was first met in PI and PID controllers whose actuators saturate. An early remedy, an algorithm for direct digital control, has been called anti-reset windup, back calculation and tracking, or integrator resetting. Integrator windup was later seen as a special case: any controller with slow or unstable modes winds up when its actuator is limited. An override adds a second cause. When the plant input is chosen among parallel controllers, those not chosen are wrongly updated, and the switch bumps the plant input. A framework published in 1994 showed every known linear anti-windup and bumpless transfer scheme to be a special case of one design: the linear controller first, then a compensation for the limits [12].

1.3 Notation and Conventions

The chapters share the symbols of the notation page. The setpoint is SP, the process variable is PV and the manipulated variable is MV. Controller settings are K_c and τ_I, and a process model has a gain K_p, a time constant τ_p and a dead time θ_p. A tag such as TIC or FIC names a temperature or flow controller, and TC-2 names the override of this chapter. Following the shared framework, v is the demand a controller computes before any limit and u is the value applied to the actuator, after the clamp, the rate limit or the selector. The tracking time of back-calculation is T_t, equal to τ_I unless a chapter says otherwise. A split-range controller divides its output c at the split point s, both in percent. A valve-position controller (VPC) holds a fast actuator at its resting position r. Heat duties are Q, in kW, and the charge of a store is E.

Each chapter keeps the time unit of its plant, from a hundredth of a second in the turbofan to weekly decisions in the currency peg. The running example of this chapter uses minutes. A few symbols need care across chapters. The shared framework uses a for the controller action and h for

the integration step, so neither letter is used here for anything else. Several chapters use v for a speed, the ground speed in Chapter 4, the line speed in Chapter 6 and the advance in Chapter 11. Chapter 2 writes s_k for the accumulator of its controller, and Chapter 8 writes z for an override level, which is a split-range input.

1.4 Limits, Selectors, and Shared Actuators

1.4.1 The Running Example

The running example is a continuous reactor that converts a liquid feed in an exothermic reaction. The feed F, in kg/min, is the product the plant sells, and each kg/min releases 25 kW. Two coolants remove the heat. Cooling water passes a jacket and removes up to 400 kW, 4 kW per percent of its valve stroke. It is cheap, but it reaches the reactor 1.5 min after its valve moves, the time the water takes to cross the jacket circuit. Chilled water from a cold store removes up to 100 kW, 1 kW per percent of stroke, and acts at once. The store holds 100 kWh, enough for one hour at full flow. The reactor temperature T follows Equation (1.1).

$$\tau_r \frac{\mathrm{d}T}{\mathrm{d}t} = -(T - 80) + \frac{\Delta H\, F - Q_w - Q_c + d}{UA} \tag{1.1}$$

where $\tau_r = 10\,\mathrm{min}$, $\Delta H = 25\,\mathrm{kW}$ per kg/min, $UA = 5\,\mathrm{kW/^\circ C}$, Q_w and Q_c are the cooling-water and chilled-water duties, and d is a load disturbance in kW. The thermometer adds 0.5 min of dead time. Each path to the measured temperature is therefore first order with $\tau_p = 10\,\mathrm{min}$. The feed path has a gain of 5 °C per kg/min and a dead time of 0.5 min, the cooling-water path −0.8 °C/% and 2 min, and the chilled-water path −0.2 °C/% and 0.5 min.

At design the feed is 10 kg/min, its 250 kW is carried by the cooling water at 62.5 % of stroke, and the temperature is 80 °C. The reactor trips at 90 °C. Two controllers act on the temperature, as shown in Figure 1.2. TIC-1 holds 80 °C with the coolants. TC-2 is an override on the feed with a setpoint of 85 °C. It does nothing while the cooling copes and cuts the feed when the cooling cannot. Both are PI controllers tuned by the SIMC rule of Volume 2 with $\tau_c = \theta_p$ and executed every 0.1 min in the velocity form of Section 1.6.2. TIC-1 has $K_c = 2.5\,\%/^\circ\mathrm{C}$ and $\tau_I = 10\,\mathrm{min}$ on the split range of Section 1.4.6. TC-2 has $K_c = 2\,\mathrm{kg/min}$ per °C and $\tau_I = 4\,\mathrm{min}$.

1.4.2 Constraints as Objectives

The reactor earns on its feed, so the operator would run it as fast as the cooling allows. At steady state Equation (1.1) gives the feed that a temperature T and the two duties can carry, which is Equation (1.2) with the cooling water at full stroke and the store unused.

$$F = \frac{Q_w + Q_c + UA\,(T - 80)}{\Delta H} = 16 + 0.2\,(T - 80)\ \mathrm{kg/min} \tag{1.2}$$

Two limits hold the best steady point. The cooling-water valve is wide open, and the temperature sits as high as the trip allows. Neither value is chosen for its own sake. Each is a constraint that the economics press the plant against, so the constrained variable becomes a controlled variable and its limit, less a margin, becomes its setpoint. Such a constraint is called active. Which constraints are active depends on prices, feed and disturbances, and a control structure must switch its controlled variables when the active set changes [17, 18].

The temperature cannot be held at the trip itself, because every upset would trip the reactor. Its setpoint is backed off from the limit by the largest excursion the control leaves, plus a margin. Each degree of back-off costs 0.2 kg/min here, 1.2 % of the 17 kg/min that the override setpoint of 85 °C allows. A structure that overshoots less permits a setpoint nearer the limit, which is how better control is paid for in throughput [19]. The runs of Section 1.4.3 show the size of the effect. With external reset the override overshoots its setpoint by 0.95 °C, and without it by 2.87 °C. With the setpoint

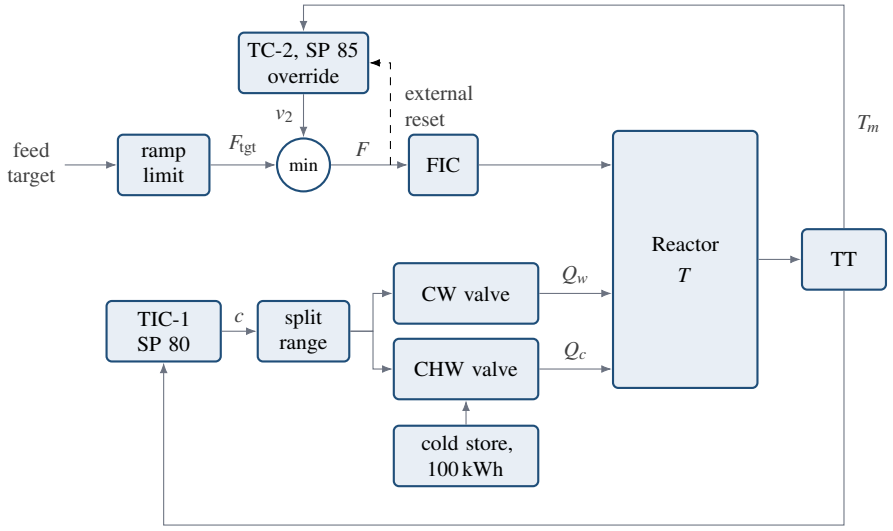

Figure 1.2 The Running Example. A Low Selector Passes the Smaller of the Ramped Feed Target and the Demand of the Temperature Override TC-2 to the Feed Flow Controller. TIC-1 Moves Cooling Water (CW) and Chilled Water (CHW) Through a Split Range. The Dashed Line Returns the Applied Feed to TC-2 (External Reset). Section 1.4.7 Replaces the Split Range With Mid-Ranging.

moved to 88 °C, the first peaks at 88.67 °C and earns 17.6 kg/min, while the second crosses the trip for 2.2 min.

The reactor shows the two shapes that every constraint problem in this volume takes. With one feed and several limits to respect, there are more variables to protect than actuators to move, and a selector decides which variable gets the actuator. Sections 1.4.3 to 1.4.5 treat that case. With two coolants for one temperature, there are more actuators than controlled variables, and split range, mid-ranging or an allocation rule decides which actuator moves. Sections 1.4.6 to 1.4.9 treat that case.

Three chapters push a throughput against its active constraints. Chapter 5 ramps its feed while every predicted constraint keeps a slack in units of its own margin, and declares a target infeasible with the binding constraint named. Chapter 6 pushes the line speed to a vapor target of 20 % of the lower explosive limit, under an override at 21 and a cap at 25 %. Chapter 7 finds that the soil chooses the binding limit, torque in clay, the mud ratio in sand and thrust in gravel. Chapter 3 measures the price of a surge margin in seconds of acceleration, and Chapter 10 argues that a plant should be run until every constraint binds, so that the operator knows which one comes first.

1.4.3 High and Low Selectors and Override

A selector passes one of several demands to a single actuator. A low selector passes the smallest, so each of its inputs acts as a cap: the actuator may move no higher than any input allows. A high selector passes the largest, so each input acts as a floor. On the reactor the feed is capped by the operator's target and by the override, as shown in Equation (1.3).

$$F = \min\left(F_{\text{tgt}},\ v_2\right), \qquad v_2 = K_c\, e_2 + I_2, \qquad e_2 = 85 - T_m \tag{1.3}$$

where T_m is the measured temperature and I_2 is the override's integral term. The override is reverse acting, since more feed raises the temperature. Its setpoint lies inside the hard limit and above the normal operating temperature. In healthy operation its demand is large and the selector ignores it, and its quality is judged by the worst excursion it allows, not by an integral of error that it never intends to remove.

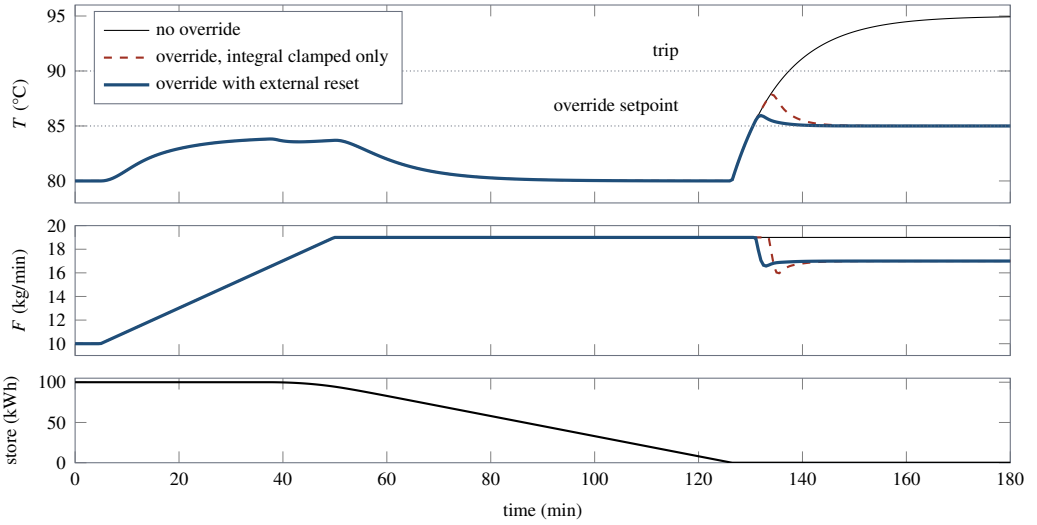

Figure 1.3 One Shift on the Reactor. The Feed Target Ramps From 10 to 19 kg/min Between 5 and 50 min, and the Cold Store Empties at 126.4 min. Thin Solid: No Override. Dashed: Override Whose Integral Is Only Clamped to Its Output Range. Thick Solid: Override With External Reset. The Store Charge Is the Same in All Three Runs.

When caps and floors act on one actuator, their order matters. The floors are applied after the caps when what a floor protects outranks what the caps protect, as in $u = \max\left[u_{\text{floor}}, \min(v_1, \ldots, v_n)\right]$. A selector can also be realized without a separate block, by setting the override's upper output limit to the other demand, so that the clamp of Equation (1.14) performs the selection. Where two demands can sit close together, a small margin that an incumbent must be undercut by keeps noise from passing the actuator back and forth.

Figure 1.3 runs the reactor through one shift. From 5 min the feed target ramps from 10 to 19 kg/min at 0.2 kg/min per minute. The cooling water reaches full stroke at 39 min, and the split range moves the chilled water in to carry the remaining 75 kW. The store empties at 126.4 min. From then on the cooling falls 75 kW short. Without an override the temperature passes the trip and heads for 95 °C, above 90 °C for 42.6 min of the run. With the override and external reset, TC-2 takes the feed at 131.0 min, when the measured temperature reaches 85.02 °C, and settles the feed at the 17 kg/min that Equation (1.2) predicts. Table 1.1 gives the measures. Without external reset, the override's integral has run to its upper limit of 25 kg/min during the hours of healthy operation. It must fall below the applied 19 kg/min before the selector passes it, so the override takes over 2.6 min later, at 87.31 °C, and the temperature peaks at 87.87 °C.

Most chapters of this volume have a selector or a clamp that acts as one. Chapter 3 takes the least of a speed governor, an acceleration schedule, a turbine temperature limiter, two overspeed limiters, a fuel rate limit and a stall cut, and then applies two floors that keep the flame lit, because a flameout outranks every limit a cap protects. Chapter 4 passes the lowest of five ground-speed demands, with an incumbent margin of 0.072 km/h. Chapter 5 realizes a high selector on the screw speed and a low selector on the feed as output limits of its two overrides. Chapter 7 and Chapter 11 take the least of several limits on the thrust and the advance, and Chapter 9 low-selects its power loop against an exhaust limiter on one fuel valve. Chapter 2 has the simplest form, a single protective loop whose output rests on its closed clamp while the machine is healthy.

Table 1.1 The Override After the Store Empties, by Anti-Windup Scheme and Tracking Time T_t. Takeover Time and the Measured Temperature Then, Peak Temperature, and Time Above 85.5 °C in the 180 min Run

Override	Takeover (min)	T_m then (°C)	Peak (°C)	Above 85.5 (min)
None	—	—	94.93	49.0
Integral clamped only	133.6	87.31	87.87	9.0
External reset, $T_t = 0.1\,\tau_I$	127.5	80.83	85.00	0
External reset, $T_t = 0.25\,\tau_I$	128.3	81.92	85.00	0
External reset, $T_t = 0.5\,\tau_I$	129.3	83.17	85.00	0
External reset, $T_t = \tau_I$	131.0	85.02	85.95	3.0
External reset, $T_t = 2\,\tau_I$	133.7	87.38	87.97	9.3
External reset, $T_t = 4\,\tau_I$	138.1	90.09	90.40	15.5

✦ DID YOU KNOW? A Hard and a Soft Override on One Well Pump

A well pump on a variable-speed motor holds the water pressure for its customers. A level switch that stops it when the well runs low saves the seals but cuts off the water, a hard override. A level controller low-selected against the pressure controller only slows the pump, a soft override [16].

1.4.4 External Reset

A controller that the selector does not pass still computes a demand, and its integral still sees an error. Left alone, the integral of the override in the reactor runs upward for as long as the temperature sits below 85 °C, which is the windup of Table 1.1. The cure is to feed the integral with the output actually applied. The shared framework already does this through back-calculation, Equation (1.15), once the applied output u is taken after the selector. For a PI with no bias, where the running sum equals $K_c e$, the integral obeys Equation (1.4).

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \frac{K_c}{\tau_I}\,e + \frac{u - (K_c e + I)}{T_t} \xrightarrow{T_t = \tau_I} \frac{\mathrm{d}I}{\mathrm{d}t} = \frac{u - I}{\tau_I} \tag{1.4}$$

With $T_t = \tau_I$ the error drops out. The integral is the applied output passed through a first-order lag of time constant τ_I, and the demand is $v = K_c e + I$. This is external reset feedback [3, 20]. While the controller holds the actuator, $u = v$, and the lag of its own output reproduces integral action. While another controller holds it, the integral follows the applied output and the demand rests at its proportional distance above it, $v - u \to K_c e$. The override therefore takes over when its own error changes sign, at its setpoint, without a bump and without delay. On the reactor at 83 °C, for example, TC-2 asks for $19 + 2 \times 2 = 23\,\mathrm{kg/min}$ while 19 kg/min is applied.

A different tracking time shifts the takeover. For a steady applied output the demand rests at $v - u = K_c e\,T_t/\tau_I$, and while the error falls the integral trails its own target by about T_t. A short T_t lets the override take over on its rate of approach, before its limit is reached. A long one lets it wind partly and take over late. Table 1.1 measures both. At $T_t = 0.1\,\tau_I$ the override cuts the feed when the measured temperature is still 80.83 °C, which here costs 1.9 kg of feed and removes the overshoot. At $4\,\tau_I$ it waits until 90.09 °C and the reactor trips. Only $T_t = \tau_I$ takes over at the limit itself. Anti-windup schemes of this kind are compared in [12] and treated for selector structures in [4].

Chapter 3 measures the same trade on a jet engine. Its plan asked for every integrator to be reset to the applied fuel each scan, and with that reset the temperature limiter won the valve on its rate of approach 215 K under its limit. Tracking times of 1.5 and 2 integral times let the challenge's turbine temperature reach 973 and 1012 °C, so the chapter keeps $T_t = \tau_I$. Chapter 4 writes the lag of Equation (1.4) explicitly, and Chapter 11 and Chapter 7 build each limit as an integrator on the request actually selected. In Chapter 2 the clamp is the other input of the selector. With the valve shut, Equation (1.4) gives $I_k = (1 - \Delta t/\tau_I)\,I_{k-1}$, so the integral decays to zero and the valve opens

at the first execution after the flow crosses its setpoint. Chapter 8 documents the opposite case. Its runway override acts on the request after the growth loop, so that loop's integral keeps running while the override holds the spending down. A cascade needs the same care one level up. In Chapter 4 the loss supervisor may lower the feed setpoint while an override holds the speed but may not raise it.

1.4.5 Saturation and Rate Limits

An actuator at its limit has no more to give, and no tuning changes that. After the store empties, TIC-1 asks for more cooling than exists, its output reaches 100 % and back-calculation holds its integral there. The plant then needs a different actuator, here the feed. A protective actuator that saturates is a sizing statement, as Chapter 2 shows with a recycle valve whose full opening cannot hold the flow, and one that strokes too slowly fails even when the controller's decision is right.

The empty store hides a subtler case. The chilled-water valve still moves, but it no longer cools, so the controller sees an actuator whose gain has fallen to zero without reaching its clamp. Back-calculation cannot react to a limit it is not told about. Chapter 13 handles the same problem by back-calculating each loop from the force that its feet actually delivered, and Chapter 12 reads the thrust its tank can still supply from the tank pressure.

A rate limit restricts how fast an actuator or a setpoint may move. A PI loop that must deliver an output rising at rate ρ lags the ramp by the error of Equation (1.5), the steady ramp error of a loop with integral action.

$$e_\rho = \frac{\rho\,\tau_I}{K_c} \tag{1.5}$$

A feed ramp of r kg/min per minute asks TIC-1 for $25r/5 = 5r$ percent of its output per minute, so $e_\rho = 20r$ °C. The ramp of 0.2 kg/min per minute lags by 4 °C in principle, and the simulated temperature reaches 83.82 °C before the ramp ends. A ramp of 0.5 kg/min per minute would lag by 10 °C, and without an override the temperature would reach 87.34 °C during the ramp, past the override's setpoint. The ramp limit is a constraint handled open loop. It should be chosen from the margin to the override, here $r \le (85-80)/20 = 0.25\,\mathrm{kg/min}$ per minute, and a faster demand hands the job to the selector.

Chapter 3 places its fuel rate limit among the caps of its selector. Chapter 4 applies acceleration limits after the selector, where no PI loop sees them, and Chapter 2 applies one inside the controller and again at the valve. Chapter 8 and Chapter 14 shape every decision with a quantum and a calendar, and the currency peg also with a move cap, so that a request can change only at a board meeting or a weekly decision and only by whole steps.

1.4.6 Split Range

A split-range controller drives several actuators in sequence from one output. TIC-1 uses the first part of its output c for the cooling water and the rest for the chilled water, as in Equation (1.6).

$$u_w = \frac{100\,c}{s}, \qquad u_c = \frac{100\,(c-s)}{100-s}, \qquad \text{each clipped to 0–100 \%} \tag{1.6}$$

The cheap actuator comes first, so the store is touched only when the cooling water is wide open. The split point sets the gain the loop sees on each segment. One percent of output moves the cooling water by $100/s$ percent and the chilled water by $100/(100-s)$ percent. The two segments give the same duty per percent of output when the split point is chosen by Equation (1.7).

$$s = 100\,\frac{Q_{w,\max}}{Q_{w,\max}+Q_{c,\max}} = 100 \times \frac{400}{500} = 80\,\% \tag{1.7}$$

Both segments then give 5 kW per percent. A split at 50 % would give 8 and 2 kW per percent, as shown in Figure 1.4, and a loop tuned on the first segment would be four times too slow on the second. A gap between the segments leaves a band of output that moves nothing, a dead zone in which the loop can cycle, and an overlap doubles the gain where both move.

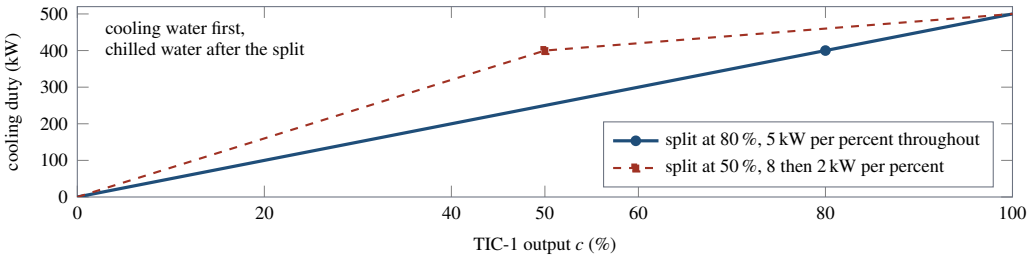

Figure 1.4 Cooling Duty Against the Output of TIC-1. Thick Solid With a Circle: Split Point at 80 % From Equation (1.7). Dashed With a Square: Split Point at 50 %. The Marker Is the Split Point.

On the shift of Figure 1.3, the split range brought the chilled water in by itself when the cooling water reached full stroke, and the output of TIC-1 settled at 95 %. That is the strength of the structure. It extends the range of one controller across several actuators in a fixed order of preference [21]. Its weakness is that it moves one actuator at a time, always the first one first, and assumes they act at similar speed. Here they do not. The cooling water reaches the reactor 1.5 min later than the chilled water, and the controller crosses from one to the other with the same tuning. Generalized forms give each actuator its own controller and pass a baton between them [22].

Three chapters use the structure. Chapter 5 heats or cools each barrel zone from one output between -1 and 1, bands above zero and cooling water below it. Chapter 8 splits the level of its runway override over three ranges, so that marketing and investment go first, then the credit line is drawn, and only then is research cut. Chapter 12 launches its car as a split range in time, with the fans spooled first, the tires pushing next and the thrusters supplying only what the tires cannot.

1.4.7 Mid-Ranging and Valve-Position Control

When the cheap actuator is slow, it should not be the first to answer. In mid-ranging the fast actuator holds the controlled variable, and a slow valve-position controller moves the second actuator until the fast one is back at a resting position r [13, 23]. On the reactor, TIC-1 moves the chilled water alone, with $K_c = 50\,\%/^\circ\mathrm{C}$ and $\tau_I = 4\,\mathrm{min}$ from the same SIMC rule, and the VPC moves the cooling water by integral action on the chilled-water demand v_c, as shown in Equation (1.8).

$$\frac{\mathrm{d}u_w}{\mathrm{d}t} = K_I\,(v_c - r), \qquad r = 10\,\%, \qquad K_I = 0.05\,\mathrm{min}^{-1} \tag{1.8}$$

Through the closed temperature loop, each percent that the cooling water opens is cancelled by four percent less chilled water, so the VPC sees a static gain of 4. An integrator on a static gain closes with the time constant $\tau_h = 1/(4K_I) = 5\,\mathrm{min}$, chosen well above the 2 min that the cooling water takes to act.

The handover has an exact cost. After a load step D that the cooling water can carry, the VPC must move it by $D/4$ percent, and by Equation (1.8) that move equals K_I times the integral of $v_c - r$. While the chilled water is not clamped, its energy above the resting draw is therefore fixed by Equation (1.9), whatever the tuning of TIC-1.

$$E_h = \int \left(Q_c - Q_{c,r}\right) \mathrm{d}t = \frac{K_{ch}}{K_w K_I}\, D = \tau_h D \tag{1.9}$$

where $K_w = 4$ and $K_{ch} = 1\,\mathrm{kW}$ per percent and $Q_{c,r}$ is the resting draw. A step of 50 kW costs 250 kW min, about 4.2 kWh, and the simulation gives the same 250.0 kW min.

Figure 1.5 and Table 1.2 compare the two structures. Upward, mid-ranging cuts the peak from 2.96 to 0.95 °C and the IAE nearly eightfold, because the fast actuator answers first and the slow one takes over behind it. Downward the gain is smaller. Resting at 10 kW, the chilled water can give back

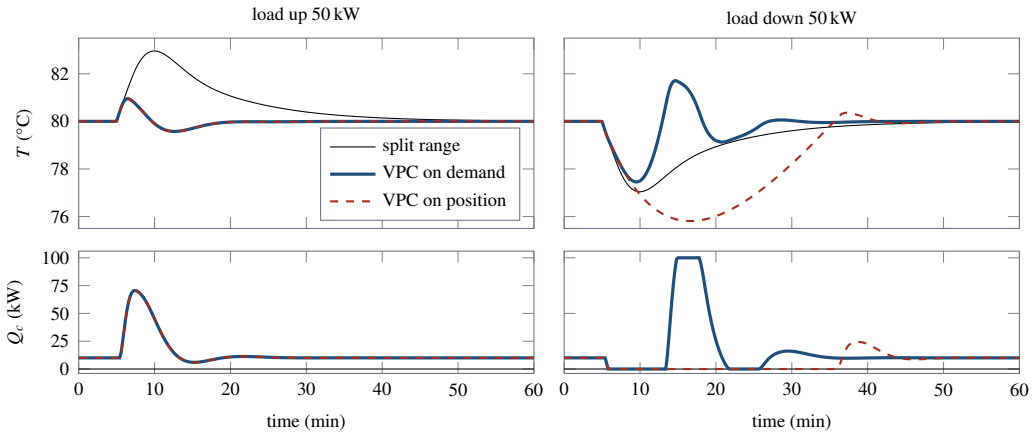

Figure 1.5 A Load Step of 50 kW at 5 min, Up (Left) and Down (Right), at the Design Feed. Top: Reactor Temperature. Bottom: Chilled-Water Duty. Thin Solid: Split Range. Thick Solid: Mid-Ranging With the VPC Watching the Unclamped Chilled-Water Demand. Dashed: VPC Watching the Clamped Valve Position, Identical to Thick Solid on the Upward Step.

Table 1.2 Split Range Against Mid-Ranging on the Steps of Figure 1.5, Over 60 min. Peak Deviation in °C, IAE in °C min, Chilled Water Above the Resting Draw in kW min

	Load up 50 kW			Load down 50 kW	
Structure	Peak	IAE	Chilled	Peak	IAE
Split range	2.96	39.8	0	−2.96	39.8
VPC on the unclamped demand	0.95	5.2	250	−2.55	24.1
VPC on the clamped position	0.95	5.2	250	−4.19	85.3

only 10 kW at once, then sits at its lower clamp, and the cooling water must do the rest. What the VPC watches decides how fast. Watching the unclamped demand, it sees how far below zero the temperature loop would like to go and moves the cooling water accordingly, fast enough to overshoot. The temperature rises to 81.71 °C and the chilled water runs at full stroke for 3 min, yet the IAE is 24.1 against 85.3. Watching the clamped position, it sees $0 - 10$ for any deficit and walks the cooling water down at 0.5 % per minute. A clamped output carries no information about how far past the clamp a controller wanted to go.

Table 1.3 A Cooling Demand Shared Between Cooling Water (CW, 400 kW at Most) and Chilled Water (CHW, 100 kW at Most) by Priority and by Equation (1.10) With Prices 1 and 4. Duties in kW

	Priority, CW first			Quadratic prices 1 and 4		
Demand	CW	CHW	Short	CW	CHW	Short
250	250	0	0	200	50	0
375	375	0	0	300	75	0
450	400	50	0	360	90	0
480	400	80	0	384	96	0
520	400	100	20	400	100	20

✦ DID YOU KNOW? **Mid-Ranging Can Lose to Split Range**

From the simulator: on the downward load step of Table 1.2, split range leaves an IAE of 39.8 and a dip of 2.96. Mid-ranging whose VPC watches the clamped valve position leaves 85.3 and 4.19, worse than the simpler structure it was meant to improve.

The resting position buys authority in both directions with the stock it draws. At $r = 10\,\%$ the store loses 10 kW at all times, 600 kW min an hour, more than two handovers of a 50 kW step, and that draw alone would empty it in ten hours. At 5, 10 and 20 %, the dip after the downward step is 2.86, 2.55 and 1.96 °C, and the chilled water drawn over the 60 min upward run is 550, 850 and 1450 kW min. When the slow actuator itself saturates, as the cooling water does on the shift of Figure 1.3, the VPC must stop integrating, or it winds up exactly as an override does [24]. A feedforward from the fast actuator's demand to the slow one can speed the handover [25].

Chapter 9 returns cooling from fast electric chillers to a slow absorption chiller with a VPC whose setpoint is an electric reserve of 0.5 MW, and lets the VPC's tracking stop it from winding up at a wide-open valve. Chapter 10 stacks three mid-ranging layers over a turbine bypass, the compressor speed and the firing rate, and finds that the order by cost need not match the order by speed. It also built the handover from the clamped bypass first, and the plant shed load in seconds and picked it up in hours. Chapter 11 uses mid-ranging as a limit inside a selector. The fast screw holds the face pressure, and the slow advance yields when the screw nears the top of its range.

1.4.8 Control Allocation

Split range and mid-ranging share a demand by a fixed rule. Control allocation states the rule as an optimization. A vector of virtual demands $\mathbf{w}$, such as a total force and moment, must be produced by actuators $\mathbf{u}$ through an effectiveness matrix B, within limits and at least cost [11, 26]. For one demand and quadratic prices w_i the answer has a closed form, Equation (1.10).

$$\min_{\mathbf{u}} \sum_i w_i u_i^2 \ \text{ s.t. } \sum_i b_i u_i = v \quad \Rightarrow \quad u_i = \frac{b_i / w_i}{\sum_j b_j^2 / w_j}\, v \tag{1.10}$$

When an actuator reaches a limit, it is fixed there and the remainder is shared among the others by the same formula. When every actuator is at a limit, a shortfall remains that no allocation removes.

Table 1.3 shares a cooling demand on the reactor by two rules. The priority rule is the split range of Section 1.4.6, a linear price that ranks cooling water before chilled water. The quadratic rule with prices of 1 and 4 shares the demand four to one until the cooling water reaches its limit at a demand of 500 kW. Both meet every demand up to 500 kW and leave the same 20 kW short at 520 kW. They differ in what they spend. At the design load of 250 kW the quadratic rule draws 50 kW of chilled water and would empty the store in two hours, while the priority rule draws none. A quadratic price suits actuators that wear or saturate, which should share a load. It does not suit an actuator that depletes, which should be kept for what the others cannot do.

A shortfall must reach the loops that asked for the demand, or their integrals wind up against a force that was never delivered. The same holds when demands compete for actuators that serve several of them. Weighting the residual of each virtual demand ranks them, so that what matters most is met first.

Chapter 12 allocates by fixed priority, fans before tires before thrusters, and its thruster mixer gives the attitude torques priority over the total lift. Chapter 13 solves a quadratic program every period for the forces of the planted feet, inside friction cones cut off by the leg ratings, with the pitch moment weighted thirty times the forces and the tangential pushes kept toward the middle of each cone, and feeds the delivered wrench back to its loops. Chapter 9 offers a predictive allocator that prices absorption and electric cooling, and Chapter 7 shares one electric bus among the spindle, the carriage and the mud pump.

1.4.9 Actuators That Deplete

The cold store is an actuator with a stock behind it. Every kilowatt it delivers is drawn from a charge E that the run cannot replenish. At a steady shortfall Q_{short}, the time it can still act, its runway, is given by Equation (1.11).

$$t_R = \frac{E - E_{\text{res}}}{Q_{\text{short}}} \tag{1.11}$$

where E_{res} is a reserve held back. The 100 kWh store carries the reactor's 75 kW shortfall for 80 min. On the shift of Figure 1.3 it began to drain as the cooling water saturated and was empty at 126.4 min. A store that is spent to the last kilowatt leaves nothing for the next fast disturbance. By Equation (1.9) one handover of a 100 kW step needs $\tau_h \times 100\,\text{kW} = 500\,\text{kW min}$, so a reserve of that size shortens the runway at 75 kW to 73.3 min.

A depleting actuator turns the stock into a constraint of its own. Rather than wait for the temperature override, a third input to the selector of Equation (1.3) could cut the feed as the runway falls below the time a slower remedy needs, and keep the reserve for the transients only the store can meet. The runway of Equation (1.11) is also optimistic whenever the drain grows as the stock falls, and in two chapters it does.

Chapter 8 runs a PI on the stressed runway of its cash, with a target of twelve months, whose output cuts spending, draws credit and cuts research in that order. Chapter 12 finds that its nitrogen tank runs out by pressure before mass. The gas cools as it expands, so a twelve-second hover ends in a landing at 2.83 m/s with 22.6 % of the nitrogen still in the tank. The car therefore lands when the gauge reads 8 % of the fill or, as an option, when the pressure reaches the regulator's setpoint. Chapter 14 defends a band with a vault of reserves above a floor. Below a threshold the market expects a devaluation and the drain grows by itself, and with a monetized deficit the date of the attack is set by credit growth alone, so no intervention policy moves it. Chapter 9 discharges a cold store in proportion to the chilled-water error, and Chapter 7 compares a battery with shore power.

1.4.10 What Each Chapter Protects and Shares

The section reduces to two questions that can be asked of any plant. Which limits must be protected, and which selector or clamp protects each? Which actuators can move the same variable, and by what rule, priority, mid-range or price, do they share it? Table 1.4 answers both for every chapter, together with any stock that depletes. The final chapter returns to this table with the measured result of each structure.

Three patterns run through the table. The machines of Chapters 2 to 7 have one main actuator and several limits, so their structure is a selector and their lesson is how the override takes over. The plants of Chapters 9 to 13 have more actuators than controlled variables, so their structure is a sharing rule and their lesson is what each actuator costs. The stocks of Chapter 8, Chapter 12 and Chapter 14 join the two, because the actuator that is spent becomes a limit of its own. A systematic treatment of how these simple elements combine is given in [27], and Volume 9 develops the predictive alternative that handles limits and prices in one optimization.

Table 1.4 What the Selectors of Each Chapter Protect, How Its Actuators Are Shared, and Which Stock Depletes

Ch.	Limits protected, and how	Actuators shared or limited	Stock
2	Surge margin, by one PI resting on its closed clamp	One recycle valve, by stroke time and capacity	—
3	Surge margin, turbine temperature, spool speeds, fuel rate, by six caps, then flame-out by two floors	One fuel valve for every limit	—
4	Engine power, rotor torque, speed and a preview cap, by a low selector on the ground speed	One engine for every drive	—
5	Torque, die pressure and dryer air, by overrides realized as clamps	Barrel zones split between heater bands and cooling water	—
6	Solvent vapor in each zone, by an override with three handles in order	Exhaust, then line speed or zone heat	—
7	Mud pressure against the ground's limit and torque, by a low selector on thrust	One electric bus for spindle, carriage and pump	Battery
8	Cash runway, by an override whose level is split over three ranges	Six decisions on a quarterly calendar	Cash above a zero floor
9	Exhaust temperature, by low selection on the fuel, and grid import	Electric chillers handed back to absorption by a VPC	Cold store
10	Turbine inlet metal and pressure rating	Bypass, compressor speed and firing in three mid-ranging layers	—
11	Torque, thrust, screw room and ramp, by the least of five limits on the advance	Screw holds the face, advance yields near the screw's top	—
12	Tire grip and attitude, by priority in the launch and in the thruster mixer	Fans, tires, then thrusters	Nitrogen, by pressure
13	Friction cones and leg ratings, feet that push but cannot pull	Weighted least squares among planted feet	—
14	Band edges and the reserve floor, with a capital-controls override	Desk sales and a policy rate capped by politics	Reserves

1.5 Instruments and Final Elements in This Volume

An override acts only on what its transmitter reports, and only as fast as its final element moves. Figure 1.6 lists what the controllers of each chapter read and what they move. Temperature and pressure lead the readings, with six chapters each, because most of the limits that the selectors protect are a temperature or a pressure. A motor or drive leads the final elements with eight chapters, ahead of the valve or flow loop with six.

In four chapters a pressure is the very variable the structure exists to hold: the die pressure of Chapter 5, the mud pressure against the ground in Chapter 7, the face pressure of Chapter 11 and the discharge pressure of Chapter 10. The top of Figure 1.7 shows the principle of the transmitter.

A variable-speed motor is often a better final element than a control valve, because slowing a pump in low-flow conditions saves the energy that a constant-speed pump and a throttling valve would waste [16]. Drives move the ground speed of Chapter 4, the screws of Chapter 5 and Chapter 11, the compressor of Chapter 10 and the legs of Chapter 13. The bottom of Figure 1.7 shows a drive on a pump.

The selector itself is a standard function block. In one fieldbus standard the control selector block sends a back-calculation signal to each upstream controller, so that a controller it de-selects knows at once to halt its integral action [16]. In a process plant, programmable logic controllers serve as local controllers under a distributed control system, and nested loops run with cycle times from milliseconds to minutes [28]. A jet engine's computer has two channels, each an independent computer that can control the engine completely [29], and a car executes its control in electronic control units [1]. Figure 1.8 places the loops of this volume on their platforms. The apps run from the

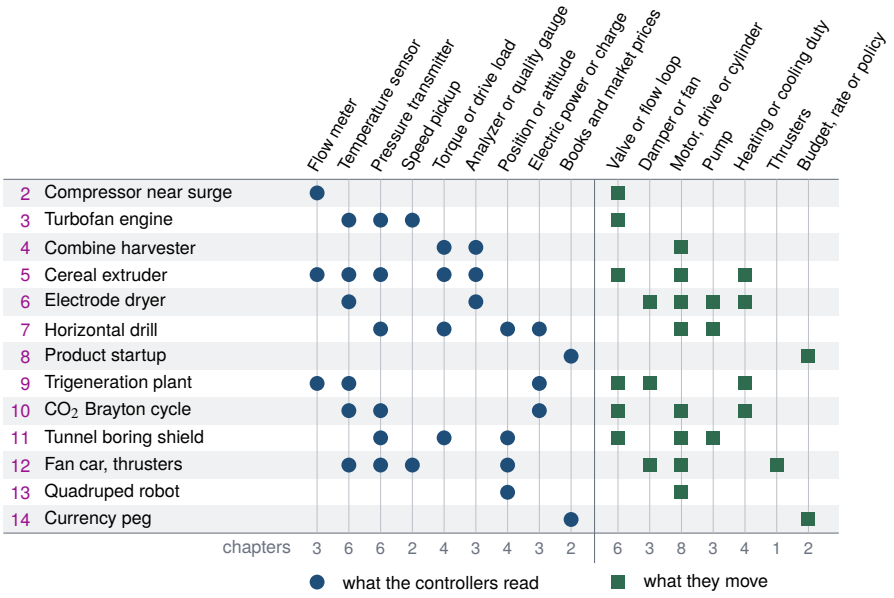

Figure 1.6 What the controllers of each chapter read and move: a round mark where a controller reads the instrument and a square mark where it moves the element, as the chapter's variables table and its technology section name them. A flow set through a flow station counts as a valve or flow loop, and a computed value is marked under the instruments it is computed from. Notice that temperature and pressure lead the readings and a motor or drive leads the final elements.

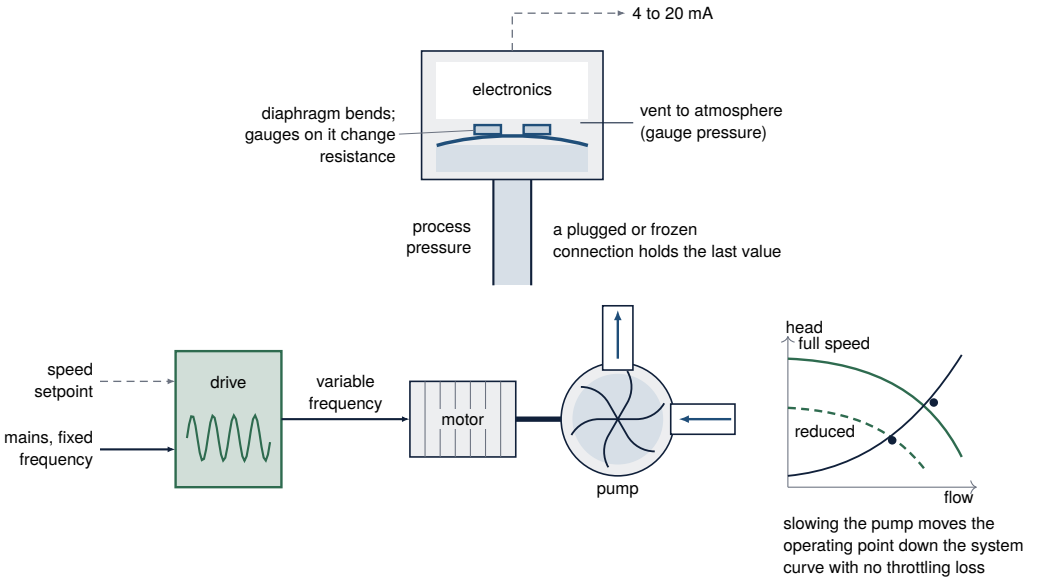

Figure 1.7 Top: principle of a gauge pressure transmitter. Bottom: a variable-speed drive on a centrifugal pump, with the pump curves at two speeds on one system curve. Notice that a plugged connection freezes the pressure signal, so an override that watches it never takes over, and that the drive moves the operating point by speed instead of by throttling.

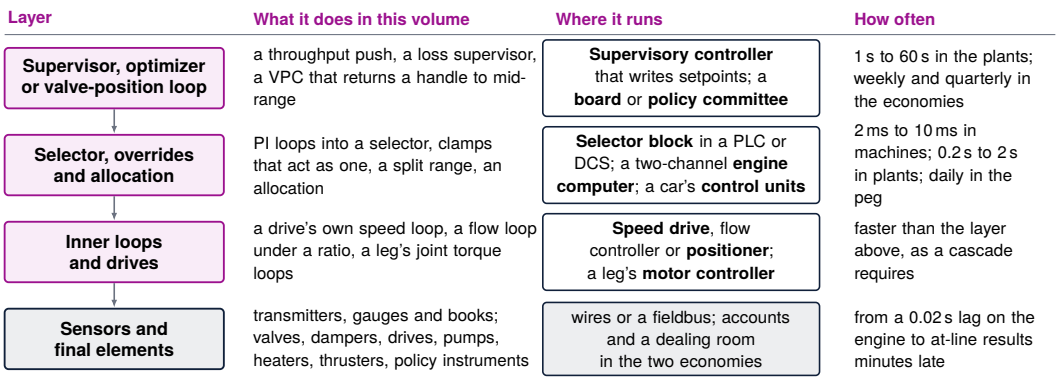

Figure 1.8 The control layers of this volume, the platform that hosts each and how often each acts. Platforms are from [16], [28], [29] and [1], and the periods from the chapters' technology sections. Notice that the selector layer runs in every chapter, from milliseconds in the machines to a day in the currency peg.

Table 1.5 What Every App Provides. Standalone Apps Are Described in Their Own Chapters.

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

2 ms period of Chapter 13 to the quarterly calendar of Chapter 8.

1.6 The Simulator Framework

Every chapter of this book has an app that runs in a web browser, collected with the others at https://apmonitor.com/apps/. About 240 apps exist across the series. All but a few are built on one shared framework, so the engine, the PID algorithm, the trends and the data export behave the same way in every chapter. The exceptions are standalone studios, such as the cascade studio and the two Temperature Control Lab apps, which follow the same ideas in their own code. This section describes the framework once and states where the standalone apps differ. Table 1.5 lists what a reader meets.

1.6.1 The Simulation Engine

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

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

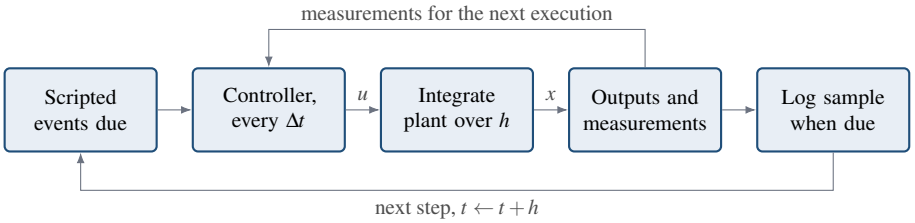

Figure 1.9 Work of the Engine in One Integration Step

executions. The log takes one sample at a fixed interval of simulated time for the trends, the scorecard and the exported file. The order of work in one step is shown in Figure 1.9. The controller also executes once when a scenario is loaded, so every run starts from an initialized controller.

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

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

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

1.6.2 The Velocity-Form PID Algorithm

One PID implementation serves the PID loops of every framework app. It is the velocity (incremental) form: each execution adds the change in the proportional and derivative terms to a running sum, and the integral is a separate term, as shown in Equations (1.12) to (1.15) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.12}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.13}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.14}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.15}$$

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 (1.15) drives I toward the

value that reproduces the output applied, which is back-calculation [30]. With no integral action, Equation (1.15) 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 $I_k = u_{\text{man}} - u_b - S_k$, 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 (1.16).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.16}$$

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 (1.17) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q\,(y_k - r_k) + \sum_{j} \Delta u_j^{\top} R\,\Delta u_j + \sum_{i} \rho_i\, s_i^2 \tag{1.17}$$

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 [31]. 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 [32]. 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 [33].

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.

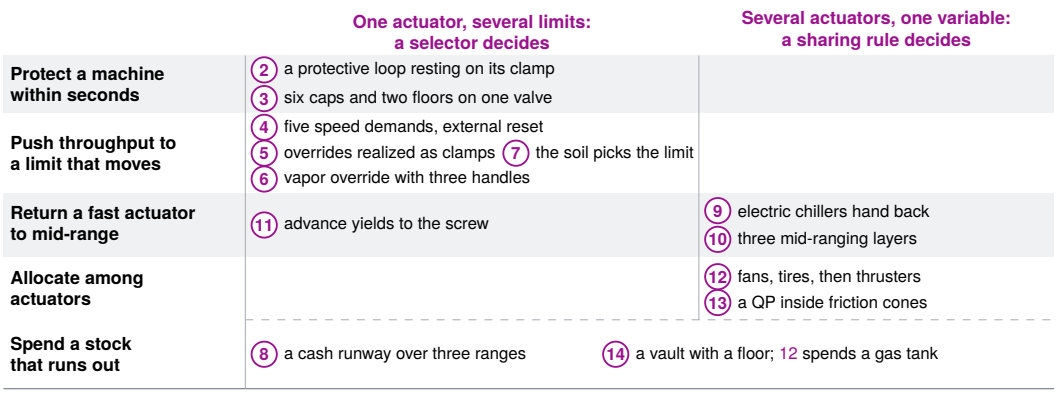

Figure 1.10 Method map of the volume. Each chapter is placed by the shape of its constraint problem, from Section 1.4.2 and Table 1.4, and by the lesson it centers on, shortened from Table 1.6. The placement is the editor's reading of the two tables. Notice that the stocks of the bottom row join the two columns, because an actuator that is spent becomes a limit of its own.

- *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 15.

1.8 Chapter Roadmap

Table 1.6 lists the chapters with the lesson each adds to this one. The level says how much of Section 1.4 a chapter assumes. The time is an estimate for reading the chapter once and running each case of its app once, at about eight minutes a case. Figure 1.10 places the same chapters by the shape of their constraint problem and by the lesson each centers on.

Chapter 2 and Chapter 4 are the first to read. The first isolates one protective loop, and the second shows a low selector with external reset among five demands. Chapter 3 then adds floors, a rate limit and the choice of tracking time. Chapters 5 to 7 push throughput against limits that move and can be read in any order once Sections 1.4.2 to 1.4.5 are familiar. For sharing, Chapter 9 is the gentler introduction to mid-ranging and Chapter 10 the fuller one, with Chapter 11 as a variant inside a selector. Chapter 12 and then Chapter 13 carry allocation from a fixed priority to an optimization. Chapter 8 and Chapter 14 form a path for readers from economics and finance, where the depleting stock is money.

Table 1.6 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Centrifugal compressor near surge	A protective loop resting on its clamp, and a valve too slow for a correct decision.	introductory	2.5 h
3	Two-spool turbofan engine	Six caps and two floors on one fuel valve, and what the tracking time decides.	intermediate	2.5 h
4	Combine harvester in a patchy field	Five speed demands in a low selector with external reset, under a loss supervisor.	introductory	2.5 h
5	Cereal extrusion line	Overrides realized as clamps, split-range zones and a throughput supervisor.	advanced	2.5 h
6	Battery electrode coater and dryer	An explosive-limit override with three handles, and a push to the first binding limit.	advanced	2.5 h
7	Horizontal directional drill	A rate optimizer whose binding limit changes with the soil, and pressure overrides.	advanced	2 h
8	Product startup company	A runway override over three ranges, on a quarterly calendar.	intermediate	2.5 h
9	Gas turbine trigeneration plant	An exhaust limiter on the fuel, and a VPC that returns cooling to the cheap chiller.	intermediate	2.5 h
10	Supercritical CO_2 Brayton cycle	Three mid-ranging layers, and why the handover reads the unclamped demand.	advanced	3 h
11	Earth pressure balance shield	An inventory held through forced stops, with a mid-ranging limit on the advance.	intermediate	2.5 h
12	Car with fans and cold-gas thrusters	Allocation by priority on a tank that runs out by pressure.	intermediate	2.5 h
13	Planar quadruped robot	Allocation by a quadratic program inside friction cones that switch with each step.	advanced	2.5 h
14	Central bank defending a peg	A vault with a floor, a drain that grows as it empties, and too few instruments.	intermediate	2.5 h

KEY POINTS

- When several limits share one actuator, a selector passes the most cautious demand, and the guard of a limit takes over only as that limit comes near.
- External reset feeds each integral from the output actually applied, so with a tracking time equal to the integral time the override takes over at its setpoint, without a bump.
- Split range moves one actuator at a time in a fixed order. Mid-ranging lets a fast actuator answer first and hands the load to a slow one.
- An actuator that draws on a stock is also a limit. Its runway is the stock above a reserve divided by the shortfall.

Where else this applies. A low selector with external reset on a treadmill belt in Volume 19, Chapter 6, a four-way low selector on a fast charger in Volume 17, Chapter 9, and predictive control of limits in Volume 9.

References

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

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

[3] F. G. Shinskey and J. P. Shunta, “Selective, override, and limit controls,” in *Process Control*. Elsevier, 1995, pp. 112–118.

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

[5] M. Pitt, “Taking a look back at control: Part 2,” *The Chemical Engineer*, no. 989, Nov. 2023, thechemicalengineer.com/features/taking-a-look-back-at-control-part-2.

[6] ——, "Taking a look back at control: Part 1," *The Chemical Engineer*, no. 988, pp. 51–54, Oct. 2023, thechemicalengineer.com/features/taking-a-look-back-at-control-part-1.

[7] L. C. Jaw and S. Garg, "Propulsion control technology development in the United States: A historical perspective," NASA Glenn Research Center, Cleveland, OH, Technical Memorandum NASA/TM-2005-213978, 2005, nTRS 20050238468.

[8] 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.

[9] A. H. Glattfelder, W. Schaufelberger, and H. P. Fässler, "Stability of override control systems," *International Journal of Control*, vol. 37, no. 5, pp. 1023–1037, 1983.

[10] B. Eichengreen and C. Wyplosz, "The unstable EMS," *Brookings Papers on Economic Activity*, vol. 1993, no. 1, pp. 51–144, 1993.

[11] W. C. Durham, "Constrained control allocation," *Journal of Guidance, Control, and Dynamics*, vol. 16, no. 4, pp. 717–725, 1993.

[12] M. V. Kothare, P. J. Campo, M. Morari, and C. N. Nett, "A unified framework for the study of anti-windup designs," *Automatica*, vol. 30, no. 12, pp. 1869–1883, 1994.

[13] B. J. Allison and A. J. Isaksson, "Design and performance of mid-ranging controllers," *Journal of Process Control*, vol. 8, no. 5–6, pp. 469–474, 1998.

[14] S. Hyon, J. G. Hale, and G. Cheng, "Full-body compliant human–humanoid interaction: Balancing in the presence of unknown external forces," *IEEE Transactions on Robotics*, vol. 23, no. 5, pp. 884–898, 2007.

[15] A. Moisseytsev and J. J. Sienicki, "Development of the ANL plant dynamics code and control strategies for the supercritical carbon dioxide Brayton cycle and code validation with data from the Sandia small-scale supercritical carbon dioxide Brayton cycle test loop," Argonne National Laboratory, Argonne, IL, USA, Tech. Rep. ANL-ARC-218, 2011.

[16] 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.

[17] D. Krishnamoorthy and S. Skogestad, "Systematic design of active constraint switching using selectors," *Computers & Chemical Engineering*, vol. 143, p. 107106, 2020.

[18] A. Reyes-Lúa and S. Skogestad, "Systematic design of active constraint switching using classical advanced control structures," *Industrial & Engineering Chemistry Research*, vol. 59, no. 6, pp. 2229–2241, 2020.

[19] C. Loeblein and J. D. Perkins, "Structural design for on-line process optimization: I. dynamic economics of MPC," *AIChE Journal*, vol. 45, no. 5, pp. 1018–1029, 1999.

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

[21] A. Reyes-Lúa, C. Zotică, K. Forsman, and S. Skogestad, "Systematic design of split range controllers," *IFAC-PapersOnLine*, vol. 52, no. 1, pp. 898–903, 2019.

[22] A. Reyes-Lúa and S. Skogestad, "Multi-input single-output control for extending the operating range: Generalized split range control using the baton strategy," *Journal of Process Control*, vol. 91, pp. 1–11, 2020.

[23] B. J. Allison and S. Ogawa, "Design and tuning of valve position controllers with industrial applications," *Transactions of the Institute of Measurement and Control*, vol. 25, no. 1, pp. 3–16, 2003.

[24] S. Haugwitz, M. Karlsson, S. Velut, and P. Hagander, "Anti-windup in mid-ranging control," in *Proceedings of the 44th IEEE Conference on Decision and Control*. IEEE, 2005, pp. 7570–7575.

[25] T. Hägglund, "A feedforward approach to mid-ranging control," *Control Engineering Practice*, vol. 108, p. 104713, 2021.

[26] T. A. Johansen and T. I. Fossen, "Control allocation—a survey," *Automatica*, vol. 49, no. 5, pp. 1087–1103, 2013.

[27] S. Skogestad, "Advanced control using decomposition and simple elements," *Annual Reviews in Control*, vol. 56, p. 100903, 2023.

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

[29] Federal Aviation Administration, *Aviation Maintenance Technician Handbook–Powerplant*. U.S. Department of Transportation, Federal Aviation Administration, 2023, no. FAA-H-8083-32B.

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

[31] 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.

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

[33] 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

The reactor of Section 1.4.1 runs with the cooling water at full stroke and the cold store unused. TIC-1 has $K_c = 2.5\,\%/°\mathrm{C}$ and $\tau_I = 10\,\mathrm{min}$ on a split range that gives 5 kW per percent of its output, and each kg/min of feed releases 25 kW.

(a) Find the steady feed the reactor can carry if the override setpoint is raised from 85 to 87 °C, and the gain over the feed at 85 °C. *Checkpoint:* 17.4 kg/min, 0.4 kg/min more.

(b) TIC-1 holds 80 °C. Find the fastest feed ramp, in kg/min per minute, whose steady ramp lag keeps the temperature at least 2 °C under an override setpoint of 85 °C. *Checkpoint:* 0.15 kg/min per minute.

1.2 Concept check · 2 min

A gas turbine's fuel demand passes several caps, among them a temperature limiter and a rate limit, and a minimum fuel flow keeps the flame lit. Which arrangement keeps the flame lit even when a cap asks for less than the minimum fuel?

(a) One low selector over the caps and the minimum fuel together

(b) A high selector with the minimum fuel, applied after a low selector over the caps

(c) A low selector over the caps, applied after a high selector with the minimum fuel

(d) The average of all caps and the minimum fuel, sent to the valve

1.3 Theory and calculation · 3 min

The override TC-2 of the reactor has $K_c = 2\,\mathrm{kg/min}$ per °C, setpoint 85 °C, external reset with $T_t = \tau_I$, and the selector applies a feed of 19 kg/min. Find the override's demand once its integral has settled, at measured temperatures of 83 and 84.5 °C, and the temperature at which it takes over.

1.4 Theory and calculation · 4 min

A split-range controller moves a cooling-water valve worth 6 kW per percent of stroke (600 kW at full stroke) and then a chilled-water valve worth 2 kW per percent (200 kW). Find the split point that gives the same duty per percent of controller output on both segments, and that duty. Then find the duty per percent on each segment for a split at 50 %, and the ratio of the two loop gains.

1.5 Concept check · 2 min

A valve-position controller returns a fast bypass valve to a resting opening of 5 %. The bypass is clamped at zero. When the plant falls short of its target, the fast controller asks for a negative opening. What happens if the VPC watches the clamped bypass position instead of the fast controller's unclamped demand?

(a) The VPC moves faster, because the clamped position is less noisy

(b) Nothing changes, because the two signals are equal in steady state

(c) The VPC sees the same error for a small and a large shortfall and moves at one slow rate

(d) The VPC winds up and drives the slow actuator to its opposite limit

1.6 Theory and calculation · 4 min

A reactor uses mid-ranging with chilled water as the fast actuator, 1 kW per percent, and cooling water as the slow one, 5 kW per percent. The VPC moves the cooling water at $K_I = 0.05\,\mathrm{min}^{-1}$ times the chilled-water demand less its resting position.

(a) Find the static gain the VPC sees and its closed-loop time constant. *Checkpoint:* 5, and 4 min.

(b) A load step of 30 kW arrives that the cooling water can carry. Find how far the cooling water moves and the chilled-water energy drawn above the resting draw during the handover. *Checkpoint:* 6 %, and 120 kW min (2 kWh).

1.7 Concept check · 2 min

An override's integral is corrected by back-calculation from the selector's output, with a tracking time T_t of one tenth of its integral time instead of the integral time itself. Compared with $T_t = \tau_I$, what does the override do as its constraint approaches the limit?

(a) It takes over earlier, on the rate of approach, before the limit is reached

(b) It takes over later, after the constraint has passed its limit

(c) It takes over at the same moment, because the tracking time only affects noise

(d) It never takes over, because its integral cannot leave the applied output

1.8 Theory and calculation · 4 min

A cooling demand is shared between cooling water (at most 400 kW) and chilled water (at most 100 kW) by Equation (1.10) with prices of 1 and 9 and unit effectiveness. Find the two duties for demands of 300 and 450 kW, and compare each with the priority rule that uses cooling water first.

1.9 Concept check · 2 min

A protective PI controller rests on its lower output clamp while the plant is healthy, with back-calculation at $T_t = \tau_I$ and an error that stays negative. What does its integral term do while it rests there?

(a) It integrates the negative error without bound

(b) It stays frozen at the value it had when the clamp was reached

(c) It is reset to zero at every execution

(d) It decays toward the clamp value with time constant τ_I, whatever the error

1.10 Theory and calculation · 4 min

A cold store holds 100 kWh (6000 kW min).

(a) The store must cover a steady shortfall of 60 kW. Find how long it lasts. *Checkpoint:* 100 min.

(b) The store keeps a reserve of 500 kW min for one handover. Find the runway at the same shortfall. *Checkpoint:* 91.7 min.

(c) With no shortfall, a VPC rests the chilled-water valve at 5 %, worth 5 kW. Find how long the resting draw alone takes to empty the full store. *Checkpoint:* 20 h.

CHAPTER 2

Teaching Limit Cycles and Protective Override with a Compressor Surge Simulator

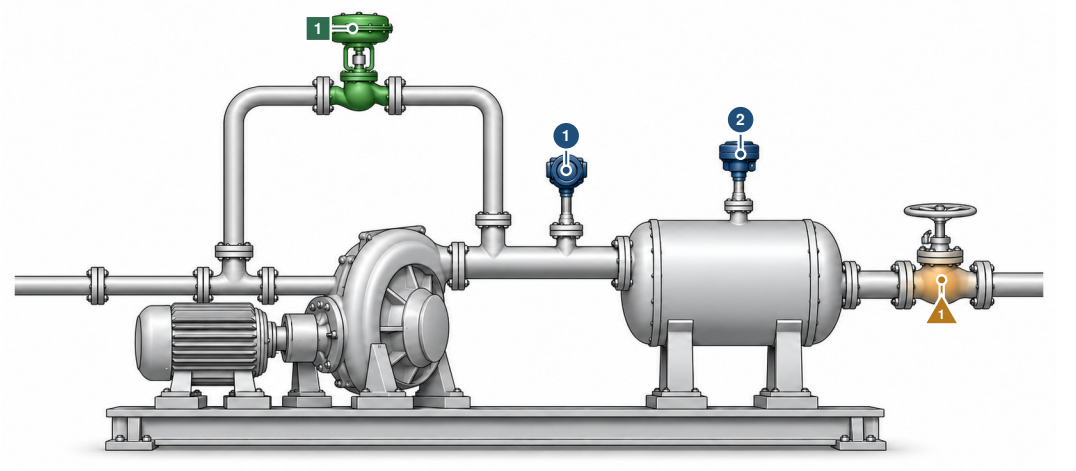

The controller moves 1 valve that sends gas back to the inlet

The controller reads 1 flow from the compressor 2 pressure in the tank, for the surge alarm

What disturbs it 1 throttle closed by the users

Notice that the only path the controller can open is the overhead line back to the compressor inlet, while the users' throttle on the right is not its to move. Illustration, not to scale.

IN PLAIN TERMS Keeping a Compressor Out of Surge

The main idea. A compressor pushes gas into a tank and on to the users. If the users take too little gas, the flow through the machine can collapse and run backward, again and again, in a violent cycle called surge.

The challenge. A protective controller opens a valve that sends gas back to the compressor inlet, but only when the flow nears the danger line. If that valve is too slow or too small, no setting of the controller can save the machine.

Why it matters. One surge can damage the bearings and blades, yet gas sent back around the machine wastes the power used to compress it.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on what protection costs (Section 2.1), the two history segments (Section 2.2), the two try-it segments and the fact on valve speed (Section 2.7) and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 2.7), Technology in Practice (Section 2.8), Limitations (Section 2.9) and the two open challenges

The control question A centrifugal compressor throttled past the peak of its characteristic does not diverge but settles into a deep-surge limit cycle that reverses the flow through the rotor. How does a protective override whose valve rests shut keep the machine off that orbit, and when do stroke speed, valve capacity or noise defeat it whatever the tuning?

Learning objectives

- Derive the trace and determinant conditions and locate the static and dynamic surge boundaries.
- Tell a limit cycle from a divergence, and an aerodynamic surge cycle from a control-induced one, by experiment.
- Explain why the integral of an override resting on its clamp decays instead of winding up.
- Judge a protective loop by surge margin and recycled flow, not by IAE.
- Decide whether a failed closure is a tuning, a stroke-speed or a capacity problem, and what a throttle signal adds.

Prerequisites Sections 1.4.3 (override), 1.4.4 (external reset), 1.4.5 (saturation and rate limits), 1.6.2 (velocity form) and 1.6.4 (scenarios and exported files). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/compressor-surge/

See also Volume 7, Chapter 2 for a Hopf bifurcation and a loop left cycling, and Volume 7, Chapter 4 for a setpoint at the peak of a characteristic. Volume 10, Chapter 11 for override anti-surge loops on parallel machines that share one load. In this volume, Chapter 3 for a compressor protected without a measured surge margin.

2.1 Introduction

Compressor surge is a control problem in which the regulated steady state itself disappears. As the downstream valve closes, the operating point moves toward lower flow along the machine's pressure-rise curve, and near the peak of that curve the steady state the loop is regulating ceases to be attracting. What replaces it is not a new steady state and not an exponential divergence. It is a closed orbit. Flow collapses, the discharge volume blows down backward through a rotor still turning at full speed, flow re-establishes, the volume refills, and the sequence repeats at a period set by the acoustics of the pipework. In the simulator of this chapter that period is 7.5 nondimensional time units and the flow swings from -0.175 to $+0.711$, and the oscillation persists indefinitely because the amplitude is a property of the system rather than of the disturbance that started it.

The mechanism follows from one momentum balance. The compressor accelerates the gas in its duct according to the difference between the pressure rise it can produce at the flow currently passing and the pressure already stored downstream. Where the characteristic slopes down, a dip in flow moves the machine to a part of the curve that delivers *more* pressure rise than the plenum holds, and the flow is pushed back up. Where the characteristic slopes up, the same dip delivers *less*, and the flow falls further. The sign of one derivative determines whether the aerodynamics contribute negative feedback or positive feedback to the compression system. Volume 7, Chapter 4 meets the same reversal of the process gain at the peak of a tire's friction curve.

The consequence for control is unusual. The correct response to an approaching surge line is to open a recycle valve, and the correct response at every other time is to leave that valve shut, because recycled gas is compressor power converted directly into heat. An anti-surge controller therefore spends almost its entire service life at the closed end of its output range and is judged on an event that may occur once a year. That is an override of the kind a process control text treats alongside selectors and constraint loops [1] (Section 1.4.3), not a regulator, and it changes what the usual performance measures mean. Integral of absolute error accumulates distance to a setpoint the loop has no intention of reaching. The fraction of time spent saturated is meaningless at one end of travel and critical at the other. The numbers that matter are the smallest surge margin reached and the quantity of gas recycled.

The model behind the application is built from two conservation statements, a momentum balance on the gas in the duct and a mass balance on the plenum, closed by a static compressor characteristic and two valve laws [2]. The experiments built on it address process dynamics and control principles that a course built on linear plants reaches late or not at all: nonlinear stability and limit cycles, protective override structure, integral anti-windup, actuator rate and capacity limits, and the effect of measurement noise and filtering on a loop with no lag to spare [1, 3]. The behavior matters industrially because a single deep-surge event can damage bearings and blading, while a conservative surge control line recycles gas continuously, so both the protection and its operating cost are plant-level concerns [4].

> **DID YOU KNOW? Protection Paid in Megawatts**
>
> An instrumentation textbook calls surge a violent and potentially self-destructing action of a centrifugal compressor [5]. Avoiding it has a price. On the refrigerant compressor of Volume 10, Chapter 14, turned down to 20 kg/s, the margin costs 10.6 % recycle valve opening and 1.8 MW of shaft power.

This chapter gives every equation, parameter and measured number of the simulated compression system, so that its four states, its characteristic, its two valves and its override can be reimplemented from the text, and it separates three effects that one surge event presents together. The analysis yields a pair of stability boundaries computed from the implemented equations, a geometric one at the peak of the characteristic and a dynamic one two percentage points of throttle travel inside it. At the default parameters the classical static criterion never fails at any throttle position, so the Greitzer parameter B decides the outcome by itself. The model has stroke-limited valves, a windowed surge detector and twelve scenarios. The override's lower clamp and back-calculation anti-windup are given as implemented rather than as intended, and the measured advantage of the predictive controller is divided between its prediction and the throttle signal it reads. The result is a setting where a limit cycle can be measured instead of described, and where two oscillations that look alike on a trend can be distinguished by experiment.

Section 2.2 traces how surge protection reached its present form. Section 2.3 sets out the compression system, its instruments and its assumptions, and Section 2.4 derives the two boundaries and the role of B. Section 2.5 gives the override as implemented and, in Section 2.5.1, the predictive controller. Section 2.6 explains which performance measures mean anything on a protective loop. Section 2.7 reports the twelve scenarios: the staged closure, the deep-surge limit cycle, valve stroke speed and capacity, a control-induced oscillation, measurement noise, the predictive controller and the challenge. Section 2.8 sets the loop on plant hardware, Sections 2.9 and 2.10 state the limitations and conclude, and Section 2.11 holds the exercises.

2.2 Background

Surge was documented as mechanical damage long before it was described as a system dynamic. Industrial blowers and early aircraft engine compressors were reported to hammer, to reverse, and to shed blading when the discharge was restricted, and the practical remedy adopted throughout the first half of the twentieth century was a blow-off valve set by experience to keep the machine away from the region where the damage started. The phenomenon acquired a physical description once compressor instabilities were separated into two distinct objects. A circumferential cell of stalled passages that travels around the annulus at a fraction of rotor speed leaves the annulus-averaged flow steady. A one-dimensional oscillation of the whole compression system moves the annulus-averaged flow and depends on what is bolted downstream. The first is rotating stall and the second is surge.

The decisive step for control was the recognition that the second object is a property of the system rather than of the machine, as the history segment below tells. A smooth cubic representation of the compressor characteristic, valid on both sides of the peak and continued into reverse flow, arrived

with the post-stall transient theory of Moore and Greitzer [6], and it is that cubic which the present application uses.

> ⧖ HISTORY **The Installation Decides Between Stall and Surge, 1976**
>
> A two-part study published in 1976 posed a nonlinear lumped model of a compression system to predict its transient response after a disturbance. It found one nondimensional parameter on which that response depends. Above a critical value the system falls into the large oscillation of surge, and below it the machine moves toward rotating stall at a much reduced flow and pressure ratio [7]. The second part tested the model on a three-stage axial compressor, computed the instantaneous flow from a mass balance on the plenum, and found the predicted criterion in good accord with the data [8]. The 1980 Freeman Scholar Lecture then reviewed the instabilities of pumping systems of many kinds and found the idea of negative damping common to them [9]. That parameter is the B of this chapter.

Industrial practice and research followed separate paths. Plant protection settled on a recycle line around the machine, a fast valve in it, and a control line drawn a fixed margin to the right of the measured surge limit line in a corrected-flow against corrected-head plane, with the valve opened by a controller that acts only when the operating point crosses that line. A technology assessment of centrifugal compressor surge control drew together the sensing options, the valve requirements and the field experience of that arrangement [10], and the practical consequences for gas transmission service, where a station may be tripped or recycled many times a day, are set out with the associated economics [4]. Volume 10, Chapter 11 places three such machines in parallel on one pair of headers, each with its own override anti-surge loop. Research instead asked whether feedback could move the surge line rather than merely respecting it, the line of work the second history segment follows. Active stabilization has not displaced the recycle valve in production plants, so the loop found in industrial service is still the override, and it is the override that this application implements.

> ⧖ HISTORY **Feedback Aimed at the Surge Line Itself, 1989 to 1999**
>
> A strategy proposed in 1989 drove an extra disturbance in a turbomachine from data measured inside it in real time, so that the controlled machine behaved differently from the bare one. For the compressor analyzed, the control widened the stable operating range by 20 percent of the mean flow, at a control power that can be several orders of magnitude below the machine's [11]. In 1993 a low-speed axial research compressor delayed rotating stall with a ring of hot wires and individually driven inlet guide vanes, and lowered its stalling flow by 23 percent [12]. A study of sensor and actuator choices for surge found the actuators most closely coupled to the compressor the most effective [13], and in 1999 a close-coupled valve combined with speed control was given a proof of stability [14].

2.3 Compression System Model

The application simulates a compressor discharging through a short duct into a plenum, with a throttle taking gas out of the plenum to the process and a recycle line returning gas from the plenum to the compressor inlet through the anti-surge valve. The layout and the symbol for each quantity in the equations below are given in Figure 2.1. Nothing in the model is dimensional. Flow and pressure rise are the usual compressor coefficients and time is scaled by the Helmholtz period of the duct and plenum pair, so the trends and the exported time column both read $\hat{t}$.

2.3.1 Variable Classification

The one controlled variable, the manipulated variable, the disturbances, the four states and the derived outputs worth watching are collected in Table 2.1. The throttle demand requires one clarification. It is an operator-set input that stays active in both modes, but it is not a controlled variable and it is not the controller's actuator. It is the load, and it is the disturbance that the manual experiments of Section 2.7 drive directly.

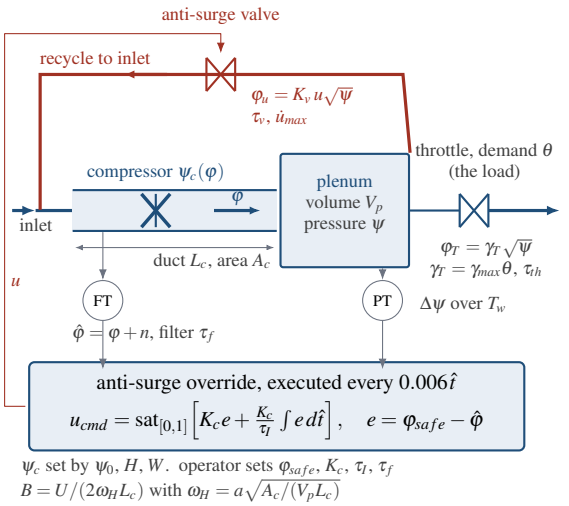

Figure 2.1 Compressor, duct, plenum, throttle and the anti-surge recycle line. Flow is measured on the duct and is the only signal the override uses. Plenum pressure feeds the surge detector and never reaches the controller. The predictive controller of Section 2.5.1 also reads the throttle demand θ.

Table 2.1 Classification of Every Signal in the Model

Symbol	Role	Range	Unit
φ	CV, state	−0.2 to 0.80	—
$\hat{\varphi}$	measured CV	$\varphi + n$	—
φ_{safe}	setpoint	0.500–0.720 (0.575)	—
u_{cmd}	MV, controller output	0–100	%
u	state, valve position	0–1	—
ψ	state, plenum pressure rise	0–0.72	—
γ_T	state, throttle conductance	0–1.40	—
θ	DV, throttle demand	0–100 (70)	%
σ_n	DV, flow transmitter 1σ	0–0.040 (0)	—
$\Delta\theta_b$	DV, restriction drift	−25 to +5 (0)	%
$\psi_c(\varphi)$	output, characteristic	—	—
ψ_c'	output, slope	−3.5 to +1.08	—
φ_T	output, throttle flow	—	—
φ_u	output, recycled flow	—	—
SM	output, surge margin	—	%
$\Delta\psi$	output, detector swing	—	—

2.3.2 Aerodynamic States

The gas in the duct between the compressor and the plenum is treated as a slug of fixed inertia, and Newton's law applied to that slug gives the first state equation, as shown in Equation (2.1).

$$\frac{\mathrm{d}\varphi}{\mathrm{d}\hat{t}} = B\left[\psi_c(\varphi) - \psi\right] \tag{2.1}$$

where φ is the nondimensional flow through the compressor, ψ is the nondimensional pressure rise held in the plenum, $\psi_c(\varphi)$ is the pressure rise the compressor produces at the flow currently passing, and B is the Greitzer parameter. The right-hand side is a difference of two pressures and the equation is a force balance, so the acceleration of the flow is zero whenever the machine and the plenum agree, whatever value they agree on. Nothing in Equation (2.1) restricts φ to positive values, and

the reverse-flow excursion reported later is a solution of this equation rather than an artifact of the numerics.

Conservation of mass for the plenum gives the second state equation, as shown in Equation (2.2).

$$\frac{\mathrm{d}\psi}{\mathrm{d}\hat{t}} = \frac{1}{B}\left[\varphi - \varphi_T(\psi) - \varphi_u(\psi)\right] \tag{2.2}$$

where φ_T is the flow the throttle passes to the process and φ_u is the flow the anti-surge valve recycles to the inlet. The plenum is a capacity and Equation (2.2) is an inventory balance, so the pressure rises when the compressor delivers more than the two valves remove. Opening the recycle valve therefore acts on both states at once. It lowers ψ, which raises the imbalance driving Equation (2.1), and it adds a second outlet, which raises the flow the compressor has to pass.

The compressor characteristic is the Moore-Greitzer cubic, given by Equation (2.3).

$$\psi_c(\varphi) = \psi_0 + H\left[1 + \tfrac{3}{2}\left(\frac{\varphi}{W} - 1\right) - \tfrac{1}{2}\left(\frac{\varphi}{W} - 1\right)^3\right] \tag{2.3}$$

where ψ_0 is the shut-off pressure rise, H is the semi-height of the curve and W is its semi-width. Differentiating gives $\psi_c' = (3H/2W)\left[1 - (\varphi/W - 1)^2\right]$, which vanishes at $\varphi = 0$ and at $\varphi = 2W$, is positive between them and negative beyond. The peak of the curve is therefore at exactly $\varphi = 2W = 0.500$ with $\psi = \psi_0 + 2H = 0.560$, and that peak is the conventional surge line. The cubic is used rather than a fitted rig curve because it is smooth, it is bounded, and its continuation to $\varphi < 0$ turns upward again, which is the feature that closes the deep-surge orbit instead of letting the solution run away.

Both valves are square-root devices with no downstream pressure, as shown in Equation (2.4).

$$\varphi_T = \gamma_T\sqrt{\max(\psi, 0)}, \qquad \varphi_u = K_v\, u\sqrt{\max(\psi, 0)} \tag{2.4}$$

where γ_T is the conductance of the throttle actually in place, K_v is the capacity of the anti-surge valve wide open and u is its position on 0 to 1. The max keeps the square root real and prevents either valve from passing flow backward during the part of the surge cycle in which the plenum pressure falls below zero. The slope of the throttle line in the pressure-against-flow plane follows from the same expression as $\psi_T' = 2\sqrt{\psi}/\gamma_T$, and that slope appears in the stability conditions of Section 2.4.

2.3.3 Actuator States

Neither valve moves instantly. Both are first order on their commands, as shown in Equation (2.5).

$$\tau_v\frac{\mathrm{d}u}{\mathrm{d}\hat{t}} = u_{app} - u, \qquad \tau_{th}\frac{\mathrm{d}\gamma_T}{\mathrm{d}\hat{t}} = \gamma_{cmd} - \gamma_T \tag{2.5}$$

where u_{app} is the controller output after rate limiting, τ_v is the anti-surge valve time constant, $\gamma_{cmd} = \gamma_{max}\,\mathrm{clamp}(\theta + \Delta\theta_b/100, 0, 1)$ is the commanded throttle conductance including the unmeasured restriction drift, and τ_{th} is the throttle stroke time constant. After each integration step u is clamped to $[0, 1]$ and γ_T is floored at zero. Rate limiting is separate from τ_v and is described in Section 2.5, because it is applied inside the controller as well as on the applied position.

2.3.4 Instrumentation and Surge Detection

Flow is the only measurement the controller receives, and Equation (2.6) is what reaches it.

$$\hat{\varphi} = \mathscr{F}_{\tau_f}\left[\varphi + \sigma_n\,\xi(\hat{t})\right] \tag{2.6}$$

where ξ is unit-variance noise passed through a first-order coloring filter of time constant $0.02\,\hat{t}$, σ_n is the operator-set transmitter standard deviation, and $\mathscr{F}_{\tau_f}$ is an optional first-order measurement filter that is switched off by default. There is no transmitter bias and no transport delay. The noise sequence is generated by a seeded mulberry32 stream, so a scenario with noise replays identically.

Table 2.2 Parameter Defaults and Adjustable Ranges

Symbol	Meaning	Value	Range	Unit
B	Greitzer parameter	1.50	0.30–3.00	—
H	characteristic semi-height	0.180	0.080–0.300	—
W	characteristic semi-width	0.250	0.150–0.400	—
ψ_0	shut-off pressure rise	0.20	0.00–0.40	—
γ_{max}	throttle wide open	1.40	0.80–2.20	—
τ_{th}	throttle stroke constant	0.50	0.05–4.00	$\hat{t}$
K_v	anti-surge valve capacity	0.60	0.05–1.40	—
τ_v	anti-surge valve constant	0.30	0.02–3.00	$\hat{t}$
$\dot{u}_{max}$	valve stroking rate limit	4.00	0.10–20.0	$1/\hat{t}$
Fixed detector settings				
T_w	detector window	8.0	2–20	$\hat{t}$
$\Delta\psi_{on}$	surge alarm threshold	0.25	0.05–0.60	—
$\Delta\psi_{off}$	alarm clear threshold	0.12	0.02–0.40	—
Derived from the parameters above				
$2W$	surge line flow	0.500	—	—
$\psi_0 + 2H$	pressure rise at the peak	0.560	—	—
φ_z	flow at zero pressure rise	0.8037	—	—

Plenum pressure is instrumented separately and never reaches the controller. It feeds a detector that holds the last $T_w = 8.0\ \hat{t}$ of ψ in a ring buffer sampled on a fixed $0.02\ \hat{t}$ grid and reports the peak-to-peak swing $\Delta\psi$ over that window. The surge alarm latches when $\Delta\psi$ exceeds $\Delta\psi_{on} = 0.25$ or when φ goes negative, and clears when $\Delta\psi$ falls below $\Delta\psi_{off} = 0.12$ with the flow positive. The window is deliberately about one surge period long so that the single transient following an ordinary throttle move does not trip it.

2.3.5 Parameter Values, Start-Up State and Numerics

Every parameter, its default and its adjustable range are given in Table 2.2. The three detector settings are excluded from the adjustable set because they are instrument configuration rather than controller tuning.

Every run begins at an exact steady state. The initializer takes the total downstream conductance $\gamma = \gamma_T + K_v u_0$, with u_0 the operator's valve position in MANUAL and zero in AUTO, and solves $\varphi = \gamma\sqrt{\psi_c(\varphi)}$ by bisection. The bisection interval is split at $\varphi = 2W$ and the branch is chosen by the sign of the residual at the peak, so a requested operating point on the positive-slope side is placed there deliberately rather than by accident. The states are then $[\varphi^*,\ \psi_c(\varphi^*),\ u_0,\ \gamma_T]$, which makes every derivative zero to within 1×10^{-9}, so any response after the first event is attributable to the model rather than to an initialization transient.

The engine and its three clocks are described in Section 1.6.1. Here the states advance by fourth-order Runge-Kutta at the fixed step $\Delta\hat{t} = 0.002$, and the controller executes every $0.006\ \hat{t}$, three physics steps, after a priming execution at $\hat{t} = 0$. The value 0.006 is also the Δt handed to the integral term and the rate limiter. The throttle demand reaches the model through the controller as well, so a throttle move takes effect at the next execution (a move scheduled at $\hat{t} = 8$ reaches the throttle actuator at $\hat{t} = 8.004$). The historian records 50 samples per unit of $\hat{t}$, and the exported file puts those samples in comma-separated rows under a time column named `time_hat`. The default seed is 20260824 and the application starts in AUTO, so every run reported below is reproducible.

2.3.6 Assumptions

The model is a lumped two-state representation of the aerodynamics with two actuator states attached, and each simplification has a regime in which it fails.

One flow and one pressure. The compressor is one duct with one inertia, so a circumferential disturbance pattern cannot be represented at all and rotating stall is absent. Real machines frequently enter rotating stall first and surge afterwards, and a compressor that recovers into a stalled branch rather than into surge exhibits behavior this model cannot represent.

Quasi-steady characteristic. The machine is assumed to deliver its steady-state pressure rise instantly at whatever flow is passing, including during a reversal. The reverse-flow branch of a real characteristic is measured with difficulty on a rig and is only approximately the continuation of the cubic, so the deep-surge orbit reported here has the correct qualitative shape, but its exact amplitude is a property of the cubic.

One speed line. Shaft speed, inlet temperature and inlet pressure are constant, so there is one characteristic and the surge line is fixed at $\varphi = 2W$. An industrial anti-surge controller works in a corrected-flow against corrected-head plane precisely because none of that holds, and the surge control line it uses is a curve fitted across a family of speed lines. Chapter 3 protects a jet engine compressor on corrected schedules of this kind, with no measurement of the surge margin at all.

Incompressible valve law. Both valves obey $\varphi = \gamma\sqrt{\psi}$ with no choking, no downstream pressure and no dead time. A recycle valve on a high-pressure machine is choked for much of its travel, which flattens its installed characteristic, and the transmitter and signal path ahead of it add several milliseconds of delay that this model does not represent.

Cold recycle. Gas returned through the anti-surge valve arrives at the inlet at the conditions it left. A real recycle line returns hot gas, so inlet temperature climbs while the valve is open, the characteristic moves, and a cooler is required. This omission understates the cost of extended operation with the recycle valve open.

Detection by pressure swing. Surge is declared from a windowed peak-to-peak of ψ, which is a plausible detector but a slow one. It is set not to trip on a single transient, an appropriate trade for instruction but not for a machine that must be protected within the first surge cycle.

2.4 Stability Analysis

With the anti-surge valve shut, Equations (2.1) and (2.2) linearized about an equilibrium (φ^*, ψ^*) give the Jacobian of Equation (2.7).

$$J = \begin{bmatrix} B\,\psi_c' & -B \\ \dfrac{1}{B} & -\dfrac{1}{B\,\psi_T'} \end{bmatrix}, \qquad \psi_T' = \frac{2\sqrt{\psi^*}}{\gamma_T} \tag{2.7}$$

where ψ_c' is the slope of the characteristic at the operating point and ψ_T' is the slope of the throttle line through it. Both slopes are pressure rise per unit of flow, so the two off-diagonal terms carry B and $1/B$ and the trace and determinant follow directly, as shown in Equation (2.8).

$$\operatorname{tr} J = B\,\psi_c' - \frac{1}{B\,\psi_T'}, \qquad \det J = 1 - \frac{\psi_c'}{\psi_T'} \tag{2.8}$$

Stability of the equilibrium requires a negative trace and a positive determinant. The determinant condition is the classical static criterion that the throttle line must be steeper than the characteristic, and it is independent of B. The trace condition is dynamic, it involves no comparison of the two curves, and B is the only quantity in it that is not fixed by where the operating point sits. On the negative-slope branch $\psi_c' < 0$ and both conditions hold for any B. On the positive-slope branch the trace turns positive as soon as $\psi_c' > 1/(B^2\psi_T')$, so the machine can lose stability while the static criterion is still satisfied.

At the default parameters that is what happens, and the static criterion never fails at any throttle position. At an equilibrium the throttle-line slope reduces to $\psi_T' = 2\psi^*/\varphi^*$, which is at least 2.24 at every equilibrium on the positive-slope branch, while the largest slope the cubic reaches is $3H/2W = 1.08$ at $\varphi = W$. Sweeping the throttle across its whole travel gives a minimum determinant

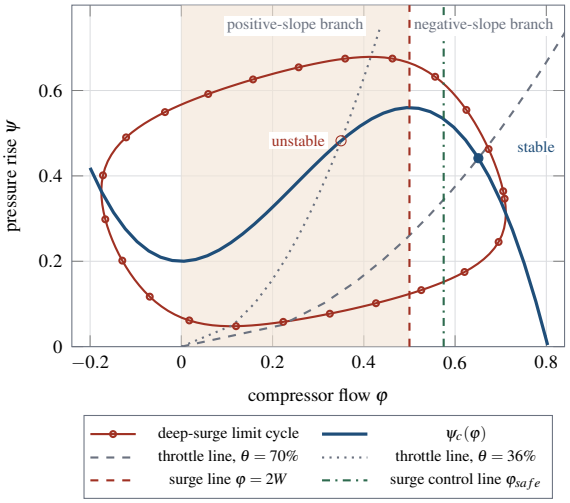

Figure 2.2 Compressor map with both stability boundaries and the measured orbit. The orbit carries open circles, the characteristic is the heavy solid curve, and the throttle lines at 70 % and 36 % are dashed and dotted. The filled marker is the stable point at 70 % throttle and the open marker is the unstable equilibrium at 36 % throttle that the deep-surge cycle encircles.

of 0.639 at $\theta = 31.2\,\%$, where the local slopes are 1.06 for the characteristic and 2.94 for the throttle line. The trace, by contrast, changes sign at $\theta = 45.62\,\%$, where $\gamma_T = 0.6386$, $\varphi^* = 0.4769$, $\psi^* = 0.5578$, $\psi_c' = 0.1900$ and $\psi_T' = 2.3390$. A pair of complex eigenvalues crosses the imaginary axis there with $\det J = 0.9188$, so the linearized oscillation frequency is $\sqrt{\det J} = 0.9585$ and its period is 6.56 $\hat{t}$. The crossing is a Hopf bifurcation, the same event that Volume 7, Chapter 2 locates on the hot branch of an exothermic reactor.

Two boundaries therefore exist and they are not in the same place. The surge line an operator is taught, the peak of the characteristic at $\varphi = 0.500$, is reached at $\theta = 47.73\,\%$. The dynamic boundary at which the equilibrium actually stops attracting is at $\varphi = 0.4769$ and $\theta = 45.62\,\%$, which is 2.11 percentage points of throttle travel further in. The healthy operating point at the default $\theta = 70\,\%$ is far from both boundaries, with $\gamma_T = 0.980$, $\varphi^* = 0.6511$, $\psi^* = 0.4414$, $\psi_c' = -1.700$, $\operatorname{tr} J = -3.042$ and $\det J = 2.254$.

Pause. Try Exercise 2.1 before reading on.

The response past the dynamic boundary is shown in Figure 2.2. The trajectory spirals away from the equilibrium, and the two ends of the cubic then stop the growth. Toward high flow the characteristic falls steeply, so the imbalance driving Equation (2.1) becomes strongly negative. Toward negative flow the cubic turns upward again, so a reversed flow is pushed forward. The result is a closed orbit, whose amplitude is set by the nonlinearity and not by the disturbance that displaced the equilibrium. A throttle closure to 36 % produces an orbit with a period of 7.51 $\hat{t}$, φ running from -0.175 to 0.711 and ψ from 0.048 to 0.679. The flow amplitude changes by 3×10^{-9} between the successive 60 $\hat{t}$ windows that begin at $\hat{t} = 68$ and $\hat{t} = 128$, which is the test that separates a limit cycle from a slow divergence. The application's verification suite applies the same test to the orbit at 40 % throttle with a tolerance of 2 % and measures a change of 7×10^{-6}. The nonlinear period is 15 % longer than the 6.56 $\hat{t}$ predicted by the linearization at onset, because most of the orbit is spent far from the equilibrium where the linear model does not apply.

2.4.1 Physical Interpretation of the Greitzer Parameter

The Greitzer parameter is the ratio of a pressure-rise time scale to a flow-acceleration time scale, and it is built from the geometry of the installation as much as from the machine, as shown in

Equation (2.9).

$$B = \frac{U}{2\,\omega_H L_c}, \qquad \omega_H = a\sqrt{\frac{A_c}{V_p L_c}} \tag{2.9}$$

where U is rotor tip speed, a is the speed of sound, A_c and L_c are the flow area and length of the duct and V_p is the plenum volume. Combining the two gives $B = (U/2a)\sqrt{V_p/(A_c L_c)}$, so B grows with rotor speed and with the square root of the volume bolted to the discharge, and shrinks with duct length. The Helmholtz frequency ω_H is the resonance of that duct and volume acting as a mass and a spring, and $\hat{t} = \omega_H t$ is the nondimensional time used throughout. On a real machine one unit of $\hat{t}$ is typically tens of milliseconds.

The trace condition puts a critical value on B at every operating point, $B_{crit} = 1/\sqrt{\psi_c'\,\psi_T'}$ wherever $\psi_c' > 0$. At $\theta = 36\%$ that value is 0.633, so the default $B = 1.5$ surges and $B = 0.5$ does not. Along the whole positive-slope branch B_{crit} never falls below 0.538, the limit it approaches as the throttle closes completely, so at $B = 0.5$ no throttle position produces an unstable equilibrium. Scenario 8 demonstrates this result. Two machines with identical aerodynamics, identical characteristics and identical final operating points can behave differently depending on the volume connected downstream, because a large discharge drum raises B, and a compressor that is stable on a test stand can surge in the plant installation.

2.5 Controller Design

The override described here is the default, and a model predictive controller (Section 2.5.1) is offered as a second option. The anti-surge loop is a proportional-integral controller with a positive gain and the ordinary sign convention, so a measured flow below the setpoint produces a positive error and opens the valve. Derivative action is set to zero and is not offered. The algorithm is the velocity form of Section 1.6.2, which computes an increment to the valve signal at each control interval $\Delta t = 0.006\,\hat{t}$ rather than the signal itself. For this loop it reads as Equation (2.10), with subscript k the execution index.

$$e_k = \varphi_{safe} - \hat{\varphi}_k \tag{2.10a}$$

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

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

$$s_k = s_{k-1} + \Delta u_k^P$$

$$c_k = s_k + I_{k-1} \tag{2.10b}$$

$$u_{cmd,k} = \mathrm{slew}\big(\mathrm{sat}_{[0,1]}(c_k),\ u_{cmd,k-1},\ \dot{u}_{max},\ \Delta t\big) \tag{2.10c}$$

$$I_k = I_{k-1} + \Delta u_k^I + \frac{\Delta t}{T_t}\,\big(u_{cmd,k} - c_k\big) \tag{2.10d}$$

$$u_{app,k} = \mathrm{slew}\big(u_{cmd,k},\ u_{app,k-1},\ \dot{u}_{max},\ \Delta t\big) \tag{2.10e}$$

where s_k is the accumulator that carries the proportional increments without limit, c_k is the unclamped output the algorithm demands, built on the integral accumulated through the previous execution, sat clips to the valve travel, slew limits the change to $\dot{u}_{max}\Delta t$ per execution, I is the integral state in valve fraction and $T_t = \tau_I$ is the anti-windup tracking time. In the symbols of Equations (1.13) to (1.15), $s_k = S_k$, $c_k = v_k$ and $u_{cmd,k} = u_k$, with the output bias u_b zero, so $s_k = K_c e_k$ at every execution. Anti-windup is back-calculation, the last term of (2.10d), which drives the integral toward the value that reproduces the command actually issued after both the clamp and the rate limit. Away from both limits the recursion is the positional PI law with the same gains. Every term in (2.10d) carries units of

Table 2.3 Controller Settings and Their Ranges

Symbol	Meaning	Default	Range
φ_{safe}	anti-surge setpoint	0.575	0.500–0.720
K_c	proportional gain, valve fraction per φ	8.00	0–60
τ_I	integral and tracking time ($\hat{t}$)	4.00	0.2–30
τ_f	flow measurement filter ($\hat{t}$)	0	0–1.0

valve fraction, because K_c is valve fraction per unit of flow, e is a flow, τ_I, T_t and Δt are all in $\hat{t}$, and the tracking difference is one valve fraction subtracted from another. The rate limit appears twice, in (2.10c) inside the controller and again in (2.10e) on the position that reaches the valve model, which is why the applied position lags the command even before the first order lag τ_v of Equation (2.5) is applied.

The lower clamp in (2.10c) is what makes this an override rather than a regulator, and the tracking term in (2.10d) is what makes the clamp usable. At the default operating point the flow is 0.6511 against a setpoint of 0.575, so $e = -0.0761$, the proportional term is -0.61 and the unclamped output is far below zero. A textbook integrator fed that error would integrate downward without bound for as long as the machine ran well, and would then need the same time to unwind before the valve could move at all. With the command resting on the lower clamp, $u_{cmd,k} = 0$ and $c_k = K_c e_k + I_{k-1}$, and because $T_t = \tau_I$ the error terms in (2.10d) cancel exactly, leaving $I_k = (1 - \Delta t/\tau_I) I_{k-1}$ [15]. This is the external reset of Equation (1.4) with the lower clamp as the applied output. The integral therefore decays toward zero with time constant τ_I whatever the size of the error, and the demand rests on the proportional term alone. In scenario 1 the integral starts at zero and stays there, so the demand is –0.609 for the whole run. At the wide-open clamp the same algebra relaxes I toward 1, so the command leaves either stop as soon as the error changes sign. The resulting behavior is the one a protective loop requires. The valve stays shut while the flow is above the setpoint, and the command leaves zero at the first execution after the measured flow crosses φ_{safe}, with no dead period. In scenario 3 the flow falls below 0.575 at $\hat{t} = 9.124$ and the command opens at the execution at $\hat{t} = 9.126$.

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

In MANUAL the integral state is seeded, $I_k = u_{man} - s_k$, so the transfer in either direction is bumpless, as in Section 1.6.2. On request in MANUAL the application also sets the valve to the opening that holds the declared machine at the setpoint in steady state with the throttle where it is set, the opening u_{ss} of Section 2.5.1 at $\hat{d} = 0$. The required opening changes as the throttle moves.

Two defenses that a real installation carries are deliberately absent. There is no feedforward from the rate of throttle movement, so the loop cannot begin opening before the flow has already started to fall, and there is no close-coupled valve or speed manipulation of the kind studied in the active stabilization literature [13, 14]. The setpoint itself is the third defense and it is the one left open to adjustment, because placing φ_{safe} above the surge line provides anticipation that no tuning can supply. The default $\varphi_{safe} = 0.575$ is a surge margin of 15.0 %, and the surge control line it draws is the vertical line in Figure 2.2.

Table 2.3 lists the default settings of the four operator-adjustable loop quantities.

2.5.1 Model Predictive Control

The second option replaces the proportional-integral law by a receding-horizon optimization on the same valve [16, 17], the option set out in general in Section 1.6.3. In AUTO, at every third execution, $T_s = 3\Delta t = 0.018\ \hat{t}$, the controller plans the valve command over the next $N = 111$ steps (2.0 $\hat{t}$) in eight move blocks of 1, 2, 3, 6, 12, 24, 30 and 33 steps, and the first move of each plan is applied and held for one decision period. Its model is the compression system of Section 2.3 with the declared

defaults of Table 2.2, marked with an overbar, and never with the values of the running plant, as shown in Equation (2.11).

$$\begin{aligned} \frac{\mathrm{d}\varphi}{\mathrm{d}\hat{t}} &= \bar{B}\big[\bar{\psi}_c(\varphi) - \psi\big], & \bar{\tau}_v \frac{\mathrm{d}u}{\mathrm{d}\hat{t}} &= u_c - u, \\ \frac{\mathrm{d}\psi}{\mathrm{d}\hat{t}} &= \frac{\varphi - (\gamma_T + \bar{K}_v u)\sqrt{\psi} + d}{\bar{B}}, & \bar{\tau}_{th} \frac{\mathrm{d}\gamma_T}{\mathrm{d}\hat{t}} &= \bar{\gamma}_{max}\theta_m - \gamma_T \end{aligned} \tag{2.11}$$

where u_c is the valve command, $\bar{\psi}_c$ is the cubic of Equation (2.3) with the declared ψ_0, H and W, d is an unmeasured flow imbalance held constant over the horizon, and θ_m is the commanded throttle demand, or the design value of 0.70 when the throttle signal to the controller is switched off. The restriction drift is not part of θ_m and reaches the model only through d. At each decision the state $x = [\varphi,\ \psi,\ u,\ \gamma_T,\ d]^T$ is linearized about the current estimate, with the remainder carried as an affine term, and sampled by an exact zero-order hold over T_s, so the model follows the slope ψ_c', which changes sign at the peak, and the valve-to-flow gain, which step tests on the plant put between 0.151 and 0.238 over the operating range.

The model is checked rather than identified. At the 70 % design point its Jacobian agrees with central differences of the plant's derivatives to 1.5×10^{-10}, and its steady valve-to-flow gain is 0.1769 against 0.1697 from a 0.05 valve step on the plant, to which a first-order-plus-dead-time fit (0.170, 1.15 $\hat{t}$, 0.55 $\hat{t}$), never used by the controller, leaves a residual of 1.4 % of the step. On scenario 11 the root-mean-square error of the predicted flow is 4.5×10^{-4} one decision ahead and 0.011 at 2 $\hat{t}$. Each plan solves Equation (2.12).

$$\begin{aligned} \min_{u_0^b,\ldots,u_7^b,\ s} \ & \sum_{i=0}^{N-1} \ell\, u_i + S \sum_{j=0}^{7} \big(u_j^b - u_{j-1}^b\big)^2 + \rho s^2 + \sigma s \\ \text{s.t.}\ \ & \varphi_i \geq \varphi_{safe} - s, \quad i = 1,\ldots,N, \qquad s \geq 0, \\ & 0 \leq u_j^b \leq 1, \qquad \big|u_j^b - u_{j-1}^b\big| \leq \dot{u}_{max} T_s, \\ & u_7^b \geq \min\big(u_{ss},\ u_{-1}^b + 8\,\dot{u}_{max} T_s\big) \end{aligned} \tag{2.12}$$

where u_j^b is the command of block j, u_i is the command in force at step i, u_{-1}^b is the command applied when the plan is made, φ_i is the predicted flow, $\ell = 0.01$ is the price of valve opening per step, $S = 10$ is the move weight, $\rho = 10^6$ and $\sigma = 10^3$ are the prices of the slack s, and $u_{ss} = [(\varphi_{safe} + \hat{d})/\bar{\psi}_c(\varphi_{safe})^{1/2} - \bar{\gamma}_{max}\theta_m]/\bar{K}_v$, clipped to $[0, 1]$, is the opening that holds the declared model at φ_{safe} in steady state. The recycle is a linear cost, so the plan opens the valve only as far as the predicted path requires, and the setpoint becomes a soft floor with one slack for the horizon (one per step changed no minimum margin on scenarios 6, 7 and 11 at 41 to 53 times the solve time). The valve travel and the stroking rate the override slews at are hard, the terminal row stops the final block from closing the valve where its effect would fall past the horizon, plenum pressure carries no weight, and $S = 10$ is the smallest move weight within 10 % of the least command travel found on scenario 11 (85.6 % against 79.1 % of stroke).

The estimator is updated at each execution regardless of which controller has the valve, as given by Equation (2.13).

$$\hat{x}_k^- = \Phi\big(\hat{x}_{k-1},\, u_{app,k-1},\, \theta_m\big), \qquad \hat{x}_k = \hat{x}_k^- + L\big(\hat{\varphi}_k - \hat{x}_{k,1}^-\big) \tag{2.13}$$

where Φ is one fourth-order Runge-Kutta step of Equation (2.11) over Δt with the applied command, $\hat{x}_{k,1}^-$ is the predicted flow, and L is the steady-state Kalman gain [3] of the model linearized at 70 % throttle with the valve shut, for a process noise covariance of $\mathrm{diag}(10^{-7}, 10^{-7}, 10^{-9}, 10^{-9}, 10^{-6})$ and a measurement variance of 10^{-5}, which gives $L = [0.0957,\ -0.161,\ 6.5 \times 10^{-6},\ 2.14 \times 10^{-5},\ -0.301]^T$. The gain is fixed, because recomputing it at every decision costs 6.35 to 6.42 ms in Node, about 35 % of a processor core at real time, so the estimator is a constant-gain extended

Kalman filter. The imbalance d is an input disturbance model that makes the controller offset free [18], and on a plant with $K_v = 0.45$, $B = 1.2$ and $H = 0.19$ the flow rests 1.8×10^{-9} from φ_{safe} at $\hat{t} = 60$.

A plan is accepted only if the solver reports an optimal solution, the first move equals the first entry of the plan, every planned command lies in $[0, 1]$, and the predictions, cost and slack are finite and of the expected sizes. Otherwise the override, which tracks the applied command with no preload, takes the valve with a step of 0.25 % of stroke when forced in the middle of the closure of scenario 11, until a later plan is accepted. Tracking leaves the override's demand at zero, so a fallback forced from $\hat{t} = 8$ to 20 opens at the first fall in flow and reaches +7.70 % rather than −8.6 %. Under the plan the override's gain and integral time have no effect, while φ_{safe} and τ_f stay in force.

Three models are kept apart, the plant of Equations (2.1) to (2.5) with the run's parameters and the drift, the instrument of Equation (2.6), which is the only measurement, and the controller's model of Equation (2.11), which takes the throttle command as an operator setting and not as a measured position. At normal speed, one unit of $\hat{t}$ a second, the controller solves 55.6 plans a second, and in Chrome 157 its whole execution took 1.30 to 2.00 % of one processor core and 5.19 to 6.49 % at the fastest speed of four times, with a worst execution of 2.1 ms.

2.6 Performance Metrics for Override Control

The application reports the five standard control-performance measures that every framework app shows on its scorecard (Section 1.6.4) together with three measures specific to this loop, and four of the five standard measures require careful interpretation on an override.

The integral of absolute error is the clearest case. In healthy operation the valve is clamped shut, the flow sits well above the anti-surge setpoint, and the accumulating error is simply the distance to a target the controller has no intention of reaching. Standing still at the default operating point for 29.9 $\hat{t}$ produces an IAE of 2.28 while the valve has not moved once and the surge margin has never fallen below 30.2 %, and a longer run would produce a larger number for the same perfect behavior. Every IAE quoted in this chapter is reported with the run length beside it and is never used to compare one tuning against another.

The saturated fraction is reported for the wide-open end of travel only. A regulatory loop that sits on either limit has lost authority, but a protective override resting shut is doing its job, so counting the closed end would report 100 % saturation for a machine running perfectly. Saturation is therefore counted at the wide-open limit only, and these percentages are not comparable with those of the other apps, where both limits count.

Overshoot and settling time are step-response measures, and the application reports them only after a setpoint move. No scenario moves φ_{safe}, so both read blank throughout, which is the correct outcome for an override. Computed instead against the standing difference between the initial flow and the setpoint, 0.0761 at the default operating point, scenario 4's excursion to −0.175 would read as an overshoot near 986 %, which restates the minimum surge margin rather than judging the tuning. Settling has a second defect on this loop, because a surging flow passes through any settling band twice per cycle. The application commits a settling time only after the flow has stayed inside the band for a dwell, which keeps a limit cycle from reporting a finite value, but the loop still has no target to settle to. Neither measure appears in Table 2.5.

Three measures are meaningful without qualification. Minimum surge margin, $100(\varphi_{min} - 2W)/2W$, is negative if and only if the operating point crossed the surge line at any moment. Cumulative recycled flow, $\int \varphi_u \, d\hat{t}$, is the gas compressed and thrown straight back to the inlet and is the running cost of protection. Peak detector swing is the largest windowed $\Delta\psi$ reached and is the quantity the surge alarm acts on. Those three and the total variation of the controller output are what the scenarios below are judged by.

Table 2.4 The Twelve Scenarios and the Quantity Each One Moves

#	Name	Setting
1	Stable operation	AUTO at 70 % throttle, valve shut, nothing scheduled
2	Close the throttle by hand	MANUAL, operator steps the throttle down
3	Rapid closure, controller on	AUTO, $\tau_{th} = 0.12$, $\theta \to 36\,\%$ at $\hat{t} = 8$
4	Same closure, controller off	MANUAL, $\tau_{th} = 0.12$, same event
5	Fast anti-surge valve	$\tau_v = 0.08$, $\dot{u}_{max} = 14$, $\varphi_{safe} = 0.60$, $\tau_{th} = 0.25$, $\theta \to 36\,\%$ at $\hat{t} = 8$
6	Slow anti-surge valve	$\tau_v = 2.2$, $\dot{u}_{max} = 0.5$, $\varphi_{safe} = 0.60$, $\tau_{th} = 0.25$, same event
7	Undersized valve	$K_v = 0.18$, $\tau_{th} = 0.40$, $\theta \to 30\,\%$ at $\hat{t} = 8$
8	A smaller plenum	MANUAL, $B = 0.5$, $\tau_{th} = 0.40$, $\theta \to 50, 44, 38, 34\,\%$ at $\hat{t} = 8, 22, 36, 50$
9	Noisy flow transmitter	$\sigma_n = 0.014$, $\tau_{th} = 0.40$, $\theta \to 42\,\%$ at $\hat{t} = 8$
10	Aggressive PI tuning	$K_c = 34$, $\tau_I = 0.8$, $\tau_{th} = 0.40$, $\theta \to 40\,\%$ at $\hat{t} = 8$
11	Predictive control against the PI	As 3, predictive controller
12	Challenge	$K_c = 1.5$, $\tau_I = 15$, $\varphi_{safe} = 0.55$, $\tau_{th} = 0.40$, $\theta \to 38\,\%$ at $\hat{t} = 10$

2.7 Results

The twelve scenarios are given in Table 2.4. A scenario sets the quantities under study together with the throttle stroke constant τ_{th} of its closure (0.12 to 0.40 against the default 0.50) and leaves every other setting at its default. Every scenario that ends in surge begins at a stable point on the negative-slope branch and is driven across the boundary by a scheduled throttle event, because an unstable equilibrium is still an equilibrium and a run that merely started there would sit motionless. Each scenario also carries a completion rule (Section 1.6.4), and several of the exercises at the end of the chapter follow those rules.

The measured performance of each run with the scenario as supplied, read once it had reached the $\hat{t}$ shown, is collected in Table 2.5. Scenario 2 is driven by hand and has no scheduled event, so it carries no row and is described separately below, and scenario 11 is compared with the override in Section 2.7.6. MV travel is the total variation of the controller output in percent of valve travel, not of the valve position, because the sixth scenario exists to show that those are different signals. The overshoot and settling measures read blank for the reasons given in Section 2.6. Figure 2.3 sets the minimum margin and the recycled flow of that table side by side.

2.7.1 Staged Manual Throttle Closure

The second scenario locates the stability boundary by measurement. Stepping the throttle down by one percent every 20 $\hat{t}$ from 70 %, with the anti-surge valve locked shut, leaves the plenum pressure swing immeasurably small down to 54 %, then growing slowly in the last few steps of the settled part of the run. Twelve to twenty units after each step the swing is 0.0014 at 48 % and 0.0100 at 46 % and still decaying, 0.046 at 45 % and growing, and 0.537 at 44 %, where the surge alarm latches 12.0 $\hat{t}$ after the step. The last position that holds and the first that does not are one percent of throttle travel apart, and the analysis of Section 2.4 puts the boundary between them at 45.62 %. Crossing the peak

Table 2.5 Measured Performance of Each Scenario

#	$\hat{t}$	IAE	MV travel	Sat. (%)	Min margin (%)	Recyc. flow	Peak $\Delta\psi$
1	29.9	2.276	0	0.0	+30.2	0.00	0.000
3	30.0	1.138	624	2.0	−8.6	4.19	0.222
4	59.7	17.808	0	0.0	−135.1	0.00	0.631
5	29.9	0.717	182	0.0	+3.3	4.75	0.171
6	59.3	2.649	387	5.2	−105.4	11.23	0.624
7	59.9	16.394	1497	56.2	−129.0	3.60	0.645
8	95.5	17.117	0	0.0	−43.8	0.00	0.211
9	30.5	0.987	4617	0.0	+1.0	2.95	0.164
10	59.9	2.207	3500	27.7	+6.0	9.01	0.170
12	60.0	4.604	402	0.0	−93.5	5.56	0.595

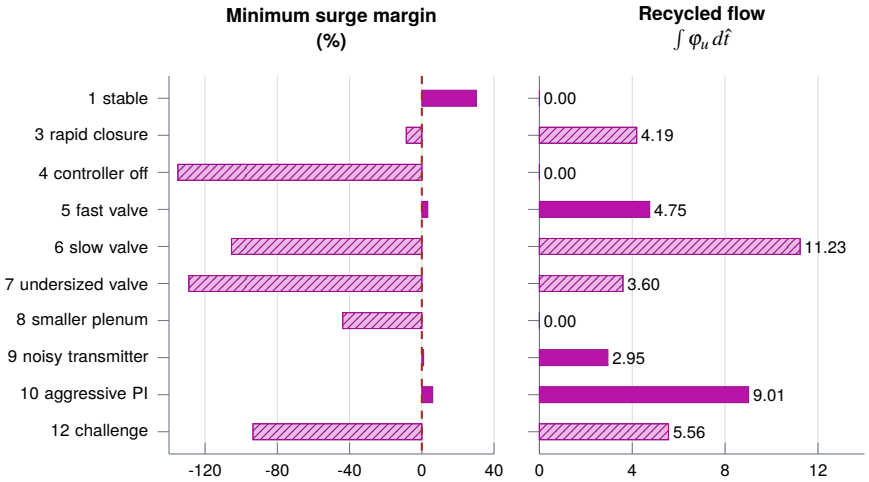

Figure 2.3 Table 2.5 redrawn: minimum surge margin and recycled flow of the ten scenarios with a row. Solid bars keep a positive margin, hatched bars cross the surge line (dashed, zero margin). Notice that recycling more does not by itself protect the machine: the slow valve of scenario 6 recycles the most gas and still surges, while the fast valve of scenario 5 holds with less than half as much.

of the characteristic at 47.73 % produced no visible response, which demonstrates that the static surge line and the dynamic stability boundary are different objects.

2.7.2 The Deep-Surge Limit Cycle

Scenario 4 is the central demonstration, and it is the run captured in Figure 2.5. The throttle steps from 70 % to 36 % in about a tenth of a Helmholtz period with the valve locked shut, and the machine settles into the deep-surge orbit of Figure 2.2. The surge alarm latches at $\hat{t} = 11.21$, the instant the flow first reverses. After 59.7 $\hat{t}$ the minimum margin is −135.1 %, meaning the flow reached −0.175, and the detector swing is 0.631 against an alarm threshold of 0.25. The flow trend shows a clean periodic waveform with negative excursions once per cycle and the pressure trend shows ψ and $\psi_c(\varphi)$ separating, which is the visual signature that the two states are no longer in balance at any instant.

TRY IT IN THE APP 4 · Same Closure, Controller Off

Load the scenario and press Run. Watch the Compressor flow trend. After the throttle closes at $\hat{t} = 8$ the flow falls, reverses to about −0.175 and recovers, again and again with the same amplitude. The surge alarm latches at $\hat{t} = 11.21$, the instant the flow first reverses. Nothing grows without bound.

The oscillation neither settles nor grows. That combination is what makes the object a limit cycle rather than an instability in the usual sense. It has no dependence on the size of the disturbance that started it, it cannot be described by a growth rate, and it returns to the same orbit from any nearby initial condition. The eighth scenario is the control against which that claim is tested. With B lowered to 0.5 and the throttle stepped to 34 % in four stages, the machine oscillates and decays after each step, dips past the surge line on the way, and then settles toward the equilibrium at $\varphi = 0.3209$, reading 0.3207 at $\hat{t} = 95.5$ with a residual swing of ± 0.002. The minimum margin over the run is −43.8 % and the peak detector swing is 0.211, below the alarm threshold. Identical aerodynamics, identical final flow, and no surge, because B_{crit} at that operating point is 0.596 and the plenum is now too small to sustain the oscillation. Run again with $B = 1.5$, the same sequence trips the surge alarm at $\hat{t} = 34.3$, 12.3 units after the step to 44 %, and ends in deep surge with a minimum margin of −138.0 %.

2.7.3 Effect of Valve Stroke Speed and Capacity

Scenarios 3, 5, 6 and 7 take a fast closure and change the final element. With the standard valve and a closure to 36 % at $\tau_{th} = 0.12$, about a tenth of a Helmholtz period, the controller is too late, reaching a minimum margin of −8.6 % after spending 624 points of command travel and 4.19 of recycled flow. Scenario 5 gives the valve a stroke constant of 0.08 instead of 0.30 and a rate limit of 14 instead of 4, moves the setpoint out to 0.60, and strokes the throttle with $\tau_{th} = 0.25$. It holds the closure at +3.3 % minimum margin for 182 points of travel and 4.75 of recycle. Both changes to the loop are needed on that closure, since the standard valve at a setpoint of 0.60 reaches −1.03 % and the fast valve at 0.575 reaches −3.03 %. The slower throttle is not what saves it, because the fast valve with $\tau_{th} = 0.12$ still holds +2.95 %.

The reverse change is the more instructive one. Scenario 6 has the same capacity, the same tuning and the same setpoint as scenario 5, and only the stroke speed is different at $\tau_v = 2.2$ and $\dot{u}_{max} = 0.5$. The controller acts correctly. Its command leaves zero 1.0 $\hat{t}$ after the closure and reaches 100 % at $\hat{t} = 11.0$, when the valve is 34 % open, and the machine surges anyway. The alarm latches at $\hat{t} = 11.7$, the flow reverses once, to a minimum margin of −105.4 % at $\hat{t} = 12.4$, and the alarm clears at $\hat{t} = 24.6$ once enough recycle is open. No proportional gain and no integral time changes that outcome, because the decision was correct while the actuator had not reached the required position in time. This is why anti-surge valves are specified by stroke time as well as by capacity [10].

Scenario 7 fails for a third reason. Holding the flow at $\varphi_{safe} = 0.575$ requires a total downstream conductance of 0.7874, and at 30 % throttle the throttle itself supplies only $\gamma_T = 0.420$, so the recycle valve must contribute 0.3674. At $K_v = 0.18$ the most it can contribute wide open is 0.18. The command pins at 100 % for 56.2 % of the run, the reachable equilibrium at $\varphi = 0.444$ is past the dynamic boundary on the positive-slope branch, and the machine surges to −129.0 % minimum margin. A saturated protective actuator is a sizing statement, not a tuning statement (Section 1.4.5).

DID YOU KNOW? Same Valve Size, Opposite Outcomes

From the simulator: scenarios 5 and 6 share valve capacity, tuning and setpoint and differ only in how fast the valve strokes. The fast valve holds the closure at a minimum surge margin of +3.3 %. The slow one lets the flow reverse, to −105.4 % (Table 2.5).

2.7.4 Control-Induced Oscillation

Scenario 10 raises the gain to $K_c = 34$ and shortens the integral time to $\tau_I = 0.8$, and the first flow dip is shallower. Running the identical closure to 40 % throttle with the default $K_c = 8$ and $\tau_I = 4$ gives a minimum margin of −2.13 % for 476 points of command travel, and the aggressive tuning gives +5.97 % for 3500 points. The added margin costs a factor of seven in valve activity.

What the aggressive tuning leaves behind is a sustained oscillation, and it is not surge. The peak detector swing is 0.170 against a threshold of 0.25, so the alarm never trips. The flow cycles between 0.534 and 0.627 with a period of 2.92 $\hat{t}$, entirely above the surge line and astride the setpoint, against the 7.51 $\hat{t}$ and the excursion to −0.175 of the aerodynamic cycle. Period and amplitude are set here by the loop gain and the valve lag rather than by the Helmholtz dynamics, and one diagnostic separates the two cases, switching the loop to MANUAL. Aerodynamic surge continues and a control-induced cycle stops.

Pause. Try Exercise 2.2 before reading on.

TRY IT IN THE APP 10 · Aggressive PI Tuning

Run the scenario and watch the flow after the closure to 40 % throttle. It cycles between 0.534 and 0.627, astride the setpoint and above the surge line, with a period of 2.92 $\hat{t}$, and the surge alarm never trips. Then switch the loop to Manual. The chapter's test is that a control-induced cycle stops there while aerodynamic surge would continue.

2.7.5 Effect of Measurement Noise

Scenario 9 adds 0.014 of transmitter noise, about 3 % of the surge flow, and closes the throttle to 42 %. Over 60.0 $\hat{t}$ the controller output travels 11 234 points with the noise present and 399 points with the noise removed, so the same duty costs twenty-eight times as much command activity. The travel with noise accumulates linearly with run length and the travel without it does not, so that ratio is only meaningful with the run length attached. Adding a measurement filter of 0.10 $\hat{t}$ brings the travel back to 3571 and costs 1.78 points of minimum surge margin, from +1.03 % to −0.74 %, because the filter adds phase lag to a loop that has none to spare.

The valve position travels far less, 1090 points with the noise against 352 without it, a factor of 3.1 rather than 28, because the rate limit and the valve lag absorb most of the chatter. This is why the failure mode is hard to see from the process side until the positioner is worn out. A noisy protective loop is judged on the controller output rather than on what the process ends up seeing.

2.7.6 Predictive Control Against the Override

Scenario 11 repeats the closure of scenario 3 under the predictive controller, as shown in Figure 2.4. The plan starts opening the valve at its first decision after the throttle command changes ($\hat{t} = 8.01$), strokes the command to 64.5 % by $\hat{t} = 9.43$ and eases it to 47.2 % by $\hat{t} = 20$, and the flow comes down onto the floor at $\hat{t} = 14.50$ without crossing it, a minimum margin of +15.00 %. The override opens only once the measured flow crosses its setpoint and reaches −8.61 % at $\hat{t} = 10.33$, after 1.12 $\hat{t}$ under the surge line. To $\hat{t} = 20$ the plan recycles 2.459 against 2.114, with 85.6 % of command travel against 529.4 %.

The comparison over nine cases is collected in Table 2.6, every row of a case with the same seed, duration, instruments and noise, and the predictive controller at scenario 11's settings throughout. Two classical rows receive the throttle command the plan reads, measured outside the application on

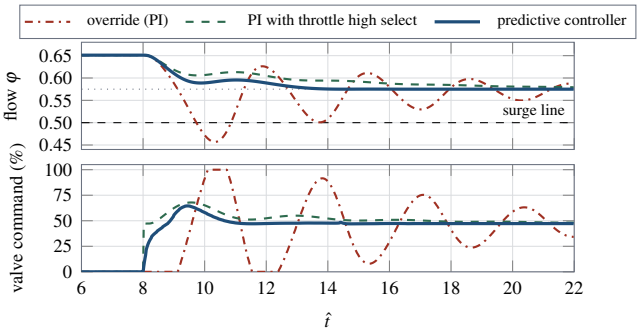

Figure 2.4 Flow and anti-surge valve command on the rapid closure of scenario 11 under three structures. The throttle is commanded from 70 % to 36 % at $\hat{t} = 8$. The override is dash-dotted, the override with the throttle high select dashed and the predictive controller thick solid. The thin dashed line is the surge line and the dotted line is the anti-surge setpoint at 0.575.

the same engine. As a feedforward of u_{ss} at $\hat{d} = 0$ added to the override's output it raises the margin on the closure only to +3.77 %, because the proportional term on a flow above the setpoint cancels it, while as a minimum opening through a high select, the open-loop step of industrial anti-surge systems [10], it gives the second group of Table 2.6.

Most of the margin on the closure is the throttle signal. Without it the plan bottoms at +4.67 %, near the additive feedforward's +3.77 %, while the high select holds +16.16 % at 8.6 % more recycle than the plan to $\hat{t} = 20$ and 1.4 % more over 90 $\hat{t}$, so what the plan adds is recycle and travel. The override alone needs a setpoint of 0.661 to hold +15 %, with 58 % more recycle over 90 $\hat{t}$ than the plan (26.762 against 16.946), and three retunings without the signal reach −10.23 % to +1.85 %. Where the command says nothing, the prediction from the flow still helps, at +11.67 % against +3.81 % for an unannounced drift (the most aggressive retuning reaches +8.97 %) and +15.00 % against +12.10 % for deep surge caught at $\hat{t} = 30$, with less travel in both. With a command that overstates the opening by 10 points, the plan holds +11.86 % and the high select +8.36 %.

The plan loses on recycle wherever its floor binds and the override's does not, by 2.0 % over 90 $\hat{t}$ of scenario 3, where the override saves gas by letting the flow sag under its setpoint, and by 16.7 % over 60 $\hat{t}$ of scenario 9, where the floor binds on the low excursions of a noisy estimate. On the slow valve of scenario 6 the plan (−4.09 %) and the high select (−3.91 %) avoid the alarm the override trips but cannot hold the surge line, because no plan strokes the valve faster than $\dot{u}_{max} = 0.5$ allows, and the undersized valve of scenario 7 defeats every structure (−118.03 % against −128.96 % over 90 $\hat{t}$). With the challenge's supplied settings the plan meets the challenge's target, +10.00 % with no alarm and 7.72 of recycle, because it does not use the weak gains.

2.7.7 Challenge Scenario

The challenge closes the throttle from 70 % to 38 % at $\hat{t} = 10$ and stops itself at $\hat{t} = 60$ so that the two scoring measures of different runs can be compared. The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt. The supplied tuning is deliberately weak at $K_c = 1.5$ and $\tau_I = 15$ on a loop whose work is over within two or three units, and it fails, reaching −93.5 % minimum margin with the surge alarm on from $\hat{t} = 14.1$ to 36.1. The target is a minimum margin at or above zero, no surge alarm, and cumulative recycled flow below 12.0. The suggested answer moves the setpoint to 0.60, raises K_c to 12 so that most of the valve travel happens inside the first unit, and leaves τ_I at 4. It finishes at +5.34 % minimum margin with a peak detector swing of 0.162, 10.59 of recycled flow (1.41 under the limit) and 544 points of command travel.

The two requirements pull against each other by design. A conservative setpoint makes the closure easy to survive and recycles gas every hour the plant runs, while a tight setpoint is inexpensive and

Table 2.6 Predictive Controller Against the Override and the Override Given the Throttle Signal

		Override (PI)			PI with throttle high select			Predictive controller				Predictive, signal off		
Case	$\hat{t}$	SM	Rec.	Trav.	SM	Rec.	Trav.	SM	Rec.	Trav.	Plans	SM	Rec.	Trav.
Rapid closure, sc. 3 and 11	0–20	−8.61	2.114	529.4	+16.16	2.670	96.4	+15.00	2.459	85.6	1112	+4.67	2.268	167.6
Rapid closure, sc. 3 and 11	0–90	−8.61	16.612	638.1	+15.00	17.190	98.2	+15.00	16.946	85.6	5000	+4.67	16.720	167.9
Signal 10 points wrong, sc. 3	0–90	−25.80†	24.757	738.9	+8.36	25.260	243.9	+11.86	25.333	140.8	5000	−3.55	25.100	192.2
Unannounced drift, sc. 1	0–90	+3.81	9.121	288.0	+3.81	9.121	288.0	+11.67	9.275	206.5	5000	+11.67	9.275	206.5
Deep surge caught, sc. 4	30–90	+12.10	12.339	197.0	+15.00	12.782	97.8	+15.00	12.266	57.3	3334	+12.28	12.200	181.6
Slow valve, sc. 6	0–90	−105.43†	18.593	387.5	−3.91	19.603	154.8	−4.09	19.282	234.8	5000	−96.64†	18.134	376.8
Noisy flow, sc. 9	0–60	+1.03	7.226	11234	+22.39	11.017	11731	+15.10	8.433	6205	3334	+11.47	8.257	6276
Challenge, supplied	0–60	−93.49†	5.556	402.3	+12.12	8.831	51.6	+10.00	7.718	49.6	3334	+3.04	7.633	200.9
Challenge, suggested	0–60	+5.34	10.588	544.2	+20.00	11.062	77.1	+20.00	10.861	59.3	3334	+11.83	9.995	185.1

SM is the minimum surge margin in the window (%), Rec. the recycled flow $\int \varphi_{rt}\, d\hat{t}$ in the window, Trav. the command travel (% of stroke), Plans the decisions solved. †Surge alarm. The wrong signal is a restriction drift of −10 % from $\hat{t} = 0$, the unannounced drift a step of −25 % at $\hat{t} = 8$, and the caught deep surge scenario 4 switched to AUTO at $\hat{t} = 30$, its earlier alarm unmarked. No predictive run had a fallback, and the mean solve time was 122 to 184 µs a plan in Node.

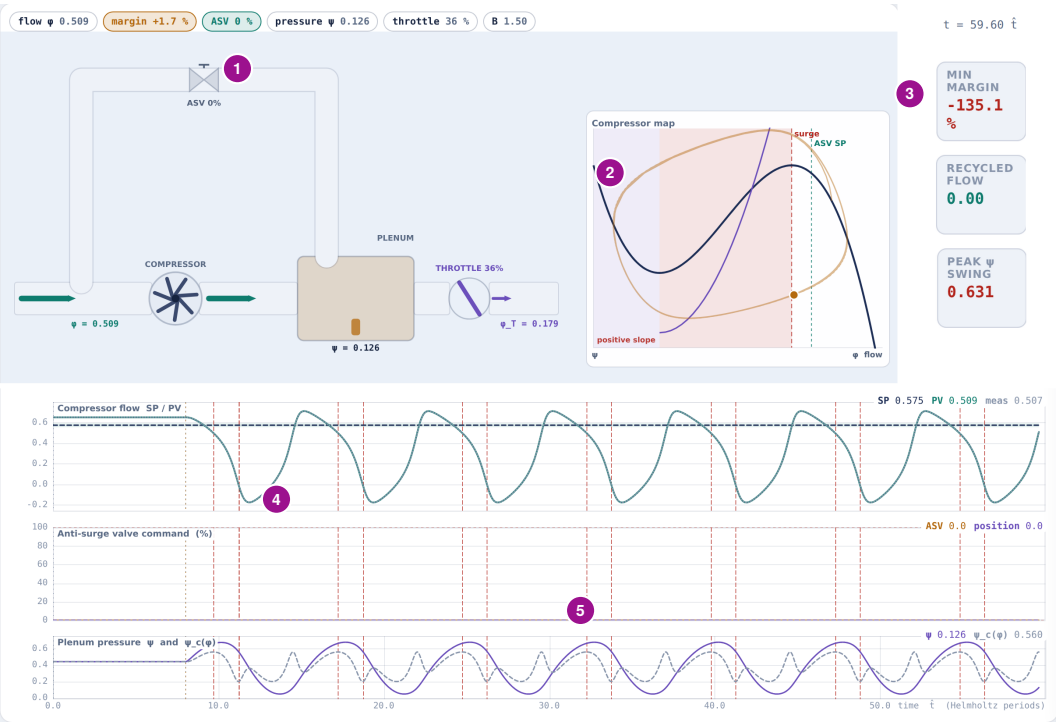

Figure 2.5 Deep surge after a throttle slam from 70 % to 36 % with the anti-surge valve locked shut, in a composite of the application window: the plant view, with the run clock and three scorecard tiles moved beside it, above the three trend strips. (1) The anti-surge valve rests at 0 %. (2) The compressor map shows the closed orbit crossing into reverse flow. (3) The tiles give a minimum surge margin of −135.1 % and a peak detector swing of 0.631. (4) The flow strip shows the periodic collapse and recovery over the whole run. (5) The valve command strip stays at zero on its 0 to 100 % scale. Constant actuator traces retain both declared limits, while moving traces can autoscale to show small movements. The pressure strip shows ψ and $\psi_c(\varphi)$ separating. This is scenario 4 at $\hat{t} = 59.6$.

leaves no margin when the load moves. The same trade-off drives anti-surge setpoint selection in operating plants, and the recycled-flow measure carries the operating-cost side of it that a study of tuning alone omits.

2.8 Technology in Practice

An anti-surge loop in a plant reads the same flow and moves the same valve as the override of Section 2.5. Figure 2.6 draws that loop with field devices in place of the schematic's symbols.

What Measures It

Flow is the most direct sign of surge, which is a breakdown of stable flow, but a 2025 review notes that accurate and fast flow measurement in the harsh compressor environment is demanding and that flow sensors often need an intrusive installation. Pressure-based monitoring is mature and easy to implement, though its accuracy depends on where the sensor sits and it may suffer from noise [19]. The model keeps the two apart as Section 2.3 does: the flow transmitter carries only the noise σ_n and the filter τ_f, with no bias and no transport delay, and the plenum pressure reaches only the detector.

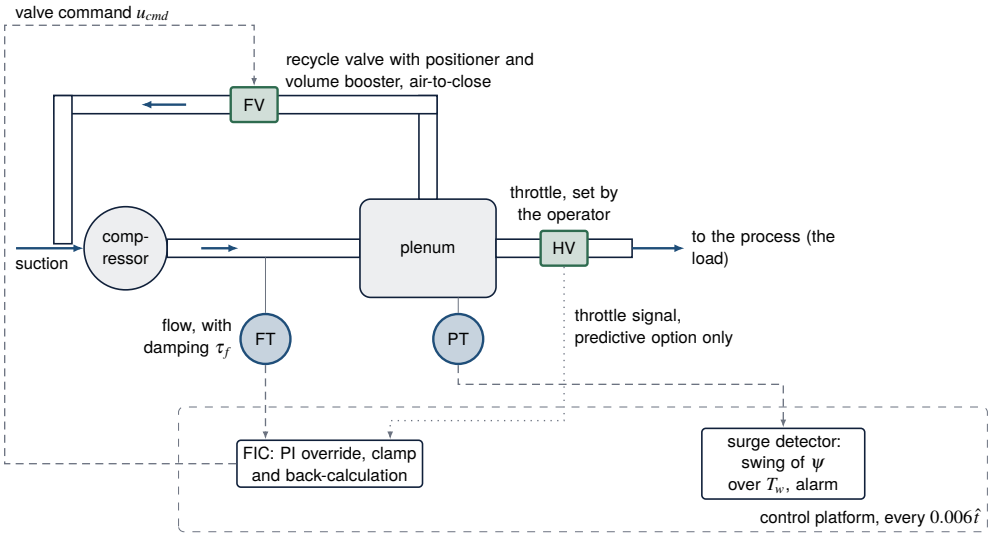

Figure 2.6 The anti-surge loop of Figure 2.1 on plant hardware. A flow transmitter (FT) with its damping feeds the override (FIC) on the control platform, which moves the recycle valve (FV) through its positioner and volume booster. The plenum pressure transmitter (PT) feeds only the surge detector, and the throttle (HV) is the load. Notice that the override never sees pressure, so the detector can report surge but cannot prevent it, and that only the predictive option reads the throttle (dotted).

> **TECHNOLOGY** **Sensor: A Flow Transmitter and Its Damping**
>
> Nearly every modern process transmitter, the flow transmitter of an anti-surge loop among them, carries a damping setting: a low-pass filter in line with its signal that keeps turbulence noise, such as that at a machine's discharge, out of the control system. Too little damping passes noise as ragged trends and erratic control. Too much makes the transmitter report a process that changes more slowly than it does, and an aggressively tuned loop then overshoots or oscillates while its own trend looks smooth [5]. The rule given is to use as little damping as good control needs. No open source read gives a range, accuracy or speed for such a transmitter. In the model the damping is the filter τ_f, and at 0.10 it cost scenario 9 1.78 points of minimum margin.

What Moves It

The final element is the recycle valve from the discharge back to the suction. Opening it keeps enough flow through the compressor and unloads the pressure across it, and a textbook loop diagram makes it air-to-close, so that a loss of air opens it [5]. The model gives the valve a capacity K_v, a lag τ_v and a stroking rate limit $\dot{u}_{max}$, and scenarios 6 and 7 show that stroke speed and capacity each bound what any tuning achieves.

TECHNOLOGY Final Element: A Recycle Valve Built to Stroke Fast

A pneumatic valve moves only as fast as air can fill and vent its actuator, and a large actuator is slow to fill. A volume booster repeats the signal pressure one to one with far greater air flow, which shortens the delay, although no finite flow removes it. A positioner places the stem precisely and also sources and vents high air flow, so it gives a faster stem as well [5]. These are the means of giving a recycle valve the stroke speed this chapter asks of it. A slow valve fails as scenario 6 does on its trend: the command reaches 100 % while the valve is 34 % open, and the machine surges. In the model τ_v and $\dot{u}_{max}$ stand for the actuator and its air supply.

Where the Controller Runs

Plant control loops run with cycle times from milliseconds to minutes [20]. Compressor surge control is given as an industrial example of a safety function that must respond within a set time, which asks for a network with hard determinism [5]. The application executes the override every 0.006 $\hat{t}$, and on a real machine one unit of $\hat{t}$ is typically tens of milliseconds. In a function-block system the override's selector is a control selector block whose back-calculation output tells a deselected controller to stop integrating, the role external reset played in pneumatic controllers [5]. Here the lower clamp and the tracking term of Equation (2.10d) play that part.

IN PRACTICE A Soft Override With a Hard Trip Behind It

Practice separates two layers of protection. A hard override, such as a switch that stops a well pump on low level, protects the machine at the cost of the service. A soft override is a controller whose output a selector passes only when its limit nears, so the service is curtailed rather than cut. A prudent design keeps the hard trip behind the soft loop, in case a setpoint is badly set or the tuning is poor, with the soft loop's setpoint placed so that it acts before the trip does. A controller with integral action that the selector deselects needs its windup limited [5]. The anti-surge loop of this chapter is the soft layer, and its surge alarm only reports.

What Is Modernizing It

Three developments bear on the loop. The first is predictive control. Nonlinear and simplified predictive controllers for anti-surge and process control have been timed on an industrial programmable logic controller and outperformed standard industrial control on surge margin, tracking and actuation without a large increase in computation time [21], the comparison Section 2.7.6 makes on scenario 11. The second is faster hardware: research is exploring high-speed pressure sensors, solenoid valves and piezoelectric actuators, because sensor delay and slow valve response often limit what a control law can do. The third is fusion and integration: pressure, vibration and temperature read together, data-driven methods combined with classical ones, and surge control coordinated with the combustion and cooling around the machine [19].

2.9 Limitations

The limitations of the model are stated in Section 2.3.6 and are not repeated. Four limitations belong to the application rather than to the physics.

The override sees a perfect throttle position in the sense that no feedforward path exists to use it, so under the override the application demonstrates the inadequacy of purely reactive protection without offering the standard cure, and the minimum valve opening computed from the throttle command, which holds +16.16 % on scenario 3, was measured outside the application. A rate-of-closure feedforward term would change the outcome of scenario 3 under the override and is the most valuable classical extension.

The surge margin is computed against a fixed line at $\varphi = 2W$. On a real machine the surge limit line is a curve in a corrected plane that moves with speed, suction temperature and gas composition, and much of the engineering in a commercial anti-surge system is in constructing and maintaining that curve. Nothing in this application exercises that part of the problem.

⚠ OPEN CHALLENGE Seeing Surge Before It Arrives

Surge monitoring aims at precursors seen early enough to intervene, not at the surge itself. A 2025 review finds those precursors weak, noisy and different from one machine and operating condition to the next, with no universally accepted set of features that reliably foretells surge. A change of fuel flow or of speed can raise pressure swings that resemble surge, and sensor bandwidth, filtering and data transmission add delay [19]. The chapter's detector, a windowed pressure swing set not to trip on one transient, shows the same trade. Progress would be a warning that holds across machines and arrives before the first reversal.

The pressure pulse event adds 0.04 directly to the plenum state rather than acting through an inlet or a valve, which is convenient for probing the basin of attraction near the peak but is not a physical disturbance. The application labels it a non-physical probe and offers it only among the advanced settings for that reason.

The predictive controller's model is the declared machine, linearized at each decision, so it knows neither a parameter the operator changes nor a drifted restriction, which the estimated imbalance absorbs only after the fact, and it reads the throttle command, not the position, which strokes in $0.12\,\hat{t}$ on scenario 3 against the model's 0.50. Its estimator gain is fixed at the 70 % design point, and the signal carries no preview of a closure before it is commanded.

⚠ OPEN CHALLENGE Active Control Has Yet to Leave the Test Bench

Active stabilization widens the stable flow range of research compressors [13], yet engineering hurdles remain. Sensors may suffer noise, delay or drift. Valves, guide vanes or pistons must answer fast enough, and the fastest add cost and complexity. Most strategies have been shown on simplified test benches or in simulation and must still be validated on full-scale machines, with a real-time platform and safety mechanisms for control failures [19]. A close-coupled valve also costs a pressure drop that may become unacceptably large when the valve is used as a steady device [14]. Progress would be a full-scale demonstration whose gain in range outweighs those losses.

2.10 Conclusion

This chapter specifies a four-state Greitzer compression system and the protective override that guards it, at the level of detail a second implementation needs. The two aerodynamic states carry a Moore-Greitzer cubic whose peak at $\varphi = 2W = 0.500$ is the surge line, and linearizing about the equilibrium gives a trace condition that fails at $\theta = 45.62\,\%$ and $\varphi = 0.4769$ while the static determinant condition never fails at any throttle position, so the Greitzer parameter B decides the outcome alone. Past that boundary the trajectory does not diverge but closes into a limit cycle of period $7.51\,\hat{t}$ with the flow running from −0.175 to 0.711, and lowering B from 1.5 to 0.5 removes the cycle at the same final flow. The default controller is a velocity-form proportional-integral override with back-calculation anti-windup against a lower clamp, which drives the integral toward zero while the valve rests shut, and the scenarios show that its outcome is bounded by valve stroke speed at −105.4 % minimum margin and by valve capacity at −129.0 %, neither of which any tuning repairs. A predictive controller holds the closure of scenario 3 at +15.00 % against −8.61 %, although the override given its throttle signal as a minimum opening holds +16.16 % at 8.6 % more recycle. Reported integral error on this loop measures distance to a setpoint the controller never intends to reach and is quoted only with its run length attached, while minimum surge margin and cumulative recycled flow carry the comparison. A fuller model would add a rate-of-closure feedforward path, a family of speed lines with a surge limit line fitted across them, a hot recycle temperature state, and a close-coupled valve for comparison with active stabilization. Within those limits the simulator lets a limit cycle be measured, and a protective override be operated, before either is met in a plant.

KEY POINTS

- Closed too far, a compressor does not run away. It settles into a repeating cycle in which the flow reverses through the machine.
- The dynamic boundary at $\theta = 45.62\,\%$ lies 2.11 points of throttle travel inside the static surge line, and the Greitzer parameter B alone decides it.
- Back-calculation against the lower clamp lets the integral decay while the valve rests shut, so the valve opens as soon as the flow crosses its setpoint.
- Valve stroke speed and capacity bound the outcome, at −105.4 % and −129.0 % minimum margin, and no tuning repairs either.
- A protective loop is judged by minimum surge margin and recycled flow, not by IAE.

Where else this applies. Volume 10, Chapter 11 runs the same override on three machines that share one load, and Chapter 3 protects a jet engine compressor with limiters that vote on one fuel valve.

References

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

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

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

[4] R. Kurz and R. C. White, "Surge avoidance in gas compression systems," *Journal of Turbomachinery*, vol. 126, no. 4, pp. 501–506, 2004.

[5] 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.

[6] 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.

[7] 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.

[8] ——, "Surge and rotating stall in axial flow compressors—part II: Experimental results and comparison with theory," *Journal of Engineering for Power*, vol. 98, no. 2, pp. 199–211, 1976.

[9] ——, "The stability of pumping systems—the 1980 Freeman Scholar Lecture," *Journal of Fluids Engineering*, vol. 103, no. 2, pp. 193–242, 1981.

[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] A. H. Epstein, J. E. Ffowcs Williams, and E. M. Greitzer, "Active suppression of aerodynamic instabilities in turbomachines," *Journal of Propulsion and Power*, vol. 5, no. 2, pp. 204–211, 1989.

[12] J. D. Paduano, A. H. Epstein, L. Valavani, J. P. Longley, E. M. Greitzer, and G. R. Guenette, "Active control of rotating stall in a low-speed axial compressor," *Journal of Turbomachinery*, vol. 115, no. 1, pp. 48–56, 1993.

[13] J. S. Simon, L. Valavani, A. H. Epstein, and E. M. Greitzer, "Evaluation of approaches to active compressor surge stabilization," *Journal of Turbomachinery*, vol. 115, no. 1, pp. 57–67, 1993.

[14] 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.

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

[16] 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.

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

[18] 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] J. Du, Y. Zhang, M. Martínez García, and A. Spencer, "Review on compressor surge monitoring, modeling, and anti-surge control approaches," *Journal of Dynamics, Monitoring and Diagnostics*, vol. 4, no. 4, pp. 292–307, 2025.

[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] G. Torrisi, S. Grammatico, A. Cortinovis, M. Mercangoz, 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.

2.11 Exercises

2.1 Concept check · 2 min

With the default $B = 1.5$ and the anti-surge valve shut, the throttle is closed slowly from 70 %. Which stability condition of Equation (2.8) fails first, and where?

(a) The determinant condition, at the peak of the characteristic, 47.73 % throttle
(b) Both conditions together, at the peak, where the slope ψ_c' changes sign
(c) The trace condition, at $\varphi = 0.4769$ and 45.62 % throttle
(d) The determinant condition, at its minimum of 0.639 near 31.2 % throttle

2.2 Concept check · 2 min

In scenario 10, Aggressive PI tuning, the flow keeps oscillating long after the closure, and an operator reports that the machine is surging. Which single test tells a control-induced cycle from aerodynamic surge?

(a) Read the Peak ψ swing tile: any swing above zero means the machine surges
(b) Compare the IAE tile with scenario 4's: the larger IAE belongs to surge
(c) Close the throttle 5 points further: surge grows and a tuning cycle dies out
(d) Switch the loop to Manual: surge continues and a tuning cycle stops

2.3 Concept check · 2 min

The machine runs at 70 % throttle with the valve shut and the flow 0.0761 above the anti-surge setpoint. What does the integral state I of the override do?

(a) It decays toward zero with time constant τ_I, so the demand rests on $K_c e$
(b) It integrates the negative error without bound, so the valve opens late
(c) It is held just below the lower clamp, so the valve opens a little early
(d) It is frozen at its last value until the flow falls below the setpoint

2.4 Theory and calculation · 6 min

Use Equations (2.3) and (2.8) with the defaults of Table 2.2, so $B = 1.5$.

(a) At the healthy point, $\varphi^* = 0.6511$ and $\psi^* = 0.4414$, find ψ_c', the throttle-line slope $\psi_T' = 2\psi^*/\varphi^*$, $\operatorname{tr} J$ and $\det J$. *Checkpoint:* -1.700, 1.356, -3.042 and 2.254.

(b) At the trace crossing, $\varphi^* = 0.4769$ and $\psi^* = 0.5578$, check that $\psi_T' \approx 2.339$. With $\psi_c' = 0.1900$, show that the trace is zero to three decimals and find $\det J$. *Checkpoint:* $1.5 \times 0.1900 = 0.285$ and $1/(1.5 \times 2.339) = 0.285$, so $\operatorname{tr} J \approx 0$. $\det J = 0.919$.

(c) Find the frequency and the period of the linearized oscillation there, and compare the period with the 7.51 $\hat{t}$ of the deep-surge orbit. *Checkpoint:* 0.9585 and 6.56 $\hat{t}$. The orbit is about 15 % slower.

2.5 Theory and calculation · 6 min

Scenarios 3 and 4 close the throttle to 36 %, where $\gamma_T = 1.40 \times 0.36 = 0.504$.

(a) Show that $\varphi^* = 0.350$ is an equilibrium, $\varphi^* = \gamma_T \sqrt{\psi_c(\varphi^*)}$, on the positive-slope branch, and find ψ^*. *Checkpoint:* $\psi^* = \psi_c(0.350) = 0.4822$, and $0.504 \times \sqrt{0.4822} = 0.3500$.

(b) Find ψ_c', ψ_T' and $B_{crit} = 1/\sqrt{\psi_c' \psi_T'}$. Do $B = 1.5$ and $B = 0.5$ surge there? *Checkpoint:* 0.907, 2.756 and 0.633. $B = 1.5$ surges and $B = 0.5$ does not.

(c) With rotor speed and duct unchanged, by what factor must the plenum volume shrink to bring B from 1.5 down to 0.5? *Checkpoint:* A factor of 9.

2.6 Theory and calculation · 5 min

Scenario 1 holds the default point for 29.9 $\hat{t}$ with $e = -0.0761$ at every execution, $K_c = 8$ and $\tau_I = 4$.

(a) Find the demand c_k with the shipped back-calculation, and the IAE the scorecard shows at $\hat{t} = 29.9$. *Checkpoint:* -0.609 and about 2.28.

(b) A positional PI with no anti-windup sums the same error. Find its integral term and its demand at $\hat{t} = 29.9$. *Checkpoint:* -4.55 and -5.16.

(c) A closure then holds the flow at the surge line, $\varphi = 0.500$. How long does the demand of part (b) take to climb back to zero so that the valve can move? Compare with the deep-surge period. *Checkpoint:* About 26 $\hat{t}$, three and a half surge periods.

2.7 App challenge · 8 min

Load scenario 4, Same closure, controller off. Change nothing, run to $\hat{t} = 59.7$ at the default seed and export the CSV file.

(a) Read the Min margin and Peak ψ swing tiles and find the lowest flow from the margin. In the file find the first row with `in_surge` equal to 1 and the times of the minima of `flow_phi` after $\hat{t} = 15$. *Checkpoint:* $-135.1\,\%$, 0.631 and $\varphi = -0.175$. Alarm at $\hat{t} = 11.22$ and minima 7.51 $\hat{t}$ apart, both within 0.02.

(b) Load scenario 8, A smaller plenum (B = 0.5), run to $\hat{t} = 95.5$, and read the same two tiles and the last `flow_phi`. Then set Greitzer parameter to 1.5 in the Parameters tab, press Reset, run to $\hat{t} = 60$ and find the first row with `in_surge` equal to 1. *Checkpoint:* $-43.8\,\%$, 0.211, no alarm and 0.3207 within 0.0005. With $B = 1.5$ the alarm comes at $\hat{t} = 34.32$ within 0.02, and Min margin reads $-137.9\,\%$.

Explain why the run at $B = 0.5$ dips past the surge line but never trips the alarm.

2.8 App challenge · 8 min

Load scenario 5, Fast anti-surge valve, run to $\hat{t} = 29.9$ at the default seed and read the Min margin, Recycled flow and MV travel tiles. Then load scenario 6, Slow anti-surge valve (it does not make it), run to $\hat{t} = 59.3$, read the same tiles and export the CSV file.

(a) Report the three tiles for both scenarios. *Checkpoint:* Scenario 5: $+3.3\,\%$, 4.75 and 182. Scenario 6: $-105.4\,\%$, 11.23 and 387. Travel within 1.

(b) In the scenario 6 file find the first row where `asv_cmd` is above zero, the first where it reaches 100, the `asv_position` in that row, and the rows where `in_surge` turns on and off. *Checkpoint:* 9.02 and 11.02, with the valve about 34.5 % open (within 1). Alarm from 11.74 to 24.62, within 0.02.

Explain why no change of K_c or τ_I would save the scenario 6 closure.

2.9 App challenge · 8 min

Load scenario 11, MPC against the anti-surge PI (rapid closure), and use the default seed. For each run below go to $\hat{t} = 20$ and read the Min margin, Recycled flow and MV travel tiles.

(a) Run the scenario as loaded. *Checkpoint:* $+15.0\,\%$, 2.46 and 85.6.

(b) Under Advanced set Throttle signal to the MPC to Off, press Reset and run again. *Checkpoint:* +4.7 %, 2.27 and 168.

(c) Set Controller to PI, press Reset and run again. *Checkpoint:* −8.6 %, 2.11 and 529.

How much of the predictive controller's margin does the throttle signal supply? The PI given the same signal as a minimum valve opening through a high select, measured outside the app, holds +16.16 % to $\hat{t} = 20$ but recycles 8.6 % more than the MPC. What does the plan add?

2.10 Application area · 4 min

Scenario 7 closes the throttle to 30 % with an anti-surge valve of capacity $K_v = 0.18$.

(a) Holding $\varphi_{safe} = 0.575$ needs a total conductance $\gamma = \varphi_{safe} / \sqrt{\psi_c(\varphi_{safe})}$. Find it, the throttle's share, the valve's share, the smallest K_v that supplies that share wide open, and the opening that the standard $K_v = 0.60$ needs. *Checkpoint:* 0.7874, 0.420 and 0.3674. $K_v \approx 0.37$, and 61 % open.

(b) The command of scenario 7 sits at 100 % for 56.2 % of the run. What should an engineer conclude, and what are the remedies?

2.11 Application area · 4 min

After the challenge's closure to 38 % throttle the PI settles with the flow on its setpoint, so the recycle valve carries $\varphi_u = \varphi_{safe} - \gamma_T \sqrt{\psi_c(\varphi_{safe})}$.

(a) Find φ_u and the fraction of the compressor flow recycled for the setpoints 0.575 and 0.60. *Checkpoint:* 0.1865 and 0.2197, or 32.4 % and 36.6 % of the flow. The wider margin recycles about 18 % more.

(b) The model returns the recycle cold (Section 2.3.6). What does hot recycle change in the cost of part (a), and in the plant needed?

2.12 Review the agent · 5 min

An AI assistant reviews a run of scenario 6, Slow anti-surge valve, and replies: "The controller is too sluggish. Its command does not leave zero until 1.0 $\hat{t}$ after the closure, and the machine surges. Raise the proportional gain to 34 and cut the integral time to 0.8, as in scenario 10, and the valve will open in time." Find the flaw and say how you would check it.

CHAPTER 3

Limiters That Vote on One Fuel Valve in a Turbofan Simulator

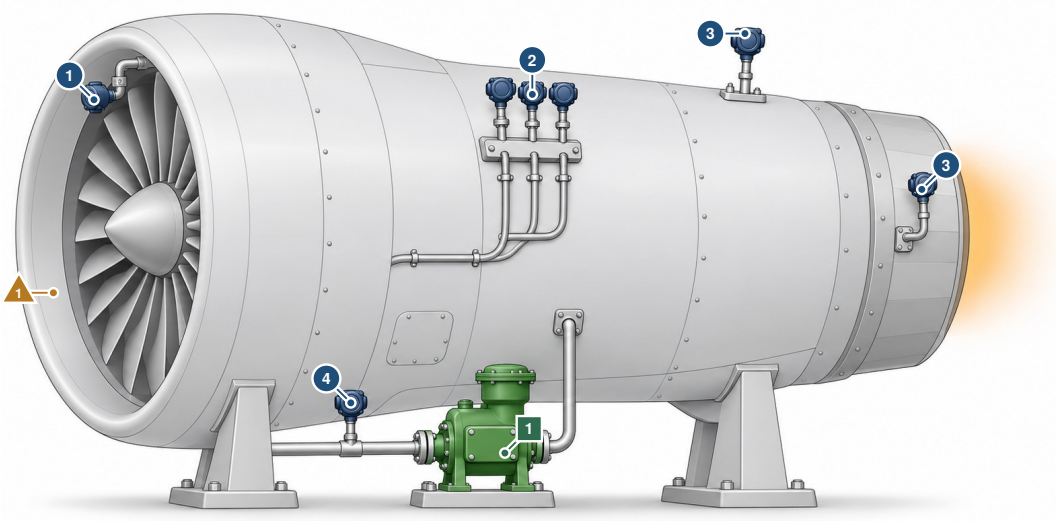

The controller moves 1 valve that meters fuel to the flame

The controller reads 1 inlet air temperature and pressure 2 fan and core speeds, compressor pressure 3 exhaust temperature, read at several points 4 fuel flow

What disturbs it 1 a hot day or thin air at altitude

Notice that every head reads speed, pressure, temperature or fuel, and that none reads how close the compressor is to surge, so the one valve is moved by limits on what can be measured. Illustration, not to scale.

IN PLAIN TERMS One Fuel Valve and Many Guards

The main idea. A jet engine is steered by one fuel valve. Several small controllers each ask for an amount of fuel, and the valve follows the most cautious request.

The challenge. Fuel added faster than the spinning blades can speed up pushes the engine toward surge, a violent backflow of air that nothing on the engine measures. A guard not in charge must keep its running total matched to the real fuel flow, or it takes over late and with a jolt.

Why it matters. These rules decide whether an engine reaches takeoff power in seconds without surging, overheating or losing its flame.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page, the fact on how controls grew (Section 3.1), the two history segments (Section 3.2), the two try-it segments and the fact on the challenge (Section 3.5) and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 3.5), Technology in Practice (Section 3.6), Limitations (Section 3.7) and the open challenge

The control question A jet engine has one handle, its fuel flow, and the fuel the five-second thrust rule demands drives the compressor toward a surge line that nothing on the engine measures. How does a min/max selector of a speed governor, schedules and limiters share that one valve, and what must every loop that is not in charge do so that it can take over cleanly?

Learning objectives

- Explain why a fuel step moves the compressor toward surge before the spools can speed up.
- Trace which cap or floor holds the valve through a slam, a chop and a stall, and why the floors come after the caps.
- Derive the back-calculation trail of a loop that is not in charge and the distance from its limit at which it takes over.
- Trade an acceleration schedule's design margin against the time to 95 % thrust, on a new and on a worn engine.
- Use corrected quantities to compare the engine at altitude and on a hot day with the engine at sea level.

Prerequisites Sections 1.4.3 (selectors and override), 1.4.4 (external reset), 1.4.5 (saturation and rate limits), 1.4.2 (back-off) and 1.6.4 (scenarios and exported files); Chapter 2 for surge on a compressor map. Section 1.6.3 is needed only for the predictive governor.

The app https://apmonitor.com/apps/turbofan-fadec/

See also Volume 7, Chapter 3 for a gain schedule that keeps one linearized closed loop across operating points, and Volume 7, Chapter 4 for a back-calculation PI under an actuator limit. In this volume, Chapter 9 for an exhaust-temperature limit that derates a gas turbine on a hot day.

3.1 Introduction

At flight idle the simulated engine burns 0.164 kg/s of fuel and makes 12.0 kN, and ten times that thrust is a lever stroke away. Fuel delivered faster than the two spools can speed up heats the gas at once, and because the first turbine nozzle is choked the compressor has to deliver the same flow at a higher pressure. In case 1 of the simulator, takeoff fuel set in one move reaches the surge line 0.020 s after the move, at 0.51 kg/s of fuel on its way through the valve, while the core speed has risen from 60.6 to only 60.8 %. The same engine, slammed under an acceleration schedule designed for an 8 % margin, reaches 95 % of its rated thrust in 4.74 s with the surge margin never under 9.5 %.

The engine controller that makes the second run possible is not one loop. A fan-speed governor asks for fuel, an acceleration schedule caps it so that the compressor does not surge, a turbine temperature limiter caps it on a hot day, overspeed limiters cap it at the top, and a deceleration schedule floors it so that the combustor does not blow out. The lowest cap and then the highest floor win, and every loop that does not win must keep its integrator on the fuel actually commanded or its hand-over overshoots. This min/max architecture with integral wind-up protection is the structure of the public generic engine controller of a NASA study [1], and engine control texts and surveys treat it at length [2–4]. Selector and override control of the same shape appears in process plants wherever one actuator protects several limits [5] (Section 1.4.3), and anti-windup by back-calculation is the standard remedy for an integrator that is not in charge of its output [6] (Section 1.4.4).

+ DID YOU KNOW? From One Controlled Variable to Eleven

A jet engine flight-tested in 1942 had one controlled variable and three sensors. A large commercial turbofan of 1993 had three controlled variables and seven sensors, and a military engine of 1989 had eleven of each [7].

One belief this engine is meant to dislodge is that a controller protects a limit by measuring it. Nothing on an engine measures its surge margin. The schedule protects it by holding a ratio of measured fuel to measured compressor delivery pressure below a table drawn against corrected core

speed, and a student who tightens that table finds the margin spent by how fast fuel arrives rather than by how much arrives. A second belief, that gain scheduling is what makes an engine controller work across its envelope, is also tested. In this engine a fixed governor gain is stable at every point tried, and the scheduled gains buy uniform damping (overshoot of 0 to 8 % on 2 % lever steps at nine points, against 7 to 21 % fixed) without being what passes the graded challenge.

Teaching material usually separates these pieces. Compressor stability is taught on lumped surge models with a throttle and no fuel [8], the kind of model Chapter 2 works with, gas turbine performance on steady cycle calculations [9, 10], and the public component-level engine simulations that pair a turbofan with its controller, NASA's C-MAPSS40k and the T-MATS toolbox [11, 12], are research codes run inside a desktop modeling environment. This chapter works with a generic turbofan and its min/max fuel selector as a browser simulator runs them, and measures on that simulator where the surge margin, the turbine temperature, the flame and the five-second rule meet in one valve. The selector, the ratio-unit schedule, back-calculation and corrected-speed scheduling are established practice, and none of the results is a research result about engines. They are measurements on one teaching engine and one seed, among them that the graded re-slam can be bought from the chop as well as from the slam. The engine is a component-level model on declared maps, written for this application, whose controller reads only declared measurements, and the chapter states every equation and constant needed to build it again.

Section 3.2 follows the domain, from hydromechanical fuel controls to the five-second certification rule. Section 3.3 sets out the engine, its maps, spools, plenums, instruments and assumptions. Section 3.4 gives the controller as implemented: the governor and its gain schedule, the caps, floors and selector, integrator tracking, stall recognition and, in Section 3.4.6, an optional predictive fuel governor. Section 3.5 measures the ten cases of the simulator, the challenge, the verification of the code and the predictive governor against the selector. Section 3.6 sets the selector on engine hardware, Sections 3.7 and 3.8 state the limitations and conclude, and Section 3.9 holds the exercises.

3.2 Fuel Schedules, Limiters and the Five-Second Rule

The control problem of a jet engine was settled in its present shape before electronics could solve it. A pilot's lever had to be turned into a fuel flow that brought the engine to the requested speed without stalling the compressor, overheating the turbine or blowing out the combustor, and the early answer was a hydromechanical fuel control in which cams, levers and servo valves computed the governor and the limits mechanically. Surveys of the field follow that lineage from those fuel controls to the full-authority digital electronic engine control that replaced them [2, 3]. The acceleration and deceleration limits are statements of physics, a fuel flow allowed for a given compressor pressure and core speed, so they mean the same thing whether a cam or a processor computes them.

> ⧖ HISTORY **A Governor With a Stop and a Cap, 1942 to 1981**
>
> The first jet engine built in the United States, in 1942, had a hydromechanical governor that metered fuel in proportion to the difference between the set and the actual turbine speed. A minimum-flow stop on the metering valve kept the flame from going out, and a maximum-flow schedule kept the engine from overheating. The first full-authority electronic engine control entered service in 1972. It was a single analog channel that computed every schedule and limit, backed by a hydromechanical control with conservative, fixed fuel-ratio maximum and minimum schedules. Flight tests of the first full-authority digital control began in 1981, and between 1990 and 2002 dual-channel digital controls became the standard for jet engines [7]. The floor and the cap of this chapter's selector were there from the start.

Compressor surge entered the control problem as a stability question. A lumped compression-system model of 1976, a compressor feeding a duct, a plenum and a throttle, showed that whether an installation falls into rotating stall or into surge depends on one dimensionless group [8], and the engine's operating point, its steady line and its surge line became the vocabulary of every acceleration

schedule since. Engine performance texts give the corrected quantities that make one schedule serve every altitude and day, speed over $\sqrt{\theta}$, flow times $\sqrt{\theta}/\delta$ and fuel over $\delta\sqrt{\theta}$ [9], and the dynamic modeling of gas turbines for control has its own book-length treatment [10].

> ⧖ HISTORY **A Speed Loop That Cycled at Altitude, Before 1950**
>
> The first afterburning turbojet, tested in 1948, had a speed governor designed by frequency response alone. In altitude tests at the national aeronautics laboratory, noise in the speed sensor together with the governor's high gain made the engine behave like a limit cycle. Engineers of the maker and the laboratory turned to step-response analysis, and the fix was to reduce the control gain at altitude. In the same years corrected flows, speeds and pressure ratios, derived from dimensional similarity, cut the number of operating points that a calculation by slide rule had to cover. The corrected approach is still used because it simplifies the design and implementation of control laws [7]. A gain that changes across the envelope and tables written in corrected quantities are both part of this chapter's controller.

The certification requirement that makes the schedule tight is old and short. In the text in force, last amended in March 1971, 14 CFR 33.73 asks in paragraph (b) for an increase from the fixed minimum flight idle lever position, or from not more than 15 % of rated takeoff thrust, to 95 % of rated takeoff thrust in not over 5 s, from a stabilized static condition with only the bleed and accessory loads needed to run the engine. Paragraph (a) asks that a lever moved from minimum to maximum in not more than 1 s, with the most bleed and power extraction to be permitted, cause no overtemperature, surge or stall [13]. Neither paragraph names an altitude. A schedule wide enough to be safe on a worn engine has to be narrow enough to meet paragraph (b) on a new one.

Research on the controller itself moved to public generic engines. NASA's propulsion controls research has been reviewed in its own right [14]. Its non-proprietary C-MAPSS40k simulation of a generic 40,000 lbf two-spool high-bypass turbofan carries a controller built on common industry architecture [11], and the design of that controller, a gain-scheduled proportional-integral governor with logic that protects the engine's limits, was published with its acceleration schedule and its response to the Federal Aviation Administration's thrust requirement [1]. The open T-MATS toolbox later made the thermodynamic building blocks of such models available for reuse [12]. Advanced methods of turbofan control, among them alternatives to the min/max selector for limit protection, are collected in a monograph [4], while gain scheduling as a design method has general surveys of its own [15, 16]. Volume 7, Chapter 3 schedules a gain to preserve one linearized closed loop across operating points. The selector with tracking integrators is the structure this application implements.

3.3 The Engine as Modeled

The simulated engine is a component-level teaching model of a generic 120 kN-class two-spool high-bypass turbofan, on a test stand or in flight, as shown in Figure 3.1. It was written for this application from textbook relations on closed-form maps declared in its source. It is not a port of C-MAPSS40k, T-MATS or any licensed engine simulation, and no map, table or coefficient is taken from one or from any engine maker. The fan's outer stream passes a bypass duct to its own nozzle, while its root and a booster on the low spool feed the high-pressure compressor (HPC), the combustor, the high-pressure turbine (HPT) that drives the HPC on the high spool, and the low-pressure turbine (LPT) that drives the fan and booster. Five plenum volumes sit between the components. Station numbers follow the usual convention, 2 the fan face, 13 the bypass duct, 25 the HPC inlet, 3 its exit, 4 the combustor exit, 45 between the turbines and 5 the LPT exit.

The variables are classified in Table 3.1. Ten states are integrated, the two spool speeds, the five plenum pressures, the fuel actually flowing through the metering valve, the thermocouple's reading and the bleed valve's position. Units in the model are seconds, kilograms, pascals, kelvin, watts and radians per second, and the interface shows percent of gauge speed, bar, degrees Celsius and kilonewtons.

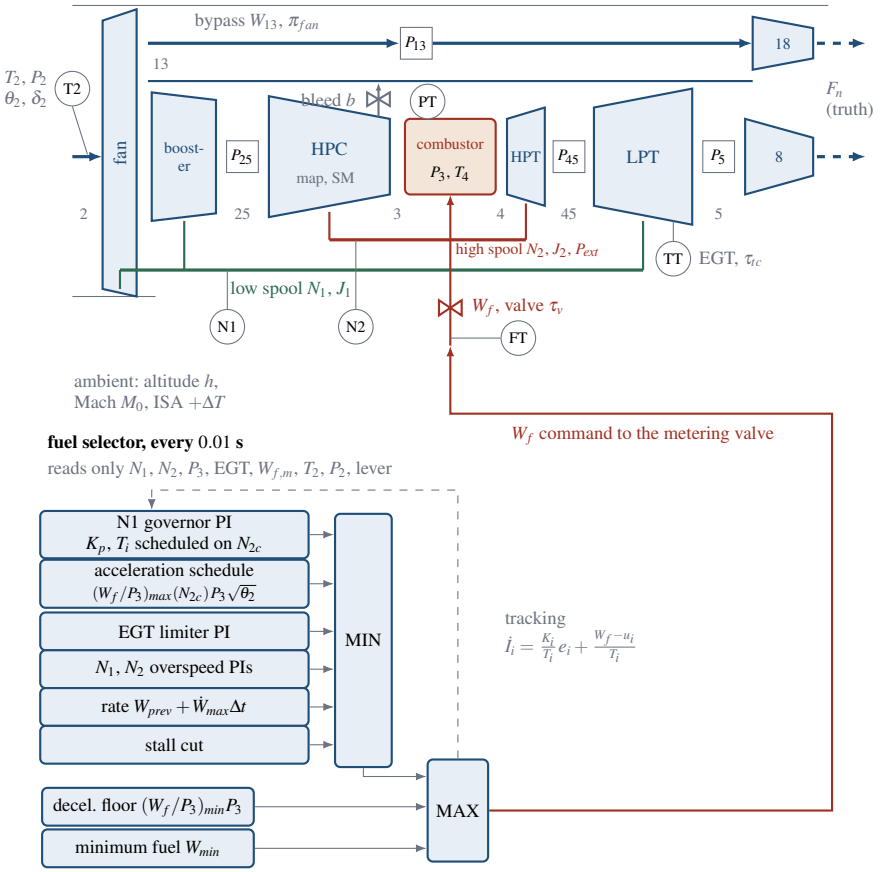

Figure 3.1 The engine as the Engine view draws it (top) and the fuel selector of the Ladder view (bottom). The fan and booster on the low spool (N_1, J_1) and the HPC on the high spool (N_2, J_2) feed five plenums, P_{13}, P_{25}, P_3, P_{45} and P_5. Fuel W_f through the metering valve sets T_4, the handling bleed b opens at the HPC exit, and the thrust F_n is truth. Circles are the measurement packet: speeds, P_3, the EGT thermocouple, the fuel flowmeter and the fan-face probes. Below, the minimum of six caps meets the maximum of two floors, the thick solid path is the fuel command to the valve, and the dashed path is integrator tracking. Process model only, not to scale.

3.3.1 Ambient Conditions and Corrected Quantities

The fan face sees the standard atmosphere shifted by the day offset ΔT at constant pressure, raised by the ram of the flight Mach number M_0 with a recovery of 0.995, as shown in Equation (3.1).

$$\begin{aligned} T_2 &= \left(T_{ISA}(h) + \Delta T\right)\left(1 + 0.2M_0^2\right) \\ P_2 &= 0.995\, p_{ISA}(h)\left(1 + 0.2M_0^2\right)^{3.5} \\ \theta_2 &= T_2/288.15\,\mathrm{K}, \qquad \delta_2 = P_2/101.325\,\mathrm{kPa} \end{aligned} \tag{3.1}$$

where T_{ISA} falls by 6.5 K/km to 11 km and is 216.65 K above it, and the flight velocity is $V_0 = M_0\sqrt{\gamma_a R T_0}$. Every map below is written in corrected speed $N/\sqrt{\theta}$ and corrected flow $W\sqrt{\theta}/\delta$ at its own inlet, so corrected similarity holds by construction except where a fixed offtake or the nozzles' exhaust to ambient pressure break it.

Table 3.1 Variables by Role, With Their Values at Flight Idle and Takeoff at Sea-Level Static ISA

Variable	Role	Idle, takeoff	Range
N_1	CV, fan speed[a]	34.5, 96.1 %	max 104
lever	SP, corrected N_1	0, 100 %	0–100
W_f	MV, fuel	0.164, 1.241 kg/s	0–1.6
N_2	state, core speed[a]	60.6, 97.2 %	max 105
EGT	limited, indicated	444, 872 °C	max 950
P_3	state, measured	5.8, 30.2 bar	—
SM	truth, HPC margin	22.8, 22.0 %	min 8
F_n	truth, net thrust	12.0, 120.2 kN	—
P_{13}, P_{25}, P_{45}, P_5	states, plenums	—	—
$W_{f,v}$, T_{tc}, b	states, valve, TC, bleed	—	—
h, M_0	DV, altitude, Mach	0, 0	12 km, 0.85
ΔT	DV, ISA offset	0 K	−30 to 40
d	DV, deterioration	0	0–100 %
W_{cust}	DV, customer bleed	0	0–8 %
events	DV, slam, chop, bodie, stall	—	—

[a] Percent of gauge speeds of 5100 and 14 500 rpm.

3.3.2 A Closed-Form Compressor Map

One closed form serves the fan's bypass stream, the fan root with the booster, and the HPC. With ψ the normalized pressure rise and n the corrected speed over its design value, the normalized corrected flow m is an explicit function of the pressure ratio, as shown in Equation (3.2).

$$\begin{aligned} \psi &= \frac{\pi - 1}{\pi_d - 1}, \qquad W_c = M_{ref}\, m \\ m &= n^a \left[1 - k\left(\frac{\psi}{S n^b}\right)^q\right], \qquad \psi_s(m) = S\left(\frac{m}{1-k}\right)^{b/a} \\ \mathrm{SM} &= \frac{1 + (\pi_d - 1)\psi_s(m)}{\pi} - 1 \end{aligned} \tag{3.2}$$

where each speed line runs from a choked top at $\psi = 0$ to its surge point $\left((1-k)n^a,\ Sn^b\right)$, so eliminating n gives the surge line ψ_s in closed form, and SM is the margin at constant corrected flow over the operating pressure ratio (some definitions divide by the surge line's ratio instead). For the HPC $a = 2.1$, $b = 2.9$, $k = 0.3$, $q = 3$ and $\pi_d = 12.5$, and $S = 1.0898$ is solved by bisection so that the design point has 22 % margin. At that point $m = 0.768$, the surge line sits at $\pi = 15.25$ and $M_{ref} = 35.33\,\mathrm{kg/s}$. The efficiency is a declared island, $\eta = \eta_{max} - e_1(n - n_e)^2 - e_2(r - 0.72)^2$ clipped to 0.30 to 0.95, with $r = \psi/(Sn^b)$ the fraction of the way up a speed line, shifted so that each component has its declared efficiency at design. Because the flow follows from the pressure ratio, the plenum pressures can be states and no map is ever inverted.

3.3.3 Compression, Combustion and Expansion

Each compression and expansion uses a constant-property isentropic relation with a map or turbine efficiency, and each shaft power is a flow times a temperature change, as shown in Equation (3.3).

$$\begin{aligned} T_{out} &= T_{in}\left[1 + \frac{\pi^{\kappa_a} - 1}{\eta_c}\right] \quad \text{(compressor)} \\ T_{out} &= T_{in}\left[1 - \eta_t\left(1 - \pi^{-\kappa_g}\right)\right] \quad \text{(turbine)} \\ P &= W c_p |T_{out} - T_{in}|, \qquad \eta_t = \eta_{t,d} - c_t(u-1)^2 \end{aligned} \tag{3.3}$$

where $\kappa_a = 0.2857$ for air with $c_{pa} = 1004.5\,\mathrm{J/(kg\,K)}$, $\kappa_g = 0.25$ for gas with $c_{pg} = 1148\,\mathrm{J/(kg\,K)}$, u is the blade-speed ratio $N/\sqrt{T_{in}}$ over its design value, and c_t is 0.35 for the HPT and 0.30 for the LPT. The combustor holds no gas. Its exit flow must pass the HPT, whose nozzle guide vanes are

choked above a pressure ratio of 2.2, and its exit temperature follows from an energy balance on one specific heat, as shown in Equation (3.4).

$$\begin{aligned} c_{pg} W_4 T_4 &= c_{pg} W_3 T_3 + \eta_b \,\mathrm{LHV}\, W_f, \qquad W_4 = W_3 + W_f \\ \frac{W_4\sqrt{T_4}}{P_4} &= \Phi_H\, g(\pi_H), \qquad g = \sqrt{1-x^2},\; x = \frac{2.2-\pi_H}{1.2} \\ \frac{W_L\sqrt{T_{45}}}{P_{45}} &= \Phi_L \sqrt{1-\pi_L^{-2}} \end{aligned} \tag{3.4}$$

where $P_4 = 0.95P_3$, $g = 1$ above $\pi_H = 2.2$, LHV is 43.0 MJ/kg, $\eta_b = 0.995$, and the LPT follows Stodola's ellipse law. With $G = \Phi_H P_4 g$ the first two lines close as one quadratic, $T_3 W_4^2 + (Q/c_{pg} - W_f T_3) W_4 - G^2 = 0$ with $Q = \eta_b \mathrm{LHV}\, W_f$, and then $T_4 = G^2/W_4^2$. This is the mechanism of the whole lesson. More fuel at the same core speed raises T_4, the choked HPT passes the same $W_4\sqrt{T_4}$ only at a higher P_4, and the HPC's point climbs its speed line. Fuel richer than a fuel-air ratio of 0.045 releases no more heat (a teaching cap that keeps T_4 in the flame range with a constant c_p). The exhaust gas temperature (EGT) is read a quarter of the way through the LPT's temperature drop, $T_{EGT} = T_{45} - 0.25(T_{45} - T_5)$, as many engines of the class measure it in the LPT's front stages, and the combustor blows out when W_f/P_3 falls under 9 g/s per bar.

Pause. Try Exercise 3.1 before reading on.

3.3.4 Spools, Plenums and Thrust

The two spools and the five plenums carry the dynamics, as shown in Equation (3.5).

$$\begin{aligned} J_1 N_1 \frac{\mathrm{d}N_1}{\mathrm{d}t} &= \eta_m P_{LPT} - P_{fan} - P_{boost} \\ J_2 N_2 \frac{\mathrm{d}N_2}{\mathrm{d}t} &= \eta_m P_{HPT} - P_{HPC} - P_{ext} \\ \frac{\mathrm{d}P_i}{\mathrm{d}t} &= \frac{RT_i}{V_i}\left(W_{in,i} - W_{out,i}\right),\; i \in \{13, 25, 3, 45, 5\} \\ F_n &= c_v \textstyle\sum_e A_e \left[\rho_e V_e^2 + p_e - p_0\right] - W_2 V_0 \end{aligned} \tag{3.5}$$

where N is in rad/s, e runs over the core and bypass nozzles, $W_2 = W_{13} + W_{25}$, $\eta_m = 0.99$, $P_{ext} = 150\,\mathrm{kW}$ is the accessory offtake, $R = 287.05\,\mathrm{J/(kg\,K)}$, and each volume V_i is sized to a 25 ms residence time at design (a real combustor is nearer 10 ms). The plenum outflows are the next component's map flow, and the HPC exit also loses the handling bleed $W_{bl} = bC_{bl}P_3/\sqrt{T_3}$ and the customer bleed on the same orifice law, sized so that the handling bleed passes 15 % of design core flow fully open. Both nozzles are convergent, choked or not, with $c_v = 0.985$ and pressure losses of 2 and 1 % upstream of them. The spool inertias $J_1 = 75\,\mathrm{kg\,m^2}$ and $J_2 = 10\,\mathrm{kg\,m^2}$ are geometric estimates for a 1.55 m fan, 24 blades and their disks on the low spool and nine HPC disks with one HPT stage on the high spool. The build's first guesses, 32 and 2.4 kg m^2, accelerated the engine from idle to takeoff in 1.2 to 1.5 s, faster than the class.

The metering valve, the thermocouple and the bleed valve are first-order lags, $\tau_v \dot{W}_{f,v} = W_f - W_{f,v}$, $\tau_{tc}\dot{T}_{tc} = T_{EGT} - T_{tc}$ and $\tau_b \dot{b} = b_{cmd} - b$, with $\tau_v = 0.05\,\mathrm{s}$, $\tau_{tc} = 1.0\,\mathrm{s}$ and $\tau_b = 0.3\,\mathrm{s}$.

3.3.5 Surge, Stall and Deterioration

A surge is the operating point reaching the surge line, SM ≤ 0, detected once at the end of a step. It starts a stall, and a stall started by an inlet distortion event behaves the same way. A stalled HPC delivers its flow only at a lower pressure rise, read from the map at $\psi/(1-d_s)$ with $d_s = 0.35$, its efficiency falls by $0.6d_s$, and the stall clears once the point has stayed below $r_u = 0.65$ of its speed line's height to surge for 0.2 s. These values were chosen from measured line positions: steady running sits at 0.90 to 0.92 of the way to surge, a stall held at constant fuel at 0.72, and fuel on the

Table 3.2 Declared Parameters. Fitted to the Class (F), Estimated (E) or a Teaching Value (T). All Below the Line Are Adjustable in the Parameters Panel

Symbol	Meaning	Value	Unit	
W_2, BPR	takeoff airflow, bypass ratio	355, 5.3	kg/s, —	F
π_{fan}, π_{bst}, π_{HPC}	design ratios	1.72, 2.40, 12.5	—	T
T_4	takeoff turbine inlet	1640	K	F
η fan, booster, HPC	design	0.89, 0.88, 0.85	—	T
η HPT, LPT	design	0.88, 0.90	—	T
$N_{1,d}$, $N_{2,d}$	design speeds	4900, 14,100	rpm	T
η_m, η_b, c_v	shaft, burner, nozzle	0.99, 0.995, 0.985	—	T
LHV, FAR_{max}	fuel, rich cap	43.0, 0.045	MJ/kg, —	T
fan a, b, k, q, SM_d	bypass map	1.05, 2.0, 0.3, 2.5, 18	—, %	T
booster, same	core map	1.30, 2.2, 0.3, 2.5, 20	—, %	T
HPC a, k, q	HPC map	2.1, 0.3, 3	—	T
π_{ch}, f_{EGT}	HPT choke, EGT plane	2.2, 0.25	—	T
τ_i	plenum residence	0.025	s	T
C_{bl}	bleed, fully open	15	% core	T
J_1, J_2	spool inertias	75, 10	kg m^2	E
P_{ext}	accessory offtake	150	kW	T
SM_d, b	HPC margin, exponent	22, 2.9	%, —	T
τ_v, τ_{tc}, τ_b	valve, TC, bleed lags	0.05, 1.0, 0.3	s	T
$(W_f/P_3)_{LBO}$	lean blow-out	9	g/s per bar	T
d_s, r_u	stall depth, clearing line	0.35, 0.65	—	T
wear, 100 %	HPC η, flow, surge, HPT η	3, 3, 4, 2	pt, %	T

deceleration floor with the bleed open at 0.64, so a hung stall and a recovered one are separated. Deterioration at 100 % removes 3 points of HPC efficiency, 3 % of its flow capacity, 4 % of its surge line's pressure rise and 2 points of HPT efficiency. The surge-line loss is a teaching value with no source. Without it, efficiency and flow loss alone moved the takeoff margin from 22.0 to 22.3 %, the wrong way for the lesson, and with it the margin falls to 17.8 % while the EGT rises 59.8 K at the same thrust. Variable stator vanes are folded into the HPC map on their nominal schedule, and an offset from it moves flow and the surge line by declared sensitivities of 1.2 and 1.0 % per degree.

3.3.6 Instruments

The controller receives a packet and nothing else. Speed pickups report N_1 and N_2 in percent with 0.02 % of colored noise, the thermocouple state carries 1.5 K, P_3 passes a 0.02 s lag with 0.2 % noise, the fuel flowmeter a 0.1 s lag, and the fan-face probes for T_2 and P_2 a 0.5 s lag. The noise is first-order filtered Gaussian, seeded from the case, so a run replays identically after Reset (Section 1.6.1). The surge margin, the thrust, T_4, the true EGT and the stall and flame states are truth, logged with a prefix and drawn dashed, and an automated check rejects any read of them in the controller's source.

3.3.7 Parameters, Initialization and Numerics

The declared values are collected in Table 3.2. A calibration at load sizes every flow capacity, nozzle area and volume so that the sea-level static takeoff point of the table is an exact steady state, and an independent hand calculation from the declared values reproduces the engine's 120.18 kN, 1.2405 kg/s of fuel, 10.322 g/(kN s), $T_3 = 855.6$ K, $T_{45} = 954.3$ °C and an EGT of 872.0 °C, with both nozzles unchoked at takeoff (core 525 m/s, bypass 307 m/s). Every run starts at a genuine steady state found by a damped Newton method on the seven gas-path derivatives with the fuel as an eighth unknown, for the lever's corrected N_1 in Auto, the hand-set fuel in Manual, or the EGT limit if it binds first. A start the corrected-similarity guess misses walks down the steady line from takeoff, and a start with no running steady state, or a step that turns non-finite, holds the run with a note saying it has left the model's range. The state then advances by fourth-order Runge-Kutta at $\Delta t = 0.005$ s, the controller runs every 0.01 s after a priming execution at $t = 0$, and a row is logged every 0.05 s, on the engine of Section 1.6.1.

Halving the step to 0.0025 s moves case 3's minimum surge margin at 8 % from 9.555 to 9.565 % and its time to 95 % from 4.745 to 4.750 s. Headless, a simulated second costs 1.0 ms, and the browser's fastest setting, five times real time, plays the 20 s challenge in 4.0 s.

3.3.8 Assumptions

The specific heats of air and gas are constant, the plenums isothermal and the combustor without volume or combustion delay, and there is no heat soak into casings and disks, no tip-clearance transient, no Reynolds-number effect at altitude and no cooling air. Each of these changes real transient margins by a point or more. The fan and booster have maps but no surge model, so only the HPC can stall. The stall is a declared loss of pressure rise, not a post-stall model of rotating stall or deep surge. The lever's flight idle is a corrected fan speed and the minimum fuel a fixed 0.06 kg/s, both sized for the low, slow part of the envelope, and there is no minimum core speed, so in fast, thin air at idle the core windmills and the run holds with a note. Every limit, the 8 % margin floor, the 950 °C EGT line, the overspeed limits of 104 and 105 % and the blow-out ratio, is a teaching value, and a real engine's come from its type certificate data sheet and its maker's operating limits.

3.4 The Controller as Implemented

The controller executes every $\Delta t_c = 0.01$ s from the measurement packet alone and emits one fuel command. Its loops are written in the application rather than taken from the shared velocity-form PID of Section 1.6.2, so that the tracking can be switched off and read. Each proportional-integral (PI) loop has no derivative term.

3.4.1 Power Management and the Governor

The lever sets a corrected fan-speed target from a table that makes sea-level static thrust linear in the lever, from 34.5 % corrected N_1 at flight idle to 96.08 % at takeoff, and the physical target is $N_{1t} = N_{1c}(\text{lever})\sqrt{\theta_2}$ with θ_2 from the fan-face probe. The governor's demand is $u_{gov} = Ke + I$ with $e = N_{1t} - N_{1,m}$. Its default gain is 0.2 kg/s of fuel per percent of N_1 with an integral time of 1.0 s. Scheduled on corrected core speed $N_{2c} = N_{2,m}/\sqrt{\theta_2}$, the gain becomes $K = K_p g_K(N_{2c})$ and the integral time $T = T_i g_T(N_{2c})\sqrt{\theta_2}/\delta_2$. The factors come from the eigenvalues of the seven-state Jacobian of the gas path along the steady line, g_K the design point's high-frequency gain $K_{ss}|\lambda_1|$ over its local value and g_T the local spool time constant $1/|\lambda_1|$ over the design one, with $K_{ss} = \mathrm{d}N_1/\mathrm{d}W_f$ and λ_1 the slowest eigenvalue. Over N_{2c} from 56 to 104 %, g_K runs from 0.55 to 1.11 and g_T from 7.48 down to 0.87. A first schedule built on the steady gain alone made the loop slower without damping it better, because K_{ss} varies sevenfold along the line while $K_{ss}|\lambda_1|$ varies only between 44 and 95.

3.4.2 Caps, Floors and the Selector

Each cap asks for at most a fuel and each floor for at least one, and the selector combines them as shown in Equation (3.6).

$$
\begin{aligned}
W_{acc} &= R_{acc}(N_{2c}, \mathrm{SM}_{acc})\, P_{3,m}\sqrt{\theta_2} \\
W_{dec} &= R_{dec} P_{3,m} \\
W_{rate} &= W_{f,k-1} + \dot{W}_{max}\Delta t_c \\
\bar{W} &= \min\left(u_{gov}, W_{acc}, u_{EGT}, u_{N1}, u_{N2}, W_{rate}, W_{stall}\right) \\
W_{f,k} &= \mathrm{clip}_{[0,\,1.6]}\left(\max(W_{dec}, W_{min}, \bar{W})\right)
\end{aligned}
\tag{3.6}
$$

where R_{acc} is the acceleration schedule's table of $W_f/(P_3\sqrt{\theta_2})$ against N_{2c} for a design surge margin SM_{acc} (interpolated between tables for 4 to 20 %), $R_{dec} = 13$ g/s per bar is 44 % above the blow-out ratio, $W_{min} = 0.06$ kg/s, $\dot{W}_{max} = 2$ kg/s per second, W_{stall} is $-\infty$ while a stall recovery cuts the fuel and $+\infty$ otherwise, and 1.6 kg/s is the valve's full flow. The EGT limiter is a PI on $950 - \mathrm{EGT}_m$ at

6 g/s per kelvin and 1.0 s, and the overspeed limiters are PIs on 104 and 105 % $-\,N_m$ at 0.2 kg/s per percent and 0.5 s, scheduled with the governor. The floors come last, because keeping the flame lit outranks every limit a cap protects, the order of floors after caps of Section 1.4.3. The fuel rate limit is one of the caps, inside the selector where every PI loop sees it, one of the placements Section 1.4.5 compares. The loop whose demand wins holds the valve, and the Ladder view lights it.

The schedule's tables were designed on the nominal engine at sea-level static ISA, as the ratio that puts the HPC at the design margin with the spools held at their steady speeds and the pressures solved. At 8 % the ratio rises from 38.5 to 56.0 g/s per bar over N_{2c} from 54 to 104 %. The table is corrected as fuel over $\delta_2\sqrt{\theta_2}$ divided by pressure over δ_2, which is why $\sqrt{\theta_2}$ multiplies it. A usability review found the first, uncorrected form wanting. On W_f/P_3 alone, a 70 to 80 % lever step at 10 668 m and Mach 0.8 dipped the margin to 6.5 %, under the floor, and the corrected table keeps 10.1 to 10.3 % there. A technical review then confirmed θ_2 against θ_3, since at a corrected operating point every station's θ scales with θ_2, T_3 is not measured, and a θ_3 correction would loosen the cap by up to 1.0 % during case 3's slam, the wrong way.

3.4.3 Integrator Tracking

With tracking on, every PI's integrator is pulled toward the fuel actually commanded by back-calculation [6], with the tracking time equal to its integral time, as shown in Equation (3.7).

$$\begin{aligned} u_{i,k} &= K_i e_{i,k} + I_{i,k} \\ I_{i,k+1} &= I_{i,k} + \left[\frac{K_i}{T_i} e_{i,k} + \frac{W_{f,k} - u_{i,k}}{T_i}\right]\Delta t_c \end{aligned} \tag{3.7}$$

where i is the governor, the EGT limiter or an overspeed limiter. The loop that holds the valve has $u_i = W_f$ and sees no correction. For any other the $K_i e_i$ terms cancel, so its integrator follows the commanded fuel one integral time behind it, $\dot{I}_i = (W_f - I_i)/T_i$, and its demand rests at its proportional distance $K_i e_i$ above the fuel less that trail. A limiter therefore takes over from the fuel in the line, bumplessly, and only as it nears its limit. This is the external reset of Section 1.4.4 with $T_t = T_i$, Equation (1.4), applied to every loop of the selector, and Volume 7, Chapter 4 uses the same back-calculation on a slip loop under a torque limit. With tracking off each integral runs on its own error, $I_i \leftarrow \mathrm{clip}_{[0,1.6]}(I_i + K_i e_i \Delta t_c / T_i)$, and a loop that does not hold the valve winds up against that clamp. At $t = 0$ and in Manual the governor's integrator is set to $W_f - Ke$ and every limiter's to W_f, so a return to Auto is bumpless. The plan of the application asked for every integrator to be reset each scan to the applied fuel, the one-scan limit of this law, but with that reset the EGT limiter won the valve on its rate of approach 215 K under its limit on a sea-level slam, while longer tracking times of 1.5 and 2 integral times let the challenge's EGT overshoot to 973 and 1012 °C.

3.4.4 Bleed, Stall Recognition and Recovery

The handling bleed opens fully below 55 % corrected core speed and shuts 5 % above it. A stall is recognized from measurements alone: $P_{3,m}$ down 6 % or more over the last five scans (50 ms) while the indicated EGT has not fallen by more than 1 K over the same window, with $N_{2,m}$ above 25 %. A stall drops P_3 and heats the gas, while a chop, a deceleration or a flameout drops P_3 because the gas cooled first. The recovery holds the fuel on the deceleration floor for 1 s and the bleed open for 3 s, a timed sequence that a second stall restarts and that never latches.

This rule replaced one that asked for the fuel command not to have been cut over the window. The earlier rule recognized 1 of 30 stalls injected during case 3's slam and left 29 still stalled 8 s later, because the ratio-unit schedule cuts the fuel with P_3, and it read case 5's chop without its floor as a stall 0.055 s after the chop. The present rule recognizes and clears all 30 mid-slam stalls and makes no false recognition in 936 ordinary maneuvers (four altitudes, three days, two Mach numbers, four schedule margins, both trackings, five lever moves). Its blind spot follows from the signature it uses. A stall within 0.3 s of a lever pull from 100 to 80 %, or within 0.7 s of a pull to 50 %, arrives while

the gas is still cooling, goes unrecognized and hangs. After a chop to idle such a stall is also missed but clears by itself as the operating point falls down its speed line.

3.4.5 Three Models

The plant is Section 3.3. The identified model is the linearization along the steady line that sets the gain schedule. The controller's model is its tables, designed once on the nominal engine at sea-level ISA, so a worn engine, a hot day or a changed map parameter meets schedules that do not know it.

3.4.6 Predictive Fuel Governor

The optional governor solves a fuel plan subject to the same current limiter votes, while the rate limiter, stall cut and flame floors still act every controller scan. This is receding-horizon optimization in the sense of industrial MPC [17], set up as in Section 1.6.3, not a replacement for the safety selector. Only the measured N_1, N_2 and indicated EGT enter the model. Three local fuel-to-output paths have the form

$$G_i(s) = \frac{K_i}{(1+\tau_{i1}s)(1+\tau_{i2}s)}. \tag{3.8}$$

At 0.6 kg/s, symmetric 0.003 kg/s steps on the nominal sea-level engine give $(K_i, \tau_{i1}, \tau_{i2}) = (53.725, 0.714, 0.200)$ for N_1, $(28.060, 0.433, 0.093)$ for N_2, and $(398.957, 0.347, 0.119)$ for EGT. Gains use percentage points or kelvin per kg/s, with times in seconds. After 30 seconds settling and 20 seconds observation, RMS fit residuals are 0.000635 and 0.000420 percentage points and 0.00987 K, respectively. The separate idle experiment has much larger gain and lag, so these paths are retained as one local nominal model rather than an envelope model.

Six states represent the two lags of each path, anchored at the first measured packet and known command. Each 0.01 second scan propagates the model and filters the measured-minus-predicted residual with gain 0.05 on each output. This bias adaptation is distinct from proving offset-free operation under general mismatch [18]. No true temperature, surge margin, thrust, plant state or changed kinetic or map parameter enters the prediction. Plans execute every 0.05 seconds with 40 predictions over 2 seconds and eight graduated fuel blocks. The cost is

$$\begin{aligned} J = {} & \sum_{k=1}^{40} (N_{1,k} - N_{1,\text{target}})^2 + 100 \sum_j (\Delta W_j)^2 \\ & + 10^3 (s_1^2 + s_2^2) + 0.1 s_E^2 + 10^3 (s_1 + s_2) + 10 s_E. \end{aligned} \tag{3.9}$$

Here $s_1, s_2, s_E \geq 0$ soften the predicted speed and indicated-EGT ceilings when their corresponding limiters are enabled. Each slack is shared over its full horizon. Current limiter votes form a hard fuel box, with the shipped floor precedence if a cap conflicts with a flame floor. The physical valve remains bounded between 0 and 1.6 kg/s. Fuel-block changes are bounded by the requested rate times 0.05 seconds, with the existing 0.01 second rate selector still above the result. The governor tracks applied fuel exactly. Limiter PIs retain their original back-calculation law, always enabled under MPC, because their proportional distance from a limit gives the safety vote its meaning. Failed, nonoptimal, nonfinite, malformed or out-of-box plans restore the tracked classical governor.

3.5 Ten Cases, Measured

The figures of this section are headless runs at seed 20261027, the default, stepped exactly as the browser steps the engine, each case following its own instructions. Settings named in a row are entered before Run, so they are part of the starting state, and a slam is the lever moved to 100 % at the time named. The time to 95 % runs from the slam to 95 % of the rated thrust at that condition, the nominal engine's steady takeoff thrust or, if its EGT would pass 950 °C first, its thrust at that line. The minimum surge margin is the truth's, with stalled steps left out because a stall is counted instead. Each case runs the same whether it is opened first, picked from the menu with another case

Table 3.3 Ten Scenario Comparisons and Their Completion Criteria

#	Case and setting	Outcome as measured	Required to complete (kind)
1	Manual at flight idle, fuel by hand	1.2 kg/s in one move surges 0.020 s later, ramps of 0.1 and 0.2 kg/s per second dip to 8.7 and 2.0 %, 0.3 surges	set the fuel to 0.8 kg/s or more in one move, watch until SURGE, confirm (demonstration)
2	Governor, schedule off, slam	surge 0.125 s after the slam, 9 surges in the next 10 s, N_1 down to 19.5 %	slam with the schedule off, watch until SURGE, confirm the governor's demand (demonstration)
3	Schedule at 10 %, slam	5.37 s at 11.3 %, at 8 % 4.74 s at 9.5 %, tracking off N_1 peaks at 97.98 %	slam at 10, then at 8 % within 5 s above 8 %, then with tracking off see N_1 pass its target (experiment)
4	ISA + 30 K, slam	EGT limiter from 3.94 s, 111.3 kN against 120.2 kN at ISA	slam, watch the EGT limiter hold the valve 3 s, confirm the thrust (experiment)
5	Takeoff, chop, floor off	flameout 0.105 s after the chop, with the floor N_1 near idle at 3.26 s	see FLAMEOUT, then with the floor on chop to within 2 points of idle lit (experiment)
6	10 668 m, Mach 0.8, lever 70, gains fixed	N_1 77.6 % against 83.9 at sea level, 14.6 kN against 87.7, 70 to 80 % overshoots 0.26 points fixed, 0.01 scheduled	step 70 to 80 % fixed and read the Overshoot tile, repeat scheduled (experiment)
7	Takeoff, inlet stall	recognized 0.025 s, cleared 0.39 s, back to 95 % in 4.88 s, recovery off hung at 89.5 kN	press Stall, see it clear and the thrust back to 95 %, then with the recovery off see it hang 5 s (experiment)
8	ISA + 15 K, wear 100 % at 2 s	EGT limiter 0.59 s later, 110.0 kN, 91.5 % of the rating, SM 21.9 to 18.4 %, at ISA + 30 K 86.6 %	set wear to 100 %, watch the limiter trim the thrust under 95 % for 5 s, confirm (demonstration)
9	Predictive fuel governor, sea-level slam	MPC 5.21 s at 9.69 % margin, peak EGT 884.3 °C	select PI selector once, return to MPC, run the 15 s comparison (experiment)
10	Challenge: bodie at 3048 m, ISA + 30 K	shipped fails (8.49 s, 996 °C), suggested passes (5.86 s, 9.4 %, 938 °C)	meet the target of Table 3.5 at 20 s with the parameters and disturbances as given (challenge)

Table 3.4 Measured Slams. Time to 95 % of the Rating, Smallest True Surge Margin, and Peak Indicated (True) EGT

Case	Setting	t_{95} (s)	SM (%)	EGT (°C)
2	schedule off	—	surge	709 (1525)
3	6 %	4.30	7.8	902 (1027)
	8 %	4.74	9.5	899 (1022)
	8 %, tracking off	4.66	9.5	927 (1060)
	10 % (shipped)	5.37	11.3	896 (1011)
	12 %	6.27	13.0	889 (1005)
	14 %	7.55	14.8	879 (985)
4	ISA + 30, tracking on	6.25	9.3	955 (1020)
	ISA + 30, tracking off	4.75	9.3	996 (1140)
7	stall at takeoff, recovery on	4.88[a]	12.7	905 (1080)

[a] From the stall, not a slam. Case 2's figures are over the 10 s after its first surge.

already loaded, or restarted by Reset. The cases are listed in Table 3.3 and the measured comparisons in Table 3.4, which Figure 3.2 redraws as bars against the three limits.

3.5.1 Fuel Faster Than the Spools

Case 1 is the plant without a controller. Takeoff fuel in one move surges the compressor 0.020 s after the move (0.018 to 0.027 s over 20 instants of the move), and the path on the map climbs nearly straight up its speed line. In case 2's slam the corrected flow even falls, from 12.9 to 11.7 kg/s, while the pressure ratio rises from 4.98 to 5.48, as shown in Figure 3.3. Added slowly, 0.1 kg/s each second, the fuel lets the spools keep up and the margin dips only to 8.7 %, while 0.2 kg/s each second dips it to 2.0 % and 0.3 kg/s surges. The margin is spent by the rate of fuel. Case 2 gives the same engine its governor and no schedule. At the slam the governor's error is 62 points of N_1, so only the 2 kg/s-per-second rate limit stands between its demand and the valve, and 0.33 kg/s has passed the valve when the path reaches the surge line 0.125 s after the slam (0.123 to 0.132 s over 20 instants). Run on, the stall recovery ends each surge and the governor causes the next, 9 in 10 s, while N_1 falls to 19.5 % because each recovery puts the fuel on the floor.

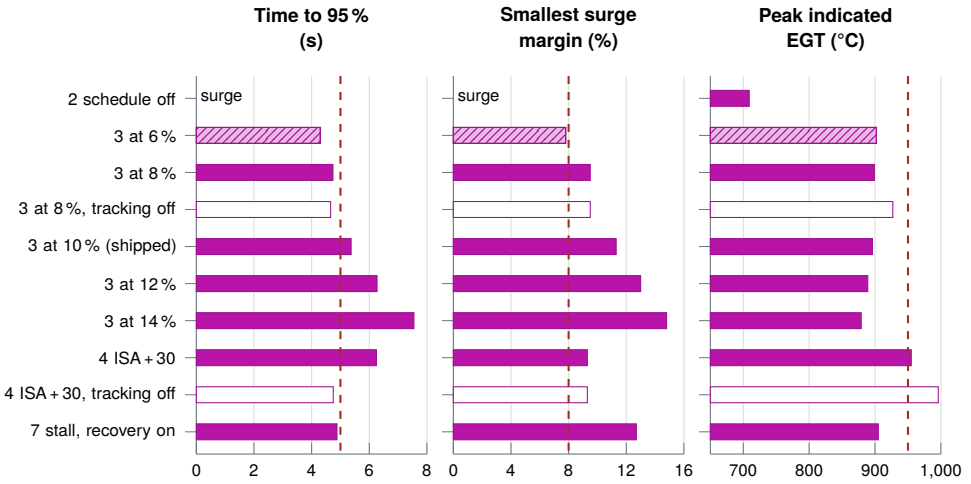

Figure 3.2 Table 3.4 redrawn. Dashed lines mark the 5 s of paragraph (b), the 8 % margin floor and the 950 °C indicated limit, and case 7 is timed from its stall. White bars are runs with integrator tracking off, and the hatched bar is the 6 % schedule, under the floor. Notice that the design margin trades time for margin row by row, that only the 8 % schedule meets both lines of a sea-level slam, and that switching tracking off shortens the time but raises the peak indicated temperature, to 996 °C on the hot day.

> **TRY IT IN THE APP** **1 · Fuel by Hand: Slam It (Start Here)**
>
> Load the scenario, press Run on the Map view and find the operating point below the surge line. Set Fuel flow to 1.2 kg/s in one move and watch the dashed path. It climbs nearly straight up its speed line and reaches the surge line 0.020 s after the move, while the core speed has hardly changed.

3.5.2 The Schedule Against the Five Seconds

Case 3 ships with the schedule designed for 10 %. The path hugs the margin line, the schedule holds the valve for 5.18 s, and 95 % comes 5.37 s after the slam (5.36 to 5.39 s over the slam instant and three seeds), too late for paragraph (b). At 8 % the schedule holds for 4.40 s, 95 % comes at 4.74 s, and the margin never falls under 9.5 %, a little above the line because the low spool lags the high one, as shown in Figure 3.3. Over 40 slam instants the time stays between 4.735 and 4.760 s. A 6 % schedule would take 4.30 s but dip to 7.8 %, under the floor, while 12 and 14 % take 6.27 and 7.55 s, as listed in Table 3.4. The hand-overs are continuous. In the 8 % slam the valve passes from the governor to the rate limit, the schedule, the EGT limiter and back to the governor, and the largest step in the command at any hand-over is 0.020 kg/s, one scan of the rate limit.

Near the top of that slam the EGT limiter holds the valve for 0.43 s, taking it 4.43 s after the slam with the gauge at 861 °C, 89 K under its line. This looks like a fault and is not one. The gauge is climbing 136 K/s on the EGT card and lags the gas by 1 s, and the true EGT is already 1021 °C. With tracking the limiter's integrator trails the rising fuel by about one integral time, and the limiter's demand meets the fuel when its proportional term has shrunk to that trail, 0.539 kg/s against a trail of 0.534 kg/s at the takeover. The limiter is protective by rate through the back-calculation trail, and it slows the last of the acceleration. With tracking off, the same 8 % slam reaches 95 % sooner, at 4.66 s, but N_1 overshoots to 97.98 % against its 96.08 % target, because the governor's integrator wound up while the schedule held the valve. With tracking on it peaks at 96.10 %. Figure 3.4 shows the application at the 95 % pause of the 8 % slam with tracking on.

Pause. Try Exercise 3.2 before reading on.

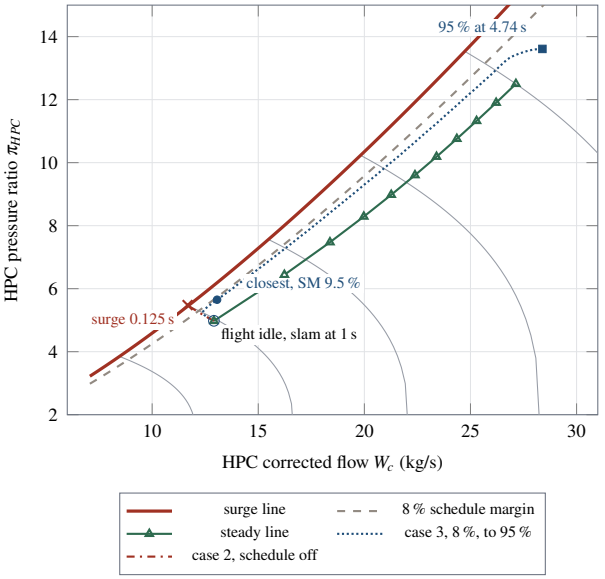

Figure 3.3 HPC map from the application's own map functions, with two runs at the default seed. Speed lines at $n = 0.6$ to 1.0 (thin gray), the surge line (very thick solid), the 8 % schedule-margin line (dashed) and the steady line at sea-level ISA (solid with triangles). Case 2 (schedule off, dash-dotted, its surge marked by a cross) climbs its speed line to the surge line 0.125 s after the slam, and case 3 (schedule at 8 %, dotted) runs along the margin line to 95 % thrust.

3.5.3 Heat, Chop, Altitude, Stall and Wear

On an ISA + 30 K day at sea level the rating is EGT-limited at 111.2 kN. In case 4 the EGT limiter first holds the valve 3.94 s after the slam and keeps it, and 20 s after the slam the engine makes 111.3 kN against 120.2 kN at ISA, 7.4 % less, with N_1 at 97.7 % short of its 101.0 % target. The indicated EGT peaks at 955 °C while the gas reaches 1020 °C, and with tracking off the limiter's wound-up integrator lets the indication overshoot to 996 °C. The flat rating ends where the takeoff lever's EGT reaches the line, at ISA + 20.0 K. Chapter 9 meets the same exhaust-temperature limit and ambient derating on a constant-speed gas turbine.

> **TRY IT IN THE APP** **4 · A Hot Day: The EGT Limiter**
>
> Run the scenario, slam the Thrust lever to 100 % and watch the Ladder view. The EGT limiter first holds the valve 3.94 s after the slam and keeps it, and the engine makes 111.3 kN instead of 120.2 kN. Then set the lever back to 0 %, set Integrator tracking Off, press Reset, Run and slam again: the indicated EGT now overshoots to 996 °C.

In case 5 the governor answers a chop from takeoff by cutting to the 0.06 kg/s minimum, which is 2 g/s per bar at takeoff pressure. With the floor off, the ratio crosses the blow-out line 0.105 s after the chop (0.103 to 0.112 s over 20 instants). With the floor on, the fuel follows P_3 down at 13 g/s per bar, the ratio never falls under 13.2, and N_1 is within 2 points of idle 3.26 s after the chop.

Case 6 flies the lever at 70 % at 10 668 m and Mach 0.8. The lever asks for the same corrected speed, 83.9 %, and the cold fan face ($\theta_2 = 0.86$) brings the physical N_1 down to 77.6 %. The thrust is 14.6 kN against 87.7 kN at sea level, the EGT 588 against 748 °C and the fuel 0.267 against 0.848 kg/s. At the same corrected N_1 of 95 % every corrected quantity matches sea level within 0.32 % at 3048 m and 0.19 % on an ISA + 30 K day, and within 2.06 % at 10 668 m, mostly because the fixed 150 kW offtake is a larger share of the smaller core power there (0.41 % without it). A step from 70 to 80 %

Table 3.5 The Challenge Over 20 s. Target: 95 % Within 7.0 s, SM at Least 8 %, Indicated EGT at Most 950 °C

Settings	t_{95} (s)	SM (%)	EGT (°C)	Verdict
shipped (fixed, 14 %, tracking off)	8.49	14.8	996	fails
8 % schedule alone	5.17	9.4	1019	fails
tracking on alone	8.55	14.8	934	fails
8 % + tracking, scheduled (suggested)	5.86	9.4	938	passes
8 % + tracking, fixed gains	5.86	9.4	936	passes
6 % + tracking	5.44	7.7	936	fails
10 % + tracking	6.46	11.2	935	passes
12 % + tracking	7.30	13.0	935	fails
shipped + tracking + floor 20 g/s per bar	6.09	15.0	935	passes

overshoots by 0.26 points of N_1 with fixed gains, 6 to 7 % of the step, and by 0.01 points scheduled. Both are stable. At the shipped gain, 2 % lever steps at nine points from sea level to cruise overshoot 7 to 21 % fixed and 0 to 8 % scheduled, so what the schedule buys is uniform damping. The corrected schedule keeps the margin of lever steps at 8.84 % or more over 7 altitudes, 3 Mach numbers and 5 days, and of slams from flight idle at 8.1 % or more up to 9144 m. Above that the envelope has a corner. At 10 668 m and higher with Mach 0 to 0.4, slams from idle fall to 0 to 7.3 % or surge, identically with and without $\sqrt{\theta_2}$ at Mach 0, because the 0.06 kg/s minimum fuel there exceeds the flight-idle fuel and, as a floor, outranks the schedule. That corner is a test-cell condition outside every case, and it shows that a floor outranks every cap. In the same sweep the uncorrected table took 115 of 236 transients under 8 %.

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

In case 7 an inlet distortion stalls the HPC at takeoff. The controller recognizes it 0.025 s later, the fuel goes to the floor and the bleed opens, the stall clears 0.39 s after it began, the thrust dips to 38.1 kN, and the schedule brings it back to 95 % of the rating 4.88 s after the stall. With the recovery off the EGT limiter holds the indication at 951 °C but the compressor stays stalled, and 20 s later the engine makes only 89.5 kN with N_1 at 85.5 %.

Case 8 wears the engine fully at 2 s on an ISA + 15 K day. The EGT limiter takes the valve 0.59 s later and trims the thrust to 110.0 kN, 91.5 % of the rating, while the margin falls from 21.9 to 18.4 %, and nothing on the flight deck but a lower N_1 says so. On an ISA + 30 K day the worn engine settles at 96.3 kN, 86.6 % of its 111.2 kN rating, and that variant is infeasible by design. A pass would hold 95 % of the rating for 2 s with the indicated EGT never over 950 °C in 25 s. None of 65 logged settings in the application's tests (16 automatic, 13 hand-set fuels, 36 hand pulses) passes, nor do 1,581 more drawn over the controller's settings and hand-set fuels, and the closest reaches 86.5 % of the rating.

3.5.4 The Challenge

The challenge starts at takeoff at 3048 m on an ISA + 30 K day, chops the lever at 2 s and slams it back 3 s later (a bodie). Over 20 s it asks for 95 % of the 82.6 kN rating within 7.0 s of the second slam, the true margin never under 8 %, no surge or flameout, and the indicated EGT never over 950 °C. It ships with fixed gains at the scheduled gains' idle values (0.16 kg/s per percent, 5.5 s), a 14 % schedule and tracking off. Table 3.5 gives the verdicts, whose verdicts repeat at seeds 1 and 2 with every figure within 0.04 s, 0.1 % and 4 K.

Two settings pass it together and neither alone. The 8 % schedule alone makes 5.17 s but lets the indication reach 1019 °C, because the EGT limiter's integrator wound up while the engine was cold, and tracking alone holds 934 °C but takes 8.55 s. The gain schedule makes no difference here (5.86 s either way), because the acceleration schedule holds the valve until 95 %. There is a second way through. Tracking with the deceleration floor raised to 20 g/s per bar passes on the shipped 14 % schedule in 6.09 s, because the floor keeps fuel in through the chop and the re-slam starts from 28.5 %

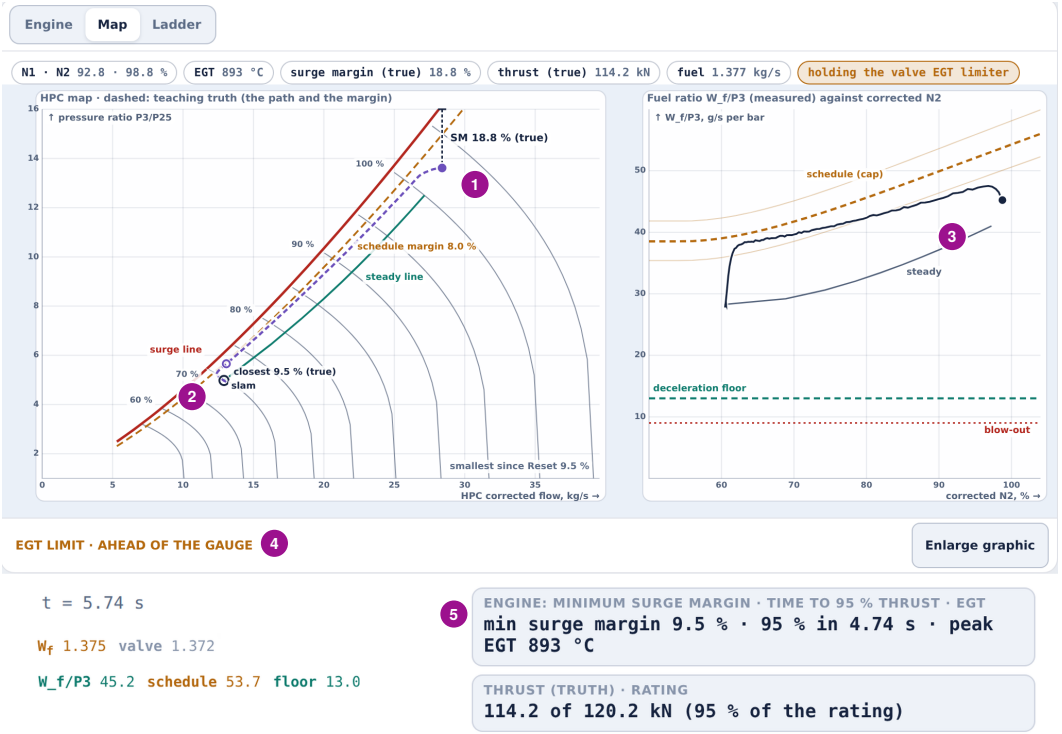

Figure 3.4 Case 3 with the schedule designed for an 8 % margin, at its 95 % pause on the Map view, $t = 5.74$ s after a slam at 1.0 s, in a composite of the application window: the plant view, with the run clock, two trend labels and two scorecard tiles set below it. On the HPC map the dashed violet path runs from the slam ring along the amber dashed 8 % line, with its closest approach marked at 9.5 % (2), and the operating point now sits at 18.8 % true margin (1). Beside it the measured W_f/P_3 of 45.2 rides under the schedule's ceiling of 53.7 g/s per bar, above the 13 g/s-per-bar floor (3), as the trend label below reads. The EGT limiter holds the valve at 1.375 kg/s with the gauge at 893 °C (status EGT LIMIT, AHEAD OF THE GAUGE (4)). The scorecard reads 95 % in 4.74 s, a minimum margin of 9.5 % and 114.2 kN of 120.2 kN (5).

of the rating instead of 20.2 %. It buys the time from the chop, not from the slam, and the verdict does not grade that.

> **DID YOU KNOW? Two Settings Pass Together, Neither Alone**
>
> From the simulator: in the challenge, the 8 % schedule alone reaches 95 % thrust in 5.17 s but lets the indicated exhaust temperature reach 1019 °C, and integrator tracking alone holds 934 °C but takes 8.55 s. Together they pass, in 5.86 s at 938 °C (Table 3.5).

The 7.0 s is a teaching value, not paragraph (b)'s 5 s, which times a slam from a stabilized idle. With the shipped floor and governor no schedule margin from 4 to 8 % gets the bodie to 95 % within 5 s with 8 % of margin, the fastest being 5.63 s at a 7 % design margin. A sweep of 3,300 settings over every control found 907 that do reach 5 s with the margin and EGT kept, as fast as 2.87 s, all by holding fuel through the chop with a higher floor or minimum fuel. With the floor and minimum fuel as shipped, 2 of 4,450 settings reach it, through a governor at a quarter of its gain with a 5 s integral time. From a stabilized idle in this air the 8 % schedule takes 7.64 s, against 4.74 s at sea level.

Table 3.6 Matched Fuel-Governor Comparisons

Case	Governor	t_{95} s	Min SM %	Peak EGT °C	Travel kg/s
Sea-level slam	PI	4.74	9.51	899.3	4.50
	MPC	5.21	9.69	884.3	1.81
Hot-day slam	PI	6.25	9.26	954.6	4.38
	MPC	6.32	9.43	954.4	3.24
Bodie, suggested	PI	5.86	9.45	938.3	4.17
	MPC	5.71	9.66	938.7	2.11
Bodie, shipped	PI	8.49	14.78	996.4	3.44
	MPC	8.34	14.89	936.1	1.93

Travel sums absolute fuel-command changes, not fuel mass burned. Surge margin and thrust are evaluated after the run, never supplied to MPC.

3.5.5 Verification and Comparison With the Class

The following checks concern the code, not the quality of control. The takeoff design point and flight idle are exact steady states, with no derivative above 1.6×10^{-14} and 1.4×10^{-13} of its state per second. At takeoff the air in equals the air out of both nozzles and the bleeds to 9.7×10^{-14}. Over a slam each spool's kinetic energy change equals the integral of its net shaft power to 2×10^{-5} on the low spool and 4×10^{-5} on the high spool, by the trapezoid rule on the 0.005 s step. The engine's design point equals an independent hand calculation from the declared values to the last digit quoted, and the combustor's convention of one specific heat on both sides of the burner needs 1.6 % less fuel than air-table enthalpies would, against 13.6 % more for a balance with c_{pa} on the inlet side. The surge line in closed form reproduces each speed line's surge point at $n = 0.6$, 0.8 and 1.0. The selector hands over without a step larger than one scan of the rate limit in every case except the commanded cut of a chop, and switching to Manual and back at 40 instants of a slam moves the valve by at most 0.014 kg/s beyond one scan of the rate limit. Across 384 corners of the sliders at the default controls no state goes non-finite, and 18 starts, all at idle in fast, thin air, hold the run as outside the model's range, where 10 had gone non-finite before that hold was added. All 46 of the application's checks and all 10 completion walks pass.

Agreement with published data is a narrower claim. The takeoff point, 355 kg/s of air, a bypass ratio of 5.30, an overall pressure ratio of 30.0, 120.2 kN and 10.32 g/(kN s), and the flight idle at 10.0 % of takeoff thrust with N_2 at 60.6 %, were fitted inside ranges that the build collected for this thrust class from an airworthiness type-certificate data sheet and public specification tables, which are not cited here so that no engine maker is named. Flight idle at 10 % is inside paragraph (b)'s 15 % [13], and the NASA generic engine's controller reached 95 % from 15 % in 4.07 s on a larger engine [1], against 4.74 s here from 10 %. The maps are teaching shapes, no transient has been compared with an engine's, and no learner study has been made.

3.5.6 Predictive Governor Comparison

Table 3.6 compares complete 20 second runs at the same seed and settings, changing only the governor. Case 9 slams at 1 second with an 8% acceleration schedule. The hot-day case also slams at 1 second. Wear is applied at 2 seconds and the warmer day at 10 seconds in the wear comparison. The bodie rows use the challenge's suggested settings of Table 3.5 and, separately, its shipped settings.

The sea-level loss is explicit: 0.47 seconds more to reach 95% thrust in exchange for 2.68 kg/s less accumulated command movement. The suggested bodie passes under both laws, with MPC gaining 0.15 seconds but adding 0.4 K to peak indicated EGT. The hot-day test exceeds the 950 °C indicated-EGT limit under both laws, by 4.6 and 4.4 K. Prediction does not make the protection a hard guarantee. The wear-and-warmer-day test reaches minimum margins of 0.37% with PI and 0.36% with MPC, without a modeled surge. A fixed local model is not a margin estimator.

Thirty comparison rows include four schedule margins, six hindsight PI gain settings and half-step refinement. The latter keeps the sea-level ordering at 4.758 versus 5.220 seconds and the movement

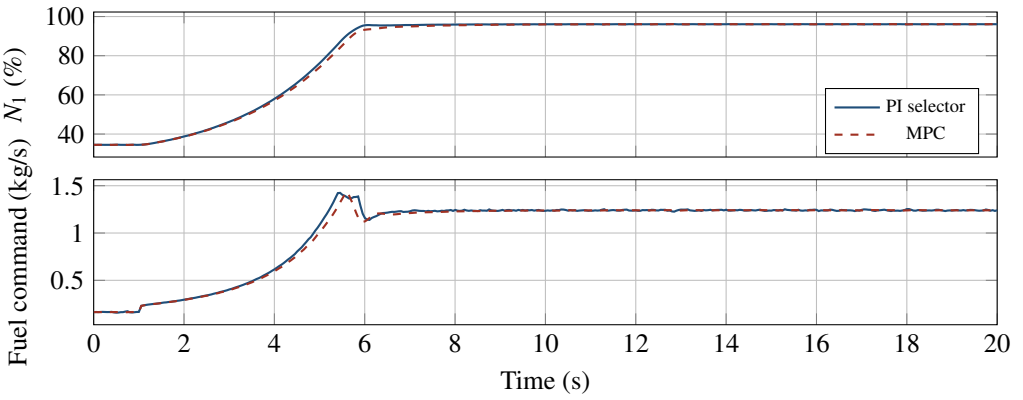

Figure 3.5 Same sea-level slam and protection settings, PI selector solid and MPC dashed. MPC suppresses fuel movement but reaches the thrust criterion 0.47 seconds later. The plotted speed is measured, while the thrust criterion is scored on teaching truth.

ordering at 4.53 versus 1.79 kg/s. Independent exact second-order responses agree with the sampled model within 3.0×10^{-11}, and an independent stepping harness reproduces both nominal travel totals exactly. Eight injected solver faults are rejected. Chrome averages 201 microseconds per plan over 400 plans, 0.402% of one core at unit playback and 2.01% at the maximum fivefold speed. Figure 3.5 shows the command tradeoff.

3.6 Technology in Practice

An engine control on a real turbofan reads the same kinds of measurement and moves the same valve as the selector of Section 3.4. Figure 3.6 draws that loop with engine hardware in place of the schematic's symbols.

What Measures It

Thrust cannot be measured in flight, so engines are controlled on shaft speed or on engine pressure ratio, both good indicators of it [7]. Many newer engines read shaft speed with a magnetic pickup that counts passing gear-tooth edges as pulses, which needs little maintenance beyond setting its gap to the teeth, and a maintenance check compares the speed indicator with an analyzer that reads engine speed within 0.1 %. The fuel control also senses compressor inlet pressure and temperature, burner pressure and compressor discharge pressure [19]. The model gives the speed pickups 0.02 % of colored noise and the discharge pressure P_3 a 0.02 s lag, and the ratio schedule divides the fuel by that P_3.

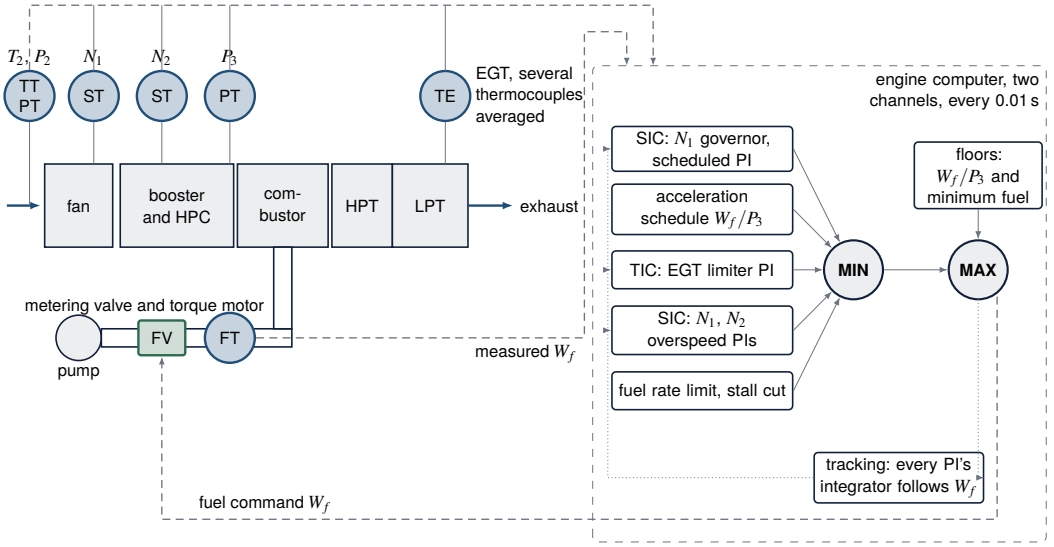

Figure 3.6 The selector of Figure 3.1 on engine hardware. Speed pickups (ST) on both spools, the discharge pressure (PT), thermocouples spaced around the exhaust and averaged (TE), the fan-face probes (TT, PT) and the fuel flowmeter (FT) feed the engine computer. There the governor (SIC) and the limiters (TIC, SIC) vote: the least of the caps, then the greatest of the floors, goes to the metering valve (FV). Notice that no sensor reads surge margin or thrust, and that the dotted tracking path returns the selected fuel to every loop that does not hold the valve.

TECHNOLOGY **Sensor: Exhaust Thermocouples Averaged Around the Duct**

The gas entering the first turbine is the most critical engine variable, but in most engines it is impractical to measure. Thermocouples at the turbine discharge give a relative indication of it instead: several are spaced around the exhaust duct near the turbine exit, the gauge shows their average, and they are made of chromel and alumel [19], the type K pair, usable to about 2300 °F [20]. In the oxidizing exhaust they drift from calibration, and a harness whose resistance has risen reads low, which the handbook calls dangerous because the engine then runs hot behind a normal gauge [19]. No open source read gives the probe's speed. The model reads the gas a quarter of the way through the low-pressure turbine's temperature drop, behind a 1.0 s lag.

What Moves It

The fuel metering unit sits on the front of the accessory gearbox, attached to the fuel pump, and the engine computer commands both its metering valve and the minimum-pressure and shutoff valve that starts and stops the flow [19]. A generic engine simulation of the kind this chapter cites models three actuators, the fuel metering valve, the variable stator vanes and the variable bleed valve [1]. The model has the first and the last, and folds the vanes into the compressor map on their schedule.

TECHNOLOGY **Final Element: A Metering Valve Behind a Fixed Pressure Drop**

In a full-authority control the computer decides the fuel flow and the metering valve only meters it. A differential pressure valve holds a constant drop across the metering valve by returning surplus pump flow to the pump inlet, so the metered flow is proportional to the valve's open area, and a torque motor driven by the control's current sets that area. In a hybrid control with a hydromechanical backup, an electrical failure, or the pilot's choice, puts the engine in manual mode: the torque motor current is zero, which sets a fixed ratio of fuel to compressor discharge pressure, and the engine runs at reduced power [19]. The handbook gives no stroke time or accuracy. The model's valve is a 0.05 s lag with a full flow of 1.6 kg/s.

Where the Controller Runs

The engine computer of a large turbofan has two channels, each an independent computer that can control the engine completely, and cross-talk logic picks the better channel for each bank of output drivers [19]. On commercial engines it is mounted on the fan casing and cooled by air [21]. The NASA generic controller samples every 0.015 s, beside the application's 0.01 s, and it uses the min/max structure that is common in commercial engines. Its designers rejected zeroing a deselected integrator's error, which steps the fuel at the hand-over and slows the response, and fed back instead the difference between the selected command and each loop's own output through an integral feedback gain tuned empirically. Too high a gain makes the loop act like a proportional controller or even destabilizes it, and too low a gain leaves windup in place [1]. Equation (3.7) is that scheme with the gain set by $T_t = T_i$.

> **IN PRACTICE** **Hand-Overs That Certification Asks For**
>
> The United States rule for engine control systems, 14 CFR 33.28, asks that the control keep the engine within its approved limits as conditions change. When a fault moves the engine from one control mode, channel or backup system to another, the engine must not exceed a limit, surge or stall, and the size and duration of the thrust change must be stated in the instructions. The full system must tolerate any single electrical or electronic failure without losing thrust control, an electronic overspeed protection must be tested once per start and stop cycle, and the software must be verified by an approved method [22]. The bumpless return from Manual in this chapter answers the first of these in miniature.

What Is Modernizing It

Three developments bear on the selector. The first is the onboard model: a flight program that estimated engine parameters with a Kalman filter led to a self-tuning engine model implemented in the engine computer of a military engine [7]. The second is distributed control. Today's architecture wires every sensor and actuator to one hardened box, because commercial electronics cannot survive the heat around the engine, and it suffers from parts obsolescence and heavy harnesses. A distributed architecture would replace them with nodes on a standard digital network, for less weight, better reliability and room for new functions [21]. The third is electric actuation: electromechanical actuators have joined stepper motors on engines [7].

3.7 Limitations

The engine is a teaching model of a class, with declared maps, constant properties, isothermal plenums of a declared residence time and a combustor without volume, and its transient margins are plausible for the class rather than predicted for any engine. The stall is a declared loss of pressure rise that clears low on its speed line, so neither rotating stall nor deep surge is represented, and only the HPC can stall. The schedules were designed on the model itself, so they fit it better than a real schedule fits a real engine, and the deterioration's surge-line loss has no source. The stall recognition has a stated blind spot after lever pulls, the controller has no minimum core speed, no thrust rating schedule against the ambient, no duplicated sensors and no fault accommodation, and the protection it provides is an educational guardrail, not certified logic. A pass of the challenge is not compliance with 14 CFR 33.73 or any other requirement.

The MPC's single local fuel model omits nonlinear gas-path transients, ambient scheduling and wear. Its soft speed and indicated-temperature ceilings do not constrain true EGT or actual surge margin. In particular, the 954.4 °C hot-day peak stands as a counterexample to an unconditional protection claim.

OPEN CHALLENGE **Closing the Loop on What No Sensor Reads**

Thrust, the variable an engine is flown by, cannot be measured, and neither can turbine inlet temperature or stall margin, two of the limits that matter, while the components degrade with use. Model-based control would run an onboard engine model, tuned to the individual engine by a tracking filter, and close loops on its estimates of thrust and stall margin. The difficulty is that the health parameters to be estimated generally outnumber the sensors, so the unmeasured outputs can carry significant errors [21], and any model used inside a control loop will itself need to be certified [23]. Progress would be estimates good enough on a worn engine to replace fixed schedules designed once on the nominal engine.

3.8 Conclusion

This chapter studied a browser simulator of a generic 120 kN-class two-spool turbofan with its min/max fuel selector. Takeoff fuel in one move surges its compressor 0.020 s later, and the governor alone surges 0.125 s after a slam. An acceleration schedule on the corrected ratio $W_f/(P_3\sqrt{\theta_2})$ trades margin for time, 5.37 s at a 10 % design margin and 4.74 s at 8 % with the margin never under 9.5 %, and it keeps 8.1 % or more on slams up to 9144 m except where the minimum fuel outranks it. With integrator tracking by back-calculation each limiter takes the valve from the fuel in the line, the EGT limiter 89 K under its line while the gas is already at 1021 °C, and without it the governor overshoots by 1.9 points and the challenge's EGT reaches 1019 °C. The challenge passes with the schedule tightened and tracking on, in 5.86 s, whatever the gain schedule, or with tracking and a higher floor that buys the time from the chop. The model's maps are teaching shapes and its limits teaching values. A fuller model would add a minimum core speed, a thrust rating schedule against the ambient, heat soak and tip clearance, a fan stall model, and a model predictive limit protection for comparison with the selector. The simulator shows that an engine controller is a set of limits voting on one valve before it is a speed loop.

The predictive fuel governor reduces the nominal fuel-command travel from 4.50 to 1.81 kg/s while slowing 95% acquisition from 4.74 to 5.21 seconds.

KEY POINTS

- An engine controller is a set of limits voting on one fuel valve before it is a speed loop: the lowest cap and then the highest floor win.
- Takeoff fuel in one move surges the compressor 0.020 s later, because the margin is spent by the rate of fuel rather than its amount.
- The acceleration schedule trades margin for time, 5.37 s at a 10 % design margin and 4.74 s at 8 % with the margin never under 9.5 %.
- With integrator tracking by back-calculation each limiter takes the valve from the fuel in the line, and without it the challenge's indicated EGT reaches 1019 °C.

Where else this applies. Chapter 9 meets the same exhaust-temperature limit on a constant-speed gas turbine, and Volume 7, Chapter 4 uses the same back-calculation on a slip loop under a torque limit.

References

[1] J. Csank, R. May, J. Litt, and T.-H. Guo, "Control design for a generic commercial aircraft engine," in *46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. American Institute of Aeronautics and Astronautics, 2010, paper AIAA 2010-6629.

[2] H. A. Spang, III and H. Brown, "Control of jet engines," *Control Engineering Practice*, vol. 7, no. 9, pp. 1043–1059, 1999.

[3] L. C. Jaw and J. D. Mattingly, *Aircraft Engine Controls: Design, System Analysis, and Health Monitoring*. Reston, Virginia: American Institute of Aeronautics and Astronautics, 2009.

[4] H. Richter, *Advanced Control of Turbofan Engines*. New York: Springer, 2012.

[5] 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.

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

[7] L. C. Jaw and S. Garg, "Propulsion control technology development in the United States: A historical perspective," NASA Glenn Research Center, Cleveland, OH, Technical Memorandum NASA/TM-2005-213978, 2005, nTRS 20050238468.

[8] 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.

[9] P. P. Walsh and P. Fletcher, *Gas Turbine Performance*. Wiley, 2004.

[10] G. G. Kulikov and H. A. Thompson, Eds., *Dynamic Modelling of Gas Turbines*, ser. Advances in Industrial Control. London: Springer, 2004.

[11] R. May, J. Csank, T. Lavelle, J. Litt, and T.-H. Guo, "A high-fidelity simulation of a generic commercial aircraft engine and controller," in *46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. American Institute of Aeronautics and Astronautics, 2010, paper AIAA 2010-6630.

[12] J. W. Chapman, T. M. Lavelle, R. May, J. S. Litt, and T.-H. Guo, "Propulsion system simulation using the Toolbox for the Modeling and Analysis of Thermodynamic Systems (T-MATS)," in *50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*. American Institute of Aeronautics and Astronautics, 2014, paper AIAA 2014-3929.

[13] U.S. Code of Federal Regulations, "Title 14, Part 33, Section 33.73: Power or thrust response," Electronic Code of Federal Regulations, https://www.ecfr.gov/current/title-14/part-33/section-33.73, 2026, amdt. 33-1, 36 FR 5493, Mar. 24, 1971; text as in force on 2026-10-01.

[14] S. Garg, "Aircraft turbine engine control research at NASA Glenn Research Center," *Journal of Aerospace Engineering*, vol. 26, no. 2, pp. 422–438, 2013.

[15] W. J. Rugh and J. S. Shamma, "Research on gain scheduling," *Automatica*, vol. 36, no. 10, pp. 1401–1425, 2000.

[16] D. J. Leith and W. E. Leithead, "Survey of gain-scheduling analysis and design," *International Journal of Control*, vol. 73, no. 11, pp. 1001–1025, 2000.

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

[18] 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] Federal Aviation Administration, *Aviation Maintenance Technician Handbook–Powerplant*. U.S. Department of Transportation, Federal Aviation Administration, 2023, no. FAA-H-8083-32B.

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

[21] S. Garg, "Aircraft turbine engine control research at NASA Glenn Research Center," NASA Glenn Research Center, Cleveland, OH, Technical Memorandum NASA/TM-2013-217821, 2013, nTRS 20130013439.

[22] "Title 14, Part 33, Section 33.28: Engine control systems," Code of Federal Regulations, 14 CFR Ch. I (1–1–25 Edition), govinfo.gov, 2025, amdt. 33–26, 73 FR 48284, Aug. 19, 2008.

[23] S. Garg, K. Schadow, W. Horn, H. Pfoertner, and I. Stiharu, "Sensor and actuator needs for more intelligent gas turbine engines," NASA Glenn Research Center, Cleveland, OH, Technical Memorandum NASA/TM-2010-216746, 2010, also ASME paper GT2010-22685; NTRS 20100029602.

3.9 Exercises

3.1 ○ Concept check · 2 min

In scenario 1, Fuel by hand, the fuel at flight idle is raised from 0.164 to 1.2 kg/s in one move. The HPC operating point reaches the surge line 0.020 s later, while the core speed has moved only from 60.6 to 60.8 %. Why does the point climb its speed line before the spools speed up?

(a) The governor raises the core speed first, and the pressure ratio follows it along the steady line

(b) The hotter gas must pass the choked HPT nozzle at the same flow, which takes a higher P_3

(c) The handling bleed shuts as the fuel rises, so the air at the HPC exit has nowhere to go

(d) The fan speeds up first and forces more air into the HPC than its speed line can pass

3.2 Concept check · 2 min

In scenario 3 with the schedule designed for 8 % and integrator tracking on, the EGT limiter takes the valve 4.43 s after the slam with the gauge at 861 °C, 89 K under its line. What makes it take over that far from its limit?

(a) Thermocouple noise now and then carries the gauge over the limiter's setpoint

(b) The limiter's setpoint is lowered by a fixed margin while the schedule holds the valve

(c) Its integrator trails the rising fuel, and its demand meets the fuel when $K_E e$ equals that trail

(d) The selector gives the valve to the limiter whose error is falling fastest

3.3 Concept check · 2 min

At 10 668 m and Mach 0 to 0.4, slams from flight idle fall to 0 to 7.3 % of surge margin or surge, and they do so identically with and without the $\sqrt{\theta_2}$ correction at Mach 0. Which property of the selector, Equation (3.6), explains it?

(a) The $\sqrt{\theta_2}$ correction loosens the schedule too far in cold, thin air

(b) The fuel rate limit of 2 kg/s per second is too fast for the thin air

(c) The gain schedule raises the governor's gain at low corrected core speed

(d) The 0.06 kg/s minimum fuel exceeds the idle fuel there and, as a floor, outranks the schedule

3.4 Theory and calculation · 5 min

Scenario 6 flies the lever at 70 % at 10 668 m and Mach 0.8 on a standard day. The lever asks for a corrected fan speed of 83.9 %.

(a) From Equation (3.1), with the standard atmosphere falling 6.5 K/km from 288.15 K, find T_0, T_2 and θ_2. *Checkpoint:* $\theta_2 = 0.857$

(b) Find the physical fan-speed target N_{1t}. *Checkpoint:* 77.6 %

(c) The acceleration cap is $W_{acc} = R_{acc} P_{3,m} \sqrt{\theta_2}$. By what fraction is it lower here than a cap on the uncorrected ratio W_f/P_3 with the same table value and P_3, and which way does that move the surge margin? *Checkpoint:* 7.4 % lower, more margin

3.5 Theory and calculation · 5 min

Integrator tracking, Equation (3.7), for a loop i that does not hold the valve.

(a) Substitute $u_{i,k} = K_i e_{i,k} + I_{i,k}$ into the integrator update and show that I_i is a first-order lag of the commanded fuel with time constant T_i. *Checkpoint:* $\dot{I}_i = (W_f - I_i)/T_i$

(b) The fuel rises at a steady rate ρ. Show that the trail $W_f - I_i$ settles at ρT_i, and find the error e_i at which the loop's demand meets the fuel. Evaluate it for the EGT limiter in the 8 % slam of scenario 3, with $K_E = 6\,\mathrm{g/s}$ per kelvin, $T_E = 1.0\,\mathrm{s}$ and a trail of 0.534 kg/s. *Checkpoint:* $e = \rho T_i / K_i$, about 89 K

(c) Estimate the distance from the limit at takeover if the EGT limiter's gain were doubled to 12 g/s per kelvin with the same trail, and say what the limiter's integrator does instead with tracking off. *Checkpoint:* about 45 K

3.6 Theory and calculation · 6 min

The HPC map of Equation (3.2) has $a = 2.1$, $b = 2.9$, $k = 0.3$, $\pi_d = 12.5$ and $S = 1.0898$, and the design point sits at $m = 0.768$ with $\pi = 12.5$.

(a) Find the normalized pressure rise ψ_s on the surge line at $m = 0.768$ and the surge pressure ratio. *Checkpoint:* $\psi_s = 1.239$, $\pi_s = 15.24$

(b) Find the design surge margin by the chapter's definition, and by the other common definition that divides by the surge line's ratio. *Checkpoint:* 22.0 % and 18.0 %

(c) At the same corrected flow, at what pressure ratio does the margin fall to the 8 % floor? *Checkpoint:* 14.12

3.7 App challenge · 8 min

Scenario 3, The acceleration schedule, default seed. In each run press *Run* and, about 1 s later, move *Thrust lever* to 100 % (or press *Throttle slam*). The run pauses at 95 % of the rated thrust.

(a) As shipped (*Schedule design margin* 10 %), read the time to 95 % and the minimum surge margin on the scorecard at the pause. *Checkpoint:* 5.37 s, 11.3 %

(b) Reload the case, set *Schedule design margin* to 8 % before Run, and slam again. Read the same two values and the *Holding the valve* tile at the pause. *Checkpoint:* 4.74 s, 9.5 %, EGT limiter

(c) Reload the case, set the margin to 8 % and *Integrator tracking* to Off before Run, and slam again. Read the time to 95 %, then press Run at the pause and read the *Peak N1* tile about 15 s later. Explain the difference from (b). *Checkpoint:* 4.66 s, 98.0 %

Tolerance: 0.03 s on the times, which move that much with the instant of the slam, and 0.1 point on the percentages.

3.8 App challenge · 8 min

Scenario 4, A hot day: the EGT limiter (ISA + 30 K at sea level), default seed. Press *Run*, move *Thrust lever* to 100 % about 1 s later, press Run again at the 95 % pause, and let the run end at 21 s.

(a) Export the CSV file. Take the time of the first row with `lever_pct` at 100 and of the first row whose `active` column reads `egt`. How long after the slam does the EGT limiter first hold the valve? *Checkpoint:* about 3.9 s

(b) At the end of the run read the thrust and the rating on the *Thrust (truth)* tile, the peak indicated EGT on the scorecard, and the largest `truth_egt_c` in the file. *Checkpoint:* 111.3 kN of 111.2 kN, 955 °C, about 1020 °C

(c) Set the lever back to 0 %, set *Integrator tracking* Off, press Reset and repeat the run. Read the peak indicated EGT. *Checkpoint:* 996 °C

Tolerance: 0.1 s on (a), 0.2 kN on the thrust, 2 K on the indicated EGT and 3 K on the true one.

3.9 App challenge · 8 min

The challenge, *the bodie at 10,000 ft on a hot day*, default seed. Run each setting for the full 20 s, then change it, press Reset and run again. Leave every parameter and disturbance as shipped.

(a) Run as shipped and read the verdict: time to 95 %, minimum surge margin, peak indicated EGT, and which limits are missed. *Checkpoint:* 8.49 s, 14.8 %, 996 °C; time and EGT

(b) Run with only *Schedule design margin* at 8 %, then with only *Integrator tracking* On. *Checkpoint:* 5.17 s and 1019 °C; 8.55 s and 934 °C

(c) Run with both changes and the shipped fixed gains, then with *Gain schedule* On N2c, governor gain 200 g/s per percent and integral time 1.0 s. Does the gain schedule pass the challenge? *Checkpoint:* 5.86 s both times

Tolerance: 0.02 s, 0.1 point and 2 K.

3.10 Application area · 4 min

14 CFR 33.73(b) asks for 95 % of rated takeoff thrust within 5 s of a slam. The schedules of this engine were designed on the nominal engine. On the sea-level slam of scenario 3 they give 4.74 s and a minimum margin of 9.5 % at an 8 % design margin, 5.37 s and 11.3 % at 10, 6.27 s and 13.0 % at 12, and 7.55 s and 14.8 % at 14. Full deterioration lowers the steady takeoff margin from 22.0 to 17.8 %.

(a) Suppose a worn engine's slam loses the same 4.2 points as its steady takeoff point. Which design margins keep a worn engine's slam above the 8 % floor, and which of those meet the 5 s? *Checkpoint:* 12 and 14; neither

(b) Why is that supposition optimistic, and what would an engineer do before choosing the design margin?

3.11 Application area · 4 min

At sea-level ISA the new engine's takeoff EGT is 872.0 °C, its rating is flat to ISA + 20.0 K, and full deterioration raises the EGT by 59.8 K at the same thrust. The indicated limit is 950 °C.

(a) What is the new engine's EGT margin at sea-level ISA, and about how many kelvin of EGT does each kelvin of day temperature add at takeoff? *Checkpoint:* 78 K, about 3.9 K per kelvin

(b) Estimate the day temperature at which the fully worn engine's rating stops being flat, and compare with scenario 8. *Checkpoint:* about ISA + 5 K

3.12 Review the agent · 5 min

An AI assistant reviews the trends of scenario 3 and writes: "The EGT limiter takes the valve at 861 °C, 89 K under its 950 °C limit, which shows it is tuned too aggressively. Lengthen its tracking time to two integral times, or switch integrator tracking off, so that it acts only at the limit. The slam then reaches 95 % sooner, which is what the five-second rule asks for." Find the flaw and say how you would check it in the application.

CHAPTER 4

Feed Rate, Grain Loss and a Rain Deadline in a Combine Harvester Simulator

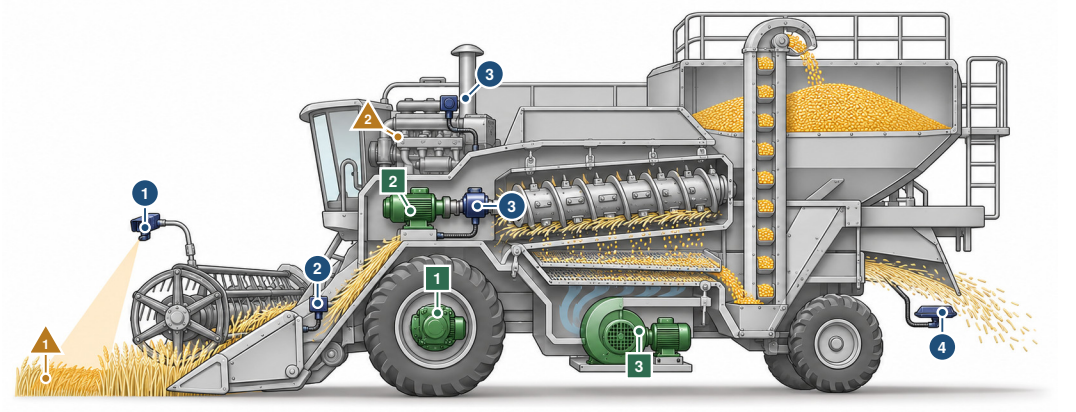

The controller moves 1 ground speed, set through the hydrostatic drive 2 rotor speed 3 cleaning fan speed

The controller reads 1 crop ahead of the header 2 crop flow into the feeder 3 engine load and rotor torque 4 grain lost behind the machine

What disturbs it 1 denser, wetter or flattened crop 2 engine power lost

Notice that the scanner reads the crop meters ahead of the header, while the loss pad reads grain only after it has passed through the whole machine. Illustration, not to scale.

IN PLAIN TERMS **Driving at the Pace the Machine Can Take**

The main idea. A combine cuts, threshes and cleans grain as it drives across a field. Its speed sets how much crop goes in. A controller holds that intake steady by changing the speed.

The challenge. The intake reading cannot show what really limits the machine: wet straw that clogs the spinning drum, an engine running out of power, or grain thrown out the back with the straw. Guard controllers watch those limits and may slow the machine, while the reading of lost grain arrives late and can be wrong.

Why it matters. Drive too slowly and the rain finds crop still standing, drive too fast and grain is lost or the machine jams.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page, the fact on overload (Section 4.1), the two history segments (Section 4.2), the two try-it segments (Sections 4.6.3 and 4.6.6), the fact on the challenge (Section 4.6.7) and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 4.6), Technology in Practice (Section 4.7), Limitations (Section 4.8) and the open challenge

The control question A combine's ground speed sets its feed, but grain loss, a plugging rotor and the engine's power limit do not show in the feed rate and are measured late or not at all. Which layers of override, preview and supervisory control keep the machine and the grain safe, and what does each one cost in grain brought in before the rain?

Learning objectives

- Compute the process gain and the delay budget of the feed loop, and predict when a later feed signal makes it hunt.
- Explain how a low selector with external reset lets the engine and rotor overrides take the ground speed without a bump.
- Size the range of a crop preview from its braking curve and the drive lag.
- Tell which disturbances the feed loop, the preview and the loss supervisor can each see, and judge a supervisor that trusts a biased monitor.
- Choose and defend a configuration that meets a grain target under loss, plug and overload limits.

Prerequisites Sections 1.4.2 (constraints as objectives), 1.4.3 (selectors and override), 1.4.4 (external reset), 1.4.5 (rate limits), 1.6.2 (velocity form) and 1.6.4 (scenarios and exported files). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/combine-harvester/

See also Volume 5, Chapter 3 for a fixed tuning eroded by a slow analyzer, and Volume 5, Chapter 2 for control through a transport delay. Volume 3, Chapter 2 for matching a cascade to where a disturbance enters, and Volume 3, Chapter 6 for feedforward from an upstream measurement with feedback demoted to a slow trim.

4.1 Introduction

The combine harvester cuts a swath of crop, threshes the grain out of the heads, separates it from the straw and cleans it from the chaff, all while it drives across the field. The ground speed is the only handle on how much material enters, because the header takes whatever crop stands in its path. The rotor and the sieves can pass only so much straw and chaff, and the grain that leaves with them rises steeply as the feed approaches that capacity [1, 2]. The harvest window is short as well, and crop still standing when the weather turns carries a cost of its own, so studies of automatic combine control have treated the forward speed as a balance between the grain the machine loses and the grain that time loses [3, 4].

> ✦ DID YOU KNOW? **Five Bushels an Acre to a Slug of Straw**
>
> Combines were designed for wheat, and the first was built in 1838. An extension guide on harvesting wheat warns that many growers forfeit as much as 5 bushels per acre in the heaviest-yielding parts of a field when the header is not raised for tall straw and the straw overloads the combine for a moment [5].

The loop that sets the speed is not difficult by itself. Feed rate is crop density times header width times speed, so the process gain is known, and the dynamics are a drive lag and a transport delay. The difficulty is that the constraints that matter do not appear in the feed rate. Wet crop loses more grain at the same feed, a rotor fed faster than its capacity fills until it plugs, an engine asked for more power than it has slows everything it drives, and the loss monitor reports the result seconds late through a calibration that may be wrong. The simulator of this chapter is built to correct the belief that a throughput loop on its setpoint is protecting the machine and the grain. Section 4.6 measures a feed loop holding 70 t/h while the loss rises from 1.00 to 2.69 %, and one holding 80 t/h while the rotor plugs 27 s into a patch of lodged crop.

The same structure appears wherever a throughput loop feeds equipment whose limit it cannot see. Grinding mills are fed ore against the load and the power they can take, and the choice of controlled and manipulated variables around them has been surveyed across 68 circuits [6]. Volume 15, Chapter 3 treats one such mill, with the ore feed moved against the bearing pressure. Crusher feeders are

regulated to keep the crushing chamber full while the crusher power stays inside a safe margin [7]. The established answer has several parts. Override controllers on the equipment limits share the manipulated variable with the throughput loop through a selector [8], and the controllers that lose the selector need anti-windup so that they can take over without delay [9] (Sections 1.4.3 and 1.4.4). A slower loop on a quality measurement moves the throughput setpoint as the outer loop of a cascade, and a measurement of the incoming material, where one exists, supplies feedforward [10, 11]. The combine carries every one of these parts.

This chapter weighs, on one simulated combine, what each layer of that structure contributes, using patches of crop that each defeat a different layer and a rain deadline that prices every lost minute. The strategies are established practice. The model integrates the feed exactly from a mapped field, has a rotor whose capacity falls with moisture and with engine lugging and a plug with a real cost, and carries a measurement chain in which every signal is late, noisy or biased. Nine case studies and a graded challenge take the feed loop, its transport delay, the preview, the overrides and the supervisor one at a time before setting them together against the weather.

Section 4.2 reviews the automation of combines. Section 4.3 sets out the field, the machine, its instruments and the assumptions, and Section 4.4 derives the characteristics that shape the control problem: the loss curve, a process gain set by the crop and the delay budget of the feed loop. Section 4.5 gives the speed selector, the preview, external reset, the loss supervisor and the trims as implemented, and in Section 4.5.7 the optional predictive controller. Section 4.6 reports the measured scenarios, from the uniform field through the dense patch, the wet strip, the delayed and the biased signals and the engine and rotor limits to the harvest against the rain and the predictive comparison. Sections 4.8 and 4.9 state the limitations and conclude after Section 4.7 has set the loop on real machine hardware, and Section 4.10 holds the exercises.

4.2 Automation of the Combine Harvester

The capacity a combine works against is set inside the machine. Grain must be detached from the heads and must then find its way through a moving mat of straw to the concave. Stochastic models of threshing and separation along the rotor express the unthreshed and unseparated grain, which become the losses, as functions of the operating and design parameters, and they have been validated on extensive tests of axial and tangential threshing units [2]. The theory, modeling and design of the machine from the header to the cleaning shoe have been collected in one monograph [1].

Automatic control of the forward speed was first weighed by simulation in a 1983 thesis, whose findings the history segment below gives. Timeliness costs vary less from year to year for machinery with a high daily capacity, and they respond nonlinearly to a loss of capacity because more fields are then left partly or wholly unfinished [4].

> ⧖ HISTORY **Holding the Loss or the Feed Constant, 1983**
>
> By 1983 a doctoral thesis on combine automation could survey several automatic controls that had been built and tested. Each varied the ground speed to keep either the grain loss or the straw feed rate constant as the crop density changed across the field. The thesis simulated five ways of setting the forward speed and the threshing cylinder speed, from the driver alone through loss control and combined loss and feed-rate control to two schemes that also set the cylinder. Following the slowly varying crop brought the main savings, following its rapid variations brought none, and timeliness losses mattered greatly to the optimum. The disturbances could not be measured accurately while the machine worked, and the effect of control was seen only after a delay. A reliable loss measurement was a necessary condition, and the meters of the day were not accurate enough [3].

A quarter century later, model-based predictive control was applied to the same problem. A cruise controller set the hydrostatic pump and the engine speed from a nonlinear model of the propulsion system and was validated on a production machine [12], and a throughput controller regulated the

Table 4.1 Variables at the Uniform-Field Operating Point

Symbol	Role	Nominal	Range
Q_m	CV, feed rate as measured	70 t/h	SP 20–115
λ_m	CV (override), engine load	71 %	SP 60–100
T_m	CV (override), rotor torque	38 %	SP 40–95
L_m	CV (supervisor), loss monitor	1.02 %	SP 0.5–3
v_{cmd}	MV, ground speed command	7.56 km/h	0–10
Q_{SP}	MV (supervisor), feed setpoint	70 t/h	14.4–115.2
ω_{set}	MV (trim), rotor speed	950 rpm	700–1200
n_{set}	MV (trim), fan speed	1000 rpm	700–1300
ρ, m	DV, crop density, moisture	1.045, 14 %	patches
δ	DV, engine power available	100 %	50–100
b, b_p	DV, loss, preview calibration	0, 0	±60, ±30 %
m_{off}, $\dot{m}_d$	DV, moisture offset, drift	0, 0	−4 to 6, ±0.5/min
σ_Q	DV, feed-signal noise	3 %	0–10
s, v	states, route position, speed	0, 7.40 km/h	—
n_e	state, engine speed of rated	0.986	0.5–1
M_r, G_r	states, rotor straw, grain	26.5, 29.4 kg	—
Ψ_r	state, rotor straw times moisture	369 kg %	—
G_s, C_s	states, shoe grain, chaff	15.2, 2.77 kg	—
G_t, G_d	states, tank grain, cracked grain	0, 0	—
G_{ls}, G_{lc}, V_f	states, lost grain, fuel	0, 0, 0	—

Density in kg/m^2. Engine load is a share of the power available, and rotor torque a share of the plug torque.

ground speed from measurements of feed rate, engine load and grain mass flow so that the capacity of the machine was used at all times [13]. Loss measurements entered the loop in the same years. A fuzzy controller built on expert knowledge adjusted the cylinder speed, the concave clearance, the fan speed and the forward speed from the losses measured at the straw walkers and the sieve [14], and a fuzzy controller for the cleaning section set the fan speed from the grain loss sensors and a differential air pressure that expresses the working load, and lost less grain than a sub-optimal fixed fan speed [15].

The second history segment follows the sensors that have tried to measure the crop before the header reaches it. LiDAR has been shown to estimate the crop density of small grains [16]. Control that uses such future information is preview control, whose optimal continuous-time form was derived in 1975 [11].

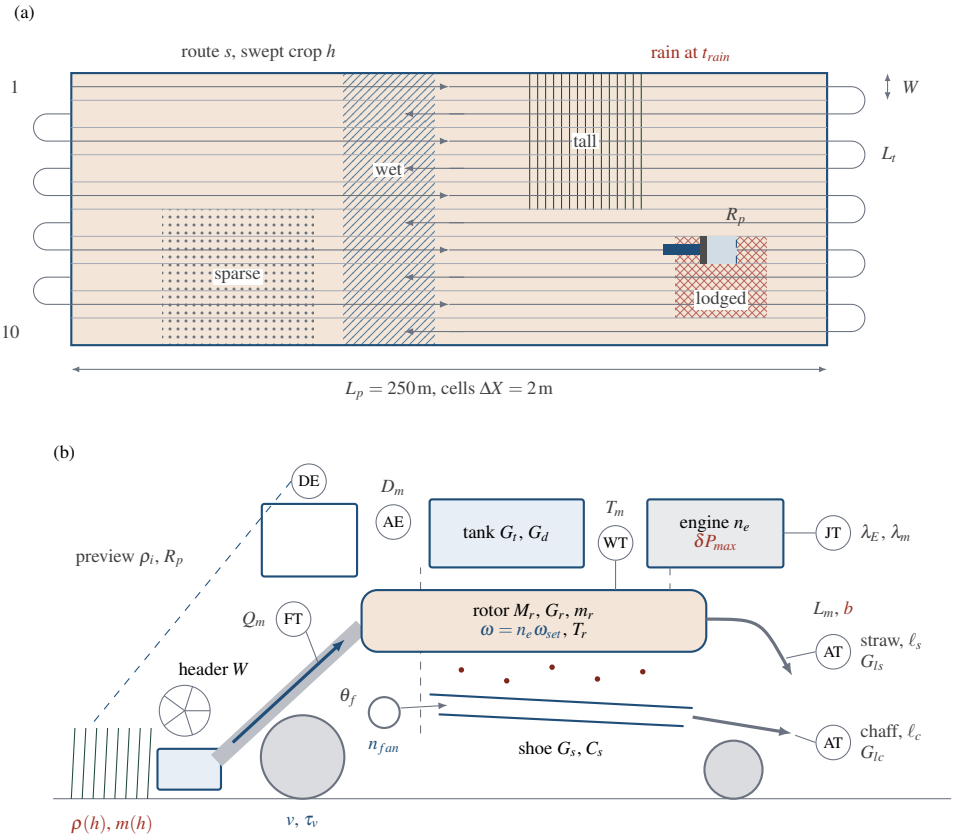

Figure 4.1 The simulated field and machine. (a) The challenge field in plan view, ten passes of width W and length L_p joined by headland turns of length L_t, with the tall, sparse, wet and lodged patches and the combine on pass 7 with its preview range R_p ahead. The route position is s and the crop swept so far is h. (b) The machine in side view. Crop of density $\rho(h)$ and moisture $m(h)$ is cut by the header and reaches the rotor after θ_f. The rotor holds straw M_r, grain G_r and moisture m_r and turns at ω, the shoe holds grain G_s and chaff C_s under the fan at n_{fan}, and the clean grain rises to the tank. Straw and chaff leave with the separation and cleaning losses. One engine at speed n_e drives the wheels at v, the rotor and the fan. The instruments are the feed-rate signal FT, the rotor torque WT, the engine load JT, the loss sensors AT with calibration error b, the damage camera AE and the forward crop sensor DE. The disturbances ρ, m, δP_{max}, b and the rain are printed in red, and the manipulated quantities v, ω and n_{fan} in blue.

HISTORY Seeing the Crop Before the Header, 1980 to 2020

A feed-rate control that weighs the crop on its way through the header and the feeder measures crop already inside the machine. With no time to anticipate a sudden change, it keeps a safety margin and leaves part of the capacity unused. Researchers therefore moved the measurement forward. An ultrasonic beam across the header, tried in 1980, did not follow the crop density closely enough. A laser beam ahead of the header followed in 1986, capacitance plates on the header table in 1994, and in 1996 a pendulum bar pushed by the crop, whose reading a speed change of 0.5 km/h shifted by up to 2 %. Laser scanners came next, and by 2020 crop height and density estimation ahead of the combine had been studied for 30 years and remained a challenge [17].

4.3 Machine and Field Model

The combine carries a 9 m header along a fixed serpentine route of ten 250 m passes, 2.25 ha of wheat yielding 5.5 t/ha of grain with 0.9 kg of straw and chaff for every kilogram of grain. The field, the machine and its instruments are shown in Figure 4.1. Thirteen states are integrated, and they are listed with the other variables in Table 4.1.

4.3.1 Route, Field and Feed

Because the route is fixed, the position s along it is the only coordinate, and the harvest coordinate h counts the meters of crop the header has swept, as shown in Equation (4.1).

$$\dot{s} = \max(0, v), \qquad h = kL_p + \min(u, L_p) \tag{4.1}$$

Here $k = \lfloor s/(L_p + L_t) \rfloor$ is the pass index, $u = s - k(L_p + L_t)$, $L_p = 250\,\mathrm{m}$ is the pass length and $L_t = 2 \times 8\,\mathrm{m} + \pi W/2 = 30.1\,\mathrm{m}$ is the headland turn. The coordinate h stops during a turn, after the last pass and once the rain has come, so nothing is cut then.

The field is stored in the order the header meets it, in cells $\Delta X = 2\,\mathrm{m}$ long and one header width wide. Cell i carries grain $\rho_{g,i} = Y f_{g,i}$, straw and chaff (material other than grain, MOG) $\rho_{m,i} = Y r_M f_{m,i}$ and moisture $m_i = m_0 + \Delta m_i$, with $Y = 0.55\,\mathrm{kg/m^2}$, $r_M = 0.9$ and $m_0 = 14\,\%$ on a wet basis. A smooth texture $f_g = 1 + a\xi$, $f_m = 1 + 1.3a\xi$ and $\Delta m = 0.3\zeta$ comes from bilinear interpolation of standard normal deviates ξ and ζ drawn by the seeded generator on a grid 25 m long and one pass wide, with $a = 0.03$, or 0.02 in case 1. A patch multiplies f_g and f_m and adds to Δm by $(1.25, 1.9, 0)$ for tall crop, $(0.55, 0.5, 0)$ for sparse crop, $(1.0, 1.1, 6)$ for a wet strip and $(1.1, 1.4, 7)$ for lodged crop. Uniform crop therefore stands at $\rho = \rho_g + \rho_m = 1.045\,\mathrm{kg/m^2}$, and the four patches at 1.56, 0.53, 1.05 and 1.24 times that density.

The crop enters the header at $Q_c = W\rho(h)v$, and the mass cut in each physics step is taken exactly from cumulative sums along the route, as shown in Equation (4.2).

$$\Delta G_c = W\left[C_g(h_{k+1}) - C_g(h_k)\right], \qquad C_g(h) = \int_0^h \rho_g \,\mathrm{d}h' \tag{4.2}$$

The MOG and the MOG weighted by its moisture are cut the same way, so what leaves the field is exactly what enters the machine. A moisture offset and a moisture drift, both disturbances, are added to the crop as it is cut. The header auger and the feeder chain carry the crop to the rotor through a ring buffer of $\theta_f/\Delta t = 30$ physics steps, with $\theta_f = 1.5\,\mathrm{s}$ and $\Delta t = 0.05\,\mathrm{s}$. The cell leaving the buffer is fed to the rotor over the following step, so the applied delay is one step longer, and the rotor inflow $Q_i(t) = Q_c(t - \theta_f - \Delta t)$, 1.55 s late, holds for grain, MOG and moisture alike.

4.3.2 Rotor and Plugging

The rotor is one well-mixed inventory of straw M_r, grain G_r and moisture-weighted straw $\Psi_r = M_r m_r$, emptied at the specific rate r, as shown in Equation (4.3).

$$\begin{gathered}
\dot{M}_r = Q_{i,m} - rM_r, \qquad \dot{G}_r = Q_{i,g} - rG_r \\
\dot{\Psi}_r = Q_{i,w} - r\Psi_r, \qquad rM_r = \bar{Q}\tanh\frac{M_r}{\bar{Q}\,\bar{\tau}_r} \\
\bar{Q} = Q_{cap}\left(\frac{\omega}{\omega_{ref}}\right)^{0.8} \mathrm{e}^{-0.04(m_r - 14)}, \qquad \bar{\tau}_r = \tau_r\left(\frac{\omega_{ref}}{\omega}\right)^{1/2}
\end{gathered} \tag{4.3}$$

Here $\tau_r = 2.5\,\mathrm{s}$, $Q_{cap} = 16\,\mathrm{kg/s}$, $\omega = n_e\omega_{set}$ is the rotor speed, $\omega_{ref} = 950\,\mathrm{rpm}$, and $m_r = \Psi_r/M_r$ is in percent. Lightly loaded, the straw leaves in $\bar{\tau}_r$. Near capacity the outflow saturates at $\bar{Q}$, and any faster feed piles up. The capacity is 16.0 kg/s of straw at 14 % moisture, 12.6 kg/s at 20 and 12.1 kg/s at 21 %.

The rotor torque grows with the straw mat and with its moisture, as shown in Equation (4.4).

$$T_r = T_0 \frac{\omega}{\omega_{ref}} + k_T M_r \max\left[0.7,\ 1 + 0.06(m_r - 14)\right] \tag{4.4}$$

with $T_0 = 200\,\mathrm{N\,m}$ and $k_T = 26.8\,\mathrm{N\,m/kg}$. At 70 t/h of dry crop the rotor holds 26.5 kg of straw at 908 N m, which is 38 % of the plug torque $T_{plug} = 2.4\,\mathrm{kN\,m}$. The plug is reached at 82 kg of dry straw and at 58 kg at 21 % moisture, and it is a hard event. When $T_r \geq T_{plug}$ the drive command drops to zero at once, the rotor and shoe flows stop, the crop still arriving piles onto the conveyor, and after $t_{clear} = 120\,\mathrm{s}$ the grain in the rotor is pulled out onto the ground as dumped grain G_p and the rotor restarts empty.

4.3.3 Losses, Cleaning and Damage

Grain leaves the rotor at the same specific rate as the straw, and a fraction ℓ_s of it goes out of the back with the straw as separation loss, as shown in Equation (4.5).

$$\ell_s = \min\left[0.5,\ 0.0065\,\mathrm{e}^{2(\lambda_r - 1)}\,\mathrm{e}^{0.10(m_r - 14)}\,\frac{\omega_{ref}}{\omega}\right] \tag{4.5}$$

Here $\lambda_r = M_r / M_{ref}$ is the rotor load, and the reference inventory $M_{ref} = 25.5\,\mathrm{kg}$ is what the rotor holds at 9 kg/s of dry straw and 950 rpm. The loss rises exponentially with the rotor load and with moisture and falls with rotor speed, which are the directions reported for threshing and separating units [1]. Its constants were chosen so that dry crop reaches 2 % of loss at a plausible feed rate, and they are not calibrated to any machine.

The rest of the grain falls through the concave onto the cleaning shoe with a share $\chi = 0.2$ of the straw as chaff, and the fan lifts the chaff away, as shown in Equations (4.6) and (4.7).

$$\dot{G}_s = (1 - \ell_s) r G_r - \frac{G_s}{\tau_s}, \qquad \dot{C}_s = \chi r M_r - \frac{C_s}{\tau_s} \tag{4.6}$$

$$\ell_c = \min\left[0.5,\ 0.003 \max(0.2, \lambda_s) + 0.008\,\varepsilon_f^2\right] \tag{4.7}$$

$$n_{opt} = 1000\left[1 + 0.1(\lambda_s - 1) + 0.012(m_r - 14)\right]$$

Here $\tau_s = 1.5\,\mathrm{s}$, $\lambda_s = C_s / C_{ref}$ with $C_{ref} = 2.7\,\mathrm{kg}$ the chaff at the reference straw rate, $\varepsilon_f = (n_{fan} - n_{opt})/150$ with $n_{fan} = n_e n_{set}$, and n_{opt} is the fan speed in rpm at which the shoe loses least. Too little air leaves grain in a chaff mat that rides off the sieve, and too much blows it out. At 70 t/h of dry crop the shoe loses 0.31 % of the grain leaving the machine at 1000 rpm, 1.26 at 850 and 0.91 % at 1150 rpm.

Clean grain reaches the tank, and a fraction D of it is cracked, as shown in Equation (4.8).

$$\begin{aligned} \dot{G}_t &= (1 - \ell_c)\frac{G_s}{\tau_s}, \qquad \dot{G}_d = D\dot{G}_t \\ \dot{G}_{ls} &= \ell_s r G_r, \qquad \dot{G}_{lc} = \ell_c \frac{G_s}{\tau_s} \\ D &= \min\left[0.2,\ 0.008\,\mathrm{e}^{4(\omega/\omega_{ref} - 1)}\,\mathrm{e}^{0.15(14 - m_r)}\right] \end{aligned} \tag{4.8}$$

The damage is 0.80 % at 950 rpm and 14 % moisture, 1.50 % at 1100 rpm and 1.08 % at 12 % moisture, because a faster rotor threshes harder and dry grain is brittle. The grain balance of Equation (4.9) holds at every step,

$$G_{cut} = G_{conv} + G_r + G_s + G_t + G_{ls} + G_{lc} + G_p \tag{4.9}$$

where G_{conv} is the grain on the conveyor. The verification suite checks it through turns and two plugs, and the largest imbalance over the whole challenge as shipped is 6.0×10^{-11} kg.

4.3.4 Engine and Drive

One engine drives the wheels, the rotor and the fan, which all turn in proportion to its speed n_e, a fraction of rated. The power drawn is the sum in Equation (4.10).

$$\begin{aligned} P = {} & \frac{(m_{mc} + G_t)\, g\, c_{rr}\, v}{\eta_h} + T_r \omega + e_{ch}(1-\chi) r M_r \\ & + P_h n_e + e_h Q_c + P_f \left(\frac{n_e n_{set}}{1000} \right)^3 + P_a n_e \end{aligned} \tag{4.10}$$

The machine mass is $m_{mc} = 18\,\mathrm{t}$, the rolling resistance $c_{rr} = 0.08$, the drive efficiency $\eta_h = 0.75$, the chopper energy $e_{ch} = 3\,\mathrm{kJ/kg}$ of straw, the header power $P_h = 8\,\mathrm{kW}$ plus $e_h = 600\,\mathrm{J/kg}$ of crop, the fan power $P_f = 12\,\mathrm{kW}$ at 1000 rpm and the auxiliaries $P_a = 55\,\mathrm{kW}$. At 70 t/h of dry crop the machine draws 235 kW of its 330 kW, of which 39 go to the drive, 89 to the rotor, 22 to the chopper, 20 to the header, 11 to the fan and 54 to the rest.

The governor holds the engine within 2 % of rated speed up to full power, and beyond it the engine lugs down 1.5 % of speed per percent of overload, as shown in Equation (4.11).

$$\begin{aligned} \tau_e \dot{n}_e &= n_{ss} - n_e \\ n_{ss} &= \begin{cases} 1 - 0.02\lambda_E, & \lambda_E \le 1 \\ \max\left[0.5,\ 0.98 - 1.5(\lambda_E - 1)\right], & \lambda_E > 1 \end{cases} \end{aligned} \tag{4.11}$$

Here $\lambda_E = P/(\delta P_{max})$ is the engine load, $P_{max} = 330\,\mathrm{kW}$, δ is the share of power available and $\tau_e = 0.8\,\mathrm{s}$. A 10 % overload lugs the engine 17 % below rated speed, and the rotor slows with it.

The hydrostatic drive follows the applied command v_a through a lag, and the command is limited in acceleration and deceleration at every control step $\Delta t_c = 0.2\,\mathrm{s}$, as shown in Equation (4.12).

$$\begin{aligned} \tau_v \dot{v} &= n_e v_a - v \\ \Delta v_k &= \mathrm{clamp}\left(v_{cmd,k} - v_{a,k-1},\ -d_{max}\Delta t_c,\ a_{max}\Delta t_c\right) \\ v_{a,k} &= \mathrm{clamp}\left(v_{a,k-1} + \Delta v_k,\ 0,\ v_{max}\right) \end{aligned} \tag{4.12}$$

with $\tau_v = 1.2\,\mathrm{s}$, $a_{max} = 0.25\,\mathrm{m/s^2}$, $d_{max} = 1.0\,\mathrm{m/s^2}$ and $v_{max} = 11.5\,\mathrm{km/h}$. A plug sets v_a to zero at once. Fuel follows the power at 220 g/kWh of diesel at 0.84 kg/L.

4.3.5 Measurements

The feed-rate signal comes from the feeder house. The mean feed cut over each control interval, $\bar{Q}_c$, is delayed by $\theta_m = 1.5\,\mathrm{s}$, multiplied by $1 + \sigma_Q \nu_Q$ with $\sigma_Q = 3\,\%$ and passed through a first-order filter of 0.4 s. The loss monitor divides delayed, averaged loss flows, scaled by a calibration gain $1 + b$, by the flow of grain leaving the machine, as shown in Equation (4.13).

$$\begin{aligned} Q_m &= F_{0.4}\left[(1 + \sigma_Q \nu_Q)\, \bar{Q}_c(t - \theta_m)\right] \\ L_m &= 100\, \frac{\hat{S} + \hat{C}}{\hat{S} + \hat{C} + \hat{T}} \\ \hat{S} &= F_{\tau_L}\left[(1+b)(1 + 0.05\nu_s)\, \dot{G}_{ls}(t - \theta_L)\right] \end{aligned} \tag{4.13}$$

Here $\hat{C}$ is formed like $\hat{S}$ from $\dot{G}_{lc}$ with its own noise, $\hat{T} = F_{\tau_L}[\dot{G}_t(t - 2\,\mathrm{s})]$ is the yield sensor at the top of the clean grain elevator, $\theta_L = 1\,\mathrm{s}$, $\tau_L = 4\,\mathrm{s}$, F_τ is a first-order filter of time constant τ in its exact discrete form, and each ν is unit-variance Gaussian noise passed through a first-order filter, correlated over 0.5 s on the feed and over 1.5 s on the losses. The rotor torque T_m is measured with 1.5 % noise and a 0.3 s filter, the engine load λ_m with a 0.3 s filter, and a camera reports the cracked share of the grain, D_m, 3 s after it passes, filtered over 6 s. With the preview on, a forward sensor reports the density of each 2 m cell ahead of the header averaged over two cells, with a calibration gain $1 + b_p$ and 2 % of noise. The true loss share L that the trends show is $100(\dot{G}_{ls} + \dot{G}_{lc})/(\dot{G}_{ls} + \dot{G}_{lc} + \dot{G}_t)$, held at its last value through turns and stops.

Table 4.2 Machine and Field Parameters, With Hidden Ones Marked

Symbol	Parameter	Value	Class
Y, r_M	grain yield, MOG per kg grain	5.5 t/ha, 0.9	adjustable
m_0	base grain moisture	14 %	adjustable
W	header width	9 m	hidden
L_p, N_p	pass length, passes	250 m, 10	hidden
v_{max}	drive speed limit	11.5 km/h	adjustable
a_{max}, d_{max}	acceleration, deceleration	0.25, 1.0 m/s^2	adjustable
τ_v	hydrostatic drive lag	1.2 s	adjustable
v_{turn}	headland turn speed	5.4 km/h	adjustable
θ_f	cutter bar to rotor	1.5 s	hidden
m_{mc}, c_{rr}	empty mass, rolling resistance	18 t, 0.08	hidden
τ_r, Q_{cap}	rotor residence, straw capacity	2.5 s, 16 kg/s	adjustable
T_{plug}, t_{clear}	plug torque, time to clear	2.4 kN m, 120 s	adjustable
k_T, T_0	torque per kg, empty torque	26.8, 200 N m	hidden
τ_s	shoe residence time	1.5 s	adjustable
χ	MOG share reaching the shoe	0.2	hidden
P_{max}, τ_e	rated power, speed response	330 kW, 0.8 s	adj., hidden
θ_m	feed-rate signal delay	1.5 s	adjustable
θ_L, τ_L	loss sensor delay, averaging	1.0, 4 s	adjustable
t_{rain}	rain arrives	20 min	adjustable
ω_{ref}, M_{ref}, C_{ref}	reference speed, straw, chaff	950, 25.5, 2.7	fixed
η_h, e_{ch}, e_h	drive efficiency, chopper, header	0.75, 3, 0.6 kJ/kg	fixed
P_h, P_f, P_a	header, fan, auxiliary power	8, 12, 55 kW	fixed

The references are in rpm and kg. Fixed constants are not in the panel.

4.3.6 Parameters, Initialization and Numerics

The parameters are listed in Table 4.2. Those marked adjustable are sliders in the model panel, those marked hidden are part of the model but not offered there, and the constants of Equations (4.3) to (4.10) are fixed in the code. Each scenario starts at the head of pass 1, already harvesting steadily at the speed that meets the feed setpoint in the mean crop of the first 12 m, or at the manual speed. The rotor, the shoe and the engine start at the equilibrium of Equations (4.3) to (4.11), found by iterating the engine speed, with the conveyor full and the tank empty. The states are integrated by the classical fourth-order Runge-Kutta method with $\Delta t = 0.05$ s. The harvest, the conveyor, the plug and the pass bookkeeping run after each step, and the historian logs at 4 Hz. Halving Δt changes the tank at 200 s in case 2 by 0.0003 %. The texture and the noise come from the framework's seeded generator (Section 1.6.1).

4.3.7 Assumptions

One well-mixed inventory stands for the whole rotor and another for the shoe. A real axial rotor threshes at the front and separates along its length, and the grain spends less time in it than the straw does [2]. The lumped model keeps the capacity, the load dependence and the transport delay, which are what the control problem is made of.

The separation, cleaning and damage relationships are illustrative. Their shapes, exponential in the rotor load, a U in fan speed and rising with rotor speed and dryness, follow published combine behavior, while their constants were chosen [1]. They are not calibrated to any machine or crop variety.

One moisture stands for both the grain and the straw. Real straw is usually wetter than the grain, green weeds add moisture without grain, and dew wets the straw first.

The route is fixed and the field flat. The combine does not steer, the header height is constant, the reel and the header add no losses of their own, and the tank is never unloaded. No scenario field carries more than 12.9 t of grain, less than the 13 t the scene draws as a full tank.

The engine is a quasi-static governor with a first-order speed response, with no turbocharger lag and no fuel map, and the hydrostatic drive is a first-order lag with separate acceleration and

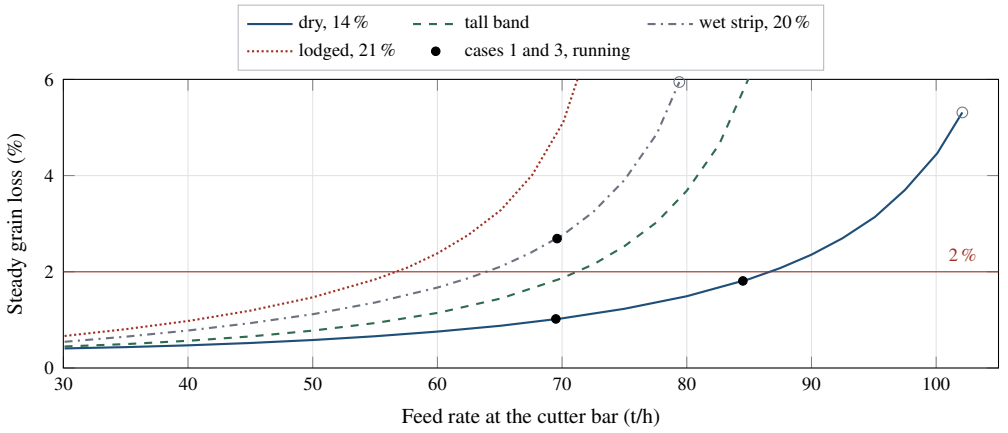

Figure 4.2 Steady loss against feed rate at 950 rpm of rotor and 1000 rpm of fan for the four crops of the scenarios: dry crop solid, the tall band dashed, the wet strip dash-dotted and lodged crop dotted, with the 2 % level as a thin line. Each curve is drawn to the feed at which the engine reaches its rated 330 kW (open circles) or to 6 %. The filled points are the running machine, case 1 at 70 and 85 t/h and the peak of case 3 in the wet strip at the mean feed of the crossing.

deceleration limits. On a real machine the speed depends nonlinearly on the pump setting and the engine speed [12].

The loss monitor is a delayed, averaged ratio of flows with a steady calibration error, and the preview measures mass with 2 % of noise. Real impact sensors count kernels and saturate at high loss rates, and real forward sensors see height and volume and remain hard to calibrate [17].

4.4 What Shapes the Control Problem

4.4.1 Throughput Against Loss

Setting every derivative of Equations (4.3) to (4.11) to zero at a fixed drive command gives the steady loss against feed, the curve on which every choice of speed sits, as shown in Figure 4.2. In dry crop at 950 rpm with the fan at 1000 rpm, the loss is 1.03 % at 70 t/h, 1.49 % at 80 and 2.36 % at 90 t/h, so a 21 % rise from 70 to 85 t/h costs 79 % more loss (1.03 to 1.85 %). The same 70 t/h loses 1.87 % in the tall band, whose crop carries 1.37 kg of straw for every kilogram of grain, 2.77 % in the wet strip and 5.11 % in lodged crop. The loss reaches 2 % at 86.7 t/h in dry crop, 71.2 in the tall band, 64.0 in the wet strip and 56.7 t/h in lodged crop. At a given feed the feed-rate signal is the same in all four crops, so nothing the feed loop measures tells them apart.

Each curve ends where the engine reaches its rated power, at 102.1, 89.0, 79.4 and 72.6 t/h, with the rotor at about 65, 75, 75 and 80 % of its plug torque. Past that point the engine lugs. In lodged crop the feed then rises to no more than 72.8 t/h, at a command of 6.5 km/h, and it falls as the command rises further while the slower rotor fills. At a command of 6.9 km/h the lugging steady state already puts the rotor above its plug torque, so in lodged crop an engine overload is the first stage of a plug.

4.4.2 A Process Gain Set by the Crop

The process gain of the feed loop is $\partial Q_c / \partial v = W \rho$, which is 9.4 t/h per km/h in uniform crop. It is 14.9 in the tall band of case 2, 1.59 times the uniform value, 11.7 in lodged crop and 4.95 in the sparse stretch of the challenge. A tuning chosen in average crop is therefore hotter in heavy crop and more sluggish in thin crop, by a factor of three between the extremes.

4.4.3 The Delay Budget

Around the uniform operating point the feed loop is a PI controller acting through the drive lag, the feed-signal filter and the transport delay, as shown in Equation (4.14).

$$L(s) = K_c \left(1 + \frac{1}{\tau_I s}\right) \frac{n_e W \rho \, \mathrm{e}^{-\theta_m s}}{(\tau_v s + 1)(\tau_f s + 1)} \tag{4.14}$$

Here $n_e = 0.986$ and $\tau_f = 0.4\,\mathrm{s}$, and the rate limits, the 0.2 s sampling and the averaging of the signal are neglected. With the shipped tuning, $K_c = 0.08$ km/h per t/h and $\tau_I = 2\,\mathrm{s}$, the loop gain $K_c n_e W \rho$ is 0.74. With the signal 1.5 s late the gain margin is 1.95 and the phase margin 57°. The gain margin falls to one at a delay of 3.82 s, and at 4 s the loop gain at the phase crossover is 1.04, with a crossover period of 15.3 s. Halving K_c restores a gain margin of 1.93 at 4 s, and lengthening τ_I to 5 s restores 1.50. In the tall band of case 2 the shipped tuning keeps a gain margin of only 1.23 and a phase margin of 23° at the nominal 1.5 s of delay. Section 4.6.4 compares these figures with the running loop.

4.5 Implemented Controller

Control executes every 0.2 s, once every four physics steps and before the integration, and the loss supervisor and the trims run on slower clocks within it. Four PI controllers share the ground speed and the feed setpoint, as shown in Figure 4.3, and their settings are listed in Table 4.3.

4.5.1 Discrete PI Algorithm

All four PI loops use the velocity-form PID of the framework (Section 1.6.2). The running sum S_k of the proportional and derivative increments, Equation (1.13), is never clamped, the output before limiting is $c_k = u_{bias} + S_k + I_{k-1}$, Equation (1.14), with the bias u_{bias} written u_b and c_k written v_k in Chapter 1, and the integral I_k is unwound by back-calculation, Equation (1.15), with $T_t = \tau_I$. The derivative term is zero here because $\tau_D = 0$ in all four loops. The acceleration and deceleration limits of Equation (4.12) act after the selector in PI operation, and no PI loop sees them, so the loop in charge keeps integrating while the drive is rate limited. In MANUAL, and whenever a schedule rather than the selector sets the speed, each speed controller tracks the command with $I_k = v_{cmd,k} - u_{bias} - S_k$, so its output equals that command exactly and the return to automatic is bumpless. All four loops set $a = +1$ because every process gain is positive. More speed means more feed, more engine load and more torque, and more feed means more loss. Each output therefore falls as its measurement rises, which is reverse action in the usual process-control convention [10].

4.5.2 The Speed Selector

Under PI control the ground speed command is the lowest of five demands, as shown in Equation (4.15).

$$v_{cmd} = \min\left(v_{feed},\ 1.1\,v_{ff},\ v_{eng},\ v_{rot},\ v_{hi}\right) \tag{4.15}$$

The feed PI acts on $Q_{\mathrm{SP}} - Q_m$, the engine PI on $E_{\mathrm{SP}} - \lambda_m$ with $E_{\mathrm{SP}} = 95\,\%$, and the rotor PI on $T_{\mathrm{SP}} - T_m$ with $T_{\mathrm{SP}} = 70\,\%$ of the plug torque. The preview speed v_{ff} is that of Section 4.5.3, and $v_{hi} = \min(v_{lim}, v_{max}) = 10\,\mathrm{km/h}$. Every PI is limited to $[0, v_{hi}]$. The incumbent keeps the drive until another demand undercuts it by 0.072 km/h (0.02 m/s), so noise cannot make two nearly equal demands trade it back and forth, except that the speed limit takes over as soon as it is lower. This incumbent margin is the one of Section 1.4.3, and Volume 19, Chapter 6 shares a belt speed between two limits through the same low selector with external reset. With the overrides off the engine and rotor controllers still run under external reset, but they are not candidates, and with the preview off neither is its cap.

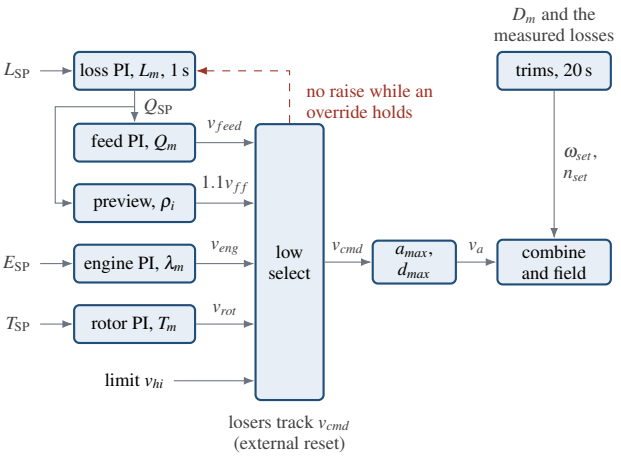

Figure 4.3 Control structure. Five speed demands meet in a low selector, and the controllers that lose track the applied command through external reset. The loss PI sets the feed setpoint once a second as the outer loop of a cascade and may not raise it while the engine limit, the rotor limit or the speed limit holds the drive. The trims step the rotor and fan speeds every 20 s.

4.5.3 Preview Braking Curve

With the preview on, the forward sensor scans the crop from the header to the range $R_p = 10\,\text{m}$ or the end of the pass. For each cell i in view, at distance d_i ahead, the speed that would meet the feed setpoint there is $v_i = Q_{SP}/(\rho_i W)$, and the preview speed is the fastest one from which the drive could still slow to every v_i in time, as shown in Equation (4.16).

$$v_{ff} = \min_i \sqrt{v_i^2 + 2d_{max}\max\left(0,\ d_i - 2.5\,v\,\tau_v\right)} \tag{4.16}$$

The term $2.5v\tau_v$ is the distance the header covers while the drive lag settles, and the square root is the braking curve at the deceleration limit. The selector sees $1.1v_{ff}$, a 10 % margin that keeps the cap clear in crop that suits the present speed, and at each pass start the uncapped v_{ff} is the entry speed. The preview is therefore feedforward that acts as a moving limit rather than as an additive bias. Placed in the selector, it slows the machine while the feed PI, which cannot yet see the heavier crop, is held on external reset instead of fighting the slowdown.

Pause. Try Exercise 4.2 before reading on.

4.5.4 External Reset

Each losing controller, after its own update, has its integral pulled toward the applied command through a first-order lag, as shown in Equation (4.17).

$$I \leftarrow I + \left(1 - \mathrm{e}^{-\Delta t_c/\tau_I}\right)\left(v_{cmd} - u_{bias} - I\right) \tag{4.17}$$

This is the lag of the applied output of Equation (1.4), with $T_t = \tau_I$, executed in its exact discrete form. Its output therefore settles at the applied speed plus its own proportional term. An override whose constraint is comfortable asks for more than the applied speed and stays out of the way, and one whose constraint is threatened drops below it at once and takes the drive without a bump [8, 9]. The feed PI is treated the same way whenever any other demand holds the drive.

4.5.5 Loss Supervisor and Trims

The loss controller is the outer loop of a cascade. Once a second it moves the feed setpoint Q_{SP} with a PI on $L_{SP} - L_m$, with $K_L = 7$ t/h per percent and $\tau_{I,L} = 20\,\text{s}$, within 14.4 t/h to 115.2 t/h and at no

Table 4.3 Controller Settings as Shipped

Loop	Measurement, SP	Gain, τ_I	Output and limits
feed	Q_m, 70 t/h	0.08 km/h per t/h, 2 s	v_{feed}, 0–10 km/h
engine	λ_m, 95 %	0.1 km/h per %, 2 s	v_{eng}, 0–10 km/h
rotor	T_m, 70 %	0.1 km/h per %, 2 s	v_{rot}, 0–10 km/h
loss	L_m, 1.5 %	7 t/h per %, 20 s	Q_{SP}, 14.4–115.2
rotor trim	D_m, ℓ_s	25 rpm per 20 s	ω_{set}, 700–1200
fan trim	ℓ_c	30 rpm per 20 s	n_{set}, 700–1300

Periods 0.2, 1 and 20 s. Selector margin 0.072 km/h, preview margin 1.1, range 10 m. Overrides on, preview, loss controller and trims off, except where a scenario sets them.

more than 0.54 t/h per second. It acts only while the machine is harvesting and more than 25 s after a pass has started, so that the rotor has refilled after the turn, and otherwise it tracks the setpoint in force. The first pass starts with the rotor full. The priming call tracks the starting feed setpoint, and the first update occurs after one simulated second. Paused pre-Run selections therefore cannot ratchet the setpoint. While the engine limit, the rotor limit or the speed limit holds the drive the feed loop is not delivering its setpoint, and the supervisor may then lower Q_{SP} but not raise it. A raise is refused by tracking the previous value, which is conditional integration in the outer loop of the cascade. The preview cap does not count as a hold, because it lets go as soon as the heavier crop is under the header.

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

Every 20 s of settled harvesting a rule steps the rotor speed ω_{set} by 25 rpm and a perturb-and-observe search steps the fan n_{set} by 30 rpm. The rotor steps down when the damage camera reads above the damage limit $D_{max} = 1.5\,\%$, and it steps up when the measured separation loss exceeds half the loss setpoint while the damage is below 85 % of its limit, within 700 rpm to 1200 rpm. The fan reverses its direction of search whenever the cleaning loss averaged over the last 12 s of a window exceeds that of the previous window, within 700 rpm to 1300 rpm. The variator and the fan follow at no more than 8 and 15 rpm/s.

4.5.6 Speed Outside Automatic Harvesting

Outside automatic harvesting a schedule or the operator sets the speed, and every speed controller tracks it. Through a headland turn the command is 5.4 km/h. For $\theta_m + 0.6\,\mathrm{s} = 2.1\,\mathrm{s}$ after a pass starts, until the feed signal carries crop again, the command stays at the turn speed or, with the preview on, at v_{ff}. A plug stops the drive at once, and the rain or the end of the field brings it to rest at the deceleration limit. In MANUAL the operator sets the speed from 0 to 11.5 km/h with a slider, or in 0.5 km/h steps from the keyboard or two buttons. The supervisor and the trims then stop, and the drive limits of Equation (4.12) still apply.

4.5.7 Model Predictive Ground-Speed Control

The optional model predictive controller (MPC) replaces the normal speed request, applying the first planned command [18], as the framework's MPC option does (Section 1.6.3). It replans at initialization and every 1 s over 30 samples in ten blocks of lengths $(1, 1, 1, 2, 2, 3, 4, 4, 6, 6)$. Integration, control and recording remain 0.05 s, 0.2 s and 4 Hz.

The ten-state predictor approximates the nominal dry machine at 950 rpm rotor and 1000 rpm fan

speed, with exact zero-order-hold discretization of Equation (4.18).

$$\begin{gathered} 1.2\dot{\hat{v}} = 0.985739u - \hat{v}, f = 9\rho\hat{v} + w \\ 0.75\dot{z}_1 = f - z_1, 0.75\dot{z}_2 = z_1 - z_2 \\ 0.4\dot{\hat{q}} = z_2 + b_q - \hat{q}, 0.775\dot{z}_3 = f - z_3 \\ 0.775\dot{z}_4 = z_3 - z_4, 1.5\dot{z}_5 = z_4 - z_5 \\ 4.038075\dot{\hat{T}} = 2.022643z_4 - 1.496689 + b_T - \hat{T} \\ 1.1\dot{\hat{E}} = 2.546671z_4 + 21.799254 + b_E - \hat{E} \\ \dot{\hat{G}} = 0.521053z_5 \end{gathered} \tag{4.18}$$

Units are m/s for $u, \hat{v}$, kg/s for $z_1, \ldots, z_5, \hat{q}$, percentage points for $\hat{T}, \hat{E}$, kg for $\hat{G}$ and seconds for lags. Nominal $\rho = 1.045\,\mathrm{kg/m^2}$ gives feed gain 9.270876 kg/s per m/s. Load slopes use steady balances. The rotor lag approximates capacity-dependent residence and instrument lag, 1.72 % below the 4.108 870 s operating-speed value. Texture-free $\pm 0.15\,\mathrm{m/s}$ steps give feed RMSE 0.0417/0.0566 kg/s, engine 0.427/0.339 points and torque 0.376/0.457 points, with feed noise off but 1.5 % torque noise retained.

Instruments and applied-command history drive the observer. Each scan, propagation gives x^- and residuals $e_i = y_{m,i} - x_i^-$ update $b_i^+ = \mathrm{clip}(b_i + 0.012\Delta t_c e_i)$, and the observed feed, engine and torque states receive $(1 - \exp(-\Delta t_c/0.6))e_i$. The integration gain is $0.012\,\mathrm{s^{-1}}$. Bias limits are $\pm 8\,\mathrm{kg/s}$ for feed and ± 60 points for each load. Integrating disturbances address constant offsets [19], without a guarantee of zero error for this nonlinear machine. Measured speed replaces $\hat{v}$, and measured yield integrates $\hat{G}$. Density comes from the scanner's first sample, or $Q_m/[9\max(0.3, v_{old})]$, bounded to 0.2–3 $\mathrm{kg/m^2}$, with v_{old} about 2 s old to align the delayed feed. Propagation uses density rounded to 0.01 $\mathrm{kg/m^2}$. Observation continues under PI and Manual. Initial flows, speed and loads equal measurements, with zero tank integral and biases.

The scanner model, range, calibration and seeded noise mechanism match. Each run's profile and braking cap share one scan without extra draws. With v_m the measured speed, the last scanner sample before distance $\max(0.25, v_m)j$ supplies ρ_j, held beyond range, with $w_j = 9(\rho_j - \rho)v_m$ in kg/s. A 10 m range covers about 4.8 s at nominal speed. After estimated arrival at a visible pass end, reference, reward and w_j become zero. The current density gain remains, so this is not an exact zero-crop or turn model.

The shared active-set quadratic solver minimizes Equation (4.19) with ten speed decisions, three slacks and a 600-iteration ceiling.

$$\begin{aligned} J = &\sum_{j=1}^{30}[(\hat{q}_j - r_j)^2 - 0.5\chi_j\hat{q}_j] + 8\sum_{b=1}^{10}(\Delta u_b)^2 \\ &+ 20s_q^2 + 10s_E^2 + 10s_T^2 \\ 0 \le u_b \le \bar{u}, &\quad -d_{max}T_s \le \Delta u_b \le a_{max}T_s \\ \hat{q}_j \le Q_{\mathrm{SP}} + 1 + s_q, &\quad \hat{E}_j \le E_{\mathrm{SP}} + s_E, \quad \hat{T}_j \le T_{\mathrm{SP}} + s_T \end{aligned} \tag{4.19}$$

Here $T_s = 1\,\mathrm{s}$, $r_j = Q_{\mathrm{SP}}$ in kg/s and $\chi_j = 1$ except after the visible end, where both are zero. The hard command ceiling $\bar{u}$ is v_{hi}, reduced by $1.1v_{ff}$ when preview is on. Rate bounds apply at block transitions from the previous applied command. The three nonnegative slacks span the horizon. Tank and terminal weights are zero. Weights use numerical SI values: $(\mathrm{kg/s})^{-2}$ for feed, $(\mathrm{kg/s})^{-1}$ for reward, $(\mathrm{m/s})^{-2}$ for moves and point^{-2} for load slacks. A 10^{-7} Hessian ridge regularizes all 13 variables. Nominal gains remain frozen when plant parameters or power change.

Loss supervision, trims, stops, turns, pass-start hold and slew remain. Armed load PIs override when measured limits are breached and all speed PIs track post-slew applied command. Nonoptimal,

Table 4.4 Headline Measurements by Case

Case, setting	Measured result
1, 70 t/h, 20 s to 100 s	7.43 km/h, loss 1.02 %
1, 85 t/h	9.03 km/h, loss 1.81 %
1, 300 s, 70 / 85 t/h	2.61 / 3.13 t, harvest loss 0.99 / 1.73 %
2, reactive	feed peak 107.8 t/h, speed down to 2.55 km/h
2, preview 10 m	feed peak 78.8 t/h, speed down to 4.12 km/h
2, range 2 / 3–4 / 5–6 / ≥7 m	feed peak 101.8 / 86.2 / 79.8 / 78.8 t/h
3, 70 t/h	loss 1.00 to 2.69 %, torque 37 to 57 %
3, preview / loss controller	peak 2.88 / 3.04 %
3, 60 t/h / 1050 rpm	peak 1.71 / 2.10 %
4, loss controller	Q_{SP} 76.3 to 65.7 t/h, loss 1.37–1.76 %
4, Manual / trims on	loss to 1.92 % / 8.74 against 8.10 t
5, $\theta_m = 4$ s	speed 6.40–8.75 km/h, period 15.8 s
5, K_c 0.04 / τ_I 5 s / θ_m 1.5 s	7.08–7.75 / 7.10–7.75 / 7.07–7.75 km/h
6, $b = -40$ %, 250 s to 500 s	monitor 1.44 %, true 2.40 %
6, Q_{SP} biased / calibrated	91.3 / 79.6 t/h
7, overrides on, 80 s to 120 s	feed 55.1 t/h, engine 95 %, loss 0.67 %
7, overrides off	engine 105 %, rotor 854 rpm, loss 1.64 %
8, overrides off	plugs at 53.15 and 203 s, 1.29 t at 400 s
8, overrides on	no plug, torque peak 88 %, 3.55 t at 400 s

malformed, nonfinite, out-of-bounds or rate-violating plans, negative feed and decreasing tank reject. First failure holds applied speed for one 0.2 s scan under current caps, then tracked PI selection acts until recovery. Noise-free first-failure command jump is zero.

4.5.8 Plant, Analysis and Controller Models

The plant as simulated is Equations (4.1) to (4.13). No model is identified in the application, and the control model used for analysis is the linearization of Equation (4.14). The classical controller model is the preview, which uses the static feed relation $Q = \rho W v$, the deceleration limit and a settling distance of $2.5 v \tau_v$, and knows nothing of moisture. The PI controllers and the supervisor are model-free. The MPC model is Equation (4.18), distinct from the nonlinear plant and the pure-delay model used for loop analysis.

4.6 Measured Scenarios

The ten scenarios are the uniform field (case 1), a dense patch (2), a wet strip (3), a moisture drift (4), a sensor delay (5), a biased loss sensor (6), a power reduction (7), a plugging threat (8), the predictive controller against the preview and the overrides (9) and the challenge against the rain (10). Each carries a completion rule (Section 1.6.4). The challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt. Table 4.4 collects the measured results. Every value was read from headless runs of the application's engine at its default seed, with each setting made after the case is loaded and before the run starts unless a time is named, over the window stated with it. The scorecard's harvest loss is the grain lost at separation, at cleaning and in plugs as a share of all the grain that has left the machine. The worst pass is the largest share lost on any one pass, counted from its start to the start of the next pass so that the turn after it, where the rotor empties, belongs to it. The pass in progress counts with its share so far, which can still fall as the pass goes on. Engine overload is the total time with $\lambda_E > 1$.

4.6.1 Throughput Against Loss

In the uniform field at the default 70 t/h, the combine drives 7.43 km/h over 20 s to 100 s, takes in 69.5 t/h at the cutter bar, draws 71 % of its engine power, runs its rotor at 38 % of the plug torque and loses 1.02 % of its grain. The first pass takes 121 s and the headland turn 19.7 s. At 85 t/h the speed is 9.03 km/h and the loss 1.81 %, so 21 % more feed costs 78 % more loss. Over 300 s the faster machine puts 3.13 t in the tank and cuts 0.58 ha, against 2.61 t and 0.48 ha, at a harvest loss of 1.73 %

against 0.99 %. Whether the extra half tonne is worth the extra loss depends on the price of grain and on how soon the rain comes, which is the decision the challenge poses. It is the trade of Section 1.4.2, where better control is paid out in throughput.

4.6.2 A Dense Patch Seen Late and Seen Early

In case 2 a 40 m band of tall wheat, 1.56 times as dense as the crop just before it, crosses every pass. The header enters it at 53.70 s. The feed at the cutter bar jumps at once and peaks at 107.8 t/h at 55.20 s, while the measured feed still reads near 70, because the feeder-house signal is 1.5 s late and filtered. By the time the loop answers, its process gain is 1.59 times the one it was tuned in, which leaves the gain margin of 1.23 found in Section 4.4.3, and it overcorrects, taking the speed down to 2.55 km/h before it settles near 4.8 km/h in the band. The rotor torque peaks at 55 % of the plug torque and the loss at 2.90 %. With the 10 m preview the machine is already slowing as the header reaches the band. The feed then peaks at 78.8 t/h, the speed bottoms at 4.12 km/h, the torque at 49 % and the loss at 2.06 %. Inside the band the loss averages 2.05 % without the preview and 1.89 % with it over 65 s to 82 s, because tall crop brings more straw per tonne whatever the speed. The preview removes the surge, not the heavier crop. The same division of labor, feedforward from a measurement upstream with feedback left to trim, is the subject of Volume 3, Chapter 6.

The range the preview needs is set by the drive. At 7.4 km/h the header covers $2.5 v \tau_v = 6.2$ m while the drive lag settles and another 1.3 m while the drive brakes at 1 m/s^2 to the 4.7 km/h the band needs, and with the 10 % margin the cap starts to bind about 7 m before the band. The feed peaks at 101.8 t/h with a 2 m range, 86.2 with 3 or 4 m, 79.8 with 5 or 6 m and 78.8 from 7 m on. The peaks change in pairs of ranges because the sensor resolves the crop in 2 m cells.

4.6.3 Moisture That No Sensor Sees

Case 3 crosses a 50 m strip that is 6 points wetter than the rest and carries 10 % more straw. At a constant 70 t/h the rotor torque rises from 37 to 57 % of the plug torque, the engine load from 71 to 85 % and the true loss from 1.00 to 2.69 %, while the feed trend shows nothing. The preview cannot help, because the strip is only 5 % denser, and the peak is 2.88 % with it. The loss controller makes the peak worse, 3.04 %. It enters the strip at 46.85 s with the setpoint already raised to 75.6 t/h by the dry crop before it, its monitor peaks at the scan at 65.25 s, 5.25 s after the true loss, and the setpoint bottoms at 62.6 t/h at the scan at 75.05 s, 2.95 s after the header has left the strip at 72.10 s. A strip that passes under the header in about 25 s is gone before a loop with a 5 s measurement lag and a 20 s integral time has finished answering it. What works is a decision made before the strip. A feed of 60 t/h holds the peak to 1.71 %, and a rotor at 1050 rpm holds it to 2.10 % while cracking 0.56 % of the grain that reaches the tank during the crossing against 0.37 %.

> **TRY IT IN THE APP** **3 · Wet Strip**
>
> Load the scenario and press Run. Watch the Feed rate trend and the Grain loss trend as the header crosses the wet strip. The feed stays at 70 t/h while the true loss climbs from 1.00 to 2.69 % and the rotor load rises from 37 to 57 % of the plug torque.

Case 4 is the disturbance the supervisor is for. Dew wets the crop by 0.3 points a minute, and the loss controller lowers the feed setpoint from 76.3 t/h at 60 s to 65.7 t/h at 880 s while the moisture at the header rises from 14.4 to 18.5 %. From 300 s on, the settled true loss stays within 1.37 % to 1.76 % and averages 1.58 % over the last 200 s against its 1.5 % setpoint. With the supervisor in Manual the setpoint stays at the 70.0 t/h it held when the case was loaded, which the switch to Manual copies into the setpoint and Reset keeps, and the loss creeps from 1.05 % at 60 s to 1.92 % at 900 s. With the trims on as well, the rotor steps up to its 1200 rpm limit, turning at 1178 rpm by 600 s because wet grain cracks less, and the machine brings in 8.74 t in 15 min against 8.10 t, over 1.61 ha against 1.50 ha.

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

4.6.4 Hunting on a Later Signal

Case 5 infers the feed from the rotor drive pressure, 4 s after the crop is cut instead of 1.5 s, with the tuning unchanged. Over 55 s to 115 s the speed hunts between 6.40 and 8.75 km/h with a period of 15.8 s, and the feed between 60 and 81 t/h, where the linearization of Equation (4.14) gives a loop gain of 1.04 at a phase crossover period of 15.3 s. Halving the gain to 0.04 keeps the speed within 7.08 km/h to 7.75 km/h, a 5 s integral time keeps it within 7.10 km/h to 7.75 km/h, and with the signal back at 1.5 s the shipped gains give 7.07 km/h to 7.75 km/h. The limit of the running loop was located in a counterfactual run with the crop texture removed, the feed noise off and the pass lengthened to 1000 m, from the oscillation left by a 2 t/h setpoint step. Its amplitude falls by 4.4 % per cycle at a delay of 3.7 s and grows by 2.0 % per cycle at 3.8 s, so the running loop reaches its limit at 3.77 s, with a period of 15.1 s, close to the linear estimate of 3.82 s. Volume 5, Chapter 3 measures the same erosion of a fixed tuning as an analyzer lag grows.

4.6.5 A Monitor That Reads Low

In case 6 the loss monitor reads 40 % low and the supervisor is in Auto at a 1.5 % setpoint. Over 250 s to 500 s of settled harvesting the monitor reads 1.44 % while the true loss is 2.40 %, as the calibration says it must be, and the supervisor has raised the feed setpoint to 91.3 t/h, against 79.6 t/h with a calibrated monitor. By 500 s the harvest loss is 1.97 % and the worst pass has lost 2.23 %. A kernel count behind the machine at 300 s, which is the recalibration event, brings the setpoint back down to 80.3 t/h at 436 s and 80.7 t/h at 500 s, with the true loss at 1.56 % over 430 s to 500 s. The loop held its measurement on setpoint throughout, and the measurement was wrong.

4.6.6 The Engine and the Rotor as Constraints

At 60 s in case 7 the engine derates to 65 % of its 330 kW. With the overrides on, the engine limit takes the drive 0.8 s later and holds the engine at 95 %. Over 80 s to 120 s the feed is 55.1 t/h against 69.5 t/h before the derate, the speed 5.99 km/h and the loss 0.67 %, because the rotor is lightly loaded, and on the way down the engine is over its available power for 6.05 s. Without the overrides the feed loop holds 70.4 t/h by demanding power the engine does not have. The engine runs at 105 % and lugs to 89.9 % of rated speed, the rotor slows from 937 to 854 rpm with it, the loss rises to 1.64 %, and the engine is still at 106 % of its available power at 180 s.

Case 8 feeds 80 t/h into 140 m of lodged crop across the first two passes with the overrides off, beyond the 72.6 t/h at which the engine reaches full power in that crop. The header enters it at 25.75 s with the rotor at 45 % of its plug torque. The feed loop holds a mean of 80.2 t/h, the engine is overloaded from 36.80 s and lugs, the rotor slows from 935 to 847 rpm and loses capacity, and the torque climbs to the plug at 53.15 s, 27 s into the patch, with 51 kg of grain in the rotor. The machine stands for 120 s, restarts in the same crop and plugs again at 203 s, and at 400 s the tank holds 1.29 t. The preview alone does not help, because lodged crop is only 1.24 times as dense, and the rotor plugs at 53.80 s. With the overrides on, the rotor limit holds the drive for 156 s of the first 300 s, the torque overshoots its 70 % limit to 88 % before the override catches it, because the crop already on the feeder must still pass, and the machine never stops. The feed in the lodged crop averages 68.4 t/h and the tank holds 3.55 t at 400 s. The override protects the rotor, not the grain. After 300 s the harvest loss is 3.49 %, because a torque limit of 70 % admits a steady loss of about 4.5 % in lodged crop.

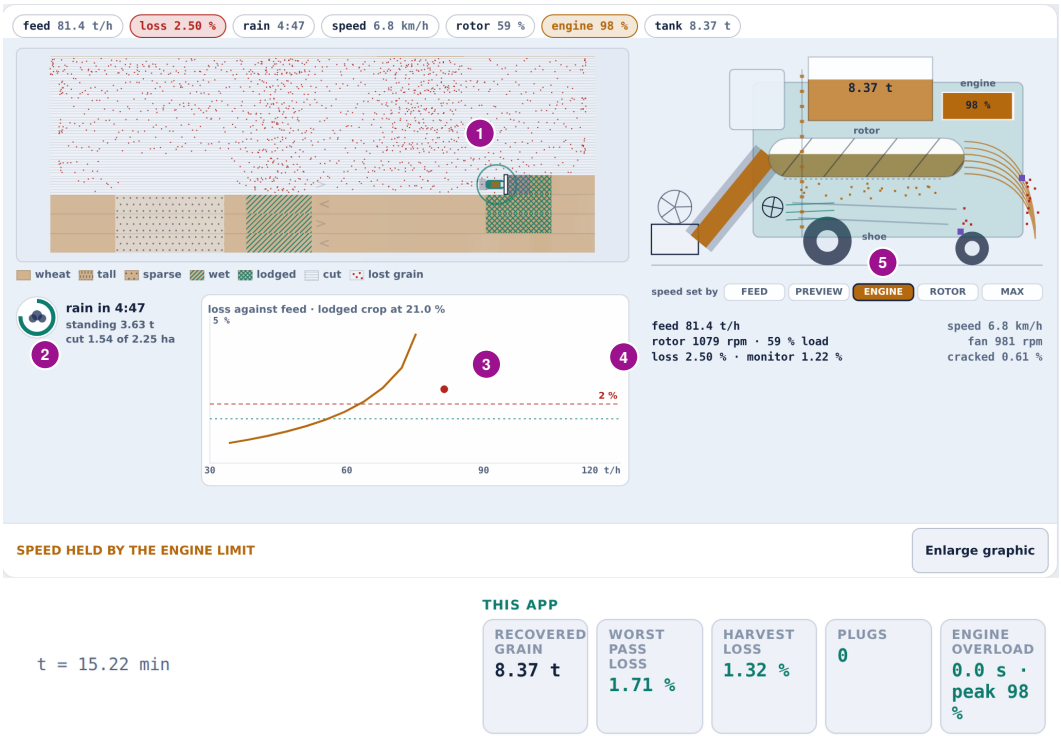

Figure 4.4 The combine harvester application 913 s into the challenge (the clock reads 15.22 min) after the suggested configuration was revealed, with the loss controller in Auto, the overrides armed at an engine limit of 90 % and the preview and the trims on. The figure is a composite: the plant view of the window and, enlarged below it, the run clock and the scorecard tiles. (1) The field is 1.54 of 2.25 ha cut, with red dots of lost grain on the stubble, and the combine is 210 m into pass 7, just inside the cross-hatched lodged patch, with the violet preview wedge ahead and the dotted sparse, hatched wet and cross-hatched lodged crop still standing in the last passes. (2) Below the field the weather clock shows the rain 4 min 47 s away with 3.63 t standing. (3) The steady loss curve is drawn for the lodged crop at 21.0 % moisture, and the operating point, at 81.4 t/h and 2.50 %, lies above the 2 % line but below that curve, because the rotor still holds crop cut before the patch, at 19.4 % moisture. (4) The readout gives the true loss against the lagging monitor (2.50 % against 1.22 %), and the cutaway shows 8.37 t in the tank and the engine at 98 %. (5) The ENGINE chip says that the engine override has the speed. The worst pass so far has lost 1.71 %.

> **TRY IT IN THE APP** **8 · Plugging Threat**
>
> Run the scenario with the Overrides off and watch the rotor load after the header enters the lodged crop. The rotor plugs at 53.15 s, 27 s into the patch. Press Reset, set Overrides to On before pressing Run, and repeat: the torque overshoots to 88 % of the plug torque, the machine never stops, and at 400 s the tank holds 3.55 t against 1.29 t.

4.6.7 The Harvest Against the Rain

The challenge field carries a tall band across passes 1 to 5, a sparse stretch in passes 6 to 10, a wet strip across every pass and a lodged patch in passes 7 to 9, 12.10 t of grain in all, with the rain at 20 min. The target is at least 10.5 t in the tank, no pass above 2 % of loss, no plug and at most 10 s of engine overload, and the run stops itself 12 s after the rain. The shipped configuration drives

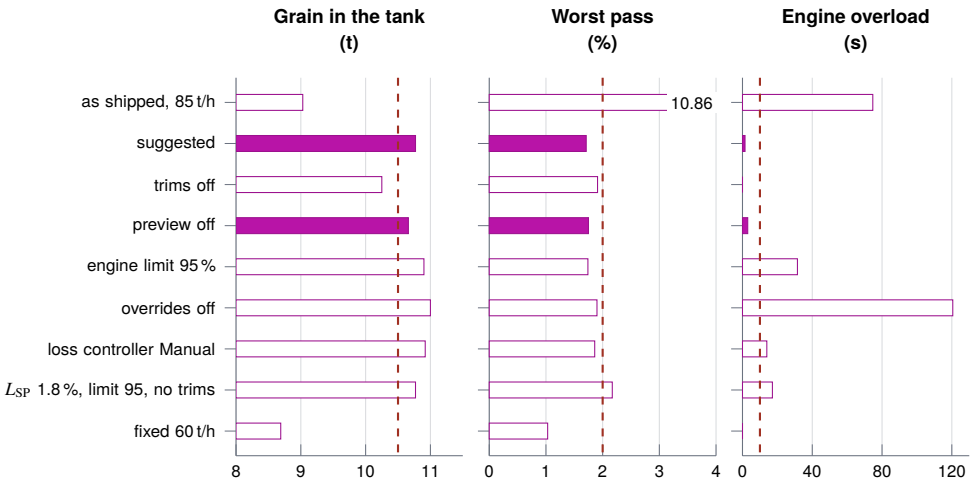

Figure 4.5 Table 4.5 redrawn against the three limits of the target, dashed: at least 10.5 t, no pass above 2 % and at most 10 s of engine overload. Solid bars meet the whole target and white bars do not. The shipped configuration's worst pass, 10.86 %, runs off its axis. Notice that each of the three configurations that bring in more grain than the suggested one fails on overload, and that the clean fixed feed fails on grain.

Table 4.5 Challenge Configurations, Read When the Run Stops Itself

Configuration	Grain (t)	Worst pass (%)	Plugs	Overload (s)	Met
as shipped, 85 t/h, overrides off	9.03	10.86	2	74.70	no
suggested	10.77	1.71	0	1.45	yes
trims off	10.25	1.91	0	0.00	no
preview off	10.66	1.75	0	3.00	yes
engine limit 95 %	10.90	1.74	0	31.50	no
overrides off	11.00	1.90	0	120.55	no
loss controller Manual	10.92	1.86	0	13.95	no
L_{SP} 1.8 %, limit 95 %, trims off	10.77	2.17	0	17.15	no
fixed 60 t/h, overrides, preview	8.69	1.03	0	0.00	no

85 t/h with the overrides off. It plugs twice in the lodged patch, at 865.5 s and 1044.8 s, losing 240 s and 104 kg of grain dumped from the rotor. It overloads the engine for 74.7 s with a peak of 109 % and brings in 9.03 t with 2.78 t still standing, as shown in Table 4.5 and drawn against the target in Figure 4.5. Its worst pass reads 10.86 % when the run stops, and that figure belongs to pass 8, still in progress and carrying the second plug, while the worst completed pass lost 7.20 %.

The suggested configuration puts the loss controller in Auto, arms the overrides with the engine limit at 90 %, and turns on the preview and the trims. It brings in 10.77 t from 2.00 ha, loses at most 1.71 % on any pass and 1.36 % over the harvest, never plugs and overloads the engine for 1.45 s, at 2.13 L/t with 1.24 % cracked grain. It passes with 0.27 t of margin on the grain, 0.29 points on the worst pass and 8.1 s on the overload. Its worst pass is the first, driven while the supervisor brings the setpoint down from the shipped 85 t/h (to 72.4 t/h by 90 s) and the trims raise the rotor from 950 to 1075 rpm. Removing the trims leaves 10.25 t, short of the target, and removing the preview still passes at 10.66 t. The default engine limit of 95 % brings in 10.90 t but overloads the engine for 31.50 s, because the load overshoots the limit its override holds in every transient, and removing the overrides reaches 11.00 t with 120.55 s of overload, although its worst pass, 1.90 %, stays inside

Table 4.6 Matched 1260 s Comparison. Paired Entries Give PI / MPC

Case, setting	Tank (t)	Loss (%)	Plugs	Overload (s)	Feed IAE (10^3 t/h s)	Travel (km/h)	Sat. (%)	Attempts/ rejects	Solve (μs)
Dense, scan On	10.150 / 10.576	1.134 / 1.328	0 / 0	0.00 / 0.00	17.06 / 20.35	342 / 306	4.76 / 4.76	1260/91	104
Dense, scan Off	10.193 / 10.229	1.190 / 1.277	0 / 0	0.00 / 0.00	19.43 / 19.88	446 / 415	4.76 / 4.76	1260/56	80
Dense, scan −30%	10.211 / 10.682	1.206 / 1.961	0 / 0	0.00 / 0.00	19.96 / 31.89	479 / 453	4.76 / 5.63	1260/0	75
Dense, power 65%	8.111 / 8.917	0.708 / 0.898	0 / 0	55.85 / 705.45	30.55 / 24.45	119 / 134	4.76 / 4.76	1260/91	109
Plug threat, scan On	11.247 / 11.890	1.877 / 2.247	0 / 0	19.90 / 24.85	23.13 / 25.34	378 / 352	4.76 / 4.76	1260/89	89
Plug threat, scan Off	11.122 / 11.474	1.903 / 2.028	0 / 0	24.45 / 3.50	24.60 / 23.64	378 / 381	4.76 / 4.76	1260/55	72
Plug threat, scan −30%	11.252 / 11.292	1.888 / 2.882	0 / 0	17.65 / 41.45	23.75 / 36.09	383 / 360	4.76 / 5.71	1260/0	87
Plug threat, power 65%	8.051 / 9.054	0.934 / 1.633	0 / 0	58.75 / 937.10	44.50 / 37.05	112 / 101	4.76 / 4.76	1260/92	118
Challenge, scan On	11.209 / 11.444	2.336 / 2.843	0 / 0	32.65 / 106.80	28.45 / 29.82	410 / 373	4.76 / 4.84	1260/84	90
Challenge, scan Off	11.120 / 11.239	2.404 / 2.538	0 / 0	16.75 / 62.95	31.11 / 30.59	446 / 467	4.76 / 4.76	1260/55	81
Challenge, scan −30%	11.266 / 10.661	2.445 / 3.720	0 / 0	22.50 / 173.00	29.63 / 37.92	429 / 394	4.76 / 6.67	1260/0	95
Challenge, power 65%	8.058 / 8.909	0.893 / 1.422	0 / 0	109.50 / 807.85	49.77 / 43.32	172 / 162	4.76 / 4.76	1260/95	116
Challenge, suggested	10.770 / 10.913	1.362 / 1.405	0 / 0	1.45 / 48.60	23.37 / 25.39	360 / 309	4.76 / 4.84	1260/88	90
Challenge, SIMC suggested	10.833 / –	1.382 / –	0 / –	0.00 / –	23.41 / –	300 / –	4.76 / –	–	–

PI / MPC entries share settings, seed and the same 10 m scan, or both turn it Off. Wrong scan has −30% calibration bias. Power falls unannounced to 65% at 60 s. PI uses shipped feed tuning. The fixed SIMC row has no paired MPC entry. Overload counts actual engine load above 100%. Scorecard IAE, travel and command saturation include stops and turns. MPC attempts/rejects include physical sequences. Solve is the mean QP time in Node.js on the measuring host, excluding observation and model construction. The dense rows are case 9, whose physics duplicates the dense case.

the limit. A fixed 60 t/h with the overrides and the preview is clean and brings in only 8.69 t. The verdict does not depend on the seed. Over five other seeds of the texture and the noise, the suggested configuration brings in 10.78 to 10.85 t with a worst pass of 1.688 to 1.692 %, and the shipped one fails every time. It plugs twice at four of the seeds, bringing in 9.01 to 9.05 t, and once at the fifth, which brings in 10.11 t but overloads the engine for 126.1 s and loses 7.47 % on its worst pass.

> **+ DID YOU KNOW? The Most Grain, and a Failed Harvest**
>
> From the simulator: against the rain, the suggested configuration with its overrides switched off brings in the most grain of any configuration, 11.00 t, yet it fails on 120.55 s of engine overload. A fixed 60 t/h with the overrides and the preview is clean and brings in only 8.69 t (Table 4.5).

With the rest of the suggested configuration fixed, the loss setpoint sets the trade. The grain rises from 9.93 t at 1.0 % to 10.45 at 1.2, 10.75 at 1.4, 10.77 at 1.5 and 10.90 at 1.6 %, and it stays between 10.89 and 10.90 t from 1.7 to 2.0 %, while the worst pass rises from 1.47 to 1.93 %. The target is met at every setpoint tested from 1.3 to 2.0 %, with 10.57 t at 1.3, and missed at 1.2 and 1.0 %. Above about 1.7 % the loss is no longer what limits the machine. The engine limit holds the drive for 264 s of the run at 1.5 %, 22 % of it, for 428 s at 1.8 % and for 524 s at 2.0 %, against only 84 s without the trims. The trims let the rotor separate better at the same feed, the supervisor spends that margin on feed, and the engine becomes the active constraint, which is how the trims earn their 0.52 t. The machine at 913 s of the suggested run, with the engine override holding the drive as the header enters the lodged patch, is shown in Figure 4.4.

4.6.8 Predictive Control with the Same Crop Scanner

Case 9 repeats case 2's tall band. Table 4.6 uses 1260 s, seed 20261009, matched instruments and protections, original suggested settings and frozen MPC tuning. A fixed SIMC feed PI [20] uses the nominal half-rule values $\tau = 1.4$ s, $\theta = \lambda = 1.7$ s and $k = 9.270876$: $K_c = \tau/[k(\lambda + \theta)] = 0.044415$ m/s per kg/s, $\tau_I = \min(\tau, 4(\lambda + \theta)) = 1.4$ s. Only feed tuning changes, with the rule frozen before comparison.

Case 9's real tank holds 1786.038 versus 1743.042 kg at 200 s, and 10576.485 versus 10 150.144 kg at 1260 s, for MPC versus PI. Harvest loss rises from 1.134 to 1.328 %. Feed

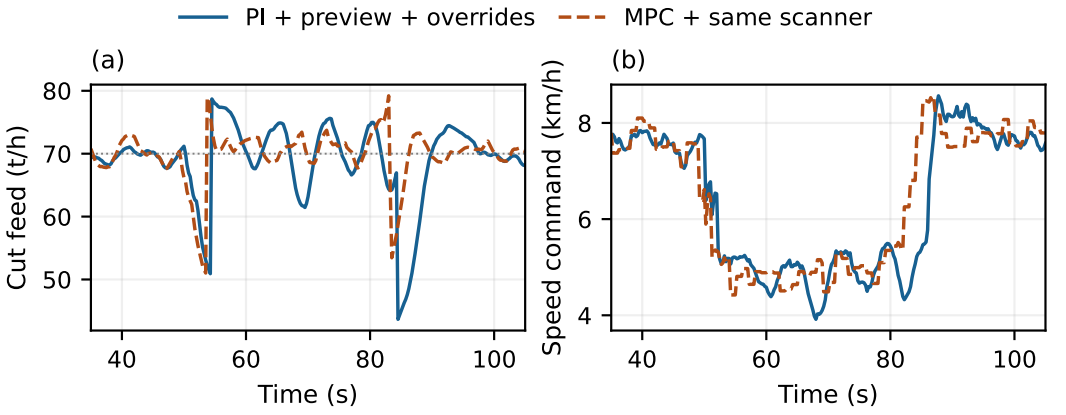

Figure 4.6 Dense-band passage with the same scanner model and range, PI with preview and overrides solid and MPC dashed. Feed rate and commanded speed over the first band show the local tracking improvement, with the whole-run throughput and loss trade reported in Table 4.6.

IAE over 45 s to 95 s falls from 262.88 to 147.57 t/h s, while whole-run IAE, including turns and stops, rises from 17060.39 to 20 354.86 t/h s. Figure 4.6 shows the local improvement.

In the plug case both protected laws avoid plugging, but MPC raises overload duration from 19.90 to 24.85 s, depth from 0.767 to 1.715 points and area from 7.624 to 16.792 point seconds. Suggested-challenge MPC collects 10912.76 versus 10 769.84 kg, with worst-pass losses 1.724 versus 1.709 %, yet fails with 48.60 s overload against PI's 1.45 s and the 10 s allowance. Fixed SIMC PI passes with 10 832.90 kg and zero overload. With unannounced power loss, dense-case MPC overloads for 705.45 versus 55.85 s.

The 76 primary, 52 zero-sensor-noise and 34 classical rows include half-step repeats. No paired ranking or challenge grade reverses, but suggested MPC overload becomes 27.45 s, without numerical convergence. Four independent grain ledgers close within 1.7×10^{-10} kg. Case 9's 1260 attempts reject 91 plans (36 visible-end, 55 after rain). The 91 s fallback flag includes 57.2 s of physical sequences and 33.8 s automatic harvesting, distinct from 992.4 s of actual MPC actuation. Chrome's full control callback averages 418.3 μs at real time (0.209% core at five scans/s). Requested 50× reaches 50.33× at 60.16 frames/s. Both guard settings have zero stops, with no observed pacing benefit. The 25 ms wrapper preserves reached steps and decisions. Four browser playback variants produce 961 identical non-wall rows, with zero Node/Chrome numerical difference in this test. Timing is host dependent.

4.7 Technology in Practice

A real combine reads the same kinds of signal and moves the same drive as the selector of Section 4.5.2. Figure 4.7 draws that loop with machine hardware in place of the schematic's symbols.

What Measures It

The automatic feed-rate controls built into combines rely on a sensor of the crop intake. The drive torque of the conveyors that carry the crop, the header auger and the feeder, gives the throughput of biomass [17]. The model's feed signal stands for such a measurement at the feeder house, 1.5 s after the cut and filtered over 0.4 s, and case 5 moves it to the rotor drive pressure, 4 s late. Grain loss is read behind the rotor and the sieves, by the pads of the next segment.

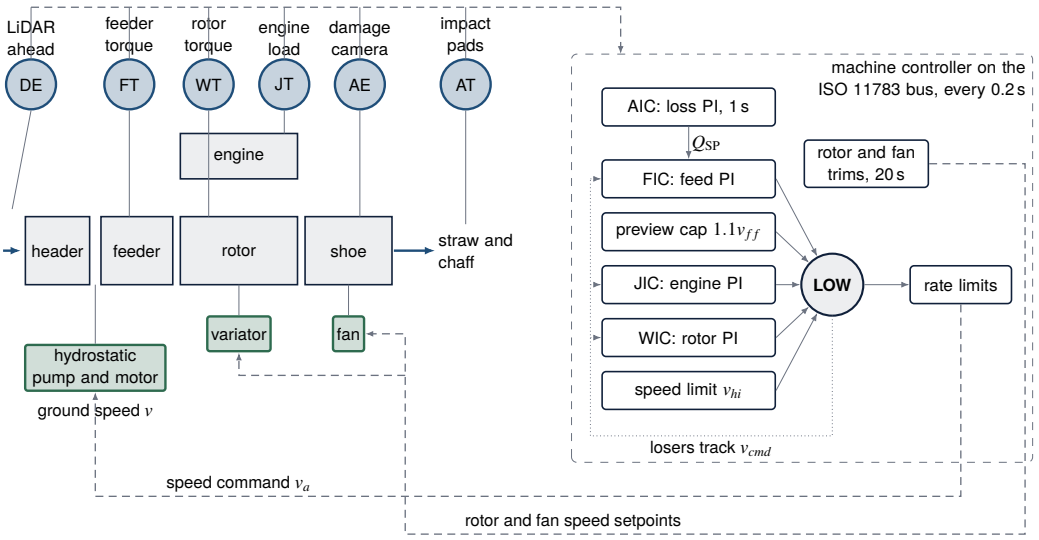

Figure 4.7 The feed loop and its overrides on machine hardware. Sensors (blue circles): DE, a LiDAR or other forward crop sensor; FT, the feed rate from the feeder drive torque; WT, the rotor torque; JT, the engine load; AE, the damage camera on the clean grain; AT, the kernel-impact loss pads behind the rotor and the sieves. Final elements (green): the hydrostatic pump and motor of the ground drive, the rotor variator and the fan. Five speed demands meet in the low selector, the losers track the applied command (dotted), and the loss PI moves the feed setpoint. Notice that only DE sees the crop before it enters, and only AT sees the grain that leaves.

TECHNOLOGY **Sensor: Kernel-Impact Loss Pads**

Grain loss sensors are plates or films of piezoelectric material that turn the impact of material into an electrical signal, which is amplified, filtered and rectified before the impacts are counted [21]. They sit where the straw and the chaff leave the machine. An extension guide asks that the sensitivity be set so that grain, not pieces of straw, triggers them, that the monitor be calibrated for the crop, and that the operator check now and then that the losses on the ground follow the signal in the cab [5]. Once the losses have been counted and the machine adjusted, the monitor serves as an early warning of change [22]. The model's monitor is a delayed ratio of flows, and its calibration error b stands for a pad that is set wrong.

What Moves It

The selector's output is a ground-speed command, and the drive that follows it is hydrostatic. A wheat guide asks the operator to use the slowest rotor speed and the widest concave spacing that still thresh the crop, and to judge threshing by the damaged kernels, the kernels left in the heads and how finely the straw is broken [5].

> **TECHNOLOGY** **Final Element: A Hydrostatic Ground Drive**
>
> The operator sets a combine's ground speed with the hydrostat, the speed control of its hydrostatic drive [22]. The heart of such a drive is a variable-displacement pump. It delivers a set volume of oil per shaft revolution, and moving a part inside the pump changes that volume. Coupled closely to a hydraulic motor, such a pump works as a continuously variable transmission between a diesel engine and a drive axle [23]. No open source read gives the drive's speed of response. The model draws it as a 1.2 s lag on the command scaled by the engine speed, behind acceleration and deceleration limits.

Where the Controller Runs

Agricultural machines share an electronic communication bus [22]. Its standard, the 14 parts of ISO 11783, builds on a controller area network at 250 kbit/s. One of its functions, Tractor-Implement Management, lets an implement set the tractor's speed and steering [24]. Such control loops run with cycle times from milliseconds to minutes and are often nested, the setpoint of one taken from the measurement of another [25]. In the application the speed loops run every 0.2 s, the loss supervisor every 1 s and the trims every 20 s. The low selector itself is textbook override control. An instrumentation text explains it with a truck driver who holds the speed limit until the exhaust temperature nears its limit on a mountain pass: whichever demand calls for less engine power wins. The same text warns that a deselected controller must not wind up [23], and external reset is this chapter's answer to that warning.

> **IN PRACTICE** **Counting Kernels Behind the Machine**
>
> Extension guides check a combine by counting the kernels on the ground behind it. In wheat, 20 to 22 kernels in one square foot are a loss of one bushel per acre [22]. The combine's own losses should not exceed one bushel per acre, and one adjustment is made at a time, each followed by a check under normal load of the losses in the field, the bin sample and the tailings [5]. The loss monitor is put to use only after such a count, as an early warning of changes in the crop or in the load [22]. The recalibration event of case 6, a kernel count behind the machine, is this check.

What Is Modernizing It

Three developments bear on the selector. The first is forward sensing. A scanning LiDAR ahead of the header estimated crop height with a root mean squared error of 0.082 m, and crop density with a coefficient of determination of 0.77, and its mounting height mattered more than its angle [17]. The second is predictive throughput control, which sets the ground speed from the feed rate, the engine load and the grain mass flow so that the capacity of the machine is used at all times [13]. The third is the bus itself. A high-speed version built on gigabit single-pair Ethernet is being developed, and functional safety is named as one factor holding back automation in agriculture [24].

4.8 Limitations

The rotor is one well-mixed inventory, so the distribution of threshing and separation along its length, which stochastic separation models resolve [2], is absent, and grain and straw share one residence time.

The loss, cleaning and damage constants are chosen rather than calibrated. Their shapes follow published behavior of rotors and shoes, but the numbers in Tables 4.4 and 4.5 belong to a teaching model and should be read as trends and ratios rather than as a machine's performance.

One moisture describes the grain and the straw, the preview sees only mass, and the loss monitor is a ratio of flows with a steady bias. The route is fixed and flat, and the engine and the drive are first-order models without the nonlinear pump and engine relation of a real propulsion system [12].

The MPC omits moisture, nonlinear capacity, power changes, losses, turns and the rain deadline. Its 30 s horizon extrapolates beyond 4.8 s of density information. Its nominal tank fraction differs from the physical ledger, and soft bands permit the suggested challenge's 48.60 s overload.

OPEN CHALLENGE **A Loss Signal Good Enough to Close a Loop On**

A 1983 study found a reliable loss measurement necessary for cheaper harvests and called better meters a pressing need [3]. A 2022 review of loss sensing still lists low resolution, impact signals that overlap when material arrives densely, interference from machine vibration, grain taken for straw and straw for grain, and changes with grain moisture. Most of the systems it reviewed had been tested in the laboratory and only a few in field harvesting, and it notes that readings in units of absolute grain loss would help operators decide [21]. Progress would be a field-proven monitor whose calibration holds across crops and moisture, so that a supervisor like that of case 6 could trust it.

4.9 Conclusion

A combine's ground speed sets its feed, and the simulator of this chapter shows how little the feed rate says about the constraints that decide a harvest. In dry crop 21 % more feed costs 78 % more loss, the same 70 t/h loses 2.69 % instead of 1.00 % in a wet strip that no sensor on the machine detects, and a loss supervisor that corrects late makes that strip worse, at 3.04 %. A dense patch seen 1.5 s late drives the feed to 107.8 t/h, and a preview that sees 7 m ahead, as far as the machine travels while its drive settles and brakes, holds it to 78.8 t/h. Moving the feed signal to 4 s makes the shipped tuning hunt between 6.40 and 8.75 km/h, and the running loop reaches its stability limit at 3.77 s, close to the 3.82 s of a linearization.

A monitor reading 40 % low lets the true loss run to 2.40 %, and a feed loop holding its setpoint in lodged crop plugs the rotor 27 s after entering it, which the rotor override prevents. Against the rain, the supervisor, the overrides with a 90 % engine limit, the preview and the trims together bring in 10.77 t with no pass above 1.71 %, against 9.03 t and two plugs as shipped. The optional MPC gains 426 kg on the dense-field comparison but fails the suggested challenge with 48.60 s of overload, while a fixed model-derived PI passes with zero overload.

The constants of the losses are chosen rather than calibrated, and the rotor is one mixed inventory. A fuller model would add a distributed rotor, separate grain and straw moisture, an impact-sensor loss monitor and a predictive controller that also plans against the time left before the rain. Within those limits the simulator poses throughput control against limits that the throughput measurement cannot see.

KEY POINTS

- A combine's ground speed sets its feed, but the feed rate says little about the limits that decide a harvest.
- In a wet strip the same 70 t/h loses 2.69 % instead of 1.00 %, and a loss supervisor that corrects late makes it worse, at 3.04 %.
- Five speed demands meet in a low selector, and the controllers that lose track the applied command through external reset, so an override takes the drive without a bump.
- Against the rain, the supervisor, the overrides, the preview and the trims together bring in 10.77 t with no pass above 1.71 %, against 9.03 t and two plugs as shipped.

Where else this applies. Volume 15, Chapter 3 moves an ore feed against the bearing pressure of a mill, and Volume 19, Chapter 6 shares a belt speed between two limits through the same low selector with external reset.

References

[1] P. Miu, *Combine Harvesters: Theory, Modeling, and Design*. CRC Press, 2015.

[2] P. I. Miu and H.-D. Kutzbach, "Modeling and simulation of grain threshing and separation in threshing units—Part I," *Computers and Electronics in Agriculture*, vol. 60, no. 1, pp. 96–104, 2008.

[3] W. Huisman, "Optimum cereal combine harvester operation by means of automatic machine and threshing speed control," Ph.D. dissertation, Landbouwhogeschool Wageningen, 1983.

[4] A. de Toro, "Influences on timeliness costs and their variability on arable farms," *Biosystems Engineering*, vol. 92, no. 1, pp. 1–13, 2005.

[5] University of Arkansas Cooperative Extension Service, “Harvesting wheat,” University of Arkansas Division of Agriculture, Little Rock, AR, Fact Sheet FSA1011-PD-3-06RV.

[6] D. Wei and I. K. Craig, “Grinding mill circuits—A survey of control and economic concerns,” *International Journal of Mineral Processing*, vol. 90, no. 1–4, pp. 56–66, 2009.

[7] P. Airikka, “Automatic feed rate control with feed-forward for crushing and screening processes,” *IFAC-PapersOnLine*, vol. 48, no. 17, pp. 149–154, 2015.

[8] F. G. Shinskey, *Process Control Systems: Application, Design, and Tuning*, 4th ed. New York: McGraw-Hill, 1996.

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

[10] 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.

[11] M. Tomizuka, “Optimal continuous finite preview problem,” *IEEE Transactions on Automatic Control*, vol. 20, no. 3, pp. 362–365, 1975.

[12] T. Coen, W. Saeys, B. Missotten, and J. De Baerdemaeker, “Cruise control on a combine harvester using model-based predictive control,” *Biosystems Engineering*, vol. 99, no. 1, pp. 47–55, 2008.

[13] T. Coen, J. De Baerdemaeker, and W. Saeys, “Throughput control on a combine harvester using model-based predictive control,” in *2010 Pittsburgh, Pennsylvania, June 20–23, 2010*. American Society of Agricultural and Biological Engineers, 2010, paper No. 1009078.

[14] M. Omid, M. Lashgari, H. Mobli, R. Alimardani, S. Mohtasebi, and R. Hesamifard, “Design of fuzzy logic control system incorporating human expert knowledge for combine harvester,” *Expert Systems with Applications*, vol. 37, no. 10, pp. 7080–7085, 2010.

[15] G. Craessaerts, J. de Baerdemaeker, B. Missotten, and W. Saeys, “Fuzzy control of the cleaning process on a combine harvester,” *Biosystems Engineering*, vol. 106, no. 2, pp. 103–111, 2010.

[16] W. Saeys, B. Lenaerts, G. Craessaerts, and J. De Baerdemaeker, “Estimation of the crop density of small grains using LiDAR sensors,” *Biosystems Engineering*, vol. 102, no. 1, pp. 22–30, 2009.

[17] J.-E. Blanquart, E. Sirignano, B. Lenaerts, and W. Saeys, “Online crop height and density estimation in grain fields using LiDAR,” *Biosystems Engineering*, vol. 198, pp. 1–14, 2020.

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

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

[21] M. I. Bomoi, N. M. Nawi, S. Abd Aziz, and M. S. Mohd Kassim, “Sensing technologies for measuring grain loss during harvest in paddy field: A review,” *AgriEngineering*, vol. 4, no. 1, pp. 292–310, 2022.

[22] D. Humburg, “Minimizing combine harvest losses in wheat,” in *iGrow Wheat: Best Management Practices for Wheat Production*. Brookings, SD: South Dakota State University Extension, 2019, ch. 28.

[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] S. Brodie, T. Oksanen, and H. Auernhammer, “Buzzword ISOBUS,” *Informatik Spektrum*, vol. 46, pp. 46–50, 2023.

[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.

4.10 Exercises

4.1 Concept check · 2 min

In scenario 3, Wet strip, the feed trend stays at 70 t/h while the true loss climbs from 1.00 to 2.69 %, and no controller moves the ground speed. Why does the feed loop not react?

(a) The feed signal arrives 1.5 s late, so the strip has passed before the loop can see it

(b) At a given feed the feed signal reads the same in dry and wet crop, so nothing it measures changes

(c) The speed command sits at its 10 km/h limit in the strip, so the loop has no room to slow down

(d) The rotor override holds the drive in the strip and keeps the ground speed constant

4.2 Concept check · 2 min

The crop preview enters the low selector of Equation (4.15) as a cap $1.1\,v_{ff}$ instead of being added to the feed PI's output as a bias. What does the cap gain?

(a) Selector inputs bypass the acceleration and deceleration limits, so the cap slows the machine faster

(b) The cap can use the crop moisture, which an additive bias on the feed loop could not use

(c) While the cap slows the machine, the feed PI tracks the applied speed instead of fighting the slowdown

(d) The cap leaves the feed setpoint to the loss supervisor, while a bias would have to move the setpoint

4.3 Concept check · 1 min

While the engine override holds the drive, the loss supervisor may lower the feed setpoint Q_{SP} but not raise it. Why is a raise refused?

(a) A raise would lift the engine load limit E_{SP} above 95 %

(b) The loss monitor reads low while an override holds, so a raise would rest on a false reading

(c) The supervisor executes once a second, too slowly to act while the override holds

(d) The feed loop is not delivering its setpoint, so a raise would wind up the outer loop without moving the feed

4.4 Theory and calculation · 6 min

Use Equations (4.3) and (4.4) with $\omega = \omega_{ref}$, so the rotor turns at 950 rpm.

(a) Find the straw capacity $\bar{Q}$ at a rotor moisture of 14, 20 and 21 %. *Checkpoint:* 16.0, 12.6 and 12.1 kg/s.

(b) Find the straw inventory M_r at which the torque reaches the plug torque $T_{plug} = 2.4\,\mathrm{kN\,m}$ in dry crop and at 21 % moisture. *Checkpoint:* 82.1 and 57.8 kg.

(c) In dry crop the steady inventory solves $rM_r = Q$. Show that 9 kg/s of straw gives the reference inventory M_{ref}, then find M_r at 12 kg/s and the factor by which Equation (4.5) raises the separation loss share. *Checkpoint:* 25.5 and 38.9 kg, $\lambda_r = 1.53$, a factor of 2.88.

4.5 Theory and calculation · 7 min

Take the feed loop of Equation (4.14) with the shipped tuning, $K_c = 0.08$ km/h per t/h, $\tau_I = 2\,\mathrm{s}$, $n_e = 0.986$, $W = 9\,\mathrm{m}$, $\tau_v = 1.2\,\mathrm{s}$ and $\tau_f = 0.4\,\mathrm{s}$.

(a) Find the process gain $W\rho$ in t/h per km/h for uniform crop, $\rho = 1.045\,\mathrm{kg/m^2}$, and for the tall band at its entry, $\rho = 1.657\,\mathrm{kg/m^2}$, and the loop gain $K_c n_e W \rho$ in uniform crop. *Checkpoint:* 9.41 and 14.9, a ratio of 1.59. Loop gain 0.742.

(b) With the signal delay moved to $\theta_m = 4\,\mathrm{s}$, evaluate the phase and the magnitude of $L(j\omega)$ at $\omega = 0.4095\,\mathrm{rad/s}$. *Checkpoint:* -50.7, -26.2, -9.3 and -93.9 degrees, about -180 in all, and $|L| = 1.04$.

(c) What gain margin does halving K_c give at $\theta_m = 4\,\mathrm{s}$? Why does lengthening τ_I not have so simple an answer? *Checkpoint:* 1.93.

4.6 Theory and calculation · 5 min

In case 2 the header approaches the tall band at $v = 7.43\,\mathrm{km/h}$ with $Q_{\mathrm{SP}} = 70\,\mathrm{t/h}$, $W = 9\,\mathrm{m}$, $\tau_v = 1.2\,\mathrm{s}$ and $d_{max} = 1\,\mathrm{m/s^2}$. The forward sensor reads $\rho_i = 1.657\,\mathrm{kg/m^2}$ in the band.

(a) Find the speed v_i that meets the setpoint in the band. *Checkpoint:* 1.304 m/s, or 4.69 km/h.

(b) Find the settling distance $2.5v\tau_v$ and the distance needed to brake from v to v_i at d_{max}. *Checkpoint:* 6.19 and 1.28 m.

(c) From Equation (4.16), at what distance ahead of the band does $1.1\,v_{ff}$ fall below the present speed, so that the cap takes the drive? Compare with the measured feed peaks for preview ranges of 2 and 7 m. *Checkpoint:* About 7.1 m.

4.7 App challenge · 8 min

Load scenario 1, Uniform field, at the default seed. Run at 50× until the run stops itself, shortly after the rain at 20 min, and read the Recovered grain, Harvest loss and Standing at the rain tiles. Reload the scenario with the reload button, set the Feed-rate setpoint to 85 t/h before Run and repeat.

(a) Report the three tiles for both runs. *Checkpoint:* At 70 t/h: 10.48 t, 1.00 % and 1.84 t. At 85 t/h: 12.23 t, 1.73 % and 0.00 t, within 0.01.

(b) Find the grain lost behind the machine in each run, $G_t L/(100 - L)$ with G_t the tank and L the harvest loss in percent, and compare the extra loss with the extra grain in the tank. *Checkpoint:* 0.106 and 0.215 t. About 0.11 t more lost for 1.75 t more in the tank.

4.8 App challenge · 8 min

Load scenario 8, Plugging threat (80 t/h, overrides off), at the default seed.

(a) Run to 400 s at 20×, read the Plugs and Recovered grain tiles, export the CSV file and find the first row with `plugged` equal to 1. *Checkpoint:* 2 plugs, 240 s lost, and 1.29 t. First plugged row at 53.25 s, within 0.1.

(b) Press Reset, set Overrides to On, run to 300 s and read the Plugs and Harvest loss tiles. In a new CSV file find the largest `rotor_torque_pct`. Then run on to 400 s and read Recovered grain. *Checkpoint:* No plug, 3.49 %, torque peak 88.2 % within 0.5, and 3.55 t.

Why does the torque overshoot a limit of 70 % by 18 points, and what does the override protect?

4.9 App challenge · 9 min

Load the challenge, Beat the rain, at the default seed. For each run below, run at 50× until the run stops itself and read the Recovered grain, Worst pass loss, Plugs and Engine overload tiles.

(a) Run the challenge as loaded (85 t/h, overrides off). *Checkpoint:* 9.03 t, 10.86 %, 2 plugs and 74.7 s.

(b) Press Reveal suggested tuning, then Reset, and run again. *Checkpoint:* 10.77 t, 1.71 %, no plug and 1.5 s.

(c) Set Rotor and fan trims to Off, press Reset and run again. *Checkpoint:* 10.25 t, 1.91 %, no plug and 0.0 s.

Which target does run (c) miss, and how do the trims earn the difference?

4.10 Application area · 3 min

In scenario 6 the loss sensors read 40 % low, $b = -0.4$, and the monitor of Equation (4.13) forms $L_m = 100\,(1+b)X/[(1+b)X+T]$ from the loss flow X and the tank flow T.

(a) What true loss does a supervisor see as its 1.5 % setpoint? What does the monitor read when the true loss is 2.40 %? *Checkpoint:* 2.48 %, and 1.45 %.

(b) How would you keep a loss-limited speed controller honest on a real machine?

4.11 Application area · 4 min

In scenario 7 the engine derates at 60 s to 65 % of its 330 kW. At 70 t/h the machine draws 235 kW: 39 to the drive, 89 to the rotor, 22 to the chopper, 20 to the header, 11 to the fan and 54 to the rest.

(a) Treat the fan, the rest, the 8 kW header base and the empty-rotor term $T_0(\omega/\omega_{ref})\,\omega \approx 19\,\mathrm{kW}$ as fixed, and the remainder as proportional to the feed. Estimate the feed at which the engine override holds the engine at 95 % of its available power. *Checkpoint:* Available 214.5 and limit 203.8 kW. Fixed 92 kW, 2.04 kW per t/h, about 55 t/h.

(b) With the overrides off the engine runs at 105 %. What engine and rotor speeds does Equation (4.11) give, and what does a slower rotor do to the grain? *Checkpoint:* 90.5 % and about 860 rpm.

4.12 Review the agent · 5 min

An AI assistant reviews scenario 3, Wet strip, run with the Loss controller in Auto, where the true loss peaks at 3.04 %, and replies: “The loss controller is too slow for a strip that lasts 25 s. Cut the Loss controller integral time from 20 to 5 s and it will catch the strip and hold the loss near 1.5 %.” Find the flaw and say how you would check it.

CHAPTER 5

Shear Heating, Overrides and Delayed Density Feedback in a Cereal Extrusion Simulator

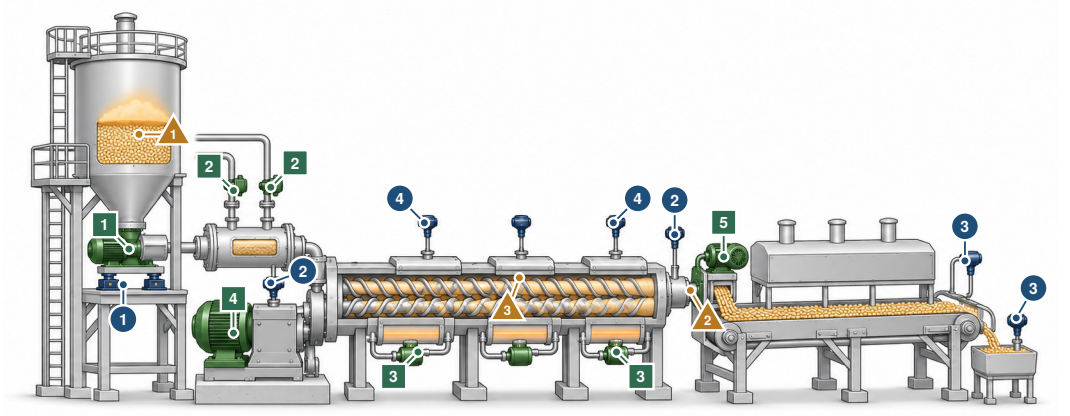

The controller moves 1 flour feeder 2 water and steam valves 3 barrel heating and cooling 4 screw speed 5 cutter knife speed

The controller reads 1 flour weight in the feeder 2 drive torque and die pressure 3 final moisture and bulk density 4 barrel temperatures

What disturbs it 1 a new flour lot 2 deposits in the die holes 3 worn screws

A twin-screw cereal extruder with its preconditioner, die, cutter and belt dryer, in which the torque and the die pressure that the overrides guard are read on the machine while the density that decides the product is known only after the die. Illustration, not to scale.

IN PLAIN TERMS **Cooking Cereal With the Push of Two Screws**

The main idea. An extruder pushes damp flour along two turning screws and out through small holes. The rubbing of the screws, not the heaters, heats the dough, and the water in it turns to steam at the holes and puffs each piece.

The challenge. A wetter or stiffer flour changes the push on the screws, the pressure at the holes and the puff of the pieces all at once. Guard controllers keep the screws and the holes inside safe limits, but the weight of a cup of pieces, the real test, comes back only minutes later.

Why it matters. A line pushed for more output can protect its machine and still fill the reject bin.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the first history segment (Section 5.1), the second history segment (Section 5.2), the fact on maintenance (Section 5.3), the two try-it segments and the fact on the challenge (Section 5.6) and the key points
Study	75 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Section 5.6), Technology in Practice (Section 5.7), Limitations (Section 5.8) and the open challenge

The control question The screws of a cereal extruder heat the melt they carry, so a wetter flour or a stiffer starch moves torque, die pressure and expansion together, while the product's density is known only minutes later. How do overrides realized as clamps, a sampled density trim and a throughput supervisor keep the line inside its limits?

Learning objectives

- Trace the melt's heat to viscous dissipation, not the barrel zones.
- Explain how two overrides realize a maximum and a minimum selector as output limits.
- Compute the error a wrong lot certificate puts into a ratio station.
- Predict where a throughput supervisor stops from its predicted constraints.
- Show that the predictive controller's gain on a wet lot comes from its estimator.

Prerequisites Sections 1.4.3, 1.4.4, 1.4.2, 1.4.6, 1.6.2 and 1.6.4, and Section 1.6.3 for the predictive controller.

The app https://apmonitor.com/apps/food-extrusion/

See also Volume 5, Chapter 10 for a soft sensor corrected by a delayed laboratory, Volume 5, Chapter 3 for control through a slow analyzer and Volume 5, Chapter 4 for estimation ahead of a late analyzer. In this volume, Chapter 4 for throughput pushed against overrides.

5.1 Introduction

Extrusion cooking makes breakfast cereals, expanded snacks and many other foods in one continuous thermal and mechanical process [1, 2]. Flour and water are metered into a preconditioner, sheared, cooked and pressurized by the screws, and forced through a die, where the water superheated in the melt flashes to steam and blows it into a foam [3]. How long this takes depends on the throughput, the screw speed and the screw profile [4]. In the simulator of this chapter the dough spends 31.6 s in the barrel at a recipe of 300 kg/h of flour and 400 rpm, and the pieces then dry on a belt for 12 min before they are sold by their bulk density and their final moisture.

> HISTORY **From Ram to Twin Screw, 1930s to 1980s**
>
> The first extruders used for food were pistons and rams that pushed meats and sausages. In the 1930s a single screw replaced the piston and turned a batch process into a continuous one, and single screws were put to work making pasta. Twin-screw extruders were established in the 1960s, and their wider scope led in the 1980s to a large variety of new snack and cereal products. An extension overview of 2017 describes the modern machine as one that mixes, kneads, cooks, forms and cuts in a single piece of equipment, with twin-screw systems the more widely used because of their flexibility [5].

The heat is what separates an extruder from a heated vessel. The power the drive puts into the screws is dissipated in the melt, so the specific mechanical energy (SME), the shaft power per unit of melt flow, measures the effort of the drive and the self-heating of the melt at once. At the recipe the screws deliver 14.8 kW, 43.0 kWh/t of melt, and carry the dough from 88.0 °C at the preconditioner to 158.0 °C at the die exit, while the heater bands and cooling water of the five barrel zones exchange a net −0.06 kW with it. That heat depends on a viscosity that falls with temperature and moisture [6], and the expansion depends on the same moisture and melt temperature [7], so moisture alone moves every output. One point of process moisture lowers the torque by 5.0 % of its rating, the die pressure by 0.51 MPa and the melt temperature by 2.3 °C, and it raises the bulk density by 11.5 g/L.

The simulator addresses the belief that an extruder whose barrel zones hold their setpoints is still making the product it made an hour ago. When a flour lot arrives three points wetter than its certificate, every heated zone reads within 0.91 °C of its setpoint over the first 20 min while the density at the cutter climbs from 135 to 164.6 g/L, out of its 120 to 150 g/L window. The coupling also runs the other way, since a falling torque is the first sign of a worn screw. The class of problem is general. In polymer extrusion the rheology of the material decides both the energy the screw consumes and

the thermal quality of the melt it delivers [8], and any process whose quality is set by a nonlinear, self-heating step and measured downstream has the same structure. Inferential control of product quality attributes that are not measured on line has been applied to food cooking extrusion [9].

Control studies of cooking extruders identify dynamic models of the process and design predictive and inferential controllers on them [10, 11]. However, a model identified at one recipe carries its gains as numbers, and a student who tunes a loop against it cannot see that the gains are the heat balance of a melt heating itself, or why they move with every lot of flour.

This chapter accounts for the heat, the pressure and the product of a twin-screw cereal extruder in a simulator whose loops, overrides and quality trim can each be switched, retuned and measured. Extrusion cooking, ratio control, override selection, sampled trims and cascade control are all standard practice [12, 13]. The simulator rests on a 105-state model in which the shaft power equals the heat released in the melt by construction, the dry-solids, water and energy ledgers close to round-off, and every parcel of product keeps its true density and moisture apart from what the instruments report. The model is set down in full so that it can be rebuilt. Its cook, expansion and density correlations are calibrated teaching proxies, and its thermal treatment is not a validated food-safety lethality model. The optional model predictive controller (MPC) treats the overrides as its constraints, as industrial practice does [14]. Three of its elements are specific to this line. The at-line density, reported 90 s after its cup is filled, corrects the prediction made for the moment the cup was taken. The soft sensor's error, shared by every controller that trusts the certificate, is estimated as an unknown moisture entering with the water, so that a torque change the plan did not cause is read as moisture. The torque and the die pressure pay for an excess at every step of the plan's first minute, which is what kept the plan from crossing the 90 % torque limit. Nine cases and an hour-long challenge set a wrong certificate, a stiffer starch, a push for throughput and a worn screw against the same coupling, each from its own side.

Section 5.2 traces how extrusion cooking came to be modeled, measured and controlled. Section 5.3 sets out the model of the line and Section 5.4 its heat balance and steady-state gains. Section 5.5 describes the implemented loops, the two overrides, the density trim, the throughput supervisor and, in Section 5.5.9, the predictive controller. Section 5.6 gives the cases as measured, and Section 5.6.8 compares the predictive controller with the loops. Section 5.7 sets the loops on real line hardware, Section 5.8 states the limitations, Section 5.9 concludes, and Section 5.10 holds the exercises.

5.2 Extrusion Cooking, Its Measurement and Its Control

A review of food extrusion ran to 61 pages by 1979 [1], and the processing studies that followed took the machine apart one mechanism at a time. The mean residence time measured on a co-rotating twin-screw cooking extruder in 1986 depended strongly on the throughput, the screw speed and the screw profile and not on the water content, the die size or the temperature, and a combination of plug flow and perfect mixing described the dimensionless distributions [4]. The rheology of low-moisture molten maize starch was measured the following year [6].

Expansion, the property that decides the product, was modeled next, and the instrument that reads the melt was found wanting, as the history segment below tells. The mechanism of expansion was later summarized for cereal foods. Moisture turns the glassy starch matrix rubbery at process temperatures, superheated steam bubbles form at nuclei and grow, and the biaxial extensional viscosity of the matrix governs how far they expand [3]. Expansion and bulk density can be read only from product that has already left the die.

Control followed the modeling. The modeling and process control of twin-screw cooking extruders were reviewed in 1994 [10]. By 2000 on-line system identification was used to design model predictive and inferential controllers for an extrusion cooking process [11], and in 2001 inferential control of product quality attributes was applied to food cooking extrusion [9], the year a chapter on effective process control appeared in a book on extrusion cooking [12]. The sequence runs from modeling the

process toward inferring what cannot be measured on it, which is also the order of the simulator's layers.

> ⧖ HISTORY **Modeling the Puff, Misreading the Melt, 1988 to 1997**
>
> A 1988 study of corn extrudates from a single-screw extruder put expansion into one model. Lengthwise and volumetric expansion indices were defined beside the sectional one, an inverse relation between radial and lengthwise expansion was confirmed, and one model in the dough moisture, the melt temperature and the die and screw shear strains described all three [7]. Measuring those inputs proved harder. In 1997 infrared and extended melt probes agreed on the bulk temperature of the extrudate, while the standard flush-mounted probe differed substantially from both and may not read it accurately. The heat lost from the surfaces of a twin-screw barrel varied between 2 and 19 % of the total heat input [15].

An extruder also does things that no loop is asked to control. Extrusion cooking destroys antinutritional factors, gelatinizes starch and raises the soluble dietary fiber, while Maillard reactions lower the value of the protein by amounts that depend on the raw material and the conditions, and heat-labile vitamins may be lost to varying extents [2]. Where an extruder is the kill step of a process, its lethality is established by experiment, for example for *Enterococcus faecium* in a balanced carbohydrate-protein meal over a response surface design [16]. The simulator models neither, and its cook index describes starch conversion only.

5.3 Model of the Line

At its recipe the simulated line turns 300 kg/h of flour at 12 % moisture into an expanded cereal. Moisture is a wet-basis mass fraction throughout, except for the dry-basis ratio X on the dryer belt. A loss-in-weight feeder meters the flour into a steam preconditioner, a ratio station adds the water that makes the dough 23.3 % water, and two co-rotating screws carry it through five barrel zones of 220 mm, a feed zone, a conveying zone, a kneading block, a second conveying zone and a metering zone. A die of 20 holes of 2.8 mm diameter sets the pressure, a four-knife cutter sets the piece length, and a belt carries the pieces through the dryer for 12 min. Every actuator, instrument and equation symbol has its place on the drawing of the line in Figure 5.1. A feeder of this kind reads its flow as the falling weight of its hopper, and Volume 12, Chapter 12 shows what such a weighing loop can and cannot see of the product.

> ✦ DID YOU KNOW? **Two Screws Cost More to Keep**
>
> A twin-screw extruder costs 1.5 to 2.0 times as much to maintain as a single-screw one, an extension overview of 2017 notes, yet twin screws are the more common in the food industry for their wide range of operating conditions and products [5].

The variables are sorted by role in Table 5.1, and 105 states are integrated. Each of the two preconditioner cells and ten barrel cells carries six, dry solids S, water W, enthalpy H, cooked solids G, amylose-weighted solids A and a dye R, and there are five barrel-metal temperatures, the filled length, four drives, the dryer air, the die deposit and 21 cumulative ledgers of what has entered and left. The product on the belt is carried as discrete parcels, updated after each integration step.

5.3.1 Feed, Preconditioner and the Chain of Cells

The feeder, the two valves and the main drive follow their commands through first-order lags, as shown in Equation (5.1).

$$\begin{aligned} \tau_F \dot{F} &= u_F F_{max}(1+\varepsilon_F) - F, \qquad \tau_N \dot{N} = N_{cmd} - N \\ \tau_W \dot{W} &= u_W W_{max} - W, \qquad \tau_S \dot{S} = u_S S_{max} - S \end{aligned} \tag{5.1}$$

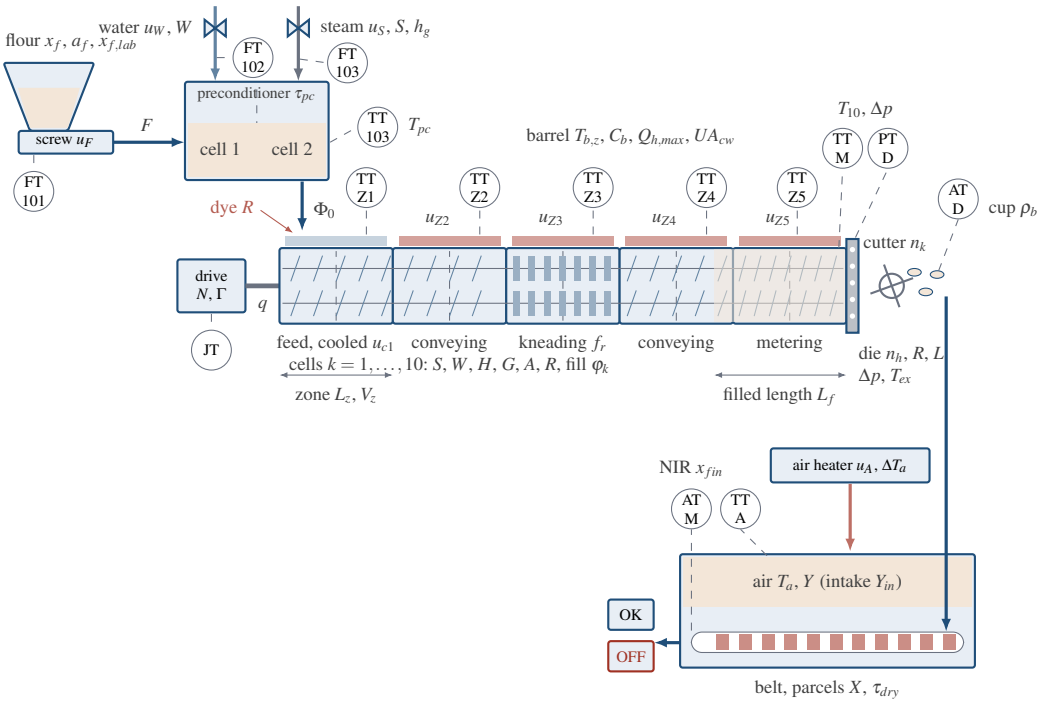

Figure 5.1 Cereal extrusion line as modeled. The loss-in-weight feeder meters flour F of true moisture x_f and amylose a_f, whose certificate states $x_{f,lab}$, and valves add water W and steam S to a preconditioner of two mixed cells. The two screws, turned at speed N against torque Γ, carry the dough through ten mixed cells in five zones, each zone with one split-range output u_Z to its heater bands or cooling water, and the metering cells in front of the die run full over the filled length L_f. The die of n_h holes sets the pressure Δp, the cutter at n_k sets the piece length, and the belt of the dryer, drawn below the barrel and running from right to left, carries 10 s parcels through air at T_a and humidity Y. AT-D weighs a cup at the cutter every minute and AT-M reads the moisture at the belt exit.

where u_F, u_W and u_S are fractions of full output, N is in rpm, and ε_F is a pulsation of the feeder's real flow, colored noise of 0.6 % standard deviation and 2 s correlation time. The two preconditioner cells and the ten barrel cells form one chain. For each cell variable $Z \in \{S,W,H,G,A,R\}$ the balance is Equation (5.2).

$$\frac{\mathrm{d}Z_j}{\mathrm{d}t} = r_{j-1}Z_{j-1} - r_j Z_j + \sigma_j^Z, \qquad T_j = \frac{H_j}{c_s S_j + c_w W_j} \tag{5.2}$$

Here r_j is the discharge rate of cell j, $x_j = W_j/(S_j+W_j)$ its moisture, $c_j = G_j/S_j$ its cook index and $a_j = A_j/S_j$ its amylose share. The first cell receives dry solids $F(1-x_f)$, water Fx_f+W+S, enthalpy $Fh_{fl}+Wc_wT_w+Sh_g$ and amylose $a_fF(1-x_f)$, where $h_{fl} = [(1-x_f)c_s + x_f c_w]T_{fl}$, so the steam condenses completely. The dye enters with the dough at the first barrel cell. In the preconditioner $r_j = o/(\tau_{pc}/2)$, with $o = \mathrm{clamp}(N/N_{min},0,1)$ so that a stopped screw also stops the discharge, the enthalpy source is the heat loss $-(UA_{pc}/2)(T_j - T_{amb})$, and the cook source is thermal only.

5.3.2 The Starve-Fed Barrel

The feeder, not the screws, sets the throughput. A barrel cell passes on what it holds at the drag capacity of the screws, except the fraction f_k of its volume that is completely full, which passes on

Table 5.1 Roles and Recipe Values of the Line Variables

Variable	Role	Recipe value	Range
ρ_b (AT-D)	CV, QIC, sampled	135 g/L SP	120–150
x_{fin} (AT-M)	CV, MIC, delayed	3.00 % SP	2–4
$\hat{x}$	estimate, ratio	23.3 % SP	19–28
F (FT-101)	CV, FIC-101	300 kg/h SP	150–540
W (FT-102)	CV, FFC-102	22.5 kg/h	0–100
T_{pc} (TT-103)	CV, TIC-103	88 °C SP	60–98
$T_{b,2..5}$ (TT-Z)	CVs, split range	115/140/140/140 °C	60–180
Γ (JT)	constraint, TQC	65.4 %	85, trip 100
Δp (PT-D)	constraint, PIC	6.77 MPa	9, trip 12
T_{ex}	constraint	158.0 °C	up to 170
T_a (TT-A)	inner CV, TIC-A	133.7 °C	90–149
u_F, u_W, u_S	MVs	50.0/22.5/54.3 %	0–100
$u_{Z2..Z5}$	MVs, heat or cool	44/40/14/−19 %	±100
N	MV of TQC	400 rpm	150–450
F_{SP}	MV of PIC	300 kg/h	60–540
Δx	MV of QIC	0	±4 points
u_A	MV of TIC-A	72.4 %	0–100
n_k	knife, ratio	755 rpm	100–3000
x_f, a_f	DV, unmeasured	12, 25 %	8–18, 15–40
$x_{f,lab}$	certificate	12 %	8–18
w, h_3	DV, wear, heaters	0, 100 %	0–80, 0–100
r_δ, Y_{in}	DV, deposit, air	0, 12 g/kg	0–2, 3–60
S, W, H, G, A, R	states, 12 cells	see text	—
$T_{b,z}$, L_f, T_a, δ	states	74.6–140, 163 mm	—

only what arrives, as shown in Equation (5.3).

$$\begin{aligned} r_k^{-1} &= \frac{V_c(1-f_k)}{qN/60} + \frac{\rho V_c f_k}{\Phi_{k-1}} \\ f_k &= \max\Big[f_{r,z},\ \mathrm{clamp}\Big(\frac{L_f-(10-k)L_c}{L_c},0,1\Big)\Big] \end{aligned} \tag{5.3}$$

where $k = 1,\ldots,10$ counts the barrel cells, $V_c = V_z/2$ and $L_c = L_z/2$ are the volume and length of a cell, $qN/60$ is the volumetric drag capacity, Φ_{k-1} is the mass flow into the cell, $f_{r,z}$ is the fraction a restrictive element holds full (0.35 in the kneading block, zero elsewhere) and L_f is the filled length in front of the die. At steady state the fill $\varphi_k = (S_k + W_k)/(\rho V_c)$ is $\varphi_{st} + (1-\varphi_{st})f_k$ with $\varphi_{st} = \Phi/(\rho qN/60)$. At the recipe the drag capacity is 1092 kg/h against a flow of 344.2 kg/h, so the conveying cells run 31.5 % full, the kneading block 55.5 % and the last cell full. A faster screw holds less dough, which is how screw speed sets the residence without changing the throughput.

5.3.3 Viscosity, Dissipation and Torque

The melt is a power-law fluid of index n whose consistency is the product of four factors, as shown in Equation (5.4).

$$\begin{aligned} K = K_{ref}\,\mathrm{e}^{E_\eta\left(\frac{1}{T^*}-\frac{1}{T^*_{ref}}\right)}\,\mathrm{e}^{-\beta_X(x-x_{ref})} \\ \times\left[1-\kappa_c(1-c)\right]\mathrm{e}^{\beta_A(a-a_{ref})} \end{aligned} \tag{5.4}$$

with T^* the absolute temperature, floored at 30 °C. The consistency falls with temperature and moisture and rises with the cook and with the amylose share of the starch. Each barrel cell dissipates

Table 5.2 Screw Profile, Fixed in the Source

Zone	Element	C_γ	χ	f_r
1	feed, cooled at u_{c1}	20	0.10	0
2	conveying	20	0.30	0
3	kneading block	40	0.55	0.35
4	conveying	25	0.40	0
5	metering	30	0.60	0

the power of Equation (5.5).

$$P_k = \chi_z K_k \dot{\gamma}_z^{n+1} V_c \varphi_k, \qquad \dot{\gamma}_z = C_{\gamma,z} m_z \frac{N}{60} \tag{5.5}$$

where χ_z and $C_{\gamma,z}$ are the dissipation factor and the shear rate per revolution per second of the screw elements in zone z, listed in Table 5.2, and $m_z = \max(0.1, 1 - w_s w)$ in zones 3 to 5 and 1 elsewhere, with w the wear fraction. The same P_k heats the cell and loads the shaft, as shown in Equation (5.6).

$$P_{sh} = \sum_{k=1}^{10} P_k + \Delta p\, Q_v, \qquad \Gamma = \frac{P_{sh}}{2\pi N/60}, \qquad \mathrm{SME} = \frac{P_{sh}}{\Phi_{10}} \tag{5.6}$$

with Δp the die pressure, Q_v the volumetric flow through the die and Φ_{10} the melt flow. Shear heating and mechanical work therefore cannot disagree. The application checks $2\pi N\Gamma/60$ against the sum to 10^{-12} and the integral of $\mathrm{SME}\,\Phi_{10}$ against the energy ledger to 1.8×10^{-5} over 600 s.

5.3.4 Barrel Metal and the Split-Range Zones

Each zone's barrel section is one heat capacity, heated by its bands or cooled by its water through a single output $u_z \in [-1, 1]$, a split range at zero output in the sense of Section 1.4.6, as shown in Equation (5.7).

$$\begin{aligned} C_b \frac{\mathrm{d}T_{b,z}}{\mathrm{d}t} &= u_z^+ h_z Q_{h,max} - u_z^- U A_{cw} (T_{b,z} - T_{cw}) \\ &\quad - \sum_{k \in z} q_k - U A_{amb} (T_{b,z} - T_{amb}) \\ q_k &= \tfrac{1}{2} U A_b \left[0.3 + 0.7 \min(1, \varphi_k)\right] (T_{b,z} - T_k) \end{aligned} \tag{5.7}$$

where $u^+ = \max(0, u)$, $u^- = \max(0, -u)$, h_z is the share of heater capacity still working (1 except in zone 3, where it is a disturbance) and q_k is the heat passed to the dough in cell k, scaled by how much of the wall the dough wets. The enthalpy source of a barrel cell is $q_k + P_k$. The feed zone has no loop, and its cooling valve is held at $u_1 = -u_{c1}$.

5.3.5 Die and Filled Length

The die is n_h capillaries of radius $R = R_h(1 - \delta)$, where the deposit δ grows at the disturbance rate r_δ in percent of the radius per minute and is capped at 0.6. The pressure drop of a power-law melt with the end correction of [17] is Equation (5.8).

$$\begin{aligned} \Delta p &= \frac{2\tau_w (L + eR)}{R}, \qquad \tau_w = K_{10} \dot{\gamma}_w^n \\ \dot{\gamma}_w &= \frac{3n+1}{n} \frac{Q_v}{n_h \pi R^3} \end{aligned} \tag{5.8}$$

where $Q_v = \Phi_{10}/\rho_m$ and e is the entrance correction in radii. At the recipe the wall shear rate is 2259/s, the apparent viscosity 58.9 Pa s and the pressure 6.77 MPa. The screws must run completely

full over the length that builds this pressure, which follows its target through a one-second lag, as shown in Equation (5.9).

$$\begin{aligned} \tau_L \dot{L}_f &= \mathrm{clamp}(\Delta p / G_p, 0, 3L_z) - L_f \\ G_p &= \frac{c_G K_{10} \dot{\gamma}_5^n}{h} (1-\varphi_{st})(1-w) \end{aligned} \tag{5.9}$$

with h the channel depth, $\varphi_{st} \le 0.95$ and both bracketed factors floored at 0.05. A screw turning below $N_{min}/2$ holds its filled length. The die adds the heat of its own pressure work, so the melt leaves at $T_{ex} = T_{10} + \Delta p/(\rho_m c_p)$ with $c_p = (1-x)c_s + x c_w$, 2.28 °C above the melt in front of the die at the recipe. A filled length beyond $2L_z$ reaches the kneading block and raises an alarm that clears below $1.8L_z$.

5.3.6 Cook

The cook index is first order in the uncooked starch, as shown in Equation (5.10).

$$\begin{aligned} \sigma_j^G &= \nu_j \, (S_j - G_j) \\ \nu_j &= k_c \, \mathrm{e}^{-E_c \left(\frac{1}{T^*} - \frac{1}{T^*_{c,ref}} \right)} \, \mathrm{e}^{\beta_c (x - x_{ref})} + k_M \frac{P_j}{S_j + W_j} \end{aligned} \tag{5.10}$$

Its thermal part, with the temperature floored at 20 °C, rises steeply with temperature and with moisture, and its mechanical part exists only in the barrel, so a screw that shears less also cooks less. Because it is evaluated in every mixed cell, the die sees the cook averaged over the whole residence-time distribution, 0.945 at the recipe, with 0.196 already done in the preconditioner.

5.3.7 Flash, Expansion and Bulk Density

This part of the model is a calibrated teaching proxy, chosen to give the directions reported for expanded cereals [3, 7] and labeled illustrative in the application. It is not a validated texture or density model for any flour or extruder. At the die exit the melt flashes adiabatically to 100 °C, losing the mass fraction and keeping the moisture of Equation (5.11).

$$e = \min\left[\frac{c_p (T_{ex} - 100)}{\lambda}, \, x - 0.03 \right], \qquad x_{ext} = \frac{x - e}{1 - e} \tag{5.11}$$

The sectional, lengthwise and volumetric expansion indices and the bulk density of the dried pieces follow from Equation (5.12).

$$\begin{aligned} \mathrm{SEI} &= 1 + (\mathrm{SEI}_{max} - 1) \, g_T \, \mathrm{e}^{-\beta_E (x - x_{ref})} \left(\frac{\Delta p}{p_{ref}} \right)^{p_E} c^3 \\ g_T &= \left(\frac{T_{ex} - 100}{\theta_T} \right)^{m_T} \left[1 + \mathrm{e}^{(T_{ex} - T_{col})/w_{col}} \right]^{-1} \\ \mathrm{LEI} &= \mathrm{LEI}_{ref} \, (\mathrm{SEI}/4.5)^{-\lambda_L} \\ \rho_b &= \frac{k_{pack} \, \rho_m (1-x)}{\mathrm{SEI} \cdot \mathrm{LEI} \, (1 - x_{fin})} \end{aligned} \tag{5.12}$$

where Δp is floored at 0.2 MPa, $g_T = 0$ below 100 °C and x_{fin} is the final moisture. The expansion rises with the superheat, the cook and weakly with the die pressure, falls with moisture, collapses above T_{col}, and trades length against section, which is the inverse relation of [7]. At the recipe the melt flashes 5.66 % of its mass, leaves at 18.70 % moisture with an SEI of 4.53 and an LEI of 1.047, and makes 135.0 g/L at the 3.0 % target moisture. The pieces are $D = 2R\sqrt{\mathrm{SEI}}$ across, 5.96 mm, and $\ell = v\,\mathrm{LEI}/f_{cut}$ long, with v the velocity in a hole and $f_{cut} = n_{knife} n_k / 60$ the cutting frequency, 12.0 mm at a knife speed of 755 rpm.

5.3.8 Belt Dryer and the Product Record

The cutter forms a parcel every 10 s from the dry solids and the water that left the die, less the flashed vapor, with flow-weighted means of the expansion, density, cook and piece size. Each parcel rides the belt for τ_{dry} and dries first-order toward the equilibrium of the air it meets, which is integrated exactly across each step, as Equation (5.13) states.

$$\begin{aligned} \dot{X} &= -k(X - X_{eq}), \qquad X_{eq} = a_E \left(\frac{RH}{1 - RH} \right)^{b_E} \\ k &= k_d \, \mathrm{e}^{(T_a - T_{d,ref})/\theta_D} \, \frac{\max(0, T_a - T_{wb})}{T_{d,ref} - 42} \end{aligned} \tag{5.13}$$

where X is the dry-basis moisture, X_{eq} follows the isotherm of [18], $RH = p_v / p_{sat}(T_a)$ is capped at 0.95, $p_v = 101.325\,Y/(0.622 + Y)$ kPa, p_{sat} comes from the Antoine equation, and the wet-bulb temperature is a linear fit to the psychrometric chart between 100 and 160 °C, $T_{wb} = 42 + 0.12(T_a - 130) + 250(Y - 0.020)$. The air follows its heater through $\tau_a \dot{T}_a = T_{in} + u_A \, \Delta T_a - T_a$, and it carries the intake humidity plus the water just evaporated, $Y = Y_{in} + E/\dot{m}_a$, with E the evaporation rate through a 30 s lag. A parcel leaving the belt is judged on its true properties: dense above 150 g/L, light below 120 g/L, wet above 4.0 % and dry below 2.0 %. The controllers never read these records.

5.3.9 Numbers, Initialization and Numerics

Every parameter is listed in Table 5.3. Fifty are shown in the parameter panel and may be moved while the line runs, and 35 are hidden. Each scenario opens on the steady state of the line at its recipe. When a setting chosen before Run would put the torque or the die pressure past an override's limit, the start is the steady state with that override holding it, the screws faster or the feed lower, found by false position. The solver sets the preconditioner analytically, with the water from the certificate and the steam that holds its setpoint, solves each barrel cell's energy balance by bisection, which is monotone because a hotter melt dissipates less, iterates the cook and the dissipation together, and closes the filled length by substitution damped by 0.6. At the recipe it returns 22.48 kg/h of water, 21.71 kg/h of steam, a filled length of 163.2 mm, a torque of 353.1 N m (65.38 %), the feed-zone metal at 74.6 °C and the zone outputs at 44.0, 39.9, 14.2 and −19.1 %. The dryer air, 133.7 °C, is found by bisection for the 3.0 % target, and the belt is filled with 72 parcels at the steady drying profile. The largest relative derivative at the stored state is 1.7×10^{-16}.

On the three clocks of Section 1.6.1, the integrator is fourth-order Runge-Kutta at a fixed 0.1 s step, the controllers execute every 0.5 s, and the historian logs once per simulated second. Halving the step moves the cutter density at 900 s of the wetter-flour case by 2.6×10^{-12} g/L, because the controllers then execute at the same instants, and parcels of 5 s instead of 10 s move the final moisture of the humid-air case at 1200 s by 1.4×10^{-3} points. Over 30 min of the wetter-flour case with a forced trip and a restart the dry-solids and water ledgers close to 1.9×10^{-15} of their throughput, and over 20 min of the heater case the energy ledger closes to 1.2×10^{-15}. These are among the application's 55 passing checks. The random streams start from the seed 20261204, the default seed, so a run replays to the digit.

5.3.10 Assumptions

The barrel is ten well-mixed cells in series with residences from a starve-fed conveying law. The cells are a residence-time model, not a discretization, and they spread a pulse from its first appearance at the die, while a real screw adds a plug-flow delay [4] and narrower tails.

The melt is a power-law fluid with one index everywhere, and its consistency is an exponential product of temperature, moisture, cook and amylose. Real starch melts are viscoelastic, and the expansion mechanism depends on their extensional viscosity [3], which the model does not carry. The correlations are meant for about 19 to 28 % moisture and 100 to 180 °C.

Table 5.3 Parameters of the Line. A Marks a Value Shown and Adjustable in the Parameter Panel and H a Hidden One

Symbol	Meaning	Value		Symbol	Meaning	Value	
q	drag volume per revolution	35 cm^3	A	τ_{pc}	preconditioner residence	150 s	A
V_z, L_z	free volume, length per zone	1.0 L, 220 mm	H	UA_{pc}	preconditioner loss	25 W/K	H
ρ, ρ_m	dough in barrel, melt at die	1300, 1350 kg/m^3	H	S_{max}, W_{max}	steam, water valve capacity	40, 100 kg/h	A
N_{min}, N_{max}	screw speed range	150, 450 rpm	H, A	h_g	steam enthalpy, 3 bar	2725 kJ/kg	H
Γ_{rated}	rated torque, both shafts	540 N m	A	T_{fl}, T_w, T_{amb}	flour, water, room	20, 20, 25 °C	H
τ_N	drive speed lag	1.5 s	H	τ_W, τ_S, τ_F	water, steam, feeder lags	2, 3, 1 s	H
w_s	shear lost per unit wear	0.6	H	F_{max}	feeder capacity	600 kg/h	A
n	power-law index	0.40	A	σ_F	feeder pulsation, 1σ	0.6 %	A
K_{ref}	consistency at reference	7000 Pa s^n	A	k_c	cook rate at 140, 23.3	0.06/s	A
T_{ref}, x_{ref}, a_{ref}	consistency reference	150 °C, 23.3, 25 %	H	$T_{c,ref}$	cook reference temperature	140 °C	H
E_η	viscosity activation	3500 K	A	E_c, β_c	cook activation, moisture	11 000 K, 10	A
β_X, β_A	moisture, amylose exponents	12, 4	A	k_M	mechanical cook rate	10 kg/MJ	A
κ_c	consistency lost uncooked	0.6	A	SEI$_{max}$	expansion scale	4.1	A
UA_b	barrel to dough, wetted	200 W/K	A	θ_T, m_T	reference superheat, exponent	50 K, 2.0	A
$Q_{h,max}$	heater capacity per zone	6.0 kW	A	T_{col}, w_{col}	bubble collapse, width	185 °C, 6 K	A, H
UA_{cw}	cooling water, valve open	60 W/K	A	β_E, p_E	moisture, pressure exponents	10, 0.10	A
C_b	barrel capacity per zone	20 kJ/K	A	p_{ref}	expansion reference pressure	6 MPa	H
UA_{amb}, T_{cw}	barrel loss, cooling water	8 W/K, 25 °C	H	LEI$_{ref}$, λ_L	lengthwise index, exponent	1.05, 0.4	A, H
u_{c1}	feed-zone cooling valve	25 %	A	k_{pack}	packing fraction in a cup	0.60	A
n_h, R_h, L	die holes, radius, land	20, 1.4, 30 mm	A	τ_{dry}, Δt_p	belt residence, parcel	720, 10 s	A, H
e	entrance correction	4 radii	H	k_d, $T_{d,ref}$	drying rate, reference air	0.004/s, 130 °C	A, H
h, c_G	channel depth, pressure factor	5 mm, 6	H	θ_D	air scale of drying rate	40 K	A
τ_L	filled-length settling	1.0 s	H	a_E, b_E	Oswin coefficient, exponent	0.13, 0.40	A
p_{trip}	die-pressure trip	12 MPa	A	T_{in}, ΔT_a	intake air, heater rise	25 °C, 150 K	H, A
Γ_{trip}, Γ_{max}	torque trip, hard limit	100, 90 %	A	τ_a, $\dot{m}_a$	air lag, dry-air flow	45 s, 1.2 kg/s	H
p_{max}, T_{max}	die, melt hard limits	10 MPa, 170 °C	A	P_{air}, n_{knife}	air heater rating, knives	240 kW, 4	H
$T_{a,max}$	dryer air hard limit	150 °C	A	ρ_{lo}, ρ_{hi}, x_{lo}, x_{hi}	specification	120, 150 g/L, 2, 4 %	A
c_s, c_w	solids, water heat capacity	1600, 4180 J/(kg K)	fixed	λ	latent heat at 100 °C	2257 kJ/kg	fixed

Dissipation uses one characteristic shear rate per zone, proportional to the screw speed. Wear is one number that lowers the shear in zones 3 to 5 and the pressure built per filled meter, while real wear is uneven and opens leakage paths that also change the residence distribution.

Heat transfer between barrel and dough is one conductance per zone scaled by the wetted wall. The heater bands act at once, there is no conduction between zones, and the melt thermocouple reads the melt in front of the die exactly apart from its lag and noise, although a real flush probe may not [15].

The preconditioner is two mixed cells of fixed residence in which the steam condenses completely. Expansion, bulk density and cook are the illustrative proxies of Equations (5.10) to (5.12), and there

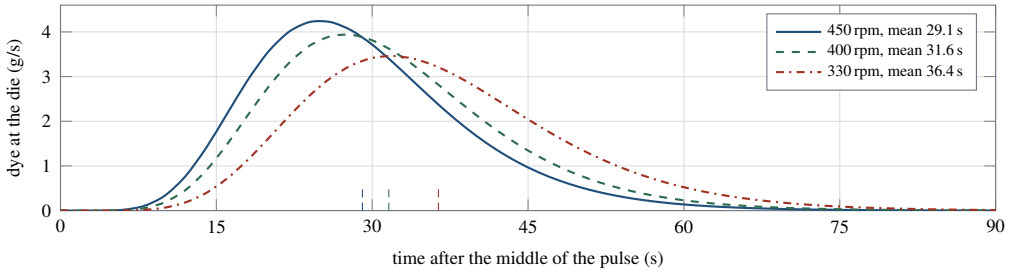

Figure 5.2 Dye leaving the die after a 0.1 kg pulse at three screw speeds, with the feed held at 300 kg/h and the noise switched off: 450 rpm solid, 400 rpm dashed and 330 rpm dash-dotted. The short vertical ticks on the axis, in the order 450, 400 and 330 rpm from left to right, mark the first moment of each curve, which equals the barrel holdup over the throughput.

Table 5.4 Steady-State Gains at the Recipe, by Central Differences on the Model's Own Steady-State Solver

Step	Γ (%)	Δp (MPa)	T_{10} (°C)	ρ_b (g/L)
process moisture, +1 point	−5.04	−0.505	−2.32	+11.48
amylose, +1 point	+2.11	+0.194	+0.75	−2.75
feed, +10 kg/h	+1.53	+0.103	−0.17	+0.59
screw speed, +10 rpm	−0.64	−0.048	+0.41	−0.97
zone 3 setpoint, +1 °C	−0.21	−0.008	+0.07	−0.22
zone 5 setpoint, +1 °C	−0.15	−0.050	+0.41	−0.91
preconditioner setpoint, +1 °C	−0.10	−0.004	+0.08	−0.47

is no protein texturization, no microbial kinetics and no lethality. The pieces keep the volume they expand to and dry on the belt with no shrinkage, case hardening or temperature history, in one air temperature and humidity. The limits and trips are teaching values for an illustrative line, not ratings of any machine.

5.4 Where the Heat Comes From

The heat budget of the recipe is the first result, averaged over 300 to 600 s of the first case. The steam brings 15.94 kW relative to water at 20 °C, and apart from the 1.62 kW the preconditioner loses, all of it goes into raising the dough from 20 to 88 °C. The shaft then delivers 14.80 kW. Of that, 0.48 kW is the pressure work of the die, and the 14.32 kW dissipated in the barrel splits into 0.49, 1.22, 6.27, 1.17 and 5.16 kW across the five zones, so the kneading block and the metering zone release 80 % of it. The heater bands deliver 5.88 kW to the barrel metal, the cooling water removes 2.06 kW and the room takes 3.88 kW. Zones 2 and 3 pass 1.92 and 1.47 kW to the dough, and zones 1, 4 and 5 take 1.14, 0.07 and 2.24 kW back, a net −0.06 kW. The whole rise from 88.0 °C to 158.0 °C is mechanical work, and the barrel only shapes the temperature profile along the way.

Screw speed does two things at once, as shown in Figure 5.2 for the dye pulse that case 1 adds at 120 s. At the fixed feed of the recipe a faster screw runs emptier, so the mean residence falls from 36.4 s at 330 rpm to 31.6 at 400 rpm and 29.1 s at 450 rpm, with the conveying cells 38.2, 31.5 and 28.0 % full. It also shears harder, so the SME rises from 38.2 to 43.0 and 46.1 kWh/t and the melt from 152.8 to 155.7 and 157.7 °C, at a torque that falls from 70.5 to 65.4 and 62.4 %. The two effects pull the cook in opposite directions, and the heating wins in this model, with the cook index rising from 0.938 to 0.945 and 0.950 while the density falls from 142.3 to 135.0 and 130.4 g/L. Measured from the middle of its 2 s pulse at 400 rpm, the dye reaches 1 % of its peak rate after 7.7 s, peaks at 27.4 s and is 90 % out by 46.2 s. Its first moment, 31.59 s, matches the 3.021 kg holdup over the 344.2 kg/h throughput, both averaged over the two minutes of its passage, to 0.001 %.

Table 5.5 Instruments. Noise Is Colored With the Correlation Time Given, and the Feeder's Is a Fraction of the Reading

Tag	Measures	Lag or delay	Noise 1σ, time
FT-101	feeder rate	2 s	0.5 %, 1 s
FT-102, FT-103	water, steam	1 s	0.3, 0.2 kg/h, 0.5 s
TT-103	preconditioner discharge	5 s	0.1 °C, 2 s
TT-Z1 to Z5	barrel metal	3 s	0.1 °C, 1 s
TT-M	melt in front of the die	2 s	0.2 °C, 1 s
PT-D	die pressure	0.3 s	0.03 MPa, 0.5 s
JT	drive torque	0.5 s	0.4 % rated, 0.5 s
TT-A	dryer air	5 s	0.2 °C, 1 s
AT-M	NIR at the belt exit	30 s + 10 s	0.05 points, 5 s
AT-D	at-line bulk density	60 s, 90 s	1.5 g/L, white

Moisture is the master variable, as the gains of Table 5.4 show. One point of process moisture moves the density by 11.48 g/L, while the metering zone's setpoint needs 12.7 °C, the screw speed 118 rpm and the kneading zone 52 °C to do the same at the local gains. The kneading zone, where most of the cook happens, moves the melt at the die by only 0.07 °C per degree of its barrel. The melt also pulls itself back toward where it was. Raising the moisture from 23.3 to 24.3 % lowers the consistency of Equation (5.4) by $1 - e^{-0.12} = 11.3\,\%$ at a fixed state, but the thinner melt dissipates less and runs 2.25 °C cooler, so the consistency at the die falls by only 7.8 %. A 32 instead of 25 % amylose flour raises the consistency law by 32.3 % at a fixed state, and the realized rise at the die is 21.1 % because the melt runs 5.58 °C hotter. Every gain in Table 5.4 is therefore a property of the heat balance, and it moves with the recipe. Across 21.3 to 25.3 % moisture the density gain runs from 10.6 to 11.8 g/L per point and the torque gain from −5.50 to −4.58 % per point.

5.5 Implemented Controller

All loops and the supervisor execute every 0.5 s, at 0.5, 1.0, 1.5 s and so on after a priming call at $t = 0$, in the order of work of Section 1.6.1, and each call receives $\Delta t_c = 0.5\,\text{s}$. The density trim executes once per at-line result, every 60 s. The speed selector plays the line at 1 to 60 times real time, 10 by default, so a 30 min case takes three minutes to watch.

5.5.1 Instruments and the Soft Sensor

No controller reads a model state directly. Each measurement is the true value filtered by an exact first-order lag, with colored noise added at the levels of Table 5.5, and only the screw speed is read exactly. The NIR analyzer reads the moisture of the parcel leaving the belt 30 s late through a 10 s lag, and the at-line check weighs a cup of pieces at the cutter every minute and reports it 90 s later, referred to the target moisture. The process moisture entering the extruder has no instrument. A soft sensor builds it from the metered flows and the certificate, as shown in Equation (5.14).

$$\hat{x}_{in} = \frac{F_m x_{f,lab} + W_m + S_m}{F_m + W_m + S_m} \tag{5.14}$$

The estimate $\hat{x}$ passes through two lags of $\tau_{pc}/2$, the preconditioner's two cells at their declared 150 s whatever the run's residence, starts at reset from its own reading of the metered flows, and is exact only while the certificate is right.

Table 5.6 Implemented Loops and Default Tuning. Action $a = +1$ Is Reverse Acting and $a = -1$ Direct Acting in the ISA Sense

Loop	PV	Output	K_c	τ_I (s)	a, limits
FIC-101	F_m	u_F	0.080	4	+, 0–100 %
FFC-102	W_m	u_W	0.50	4	+, 0–100 %
TIC-103	T_{pc}	u_S	1.0	110	+, 0–100 %
TIC-Z2..Z5	$T_{b,z}$	u_Z	8	120	+, ±100 %
TQC	Γ_m	N_{cmd}	5	15	−, N_{SP}–N_{max}
PIC	Δp_m	F_{SP}	100	150	+, 60–F_{base}
QIC	ρ_m	Δx	0.05	240	+, ±4 points
MIC	NIR	$T_{a,SP}$	10	300	−, 90–149 °C
TIC-A	$T_{a,m}$	u_A	1.5	45	+, 0–100 %
TSH-A	$T_{a,m}$	u_A	4	20	+, 0–request

Units of K_c follow the pairing, % per kg/h (FIC, FFC), % per °C (TIC), rpm per % (TQC), kg/h per MPa (PIC), points per g/L (QIC) and °C per point (MIC).

5.5.2 One Proportional-Integral Algorithm

All thirteen controllers share the framework's proportional-integral code, the velocity form of Section 1.6.2 without its derivative term, written in this chapter's symbols as Equation (5.15).

$$\begin{aligned} e_k &= a(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c e_k \\ S_k &= S_{k-1} + (P_k - P_{k-1}), \qquad c_k = S_k + I_{k-1} \\ u_k &= \mathrm{clamp}(c_k, u_{min}, u_{max}) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{\tau_I}\right]\Delta t_c \end{aligned} \tag{5.15}$$

with $a = \pm 1$ the action and a zero output bias. Here c_k is the output before limiting, v_k of Equation (1.14), and the integral is corrected by back-calculation at a tracking time of τ_I, only while τ_I is finite. The shared code's derivative and output rate limit are not used here. Manual mode and the switch back to automatic are bumpless, as Section 1.6.2 describes. Table 5.6 gives each loop's pairing, tuning and limits.

5.5.3 Regulatory Loops and the Ratio Station

FIC-101 holds the feed with the feeder screw, TIC-103 holds the preconditioner with the steam, and the four zone loops, which share one tuning, hold the barrel metal through their split-range outputs. The water setpoint comes from a ratio station that solves the process-moisture balance for the water, as shown in Equation (5.16).

$$W_{SP} = \frac{x_{SP}(F_m + S_m) - F_m x_{f,lab} - S_m}{1 - x_{SP}} \tag{5.16}$$

The result is clamped to the valve's range and slewed at 60 kg/h per minute before FFC-102 receives it. The ratio station is a feedforward from ingredient data, exact while the certificate is right and silently wrong when it is not, because Equations (5.14) and (5.16) trust the same number. In ratio mode the knife follows the flow, with $n_k = 60\,\hat{v}\,\mathrm{LEI}_{ref}/(n_{knife}\,\ell_{SP})$ and $\hat{v}$ the velocity in a clean hole of the melt flow $F_m(1 - x_{f,lab})/(1 - \hat{x})$, both LEI_{ref} and the melt density at their declared values.

5.5.4 Two Overrides by External Reset

The torque override raises the screw speed above the operator's setpoint when the torque passes its 85 % limit, because a starve-fed screw carries the same flow with less fill and less torque at a higher speed. It is a direct-acting PI whose lower output limit is the speed setpoint, so the selector

$N_{cmd} = \max(N_{SP}, \mathrm{PI})$ is realized by the clamp, the form of Section 1.4.3, and the integral that back-calculation holds on the clamp is the external reset of Section 1.4.4. Below the limit the output rests on N_{SP} and back-calculation holds the integral at N_{SP}, with the unclamped output K_c times the error below it. At the limit the error changes sign and the output rises without a bump. The die-pressure override is the mirror image, a reverse-acting PI on the die pressure whose upper output limit is the feed setpoint F_{base} from the supervisor or the operator, so $F_{SP} = \min(F_{base}, \mathrm{PI})$ above a floor of 60 kg/h. When an override is switched off, in manual or after a trip, it tracks the selected value plus its own proportional action, $F_{base} + K_c(p_{lim} - p_m)$ or $N_{SP} - K_c(\Gamma_{lim} - \Gamma_m)$, the value its back-calculation would hold, so switching it on does not engage it on noise.

5.5.5 The Density Trim and the Dryer Cascade

QIC adds a bias of up to ± 4 points to the process-moisture setpoint from the at-line density, executing Equation (5.15) once per result with $\Delta t_c = 60\,\mathrm{s}$. The setpoint is then clamped to 19 to 28 %, a limit outside the trim's own, so with the recipe moved near either end the setpoint can stop while the trim's integral runs on toward ± 4 points. The supervisor counts either limit as a saturated trim. A denser product lowers the moisture, since a drier dough is stiffer, heats itself more and puffs more. With the lot restart on, a new flour lot resets the trim to zero because the correction it learned belongs to the lot that ran out. In manual or after a trip the trim holds what it learned. The dryer is a cascade. MIC, direct acting, sets the air setpoint from the NIR reading, and TIC-A holds the air with the heater. With the setpoint fixed, MIC tracks the operator's air setpoint. The air setpoint is limited to 90 to 149 °C, 1 K under the air's hard limit. A high-limit override, TSH-A, is a PI on the measured air at 149.5 °C whose upper output limit is the heater request, so the heat is the lower of the two, in automatic and in manual and under either controller. Until the air passes its limit it tracks the request plus its own proportional action. With the heater stepped to 100 % in manual in case 8, the air peaks at 154.0 °C through the thermocouple's 5 s lag and stays at or below 149.8 °C after 120 s, against 175.0 °C without it.

5.5.6 The Throughput Supervisor

The supervisor ramps the feed setpoint toward a production target while every constraint keeps its margin, the back-off of Section 1.4.2, as shown in Figure 5.3. Chapter 4 pushes the throughput of a harvester against engine and rotor overrides through a low selector. Because the preconditioner holds 2.5 min of dough, the torque and die pressure respond to a feed change late, so the supervisor predicts where they will settle at the feed it is asking for, by Equation (5.17).

$$\hat{\Gamma} = \Gamma_m \rho^{0.72} \left(\frac{N_{fin}}{N_m} \right)^{-0.4}, \qquad \hat{p} = p_m \rho^{0.46}$$
$$\rho = \frac{F_{base}}{F_m} \frac{F_{mix}}{\hat{F}_x} \tag{5.17}$$

with $F_{mix} = F_m + W_m + S_m$ the metered flow into the preconditioner, $\hat{F}_x$ the same flow through the same two declared lags, and N_{fin} the top speed when the torque override is on. Each constraint is given a slack in units of its own margin, $[(\Gamma_{lim} - m_\Gamma) - \hat{\Gamma}]/m_\Gamma$ for torque with $m_\Gamma = 3\,\%$, the same form for pressure with 0.3 MPa, $[(T_{max} - 4) - \hat{T}_{ex}]/4$ on the exit temperature predicted from the melt thermocouple and the die pressure, $\hat{T}_{ex} = T_{10,m} + p_m/(\rho_m c_{p,0})$ with the declared ρ_m and $c_{p,0}$ taken at 23.3 % moisture, $\min(\rho_{hi} - 3 - \rho_m,\ \rho_m - \rho_{lo} - 3)/3$ on the at-line density, $(x_{hi} - 0.2 - \mathrm{NIR})/0.2$ on the final moisture in points, and -1 while the trim is saturated. A measured torque above its limit at top speed or a die pressure above its limit forces the slack to -1. With the smallest slack above 0.25 the supervisor raises the feed at up to the ramp rate (10 kg/h per minute), between -0.25 and 0.25 it holds, and below -0.25 it backs off at the ramp rate, down to 60 kg/h. Above the target it comes down at twice that rate. A leaky timer counts the time spent more than 5 kg/h short of the target while not pushing, and at 120 s the target is declared infeasible, with the binding constraint named. No limit is relaxed to reach a target.

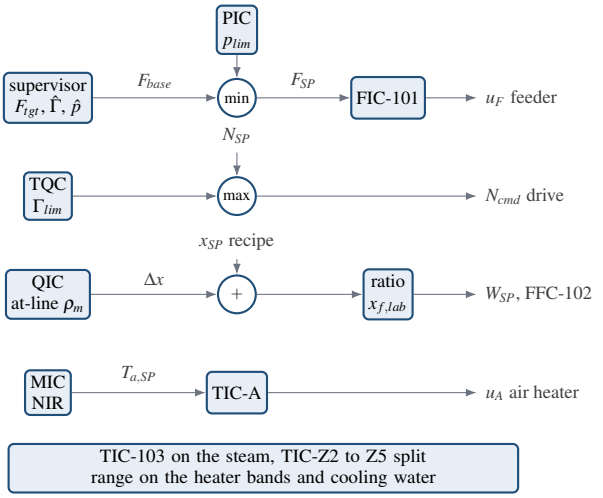

Figure 5.3 Control structure. The supervisor's feed setpoint passes a minimum selector with the die-pressure override, the speed setpoint a maximum selector with the torque override, and both selectors are realized as output limits of the overriding PI. The density trim biases the recipe moisture that the ratio station turns into a water setpoint.

5.5.7 Trips

A torque of 100 % of rating held for 1 s, or a die pressure of 12 MPa, trips the extruder. The drive, the feeder, the water and the steam go to zero, the preconditioner stops discharging and the filled length freezes, while the zone and dryer loops continue, and the run pauses 5 s later. Only the operator's trip reset restarts the line, an action separate from the Reset that starts a case again from $t = 0$.

5.5.8 Simulation, Identified and Internal Models

Three models have to be kept apart. The 105-state plant of Section 5.3 is the simulation model. Step tests on that plant give the identified control model, the gains of Table 5.4 and the dynamics behind them. A one-point step down of the moisture setpoint at 300 s of the first case, taken as the difference from an otherwise identical unstepped run, lowers the cutter density by 11.25 g/L, reaching 10 % of the change after 77 s, 63 % after 187 s and 90 % after 295 s, which a two-point fit reduces to a dead time of 77.5 s and a time constant of 109 s. The at-line check adds its 90 s report delay and its 60 s sampling to that, and the trim's loop gain is $0.05 \times 11.25 = 0.56$. The internal models of the loops and the supervisor are three. The ratio station and the soft sensor are one mass balance built on the certificate. The supervisor's are the exponents of Equation (5.17), which agree with the simulation to within 0.02 at the recipe, where the steady-state solver gives 0.70 to 0.72 for the feed exponent of torque between 285 and 437 kg/h, -0.39 to -0.40 for its speed exponent between 330 and 450 rpm, and 0.46 to 0.47 for the feed exponent of pressure. The knife assumes the design lengthwise expansion. The trim, the overrides, the zone loops and the dryer cascade carry no model. The predictive controller of Section 5.5.9 carries a fourth internal model, identified by step tests at a higher feed than the recipe.

5.5.9 A Predictive Controller on Three Setpoints

A controller option on the supervisor's card switches the line from the loops to a model predictive controller (MPC) in automatic, the option of Section 1.6.3, and the loops stay the default everywhere except case 9. The MPC moves three setpoints that the loops already own, the feed setpoint below the die-pressure override's minimum selector, the process-moisture setpoint that the ratio station of Equation (5.16) turns into water, and the screw speed setpoint below the torque override's maximum

selector. Both overrides, the trips, the zone loops, the preconditioner loop, the ratio station and the dryer cascade act beneath it unchanged, so the overrides remain the protection and the plan only respects their limits, as an optimizer above selectors does in industrial practice [14]. The production target, the ramp rate and the two margins become the plan's own bound, ramp and ceilings.

The controller's model was identified on the plant by stepping the three setpoints in automatic, ±20 kg/h, ±0.5 points and ±20 rpm in both directions with the noise off, at 400 kg/h, 23.3 % and 400 rpm, where the limits bind, rather than at the 300 kg/h recipe. The regulatory loops stay closed during the steps because the MPC moves their setpoints. The feed and the moisture each pass the preconditioner's two cells at the declared $\tau_{pc}/2 = 75\,\mathrm{s}$ and then two parallel lags, and the speed passes one lag and then two more, as Equation (5.18) states.

$$\begin{aligned} \tfrac{1}{2}\tau_{pc}\dot{\phi}_{m,1} &= \Delta u_m - \phi_{m,1}, & \tfrac{1}{2}\tau_{pc}\dot{\phi}_{m,2} &= \phi_{m,1} - \phi_{m,2} \\ \vartheta_{m,j}\dot{\phi}_{m,j} &= \phi_{m,2} - \phi_{m,j}, & \vartheta_{N,1}\dot{\nu}_1 &= \Delta N_{SP} - \nu_1 \\ \vartheta_{N,j}\dot{\nu}_j &= \nu_1 - \nu_j, & \hat{y}_i &= y_{0,i} + \textstyle\sum_m \mathbf{c}_{i,m}^{\mathsf{T}} \xi_m \end{aligned} \tag{5.18}$$

Here $m \in \{F,x\}$ in the first two lines, $j \in \{a,b\}$, Δu is a setpoint's deviation from $\mathbf{u}_0 = (400\,\mathrm{kg/h}, 23.3\,\%, 400\,\mathrm{rpm})$, $\xi_F = (\phi_{F,2}, \phi_{F,a}, \phi_{F,b})$, ξ_x likewise, $\xi_N = (\nu_1, \nu_a, \nu_b)$, and the four outputs $\hat{y}_i$ are the torque, the die pressure, the melt at the die exit and the bulk density at the cutter, with $\mathbf{y}_0 = (80.61\,\%, 7.735\,\mathrm{MPa}, 156.74\,^\circ\mathrm{C}, 140.90\,\mathrm{g/L})$. The lags, fitted by a grid, are $\vartheta_{F,a}, \vartheta_{F,b} = 45$, 90 s, $\vartheta_{x,a}, \vartheta_{x,b} = 30$, 90 s and $\vartheta_{N,1}, \vartheta_{N,a}, \vartheta_{N,b} = 3$, 12, 150 s, and the twelve coefficient vectors $\mathbf{c}_{i,m}$ by least squares, with fit residuals of 0.17 to 2.5 % of each response. The steady gains are +0.145, −6.13 and −0.079 % of torque per 1 kg/h, per point and per rpm, and +0.062, +12.6 and −0.110 g/L of density. The melt output is the melt thermocouple plus the die's pressure work $\Delta p_m/(\rho_m c_{p,0})$ with the declared ρ_m, the supervisor's own estimate. The eleven states are discretized exactly at the decision period. The model is linear at one point, and it omits the amylose share, the wear, the die deposit, the zone heaters' capacity and the dryer.

The NIR analyzer is not one of its outputs, and its 40 s of delay and lag are therefore not in the model. Within the 5 min horizon the three setpoints move the NIR reading by 7.8×10^{-5} points per 1 kg/h and 0.0026 points per point of moisture setpoint, because the product reaches the analyzer only after the 12 min belt and the dryer cascade holds it. The delayed measurement the plan does use is the at-line density, whose cup is sampled every 60 s and reported 90 s later.

A steady-state Kalman filter estimates the model's state from the torque, the die pressure and the melt at every decision, propagating it with the setpoints that were applied over the last period, the feed after the die-pressure override and the speed after the torque override. In its default form the soft sensor's error enters where it arises, with the flour, as an integrating disturbance d_x added to Δu_x in the first line of Equation (5.18), while the die pressure and the melt carry output offsets b_p and b_T and the torque carries none, which is one of the disturbance models that make a plan offset free [19, 20]. A torque change that the planned setpoints do not explain is therefore read first as moisture, and the same model turns that moisture into a predicted density change of 12.6 g/L per point. The measurement variances are the instruments' of Table 5.5, and the random walks per decision are 0.02 points for the cells and for d_x, 0.004 MPa and 0.03 °C. At a new flour lot d_x restarts from zero, because the old lot's error leaves the cells with its dough. The at-line density corrects the predicted density through a bias b_ρ, as Equation (5.19) states.

$$b_\rho \leftarrow \mathrm{clamp}\Big(b_\rho + \lambda\left[\rho_{lab} - \hat{\rho}(t_s) - b_\rho\right],\ \pm 40\,\mathrm{g/L}\Big) \tag{5.19}$$

Here $\lambda = 0.4$ and $\hat{\rho}(t_s)$ is the model's density at the time t_s the cup was taken, interpolated from the decision history, so a result that arrives 90 s late is compared with the prediction it can confirm. The correction applies on arrival and enters the next plan. Volume 5, Chapter 10 corrects the bias of a soft sensor from a sampled and delayed laboratory result as well. A second estimate, chosen as an

advanced option, replaces d_x by an output offset on the torque, the ordinary form, and serves as the ablation of Section 5.6.8.

Every 5 s, ten executions of the loops, the plan is solved over $N = 60$ steps (5 min), its three setpoints held over ten blocks that grow from one step each to 14 (1, 1, 2, 2, 4, 6, 8, 10, 12, 14), 30 decision variables, for the objective of Equation (5.20).

$$\begin{aligned} J = \sum_{k} \Big[& Q_\rho \,(\hat{\rho}_k - \rho_{SP})^2 - w_F F_{SP,k} \\ & + R_N (N_{SP,k} - N_{op})^2 \Big] + \sum_{j} \Delta \mathbf{u}_j^{\top} \mathbf{S}\, \Delta \mathbf{u}_j \\ & + \sum_{i,\ell} \left(r_i s_{i,\ell}^2 + q_i s_{i,\ell} \right) \end{aligned} \tag{5.20}$$

The feed earns $w_F = 1$ per 1 kg/h and step, the density is held on its setpoint with $Q_\rho = 5$ per $(\mathrm{g/L})^2$, the speed is drawn toward the operator's setpoint N_{op} with $R_N = 0.005$ per rpm^2, and the block moves are suppressed by $\mathbf{S} = \mathrm{diag}(2, 200, 0.05)$ per $(\mathrm{kg/h})^2$, point squared and rpm^2. The soft rows keep the torque under $\Gamma_{lim} - m_\Gamma = 82\,\%$, the die pressure under $p_{lim} - m_p = 8.7\,\mathrm{MPa}$, the exit melt under $T_{max} - 4 = 166\,^\circ\mathrm{C}$ and the density between 123 and 147 g/L, each with its slack $s_{i,\ell} \geq 0$ priced at $r = (10^3, 10^5, 10^3, 100)$ and $q = (2000, 20\,000, 2000, 200)$ per unit of %, MPa, °C and g/L. The linear price exceeds what the room is worth, which makes the penalty exact when the rows can be met [21]. A percent of torque is worth 6.9 kg/h of feed at the model's gain, and its price of 2000 is 4.8 times the reward that room would earn over all 60 steps. The torque and the die pressure have one slack for each of the first 12 steps (1 min) and one for their worst step beyond, and the melt and the density one for the whole horizon, a choice explained in Section 5.6.9. The hard bounds are the feed between the overrides' 60 kg/h floor and the production target, the moisture setpoint within the trim's ± 4 points and 19 to 28 %, and the speed between N_{min} and N_{max}. The first feed move may differ from the feed in force by the supervisor's ramp times 5 s, and each later block boundary by the ramp times the time between block centers, in both directions.

A solve is accepted only if it is optimal, its plan, predictions, cost and slacks are finite and of the right size, every planned setpoint is inside its bounds and the plan's first move is the move applied. Any other outcome hands the three setpoints back to the loops from where the plan left them, marks the historian and posts a note, and control returns to the plan at the next solve that passes. Under the MPC the supervisor's feed and the trim's integral follow the applied setpoints, so a switch to manual moves no actuator, and with the noise off the return to automatic moved the speed by 0.29 rpm, the water valve by 0.21 % and the feeder by 0.04 %, against 0.09 rpm, 0.08 % and 0.04 %, the largest change in one 0.1 s step of ordinary running from 1400 to 1500 s. The plan reads only the torque, die pressure and melt readings, the at-line results with their cup times, the setpoints applied and on the panel, the limits, and its own state. Its model constants are frozen from the declared parameters, so a parameter the learner changes on the plant becomes model mismatch. A rule written before any search and judged on case 9 only set the weights, and the final 243-setting grid was flat, 212.6 to 212.9 kg. In Chrome a decision takes 2.26 ms on average and 16.1 ms at most over the 360 decisions of 1800 s of case 9, measured under a load average of about 3, which is 0.045 % of a core at real time and 2.7 % at the top speed of 60 times real time. The decision exceeds 2 ms, so a frame guard skips a decision rather than a frame when the browser is behind.

5.6 The Cases as Measured

The cases are taken below in the order of the application's selector, from case 1, the steady recipe with a dye pulse at 120 s, to case 10, the graded challenge. Each has a completion rule of the kind Section 1.6.4 describes, and the challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt.

Table 5.7 Measured Outcomes, Default Seed

Case	Measured outcome
1, 330/400/450 rpm	residence 36.4/31.6/29.1 s, SME 38.2/43.0/46.1
2, trim off	true 25.9, estimate 23.3 %, torque 51.9 %
2, density crossings	cutter 208.8, at-line 330.0, exit 939.9 s
2, by 1800 s	63.8 kg dense, air 133.7 to 135.6 °C
2, trim on	−2.67 points, below 150 g/L at 540.3 s, 24.3 kg
2, certificate fixed	17.7 kg off-spec
3, trim off	torque 81.3 %, 118.0 g/L, 56.7 kg light
3, trim on	+1.55 points, torque 72.6 %, nothing lost
4, 380 kg/h	torque 77.6 %, 140.1 g/L, 40.3 kWh/t
4, 480 kg/h target	held at 437.2 kg/h, 409 rpm, infeasible 1657 s
5, override on	takes over 740 s, peak 9.40 MPa, 112.1 kg/h
5, override off	trips at 1344.0 s
6, zone 3	132.0 °C at 100 %, melt 155.1 °C
7, 60 % wear	torque 52.2 %, L_f 487.7 mm, 145.6 g/L
8, cascade	peak 3.56 % at 360 s, air 145.3 °C
8, air fixed	4.11 %, 89.5 kg wet

Each case begins from the recipe of Section 5.3, with both overrides on, the dryer in cascade, the trim and the supervisor off and the lot restart on, unless the case says otherwise. A setting changed for a run was changed after the case was selected and before Run. The application treats such a change exactly as the same setting carried through Reset, so each configuration has one run. The numbers in this section were measured headlessly on the engine at the default seed, in seconds of simulated time. Plant events carry the resolution of the 0.1 s step, and a controller action is timed at the execution that made it. Table 5.7 collects the main outcomes.

5.6.1 A Wrong Certificate

The second case loads a lot of 15 % moisture at 60 s whose certificate still says 12 %. The ratio station keeps adding the water the certificate implies, so the dough becomes $300 \times 0.03/344.2 = 2.61$ points wetter, 25.9 % water, while the estimate built on the same certificate reads 23.3 %. The thinner melt needs 51.9 % torque and 5.45 MPa at the die, it runs 150.3 °C instead of 155.7 °C, and it expands to an SEI of 3.07 instead of 4.53. The density at the cutter passes 150 g/L at 208.8 s and settles at 164.6 g/L. The at-line check reports the crossing at 330.0 s, after 31 s of waiting for the next cup and 90 s in the laboratory, and the first off-spec parcel leaves the dryer at 939.9 s. By then 12 min of dense product is already on the belt, as shown in Figure 5.4. The dryer's cascade copes with the extra water by raising the air from 133.7 to 135.6 °C, and the final moisture stays in its window, because a density cannot be dried away. By 1800 s 63.8 kg are off-spec, all of them dense.

> **TRY IT IN THE APP** **2 · Wetter Flour**
>
> Load the scenario and press Run with the density trim off. Watch the process-moisture trend: the estimate stays on its 23.30 % setpoint while the truth climbs to 25.93 %. The true density at the cutter passes 150 g/L at 208.8 s, and the at-line result reports the crossing only at 330.0 s.

With the trim on, the at-line results lower the moisture setpoint by 2.67 points by 1800 s, close to the 2.61-point certificate error expressed in process moisture, the cutter density is back under 150 g/L at 540.3 s, and 24.3 kg are lost. Correcting the certificate at 300 s restores the water exactly, and 17.7 kg are lost. The feedforward is exact but only as good as its data, while the feedback measures the product itself and arrives late, after the preconditioner, the barrel and the laboratory. An error of three points in a certificate is large, but the case serves to separate the two kinds of correction.

Pause. Try Exercise 5.1 before reading on.

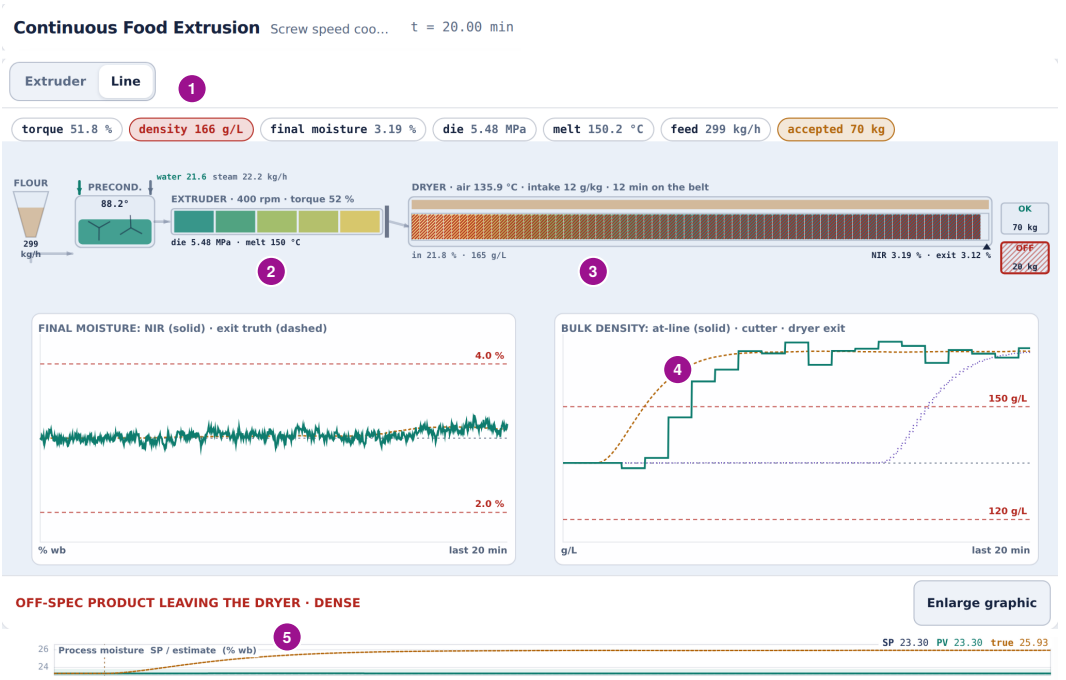

Figure 5.4 The extrusion application in case 2, wetter flour, at $t = 20.00$ min (1200 s) in the Line view, 1140 s after a 15 % lot on a 12.0 % certificate entered the feeder, run with nothing changed and a 20 min trend window. The picture is a composite: the run clock of the header above the Line view, and below it the process-moisture strip of the trends, scaled to the same width. Every parcel on the belt (3) is hatched as too dense, the OK bin holds 70 kg and the OFF bin 20 kg, and the badges read an at-line density of 166 g/L in alarm (1), a torque of 51.8 %, 5.48 MPa at the die and a melt of 150.2 °C, and the extruder bar (2) shows the screws at 400 rpm. The density history (4) shows the cutter truth rising first, the at-line staircase following it and the dryer-exit truth rising 12 min later, while the final moisture, NIR 3.19 % against an exit truth of 3.12 %, stays between its 2.0 % and 4.0 % limits with the air at 135.9 °C. The process-moisture strip (5) holds the estimate on its 23.30 % setpoint while the truth reads 25.93 %, and the status line reports off-spec product leaving the dryer as dense.

5.6.2 A Stiffer Starch

The third case switches to a flour with 32 % amylose at the same moisture at 60 s. The torque rises from 65.4 to 81.3 %, 3.7 points under its override limit, the die pressure from 6.77 to 8.20 MPa, and the SME from 43.0 to 53.5 kWh/t. The melt reaches the die at 161.3 °C instead of 155.7 °C, and the extra superheat, a fuller cook of 0.971 and the higher die pressure expand the pieces to an SEI of 5.67, and the density falls to 118.0 g/L, under the 120 g/L limit at the cutter from 330 s on. By 1800 s 56.7 kg are off-spec, all too light. The fix is the same handle as for the wetter flour, turned the other way. By 1800 s the trim adds 1.55 points of moisture, the torque comes down to 72.6 % and the density to 134.2 g/L. It acts early enough that the cutter density bottoms at 120.2 g/L at 372 s, just inside the limit, and nothing is lost.

5.6.3 Throughput Against Torque

The fourth case turns the supervisor on and raises its target from 300 to 380 kg/h at 60 s. The feed setpoint reaches 380 kg/h at 539.5 s, the torque settles at 77.6 % and the density rises to 140.1 g/L, because each kilogram now gets 40.3 kWh/t instead of 43.0. At 900 s the target becomes 480 kg/h.

The predicted torque at the top speed reaches its margin at 424.8 kg/h at 1200 s, and from there the supervisor alternates between holding and slow pushing. At 1422.5 s, with the setpoint at 433.0 kg/h, a noise excursion of the torque transmitter first touches 85 % and the override takes its first increment, 129.5 s before the true torque reaches the limit at 1552.0 s. At 1657 s the supervisor declares the target infeasible, with the torque at maximum speed as the binding limit. The line settles with the feed setpoint at 437.2 kg/h, the screws at 409 rpm and the torque at 85.0 %, the worst case over the run being 85.5 % against the 90 % hard limit, and the density at 142.4 g/L. With the trim on as well, the trim lowers the moisture 0.41 points to hold 135.4 g/L at 2400 s, and the torque then caps the feed at 416.0 kg/h. At the top speed the recipe would need 87.6 % torque for 480 kg/h. Half a point more water brings that to 84.4 % but the density to 147.2 g/L, at the supervisor's quality margin, while a drive rated at 600 instead of 540 N m would carry it at 78.9 %.

5.6.4 A Restricted Die

The fifth case grows a deposit in the die holes at 1.2 % of their radius per minute from 60 s. The die pressure passes 8.76 MPa at 699 s and the override takes over at 740 s. It holds the pressure but overshoots its 9.0 MPa limit to 9.40 MPa at 913 s, because a feed change reaches the die through the preconditioner's 2.5 min. By 1500 s the feed is down to 112.1 kg/h. The screws still turn at 400 rpm, so each kilogram takes 68.2 kWh/t, the melt reaches 162.2 °C, the dough stays 50.8 s, and the pieces over-expand to 111.7 g/L at the cutter. The override protects the die, not the product. After the clean plate at 1500 s the feed is back at 300 kg/h by 1545 s. With the override off the pressure reaches the 12 MPa trip at 1344.0 s and the extruder stops.

> **TRY IT IN THE APP** **5 · Restricted Die**
>
> Run the scenario with the Die-pressure override on and watch the die pressure and the feed. The override takes over at 740 s, overshoots to 9.40 MPa and has cut the feed to 112.1 kg/h by 1500 s, while the pieces over-expand to 111.7 g/L. Press Reset, set Die-pressure override to Off on the PIC card and run again: the extruder trips at 1344.0 s.

5.6.5 A Failed Heater and a Worn Screw

In the sixth case zone 3 keeps a quarter of its heater capacity from 60 s on. Its loop pins the output at 100 % from 161.5 s, and the barrel falls from 140 to 132.0 °C and stays there, with back-calculation keeping the integral from winding up. The dough in zone 3 runs 134.0 °C instead of 136.6 °C, which makes it stiffer, the torque rises from 65.4 to 67.0 %, and the melt at the die falls only 0.6 °C, to 155.1 °C, with the density at 136.9 g/L at 1200 s. A saturated loop is not always a loop that matters, since the kneading zone's barrel is a small part of the melt's heat. Raising zone 4 to 145 °C at 300 s, where its loop runs at 24.4 % heat, returns the density to 135.1 g/L.

Pause. Try Exercise 5.2 before reading on.

The seventh case wears the screws to 20, 40 and 60 % at 60, 240 and 420 s. By 1200 s the torque is 52.2 %, the SME 34.3 kWh/t, the melt 150.2 °C, the SEI 3.99 and the density 145.6 g/L, 4.4 g/L under the upper limit. The worn flights build less pressure per filled meter, so the filled length grows from 163.2 to 487.7 mm, reaches the kneading block at 444.2 s, and raises the residence from 31.6 to 43.7 s. The die pressure rises from 6.77 to 7.53 MPa, because the cooler melt is stiffer. The falling torque is the symptom of a machine that has lost shear, not of an easier product. Raising the speed to 450 rpm at 600 s brings back part of the work, 37.9 kWh/t, and the density to 139.5 g/L, but the filled length stays at 442.7 mm.

Pause. Try Exercise 5.3 before reading on.

5.6.6 Humid Dryer Air

The eighth case turns the dryer intake to 50 g/kg at 60 s. Every parcel on the belt dries more slowly and toward a higher equilibrium at once, so the final moisture starts to rise within a minute. The

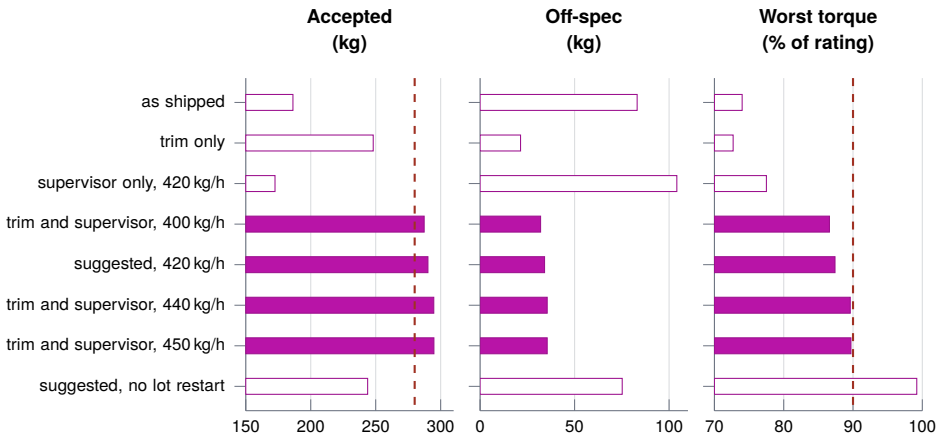

Figure 5.5 Table 5.8 redrawn against the challenge's requirement of at least 280 kg accepted and its 90 % hard limit on torque, dashed. Solid bars pass and white bars fail. Notice that the supervisor alone accepts less than the line as shipped, and that without the lot restart the suggested configuration runs its torque to 99.2 %.

Table 5.8 Challenge Configurations Over 3600 s. Mass in kg, Worst Torque in Percent of Rating

Configuration	Accepted	Off-spec	Torque	Result
as shipped	186.3	83.1	74.0	fail
trim only	248.1	21.4	72.7	fail
supervisor only, 420 kg/h	172.5	104.1	77.5	fail
trim and supervisor, 400 kg/h	287.3	32.0	86.6	pass
suggested, 420 kg/h[a]	290.2	34.1	87.4	pass
trim and supervisor, 440 kg/h	294.7	35.5	89.6	pass
trim and supervisor, 450 kg/h	294.8	35.5	89.7	pass
suggested, no lot restart	243.8	75.2	99.2	fail

[a] trim on with the lot restart, supervisor on at 420 kg/h

cascade raises the air, the final moisture peaks at 3.56 % at 360 s, is back to 3.05 % at 1200 s and stays within 3.04 to 3.06 %, and the air reaches 145.3 °C by 1800 s, 4.7 °C under its limit, with no wet product. With the air fixed at the recipe's 134 °C, the run starting at that air temperature, the final moisture passes 4.0 % at 640 s and settles at 4.11 %, and 89.5 kg are off-spec by 1800 s.

5.6.7 The Shift of Flour Lots

The challenge runs 3600 s through lot B at 600 s, with 14.5 % moisture on a 12.0 % certificate, lot C at 1800 s, a 29 % amylose flour at 12.5 % moisture on a correct certificate, and 50 g/kg dryer air at 2700 s. It is met with at least 280 kg of accepted product, the torque never above 90 %, the die never above 10 MPa, the melt never above 170 °C, the dryer air never above 150 °C and no trip. The measured configurations are listed in Table 5.8, and Figure 5.5 draws them against the requirement.

As shipped the line accepts 186.3 kg, 93.7 kg short of the requirement, and all 83.1 kg of its off-spec product is lot B's dense product, with every hard limit respected. The trim alone accepts 248.1 kg. The supervisor alone accepts 172.5 kg, less than doing nothing. It pushes toward 420 kg/h just as lot B arrives and makes 104.1 kg of dense product, and in the last part of the hour its quality condition backs the feed from 249.8 kg/h at 2250 s to 79.8 kg/h at 3600 s while the at-line density sits between 117.9 and 124.7 g/L, although less feed puts more work into each kilogram and makes the pieces lighter still. The suggested configuration, the trim on with its lot restart and the supervisor

Table 5.9 Case 9 and the Challenge at 420 kg/h Under Five Structures, Default Seed

Structure	Case 9, wet lot, 2400 s					Challenge, 3600 s				
	Acc. (kg)	Off (kg)	Feed (kg/h)	Γ_{max} (%)	ρ_{max} (g/L)	Acc. (kg)	Off (kg)	Feed (kg/h)	Γ_{max} (%)	$t_{>85}$ (min)
Loops, supervisor and trim	169.4	36.8	389.3	84.7	171.2	290.2	34.1	383.9	87.4	12.5
Loops retuned once	182.3	26.0	389.6	78.6	166.4	302.0	23.1	388.0	83.5	0.0
Loops with the MPC's estimate[a]	186.0	20.9	388.2	84.9	161.5	310.8	20.7	396.1	84.4	0.0
MPC, bias estimate	177.6	35.4	401.7	82.0	156.1	308.4	18.3	384.3	87.3	4.9
MPC, moisture estimate	212.9	0.0	401.6	82.3	147.8	336.9	0.0	401.5	84.3	0.0

Accepted and off-spec mass at the dryer exit, mean feed, worst true torque, peak true density at the cutter and minutes with the true torque above the 85 % override limit. Every row respects every hard limit and none trips. Each MPC run makes 480 or 720 decisions with no fallback.
[a] At 400 rpm on case 9 and 450 rpm on the challenge, the better of the two speed setpoints measured for each (172.5 and 34.0 kg at 450 rpm on case 9, 300.7 and 27.8 kg at 400 rpm on the challenge).

aimed at 420 kg/h, accepts 290.2 kg, 10.2 kg above the requirement, with 34.1 kg off-spec, a worst torque of 87.4 %, the die at 8.36 MPa, the melt at 162.9 °C and the air at 147.2 °C. Its trim reaches −2.78 points for lot B by 1800 s. Without the lot restart that correction carries into lot C, whose stiffer starch needs more water rather than less, and the torque reaches 99.2 %, less than a point from the coupling's trip, and 44.1 kg come out too light. A target of 440 kg/h passes at a worst torque of 89.6 %, and every target from 445 kg/h up gives one and the same run, 294.8 kg at 89.7 %, because the supervisor's margins then set the rate and hold the feed setpoint at or below 440.9 kg/h. Over six seeds the suggested configuration accepts 290.2 to 293.5 kg, the 450 kg/h target reaches a worst torque of 89.1 to 89.7 %, and the supervisor alone accepts 170.4 to 177.0 kg, always less than the 186.3 kg of the shipped line.

> ✦ **DID YOU KNOW? Pushing for Throughput, Accepting Less**
>
> From the simulator: through the hour of flour lots, the throughput supervisor alone accepts 172.5 kg of product, less than the 186.3 kg of the line as shipped. With the density trim and its lot restart added, the same supervisor aimed at 420 kg/h accepts 290.2 kg (Table 5.8).

5.6.8 Predictive Control Through a Wet Lot

Case 9 asks for 420 kg/h from the recipe's 300 and at 600 s refills the feeder from a lot of 15 % moisture whose certificate says 12 %, the error of case 2 at full rate. It runs 2400 s under the MPC, and the comparison runs it again on the same seed under the loops with the supervisor and the trim on, as the case pins them. Table 5.9 adds two stronger classical rows. The first is the loops retuned once by the MPC's own rule over every setting they have, 864 combinations of the ramp, both margins, the trim's gain and integral time and the speed setpoint, which kept the ramp and the torque margin, narrowed the die margin to 0.1 MPa, doubled the trim's gain and raised the speed setpoint to 450 rpm. The second gives the loops the MPC's own moisture estimate, the same filter on the same instruments and applied setpoints, with d_x subtracted from the moisture setpoint ahead of the trim and restarted at each lot. No signal the MPC uses is withheld from the loops.

The wet dough reaches the screws through the preconditioner, and the torque falls at the same feed and speed, to 71.2 % under the MPC and 68.5 % under the loops. The MPC's moisture estimate passes one point at 720 s and reaches 3.09 points by the end. By 755 s the plan has cut the moisture setpoint from 22.8 to 19.9 %, it runs the screws to their 450 rpm top and holds the feed near 420 kg/h, and the density at the cutter peaks at 147.8 g/L at 771 s, inside the 150 g/L limit, as shown in Figure 5.6. The first at-line result to show the wet lot arrives at 810 s under either structure, and under the MPC it only confirms the estimate. The loops read the same torque fall as room, the supervisor pushes the feed into it, and the trim waits for the at-line density, so the cutter density passes 150 g/L at 712 s and peaks at 171.2 g/L at 876 s, and 36.8 kg come out dense.

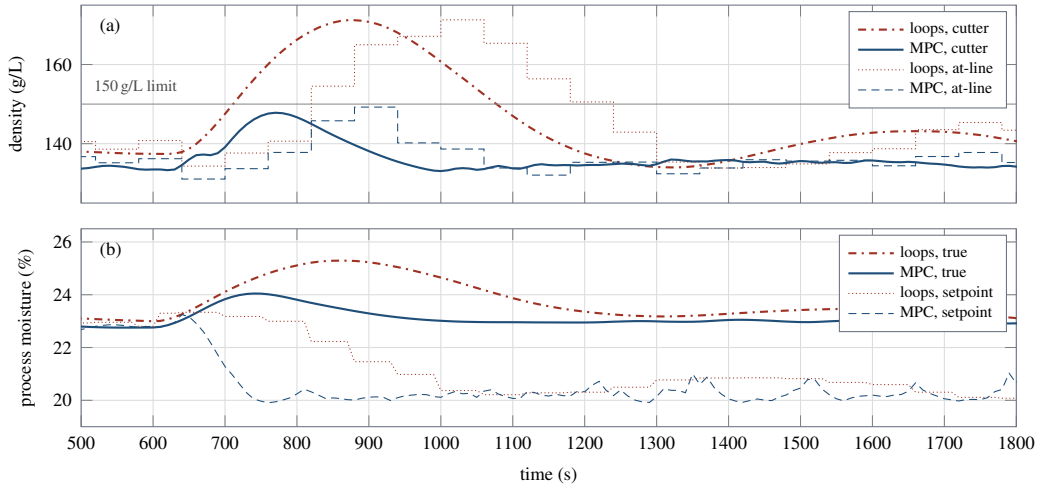

Figure 5.6 Case 9 from 500 to 1800 s, the MPC against the loops with the supervisor and the trim on. (a) True bulk density at the cutter and the at-line result, which lags it by the cup's sampling and the 90 s report. (b) True process moisture and its setpoint. The MPC's lines are thick solid (truth) and dashed (at-line result or setpoint), the loops' dash-dotted and dotted, and the thin gray line in (a) is the 150 g/L limit. The MPC lowers the setpoint within a minute of the wet dough's arrival, before any at-line result shows it.

The ablation in Table 5.9 locates the gain. With the bias estimate, the same plan, horizon and weights make 35.4 kg of dense product against the loops' 36.8 kg, because an offset on the torque explains the symptom and leaves the predicted density unchanged, so on the lot change the product gain is the estimator, not the optimizer. Given the same estimate, the loops reject 20.9 kg, 43 % less than as shipped, and the remaining 20.9 kg of off-spec product and 13.4 kg/h of mean feed belong to moving the three handles together against predicted rather than measured limits. The claim holds for the lot change only. On the challenge, whose third change is a stiffer starch with a correct certificate, the bias-estimate MPC also beats the shipped loops, 308.4 against 290.2 kg. The MPC with the moisture estimate accepts 336.9 kg on the challenge with nothing off-spec and no minute above the torque override's limit, 26.1 kg more than the strongest classical row. With the shipped 300 kg/h target it accepts 269.5 kg, also with nothing off-spec, against the shipped line's 186.3 kg.

5.6.9 One Slack for the Whole Horizon

The first build of the controller judged the torque with one slack for the whole horizon and crossed the 90 % hard limit. With the torque override off and a 32 % amylose lot at 600 s and 420 kg/h, the stiffer starch raises the torque, the estimator reads it as a drier dough (d_x reaches -2.73 points by 900 s), and an excess over the 82 % ceiling within the next minute is already unavoidable. Paid once for the horizon, the excess left the rest of the plan free under it, and the plan lowered the screw speed from 410 to 305 rpm to hold the density while the torque rode at 89.9 to 90.9 % from 780 to 960 s. The worst torque was 91.1 %, 3.1 min above the hard limit, with 28.6 kg of dense product. Judging the torque and the die pressure at every step of the first minute, where an excess may already be on its way through the preconditioner, and once beyond, removes the breach. On every row of Table 5.11 but the two die-deposit trips, where the dense product differs by 2.1 kg, it gives the same result as a slack at all 60 steps, at 1.8 to 2.8 ms a decision in Node against 16 to 21 ms. The same run then keeps the speed between 394 and 450 rpm and the torque at or below 85.3 %, as Table 5.10 shows.

The change cost one row. With die deposits growing at 2 % of the radius per minute at 420 kg/h, both versions trip at 2329 s, and the dense product before the trip grows from 26.5 to 33.8 kg. Case 9

Table 5.10 The MPC With One Torque Slack for the Horizon and With a Slack per Step Over the First Minute. Accepted/Off-Spec kg, Worst Torque %

Case	One slack		First minute	
Amylose at 420, torque override off	178.6/28.6	91.1	210.2/0.0	85.3
Challenge 420, torque override off	325.9/9.1	86.9	336.9/0.0	84.3
Challenge 420	336.1/0.0	85.7	336.9/0.0	84.3
Case 3, amylose at 300	131.1/4.9	84.5	136.0/0.0	85.1
Amylose at 420	210.6/0.0	87.3	210.2/0.0	85.3
Drier lot at 420	215.2/0.0	85.9	216.1/0.0	84.3
Die deposits 2 %/min, trip	168.8/26.5	82.9	162.3/33.8	82.4

moved only in its mean feed, from 401.3 to 401.6 kg/h, and its tuning rule gave the same weights. Judging the rows only at every second or third step degrades the plan, with 22.3 kg of dense product on case 9 at every second step and a trip with the torque override off at every third.

Two other parts of the design were changed after the worst cases had been run, and both are stated here with the results without them. The speed setpoint was first unpriced once the feed reached its target, and the QP's small ridge, which acts on the absolute inputs, drifted it toward the torque ceiling. Without the pull toward N_{op} case 3 reaches a worst torque of 89.2 % and case 5 a die pressure of 9.02 MPa. The die pressure was first judged with one slack for the horizon too, and without a slack per step case 5 peaks at 9.58 MPa and spends 8.9 min over its 9.0 MPa limit, against 8.94 MPa and no minute with it. Without both changes, and with the move suppression of the first grid, case 9 still accepts 212.9 kg with nothing off-spec and the challenge 337.9 kg, so the headline does not depend on either change.

5.6.10 Worst Cases Under Every Structure

Table 5.11 runs every case and every disturbance the application offers under the loops as the case ships them, the loops with the supervisor and the trim on, and the MPC with either estimate, with every limit checked. The melt never exceeds 165.3 °C against its 170 °C limit, the dryer air never exceeds 147.4 °C against 150 °C, and only case 7 backs the filled length into the kneading block, for 12.6, 11.8 and 12.0 min under the loops, the MPC and its bias form. After the change of Section 5.6.9, no MPC run spends any time above the 90 % torque limit or the 10 MPa die limit outside the die-deposit trips, with either estimate. The loops cross the torque limit with the override on in the amylose lot at 420 kg/h, at 90.8 % for 1.9 min, and reach 93.3 % with the override off, and the supervisor and trim reach 90.9 % for 2.2 min on the challenge with the torque override off.

The MPC is worse in one row. With die deposits at 2 % per minute every structure trips, because the overrides' and the plan's feed sits at its 60 kg/h floor for more than five minutes before the pressure reaches 12 MPa, from 1980 s under the MPC and 1860 s under the loops, so no handle can prevent the trip. Before it, the moisture estimate reads the deposits as a drier dough, d_x reaches −5.49 points and the moisture setpoint sits at its 27.3 % ceiling, so the MPC makes 33.8 kg of dense product where the loops make 20.3 kg of light product and the bias form none. The same misreading appears without a loss in case 3, where the stiffer starch is taken for a dough 2.83 points drier and the cutter density climbs to 149.3 g/L against the loops' 135.1 g/L. Amylose, dryness and deposits move torque, die pressure and melt in the same directions, and only the delayed density tells them apart.

The model was identified at 400 kg/h. Stepped again on the plant, its steady gains agree within 2.1 % at 400 kg/h, while at 300 kg/h the plant's differ from them by 0.4 to 18.6 % of the model's value. The model's torque per point of moisture, −6.13 %, is then 21.7 % above the plant's −5.04 %, the value of Table 5.4, and its torque per rpm 22.9 % above. Run on case 9 with targets of 150, 200, 300, 360, 420, 480 and 540 kg/h, the MPC makes no off-spec product and spends no time above an override limit at any of them, while the loops make 8.1 to 38.3 kg off-spec. With the recipe's moisture

Table 5.11 Worst Cases Under Every Structure. Accepted/Off-Spec kg (Worst Torque %, Worst Die Pressure MPa). A Dagger Marks a Die-Pressure Trip

Case	Loops as shipped	Loops, supervisor and trim	MPC, moisture	MPC, bias
2, wetter flour, 300	70.3/63.8 (65.4, 6.77)	106.8/23.4 (65.4, 6.77)	133.9/0.0 (66.7, 6.93)	114.3/19.8 (65.6, 6.79)
3, amylose, 300	79.4/56.7 (81.5, 8.20)	134.8/0.0 (79.0, 7.97)	136.0/0.0 (85.1, 8.29)	133.7/0.0 (85.6, 8.39)
4, targets 380 and 480	211.5/0.0 (85.5, 8.08)	210.1/0.0 (85.6, 8.14)	211.8/0.0 (82.3, 7.96)	211.5/0.0 (82.3, 7.95)
5, die deposits	129.1/5.0 (70.8, 9.40)	135.0/0.0 (68.8, 9.33)	135.4/0.0 (69.6, 8.94)	133.7/0.0 (74.1, 9.04)
5, die override off	101.2/0.0† (74.7, 12.00)	135.2/0.0 (68.8, 9.66)	135.4/0.0 (69.6, 8.94)	133.7/0.0 (74.1, 9.04)
Challenge, 420	186.3/83.1 (74.0, 7.56)	290.2/34.1 (87.4, 8.36)	336.9/0.0 (84.3, 8.06)	308.4/18.3 (87.3, 8.33)
Challenge, 420, torque override off	186.3/83.1 (74.0, 7.56)	285.5/34.5 (90.9, 8.60)	336.9/0.0 (84.3, 8.06)	308.3/18.3 (87.3, 8.33)
Case 9 lots at 420 kg/h: wet, 15 %	—	169.4/36.8 (84.7, 8.08)	212.9/0.0 (82.3, 7.91)	177.6/35.4 (82.0, 7.87)
amylose, 32 %	—	182.4/20.5 (90.8, 8.67)	210.2/0.0 (85.3, 8.16)	202.5/0.0 (85.5, 8.19)
amylose, torque override off	—	198.9/0.0 (93.3, 8.87)	210.2/0.0 (85.3, 8.16)	202.5/0.0 (85.5, 8.19)
drier, 10 %	—	185.1/21.0 (87.0, 8.32)	216.1/0.0 (84.3, 8.09)	208.6/0.0 (84.4, 8.12)
die deposits 2 %/min	—	153.6/20.3† (85.5, 12.00)	162.3/33.8† (82.4, 12.00)	190.2/0.0† (83.0, 12.00)
same, die override off	—	103.8/0.0† (85.5, 12.00)	144.8/0.0† (82.4, 12.00)	151.0/0.0† (83.0, 12.00)
humid air	—	215.2/0.0 (85.6, 8.14)	216.4/0.0 (82.5, 7.95)	216.3/0.0 (82.5, 7.96)

Cases 1 to 8 over their own duration, the challenge over 3600 s, the lots of case 9 changed at 600 s and run 2400 s. Case 9 ships with the supervisor and the trim on. The challenge's "as shipped" column is its 300 kg/h default, and its other columns aim the supervisor or the MPC at 420 kg/h. A lot at 480 kg/h gives the same rows as at 420. Cases 6, 7 and 8 are ties, 90.7 to 90.8, 89.8 to 89.9 and 136.4 kg with nothing off-spec under every structure, and case 9 with both overrides off or a 3 % feeder noise gives the MPC's rows unchanged and the loops 168.3/37.3 and 169.4/36.8.

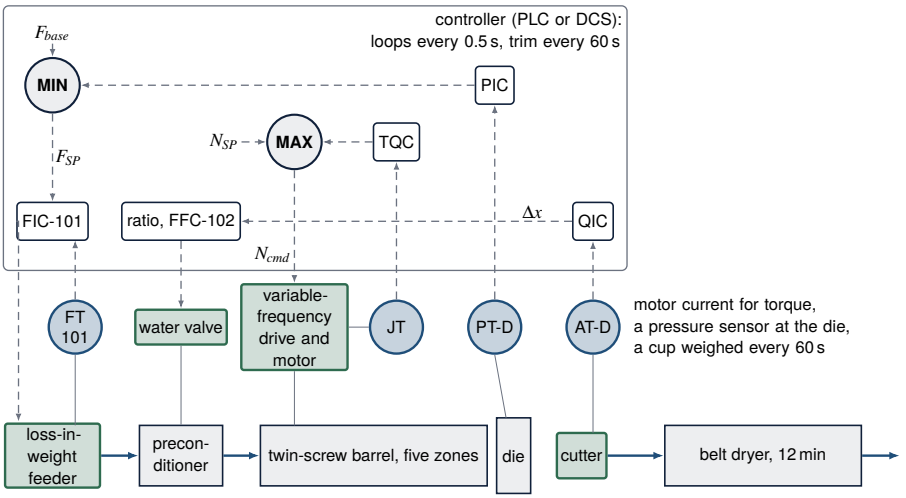

Figure 5.7 The overrides and the density trim on real hardware. The torque JT is read from the main drive's motor current and feeds TQC, whose output meets the speed setpoint in the maximum selector ahead of the variable-frequency drive. The die pressure PT-D feeds PIC, whose output meets the supervisor's feed in the minimum selector ahead of the feeder loop FIC-101. The at-line density AT-D trims the moisture that the ratio station turns into water. Notice that both selectors sit in the controller beside the loops they limit.

set to 21 or 26 %, outside the identified neighborhood, both structures make off-spec product, the MPC 57.9 to 91.2 kg and the loops 81.4 to 107.8 kg.

5.7 Technology in Practice

A real extrusion line reads the same signals and moves the same handles as the overrides of Section 5.5. Figure 5.7 draws the two overrides and the density trim with line hardware in place of the schematic's symbols.

What Measures It

An extension overview counts the residence time, the back pressure, the motor torque and the specific mechanical energy among the system parameters of an extruder: they can be measured, but the operator can change them only through the raw material, the feed rate, the barrel temperatures and the screw speed [5]. Those are the two constraints of this chapter, the torque on the drive and the pressure in front of the die. The melt temperature is harder. Accurate monitoring of the melt temperature remains one of the main challenges of hot-melt extrusion, and wall-mounted thermocouples in a die captured its variations poorly in a polymer study [22]. The model's TT-M reads the bulk melt exactly apart from a 2 s lag and its noise, which Assumptions already flags.

> **TECHNOLOGY Sensor: Drive Torque From the Motor**
>
> Extruders carry standard monitoring of barrel and die temperatures, throughput, screw speed, torque and drive amperage, melt pressure and melt temperature. Torque needs no change to the barrel or the die, and it follows properties that cannot be measured directly during extrusion, such as the rheology of the melt [22]. It is usually read from the drive, because the current of a motor is usually proportional to its torque [23]. In a polymer extruder the screw load torque inferred from the motor current was dominated by the friction of the solids and did not follow melting fluctuations well [22]. No open source read gives a range or an accuracy. The model's JT is a 0.5 s lag with noise of 0.4 % of rating.

What Moves It

Three handles carry the chapter's structure. The feeder's screw meters the flour, the water valve takes its setpoint from the ratio station, and the main drive turns the screws. Feeders deliver gravimetrically or volumetrically, and more than one may feed one extruder [5]. The torque override acts on the main drive.

> **TECHNOLOGY Final Element: A Variable-Frequency Screw Drive**
>
> The motor of an extruder provides the energy that turns the screws [5]. An alternating-current motor is run by a variable-frequency drive, which sets the speed through the frequency of the power it feeds the motor and the torque through the voltage and the current. Such drives offer speed, current and torque limiting, limits on acceleration and deceleration and shutdown on overcurrent, and a host controller can update these settings over a network while the motor runs. The drive must first be programmed with the motor's nameplate data, or it may put out more current than the motor is rated for [23]. The model's drive follows its speed command through a 1.5 s lag.

Where the Controller Runs

Programmable logic controllers serve as local controllers within a supervisory scheme, loops are often nested so that the setpoint of one comes from another, and cycle times range from milliseconds to minutes [24]. In the application the loops, both overrides and the supervisor execute every 0.5 s, the density trim once per at-line result every 60 s, and the predictive controller every 5 s. The selectors are standard function blocks. In one fieldbus standard the control selector block picks among controller outputs and carries back-calculation facilities that modify the response of the deselected controllers, which makes it the block for limit and override strategies [23]. The clamps of Section 5.5, with back-calculation at a tracking time of τ_I, do the same work inside one PI.

What Is Modernizing It

Two developments bear on the overrides and the trim. The first is in-line measurement. Spectroscopic techniques have been explored for monitoring and control of hot-melt extrusion, because control on the conventional temperatures, pressures and torque is in many cases not sufficient to predict the quality of the extrudate [22]. The second is the data-driven model. A review of automated food processing names extrusion as a process that has challenged food technologists for decades,

with operators struggling to relate product quality to settings such as the barrel temperature, and describes machine learning as a way to capture those relations from data [25]. A review of expanded snacks proposes a staged digital twin, from curated data through state variables and inline sensors to optimization or model predictive control that adjusts the inputs toward the product target [26], the role the predictive controller of Section 5.5.9 plays on the simulated line.

> ⚠ **IN PRACTICE** **Limits a Food Safety Plan Writes Down**
>
> United States rules for human food require preventive controls. Process controls ensure the control of parameters during operations such as heat processing, and they name the parameters associated with the hazard and the maximum or minimum value to which each must be controlled (21 CFR 117.135). Verification includes calibration of the process monitoring instruments, or checks of their accuracy, and a review of the monitoring and corrective action records within 7 working days after they are made (21 CFR 117.165) [27]. Such limits come from the hazard, not from the machine. The torque and die-pressure limits of this chapter protect the drive and the die, and its model carries no lethality.

5.8 Limitations

The correlations for cook, expansion and bulk density are calibrated teaching proxies. They give the reported directions but have not been fitted to any flour or extruder, and there is no protein texturization. The thermal treatment is not a food-safety lethality model, and no microbial kinetics are modeled, while a real kill step is validated by experiment [16]. Volume 8, Chapter 13 follows the thermal history of every parcel of a certified heat treatment.

The barrel is ten mixed cells with algebraic fills, so a pulse reaches the die without the delay and with wider tails than a real screw [4], and the melt is a power-law fluid with one index. The instruments are lags with colored noise and fixed delays, the melt thermocouple reads the bulk melt, and there are no failed-instrument modes beyond the wrong certificate. The feeder never refills, the piece length is displayed but not measured, and a trip restarts only through the trip reset, with no purge or die-change procedure.

The supervisor's quality condition is two-sided, so it backs off on a light product, which lowering the feed makes lighter. The application keeps this behavior by choice and records it with its defects. The melt condition guards the exit temperature on which the hard limit is scored, 2.3 °C above the thermocouple at the recipe and about 3 °C above it at 9 MPa, through a prediction that uses the heat capacity of the recipe's moisture whatever the moisture is.

The predictive controller's model is linear at 400 kg/h, 23.3 % and 400 rpm, and at the 300 kg/h recipe its torque gain per point of moisture is 21.7 % above the plant's. It carries no amylose, wear, die deposit or dryer, so every torque change it did not cause is first read as moisture, which is wrong for a stiffer starch (149.3 g/L at the cutter in case 3) and for die deposits (33.8 kg of dense product before the trip). The final moisture is left to the dryer cascade, and no lot change is announced to the plan in advance.

OPEN CHALLENGE **Reading the Product While It Is Made**

A 2026 review of models that link extrusion conditions to the quality of expanded snacks finds quality measurement not standardized: expansion, density, texture and crispness are measured by different protocols. Melt rheology and the behavior of water under real extrusion moisture, temperature and shear remain poorly quantified, and die-exit expansion models rarely bring die geometry, pressure drop and bubble growth into one validated framework. Its proposed path ends with inline sensors and then model predictive control that closes the loop [26]. In this chapter the density still comes from a cup reported 90 s late. Progress would be a model validated across extruders and ingredient lots, fed by a sensor that reads the pieces as they leave the die.

5.9 Conclusion

A cooking extruder is a plant whose heat source is the viscosity of what it carries, and the simulator of this chapter keeps that property exact by making the shaft power and the heat released in the melt one quantity. At the recipe the screws deliver 14.8 kW while the barrel exchanges a net −0.06 kW with the dough, one point of moisture moves the density by 11.48 g/L, and the kneading zone's barrel moves the melt at the die by 0.07 °C per degree. A certificate three points wrong keeps the moisture estimate on 23.3 % while the density climbs to 164.6 g/L, and a trim on the at-line density cuts the loss from 63.8 to 24.3 kg. The torque override and the supervisor hold a 480 kg/h request at 437.2 kg/h and say why, and the die-pressure override saves the die by cutting the feed to 112.1 kg/h and over-expanding the product. An optional predictive controller that moves the feed, the moisture setpoint and the speed together accepts 212.9 kg through a wet lot with nothing off-spec, and its bias-estimate form, at 177.6 and 35.4 kg against the loops' 169.4 and 36.8 kg, shows that the gain lies in where its estimator places the unknown. The hour of flour lots goes from 186.3 to 290.2 kg of accepted product with a trim that restarts at each lot and a supervisor aimed at 420 kg/h. The proxies for cook and expansion are illustrative, and the thermal treatment is not a lethality model. A fuller model would add a plug-flow component to the residence distribution, an extensional viscosity for the expansion, feeder refills, a sign-aware quality condition in the supervisor and microbial kinetics for the thermal treatment. What the application offers is a browser plant, open to anyone, on which the loops of a self-heating process can all be right while its product is wrong.

KEY POINTS

- The work of the screws, not the barrel heaters, heats the dough, so a wetter or stiffer flour moves the torque, the die pressure and the product together.
- The torque and die-pressure overrides realize a maximum and a minimum selector as output limits of the overriding PI.
- A certificate three points wrong keeps the moisture estimate on 23.3 % while the density climbs to 164.6 g/L, and a trim on the at-line density cuts the loss from 63.8 to 24.3 kg.
- The die-pressure override saves the die by cutting the feed to 112.1 kg/h and over-expanding the product: it protects the die, not the product.

Where else this applies. Chapter 4 pushes a harvester's throughput against engine and rotor overrides, and Volume 5, Chapter 10 corrects a soft sensor from a delayed laboratory result.

References

[1] J. M. Harper and J. P. Clark, "Food extrusion," *CRC Critical Reviews in Food Science and Nutrition*, vol. 11, no. 2, pp. 155–215, 1979.

[2] S. Singh, S. Gamlath, and L. Wakeling, "Nutritional aspects of food extrusion: a review," *International Journal of Food Science & Technology*, vol. 42, no. 8, pp. 916–929, 2007.

[3] C. I. Moraru and J. L. Kokini, "Nucleation and expansion during extrusion and microwave heating of cereal foods," *Comprehensive Reviews in Food Science and Food Safety*, vol. 2, no. 4, pp. 147–165, 2003.

[4] R. E. Altomare and P. Ghossi, “An analysis of residence time distribution patterns in a twin screw cooking extruder,” *Biotechnology Progress*, vol. 2, no. 3, pp. 157–163, 1986.

[5] B.-J. Gu, R. J. Kowalski, and G. M. Ganjyal, “Food extrusion processing: An overview,” Washington State University Extension, Pullman, WA, Fact Sheet FS264E, 2017.

[6] B. Vergnes and J. P. Villemaire, “Rheological behaviour of low moisture molten maize starch,” *Rheologica Acta*, vol. 26, no. 6, pp. 570–576, 1987.

[7] L. Alvarez-Martinez, K. P. Kondury, and J. M. Harper, “A general model for expansion of extruded products,” *Journal of Food Science*, vol. 53, no. 2, pp. 609–615, 1988.

[8] C. Abeykoon, P. Pérez, and A. L. Kelly, “The effect of materials’ rheology on process energy consumption and melt thermal quality in polymer extrusion,” *Polymer Engineering & Science*, vol. 60, no. 6, pp. 1244–1265, 2020.

[9] L. Wang, C. Chessari, and E. Karpiel, “Inferential control of product quality attributes—application to food cooking extrusion process,” *Journal of Process Control*, vol. 11, no. 6, pp. 621–636, 2001.

[10] B. Singh and S. J. Mulvaney, “Modeling and process control of twin-screw cooking food extruders,” *Journal of Food Engineering*, vol. 23, no. 4, pp. 403–428, 1994.

[11] T. A. Haley and S. J. Mulvaney, “On-line system identification and control design of an extrusion cooking process: Part II. Model predictive and inferential control design,” *Food Control*, vol. 11, no. 2, pp. 121–129, 2000.

[12] C. J. Chessari and J. N. Sellahewa, “Effective process control,” in *Extrusion Cooking*, R. Guy, Ed. Woodhead Publishing, 2001, pp. 83–107.

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

[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. V. Karwe and S. Godavarti, “Accurate measurement of extrudate temperature and heat loss on a twin-screw extruder,” *Journal of Food Science*, vol. 62, no. 2, pp. 367–372, 1997.

[16] A. Bianchini, J. Stratton, S. Weier, T. Hartter, B. Plattner, G. Rokey, G. Hertzel, L. Gompa, B. Martinez, and K. M. Eskridge, “Validation of extrusion as a killing step for *Enterococcus faecium* in a balanced carbohydrate-protein meal by using a response surface design,” *Journal of Food Protection*, vol. 75, no. 9, pp. 1646–1653, 2012.

[17] E. B. Bagley, “End corrections in the capillary flow of polyethylene,” *Journal of Applied Physics*, vol. 28, no. 5, pp. 624–627, 1957.

[18] C. R. Oswin, “The kinetics of package life. III. The isotherm,” *Journal of the Society of Chemical Industry*, vol. 65, no. 12, pp. 419–421, 1946.

[19] 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.

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

[21] P. O. M. Scokaert and J. B. Rawlings, “Feasibility issues in linear model predictive control,” *AIChE Journal*, vol. 45, no. 8, pp. 1649–1659, 1999.

[22] L. Saerens, C. Vervaet, J. P. Remon, and T. De Beer, “Process monitoring and visualization solutions for hot-melt extrusion: a review,” *Journal of Pharmacy and Pharmacology*, vol. 66, no. 2, pp. 180–203, 2014.

[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. Zhang, R. M. Boom, and Y. Ma, “Machine learning in automated food processing: A mini review,” *Annual Review of Food Science and Technology*, vol. 16, no. 1, pp. 25–37, 2025.

[26] D. Ying, “Modelling relationships between extrusion conditions and quality attributes of expanded snacks,” *Foods*, vol. 15, no. 12, p. 2118, 2026.

[27] “Title 21, Part 117: Current good manufacturing practice, hazard analysis, and risk-based preventive controls for human food, Sections 117.135 and 117.165,” Code of Federal Regulations, 21 CFR Ch. I (4–1–24 Edition), govinfo.gov, 2024.

5.10 Exercises

5.1 Concept check · 2 min

In scenario 2, Wetter flour, with the density trim off, the process-moisture trend holds the estimate on its 23.30 % setpoint for the whole run, and every heated barrel zone stays within 0.91 °C of its setpoint over the first 20 min. Meanwhile the density at the cutter climbs from 135 to 164.6 g/L. What does the estimate on its setpoint show?

(a) Only that the metered water matches what the certificate implies, since the soft sensor and the ratio station trust the same certificate

(b) That the dough entering the extruder holds 23.3 % water, within the noise of the flow transmitters

(c) That the density error must come from the barrel zones, since the process moisture is held on its setpoint

(d) That the at-line density is in error, since it reports 90 s after its cup is filled

5.2 Concept check · 2 min

In scenario 6, Heater limitation, zone 3, the kneading block, keeps a quarter of its heater capacity. Its loop pins the output at 100 % and the barrel falls from 140 to 132.0 °C, yet the melt at the die falls only 0.6 °C. Why so little?

(a) Back-calculation keeps zone 3's integral from winding up, so its loop makes up the lost heat further along the barrel

(b) The melt's heat comes mostly from viscous dissipation, and the cooler, stiffer dough in zone 3 dissipates more

(c) The melt thermocouple reads the barrel metal of zone 5, which its own loop still holds at 140 °C

(d) The preconditioner's steam loop raises the steam to supply the heat that zone 3 can no longer give

5.3 Concept check · 1 min

In scenario 7, Screw wear, the torque falls from 65.4 to 52.2 % by 1200 s at the same feed and screw speed. What does the falling torque tell the operator?

(a) The drive has gained headroom, so the feed can rise until the torque returns to its old value

(b) The dough has become wetter, since a thinner melt needs less torque

(c) The worn flights shear less, so less work goes into each kilogram and the melt is cooler and expands less

(d) The torque override has raised the screw speed, which lowers the torque at the same flow

5.4 Theory and calculation · 5 min

Scenario 2 feeds 300 kg/h of flour that holds 15 % moisture on a 12 % certificate. The recipe adds 22.48 kg/h of water and 21.71 kg/h of steam, for 344.2 kg/h of dough at a process-moisture setpoint of 23.3 %. Take the steam as unchanged.

(a) How much more water enters with the flour than the ratio station assumes, and how far above 23.3 % does the true process moisture run? *Checkpoint:* 9 kg/h, 2.61 points

(b) With the density gain of Table 5.4, estimate the density at the cutter and compare it with the 164.6 g/L measured. *Checkpoint:* about 165.0 g/L

(c) Use Equation (5.16) to find the water setpoint once the certificate is corrected to 15.0 %. Why does the water fall by more than the 9 kg/h of part (a)? *Checkpoint:* 10.75 kg/h, 11.73 kg/h less

5.5 Theory and calculation · 6 min

The supervisor predicts the torque at the feed it asks for by Equation (5.17). At steady state the two lagged flows agree, so $\rho = F_{base}/F_m$. At the recipe the torque is 65.4 % at 300 kg/h and 400 rpm, the torque limit is 85 % and the margin m_Γ is 3 %.

(a) Predict the steady torque at 380 kg/h and 400 rpm, and compare it with the 77.6 % of case 4. *Checkpoint:* 77.5 %

(b) Predict the torque at 480 kg/h with the screws at their 450 rpm top speed, and compare it with the 87.6 % of the steady-state solver. What is the supervisor's torque slack there, and what does the supervisor do? *Checkpoint:* 87.5 %, slack -1.84

(c) Taking the recipe as the base, with the screws at their 450 rpm top speed as in part (b), find the feed at which the slack falls to 0.25, where the supervisor stops raising the feed, and the feed at which it reaches zero. Compare them with case 4. *Checkpoint:* about 433 and 439 kg/h

5.6 Theory and calculation · 5 min

QIC executes Equation (5.15) once per at-line result, with $\Delta t_c = 60$ s, $K_c = 0.05$ points per g/L and $\tau_I = 240$ s. The step test of Section 5.5 found 11.25 g/L per point of moisture, 10 % of the change after 77 s and 63 % after 187 s, and the at-line check reports 90 s after its cup is filled, one cup a minute.

(a) Find the loop gain of the trim. *Checkpoint:* 0.56

(b) After a run of results on the setpoint, a result arrives 10 g/L above it. How far does that execution move the moisture setpoint, and how much of the move is proportional and how much integral? *Checkpoint:* 0.625 points down

(c) After a move of the moisture setpoint, about how many trim executions pass before an at-line result shows 63 % of its effect? What does that say about a larger gain? *Checkpoint:* about four

5.7 App challenge · 8 min

Scenario 2, Wetter flour, default seed, played at 60 times real time unless stated.

(a) Run the case as loaded, with the density trim off, to 1800 s. Export the CSV file and find the first rows in which `rho_cutter_gL` and `rho_atline_gL` exceed 150. Read the Off-spec product tile at 1800 s. *Checkpoint:* 209 s and 331 s, 64 kg

(b) Press Reset, set *Density trim of the moisture setpoint* On before Run, and run to 1800 s. Read the Off-spec product tile and the trim on the QIC card, and compare the trim with Exercise 5.4(a). *Checkpoint:* 24 kg, -2.67 points

(c) Reload the case. At 20 times real time run to about 300 s, pause, set *Flour moisture on the lot certificate* to 15.0 %, then run to 1800 s and read the Off-spec product tile. *Checkpoint:* 18 kg

Tolerance: 2 s on the crossings, 1 kg on the tiles and 0.02 points on the trim.

5.8 App challenge · 8 min

Scenario 4, Feed-rate increase, default seed, played at 60 times real time. Run each setting to 2400 s and export the CSV file.

(a) Run the case as loaded. Find the first row with `ctrl_infeasible` at 1 and the binding limit in the `limit` column, and at 2400 s read `feed_sp_kgh`, `screw_rpm` and the Worst torque tile. *Checkpoint:* about 1657 s, torque at maximum speed, 437.2 kg/h, 409 rpm, 85.5 %

(b) In the same file, find the first row with `torque_override_active` at 1 and the first row with the true torque `torque_pct` at or above 85. Why does the override act first? *Checkpoint:* 1423 s and 1552 s

(c) Reload the case, set *Density trim of the moisture setpoint* On before Run, and read `feed_sp_kgh` and the trim at 2400 s. Why does a trim that holds the density lower the feed the torque allows? *Checkpoint:* 416.0 kg/h, −0.41 points

Tolerance: 2 s, 0.5 kg/h, 1 rpm, 0.1 point of torque and 0.02 points of trim.

5.9 App challenge · 9 min

The challenge, a shift of flour lots, default seed, played at 60 times real time. Run each setting for the full 3600 s, then change the settings, press Reset and run again. Leave every parameter and disturbance as shipped.

(a) Run as shipped. Read the verdict, the Accepted and Off-spec product tiles and the Worst torque tile. *Checkpoint:* 186 kg, 83 kg, 74.0 %, fails

(b) Set *Density trim of the moisture setpoint* On, *Throughput supervisor* On and *Production target* 420 kg/h. *Checkpoint:* 290 kg, 34 kg, 87.4 %, passes

(c) Keep (b) and set *Restart the trim at a new lot*, an advanced option on the QIC card, Off. Read the same tiles, and the dense and light shares of the off-spec product in the CSV columns `metric_off_dense_kg` and `metric_off_light_kg`. Explain the result. *Checkpoint:* 244 kg, 75 kg with 44 kg light, 99.2 %, fails

Tolerance: 1 kg and 0.1 point.

5.10 Application area · 4 min

Where an extruder is the kill step of a food process, the plant must show that every particle received a lethal treatment. In case 1 at 400 rpm, measured from the middle of the dye pulse, the dye reaches 1 % of its peak rate after 7.7 s, peaks at 27.4 s and is 90 % out by 46.2 s, with a mean residence of 31.6 s.

(a) What fraction of the mean residence does the earliest dough spend in the barrel? Why is a treatment judged on the mean residence and the melt thermocouple not a lethality record? *Checkpoint:* about 0.24

(b) Name two ways in which this model could misstate what the fastest real particle sees, and say what a plant does instead.

5.11 Application area · 4 min

At the 450 rpm top speed the recipe needs 87.6 % of the 540 N m rating for 480 kg/h. Half a point more water brings that to 84.4 % but the density to 147.2 g/L, and a drive rated at 600 N m would carry it at 78.9 %.

(a) What torque and shaft power does 480 kg/h need at the top speed? Check the 78.9 %. *Checkpoint:* about 473 N m, 22.3 kW

(b) Compare the three ways to 480 kg/h: run at 87.6 % with the override limit raised, add water, or buy the larger drive. Which would you put to a plant manager, and what would you measure first?

5.12 Review the agent · 5 min

An AI assistant reads Table 5.9 and writes: “The MPC accepts 212.9 kg through the wet lot with nothing off-spec because it optimizes the feed, the moisture setpoint and the screw speed together over a five-minute horizon. The estimator is a detail, since any offset-free estimate would do. To match it with the loops, replace the supervisor and the trim by a multivariable optimizer.” Find the flaw and say how you would check it in the application.

CHAPTER 6

The Drying Front Sets the Line Speed of a Battery Electrode Coater

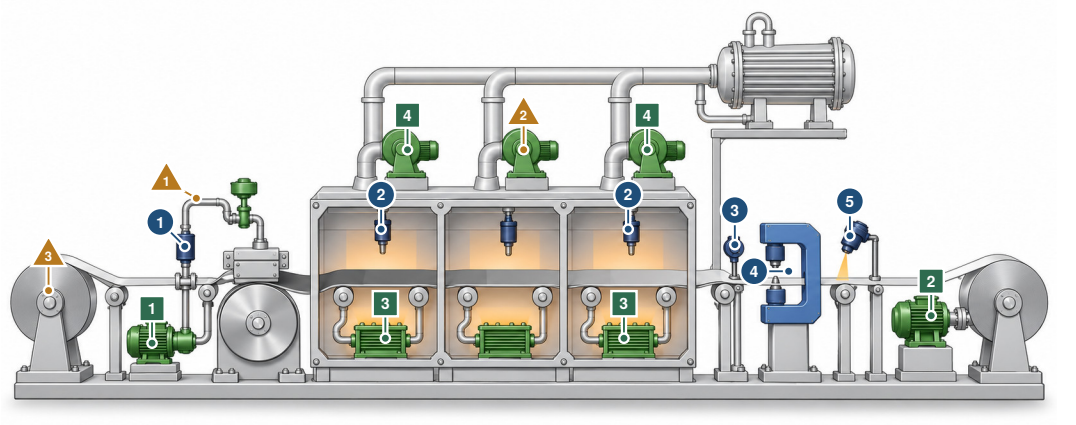

The controller moves 1 slurry pump 2 line drive (speed) 3 zone air heaters 4 exhaust fans and dampers

The controller reads 1 slurry density 2 solvent vapor in each zone 3 line speed 4 coating weight 5 web temperature at the exit

What disturbs it 1 a thinner slurry 2 a weakening exhaust fan 3 a splice in the foil

An electrode coater whose dryer, drawn with three of its six zones, keeps each zone's solvent vapor under a quarter of its explosive limit: notice the vapor monitor in every zone, because the zone that limits the speed moves with the drying front. Illustration, not to scale.

IN PLAIN TERMS Drying a Battery Coating Without a Fire Risk

The main idea. A coating line spreads a wet paste on thin metal foil and dries it in a long oven of six heated rooms. The paste gives off a vapor that can burn, so fresh air blown through each room keeps that vapor far below the level that could catch fire.

The challenge. A faster line raises the vapor in only one room, the one where the coating is just drying out. Drying too hard also weakens how well the coating holds the foil, and only a test an hour later finds it.

Why it matters. The vapor limit, not the coating, sets how fast the line can run.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page, the two history segments (Section 6.2), the fact on the air heater (Section 6.4), the two try-it segments and the fact on the predictive controller (Section 6.7) and the key points
Study	70 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 6.7), Technology in Practice (Section 6.8), Limitations (Section 6.9) and the open challenge

The control question The solvent vapor in a six-zone electrode dryer must stay under a quarter of its lower explosive limit, and that cap, not the coating, bounds the line speed. Which zone binds, which handle brings it back, and what do an override, a push and a better profile buy when binder damage is found only an hour later?

Learning objectives

- Compute the coat weight, the wet film and the solvent load from the die's ratio law, and the gauge's transport delay L/v.
- Explain why a faster line raises the vapor only in the zone that the drying front enters, and derive the exhaust bound on any profile.
- Choose the override's handle (exhaust, zone heat or speed) for a zone from where it sits against the drying front.
- Measure what each layer (override, push, feedforward, predictor, profile) adds to the challenge, and state what the margins cost.
- Judge a predictive controller that holds unmeasured quality predictions at their limits.

Prerequisites Sections 1.4.2 (active constraints and back-off), 1.4.3 (selectors), 1.4.4 (external reset) and 1.6.2 (velocity form). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/electrode-coating/

See also Volume 5, Chapter 5 for a predictor on a moving line, Volume 5, Chapter 10 for a throughput push on inferred limits and Volume 5, Chapter 11 for parcels through a long delay. Volume 3, Chapter 6 for feedforward with a slow ratio trim and Volume 3, Chapter 4 for selectors that protect a ratio.

6.1 Introduction

A cathode coater turns a slurry into an electrode in about two minutes. In the line simulated here a pre-metered slot die lays an NMP slurry of 70 % solids on aluminum foil 1.2 m wide, a wet film 115.7 µm thick that carries 77.1 g/m^2 of solvent and dries to 180 g/m^2 of coating. At 35 m/min the six 9 m zones of the dryer must take about 0.054 kg/s of NMP out of the web, and a beta gauge 60 m from the die weighs the result 102.9 s after it was coated. Raising the speed to 40 m/min leaves the vapor in zones 2 and 3 almost where it was, near 21 % LEL, and carries zone 4 from 18.57 to 29.7 % LEL, past a cap of 25. A gentler profile, with zones 2 to 4 cooler by 10 to 25 °C, nearly triples at the same speed the strength with which the coating holds the foil, from 8.0 to 23.3 N/m, and nothing measured on the line tells the two runs apart.

Two misconceptions meet in this plant. The first is that a dryer limits a line by its heat, so that more speed loads every zone a little more. A zone that is wet from end to end evaporates at the rate its air allows whatever the speed, and a faster web only carries the wet coating into the next zone, so the safety constraint moves with the drying front and the zone that binds depends on the speed and the profile. The second is that a constraint is held by cutting the throughput. Cutting the speed lowers the vapor only in the zone the front is entering, and in a wet zone a slower web is wetter for longer, so the right handle differs from zone to zone. The quality side has the same shape. Binder carried to the surface by fast drying lowers the adhesion of the electrode [1, 2], and its cause lies in a stretch of the profile that no instrument on the line watches.

The class is wide. Any plant whose throughput is bounded by a safety or quality limit measured at many points, with the active limit moving between those points as the operating point changes, has this form, among them solvent ovens for coatings and printing, spray dryers, kilns and multi-bed reactors. The spray dryer of Volume 5, Chapter 10 pushes its throughput under overrides on margins that are inferred, much as the push of this chapter predicts limits that no instrument on the line measures. Selector logic has switched between such constraints in industry for decades with designs that rest on experience [3], a model of the constraint corrected by measured biases is the basis of constraint adaptation in real-time optimization [4], and dead-time compensators, surveyed in [5], address a measurement that arrives a transport delay after its cause.

These pieces are rarely seen together on one plant whose truth can be read beside its instruments. This chapter specifies a six-zone NMP cathode dryer, built as a browser simulator, so completely that it can be rebuilt, and reports measured runs on it. Its coat weight loop, explosive-limit override, push toward the binding constraint and offline profile search can each be enabled separately, and the runs ask where the constraint binds and what each layer buys. Ratio control, the Smith predictor, override selection, bias-corrected constraint models and direct search are established methods [3–6], and the drying and migration laws are simplified from published work [2, 7]. The optional model predictive controller of Section 6.5.6 is a standard receding-horizon design [8], and its comparison with the push is reported as measured, losses included. A focused search found no published treatment of constraint pushing against a vapor limit that moves with a drying front, or of the release logic a multi-zone override needs, so the measured findings below, including two defects of such logic found and fixed during review, are results of this model and not general claims about override design. The solvent ledgers close to within 10^{-13} of the solvent fed, the drying law is checked against an independent integration, and the equations and constants listed here are enough for a second implementation.

Section 6.2 traces the industry, its safety rule for solvent ovens and the evidence on binder migration. Section 6.3 sets out the line as modeled: the die, the web as parcels, the zones and their vapor, the migration index, the instruments and every parameter with its basis. Section 6.4 explains why the vapor cap binds at the drying front and derives an exhaust bound on any profile. Section 6.5 describes the control layers as implemented, from the zone loops through the override and the push to the optional MPC. Section 6.6 keeps code verification apart from agreement with sources. Section 6.7 measures the nine cases, the comparison of the MPC with the push and the challenge, Section 6.8 sets the devices of a real line beside the model, and Sections 6.9 and 6.10 state the limitations and the conclusions.

6.2 Electrode Lines, Solvent Ovens and Binder Migration

The inventor of the lithium-ion cell has described how it was first developed and made commercially practical, with successful safety tests as the turning point [9]. Its use has since moved from portable electronics into vehicles, and the production technologies of automotive cells, from electrode making to cell assembly, have been reviewed for their current status and challenges [10]. Within that chain the dryer and the recovery of its solvent stand out. An economic analysis of NMP-based and water-based electrode drying found that electrode manufacturing makes up about 8 to 9 % of the pack cost either way, and that water could cut the cost of drying and solvent recovery by up to a factor of two and save 3 to 6 million dollars of capital equipment for a plant making 100,000 packs of 10 kWh [11]. A continuum model of cathode drying showed that about 90 % of the solvent leaves in less than half the drying time and that splitting the oven into zones can cut the drying energy by at least half and the peak solvent flux by at least 40 % [12].

The safety rule came from older industries, the ovens that dry solvent coatings, and the history segment that follows tells its story.

⧖ HISTORY Solvent Ovens and Their Margin, 1969 to 2007

Industrial ovens evaporate flammable solvents, and because the vapor can burn, the exhaust air carries it out before it builds up. Fire-protection guidelines set how much air that takes as a ventilation ratio over the solvent load [13]. United States regulation requires the air exhausted through a spray booth to dilute solvent vapor to at least 25 percent of its lower explosive limit, and it sends the drying and curing apparatus of spray finishing to the oven standard NFPA 86A-1969 [14]. A federal tip sheet of 2007 showed what continuous monitoring changes. With lower flammable limit monitors an oven may run a smaller safety ratio, and the monitor can set the ventilation through an adjustable-speed drive on the exhaust fan or through its dampers. In the tip sheet’s example of a continuous strip coating process the ratio fell from 4:1 to 2:1, halving the exhaust and the energy [13].

Table 6.1 Roles of the Variables and Their Design Values in Case 3 (Standard Profile, 35 m/min)

Variable	Role	Design	Range
g	CV, beta gauge	180 g/m^2	SP 150–220
$\max_z$ %LEL$_z$	CV of the push and override	21.1	cap 25
T_z	CV, zone air (TIC 1–6)	125/105/110/125/125/125	70–170 °C
h	MV, wet-film ratio	115.7 µm	40–220
v	MV, line speed	35 m/min	10–80
$Q_{h,z}$	MV, zone heater	—	0–300 kW
F_z	MV, zone exhaust	1.53 kg/s	30–100 % of fan
x	DV, slurry solids	70 %	64–76
—	DV, pump slip, fan capacity	0, 100 %	0–8, 40–100
—	DV, events	—	drift, fan, heater, splice
s, T, M	states, per parcel (120)	—	—
$m_{v,z}$, T_z	states, per zone	—	—
v, Q, F_z	states, actuators	—	—
s_{exit}, M	outputs, residual and migration	0.068, 6.13 g/m^2	truth only

The current NFPA 86 could not be consulted for this chapter and is not quoted here. A ratio of 4:1 is the quarter of the limit used as the cap in this simulator, and the halved ratio corresponds to the half of the limit allowed with continuous monitors, a margin the simulator does not use.

Binder migration is the quality side of the same oven, and the second history segment follows the laboratory evidence on it. The tip sheet's monitor sets the exhaust, and no source found in a focused search treats the line speed as a variable pushed against the explosive-limit constraint of a multi-zone dryer.

> ⧖ HISTORY **Binder Migration and the Peel Test, 2016 to 2018**
>
> Graphite anodes dried at higher temperatures in a laboratory impingement dryer held their copper foil less strongly in a 90° peel test, which pointed to a change in the binder at the interface between the foil and the coating [1]. Film solidification and binder migration were then studied together during anode drying [15], and a three-stage profile, a mild drying stage between two fast ones, kept the adhesion while cutting the drying time by about 40 % against a reference process [16]. A mathematical model of binder transport agreed qualitatively with such experiments: low drying rates left the binder even across the film and high rates piled it near the evaporating surface. For the film it studied the rate had to stay near 1 µm/s or below, and drying fast at the start and slowly near the end gave the most even binder [2].

6.3 The Coating Line as Modeled

The line is a generic teaching design assembled from public values, as shown in Figure 6.1. Foil from an unwind passes a slot die 3 m before the dryer, six zones of 9 m each with their own air heater, nozzles on both sides of the web, an exhaust damper and a vapor monitor, then 3 m of open web to the beta gauge and the rewind, so the gauge is $L = 60$ m of web from the die. The exhaust of every zone is fresh-air dilution that leaves for a common condenser. A densitometer reads the slurry before the pump, a pyrometer reads the web at the dryer exit, and a laboratory cuts a strip at the gauge every half hour and reports its peel strength half an hour later. The variables are classified in Table 6.1.

6.3.1 The Die and the Coat Weight

The slot die is pre-metered, so everything the pump delivers lands on the foil, and the slurry density follows a volume-additive mixing rule, as shown in Equation (6.1).

$$\frac{1}{\rho} = \frac{x}{\rho_s} + \frac{1-x}{\rho_l}, \qquad h_{wet} = \frac{Q}{wv}$$
$$m_A = \frac{Q\rho x}{wv}, \qquad s_0 = m_A \frac{1-x}{x} \tag{6.1}$$

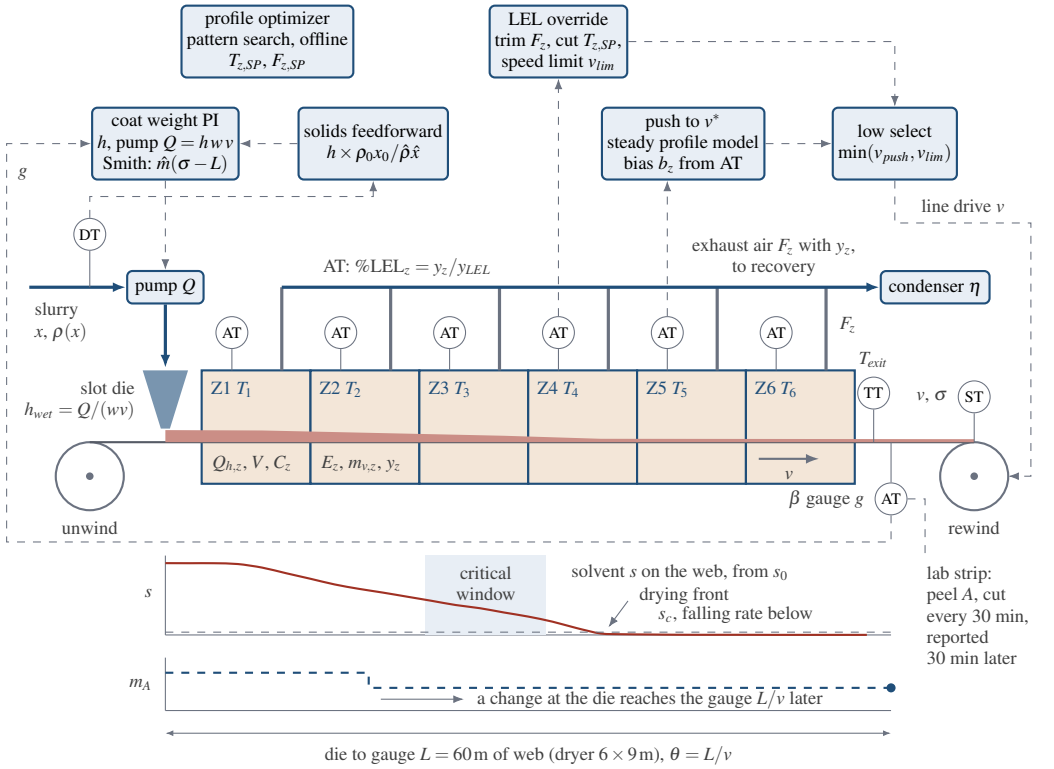

Figure 6.1 The cathode coater as modeled, mirroring the Line view of the simulator. The pump delivers Q of slurry of solids x and density $\rho(x)$, read by the densitometer DT, to the slot die, which lays a wet film h_{wet} on the foil moving at v. Each zone z holds air at T_z heated by $Q_{h,z}$, with volume V and heat capacity C_z, gathers the solvent E_z the web evaporates into its vapor $m_{v,z}$ (mole ratio y_z), and exhausts F_z to the condenser of recovery η, while its monitor AT reads $\%\text{LEL}_z$. TT reads the exit web, the beta gauge reads g and ST the speed and the web length σ, and the laboratory strip gives the peel strength A. Below, the solvent s on the web against position from case 4 (Standard profile) at 3600 s, with the critical window and the critical content s_c, and the coat weight m_A along the web, whose change at the die reaches the gauge L/v later. Above, the control layers as implemented, each fed by instruments only. Positions along the line are to scale. Process model only, not an equipment or safety drawing.

where ρ_s and ρ_l are the densities of the solids and of NMP, w the coated width and s_0 the solvent laid per unit area. At 70 % solids the slurry is 2223 kg/m^3, and the pump that lays 180 g/m^2 at 35 m/min delivers 4.86 L/min. In Auto the pump is geared to the line, $Q = hwv$ at every step with h the controller's wet-film ratio, and in Manual its flow is held at the operator's value whatever the speed. The delivered flow lags its command with $\tau_Q = 0.3$ s and loses the slip fraction set by the viscosity slider, and the line drive moves toward its command at no more than 0.5 m/min per second.

6.3.2 The Web as Parcels

The web is a Lagrangian line of parcels $\Delta\ell = 0.5$ m long, born at the die and retired at the gauge, so a parcel keeps the coat weight it was given and the gauge sees it exactly L/v later at any speed history. Each parcel carries its solvent s, its temperature T and its migration index M, and dries with the flux

of Equation (6.2).

$$\begin{aligned} \frac{\mathrm{d}s}{\mathrm{d}t} &= -\dot{m} \\ \dot{m} &= \beta\,\frac{M_w}{R\,T_f}\left[p_{sat}(T)\,f(s) - y_z P\right] \\ f(s) &= \begin{cases} 1 & s \ge s_c \\ (s/s_c)^n & s < s_c \end{cases}, \qquad \beta = \frac{h_t}{\rho_{air}\,c_p\,Le^{2/3}} \end{aligned} \tag{6.2}$$

where $T_f = (T + T_z)/2$ is the film temperature, $\rho_{air} = PM_{air}/(RT_f)$, P is atmospheric, h_t is the nozzle coefficient on the coated side and the mass coefficient β follows from it by the Chilton and Colburn analogy [7]. Above the critical content s_c the surface is wet and the rate is set by the air, the constant-rate period. Below it the solvent's activity on the web falls as $(s/s_c)^n$ with $n = 2$, so the web can only dry toward the content in equilibrium with its zone's vapor. The parcel's energy balance, Equation (6.3), carries the foil, the dry coating and the solvent left.

$$\begin{aligned} C_w(s)\,\frac{\mathrm{d}T}{\mathrm{d}t} &= (h_t + h_b)(T_z - T) - \Delta H_v\,\dot{m} \\ C_w &= c_{foil} + m_A c_s + s\,c_l \end{aligned} \tag{6.3}$$

In the constant-rate period the web sits near the temperature at which the nozzles' heat balances the evaporation, a wet-bulb temperature. Outside the dryer a parcel sees room air at 25 °C with $h = 8\,\mathrm{W/(m^2\,K)}$ and no vapor, and what it loses there is booked to the hall. The vapor pressure of NMP is a Clausius–Clapeyron line through two points of the NIST Chemistry WebBook, the normal boiling point 475.2 K and 354.7 K at 0.013 bar [17], as shown in Equation (6.4).

$$p_{sat}(T) = P_0 \exp\!\left[-B\left(\frac{1}{T} - \frac{1}{T_b}\right)\right], \qquad B = 6093\,\mathrm{K} \tag{6.4}$$

which gives 6.98 kPa at 120 °C and 14.8 kPa at 140 °C, a doubling every 17.6 K near 120 °C.

6.3.3 Zones, Vapor and the Explosive-Limit Fraction

Each zone is a well-mixed volume of air at T_z, supplied with fresh air equal to its exhaust, so its vapor and air temperature obey Equation (6.5).

$$\begin{aligned} \frac{\mathrm{d}m_{v,z}}{\mathrm{d}t} &= E_z - \frac{F_z}{m_{air,z}}\,m_{v,z} \\ y_z &= \frac{m_{v,z}/M_w}{m_{air,z}/M_{air}} \\ C_z\,\frac{\mathrm{d}T_z}{\mathrm{d}t} &= Q_{h,z} - (F_z c_p + UA)(T_z - T_{amb}) - Q_{web,z} \\ \%\mathrm{LEL}_z &= 100\,y_z/y_{LEL} \end{aligned} \tag{6.5}$$

where $m_{air,z} = PVM_{air}/(RT_z)$, $E_z = \sum w\Delta\ell(s_A - s_B)/\Delta t$ over the parcels in the zone during the step, $Q_{web,z}$ the heat they take, and $y_{LEL} = 0.013$, the lower explosive limit of NMP, 1.3 to 9.5 vol% in air on the International Chemical Safety Card [18]. In steady state the vapor is the evaporation over the exhaust, $y_z = (E_z/M_w)/(F_z/M_{air})$. The exhaust flow follows its damper command with a 5 s lag, limited by the fan's capacity F_{max} times the fan's health, and the heater delivers its command between zero and its rating, which a heater fault lowers to 40 %. Three ledgers are integrated alongside, as shown in Equation (6.6).

$$\begin{aligned} \text{fed} &= \text{evaporated} + \text{hall} + \text{past gauge} + \Delta\text{web} \\ \text{evaporated} &= \text{exhausted} + \Delta\text{vapor} \\ \text{exhausted} &= \text{recovered} + \text{emitted} \end{aligned} \tag{6.6}$$

with the condenser recovering a fixed $\eta = 0.95$ of what is exhausted.

6.3.4 The Migration Index and the Peel Test

The migration index counts how far the drying rate exceeds a safe rate while the film passes through a critical window around the end of its shrinkage, and the laboratory converts it into a peel strength, as shown in Equation (6.7).

$$\frac{\mathrm{d}M}{\mathrm{d}t} = \left(\dot{m} - r_{safe}\right)^{+} \text{ for } 0.5 \le 1 - \frac{s}{s_0} \le 0.8$$
$$A = \frac{A_0}{1 + M/M_{max}} \tag{6.7}$$

with $r_{safe} = 1.0\,\mathrm{g/(m^2\,s)}$, $A_0 = 30\,\mathrm{N/m}$ and the threshold $M_{max} = 2\,\mathrm{g/m^2}$ where the peel strength halves. The safe rate is the source's estimate, 1 µm/s of NMP recession for a 114 µm film [2], which is 1.03 g/(m^2 s), and the window's start, about half the solvent gone, lies where that model's binder at the surface rises fastest when a high rate lasts too long. The window's end at 80 %, the critical content, the peel law and its threshold are teaching values. The index is a proxy that stands for binder carried to the surface, and it answers to where the window lands. A hotter first zone moves the window earlier, so with zones 3 and 4 held, zone 1 alone at 115 °C on the Standard profile tests worse than at 125 °C (7.2 against 8.0 N/m), while zones 3 and 4 alone at 95 and 100 °C test 28.6 N/m. The sources make a fast first stage harmless [2, 16], not helpful, so the first-zone result belongs to the proxy.

6.3.5 Instruments

The controller sees only a packet of measurements, and one of the application's tests searches its source for any truth field and fails on a match. The zone thermocouples have a 5 s lag and 0.3 K of colored noise, the vapor monitors an 8 s lag and 0.2 % LEL, the exhaust transmitters a 2 s lag and 1 %, the densitometer a 10 s lag and 1 kg/m^3, and the pyrometer a 2 s lag. The beta gauge weighs everything passing it, the dry coating, any solvent left and the second foil of a splice, as shown in Equation (6.8).

$$\tau_g \frac{\mathrm{d}g}{\mathrm{d}t} = m_A(\sigma - L) + s_{exit} + m_{spl} + \sigma_g n(t) - g \tag{6.8}$$

where $\tau_g = 2\,\mathrm{s}$, σ is the web length the encoder has counted, $\sigma_g = 0.6\,\mathrm{g/m^2}$ and n a unit colored noise with a 1 s correlation time. The laboratory is the framework's sampled analyzer (Section 1.6), which cuts a strip at the gauge every 1800 s, reports Equation (6.7) with 0.8 N/m of noise 1800 s later and holds the report until the next. The residual solvent and the migration index are truth only.

6.3.6 Parameters and Their Evidence

Table 6.2 lists every constant with its basis. None is plant data. Each is a public property, a value fitted to two public points, or a teaching value chosen to make a consistent design point, and the safety and quality limits are teaching values whose real counterparts come from the oven standard and from the electrode specification.

6.3.7 Integration, Line-Out and the Declared Stream

The application integrates with its own scheme inside the framework's engine (Section 1.6.1). In each 0.5 s step every parcel takes one fourth-order Runge–Kutta step in the zone it is in, with that zone's air and vapor frozen, and every zone is then advanced exactly over the step, because Equation (6.5) is linear in its own state once the evaporation is known. Parcels past the gauge leave and new ones are born at the die. The controller executes every 2 s and the historian logs every 2 s. A case begins in the steady state of its own settings. A steady march of one parcel through the line, with the zone vapor found by fixed-point iteration, fills the line with parcels in the state the march gives at their positions, and in Auto the pump is set so the gauge, residual included, reads the target. Over 120 s open loop from that start the zone vapor moves by at most 0.05 points of LEL and the air by 0.007 K. A steady

Table 6.2 Parameters, With Value and Basis. Source: Cited Public Value. Fitted: Through Cited Points. Teaching: Chosen. Adj.: Adjustable in the Application

Symbol	Meaning	Value	Basis
w	coated width	1.2 m	teaching
L	die to gauge (3 + 6×9 + 3)	60 m	teaching, adj. exit
$\Delta\ell$	parcel length	0.5 m	numerical
ρ_s, ρ_l	solids, NMP density	4430, 1028 kg/m^3	teaching, source
x_0	recipe solids	70 %	teaching, adj.
$m_{A,SP}$	coat weight per side	180 g/m^2	teaching, adj.
$h_t = h_b$	nozzle coefficient, each side	30 W/(m^2 K)	teaching, adj.
$Le^{2/3}$	Lewis number term	1.9	teaching
c_{foil}	15 µm aluminum	36.5 J/(m^2 K)	derived
c_s, c_l	solids, NMP heat capacity	800, 1700 J/(kg K)	teaching
T_b, B	boiling point, slope	475.2 K, 6093 K	fitted [17]
ΔH_v	heat of vaporization	500 kJ/kg	source [17]
y_{LEL}	lower explosive limit	1.3 vol%	source [18]
s_c, n	critical content, exponent	3 g/m^2, 2	teaching, adj.
r_{safe}	safe rate in the window	1.0 g/(m^2 s)	source [2], adj.
—	window, dried fraction	0.5–0.8	start source, end teaching
M_{max}, A_0	migration threshold, peel scale	2 g/m^2, 30 N/m	teaching, adj.
s_{max}	residual limit at the gauge	0.5 g/m^2	teaching, adj.
—	LEL cap	25 %	class [13, 14]
V, C_z, UA	zone air, capacity, loss	11.5 m^3, 80 kJ/K, 200 W/K	teaching
Q_{max}	heater rating per zone	300 kW	teaching, adj.
F_{max}	fan capacity per zone	1.8 kg/s	teaching, adj.
τ_F, τ_Q	exhaust, pump lags	5, 0.3 s	teaching
η	condenser recovery	0.95	teaching
T_{amb}	room	25 °C	teaching

model at the Standard profile and 35 m/min gives 18.25, 20.72, 21.09 and 18.61 % LEL in zones 1 to 4, a residual of 0.068 g/m^2 and an energy of 1.475 MJ/m^2 at the heaters.

The coater is one link of a battery line of applications, from the `lithium-refinery` and `cathode-kiln` applications to `ev-fast-charge`. The kiln of Volume 5, Chapter 11 fires the cathode material that this line coats, and it too carries its product as Lagrangian parcels through a long transport delay. The coater declares its product in a stream table, 180 g/m^2 per side with 96 % active material at 170 mAh/g, 2.938 mAh/cm^2 per side, all teaching values. The table is a declaration that a later application may copy and cite, not a coupling, and no other application reads it.

6.3.8 Assumptions

One side is coated per pass, and the second side would repeat it on a second pass that is not simulated. The die stays inside its coating window, so bead breakup, edge build-up and the cross-web profile are absent. Each zone is well mixed, and the web sees its zone's air temperature and vapor with fixed nozzle velocities. The falling rate is an activity law with no diffusion limit inside the film, so binder sorption of NMP, skin formation and cracking of thick films are not represented, and the residual is least reliable at the smallest values. The LEL is a room-temperature value, although explosive limits fall slightly in hot air, and y_z is the ratio of NMP to air, which differs from the mole fraction by 0.3 % at the cap. One fan serves the six dampers and loses capacity evenly. The migration index and the peel law are teaching proxies, and a real electrode's adhesion also depends on its binder, carbon and calendering.

6.4 Why the Drying Front Carries the Cap

6.4.1 Wet Zones Do Not Answer to Speed

A zone the wet film fills from end to end evaporates at the constant rate its air sets, $wL_z\dot{m}$, whatever the speed, so its vapor at a fixed exhaust is fixed. On the steady model of the Standard profile at 85 % exhaust, zones 2 and 3 sit at 20.4 and 21.2 % LEL at 30 m/min and at 20.9 and 20.9 % LEL at

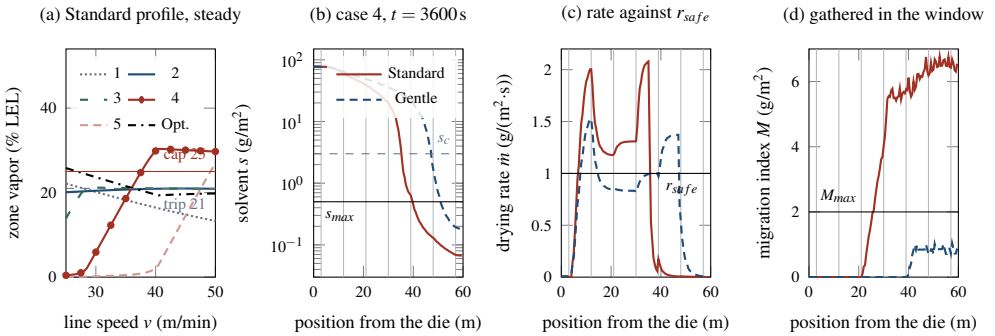

Figure 6.2 Where the vapor cap and the migration are decided, from the simulator's own functions. (a) Zone vapor against line speed on the steady model, Standard profile at 85 % exhaust, with zone 1 dotted, zone 2 solid, zone 3 loosely dashed, zone 4 solid with dots, zone 5 densely dashed and the worst zone of the Optimized profile dash-dotted. The cap is the thin solid line at 25 and the trip the thin dashed line at 21. The wet zones 2 and 3 stay near 21 % LEL while the zone the drying front enters climbs past the cap. (b to d) Case 4 at 3600 s and 35 m/min, every parcel on the line, on the Standard (solid) and Gentle (dashed) profiles: the solvent against position with the critical content and the residual limit, the drying rate against the safe rate, and the migration index against its threshold. Thin vertical lines are the zone boundaries. The Standard profile's zone 4 at 125 °C dries the window at up to 1.82 g/(m^2 s) and the web reaches the gauge with an index of 6.2 g/m^2, while the Gentle profile holds its window to 1.29 g/(m^2 s) at the peak and reaches the gauge with 0.87 g/m^2. The strips the laboratory cut at 1800 s tested 8.0 and 23.3 N/m.

50 m/min, as shown in Figure 6.2a. The zone that the drying front enters is only partly wet, and a faster web carries more wet coating into it, so its vapor climbs. Zone 4 rises from 5.9 to 29.9 % LEL between 30 and 40 m/min and first passes 25 at 38 m/min, after which it is wet throughout and levels off near 30. Zone 1 falls with speed, from 22.1 at 25 m/min to 13.3 % LEL at 50 m/min, because a faster web enters it colder and stays colder longer. Above 45 m/min the front reaches zone 5. The constraint therefore binds on whichever zone holds the front, and the profile decides which one that is. In the language of Section 1.4.2, the active constraint is a different variable at different operating points.

🕮 **Pause.** Try Exercise 6.1 before reading on.

6.4.2 An Exhaust Bound on Every Profile

All the solvent the die lays must leave through the exhaust, so the fans bound the speed of any profile, as shown in Equation (6.9).

$$w\, v\, s_0 = \sum_z E_z \le \frac{M_w}{M_{air}}\, 0.25\, y_{LEL} \sum_z F_{max,z} \qquad (6.9)$$
$$\Rightarrow\ v \le 77.86\,\mathrm{m/min}$$

Six fans of 1.8 kg/s at a quarter of the LEL carry at most 0.120 kg/s of NMP. Real profiles stay well under the bound because the first zone heats the web and the last finishes it at low rates. None of 3000 random profiles at 80 m/min, with air from 80 to 160 °C and exhaust from 50 to 100 %, kept both the cap and the residual limit. With the residual kept the lowest worst zone was 34.4 % LEL, and with every zone under the cap the lowest residual was 16.72 g/m^2. Eight pattern searches with only the cap and the residual as limits reached 70.9 m/min at best, so no profile reaches the top of the speed range, 80 m/min, and the application says so.

> **DID YOU KNOW? The Air, Not the Solvent, Takes the Heat**
>
> A process model of solvent recovery for a battery plant using 4.1 million kilograms of NMP a year put the dryer's own heat load near 130 kW and the air heater's near 3800 kW. The air heater works that hard because the dryer air must keep the vapor far below its flammability limits [11].

6.4.3 The Profile Decides the Window

The critical window is a stretch of the dried fraction, not of the line, and its position along the dryer moves with the speed and the profile. At 35 m/min on the Standard profile it lies between 21.5 and 31.5 m from the die, in zones 3 and 4, and zone 4 at 125 °C drives the rate there to 1.82 g/(m^2 s) against a safe rate of 1.0, as shown in Figure 6.2c. The Gentle profile's 95 and 100 °C in zones 3 and 4 place its window between 28.5 and 41.5 m and hold the rate there to 1.29 g/(m^2 s). The residual answers to the last zone in the same way. Below s_c the rate is the web's vapor pressure times its activity, so on the Optimized profile at 50 m/min the residual is 1.377, 0.615, 0.324 and 0.244 g/m^2 with zone 6 at 105, 125, 147.5 and 160 °C, while zone 6's vapor moves only from 13.7 to 15.7 % LEL. Three objectives thus live in three places along the line, the cap at the front, the migration in the window and the residual at the exit.

6.5 The Control Layers as Implemented

All loops execute every 2 s on the instrument packet, in the order zone air, exhaust, override, push and coat weight, and the plant receives the commands at the next physics step.

6.5.1 Zone Air and Exhaust

Each zone's air is held by the framework's velocity-form PID of Section 1.6.2, Equations (1.13) to (1.15), with a control period $\Delta t_c = 2\,\mathrm{s}$. Its derivative term is off in every loop of this application, which leaves a PI law whose running sum is never clamped and whose integral alone carries back-calculation, with a tracking time equal to τ_I. The zone loops output the heater's share of its rating, 0 to 1, with $K_c = 2\,\%$ per kelvin and $\tau_I = 45\,\mathrm{s}$, and their setpoint is the profile's less any heat cut the override has made. Each exhaust damper takes its setpoint as a fraction of the fan's nameplate. With the override on, an integral trim $\delta_z \leftarrow \delta_z + 0.5\,(\%\mathrm{LEL}_z/100 - 0.20)\Delta t_c$, clipped between zero and the damper's room to full flow, adds exhaust while a zone's monitor reads over the push target of 20 % LEL and gives it back below.

6.5.2 Coat Weight: Ratio, Predictor and Feedforward

The coat-weight loop is the same PI on the gauge, so a light gauge raises the wet-film ratio, with $K_c = 0.3\,\mu\mathrm{m}$ per g/m^2, $\tau_I = 90\,\mathrm{s}$ and the ratio limited to 40 to 220 µm. Its output is the ratio h, and the pump flow is hwv, so a speed change carries the pump with it. With the predictor on, the loop's measurement is the model coat weight now, corrected by the gauge against what the model laid down for the web under the gauge, as shown in Equation (6.10).

$$\mathrm{PV} = \hat{m}(t) + g(t) - \hat{m}(\sigma - L), \qquad \tau_m \frac{\mathrm{d}\hat{m}}{\mathrm{d}t} = h\hat{\rho}\hat{x} - \hat{m} \tag{6.10}$$

The model's lag is $\tau_m = 2.3\,\mathrm{s}$, and its value is stored against the encoder's web length every 0.25 m and read back exactly 60 m later, so the predictor's delay is L/v at every speed, which a predictor keyed to a fixed time would not be. The float glass line of Volume 5, Chapter 5 runs a predictor of the same family on a ribbon whose thickness is measured downstream of where it is formed. The lag stands for the pump and the gauge's average. Without it a gain above 0.64 µm per g/m^2 made the sampled loop through the algebraic model term unstable. With the solids feedforward on, $\hat{\rho}\hat{x}$ comes from the densitometer through the inverse of Equation (6.1), and the ratio is scaled by $\rho_0 x_0/\hat{\rho}\hat{x}$, the recipe's solids load over the estimate's. The feedforward acts on a measured disturbance at the input, and the gauge loop remains as a slow trim, the arrangement that the water plant of Volume 3,

Chapter 6 uses for its coagulant dose. A change of either option, and Manual to Auto, re-bases the integral on the output the loop has (Section 1.6.2), and the measured jump in the next period is 0.000 µm. The gauge reading is held while a splice that the encoder tracked passes it.

6.5.3 The Explosive-Limit Override

The override uses three handles in order. The exhaust trim is the first. A zone whose monitor reads at or over the trip, 21 % LEL, while its exhaust is already at capacity gets a second handle chosen by the push's profile model on measured inputs. If a line 2 m/min slower lowers that zone's predicted vapor by more than 0.3 points, which is the case for a zone the front is entering, the zone joins the speed-limited set. Otherwise, as for a wet zone or the first zone, its own air setpoint is cut at 8 °C/min, up to 40 K, and the heat returns at 1 °C/min once the zone reads 1.5 points under the trip. The speed limit starts at the requested speed when the set first fills, falls at up to 15 m/min per minute (from 40 % of that rate at the trip to all of it 1.5 points over) while any zone of the set reads at or over the trip, holds within 2 points under it, and rises at the push rate, 2 m/min per minute, below that, until the request falls under the limit and the override lets go. A low selector applies $\min(v_{push}, v_{lim})$, and the push's own speed tracks the applied one while the limit acts, which is the external reset of Section 1.4.4, so the hand-back is bumpless. The selector itself is the cap of Section 1.4.3, with the twist that the protected variable is whichever zone holds the front. Over 25 releases in case 8 the largest step in the speed request was 0.000 m/min.

Two release rules were added during the review of the application, each after a measured defect, and both are part of the controller of this chapter. In the first, the exhaust trim held a speed-limited zone at its 20 % LEL target, inside the hold band from 19 to 21, so a zone that had joined the set at a low speed kept the limit down for good. Case 3 with the setpoint raised at 50 s rested at 28.1 m/min, and in sweeps of the moment of the change, 10 of 81 runs and 22 of 82 runs rested under 31 m/min. A zone now leaves the set after 60 s under the trip with 3 % of room in its exhaust, after which its trim holds it, and the same sweeps rest under 31 m/min in 0 of 81 and 0 of 82 runs, with the case 3 run at 35.5 m/min. In the second, after a deep cut in case 7 the zone upstream of the front took the speed handle as well and the line settled near 21.5 m/min instead of 27.6, in 3 of 8 seeds at the shipped step and 6 of 8 at half of it, so halving the step changed the lesson. A zone upstream of one already speed-limited now takes the heat handle, and all 16 runs average 27.2 to 27.9 m/min. An alternative rule that chose the handle at the requested speed was tried and dropped because it let the line run at 44.7 m/min on cut heat at a setpoint of 45 m/min. The challenge's runs were unchanged by both rules, while the predictor-off run of case 2 moved from a gauge range of 119.9 to 276.8 to one of 122.5 to 271.0 g/m^2.

6.5.4 The Push

Every 10 s the push runs the plant's own steady model with the declared parameters of Table 6.2, Equations (6.1) to (6.7) in steady state, on the measured zone air, the exhaust each zone can get, the gauge and the solids estimate, and finds the fastest speed that meets every limit, as shown in Equation (6.11).

$$v^{*} = \max\{v : b_z\,\%\mathrm{LEL}_z(v) \le 20,\; s_{exit}(v) \le 0.8\,s_{max},\; M(v) \le 0.9\,M_{max}\} \tag{6.11}$$

The bias b_z is each monitor's reading over the model's prediction at the present speed, filtered with a 120 s time constant and clipped between 0.5 and 2, the constraint-adaptation idea of real-time optimization [4]. The search scans down from 80 m/min in steps of 2 and bisects to 0.1 m/min. A limit broken at every speed, a zone whose vapor the speed moves by less than 1.5 points over the range, and a zone the override holds by its heat are reported and left out, because no speed can meet them. The push's speed moves toward v^{*} at 2 m/min per minute in either direction. The push's model has the declared parameters and not the run's, so a learner's edit of the plant is model error that only the bias corrects, and only for the vapor. The margins of Equation (6.11), a target of 20 under a trip

of 21 and a cap of 25, and 80 % and 90 % of the quality limits, are the back-off of Section 1.4.2, paid for in line speed. The extruder of Chapter 5 runs a throughput supervisor of the same kind against its torque, pressure and dryer limits.

6.5.5 The Profile Optimizer

The Optimized profile is the result of a coordinate pattern search in the manner of direct search [6] on the same steady model, over six zone temperatures from 80 to 160 °C and six exhaust fractions from 0.5 to 1, with steps from 10 down to 2.5 °C and 0.1 down to 0.05, as shown in Equation (6.12).

$$\max_{T_{1..6},F_{1..6}} J = v_{max}(T,F) - w_E\,E(T,F,v_{max}) \tag{6.12}$$

where v_{max} is the fastest speed with every zone at 20 % LEL, the residual at 0.40 g/m^2 and the migration at 1.8 g/m^2, found by a scan and a bisection, and E the heater energy per square meter. From the Standard profile with $w_E = 0$ it ran 4516 model evaluations and found 137.5, 105, 102.5, 100, 105 and 147.5 °C with exhausts of 100, 100, 95, 85, 95 and 100 %, which allows 51.88 m/min at 1.10 MJ/m^2, a hot first zone, a mild window and a hot last zone. The search is local and runs offline, and the learner applies its result as a profile. The tests run it again and require the same answer. In the infeasibility sweep one of eight starts stopped at 34.3 m/min.

6.5.6 Model Predictive Control

A model predictive controller (MPC) is offered as an option for the speed and the zones. The push, with the zones on their profile, is what the application loads and what every case except case 9 uses, and case 9 pins the MPC. Every $T_s = 10$ s, or every fifth control period, one quadratic program sets the line speed request and the six zone air setpoints together over $N = 60$ steps, or 10 min. The horizon is cut into seven move blocks whose lengths are 1, 1, 2, 2, 4, 6 and 44 steps. The plan is applied in receding horizon as Section 1.6.3 describes [8, 19], and the first plan is made at $t = 0$, so the historian records the MPC acting from the first row. Everything below the plan runs as shipped. The zone loops receive $T_{p,z} + \delta_{T,z} - c_z$, the profile's air plus the plan's trim less any heat cut c_z the override has made, the exhaust dampers keep their trim and the coat-weight loop stays geared to the speed. The plan's speed takes the push's place in the override's low selector, $\min(v_{mpc}, v_{lim})$, before the clamp of 10 to 80 m/min, so the override, its heat cut, the exhaust trim and the clamp act above the plan exactly as they act above the push.

The controller moves speed v and six zone requests $T_{SP,z}$, with eight outputs: the residual s_{exit} and the migration index M at the gauge and the six monitors' $\%\mathrm{LEL}_z$. Its gains are the push's steady model, Equations (6.1) to (6.7) in steady state, computed from the declared values listed in Table 6.2 and not from the run's, so that a learner's edit of the plant is model mismatch for the MPC as it is for the push. At each decision the model is relinearized at the operating point the controller reads, the applied speed and setpoints, the exhaust each zone can get as the push reckons it, the gauge's coat weight and the recipe's solids, or the densitometer's with the feedforward on, by forward differences of 0.5 m/min and 1 K, eight steady marches in all. A zone whose heater is at its rating is linearized at its thermocouple when that reads under the setpoint, and the plan may not raise that zone's setpoint. The dynamics of each pair are a first-order lag with a dead time, Equation (6.13), in which x_z is the distance from the die to the middle of zone z, the speed v is in m/s and the transport parts scale with it.

$$\begin{aligned}
v \to s_{exit}, M:&\ \theta = 0.185\,L/v, \quad \tau = 0.4\,L/v \\
v \to \%\mathrm{LEL}_z:&\ \theta = \max(8,\, 0.38\,x_z/v) \\
&\ \tau = \max(10,\, 0.45\,x_z/v) \\
T_{SP,z} \to \%\mathrm{LEL}_z:&\ \theta = 10, \quad \tau = 16 \\
T_{SP,j} \to \%\mathrm{LEL}_{z>j}:&\ \theta = 8 + (z-j)L_z/v, \quad \tau = 15 \\
T_{SP,j} \to s_{exit}, M:&\ \theta = (L - x_j)/v - 4, \quad \tau = 9
\end{aligned} \tag{6.13}$$

Times are in seconds and $L_z = 9\,\text{m}$. The shapes come from steps of 1 m/min and 2 K, both ways, on the application in Auto at 50 m/min on the Optimized profile, override off and gauge noise zero, against an unstepped twin on the same seed. An independent refit of the same kind returns 10 s and 16 s for zone 3 on its own vapor and 62 and 14.5 s for zone 1 on zone 6, with residuals of 1.6 to 1.7 % of the change, and 42 and 9.5 s for zone 3 on the residual against the shape's 37.4 and 9 s. Its gains agree with the linearization within 2 % for every vapor pair and every zone's effect on the residual, 2.21 against 2.17 points of LEL per m/min from the speed to zone 6, and within 6 % for the speed's effect on the residual. The speed's effect on the residual, a ramp over the web's transit, fits with 30 % RMS error and the migration index with 25 to 72 %, so the index borrows the residual's shapes. Each pair's step response, sampled at T_s, becomes impulse weights on a delay line of the inputs applied in each period, at most 12 periods long, which gives up to 77 states. At steady state the model equals the linearization.

The decisions are the moves of the seven handles from the inputs applied now, constant within each block. Each plan minimizes Equation (6.14), in which $\hat{y}_{i,k}$ is a predicted output, j counts the blocks, $T_{p,z}$ is the learner's profile and Δu_0 is measured from the input applied now.

$$
\begin{aligned}
\mathscr{J} = &- \sum_{k=1}^{N} \frac{v_k}{\text{m/min}} + \rho_T \sum_{k=1}^{N} \sum_{z=1}^{6} \left(T_{SP,z,k} - T_{p,z}\right)^2 \\
&+ \lambda \sum_{j=0}^{6} \left[\left(\frac{\Delta v_j}{0.5\,\text{m/min}} \right)^2 + \sum_{z=1}^{6} \left(\frac{\Delta T_{SP,z,j}}{1\,\text{K}} \right)^2 \right] \\
&+ \sum_{i=1}^{8} \left(\nu_i \varepsilon_i + \kappa_i \varepsilon_i^2 \right) \\
\text{s.t. } & \hat{y}_{i,k} \le \bar{y}_{i,k} + \varepsilon_i, \quad \varepsilon_i \ge 0 \\
& 20\,\text{m/min} \le v_j \le 80\,\text{m/min}, \quad |\Delta v_j| \le r_{push} T_s \\
& |T_{SP,z,j} - T_{p,z}| \le \Delta T_{max}, \quad |\Delta T_{SP,z,j}| \le 2.5\,\text{K}
\end{aligned}
\tag{6.14}
$$

Speed earns one unit per m/min per step, and $\lambda = 10$ prices moves. The pull $\rho_T = 0.002$ per K^2 per step keeps a zone near its profile when nothing else decides it. The limits $\bar{y}_{i,k}$ are soft and judged on every second step, with one slack per output, eight in all. They are the push's own target for each zone's vapor, 20 % LEL, the residual limit $s_{max} = 0.5\,\text{g/m}^2$ and the migration threshold $M_{max} = 2\,\text{g/m}^2$ in the default setting, At the limits. The alternative setting, Push margins, moves the last two to the push's $0.8\,s_{max}$ and $0.9\,M_{max}$ of Equation (6.11). The slack prices are $\nu = 2 \times 10^3$ and $\kappa = 2 \times 10^4$ per point of LEL, 2×10^5 and 2×10^6 per g/m^2 of residual and 2×10^4 and 2×10^5 per g/m^2 of migration, each far above what a m/min earns over the horizon. A limit that the free response already crosses is re-entered along a straight line over 60 s. The handle limits are hard. The speed stays in the push's range and moves no faster than the push rate, at most $r_{push} T_s = 0.33\,\text{m/min}$ from one block to the next, with r_{push} read as the push reads it, so neither structure can change the speed faster than the other. Each zone may move $\Delta T_{max} = 10\,\text{K}$ from the profile, and at most 2.5 K from one block to the next, within the setpoint range of 70 to 170 °C. Heater energy is not priced. The move price was chosen once, before any comparison, on case 9 with seeds 1 to 3 over 60 min. Of λ from 0.1 to 30, all six values moved the speed request no more than the push did (98.5 m/min), and the rule, the largest λ within 1 % of the best throughput of 3693 m^2/h, chose $\lambda = 10$ at 3688 m^2/h. Nothing was retuned afterwards for any case.

Each zone's vapor carries the push's bias b_z of Equation (6.11), the monitor's reading over the declared model's value at the measured speed, air and flows, filtered over 120 s and clipped between 0.5 and 2. This output-disturbance estimate makes the vapor limits offset-free on what the monitors read [20]. The residual and the migration index have no on-line measurement, so the plan predicts them from the declared model and holds the prediction at its limit, as the push holds its own prediction under a margin. No preview of the solids drift, the fan loss or any other event reaches either structure.

The program has 57 variables (49 input values and 8 slacks) and 444 rows, past the 30 states and 6 inputs for which the shared solver class is documented, and a dual active-set solver [21] handles it. A plan is discarded if the solver returns it as anything but optimal, if a value in it is infinite or undefined, or if it breaks a bound. The push, or the setpoint ramp when the push is off, then takes the line from the applied speed, which it has tracked as external reset, while each zone’s trim returns to the profile at the override’s heat-back rate, 1 °C/min. A switch from the MPC to the push leaves the first period unmoved and then moves the speed request 0.017 m/min and the setpoints 0.033 K per period. Over 8786 plans, on every case as loaded, on four more seeds of case 9, on the challenge with its suggested settings and at one extreme setting, none was rejected. At the shipped settings the Hessian’s condition number was 4.28×10^{6}, a plan took at most 124 iterations and the stationarity residual of the optimality conditions was at most 3×10^{-7} of the largest linear term. The MPC reads what the push reads, the instruments, the applied inputs, the limits and the declared values. Replay tests in which two plants differ in their nozzle coefficient, heater rating, critical content, safe rate, gauge distance, solids and pump slip, with the instruments of the first copied into the second, find identical plans, and eight mutated copies of the controller that each read one truth field or run parameter all fail the application’s tests. Over the 540 plans of case 9 in Chrome a plan takes 11.9 ms on average and 16.5 ms at the most, 0.12 % of one core at real time and 36 % at the top speed of 300 times real time, where the page draws 59.6 frames per second against the push’s 59.9.

6.5.7 Three Models

Section 6.3 is the simulation model. The application identifies no control model except the delay-line shapes of the MPC, and the coat-weight tuning ships with the cases. The controller’s internal models are the predictor’s coat weight, $h\hat{\rho}\hat{x}$ through a lag, and the steady profile model of the push, the override’s handle choice and the optimizer, which shares the plant’s equations. The push and the override’s handle choice evaluate it with the declared parameters and meet model error only through the monitors’ bias. The MPC’s internal model is the same steady model evaluated with the declared parameters and joined to the identified shapes, so a learner’s edit of a parameter reaches neither the push nor the MPC.

6.6 Code Verification and Agreement With Sources

Code verification asks whether the program solves the equations above. The three ledgers of Equation (6.6) close over a disturbed run with speed steps, the solids drift, the fan loss and a heater fault, to residuals of 5.3×10^{-12}, 3.1×10^{-13} and 5.6×10^{-14} kg against 115.432 kg of solvent fed. Independent quadratures over the challenge’s suggested run of 5000 m give 495.387 kg pumped against the ledger’s 495.358 and 491.497 kg exhausted against 491.488. A separate explicit integration of the steady march with 2 mm steps, written without the application’s code, gives 18.24, 20.73, 21.09 and 18.61 % LEL in zones 1 to 4 of the Standard profile at 35 m/min against the application’s 18.27, 20.71, 21.09 and 18.64, and at 40 m/min the two agree to 0.05 points. Equation (6.1) holds exactly for the coating laid at the die and for the coating weighed at the gauge, and the gauge’s delay equals L/v within one logged row (104.0 against 102.9 s at 35 m/min). Halving the step to 0.25 s keeps every challenge verdict and moves the suggested throughput from 3076 to 3073 m²/h and the mean residual by 0.75 %, and halving the parcel to 0.25 m leaves the zone vapor unchanged to 0.01 points and moves the Standard profile’s migration index from 6.13 to 6.23 g/m². Loaded fresh, reached through the case menu from a different case, or restarted with Reset, each case logs the same run bit for bit.

Agreement with sources is a separate and weaker claim. The model’s constant drying rates, 1.17 and 1.36 g/(m² s) in the wet zones at 105 and 110 °C, lie near the 1.19 g/(m² s) of the experiments the binder model reproduces, and the 115.7 µm wet film is close to its 114 µm [2]. The vapor pressure’s slope, 50.7 kJ/mol, lies among the WebBook’s enthalpies of vaporization of 47.7 to 53.4 kJ/mol between 345 and 425 K, and at 25 °C, outside the two fitted points, the line gives 50 Pa against the safety card’s 39 Pa [17, 18]. No readable source gives a critical content or a falling-rate law for an

Table 6.3 The Nine Cases and the Challenge

#	Settings and events	Teaching point
1	Manual pump and speed, 35 m/min	coat weight is a ratio
2	Auto, $K_c = 0.6$, target 180 to 185 g/m^2 at 60 s, override on	delay limits gain, the predictor
3	Auto at 35 m/min, override off, raise to 40	the front's zone binds
4	Standard against Gentle, pause at 60 min	migration, found late
5	Optimized at 50 m/min, zone 6 cooled	residual at the exit
6	Solids drift of 2 points from 2 min	feedforward from density
7	Fan to 70 % from 5 min, override off	choosing the handle
8	Push and override on, Standard profile	constraint pushing
9	Optimized at 50 m/min, push and override on, MPC on speed and zones, drift at 10 min	speed and zones in one plan
Ch.	Challenge, 5000 m: drift at 20, fan at 40, splice at 60 min	the whole problem

Table 6.4 Measured Results of the Cases

#	Run	Result
1	30 m/min, pump held	gauge 209.7 g/m^2, first moves 122.5 s after
2	predictor off	gauge 122.5–271.0 g/m^2, swings 30 min
	predictor on	within 1.5 % from 3.0 min
	off, K_c = 0.5 / 0.4 / 0.3	settles at 26.7 / 6.0 / 3.0 min
3	override off, 40 m/min	zone 4 30.3 % LEL, over 25 at 93 s
	override on	peak 22.9 % LEL, mean 35.6 m/min
4	Standard / Gentle	first peel 8.0 / 23.3 N/m
	Optimized at 35 m/min	13.7 N/m
5	zone 6 147.5 to 125 °C	residual 0.327 to 0.630 g/m^2
6	feedforward off / on	worst weight 1.59 / 0.49 %
7	override off	zone 3 25.8 % LEL, 1491 s over
	override on	worst 24.8 % LEL, mean 27.6 m/min
8	Standard / Gentle / Optimized	36.4 / 37.2 / 51.2 m/min
	Push against MPC, means of five seeds over 90 min	
5	push on, zone 6 to 125 °C at 10 min	3419 / 3815 m^2/h, residual over 103 / 1271 s
6	push on	2405 / 2308 m^2/h, strips over 15 / 13 of 15
8	as loaded	2603 / 2484 m^2/h, strips over 15 / 7 of 15
9	MPC at the limits	3364 / 3658 m^2/h, strips over 0 / 8 of 15
	MPC, push margins	3364 / 3639 m^2/h, strips over 0 / 0 of 15

NMP cathode, so s_c and n remain teaching values, and no peel strength in this chapter is a prediction for a real electrode.

6.7 Nine Cases and a Challenge, Measured

The cases are listed in Table 6.3. All values in this section were taken without a browser, stepping the engine through the same calls the page makes (Section 1.6.1), on the case's default seed unless five seeds are named, with only the change a sentence names, and the scripts that produce them are kept with the application's records. The headline results are gathered in Table 6.4. Each case has a completion rule of the kind Section 1.6.4 describes, and the exported file places each instrument's column beside a prefixed column that holds the true value.

6.7.1 A Ratio and a Late Gauge

Case 1 runs the pump and the speed by hand. Slowing the line from 35 to 30 m/min with the pump held makes the die lay more than 185 g/m^2 within 2.0 s, but the gauge rises 2 % only 122.5 s later, because the web needs 120 s to carry it 60 m. It settles at 209.7 g/m^2, the 180 times 35/30 of the ratio with the residual added. Matching the pump from 4.86 to 4.16 L/min brings the gauge back within 1.5 % 126.0 s later. The worst zone stays near 21 % LEL throughout.

Case 2 gives the loop twice the shipped gain, $K_c = 0.6$, and steps the target to 185 g/m^2 at 60 s. With the predictor off the loop makes about 50 corrections before the gauge sees the first, and the

gauge swings between 122.5 and 271.0 g/m^2 for the whole half hour. Its heavy half-swings load the dryer, the override slows the line to 20 m/min in 73 % of the steps, the delay stretches from 103 to 180 s, and zone 4 spends 46 s over the cap. With the predictor on, the same gain holds the gauge between 179.0 and 186.2 g/m^2 and within 1.5 % of the target from 3.0 min, while the override takes the line no lower than 32.8 m/min. Without the predictor the gain must come down. At 0.5 the loop swings for 26.7 min, and with the shipped integral time of 90 s a gain of 0.4 settles at 6.0 min and the shipped 0.3 at 3.0 min, as fast as the predictor, because the transit time is the floor for both. The predictor buys gain, not time.

Pause. Try Exercise 6.3 before reading on.

6.7.2 The Constraint at the Front

Case 3 raises the setpoint from 35 to 40 m/min at 30 s with the override off. Zone 1 falls from 18.25 to 16.5 % LEL, zones 2 and 3 stay near 21, and zone 4 climbs from 18.57 to 29.7 % LEL, crossing 25 at 93 s and peaking at 30.3 % LEL. With the override on from the start, the exhaust trim opens first, and when zone 4 reaches the trip with its exhaust at capacity, at 77 s, the speed limit takes over. Zone 4 peaks at 22.9 % LEL and never reaches the cap, and the line hunts between 33.6 and 37.6 m/min from 10 to 30 min, a mean of 35.6. The hunt has a cause. Zone 4 gains about 2 points of LEL per m/min with the front in it, and a speed change reaches it only as the web carries it there, about a minute later. Over 81 moments of the change from 20 to 100 s the mean speed was 35.4 to 35.9 m/min and no zone crossed the cap.

> **TRY IT IN THE APP 3 · Raise the Speed: A Zone Crosses Its Cap**
>
> Load the scenario and press Run on the Zones view with the LEL override Off. At 30 s set Line speed setpoint to 40 m/min and watch the zone bars against the 25 % LEL line. Zones 2 and 3 stay near 21 while zone 4, which the drying front enters, climbs from 18.57 to 29.7 % LEL and crosses the cap at 93 s.

Case 7 weakens the fan to 70 % over ten minutes at the same speed. With the override off, zones 2 and 3, wet from end to end, rise with the falling exhaust, from 1.53 to 1.26 kg/s, and zone 3 crosses the cap at 14.4 min and stays over for 1491 s. A slower line would not help these zones. With the override on, no zone crosses the cap (worst 24.8 % LEL, zone 4), the dampers being wide open already, so the override cuts the air of zones 1 to 3 by 9.4, 5.5 and 7.5 K on average between 20 and 40 min and limits the speed for zone 4, which cycles eight times between 24.0 and 30.6 m/min, a mean of 27.6, about a fifth less output. A gentler cut of 10 m/min per minute lets zone 4 reach 25.5 % LEL for 17 s, and one of eight seeds sees 25.3 % LEL for 11 s even at the shipped rate, during the first cut while the fan is still weakening.

> **TRY IT IN THE APP 7 · Fan Degradation: The Caps Creep Up**
>
> Run the scenario with the LEL override Off: at the same 35 m/min zone 3 crosses the cap at 14.4 min and stays over for 1491 s. Set LEL override to On, press Reset and run again. No zone crosses the cap: the override cuts the air of zones 1 to 3 and limits the speed for zone 4, and the line averages 27.6 m/min.

Pause. Try Exercise 6.2 before reading on.

6.7.3 Migration, Residual and Solids

Case 4 runs the Standard and Gentle profiles at 35 m/min to the first peel test at 60 min. The strip cut at 30 min holds 8.0 N/m on the Standard profile, a migration index of 6.20 g/m^2, and 23.3 N/m on the Gentle one, an index of 0.66 g/m^2, against a pass at 15 N/m. The gauge reads 179.3 g/m^2 on both, the residual is 0.067 and 0.180 g/m^2, and the worst zone over the hour is 21.5 and 20.2 % LEL, so nothing on the line tells them apart, as shown in Figure 6.2. The Optimized profile tests 13.7 N/m at this speed, because it was found for 52 m/min, where its window falls in its mild zones.

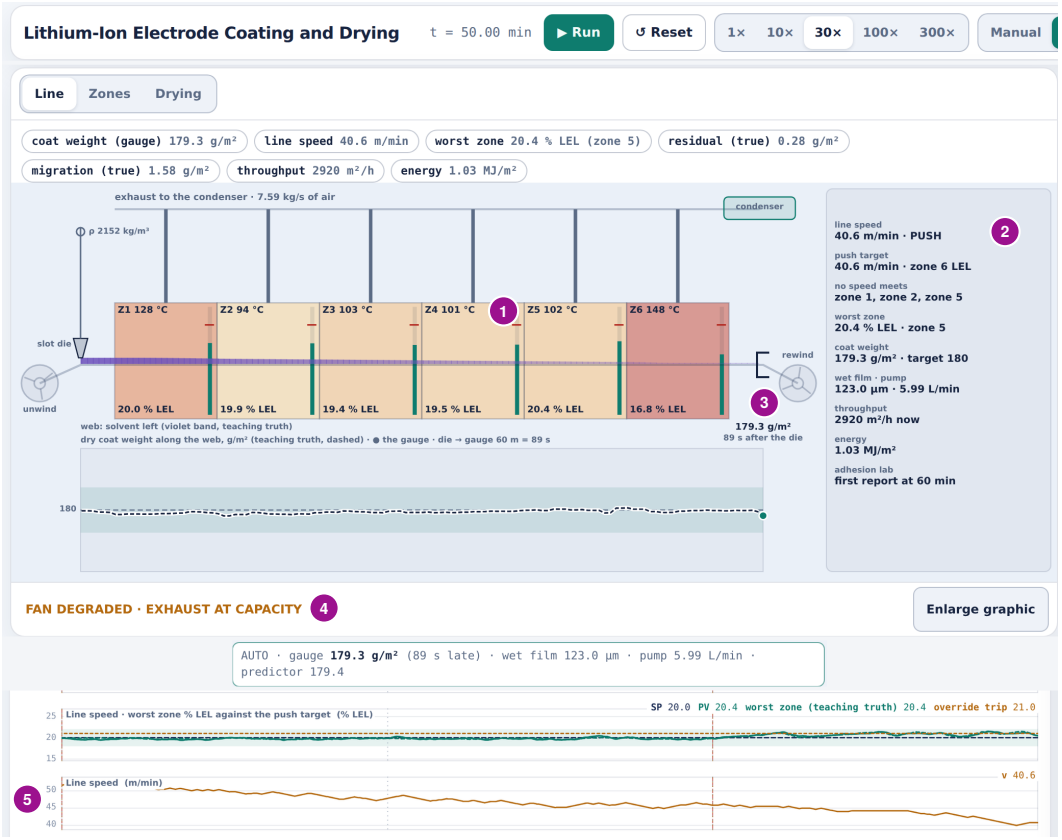

Figure 6.3 Line view of the suggested challenge run at $t = 50$ min, a composite of the application's window: the header, the Line view, the status line of the coat-weight card and, below, the trend strips of the worst zone and the line speed. The drift has thinned the slurry to 68 % solids, which the densitometer reads as 68.02 (a value this view does not print), and the fan has just reached 70 % of its nameplate, so every damper is at capacity and the status strip (4) reads FAN DEGRADED. The override has cut the air of zones 1, 2 and 5 by 10.4, 11.5 and 3.5 K, which the zone labels show only as air temperatures, and the push, leaving those zones out, holds 40.6 m/min against zone 6 (2), with the monitors (1) reading 20.0, 19.9, 19.4, 19.5, 20.4 and 16.8 % LEL under the 25 % LEL cap. The gauge (3) reads 179.3 g/m^2, 89 s after the die, with a wet film of 123.0 µm and the predictor's measurement at 179.4 on the card's status line, and the dashed coat weight along the web stays on its 1.5 % band. Over the 30-minute window of the trends (5) the line speed falls from 51.4 m/min at 20 min to 40.6, touching 39.8 (values from the run's log, of which the trend prints only the last), as the fan weakens, and the worst zone stays at the push's 20 % LEL target under the override's 21. The first laboratory report is due at 60 min.

Case 5 runs the Optimized profile at 50 m/min and cools zone 6. At 147.5 °C the web leaves with 0.327 g/m^2 at 142.5 °C. With zone 6 at 125 °C from 600 s, the residual passes its 0.5 g/m^2 limit 21 s later and reaches 0.630 g/m^2, and the exit web reads 119.6 °C, the one instrument that moves with the residual.

Case 6 thins the slurry from 70 to 68 % solids over 15 min. The die meters volume, so the wet film stays at 115.6 µm while the dry weight falls. With the feedforward off the worst error is 1.59 % at 16.9 min, and the loop then raises the film to 122.9 µm. With the densitometer's estimate (68.03 against 68.0 %) scaling the ratio, the worst error is 0.49 %. The thinner slurry carries 9.8 % more

Table 6.5 Case 9 Over 90 min, Means of Five Seeds. Strips Over: Laboratory Strips Whose True Index Exceeds M_{max}, of 15. Trip: Seconds With a Zone Over 21 % LEL. Ovr.: Override Engagements. Travel: Speed Request Travel

Structure	m²/h	Strips over	s_{exit} worst	Trip (s)	Ovr.	MJ/m²	Travel (m/min)
Push (shipped default)	3364	0	0.387	0	0.0	1.218	150
MPC, push margins	3639	0	0.436	54	3.6	1.145	58
MPC, at the limits (default)	3658	8	0.497	52	3.8	1.139	57
MPC, zones held, push margins	3363	0	0.409	0	0.0	1.219	49
Push at the limits (not shipped)	3465	1	0.423	14	1.0	1.185	131
MPC, delay left out	3660	6	0.469	213	13.4	1.141	54
Fixed 48 m/min, best in hindsight	3430	4	0.392	133	7.2	1.196	26

solvent per square meter, so the override slows the line to 30.8 m/min without the feedforward and 30.3 m/min with it.

6.7.4 The Push

Case 8 turns on the push with the Standard profile. Between 30 and 60 min it holds a mean of 36.4 m/min, between 33.4 and 37.7, bound by zone 4, and it reports that no speed meets the migration limit on this profile (an index of 7.05 g/m²). The wander is the push's 20 % LEL target and the override's 21 % LEL trip handing the speed back and forth, because at a gain of about 2 points per m/min the two are only half a meter per minute apart, and the override engages ten times in the half hour. A push target of 19 % LEL narrows the wander to 35.2 to 37.1 m/min at a mean of 36.3 with four engagements, and 18 costs six meters per minute (a mean of 30.0 m/min). The Gentle profile allows 37.2 m/min, bound by the migration index, and the Optimized profile 51.2 m/min, between 49.6 and 52.1, bound by the migration index at 90 % of its threshold with a residual of 0.363 g/m².

6.7.5 One Plan for the Speed and the Zones

Case 9, named Speed and zones in one plan, starts at 50 m/min on the Optimized profile with the push, the override, the predictor and the solids feedforward on, and with the MPC of Section 6.5.6 setting the speed and the zones. At 10 min the slurry's solids begin to fall by 2 points over 15 min. Each row of Table 6.5 changes only the structure or one of its settings, over 90 min on the case's own seed and seeds 1 to 4, with the same instruments, noise and events and no preview, and is scored on the engine's truth.

Like for like, with the push's own margins of $0.8\,s_{max}$ and $0.9\,M_{max}$, the MPC coats 3639 m²/h (3633 to 3643 over the seeds) against the push's 3364 m²/h (3362 to 3367), 8.2 % more at every seed, with none of 15 strips over the threshold, 6 % less heater energy per square meter (1.145 against 1.218 MJ/m²) and 61 % less travel of the speed request (58 against 150 m/min). The whole gain is the zones. Held on their profile at the same margins, the plan makes 3363 m²/h, the push's throughput, so the dynamic model and the delay lines add nothing to the area. What the plan does with the zones is visible in its mean trims after 30 min on the case's seed, about +10 K on zones 1 and 3 and −1 to −5 K on zones 2, 4, 5 and 6, a hotter start and a milder window that the offline search of Section 6.5 could not choose for a slurry that thins while the line runs.

The default setting, At the limits, makes 3658 m²/h, so riding the limits adds 19 m²/h, 0.5 %, for 8 of 15 strips over the threshold. Held zones at the limits make 3474 m²/h with 1 of 15 strips over, and a diagnosis copy of the push at the limits, which is not part of the application, makes 3465 m²/h, so riding the limits is worth about 3 % to either structure when only the speed moves. The transport delay in the model is worth fewer trips rather than more area. Without it the plan makes 3660 m²/h but engages the override 13.4 times instead of 3.8 and spends 213 s over the trip instead of 52 s. The best of 126 shipped settings in hindsight, 48 m/min fixed with the push off, makes 3430 m²/h and passes every strip at 2 of 5 seeds.

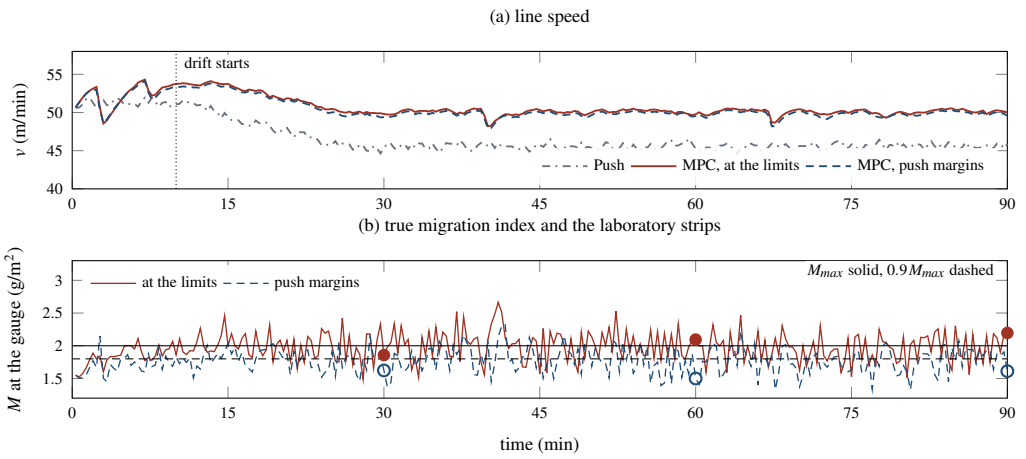

Figure 6.4 Case 9 on its own seed over 90 min, sampled every 20 s. (a) The true line speed under the push (dash-dotted) and under the MPC at its two settings (At the limits solid, Push margins dashed), after the solids drift begins at 10 min. The MPC runs about 4 m/min faster on average. (b) The true migration index at the gauge under the MPC, with the strips the laboratory cut at 30, 60 and 90 min marked (filled circles at the limits, open circles at Push margins). At the limits the index is held near M_{max} and two of three strips test over it, while Push margins holds it near $0.9\,M_{max}$ and every strip passes.

At either setting zone 6 spends about a minute in 90 min over the trip (52 and 54 s, the push never), about half of it on three seeds while the plan accelerates in the first 10 min. At the end of its horizon the plan sits within 2 % of the zone 2 vapor limit in 100 % of its decisions, of zone 6's in 99.4 % and of the migration threshold in 99.1 %, and never near the residual limit. The residual does not bind in case 9. With its limit left out of the plan the line makes 3657 m^2/h, as with it, and the residual passes the limit for under 2 s at one seed.

The failed strips are scatter, not bias. After 40 min the plan's steady prediction of the index at the gauge averages 2.03 g/m^2 and the true index 1.98 g/m^2, so the prediction errs high, while the true index scatters 0.22 g/m^2 (one standard deviation) from parcel to parcel. A plan holding the mean at the threshold fails about half the strips, as shown in Figure 6.4. Slack prices 10 and 100 times higher leave 7 and 4 of 15 strips over, the second with the residual 111 s over its limit at every seed, and the migration slack never exceeds 0.13 g/m^2, so neither the price nor the single slack per output is the cause. A back-off to 95 % of the threshold leaves 3 of 15 and to 90 %, the push's margin, none. It is reported as information, and the default was not retuned.

The other cases of Table 6.4 bound the result. In case 5 with its zone 6 step the residual binds, and the MPC, 11.6 % faster, is over its limit for 1271 s per 90 min at every seed against the push's 103 s (39 s at Push margins, 3681 m^2/h). On the Standard profile, where no setting meets the migration limit, the MPC loses 4.0 and 4.6 % at every seed in cases 6 and 8 while failing fewer strips. The push leaves the unmet limit out, and the plan chases it with a prediction that is low, 1.46 against 2.05 g/m^2 in case 8.

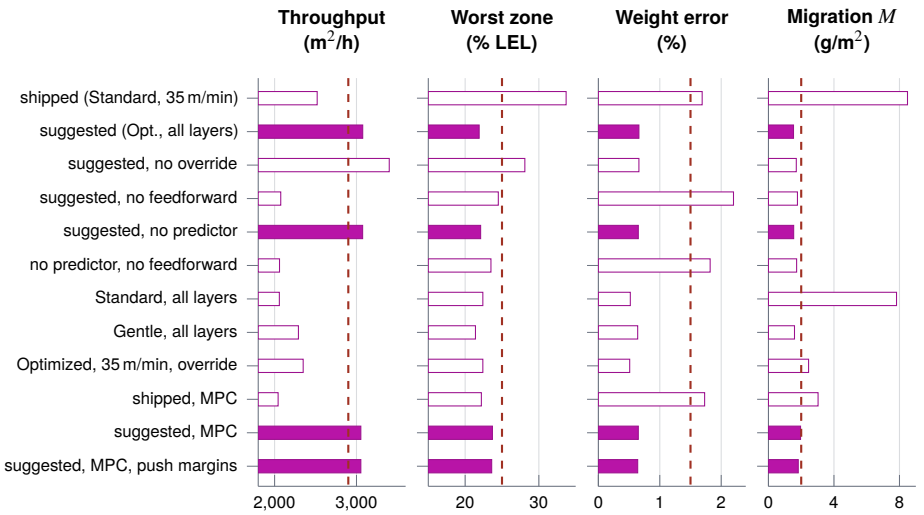

Figure 6.5 Table 6.6 redrawn against the challenge's limits, dashed: at least 2900 m^2/h, every zone at or under 25 % LEL, the dry coat weight within 1.5 % and the migration under its threshold of 2 g/m^2. Solid bars pass and white bars fail. Notice that the run without the override is the fastest and crosses only the dashed limit of the cap, and that the Gentle profile with every layer stays inside every dashed limit except throughput.

> ✦ DID YOU KNOW? **Riding the Limit Buys Half a Percent**
>
> From the simulator: in case 9 the predictive controller at the limits coats 3658 m^2/h, only 19 more than at the push's margins, yet 8 of 15 laboratory strips fail against none, because the index it holds at the threshold scatters from parcel to parcel (Table 6.5).

6.7.6 The Challenge

The challenge coats 5000 m through a solids drift at 20 min, a fan loss at 40 min and a splice at 60 min, and requires every zone at or under 25 % LEL at every instant, the dry coat weight within 1.5 % of the target after the first minute, the residual under its limit, every laboratory sample's true migration under its threshold and at least 2900 m^2/h, judged on the truth with the case's plant as shipped. The configurations are compared in Table 6.6, and Figure 6.5 draws them against the challenge's limits. As shipped, at a fixed 35 m/min on the Standard profile with the push and the override off, the line makes 2520 m^2/h and fails four ways. The suggested settings make 3076 m^2/h, 22 % more, and pass, and on five more seeds they make 3069 to 3084 m^2/h and pass every time, while the shipped settings fail on all six. Each piece earns its place except one. Without the override the line runs faster, 3400 m^2/h, and crosses the cap when the fan weakens. Without the feedforward the drift takes the weight 2.2 % off and the throughput falls to 2075 m^2/h. Without the predictor the run passes at 3075 m^2/h, so the predictor is not what this challenge needs, and its value is the gain of case 2. The line at 3000 s of the suggested run is shown in Figure 6.3. At the fastest playback speed the challenge's 5000 m take 23.4 to 28.7 s of wall-clock time in a browser tab.

For information, with its rules unchanged and over five seeds, the challenge as shipped with the MPC added fails at every seed (2038 to 2044 m^2/h) but keeps the worst zone at 21.8 to 22.2 % LEL. The suggested settings with the MPC pass at 2 of 5 seeds (3053 to 3116 m^2/h, seed 2 also with a residual of 0.503 g/m^2) and at Push margins at all five (3053 to 3115, the push 3076 to 3084 m^2/h).

Table 6.6 The Challenge, 5000 m, Default Seed

Configuration	m^2/h	LEL (%)	Wt. (%)	M (g/m^2)	Verdict
shipped (Standard, 35 m/min)	2520	33.7	1.69	8.49	fails
suggested (Opt., all layers)	3076	21.9	0.66	1.53	passes
suggested, no override	3400	28.1	0.66	1.70	fails
suggested, no feedforward	2075	24.5	2.20	1.77	fails
suggested, no predictor	3075	22.1	0.65	1.54	passes
no predictor, no feedforward	2060	23.5	1.82	1.72	fails
Standard, all layers	2058	22.4	0.52	7.82	fails
Gentle, all layers	2291	21.4	0.64	1.59	fails
Optimized, 35 m/min, override	2349	22.4	0.51	2.45	fails
shipped, MPC	2043	22.2	1.73	3.03	fails
suggested, MPC	3053	23.7	0.65	1.96	passes
suggested, MPC, push margins	3053	23.6	0.64	1.82	passes

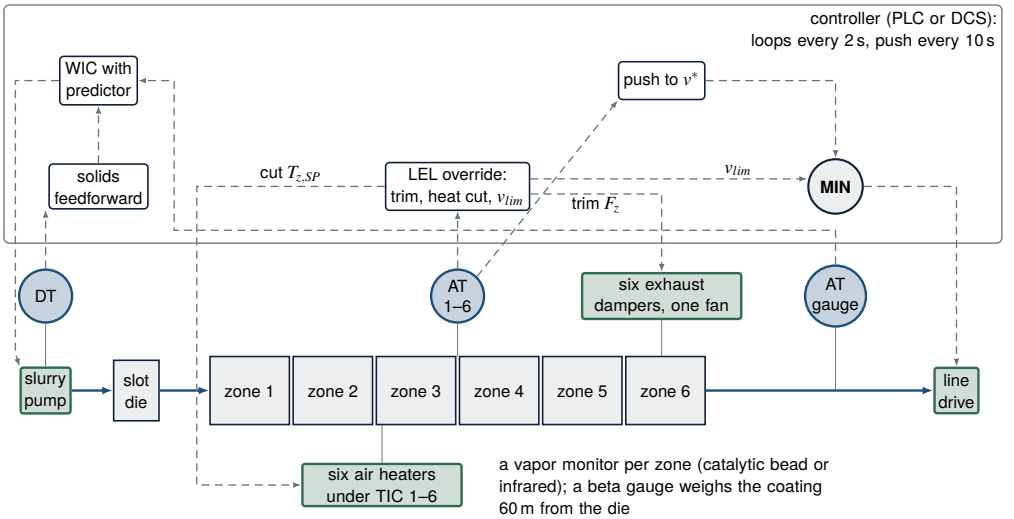

Figure 6.6 The override, the push and the coat-weight loop on real hardware. The vapor monitors AT 1–6 feed the LEL override, which trims the exhaust dampers first, then cuts a zone's air through its heater loop or limits the speed, and the push, whose speed meets the limit in the low selector ahead of the line drive. The beta gauge feeds the coat-weight loop WIC through its predictor, and the densitometer DT feeds the solids feedforward. Notice that every handle of the override is a different device: a damper, a heater and the drive.

6.8 Technology in Practice

A production coater reads the same signals and moves the same handles as the layers of Section 6.5. Figure 6.6 draws the override, the push and the coat-weight loop with line hardware in place of the schematic's symbols.

What Measures It

The constraint of this chapter is read by one vapor monitor per zone. A cost analysis of electrode drying names these instruments as a requirement of an NMP line and puts the limit they guard at a quarter of the lower explosive limit [11]. A national laboratory study mounted laser calipers at the oven inlet of a slot-die coater to read the wet thickness of cathodes and anodes, and found the precision of the in-line measurement to be within ±2 % [22]. The model's monitors read the true vapor through a lag, with no drift and no calibration error, as Limitations states.

TECHNOLOGY Sensor: A Vapor Monitor in Every Zone

Sensors for combustibles in air include the catalytic bead, the thermocouple, the infrared absorber and the flame ionization detector. Bead and thermocouple sensors burn a little of the sample on a heated element and read the heat it gives off. An infrared analyzer reads how much of a beam the vapor absorbs and must reject carbon dioxide and water vapor, which absorb too. Safety analyzers are built for fast response and ruggedness and are usually less accurate than process analyzers [23]. Oven monitors check themselves and recalibrate periodically on a test gas [13]. No open source read gives a range or a response time for an NMP oven monitor. The model's AT is an 8 s lag with 0.2 % LEL of noise and no drift, so a monitor reading low, which would let the push run faster, is not simulated.

What Moves It

The override has three handles. The exhaust trim opens each zone's damper, the heat cut lowers a zone's air setpoint through its heater loop, and the speed limit acts on the line drive through the low selector. A tip sheet on oven efficiency links the vapor monitor to the exhaust, through the fan's drive or its dampers [13], which is the first handle.

TECHNOLOGY Final Element: Exhaust Dampers and the Fan Drive

A damper throttles a large flow of air at low pressure with several vanes turned together, parallel vanes like a Venetian blind across a rectangular duct or radial vanes like petals in a round one [23]. An oven whose vapor monitor sets its ventilation can act through the dampers or, with further savings, through an adjustable-speed drive on the exhaust fan [13]. No open source read gives a stroke time or a rangeability for an oven damper. In the model each zone's exhaust follows its damper command through a 5 s lag, limited by the fan's capacity times its health, so a weakening fan looks on the trend as it does in case 7: every damper wide open and the vapor of the wet zones climbing.

Where the Controller Runs

Programmable logic controllers serve as local controllers within a supervisory scheme, loops are often nested so that the setpoint of one comes from another, and cycle times range from milliseconds to minutes [24]. In the application every loop and the override execute every 2 s, the push every 10 s, and the predictive controller every $T_s = 10\,\mathrm{s}$, so the push and the plan sit naturally in a supervisory layer above the zone loops. The low selector that applies $\min(v_{push}, v_{lim})$ is a standard function block. In one fieldbus standard the control selector block picks among controller outputs and carries back-calculation facilities that modify the response of the deselected controllers, which makes it the block for limit and override strategies [23].

What Is Modernizing It

Two developments bear on the dryer and its limits. The first is ventilation set by measurement rather than by the worst case. Continuous monitoring equipment tracks the solvent load in real time and controls the ventilation air to the real need, and linking it to an adjustable-speed drive on the exhaust fan improves the efficiency further [13]. The second is in-line inspection of the product. The same laboratory study read the homogeneity of the dried electrodes by infrared thermography at the oven outlet, where a rise or fall in the temperature profile showed flaws that optical images missed, and proposed such methods to find flawed electrode before cell assembly [22]. Neither measures the binder in the critical window, which the laboratory peel test still reports half an hour after the strip is cut. The predictive controller of Section 6.5.6 shows what a plan that moves the zones as well as the speed adds once such limits are predicted.

IN PRACTICE One Monitor per Zone, a Quarter of the Limit

A cost analysis of electrode drying from two national laboratories states the rule for NMP. The vapor in the drying chamber must be monitored to stay at or below 25 % of the lower explosive limit, 0.325 volume percent, and LEL monitors in each dryer zone read it in real time, with control logic to keep the operation safe. NMP is a combustible rather than a flammable liquid with a high boiling point, so its slurry mixers need no explosion-proof design, yet a line that coats with it must still be designed to prevent a fire or an explosion [11]. The override of this chapter is that control logic, and its trip at 21 % LEL sits under the cap.

6.9 Limitations

The line is a declared teaching design, not a fit to any coater, and the drying law is a constant rate and an activity law with no transport inside the film, so the residual and the falling-rate tail are plausible rather than predicted. The migration index is a proxy whose window and threshold are teaching values, and it rewards a hot first zone for a reason the sources do not give. The cap, the trip, the residual limit and the peel pass are teaching values. A real oven's limits come from the oven standard and the plant's own safety analysis, and a simulated pass is neither compliance nor a safety case. The zones are well mixed, the monitors read the true vapor through a lag without drift or calibration error, and one fan serves all six dampers.

The push, the override's handle choice and the optimizer share the plant's equations and its declared parameters, so model error reaches them only through a case that introduces it, and the monitors' bias then corrects only the vapor. The override hunts in cases 3, 7 and 8 because its speed handle acts through the web's transit, and the release rules that removed its two resting points were found by measurement on this model and are not proven for other profiles. The second side of the electrode, the calender and the cell that would use the declared coat weight are outside the simulator.

The MPC holds predictions of the residual and the migration index, which are never measured on line, at their limits, and the index's scatter of 0.22 g/m^2 is outside its model, so its default fails half the strips in case 9 and exceeds the residual limit for 1271 s in case 5 with its step. On the Standard profile its migration prediction is about 0.6 g/m^2 low. The push, the override's handle choice and the MPC all use the declared configuration with the measured temperatures, exhaust flows, gauge and recipe or measured solids. A parameter edit therefore introduces model mismatch for each structure. No shipped case edits a parameter, so the nominal comparisons are unchanged. With the critical content edited to 4 g/m^2 in case 9, the push exceeds the residual limit for 815 s at 3364 m^2/h, against 5233 s at 3672 m^2/h under the MPC. The MPC still loses this mismatch case. The override's stepped release when switched off remains the learner's act. Manual starts from the override's output under every structure. The program is larger than the shared solver class is documented for, which 8786 plans without a rejection support but do not prove.

OPEN CHALLENGE Taking the Solvent Out of the Cathode

NMP is toxic, and processing electrodes with water has become an expanding field of research. It has succeeded widely for anodes, but serious challenges remain for the cathode, and a 2024 study found water-based cathodes more sensitive to the drying temperature than NMP-based ones. It points to drying regimes tuned to that sensitivity as the way to capture the savings [25]. A line using water would also do without the LEL monitors and their control logic, and its exhaust could go to the atmosphere [11]. Progress would be a water-based cathode that holds its adhesion under a drying profile fast enough for production, which would trade this chapter's vapor cap for a narrower temperature window.

6.10 Conclusion

This chapter modeled a six-zone NMP cathode dryer with parcels of web that carry their own solvent, temperature and migration, and zones that carry their own air and vapor, with solvent ledgers that close to within 10^{-13} of the solvent fed.

On the Standard profile the wet zones stay near 21 % LEL at any speed while zone 4, where the drying front enters, climbs from 18.57 to 29.7 % LEL between 35 and 40 m/min, so the explosive-limit cap binds at the front. An override that trims the exhaust first and then chooses heat or speed zone by zone holds every zone under the cap through a fan loss, at the cost of a fifth of the output, once two release rules found in review remove its resting points. The same speed and coat weight tested 8.0 or 23.3 N/m depending on the air in the critical window.

With the push, the override, the feedforward and an optimized profile the challenge passes at 3076 m^2/h against 2520 m^2/h as shipped, while the predictor is not needed there. An optional MPC that also moves the zones coats 8.2 % more than the push at the push's margins and fails half the strips at its default limits.

The model is a declared teaching design, and its migration index is a proxy. A fuller model would add transport inside the film, a fitted falling-rate law, drifting monitors and the second side. Within those limits the simulator shows that a vapor limit is held at the drying front, and that the handle for it depends on where that front is.

KEY POINTS

- A faster line does not load every dryer zone a little more. It carries the wet coating into the next zone, so the vapor limit binds where the drying front is.
- An override that trims the exhaust first and then chooses heat or speed zone by zone holds every zone under the cap through a fan loss, at the cost of a fifth of the output.
- The same speed and coat weight tested 8.0 or 23.3 N/m depending on the air in the critical window.
- With the push, the override, the feedforward and an optimized profile the challenge passes at 3076 m^2/h against 2520 m^2/h as shipped.

Where else this applies. Chapter 5 runs a throughput supervisor against torque, pressure and dryer limits, and Volume 5, Chapter 10 pushes a dryer's throughput on inferred limits.

References

[1] M. Baunach, S. Jaiser, S. Schmelzle, H. Nirschl, P. Scharfer, and W. Schabel, "Delamination behavior of lithium-ion battery anodes: Influence of drying temperature during electrode processing," *Drying Technology*, vol. 34, no. 4, pp. 462–473, 2016.

[2] F. Font, B. Protas, G. Richardson, and J. M. Foster, "Binder migration during drying of lithium-ion battery electrodes: Modelling and comparison to experiment," *Journal of Power Sources*, vol. 393, pp. 177–185, 2018.

[3] D. Krishnamoorthy and S. Skogestad, "Systematic design of active constraint switching using selectors," *Computers & Chemical Engineering*, vol. 143, p. 107106, 2020.

[4] A. Marchetti, B. Chachuat, and D. Bonvin, "Real-time optimization of continuous processes via constraints adaptation," *IFAC Proceedings Volumes*, vol. 40, no. 5, pp. 45–50, 2007.

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

[6] R. Hooke and T. A. Jeeves, ""Direct Search" solution of numerical and statistical problems," *Journal of the ACM*, vol. 8, no. 2, pp. 212–229, 1961.

[7] T. H. Chilton and A. P. Colburn, "Mass transfer (absorption) coefficients prediction from data on heat transfer and fluid friction," *Industrial & Engineering Chemistry*, vol. 26, no. 11, pp. 1183–1187, 1934.

[8] 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.

[9] A. Yoshino, "The birth of the lithium-ion battery," *Angewandte Chemie International Edition*, vol. 51, no. 24, pp. 5798–5800, 2012.

[10] A. Kwade, W. Haselrieder, R. Leithoff, A. Modlinger, F. Dietrich, and K. Droeder, "Current status and challenges for automotive battery production technologies," *Nature Energy*, vol. 3, no. 4, pp. 290–300, 2018.

[11] D. L. Wood, J. D. Quass, J. Li, S. Ahmed, D. Ventola, and C. Daniel, "Technical and economic analysis of solvent-based lithium-ion electrode drying with water and NMP," *Drying Technology*, vol. 36, no. 2, pp. 234–244, 2018.

[12] N. Susarla, S. Ahmed, and D. W. Dees, "Modeling and analysis of solvent removal during Li-ion battery electrode drying," *Journal of Power Sources*, vol. 378, pp. 660–670, 2018.

[13] U.S. Department of Energy, "Use lower flammable limit monitoring equipment to improve process oven efficiency," National Renewable Energy Laboratory for the U.S. Department of Energy, Industrial Technologies Program, Golden, CO, Process Heating Tip Sheet #11 DOE/GO-102007-2511, Oct. 2007.

[14] U.S. Occupational Safety and Health Administration, "Occupational safety and health standards, 29 CFR 1910.94 (Ventilation) and 1910.107 (Spray finishing using flammable and combustible materials)," Electronic Code of Federal Regulations, https://www.ecfr.gov/current/title-29/part-1910, 2026, sections 1910.94(c)(6)(ii) and 1910.107(j)(1), read 2026-10-26.

[15] S. Jaiser, M. Müller, M. Baunach, W. Bauer, P. Scharfer, and W. Schabel, "Investigation of film solidification and binder migration during drying of Li-ion battery anodes," *Journal of Power Sources*, vol. 318, pp. 210–219, 2016.

[16] S. Jaiser, A. Friske, M. Baunach, P. Scharfer, and W. Schabel, "Development of a three-stage drying profile based on characteristic drying stages for lithium-ion battery anodes," *Drying Technology*, vol. 35, no. 10, pp. 1266–1275, 2017.

[17] National Institute of Standards and Technology, "2-pyrrolidinone, 1-methyl-: phase change data," NIST Chemistry WebBook, NIST Standard Reference Database 69, https://webbook.nist.gov/cgi/cbook.cgi?ID=C872504&Mask=4, 2026, read 2026-10-26.

[18] ILO and WHO, "International chemical safety card 0513: N-methyl-2-pyrrolidone," International Labour Organization and World Health Organization, https://chemicalsafety.ilo.org/dyn/icsc/showcard.display?p_card_id=0513, 2014, card dated April 2014, read 2026-10-26.

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

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

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

[22] D. Mohanty, J. Li, R. Born, L. C. Maxey, R. B. Dinwiddie, C. Daniel, and D. L. Wood, III, "Non-destructive evaluation of slot-die-coated lithium secondary battery electrodes by in-line laser caliper and IR thermography methods," *Analytical Methods*, vol. 6, no. 3, pp. 674–683, 2014.

[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] A. Kukay, G. Polizos, E. Bott, A. Ielvev, R. Tao, J. Sharma, and J. Li, "Mechanical and electrochemical implications of drying temperature on lithium-ion battery electrodes," *Batteries & Supercaps*, vol. 7, no. 7, p. e202400113, 2024.

6.11 Exercises

6.1 Concept check · 2 min

The coater runs at 35 m/min on the Standard profile with the LEL override off, and the line speed setpoint is raised to 40 m/min. Which zone's vapor rises past the 25 % LEL cap?

(a) Every zone a little, since 14 % more solvent enters the dryer each minute

(b) Zone 1, because the faster web arrives hotter and evaporates sooner

(c) Zones 2 and 3, because they evaporate the most solvent of the six zones

(d) Zone 4, because the faster web carries wet coating into the zone the front enters

6.2 Concept check · 2 min

In case 7 the exhaust fan loses 30 % of its capacity at a fixed 35 m/min, every damper is already wide open, and zones 2 and 3, wet from end to end, approach the cap. Which handle lowers their vapor?

(a) A cut in each zone's own air temperature, which lowers its constant drying rate
(b) A lower line speed, which reduces the solvent each zone receives per minute
(c) A wider damper opening, which restores the exhaust each zone receives
(d) A higher push target, which gives the zones more room before the trip

6.3 Concept check · 2 min

Case 2 doubles the coat-weight gain to 0.6 µm per g/m^2. Without the predictor the gauge swings for half an hour, and with it the gauge holds within 1.5 % from 3.0 min. With the predictor off and the shipped gain of 0.3, the loop also settles at 3.0 min. What does the predictor buy?

(a) It makes the gauge report each change sooner than the web can carry it
(b) It allows a higher gain without oscillation, but no faster settling than the detuned loop
(c) It replaces the solids feedforward, since it also corrects a thinning slurry at once
(d) It shortens the transport delay by keying the model to the encoder's web length

6.4 Theory and calculation · 6 min

The slurry has 70 % solids, a solids density of 4430 kg/m^3 and an NMP density of 1028 kg/m^3. The coated width is 1.2 m, the coat weight target is 180 g/m^2 and the gauge is 60 m of web from the die.

(a) Find the slurry density and the wet film thickness at the target. *Checkpoint:* 2223 kg/m^3 and 115.7 µm.

(b) At 35 m/min, find the solvent laid per square meter, the NMP the dryer must remove each second and the pump flow. *Checkpoint:* 77.1 g/m^2, 0.054 kg/s and 4.86 L/min.

(c) The line slows to 30 m/min with the pump held, as in case 1. Find the dry coat weight the die now lays and the time the gauge needs to see it, against the time at 35 m/min. *Checkpoint:* 210 g/m^2, after 120 s against 102.9 s.

6.5 Theory and calculation · 6 min

All solvent the die lays must leave through the exhaust. Use the molar masses 99.13 g/mol for NMP and 28.965 g/mol for air, the lower explosive limit $y_{LEL} = 0.013$, six fans of 1.8 kg/s of air and the solvent load $s_0 = 77.14\,\mathrm{g/m^2}$ on a coated width of 1.2 m.

(a) Find the largest NMP flow the six exhausts can carry with every zone at the 25 % LEL cap. *Checkpoint:* 0.120 kg/s.

(b) Find the line speed above which no profile can keep every zone under the cap. *Checkpoint:* 77.86 m/min.

(c) The fan loses 30 % of its capacity, as in case 7 and the challenge. Find the new bound. *Checkpoint:* 54.5 m/min.

6.6 Theory and calculation · 5 min

The vapor pressure of NMP in the model is the Clausius–Clapeyron line of Equation (6.4), with $P_0 = 101\,325$ Pa, $T_b = 475.2$ K and $B = 6093$ K.

(a) Find the vapor pressure at 120 and 140 °C. *Checkpoint:* 6.98 kPa and 14.8 kPa.

(b) From the slope of $\ln p_{sat}$ at 120 °C, find the temperature rise that doubles the vapor pressure there. *Checkpoint:* 17.6 K.

(c) Find the molar enthalpy of vaporization that B implies, and compare it with the range of 47.7 to 53.4 kJ/mol that the NIST WebBook gives between 345 and 425 K. *Checkpoint:* 50.7 kJ/mol, inside the range.

6.7 App challenge · 8 min

Load scenario 3, Raise the speed: a zone crosses its cap, at the default seed. The LEL override is off as loaded.

(a) Press Run at 10×. When the clock passes about 30 s (any time from 20 to 45 s will do), set Line speed setpoint to 40 m/min, switch to 100× and run to 30 min. Read the Worst zone (true) tile. Export the CSV file and find the time from the first row with `speed_req_mpm` above 35 to the first row with `truth_lel_z4_pct` above 25. *Checkpoint:* 30.3 % LEL in zone 4, within 0.1. About 62 s, within 4 s.

(b) Press the reload button, set LEL override to On, and repeat the same steps. Read the Worst zone (true) tile, and in a new CSV file find the mean of `speed_mpm` from 10 to 30 min. *Checkpoint:* 22.9 % LEL in zone 4, within 0.3, and 35.6 m/min, within 0.2.

Why does the line not settle at a fixed speed in run (b)?

6.8 App challenge · 8 min

Load scenario 4, The profile and the adhesion test, at the default seed. Run at 300× until the run pauses at 60 min. Read the peel test in the Residual · adhesion lab tile, export the CSV file and read `gauge_gm2` at 3600 s and `truth_migration_gm2` at 1800 s, the time the tested strip was cut. Then set Dryer profile to Gentle, press Reset and repeat.

(a) Report the four values for the Standard profile. *Checkpoint:* 8.0 N/m, 179.3 g/m^2 and 6.20 g/m^2, within 0.1.

(b) Report them for the Gentle profile. *Checkpoint:* 23.3 N/m, 179.3 g/m^2 and 0.66 g/m^2, within 0.1.

(c) Use the peel law of Equation (6.7) to find the true peel strength of each strip, and explain why the reports differ from it. *Checkpoint:* 7.32 N/m and 22.56 N/m.

6.9 App challenge · 9 min

Load the challenge, Coat 5000 m at the greatest throughput, at the default seed. For each run below, run at 300× until the run pauses at 5000 m and read the Challenge target tile: the throughput, the worst zone, the coat weight error, the worst migration and the verdict.

(a) Run the challenge as loaded (Standard profile at a fixed 35 m/min, push and override off). *Checkpoint:* 2520 m^2/h, 33.7 % LEL, 1.69 % and 8.49 g/m^2, not met.

(b) Press Reveal suggested tuning, then Reset, and run again. *Checkpoint:* 3076 m^2/h, 21.9 % LEL, 0.66 % and 1.53 g/m^2, met.

(c) Set Solids feedforward (inline density) to Off, press Reset and run again. *Checkpoint:* 2075 m^2/h, 24.5 % LEL, 2.20 % and 1.77 g/m^2, not met.

Name the targets that runs (a) and (c) miss, and explain why run (c) is also slower.

6.10 Application area · 4 min

At 35 m/min the dryer removes 0.054 kg/s of NMP. Each kilogram of air at the 25 % LEL cap carries 0.011 12 kg of NMP, and twice that at the 50 % of the limit that continuous monitors allow under the oven standard.

(a) Find the least total exhaust that could carry this solvent at each margin, and compare it with the design exhaust of 1.53 kg/s per zone. *Checkpoint:* 4.85 kg/s and 2.43 kg/s, against 9.18 kg/s.

(b) Give two reasons why a plant should not simply move to the larger margin to save the exhaust energy, using the chapter's limitations.

6.11 Application area · 4 min

The laboratory cuts a strip at the gauge every 30 min and reports its peel strength 30 min later. The line runs at 35 m/min on a coated width of 1.2 m, and the declared product carries 2.938 mAh/cm^2 per side.

(a) A profile change that harms the adhesion is made just after a strip is cut. Find the longest time before a report can show it and the electrode area coated in that time. *Checkpoint:* 60 min and 2520 m^2.

(b) Find the capacity of one side of that area, and name the one on-line instrument that moves with a quality the line cannot weigh. *Checkpoint:* About 74.0 kAh.

6.12 Review the agent · 5 min

An AI assistant reads Table 6.5 and writes: “At its default setting, At the limits, the MPC coats 3658 m^2/h against the push’s 3364 m^2/h, 8.7 % more. It keeps every zone under the 25 % LEL cap and the residual under its limit, and it uses 6 % less heater energy per square meter. Switch the line to the MPC at its default setting.” Find the flaw in the recommendation and say how you would check it.

CHAPTER 7

Horizontal Directional Drilling Where the Soil Chooses the Binding Limit

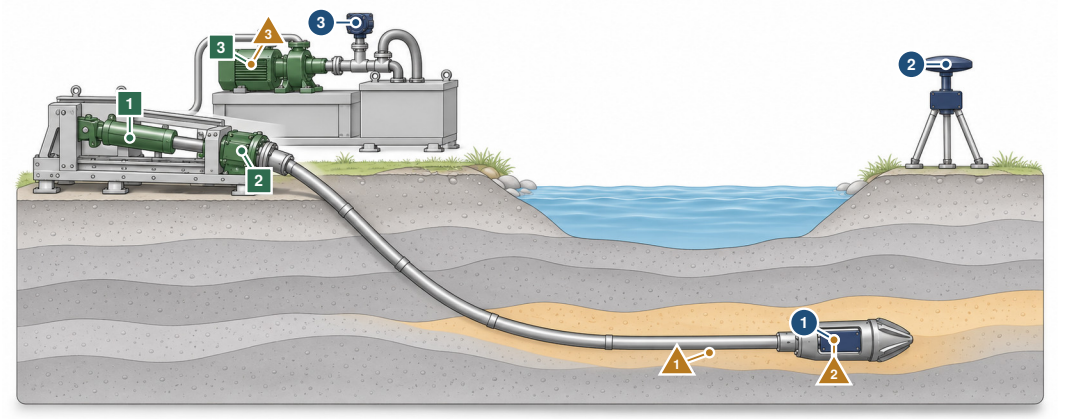

The controller moves 1 carriage thrust 2 spindle speed 3 mud pump flow

The controller reads 1 beacon behind the bit: pitch, roll, mud pressure 2 walkover locate of the depth 3 mud pump pressure

What disturbs it 1 a change of soil at the bit 2 a bias in the beacon's pitch 3 a mud pump fault

The drill on the bank pushes, turns and feeds mud to a bore it cannot see, and only the beacon at the head reads the mud pressure that the ground above must hold. Illustration, not to scale.

IN PLAIN TERMS **Drilling Under a River Without Breaking Out**

The main idea. A drill on one bank pushes a string of steel rods under a river and steers the cutting head by turning or pushing each rod. Mud pumped down the rods carries the cut soil back out, and the drill should cut as fast as the soil allows.

The challenge. Each kind of soil sets a different limit on that speed. The mud pressure grows with every meter, and where the ground above is thin it can split the soil and burst out at the surface.

Why it matters. A safety layer that eases off the push when the pressure nears that limit keeps the mud underground.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on bore sizes (Section 7.1), the two history segments (Section 7.2), the two try-it segments and the fact on the override (Section 7.5) and the key points
Study	55 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 7.5), Technology in Practice (Section 7.7), Limitations (Section 7.8) and the open challenge

The control question A steered pilot bore drills through soils that each hand the drilling rate to a different limit, while the mud pressure along the bore meets the thin ground of the exit climb. How do a rate optimizer, a pressure override and an estimator on sparse locates carry a river crossing to its exit within every limit?

Learning objectives

- Compute the turn of one pushed rod and explain why steering is a duty-cycle actuator that costs rate.
- Identify the constraint that holds the rate in clay, sand and gravel, and derive the mud-ratio ceiling on the rate.
- Predict where along a bore the annular pressure first meets the ground's limit, and what a pressure override takes away.
- Estimate the depth error a pitch bias builds between locates, and explain how an estimator that learns the bias carries the depth across a river.

Prerequisites Sections 1.4.2, 1.4.3, 1.4.4, 1.4.8 and 1.6.4, and Chapter 5 for a throughput supervisor on predicted constraints.

The app https://apmonitor.com/apps/horizontal-directional-drilling/

See also Volume 9, Chapter 8 for an optimum on a constraint that moves with the conditions, Volume 5, Chapter 13 for dead reckoning between sparse fixes with a Kalman filter, and Volume 10, Chapter 2 for a rate of penetration pushed against an annular pressure window and a cuttings load.

7.1 Introduction

A utility drill steers without a rudder. In the generic 30 m driveway service simulated here, a stiff clay takes the pilot bit down at a pitch of −14 %, and every 3 m rod is either rotated, which holds the pitch within 0.3 point, or pushed without rotation with its slanted face turned to a clock position. One rod pushed at 12 o'clock turns the bore from −14.2 to −4.7 % (9.4 points), but the bit cuts at 1.19 m/min pushed against 2.43 m/min rotating and the rod takes 206 s. Steering is therefore a duty-cycle actuator sampled once per rod, and every meter it costs is a meter drilled slower.

The same drill on a 180 m river crossing meets four further problems at once. The rate is limited by the spindle torque in clay, by the mud that must carry the cuttings out of sand, by the thrust in gravel and by the bus on a shore supply, so the binding constraint migrates along one bore. Rotary drilling in the oilfield has long been run against such limits through a rate law and a measure of drilling efficiency [1, 2], and Volume 10, Chapter 2 pushes an oil well's rate of penetration against limits of the same kind. Depth is known only by integrating the beacon's pitch between walkover locates, and nobody can walk over a river. The mud pressure in the annulus grows with every meter drilled, while the pressure the ground can hold is a cavity-expansion limit that falls with the cover [3–5]. Mud that breaks through to the surface, an inadvertent return, is the failure regulators and owners watch for [5], and surface heave is the related damage mechanism that field trials have measured [6].

> **DID YOU KNOW? From a Pilot Bore to a Wide Crossing**
>
> The first pass of a crossing cuts a guided bore 73 to 114.3 mm across, which later passes ream to the size of the pipe. The industry's equipment now accommodates bores up to 1524 mm across and lengths beyond 2000 m [7].

Two common beliefs are contradicted by the simulator. The first is that the bit sets the drilling rate. In the sand of case 4 the drill set for speed cut at 8.32 m/min and averaged 126.0 m/h, but it outran its 100 L/min of mud, laid a bed of cuttings and broke out at 114.6 m. The second is that the deepest part of a crossing is where the mud pressure is most dangerous. On the river crossing as shipped, at 6080 s the annular pressure was 0.20 to 0.31 of the ground's limit under the river and 0.996 of it at 172.6 m on the exit climb, where the cover had thinned to 1.11 m, and the mud broke out there 14.5 s later.

However, each part of this plant has its own literature, and the parts rarely meet. The geotechnical literature treats the allowable mud pressure [4, 5, 8], the construction literature treats the industry and its practice [9, 10], the drilling literature treats the rate and its efficiency [1, 2], and design guidance treats the pipe that is pulled back afterward [11, 12]. This chapter assembles the steering, the guidance, the drilling rate, the cuttings and the mud-pressure window of one pilot bore on a single simulated electric drill, and measures which control layers a river crossing needs and why. The cavity-expansion limit, the specific energy, the rate law and the Kalman filter are established methods, and the chapter works with an implementation of them and the numbers measured on it: a fifteen-state model whose cuttings balance closes to 1.2×10^{-12} of the soil cut, a rate optimizer within 0.5 % of a dense grid, and a controller boundary that admits only instrument readings. Eight cases and a graded crossing each carry a completion rule.

Section 7.2 follows the domain. Section 7.3 sets out the model, from the path and the rate law to the mud pressure against the ground's limit and the electric bus, and Section 7.4 the rod sequencer, the guidance, the steering autopilot, the rate optimizer and the two overrides. Section 7.5 measures the cases, from one pushed rod to the graded river crossing, and Section 7.6 separates the kinds of evidence, Section 7.7 sets the loop on a real drill's hardware, Section 7.8 states the limitations, Section 7.9 concludes, and Section 7.10 holds the exercises.

7.2 Trenchless Crossings, Mud Pressure and the Pilot Bore

Horizontal directional drilling installs a product pipe in three passes, a steered pilot bore, one or more reaming passes and a pullback. How the method learned to steer and grew into an industry is told in the first history segment below. The pullback that follows the pilot has its own load estimates [12], and the design of pipelines installed by the method is the subject of an engineering society's design volume [11]. The simulator stops at the end of the pilot bore.

> ⧖ HISTORY **1960s to 2000: Learning to Steer Under a River**
>
> Early attempts at horizontal drilling in the 1960s on the West Coast of the United States had no proven way to steer the bit. A pioneer of the method learned that some tools climb toward the surface on their own and set the length of a crossing by the angle of entry and the type of bit, which gave vertical control but no horizontal steering. By 1971 downhole survey methods and drilling techniques had been brought over from oil exploration [13], and a first crossing was drilled under a river in California that year [7]. By 2000 the method was described as perhaps the fastest growing segment of trenchless construction, with 17 makers of rigs and accessories and thousands of rigs in North America, and a survey drew 49 responses from contractors in the United States and Canada [9]. Suggested practices for urban work appeared the same year [10].

The mud-pressure problem came from soil mechanics, and the second history segment below follows the cavity-expansion limit from 1972 to the step-by-step evaluation of 2008. A 2020 alignment optimization compared the drilling pressure with the allowable every 0.3 m along candidate designs and described the cavity-expansion equation known as the Delft method as the most commonly accepted formulation [14]. A newer design equation has since been proposed for sand [8]. A full-scale field experiment on four bore paths related surface heave to backream rate, cover, mud flow and reamer type [6].

The drilling rate came from the oilfield. Specific energy, the work spent per volume of rock removed, was proposed as a drilling efficiency measure in 1965 [1], and a multiple regression of drilling variables on offset-well data gave the multiplicative rate law of 1974 [2]. The simulator borrows the multiplicative form and the efficiency measure, not either paper's coefficients.

> ⧖ HISTORY **1972 to 2008: How Much Mud Pressure the Ground Holds**
>
> The expansion of a cylindrical cavity in an elastic-plastic soil was solved in general form in 1972, with the ultimate pressure set by the initial ground stress, the strength and the rigidity index, the ratio of the shear modulus to the strength [3]. Applied to a borehole, the pressure at which the plastic zone reaches an allowable radius is the largest the ground will hold. For purely cohesive ground, a 2006 analysis introduced the earth pressure at rest explicitly and found that holding the zone of shear failure to half the cover leaves a reserve of 20 to 30 % against blowout [4]. A 2008 step-by-step guide to hydrofracture evaluation called the cavity expansion model the appropriate one for most bores in soil, kept a tensile strength model for shallow rock, and paired the maximum allowable pressure with a minimum required pressure that ensures the cuttings leave the hole [5].

The electric drive is the newest element. A 2018 assessment of non-road mobile machinery found that electrification had started slowly, through hybrids of some specific machines, with component cost the main short-term barrier and electric powertrains expected to become mainstream in the long term [15]. A battery-electric drill turns the spindle, the carriage and the mud pump into three loads on one bus, which makes energy per meter a product number and power a constraint on the rate. Guidance has changed in the same direction. In conventional practice the beacon behind the bit reports its pitch and roll to a walkover locator, and a drill whose console receives the beacon's data directly (integrated guidance) still needs the locator for depth wherever the surface can be walked. The console's estimate is then a filtering problem [16], as it is for the survey vehicle of Volume 5, Chapter 13, which dead reckons between sparse acoustic fixes.

7.3 The Crossing as Modeled

The simulated drill is a generic mid-size, battery-electric, rack-and-pinion utility drill with 180 kN of thrust and pullback, a 7.5 kN m spindle on a 90 kW motor from 0 to 220 rpm, a mud pump of 380 L/min at 8 MPa, 3 m rods of 60 mm outside and 44 mm inside diameter, and a 115 mm slant-face pilot bit with the beacon housing behind it. A 150 kWh battery and a 40 kW charger feed a 120 kW inverter. Every value is a teaching value of that class and none is a maker's figure. The river crossing, shown in Figure 7.1, runs 180 m from the entry at station $x = 0$ down a -14 % tangent, through a sag bend of radius 90 m to a horizontal 8 m deep, under a river from 55 to 120 m whose bed lies 3.5 m down, and up a $+12$ % exit tangent. The ground is a soft silty clay cap 1.5 m thick on land over stiff clay, sand from 58 m, a gravel lens from 124 to 146 m and sand again, with the water table 1.5 m down. Units are meters, seconds, newtons and pascals. The console shows US field units by default, and its unit switch selects the SI units in which the exercises of this chapter are stated.

Each variable in Table 7.1 carries its role. Fifteen states are integrated. Six are the path, the drilled length s, the position (x, z, y), the pitch θ (nose-up positive) and the azimuth ψ. Four are the annulus, the soil fraction of the returns C, its product with the required mud ratio C_R, the bed B and the soil returned to the pit. Two are the battery energy E and the energy drawn, and three are the lagged spindle, pump and carriage commands.

7.3.1 The Path and How a Slant Face Steers

The bore advances along its own length at the drilling rate, as shown in Equation (7.1).

$$\frac{\mathrm{d}x}{\mathrm{d}s} = \cos\theta\cos\psi, \quad \frac{\mathrm{d}z}{\mathrm{d}s} = \sin\theta, \quad \frac{\mathrm{d}y}{\mathrm{d}s} = \cos\theta\sin\psi$$
$$\frac{\mathrm{d}\theta}{\mathrm{d}s} = \begin{cases} \kappa_s\cos\varphi & \text{pushed at } \varphi \\ -\kappa_d & \text{rotated} \end{cases} \qquad \frac{\mathrm{d}\psi}{\mathrm{d}s} = \frac{\kappa_s\sin\varphi \text{ or } \kappa_w}{\max(0.2, \cos\theta)} \tag{7.1}$$

where $\frac{\mathrm{d}s}{\mathrm{d}t} = \mathrm{ROP}$, κ_s is the soil's push curvature, κ_d its droop and κ_w its walk while rotating, and $\varphi = 0$ is 12 o'clock. A push fraction d of a rod is pushed and the rest rotated, so the rod's mean

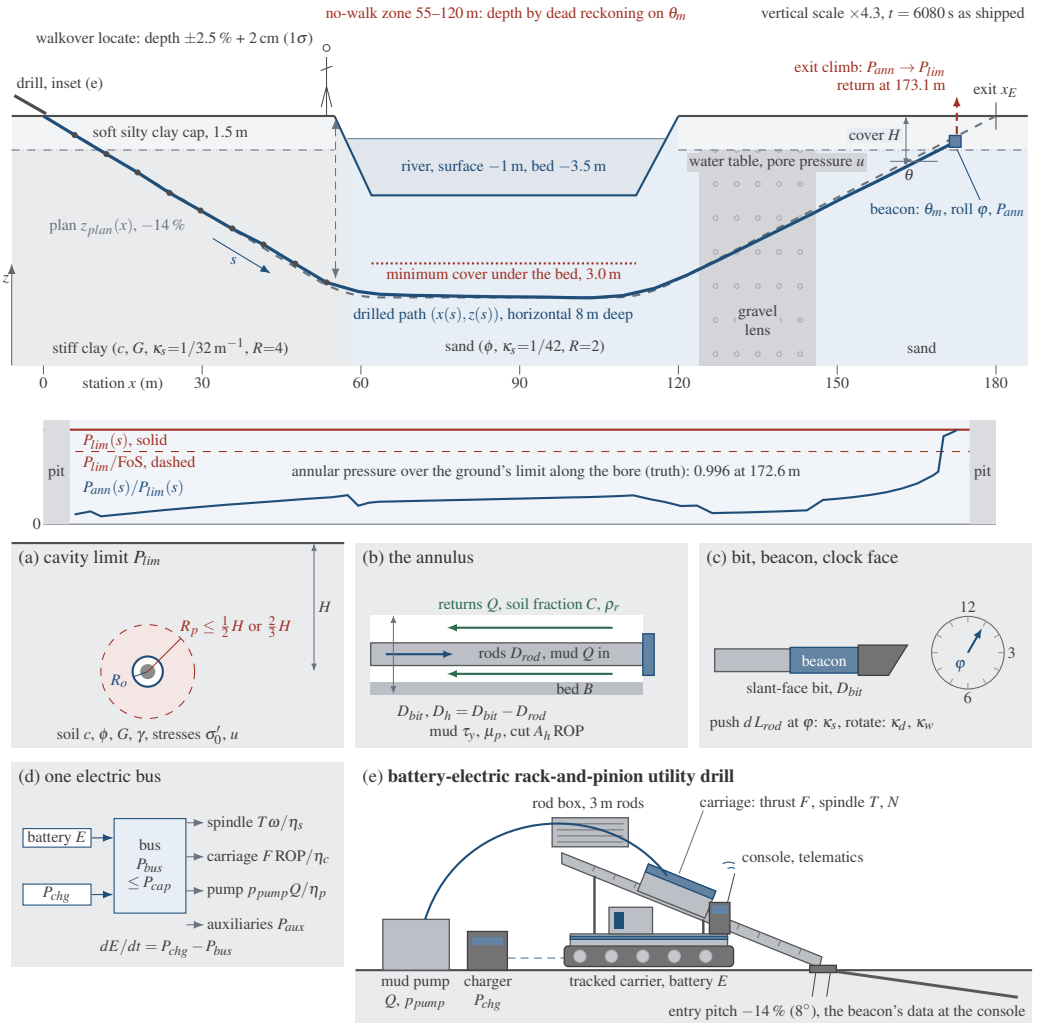

Figure 7.1 The river crossing as modeled, mirroring the simulator's Section view at $t = 6080$ s as shipped. The drill at the entry pushes rods along the drilled length s through soils of strength c, ϕ, G, γ and push curvature κ_s, on a path (x, z) of pitch θ (thick solid) against the plan z_{plan} (dashed), with locates on land (dots) and none over the river. The beacon reports θ_m, the roll φ and the annular pressure P_{ann}, and the strip under the section shows P_{ann}/P_{lim} along the bore against P_{lim} (solid) and P_{lim}/FoS (dashed), 0.996 at 172.6 m on the exit climb where the cover H is thin. Insets: (a) the cavity of radius R_o and its allowable plastic radius R_p (dashed circle), (b) the annulus with returns of soil fraction C over a bed B, (c) the bit, beacon and clock face, (d) the bus, (e) the drill at the entry, drawn at its own scale. The lateral offset y and azimuth ψ lie normal to the page. Not to scale.

curvature is about $d\kappa_s$. The console shows the pitch as $100 \tan\theta$ in percent. A rod whose curvature passes $1/R_{min}$, with a teaching bend radius of 30 m, counts as bent past its radius. At that radius a 3 m rod turns 5.73°, 10.03 points of pitch from level and 10.09 from −14 %, while the stiff clay's $\kappa_s = 1/32\,\text{m}^{-1}$ gives 9.47 points by hand from −14.2 % against the 9.40 measured in case 1, the difference being the 5 s spent turning the face before the push.

Table 7.1 Variables by Role, With Their Ranges

Variable	Role	Range or default	Unit
$z - z_{plan}$, x_{exit}	CV, path and exit	exit ±1.0	m
ROP, average rate	objective	—	m/min, m/h
P_{ann}/P_{lim}	constraint (truth)	≤ 1/FoS	—
T, P_{bus}, C	constraints	85 %, cap, ratio	—
d, φ	MV, push, clock face	0–100 %, 24	—
F	MV, carriage thrust	0–180 (100)	kN
N	MV, spindle speed	0–220 (100)	rpm
Q	MV, mud flow	0–380 (120)	L/min
soil, boundaries	DV, geology	clay to soft rock	—
b	DV, beacon pitch bias	±2 (0)	%
cobble, pump, power	DV, events	—	—
s, x, z, y	states, path	0–190	m
θ, ψ	states, pitch, azimuth	—	rad
C, C_R, B	states, annulus	≥ 0, $B \leq 0.6$	—
E	state, battery	0–150	kWh
θ_m, φ_m, $P_{ann,m}$	beacon, 1 Hz	0.1 %, 24, 1 kPa	—
depth, offset	walkover locate	2.5 % + 2 or 5 cm	m

7.3.2 Drilling Rate, Torque and the Drive

The rate saturates in the thrust that reaches the bit and has a best spindle speed past which the face balls in clay or washes in sand, in the multiplicative spirit of the oilfield rate laws [2], as shown in Equation (7.2).

$$\begin{aligned} \mathrm{ROP}_{rot} &= K\left(1 - \mathrm{e}^{-F_b/F_s}\right)\left[\tfrac{N}{N_f}\,\mathrm{e}^{1-N/N_f}\right]^{\gamma} f_{pk} \\ \mathrm{ROP}_{push} &= \min(K_p F_b,\ K)\, f_{pk} \\ f_{pk} &= \left[1 + (B/0.25)^2\right]^{-1} \\ F_b &= F - c_{rot}\left(\mu_r w_b L + \textstyle\sum d_L \Delta s\right)\mathrm{e}^{\mu_r \Theta} \end{aligned} \tag{7.2}$$

where $\gamma = 0.6$, $L = s$ is the string length, w_b the rod's buoyed weight per meter, $\mu_r = 0.3$ its axial friction, d_L the soil's drag per meter of hole, Θ the total turn of the path so far (a capstan factor), and $c_{rot} = 0.15$ while the string rotates and 1 while it is pushed. The spindle torque and the drive's limits are given by Equation (7.3).

$$\begin{aligned} T &= T_0 + \mu F_b r_e + \textstyle\sum t_L \Delta s + c_v\, \omega L + k_{pk} B L \\ \omega &= \max\left\{\omega' \leq \omega_{ref} : T(\omega') \leq T_{rated},\right. \\ &\qquad \left. T(\omega')\,\omega' \leq \min(P_{sp},\ \eta_s P_{left})\right\} \end{aligned} \tag{7.3}$$

where $\omega = 2\pi N$, r_e is two thirds of the bit radius, t_L the soil's string torque per meter, ω_{ref} the drive's ramped reference (a 2 s lag), and P_{left} the bus power left after the pump and the carriage. The motor's 90 kW meets the 7.5 kN m rating at 114.6 rpm, and a torque over the rating at standstill stalls the spindle. A cobble adds 1500 N m for 0.3 m and knocks the head 0.4 to 0.8° in pitch.

7.3.3 Cuttings, the Bed and the Mud Ratio

The annulus between the rods and the hole, of area A_{ann} and volume $V = A_{ann}\max(3, s)$, is one well-mixed volume. The bit cuts A_hROP of soil, the mud carries a soil fraction up to $C_{cap} = \eta/(1 + R_{mix})$ with $\eta = v/(v + v_{set})$ and $v = Q/A_{ann}$, and what it cannot carry settles into a bed that clean mud

erodes slowly, as shown in Equation (7.4).

$$\begin{aligned} V\frac{\mathrm{d}C}{\mathrm{d}t} &= A_h\mathrm{ROP} - C(Q + A_h\mathrm{ROP}) - V\,r_b \\ V\frac{\mathrm{d}C_R}{\mathrm{d}t} &= RA_h\mathrm{ROP} - C_R(Q + A_h\mathrm{ROP}) - V\,r_b R_{mix} \\ V\frac{\mathrm{d}B}{\mathrm{d}t} &= V\,r_b - BA_{ann}\mathrm{ROP} \\ r_b &= k_d\,g(Q)(C - C_{cap})^+ - k_e\frac{Q}{V}B\left(1 - \frac{C}{C_{cap}}\right)^+ \end{aligned} \tag{7.4}$$

where R is the required mud-to-soil volume ratio of the soil at the bit, $R_{mix} = C_R/C$ that of the soil the returns carry, $k_d = 1/30\,\mathrm{s}^{-1}$, $k_e = 0.3$, and $g(Q)$ rises from 0 at 30 L/min to 1 at 60 L/min, so the mud's gel holds what it carries while the pump stops and starts. The bed is capped at 0.6. The soil leaving at the pit is $C(Q + A_r\mathrm{ROP})$, which closes the balance.

7.3.4 The Annular Pressure and the Ground's Limit

The mud pressure at a node i of the path, kept every 0.5 m of drilled length, is the weight of the returns below the entry pit plus the friction of the annulus between the node and the pit, as shown in Equation (7.5).

$$\begin{aligned} P_i &= \rho_r g\,(z_{pit} - z_i) + G_f\,s_i, \qquad \rho_r = \rho_m + C(\rho_s - \rho_m) \\ G_f &= \left[\frac{6\tau_y}{D_h} + \frac{48\mu_p v}{D_h^2}\right]\frac{1 + a_s C}{(1 - B)^3} \end{aligned} \tag{7.5}$$

where $D_h = D_{bit} - D_{rod} = 55\,\mathrm{mm}$, $\tau_y = 4\,\mathrm{Pa}$, $\mu_p = 12\,\mathrm{mPa\,s}$ and $a_s = 2$. The bracket is the laminar Bingham slot form of drilling hydraulics, and at 115 L/min its yield term gives 436 of the 485 Pa/m, so less flow barely lowers the friction. The bed's cubic factor is a teaching closure that makes a pack-off felt. Because every node carries the friction of all the annulus behind it, the pressure grows with the bore's length. The ground's limit at each node is the Delft cavity-expansion pressure [3, 14], as shown in Equation (7.6).

$$\begin{aligned} P_{lim} &= u + (p'_f + c\cot\phi)\,X^{-\sin\phi/(1+\sin\phi)} - c\cot\phi \\ X &= \left(\frac{R_o}{R_p}\right)^2 + \frac{\sigma'_0\sin\phi + c\cos\phi}{G}, \quad R_p = r_p H \\ p'_f &= \sigma'_0(1 + \sin\phi) + c\cos\phi \\ P_{lim}\big|_{\phi\to 0} &= \sigma_0 + c_u\left[1 + \ln\frac{1}{(R_o/R_p)^2 + c_u/G}\right] \end{aligned} \tag{7.6}$$

where $R_o = 57.5\,\mathrm{mm}$ is the hole's radius, H the cover, σ_0 the total vertical stress (the river's water included), u the pore pressure from the local water table and $\sigma'_0 = \sigma_0 - u$. The allowable plastic radius $r_p H$ is half the cover in cohesive ground, the fraction the cohesive analysis supports [4], and two thirds in granular ground, both teaching values. The far field is taken as isotropic ($K_0 = 1$). An inadvertent return occurs when $P_i > P_{lim,i}$ at any node outside the 5 m pits at each end. It is latched, and the vent fixes the pressure of the nodes beyond it from then on. A return under the river stops the crossing.

7.3.5 The Electric Bus and the Rod Cycle

The three drives draw from one bus, the pump first, then the carriage, and the spindle gets the rest, a fixed priority in the sense of Section 1.4.8, as shown in Equation (7.7).

$$\begin{aligned} P_{bus} &= \frac{T\,\omega}{\eta_s} + \frac{F\,\mathrm{ROP}}{\eta_c} + \frac{p_{pump} Q}{\eta_p} + P_{aux} \le P_{cap} \\ P_{cap} &= \min(P_{inv},\; P_{chg} + P_{batt}), \qquad \frac{\mathrm{d}E}{\mathrm{d}t} = P_{chg} - P_{bus} \\ p_{pump} &= f_{in}(Q, L) + k_n Q^2 + P_{bit} \le 8\,\mathrm{MPa} \end{aligned} \tag{7.7}$$

where $\eta_s = 0.90$, $\eta_c = 0.85$, $\eta_p = 0.80$, $P_{aux} = 3\,\mathrm{kW}$, $P_{batt} = 100\,\mathrm{kW}$ and f_{in} is the turbulent loss inside the rods. The bit's jets alone pass 170 L/min at 8 MPa and 168.1 L/min with 180 m of rods, so the pump's 380 L/min rating is never usable with this bit. Without the battery the cap is the charger's 40 kW, and the battery is the depleting stock of Section 1.4.9. The rod cycle is the forced periodic stop. When 3 m of a rod is in, thrust and rotation stop (4 s), the rod is changed with the pump off (42 s), and the pump restarts (4 s), so every rod costs 50 s, and turning the face before a push costs another 5 s.

7.3.6 The Instruments

The drill measures its own drilled length exactly, the thrust, the spindle speed, the torque with 1 % colored noise, the flow, the pump pressure and the bus. The beacon reports once a second the pitch rounded to 0.1 % with a declared bias b and 0.05 % of colored noise, the roll in 24 clock positions, and the annular pressure at the head rounded to 1 kPa with 2 kPa of noise, and an interference event holds its last readings for 60 s. A walkover locate is sampled when a rod change starts wherever the surface can be walked and reported 25 s later, depth with a standard deviation of 2.5 % of the depth plus 2 cm and the lateral offset with 2.5 % plus 5 cm. None is taken over a road or a river, and past one only when far-side locates are switched on. The beacon gives no azimuth. True quantities are logged with a `truth_` prefix and drawn dashed, and a source test confirms that no controller file reads one.

7.3.7 Soils, Parameters and Numerics

Table 7.2 gives the five soils and Table 7.3 the drill and mud. Each soil was calibrated so that it hands the rate to its intended limit, and every value is a teaching value whose real counterpart comes from the site investigation. The plant is integrated by a fourth-order Runge-Kutta step of 0.5 s with everything that changes discretely (the soil at the bit, the path's sums, the pump's pressure ceiling) frozen over the step, and each step is cut exactly at the end of a push and of a rod's stroke. Control decisions are made each second, and a row is logged every two seconds (Section 1.6.1).

7.3.8 Assumptions

The annulus is one well-mixed volume with one bed, so a real pack-off, which is local and can be sudden, is spread over the whole hole. The friction is the laminar Bingham slot form with a fixed yield point, and the bed's cubic factor overstates what a narrowed slot does to a yield-dominated friction, which would scale as $(1 - B)^{-1}$. Each soil steers at one curvature and drills by one law, boundaries are sharp, and a slant face's deflection does not depend on its angle, the thrust or the jets. The cavity limit assumes an isotropic far field, one homogeneous soil at each node and a circular plastic zone, so layering, a free face near the exit and fractures, which all lower a real limit, are absent. No mud is lost into permeable ground. The rod's 8.5 kg/m is lighter than a plain 60/44 mm steel tube's 10.3 kg/m, which changes the string's axial friction at 180 m from 3.90 to 4.71 kN, under 0.5 % of the rating. The bus has fixed efficiencies and no battery ageing or charger losses. The safety numbers (the factor of safety 1.3, the allowable plastic radius, the 3.0 m minimum cover under the river, the 90 % torque limit, the 30 m bend radius) are teaching values. The real ones come from

Table 7.2 The Five Soils (Teaching Values)

	cap	clay	sand	gravel	rock	unit
K	3.0	3.0	11.0	5.5	1.2	m/min
F_s	45	60	35	150	160	kN
N_f	120	110	160	120	150	rpm
K_p	1.4	1.2	2.6	0.8	0.03	m/min per 100 kN
μ	0.85	1.10	0.35	0.65	0.90	—
t_L, d_L	12, 70	16, 80	7, 90	11, 250	9, 40	N m/m, N/m
$1/\kappa_s$	22	32	42	70	2000	m
$1/\kappa_d$	1400	2000	1500	—	—	m
$1/\kappa_w$	—	—	20000	2500	—	m
R	3.5	4.0[a]	2.0[a]	2.5[a]	1.5	—
v_{set}	0.012	0.010	0.045	0.090	0.060	m/s
c	30	60	0	0	400	kPa
ϕ	0	0	33	38	30	°
G	4.5	9	15	40	300	MPa
γ	18.5	19.5	19.5	20.5	22.5	kN/m^3
r_p	1/2	1/2	2/3	2/3	2/3	of cover

[a] Adjustable in the Model dialog.

Table 7.3 Drill, Mud, Guidance and Controller Parameters. A = Adjustable in the Model Dialog, T = Teaching Value, D = Derived

Symbol	Meaning	Value	Unit	Status
F_{max}	thrust rating	180	kN	T
T_{rated}, P_{sp}	spindle torque, motor	7.5, 90	kN m, kW	T
N_{max}	spindle speed range	220	rpm	T
Q_{rated}, p_{max}	pump	380, 8	L/min, MPa	T
k_n	bit jets, $\Delta p = k_n Q^2$	1×10^{12}	Pa s^2/m^6	T
L_{rod}, D_{rod}, d_{rod}	rod length, diameters	3.0, 60, 44	m, mm	T
w_{rod}	rod mass per meter	8.5	kg/m	T
R_{min}	rod bend radius	30	m	A
D_{bit}, r_e	bit, torque arm	115, 38.3	mm	T, D
T_0, c_v, k_{pk}	torque terms	150, 0.3, 30	SI	T
μ_r, c_{rot}	string friction, rotation cut	0.3, 0.15	—	T
τ_N, τ_Q, τ_F	drive lags	2, 1.5, 1	s	T
E_{max}, P_{chg}	battery, charger	150, 40	kWh, kW	T, A
P_{inv}, P_{batt}	inverter, battery power	120, 100	kW	A, T
η_s, η_c, η_p	drive efficiencies	0.90, 0.85, 0.80	—	T
t_{rod}	rod change, stop to restart	50	s	A
ρ_m, ρ_s	mud, soil density	1050, 2000	kg/m^3	T
τ_y, μ_p	yield point, plastic viscosity	4, 0.012	Pa, Pa s	A, T
a_s, B_{max}	solids factor, bed cap	2, 0.6	—	T
k_d, k_e	settling, erosion	1/30, 0.3	s^{-1}, —	T
FoS	factor of safety on P_{lim}	1.3	—	A
pits	returns expected	5	m	T
σ_{loc}	locate depth accuracy	2.5 % + 0.02	m	A
t_{loc}	locate report delay	25	s	T
q_z, q_b	depth, bias filter noise	1×10^{-4}, 1×10^{-10}	per m	T
q_y, q_ψ	lateral, azimuth filter noise	1×10^{-4}, 2×10^{-6}	per m	T
L_a	autopilot look-ahead	9	m	A
T_{opt}	optimizer period	10	s	A
T_{cap}	optimizer torque ceiling	85	%	A
m	mud-ratio margin	10	%	T
r_{ov}, F_{rel}	override target, release	0.97, 20	—, kN	T
limits	torque, cover, exit, time	90 %, 3, 1, 2.75	m, h	T

the crossing's design report and permit, the hydrofracture evaluation [5], the drill maker's operating limits and the rod maker's specification, and a simulated pass is not a design or a permit.

7.4 Sequencer, Guidance, Autopilot and Optimizer as Coded

Every layer runs each $\Delta t_c = 1\,\mathrm{s}$ from the measurement packet, the declared geotechnical log and the bore plan. The log has the same soils and boundaries as the ground unless a case moves a true boundary earlier.

7.4.1 The Rod Sequencer and the Guidance

The sequencer reads only the drilled length and steps through drill, stop, change and start. At each start it takes the next rod's push fraction d and clock face φ from the operator or the autopilot. The console's position is dead reckoned from the measured pitch for every meter drilled, and three guidance modes differ in how a locate enters, as shown in Equation (7.8).

$$\begin{aligned} \hat{z}^{+} &= \hat{z} + \Delta s\,\sin(\theta_m - \hat{b}) \\ \hat{\psi}^{+} &= \hat{\psi} + \frac{\kappa_{log}\,\Delta s_{push}\,\sin\varphi}{\max(0.2, \cos\theta_m)} \\ \begin{bmatrix} \hat{z} \\ \hat{b} \end{bmatrix}^{+} &= \begin{bmatrix} \hat{z} \\ \hat{b} \end{bmatrix} + \mathbf{K}\,(z_{loc} - \hat{z}), \qquad \mathbf{K} = \frac{\mathbf{P}\mathbf{h}^{\top}}{\mathbf{h}\mathbf{P}\mathbf{h}^{\top} + \sigma_{loc}^{2}} \end{aligned} \tag{7.8}$$

where $\mathbf{h} = [1\ 0]$ and the covariance grows with $\partial\hat{z}^{+}/\partial\hat{b} = -\Delta s\cos\theta_m$ and the process noise of Table 7.3. Dead reckoning ignores the locates and keeps $\hat{b} = 0$. The walkover practice (Locate reset) replaces $\hat{z}$ and $\hat{y}$ at each locate. The estimator is two two-state Kalman filters [16], one on depth and pitch bias and one on lateral offset and azimuth, started with a pitch-bias standard deviation of 1 %. A constant pitch bias b in percent grows a depth error of $b\cos^3\theta$ per meter, 1.000 cm/m level and 0.973 cm/m at −14 %, so a learned bias carries the depth across a reach nobody can walk.

7.4.2 The Steering Autopilot

At each rod start the autopilot aims at the plan a look-ahead L_a ahead of the console's head, in the manner of pure-pursuit path tracking, as shown in Equation (7.9).

$$\begin{aligned} \alpha_v &= \operatorname{atan}\frac{z_{plan}(\hat{x} + L_a) - \hat{z}}{L_a} - \hat{\theta}, \quad \alpha_h = \operatorname{atan}\frac{-\hat{y}}{L_a} - \hat{\psi} \\ \mathbf{t} &= \left(\frac{2\sin\alpha_v}{L_a} + \kappa_d,\ \frac{2\sin\alpha_h}{L_a} - \kappa_w\right), \quad \varphi = \operatorname{atan2}(t_h, t_v) \\ d &= \min\left(\frac{|\mathbf{t}|}{\kappa_s},\ \frac{0.9}{R_{min}\,\kappa_s},\ 1\right) \end{aligned} \tag{7.9}$$

where the clock face is rounded to 24 positions, a push fraction under 8 % becomes a rotated rod, and the soil's curvatures come from the log. Where the soil gives no authority ($\kappa_s < 1/500\,\mathrm{m}^{-1}$) the autopilot rotates and reports that it is short of steering authority instead of pushing a rod that would only drill slower.

7.4.3 The Soil Read and the Rate Optimizer

With the bit rotating and cutting, the measured specific energy [1] is compared with each soil's model at the same thrust and speed, as shown in Equation (7.10).

$$\begin{aligned} \mathrm{MSE} &= \frac{F_b}{A_h} + \frac{2\pi N T}{A_h\,\mathrm{ROP}} \\ \mathrm{score}_j &= \left\langle \left| \ln\frac{\mathrm{MSE}}{\mathrm{MSE}_j} \right| \right\rangle_{15\,\mathrm{s}} + 0.25\,[j \neq j_{log}] \end{aligned} \tag{7.10}$$

where the candidates are the soils the log shows within 12 m of the head and the reading changes only on an improvement of 0.15. Every $T_{opt} = 10\,\mathrm{s}$ the optimizer then solves the problem of Equation (7.11)

for the soil it reads.

$$\begin{aligned} &\max_{F,N,Q} \ \mathrm{ROP}(F,N) \\ &\text{s.t. } T + T_b \le 0.85\min(T_{rated},\ P_{sp}/\omega) \\ &\quad P_{bus} \le 0.97 P_{cap}, \quad F \le F_{max} \\ &\quad \frac{A_h \mathrm{ROP}}{Q + A_h \mathrm{ROP}} \le (1-m)\frac{\eta(Q)}{1+R}, \quad p_{pump} \le 8\,\mathrm{MPa} \\ &\quad \rho_r g(z_{pit} - \hat{z}_i) + \kappa_f G_f(Q, C_b)\, s_i \le \frac{P_{lim,i}}{\mathrm{FoS}} \ \ \forall i \end{aligned} \tag{7.11}$$

where the nodes i are every point of the console's path outside the pits and twelve points every 0.5 m up to 6 m ahead on the plan, C_b is the larger of the steady returns and a cuttings observer's estimate (the plant's balance driven by the measured rate and flow), and the load ahead relaxes from the observer's value toward the steady one over the time the bit takes to get there. The friction factor κ_f and the torque offset T_b are learned by the framework's shared bias updater from the beacon's annular pressure and the measured torque, and the bit's torque is taken at the stickiest soil the log shows over the next rod. The search is a $19 \times 15 \times 8$ grid over thrust, speed and flow followed by a pattern search that halves its steps seven times, and among points within 0.3 % of the best rate it keeps the one with the least bus power. When the best point is held by the pressure at less than 97 % of the full flow, the optimizer drills at the full flow instead, or circulates there if no rate is legal, because less flow lowers the yield-dominated friction only slightly and leaves the cuttings in the hole. The binding constraint is the one with the highest use at the chosen point, at least 0.97, the active constraint of Section 1.4.2. The torque ceiling of 85 %, the factor 0.97 on the bus, the margin m on the mud ratio and the factor of safety on the pressure are back-offs in the sense of that section, each paid in rate.

7.4.4 Overrides, the Soft Start and the Three Models

Two min-select overrides act on the thrust with external reset, each an integrator on the request actually in force [17] (Sections 1.4.3 and 1.4.4), as shown in Equation (7.12). Chapter 11 builds the limits on its advance the same way.

$$\begin{aligned} \dot{F}_{ov} &= -0.25 F_{max} \frac{\hat{r} - 0.97}{0.03} \\ \hat{r} &= \max_i \frac{\rho_r g(z_{pit} - \hat{z}_i) + \hat{G} s_i}{P_{lim,i}/\mathrm{FoS}} \\ \dot{F}_{tq} &= -0.5 F_{max} \frac{T_m/T_c - 1}{0.02} \quad (\times 0.2 \text{ recovering}) \\ F_{cmd} &= \min(F_{req},\ F_{ov},\ F_{tq},\ F_{max}) \end{aligned} \tag{7.12}$$

where $\hat{G}$ is the gradient the measured head pressure implies over the drilled length, T_c is the torque ceiling at the present speed, and each limit snaps back to the request once it is clear. The pressure override takes thrust off and keeps the mud flowing. In Auto the thrust command climbs at 40 kN/s whenever rotation or a push starts, and after a push the spindle waits until the measured thrust has fallen under 20 kN, so a bit still loaded with the push's thrust never starts against the torque rating. The plant model is Section 7.3. The controller's model is the declared soil table, the log and the plan, which does not know the soil at the bit, the bed or the friction and torque offsets except through what it learns. No model predictive controller was built, so the app does not offer the option of Section 1.6.3. The rate problem is solved again every 10 s on a steady evaluator, and the time of a crossing is set by the rod changes and the pressure window.

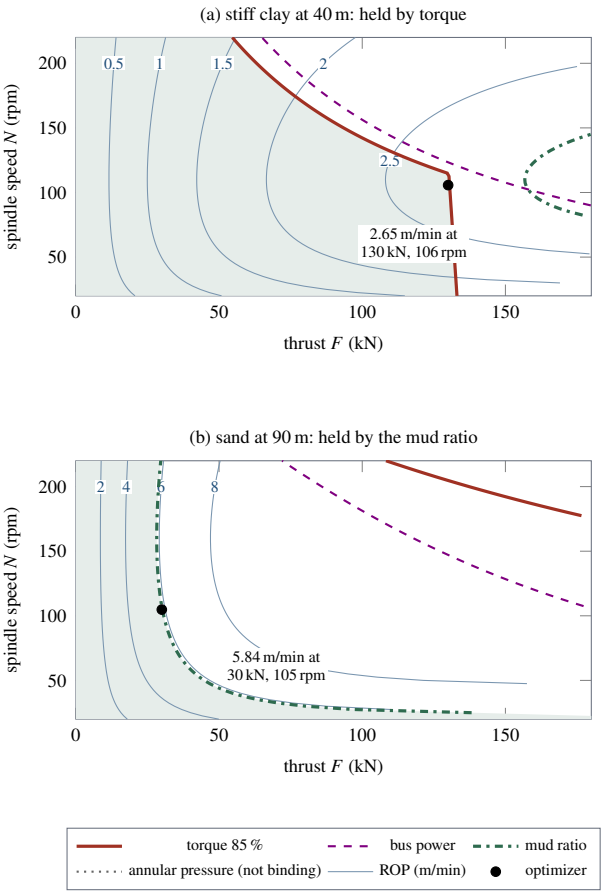

Figure 7.2 The optimizer's rate map on the river crossing with the suggested settings, from the app's own evaluator at the optimizer's flow (137 L/min in the clay, 167 L/min in the sand). Shaded points are legal, thin lines are contours of the rate, the dot is the optimizer's point, and the limits are the torque ceiling (very thick solid), the bus power (dashed) and the mud ratio (very thick dash-dotted). In the stiff clay at 40 m the torque ceiling holds the best point, and in the sand at 90 m the mud ratio does, an edge that is itself a contour of the rate.

7.5 What the Nine Runs Show

The cases appear in Table 7.4. All figures in this section come from the engine of Section 1.6.1 stepped without a browser at the default seed and the app's 0.5 s step. A setting that a sentence changes was entered before the run began, and loading a case first, choosing it from the list after another case, or pressing Reset gives the same run.

7.5.1 Steering by the Rod and Finding the Head

Rotated, the first two rods of case 1 held the pitch within 0.3 point and the rod after the push within 0.2. Left rotating throughout, the bore stayed at −14 % and was 2.96 m under the plan at 25 m, while left on Push it surfaced at 14.3 m after four rod changes. In case 2 the look-ahead is the steering loop's gain, as listed in Table 7.5. Aiming 9 m ahead, the autopilot reached the exit 0.15 m from the target in 0.80 h, never more than 0.18 m off the plan, with 14 of 20 rods pushed. Aiming 3 m ahead it held within 0.11 m but pushed 17 of 19 rods, took 1.04 h and came out 0.90 m off, and at 25 m it cut the bends, 0.59 m off the plan and 1.79 m off at the exit.

Table 7.4 The Cases, as Shipped. The App Lists Case 9 as the Challenge

#	Site and settings	Teaching point
1	30 m driveway, stiff clay, Manual, rotating	rotate to go, push to steer
2	60 m lawn, clay, Manual, then the autopilot in Auto	following a plan
3	66 m road, no walking 21–45 m, 1 % pitch bias, Locate reset	dead reckoning, the estimator
4	120 m of sand from 34 m, 100 kN, 120 rpm, 100 L/min	the mud-ratio edge
5	150 m, clay, sand from 52, gravel from 104 m, clay settings	the binding limit moves
6	160 m from 100 m, loaded annulus, 120 kN, 120 rpm, 100 L/min	pressure override
7	90 m of clay, battery out, 40 kW	the shared bus
8	70 m, cobbles, soft rock from 17 m under the sag bend	infeasible steering
9	180 m river, clay's 125 kN, 110 rpm, 115 L/min, Locate reset, no far-side locates, 0.3 % bias	the whole problem

Table 7.5 Measured Results on the Cases. Average Rate Includes Every Rod Change

Case	Setting	Rate (m/h)	Other
2	look-ahead 9 m	75.3	exit 0.15 m, off plan 0.18 m
	look-ahead 3 m	57.1	exit 0.90 m, off plan 0.11 m
	look-ahead 25 m	76.5	exit 1.79 m, off plan 0.59 m
3	Locate reset (shipped)	74.8	depth error 0.246 m, exit 0.49 m
	estimator	75.7	depth error 0.141 m, exit 0.37 m
	dead reckoning	77.8	0.494 m at 50 m, exit 5.39 m
4	set for speed (shipped)	126.0	return at 114.6 m
	optimizer	108.4	worst ratio 0.71
	20 kN, 60 rpm	55.2	worst ratio 0.67
5	clay settings (shipped)	87.6	worst ratio 0.84
	optimizer	93.8	worst ratio 0.73
6	shipped	105.9	return at 149.8 m
	pressure override	44.2	worst 0.77, acting 54.6 min
	optimizer, override Off	76.6	worst 0.76
	optimizer and override	66.6	worst 0.74
	150 L/min	111.3	return at 154.8 m
7	shipped, 40 kW	68.5	37 rpm, 0.348 kWh/m
	optimizer	72.6	101 L/min, 0.312 kWh/m
	battery On	76.3	96 rpm, 0.505 kWh/m
8	shipped	48.7	1.50 m off plan at 39.2 m

Pause. Try Exercise 7.1 before reading on.

Under the road of case 3 the last locate before the road came at 20.8 m and the first after it at 47.8 m, 27.0 m of drilling later. With Locate reset the console was 0.246 m off the truth just before that locate, which brought it to 0.010 m. The estimator learned a bias of 0.90 % from the locates before the road and was never more than 0.141 m off, and dead reckoning alone drifted 0.494 m by 50 m and exited 5.39 m from the target. On seeds 1 and 2 the worst errors were 0.367 and 0.405 m with Locate reset against 0.064 and 0.074 m with the estimator, so the estimator's advantage is real but its size depends on the noise of the locates that taught it.

7.5.2 The Rate and the Limit That Holds It

In the sand of case 4 the drill set for speed cut at 8.32 m/min on the level run at 100 L/min, built a bed to 0.110 of the annulus and broke out at 114.6 m, just short of the exit pit. The optimizer opened the flow to the jets' limit of 168 L/min, slowed the bit to 5.36 m/min on the mud-ratio edge, and averaged 108.4 m/h with no return. It was held by the mud for 948 s of its drilling and by the pressure for 377 s on the climb to the shallow exit. Because the mud constraint bounds the returns' soil fraction, it bounds the rate itself, ROP $\leq Qc/[A_h(1-c)]$ with $c = (1-m)\eta/(1+R)$, which at 167 L/min in sand is 5.86 m/min whatever the thrust or speed, as shown in Figure 7.2b. The shipped bit cut 55 % faster than the optimizer's on the level run and averaged only 16 % more, because the 50 s rod change dominates a fast rod.

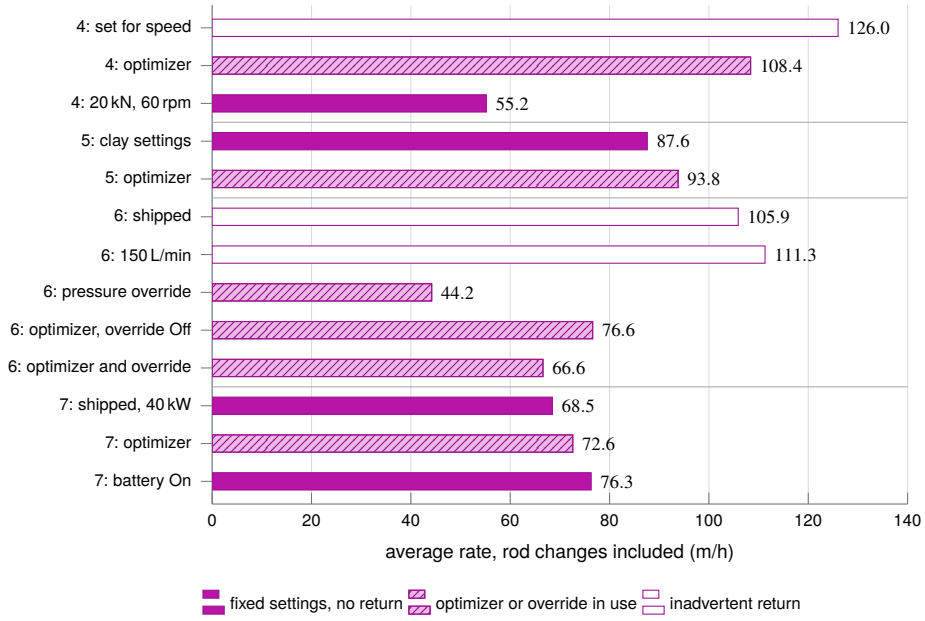

Figure 7.3 Average rates of cases 4 to 7 from Table 7.5, grouped by case (thin rules). White bars broke out, hatched bars ran with the rate optimizer or the pressure override. Notice that the fastest bar of cases 4 and 6 is a run that broke out, and that the override alone is the slowest way through case 6.

Pause. Try Exercise 7.2 before reading on.

In case 5 the limit moved with the soil, as the optimum of the data hall in Volume 9, Chapter 8 moves with the weather. With the optimizer the clay ran at 2.55 m/min held by the torque for 99 % of its drilling, the sand at 5.53 m/min held by the mud ratio for 97 %, and the gravel at 3.17 m/min held by the thrust for 64 %. The fixed clay settings drove the sand at 8.66 m/min, past its mud, and the gravel at 2.58 m/min, short of its thrust, and averaged 87.6 m/h against 93.8 m/h with a worst pressure ratio of 0.84 against 0.73. On the river crossing the same search, checked against a dense $61 \times 51 \times 24$ grid, found 2.65 m/min in the clay at 40 m held by torque and 5.84 m/min in the sand at 90 m held by the mud, as shown in Figure 7.2.

TRY IT IN THE APP 5 · The Soils Change, and So Does the Limit

Load the scenario, set Rate optimizer to On, press Reset and Run on the Map view. Watch the Drilling rate card name the limit that holds the optimizer. It reads torque in the clay, the mud ratio in the sand and thrust in the gravel, as the outlines on the map change with the soil.

On the 40 kW shore supply of case 7 the bus held the spindle at about 37 rpm for 96 % of the rotating time, and the bore averaged 68.5 m/h. The optimizer cut the flow from 120 to 101 L/min, which the clay's mud ratio still allowed, raised the thrust from 97 to 114 kN, turned at 35 rpm and made 72.6 m/h, held by the bus for 1730 s and the mud for 899 s. With the battery the same fixed settings made 76.3 m/h at 96 rpm and 75.1 kW, at 0.505 kWh/m against 0.312 kWh/m for the optimizer on shore power. Figure 7.3 sets the rates of cases 4 to 7 side by side.

7.5.3 The Pressure Window and the Override

The annular pressure of the river crossing grows along the bore while the limit follows the cover, as shown in Figure 7.4. At 6080 s as shipped, with a bed of 0.074 and returns of soil fraction 0.244

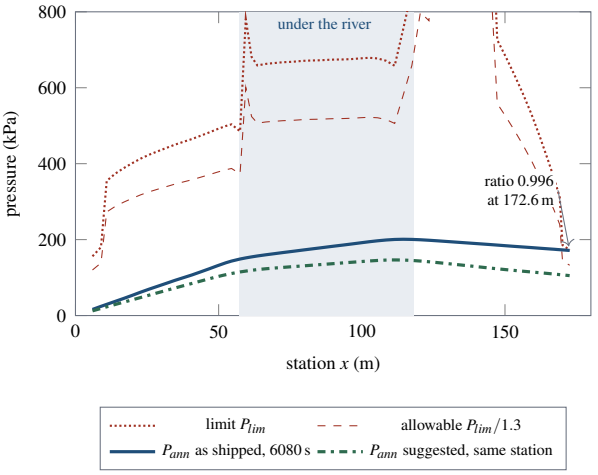

Figure 7.4 Annular pressure along the river crossing against the Delft limit (dotted) and the allowable (dashed), with the head at 172.6 m. As shipped (6080 s, very thick solid) the loaded annulus reaches 0.996 of the limit in the soft cap of the exit climb, and the mud breaks out at 173.1 m. With the suggested settings (very thick dash-dotted) the worst ratio at the same station is 0.639.

at 115 L/min, the friction gradient was 908 Pa/m and the pressure 183 kPa at 90 m against a limit of 674 kPa under the river. On the climb the limit fell from 511 kPa at 160 m in the sand to 185 kPa at 170 m in the soft cap, and the pressure there was 173 kPa. With the suggested settings at the same station the gradient was 545 Pa/m, the returns 0.054 and the bed empty, and the worst ratio 0.639.

Pause. Try Exercise 7.3 before reading on.

Case 6 isolates the exit climb. Set to keep going, the loaded annulus broke out at 149.8 m under 1.16 m of cover, and more flow alone (150 L/min) still broke out at 154.8 m. The pressure override held the thrust back for 54.6 min, let the flowing mud clean the hole, and finished with a worst ratio of 0.77 at 44.2 m/h. The optimizer alone also stayed clear (worst 0.76, 76.6 m/h) because its search keeps a predicted pressure under the allowable, and with both the rate fell to 66.6 m/h. Seven perturbations the optimizer cannot predict (a pump fault at 900, 1500 or 5400 s, soil boundaries 10 m earlier than the log, and a 60 s optimizer period, on case 6 and the crossing) never made the optimizer alone break out (worst 0.73 to 0.79), so the override is taught as the backstop for a wrong prediction or a drill run without the optimizer. With its fixed settings and no override, case 6 also broke out on seeds 1 and 2, at 150.3 and 147.4 m.

TRY IT IN THE APP 6 · Shallow Cover at the Exit: The Pressure Override

Run the scenario on the Section view with Pressure override Off and watch the strip under the section: the mud breaks out at 149.8 m on the exit climb. Set Pressure override to On, press Reset and run again. The override holds the thrust back for 54.6 min and the bore exits with no return at 44.2 m/h.

7.5.4 Soft Rock and an Honest Refusal

Case 8 places the plan's sag bend in soft rock, whose push curvature is $1/2000\,\mathrm{m}^{-1}$. A sweep of push fractions from 0 to 100 % at full thrust turned a rod by at most 0.15 point against the 6 points per rod the plan's 50 m bend needs, so no setting follows the plan. Run as shipped, the bit met four cobbles, entered the rock at 19.0 m, and the autopilot reported it was short of steering authority from 20.8 m and rotated on. The true path left the plan by 1.50 m at 39.2 m, where the crew stops. The remedy is a plan drawn for the ground or rock tooling, not a tuning.

Table 7.6 The River Crossing on the Default Seed

Configuration	Time (h)	Exit (m)	Ret.	T_{max} (%)	Cover (m)	Verdict
shipped	1.83	2.54	1	91.1	4.32	fails
suggested	2.13	0.60	0	85.8	4.38	passes
suggested, override Off	1.98	0.45	0	85.8	4.38	passes
suggested, no far-side locates	2.18	1.70	0	85.8	4.38	fails
suggested, Locate reset	2.23	0.54	0	85.9	4.34	passes
shipped, estimator and far side	1.89	0.56	1	89.1	4.49	fails
shipped, optimizer and override	2.38	2.67	0	85.9	4.34	fails

DID YOU KNOW? The Safety Layer Alone Is the Slowest Way Through

From the simulator: on case 6 the pressure override alone averaged 44.2 m/h, while the optimizer alone, whose search keeps its predicted pressure under the allowable, averaged 76.6 m/h with a worst ratio of 0.76 against 0.77 (Table 7.5). Both ran without a return.

7.5.5 The River Crossing

The challenge asks for the exit within 1.0 m of the target, at least 3.0 m of cover under the bed, no inadvertent return, torque at most 90 %, no rod bent past its radius and the bore in at most 2.75 h, judged when the head breaks the surface. Table 7.6 lists the configurations. As shipped the bore took 1.83 h and missed three targets. The clay settings outran the mud in the sand and left a bed whose pressure broke out on the exit climb at 173.1 m (6094.5 s), the torque reached 91.1 % where the bit climbed from the sand into the soft cap at 169.4 m with the clay's 125 kN, and with no locate after the river the 0.3 % pitch bias put the exit 2.54 m off. That last figure follows from the bias, since $0.003 \times 125\,\text{m} = 0.375\,\text{m}$ of depth at a 12 % exit tangent moves the surfacing point about 3.1 m. The suggested settings came out 0.60 m from the target in 2.13 h with 4.38 m of cover and no return, and without the override they also pass in 1.98 h. Each half of the suggestion is necessary. Guidance alone broke out at 174.5 m, and the rate layers alone exited 2.67 m off. Over ten seeds the shipped settings failed every time (exit 1.96 to 5.30 m off, a return on 7, torque over 90 % on 9), and the suggested settings passed every time (exit 0.12 to 0.75 m, 2.13 to 2.54 h, worst ratio 0.72 to 0.77). The return as it happens is shown in Figure 7.5.

The suggested crossing is plausible for the drill's class. The bus peaked at 119.7 kW (spindle 83.2 kW, pump 27.9 kW), the pump at 8.00 MPa, the annulus at 177 kPa at the bit, the torque at 6.43 kN m and the flow at 168 L/min. Sixty rods each cost 50.0 s to change, the bit cut 46 % of the time, and the bus drew 89.6 kWh for 181.6 m, 0.494 kWh/m, while the battery fell only from 85 to 81.9 %.

7.6 Code Checks, Correlations and the Limits of the Evidence

Code verification is the application's suite of 38 checks. The cuttings balance over a disturbed run (rate, flow, pump stops, a bed laid and eroded) closes to $6.1 \times 10^{-12}\,\text{m}^3$ of $5.186\,\text{m}^3$ cut, a stopped drill has every derivative exactly zero, the drilled length equals 3 m per rod at every change, and a full push over 3 m gives each soil's declared radius to 0.00 %. The optimizer reaches 100.2, 99.7 and 99.5 % of the dense grid's best at its three test points, inside the declared 0.3 % tie-break and 0.5 % search tolerance. The bus never exceeds its cap at any step of any case (closest approach 6.8 W under it). Halving the step to 0.25 s on a replayed rod program moves case 2's exit by 1.1 cm and its time by 0.18 %. In closed loop the per-rod and per-10 s decisions can flip with the step, so every verdict and ranking holds while numbers move. The suggested crossing took 2.13 and 2.27 h at 0.5 and 0.25 s, and case 4's optimizer 108.4 and 99.5 m/h. Playback at 60 and 600× logs the same row, and a crossing costs 629 to 829 ms of headless computing per simulated hour, so the suggested 2.13 h plays in 12.8 s at the top speed.

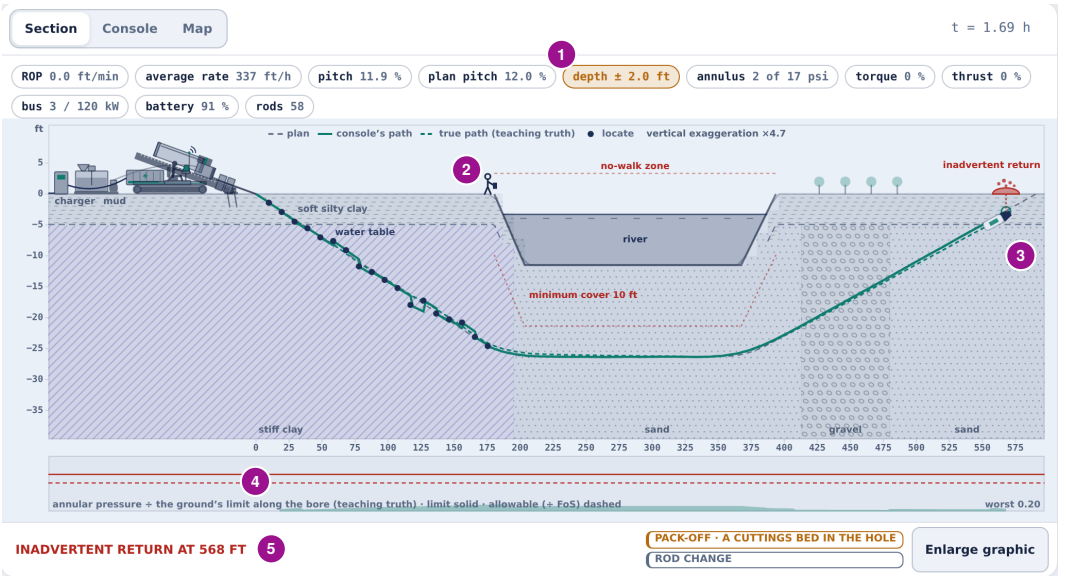

Figure 7.5 The river crossing as shipped at $t = 6100$ s on the Section view, 5.5 s after the inadvertent return, cropped to the view, with the run clock of the app's header moved into the strip beside the view tabs. (1) The readings: the beacon's pitch of 11.9 % against the plan's 12.0, the console's depth dead reckoned with ±2.0 ft (one standard deviation) after no locate past the river, an average rate of 337 ft/h (102.7 m/h) at 1.69 h, and 58 rods in the ground. (2) The no-walk zone over the river, where no locate is taken. (3) The plume over the vent at 173.1 m (568 ft) on the exit climb, with the head 1.06 m under the surface, a value from the run's log that the view does not print. (4) The strip of annular pressure over the ground's limit, the vented annulus at 0.20. (5) The status: the return, a cuttings bed in the hole, and a rod change under way. The torque had reached 91.1 % before the exit (Table 7.6).

Agreement with the literature is limited to forms and published ranges. An independent coding of the Delft equation reproduces the app's limit at 1, 3 and 6 m of cover to 0.1 kPa (stiff clay 314.5, 407.2 and 474.4 kPa, sand 164.2, 580.8 and 932.5 kPa), and the general form tends to the cohesive one as $\phi \to 0$ (3.96 % off at 1°, 0.039 % at 0.01°). For stiff clay at 6 m the app's 474 kPa lies between the rigid-plastic $2\ln(R_p/R_o)$ form at 652 kPa and the cavity limit $c_u[1 + \ln(G/c_u)]$ at 478 kPa [3]. The half-cover plastic radius in clay is the one the cohesive analysis recommends [4], and the far field is isotropic where that analysis treats K_0 explicitly. The Bingham slot gradient equals the drilling field form at 478, 485 and 507 Pa/m for 100, 115 and 168 L/min. Specific energy is Teale's definition [1], 108.52 MPa by hand at 80 kN, 100 rpm, 3 kN m and 1.8 m/min. The mud-to-soil ratios (sand 2, stiff clay 4, gravel 2.5) sit inside ranges found in trade guidance during the app's review, which no peer-reviewed source fixes, so they remain teaching values with the soil rates, curvatures and the rock's ratio of 1.5. No magnitude was fitted to a field record, and none of the returns, rates or locates reported here predicts a real crossing. The cases have not been tried with students, so nothing is claimed about what a learner gains from them.

7.7 Technology in Practice

A real drill reads the same signals and moves the same handles as the layers of Section 7.4. Figure 7.6 draws the low selector, its two overrides and the optimizer with field hardware in place of the schematic's symbols.

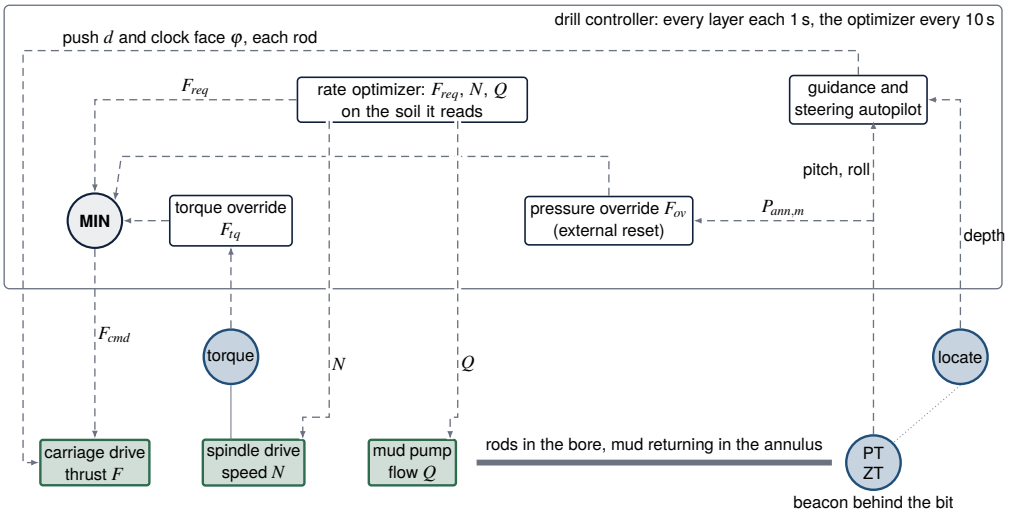

Figure 7.6 The pilot bore's control on real hardware. The beacon behind the bit sends its pitch, roll and annular pressure (PT, ZT) to the drill, and a walkover locate gives the depth where the surface can be walked. Three thrust requests meet in the low selector ahead of the carriage drive: the optimizer's request, the pressure override's limit and the torque override's limit. The optimizer also sets the spindle speed and the mud flow. Notice that both overrides act on the thrust alone and leave the mud flowing.

What Measures It

The constraint of this chapter is read at the head. During the pilot bore a downhole steering probe measures where the bit is against the designed path, and the information gathered includes the survey data, the rate of penetration, the rotary torque of the string and any loss of drilling fluid into the ground [13]. High annular pressure can come from cuttings left in the hole, a collapse or an excessive rate of penetration, and a blocked annulus can over-pressurize and fracture the ground [13], which the loaded annulus of case 6 imitates. The model's beacon stands for the probe and its pressure tool, with the declared pitch bias b of Table 7.1 as its one systematic error.

> **TECHNOLOGY** **Sensor: An Annular Pressure Tool Behind the Bit**
>
> Contractors monitor the annular pressure during the pilot bore with a downhole annular pressure tool [13]. The sensor sits several meters behind the bit as an integral part of the downhole steering system and shows the pressure live at the driller's console [18], and the steering probe that carries it typically sits 6 to 30 ft behind the bit [13]. In a field study a calibrated sensor in the orienting sub read the total annular pressure, the hydrostatic pressure and the friction loss together, and its readings reached a data logger every 10 s [7]. No open source read gives its accuracy. The model's beacon reads the pressure once a second, rounded to 1 kPa, and an interference event holds its last reading for 60 s, a flat line on the trend.

What Moves It

The overrides have one handle, the thrust. When cuttings build up in the annulus, the rotary torque on the string typically rises along with the risk of inadvertent returns [13], so the torque limit and the pressure limit often meet the same event. Both act through the carriage, and neither cuts the pump, because the mud that keeps flowing is what cleans the hole.

> **TECHNOLOGY** **Final Element: The Drill's Push, Turn and Pump**
>
> The drill moves three things from the surface: the thrust on the string, its rotation and the mud pumped down it. Hydraulically powered rigs improved the safety and efficiency of the early drills [13]. Hydraulic drives give a good ratio of power to weight but are inherently less efficient than mechanical or electric ones, and an electric drive needs an inverter that takes space beside its motor [15]. Some utility companies require a breakaway swivel, a mechanical limit on the force used when pulling polyethylene pipe [19]. No open source read gives the response time of a drill's drives. The model's carriage, spindle and pump follow their commands through lags of 1, 2 and 1.5 s.

Where the Controller Runs

A drill's console holds its own controller. A programmable logic controller can be the primary controller of such a smaller system, providing closed-loop control without a central control server, and supervisory and lower loops run with cycle times from milliseconds to minutes [20]. In the application every layer runs each 1 s and the optimizer every 10 s, so the optimizer sits in a supervisory layer above the overrides. Each override is an integrator on the request actually in force. In a PID function block this is the feedback input, once called external reset: connected to the output of the limiting block, it tells a deselected controller that winding up has no effect, and its integration stops [21].

What Is Modernizing It

Two developments bear on the drill and its limits. The first is the electric drive. Hybrid mobile machines can use about 10 to 50 % less fuel than diesel-mechanical or diesel-hydraulic ones, full electric machines could save more, and the limited specific energy of lithium-ion batteries forces a compromise between battery capacity, availability and charging [15]. Case 7 meets that compromise on a shore supply. The second is the downhole probe. Advanced survey probes let contractors determine the path of the pilot hole more quickly and more accurately [13], and a pressure sensor built into the steering system gives early warning of hydraulic fracture in real time [18]. An automatic override needs exactly that reading, and the estimator of Section 7.4 shows what a filter adds to the probe's pitch once the walkover locates run out.

> **IN PRACTICE** **Pressure on a Display Beside Its Limits**
>
> In practice the annular pressure is often shown together with the maximum allowable and minimum required pressures, so that the operator can adjust the drilling quickly and keep it in range. If it leaves that range, the operation should be stopped and investigated [18]. A state permit guide asks that drilling fluid pressures not exceed what the overburden can support, that the bore be plotted from readings on every rod, and that the plan state the pumping capacity, pressures and flow rates [19]. Crews walk the alignment to watch for released fluid [13]. The pressure override of this chapter makes the operator's adjustment automatic.

7.8 Limitations

The drill is a teaching design. The soils, the rate, torque and drag coefficients, the push curvatures, the mud ratios, the bed closure and the safety numbers follow the cited work in sign and form, and their sizes were picked so that every case shows its effect. The annulus is lumped, so a pack-off is spread over the whole hole rather than local, and the bed's cubic factor overstates its effect on a yield-dominated friction. The cavity limit is the isotropic Delft form at one soil per node, so layered ground, the free face of the exit pit and fractures are absent, and the allowable radius is a declared fraction of the cover. The optimizer's soil models are the plant's own soil table, so its model error is only what the soil read, the bed and the learned offsets leave. Reaming, pullback, mud recycling, losses into the ground and diesel-hydraulic drills are outside the model. The challenge's limits, the overrides and the stop under the river exist to teach. They are not certified trip logic, and nothing simulated here can stand in for a crossing's design, its permit or a safety case.

OPEN CHALLENGE **Both Edges of the Pressure Window Are Uncertain**

The methods used in industry predict both edges of the window poorly. The Bingham plastic model, the usual estimate of the annular pressure, overestimates it, and the Delft cavity expansion method, the usual estimate of the allowable pressure, significantly overestimates that and can lead to hydraulic fracture of the ground [7]. Ground uncertainty makes both pressures hard to predict, and an extensive set of data from field projects and laboratory tests is recommended [18]. Research to measure actual pressures on projects and compare them with the predicted limits is encouraged [5]. This chapter uses both models, so progress would show directly in its override's target.

7.9 Conclusion

This chapter simulated a battery-electric utility drill on a steered pilot bore through changing soils, with the annular mud pressure tracked along the whole bore against a cavity-expansion limit that follows the cover. A pushed rod in clay turns the bore 9.4 points while cutting at half the rotating rate, so steering is a sampled duty cycle that costs rate. A different limit holds the rate in each soil, torque in clay, the mud ratio in sand and thrust in gravel, and on a shore supply the optimizer gains 6 % by cutting the flow and spending the power on thrust. The pressure risk peaks on the exit climb, where the river crossing as shipped breaks out at 173.1 m with a ratio of 0.996 a moment before, while the suggested layers exit 0.60 m from the target in 2.13 h with a worst ratio of 0.73. The model is lumped and its coefficients are teaching values. A fuller model would add a distributed bed, layered ground in the cavity limit and a reaming pass. The crossing gives one machine on which constrained rate optimization, guidance between sparse fixes and a pressure window that moves along the path are practiced together.

KEY POINTS

- The drill cuts only as fast as its tightest limit allows, and that limit changes with the soil: torque in clay, the mud ratio in sand, thrust in gravel.
- The mud pressure carries the friction of the whole annulus, so the risk of a return peaks on the exit climb, where the cover is thin.
- A pressure override and a torque override cap the thrust in a low selector, and the pressure override takes thrust off while the mud keeps flowing.
- On the river crossing the suggested layers exit 0.60 m from the target in 2.13 h with no return, while the shipped settings break out at 173.1 m.

Where else this applies. Chapter 11 builds the limits on its advance the same way, and Volume 10, Chapter 2 pushes an oil well's rate of penetration against an annular pressure window.

References

[1] 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.

[2] 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.

[3] A. S. Vesić, "Expansion of cavities in infinite soil mass," *Journal of the Soil Mechanics and Foundations Division*, vol. 98, no. 3, pp. 265–290, 1972.

[4] H. W. Xia and I. D. Moore, "Estimation of maximum mud pressure in purely cohesive material during directional drilling," *Geomechanics and Geoengineering*, vol. 1, no. 1, pp. 3–11, 2006.

[5] D. Bennett and K. Wallin, "Step by step evaluation of hydrofracture risks for horizontal directional drilling projects," in *Pipelines 2008*. American Society of Civil Engineers, 2008, pp. 1–10.

[6] J. S. Lueke and S. T. Ariaratnam, "Surface heave mechanisms in horizontal directional drilling," *Journal of Construction Engineering and Management*, vol. 131, no. 5, pp. 540–547, 2005.

[7] A. Rostami, "Prediction and evaluation of annular pressure in horizontal directional drilling," Ph.D. dissertation, University of Alberta, Edmonton, AB, 2017.

[8] H. Lan and I. D. Moore, "New design equation for maximum allowable mud pressure in sand during horizontal directional drilling," *Tunnelling and Underground Space Technology*, vol. 126, p. 104543, 2022.

[9] E. N. Allouche, S. T. Ariaratnam, and J. S. Lueke, "Horizontal directional drilling: Profile of an emerging industry," *Journal of Construction Engineering and Management*, vol. 126, no. 1, pp. 68–76, 2000.

[10] S. T. Ariaratnam and E. N. Allouche, "Suggested practices for installations using horizontal directional drilling," *Practice Periodical on Structural Design and Construction*, vol. 5, no. 4, pp. 142–149, 2000.

[11] American Society of Civil Engineers, *Pipeline Design for Installation by Horizontal Directional Drilling*. Reston, VA: American Society of Civil Engineers, 2024.

[12] M. E. Baumert and E. N. Allouche, "Methods for estimating pipe pullback loads for horizontal directional drilling (HDD) crossings," *Journal of Infrastructure Systems*, vol. 8, no. 1, pp. 12–19, 2002.

[13] American Petroleum Institute, "Horizontal directional drilling (HDD) operations," White paper, American Petroleum Institute, Washington, DC, https://www.api.org, 2019.

[14] F. Patino-Ramirez, C. Layhee, and C. Arson, "Horizontal directional drilling (HDD) alignment optimization using ant colony optimization," *Tunnelling and Underground Space Technology*, vol. 103, p. 103450, 2020.

[15] A. Lajunen, P. Sainio, L. Laurila, J. Pippuri-Mäkeläinen, and K. Tammi, "Overview of powertrain electrification and future scenarios for non-road mobile machinery," *Energies*, vol. 11, no. 5, p. 1184, 2018.

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

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

[18] B. Gao, "Pressure management of drilling fluid and study of overcut ratio in horizontal directional drilling," Master's thesis, University of Alberta, Edmonton, AB, 2018.

[19] California Department of Transportation, "Encroachment permits: Guidelines and specifications for trenchless technology projects," California Department of Transportation, Sacramento, CA, https://dot.ca.gov/-/media/dot-media/programs/traffic-operations/documents/trenchless-booklet-a11y.pdf, Aug. 2018.

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

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

7.10 Exercises

7.1 Concept check · 2 min

In scenario 2, Follow the plan, the steering autopilot aiming 9 m ahead reached the exit in 0.80 h with 14 of 20 rods pushed. Aiming 3 m ahead it held the plan more closely, within 0.11 m, but took 1.04 h with 17 of 19 rods pushed. Why did the short look-ahead take longer?

(a) It bent rods past their radius, and every bent rod had to be drilled again

(b) It lowered the thrust so that each correction could be placed more precisely

(c) It pushed more rods, and a pushed bit cuts this clay at about half the rotating rate

(d) It kept the walkover crew locating every rod, and each locate held the drill

7.2 Concept check · 2 min

In scenario 4, Sand, the rate optimizer holds the bit on the mud-ratio edge at about 5.4 m/min on the level run, with the torque peaking at 32 % and thrust to spare. Which change would let the drill cut this sand faster without leaving cuttings in the hole?

(a) More thrust, since the rate saturates only slowly in the thrust that reaches the bit

(b) A spindle speed closer to the sand's best speed of 160 rpm

(c) A higher torque ceiling for the optimizer, since the torque decides the cut

(d) More mud flow, which here needs a bit with larger jets, as the pump already runs at the jets' limit

7.3 Concept check · 1 min

On the river crossing as shipped, at 6080 s the annular pressure is 0.20 to 0.31 of the ground's limit under the river but 0.996 of it at 172.6 m on the exit climb. Why is the exit climb the dangerous place?

(a) The pressure there carries the friction of the whole annulus back to the entry, while the limit falls with the thin cover

(b) The river's water lowers the ground's limit under the river, so the returns lose pressure before they reach the climb

(c) The pump pressure rises with every rod added, and the bit's jets pass that rise into the annulus at the head

(d) The gravel lens behind the head forces the returns upward, so the mud is heaviest on the climb

7.4 Theory and calculation · 5 min

A whole 3 m rod pushed at 12 o'clock bends the bore on an arc of radius $1/\kappa_s$, Equation (7.1), and the console shows the pitch as $100\tan\theta$.

(a) In stiff clay ($1/\kappa_s = 32$ m), starting from -14.2 %, find the pitch after one fully pushed rod and the change in points. Compare with the 9.40 points measured in case 1. *Checkpoint:* -4.73 %, 9.47 points

(b) The autopilot caps the push fraction at $0.9/(R_{min}\kappa_s)$ with $R_{min} = 30$ m. Find the cap in stiff clay and in the soft silty clay cap ($1/\kappa_s = 22$ m). In which soil would a full push bend a rod past its radius? *Checkpoint:* 0.96 and 0.66, the cap

(c) In the soft rock of case 8 ($1/\kappa_s = 2000$ m), find the change of a level pitch after one fully pushed rod, and the change per rod that the plan's 50 m bend needs. *Checkpoint:* 0.15 and 6.0 points

7.5 Theory and calculation · 6 min

The mud-ratio constraint of Equation (7.11) bounds the rate at ROP $\le Qc/[A_h(1-c)]$ with $c = (1-m)\eta/(1+R)$, $\eta = v/(v+v_{set})$ and $v = Q/A_{ann}$. The bit is 115 mm and the rods 60 mm, A_h is the area of the hole and A_{ann} that of the annulus, and the margin is $m = 0.10$. Use Table 7.2.

(a) In sand at $Q = 167$ L/min, find v, η, c and the rate ceiling. *Checkpoint:* 0.368 m/s, 0.891, 0.267, 5.87 m/min

(b) Find the ceiling at the 100 L/min with which case 4 ships, and compare it with the 8.32 m/min the bit cut there. *Checkpoint:* 3.20 m/min, 2.6 times

(c) In stiff clay at the optimizer's 137 L/min, find the ceiling. Why does the torque, and not the mud, hold the clay at 2.65 m/min in Figure 7.2a? *Checkpoint:* 2.78 m/min

7.6 Theory and calculation · 6 min

The console dead reckons the depth from the beacon's pitch, Equation (7.8). A constant pitch bias b adds about b meters of depth error per meter drilled on a nearly level bore.

(a) In case 3 the beacon reads 1 % high, and with Locate reset the last locate before the road is at 20.8 m and the first after it at 47.8 m. Estimate the depth error just before that locate and compare it with the 0.246 m measured. *Checkpoint:* about 0.27 m

(b) On the river crossing the bias is 0.3 % and, with no far-side locates, about 125 m is drilled after the last locate. Find the depth error and how far it moves the point where the 12 % exit tangent

breaks the surface. Compare with the shipped exit miss of 2.54 m. *Checkpoint:* 0.375 m, about 3.1 m

(c) The estimator learned a bias of 0.90 % before the road. What depth error does the unlearned rest build across the road, and why was the estimator still up to 0.141 m off? Take the road's bore 3 m deep. *Checkpoint:* 2.7 cm, locate error 9.5 cm

7.7 App challenge · 8 min

Scenario 4, Sand: the rate map and the mud-ratio edge, default seed, at the 60 times real time the case loads with, with the unit switch in the header set to SI. Each run pauses when the head breaks the surface.

(a) Run the case as loaded to the exit. Read the Average rate tile and the Inadvertent returns tile. *Checkpoint:* 126.0 m/h, one return at 114.6 m

(b) Set *Rate optimizer* On, press Reset and run to the exit. Read the same two tiles, and the limit the Drilling rate card names on the level run. *Checkpoint:* 108.4 m/h, none, worst ratio 0.71, the mud

(c) Read the Energy per meter tile at the end of each run. Why does the optimizer's bore cost more energy per meter? *Checkpoint:* 0.13 and 0.24 kWh/m

Tolerance: 0.2 m/h on the rates, 0.2 m on the station, 0.01 on the ratio and the energy.

7.8 App challenge · 9 min

Scenario 6, Shallow cover at the exit: the pressure override, default seed, at the 60 times real time the case loads with, unit switch on SI. Run each part to the exit, where the run pauses.

(a) Run the case as loaded. Read the Inadvertent returns tile and the Average rate tile. *Checkpoint:* one at 149.8 m, 105.9 m/h

(b) Set *Pressure override* On, press Reset and run again. Read the same tiles. *Checkpoint:* none, worst ratio 0.77, 44.2 m/h

(c) Reload the case, set *Rate optimizer* On with the override Off, press Reset and run again. Read the same tiles. Which layer would you commission first on a real drill, and what is the other one for? *Checkpoint:* none, worst ratio 0.76, 76.6 m/h

Tolerance: 0.2 m/h, 0.2 m and 0.01.

7.9 App challenge · 8 min

Challenge, The 180 m river crossing, default seed, at the 600 times real time the case loads with, unit switch on SI. Each run pauses at the exit.

(a) Run the case as loaded. Read the Challenge target tile. *Checkpoint:* not met: exit 2.54 m, return, 91.1 %

(b) Press Reveal suggested tuning, which applies the suggested settings, then press Reset and run again. Read the Challenge target tile. *Checkpoint:* met: 0.60 m, 2.13 h, cover 4.38 m

(c) Keep the suggested settings but set *Far-side locates* Off, press Reset and run again. Which target fails, and why? *Checkpoint:* exit 1.70 m

Tolerance: 0.05 m, 0.02 h and 0.2 %.

7.10 Application area · 3 min

The optimizer and the override keep the annular pressure under 0.97 of the allowable P_{lim}/FoS, with $\mathrm{FoS} = 1.3$.

(a) What fraction of the cavity limit is that? Over ten seeds the suggested crossing's worst ratio ran from 0.72 to 0.77. Is that consistent? *Checkpoint:* 0.746

(b) Where would a real crossing take its factor of safety, its minimum cover under the river and its allowable plastic radius from? Name two features of real ground that lower the limit and that the model leaves out.

7.11 Application area · 4 min

A crew is paid for meters per hour with every rod change included. Each 3 m rod costs 50 s of rod change.

(a) Find the rate with rod changes at a cutting rate of 8.32 m/min and at 5.36 m/min, the shipped and optimized cutting rates of case 4. By how much does the faster bit raise the average? Compare with the measured 126.0 and 108.4 m/h. *Checkpoint:* 150.8 and 129.2 m/h, 17 %

(b) Faster rod handling cuts the rod change to 35 s. Which gains more in percent, the fast or the slow bit, and what does that say about where a crew should look for time? *Checkpoint:* 26 % and 22 %

7.12 Review the agent · 5 min

An AI assistant reviewing case 6 writes: “The annular pressure on the exit climb is mostly friction, and friction grows with flow. Cut the mud flow from 100 to 60 L/min as the head starts to climb: the friction gradient falls roughly in proportion, and the inadvertent return goes away without slowing the bit.” Find the flaw, and say how you would check the claim with the chapter’s equations and with the app.

CHAPTER 8

Growth Control Against a Cash Floor in a Simulated Product Startup

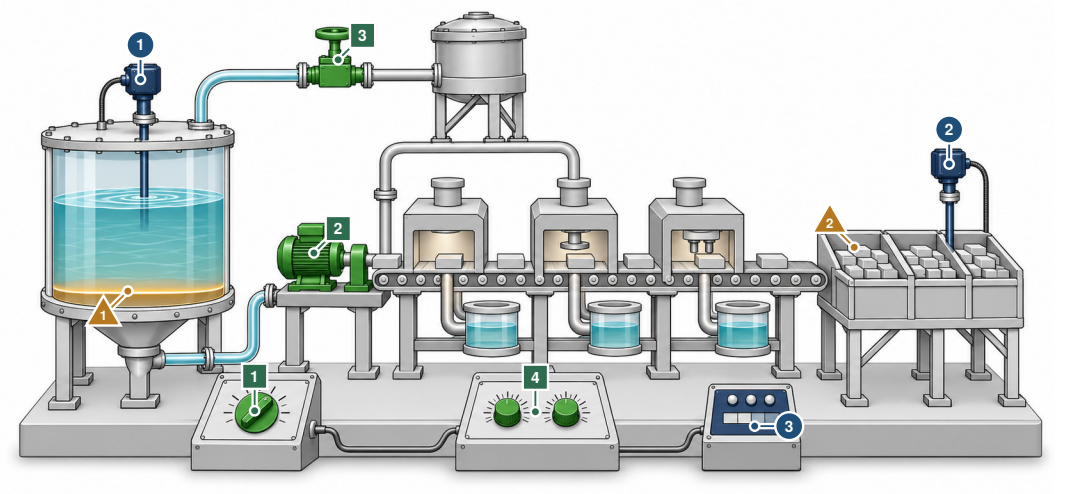

The controller moves 1 credit line draw 2 production rate 3 capital spending on the factory 4 research and marketing budgets

The controller reads 1 cash in the bank 2 weeks of stock on hand 3 the books at each quarterly close

What disturbs it 1 customers paying late 2 a recession or a rival's launch

The cash tank pays for the factory, the stock and the budgets before any sale pays it back, and the amber band near its bottom is the floor that the runway override defends. Illustration, not to scale.

IN PLAIN TERMS **Growing a Company Without Running Out of Cash**

The main idea. A young company pays for stock, a bigger factory and research long before its customers pay it back. So it can show a profit every quarter and still run out of money.

The challenge. Spending more makes the company grow faster, but every bit of growth is paid for first and earned back later. When the cash runs low, something has to give way, and the order in which things give way matters.

Why it matters. A rule written in advance, which cuts marketing and building first, borrows next and cuts research last, lets the company live through late payers and bad years.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on funded startups (Section 8.1), the history segment (Section 8.2), the two try-it segments and the fact on the runway target (Section 8.6) and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the measured cases (Section 8.6), Technology in Practice (Section 8.7), Limitations (Section 8.8) and the open challenge

The control question A product startup can book a profit every quarter and still run out of money, because its stock, its factory and its unpaid invoices are paid for before customers pay. How fast can loops, rules and an economic planner grow such a company before the floor on its cash stops them, and what should give way first when the cash runs short?

Learning objectives

- Account for where the cash of a profitable company goes, using its balance sheet and the delays between a decision and its return.
- Derive Higgins's sustainable growth rate on trailing books and the capital spending that grows a factory at that rate through its build.
- Explain how a runway override split over three bands gives up marketing and investment, then credit, then research.
- Measure what the season, a capacity held back, a late payer and a competitor launch cost the company in the app.
- Judge an economic MPC's financing verdict against its own modeling choices.

Prerequisites Sections 1.4.3 (selectors), 1.4.4 (external reset), 1.4.6 (split range), 1.4.9 (actuators that deplete) and 1.6.2 (velocity form). Section 1.6.3 is needed only for the planner.

The app https://apmonitor.com/apps/startup-growth/

See also Volume 9, Chapter 8 for setpoints chosen on a cost whose optimum sits on a moving constraint and Volume 9, Chapter 6 for a predictive controller that prices a shared resource. Volume 3, Chapter 6 for feedforward with a slow ratio trim and Volume 3, Chapter 4 for selectors that protect a ratio.



8.1 Introduction

Cash is the one stock of a company that has a floor, and every decision that grows the company draws on it before it pays anything back. In the simulator of this chapter a board that spends 15 % of revenue on research and development (R&D) and 15 % on marketing, asks its factory for 40 thousand units a year and adds \$1.5M of factory a year is insolvent after 1.58 yr. At that instant its trailing twelve-month revenue is \$8.37M, up from \$7.26M at the start, its trailing earnings before interest, taxes, depreciation and amortization (EBITDA) are \$1.11M and its net income \$0.51M. A month earlier, at 1.5 yr, its money sits in 26.1 weeks of finished stock carried at \$1.46M, a factory of 42.4 thousand units a year carried at \$4.24M and \$0.92M of receivables. The income statement shows a healthy company and the cash account shows a dead one.

Management science has written the firm as a feedback system for seventy years. Production and work-force scheduling was posed in 1955 with a quadratic approximation to the cost criterion, whose optimum is a linear decision rule [1]. Sales response to advertising was modeled two years later with a decay constant, a saturation level and a response constant [2], and advertising policy was treated as an optimization under dynamic conditions in 1962 [3]. Forrester's market growth model, a system dynamics model of a growing company, has been read as a portrait of bounded rationality, built from rules of thumb that use limited information [4]. Corporate finance supplied the bound. The growth a firm can sustain without new equity at its chosen payout and capital structure is a rate [5], and a profitable company that tries to grow faster than its cash cycle allows can run out of cash [6]. Economic model predictive control (MPC), whose stage cost is a cost in money rather than a tracking error [7], completes the vocabulary. The control frame suits this plant because that literature already writes each of its decisions as a loop, a rule or an optimization.

The structure is general. Any operation that grows by investing from a stock with a floor, and earns its return only after a delay, has this shape, as a hospital adding beds from its operating surplus, a utility building plant from its rate base and a university growing enrollment against its endowment draw all do. Its control problem combines integrating stocks, delays of a year, ratio control of spending to a measured base and an override that protects a hard constraint by taking the actuators away from the loops that normally drive them, all standard in process control [8]. Section 1.4.9 treats such a

stock as a depleting actuator whose runway is itself a limit. The belief the simulator corrects is that profit measures solvency and that sales measure the market. In its fourth case a factory held at $0.4M of investment a year sells $8.97M in year 8 while its customers ask for 2.65 times what it sells, and the growth loop, which sees only sales, raises marketing from 7 % to 20 % of revenue.

> **DID YOU KNOW? Most Funded Startups Still Shut Down**
>
> Census Bureau records show that of new U.S. firms that received venture capital from 1986 to 1997, 75 % had shut down by 2007, against 66 % of matched firms without it. The funded survivors employed 364 % of their group's starting employment, the matched survivors 67 % [9].

However, the finance literature states its bound for a company already growing steadily at constant ratios [5, 6], the system dynamics literature describes the feedback of the firm in decision rules and delays rather than in the setpoints, overrides and constraint handling of a control course [4, 10], and economic MPC is developed for process plants [7], so a student seldom has one company on which the bound, the loops and the planner can be measured together while the calendar, the rule and the financing change one at a time. This chapter quantifies how far feedback loops, two practitioner rules and an economic planner can grow a simulated product startup, whose every decision drains the same cash tank, before its cash floor stops them. The market model, the sustainable growth rate, the production rule, the override structure and economic MPC all come from the literature cited. The measurements are the chapter's own, among them an exact form of Higgins's rate for books kept on trailing twelve months, a sweep that bounds growth for every budget split that earns money, and the sensitivity of the planner's equity round to its own modeling choices. The company has 40 states, and its balance sheet, cash-flow statement and retained earnings reconcile at every step. Each of nine cases, a challenge and a reference isolates one mechanism.

Section 8.2 traces the history, from the learning curve and the linear decision rule to the venture round. Section 8.3 writes the company as stocks and flows, with each equation, constant, lag and bound it uses. Section 8.4 explains why its growth costs cash, through Higgins's rate on trailing books, a sweep of budget splits and the three delays. Section 8.5 gives the board's calendar, its loops, the production rule, the runway override, the rules that set the growth target and the economic MPC. Section 8.6 reports its eleven runs, the cases, the challenge and the reference. Section 8.7 sets the override beside a company's real accounts, budgets and credit line. Section 8.8 states the limitations and Section 8.9 summarizes. The exercises in Section 8.10 run the cases in the app.

8.2 From Learning Curves and Market Growth to the Venture Round

The cost side of the company has the oldest record. The factors that set the cost of building airplanes were analyzed in 1936 [11], the article the application cites for the learning curve on which its unit cost falls with cumulative production. Nineteen years later a research team working with a manufacturing company formalized the setting of a factory's aggregate production rate and work force with a quadratic approximation to the criterion function, and computed an optimal decision rule that is linear in the state of the factory, a rule that, unlike a rule of thumb, can claim to be the best possible in a stated sense [1]. A quadratic cost with a linear optimal rule has the form of the linear quadratic regulator of later control texts, here applied to inventories and payrolls.

Demand acquired dynamics in the same decade. Studies for major industrial concerns described the interaction of advertising and sales with three parameters, a sales decay constant, a saturation level and a response constant, consistent with controlled experiments on a large number of products and several media [2], and an optimal advertising policy under dynamic conditions followed in 1962 [3]. The diffusion of a new product was written in 1969 as a model for the timing of a consumer's first purchase, related to the number of previous buyers through innovative and imitative behavior, and tested against eleven consumer durables with a long-range forecast for color television sets [12]. A generalization added price and advertising as decision variables, reduced to the original model

as a special case, and explained why the original fits so well without them [13]. The market of the simulator is that model with marketing on the innovation coefficient and a logit share against one competitor.

The firm as a whole became a feedback system in the work of Forrester and his students. A 1983 reading of Forrester's market growth model argued that the structure of system dynamics models assumes bounded rationality, with decision functions made of simple rules of thumb and limited information content [4]. The growth trap of the fourth case, a company whose underinvestment produces the evidence for underinvestment, is the behavior the application's briefing attributes to that model. In 2002 a field study of process improvement found that delays between investing in improvement and recognizing its rewards, together with managers' attributions about the cause of poor performance, can trap an organization in a cycle that prevents useful change [10], the mechanism of the fifth case, where an R&D cut pays at once and costs share eighteen months later.

Corporate finance asked the question of this plant directly in 1977, and the runway override of Section 8.5 gives up its actuators in the order of the pecking order theory that the segment below follows.

> ⧖ HISTORY **1977 to 2001: How Fast a Company Can Grow on Its Own Cash**
>
> In 1977 corporate finance defined a sustainable growth rate, the yearly growth in sales a company can reach while it keeps its target payout ratio and capital structure and issues no new equity. Faster growth forces something in its financial objectives to give, usually to the detriment of financial soundness, and slower growth lets it raise dividends, reduce its borrowing or build up liquid assets [5]. A cash conversion cycle approach to liquidity analysis followed in 1980, because inattention to liquidity can cause severe difficulties even for a firm with favorable long-run prospects [14]. The pecking order theory of 1984 held that firms prefer internal to external funds, and debt to equity when they need outside money [15]. In 2001 a formula written for managers joined the operating cash cycle, the cash needed to finance each dollar of sales and the cash each dollar of sales generates into a self-financeable growth rate [6].

The equity round is the last decision the simulator adds. Entrepreneurship has been described as experimentation whose probabilities of success are low, extremely skewed and unknowable until an investment is made, with a few investors deciding which experiments continue [9]. A survey of 885 institutional venture capitalists at 681 firms found that they weigh the management team above the product or technology and rate deal selection as the most important source of the value they create [16]. On the control side, economic MPC was shown to have an average performance never worse than the best steady state [7], which is the property an exit objective gives up, because an acquisition at a fixed date is a terminal reward and not a steady state. The simulator's planner is a local, successively linearized version of that idea with the round as one more decision.

8.3 The Company as Stocks and Flows

The simulated company makes one hardware product and sells it to consumers against one competitor, drawn in Figure 8.1 as the tanks, pipelines and valves that it is. Money is in millions of dollars ($M), volume in thousand units (k units), time in years, and each flow is a yearly rate. Table 8.1 classifies the signals. All coefficients are the teaching values of Table 8.2, chosen for an imaginary hardware startup and not estimated for any firm. Every stock except the market's is paid for out of the cash tank, and every decision of Section 8.5 opens or closes one of the valves that fill or drain it.

8.3.1 Market

Would-be adopters arrive by Bass diffusion of the company's own installed base N toward an effective potential, and choose this product with a logit share in the quality gap and the price gap, as shown in

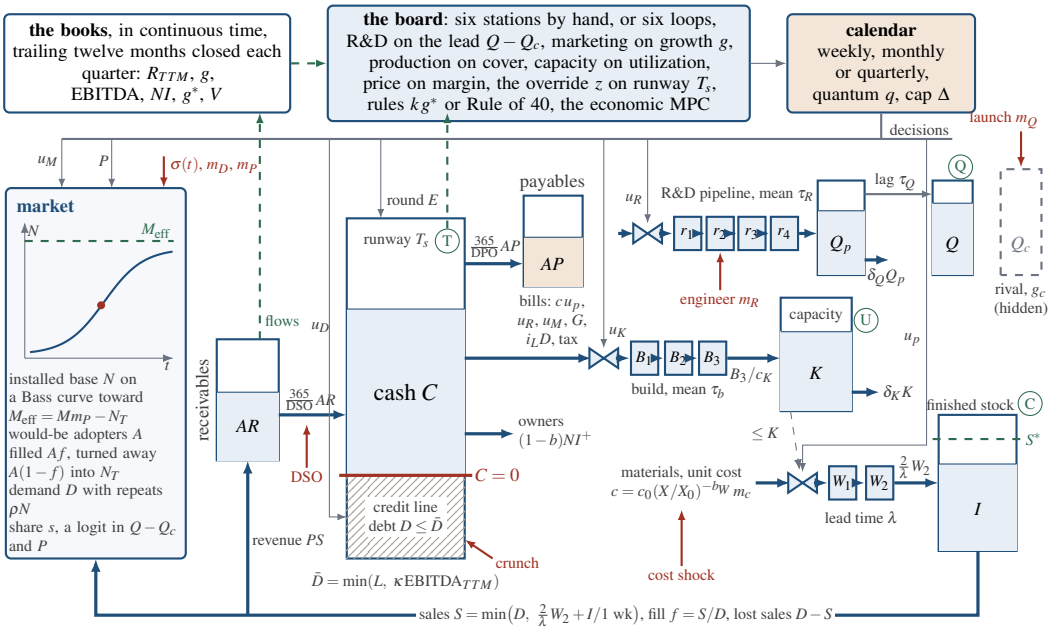

Measured live: T runway, U utilization, C cover, Q lead. The spend shares u_R and u_M are paid as shares of R_{TTM} at the last close; every station, by hand, by a loop or by the MPC, moves only on the calendar and on its grid.

Figure 8.1 The startup as stocks and flows, not to scale. Collections from the receivables AR fill the cash tank C, whose floor is the thick red line marked $C = 0$ and below which the hatched credit line holds the debt D within its ceiling $\bar{D}$. Payments of the payables AP, capital spending u_K and dividends drain it. Capital passes a three-stage build chain B_j into capacity K, which bounds the production u_p that passes the work in process W_1, W_2 into finished stock I, sold as S against demand D and booked as revenue PS. R&D spending passes a four-stage pipeline r_j into the quality potential Q_p and the product quality Q, raced against the competitor's hidden Q_c, drawn as a dashed box. Thick blue arrows carry money and goods, red arrows that start at a labeled cause are disturbances, thin gray arrows from the decision bus are the board's decisions on its calendar, and dashed green lines are the measurements and the books, with the circled letters keyed in the line under the drawing.

Table 8.1 Signals of the Company by Role

Symbol	Role	Unit
T_s, g, $I/(D/52)$	CVs, stressed runway, growth at the close, stock cover	mo, –, wk
$\hat{u}_p/K$, ℓ, m_g	CVs, utilization asked, quality lead, gross margin	–, points, –
u_R, u_M	MVs, R&D and marketing as shares of revenue, 0–0.3	–
u_p, u_K	MVs, production, 0–K, and capital spending, 0–20	k/yr, $M/yr
P, u_D	MVs, price, 210–390, and credit draw, −4–4	$/unit, $M/yr
b, E	MVs, retention and an equity round	–, $M
σ, m_D, m_P	DVs, season, demand and potential multipliers	–
m_Q, m_c, m_R	DVs, launch, supplier cost and engineer multipliers	–
DSO, m_L, i_L	DVs, collection days, credit line multiplier, rate	d, –, /yr
C, AR, AP, D	states, cash, receivables, payables, debt	$M
I, W_1, W_2, X	states, stock, work in process, cumulative output	k units
$I^\$$, $W^\$$, F	states, stock and work in process at book, fixed assets	$M
E_P, E_R	states, paid-in and retained equity	$M
K, B_1–B_3	states, capacity and the build chain	k/yr, $M/yr
Q, Q_p, r_1–r_4	states, quality, its potential, the R&D pipeline	–, $M/yr
N, N_T	states, installed base and customers turned away	k units
Σ_x	states, 15 cumulative flows the books close on	$M, k
S, D	measured sales, hidden demand	k/yr
x_{TTM}, V, ω	outputs, trailing books, valuation, founders' share	$M, –

Equation (8.1).

$$\begin{aligned} M_{\text{eff}} &= \max(1, Mm_P - N_T), \quad p = p_0 + \frac{p_M e_M}{M_s + e_M} \\ s &= \left[1 + e^{-\beta_Q(Q-Q_c)+\beta_P(P-P_c)/P_c}\right]^{-1} \\ A &= \left[(p + qN/M_{\text{eff}})(M_{\text{eff}} - N)\right]_+ s\,\sigma\, m_D \\ D &= A + \rho N s\,\sigma\, m_D \\ \dot{N} &= Af, \qquad \dot{N}_T = A(1-f) \end{aligned} \tag{8.1}$$

where $\sigma = \sigma(t)$ is the season, e_M the marketing spend in \$M/yr, $f = S/D$ the share of orders filled, N is held within $[0, M_{\text{eff}}]$ and the exponent of the logit within ± 30. Marketing raises the innovation coefficient by at most p_M, half of it at $e_M = M_s$. A would-be adopter turned away by an empty shelf buys the competitor's product and leaves the potential for good, while a repeat buyer turned away does not defect. The season is a four-term fit through mid-quarter values of 0.7, 0.85, 0.85 and 1.6, as shown in Equation (8.2).

$$\begin{aligned} \sigma(t) &= 1 + a_\sigma\big(0.375\cos\theta - 0.075\sin\theta + 0.225\cos 2\theta\big) \\ \theta &= 2\pi(t - 0.875) \end{aligned} \tag{8.2}$$

At $a_\sigma = 1$ it peaks at 1.602 at $t = 0.866$ of the year, in mid-November, falls to 0.628 at 0.187, in early March, and has quarter means of 0.789, 0.806, 0.924 and 1.481, with an annual mean of exactly one.

8.3.2 Product

R&D money passes a four-stage pipeline with a mean delay τ_R, raises a quality potential with diminishing returns, and the product on sale follows the potential through a lag, as shown in Equation (8.3).

$$\begin{aligned} \dot{r}_1 &= \tfrac{4}{\tau_R}(e_R m_R - r_1) \\ \dot{r}_j &= \tfrac{4}{\tau_R}(r_{j-1} - r_j), \qquad j = 2,3,4 \\ \dot{Q}_p &= \eta_R \frac{(1 + 1/R_s)\, r_4}{1 + r_4/R_s} - \delta_Q Q_p \\ \tau_Q \dot{Q} &= Q_p - Q, \qquad Q_c = (1 + g_c t)\, m_Q \end{aligned} \tag{8.3}$$

where e_R is the R&D spend in \$M/yr and $m_Q = 1.15$ after a competitor launch. At \$1M a year the yield is $\eta_R = 0.1$ quality point a year, at $R_s = \$3M$ a year it is 0.2, half of the 0.4 that no spend can exceed. Near parity a lead of 0.1 point moves the share by 0.124.

8.3.3 Factory and Production Line

Production starts into a two-stage work-in-process line and finishes into stock, a sale ships from stock within a week, and capital spending becomes capacity through a three-stage build, as shown in Equation (8.4).

$$\begin{aligned} \dot{W}_1 &= u_p - \tfrac{2}{\lambda} W_1, \qquad \dot{W}_2 = \tfrac{2}{\lambda}(W_1 - W_2) \\ \dot{I} &= \tfrac{2}{\lambda} W_2 - S, \qquad S = \min\big(D,\ \tfrac{2}{\lambda} W_2 + I/\tau_w\big) \\ \dot{K} &= B_3/c_K - \delta_K K, \qquad 0 \le u_p \le K \\ \dot{B}_1 &= \tfrac{3}{\tau_b}(u_K - B_1), \qquad \dot{B}_j = \tfrac{3}{\tau_b}(B_{j-1} - B_j) \\ c &= c_0\big(\max(X, X_0)/X_0\big)^{-b_W} m_c, \qquad \dot{X} = u_p \end{aligned} \tag{8.4}$$

where $\tau_w = 1/52$ year and $j = 2,3$ in the build chain. A sale the stock cannot make is lost, $D - S$. The exponent $b_W = 0.2345$ makes unit cost fall 15 % with each doubling of cumulative production, $2^{-b_W} = 0.850$. Production is booked into work in process at the unit cost when it starts, $\dot{W}^{\$} =$

$cu_p/10^3 - \frac{2}{\lambda}W_2 w$ with $w = W^{\$}/(W_1 + W_2)$, moves into stock at that average, $\dot{I}^{\$} = \frac{2}{\lambda}W_2 w - S\iota$, and is sold at the book cost of the units available over the shipment window,

$$\iota = \frac{I^{\$}/\tau_w + (2/\lambda)W_2 w}{I/\tau_w + (2/\lambda)W_2}, \qquad \text{COGS} = S\iota.$$

When the denominator is zero, sales and COGS are zero. Since sales do not exceed the available unit flow, COGS cannot exceed $I^{\$}/\tau_w + (2/\lambda)W_2 w$. Thus a zero book inventory has a nonnegative derivative, without a write-off or an artificial clipping flow. This law replaced a small-stock blending law on 2026-11-15.

8.3.4 Statements and Financing

Revenue is booked when the product ships and collected after DSO days, every cost is booked when it is incurred and paid after DPO days, and the cash tank takes what is left, as shown in Equations (8.5) and (8.6).

$$\begin{aligned} \text{Rev} &= PS/10^3, \qquad NI = \Pi - \tau_x[\Pi]_+ \\ \text{EBITDA} &= \text{Rev} - \text{COGS} - (e_R + e_M + G) \\ \Pi &= \text{EBITDA} - \delta_K F - i_L D \end{aligned} \tag{8.5}$$

$$\begin{aligned} \dot{AR} &= \text{Rev} - \tfrac{365}{\text{DSO}}AR \\ \dot{AP} &= cu_p/10^3 + e_R + e_M + G + i_L D + \tau_x[\Pi]_+ \\ &\quad - \tfrac{365}{\text{DPO}}AP \\ \dot{C} &= \tfrac{365}{\text{DSO}}AR - \tfrac{365}{\text{DPO}}AP - u_K + u'_D - (1-b)[NI]_+ \\ \dot{F} &= u_K - \delta_K F, \qquad \dot{D} = u'_D \\ \dot{E}_R &= NI - (1-b)[NI]_+ \end{aligned} \tag{8.6}$$

where Π is profit before tax, taxed at τ_x with no loss carried forward. The credit line holds debt within a ceiling that moves with the company's own earnings, and debt above it is called over a quarter, as shown in Equation (8.7).

$$\begin{aligned} \bar{D} &= \min\big(Lm_L,\ \kappa[\text{EBITDA}_{TTM}]_+\big) \\ u'_D &= \max\big\{\min(u_D, [\bar{D} - D]_+/\Delta t) \\ &\qquad - [D - \bar{D}]_+/0.25,\ -D/\Delta t\big\} \end{aligned} \tag{8.7}$$

where $\Delta t = 1/52$ year and a repayment ($u_D < 0$) passes without the first minimum. An equity round of E adds E to cash and to paid-in equity E_P and multiplies the founders' share ω, 80 % after the seed round, by $V_e/(V_e + E)$, with the pre-money $V_e = \max(1, V + C - D)$. A term loan draws up to the room the ceiling leaves at once. Both are raised after the step in which they are asked for. The balance sheet $C + AR + I^{\$} + W^{\$} + F - AP - D = E_P + E_R$ holds at every step, because each flow enters one asset or liability and its counterpart. Cash may fall below zero. The run then records INSOLVENT and pauses once, and Run carries on the company a bank would have closed, so that it can be watched.

8.3.5 Books, Runway and Valuation

The books are kept as practitioners read them. Every flow has a cumulative state Σ_x, and a ring of those states one step apart gives any trailing twelve-month total as $x_{TTM}(t) = \Sigma_x(t) - \Sigma_x(t-1)$, so the season cancels exactly. Before $t = 0$ the ring holds a flat pre-history at the opening rates with the

Table 8.2 Teaching Values by Location: M Listed in Show Model, C on a Controller Card or in the Disturbances, S Fixed in the Source. No Constant Is Hidden From the Student

Symbol	Meaning	Value	Range	
Market, product, disturbances and clocks				
M	potential adopters	2000 k	1000–3500	M
p_0, q	Bass innovation, imitation	0.01, 0.4 yr	0.002–0.03, 0.1–0.8	M
ρ	repeat purchases	0.2 yr	0.05–0.4	M
p_M, M_s	most marketing adds, spend for half	0.02/yr, \$1.5M/yr	0–0.04, 0.2–4	M
β_Q, β_P	share sensitivity to quality, price	5, 3	2–8, 1–6	M
g_c, P_c	rival's quality growth, price	0.08/yr, \$300	0–0.15, 200–400	M
η_R	quality a year from \$1M/yr	0.1	0–0.3	M
R_s	R&D for half the yield	\$3M/yr	0.5–20	M
τ_R, τ_Q	R&D pipeline mean, product lag	1, 0.5 yr	0.5–2, 0.2–1.5	M
δ_Q	obsolescence	0.02/yr	0–0.1	M
P_0	launch price	\$300	200–400	M
a_σ, m_D	seasonal amplitude, demand multiplier	1, 1	0–1.5, 0.5–1.5	C
—	recession: m_D, m_P, length	0.7, 0.8, 1.5 yr	—	S
—	credit crunch m_L, late DSO	0.5, 75 d	—	S
—	rate rise, replayed spread	+0.03, +0.05/yr	—	S
Δt	step, control period, log	1/52 yr each	—	S
Factory, finance and fixed values				
c_0	unit cost at $X_0 = 20$ k	\$150	80–200	M
b_W	learning exponent	0.2345	0–0.5	M
c_K	capital per k/yr of capacity	\$0.1M	0.05–0.2	M
τ_b, δ_K	build mean, depreciation	0.75 yr, 0.1/yr	0.25–1.5, 0.03–0.2	M
λ	production lead time	6 wk	2–16	M
DSO, DPO	collection, payment days	45, 30 d	15–100, 10–90	M
G	fixed G&A	\$1.5M/yr	0–3	M
τ_x, i_L	tax, credit-line rate	0.25, 0.08/yr	0–0.4, 0–0.2	M
L, κ	credit line, covenant multiple	\$1M, 3	0–4, 1–6	M
C_0, K_0, N_0	opening cash (\$M), capacity (k/yr), customers (k)	3, 25, 20	0.5–6, 10–80, 5–400	M
r	owners' discount rate	0.12/yr	0.05–0.2	M
—	opening stock, Q, X	6 wk, 1, 20k	—	S
—	launch m_Q, supplier m_c for 1 yr	1.15, 1.3	—	S
—	viral m_D, engineer m_R	2 for 0.25 yr, 0.6 for 1 yr	—	S
—	seed share, founders	20 %, 80 %	—	S
—	covenant call period	0.25 yr	—	S

season at its mean. The books then give the quantities of Equation (8.8).

$$
\begin{aligned}
g &= \frac{\mathrm{Rev}_{TTM}(t)}{\mathrm{Rev}_{TTM}(t-1)} - 1, \qquad m = \frac{\mathrm{EBITDA}_{TTM}}{\mathrm{Rev}_{TTM}} \\
g^* &= \frac{x}{1-x}, \qquad x = \mathrm{clip}_{[-0.9,0.9]}\left(\frac{b\,NI_{TTM}}{E_P + E_R}\right) \\
V &= [\mathrm{Rev}_{TTM}]_+ \,\mu(g,m) \\
\mu &= \mathrm{clip}_{[1.5,13]}\left(2 + 10\,\mathrm{clip}_{[0,1]}(g) + 2\,\mathrm{clip}_{[0,0.3]}(m)\right)
\end{aligned}
\tag{8.8}
$$

where g^* is Higgins's sustainable growth rate with the return on equity taken at the end of the window, and V is a teaching valuation at the multiple μ, two times revenue at no growth and twelve at 100 % a year, with the shape of the Rule of 40. The founders' value is $\omega[V + C - D]_+$. At every quarter end a close publishes these figures, and the growth and Higgins rules read the close. Runway is $T_r = 12C^+/B$ months at the trailing year's net burn $B = -\mathrm{FCF}_{TTM}$ plus any debt being called, and the override reads a stressed runway T_s whose burn is never less than a quarter of the operating outflows, $\max\big(B,\ 0.25(\mathrm{COGS}_{TTM} + e_R + e_M + G)\big)$, both held within 0 to 120 months.

8.3.6 Start and Integration

The opening company is lined out at its opening decisions. Its sales are the demand those decisions raise at the mean season, or the capacity if that is less, found by a fixed-point iteration because the spends are shares of the revenue they produce. Stock is six weeks of those sales, work in process a lead time of them, receivables and payables their DSO and DPO shares of the flows, every pipeline stage at its opening input, fixed assets $c_K K_0$, and paid-in equity whatever balances the sheet. At the default decisions the company opens with \$6.57M of trailing revenue, 21.9 thousand units a year at \$300, a trailing EBITDA of \$0.14M, a net income of $-\$0.11M$, \$0.81M of receivables, equity of \$6.54M and a credit ceiling of \$0.42M. The engine of Section 1.6.1 integrates the 40 states with classical fourth-order Runge-Kutta on a step of a week. The board's loops execute once a week ahead of each integration step, and the historian records a sample each week. A seasonal pattern and every event enter through time inside the right-hand side, so nothing is random and a case replays exactly. The default seed, 20261004, is not used by any random draw. Over six cases of six years the balance sheet closes to 1.4×10^{-14} \$M, the cash-flow statement to 7.2×10^{-14} \$M and retained earnings to zero, and the installed base follows the closed-form Bass curve to a relative 2.2×10^{-12} when every order is filled. On the balanced reference every derivative of the balance sheet is below 4.5×10^{-16}, and at half the step the challenge's growth changes from 20.7 % to 20.8 %. The application passes its 22 own checks.

8.3.7 Assumptions

One product, one competitor. The rival's quality rises on a fixed trend and never answers a price cut or a launch, there is one market and one factory, and there is no hiring, no wage bill apart from the spends and the overhead, no bad debt and no obsolete stock.

Delays as distributed lags. The R&D pipeline, the work in process and the build chain are chains of first-order stages, which reproduce a mean delay and a spread, not a delivery calendar.

Simplified finance. Tax has no loss carry-forward, the covenant is the only clause of the loan, and rounds are priced at the equity value with no preferences, options or convertible notes.

Teaching values. Every parameter, threshold, multiple and rule is a declared teaching value. Real diffusion coefficients come from fitted diffusion studies [12], real multiples from comparable transactions, real covenants from the loan agreement, and none of the numbers here is an estimate for any company.

8.4 Why Growth Costs Cash

8.4.1 Higgins's Rate on Trailing Books

Higgins's rate is the growth at which equity, raised only by retained earnings, keeps pace with the assets that the sales require. On books kept in continuous time and closed on trailing twelve months the rate is exact. A company growing at the continuous rate γ with constant ratios has $E(t) = E(0)e^{\gamma t}$ and retains $bNI = \dot{E} = \gamma E$, so its trailing net income is a year of that flow, as shown in Equation (8.9).

$$\begin{aligned} \frac{bNI_{TTM}}{E(t)} &= \frac{1}{E(t)} \int_{t-1}^{t} \gamma E(\tau)\,\mathrm{d}\tau = 1 - e^{-\gamma} \\ g^* &= \frac{x}{1-x} = e^{\gamma} - 1 \end{aligned} \tag{8.9}$$

The trailing revenue of the same company grows year over year by $e^{\gamma} - 1$, so $g^* = g$ holds with no approximation when the ratios hold. Case 2 builds a company in which they nearly do, with no learning curve, no obsolescence, no competitor race, \$0.5M of overhead, no credit line, \$1.5M of cash, and a market of 3.5 million that would buy far more than the factory makes. Growing at $k = 1$ times g^*, it settles onto the rate, 21.9 % against 22.2 % at year 4 and 22.8 % against 23.8 % at year 9, and from year 5 to 9.5 growth stays within 1.6 points of g^* and cash between 38.6 % and 42.5 % of equity, as shown in Figure 8.2(a).

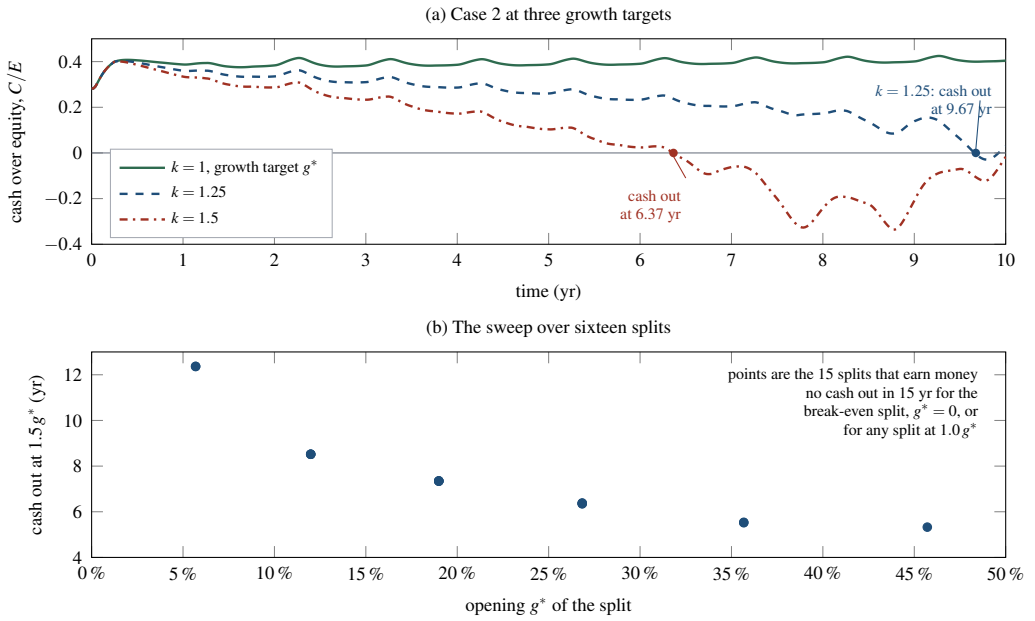

Figure 8.2 Higgins's bound measured on the application's engine. (a) Cash as a share of equity in case 2 with the growth target at 1 (solid), 1.25 (dashed) and 1.5 (dash-dotted) times g^*. (b) The year the cash runs out at $1.5\,g^*$ for each split of R&D and marketing at 5, 10, 15 and 20 % of revenue, against the split's opening g^*. Points of one g^* share a total spend.

8.4.2 The Honest Bound

Growth above g^* must be paid for by debt, new equity or the cash in the tank. At $k = 1.5$ the factory and the working capital outgrow the equity, the cash share falls and first crosses zero at 6.37 yr. At $k = 1.25$ it first crosses zero at 9.67 yr; the later positive balance is a counterfactual recovery after insolvency, not a solvent ten-year run. The application's checks run the same company for all sixteen splits of R&D and marketing at 5, 10, 15 and 20 % of revenue. At $1.5\,g^*$ every one of the fifteen splits that earns money runs out of cash, between 5.33 yr for 5 % and 5 % and 12.37 yr for 15 % and 20 %, as shown in Figure 8.2(b), and at $1.0\,g^*$ none runs out in fifteen years. The sixteenth split, 20 % and 20 %, opens at break-even, so its g^* is zero and it has no growth to finance. The more profitable a split, the higher its g^*, the faster it is asked to grow and the sooner its cash runs out, from an opening g^* of 45.7 % for the leanest split to 5.7 % for the two splits that spend 35 % of revenue.

8.4.3 Lags, Working Capital and the Season

Three delays separate a decision from its return. A step in R&D on the balanced reference reaches half its final rate of quality gain after 1.385 yr, against the 1.5 yr of pipeline and lag declared in Table 8.2, because the four-stage chain spreads the arrival around its mean. Capacity arrives through a build with a mean of nine months, and an order shipped today is cash 45 days later. Receivables are about DSO/365 of a year's revenue, so on the balanced reference a stretch from 45 to 75 days moves \$0.308M from cash into receivables one for one by year 2.5 while EBITDA is unchanged to the last digit. The season multiplies demand by 1.6 in mid-November, and a production rule that reads last quarter's sales meets each peak a season late.

Table 8.3 Board Settings, Defaults and Ranges. A Marks an Advanced Control

Symbol	Setting	Default	Range
Board			
—	decision policy	six loops	4 policies
—	calendar	quarterly	weekly, monthly
q, Δ	quantum, largest move (A)	1, 100 % of range	0–5, 5–100
Loops: setpoint / K_c / τ_I			
—	R&D on lead ℓ	0 / 1.0 / 4 yr	±0.3 / 0–2 / 0.25–20
—	marketing on growth g	0.25 / 0.2 / 3 yr	0–0.8 / 0–2 / 0.25–10
—	capacity on utilization	0.8 / 5 / 1 yr	0.4–1 / 0–20 / 0.25–10
—	price on gross margin (MANUAL)	0.5 / 100 / 1 yr	0.2–0.7 / 0–500 / 0.25–10
—	override on runway T_s	12 / 0.15 / 0.5 yr	3–30 / 0–1 / 0.1–5
c^*, a_S, β	production: cover, adjustment, line weight	6 wk, 0.5/mo, 1	1–16, 0.05–1, 0–1
θ, ff	forecast smoothing (A), feedforward	0.3/mo, on	0.05–1, off
—	count the R&D pipeline (A)	on	off
Rules and the economic MPC			
k	Higgins multiple of g^*	1	0.5–2
T, H	exit year, cash-flow horizon	6, 5 yr	3–10, 2–15
C_{min}	cash buffer, months of opex	6	0–12
w, E	departure weight, round size (A)	1, \$3M	0.1–10, 0.5–10
—	MPC may raise one round	on	off
b	earnings retained	1	0–1

K_c in share per point of lead or growth, \$M/yr per unit of utilization, \$ per unit of margin and levels per month. Opening stations: u_R 0.15, u_M 0.10, u_p 20 k/yr, u_K \$0.5M/yr, P \$300, u_D 0.

8.5 The Board, Its Loops and Its Planner

The board has six stations, R&D and marketing as shares u_R and u_M of trailing revenue, production u_p, capital spending u_K, price P and the credit draw u_D, together with retention b, which is always the learner's. Each station has its own loop with a MANUAL and AUTO switch, a setpoint and two tuning constants, and a master switch throws all six. The controller executes every week. Table 8.3 lists the settings.

8.5.1 Calendar, Quantum and Cap

Every decision, by the learner's hand, a loop or the planner, passes the calendar and shaping functions of `core/cadence.js` in the collection's framework. The same shaping governs the central bank's rate in Volume 2, Chapter 11 and the intervention desk of Chapter 14, and Section 1.4.5 sets it beside the rate limits of other chapters. Weekly makes every execution a decision, and monthly and quarterly are regular intervals of $1/12$ and $1/4$ year, quarterly as shipped. A quarter is 13 weekly executions, while a month is not a whole number of them, so a monthly decision is taken at the first execution at or after each twelfth of a year, 5, 4 and 4 weeks apart in turn. The decision at $t = 0$ is taken when the case is loaded. A setting changed before Run sets the run up again, so it acts from that decision as it does after Reset, while a setting changed between meetings reaches the stations at the next decision. At a decision each station's request u^w is shaped as shown in Equation (8.10).

$$\begin{aligned} \tilde{u} &= u_{lo} + q_u \,\mathrm{round}\big((u^w - u_{lo})/q_u\big) \\ \Delta_q &= q_u \max\big(1, \lfloor \Delta_u / q_u \rfloor\big) \\ u_k &= \mathrm{clip}_{[u_{lo},u_{hi}]}\big(u_{k-1} + \mathrm{sat}_{\Delta_q}(\tilde{u} - u_{k-1})\big) \end{aligned} \tag{8.10}$$

where ties round away from zero and the bounds snap to the grid, and the station holds until the next decision. At $q = 1$ the quantum q_u is one point of revenue for the two spends, \$3 for the price, 0.5 thousand units a year for production and \$0.05M a year for capital spending and the draw, whose grid is based at zero. The cap Δ_u is a share of the station's range, none as shipped, and production is bounded by the capacity at the decision. The spend shares are paid as dollars at the trailing revenue of

the last close, $e_R = u_R \mathrm{Rev}_{TTM}$ and $e_M = u_M \mathrm{Rev}_{TTM}$, fixed until the next decision. The application's check finds all 72 decisions of three years of case 4 on the quarterly calendar and on their grids.

8.5.2 Four Loops of the Framework's PID

The R&D, marketing, capacity and price loops, and the override's level, run the velocity-form PID of Section 1.6.2 every week with the derivative time zero, as shown in Equation (8.11).

$$\begin{aligned} e_k &= a(r_k - y_k), \qquad P_k = K_c e_k \\ S_k &= S_{k-1} + P_k - P_{k-1} \\ \tilde{u}_k &= u_{b,k} + S_k + I_{k-1}, \qquad u_k = \mathrm{clip}_{[u_{min}, u_{max}]}(\tilde{u}_k) \\ I_k &= I_{k-1} + \left[K_c e_k / \tau_I + (u_k - \tilde{u}_k)/\tau_I\right]\Delta t \end{aligned} \tag{8.11}$$

where $a = \pm 1$ is the action, $u_{b,k}$ a feedforward bias, zero for R&D and marketing, and $\tilde{u}_k$ the output before limiting, v_k in Equation (1.14), with the tracking time $T_t = \tau_I$ of Equation (1.15). The clip is the loop's own bound, the Higgins cap included, while the override, the quantum, the move cap and the calendar act on the request after it, so the back-calculation does not see them, and while the override holds marketing down the growth loop's integral keeps running on its error. Section 1.4.4 contrasts this with an override fed back by external reset. Transfer between MANUAL and AUTO is bumpless as in Section 1.6.2, and the first execution in AUTO starts from the decision in force. Between decisions the request moves every week, and the calendar samples it.

The R&D loop holds the quality lead at zero with its request within 0 and 0.3 of revenue. With its pipeline counted, as shipped, it measures the lead the R&D already spent will deliver, $\ell = Q_p + \eta'_R(\tau_R/4)\sum_j r_j - Q_c - (\tau_R + \tau_Q) g_c m_Q$, with η'_R the marginal yield at the pipeline's mean rate, which is supply-line accounting applied to quality. The marketing loop holds year-over-year growth at the last close at 25 %. The capacity loop acts in reverse on the utilization the production plan asks for, a yearly average of the production request over capacity, with the bias $u_b = \min\left(\bar{u}_K,\ c_K K(\delta_K + g_{ff}\phi)\right)$ that replaces depreciation and grows the factory at the last close's growth g_{ff}, scaled by $\phi = \mathrm{clip}_{[0,1.5]}(2y/r - 1)$ so that an idle factory does not keep growing on last year's growth. The price loop is ratio control, with $u_b = c/(1 - m^*)$ held within \$150 to \$450 and a trim on the gross margin of the units sold, smoothed over about a quarter. It is in MANUAL at \$300 as shipped, so price moves only when the learner moves it.

8.5.3 The Production Rule

Production runs the ordering rule of the milk distribution game of Volume 6, Chapter 11, anchoring and adjustment with supply-line accounting, which the application's source names the Automatic Pipeline, Inventory and Order Based Production Control System (APIOBPCS), as shown in Equation (8.12).

$$\begin{aligned} O = \max\big\{0,\ F/12 + g(a_S, \delta)(S^* - I) \\ + g(\beta a_S, \delta)(\lambda F - W) + b_P\big\} \\ g(a, \delta) = \left[1 - (1-a)^{\delta}\right]/\delta, \qquad u_p^w = 12\,O \end{aligned} \tag{8.12}$$

where O is a monthly rate, δ the months between decisions (three quarterly), $W = W_1 + W_2$, $S^* = (c^*/52)F$ the stock wanted for a cover of c^* weeks, and b_P a transfer bias that starts at the decision in force and decays at the rule's own rate. The forecast F smooths the sales the company sees, $L \leftarrow L + \gamma(S - L)$ with $\gamma = 1 - (1-\theta)^{12/52}$ each week, so lost orders are invisible to it. With the seasonal feedforward on, as shipped, the level is smoothed with the season divided out, carried forward at the last close's growth over the smoothing's lag and the lead time, and multiplied by the season over the quarter the output will sell, and S^* adds the pre-build that the factory cannot make inside the coming peak. The rule is proportional on the stock and on its own pipeline, so the quantum and the capacity bound are its only anti-saturation.

8.5.4 The Runway Override

The cash loop is a PI controller on the stressed runway against its target $T^* = 12$ months, with the measurement held below 36, whose output is an override level $z \in [0,3]$ split in range over the other loops' requests, as shown in Equation (8.13).

$$\begin{aligned} u_M &\leftarrow u_M^w\,(1-\mathrm{sat}_{[0,1]}z) \\ u_K &\leftarrow u_K^w\,(1-\mathrm{sat}_{[0,1]}z) \\ u_D^w &= (\min(z,2)-1)\,[\bar{D}-D]_+/0.25, \qquad z>1 \\ u_R &\leftarrow u_R^w\,\big(1-0.7\,\mathrm{sat}_{[0,1]}(z-2)\big) \end{aligned} \tag{8.13}$$

From 0 to 1 the level scales marketing and capital spending to nothing, from 1 to 2 it draws the room the ceiling leaves within a quarter, and from 2 to 3 it cuts R&D to 30 % of its request. The three bands are a split range in the sense of Section 1.4.6, ordered by preference rather than by gain. With z at zero and the plain runway above $1.5\,T^*$ it repays any debt over six months. The order follows the pecking order of corporate finance, internal funds before debt and debt before equity [15], and the structure is that of a pressure override in a process plant [8], which acts on the request whoever made it. In MANUAL the station is the learner's draw and no override acts.

8.5.5 Rules That Set the Growth Target

Two practitioner rules replace the growth setpoint with one computed from the plant's own state at the last close. Higgins's rule sets it to $k\,g^*$ and also caps capital spending at $c_K K\big(\delta_K + \ln(1+kg^*)\big)(1+\ln(1+kg^*)\tau_b/3)^3$, the rate that grows the factory at $k\,g^*$ through its three-stage build, so that the company finances its own growth. The Rule of 40 sets it to 0.4 less the trailing free-cash-flow margin. On case 4 with the capacity loop in AUTO, the Rule of 40 ends year 10 at 12.3 % growth with a score of 56.3 %, identically when set before Run or carried through Reset.

8.5.6 The Economic MPC and the Round

The economic MPC replaces every AUTO loop's request at a board decision and plans at most once a quarter. At each plan it simulates its model of the company, the right-hand side of Equations (8.1) to (8.7) with every parameter at its declared default and no shock, forward a week at a time with explicit Euler steps over N quarters, holding the loops' requests $\bar{u}$ and producing to keep six weeks of demand, to give a nominal path $\bar{x}_k$. It linearizes that right-hand side by central differences about its estimate of the present state, with the sales minimum smoothed by a sixth-power norm, and discretizes it exactly for a zero-order hold over a quarter with the framework's `core/linmpc.js`. The prediction is $x_{k+1} = \bar{x}_{k+1} + A_d(x_k - \bar{x}_k) + B_d(u_k - \bar{u})$, and the quadratic program (QP) of Equation (8.14) chooses the decisions of the stations in AUTO in move blocks of 1, 1, 2, 2, 4, 6 and 10 quarters, the last taking what is left. The planner reads the company's books, the credit and the rate the bank reports, its own unit cost, its sales, the quality lead and the competitor's launch, and the seasonal calendar. The states the books do not show, the installed base, the customers turned away, the quality potential and the R&D and construction pipelines, come from the declared model run beside the company on its own decisions, with the quality potential corrected through the declared lag by the measured quality. One demand factor, updated whenever the shelf did not limit sales and held over the plan, scales the model's demand to the sales seen. On the declared company, unshocked, the estimate matches the simulated state bit for bit, so in case 8 the planner starts from the company's true state. The QP keeps the move suppression and priced slacks of Equation (1.17) and replaces its tracking term by

the founders' discounted value at the exit.

$$\begin{aligned}\min_{u,\varepsilon}\ & -\omega\phi e^{-r(T-t)}\big[V_T'(R_N-R_{N-4})\\ &\quad +V_P'(R_{N-4}-R_{N-8})+C_N-D_N\big]\\ &\quad +\textstyle\sum_j\|\Delta u_j\|_S^2+\sum_k\|u_k-\bar{u}\|_W^2\\ &\quad +\textstyle\sum_i(\rho_i\varepsilon_i^2+\varsigma_i\varepsilon_i)\\ \text{s.t. } & C_k\ge C_{buf}-\varepsilon_1,\qquad D_k\le\bar{D}+\varepsilon_2\\ & 2\bar{d}-\varepsilon_3\le I_k\le 12\bar{d}+\varepsilon_3,\qquad \varepsilon\ge 0\\ & u_{min}\le u_k\le u_{max},\qquad |\Delta u_j|\le\Delta u_{max}\\ & u_{p,k}\le K_k\end{aligned}\tag{8.14}$$

where R_k is the predicted cumulative revenue, so that R_N-R_{N-4} is the trailing revenue at the exit, $N=$ round$(4(T-t))$ held within 4 to 60, ϕ the dilution of a round, C_{buf} six months of operating expense, $\bar{d}$ a week of trailing demand and $\bar{D}$ the ceiling now. The terminal gradients are $V_T'=\mu_0+10(1+g_0)$ and $V_P'=-10(1+g_0)^2$, the derivatives of Equation (8.8) at the last close's growth g_0 and multiple μ_0 when $0<g_0<1$, and μ_0 and zero otherwise. Decisions are in model units, R&D and marketing in \$M/yr with bounds of 0 and 0.3 times revenue, and a station in MANUAL is held at its value by setting both of its bounds to it. Production $u_{p,k}$ is held within K_k, the capacity predicted for the start of quarter k, by one hard row per quarter, since capacity reaches the linearized flows only through the clip on production. The move weights S, departure weights W (scaled by w) and quarterly move caps Δu_{max} are, for (e_R,e_M,u_p,u_K,P,u_D), $(0.2/\text{Rev},\ 0.2/\text{Rev},\ 2\times10^{-4},\ 0.02,\ 2\times10^{-5},\ 0.02)$, $w(1/\text{Rev},\ 1/\text{Rev},\ 10^{-3},\ 0.1,\ 10^{-4},\ 0.1)$ and $(0.05\text{Rev},\ 0.05\text{Rev},\ 10,\ 1,\ 15,\ 4)$, with Rev the trailing revenue at the close and the draw bounded between repaying all debt and filling the ceiling within a quarter. The slacks carry quadratic prices of 10^3, 10^3 and 10^2 and linear prices of 10^2, 10^2 and 10, so the bands bind as hard constraints whenever they can be met. The cash-flow objective replaces the bracket by the discounted free cash flow over a horizon H plus the discounted value at its end. With the exit objective and no round raised, the QP is solved for seven candidates, no round, or E or $2E$ now, in a year or in two, each round adding its cash at its quarter with $\phi=V_e/(V_e+E)$ at the equity value grown at g_0, and the best objective is kept. Only a round chosen for now is raised, and a later one is re-decided each quarter. The first solve of case 8 has 45 variables (42 inputs and 3 slacks) and 286 rows. The Goldfarb–Idnani dual active-set solver [17] handles the 59 QPs of the case in a mean of 91.3 iterations and at most 169, with no failure. Over its first five years the plan meets its cash floor and its ceiling in its own prediction at all 19 plans. The plan's first block replaces each AUTO loop's request, its last plan stands in for the requests between plans, and the override still acts on it.

8.5.7 Company, Rules and Planner

Equations (8.1) to (8.8), with the disturbances of Table 8.2, are the simulated company. The loops hold almost no model, only the production lead time, the R&D pipeline's delay and the build chain's lag inside their feedforward terms, each at its declared default. The planner's model is the same right-hand side with every parameter at its declared default, linearized at its estimate of the present state, with sales smoothed, production bounded by the capacity it predicts, the covenant ceiling frozen at its present value, a round priced at the present equity value grown at g_0, the valuation linearized about the present growth, and no foresight of scheduled shocks. The models they stand for, Bass diffusion, Higgins's identity, the learning curve and the anchoring-and-adjustment rule, were fitted or derived in the works cited in Section 8.2. The application fits nothing to data of its own.

8.6 Measured Cases

All measurements in this section were regenerated on 2026-11-15 using the shipment-window accounting law of Section 8.3. The sensitivity counterfactuals use the same revised company.

Table 8.4 The Cases, the Challenge and the Reference

#	Change from the defaults	Teaching point
1	all MANUAL: u_R, u_M 0.15, u_p 40, u_K 1.5	insolvent while profitable
2	constant-ratio company, Higgins's rule, $k = 1$	the sustainable rate
3	$K_0 = 45$, only production in AUTO, feedforward off	the season as feedforward
4	capacity in MANUAL at \$0.4M/yr	the growth trap
5	R&D in MANUAL, cut from 0.15 to 0.05 at year 1	a pipeline hides a cost
6	customers pay in 75 days from year 1	profit without cash
7	rival launch $\times 1.15$ at year 2	two levers, fast and slow
8	economic MPC to an exit at year 6	dilution against speed
9	$L = 3$, $C_0 = 2$, \$2.5M loan at 1.5, recession at 2, runway MANUAL	a covenant that moves
ch	12 yr, recessions at 3 and 8, launch at 4, all MANUAL	everything at once
ref	125k repeat buyers, no season, nothing moves	the identities

Table 8.4 lists the cases. Each starts from the defaults of Tables 8.2 and 8.3, selected from the dropdown, with the change its row names. Each independent variant starts from a fresh scenario with the stated settings. A setting chosen before Run is applied to the opening board decision, with the same trajectory when it is carried through Reset. A loop switched to MANUAL holds the decision then in force, while an explicitly edited station keeps the learner's value. The comparisons distinguish those two manual plans. A run that becomes insolvent pauses there, and its later values are read after Run is pressed again. Table 8.5 collects the outcomes over each run's window, all measured on the engine headless. The scorecard also shows integrated errors, which grow with the length of the run and are therefore left out here. The exported file of Section 1.6.4 holds one row a week, with columns for each loop's setpoint, measurement, station, request and saturation, the override level, the statements, the valuation, the true demand beside the sales and the forecast, and the MPC's solves and failures. Playback runs the company at up to two years a second.

8.6.1 Insolvent While Profitable

Case 1 hands the board every station. The company is insolvent at 1.58 yr and its cash bottoms at $-\$0.38M$ at 1.79 yr, while its trailing revenue climbs to \$15.38M by year 4, growing 25.7 % a year, and its cash recovers to \$4.44M only because the model lets the company run on below zero. Reloaded and set before Run to 22 thousand units a year and \$0.5M of investment, the same company never falls below \$2.20M, with the same result through Reset, but sells \$6.60M in year 4, less than the \$7.26M it opened with, and fills 69.7 % of orders. At the first insolvency sample of the shipped plan, $t = 1.577$ yr, trailing revenue is \$8.37M, EBITDA \$1.11M and net income \$0.51M. The danger is visible only in the cash, because every outflow in Figure 8.1 buys an asset that the balance sheet counts and that a credit line sized on trailing EBITDA does not lend against.

> **TRY IT IN THE APP** **1 · The Board Game by Hand**
>
> Load the scenario and press Run on the Company view. Watch the cash tank empty until the status strip reads INSOLVENT, at about 1.58 yr, while trailing revenue is still rising. Then open the Statements view and find where the cash went: stock, factory and receivables.

Pause. Try Exercise 8.1 before reading on.

8.6.2 The Season and the Trap

Case 3 runs only the production rule on a factory of 45 thousand units a year that never limits it. With the feedforward off the shelf is short for 29 weeks of three years, 91.6 % of orders are filled and 6.98 thousand units are lost, and the first quarter after each peak is overstocked, 17.1 weeks of cover at 0.27 yr against the 6 wanted. With the feedforward on every order is filled and the cover

Table 8.5 Measured Cases Over Each Run's Window: Lowest Cash, First Insolvency, Trailing Revenue, Orders Filled, Valuation and Founders' Value at the End

Case and configuration	yr	C_{min} ($M)	insolv. (yr)	Rev_{TTM} ($M)	filled (%)	V ($M)	founders ($M)
1, as shipped	4	−0.38	1.58	15.38	98.8	79.5	67.1
1, u_p 22, u_K 0.5	4	2.20	—	6.60	69.7	14.8	14.4
2, $k = 1$	10	1.54	—	40.31	53.2	196.3	170.7
2, $k = 1.5$	10	−17.12	6.37	77.19	79.3	312.2	248.8
3, feedforward off	3	2.37	—	8.97	91.6	40.9	36.0
3, feedforward on	3	2.17	—	10.45	100.0	51.9	45.1
4, u_K 0.4 by hand	8	2.53	—	8.97	59.4	24.1	26.7
4, capacity AUTO	8	0.90	—	36.11	99.9	197.7	175.6
5, R&D cut to 0.05	5	0.96	—	9.84	99.7	25.6	26.3
5, R&D held at 0.15	5	1.07	—	15.36	99.8	76.8	64.0
5, R&D loop AUTO	5	0.99	—	16.01	99.8	86.7	72.0
6, override	4	0.79	—	11.57	99.7	56.6	47.1
6, override MANUAL	4	−2.03	1.88	13.14	99.7	69.3	55.1
6, on time, override	4	0.99	—	12.49	99.7	77.5	64.0
6, on time, MANUAL	4	−1.25	2.37	13.14	99.7	69.3	56.2
7, as shipped	8	0.82	—	11.85	98.5	51.8	43.7
7, $255 at the launch	8	0.73	—	14.40	91.5	62.6	51.6
7, price loop at 40 %	8	0.40	—	7.54	74.0	19.5	16.1
7, no launch	8	0.99	—	36.32	99.9	194.4	174.1
8, MPC, one round	6	1.88	—	29.41	100.0	177.3	140.0
8, MPC, no round	6	1.47	—	22.92	100.0	128.8	108.6
8, six loops	6	0.99	—	21.69	99.8	133.4	112.0
8, loops, $3M at yr 2	6	0.99	—	22.61	99.9	132.6	104.2
8, explicit manual plan	6	2.34	—	6.00	65.2	14.1	14.7
8, hold decisions in MANUAL	6	−1.35	4.02	3.45	54.4	6.9	4.4
9, runway MANUAL	5	−0.01	2.79	13.81	97.9	94.5	76.1
9, override, T^* 24	5	1.50	—	8.80	90.7	49.9	42.3
9, override, T^* 12	5	0.62	—	12.39	97.1	90.3	74.0
ref	8	3.00	—	3.75	100.0	7.8	8.6

Founders' value is $\omega[V + C - D]_+$, and so exceeds V where cash is large.

On time, held and no launch are case 8 under the six loops, with the runway or R&D loop in MANUAL where named, which has no scripted event. The $3M round is the round button pressed at year 2.

stays between 0.84 and 8.6 weeks. Feedback alone corrects a seasonal error one season late, which is out of phase.

Case 4 holds capital spending at $0.4M a year by hand. Revenue creeps from $7.62M at year 2 to $8.97M at year 8 while demand runs 1.32 times sales at year 2 and 2.65 times at year 8, and 59.4 % of orders are filled. By year 8, 112.7 thousand would-be customers have bought elsewhere and left the potential, against 191.2 thousand served. The marketing loop, measuring stalled sales, raises marketing from 7 % to 20 % of revenue and buys demand the factory cannot serve. With the capacity loop in AUTO the company reaches $36.11M at year 8 and 99.9 % of orders filled, and its valuation is $197.7M instead of $24.1M. A loop on sales cannot see the demand it turns away, and a utilization target of 100 % never gives the capacity loop a reason to build.

🕮 **Pause.** Try Exercise 8.2 before reading on.

8.6.3 Pipelines and Receivables

Case 5 cuts R&D from 15 % to 5 % of revenue at year 1. At 1.5 yr the cash is $2.20M against $1.95M with R&D kept, and the share is 0.449 in both runs. The market notices from year 2, 0.422 against 0.442 at 2.5 yr and 0.282 against 0.475 at year 5, when the quality lead is −0.19 against −0.02. The cut company holds more cash at year 5, $7.33M against $3.21M, but sells $9.84M a year instead of $15.36M and is valued at $25.6M instead of $76.8M. The company with R&D kept is case 8 under the six loops with the R&D loop in MANUAL at 15 %, which differs from case 5 only by the cut. The R&D loop in AUTO, which counts its own pipeline, reaches $16.01M and $86.7M.

Case 6 stretches collections from 45 to 75 days at year 1. At year 2 receivables stand at \$1.96M against \$1.30M on time, in case 8 under the six loops, the same company without the late payments, while trailing EBITDA is \$1.348M against \$1.406M. The override cuts marketing to nothing, its level passes 1 at the execution of 1.69 yr, the board draws the line from its meeting at 1.75 yr, at most \$0.31M, and the level peaks at 1.31, short of the band that cuts R&D, so R&D stays at 15 % of revenue and the lowest cash is \$0.79M against \$0.99M on time. With the override in MANUAL the company is insolvent at 1.88 yr, with a trailing EBITDA at year 2 of \$1.560M, exactly what it earns with customers on time, and cash of $-\$0.29M$ against \$0.52M. Without the override and with customers on time it is insolvent at 2.37 yr, because every loop in it spends before it earns and only the override sees them all.

Pause. Try Exercise 8.3 before reading on.

> **TRY IT IN THE APP** **6 · Customers Pay Late**
>
> Run the scenario for three years and watch the Cash card: the override level passes 1, the credit line is drawn and R&D stays at 15 % of revenue. Set the Runway loop and override to Manual, press Reset and run again. The company is insolvent at 1.88 yr.

8.6.4 Two Levers on Share

Case 7 has the rival launch a product 15 % better at year 2. As shipped, share falls from 0.446 to 0.246 by 2.85 yr, the R&D loop goes to its cap of 30 % of revenue and share recovers to 0.37 by year 8, with a valuation of \$51.8M against \$194.4M without the launch, case 8 under the six loops. A price cut to \$255 decided at the meeting of year 2, the launch's quarter, which offsets about 0.09 of quality deficit at $\beta_P/\beta_Q = 0.6$, holds share at 0.338 or better and reaches 0.569 by year 8 and a valuation of \$62.6M, with the lowest cash \$0.73M. Set once the launch is announced, the cut waits for the meeting of 2.25 yr, share first falls to 0.249 and the valuation reaches \$59.9M. Letting the price loop chase a 40 % margin drives the price to its \$210 floor at the meeting of 2.50 yr, and the lower price raises demand while reducing margin. With the setting chosen before Run, the company stays solvent over eight years, with minimum cash of \$0.40M, trailing revenue of \$7.54M, 74.0 % of orders filled and a valuation of \$19.5M. Solvency alone does not make that price policy a better outcome than the unchanged \$300 price. Two such launches at once, the launch button pressed at year 2 beside the scripted one, a factor of 1.32, take share below 0.04, and all five configurations measured, among them a price of \$240 or \$210 and the R&D loop at its largest gain of 2 with $\tau_I = 2$ yr, are insolvent between 2.79 yr and 4.27 yr. The design of the application first proposed that launch, and the case uses 15 % instead.

8.6.5 An Acquisition Planned by the MPC

Case 8 hands every station in AUTO to the economic MPC with an exit at year 6, while price stays in MANUAL at \$300. The plan at 1.5 yr chooses \$3M now, raised at 1.52 yr at a \$35.24M pre-money, reducing the founders from 80 % to 73.7 %. At year 6 trailing revenue is \$29.41M, growing 34.3 %, valuation is \$177.29M and cash is \$12.56M, giving founders' value of \$139.96M (Table 8.5), and Figure 8.3 shows the cash, the growth, the valuation and the founders' value as the app rounds them. Without the round the planner reaches \$108.63M, below the six loops' \$111.95M; the loops with a \$3M round forced at year 2 reach \$104.23M. The explicit manual plan reaches \$14.68M: R&D at 15 %, marketing at 10 %, production at 20 thousand units a year, investment at \$0.5M/yr, price at \$300 and no credit draw. Switching all loops to MANUAL without editing the stations instead holds the opening production decision of 11.5 thousand units a year and is insolvent at 4.02 yr, both before Run and through Reset. The MPC's round adds \$31.33M for the founders on this model. It buys both scale and faster growth: the bootstrapped planner grows 30.2 % on \$22.92M of revenue.

The funded planner builds 126.8 thousand units a year of capacity by year 6, and the bootstrapped planner 97.3 thousand, against the loops' 110.6 thousand. The bootstrapped local plan trails the

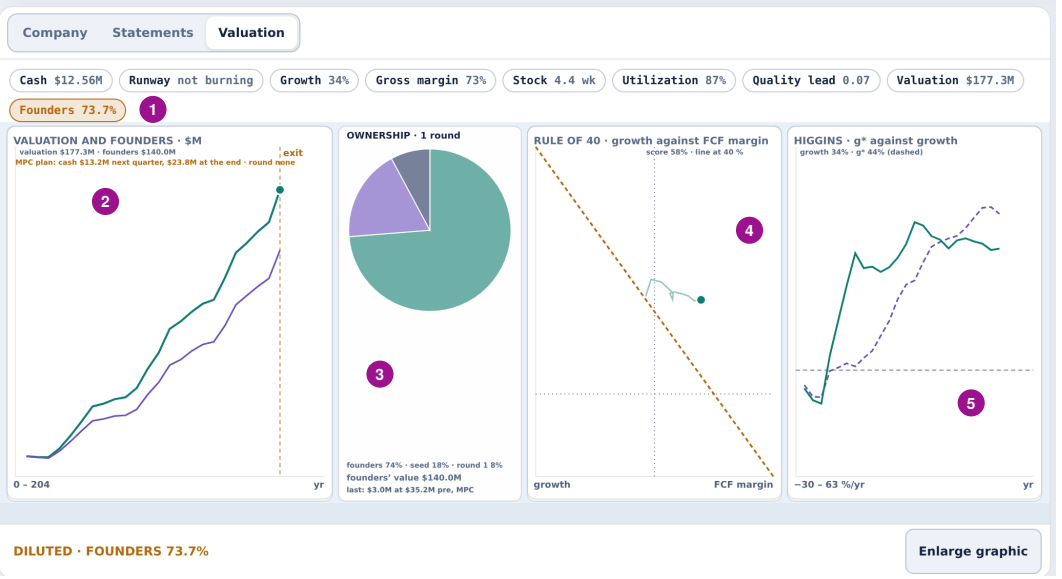

Figure 8.3 Case 8 at $t = 6.00$ yr in the Valuation view, left in AUTO as shipped, cropped to the view: the header, the control panel, the trends and the learning panel are cut. (1) The badges show cash \$12.56M, growth 34 %, gross margin 73 %, stock 4.4 weeks, utilization 87 %, quality lead 0.07, valuation \$177.3M and founders 73.7 %. (2) The valuation panel rounds founders' value to \$140.0M; its full value is \$139.96M. The last plan, made at 5.75 yr, expects cash of \$13.2M next quarter and \$23.8M at its end. (3) The ownership panel shows the round of \$3.0M at a \$35.2M pre-money. (4) Rule of 40 reads 58 % and (5) Higgins's g^* 44 %.

Table 8.6 Case 8's Founders' Value at Year 6 Under Changes to the Planner

Planner	one round ($M)	no round ($M)	round chosen
Current application			
$w = 1$, as shipped	140.0	108.6	\$3M at 1.52 yr
$w = 0.3$	168.2	134.2	\$3M at 1.52 yr
$w = 3$	148.8	124.5	\$3M at 1.52 yr
$w = 10$	70.8	94.9	\$3M at 1.77 yr
price loop in AUTO	96.6	93.9	\$3M at 3.52 yr
Counterfactuals, same revised company			
omit both fixes	153.8	165.0	\$3M at 1.52 yr
fixed production ceiling	63.6	62.3	\$6M at 1.52 yr

The six loops reach \$112.0M, or \$44.4M with the price loop in AUTO.
Each independent setting is chosen before Run; carrying it through Reset gives the same result.

loops in founders' value; predictive control does not guarantee dominance on the realized nonlinear company.

The counterfactual rows use the same accounting law and declared-model estimator. Omitting both the terminal-gradient factor $(1 + g_0)$ and the MANUAL-price hold yields \$153.8M with a round and \$165.0M without. Keeping the current formulation but bounding every future production decision by today's factory instead of its predicted capacity yields \$63.6M with a round and \$62.3M without; the funded factory reaches only 69.6 thousand units a year. The latter constraint makes an investment look like cash spent without the production it would enable.

The verdict on the round also moves with the departure weight w. It adds \$31.33M at $w = 1$,

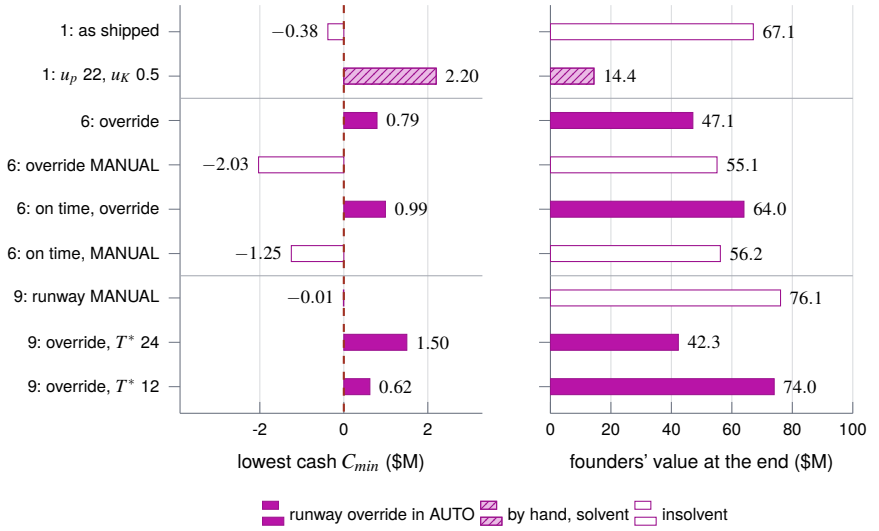

Figure 8.4 Lowest cash and founders' value of cases 1, 6 and 9 from Table 8.5, grouped by case (thin rules). The dashed line is the cash floor. Dark bars ran with the runway override in AUTO, the hatched bar by hand without insolvency, and white bars became insolvent, whose later values are counterfactual because the model runs on below zero. Notice that every run without the override in cases 6 and 9 falls below the floor, and that in case 9 the 24-month target keeps the most cash and the least founders' value.

$34.03M at 0.3 and $24.27M at 3, but loses $24.06M at 10, where both plans end below the six loops. With price in AUTO the planner reaches $96.61M with a round and $93.94M without, while the six loops reach $44.35M. An economic optimum on a linearized model belongs to that model and to the weights around it.

8.6.6 A Covenant That Moves

Case 9 opens with $2M of cash, borrows $2.5M on a $3M line at 1.5 yr and meets a recession at year 2 that takes 30 % of demand and 20 % of the potential for eighteen months. As trailing EBITDA falls the ceiling drops to $1.16M, the bank calls $0.87M from 2.75 yr, and with the runway loop in MANUAL the company is insolvent at 2.79 yr. With the override in AUTO and a 24-month target it cuts marketing and investment from its first year and repays the debt down to $0.83M by year 2, so only $0.17M is called and the Lowest cash tile rounds to $1.50M, while a 12-month target also holds with a lowest cash of $0.62M. The buffer costs growth, a founders' value at year 5 of $42.3M at 24 months and $74.0M at 12. With the credit line priced at the inflation-targeting application's replayed policy rate plus 5 points, 10.125 % at year 5, the unprotected company is insolvent at 2.73 yr. A covenant is a constraint whose bound moves with the plant's own output. Figure 8.4 sets the lowest cash and the founders' value of cases 1, 6 and 9 side by side.

Table 8.7 The Twelve-Year Challenge in Seven Configurations

Configuration	C_{min} ($M)	growth (%/yr)	filled (%)	V_{12} ($M)	Target
shipped, all MANUAL	−28.64	−23.0	100.0	0.6	not met
suggested, quarterly	0.85	20.7	97.7	272.5	met
suggested, monthly	1.00	20.6	98.8	284.5	met
suggested, weekly	1.08	20.6	99.1	286.6	met
suggested, cap 10 %	0.15	20.2	89.3	339.0	not met
Higgins's rule	2.13	17.7	79.5	306.8	not met
MPC, cash flow, $H = 5$	0.84	−5.7	99.5	6.5	not met

Insolvent at 4.40 yr as shipped, read at year 12 after Run past the pause.
Changes carried through Reset. Growth compounded from Rev_{TTM} before $t = 0$.

> ✦ **DID YOU KNOW? A Longer Runway Target Costs the Founders**
>
> From the simulator: in case 9 the company with its runway loop in MANUAL is insolvent at 2.79 yr. The override with a 24-month target keeps its lowest cash at $1.50M but ends year 5 with a founders' value of $42.3M, against $0.62M and $74.0M with a 12-month target (Table 8.5).

8.6.7 Twelve Years Without Outside Money

The challenge runs twelve years on the quarterly calendar through recessions at years 3 and 8, a launch at year 4 and the season, and asks, judged at year 12, for no insolvency, no equity round, trailing revenue compounding at least 15 % a year from the $6.57M of the year before $t = 0$, at least 95 % of orders filled and a valuation of at least $150M. Table 8.7 compares seven configurations. The shipped board holds R&D and marketing at 10 %, production at 30 thousand units a year and investment at $1M a year, and is insolvent at 4.40 yr in the first recession, as production and investment run on while demand falls. The suggested settings put every loop except price in AUTO, start R&D at 15 % with a lead setpoint of 0.1 and a gain of 2, hold twelve weeks of stock and 70 % utilization with a capacity gain of 10, and pass with a lowest cash of $0.85M, 20.7 % growth, 97.7 % of orders filled and a valuation of $272.5M, identically before Run and through Reset. Its margins over the targets are 5.7 points of growth, 2.7 points of orders filled and $122.5M of valuation.

Deciding more often helps a little. Monthly, with 144 decisions, fills 98.8 % of orders with a lowest cash of $1.00M, and weekly, with 624, 99.1 % with $1.08M, while growth stays at 20.6 %, likely because production and capacity follow the season sooner. A move cap of 10 % of each station's range per decision fills only 89.3 % and fails, while a quantum of zero or three changes the fill by at most two tenths of a point. Higgins's rule on the same loops grows 17.7 % a year with $2.13M at the lowest and fills only 79.5 % of orders, which makes it sustainable by turning customers away. The economic MPC with the cash-flow objective over five years harvests the company and shrinks it 5.7 % a year. All these independent configurations give the same result before Run and through Reset. An earlier version of the challenge asked instead for stock never below two weeks, a floor that a policy can meet by rationing sales, and it rewarded such a policy.

The balanced reference stands still. Cash holds at exactly $3.00M for eight years and equity at $5.10M, with EBITDA of $0.15M replacing the factory as it wears out and net income zero, and customers paying late there move cash into receivables one for one.

8.7 Technology in Practice

A real company reads the same stocks and moves the same handles as the board of Section 8.5, but its instruments are accounts, statements and contracts. Figure 8.5 draws the runway override with those instruments in place of the schematic's symbols.

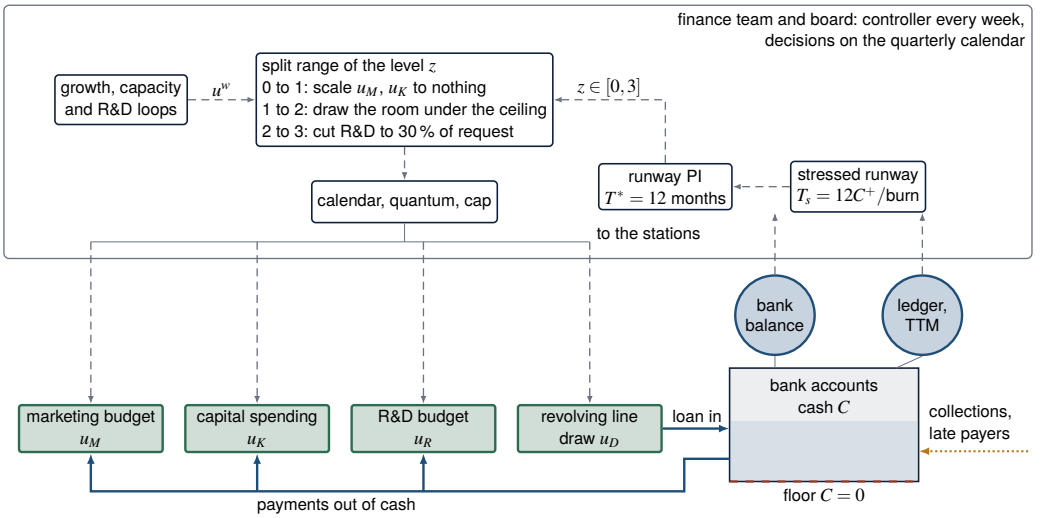

Figure 8.5 The runway override of Section 8.5 with a company's real instruments. Green boxes are the final elements, three budget lines and the draw on a revolving line of credit, blue circles the measurements, the bank balance and the ledger kept on trailing twelve months (TTM). Solid arrows carry money, dashed arrows signals, and the dotted arrow the collections that a late payer delays. Notice that one level z acts on four stations in a fixed order, and that its request waits for the board's calendar before any station moves.

What Measures It

The controlled variable of the override is a runway, cash divided by a burn, and both come from the books. Most firms use accrual accounting, and the statement of cash flows most often follows the indirect method, which starts from net income and adjusts it for noncash expenses and for changes in receivables, inventory and payables. Income statements of publicly traded companies are available quarterly [18]. Receivables are watched through an aging schedule, which sorts unpaid invoices by the days since the sale and shows which customers pay late [18]. The trailing twelve-month ring of Section 8.3 stands for that bookkeeping, kept weekly and exactly.

> **TECHNOLOGY** **Sensor: The Cash Balance and the Cash Budget**
>
> A company's cash holdings are mostly checks waiting to be deposited and the balances of its checking accounts. Profit is a different reading: under accrual accounting a sale is recorded when it occurs, not when cash moves, so a firm can be profitable and still lack an adequate flow of cash, and the statement of cash flows reconciles the two [18]. A ratio taken at one date says little about future liquidity [18], so the cash budget projects the balance forward, divided into the shortest period possible [19]. The model's sensor is its own books on trailing twelve months, from which the override reads a stressed runway. A late payer shows there first as receivables that rise while EBITDA holds.

What Moves It

The board's stations are budget lines and a draw on a bank line. The master budget has an operating part, which projects revenue and expenses, and a financial part, which projects assets, liabilities and cash, and its most common period is one year [19]. Credit terms are the disturbance side of the same ledger: net terms usually range between 30 and 90 days, depending on the industry [18], and case 6 moves the model's collections from 45 to 75 days inside that range.

TECHNOLOGY Final Element: Budgets and a Revolving Credit Line

Companies often arrange a line of credit as an emergency source of cash, much like overdraft protection on a checking account [19]. Under one federal loan guarantee program an asset-based revolving line lets a business draw continually against its existing assets and repay as its cash cycle dictates, and a revolving loan to a manufacturer is reviewed each year: an unsatisfactory review converts it to a term loan with no further draws [20]. The line may also be out of reach. Of small employer firms that applied for financing in the year before the 2025 survey, 22 % received none [21]. The model's line has a ceiling of three times trailing EBITDA, so it moves with the company's own earnings.

Where the Controller Runs

A company's bylaws determine how often its board meets [18]. The cash budget is commonly geared toward monthly or quarterly figures, and a rolling budget is adjusted each month, with a new month added as each one passes [19]. The application's controller executes every week, and its decisions pass the quarterly calendar of `core/cadence.js` as shipped, so the override's request reaches the stations at the next meeting. The structure lives in the cash budget's minimum balance, and the split range of Equation (8.13) writes down in advance what gives way first.

What Is Modernizing It

Three developments bear on this loop. Lending has moved online: among small employer firms that applied for loans, lines of credit or cash advances, the share that sought financing at online lenders rose from 17 % in the 2020 survey to 29 % in the 2025 survey, and 60 % of those that borrowed there reported costs higher than expected [21]. Artificial intelligence has reached planning: 46 % of firms in the same survey use it, 51 % of users for planning or analysis, and accuracy was the challenge users named most [21]. Faster measurement shortens the loop: electronic invoicing most likely cuts the average collection period by removing mail delays [18]. The economic MPC of Section 8.5 adds a planner whose stage cost is money [7], and case 8 shows how far its verdict depends on its own weights.

IN PRACTICE A Minimum Cash Balance in the Cash Budget

A cash budget adds the period's receipts to the beginning balance and subtracts its payments to give a potential ending balance. Organizations typically require a minimum cash balance: if the potential ending balance falls short, management plans financing to reach it, and an excess may repay loans and interest [19]. A bank may itself require a minimum balance in the company's checking account [18]. Forecasts are redone for scenarios such as a recession, and sensitivity analysis changes one input at a time, sales 10 % lower, for example [18]. The runway override turns that review into a standing rule.

8.8 Limitations

The company is one product against one rival that never responds, its delays are distributed lags, and its finance omits loss carry-forward, preferences and the clauses of a real loan. Every coefficient, multiple, threshold and rule, the challenge's 15 %, 95 % and $150M included, is a teaching value for an imaginary hardware startup, and nothing here is investment or business advice. Higgins's g^{*} is a bound only where its ratios hold, as in case 2, and in the default company the learning curve, the overhead and the rival move the ratios. The valuation multiple, the Rule of 40 and the runway are teaching rules with a practitioner's shape, not appraisals. Cash may fall below zero and the company runs on, so any recovery after INSOLVENT is a counterfactual. The economic MPC is local and successively linearized, holds the declared parameters, so that a changed parameter or a shock reaches it only through the books and the sales, and does not preview scheduled shocks, and its result in case 8 moves between $70.8M and $168.2M across the departure weights of Table 8.6.

OPEN CHALLENGE **Financing That May Not Be There**

Startup outcomes are hard to predict even for experienced investors: at one venture capital firm the scores that partners gave at the first investment correlated with the return multiple at only 0.1. Investors also speak of financing risk, the chance that an otherwise sound project cannot raise capital for its next stage. Taking larger sums less often protects against it but reduces the value of the option to abandon, and it bites hardest on experimental new technologies that must return to the market many times [9]. In case 8 a round the planner chooses is always raised. A planner that keeps its cash floor when the round may not come would be progress.

8.9 Conclusion

The simulator turns a startup into four coupled inventories, cash, stock, capacity and quality, with one floor, and measures what each decision costs in the tank. A board that grows by hand is insolvent at 1.58 yr with trailing EBITDA of $1.11M. On books kept on trailing twelve months Higgins's rate equals year-over-year growth exactly, a company held to it keeps its cash near 40 % of equity, and at 1.5 times the rate every split that earns money runs out of cash between 5.33 yr and 12.37 yr. The season's feedforward fills every order where feedback fills 91.6 %, a factory held at replacement fills 59.4 %, and the runway override survives late payments by drawing $0.31M without cutting R&D. The economic MPC reaches $139.96M at an acquisition with a round and $108.63M without; the latter trails the six loops' $111.95M, and the verdict on the round changes sign with the departure weight. In the challenge the suggested loops pass at 20.7 % growth and 97.7 % of orders filled where the board's fixed split fails at 4.40 yr. The model is one product, its finance is simplified and its values are for teaching. A fuller model would add a rival that answers, loss carry-forward and loan clauses, an estimator for the hidden demand, and a planner that previews scheduled shocks. The application makes the cost of growth a measurement in the one stock that ends the run.

KEY POINTS

- A company can book a profit and still run out of cash, because stock, factory and unpaid invoices are paid for before customers pay.
- At 1.5 times Higgins's rate every budget split that earns money runs out of cash, between 5.33 yr and 12.37 yr.
- The runway override scales marketing and capital spending to nothing, then draws the credit line, then cuts R&D.
- With customers paying late the override survives by drawing $0.31M without cutting R&D. Without it the company is insolvent at 1.88 yr.
- The verdict of the economic MPC on the equity round changes sign with the departure weight.

Where else this applies. Chapter 14 defends a reserve with a floor of its own, and Volume 6, Chapter 11 runs the same anchoring-and-adjustment ordering rule on a supply chain.

References

[1] C. C. Holt, F. Modigliani, and H. A. Simon, "A linear decision rule for production and employment scheduling," *Management Science*, vol. 2, no. 1, pp. 1–30, 1955.

[2] M. L. Vidale and H. B. Wolfe, "An operations-research study of sales response to advertising," *Operations Research*, vol. 5, no. 3, pp. 370–381, 1957.

[3] M. Nerlove and K. J. Arrow, "Optimal advertising policy under dynamic conditions," *Economica*, vol. 29, no. 114, 1962.

[4] J. D. W. Morecroft, "System dynamics: Portraying bounded rationality," *Omega*, vol. 11, no. 2, pp. 131–142, 1983.

[5] R. C. Higgins, "How much growth can a firm afford?" *Financial Management*, vol. 6, no. 3, 1977.

[6] N. C. Churchill and J. W. Mullins, "How fast can your company afford to grow?" *Harvard Business Review*, vol. 79, no. 5, 2001, PMID 11345910.

[7] D. Angeli, R. Amrit, and J. B. Rawlings, "On average performance and stability of economic model predictive control," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1615–1626, 2012.

[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] W. R. Kerr, R. Nanda, and M. Rhodes-Kropf, "Entrepreneurship as experimentation," *Journal of Economic Perspectives*, vol. 28, no. 3, pp. 25–48, 2014.

[10] N. P. Repenning and J. D. Sterman, "Capability traps and self-confirming attribution errors in the dynamics of process improvement," *Administrative Science Quarterly*, vol. 47, no. 2, pp. 265–295, 2002.

[11] T. P. Wright, "Factors affecting the cost of airplanes," *Journal of the Aeronautical Sciences*, vol. 3, no. 4, pp. 122–128, 1936.

[12] F. M. Bass, "A new product growth for model consumer durables," *Management Science*, vol. 15, no. 5, pp. 215–227, 1969.

[13] F. M. Bass, T. V. Krishnan, and D. C. Jain, "Why the Bass model fits without decision variables," *Marketing Science*, vol. 13, no. 3, pp. 203–223, 1994.

[14] V. D. Richards and E. J. Laughlin, "A cash conversion cycle approach to liquidity analysis," *Financial Management*, vol. 9, no. 1, 1980.

[15] S. C. Myers, "The capital structure puzzle," *The Journal of Finance*, vol. 39, no. 3, pp. 574–592, 1984.

[16] P. A. Gompers, W. Gornall, S. N. Kaplan, and I. A. Strebulaev, "How do venture capitalists make decisions?" *Journal of Financial Economics*, vol. 135, no. 1, pp. 169–190, 2020.

[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] J. Dahlquist and R. Knight, *Principles of Finance*. Houston, TX: OpenStax, 2022. [Online]. Available: https://openstax.org/books/principles-finance/pages/1-why-it-matters

[19] M. Franklin, P. Graybeal, and D. Cooper, *Principles of Accounting, Volume 2: Managerial Accounting*. Houston, TX: OpenStax, 2019. [Online]. Available: https://openstax.org/books/principles-managerial-accounting/pages/1-why-it-matters

[20] "SBA lenders: Types of 7(a) loans," U.S. Small Business Administration, 2026, accessed 2026-11-21. [Online]. Available: https://www.sba.gov/sba-lenders/#types-7a-loans

[21] Federal Reserve Banks, "2026 report on employer firms: Findings from the 2025 small business credit survey," Federal Reserve Banks, Small Business Credit Survey, Tech. Rep., 2026.

8.10 Exercises

8.1 Concept check · 2 min

In case 1 the run pauses at INSOLVENT at 1.58 yr. The trailing net income is then \$0.51M, nothing has been borrowed and every dollar of earnings is retained. Where did the opening \$3M of cash go?

(a) Into interest and tax, which are paid before the revenue they follow is collected

(b) Into dividends, since the board pays out the trailing net income every quarter

(c) Into a larger factory and a growing stock of finished units, bought before customers pay

(d) Into operating losses, which the trailing income statement reports as a profit

8.2 Concept check · 1 min

In case 4 the board holds capital spending at \$0.4M a year. Between year 2 and year 8 the growth loop raises marketing from 7 % to 20 % of revenue. Why does the loop keep spending more?

(a) The capacity loop, in AUTO, asks for more demand to fill the factory it builds

(b) The runway override releases marketing as cash piles up in the tank

(c) The R&D loop loses the quality race, and the growth loop makes up the lost share

(d) It measures sales, which the factory caps, so it sees stalled growth and buys demand it cannot serve

8.3 Concept check · 2 min

In case 6 the runway override's level peaks at 1.31. Which describes the stations at that moment?

(a) Marketing and capital spending at zero, the credit line being drawn, R&D at its loop's request
(b) Marketing at a third of its request, capital spending untouched, no draw, R&D at its request
(c) Marketing and capital spending at zero, no draw, R&D cut to a third of its request
(d) Every request divided by 1.31, so marketing, capital spending and R&D fall in step

8.4 Theory and calculation · 6 min

A company grows with constant ratios at the continuous rate γ, so its equity is $E(t) = E(0)e^{\gamma t}$ and it retains $bNI = \dot{E}$.

(a) Show that $bNI_{TTM}/E = 1 - e^{-\gamma}$ on books closed on trailing twelve months, and that Higgins's rate $g^* = x/(1-x)$ then equals the year-over-year growth $e^{\gamma} - 1$. *Checkpoint:* $x = 1 - e^{-\gamma}$, so $x/(1-x) = e^{\gamma} - 1$.

(b) Case 2 reports $g^* = 22.2\,\%$ at year 4. Find γ and the retained return x behind it. *Checkpoint:* $\gamma = 0.2005\,\text{yr}$ and $x = 0.1817$.

(c) At the opening close of case 2, $g^* = 26.8\,\%$ and $K = 25$ thousand units a year. With $c_K = \$0.1M$ per thousand units a year, $\delta_K = 0.1\,\text{yr}$, $\tau_b = 0.75\,\text{yr}$ and $k = 1$, find the cap that Higgins's rule puts on capital spending, and the cap without the factor $(1 + r\tau_b/3)^3$, where $r = \ln(1 + kg^*)$. *Checkpoint:* \$1.00M a year, and \$0.84M a year without the factor.

8.5 Theory and calculation · 5 min

Receivables follow $\dot{AR} = \text{Rev} - (365/\text{DSO})AR$ of Equation (8.6), and the credit ceiling is $\bar{D} = \min(Lm_L,\ \kappa[\text{EBITDA}_{TTM}]_+)$.

(a) At the default decisions the company opens with \$6.57M of trailing revenue, a trailing EBITDA of \$0.14M, DSO of 45 days, $L = \$1M$, $m_L = 1$ and $\kappa = 3$. Find the receivables in steady state and the credit ceiling. *Checkpoint:* \$0.81M and \$0.42M.

(b) The balanced reference sells \$3.75M a year. Customers stretch from 45 to 75 days. Find the cash that moves into receivables, and the time constant of the receivables before and after. *Checkpoint:* \$0.308M, with 0.123 yr (6.4 weeks) before and 0.205 yr (10.7 weeks) after.

(c) Explain why the ceiling does not rise to finance the receivables of part (b).

8.6 Theory and calculation · 6 min

The share is the logit of Equation (8.1) with $\beta_Q = 5$, $\beta_P = 3$ and $P_c = \$300$, and the competitor's quality is $Q_c = (1 + g_c t)\,m_Q$ with $g_c = 0.08\,\text{yr}$.

(a) Show that at a fixed price $\mathrm{d}s/\mathrm{d}(Q - Q_c) = \beta_Q s(1-s)$, and find the change in share for a lead of 0.1 point at $s = 0.449$. *Checkpoint:* 0.124.

(b) In case 7 the launch at year 2 sets $m_Q = 1.15$. Find the jump in Q_c and the share just after it from $s = 0.446$, with the quality and the price unchanged. *Checkpoint:* 0.174 and 0.252.

(c) Find the price that would offset the whole jump, and the quality deficit that the cut to \$255 of case 7 offsets. *Checkpoint:* \$213 and 0.09 point.

8.7 App challenge · 9 min

Load scenario 1, The board game by hand, at the default seed.

(a) Press Run and let the run go on until it pauses at INSOLVENT. Read the clock, and export the CSV file. In its last row read `revenue_ttm_musd`, `ebitda_ttm_musd` and `ni_ttm_musd`. *Checkpoint:* 1.58 yr, then \$8.37M, \$1.11M and \$0.51M, within 0.01.

(b) From the first and last rows of the same file, find the change in `cash_musd`, `fixed_musd`, `stock_book_musd`, `wip_book_musd`, `AR_musd`, `AP_musd` and `equity_musd`, and show that they account for the cash. *Checkpoint:* -3.02, $+1.82$, $+1.17$, $+0.10$, $+0.01$, $+0.09$ and -0.01 \$M.

(c) Press the reload button, set Production rate to 22 and Capacity investment to 0.5 in their cards, and run to year 4. Read the Lowest cash and Orders filled tiles and `revenue_ttm_musd` at year 4. *Checkpoint:* \$2.20M, 69.7 % and \$6.60M, within 0.01.

Is the plan of part (c) a better plan?

8.8 App challenge · 8 min

Load scenario 6, Customers pay late, at the default seed. The runway override is in AUTO with Override gain 0.15.

(a) Run three years. Read the Lowest cash tile, and in the exported CSV file find the largest `override_level` and `debt_musd` and the smallest `rd_share`. *Checkpoint:* \$0.79M, then 1.31, \$0.31M and 0.15, within 0.01.

(b) Set the Runway loop and override to Manual, press Reset and run until the run pauses at INSOLVENT. Read the clock. *Checkpoint:* 1.88 yr, within 0.02.

(c) Press the reload button, set Override gain to 0.50 and run three years. Read the same four values. *Checkpoint:* \$1.09M, then 3.00, \$0.98M and 0.05, within 0.01.

What does the higher gain buy, and what does it give up?

8.9 App challenge · 9 min

Load the challenge, Twelve years without outside money, at the default seed. For each run read the Lowest cash, Orders filled and Challenge tiles, and in the exported CSV file `revenue_ttm_musd` and `valuation_musd` at year 12. Compute the growth as $(\mathrm{Rev}_{TTM}(12)/6.57)^{1/12} - 1$.

(a) Run as loaded until the run pauses, and read the clock. *Checkpoint:* INSOLVENT at 4.40 yr.

(b) Press Reveal suggested tuning, press Reset and run twelve years. *Checkpoint:* \$0.85M, 97.7 %, met, \$62.54M and \$272.5M, growth 20.7 %.

(c) Set Decision policy to Higgins, press Reset and run twelve years. *Checkpoint:* \$2.13M, 79.5 %, not met, \$46.44M and \$306.8M, growth 17.7 %.

Which target does run (c) miss, and why does a sustainable rule miss it?

8.10 Application area · 4 min

In case 9 the company borrows \$2.5M on a \$3M line, and the bank calls any debt above three times trailing EBITDA over a quarter.

(a) Find the trailing EBITDA below which the bank starts to call the loan, and the trailing EBITDA at the bottom of the recession, when the ceiling is \$1.16M. *Checkpoint:* \$0.83M and \$0.39M.

(b) Using the chapter's results and limitations, say what the 24-month runway target buys and costs, and what a real loan agreement adds that the model leaves out.

8.11 Application area · 4 min

A seller into a holiday peak runs its factory flat out all year at the mean demand of 21.9 thousand units a year. The season of Equation (8.2) has quarter means of 0.789, 0.806, 0.924 and 1.481.

(a) Find the stock the factory must have built by the start of the fourth quarter to fill every order in it, and its book value at the opening unit cost of \$150. *Checkpoint:* 2.63 thousand units and about \$0.39M.

(b) Name two costs of such a pre-build that the model leaves out, and say what case 3 shows about feedback without the calendar.

8.12 Review the agent · 5 min

An AI assistant reads the results of case 8 in the current application and writes: “The economic MPC raised $3M at 1.52 yr and ended with founders’ value of $139.96M, against $108.63M for the same planner without a round. The round is worth $31.33M to the founders, so the company should raise it.” Find the flaw and say how you would check it.

CHAPTER 9

Valve-Position Control of Chilled Water in a Gas Turbine Trigeneration Plant

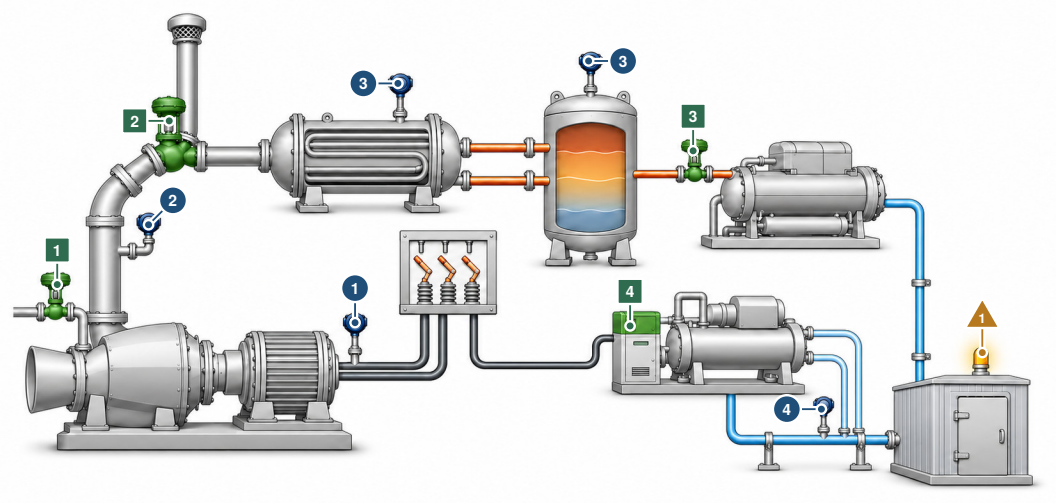

The controller moves 1 fuel valve 2 exhaust bypass 3 hot-water valve of the absorption chiller 4 electric chillers' cooling duty

The controller reads 1 generator power 2 exhaust temperature 3 hot-water temperature 4 chilled-water temperature

What disturbs it 1 heat of the computers in the data hall

Notice that the cheap cooling comes from the turbine's exhaust heat through the slow hot-water chain, while the electric chiller answers fast from the generator's power. Illustration, not to scale.

IN PLAIN TERMS **One Turbine, Power and Cooling**

The main idea. A gas turbine at a computer center makes electricity. Heat from its exhaust warms water, and the hot water drives a cooling machine that runs on heat, not on electricity.

The challenge. The heat-driven cooler is cheap but slow, while the heat of the computers comes fast, so electric coolers catch each change and then hand the work back. No setting helps once the turbine and the power lines have nothing more to give.

Why it matters. Done well, the computers stay cool at the lowest cost in fuel and electricity.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on availability (Section 9.1), the history segment (Section 9.2), the two try-it segments and the fact on the predictive allocator (Section 9.5) and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 9.5), Technology in Practice (Section 9.6), Limitations (Section 9.7) and the open challenge

The control question A data center's gas turbine makes electricity, hot water and chilled water from one fuel valve, and recovered-heat cooling is cheap but loads up over minutes while the hall's heat arrives in about a hundred seconds. How should a fast electric chiller and a slow absorption chiller share one chilled-water temperature, and where does control stop and missing capacity begin?

Learning objectives

- Trace a fuel step through power, recovered heat, hot water and absorption cooling, and name the inventory behind each delay.
- Explain why the hot-water bypass loop decouples a power change from the cooling in automatic operation.
- Compute the absorption chiller's operating point, its valve gain and the heat it rejects from the model's equations.
- Predict the handover energy of a valve-position controller and compare the electric trim with absorption cooling alone.
- Tell a tuning problem from a capacity limit by the runway of a depleting store or UPS.

Prerequisites Sections 1.4.7 (mid-ranging and valve-position control), 1.4.3 (selectors and override), 1.4.4 (external reset), 1.4.9 (actuators that deplete) and 1.6.4 (scenarios and exported files). Chapter 3 shows an exhaust-temperature limiter on a fuel valve. Sections 1.6.3 and 1.4.8 are needed only for the predictive allocator.

The app https://apmonitor.com/apps/gas-turbine-trigeneration/

See also Volume 17, Chapter 2 for gas engines on the same site model, whose power and heat also answer on two clocks, and Volume 17, Chapter 4 for one variable handed between two actuators. In this volume, Chapter 10 for three mid-ranging layers on one power block.

9.1 Introduction

Gas turbines make up 63 percent of the combined heat and power (CHP) capacity installed in the United States, 51.5 GW, according to a catalog prepared for the U.S. Environmental Protection Agency, and the thermal energy in their exhaust is generally 60 to 70 % of the fuel energy [1]. The same catalog describes a campus whose 5 MW turbine sends about 8 MW of recovered heat to space heating in winter and to single-effect absorption chillers in summer. The U.S. Department of Energy counts data centers among the sites with steady year-round cooling that suit an absorption chiller, and notes that such a chiller needs a larger cooling tower than an electric one of the same capacity because it also rejects its driving heat [2]. In the simulated plant a turbine of about 10 MW carries a 6.0 MW data hall at 6.80 MW net, burns 26.2 MW of fuel on its lower heating value (LHV), and returns 11.6 MW of exhaust heat as 95 °C hot water, from which an absorption chiller makes 5.79 MW of the hall's 6.29 MW of cooling.

DID YOU KNOW? Gas Turbines Available 93 to 97 Percent of the Year

A review of on-site generation made for a national laboratory, quoted in a 2009 report on data centers, found that gas turbines in CHP service averaged 93 to 97 % availability and reciprocating engines 96 to 98 %. More than half of the downtime of every kind of on-site generator was scheduled maintenance [3].

The simulator is meant to dislodge a common expectation, namely that more power from a gas turbine buys more cooling a few minutes later. In Manual, with every valve held, a six-point step of the fuel valve moves the net power 63 % of the way in 0.8 s, the recovered heat in 39 s, the hot water in 454 s and the absorption cooling in 573 s. In Auto the same chain is cut. The hot-water loop holds its setpoint by opening the exhaust bypass, so a 1 MW rise of the turbine's setpoint sends 2.73 MW more heat to the stack and changes the absorption cooling by −0.03 MW. Absorption cooling rises only when a controller asks the chiller for it, and the electrical coupling runs the other way, since an

absorption trip adds 0.61 MW of electric-chiller load that the turbine and an import-limited grid must carry.

The control problem has a general shape. One output, the chilled-water temperature, has two manipulated inputs, a fast, expensive electric chiller and a slow, cheap absorption chiller, and the fast one should take every disturbance and then hand the duty back. This is mid-ranging control (Section 1.4.7), which in industry is mostly built as valve-position control, with examples from pilot plants and paper mills [4, 5], and the oxygen concentration of a bioreactor and the temperature of a cooling system [6]. The fuel side has the selector shape of a regulator and a protective limit acting on one valve (Section 1.4.3), and the site carries a store and an import limit. Combined cooling, heating and power is one form of the distributed multi-generation that couples these services at one site [7], while published grid benchmarks solve a tri-generation dispatch by dynamic optimization with perfect foresight [8].

However, the usual teaching material keeps the parts apart. Gas turbine models for power system studies end at the generator terminals and the exhaust [9, 10], trigeneration and absorption-cooling studies assess energy performance of a design or an installation [11, 12], and the mid-ranging literature treats the two-actuator loop without the plant that makes one actuator slow [5]. This chapter works with a data center's gas turbine trigeneration plant simulated in a browser. It measures what a fast electric trim and a slow valve-position return buy against absorption cooling alone, where the turbine's derating and the cooling electricity meet on the bus, and where control stops and capacity begins. The methods are established ones: the cycle calculation, the capacity-lag chiller, the velocity-form PID with back-calculation, the low selector and valve-position control. The results are matched comparisons on a single plant and seed, among them the finding that in automatic operation the plant's own bypass loop decouples a power change from the cooling. The plant's ledger states close every balance independently, its site components are shared with the gas engine cogeneration plant of Volume 17, Chapter 2, and its controllers read only declared meters.

Section 9.2 traces the gas turbine from peaking duty to trigeneration. Section 9.3 builds the plant model, from the gas path at constant speed to the site bus, and Section 9.4 the controllers as implemented, with a predictive chiller allocator as an option. Section 9.5 reports the measurements: the app's cases, electricity against driving heat, where a power change goes, three matched comparisons, the capacity that no tuning adds, the challenge under five policies, and the verification of the implementation. Section 9.6 sets the loops on plant hardware, Section 9.7 states the limitations, Section 9.8 concludes, and Section 9.9 holds the exercises.

9.2 From Peaking Turbines to Trigeneration

The turbine's road from peaking duty to on-site heat and power is told in the history segment below, and the thermodynamics of a trigeneration plant are treated in the cycle literature [13] on the compressor, combustor and turbine relations of gas turbine theory [14].

> ⧖ HISTORY **From Peaking Duty to On-Site Heat and Power, 1930s to 1980s**
>
> Gas turbines have generated stationary electric power since the late 1930s, went on to transform aircraft propulsion in the 1940s, and long served utilities as peaking capacity. By the early 1980s the efficiency and reliability of machines of 1 to 40 MW made them an attractive choice for industrial and large institutional CHP. In the EPA catalog's count, close to 80 % of the gas turbine CHP capacity in the United States sits in large combined-cycle plants that export power, and the rest in simple-cycle systems typically below 40 MW. The catalog also draws a line that this chapter keeps. A turbine whose exhaust raises steam only for a steam turbine is a combined cycle and not CHP, because no end use takes the heat [1], so a plant whose recovered heat drives an absorption chiller is trigeneration and not combined-cycle generation.

The dynamic picture came from the grid. Simplified representations of heavy-duty single-shaft turbines from 18 to 106 MW, with gas and liquid fuel systems and droop and isochronous governors,

Table 9.1 Variable Roles and Starting Values at the Reference Start (25 °C Dry Bulb, 20 °C Wet Bulb, 6.8 MW Site Power Target)

Variable	Role	Start	Range
P_{net}	CV, turbine net power	6.80 MW	SP 1–12
T_4	CV, exhaust (limit)	425.9 °C	limit 490, trip 520
T_h	CV, hot-water supply	95.0 °C	SP 85–98
T_{chw}	CV, chilled-water supply	12.0 °C	SP 9–16, ±1 K
$Q_{ec,cmd}$	CV of the VPC, electric duty	0.50 MW	SP 0–4
u_f, b, v_g	MV, fuel, bypass, absorption valve	65.4, 4.9, 75.9 %	0–100
$Q_{ec,cmd}$	MV of the trim	0.50 MW	0–4.0
P_{IT}	DV, IT load (80 % critical)	6.00 MW	3–8
Q_{req}	DV, heat-export demand	3.60 MW	0–7
T_a, T_{wb}	DV, dry and wet bulb	25, 20 °C	0–45, 5–30
events	DV, trips, pulses, ramps	—	—
$\dot{m}_f$	state, fuel flow	0.5236 kg/s	$\leq$ 0.80 kg/s
T_d, T_x	states, duct, tube metal	413.8, 111.7 °C	—
Q_{cap}, Q_{ec}	states, chiller duties	5.79, 0.50 MW	map, 4.0
T_{IT}, T_{cw}	states, IT node, cooling water	27.0, 29.3 °C	trip 37
E_{cs}, E	states, cold store, UPS	80, 90 %	floors 0, 20

were published in 1983 for power system studies [9], and an IEEE working group extended them in 1994 to combined cycles with waste-heat recovery boilers and steam turbines [10]. Those models say how fast a turbine's power moves and leave the heat after the exhaust duct to the plant downstream.

Absorption refrigeration has been reviewed across its working pairs and cycles [15], and absorption chillers and heat pumps have a book-length treatment of their own [16]. The DOE fact sheet gives the present practice for CHP [2]. Single-stage machines run on hot water at 200 to 240 °F or low-pressure steam and are often paired with reciprocating engines, two-stage machines need hot water near 350 °F or steam near 115 psig and are often paired with combustion turbines, and the thermal coefficient of performance (COP) is 0.70 to 0.79 for the single-stage examples and 1.35 to 1.42 for the two-stage ones. The performance of a trigeneration system built on a micro gas turbine and an ammonia and water absorption chiller has been analyzed [11], absorption refrigeration has been evaluated for recovering the low-grade heat of data centers themselves [12], and multi-generation at one site has been reviewed as a whole [7].

Two facts about the turbine shape the coordination. A turbine reduces its power by lowering its turbine inlet temperature, so part-load efficiency can be substantially below full-load efficiency, and on a hot day its power and its efficiency both fall, for a nominal 7.5 MW machine to 80 % of its ISO power at 100 °F, while in many parts of the United States electricity costs most in summer, when the turbine performs worst [1]. A site that backs up its absorption cooling with electric chillers therefore adds load to its turbine exactly when the turbine has least to give. The control answer to two actuators of different speed on one output is mid-ranging, whose design and tuning were largely done case by case before a systematic method was proposed [5], and whose standard anti-windup schemes can degrade performance while one actuator saturates [6]. Measured-load feedforward and cascade structure are textbook material [17].

9.3 Plant Model

The plant in Figure 9.1 is an invented teaching design, set against public reference data but taken from no installation. Its exhaust passes a duct and a diverter that sends the share b around a heat-recovery unit (HRU) to a bypass stack, the HRU heats a mixed hot-water buffer that supplies a heat-export customer and, through the valve v_g, an absorption chiller, and the chilled-water loop of the data hall also carries two lumped electric chillers and a cold store. There are 13 physical states, $\dot{m}_f$, T_d, b, T_x, T_h, v_g, Q_{cap}, Q_{ec}, T_{IT}, T_{chw}, E_{cs}, T_{cw} and E, and 46 ledger states that integrate every boundary flux. Units inside the model are seconds, watts, joules, kilograms and kelvin, and Table 9.1 gives the roles of the variables.

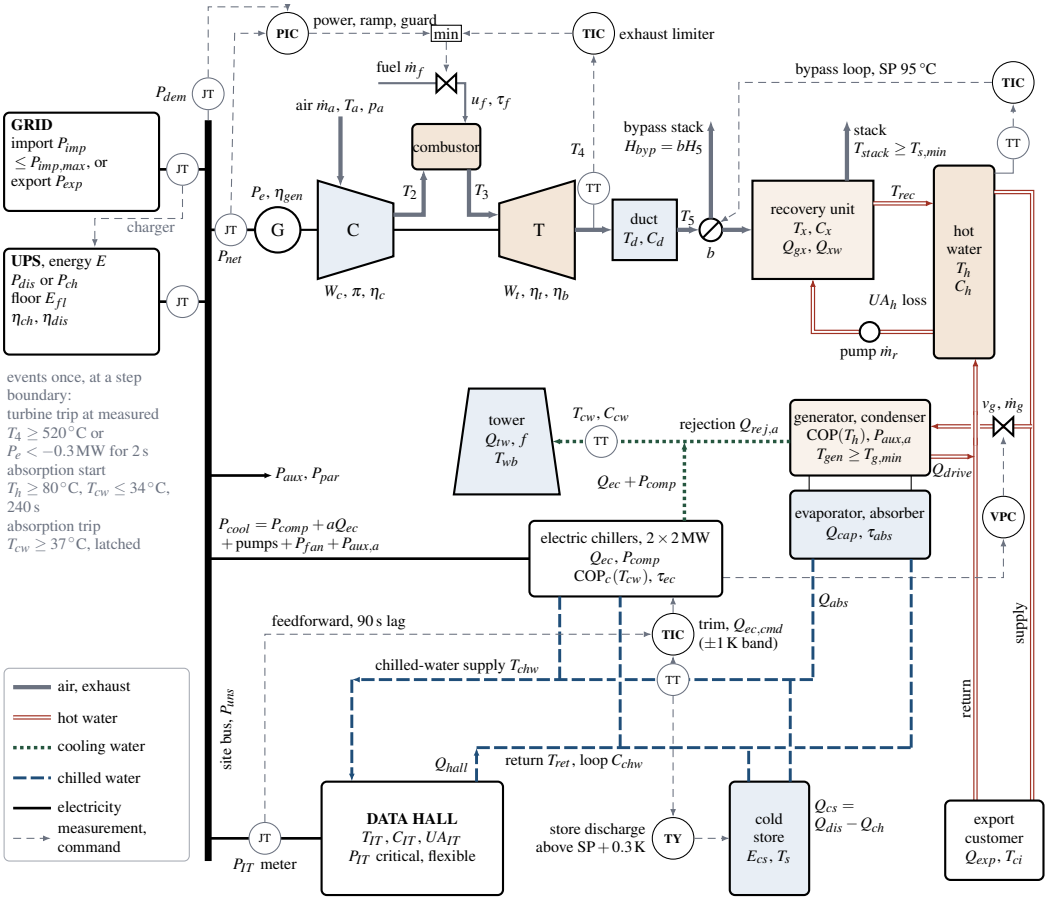

Figure 9.1 The simulated plant, after the application's Plant view. Fuel $\dot m_f$ through the valve u_f and air $\dot m_a$ pass the compressor C, the combustor and the turbine T at constant speed, and the generator G delivers P_e to the bus. The exhaust at T_4 warms the duct metal T_d, leaves it at T_5, and is split by the diverter b between the bypass stack and the recovery unit, whose tubes T_x heat the recovery water into the buffer T_h. The buffer feeds the export customer and, through v_g, the absorption generator, both chillers reject to the cooling water T_{cw} and the tower, and the electric chillers, the evaporator, the cold store and the hall sit on the chilled-water rails. Air and exhaust are thick gray arrows, hot water a hollow double line, cooling water dotted, chilled water long-dashed and electricity solid black. Thin short-dashed lines are measurements and commands: PIC and the exhaust limiter (TIC) meet at a low selector on u_f, the bypass TIC holds T_h, the trim TIC sets $Q_{ec,cmd}$ with metered-IT feedforward, VPC returns the electric duty to its reserve $Q_{ec,SP}$ through v_g, and TY discharges the store. The scripted events and the legend sit left of the bus. There is no rotor state. A process model, drawn without piping detail or scale.

9.3.1 Gas Path at Constant Speed

The grid holds the shaft speed, so the turbine has no rotor state and its shaft balance is algebraic. At fixed inlet guide vanes the compressor swallows a fixed volume, the combustor's outlet follows from the actual fuel and air flows, and the turbine's choked first nozzle ties the pressure ratio to the gas flow and the firing temperature, as shown in Equation (9.1).

$$\begin{aligned} \dot{m}_a &= \dot{m}_{a,d}\,\frac{p_a}{p_{ISO}}\,\frac{T_{ISO}}{T_a}, \qquad \dot{m}_g = \dot{m}_a + \dot{m}_f \\ T_2 &= T_a\left[1 + \frac{\pi^{(\gamma_a-1)/\gamma_a} - 1}{\eta_c}\right] \\ \dot{m}_g c_{pg}(T_3 - T_r) &= \dot{m}_a c_{pa}(T_2 - T_r) + \eta_b \dot{m}_f \,\mathrm{LHV} \\ \pi &= \pi_d\,\frac{p_{1,ISO}}{p_1}\,\frac{\dot{m}_g}{\dot{m}_{g,d}}\sqrt{\frac{T_3}{T_{3,d}}} \\ T_4 &= T_3 - \eta_t T_3\left(1 - \rho^{-(\gamma_g-1)/\gamma_g}\right) \\ \rho &= \pi(1-\Delta p_b)p_1 / \left[p_a(1+\Delta p_x)\right] \end{aligned} \tag{9.1}$$

where $T_{ISO} = 288.15\,\mathrm{K}$ and $p_{ISO} = 101.325\,\mathrm{kPa}$, $p_1 = p_a(1-\Delta p_{in})$ is the compressor inlet pressure, $T_r = 298.15\,\mathrm{K}$ is the one enthalpy reference on which the fuel's LHV is also given, and c_{pa}, γ_a and c_{pg}, γ_g are two constant specific heats, for air through the compressor and for products through the turbine. The three unknowns T_2, T_3 and π are found by eight fixed passes of substitution, whose contraction is about 0.1 per pass. The work, the generator and the first law of the gas path are Equation (9.2).

$$\begin{aligned} W_c &= \dot{m}_a c_{pa}(T_2 - T_a), \qquad W_t = \dot{m}_g c_{pg}(T_3 - T_4) \\ P_{sh} &= W_t - W_c - P_{mech} \\ P_e &= \eta_{gen} P_{sh} \;\text{ if } P_{sh} \geq 0, \quad \text{else } P_{sh}/\eta_{gen} \\ P_{net} &= P_e - P_{GT,aux}, \qquad H_4 = \dot{m}_g c_{pg}(T_4 - T_r) \\ \dot{m}_f \mathrm{LHV} + \dot{m}_a c_{pa}(T_a - T_r) &= P_e + P_{gl} + P_{mech} \\ &\quad + (1-\eta_b)\dot{m}_f \mathrm{LHV} + H_4 \end{aligned} \tag{9.2}$$

where P_{gl} is the generator loss. Because every stream's enthalpy is $c_p(T - T_r)$ on the same reference, the first law holds at every instant by construction. Load falls by lowering the firing temperature at constant air flow, as the catalog describes [1], so the exhaust cools at part load. At the reference start the turbine passes 40.01 kg/s of air at a pressure ratio of 16.20, with $T_2 = 446.5\,^\circ\mathrm{C}$, $T_3 = 947.4\,^\circ\mathrm{C}$ and $T_4 = 425.9\,^\circ\mathrm{C}$, $W_c = 16.95\,\mathrm{MW}$ and $W_t = 24.31\,\mathrm{MW}$, and of its 7.05 MW at the terminals the 6.80 MW net is 25.98 % of the fuel. The fuel follows its valve through $\tau_f\,\mathrm{d}\dot{m}_f/\mathrm{d}t = \min(u_f\dot{m}_{f,max}, F_{max}/\mathrm{LHV}) - \dot{m}_f$ while the turbine runs. The available net capacity $P_{cap}(T_a)$ is the output at which the exhaust reaches its 490 °C limit or the fuel its supply limit, 9.21 MW at 25 °C and 7.95 MW at 38 °C.

9.3.2 Exhaust, Recovery Unit and Hot Water

The duct metal and the HRU tubes are the gas path's heat stores, each stream meets an isothermal metal node with the effectiveness $\varepsilon(UA, C) = 1 - \mathrm{e}^{-UA/C}$, and the buffer and its loop are one mixed

node, as shown in Equation (9.3).

$$\begin{aligned}
Q_d &= \varepsilon(UA_d, C_g)\, C_g (T_4 - T_d), \qquad T_5 = T_4 - Q_d / C_g \\
C_d \frac{\mathrm{d}T_d}{\mathrm{d}t} &= Q_d - UA_{dl}(T_d - T_a), \qquad \tau_b \frac{\mathrm{d}b}{\mathrm{d}t} = b_{cmd} - b \\
Q_{gx} &= \min\big(\varepsilon(UA_g, C_{gx})(T_5 - T_x),\ T_5 - T_{s,min}\big) C_{gx} \\
Q_{xw} &= \varepsilon(UA_w, C_r)\, C_r (T_x - T_h) \\
C_x \frac{\mathrm{d}T_x}{\mathrm{d}t} &= Q_{gx} - Q_{xw}, \qquad H_{byp} = b(H_4 - Q_d) \\
C_h \frac{\mathrm{d}T_h}{\mathrm{d}t} &= Q_{xw} - Q_{drive} - Q_{exp} - UA_h (T_h - T_a) \\
Q_{exp} &= \min\big(Q_{req},\ Q_{perm},\ \varepsilon_c \dot{m}_c c_p (T_h - T_{ci})^+\big)
\end{aligned} \tag{9.3}$$

where $C_g = \dot{m}_g c_{pg}$, $C_{gx} = (1-b)C_g$ and $C_r = \dot{m}_r c_p$ are heat-capacity flows, the minimum caps the gas side only while it gives heat, the stack carries $(1-b)(H_4 - Q_d) - Q_{gx}$, $Q_{rec} = Q_{xw}$ leaves in water at $T_{rec} = T_h + Q_{xw}/C_r$, $C_h = M_h c_p = 209\,\mathrm{MJ/K}$, Q_{perm} is the export the priority logic permits, and $(\cdot)^+ = \max(0, \cdot)$. At the reference start the bypass is 4.9 % open and sends 0.91 MW to the stack, the HRU gives 11.61 MW to water leaving at 108.9 °C, and its stack is at 160.5 °C.

9.3.3 The Absorption Chiller Reduction and Its Sink

The absorption chiller is the capacity-lag reduction of the framework's shared site module `core/cogen.js`, which the gas-engine cogeneration application wrote and this application extracted with the engine's runs unchanged. The lag state is available capacity, not stored energy, and the chiller's internal energy is held constant, so cooling and drive heat are chosen together inside that capacity, as shown in Equation (9.4).

$$\begin{aligned}
Q_{dc} &= v_g \dot{m}_{g,max} c_p\, (T_h - T_{g,min})^+ \\
\mathrm{COP} &= \mathrm{clip}_{[0.55,\,0.76]}\big(0.70 + 0.004\,(T_h - 87.8\,^\circ\mathrm{C})/\mathrm{K}\big) \\
Q_{map} &= Q_{abs,r}\, \mathrm{clip}_{[0,\,1.15]}\big((T_h - 70\,^\circ\mathrm{C})/20\,\mathrm{K}\big) \\
&\quad \times \big(1 - 0.03\,(T_{cw} - 29.4\,^\circ\mathrm{C})^+/\mathrm{K}\big) \\
Q_{tgt} &= \mathrm{clip}_{[f_{td} Q_{map},\, Q_{map}]}\big(\mathrm{COP}\, Q_{dc}\big) \\
\tau_{abs} \frac{\mathrm{d}Q_{cap}}{\mathrm{d}t} &= Q_{tgt} - Q_{cap} \\
Q_{abs} &= s_{cool} \min\big(Q_{cap},\ \mathrm{COP}\, Q_{dc}\big) \\
Q_{drive} &= Q_{abs}/\mathrm{COP} \\
Q_{rej,a} &= Q_{drive} + Q_{abs} + s_{cool} P_{aux,a}
\end{aligned} \tag{9.4}$$

where Q_{dc} is the most heat the generator's water can give before it leaves at $T_{g,min} = 70\,^\circ\mathrm{C}$, $Q_{abs,r} = 6.0\,\mathrm{MW}$ is the rating at 90 °C, $f_{td} = 0.2$ the turndown, the target is zero unless the chiller runs, and s_{cool} is the served fraction of the cooling plant's electricity, 1 whenever the bus serves it. The valve follows its command through $\tau_v = 10\,\mathrm{s}$ and is interlocked shut unless the chiller is starting or running. The COP passes through 0.70 at 87.8 °C and 0.74 at 97.8 °C, which are the DOE fact sheet's single-stage examples on hot water [2], while the capacity map is a teaching map. No driving heat means no cooling, however large the lag state. At the reference start the valve passes 75.9 kg/s, the COP is 0.729, the chiller makes 5.79 MW against a map of 6.90 MW, takes 7.94 MW of drive heat, and rejects $7.94 + 5.79 + 0.06 = 13.79\,\mathrm{MW}$.

The sink boundary is the cooling-water node of Equation (9.5), whose tower is bounded by the wet bulb and run by its own proportional fan control.

$$\begin{aligned} C_{cw}\frac{\mathrm{d}T_{cw}}{\mathrm{d}t} &= Q_{rej,a} + s_{cool}(Q_{ec} + P_{comp}) - Q_{tw} \\ Q_{tw} &= f\,UA_{tw}(T_{cw} - T_{wb})^{+} \\ f &= \mathrm{clip}_{[0,1]}\big((T_{cw} - 28\,^{\circ}\mathrm{C})/2\,\mathrm{K}\big) \end{aligned} \tag{9.5}$$

with fan power $P_{fan} = P_{tw} f^3$. Per megawatt of cooling the absorption chiller rejects 2.37 MW to this node and an electric chiller at 29 °C 1.17 MW, which is the DOE statement about tower size in numbers. A hotter sink acts on both machines. Raising the cooling water from 29 to 37 °C cuts the absorption map at 95 °C from 6.90 to 5.33 MW, the electric COP from 6.00 to 4.80 and the electric capacity from 4.00 to 3.60 MW, while the absorption COP, read at the hot-water inlet only, does not change.

9.3.4 Electric Chillers, Cold Store and Data Hall

The two 2 MW electric chillers are one lumped machine whose duty is a state, as shown with the hall, the chilled water and the store in Equation (9.6).

$$\begin{aligned} \tau_{ec}\frac{\mathrm{d}Q_{ec}}{\mathrm{d}t} &= \mathrm{clip}_{[0,\,Q_{ec,max}(T_{cw})]}(Q_{ec,cmd}) - Q_{ec} \\ P_{comp} &= Q_{ec}/\mathrm{COP}_c(T_{cw}), \qquad P_{ec} = P_{comp} + aQ_{ec} \\ C_{IT}\frac{\mathrm{d}T_{IT}}{\mathrm{d}t} &= P_{IT,del} - UA_{IT}(T_{IT} - T_{chw}) \\ Q_{hall} &= UA_{IT}(T_{IT} - T_{chw}) + \tfrac{1}{2}s_{aux}P_{aux} \\ &\quad + UA_{env}(T_a - 22\,^{\circ}\mathrm{C}) \\ C_{chw}\frac{\mathrm{d}T_{chw}}{\mathrm{d}t} &= Q_{hall} - Q_{abs} - s_{cool}Q_{ec} \\ &\quad - (Q_{dis} - Q_{ch}) \\ \frac{\mathrm{d}E_{cs}}{\mathrm{d}t} &= Q_{ch} - Q_{dis} - UA_s(T_a - T_s) \end{aligned} \tag{9.6}$$

where $\mathrm{COP}_c = \max\big(1.5,\, 6.0[1 - 0.025(T_{cw} - 29\,^{\circ}\mathrm{C})]\big)$ excludes the auxiliaries $a = 10\,\mathrm{kW/MW}$, $Q_{ec,max}(T_{cw}) = 4.0\,\mathrm{MW}\,\mathrm{clip}_{[0.3,1]}\big(1 - 0.02(T_{cw} - 32\,^{\circ}\mathrm{C})^{+}\big)$, and $P_{IT,del}$ is the IT power actually delivered rather than requested. The store's state is its stored cooling, $E_{cs} = M_s c_p(18\,^{\circ}\mathrm{C} - T_s)$ with 1.5 MWh between 18 and 12 °C, so its temperature is derived and never integrated twice. Discharge cools the return water at $T_{ret} = T_{chw} + Q_{hall}/(\dot{m}_{chw} c_p)$ and is limited by $\varepsilon_s C_s(T_{ret} - T_s)$, charge by $\varepsilon_s C_s(T_s - T_{chw})$, with $\varepsilon_s = 0.8$ and C_s the store loop's 250 kg/s. At the reference start the hall's 6.29 MW is 6.00 MW of IT, 0.20 MW of auxiliary heat and 0.09 MW through the envelope, $T_{IT} = 27.0\,^{\circ}\mathrm{C}$, and T_{IT} lumps the equipment into one effective node that stands for no chip or inlet-air temperature. The air side of a hall, with its inlet temperature, is the subject of Volume 9, Chapter 8. The electric chillers carry their 0.50 MW reserve at a COP of 5.96 in 29.29 °C cooling water, whose tower rejects 14.38 MW at 64.5 % fan.

9.3.5 Site Bus, UPS and Import Limit

The shared module closes the electrical balance with nonnegative terms, evaluated term by term, as shown in Equation (9.7).

$$\begin{aligned} &P_{gen} + P_{imp} + P_{dis} + P_{uns} \\ &\quad = P_{IT} - P_{curt} + P_{cool} + P_{aux} + P_{par} + P_{ch} + P_{exp} \\ &\frac{\mathrm{d}E}{\mathrm{d}t} = \eta_{ch} P_{ch} - P_{dis}/\eta_{dis} \\ &P_{dis} \le \min\left(P_{dis,max},\ (E - E_{fl})\eta_{dis}/\Delta t\right) \end{aligned} \tag{9.7}$$

where $P_{gen} = P_{net}^{+}$, the turbine's auxiliaries appear as P_{par} on the demand side when $P_{net} < 0$, P_{cool} is the compressor power plus aQ_{ec}, the chilled-water, cooling-water and hot-water pumps (0.12 MW, 0.15 MW, 0.05 MW), the tower fans and the absorption chiller's 60 kW, and $P_{aux} = 0.40$ MW. The grid covers whatever remains once the UPS's charging is counted, it never imports and exports at once, and import above $P_{imp,max} = 1.0$ MW is taken back from the UPS's charging, then met by discharge inside its step limit, and the rest is unserved, flexible IT first, then auxiliaries, cooling and critical IT. Each step's limits are computed once from the inventory that the step begins with and held through the four Runge-Kutta stages, whose nonnegative weights cannot carry the stored energy past its floor, so no inventory is clipped and no step is split, and the cold store obeys the same rule. The electric chillers' electricity comes from their state in the same evaluation as the bus, so backup feeds the generation balance with no hidden delay. The UPS and the cold store are depleting stocks in the sense of Section 1.4.9. Usable UPS energy from 90 % to the floor is $0.5 \times 0.70 \times 0.95 = 0.3325$ MWh.

9.3.6 Instruments, Events and Numerics

Controllers read only a measurement packet. Every analog reading is the truth filtered by a first-order transmitter lag with noise added afterward, so the trends show the configured noise, and each scan makes a fixed count of random draws. The generator power transducer lags 0.1 s with 5 kW of noise, the averaged exhaust thermocouples 4 s with 0.3 K, the water RTDs 3 s with 0.03 K on the buffer, the chilled-water RTD 5 s with 0.01 K, the cooling-water RTD 5 s, the absorption and store heat meters 10 s, the electric chillers' duty signal 2 s, the grid meter 0.2 s with 5 kW and the UPS meter 0.1 s. The demand and IT meters publish one-second averages as each second closes, which delays a load step by 1 to 2 s before the feedforward sees it, and the UPS charge comes from its battery monitor at 1 Hz. The turbine's and the chiller's states are digital status words.

Discrete events fire once each, between integration steps, and leave a mark on the trends. The turbine trips when the measured exhaust reaches 520 °C or when P_e stays below −0.3 MW for 2 s, and a trip opens the breaker and stops the gas path at once, leaving 0.08 MW of auxiliaries on the bus. The absorption chiller starts when its measured permissives hold ($T_h \ge 80$ °C, $T_{cw} \le 34$ °C), runs 240 s later, and trips and latches at a cooling water of 37 °C, while a scripted fault clears at its time and resets itself. Every case begins lined out at the steady state its own settings imply, found by bisection for the tower and the bypass and by 40 fixed-point passes for the coupled plant. Classical fourth-order Runge-Kutta advances the states (Section 1.6.1) every 0.05 s, whose fastest retained mode is the 0.5 s fuel valve, and the historian logs one row per second. The parameters are listed in Table 9.2 and the remaining symbols in Table 9.3.

9.3.7 Assumptions and Where They Fail

The constant-speed reduction fits a grid-connected machine and leaves out the rotor, frequency, a coast-down, a restart and a cold start, and the fixed inlet guide vanes leave out the schedules that hold exhaust temperature at part load. The two constant specific heats make η_c and η_t effective values of this model rather than component efficiencies, and Section 9.5.6 gives the size of that approximation. The duct, the tubes and the buffer are single nodes, so a stratified tank or a distributed HRU would add delay of its own. No solution chemistry, crystallization risk or internal heat storage is modeled

Table 9.2 Model Parameters. The Class Is the Source of Each Value (L From the Literature, F Fitted, T Invented for Teaching) and Its Visibility in the Parameters Panel (V Shown, H Hidden)

Symbol	Meaning	Value	Unit	Class
Gas turbine				
$\dot{m}_{a,d}$	air flow at ISO	41.4	kg/s	F, V
π_d, $\dot{m}_{g,d}$, $T_{3,d}$	swallowing reference	17.7, 42.1, 1078	—, kg/s, °C	L, L, F, V/H
η_c, η_t	effective efficiencies	0.86, 0.87	—	F, V
η_b, P_{mech}	combustion, mechanical loss	0.995, 0.13	—, MW	T, H
c_{pa}, γ_a	air	1005, 1.4	J/(kg K), —	T, H
c_{pg}, γ_g	products	1150, 1.333	J/(kg K), —	T, H
LHV	fuel at 25 °C	50.0	MJ/kg	T, H
Δp_{in}, Δp_b, Δp_x	pressure losses	1, 4, 3	%	T, H
η_{gen}	generator	0.975	—	T, V
$P_{GT,aux}$	running, tripped	0.25, 0.08	MW	T, V/H
$\dot{m}_{f,max}$, F_{max}	valve open, supply	0.80, 40	kg/s, MW	T, H/V
τ_f, r	valve lag, setpoint ramp	0.5, 1.0	s, MW/min	T, V
$T_{4,lim}$	exhaust limit	490	°C	F, V
$T_{4,al}$, $T_{4,tr}$	alarm, trip	500, 520	°C	T, V
Heat recovery				
C_d, UA_d, UA_{dl}	duct metal	1.5, 15, 0.4	MJ/K, kW/K	T, V/H
C_x, UA_g, UA_w	HRU tubes, gas, water	20, 82, 1500	MJ/K, kW/K	T, V/H
$T_{s,min}$, $\dot{m}_r$	stack minimum, water flow	120, 200	°C, kg/s	T, V
M_h, UA_h, τ_b	buffer, loss, diverter lag	50, 1.0, 5	t, kW/K, s	T, V/H
$\dot{m}_c$, T_{ci}, ε_c	customer loop	60, 60, 0.8	kg/s, °C, —	T, H/V
$T_{h,SP}$	hot-water setpoint	95	°C	T, V
Absorption chiller				
$Q_{abs,r}$, τ_{abs}	rating at 90 °C, capacity lag	6.0, 180	MW, s	T, V
COP	at 87.8, 97.8 °C	0.70, 0.74	—	L
$T_{g,min}$, f_{td}	generator outlet, turndown	70, 0.2	°C, —	T, H
$\dot{m}_{g,max}$, τ_v	valve open flow, lag	100, 10	kg/s, s	T, V/H
$P_{aux,a}$, t_{start}	pumps, start sequence	60, 240	kW, s	T, H/V
Electric chillers, store and hall				
$Q_{ec,max}$, τ_{ec}	capacity, loading lag	4.0, 30	MW, s	T, V
COP_c	at 29 °C, −2.5 %/K	6.0	—	T, V
E_{cs}, ε_s	store, 18 to 12 °C	1.5, 0.8	MWh, —	T, V/H
Q_{dis}, Q_{ch} max	store ratings	3.0, 1.5	MW	T, V/H
UA_s	store heat gain	2	kW/K	T, H
C_{IT}, UA_{IT}	IT node	40, 0.4	MJ/K, MW/K	T, H
C_{chw}, $\dot{m}_{chw}$	chilled-water loop	400, 250	MJ/K, kg/s	T, V/H
UA_{env}	envelope above 22 °C	30	kW/K	T, H
Cooling water, site and UPS				
C_{cw}, UA_{tw}	loop, tower at full fan	200, 2.4	MJ/K, MW/K	T, H/V
P_{tw}	tower fans at full	0.20	MW	T, H
P_{IT} shares	critical, flexible	80, 20	%	T, V
$P_{imp,max}$	grid import limit	1.0	MW	T, V
$P_{dis,max}$, E_{cap}	UPS rating, energy	2.0, 0.5	MW, MWh	T, V
η_{ch}, η_{dis}, E_{fl}	UPS efficiencies, floor	95, 95, 20	%	T, H/V

in the absorption chiller, and its COP and capacity read the hot-water inlet only, so the DOE points anchor the COP at the inlet and not the capacity at the model's flow, since the generator cools its water from 95 to 70 °C where the DOE examples cool it by 5 and 10 K. The plant runs a single-stage machine on hot water although the DOE sheet associates two-stage machines with turbine CHP [2], a choice of one hot-water circuit for teaching. The tower carries no water balance and supports no claim about water use, and the grid is an import-limited source with a working voltage. All sizes, settings and thresholds are invented for teaching, and actual limits belong to the manufacturers' data and the governing codes.

9.4 Controllers as Implemented

One framework scan of $\Delta t_c = 0.25$ s runs the instruments, the import guard, the setpoint ramp, the power loop and the exhaust limiter. An integer counter inside it runs the hot-water, chilled-water,

Table 9.3 Symbols for Heat Flows, Temperatures and Electrical Terms

Symbol	Meaning	Symbol	Meaning
T_2, T_3, T_4	compressor, firing, exhaust	π, ρ	pressure, expansion ratio
W_c, W_t	compressor, turbine work	P_e, P_{net}	terminals, net power
H_4, H_{byp}	exhaust, bypassed heat	T_5, Q_d	after duct, to duct
Q_{gx}, Q_{xw}	gas to tubes, to water	Q_{rec}, T_{rec}	recovered heat, outlet
Q_{exp}, Q_{req}	export taken, asked	Q_{perm}	export permitted
Q_{dc}, Q_{drive}	drive capacity, drive heat	Q_{map}, Q_{tgt}	capacity map, target
Q_{abs}, $Q_{rej,a}$	absorption, its rejection	s_{cool}, s_{aux}	served fractions
$Q_{ec,cmd}$, Q_{ec}	electric command, duty	P_{comp}, P_{ec}	compressor, total
Q_{hall}, T_{ret}	hall heat, return	Q_{dis}, Q_{ch}	store out, in
T_s, S	store temperature, charge	Q_{tw}, f	tower duty, fan
P_{gen}, P_{par}	generation, parasitics	P_{cool}, P_{aux}	cooling plant, auxiliaries
P_{imp}, P_{exp}	grid import, export	P_{dis}, P_{ch}	UPS out, in
P_{uns}, P_{curt}	unserved, curtailed	P_{dem}, P_{cap}	demand meter, capacity

Table 9.4 Loops, Scans, Default Tuning and Action (+1 Reverse Acting, −1 Direct), With the Measured Response to a Manual Step From the Reference Start

Loop	PV → MV	Scan (s)	K_c	τ_I (s)	Act.	Step: response
power	$P_{net} \to u_f$	0.25	1.5 %/MW	2	+1	2 pt: +0.344 MW
limiter	$T_4 \to u_f$	0.25	0.5 %/K	10	+1	2 pt: +9.35 K
hot water	$T_h \to b$	1	10 %/K	300	−1	5 pt: −0.60 K
trim	$T_{chw} \to Q_{ec,cmd}$	1	4 MW/K	300	−1	0.2 MW: −0.145 K
abs. alone	$T_{chw} \to v_g$	1	8 %/K	900	−1	5 pt: −0.25 K
VPC	$Q_{ec,cmd} \to v_g$	5	10 %/MW	200	−1	—
store (P)	$T_{chw} \to Q_{dis}$	1	5 MW/K	—	—	0.5 MW: −0.40 K

store-discharge and UPS-charger logic every 1 s and the reserve layer, the valve-position controller, the store's recharge and the export priority, every 5 s, each given the interval elapsed since it last ran. Counters, filters and PID states are created again at every reset. Table 9.4 lists the loops with their measured plant signs.

9.4.1 One PID Implementation

All loops run the velocity-form PID of Section 1.6.2, Equations (1.12) to (1.15), with the action $a = \pm 1$ of Table 9.4, the tracking time $T_t = \tau_I$ and, as the control period, the interval since the loop last ran. Derivative action and output slew limits are unused in every loop. In Manual, and for any loop that is not in charge of its actuator, `track()` sets the integral so that the output matches the applied value, which makes every transfer bumpless, and no PID is both updated and tracked during the same scan.

9.4.2 Power, Exhaust Limit, Import Guard and Bypass

The power loop and the exhaust limiter meet at a low selector on the fuel valve, as the speed governor and the exhaust limiter of a jet engine do in Chapter 3, and the limiter waits just above the valve that was selected, so it cannot wind up while it does not limit, as shown in Equation (9.8).

$$
\begin{aligned}
u_x &= \mathrm{PI}_x\big(T_{4,lim} - T_{4,m}\big) \text{ on } [0,\ \bar{u}_x] \\
\bar{u}_x &= \min\big(u_{cap},\ u_{k-1} + 2\big) \\
u_p &= \mathrm{PI}_p\big(P_{sp} - P_{net,m}\big) \text{ on } [U_{min},\ \bar{u}_p] \\
\bar{u}_p &= \max\big(U_{min},\ \min(u_{cap},\ u_x)\big) \\
u_k &= \min(u_p,\ u_x) \\
P_{req} &= \mathrm{clip}_{[1,12]\,\mathrm{MW}}\big(\max(P_{site},\ P_{need})\big) \\
P_{need} &= P_{dem,m} + P^{-}_{UPS,m} - P_{imp,max} + m
\end{aligned}
\tag{9.8}
$$

where valves are in percent, u_{cap} is the valve the fuel supply allows, $U_{min} = 20\,\%$ holds stable combustion, $P^{-}_{UPS,m}$ is the metered UPS charging, $m = 0.2\,\text{MW}$ is the guard's margin, and the setpoint P_{sp} moves toward P_{req} at 1 MW/min. The power loop's ceiling is the limiter's output, so back-calculation holds its integral against whichever constraint binds, the external reset of Section 1.4.4. The limiter is protection and stays effective in Manual, where its ceiling sits two points above the operator's request.

A PI on the buffer's RTD moves the diverter, opening it when the water is above its 95 °C setpoint. The bypass can only reject heat, so when recovered heat falls short the loop shuts it, reads LIMITED, and the hot water falls. This loop decouples electricity from cooling in Auto, because it rejects any surplus that a power change sends down the exhaust.

9.4.3 Chilled Water: Electric Trim or Absorption Alone

With the trim on, a PI on the chilled-water RTD computes the hall's total cooling demand in megawatts, with the IT meter's reading as its bias, and the electric chillers make that demand less what the absorption chiller and the store deliver on their own heat meters, as shown in Equation (9.9).

$$\begin{aligned} B &= Q_{abs,m} + Q_{cs,m}, \qquad u_b = \hat{P}_{IT} + 0.2\,\text{MW} \\ C &= \text{PI}_t\left(T_{chw,SP} - T_{chw,m}\right) \\ &\quad \text{on } [B,\ B + Q_{ec,max}(T_{cw,m})] \\ Q_{ec,cmd} &= \text{clip}_{[0,\,Q_{ec,max}]}(C - B) \end{aligned} \tag{9.9}$$

where $\hat{P}_{IT}$ is the IT meter through a 90 s first-order lag, matched to the hall node's $C_{IT}/UA_{IT} = 100\,\text{s}$, and 0.2 MW is the commissioned estimate of the hall's other heat. Turning the feedforward on or off shifts its bias change into the integral, so the output holds. With the trim off, a PI on the same RTD moves v_g through the absorption chiller alone and the electric chillers hold their manual duty of 0.5 MW.

9.4.4 Valve-Position Control of the Electric Reserve

The valve-position controller (VPC) is a PI every 5 s whose measurement is the trim's electric command and whose setpoint is the reserve $Q_{ec,SP} = 0.5\,\text{MW}$. While the electric chillers carry more than their reserve it opens the absorption valve, $v_{cmd} = \text{PI}_v(Q_{ec,SP} - Q_{ec,cmd})$ with action -1, so the cheap recovered-heat cooling takes the load back and the trim, seeing the absorption meter rise in B, gives the electric duty up. This is the VPC law of Equation (1.8) with a proportional term added, acting on the trim's command before the electric chillers' lag. Chapter 10 stacks three such layers on one power block. The VPC runs only with the trim on and the chiller running, and otherwise it tracks the valve that is applied. Its output saturates at a wide-open valve when the absorption map is reached, and its tracking time keeps the integral from winding up there, which is the situation that the mid-ranging anti-windup literature studies [6].

9.4.5 Store, Export Priority, Charger and Curtailment

On automatic the cold store discharges in proportion to how far the measured chilled water is above its setpoint plus 0.3 K, $Q_{cs,req} = \text{clip}_{[0,3]\,\text{MW}}\left(K_{c,s}(T_{chw,m} - T_{chw,SP} - 0.3\,\text{K})\right)$, with $K_{c,s} = 5\,\text{MW/K}$ by default. It recharges at 0.6 MW toward its 80 % target when the chilled water is within 0.3 K of its setpoint, the electric chillers have more than 1 MW of room, and neither the import guard nor the exhaust limiter acts. The optional export priority cuts the permitted export by 0.5 MW per 5 s scan while the hot water is more than 2 K below its setpoint. The UPS charger asks for 0.3 MW when the charge is below 90 % and the grid meter is 0.5 MW below its limit, and it stops at 92 % or within 0.3 MW of the limit. Curtailment sheds the flexible 20 % of the IT load by plan.

Three models are kept apart. The plant is Section 9.3, no model is identified from data, and the original controllers know only commissioned constants, the 90 s feedforward lag, the 0.2 MW heat

estimate (the hall's other heat is 0.29 MW), the electric capacity against the measured cooling water, and the store's charge from its own temperature sensor.

9.4.6 Predictive Chiller Allocation

The optional MPC (Section 1.6.3) makes a joint decision on absorption valve fraction v and electric duty request q_e in MW, an allocation of one cooling demand between two actuators with prices as in Section 1.4.8. Its four-state reduction carries absorption cooling q_a, electric cooling q_{ed}, chilled water T_c and hot water T_h, with temperatures in degrees Celsius. All states initialize from the instruments every 15 s. Nominal balances are

$$\begin{aligned} \dot{q}_a &= (v\bar{q}_a(T_h, T_{cw}) - q_a)/\tau_a, \\ \dot{q}_{ed} &= (q_e - q_{ed})/\tau_e, \\ C_c\dot{T}_c &= 10^6(q_{hall} - q_a - q_{ed} - q_{cs}), \\ C_h\dot{T}_h &= 10^6(q_{rec} - q_a/\mathrm{COP}_a - q_{cust}) - Q_{hl}. \end{aligned} \tag{9.10}$$

Here $\bar{q}_a$ is the smaller of nominal full-valve map capacity and COP times the generator's nominal driving-heat limit. Capacities, COP, flows and inventories use declared defaults. The absorption approximation omits turndown, instantaneous heat clipping and startup capacity history. A shut valve has a zero target rather than free turndown cooling. At 95 °C hot and 29 °C cooling water, its sampled steady error reaches 0.545 MW at 75% valve, 7.9% of full map capacity. This approximation keeps a nonzero local valve gain for restarting.

Available return/supply and HRU-outlet/buffer temperature differences, with nominal flows, infer q_{hall} and q_{rec}. Current store flow q_{cs}, customer request, cooling-water and ambient readings are held. No future workload, fault, true capacity state or changed plant coefficient is read. Previous raw one-step thermal residuals update rate biases with gain 0.2, bounded at 0.02 K/s cold and 0.05 K/s hot. These enter the affine water balances rather than identifying physical parameters, consistent with disturbance correction in [18]. The local model is centrally linearized and exactly discretized. Sixty steps and twelve move blocks span 900 s in the shared quadratic optimizer [19]. Nominal lags are 180 and 30 s. A closed-form electric lag check agrees within 3.4×10^{-14}, and a two-control optimizer check within 8.9×10^{-16}. Raw one-step residual RMS in the new case is 0.112 and 0.143 MW for the duties, and 0.0158 and 0.0395 K for the water temperatures. Long-horizon accuracy is not implied.

Tracking weights on $[q_a, q_{ed}, T_c, T_h]$ are $[0, 0, 100, 0.05]$ and move weights are $[1, 0.05]$. Linear duty prices 0.005 for absorption and 0.2 for electric are relative resource proxies, not tariffs. Valve 0–1, zero valve when absorption is unavailable, current electric actuator capacity and block moves of 0.12 and 0.6 MW are hard. Chilled water within ±1 K and hot water 80–100 °C are soft, with quadratic prices 1000 and 20 and linear prices 10 and 1. Turbine control, bypass, store and protection retain their original logic. Cooling and valve-position PIs track the selected commands once per scan. A rejected plan holds the handover within current capacities, marks fallback and returns to those loops until a valid allocation resumes. The priming call computes the first plan.

9.5 Measured Behavior

A script drives the shipped application at its default seed (Section 1.6.1). It opens each case the way the case menu does, applies any change named in the text before time starts, which the engine then treats as part of the starting state, and integrates at the model's 0.05 s step. The chilled-water error is $e = T_{chw} - T_{chw,SP}$ on the simulator's truth at every step. Integrated absolute error (IAE) is $\int |e|\,\mathrm{d}t$ over the stated window, recovery marks the final moment inside the window when $|e|$ is outside the stated band, the reserve return is the first instant after the electric peak from which the delivered duty Q_{ec} stays within 0.1 MW of its reserve to the end of the run, travel is the total variation of a state at the physics step, saturation counts the seconds with v_{cmd} at 100 % or $Q_{ec,cmd}$ at its capacity, and energies are ledgers from $t = 0$. Table 9.5 lists the eleven cases with their completion rules (Section 1.6.4).

Table 9.5 The Eleven Cases, a Measured Outcome of Each as Shipped, and What Completes Each

#	Case and events	Outcome as shipped	Required to complete (kind)
1	Manual, 25 °C, the learner opens the fuel valve 65.4 → 71.4 %	63 % of the rise in power 0.8 s, recovered heat 39 s, hot water 454 s, absorption 573 s	after 30 s open the valve to at least 71 %, run on to 900 s, confirm the comparison of clocks (experiment)
2	30 °C, target 6.8 → 9.8 → 6.8 MW at 60, 420 s	slow loop 8.16 MW at 600 s, with $\tau_I = 2$ s limiter from 204 s at 8.7 MW, back within 0.05 MW at 648 s	set $\tau_I \leq 5$ s, see EXHAUST LIMIT, run past 660 s with power within 0.1 MW of setpoint (experiment)
3	Trim off, store off, export 3.0 MW, IT +0.6 MW at 60 s	T_{chw} peak +0.76 K at 1192 s, +0.57 K at 1800 s	see $e > 0.5$ K, raise $K_{c,a}$ to at least 20 %/K, Reset, run 1200 s, confirm (experiment)
4	As case 3 with the trim on, pulse +1 MW for 10 min pressed by the learner	$\lvert e \rvert \leq 0.012$ K, electric peak 0.97 MW, back within 0.1 MW of reserve from 806 s	see the trim above its reserve, see it back after 900 s, press the pulse, run 300 s, confirm (experiment)
5	Target 8.6 MW, 25 → 38 °C dry bulb, 20 → 26 °C wet bulb, 60–660 s	capacity 9.21 → 7.95 MW, limiter from 372 s, T_{cw} 29.3 → 32.3 °C, cooling electricity 0.52 → 0.70 MW	run to EXHAUST LIMIT and to 900 s, confirm the comparison of both sides (exploration)
6	Export customer lost at 60 s	bypass 4.8 → 52 % at 600 s, up to 9.6 MW to the stack, turbine moved 0.004 MW	run until the bypass is more than 40 % open, confirm (demonstration)
7	Absorption trip at 60 s, fault cleared at 360 s	electric at 4.0 MW by 69 s, store up to 2.4 MW, T_{chw} peak 12.76 °C, demand +0.61 MW, store 80 → 57 %	see the chillers at capacity and the store giving, run until it recharges, confirm (demonstration)
8	One electric chiller out, absorption lost at 60 s, not cleared	band lost at 207 s, store empty at 1549 s, T_{chw} 19.4 °C at 1800 s	run until $\lvert e \rvert > 1$ K, curtail flexible IT, run on to 900 s, confirm (demonstration)
9	38 °C, UPS 60 %, IT 6.0 → 8.5 MW over 60–360 s	limiter from 262 s, UPS floor at 954 s, 226 kWh of flexible IT unserved, critical 0	see the limit with the UPS giving, change K_c, run until the floor with load unserved, confirm (demonstration)
10	MPC allocation, same workload step as case 4	compressor energy increases while drive heat decreases	select MPC and run 900 s, compare both chillers and water temperatures (experiment)
11	Challenge, Table 9.8	chilled water outside ±1 K for 600 s, store 23.6 %, fails	pass the shared verdict at 1800 s with the given parameters and disturbances (challenge)

The app's menu numbers the predictive case 10 and the challenge, the last entry, 11, marked with a target.

9.5.1 Electricity Against Driving Heat

Case 10 repeats case 4's workload step with the store Off. The minimum completion run is 900 s, while the comparisons below cover 1800 s at seed 20261021. Both controllers have the same instruments and no event preview. Table 9.6 uses the framework's IAE on the measured chilled-water signal, matching its scorecard. For comparison with the original truth-based results, physical-step true IAE is 13.881 for MPC versus 5.164 K s for PI. Figure 9.2 plots the true water error and requested electric duty.

The MPC uses 0.079 49 MWh of compressor electricity against 0.049 94 MWh under the loops, and 0.304 90 MWh of total cooling electricity against 0.275 95 MWh. Absorption drive heat falls to 4.04096 from 4.287 69 MWh. Closing hot-water heat above 0 °C is 5.51923 versus 5.524 25 MWh, and cold store 1.188 22 MWh in both. Thus the relative heat price trades more compressor work and error for less driving heat, not a net efficiency gain. The 31 runs measured for this comparison also hold six hindsight trim and valve-position settings and three constant-electric alternatives with zero absorption gain. Halving the physics step preserves the new-case orderings.

After the heat customer is lost, MPC instead uses 0.02547 versus 0.041 93 MWh compressor electricity, but drive heat rises to 4.0909 from 3.9586 MWh and true IAE to 21.975 from 7.286 K s. On the hot afternoon it improves true IAE to 13.131 from 29.328 K s, but total cooling electricity rises to 0.33130 from 0.329 59 MWh. After an absorption trip, its simplified startup forecast spends more backup energy and closes with 0.809 versus 1.019 MWh cold reserve. It adds no missing chiller capacity. The suggested challenge passes with both, but MPC has greater measured IAE 435.67 versus 400.52 K s, and closes with 0.769 versus 0.872 MWh cold reserve. Both fail the shipped challenge settings.

Table 9.6 Matched 1800 s Allocation Runs, PI / MPC

Case	Meas. IAE	Compressor MWh	Drive MWh	Plans / μs
New	15.95 / 20.04	.0499 / .0795	4.288 / 4.041	121 / 1052
Hot day	35.18 / 19.41	.0498 / .0512	4.146 / 4.134	121 / 720
Heat loss	16.77 / 26.37	.0419 / .0255	3.959 / 4.091	121 / 676
Abs. trip	505.65 / 544.76	.1812 / .2980	2.630 / 1.326	121 / 933

IAE in K s. Plans and mean Node solve microseconds are MPC values, with no rejected plan in these runs.

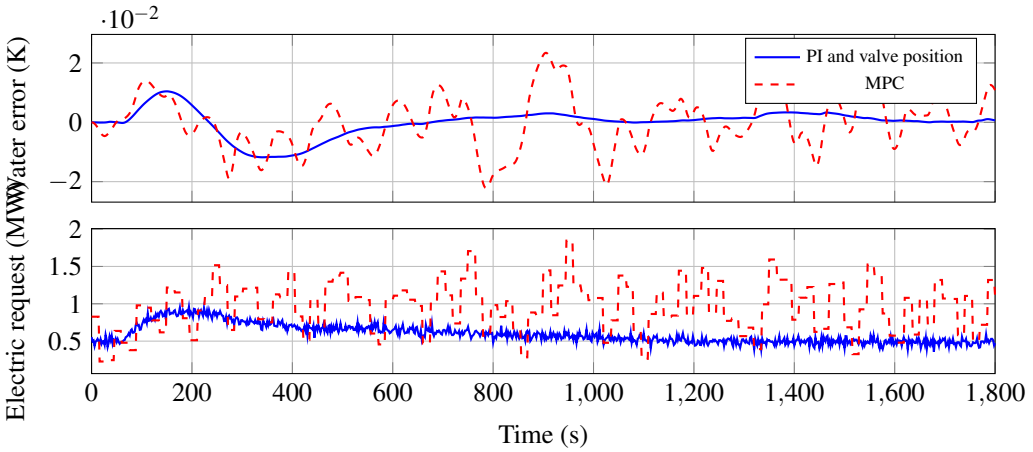

Figure 9.2 Case 10, workload increase at 60 s, PI and valve position (solid) and MPC (dashed). MPC prices driving heat and therefore keeps more electric cooling than the original valve-position loop. Both remain inside the soft band, but the larger forecast horizon does not improve tracking on this case.

Across six Chrome viewport/theme runs, forty plans per run average 2.43–4.87 ms. Planner duty is 0.016–0.032% core utilization during real-time playback, rising to 1.62–3.24% at 100-times playback, excluding state preparation.

> ✦ **DID YOU KNOW?** **A Longer Look Ahead, a Larger Error**
>
> From the simulator: planning 900 s ahead does not buy tighter control here. On the new case the predictive allocator's measured IAE is 20.04 K s against 15.95 K s for the PI and valve-position loops, because it prices driving heat and uses less of it, 4.041 against 4.288 MWh (Table 9.6).

9.5.2 Where a Power Change Goes

A power change reaches the chillers minutes later only in Manual, so the app's one-line description names what Auto does: electric chillers take the step and hand it back to the absorption chiller. In case 1, with the bypass, the absorption valve and the electric chillers held, the six-point valve step raises the measured power from 6.80 to 7.83 MW, the recovered heat from 11.6 to 12.5 MW, the hot water from 95.0 to 97.3 °C and the absorption cooling from 5.79 to 6.33 MW by 930 s, and the chilled water drifts from 12.0 to 11.5 °C, as shown in Figure 9.3a. The four clocks are the valve lag, the 21.5 MJ/K of duct and tube metal, the 209 MJ/K of hot water and the 180 s capacity lag, with no rotor store between fuel and shaft.

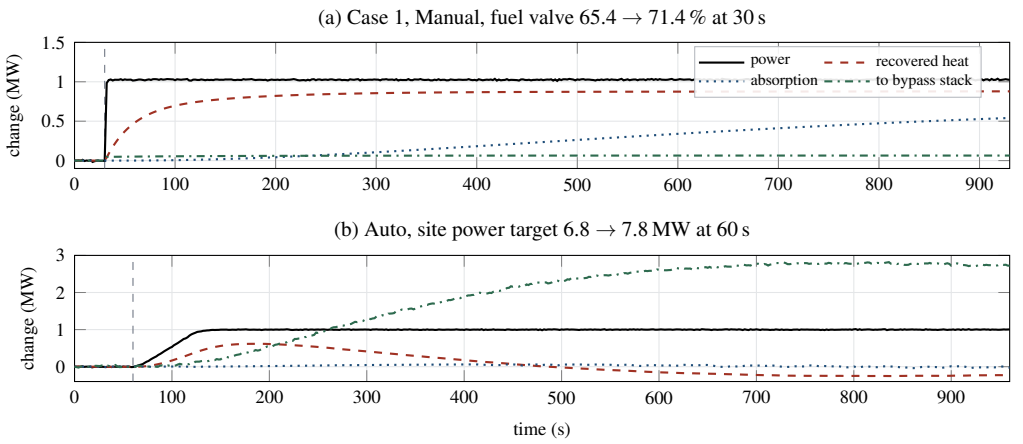

Figure 9.3 A power change in Manual and in Auto, from the exported records. Each series is its change since the first logged row: power measured (solid), recovered heat (dashed), absorption cooling (dotted) and the heat sent round the HRU to the bypass stack (dash-dotted), all true except the power. (a) Case 1, the fuel valve stepped by hand at 30 s. (b) The reference plant in Auto with its site power target raised by 1 MW at 60 s.

> **TRY IT IN THE APP** **1 · A Fuel Step: Power, Heat and Cooling on Three Clocks (Manual)**
>
> Load the scenario and press Run. After 30 s type 71.4 into Fuel valve on the turbine's card, open the Cooling response view and watch the trend Change since the start. The power covers 63 % of its rise in 0.8 s, the recovered heat in 39 s, and the absorption cooling only in 573 s.

In Auto the bypass loop answers first. With the reference plant's site power target raised from 6.8 to 7.8 MW at 60 s, the turbine gains 1.01 MW on its ramp, the bypass opens from 4.83 to 18.33 %, 2.73 MW more exhaust heat goes to the stack, the recovered heat ends 0.23 MW lower than it began, and the absorption cooling changes by –0.03 MW at 960 s, as shown in Figure 9.3b. Case 4's workload pulse confirms it from the cooling side. Pressed at 1200 s, it raises the turbine from 6.79 to 7.86 MW through the import guard and the absorption cooling from 6.38 to 6.85 MW, 63 percent of the way in 202 s, while the bypass opens from 0.25 to 9.09 % and sends 1.81 MW to the stack. With the turbine held at 6.80 MW instead (import guard off, import limit raised to 5 MW), the absorption cooling rises identically, to 6.85 MW in 203 s, with its valve wide open 369 s after the press in both runs. The rise comes from the VPC opening the valve against the electric duty above its reserve, not from the turbine's heat. In Auto the coupling that matters runs from the chillers to the bus. Case 7's absorption trip adds 0.61 MW of electric-chiller load, and in case 5 the site's room below the turbine's capacity plus the import limit falls from 3.28 to 1.85 MW, 1.25 MW of it to derating and 0.18 MW to the cooling plant.

Pause. Try Exercise 9.1 before reading on.

9.5.3 Three Matched Comparisons

Within each group of Table 9.7 the runs start identically and see the same seed, meters, limits and timed upset, changing one setting only. The trajectories are shown in Figure 9.4.

Group A tests nominal disturbance rejection. Absorption alone, at 8 %/K and 900 s, lets the water rise for twenty minutes, to 0.757 K at 1192 s, and it is still 0.57 K up at 1800 s while the valve creeps from 76 to 89 %. The electric trim holds the error within 0.012 K and cuts the IAE from 951.41 to 5.16 K s. Its command rises from 0.50 to 0.969 MW by 208 s, first on the feedforward and then on the trim, and the VPC opens the absorption valve from 76 to 85 % by 900 s, so the delivered electric duty

Table 9.7 Three Matched Comparisons on the Simulator's Truth. Grid-Connected, No Load Unserved, the UPS at 90 % at Start and End in Every Run

Setting	IAE (K s)	Max $\lvert e \rvert$ (K)	Recovery (s)	$Q_{ec,cmd}$ peak (MW)	Reserve (s)	Travel Q_{ec}, v_g (MW, %)	Sat. (s)	Store (%)	E_{cool} (MWh)
A. Rejection, 60–1800 s, IT +0.6 MW, *band* ±0.1 K									
case 3, absorption alone	951.41	0.757	not reached	0.50 held	n/a	0, 16.2	0	80 → 79.2	0.267
case 4, trim and VPC	5.16	0.012	never left	0.969	806	2.13, 54.8	0	80 → 79.2	0.276
case 4, no feedforward	45.09	0.090	never left	0.969	845	2.04, 53.7	0	80 → 79.2	0.276
B. Capacity, 0–1800 s, band ±1 K									
case 8 as shipped	5614	7.41	not reached	2.0 at cap	n/a	1.55, 77.7	1737	80 → −0.06	0.329
$K_{c,s} = 10$ MW/K	5431	7.41	not reached	2.0 at cap	n/a	1.55, 77.7	1737	80 → −0.07	0.329
curtail at 207 s	3356	3.57	not reached	2.0 at cap	n/a	1.55, 77.7	1737	80 → 2.4	0.329
C. Mismatch, 60–1800 s, case 4, band ±0.1 K									
$\tau_{abs} = 60$ s	3.80	0.008	never left	0.863	807	2.03, 54.3	0	80 → 79.2	0.276
$\tau_{abs} = 180$ s, tuned	5.16	0.012	never left	0.969	806	2.13, 54.8	0	80 → 79.2	0.276
$\tau_{abs} = 540$ s	6.83	0.013	never left	1.072	807	2.23, 55.7	0	80 → 79.2	0.277

Recovery is when the error last leaves the band inside the window. Reserve is the instant from which the delivered electric duty stays within 0.1 MW of its 0.5 MW reserve. Travel is of the delivered duty and the valve position, while the trim's command travels 77.1 MW in every group A and C run because it answers the 0.01 K RTD noise at 4 MW/K, which the 30 s loading lag smooths. E_{cool} is the cooling plant's electricity over the run. Every run burns 13.09 MWh of fuel. In B the time outside ±1 K is 1593.1, 1546.5 and 1593.1 s and 531 kWh of flexible IT is curtailed in the third run. Halving the step to 0.025 s changes the IAE values to 951.37, 5.16, 45.09, 5613.9, 5430.1, 3356.1, 3.80, 5.16 and 6.83 K s.

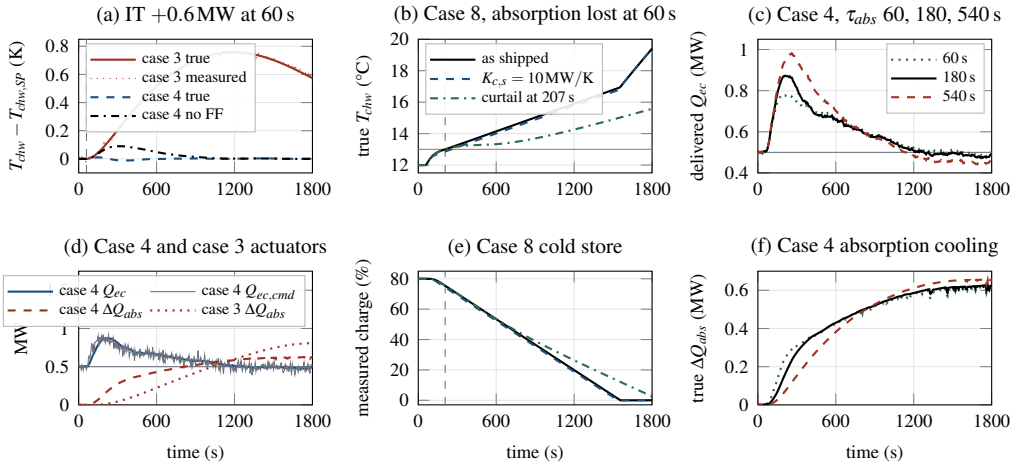

Figure 9.4 The three matched comparisons, from the exported records. (a) The chilled-water error after the 0.6 MW IT step at 60 s, case 3 true (solid) and measured (dotted), case 4 true (dashed) and case 4 without the feedforward (dash-dotted). (d) The actuators of the same runs: case 4's delivered electric duty and its command, and the rise of the absorption cooling in case 4 and in case 3, against the 0.5 MW reserve. (b) and (e) Case 8 after the absorption chiller is lost at 60 s, true chilled water against the 13 °C band edge and the measured store charge, as shipped (solid), with a store gain of 10 MW/K (dashed) and with the flexible IT curtailed at the band exit at 207 s (dash-dotted). (c) and (f) Case 4 with the absorption capacity lag at 60 (dotted), 180 (solid) and 540 s (dashed), delivered electric duty and the rise of the absorption cooling.

stays within 0.1 MW of its reserve from 806 s. Without the feedforward the IAE is 45.09 K s, nine times the value with it. The fast loop costs 0.276 MWh of cooling electricity against 0.267 MWh, while both runs end with the store at 79.21 %, the UPS at 90 % and 3.29 MWh of IT delivered. Across seeds 1 through 5 the IAE is 950.4 to 952.4 K s alone and 4.59 to 5.66 K s with the trim.

The second comparison is a capacity limit. In case 8 one electric chiller is out and the absorption

chiller is lost at 60 s for good, so 2.0 MW of electric cooling and the store's 3.0 MW rating leave the hall short by about 1.3 MW. As shipped the water leaves its band at 207 s, the store is empty at 1549 s and the water reaches 19.41 °C. A store gain of 10 MW/K only brings the store to its rating sooner, holding the band to 254 s and emptying the store at 1530 s. Curtailing the flexible IT at the band exit removes 1.2 MW of heat, so the water reaches only 15.57 °C and the store lasts to 1856.5 s in a run continued to 2400 s, at the price of 531 kWh of flexible work by 1800 s. No setting adds cooling capacity.

The third comparison is a model mismatch. The trim and the VPC were tuned for a 180 s capacity lag, and case 4 is run with the lag at 60 and 540 s. A faster chiller cuts the electric peak to 0.863 MW and the IAE to 3.80 K s, and a slower one raises them to 1.072 MW and 6.83 K s, with the error at most 0.013 K in all three. The electric command first comes back within 0.1 MW of the reserve at 429 s with 60 and 180 s and at 611 s with 540 s, while the delivered duty settles inside that band at 806 to 807 s in all three, which suggests that the VPC's own 200 s integral time, more than the chiller's lag, sets the end of the return. The valve's total travel moves only from 54.3 to 55.7 % over a ninefold range of lag, so the two actuators do not trade the load back and forth. With the hall node twice as slow as the feedforward's 90 s lag assumes, the water undershoots by 0.034 K against 0.012 K when matched.

Pause. Try Exercise 9.2 before reading on.

9.5.4 Capacity That No Tuning Adds

Case 9 is the example in which added control cannot overcome missing capacity. On a 38 °C afternoon with a 26 °C wet bulb the turbine can make 7.95 MW, the grid 1.0 MW, and the UPS from 60 % holds $0.5 \times 0.40 \times 0.95 = 0.19$ MWh above its floor. The IT load rises from 6.0 to 8.5 MW over five minutes from 60 s, the turbine reaches its exhaust limit at 262 s, the demand rises to 10.10 MW, the UPS gives up to 1.18 MW from 268 s and reaches its floor at 954.30 s, and 226.35 kWh of flexible IT goes unserved by 1800 s with none of the critical IT. A much faster power loop, 5 %/MW with 1 s set when the turbine reaches its limit, reaches the floor at the same 954.30 s with the same unserved energy, because its output is pinned. A sweep of 54 settings over power gains of 0.5, 1.5 and 5 %/MW, integral times of 1, 5 and 30 s, the import guard off or on with a margin of 0 or 0.6 MW, and the trim on or off empties the UPS between 601.4 and 1517.8 s in every run and leaves 48.9 to 680.9 kWh unserved, never critical. The trim-off runs leave less unserved only by trading cooling for IT, since as shipped with the trim off 59.72 kWh goes unserved while the chilled water peaks 0.785 K high, against 0.031 K with the trim. Curtailing the flexible IT by plan lowers the demand to 8.03 MW, the turbine stays at 7.23 MW below its limit, the UPS never discharges and nothing goes unserved, and 820.9 kWh of flexible work is curtailed instead.

Pause. Try Exercise 9.3 before reading on.

> **TRY IT IN THE APP** **9 · More Demand Than the Turbine and the Grid Can Give (Demonstration)**
>
> Run the scenario on the Margins view. The turbine reads EXHAUST LIMIT from 262 s while the UPS gives. Then set Power gain to 5 and Power integral time to 1 s and keep running: the UPS still reaches its floor at 954.30 s, the same second as with the shipped tuning, because the loop's output is pinned.

9.5.5 The Challenge Under Five Policies

The challenge runs 1800 s with the import limit at 1.0 MW, the fuel supply at 30 MW, the store at 80 % and the UPS at 90 %. The IT load rises 1.2 MW for ten minutes from 120 s, the wet bulb goes from 20 to 27 °C from 480 s and back from 1380 s, and the absorption chiller trips at 840 s and clears at 1140 s. The verdict asks for the chilled water within ± 1 K outside a 120 s allowance after each of six scripted events, no measured exhaust above 500 °C and no protective trip, no critical IT unserved, and the store at least 20 % charged at 1800 s. Table 9.8 gives five policies on identical initial states,

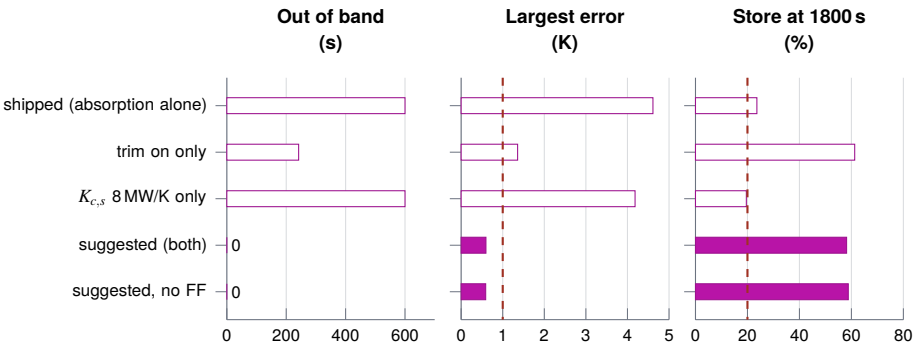

Figure 9.5 Table 9.8 redrawn. Solid bars pass the challenge's verdict and white bars fail it. Dashed lines mark the 1 K band, which the verdict applies outside a 120 s allowance after each event, and the 20 % floor for the store at the end. Notice that the trim alone shortens the time out of band but still fails, that the stronger store discharge alone ends with the store just under its floor, and that only the two together keep the water in the band and the store well above the floor.

Table 9.8 The Challenge Over 1800 s Under Five Policies. No Run Left IT Load Unserved or Tripped on Protection

Policy	Out (s)	Max $\lvert e \rvert$ (K)	Store (%)	UPS (%)	Import (MWh)	E_{cool} (MWh)	Verdict
shipped (alone, $K_{c,s}$ 2)	600.0	4.613	23.63	92.01	0.178	0.264	fails
trim on only	241.8	1.358	61.27	91.81	0.340	0.444	fails
$K_{c,s}$ 8 MW/K only	600.0	4.183	19.57	92.01	0.178	0.263	fails
suggested (both)	0	0.596	58.15	91.81	0.330	0.433	passes
suggested, no FF	0	0.590	58.80	90.71	0.328	0.438	passes

Out is the time outside ±1 K outside the allowances, store and UPS are final charges from 80 and 90 %. Every run delivers 3.200 08 MWh of IT and burns 13.314 to 13.355 MWh of fuel, and the measured exhaust peaks at 444.9 to 448.2 °C.

measurements, events and limits, and Figure 9.6 compares the shipped and suggested runs. Figure 9.5 draws the table's verdict columns as bars.

As shipped, the absorption chiller holds the water alone and the store's gain is 2 MW/K. At the trip nothing fast is left, the electric chillers hold their 0.5 MW, the store gives its 3.0 MW rating, and the water rises 4.61 K, out of band for 600 s, and the store ends at 23.63 %. The trim alone fixes the burst and the hot sink, but at the trip the weak store gain lets the water drift 1.36 K and out of band for 241.8 s. The store's gain alone cannot help without the trim and empties the store to 19.57 %. Together they pass. The electric chillers reach 4.0 MW, the store gives up to 2.56 MW, the water stays within 0.596 K, and the store ends at 58.15 %, the same verdict as without the feedforward, 0.590 K and 58.80 %, because the burst is slow enough for feedback. The suggested run imports 0.330 against 0.178 MWh and its cooling plant draws 0.433 against 0.264 MWh, but it ends with 34.5 points more stored cooling, so the comparison rewards service and the store and is not a savings estimate. At the burst the import reaches its limit and the UPS gives up to 0.32 MW, 1.45 to 1.50 kWh in every run, while the guard ramps the turbine. Across seeds 1 through 5 the shipped run is out of band for 600 s and the suggested run passes with 0.595 to 0.598 K and 58.14 to 58.17 %. Its state at 1000 s is shown in Figure 9.7.

9.5.6 Verification, Reference Data and Cost

These numbers check the implementation and say nothing about control performance. The gas-path first law closes to 5.6×10^{-16} of the fuel across fuel flows of 0.2 to 0.8 kg/s and ambient temperatures of 0 to 45 °C. Integrated from boundary fluxes with formulas written apart from the application, by the trapezoid rule over every step of nine disturbed runs, the gas path closes to

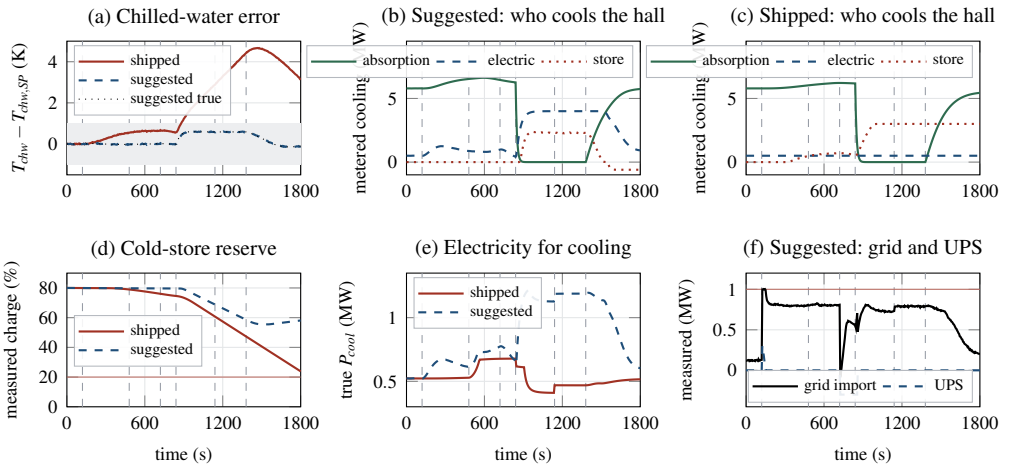

Figure 9.6 The challenge as shipped and with the suggested tuning, from the exported records. (a) Measured chilled-water error, shipped (solid) and suggested (dashed), and the suggested run's true error (dotted), with the ±1 K band shaded. (b) and (c) Metered cooling from the absorption chiller, the electric chillers and the store, suggested and shipped. (d) Measured store charge against the 20 % reserve. (e) True electricity drawn by the cooling plant. (f) Measured grid import against its 1.0 MW limit and the UPS power, suggested run. Gray dashed lines mark the events at 120, 480, 720, 840, 1140 and 1380 s.

6.1×10^{-14} of its fuel, the hot water to 2.8×10^{-5}, the chilled water to 1.5×10^{-5}, the IT node to 6.5×10^{-6} and the cooling water to 7.1×10^{-5} of their throughput, and the bus to 2×10^{-9} W at every step, where no step imports while exporting or charges the UPS while discharging it. With the turbine tripped and the UPS at its 2.0 MW rating from 90 %, the floor arrives 598.50 s later, the analytic $(0.90 - 0.20) \times 0.5\,\text{MWh} \times 0.95/2\,\text{MW}$. With the absorption valve shut by hand the cooling is 0.0001 MW at 410 s while the lag state still holds 3.85 MW, and a fuel valve shut at 60 s motors the turbine to −3.41 MW until the reverse-power relay trips it at 62.6 s. Halving the step to 0.025 s leaves the challenge verdicts unchanged, changes no other entry of Table 9.7 by more than one unit of its last digit, moves the IAE by at most 0.46 K s in 5614 K s, and moves case 9's floor from 954.30 to 954.27 s. All 71 of the application's checks pass. Headless, a simulated hour takes 0.46 s of computation, and at the browser's fastest setting, 100 times real time, the 1800 s challenge finishes in 18.0 s.

Comparison with published data is a different and more limited claim. At ISO conditions and the exhaust limit the model makes 10.28 MW net at 30.21 % LHV on 34.01 MW of fuel, with 42.08 kg/s of exhaust at 490.0 °C and a pressure ratio of 17.72, against the catalog's System 3 at 9.95 MW, 27.34 percent on the higher heating value or 30.38 percent on LHV with its ratio of 0.9, 32.76 MW of LHV fuel, 334 000 lb/h (42.08 kg/s) at 913 °F (489.4 °C) and a pressure ratio of 17.7 [1]. The model keeps 77.6 % of its ISO capacity at 100 °F with 9.1 % less efficiency, against the catalog's 80 % and 8 % for a different, nominal 7.5 MW machine. The constant specific heats carry an error of their own. With ideal-gas heat capacities from CoolProp 8.0.0 [20] and the same η_c and η_t, the ISO point gives 10.58 MW at 32.2 %. With η_t refitted to 0.854 for the model's 30.21 %, the model at 6.80 MW and 25 °C burns 3.4 % more fuel, runs its exhaust 13.4 K cooler with 2.9 % more heat above 120 °C (14.26 against 13.86 MW), credits itself with 9.21 MW of capacity at 25 °C against 8.60 MW, and derates to 77.6 against 77.2 % at 100 °F. The absorption COP passes through the DOE points at 0.6999 and 0.7399 [2]. No part of the model has been tested against field measurements, and no learner study has been made.

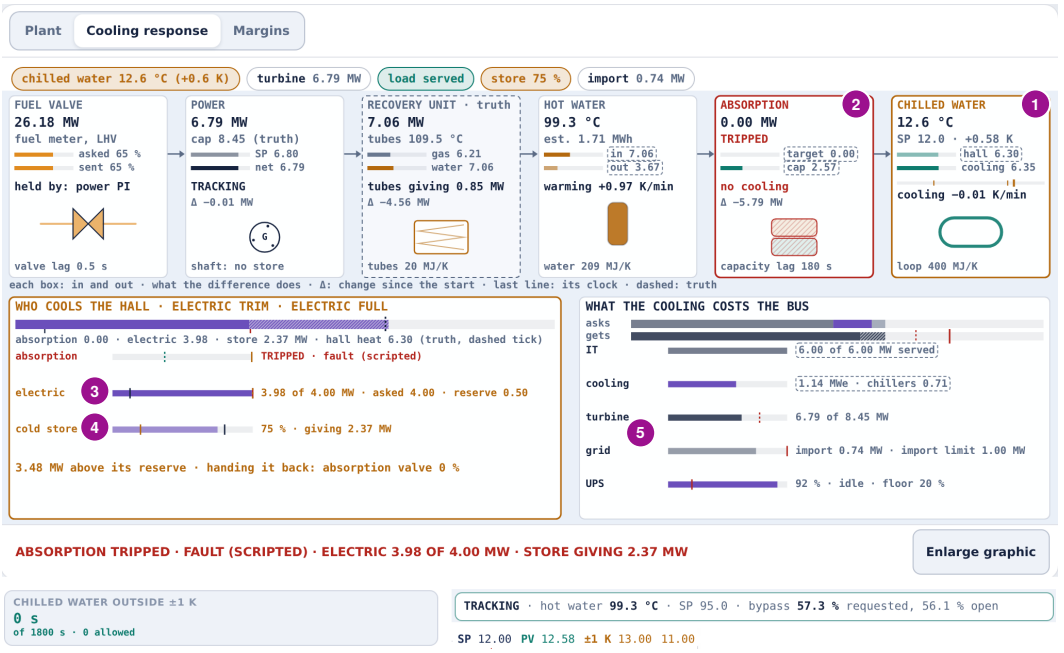

Figure 9.7 The challenge with the suggested tuning at $t = 1000\,\text{s}$ on the Cooling response view, in a composite of the application window: the view itself, with the lead scorecard tile, the hot-water loop's tracking line and the chilled-water trend label set below it. The absorption chiller has been tripped since 840 s (2), so the electric chillers deliver 3.98 MW of their 4.00 MW (3) and the store gives 2.37 MW on its meter at 75 % charge (4), while the chilled water reads 12.58 °C, 0.58 K above its setpoint and inside its band (1). The turbine makes 6.79 MW of its 8.45 MW capacity, the grid imports 0.74 MW against its 1.00 MW limit (5), the cooling plant draws 1.14 MW, and with no drive heat taken the hot water stands at 99.3 °C with the bypass 56 % open. The lead scorecard tile reads 0 s outside the band.

9.6 Technology in Practice

A real trigeneration site reads and moves the same quantities as the loops of Section 9.4. Figure 9.8 draws the fuel selector and the chilled-water loops with plant hardware in place of the schematic's symbols.

What Measures It

The chilled-water loop reads one supply temperature, and its bias reads two heat flows, which a real plant measures with thermal energy meters. The federal metering guide notes that clamp-on ultrasonic flow meters are attractive because they are completely non-intrusive, but their coupling gel must be rated for the pipe's temperature, or it may lose viscosity and melt away [21]. On the gas side, the exhaust of smaller industrial turbines reaches 800 to 900 °F [1], and the limiter reads averaged thermocouples, as in Chapter 3. The model's chilled-water RTD lags 5 s with 0.01 K of noise and its averaged thermocouples 4 s.

TECHNOLOGY Sensor: A Thermal Energy Meter

A thermal energy meter on a hot-water or chilled-water circuit finds the energy flow as the mass flow times the temperature difference across the load. A flow meter gives the flow, two temperature sensors give the supply and the return, and a flow computer integrates the three readings over each interval. Most fluid meters measure volume, so the computer also accounts for density and specific heat, which change with temperature. Orifice, turbine, vortex and ultrasonic meters serve these circuits, with turndown from under 5:1 for an orifice to 20:1 for a vortex meter, and a meter is sized for the expected flow range, never for the pipe. It is calibrated yearly, or sooner if trended data show miscalibration [21]. The model's absorption and store meters lag 10 s and form the trim's bias B in Equation (9.9).

What Moves It

The fuel valve meters gas that must already be compressed above the combustion air, because the fuel joins the air after the compressor. The pressure it needs is set by the turbine's pressure ratio, and where the site's supply pressure is low a fuel gas compressor adds cost and power consumption [1]. The model's valve answers with a 0.5 s lag behind a 40 MW supply limit. On the cooling side the valve-position controller moves the absorption chiller's hot-water valve, and the trim sends a duty command to the capacity control of the electric chillers, which load with a 30 s lag.

TECHNOLOGY Final Element: The Absorption Chiller's Thermal Compressor

An absorption chiller replaces the motor and compressor of an electric chiller with a thermal compressor made of an absorber, a solution pump and a generator. In the absorber a solution takes up the refrigerant vapor coming from the evaporator, a small pump lifts the dilute solution to the generator, and there heat boils the refrigerant out toward the condenser while the absorbent returns to the absorber. For chilled water at 40 °F or above the usual pair is water as refrigerant and lithium bromide as absorbent. Service includes periodic purging of non-condensable gases and checks of the tower and chilled-water quality [2]. The fact sheet gives no response time. In the model the valve v_g sets the drive heat, and a 180 s capacity lag stands for the machine's slow loading.

Where the Controller Runs

Supervisory and lower-level loops run with cycle times from milliseconds to minutes, and programmable logic controllers serve as local controllers under the supervision of a distributed control system [22]. The application's scans sit inside that range, 0.25 s for the power loop and the exhaust limiter, 1 s for the water and store loops and 5 s for the valve-position controller. In a function-block system the low selector of Equation (9.8) is a control selector block, whose back-calculation output tells each upstream controller whether it was selected so that a deselected one halts its integral action [23]. Here the limiter's output, set as the power loop's ceiling, does that work.

IN PRACTICE Electric Backup Behind Recovered-Heat Cooling

A 2009 report for a national laboratory describes how data centers with CHP are built. Cooling from absorption chillers must be backed up by electric chillers that can assume the load whenever the CHP plant is not running, and the favored arrangement is redundant parallel paths with redundant control systems, chilled-water pumps and electrical boards. In the most reliable tiers of data center, thermal storage may carry the cooling through a transfer. With proper design and controls an absorption chiller is as reliable as an electric chiller at comparable maintenance cost, and it can run with tower water down to 60 °F [3]. This chapter's trim and cold store are that backup in miniature.

What Is Modernizing It

The first development is inlet air cooling. Cooling the air entering a turbine by 40 to 50 °F on a hot day can raise its power by 15 to 20 %, by refrigeration, evaporative cooling or thermal storage, and

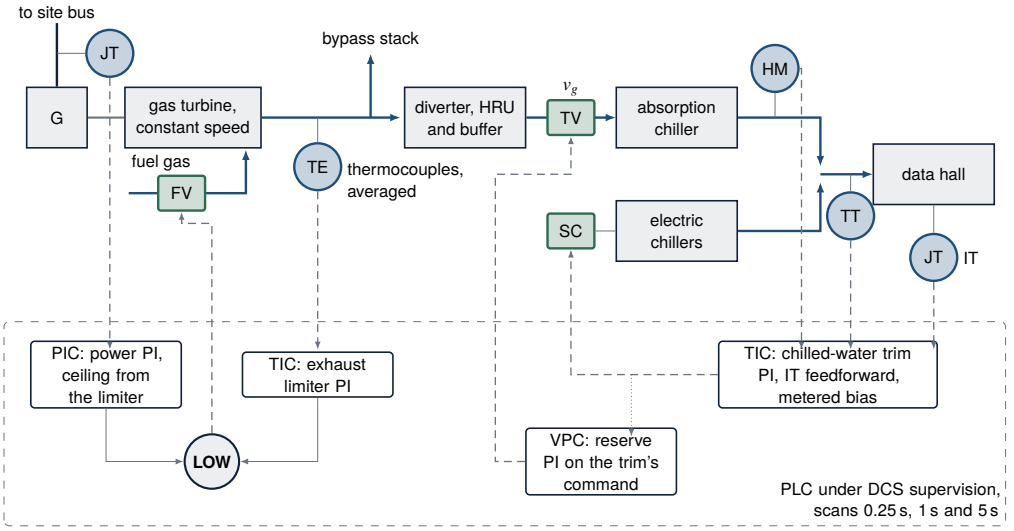

Figure 9.8 The loops of Figure 9.1 on plant hardware. The power transducer (JT) and the averaged exhaust thermocouples (TE) feed the power PI (PIC) and the exhaust limiter (TIC), whose lower output goes to the fuel valve (FV). The chilled-water RTD (TT), the IT meter (JT) and the heat meter on the absorption chiller (HM: a flow meter, two temperature sensors and a flow computer) feed the trim, which sends a duty command to the electric chillers' capacity control (SC). The valve-position controller (VPC) reads that command, dotted, and moves the absorption chiller's hot-water valve (TV). Notice that the turbine and the chillers meet only at the bus and the hot water: no signal runs from the power loops to the cooling loops. The bypass loop, the store and its meter are not drawn.

an absorption chiller on the turbine's own waste heat cuts the parasitic loss of an electric one at a cost in complexity [1]. The second is chillers in series: one developer was designing data center plants with the absorption chiller in series with a larger electric chiller with a screw compressor, which then sees cooler entering water and runs efficiently at part load, and which takes the full load when the CHP plant is down [3]. The third is supervisory dispatch. Heuristic, linear-optimization and model predictive dispatch of generators, chillers and storage proved feasible in real time on a campus plant and could outperform manual operation [24].

9.7 Limitations

Every size, map and threshold here is a teaching value chosen against the EPA and DOE documents [1, 2] and fitted to neither a turbine nor a chiller beyond $\dot{m}_{a,d}$, η_c, η_t, the swallowing reference and the exhaust limit. The gas path's constant specific heats overstate the turbine's capacity at 25 °C by 0.61 MW against a property-correct model refitted to the same ISO efficiency, and the absorption chiller's capacity at the model's hot-water flow is a teaching map. In Manual a fuel-valve step to 100 % lets the true exhaust reach 546.8 °C and stay above the 520 °C trip setting for 1.75 s, while the averaged thermocouples that the limiter, the alarm and the trip read peak at 490.8 °C, which is how protection on an averaged measurement behaves. The lumped nodes give lags of realistic size but no spatial temperature profile, so no thermocline claim is made. The added MPC uses a nominal valve-fraction approximation and held current conditions. It omits startup history and finite future store delivery, and its bias correction and soft bands do not ensure temperature or reserve limits. The trips are a teaching safeguard and not certified protection logic, the load-shed order is fixed, passing a case qualifies no equipment and estimates neither savings nor reliability, and no gain in learning has been measured.

OPEN CHALLENGE Controls for Heat and Power That Do Not Peak Together

A 2013 report to the U.S. Department of Energy named a lack of controls able to deal with highly dynamic thermal and electrical loads that do not peak together as one of three high-priority barriers to combined cooling, heating and power at smaller sites. On a campus plant whose thermal store could shift a full day of cooling, real-time dispatch saved 8 to 25 %, but the savings depended on pricing, season, forecast accuracy and grid constraints, and the store added complexity to the optimization [24]. Progress would be dispatch that keeps its benefit when the forecasts are wrong.

9.8 Conclusion

This chapter studied a gas turbine trigeneration plant for a data center whose single fuel valve makes power in 0.8 s, recovered heat in 39 s and, with every valve held, absorption cooling in 573 s. In Auto the hot-water bypass sends a power change to the stack, so absorption cooling follows the valve-position controller rather than the turbine, and a 1 MW power step moves it by −0.03 MW. On the same 0.6 MW workload step, a fast electric trim with a slow valve-position return cuts the chilled-water IAE from 951.41 to 5.16 K s and hands the duty back to its reserve by 806 s, and a capacity lag mismatched by a factor of three either way moves the electric peak between 0.863 and 1.072 MW without breaking the return. When the absorption chiller and one electric chiller are lost, or when demand exceeds the turbine and the grid, only less load helps, and under the original controllers the challenge passes only with both the trim and a stronger store discharge. Its nodes are lumped and its operating points invented. A fuller model would add a two-stage chiller on a steam circuit, a stratified buffer and store, variable gas properties, an inlet guide vane schedule, and a supervisory layer that schedules the store against a forecast. The simulator shows that cooling a data center from turbine exhaust is a two-actuator control problem before it is a heat-recovery one. The predictive allocation adds a measured tradeoff between driving heat and electricity, with benefits that depend on the current heat customer rather than the horizon alone.

KEY POINTS

- One fuel valve makes electricity within a second, but the cooling it pays for comes only after minutes, so a fast electric chiller and a slow absorption chiller share the chilled water.
- In Auto the hot-water bypass sends a power change to the stack, so a 1 MW power step moves the absorption cooling by −0.03 MW.
- A fast electric trim with a slow valve-position return cuts the chilled-water IAE from 951.41 to 5.16 K s and hands the duty back to its reserve by 806 s.
- When demand exceeds the turbine and the grid, only less load helps.

Where else this applies. Chapter 3 holds an exhaust-temperature limit with the same low selector on a jet engine's fuel valve, and Chapter 10 stacks three valve-position layers on one power block.

References

[1] U.S. Environmental Protection Agency, Combined Heat and Power Partnership, "Catalog of CHP technologies," U.S. Environmental Protection Agency, Washington, DC, Full Report, Sep. 2017. [Online]. Available: https://www.epa.gov/sites/default/files/2015-07/documents/catalog_of_chp_technologies.pdf

[2] U.S. Department of Energy, "Absorption chillers for CHP systems," U.S. Department of Energy, Advanced Manufacturing Office, Washington, DC, Combined Heat and Power Technology Fact Sheet Series DOE/EE-1608, May 2017. [Online]. Available: https://www.energy.gov/sites/default/files/2017/06/f35/CHP-Absorption%20Chiller-compliant.pdf

[3] K. Darrow and B. Hedman, "Opportunities for combined heat and power in data centers," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2009/052, Mar. 2009, prepared by ICF International.

[4] B. J. Allison and A. J. Isaksson, "Design and performance of mid-ranging controllers," *Journal of Process Control*, vol. 8, no. 5–6, pp. 469–474, 1998.

[5] B. J. Allison and S. Ogawa, "Design and tuning of valve position controllers with industrial applications," *Transactions of the Institute of Measurement and Control*, vol. 25, no. 1, pp. 3–16, 2003.

[6] S. Haugwitz, M. Karlsson, S. Velut, and P. Hagander, "Anti-windup in mid-ranging control," in *Proceedings of the 44th IEEE Conference on Decision and Control*. IEEE, 2005, pp. 7570–7575.

[7] G. Chicco and P. Mancarella, "Distributed multi-generation: A comprehensive view," *Renewable and Sustainable Energy Reviews*, vol. 13, no. 3, pp. 535–551, 2009.

[8] N. S. Gates, D. C. Hill, B. W. Billings, K. M. Powell, and J. D. Hedengren, "Benchmarks for grid energy management with Python Gekko," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 4868–4874.

[9] W. I. Rowen, "Simplified mathematical representations of heavy-duty gas turbines," *Journal of Engineering for Power*, vol. 105, no. 4, pp. 865–869, 1983.

[10] Working Group on Prime Mover and Energy Supply Models for System Dynamic Performance Studies, "Dynamic models for combined cycle plants in power system studies," *IEEE Transactions on Power Systems*, vol. 9, no. 3, pp. 1698–1708, 1994.

[11] M. Moya, J. C. Bruno, P. Eguia, E. Torres, I. Zamora, and A. Coronas, "Performance analysis of a trigeneration system based on a micro gas turbine and an air-cooled, indirect fired, ammonia–water absorption chiller," *Applied Energy*, vol. 88, no. 12, pp. 4424–4440, 2011.

[12] K. Ebrahimi, G. F. Jones, and A. S. Fleischer, "Thermo-economic analysis of steady state waste heat recovery in data centers using absorption refrigeration," *Applied Energy*, vol. 139, pp. 384–397, 2015.

[13] J. H. Horlock, "The gas turbine as a cogeneration (combined heat and power) plant," in *Advanced Gas Turbine Cycles*. Pergamon, 2003, pp. 167–181.

[14] H. I. H. Saravanamuttoo, G. F. C. Rogers, and H. Cohen, *Gas Turbine Theory*, 5th ed. Prentice Hall, 2001.

[15] P. Srikhirin, S. Aphornratana, and S. Chungpaibulpatana, "A review of absorption refrigeration technologies," *Renewable and Sustainable Energy Reviews*, vol. 5, no. 4, pp. 343–372, 2001.

[16] K. E. Herold, R. Radermacher, and S. A. Klein, *Absorption Chillers and Heat Pumps*, 2nd ed. CRC Press, 2016.

[17] 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.

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

[20] I. H. Bell, J. Wronski, S. Quoilin, and V. Lemort, "Pure and pseudo-pure fluid thermophysical property evaluation and the open-source thermophysical property library CoolProp," *Industrial & Engineering Chemistry Research*, vol. 53, no. 6, pp. 2498–2508, 2014.

[21] S. A. Parker, W. D. Hunt, K. McMordie Stoughton, B. K. Boyd, K. M. Fowler, T. M. Koehler, W. F. Sandusky, G. P. Sullivan, and R. Pugh, "Metering best practices: A guide to achieving utility resource efficiency, release 3.0," Pacific Northwest National Laboratory, Richland, WA, Tech. Rep. PNNL-23892-Rel.3.0, Mar. 2015.

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

[23] 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] S. Samuelsen and J. Brouwer, "Novel controls for time-dependent economic dispatch of combined cooling heating and power (CCHP)," University of California, Irvine, Irvine, CA, Final Report DOE-UCI-01113, 2013.

9.9 Exercises

9.1 ○ Concept check · 2 min

In Auto, with the reference plant's site power target raised from 6.8 to 7.8 MW, the turbine gains 1.01 MW, yet the absorption cooling changes by −0.03 MW fifteen minutes later. In Manual, with every valve held, a fuel step of the same size raises the absorption cooling by 0.54 MW. What makes the difference?

(a) In Auto the absorption chiller is already at the limit of its capacity map
(b) In Auto the hot-water loop opens the exhaust bypass and sends the extra heat to the stack
(c) In Auto the import guard holds the turbine below its new setpoint
(d) In Auto the absorption valve is interlocked shut while the turbine ramps

9.2 Concept check · 2 min

Case 4 is run with the absorption chiller's capacity lag at 60, 180 and 540 s. The electric peak moves from 0.863 to 0.969 to 1.072 MW, but the delivered electric duty settles within 0.1 MW of its reserve at 806 to 807 s in all three runs. What does this suggest sets the end of the return to the reserve?

(a) The electric chillers' 30 s loading lag
(b) The absorption chiller's capacity lag
(c) The valve-position controller's own 200 s integral time
(d) The feedforward's 90 s lag on the IT meter

9.3 Concept check · 2 min

In case 9 the turbine reaches its exhaust limit at 262 s and the UPS reaches its floor at 954.30 s. A much faster power loop, set when the limit is reached, reaches the floor at the same 954.30 s with the same unserved energy. Why does its tuning not matter?

(a) The import guard switches the power loop off once the UPS discharges
(b) A faster loop winds up its integral and comes back later, which cancels its speed
(c) The UPS always discharges at its 2.0 MW rating, whatever the demand
(d) The exhaust limiter holds the fuel valve through the low selector, so the power loop's output is not the one applied

9.4 Theory and calculation · 6 min

The absorption chiller at the reference start, from Equation (9.4). The hot water is at 95 °C, the cooling water at 29.29 °C, the valve passes 75.9 kg/s, and water's specific heat is 4.186 kJ/(kg K).

(a) Find the COP and the capacity map Q_{map}. *Checkpoint:* 0.729, 6.90 MW

(b) Find the drive capacity Q_{dc}, the cooling the chiller can make, and which limit binds. Then find the heat it rejects to the cooling water, with the 60 kW of auxiliaries. *Checkpoint:* 7.94 MW, 5.79 MW, the drive heat binds, 13.79 MW

(c) The cooling water rises to 37 °C. Find the absorption map at 95 °C, the electric chillers' COP and their capacity from Equation (9.6). *Checkpoint:* 5.33 MW, 4.80, 3.60 MW

9.5 Theory and calculation · 7 min

The handover of case 4. At the reference start the absorption chiller is limited by its drive heat (Exercise 9.4), so $Q_{abs} = \mathrm{COP}\, v_g\, \dot{m}_{g,max} c_p (T_h - T_{g,min})$ with COP 0.729, $\dot{m}_{g,max} = 100\,\mathrm{kg/s}$, $c_p = 4.186\,\mathrm{kJ/(kg\,K)}$ and T_h held at 95 °C. The VPC has $K_c = 10\,\%$ per MW and $\tau_I = 200\,\mathrm{s}$.

(a) Find the valve's cooling gain k in MW per percent of stroke. *Checkpoint:* 0.0763 MW per percent

(b) The workload step raises the hall's load by 0.6 MW. How far must the VPC open the valve from 75.9 % to carry it, if the hot water stays at 95 °C? *Checkpoint:* 7.9 points, to 83.8 %

(c) Following Equation (1.9), show that the energy the electric chillers spend above their reserve during the handover is $E_h = \Delta v / K_I$ with $K_I = K_c/\tau_I$, whatever the lags, and evaluate it in MJ and kWh. *Checkpoint:* about 157 MJ, 44 kWh

9.6 Theory and calculation · 6 min

Two stocks that deplete, with the runway of Equation (1.11).

(a) The turbine trips with the UPS at 90 % of its 0.5 MWh, and the UPS discharges at its 2.0 MW rating with a discharge efficiency of 95 % to its 20 % floor. How long does it last? *Checkpoint:* 598.5 s

(b) In case 8 the absorption chiller is lost at 60 s and one electric chiller is out, so the chillers give 2.0 MW against the hall's 6.29 MW. The cold store starts at 80 % of 1.5 MWh and discharges at its 3.0 MW rating. Estimate when it is empty and compare with the chapter. *Checkpoint:* about 1500 s against 1549 s

(c) The store gives $K_{c,s}(e - 0.3\,\mathrm{K})$, clipped at its rating, for a chilled-water error e. At what error does it reach its rating with the shipped 5 MW/K, and about how long does the remaining gap then take to warm the 400 MJ/K chilled-water loop to the 1 K band edge? *Checkpoint:* 0.9 K, about 30 s

9.7 App challenge · 8 min

Scenario 2, *Tune the power loop and meet the exhaust-temperature limit*, default seed. Run each case to 900 s at 20× (about 45 s each) and export the CSV file after each run.

(a) As shipped, read `P_meas_MW` at 420 and 600 s and `P_sp_MW` at 600 s. Does the turbine ever read EXHAUST LIMIT? *Checkpoint:* 8.43 MW, 8.15 MW against 6.8 MW, no

(b) Reload the case, set *Power integral time* to 2 s before Run, and run again. Find the first row with `exh_limiting` at 1, the power at 300 s and the largest `T4_meas_C`. *Checkpoint:* 204 s, 8.71 MW, 491.6 °C

(c) In the same file, find the last row after 420 s in which `P_meas_MW` differs from `P_sp_MW` by 0.05 MW or more. Why does the power come back so soon after the target falls? *Checkpoint:* about 651 s

Tolerance: 0.02 MW, 0.3 K and 3 s.

9.8 App challenge · 8 min

Scenarios 3, *Absorption cooling alone: a slow loop on a slow chiller*, and 4, *Electric trim and the slow return of its reserve*, default seed, as shipped. Run each to 1800 s at 50× (36 s each), read the chilled-water IAE on the scorecard, and export the CSV file. The chilled-water error is `truth_Tchw_C` less 12.

(a) Scenario 3: find the largest error and its time, the error at 1800 s and `abs_valve_pct` at 1800 s. *Checkpoint:* 0.757 K at 1192 s, 0.57 K, 89.0 %

(b) Scenario 4: find the largest error, the peak of `ec_cmd_MW`, and the last row in which `truth_Qec_MW` is 0.1 MW or more from 0.5 MW. *Checkpoint:* 0.012 K, 0.969 MW, 806 s

(c) The scorecard shows a chilled-water IAE of about 949 and 15.9 K s, while the chapter gives 951.41 and 5.16 K s on the simulator's truth. Explain the larger value for scenario 4.

Tolerance: 0.005 K, 0.01 MW, 3 s, 0.3 points and 1 K s.

9.9 App challenge · 7 min

The challenge, scenario 11 and listed last in the scenario menu, *Challenge · a workload burst, a hot sink and an absorption loss*, default seed. Run each policy for the full 1800 s at 100× (18 s each), then change the settings, press Reset and run again. Leave every parameter and disturbance as shipped, and read the *Challenge target* tile and the status line.

(a) Run as shipped. Read the time outside the band and the cold store at the end. *Checkpoint:* 600 s, 23.6 %

(b) Run with only *Electric backup trim* On, then with only *Store discharge gain* at 8.0 MW per K. *Checkpoint:* 242 s and 61.3 %, then 600 s and 19.6 %

(c) Run with both changes. Read the same two values and the *Grid import* tile, and compare it with the shipped run. *Checkpoint:* 0 s, 58.1 %, 0.33 MWh against 0.18 MWh

Tolerance: 2 s, 0.2 points and 0.01 MWh.

9.10 Application area · 4 min

In Manual, a fuel-valve step to 100 % lets the true exhaust reach 546.8 °C and stay above the 520 °C trip setting for 1.75 s, while the averaged thermocouples, which lag the gas by a first-order 4 s, peak at 490.8 °C. No trip occurs.

(a) What fraction of a sudden step does a 4 s first-order lag show after 1.75 s? *Checkpoint:* about 35 %

(b) What should a practicing engineer conclude about the trip, and what keeps the turbine away from such a step in Auto?

9.11 Application area · 4 min

The DOE fact sheet pairs single-stage absorption chillers on hot water at 200 to 240 °F with engines, and two-stage machines on hot water near 350 °F or steam near 115 psig, with COPs of 1.35 to 1.42, with combustion turbines. The chapter's plant uses a single-stage machine at a COP of 0.729.

(a) How much drive heat would a two-stage machine at a COP of 1.4 need for the same 5.79 MW of cooling, and how much heat would it reject to the tower per MW of cooling? *Checkpoint:* 4.14 MW, 1.71 against 2.37

(b) Why does the chapter keep the single-stage machine, and what would the plant need to run a two-stage one?

9.12 Review the agent · 5 min

An AI assistant reviews case 8 and writes: "The chilled water leaves its band at 207 s because the cold store answers too weakly and the trim is too slow. Raise *Store discharge gain* from 5 to 10 MW per K and *Trim gain* from 4 to 8 MW per K, and the water will stay in its band." Find the flaw and say how you would check it in the application.

CHAPTER 10

A Supercritical CO_2 Brayton Simulator for Three-Layer Mid-Ranging Control

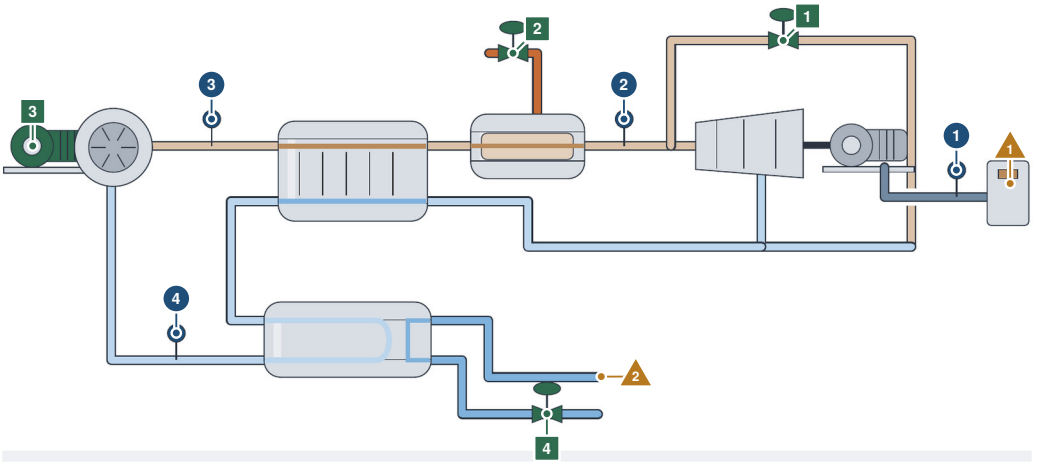

The controller moves 1 bypass valve around the turbine 2 fuel to the heater 3 compressor speed 4 cooling water valve

The controller reads 1 net electric power 2 turbine inlet temperature 3 compressor outlet pressure 4 compressor inlet temperature

What disturbs it 1 power the grid asks for 2 warm cooling water

Notice that three handles, the bypass, the heater and the compressor speed, all move one power output, while the cooler on the cold side sets how much work the compressor needs. Illustration, not to scale.

IN PLAIN TERMS **Three Handles on One Power Plant**

The main idea. This power plant pushes carbon dioxide around a closed loop. The gas is squeezed until it is almost a liquid, so the compressor needs little work. Three handles change the power: a fast valve that wastes energy, a compressor speed that costs nothing, and a slow fire that burns fuel.

The challenge. The controller uses the fast handle first and then passes the work down to the slower, cheaper ones. When the cooling water warms on a hot afternoon, the fluid swells and the power falls, and no controller can win it back.

Why it matters. Done well, the plant follows the grid quickly while it burns less fuel.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on density (Section 10.1), the two history segments (Section 10.2), the try-it segments and the fact on the dispatch cut (Section 10.6) and the key points
Study	55 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 10.6), Technology in Practice (Section 10.7), Limitations (Section 10.8) and the open challenge

The control question A supercritical carbon dioxide power block has one variable the grid asks for, net power, and three actuators that order one way by speed and another by cost. Which structure delivers a dispatch that the obvious architectures cannot reach at all, and what does a third mid-ranging layer buy that no tuning can?

Learning objectives

- Rank the turbine bypass, the compressor speed and the firing rate by response time and by what each costs, from their measured gains.
- Compute compressor work and mass flow near the critical point and explain why the compressor inlet loop is worth more than the hierarchy.
- Explain why a valve position controller reads the unsaturated demand of the loop below it, and size the error it then sees.
- Configure three-layer mid-ranging and judge the resting bypass position as a purchase of responsiveness with fuel.
- Attribute the gain of a model predictive controller to tuning, to its third handle and to its model of the inner loops.

Prerequisites Sections 1.4.7 (mid-ranging), 1.4.5 (saturation), 1.4.4 (external reset), 1.6.2 (velocity form) and 1.6.4 (scenarios). Section 1.6.3 is needed only for the MPC.

The app https://apmonitor.com/apps/sco2-brayton/

See also Volume 3, Chapter 2 for structure chosen before tuning, Volume 3, Chapter 3 for the inner-loop speed a cascade needs and Volume 3, Chapter 5 for a megawatt power cycle on one fluid inventory. Volume 17, Chapter 4 for one variable handed between two actuators, Volume 17, Chapter 2 for an exhaust bypass on two clocks and Volume 17, Chapter 6 for a heat source with a lag of its own.

10.1 Introduction

A closed supercritical carbon dioxide Brayton cycle is attractive because the fluid is compressed close enough to its critical point to behave almost like a liquid, so the compressor absorbs a small fraction of the turbine work and the cycle reaches steam-plant efficiency in a machine an order of magnitude smaller [1, 2]. That same proximity to the critical point is the source of every control problem the application poses.

> **DID YOU KNOW? Three Degrees That Halve the Density**
>
> Just above the critical point, warming carbon dioxide from 32 to 35 °C at the same pressure halves its density. In a model of a recompression cycle for a sodium-cooled reactor, raising the pre-cooler outlet from 32.0 to 33.2 °C lowered the thermal efficiency from 39.32 to 38.68 % [3].

The grid asks for net electrical power, and three actuators move it. The turbine bypass diverts flow around the turbine, so the power falls in about eight seconds and every watt of the reduction is discarded as heat into the recuperator. The compressor speed changes the mass flow and the discharge pressure together, so it raises the cycle efficiency as it raises the power, and it costs nothing. The firing rate is the actuator that adds heat to the cycle, and it reaches sixty-three percent of a turbine inlet temperature change in 16.6 min. The fastest actuator is therefore the most expensive one and the slowest is the one that buys the fuel, which is the ordering that makes mid-ranging the natural structure [4]. Section 10.6 measures what each single-actuator structure fails to do, and the structure that works applies mid-ranging three times in series. A model predictive controller (MPC) is offered as well, moving the bypass and the two inner-loop setpoints as one plan, a method with a long industrial record [5] that has also been analyzed for a supercritical carbon dioxide power converter [3]. Section 10.6.8 measures it against the hierarchy as shipped and as retuned.

The plant brings a specific list of process dynamics and control principles together. Time-scale separation across two orders of magnitude, cascade structure and the condition a secondary loop

must satisfy, valve position control and the choice of a resting position, actuator saturation with anti-windup, the handover of a saturated signal between controllers, constraint reordering rather than constraint removal, and the difference between a controlled variable that is held and a plant that is healthy. Each is standard material in a first course [6–8], and this plant supplies a setting in which they arrive together. The physical fundamentals underneath the model are equally standard, because the cycle is an energy balance around six stations, a compressor and a turbine characteristic, three exchangers with thermal mass and one strongly nonlinear property variation near the critical point. Industrially the behavior matters because a power block that follows dispatch by throwing work away earns less than one that follows it by moving pressure, and because the loop with the largest return on this plant is the one that is usually specified by a mechanical engineer and never tuned again.

This chapter works with a supercritical carbon dioxide power block simulator in which one controlled variable is reached by three actuators of very different speed and cost. It quantifies what each additional layer of the mid-ranging hierarchy buys and what it cannot buy. Section 10.2 places the model against the cycle and control literature, and Section 10.3 gives the equations and constants needed to rebuild the simulator. Section 10.4 measures the three actuators and the compressor inlet loop in open loop. Section 10.5 describes the seven PID loops in three layers, the handover of a saturated demand and the MPC option. Section 10.6 reports the measured scenarios, from the architecture comparison to the graded dispatch day and the MPC against the hierarchy. Section 10.7 describes the hardware of such a plant. Sections 10.8 to 10.10 state the limitations, the three results that carry past this plant and the conclusion.

10.2 Background

Supercritical carbon dioxide cycles have been studied since the 1960s and returned to attention as compact power conversion for nuclear, concentrating solar and waste heat duty. Reviews of the technology and of the cycle layouts that have been proposed agree on the source of the advantage, which is that compression near the critical point costs a small fraction of what compression of a gas costs, and on the price of it, which is that the same region has property gradients steep enough to make both design and control difficult [1, 2]. The reference properties themselves come from the Span and Wagner equation of state [9], and the specific heat of carbon dioxide varies by an order of magnitude across the cooler and the cold end of the recuperator.

> ⧖ HISTORY **A Test Loop at the Critical Point, 2008 to 2010**
>
> Earlier studies of the supercritical carbon dioxide Brayton cycle had been largely analytical. A United States national laboratory then built a small test loop on a skid of about 2 by 3 m, with a motor-driven radial compressor of about 50 kW, to study the stability and control of the cycle near the critical point. Assembly finished in May 2008, and commissioning ran through the 2009 fiscal year. The machine reached 65 000 rpm, flows above 4.1 kg/s and pressure ratios just over 1.65, at inlet densities near 70 % of that of water. Its compressor ran on both sides of the saturation curve, near choke and even in surge. The report named the control problem this chapter models: near the critical point a very large heat flow is needed to move the temperature by as little as 1 K [10].

Dynamic studies of these cycles concentrate on load following, because the recuperator makes the plant slow and the applications are ones where the heat source is not dispatchable. Dynamic models coupled to thermal storage have been built for solar-driven layouts [11], transient models of printed circuit recuperators show behavior a single lumped wall cannot reproduce [12], and wide-range load following strategies for recompression layouts treat the recompression split fraction as the principal handle [13, 14]. The application implements a simple recuperated cycle with no recompression branch, so the handles available to it are firing rate, compressor speed, coolant flow and a turbine bypass. An organic Rankine cycle of the same megawatt scale, with its own coupling through one fluid inventory, is the subject of Volume 3, Chapter 5.

Pause. Try Exercise 10.1 before reading on.

Mid-ranging, also called valve position control, pairs a fast actuator of limited range with a slow actuator of large range and returns the fast one toward a resting position so that its authority is available again (Section 1.4.7). The structure is old in the process industries and its tuning has been given a model-based treatment [4, 15]. Chapter 9 applies it to chilled water from a gas turbine. What this plant adds is a third layer and an explicit cost ordering, because on this plant the fastest actuator is also the only one that destroys work, and the slowest is not the cheapest. Predictive control is the usual alternative to such a ladder, because one optimization can spread a demand over several handles by speed and by cost at once [5].

> **HISTORY A Ladder of Handles for Load Following, 2011**
>
> In 2011 a national laboratory report on supercritical carbon dioxide cycles for fast reactors set out a load-following strategy that ranks its handles much as this chapter does. Moving carbon dioxide out to inventory tanks gave the highest efficiency at reduced load, but only over the range the tank volume allowed, from 50 to 90 % load. The turbine bypass was the simplest and fastest handle, almost instantaneous because it acts through pressure rather than temperature, so it made the final, fine adjustment of generator output. It was also the least efficient, because the heat entering the cycle did not fall with the output. The report therefore added a throttle valve at the turbine inlet below 50 % load, with its settings chosen so that the bypass action stays small [16].

10.3 Process Model

The cycle and the control hierarchy are shown in Figure 10.1. The model carries a specific enthalpy for each of the six cycle stations, two metal wall temperatures, a cooler effectiveness, a heat source temperature, a firing-rate actuator state and three actuator positions, which is fourteen states in all. The clock is in minutes, temperatures in °C, pressures in MPa and specific enthalpies and specific work in kJ/kg.

10.3.1 The Design Point

At design the compressor runs at 85 % speed against a 15 °C coolant supply, the firing rate is unity and the bypass is shut. That gives 5.499 kg/s of carbon dioxide circulating between 6.78 and 20.58 MPa, with station temperatures of 20.50, 38.62, 273.86, 593.10, 384.25 and 64.42 °C around the loop. The turbine produces 201.74 kJ/kg against 20.84 kJ/kg of compression, the recuperator transfers 2040.6 kW at an effectiveness of 0.9254, the primary heater supplies 2269.0 kW, and the net electrical output is 929.0 kW at 40.94 % thermal efficiency.

The recuperator carries 47.4 % of the heat the cycle needs. That is the number that makes the cycle efficient and the number that makes Section 10.6.6 uncomfortable, because a controller that holds the power hides the loss of it completely.

10.3.2 The Compressor and the Near-Critical Region

Compression work is proportional to the pressure ratio less one and to a penalty that rises steeply as the inlet temperature leaves the near-critical region, as shown in Equations (10.1) and (10.2).

$$w_c = K_{c0}\left[1 + (\Pi_{\max} - 1)\frac{u}{1+u}\right]\left(\frac{P_2}{P_1} - 1\right), \tag{10.1}$$

$$u = \left(\frac{\max(0,\ T_1 - T_c)}{\Delta T_c}\right)^{n_c} \tag{10.2}$$

where T_c is the temperature above which the fluid begins to expand, ΔT_c is the rise that produces half of the maximum penalty and $\Pi_{\max}$ is that maximum. The penalty is written as a saturating function rather than a bare power law so that a transient excursion cannot become a runaway that no real

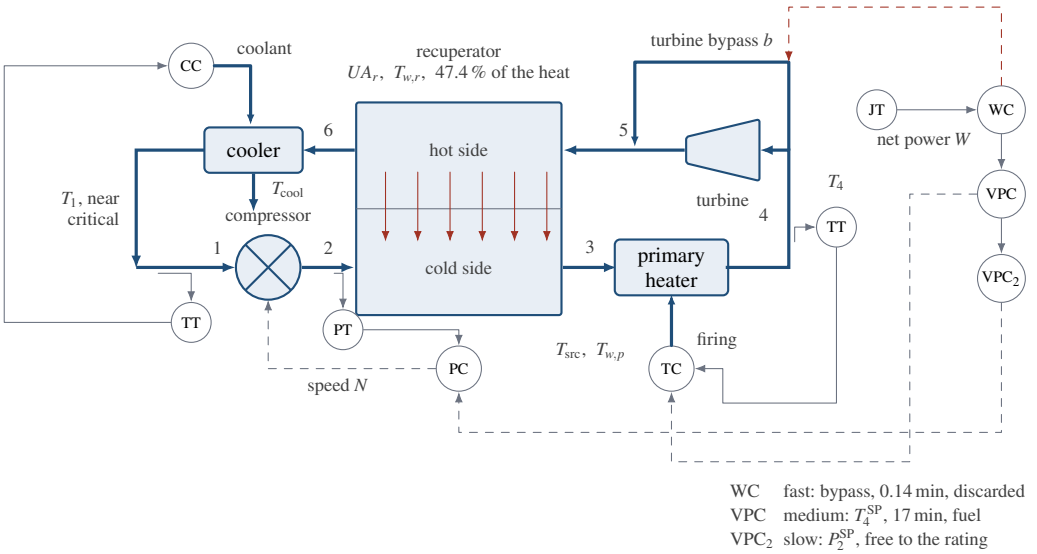

Figure 10.1 The simulated cycle and its control hierarchy. Net power is taken by the turbine bypass in seconds (WC), a valve position controller returns the bypass toward its resting position by moving the turbine inlet temperature setpoint (VPC), and a second valve position controller returns the turbine inlet temperature to its reference by raising the compressor discharge pressure setpoint (VPC$_2$). Those two setpoints go to the firing controller (TC) and the speed controller (PC). The compressor inlet temperature loop (CC) is not part of the hierarchy and is worth more than all of it.

machine has. Mass flow falls in the same variable through a density ratio, and pressure ratio rises with speed, as shown in Equations (10.3) and (10.4).

$$\frac{P_2}{P_1} = \Pi_0 \left(\frac{N}{N_0} \right)^{n_\Pi}, \tag{10.3}$$

$$\dot{m} = \dot{m}_0 \, f_i \, \frac{N}{N_0} \, \mathrm{e}^{-k_\rho (T_1 - T_{1,0})} \tag{10.4}$$

The exponential density ratio in (10.4) is clipped to 0.40 to 1.55. Speed therefore raises the mass flow and the pressure ratio together, which is why its power gain is large and its efficiency gain is positive. Inlet temperature moves (10.1) and (10.4) in opposite directions at the same time, which is why the compressor inlet loop is worth what Section 10.6.5 measures.

Turbine specific work is an ideal expansion with an isentropic efficiency, as shown in Equation (10.5).

$$w_t = f_t \, \eta_t \, c_{p,p} \, (T_4 + 273.15) \left(1 - \Pi_t^{-\gamma} \right) \tag{10.5}$$

The work is proportional to the absolute inlet temperature, which is why the turbine inlet temperature is the plant's other power handle and why the metal limit is what binds that handle.

10.3.3 Stations, Duties and the Exchanger Walls

Pressures around the loop follow from the compressor characteristic and four fixed pressure drops, as shown in Equation (10.6).

$$\begin{aligned} P_2 &= P_1 \, \Pi(N), \quad & P_3 &= P_2 - \Delta P_{23}, \quad & P_4 &= P_3 - \Delta P_{34}, \\ P_5 &= P_1 + \Delta P_{51}, \quad & P_6 &= P_5 - \Delta P_{56} \end{aligned} \tag{10.6}$$

Table 10.1 Controlled, manipulated and disturbance variables with the states that carry them. The six station enthalpies are numbered around the loop from the compressor inlet.

Symbol	Meaning	Unit	Range
Controlled variables			
W	net electrical power	kW	300–1300
T_4	turbine inlet temperature	°C	470–618
P_2	compressor discharge pressure	MPa	13.5–23.8
T_1	compressor inlet temperature	°C	15–32
Manipulated variables			
φ	firing rate	—	0.45–1.40
N	compressor speed	%	55–102
u_c	coolant valve demand	%	0–100
b	turbine bypass position	%	0–78
Disturbances			
T_{cool}	coolant supply temperature	°C	5–34
f_r	recuperator condition	—	0.40–1.05
f_p	primary heater condition	—	0.50–1.10
f_w	cooler condition	—	0.50–1.10
f_t	turbine condition	—	0.80–1.03
f_i	carbon dioxide inventory	—	0.80–1.10
ΔT_{src}	heat source temperature bias	°C	−80 to 60
States			
$h_1 \dots h_6$	station specific enthalpies	kJ/kg	50–2200
$T_{w,r}$	recuperator wall temperature	°C	—
$T_{w,p}$	heater wall temperature	°C	—
ε	cooler effectiveness	—	0.02–0.995
T_{src}	heat source temperature	°C	—
φ_a	firing rate after the actuator lag	—	0–2
N, u_c, b	actuator positions	—	—

The compressor inlet pressure P_1 is set by the carbon dioxide inventory and is treated as a constant. Station temperatures are recovered from the enthalpies by a constant specific heat per segment, as shown in Equation (10.7).

$$\begin{aligned} T_1 &= T_{1,n} + \frac{h_1 - h_{1,n}}{c_{p,w}}, & T_2 &= T_1 + \frac{h_2 - h_1}{c_{p,c}}, \\ T_3 &= T_2 + \frac{h_3 - h_2}{c_{p,rc}}, & T_4 &= T_3 + \frac{h_4 - h_3}{c_{p,p}}, \\ T_5 &= T_{5,n} + \frac{h_5 - h_{5,n}}{c_{p,rh}}, & T_6 &= T_5 - \frac{h_5 - h_6}{c_{p,rh}} \end{aligned} \tag{10.7}$$

The three exchangers are metal walls with thermal mass rather than algebraic effectiveness relations, because their inertia is a large part of what the controller is fighting. The recuperator wall is heated by the mean of the two hot-side temperatures and cooled by the mean of the two cold-side temperatures, and the heater wall is heated by the source and cooled by the cold stream, as shown in Equations (10.8) and (10.9).

$$\tau_{w,r}\dot{T}_{w,r} = \tfrac{1}{2}(T_5 + T_6) - 2T_{w,r} + \tfrac{1}{2}(T_2 + T_3), \tag{10.8}$$

$$\tau_{w,p}\dot{T}_{w,p} = (T_{\text{src}} - T_{w,p}) - g_p(T_{w,p} - T_3) \tag{10.9}$$

The primary-heater conductance slider is limited to 1.3 kJ/(kg °C) to 4 kJ/(kg °C). The lower limit allows for the heater-fouling multiplier, which can halve the conductance. A minimum of 0.7 would still permit sub-zero cycle temperatures in this surrogate model. At the 1.3 minimum, with fouling

factors of 0.5, 0.75 and 1, all twelve cases stay above 17.2 °C, whether the slider is set before Run or moved during the run. The limit bounds the slider's calibrated range and changes neither the nominal design nor the shipped scenarios.

The duties each side of those walls are conductances times the driving difference, as shown in Equation (10.10).

$$\begin{aligned} q_r &= UA_r f_r\left(T_{w,r} - \tfrac{1}{2}(T_2 + T_3)\right), \\ q_p &= UA_p f_p (T_{w,p} - T_3) \end{aligned} \tag{10.10}$$

The recuperator adds to the inertia of the firing path rather than creating it. Part of the added heat returns through the recuperator and has to cross its wall again, but the firing actuator, the source, the heater wall and the heater fluid already form a chain of lags of 2, 4, 5 and 3 min. In a counterfactual run with the recuperator wall lag at its 0.5 min minimum, the time for net power to reach sixty-three percent of a firing step falls only from 17.0 to 15.8 min.

The cooler is written differently. Its state is an effectiveness rather than a temperature, and its duty is bounded so that the outlet approaches the coolant supply temperature and can never pass it, as shown in Equation (10.11).

$$q_w = \varepsilon\,(h_6 - h_{\text{cool}}), \qquad h_{\text{cool}} = h_{1,n} + c_{p,w}(T_{\text{cool}} - T_{1,n}) \tag{10.11}$$

The coolant valve is characterized. A cooler's effectiveness saturates hard in coolant flow and this one is already 91 % effective at design, so an uncharacterized valve would hand the controller a loop gain that varies roughly twenty to one across its travel. The trim is chosen so that the controller output is the effectiveness on a linear scale and the loop gain is constant, as shown in Equation (10.12), with the physical flow and therefore the cubic pump power recovered afterward.

$$\varepsilon_v = f_w\left[\varepsilon_{\text{lo}} + (\varepsilon_{\text{hi}} - \varepsilon_{\text{lo}})\frac{u_c}{100}\right], \qquad \frac{\dot m_c}{\dot m_{c,\text{ref}}} = \frac{-\ln(1-\varepsilon)}{\text{NTU}_0} \tag{10.12}$$

That choice is not visible in the animation and it is the difference between a loop that can be tuned once and a loop that oscillates at one end of its range and crawls at the other.

10.3.4 State Equations

Each station enthalpy is a first-order tank driven by the work or the duty applied to it, as shown in Equation (10.13).

$$\begin{aligned} \tau_{\text{whx}}\dot h_1 &= (h_6 - q_w) - h_1, & \tau_{\text{comp}}\dot h_2 &= (h_1 + w_c) - h_2, \\ \tau_{rc}\dot h_3 &= (h_2 + q_r) - h_3, & \tau_{\text{phx}}\dot h_4 &= (h_3 + q_p) - h_4, \\ \tau_{\text{turb}}\dot h_5 &= h_4 - (1-b)w_t - h_5, & \tau_{rh}\dot h_6 &= (h_5 - q_r) - h_6 \end{aligned} \tag{10.13}$$

The heat source, the firing actuator, the cooler and the three actuator positions are first-order lags toward their commands, as shown in Equation (10.14).

$$\begin{aligned} \tau_{\text{fuel}}\dot\varphi_a &= \varphi_{\text{cmd}} - \varphi_a, \\ \tau_{\text{src}}\dot T_{\text{src}} &= T_{\text{src},0} + s_g(\varphi_a - 1) + \Delta T_{\text{src}} - T_{\text{src}}, \\ \tau_{w,c}\dot\varepsilon &= \varepsilon_v(u_c) - \varepsilon, \\ \tau_N\dot N &= N_{\text{cmd}} - N, \\ \tau_c\dot u_c &= u_{c,\text{cmd}} - u_c, \\ \tau_b\dot b &= b_{\text{cmd}} - b \end{aligned} \tag{10.14}$$

Electrical output is the turbine work that is not bypassed, less the compressor work and the coolant pump power, as shown in Equation (10.15).

$$W = \dot m\,\eta_g w_t (1-b) - \frac{\dot m\, w_c}{\eta_m} - P_{\text{aux}}, \qquad \eta = \frac{W}{\dot m\, q_p} \tag{10.15}$$

with $P_{\text{aux}} = P_{\text{aux},0}(\dot{m}_c/\dot{m}_{c,\text{ref}})^3$ from (10.12), where $\dot{m}_c$ is the physical coolant flow. After every step the six enthalpies are clipped to 50 kJ/kg to 2200 kJ/kg, the firing state to 0 to 2, the speed to 20 % to 110 %, the coolant demand to 0 % to 100 %, the bypass to 0 % to 100 % and the cooler effectiveness to 0.02 to 0.995.

10.3.5 Measurements

Four measurements close the loops, and each is a first-order filter on the true value with optional noise, updated at the controller scan. The turbine inlet thermocouple has a time constant of 0.35 min, the compressor inlet thermocouple 0.25 min, the pressure transmitter 0.05 min and the power transducer 0.10 min. Noise on the two thermocouples and the power transducer is band-limited with correlation times of 0.5, 0.4 and 0.2 min, the compressor inlet carrying 15 % of the thermocouple amplitude, and it is switched off by default and in every scenario. Nothing on the plant measures the mass flow, the recuperator effectiveness or the thermal efficiency, all three of which are computed for display from the states.

Pause. Try Exercise 10.3 before reading on.

10.3.6 Parameters

Table 10.2 lists the constants needed to rebuild the simulator. Those marked hidden are calibration references and are not offered for adjustment, while the five specific heat surrogates can be changed with the other constants.

10.3.7 Initialization and Numerics

Initialization solves a three-by-three damped Newton system in the compressor inlet, recuperator cold outlet and turbine inlet enthalpies, with both wall temperatures and the cooler effectiveness eliminated in closed form. Successive substitution fails here, because the recuperator returns about four times what a first pass puts into it and the iteration diverges. The reason is worth stating because it is also why the plant is slow. In the dynamic model the same feedback is broken by the wall states having their own time constants, which is what makes the plant stable and the arithmetic unstable. Three inverse calculations use the same solve, returning the firing rate that holds a chosen turbine inlet temperature, the coolant demand that holds a chosen compressor inlet temperature and the firing rate that produces a chosen net power. Manual mode offers the first two as one-press helpers.

The engine of Section 1.6.1 integrates the model by fourth-order Runge-Kutta at a fixed 0.02 min (1.2 s). All seven controllers execute on a 0.1 min (6 s) scan, and the historian records ten samples per simulated minute.

10.3.8 Assumptions

Carbon dioxide properties are constant specific heats per segment, with the near-critical behavior carried entirely by (10.1) and (10.4) as fitted functions rather than derived from an equation of state [9]. There is no recompression branch and no inventory control loop, so the flow split and the charge are parameters rather than actuators. Turbomachinery is quasi-steady, with no shaft dynamics, no surge line and no windage. Each exchanger is a single lumped wall. The heat source is a first-order lag toward a temperature proportional to firing rate, which suits a fired heater. The compressor inlet pressure is fixed by the inventory and does not move with the loop temperature. Pressure drops are constants rather than functions of flow. Section 10.8 states where each of these fails.

10.4 Actuator Characteristics and Open-Loop Response

10.4.1 Three Actuators for One Variable

Table 10.3 reports the steady-state gain and the open-loop response of net power to each actuator, measured from the design point in manual with the other actuators fixed. Response times are counted

Table 10.2 Model constants with their default values.

Symbol	Meaning	Value	Unit
$\dot{m}_0$	design mass flow	5.5	kg/s
N_0	design compressor speed (hidden)	85	%
Π_0	design pressure ratio	3.0354	—
n_Π	ratio exponent on speed	1.35	—
K_{c0}	work per unit pressure ratio	10.2388	kJ/kg
T_c	near-critical threshold	22	°C
$\Pi_{\max}$	largest compression penalty	2.6	—
ΔT_c	half-penalty rise	8.0	°C
n_c	penalty exponent	1.8	—
k_ρ	density loss per degree	0.024	1/°C
η_t	turbine efficiency	0.88	—
γ'	expansion exponent	0.21875	—
η_g, η_m	generator, motor efficiency	0.96, 0.95	—
$P_{\text{aux},0}$	pump power at design flow	15	kW
P_1	compressor inlet pressure	6.78	MPa
ΔP_{23}, ΔP_{34}	cold side, heater drop	0.20, 0.20	MPa
ΔP_{51}, ΔP_{56}	exhaust rise, hot side drop	0.30, 0.20	MPa
$T_{\text{src},0}$	source temperature at design	700	°C
s_g	source temperature per unit firing	280	°C
UA_r	recuperator conductance	10.8986	kJ/(kg °C)
UA_p	heater conductance	1.3478	kJ/(kg °C)
g_p	heater wall to fluid coupling	0.3921	—
ε_{lo}, ε_{hi}	cooler effectiveness, shut, open	0.55, 0.985	—
NTU_0	cooler NTU at design flow (hidden)	2.3772	—
$\dot{m}_{c,\text{ref}}$	coolant flow at design (hidden)	70	%
τ_{fuel}, τ_{src}	firing actuator, source lag	2.0, 4.0	min
τ_{rc}, τ_{rh}, $\tau_{w,r}$	recuperator cold, hot, wall lag	2.0, 2.0, 4.0	min
τ_{phx}, $\tau_{w,p}$	heater fluid, wall	3.0, 5.0	min
τ_{whx}, $\tau_{w,c}$	cooler fluid, wall	3.0, 4.0	min
τ_{comp}, τ_{turb}	compressor, turbine outlet	0.50, 0.20	min
τ_N, τ_c, τ_b	speed, coolant, bypass lag	0.35, 0.40, 0.15	min
$c_{p,w}$	cooler surrogate	4.1982	kJ/(kg °C)
$c_{p,c}$	compressor surrogate	1.1501	kJ/(kg °C)
$c_{p,rc}$, $c_{p,rh}$	recuperator surrogates	1.5774, 1.1602	kJ/(kg °C)
$c_{p,p}$	heater surrogate	1.2924	kJ/(kg °C)
$T_{4,\max}$	turbine metal limit	620	°C
$T_{4,\text{trip}}$	overtemperature alarm	645	°C
$P_{2,\max}$	high-pressure rating	24.0	MPa
$T_{1,\max}$	compressor inlet alarm	32.0	°C
Defaults for the operator-set inputs			
W^{SP}	dispatch setpoint	930	kW
T_4^{SP}, T_4^{ref}	turbine inlet setpoint, reference	593, 593	°C
T_1^{SP}	compressor inlet setpoint	20.5	°C
P_2^{SP}	discharge pressure setpoint	20.58	MPa
b^{ref}	bypass resting position	3	%
φ, N, u_c, b	manual actuator positions	1.0, 85, 82.6, 0	—

Table 10.3 Measured open-loop response of net power to each actuator, from the design point with the other actuators fixed.

Actuator	Steady gain	t_{63} (min)	Cost
Turbine bypass	−10.2 kW/%	0.14	all of it
Compressor speed	+22.1 kW/%	0.33	none
Firing rate	+305 kW per unit	17.02	fuel

from the controller scan that applies the step, so they exclude the up to 0.1 min a change waits for that scan.

Table 10.3 is the whole design problem. The fast actuator is the one that throws work away, because opening the turbine bypass diverts flow around the turbine and the reduction in power is discarded as heat into the recuperator rather than saved. The efficient actuator is the compressor

Table 10.4 Measured open-loop gains and response times of the four measurements, from the design point in manual.

Pairing	Gain	t_{63} (min)	t_{95} (min)
Bypass → net power	−10.2 kW/%	0.14	0.36
Speed → net power	+22.1 kW/%	0.33	0.85
Speed → discharge pressure	0.328 MPa/%	0.35	1.05
Firing → net power	305 kW per unit	17.02	35.85
Firing → turbine inlet	258 °C per unit	16.61	34.87
Coolant → compressor inlet	−0.270 °C/%	8.04	17.21

Table 10.5 Closed-loop steady state at four coolant supply temperatures with the compressor inlet controller in service. The coolant valve saturates between 19 and 22 °C.

Coolant (°C)	T_1 (°C)	Valve (%)	$\dot{m}$ (kg/s)	w_c (kJ/kg)	Net (kW)
15	20.50	82.6	5.500	20.84	929.0
19	20.50	97.4	5.500	20.84	890.9
22	22.82	100.0	5.202	21.38	807.6
29	29.80	100.0	4.400	37.13	597.4

speed, which moves the mass flow and the discharge pressure together through (10.3) and (10.4) and therefore raises cycle efficiency as it raises power, from 38.78 % at 78 % speed to 42.61 % at 92. The patient actuator is the firing rate, which reaches sixty-three percent of a power change in 17.0 min and ninety-five percent in 35.9 min. The ratio of the slowest to the fastest is 122.

Table 10.4 sets the loop hierarchy. In open loop the discharge pressure responds nearly fifty times faster than the turbine inlet temperature, which is the separation a cascade needs and the reason the two supervisory layers can be tuned independently of the inner loops they write to. The firing path contains no transport delay. Its apparent dead time comes from the chain of first-order elements between the fuel valve and the turbine inlet, which is the actuator lag, the source lag, the heater wall and the heater fluid in series. The first two percent of the turbine inlet response appears 3.3 min after the step, and a two-point first-order fit gives an apparent dead time of 5.9 min with a time constant of 10.8 min.

10.4.2 The Compressor Inlet Loop

Compression work and mass flow move in opposite directions in the same variable, so the compressor inlet temperature is the plant's most valuable number. Table 10.5 reports the closed-loop steady state at four coolant supply temperatures with the compressor inlet controller in service and everything else at design.

Fourteen degrees of coolant supply temperature, which is the difference between a spring morning and a summer afternoon on the same cooling tower, nearly doubles the compressor work and removes 331.6 kW, or 35.7 %, of the plant's output. The first four degrees cost 38.1 kW with the compressor inlet held exactly on setpoint, because holding it costs coolant flow and the pump power of (10.15) is cubic in flow. The remaining ten cost 293.5 kW, because the valve has no travel left and the inlet temperature then follows the supply. Nothing in the control hierarchy recovers the loss, because it is not a setpoint error. The cycle has become a different cycle.

10.5 Controller Design

10.5.1 Velocity Form PID Algorithm

Every loop on the plant runs the velocity-form PID law of Section 1.6.2 [17]. Derivative action is available and is set to zero in every loop, because no measurement on this plant is fast enough to

Table 10.6 The seven controllers, their actions and the tunings that Scenarios 3 to 10 load and the challenge suggests. No loop uses derivative action.

Controlled	Manipulated	a	K_c	τ_I (min)
Inner loops				
T_4	firing rate	$+1$	0.006 per °C	12
P_2	compressor speed	$+1$	1.2 %/MPa	1.0
T_1	coolant valve	-1	3 %/°C	8
Supervisory layer				
W	bypass	-1	0.09 % per kW	1.2
W	T_4^{SP} (no fast path)	$+1$	0.60 °C per kW	30
b	T_4^{SP}	$+1$	3.0 °C/%	30
T_4^{SP}	P_2^{SP}	-1	0.035 MPa/°C	60

reward it. The error and the increments are given by Equations (10.16) and (10.17).

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k) \tag{10.16}$$

$$\Delta u_k^P = K_{c,k}\,e_k - K_{c,k-1}\,e_{k-1}, \qquad \Delta u_k^I = \frac{K_c}{\tau_I}\,e_k\,\Delta t_c \tag{10.17}$$

where $\Delta t_c = 0.1$ min (6 s) is the scan period common to every loop and a is the action. The increments are combined and the demand is clipped to the actuator span, as shown in Equation (10.18).

$$\begin{aligned} s_k &= s_{k-1} + \Delta u_k^P \\ c_k &= s_k + I_{k-1} \\ u_k &= \min(u_{\max},\ \max(u_{\min},\ c_k)) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\,(u_k - c_k) \end{aligned} \tag{10.18}$$

Here s_k is the running sum S_k of Equation (1.13), c_k is the output before limiting, v_k of Equation (1.14), with the bias folded into the integral, and the integral I used at scan k was accumulated through scan $k-1$. At reset I is seeded with the starting output, which is the manual actuator position for the inner loops and the bypass loop and the present turbine inlet temperature or discharge pressure for the supervisors, so every loop starts where the plant is. Anti-windup is back-calculation with a tracking time $T_t = \tau_I$. While the output is clipped, the last term removes $\Delta t_c/\tau_I$ of the clipped excess each scan, which is $1/12$ for the bypass loop and $1/600$ for the pressure supervisor, while the running sum s_k is never truncated. Transfer between Manual and Auto is bumpless as Section 1.6.2 describes. A supervisor that the chosen architecture leaves out of service runs the same tracking routine against the value in force, so a change of architecture is bumpless too. Two quantities are needed from the algorithm rather than from the actuator. The applied output u_k drives the plant, and the unsaturated demand c_k is what one controller hands to another, for the reason Section 10.5.3 gives.

10.5.2 Seven Loops in Three Layers

Table 10.6 lists every controller, its action and its tuning. The three inner loops hold a turbine inlet temperature, a compressor discharge pressure and a compressor inlet temperature, and closed they reach sixty-three percent of a setpoint move in 10.4, 2.3 and 8.3 min. Closing the loops shrinks the ratio of the first two to 4.5, against nearly fifty in open loop (Table 10.4). The supervisors above them carry integral times of 30 and 60 min, so each is slower than the loop it writes to by a factor of nearly three or more, which is the only structural requirement a cascade makes [6]. Volume 3, Chapter 3 measures what happens when an inner loop is too slow for that requirement.

The supervisory layer is three mid-ranging controllers in series, as shown in Equations (10.19) to (10.21).

$$b = \mathrm{PID}\left(W^{\mathrm{SP}}, W\right), \tag{10.19}$$

$$T_4^{\mathrm{SP}} = \mathrm{PID}\left(b^{\mathrm{ref}}, b^{\mathrm{raw}}\right), \tag{10.20}$$

$$P_2^{\mathrm{SP}} = \mathrm{PID}\left(T_4^{\mathrm{ref}}, T_4^{\mathrm{SP}}\right) \tag{10.21}$$

The signs repay attention. A bypass demand above its resting position means the thermal side is making more power than is wanted and discarding the excess, so the turbine inlet temperature should come down, and $a = +1$ in (10.20) delivers that. A turbine inlet setpoint above its own reference means the temperature is carrying load the pressure should be carrying, so the pressure setpoint should rise, and $a = -1$ in (10.21) delivers that. Both layers are switchable, so the application offers the three architectures that Section 10.6.1 compares.

10.5.3 Handing a Saturated Demand Upward

Equation (10.20) watches b^{raw}, which is the bypass controller's unsaturated demand c_k from (10.18), and not the applied bypass position u_k, as Section 1.4.7 recommends. The distinction is the difference between a hierarchy that works and one that works in only one direction.

The bypass is clamped shut at zero. When the plant is short of power the fast controller demands a negative bypass, which it cannot have, so the applied position sits at zero and reports a shortfall of exactly b^{ref} no matter how much power is missing. A valve position controller watching that signal receives the same constant error for a 5 kW deficit and a 500 kW deficit, and it walks the turbine inlet temperature up at the same slow rate in both cases. The unsaturated demand goes properly negative and states how much power the fast path wishes it could make. While the bypass is held shut against a steady shortfall, the back-calculation of (10.18) lets the integral decay and the demand settles at the proportional term $K_c(W - W^{\mathrm{SP}})$, which is –9 % for a 100 kW shortfall, so the handover carries magnitude information in both directions. An earlier build of the application handed over the clamped position, and the symptom was a plant that shed load in seconds and picked it up in hours.

Pause. Try Exercise 10.2 before reading on.

10.5.4 Model Predictive Control

Beside manual operation, the inner loops alone and the cascade, the supervisory layer can be a model predictive controller (MPC), which is never the default and which only scenario 11 loads. Section 1.6.3 describes the option in general. It replaces the bypass loop and both valve position controllers of Equations (10.19) to (10.21). Every $T_s = 0.5$ min, or five scans, it plans the bypass command b, the turbine inlet setpoint T_4^{SP} and the discharge pressure setpoint P_2^{SP} together over $N = 60$ steps (30 min) and applies the plan's first move alone [5]. A dispatch change between decisions is answered on the scan it arrives. The three inner loops of Table 10.6 stay in service and carry the setpoints out, the commands are bounded by the clamps the cascade uses (bypass 0 % to 78 %, inlet setpoint 470 °C to 618 °C and at most $T_{4,\max} - 2\,^{\circ}\mathrm{C}$, pressure setpoint 13.5 MPa to 23.8 MPa), and every alarm stays on the run's own limits above both structures.

The controller's model describes the closed inner loops rather than the plant. Step tests of ±10 °C and ±0.5 MPa on scenario 4, with the tuned loops of Table 10.6 and the noise off, give the turbine inlet loop as a damped second-order response behind a lag and the pressure loop as a lead-lag, as shown in Equation (10.22).

$$\begin{aligned} T_4 &= \frac{\omega^2\, T_4^{\mathrm{SP}}}{(s^2 + 2\zeta\omega s + \omega^2)(\tau_3 s + 1)} + G_P(s)\, P_2^{\mathrm{SP}} + G_b(s)\, b \\ P_2 &= \frac{(as+1)\, P_2^{\mathrm{SP}}}{(\tau_1 s + 1)(\tau_2 s + 1)}, \qquad \tau_b \dot{x}_b = b - x_b \end{aligned} \tag{10.22}$$

Here s is the Laplace variable in 1/min, $\omega = 0.19668$ 1/min, $\zeta = 0.28393$, $\tau_3 = 4.4749$ min, $a = 1.0561$ min, $\tau_1 = 3.2096$ min, $\tau_2 = 0.4062$ min, and x_b is the bypass position behind the declared valve lag $\tau_b = 0.15$ min. The cross terms G_P and G_b have the form $s(\beta_2 s + \beta_1)/[(s^2 + 2\zeta_x\omega_x s + \omega_x^2)(s + 1/\tau_x)]$ with their own fitted constants, and carry the disturbance that the pressure setpoint and the bypass make on the turbine inlet before its loop removes it. Each loop has unit gain because each has integral action. The model reaches 63 % of an inlet setpoint step at 10.80 min, as the plant does, and overshoots by 26.7 % at 21.2 min against 26.3 % at 20.9 min, with an RMS residual of 0.61 % of the peak. The pressure loop reaches 63 % at 2.40 min with a residual of 0.26 %, and the cross terms fit to 1.51 and 1.78 % of their peaks.

Net power is Equation (10.15) itself, written with the declared defaults of Table 10.2, the design turbine condition and inventory, the speed that gives P_2 through (10.3) and the coolant pump's design power, never with the run's values. Its slopes change more than threefold across the handles' ranges (−3.6 to −13.7 kW/%, 0.27 to 1.54 kW/°C and 16.7 to 75.5 kW/MPa), so the map is relinearized at every decision, as shown in Equation (10.23).

$$\hat{W}_k = \bar{W} + g_b(x_{b,k} - \bar{x}_b) + g_T(\hat{T}_{4,k} - \bar{T}_4) + g_P(\hat{P}_{2,k} - \bar{P}_2) \tag{10.23}$$

The bars mark the estimates at the decision, $\bar{W}$ includes the power bias below, the compressor inlet temperature is held at its measured value over the horizon, and g_b, g_T and g_P are the partial derivatives of the map at that point. At design they are −10.65 kW/%, 1.230 kW/°C and 70.10 kW/MPa, against steady changes of −10.66, 1.223 and 70.04 in the same units on the plant with its inner loops closed.

The estimator runs every scan whichever structure is in charge, which lets the MPC take over from a current estimate. The model is driven by the commands applied, and an output bias on each of the three readings the shipped loops use is updated as shown in Equation (10.24).

$$d_{y,k} = d_{y,k-1} + \left(1 - e^{-\Delta t_c/\tau_y}\right)\left(y_k - \hat{y}_k - d_{y,k-1}\right) \tag{10.24}$$

Here y is the turbine inlet thermocouple, the pressure transmitter or the power transducer, $\hat{y}_k$ is the model's prediction of that reading for the scan (for power, through the transducer's 0.10 min lag), and τ_y is 1.0, 0.5 and 0.5 min in that order. A constant bias on each measured output removes the steady offset of a model that is wrong in a declared way [18], which is how a fouled recuperator, a worn turbine or a warm afternoon reaches the plan without being announced. The update is an output-bias filter, not a Kalman filter, and while the inner loops are open the model is reset onto the readings every scan.

A decision solves the quadratic program of Equation (10.25) on the model sampled at T_s, with the three commands held over 12 blocks of 1, 1, 1, 1, 2, 2, 3, 4, 6, 9, 15 and 15 steps, which with the slack makes 37 decision variables.

$$\begin{aligned}
\min_{\mathbf{u},s}\ & \sum_{k=1}^{N}\left(\hat{W}_k - W^*\right)^2 + \ell_b \sum_{k=1}^{N} b_k + \rho s^2 + \sigma s \\
& + \lambda \sum_{j=1}^{12}\left[\left(g_b \Delta b_j\right)^2 + \left(g_T \Delta T_{4,j}^{\mathrm{SP}}\right)^2 + \left(g_P \Delta P_{2,j}^{\mathrm{SP}}\right)^2\right] \\
& + r_T\, g_T^2 \sum_{k=1}^{N}\left(T_{4,k}^{\mathrm{SP}} - T_4^{\mathrm{ref}}\right)^2 \\
\text{s.t. } & \mathbf{u}_{\min} \le \mathbf{u}_j \le \mathbf{u}_{\max}, \qquad s \ge 0, \\
& \hat{T}_{4,k} \le T_{4,\max} - \delta_T + o\, e^{-kT_s/\tau_f} + s
\end{aligned} \tag{10.25}$$

Here $\mathbf{u}_j$ holds b, T_4^{SP} and P_2^{SP} over block j, Δ is the change from the block before (from the commands applied now for the first), $\lambda = 0.05$ prices a move by the square of the power it is worth, $r_T = 0.02$ returns the inlet setpoint toward the reference T_4^{ref} of (10.21), and $\ell_b = 2|g_b|\, e^*$ per percent prices

the bypass so that, at the margin, it opens only for the part of a power error above $e^* = 10\,\text{kW}$. The metal row holds the predicted inlet $\delta_T = 1\,^\circ\text{C}$ under the run's limit $T_{4,\max}$, o is the amount by which the estimated inlet already exceeds $T_{4,\max} - \delta_T$ and decays through a funnel with $\tau_f = 5\,\text{min}$, and $\rho = 10^5$ per $^\circ\text{C}^2$ and $\sigma = 10^4$ per $^\circ\text{C}$ price the slack s. The bounds are hard and the metal row is soft, sharing a single slack across all 60 steps, since the shared solver allots one to each soft output. The target W^* is the dispatch clipped to the steady powers the model reaches with every handle at its bound (the bypass shut, the inlet setpoint at the lower of its clamp and $T_{4,\max} - \delta_T$ and the pressure setpoint at its clamp, and the reverse), so a dispatch beyond reach does not pull the plan into the slack. With e^* fixed at 10 kW, under half the ±25 kW band, the horizon and λ were chosen once on scenario 11 (horizons of 20, 30 and 45 min give 795, 781 and 784 kW min), and nothing is retuned per case. The compressor inlet carries no row, because the three handles move it by at most 0.0064 °C per unit with the cooler saturated and no compared run came within 4.25 °C of its 32 °C alarm.

The controller keeps a plan only when the solver's status is optimal, which is tested even when the call succeeds, when every array has the expected size, when each planned command stays within its bounds and the first matches the move, and when commands, predictions, cost and slack are all finite. A specification test injects five such faults and requires each to be rejected. Otherwise the cascade takes the handles with the fast path the learner has set, because its bypass loop, power loop and both valve position controllers have tracked the applied commands with no preload, the historian marks the fallback, and the next valid plan takes the handles back. On scenario 11 without noise, forcing a rejection at 40 min moves the inlet setpoint by 5.0×10^{-4} °C and the pressure setpoint by 1.8×10^{-5} MPa on the hand-over scan. The first plan is made at $t = 0$, and none of the 5970 plans in the nine MPC runs of Table 10.11 was rejected. In Chrome a plan takes 2.05 to 2.21 ms on average over three cases, which is 0.41 to 0.44 % of one core when a simulated minute passes each second and 8.2 to 8.8 % at the top speed of twenty such minutes, and a run played at that speed logs the same rows as one stepped without the browser.

Three models stay separate, the plant on the run's parameters, the four filtered instruments, and the controller's model built from the declared defaults. The MPC reads the instruments, the commands applied, the dispatch and the operator settings, and of the run's parameters only the metal limit, which it treats as a protection limit in the row, the clip and the setpoint bound. A test replays the same instrument readings over a different hidden plant and requires the same estimate and plan. A learner's change to the plant is therefore mismatch whose steady part the biases remove, while a change to the inner-loop tunings is mismatch in the dynamics that nothing corrects, because the model assumes the loops at their identification values (K_c of 0.006 per °C and τ_I of 12 min for the turbine inlet, 1.2 %/MPa and 1.0 min for the pressure).

10.6 Results

Every number in this section was produced by driving the application's own model at its shipped constants and seed. The architecture comparisons of Sections 10.6.1 and 10.6.2 use one protocol. Each structure is selected before Run on Scenario 4, which starts from the design state and has no scripted events, run for 240 min at a dispatch of 930 kW so that the supervisory integrators are close to rest, and then stepped. Integral error is the integral of the absolute deviation of net power from the dispatch setpoint over the 185 min that follow, evaluated at the 0.02 min integration step, and the band is ±25 kW.

10.6.1 Three Architectures on a Dispatch Increase

The central experiment steps the dispatch from 930 to 1050 kW and runs the same plant under four configurations.

The first three rows of Table 10.7 do not merely track badly. They never arrive. Turbine inlet temperature is bounded by the 618 °C clamp on its own setpoint, which at the design speed and pressure is worth 959.5 kW, or 90 kW short of the request, and adding a bypass fast path in front

Table 10.7 Dispatch increase from 930 to 1050 kW. The band is ±25 kW, the turbine inlet setpoint is clamped at 618 °C and the turbine metal limit is 620 °C.

Structure	Rest (kW)	In band (min)	IAE (kW min)	Peak T_4 (°C)	Net (kW)
T_4 only, bypass shut	930.0	never	16968	624.4	959.5
T_4 only, bypass at 3 %	926.6	never	22827	618.0	926.6
Bypass and T_4	930.0	never	16751	618.3	959.5
Three deep	930.0	**6.8**	**453**	617.0	**1050.0**

Table 10.8 Dispatch reduction from 930 to 800 kW, same protocol.

Structure	In band at (min)	IAE (kW min)	Bypass loss (kWh)	Efficiency (%)
T_4 only, bypass shut	44.8	4998	0.0	38.70
T_4 only, bypass at 3 %	44.0	4999	91.8	38.43
Bypass and T_4	2.0	255	242.9	38.15
Three deep	**1.3**	**197**	107.0	**38.77**

of it changes the destination not at all, because a bypass can only subtract. What the fast path does change is the approach. Without it the temperature loop overshoots its own setpoint and takes the metal to 624.4 °C, spending 12.7 min above the 620 °C limit, while with it the peak is 618.3 °C and the limit is never reached. The second row is the same structure with three percent of the turbine work bypassed and nothing to close it, and it cannot hold even the 930 kW it started from.

The third layer supplies what the others lack, which is headroom. Raising the discharge pressure setpoint takes the compressor from 85.0 to 92.0 % speed and the discharge from 20.58 to 22.89 MPa, and (10.4) converts that into mass flow the turbine can expand. The plant reaches 1050.0 kW, enters the band in 6.8 min, peaks at 617.0 °C without reaching the metal limit, and settles at 41.76 % efficiency against 41.46 % for the temperature-only structure and 40.59 % for the bypassed one. The integral error is smaller by a factor of 37 than the two-layer structure and 50 than the bypassed one. The handover is slow to finish, because the turbine inlet temperature is still returning to its 593 °C reference five hours after the step while the power has been on setpoint since the first ten minutes.

10.6.2 The Same Comparison Downward

Reducing dispatch is a different problem, because the fast actuator now points the right way.

Table 10.8 is the ordinary mid-ranging result and then something better. The bypass fast path is twenty-two times faster into the band than the firing rate alone and it pays 242.9 kWh for the privilege, which is the trade the textbook describes. Adding the third layer is faster still, at 1.3 min, and it discards 107.0 kWh, which is less than half, because backing the compressor speed down to 80.5 % removes the power at the source and lets the bypass reclose to 2.0 %. The three-deep structure is simultaneously the fastest and the most efficient of the four, at 38.77 % against 38.70 % for the structure that discards nothing at all, which is unusual enough to be worth the extra controller. Figure 10.2 sets the two comparisons side by side.

> ✦ **DID YOU KNOW? The Fastest Structure Is Also the Most Efficient**
>
> From the simulator: on a dispatch cut from 930 to 800 kW, the three-deep structure is the fastest into the band, at 1.3 min, and also the most efficient, at 38.77 % against 38.70 % for the firing rate alone with the bypass shut. It discards 107.0 kWh, less than half the 242.9 kWh of the two-layer structure (Table 10.8).

10.6.3 Load Rejection

Scenario 9 opens the breaker at 10 min and demands 450 kW immediately. With the fast path in service the plant is within 50 kW of the new dispatch 3.4 min later and inside the ±25 kW band at

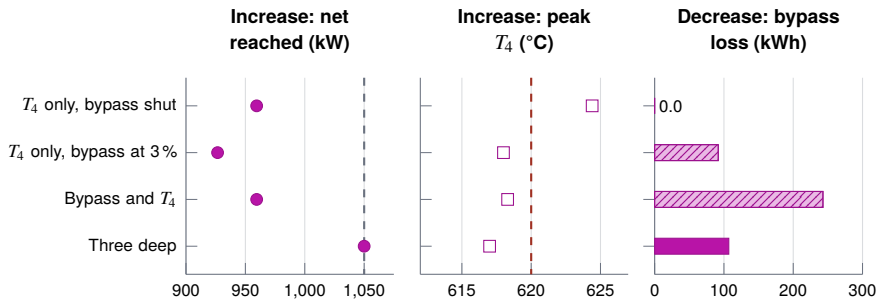

Figure 10.2 Tables 10.7 and 10.8 redrawn. Left: the net power each structure settles at after the dispatch increase, against the 1050 kW request (dashed). Middle: the peak turbine inlet temperature on the way, against the 620 °C metal limit (dashed). Right: the energy discarded through the bypass on the dispatch reduction, with the three-deep structure solid. Notice that only the three-deep structure reaches the request, that the temperature-only structure with the bypass shut is the only one to cross the metal limit, and that the three-deep structure discards less than half of what the two-layer structure discards.

4.9 min, for an integral error of 607 kW min over the following eighty minutes. With the fast path switched off before Run, the firing rate is the only handle, the plant is still at 751.4 kW eighty minutes later, and the integral error is 26 440 kW min, a factor of 44. This is the scenario in which the bypass, which the previous two subsections treat as an expense to be minimized, is the only thing standing between the plant and an overspeed trip.

10.6.4 Where the Third Layer Stops

Scenario 8 asks for 1080 kW at 15 min and then 1270 kW at 220 min. The first request is answered by the third layer as before, and at 180 min the plant holds 1080 kW with the bypass back at 4.3 %, the discharge at 23.3 MPa and the turbine inlet at 607.8 °C. The second request is not answered. All three setpoints reach their stops at the scan that receives it, and Figure 10.3 catches the plant fifteen minutes later. The discharge pressure setpoint is at its 23.8 MPa clamp against a 24.0 MPa rating, the turbine inlet setpoint is at its 618 °C clamp with the temperature at 617.4 °C and still rising, the bypass is fully shut, and the plant makes 1188 kW of a requested 1270 kW with no actuator left to move. The turbine inlet then overshoots its clamped setpoint to 620.8 °C and spends 7.2 min above the metal limit, and the plant settles at 1189.9 kW with the inlet at 618.0 °C.

That is the correct behavior and it is worth staging deliberately. A cascade of valve position controllers does not remove constraints, it reorders which one binds. The purpose of the ladder is to spend the cheap authority before the expensive authority and the expensive authority before the unavailable authority, and when all three are gone the answer is that the plant cannot deliver the dispatch.

> **TRY IT IN THE APP** **8 · Three Deep, and the Rating That Ends It**
>
> Load the scenario, press Run and let it pass 220 min, when the dispatch rises to 1270 kW. Watch the bypass, turbine inlet and discharge pressure strips. All three handles reach their stops, the turbine inlet overshoots to 620.8 °C and spends 7.2 min above the metal limit, and the plant settles at 1189.9 kW. The ladder reorders which constraint binds and creates no capacity.

10.6.5 A Loop Outside the Hierarchy

Scenario 3 warms the cooling tower from 15 to 29 °C with the compressor inlet controller in service. The controller does what it can, driving the coolant valve to 100 %, and the inlet temperature still

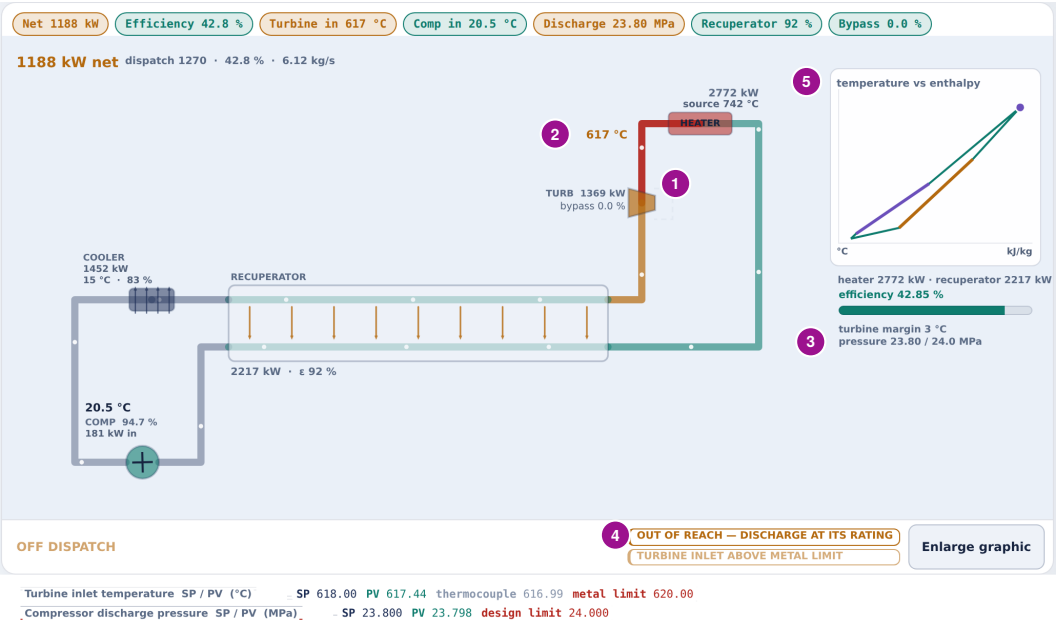

Figure 10.3 Scenario 8 at 3.92 h (235.2 min), fifteen minutes after a dispatch request of 1270 kW. The image is a composite of the application's window: the plant view, and below it the names and readings of two trend strips. All three layers of the hierarchy have run out at once. The bypass (1) is shut at 0.0 %, the turbine inlet setpoint is at its 618 °C clamp with the temperature (2) at 617.4 °C, 2.6 °C below the metal limit, and the compressor discharge (3) is 23.80 MPa against a 24.0 MPa rating. The plant makes 1188 kW, and the banner (4) reports the dispatch out of reach with the discharge at its rating. The inset (5) is the cycle drawn on temperature and enthalpy axes with the recuperator as the shared segment.

settles at 29.80 °C. Net power falls from 928.9 kW to 597.4 kW and thermal efficiency from 40.94 % to 35.33 %, and every controlled variable except the compressor inlet is exactly where it was.

> **TRY IT IN THE APP** **3 · The Cheapest Loop on the Plant**
>
> Load the scenario and press Run. At 20 min the coolant supply warms from 15 to 29 °C. Watch the compressor inlet strip. The coolant demand runs to 100 %, the inlet still settles at 29.80 °C, and net power falls from 928.9 to 597.4 kW, while every other controlled variable stays where it was.

The comparison worth making is against Section 10.6.4. The entire three-layer hierarchy, designed and tuned carefully, raises the reachable output from the 959.5 kW at which the turbine inlet alone stops to 1189.9 kW, which is about 230 kW of additional dispatch range. An afternoon costs 332 kW. On a real plant that argument is the case for evaporative or hybrid cooling, and it is an argument the control engineer is usually not in the room for.

10.6.6 Fouling, and a Controller That Hides It

Scenario 10 fouls the recuperator in two steps with the full hierarchy in service and the dispatch unchanged.

Table 10.9 is a well-behaved control system in front of a plant losing an eighth of its efficiency. Net power stays between 927.2 and 931.3 kW for four hours while the heater duty rises 14.7 % and the thermal efficiency falls 5.2 points, which is 12.8 % of the efficiency the plant started with. No

Table 10.9 Recuperator fouling at constant 930 kW dispatch. The controlled variable stays within 3 kW of its setpoint.

t (min)	ε	Recuperated	Heater (kW)	Net (kW)	η (%)
19	0.925	47.9 %	2298	930.4	40.48
60	0.869	44.8 %	2418	930.1	38.46
180	0.773	39.2 %	2638	930.0	35.25
240	0.773	39.1 %	2635	930.0	35.29

Table 10.10 The graded seven hour dispatch day, read from the application's own judgment at 420 min. The recommended configuration is the suggested setting described in the text, and each row after the second changes one element of it before Run.

Configuration	Out of band (min)	Over 620 °C (min)	Peak P_2 (MPa)	Bypass (kWh)
As loaded	420.0	0.0	20.58	0
Recommended	61.7	12.6	23.80	280
no third layer	339.1	40.2	20.58	127
$T_4^{\mathrm{ref}} = 593\,^{\circ}\mathrm{C}$	120.3	30.3	23.80	233
$b^{\mathrm{ref}} = 0\,\%$	116.3	9.7	23.62	132
$b^{\mathrm{ref}} = 8\,\%$	**17.6**	13.3	23.80	**509**

process alarm fires, because no measured variable has moved. The only warning is a calculated one. The application flags the recuperator as degraded when its computed effectiveness falls below 0.80, which happens for 0.8 min just after the first step, before the wall temperature adjusts to the lower conductance, and for good from the second step at 120 min. The first step therefore costs two points of efficiency with no lasting indication at all. The signal is in the ratio of heater duty to net power, which is a calculation nobody is required to look at, and the application draws it for that reason.

10.6.7 The Graded Dispatch Day

The challenge is seven hours of dispatch against a warming afternoon, a turbine efficiency loss and an evening peak, on a plant whose recuperator is already fifteen percent down. The target requires the plant to stay within ±25 kW for all but 80 min, spend no more than 20 min above the 620 °C metal limit, keep the discharge below 24 MPa and discard less than 400 kWh through the bypass. The suggested settings, which the recommended configuration of Table 10.10 uses, are the Cascade with the bypass and speed paths, the tunings of Table 10.6, a 3 % resting bypass and a 580 °C turbine inlet reference, with the compressor inlet setpoint and the bypass station left at the case's 20.5 °C and 0 %.

The recommended configuration passes all four clauses, inside the four limits by 18.3 min, 7.4 min, 0.20 MPa and 120 kWh, and the four ablations fail for four different reasons. As loaded, with the inner loops holding their own setpoints, weak tunings and nobody watching the power, the plant is out of band for the whole day, although the other three clauses hold. Removing the third layer fails two clauses at once, because without a route to discharge pressure the plant answers every dispatch increase with turbine inlet temperature, spends 40.2 min above the metal limit and is still out of band for 339.1 min. Leaving the turbine inlet reference at its design 593 °C rather than lowering it to 580 °C fails the same two clauses more narrowly, because the headroom the evening peak needs was spent in the morning.

The last two rows are the same change in opposite directions and they are the point of the scorecard. The resting bypass position is the amount of throughput held permanently in reserve so that the fast path has authority in both directions. Setting it to zero makes the plant maximally efficient and removes its ability to answer a load increase quickly, which costs 116.3 min out of band. Setting it to eight percent gives the best tracking on the sheet, 17.6 min, and discards 509 kWh, which fails the energy clause. Three percent passes both. The resting position is not a tuning parameter but

Table 10.11 The MPC against the cascade with the bypass and speed paths, shipped seed, noise off, full length of each case

	Power IAE (kW min)				Out of band (min)				Bypassed (kWh)				Over 620 °C (min)				MPC
Case	S	R	H	M	S	R	H	M	S	R	H	M	S	R	H	M	plans
5	943	86	78	90	11.9	0.6	0.5	1.1	147.1	0.8	156.0	0.6	0	0	0	0	600
6	42	35	7	4	0.0	0.0	0.0	0.0	79.4	0.3	83.1	0.1	0	0	0	0	300
7	180	80	30	40	1.2	0.1	0.1	0.4	152.5	0.7	146.9	0.6	0	0	0	0	600
8	24625	23075	22923	23310	309.1	281.8	281.2	281.8	115.5	1.4	144.6	0.1	15.4	11.7	11.4	0	1000
9[a]	15206	1711	592	542	136.2	23.9	4.0	3.1	116.3	16.2	133.2	12.9	15.1	0	0	0	480
10	87	41	12	46	0.0	0.0	0.0	0.0	224.9	0.4	222.3	0.1	0	0	0	0	800
11	8425	1050	633	781	101.1	9.6	8.1	10.3	143.7	8.2	152.4	2.5	12.0	3.4	0	0	510
12, suggested[b]	5098	1581	937	1482	61.7	8.6	6.4	14.7	280.2	15.2	240.3	2.5	12.6	0	0	0	840
12, as shipped[c]	66900	65709	26460	37429	420.0	375.1	284.6	309.8	47.0	37.5	228.0	72.2	0	0	0	0	840

S is the cascade at the case's own tunings, which scenarios 8 and 10 load, R the retuned cascade, H the hindsight setting and M the MPC. The MPC plans every 0.5 min, so its plans give each case's length, and none was rejected. [a]Scenario 9 loads the bypass path alone, which gives 5654 kW min, 82.3 min, 475.0 kWh and none over 620 °C. [b]The challenge with its suggested settings, the configuration recommended in Table 10.10. [c]The challenge with its own weak inner-loop tunings, which as loaded give 69 139 kW min and 420.0 min out of band.

a purchase of responsiveness with fuel, priced in this scenario at 3.8 kWh per minute of tracking recovered.

10.6.8 Model Predictive Control Against the Hierarchy

Table 10.11 sets the MPC beside three settings of the cascade with the bypass and speed paths on scenarios 5 to 11 and on the challenge, each run for its full length at the shipped seed with the noise off. All four read the same instruments and none is given a preview. The shipped setting uses each case's own tunings, the suggested ones on the challenge. The retuned setting rests the bypass at 0 % with fast valve position controllers ($K_c = 10\,°\text{C}/\%$ and $\tau_I = 10$ min for the first, 0.3 MPa/°C and 5 min for the second), and the hindsight setting is the least power error among 288 combinations of resting position, inlet reference and five tuning constants, tried on each case alone. Both stronger settings were therefore picked in hindsight, while the MPC keeps scenario 11's settings on every case. Integral error is the scorecard's, read from the logged rows.

Scenario 11 runs the tuned plant of scenario 8 through a dispatch day of 930, 820, 1050, 1180 and 930 kW, with steps at 15, 75, 135 and 195 min. Under the MPC it leaves 781 kW min of power error against 8425 for the shipped cascade, sends 2.5 kWh through the bypass against 143.7, and never passes 620 °C, where the shipped cascade spends 12.0 min above it. Most of that gain is not the optimizer. The retuned cascade, which moves the pressure within minutes and rests the bypass shut, leaves 1050 kW min and 8.2 kWh, with fewer minutes out of the band than the MPC (9.6 against 10.3). The third handle carries the plan's result, since with its pressure setpoint frozen the same controller leaves 21 050 kW min on this day, 112 220 on scenario 8 and 49 447 on the challenge.

What separates the plan from the retuned cascade is the turbine inlet (Figure 10.4). After the step to 1050 kW at 75 min both put the pressure setpoint on its 23.8 MPa clamp and are in the band for good at 77.1 min, the plan with its inlet setpoint raised only to 597.7 °C and the retuned cascade with that setpoint on its 618 °C clamp. The shipped cascade raises its pressure setpoint by 1.6 MPa in the first 20 min, overshoots to 625.7 °C at 95.7 min and stays out of the band until 113 min. After the step to 1180 kW at 135 min the plan and the retuned cascade both put the inlet setpoint on its clamp. By 142.5 min, with the inlet at 601.0 °C, the plan has begun to cut that setpoint, which stands at 600.8 °C by 145 min, ahead of the overshoot its second-order model predicts, and the inlet peaks at 613.9 °C at 150.5 min. The retuned cascade holds its clamp until 145 min, and its inlet reaches 620.0 °C at 152.7 min and 620.6 °C at 181 min after the setpoint has returned to the clamp, 3.4 min above the limit in all. The metal row is never active on this day, and a copy of the MPC whose inlet model is first order with the same 63 % time crosses 620 °C for 5.2 min, with a peak of 621.1 °C.

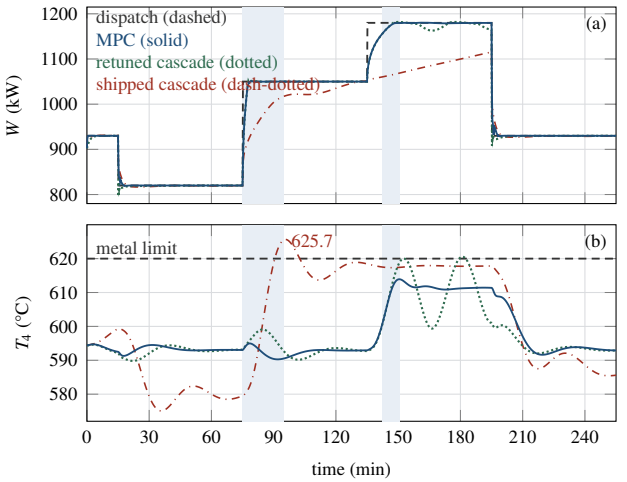

Figure 10.4 Scenario 11 run by the MPC (solid), by the shipped cascade with the bypass and speed paths (dash-dotted), and by the retuned cascade (dotted). (a) Net power against the dispatch (dashed). (b) Turbine inlet temperature against the 620 °C metal limit (dashed). The shaded windows, 75 to 95 and 142.5 to 151 min, follow the steps to 1050 and 1180 kW.

Scenario 8's request for 1270 kW is beyond every structure. The MPC spends no time above 620 °C (peak 618.3 °C), against 15.4, 11.7 and 11.4 min for the shipped, retuned and hindsight settings, at power errors within 8 % of each other. Either the clipped target or the metal row alone keeps the inlet under the limit, with peaks of 618.3 and 618.1 °C without the other, and without both the MPC spends the same 11.7 min above it as the retuned cascade.

In the load rejection of scenario 9 the plan and the shipped fast path both act on the scan the breaker opens. The MPC is in the band for good 0.86 min after the step to 450 kW, against 4.88 min for the bypass path the scenario loads, 3.42 for the shipped cascade with the speed path and 23.7 min for the retuned one, and its integral error is 542 kW min against 5654 for the bypass path. The price costs it the first scan, in which it sheds 24 kW less than the bypass path (729.0 kW at 10.1 min against 705.5), because it opens the bypass to 32.6 % rather than 43.7 and counts on the pressure setpoint it cuts to 14.8 MPa on the same scan. With $e^* = 0$ it reaches 708.0 kW on that scan, and at 10.2 min it already sheds more, at 624.4 kW against 637.9.

The MPC loses three comparisons, each with a mechanism. The hindsight setting has less power error on every case except scenarios 6 and 9 (scenario 11: 633 against 781 kW min), but each such setting rests the bypass at 3 % and throws 144.6 to 240.3 kWh away, against 0.1 to 2.5 kWh for the MPC on the same cases with tuned inner loops. The retuned cascade has less error on scenarios 5 and 10 (86 against 90 and 41 against 46 kW min) and fewer minutes out of the band on the challenge (8.6 against 14.7). When the coolant warms at 70 min, the plan, which holds the measured compressor inlet constant over its horizon, puts the pressure setpoint on its clamp by 73 min but raises the inlet setpoint only from 75 min, and it is out of the band for 6.5 min where the retuned cascade, moving both at once, is out for 1.6. The third loss is the first scan of the load rejection described above.

The plan is made on the inner loops of its identification and is only as good as that assumption. With the challenge's own weak inner-loop tunings it is out of the band for 309.8 min of 420, against 66.5 min once the pressure loop alone takes the identified tuning, 44.2 min with both inner loops and 7.5 min with the compressor inlet loop as well, and on scenario 11 doubling the inlet loop's gain takes the inlet over 620 °C for 7.2 min. Under those weak loops every structure is slow, and scenario 11 leaves 20 541, 22 279 and 20 459 kW min under the MPC and the shipped and retuned cascades.

On the challenge, the suggested settings with the supervisory cascade replaced by the MPC meet

the target, at 14.7 min out of the band, no time above 620 °C, a peak discharge of 23.800 MPa and 2.5 kWh bypassed, against 61.7, 12.6, 23.80 and 280 for the recommended cascade of Table 10.10 and 8.6 min out of the band for the retuned one. The run breaks none of the limits the rule leaves unscored. The turbine inlet peaks at 606.6 °C, the compressor inlet stays at least 7.11 °C under its alarm, and the recuperator effectiveness stays at or above 84.2 %, clear of the 0.80 at which the application flags it as degraded.

10.7 Technology in Practice

A real supercritical carbon dioxide loop reads and moves the quantities of Section 10.5 with hardware that has its own limits. Figure 10.5 draws the seven loops with plant devices in place of the schematic's symbols.

What Measures It

On the test loop of the history segments, resistance thermometers and thermocouples sat in the carbon dioxide at the inlet and exit of every component, pressure transducers at the inlet and outlet of every component, and an orifice flow meter after the compressor outlet [10]. The model keeps four of these readings as first-order lags of 0.05 to 0.35 min. The reading that matters most is the one at the compressor inlet.

> **TECHNOLOGY Sensor: A Resistance Thermometer at the Compressor Inlet**
>
> A resistance temperature detector (RTD) holds a sensing element whose electrical resistance changes with temperature. The test loop of the first history segment used RTDs with an accuracy of about 0.2 K beside thermocouples of about 1 K, and warned that heat conducted along an inserted probe can spoil the reading in small ducts. Its report judged it very difficult to know the real fluid temperature near the critical point to within a few tenths of a degree [10]. In a thermowell the response depends on heat transfer, which changes strongly with density there. An aging RTD drifts in calibration or answers more slowly, and a later study modeled the fault as a time constant growing from 2 s to 6 s [3]. The model's compressor inlet thermocouple lags 0.25 min.

What Moves It

In the reactor study of the second history segment, the settings of the slower handle were chosen so that the turbine bypass valve stays around 10 % open [16]. That is in effect the resting position of Section 10.6.7, which holds throughput in reserve so that the fast path has authority in both directions. On the cold side, the test loop moved its compressor inlet temperature with the speed of the cooling water and with a valve that recirculated warm water, against an evaporative cooler whose supply varied from about 278 to 298 K between winter and summer [10]. Pilot-plant studies keep the cooling water leaving the cooler below 50 °C to prevent scale and pump warm water back into the entering stream [19]. The model's coolant valve saturates between a supply of 19 and 22 °C (Table 10.5).

> **TECHNOLOGY Final Element: A High-Speed Motor Drive**
>
> On a test loop the compressor speed handle is the drive of a high-speed permanent-magnet motor-alternator whose stator sits in the carbon dioxide itself. The loop of the first history segment used a variable-speed controller with insulated-gate bipolar transistors and pulse-width modulation, built to motor at 50 kW and to generate up to 150 kW, and it ran the shaft to 65 000 rpm. A coil embedded in the stator replaced a speed proximity sensor that proved too noisy. Starting the motor, keeping the speed signal clean, deciding when to shift to speed control and handling internally sensed faults took considerable work, and the rotor had to accelerate quickly through its first two vibration modes at 17 000 and 32 000 rpm [10]. In the model, discharge pressure follows a speed step with a 63 percent time of 0.35 min.

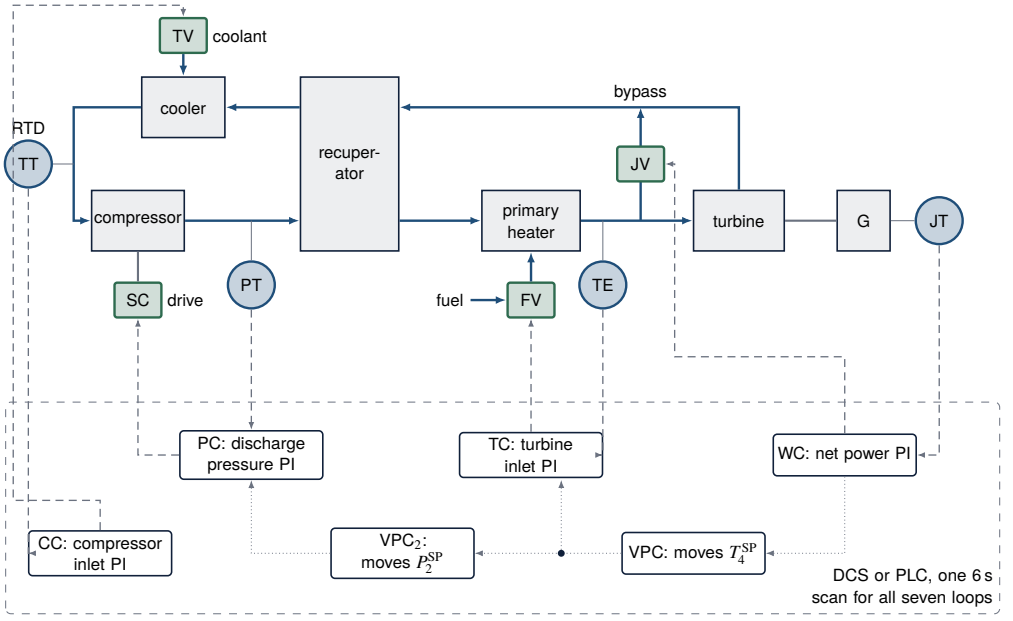

Figure 10.5 The loops of Figure 10.1 on plant hardware. A resistance thermometer (TT, RTD) at the compressor inlet feeds the inlet PI (CC), which moves the coolant valve (TV). The discharge pressure transmitter (PT) feeds the pressure PI (PC), which sets the motor drive (SC) of the compressor. The turbine inlet thermocouple (TE) feeds the firing PI (TC) and its fuel valve (FV), and the power transducer (JT) on the generator (G) feeds the power PI (WC), which moves the turbine bypass valve (JV). Dotted lines are the ladder: the valve position controller (VPC) reads the bypass loop's demand before the clamp and sets T_4^{SP}, and the second (VPC_2) reads that setpoint and sets P_2^{SP}. Notice that the coolant loop, the most valuable on the plant, has no link to the ladder.

Where the Controller Runs

Control loops are often nested, with the setpoint of one loop based on the variable of another, and supervisory and lower-level loops run with cycle times from milliseconds to minutes [20]. The test loop ran on-off and PID loops for its charge, its cooling water flow and its warm-water recirculation from a data acquisition and control system that sampled about four times a second [10]. The application's seven controllers share a 6 s scan. Each valve position controller is an ordinary PID block whose measurement is the output of another controller, which is why the block must be able to read the demand before the clamp.

> **IN PRACTICE** **Showing the Operator Where the Critical Point Is**
>
> The test loop's operators watched more than trends. Its control screen plotted the measured compressor inlet, compressor outlet and chiller states on live temperature-entropy and temperature-density diagrams, so they could see where the loop ran relative to the saturation dome and the critical point, and another tab drew the inlet on the compressor map for the last 100 recorded time steps. The aim was an inlet within about 1 to 2 °C and 100 to 400 kPa of the critical point. A Coriolis meter at the compressor inlet, which reads density directly, proved a valuable control signal while carbon dioxide was pumped in or vented [10].

What Is Modernizing It

The first development is predictive control. A study for a reactor power converter found that one predictive controller over several inputs can coordinate the components of the cycle, where separate PID loops need individual tuning and other measures against coupling [3]. The second is the move

from test loops to pilot plants. A 2017 roadmap planned a pilot stage operated off grid by scientists, then an on-grid demonstration operated by industry, and it listed dry cooling among the motivations for small reactor, concentrating solar and natural gas applications [21]. Dry cooling makes the warm afternoon of Section 10.6.5 the usual case. The third is dynamic modeling for control design before a plant runs: pilot-plant studies built dynamic models of a 10 MW recompression cycle for load changes, startup, shutdown and trips, and for the control of its cooler [22].

10.8 Limitations

There is no recompression branch. A real high-efficiency sCO_2 cycle splits the flow before the cooler and recompresses part of it to balance the recuperator, and the split fraction is the main efficiency handle at part load [14]. Its absence removes an actuator that would change several of the conclusions above, and the load-following literature for recompression cycles treats it as central [13]. The application is a simple recuperated cycle and should be read as one.

Properties are constant specific heats per segment. Real carbon dioxide near the critical point has a specific heat that varies by an order of magnitude across the cooler and the low-temperature end of the recuperator, which is precisely where the constant-c_p assumption is worst. The near-critical behavior that matters for control is carried by (10.1) and (10.4) as fitted functions, so the shapes are right and the numbers are not transferable to a specific machine.

The heat exchangers are single lumped walls. A printed circuit recuperator has a distributed temperature profile, a pinch that can move along its length and transient behavior that a one-wall model cannot reproduce [12]. The application reports a pinch and that number should be read as an indicator rather than a design quantity.

Inventory control is a slider rather than a loop. Adding or removing carbon dioxide from the loop is a standard part-load strategy on real plants and is slower than everything in Table 10.3, so it would be a fourth and slowest layer of the same ladder. Turbomachinery is quasi-steady with no shaft dynamics, no surge line and no windage, so the application cannot represent the transients an overspeed event actually produces. The heat source is a first-order lag toward a temperature proportional to firing rate, which suits a fired heater and does not represent a nuclear or solar source with its own dynamics [11]. A solar field with a transit lag of its own is the plant of Volume 17, Chapter 6. The compressor inlet pressure is held constant by assumption, so the loop cannot show the pressure excursion a real inventory responds to when the loop temperature changes. The four-minute fired-heater lag and the inventory parameter are both deliberate simplifications, so that the mid-ranging argument is made with three actuators rather than five.

The predictive controller's model is the inner loops at one tuning and a power map relinearized at each decision. It holds the measured compressor inlet constant over the horizon, which costs it the warm afternoon of the challenge, and it carries no row for that inlet, which no compared run brought within 4.25 °C of its alarm. A detuned inner loop changes dynamics that the output biases cannot correct, the last block holds the commands for the final 7.5 min of a plan, and the shared solver allots a single slack to the whole horizon instead of one to each step.

OPEN CHALLENGE Holding the Compressor Inlet Above the Critical Point

Control studies for a 10 MW pilot plant state the problem plainly. Its controllers must hold the main compressor inlet at a design 33.5 °C with a lower constraint of 32 °C, keep the bypass compressor at least 10 % above its surge limit, and allow for a cooler outlet change of 1 °C that, left uncontrolled, can move the turbine inlet by 10 °C [22]. Proportional-integral (PI) controllers have fixed gains that cannot follow the steep property changes, and sensor delays add to the difficulty. A predictive controller kept a simulated pre-cooler stable where a PI controller with a degraded sensor did not [3]. Progress would be such a controller proven on a running plant, through startup, shutdown and trips.

10.9 Discussion

Three results generalize past this plant.

The first is that mid-ranging is a statement about cost as much as about speed. The textbook presentation pairs a fast small actuator with a slow large one and recovers the fast one's range. What Tables 10.7 and 10.8 add is that the ordering by expense need not agree with the ordering by bandwidth. The bypass is the fastest and the most expensive. The firing rate is the slowest and moderately expensive. Discharge pressure is intermediate in speed and free, which is why it belongs at the bottom of the ladder and why leaving it out costs both the tracking and the metal.

The second is that a saturating actuator must be handed upward by its unsaturated demand. A clamped output carries no information about how far past the clamp the controller wanted to go, and a supervisory layer reading the clamped value inherits a plant whose gain is zero on one side of the resting position. This is a small implementation detail with a large and asymmetric consequence, and it is not specific to valve position control, because the same requirement appears in override selectors, in split-range schemes and anywhere a signal is passed between controllers [6]. Section 1.4.4 treats the companion question for an override, which signal a controller left out by a selector should track.

The third is that a hierarchy reorders constraints rather than removing them, and a plant should be run until they all bind so that the operator knows which one arrives first. Section 10.6.4 reaches the pressure rating and the metal limit within a few minutes of each other, which is a well-designed plant and also a plant with no margin left. Knowing that is worth more than another point of tracking performance [8]. The corollary is that a variable held perfectly is not evidence of a healthy plant, because the fouling of Table 10.9 costs an eighth of the efficiency behind a power trend that stays within 3 kW.

10.10 Conclusion

A supercritical carbon dioxide power block gives the control engineer one controlled variable and three actuators whose open-loop response times span 0.14 min to 17.0 min and whose marginal costs run from zero to everything. A dispatch increase of 120 kW is not slow but unreachable for the two obvious architectures, because both are stopped by the 618 °C clamp on the turbine inlet setpoint at 959.5 kW. Adding a third mid-ranging layer that converts turbine inlet temperature back into compressor discharge pressure reaches 1050.0 kW in 6.8 min with an integral error smaller by a factor of 37, a higher efficiency and 3.0 °C of metal margin in hand, and it does so with one additional PID and one careful decision about which signal to hand upward. The same layer is the fastest and the cheapest on the way down, discarding 107.0 kWh against 242.9 kWh for the two-layer structure. Behind all of it sits a compressor inlet loop that is not part of the hierarchy, cannot be improved by better tuning, and whose saturation on a warm afternoon costs 332 kW, about one and a half times the 230 kW of range the whole hierarchy adds. Offered beside the hierarchy, a model predictive controller over the same three handles leaves 781 kW min of power error on the dispatch day of scenario 11 against 8425 for the shipped hierarchy and 1050 for a retuned one, and what it keeps for itself, a turbine inlet held under its limit by the model of the inner loop and a load rejection settled in 0.86 min, depends on inner loops tuned as that model assumes.

KEY POINTS

- One power output has three handles: a fast bypass that throws work away, a free compressor speed and a slow firing rate that buys fuel.
- On a 120 kW dispatch increase the two obvious architectures never arrive, because the 618 °C clamp on the turbine inlet setpoint stops them at 959.5 kW.
- A third layer that moves the discharge pressure reaches 1050.0 kW in 6.8 min, with 37 times less integral error.
- A saturating actuator must be handed upward by its unsaturated demand.
- A warm afternoon costs 332 kW through the compressor inlet loop, more than the 230 kW of range the whole hierarchy adds.

Where else this applies. Chapter 9 returns a fast electric chiller to its reserve with a valve-position controller on chilled water, and Volume 17, Chapter 4 hands one variable between two actuators on a wind turbine.

References

[1] Y. Ahn, S. J. Bae, M. Kim, S. K. Cho, S. Baik, J. I. Lee, and J. E. Cha, "Review of supercritical CO_2 power cycle technology and current status of research and development," *Nuclear Engineering and Technology*, vol. 47, no. 6, pp. 647–661, 2015.

[2] F. Crespi, G. Gavagnin, D. Sánchez, and G. S. Martínez, "Supercritical carbon dioxide cycles for power generation: A review," *Applied Energy*, vol. 195, pp. 152–183, 2017.

[3] R. Ponciroli and R. Vilim, "Analyses of model-based predictive control for a S-CO_2 Brayton cycle power converter," Argonne National Laboratory, Argonne, IL, USA, Tech. Rep., 2016.

[4] T. Hägglund, "A feedforward approach to mid-ranging control," *Control Engineering Practice*, vol. 108, p. 104713, 2021.

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

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

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

[9] R. Span and W. Wagner, "A new equation of state for carbon dioxide covering the fluid region from the triple-point temperature to 1100 K at pressures up to 800 MPa," *Journal of Physical and Chemical Reference Data*, vol. 25, no. 6, pp. 1509–1596, 1996.

[10] S. A. Wright, R. F. Radel, M. E. Vernon, P. S. Pickard, and G. E. Rochau, "Operation and analysis of a supercritical CO_2 Brayton cycle," Sandia National Laboratories, Albuquerque, NM, USA, Tech. Rep. SAND2010-0171, 2010.

[11] G. S. Delsoto, F. G. Battisti, and A. K. da Silva, "Dynamic modeling and control of a solar-powered Brayton cycle using supercritical CO_2 and optimization of its thermal energy storage," *Renewable Energy*, vol. 206, pp. 336–356, 2023.

[12] M. Marchionni, L. Chai, G. Bianchi, and S. A. Tassou, "Numerical modelling and transient analysis of a printed circuit heat exchanger used as recuperator for supercritical CO_2 heat to power conversion systems," *Applied Thermal Engineering*, vol. 161, p. 114190, 2019.

[13] H. Wang, J. Chen, G. Yu, and M. Liu, "Research of wide-range load following control strategy for sCO_2 recompression Brayton cycle system," *Energy Conversion and Management*, vol. 349, p. 120870, 2026.

[14] F. Correa, R. Barraza, Y. C. Soo Too, R. Vasquez Padilla, and J. M. Cardemil, "Optimized operation of recompression sCO_2 Brayton cycle based on adjustable recompression fraction under variable conditions," *Energy*, vol. 227, p. 120334, 2021.

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

[16] A. Moisseytsev and J. J. Sienicki, "Development of the ANL plant dynamics code and control strategies for the supercritical carbon dioxide Brayton cycle and code validation with data from the Sandia small-scale supercritical carbon dioxide Brayton cycle test loop," Argonne National Laboratory, Argonne, IL, USA, Tech. Rep. ANL-ARC-218, 2011.

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

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

[19] J. Albright, E. Liese, S. Zitney, P. Mahapatra, and D. Bhattacharyya, "Load following for a supercritical CO_2 Brayton cycle," Presentation, Annual AIChE Meeting, Minneapolis, MN, National Energy Technology Laboratory, NETL-PUB-21179, OSTI 1409315, 2017.

[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] C. M. Mendez, G. Rochau, and M. Wilson, "Technology roadmap and systems engineering model for sCO_2 power cycle development," Poster, Sandia National Laboratories, SAND2017-12272C, OSTI 1482300, 2017.

[22] E. A. Liese, P. Mahapatra, J. T. Albright, and S. E. Zitney, "Regulatory control of a 10 MWe supercritical CO_2 recompression closed Brayton cycle," Presentation, 2019 Annual Review Meeting for Crosscutting Research, Pittsburgh, PA, National Energy Technology Laboratory, NETL-PUB-22283, OSTI 1623322, 2019.

10.11 Exercises

10.1 Concept check · 2 min

In Manual the application offers four handles on the cycle. Which of the following is *not* one of them, so that adding it would change several conclusions of this chapter?

(a) The recompression flow split fraction
(b) The compressor speed
(c) The coolant valve demand
(d) The turbine bypass position

10.2 Concept check · 2 min

The plant is 100 kW short of dispatch and the bypass is clamped shut at 0 %. The bypass position controller of Equation (10.20) reads b^{raw}, the bypass loop's demand before the clamp. What does that signal carry that the applied position does not?

(a) The resting position, which the clamp removes from the applied position
(b) The size of the shortfall, since the demand settles near $K_c(W - W^{\mathrm{SP}})$, about −9 %
(c) The integral of the shortfall, which grows without limit while the clamp holds
(d) The power thrown away through the bypass, which the shut valve reports as zero

10.3 Concept check · 1 min

Scenario 10 fouls the recuperator in two steps while every loop holds its setpoint. Which of these quantities does an instrument on the plant measure?

(a) The recuperator effectiveness
(b) The cycle thermal efficiency
(c) The net electrical power
(d) The carbon dioxide mass flow

10.4 Theory and calculation · 5 min

The compressor inlet is held at 20.5 °C, below the near-critical threshold T_c, and the speed is raised from 85 to 92 %. Use Equations (10.1) to (10.4) and Table 10.2.

(a) Compute the pressure ratio and the discharge pressure P_2 at 92 % speed. *Checkpoint:* 3.378 and 22.90 MPa

(b) Compute the specific compression work w_c at 85 and 92 %. *Checkpoint:* 20.84 kJ/kg and 24.34 kJ/kg

(c) Compute the mass flow at 92 % and the compressor power $\dot{m}w_c/\eta_m$ at both speeds. Why does net power still rise with speed? *Checkpoint:* 5.95 kg/s; 120.7 kW and 152.5 kW

10.5 Theory and calculation · 6 min

On a warm afternoon the coolant valve is saturated and the compressor inlet settles at 29.80 °C with the speed at 85 %, while the inner loops hold the turbine inlet and the discharge pressure at their design values (scenario 3). Use Equations (10.1), (10.2), (10.4) and (10.15) with Table 10.2.

(a) Compute the penalty u, the factor in square brackets of Equation (10.1) and the specific compression work. *Checkpoint:* $u = 0.955$, factor 1.782, $w_c = 37.13\,\mathrm{kJ/kg}$

(b) Compute the mass flow and the compressor power $\dot{m}w_c/\eta_m$, and compare it with the 120.7 kW it takes at design. *Checkpoint:* 4.40 kg/s, 172.0 kW

(c) The turbine specific work stays at 201.7 kJ/kg, and the exported file shows the coolant pump power rising from 15.5 to 82.7 kW. Split the fall in net power into its three parts and compare the sum with Table 10.5. *Checkpoint:* 213.0 kW, 51.3 kW and 67.2 kW, 331.5 kW in all

10.6 Theory and calculation · 7 min

The bypass loop has $K_c = 0.09\,\%$ per kW and the bypass position controller $K_c = 3.0\,^\circ\mathrm{C}/\%$, $\tau_I = 30\,\mathrm{min}$, action $a = +1$ and resting position $b^{\mathrm{ref}} = 3\,\%$ (Table 10.6). The bypass is held shut against a steady power shortfall.

(a) Using the steady state of Equation (10.18) with $T_t = \tau_I$, find the bypass demand c for shortfalls of 5, 100 and 500 kW. *Checkpoint:* –0.45 %, –9 % and –45 %

(b) Find the error of the position controller and the rate at which its integral action raises T_4^{SP} for each shortfall, first when it reads the demand and then when it reads the applied position. *Checkpoint:* 3.45, 12 and 48 %: 0.345, 1.2 and 4.8 °C/min; 3 % and 0.3 °C/min for all three

(c) At 100 kW short, how long does integral action alone take to carry T_4^{SP} from 593 to its 618 °C clamp in each case? *Checkpoint:* 20.8 min against 83.3 min

10.7 App challenge · 8 min

Load scenario 3, *The cheapest loop on the plant*, with its own settings and the default seed. Run to 140 min at speed 20 (twenty simulated minutes a second), then export the CSV.

(a) From the `cool_cmd_pct` column, find when the coolant demand first reaches 100 %. Just before 140 min, read `T1_C`, `power_kW` and the Efficiency tile. *Checkpoint:* 22.0 min (±0.2 min); 29.80 °C, 597.4 kW (±0.5 kW), 35.33 %

(b) At the same time read `mdot_kgs`, `w_comp`, `T4_C` and `P2_MPa`. Which controlled variables are still on setpoint, and why does no loop recover the lost power? *Checkpoint:* 4.40 kg/s, 37.13 kJ/kg, 593.0 °C, 20.58 MPa

10.8 App challenge · 8 min

Load scenario 8, *Three deep, and the rating that ends it*, with its own settings and the default seed. Run to 400 min at speed 20, then export the CSV.

(a) At 180 min, read the net power, the bypass position, the discharge pressure, the compressor speed and the turbine inlet temperature. Which actuator now carries the extra 150 kW? *Checkpoint:* 1080.0 kW, 4.3 %, 23.26 MPa, 93.1 %, 607.8 °C (±0.2 °C)

(b) From `metric_overtemp_min` at 219.9 and 400 min and `T4_C` after 220 min, find the minutes above 620 °C after the step to 1270 kW and the peak. *Checkpoint:* 7.2 min (±0.1 min), 620.8 °C

(c) Read the net power at 400 min and name the three stops that bind. *Checkpoint:* 1189.9 kW, 80.1 kW short

10.9 App challenge · 9 min

Load the challenge, *A dispatch day*, with the default seed and press Reveal suggested tuning. Run the full 420 min at speed 20 and export the CSV.

(a) Read the Off dispatch tile. In the last row of the file, read the minutes over 620 °C in `metric_overtemp_min` and the bypassed energy in the column `metric_bypass_loss_kWh`, then find the largest `P2_MPa`. Does the run pass, and by how much on each clause? *Checkpoint:* 61.7 min, 12.6 min, 280 kWh (±1 kWh), 23.80 MPa

(b) Reload, press Reveal suggested tuning again, set Bypass resting position (Advanced) to 8 % and repeat the run. Which clause fails? *Checkpoint:* 17.6 min, 13.3 min, 509 kWh (±1 kWh), 23.80 MPa

(c) Table 10.10 gives 116.3 min and 132 kWh for a 0 % resting position. What does a minute of tracking cost in bypassed energy between 0 and 8 %? *Checkpoint:* 3.8 kWh per minute

10.10 Application area · 3 min

The rating of a supercritical carbon dioxide plant is quoted with a cooling temperature attached.

(a) From Table 10.5, what share of the output does a coolant supply of 29 °C remove, and how much energy is lost over a six-hour afternoon at that temperature? *Checkpoint:* 35.7 %, about 1990 kWh

(b) Why can the control system not recover this loss, and what does the chapter name as the remedy?

10.11 Application area · 4 min

Table 10.9 shows the recuperator fouling while net power stays on setpoint.

(a) Compute the ratio of heater duty to net power at 19 and 240 min, and the extra heat the heater supplies. *Checkpoint:* 2.47 and 2.83; 337 kW

(b) If heat costs $0.02 per kWh and the plant runs all year (8760 h) with the fouled recuperator, what does the extra heat cost? What would you put on the operator's screen? *Checkpoint:* about $59 000 a year

10.12 Review the agent · 5 min

A student asks an AI assistant how to make the three-layer hierarchy more robust. The assistant replies: "The bypass position controller should read the applied bypass position b, not the bypass loop's internal demand. The demand can go negative, which has no physical meaning, while the applied position is what the valve really does. The change costs nothing on a load decrease and removes a source of nonsense values from the cascade." Find the flaw and say how you would check it.

CHAPTER 11

Holding a Tunnel Face With Its Own Muck Through Sixty-Six Stops

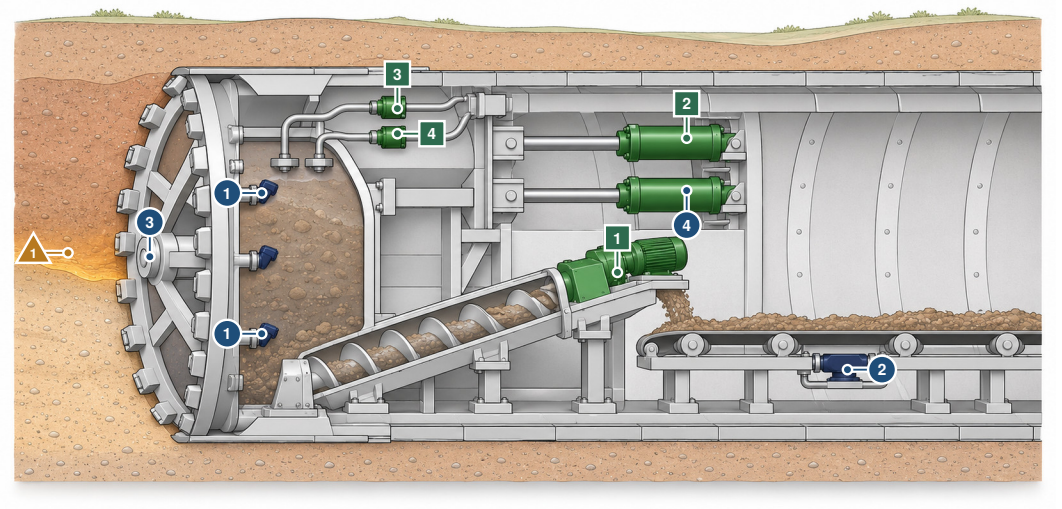

The controller moves 1 screw speed, which lets the soil out 2 push cylinders, which move the shield forward 3 foam mixed into the soil 4 clay slurry that holds the pressure while a ring is built

The controller reads 1 pressure cells in the chamber 2 belt scale, two minutes late 3 twisting load on the cutting wheel 4 push force of the cylinders

What disturbs it 1 a change from clay to sand at the face

An earth pressure balance shield holds the face with its own muck: notice that the screw lets soil out while the push cylinders, slowed by the wheel's load and their own force, bring it in. Illustration, not to scale.

IN PLAIN TERMS **Holding Up the Ground With the Dug Soil**

The main idea. A tunneling machine pushes a turning cutting wheel through soft ground. The loosened soil packs a closed chamber behind the wheel, and its pressure holds up the ground in front while a screw lets soil out as fast as it comes in.

The challenge. The machine stops for twenty minutes to line each short stretch of tunnel. The soil, a river above and buried boulders change what is safe, and the street sinks hours later, too late to act.

Why it matters. Separate control layers, one for each cause, keep the street level while the machine digs as fast as the strain on its wheel allows.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact in the Introduction (Section 11.1), the two history segments (Section 11.2), the try-it segment on the ring build (Section 11.4), the try-it segment and the fact on the drive (Section 11.6) and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 11.6), Technology in Practice (Section 11.8), Limitations (Section 11.9) and the open challenge

The control question An earth pressure balance shield holds the tunnel face with the muck in its chamber, whose pressure integrates what the cutterhead brings in and the screw takes out, and every ring of lining stops both for twenty minutes. Which control layers hold the face through sixty-six forced stops, a change of soil and a river, when the street's settlement is reported only hours after it can be prevented?

Learning objectives

- Explain why the chamber pressure is an integrator whose gain is set by the foam, and compute that gain from the chamber's stiffness.
- Compute the face's target pressure and the blow-out limit from the ground's stresses, and explain why no fixed setpoint serves both land and river.
- Describe the layers that hold the face: the screw loop with its advance feedforward, the bentonite hold, the chainage feedforward and the advance supervisor with its mid-ranging limit.
- Run the shield's cases and the challenge, and tell which layer each failure calls for.
- Judge a settlement claim against a survey that reports late.

Prerequisites Sections 1.6.2 (velocity-form PID), 1.4.3 (selectors), 1.4.4 (external reset), 1.4.7 (mid-ranging) and 1.6.4 (scenarios and exported files). Section 1.6.3 is needed only for the predictive coordinator.

The app https://apmonitor.com/apps/tunnel-boring/

See also Volume 6, Chapter 3, where a traversing gauge reports the product late, and Volume 6, Chapter 7, where inventory loops pass flows around a recycle. Volume 3, Chapter 6 for feedforward from the cause of a disturbance, and Volume 3, Chapter 2 for matching the loop structure to where a disturbance enters.

11.1 Introduction

A soft-ground tunnel is held up at its face before it is lined. In an earth pressure balance shield the cutterhead feeds the excavated soil into a closed chamber behind it, the chamber is kept full, and the pressure of that muck must balance the earth pressure on the excavation face [1]. In the generic drive simulated here a shield 6.5 m in diameter advances at 40 mm/min and cuts 1.33 m^3 of clay a minute into a 30 m^3 chamber, and a screw conveyor turning 1 rpm slower than the speed that balances it takes out 0.26 m^3 a minute too little. The axis pressure cell then rises 0.168 bar in the first minute and keeps rising, and the crown reaches its total overburden, the blow-out limit of the model, 2.9 min after the mismatch began.

> **DID YOU KNOW? A Cutterhead Almost 15 Meters Across**
>
> An earth pressure balance machine built for a road tunnel had a cutterhead almost 15.2 m in diameter [2], more than twice the 6.5 m of the shield simulated in this chapter.

Several established problems meet in the chamber. The face must be held inside a band between collapse and blow-out [3]. The ground lost into the face appears at the surface as a trough whose shape was described empirically in 1969 and has been refined with field data since [4], and the buildings over the tunnel tolerate it only to a limit [5]. The muck leaves through a screw whose pressure gradient is set by the shear of the soil plug against the casing [6], and conditioning is what makes the muck suitable for pressure control and for extraction through the screw [7]. While a ring is built the screw stands still, and in permeable ground the chamber loses pressure through the face [8, 9]. A predictive controller for the chamber identified the screw speed and the advance speed as the major factors in its pressure [1]. The arrangement that results, a fast handle that holds a variable and a slow one that the operation wants at a limit, is the mid-ranging structure of process control [10] (Section 1.4.7), and the plant shares its inventory character with a level, a drum or a bin whose content is seen only through a pressure.

The simulator corrects a specific belief, that a face pressure held on its setpoint means the street is safe. The pressure is an integrator whose gain depends on the foam, so the loop that is calm in clay

rings in sand. Once in every 1.5 m of advance the screw cannot act at all. The setpoint that keeps the street is fixed by ground conditions that change along the drive, and the settlement it produces is measured hours after anything can be done about it. In case 4 of the simulator the survey's first report of the point at 27.5 m reads 12.3 mm, under a 15 mm budget, while that point is heading for 18.8 mm.

The published work on this plant is divided among the disciplines that own its parts. Face stability, settlement prediction, screw mechanics, foam conditioning and standstill losses are geotechnical subjects [3, 4, 6, 8, 11], while the control work poses the chamber pressure as an optimal-control problem with face stability as its objective [12] or as predictive control on a model learned from field data [1]. This chapter brings the chamber's inventory, the foam's stiffness, the plug, the seepage, the ring sequence and the late survey together on one simulated drive, and measures which control layers the drive needs and why. The chamber pressure model, the standstill loss and two-handle chamber control are each established already, and the methods used here are standard ones. The simulator rests on a full-chamber model with a closed-form pressure and mass ledgers that close to 2×10^{-13} of the throughput, and on an instrument packet that is the only thing any controller is allowed to read. Its nine cases and a graded drive each carry a completion rule.

Section 11.2 follows the closed shield, face support and the settlement trough. Section 11.3 sets out the drive as modeled, from the ground's target pressure through the full chamber, the screw and its plug, the seepage and the trough to the instruments. Section 11.4 explains why the face is hard to hold: an integrator whose gain is the foam, a stop the screw cannot cover, a plug that can fail at any speed, and a target that moves while the product arrives late. Section 11.5 describes the control layers as implemented, with an optional predictive coordinator. Section 11.6 measures the cases, the drive and the predictive and classical comparisons, Section 11.7 separates the kinds of evidence, Section 11.8 sets the layers on a real shield's hardware, Section 11.9 states the limitations, Section 11.10 concludes, and Section 11.11 holds the exercises.

11.2 Closed Shields, Face Support and the Trough

A shield is a steel tube the diameter of the tunnel, pushed forward by cylinders that bear on the lining already built behind it. The reference text of the method covers open, compressed-air, slurry, earth pressure balance and combined shields [13], and of the closed-face types the slurry shield supports the face with a pressurized bentonite suspension while the earth pressure balance shield supports it with the excavated soil itself. A horizontal directional drill, the trenchless machine of Chapter 7, also crosses under roads and rivers and holds its mud pressure under a limit set by the cover. The two systems were compared for soft ground in 1991 [14]. A 2010 review of the development steps of soft-ground machines reports that they handle unstable faces against water pressures up to 15 bar [15], and the same reference text treats the planning of a drive against the geology, the groundwater and the settlement requirements together [13].

> ⧖ HISTORY **1818 to the 1970s: From an Open Shield to a Closed Face**
>
> Soft-ground tunneling did not become viable until the tunnel shield. Marc Brunel patented a circular shield in 1818, with the aim of opening no more ground than the shell of the tunnel would fill. Ground that needed more support was held with compressed air from the mid-1800s into the 1980s. From the late 1960s excavators were built into shields, but these digger shields lost ground ahead of the face, and keeping the right soil plug in the invert was always a headache. Japanese industry then developed pressurized-face machines and within a decade or so exported them to the United States. In the earth pressure balance machine, control is obtained by matching the volume the advance displaces with the volume the screw conveyor removes [2].

The band in which the face pressure must sit has two sides. Too little and the ground moves into the face, too much and the face lifts the ground above it, and the face stability conditions of an earth pressure balance shield were analyzed in 1996 [3]. What the street feels is the volume lost. How the trough it leaves came to be described is told in the history segment after this paragraph.

Analytical methods relate the loss to the gap around the tunnel and the soil [16], and a 2008 lecture on advances in earth pressure balance tunneling named face pressure, soil conditioning and effective screw conveyor operation as key factors in the volume loss [17].

> ⧖ HISTORY **1969 to 2024: Reading the Trough Over a Tunnel**
>
> The empirical method that describes the transverse trough by a Gaussian curve originated with Peck in 1969. O'Reilly and New gave its width in 1982 as $i = Kz_0$, with $K = 0.5$ for cohesive and 0.25 for granular soils, and Mair and Taylor found in 1997 that K for clays lies mostly between 0.4 and 0.6, and for sands and gravels between 0.25 and 0.45 with a mean of 0.35. A 2024 review that summarizes them back-calculated the volume loss of 78 troughs over earth pressure balance drives, from 0.03 % to 0.83 % of the face area with a median of 0.41 % [4].

Those factors are where the control problem lives. Model tests on a screw conveyor measured a casing shear stress constant along the conveyor and close to the undrained strength of the clay, so the total pressure falls linearly along the screw and its torque rises with the casing stress [6]. A laboratory screw that draws sand from a pressurized tank was built to verify the effects of foam on a noncohesive soil [7], and pressure infiltration tests found that the water permeability of foam-infiltrated sand falls as the foam injection ratio rises [11]. During a standstill for the lining the chamber pressure decreases in coarse soils, an effect linked to the permeability outside the machine that increases the settlement [8]. A simplified chamber pressure model for granular soil carries the gas volume, the liquid volume and the amount of gas as its states, through excavation and standstill, with muck compression and seepage as two of its physical processes [9]. On the control side, optimal control of the shield has been posed with face stability as its objective [12], and a predictive controller on a learned model identified the screw speed and the advance speed as the major factors in the chamber pressure [1]. The simulator of this chapter takes a layered arrangement, a pressure loop on the screw with the stops, the geology and the machine each handled by a separate layer.

11.3 The Drive as Modeled

The drive runs 100 m from a launch shaft at chainage 0 to a reception shaft, with the tunnel axis $z_0 = 15\,\text{m}$ under a street and 11.75 m of cover over the crown, as shown in Figure 11.1. The ground is a silty clay to the boundary at $x_b = 50\,\text{m}$ and sand after it, with a 1 m transition, and the water table stands $z_{wt} = 3\,\text{m}$ under the street. A building stands over 22 to 38 m. A river $h_r = 8\,\text{m}$ deep, held between flood walls over its bed at street level, crosses 62 to 82 m. Survey points every 2.5 m on the street's centerline, 39 of them, measure the settlement. Units are seconds, kilograms, meters and pascals, with pressures quoted in bar gauge, the advance in mm/min and the screw in rpm.

Each variable in Table 11.1 carries its role. Twenty-one states are integrated. Seven are physical, the condensed mass m_c and its volume w_c at zero gauge pressure, the gas mass m_g, the chainage x, the screw speed ω, the advance rate v and the chamber's conditioning c_F. Two integrate the ground lost and over-excavated, and twelve are ledgers of the mass entering and leaving by each route, the theoretical cut and the spewed volume.

11.3.1 The Ground's Target and Its Limit

The pressure the face needs is the at-rest earth pressure plus the pore water pressure at the axis, as shown in Equation (11.1).

$$\begin{aligned} \sigma_v &= \rho_w g\, h_r r(x) + \gamma z_0 \\ u &= \rho_w g \big(z_0 - z_{wt}(1-r) + h_r r\big) \\ p_t(x) &= K_0(\sigma_v - u) + u \\ \sigma_{v,crown} &= \rho_w g\, h_r r + \gamma(z_0 - R) \end{aligned} \tag{11.1}$$

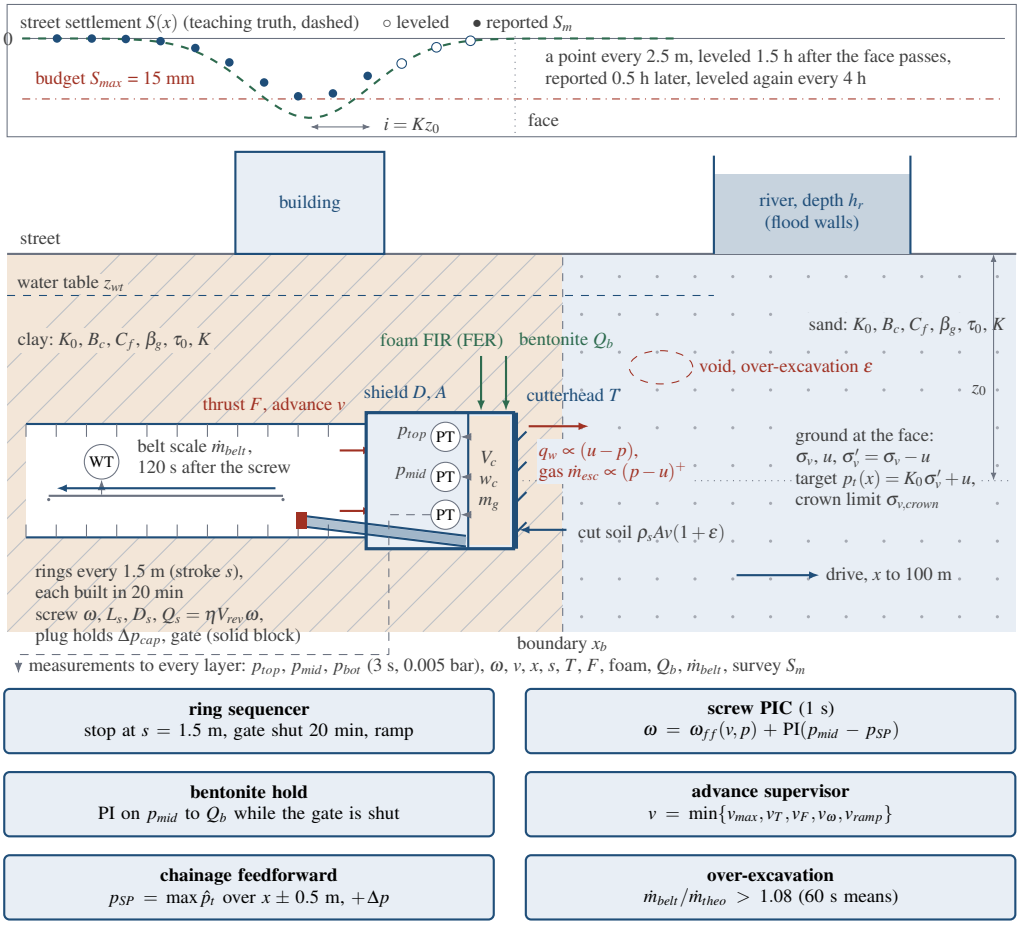

Figure 11.1 The drive as modeled, mirroring the simulator's Section view. The shield of diameter D and face area A advances at v through clay and then sand. Its full chamber of volume V_c holds condensed muck w_c and foam gas m_g at the axis pressure p, read by three cells, and receives cut soil, foam and bentonite. The screw at speed ω discharges through a gate onto a belt whose scale reads $\dot{m}_{belt}$ a conveyor delay later. The face exchanges water q_w and gas with the ground, a planted void adds ε, and the lost volume settles the street in a trough $S(x)$ of width $i = Kz_0$ that the survey reports late. In the strip the true trough is the thick dashed curve, the budget the dash-dotted line, the leveled points open circles and the reported points filled dots. The screw's gate is the solid block at its discharge. Control blocks read only the instruments, along the dashed signal line. Not to scale.

where $r(x)$ is the river's presence (0 on land, 1 under it, with 1 m cosine banks), γ the soil's unit weight, $R = D/2$, and K_0, γ and the other soil constants blend over the boundary with the sand fraction $s(x)$ (permeability terms on a logarithmic scale). The coefficient K_0 is a teaching value of the class that friction angle and stress history predict [18], 0.60 in the clay and 0.45 in the sand. The target is 2.226 bar in the clay, 1.997 bar in the sand on land and 2.944 bar under the river, where the pore pressure is 2.256 bar against 1.177 bar on land. A blow-out risk is flagged while the crown's chamber pressure $p - \rho_{mix} gR$ exceeds $\sigma_{v,crown}$, which is 2.291 bar in the clay, 2.350 bar in the sand and 3.135 bar under the river. Because ρ_{mix} is the chamber's own density, the land's limit expressed at the axis cell moves with the muck, from 2.91 to 2.95 bar in the sand before the river.

Table 11.1 Signals and States, With Values in the Clay at 12 m (Case 2 Before Its Step)

Variable	Role	Value	Range
p_{mid}	CV, axis cell	2.40 bar	SP 1.0–3.5 bar
ω	MV, screw speed	6.63 rpm	0–22 rpm
Q_b	MV, bentonite (ring builds)	0	0–8 m^3/h
v	MV, advance rate	50 mm/min	0–100 mm/min
FIR	MV, foam injection ratio	30 %	0–100 %
S	product, settlement (truth)	0	budget 15 mm
S_m	measured, survey report	—	noise 0.5 mm
T, F	constraints, torque, thrust	62, 55 %	≤ 90, ≤ 90 %
geology	DV, soil, water, river	clay	clay, sand, river
ε	DV, void (case 8)	0	0–0.3
seepage, boulder	DV, events	1, none	0–4×, +30 points
m_c, w_c, m_g	states, chamber	—	full chamber
x, c_F	states, chainage, conditioning	12 m, 0.30	0–100 m
V_L, V_o	states, lost, over-excavated	0	≥ 0
B	output, chamber stiffness	21.3 bar	—

11.3.2 A Full Chamber and Its Stiffness

The chamber is always full. Its condensed phase (soil with its pore water, foam liquid, bentonite) of volume w_c at zero gauge pressure and its gas of mass m_g occupy V_c at the one pressure that solves Equation (11.2).

$$w_c\left(1-\frac{p}{B_c}\right)+\frac{m_g R_a T}{p+p_{atm}}=V_c, \qquad \frac{1}{B}\approx\frac{1-n_g}{B_c-p}+\frac{n_g}{p+p_{atm}} \tag{11.2}$$

where B_c lumps the stiffness of the condensed muck, the bulkhead and the face (40 bar in the clay and 800 bar in the sand, teaching values), $R_a = 287.05\,\mathrm{J/(kg\,K)}$, $T = 293.15\,\mathrm{K}$ and n_g is the gas fraction. The first relation is a quadratic in p solved in closed form at every stage of the integrator, with $p = 0$ (a part-full chamber at atmospheric pressure) when no positive root exists. The second is the mixture's stiffness $B = -V\,\mathrm{d}p/\mathrm{d}V$, which the code evaluates by a finite difference of the first. Lined out 0.15 bar over the target at 50 mm/min with 40 % foam, the gas fraction is 9.9 % and $B = 18.82$ bar in the clay at 48 m, and 8.0 % and 37.72 bar in the sand at 52 m, which a hand evaluation of the second relation reproduces to within 0.1 %.

11.3.3 The Inventory

Soil enters at the advance, foam brings liquid and gas in proportion to the cut volume, and the screw and a failed plug take a fraction of the well-mixed contents, as shown in Equation (11.3).

$$\begin{aligned}
\dot{m}_c &= \rho_s A v(1+\varepsilon)+\rho_w\frac{\mathrm{FIR}\,A v}{\mathrm{FER}}+\rho_b Q_b+\rho_w q_w-f_o m_c\\
\dot{w}_c &= A v(1+\varepsilon)+\frac{\mathrm{FIR}\,A v}{\mathrm{FER}}+Q_b+q_w-f_o w_c\\
\dot{m}_g &= \mathrm{FIR}\,A v\left(1-\frac{1}{\mathrm{FER}}\right)\rho_{a,0}-\dot{m}_{esc}-f_o m_g\\
\dot{c}_F &= (\mathrm{FIR}-c_F)\,q_{in}/V_{mix}
\end{aligned} \tag{11.3}$$

where $f_o = (Q_s + Q_{sp})/V_{mix}$ is the fraction of the contents leaving per second, $A = 33.18\,\mathrm{m}^2$, ρ_s is 1990 or 2040 kg/m^3, FER = 10 is the foam expansion ratio, $\rho_{a,0} = 1.204\,\mathrm{kg/m}^3$, $\rho_b = 1030\,\mathrm{kg/m}^3$, V_{mix} is the mixture's volume (equal to V_c when full) and q_{in} the volume entering at p. The conditioning c_F is the foam ratio the chamber's contents carry, which lags the injected ratio by the chamber's turnover and is what the plug and the seepage see.

11.3.4 The Screw, Its Plug and Spewing

The screw discharges a volume per revolution, and the muck plug in it can hold only so much pressure, as shown in Equation (11.4).

$$\begin{aligned} Q_s &= g_a\,\eta V_{rev}\frac{\omega}{60}\,\phi, \qquad \Delta p_{cap} = \frac{4\,\tau_{plug} L_s}{D_s} \\ \tau_{plug} &= \left[\tau_{0,c}\mathrm{e}^{-c_F/0.8}(1-s) + \tau_{0,s}\frac{c_F^2\,s}{c_F^2+0.35^2}\right] \\ &\quad \times \left(1+0.15\frac{\omega}{\omega_{max}}\right) \\ Q_{sp} &= g_a\,k_{sp}\left(p+0.9\rho_{mix}gR-\Delta p_{cap}\right)^{+}\phi \end{aligned} \tag{11.4}$$

where g_a is the gate (shut only in ring builds), $\eta = 0.75$, $V_{rev} = 0.35\,\mathrm{m}^3$, $L_s = 12\,\mathrm{m}$, $D_s = 0.9\,\mathrm{m}$, $\tau_{0,c} = 15\,\mathrm{kPa}$, $\tau_{0,s} = 10\,\mathrm{kPa}$ and $\phi = \min(1, V_{mix}/0.25V_c)$ starves both flows in a part-full chamber. The plug law is the linear total pressure gradient set by the casing shear stress that model tests measured [6], with the sand's dependence on foam a teaching closure of the direction that laboratory screw tests show [7]. The inlet sits at the invert, so its pressure carries $0.9\rho_{mix}gR$ of muck head. The spew is a fast collapse toward the plug's capacity with $k_{sp} = 2\times10^{-6}\,\mathrm{m}^3/(\mathrm{s\,Pa})$, capped so that its time constant $V_{mix}/(B\,k_{sp})$ is never under 1 s. At 25 % foam in the sand the plug holds 1.80 bar still and 2.07 bar at full speed, while at 40 % it holds 3.17 bar.

11.3.5 Seepage Through the Face

Water moves through the face with the pressure difference and gas escapes into the ground while the chamber is over the pore pressure, as shown in Equation (11.5).

$$\begin{aligned} q_w &= \frac{\sigma_s C_f}{1+c_F/0.3}(u-p) \\ \dot{m}_{esc} &= \sigma_s\beta_g\frac{(p-u)^{+}}{1\,\mathrm{bar}}\,m_g \end{aligned} \tag{11.5}$$

where C_f is 4×10^{-11} in the clay and $1.2\times10^{-9}\,\mathrm{m}^3/(\mathrm{s\,Pa})$ in the sand, β_g is 2×10^{-6} and 3×10^{-4}/s per bar, and σ_s scales both (1 by default, 4 after a water inflow). The foam factor follows the measured fall of the infiltrated sand's permeability with the foam ratio in direction only [11]. Sealed for twenty minutes at 58 m with 60 % foam, the sand chamber falls from 2.147 to 1.529 bar, toward the pore pressure, while a clay chamber at 20 m falls from 2.353 to 2.348 bar.

11.3.6 Ground Lost and the Street's Trough

A face under its target lets the ground move in at a rate set by the relative deficit, and every lost volume settles the street in a Gaussian bowl, as shown in Equation (11.6).

$$\begin{aligned} \dot{V}_L &= A\,d\,(c_x v + c_t), \qquad d = \mathrm{clip}\left(\frac{p_t-p}{K_0\sigma_v'},0,1\right) \\ S(x_q) &= \sum_j \frac{\Delta V_j}{2\pi i_j^2}\exp\left(-\frac{(x_q-\xi_j)^2+y_q^2}{2i_j^2}\right), \quad i_j = Kz_0 \end{aligned} \tag{11.6}$$

where $c_x = 0.05$, $c_t = 1.4\,\mathrm{mm/min}$ (a standing face still moves), ΔV_j is the volume lost and over-excavated while the face crossed the 0.5 m cell centered at ξ_j, and y_q is the point's offset from the axis (0 for the survey). The trough width uses $K = 0.50$ in the clay, O'Reilly and New's cohesive value, and $K = 0.35$ in the sand, Mair and Taylor's mean [4], so $i = 7.5\,\mathrm{m}$ and 5.25 m. A void planted 1 to 2 m ahead of the face runs in at $\varepsilon = 0.3\,\mathrm{clip}((p_t+\Delta h-p)/\Delta h,0,1)$ with $\Delta h = 0.4\,\mathrm{bar}$. The face-loss law is zero whenever $p \geq p_t$. No shield-gap, tail-void or consolidation loss is modeled, so a face held on its target settles nothing, a point taken up in Section 11.7.

Table 11.2 Parameters, Their Defaults and Their Status. A = Adjustable in the Advanced Panel, T = Teaching Value, S = Inside a Cited Range. Two Values Are Clay, Sand

Symbol	Meaning	Value	Unit	Status
D, A	shield diameter, face area	6.5, 33.18	m, m^2	declared
V_c	chamber volume	30	m^3	T
z_0, z_{wt}	axis depth, water table	15, 3	m	T
h_r	river depth	8	m	T
γ	unit weight	19.5, 20.0	kN/m^3	T
ρ_s	muck density	1990, 2040	kg/m^3	T
K_0	at-rest coefficient	0.60, 0.45	—	S [18]
B_c	condensed stiffness	40, 800	bar	T, A
FER	foam expansion ratio	10	—	T
C_f	seepage conductance	4×10^{-11}, 1.2×10^{-9}	m^3/(s Pa)	T
β_g	gas loss per bar	2×10^{-6}, 3×10^{-4}	1/s	T
σ_s	seepage scale	1	—	A
ρ_b	bentonite density	1030	kg/m^3	T
$Q_{b,max}$	bentonite capacity	8	m^3/h	T
η	screw efficiency	0.75	—	T, A
V_{rev}	screw volume per revolution	0.35	m^3	T
L_s, D_s	screw length, diameter	12, 0.9	m	T
ω_{max}, τ_ω	screw range, drive lag	22, 3	rpm, s	T
$\tau_{0,c}$, $\tau_{0,s}$	plug shear strength	15, 10	kPa	T, A
k_{sp}	spew conductance	2×10^{-6}	m^3/(s Pa)	T
c_x, c_t	face loss per unit deficit	0.05, 1.4	—, mm/min	T, A
K	trough width factor	0.50, 0.35	—	S [4], A
a, T_0	torque per mm/min, at rest	1.52, 2.04 / 20	%	T
F_{skin}, k_F	skin friction, per mm/min	10, 12 / 0.08	MN	T
F_{max}, τ_v	thrust rating, cylinder lag	40, 10	MN, s	T
ring	width, build time	1.5, 20	m, min	declared, A
survey	leveling, report, round	1.5, 0.5, 4	h	T, A
limits	budget, torque	15, 90	mm, %	T

11.3.7 Machine, Instruments and Survey

The screw and the advance follow their commands through first-order lags of 3 s and 10 s. Cutterhead torque in percent of a 5 MN m rating is $T = 20 + a\,v\,\mathrm{e}^{-\mathrm{FIR}/0.5} + T_b$ with v in mm/min, $a = 1.52$ in the clay and 2.04 in the sand, and a boulder T_b that ramps 30 points over 0.15 m and back. Thrust in percent of 40 MN is $100(pA + F_{skin} + k_F v)/F_{max}$ with F_{skin} 10 and 12 MN and $k_F = 0.08$ MN per mm/min. Three cells on the bulkhead, 2.5 m above, at and below the axis, read $p \mp \rho_{mix} g\, 2.5\,\mathrm{m}$ through a 3 s lag with colored noise of 0.005 bar. The belt scale reads the screw's discharge 120 s later through a 10 s lag with 2 % noise. Each survey point is armed when the face passes it, leveled 1.5 h later, reported 0.5 h after the leveling with 0.5 mm of noise, and leveled again every 4 h. A reported value holds until the next one, and the instruments draw their noise from their own seeded stream.

11.3.8 Parameters, Initialization and Numerics

Table 11.2 lists the parameters with their status. Each case starts lined out at its chainage, advance, foam and pressure. A bisection finds the chamber turnover at which inflow fills V_c at the starting pressure, the gas is polished so the pressure is exact, and the screw starts at the speed that holds it. The plant advances by fourth-order Runge-Kutta steps of $\Delta t = 0.5$ s with every ledger carried through the same stages, the controller runs every 1 s and the log every 10 s, on the three clocks of Section 1.6.1. Each step's lost volume is assigned to the cell the face crossed and its bowl is added to every point. The chamber's mass balance closes to 2.0×10^{-13} of the mass that entered over a disturbed two-hour run.

11.3.9 Assumptions

The chamber is full and well mixed, with one pressure at the axis and a hydrostatic gradient over the bulkhead, so a chamber run part-full under compressed air is out of scope, and the model is least

reliable just after a large mismatch has emptied it. The condensed stiffness B_c lumps the muck, the bulkhead and the face. Its two values were chosen so that the stiffness about doubles at the boundary, and no measurement fixes them. The gas is ideal and isothermal and decays only by escaping into the ground. The plug, seepage and gas-loss laws carry the directions of the cited tests and teaching magnitudes. The ground loss is a closure proportional to the deficit, the face never collapses, and no loss arises at the shield gap, the tail void or by consolidation. The trough is a greenfield sum of bowls, so the building neither stiffens the ground nor changes the trough. The river is a uniform 8 m of water over a level bed, and the sand under it is the sand on land. The settlement budget, the blow-out limit at the crown's overburden and the 90 % torque limit are teaching values. Real ones come from the project's ground-movement criteria and a building damage assessment [5], from a face stability analysis [3] and from the machine's rating, and a pass here is no assessment of any real drive.

11.4 Why the Face Is Hard to Hold

11.4.1 An Integrator Whose Gain Is the Foam

With the chamber full, a screw 1 rpm faster removes $\eta V_{rev} = 0.2625\,\mathrm{m}^3$ a minute, and the pressure falls at $B\eta V_{rev}/(60V_c)$. That is 274 Pa/s per rpm in the clay at 48 m and 550 Pa/s per rpm in the sand at 52 m, at the same 40 % foam. On the plant a 1 rpm step falls at 266 and 512 Pa/s between 10 and 20 s, the gap being the screw drive's lag and, in the sand, a stiffness that falls to 31.9 bar within a minute as the pressure drops 0.27 bar. In case 1 the stiffness is 21.3 bar, which predicts 0.187 bar a minute for the 1 rpm error against the 0.168 bar measured in the first minute through the drive and cell lags. There is no self-regulation. Balancing the flows stops the pressure where it is and does not bring it back.

Pause. Try Exercise 11.1 before reading on.

The doubling at the boundary has two causes. The stiffness rises from 18.82 to 37.72 bar, a factor of 2.00. With no gas escaping into the sand the same chamber would be 30.4 bar stiff, so 1.62 of the factor comes from the two teaching values of B_c, and the rest from the gas that the sand's pores take (it keeps 78 % of its gas at this turnover). More foam softens the chamber. In the sand of case 6 the stiffness averaged over three hours 43.9 bar at 40 %, 31.1 bar at 60 % and 21.3 bar at 95 %, so the foam ratio is the loop gain as well as the plug's strength.

11.4.2 A Stop the Screw Cannot Cover

Every 1.5 m of stroke the advance comes to rest, the screw stops and its gate closes for twenty minutes. Nothing enters and nothing leaves by the screw, so the pressure moves only by seepage, as shown in Figure 11.2a. In the sand at 55.5 m with 60 % foam, the chamber starts the build 0.15 bar over its 1.997 bar target, falls under it 2.2 min into the build and reaches 1.516 bar at its end, heading for the 1.177 bar pore pressure. The face meanwhile lost enough ground to settle the 57.5 m point 1.76 mm after the first build from this stop. The 1.78 mm figure occurs only after the second build has begun. The second handle is bentonite slurry injected into the crown, which with its hold loop kept the chamber at or above 2.056 bar with 1.10 m^3 through the build and a peak of 5.78 m^3/h.

> **TRY IT IN THE APP 3 · The Ring Build: Twenty Minutes With Nothing Coming In**
>
> Load the scenario and press Run on the Chamber view. Watch the chamber pressure through the shaded ring build: with the gate shut it falls under its dashed at-rest target 2.2 min into the build and ends near 1.52 bar. Set Bentonite hold in the ring build to On, press Reset and run again. The chamber now stays at or above 2.06 bar.

11.4.3 A Plug That Fails at Any Speed

The screw's inlet pressure is the chamber's pressure plus the head of muck over the invert, while the plug's capacity depends on the muck and only 15 % on the speed. At 25 % foam in the sand the plug

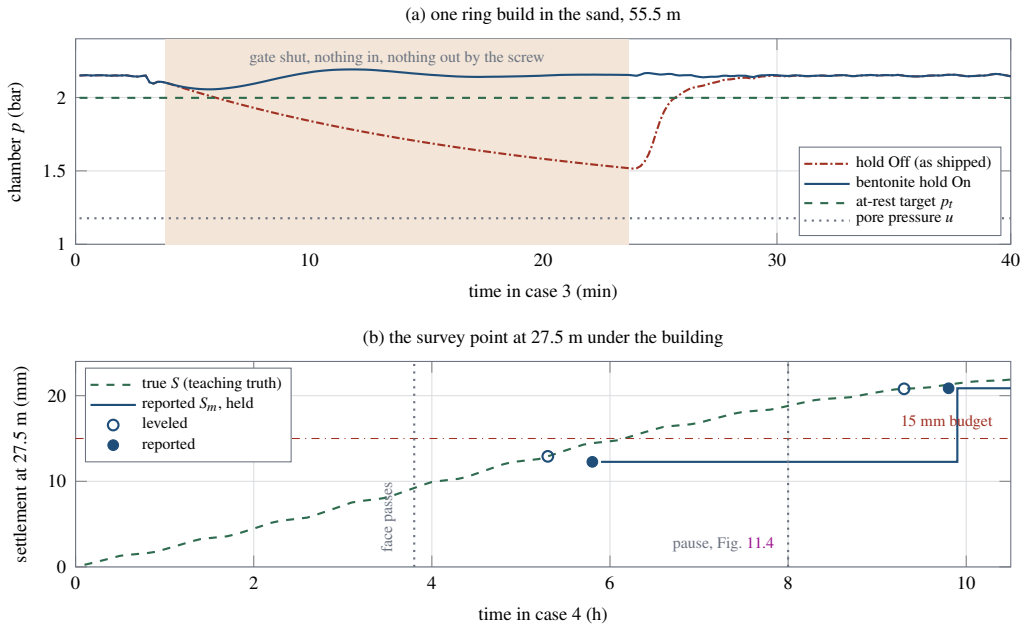

Figure 11.2 The forced stop and the late product, from the simulator. (a) Case 3 in the sand, the chamber through one ring build (shaded) without the bentonite hold (dash-dotted) and with it (solid), against the at-rest target (dashed) and the pore pressure (dotted). (b) Case 4 at the 27.5 m point, the true settlement (dashed), the leveling marks (open circles) and the held reports (solid steps, filled dots), with the budget dash-dotted. The first report, 12.3 mm at 5.80 h, is under the budget while the point is already at 14.5 mm and on its way to 21.3 mm.

holds at most 2.072 bar at full speed, and with the chamber at the land's 1.997 bar target the inlet sees at least 2.284 bar even for a muck as light as water. No screw speed, setpoint, advance or bentonite rate can then hold the chamber without spewing, which is the honestly infeasible variant of case 6. A sweep of 936 hand-set and automatic runs over four chainages including the river found none.

11.4.4 A Target That Moves and a Product That Arrives Late

The target rises by 0.947 bar from the sand on land to the river, because the water adds to the pore pressure and to the vertical stress. The river's 2.944 bar is level with the land's blow-out limit expressed at the axis, so no flat setpoint passes over one with room and under the other.

🕮 **Pause.** Try Exercise 11.3 before reading on.

The settlement is late for two reasons. A loss at the face spreads a bowl $i = 7.5$ m either way in the clay, so a point keeps settling until the face is well past, and a point's first report comes 2 h after the face passed it. In case 4 the 27.5 m point, passed at 3.80 h, was leveled at 5.30 h at 12.9 mm, reported 12.3 mm (with the survey's noise) at 5.80 h with the truth at 14.5 mm, and reported again 20.9 mm at 9.80 h, as shown in Figure 11.2b. A first report under 15 mm can still fail. A paper machine's traversing gauge, which posts a quality only after a whole scan and holds it until the next, poses the same problem of a product measured late in Volume 6, Chapter 3.

11.5 The Control Layers as Implemented

Every layer runs each $\Delta t_c = 1$ s from a measurement packet filled once per physics step with the three cells, the screw speed and torque, the cylinders' stroke and rate, the chainage, the cutterhead torque and thrust, the foam and bentonite flows, the belt scale and the survey reports, plus the declared

site investigation (the soils and target profile of the default plant). A source check in the test suite confirms that the controller and the ring sequencer contain no read of a plant state, a settlement, a score or a truth output.

11.5.1 The Screw Loop and Its Feedforward

The pressure loop is the velocity-form PI of Section 1.6.2, with the limit on the output alone and back-calculation on the integral. It acts on the axis cell in bar and speeds the screw when the pressure is over its setpoint (action $a = -1$ in Equation (1.12)), as shown in Equation (11.7).

$$
\begin{aligned}
e_k &= p_{mid,k} - p_{SP,k}, \qquad s_k = s_{k-1} + K_c(e_k - e_{k-1}) \\
\hat{u}_k &= s_k + I_{k-1} \\
u_k &= \min(\omega_{max} - \omega_{ff},\ \max(-\omega_{ff},\ \hat{u}_k)) \\
I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - \hat{u}_k}{\tau_I}\right]\Delta t_c \\
\omega_{cmd} &= \omega_{ff} + u_k \\
\omega_{ff} &= \frac{60 A v_m}{\eta V_{rev}}\left[\left(1 + \frac{\mathrm{FIR}}{\mathrm{FER}}\right)\left(1 - \frac{p_{mid}}{B_c}\right)\right. \\
&\qquad \left. + \mathrm{FIR}\left(1 - \frac{1}{\mathrm{FER}}\right)\frac{p_{atm}}{p_{mid} + p_{atm}}\right]
\end{aligned}
\tag{11.7}
$$

where the running sum s_k is never clamped, v_m is the measured advance, and B_c and FER inside ω_{ff} are the declared soil's. In the symbols of Equations (1.12) to (1.15), s_k is S_k with no derivative term, $\hat{u}_k$ is v_k and $T_t = \tau_I$. The defaults are $K_c = 15\,\mathrm{rpm/bar}$ and $\tau_I = 90\,\mathrm{s}$. The feedforward is the speed that discharges what the measured advance brings in, and the PI's range is shifted by it, so the integral never winds up against a limit the feedforward has already used. While the gate is shut the PI tracks a stopped screw, and switching the feedforward, the mode or the gate moves the command by 0.000, 0.000 and 0.043 rpm.

11.5.2 The Hold and the Ring Sequence

The ring sequencer reads only the stroke and the advance. At a stroke of 1.5 m it asks the advance to stop, keeps the screw loop running until the advance reads under 0.5 mm/min, then shuts the gate for twenty minutes, and on restart ramps the advance cap at 20 mm/min per minute. During the build a second PI on the same axis cell sets the bentonite flow, which rises when the pressure is under its setpoint, with $K_{c,b} = 18\,\mathrm{m^3/h}$ per bar and $\tau_{I,b} = 60\,\mathrm{s}$ with an output range of 0 to $8\,\mathrm{m^3/h}$, and in the original layered structure at every other time it tracks zero.

11.5.3 The Advance Supervisor and the Mid-Ranging Limit

With the supervisor on, the advance request is the least of five limits, each an integrator on the request actually selected, as shown in Equation (11.8). The minimum is a low selector (Section 1.4.3), and starting every limit from the selected request is external reset (Section 1.4.4).

$$
\begin{aligned}
b &= \min(v_{k-1}, v_{cap}) \\
v_k &= \max\left(0, \min\{v_{max}, v_T, v_F, v_\omega, v_{ramp}\}\right) \\
v_T &= b + k_T(T_{SP} - T)\Delta t_c \\
v_F &= b + k_F(90 - F)\Delta t_c \\
v_\omega &= b + k_\omega(0.85\,\omega_{max} - \omega)\Delta t_c \\
v_{ramp} &= b + r_v \Delta t_c
\end{aligned}
\tag{11.8}
$$

where v_{cap} is the ring sequence's ceiling, $k_T = 0.08\,\mathrm{mm/min}$ per percent second, $k_F = 0.02\,\mathrm{mm/min}$ per percent second, $k_\omega = 0.15\,\mathrm{mm/min}$ per rpm second, $r_v = 10\,\mathrm{mm/min}$ per minute, $v_{max} =$

100 mm/min and $T_{SP} = 82\,\%$ by default. Because every limit starts from the same selected base, none winds up while another holds the advance or a ring build stops it. The screw-room limit v_ω is the mid-ranging piece [10]. The fast screw holds the pressure, and when it nears the top of its range the slow advance yields. It is the valve-position law of Section 1.4.7 acting as one input of a selector, while Chapter 10 stacks three mid-ranging layers on one power block. A torque interlock stops the advance at 100 % and releases it under 95 %.

11.5.4 Foam, the Chainage Feedforward and the Detector

The foam ratio is fixed, or by soil from the declared consistency target of the soil half a meter ahead (30 % in clay, 60 % in sand). The chainage feedforward sets $p_{SP} = \max\{\hat{p}_t(x-0.5), \hat{p}_t(x), \hat{p}_t(x+0.5)\} + \Delta p$ from the declared target with a margin $\Delta p = 0.15$ bar, so the setpoint covers the span the face occupies. The over-excavation detector compares the belt scale with the theoretical mass rate $\hat{\rho}_s A v_m + \rho_w \,\mathrm{FIR} A v_m/\mathrm{FER} + \rho_b Q_b$ (plus the foam's gas) delayed by the same 120 s, both through 60 s averages, and alarms while their ratio exceeds 1.08 with the face advancing faster than 5 mm/min.

11.5.5 Nominal Predictive Coordination

The plant model is that of Section 11.3. A reader who identifies the screw's effect from case 1 finds an integrating gain of about 0.17 bar a minute per rpm. The shipped controllers' model is the declared soil and machine inside ω_{ff} and the declared target profile, which a change to a ground parameter in the Advanced panel does not update.

The optional MPC optimizes screw and advance [19], in the general form of Section 1.6.3, using measured pressure p_m (bar), screw speed ω (rpm) and advance v (mm/min) as states. Its coefficients use declared geometry, soil and commanded foam f. Define liquid and gas volume factors $L = (1 + f/\mathrm{FER})(1 - p_m/B_c)$ and $G = f(1 - 1/\mathrm{FER})p_{atm}/(p_m + p_{atm})$. Then $n_g = G/(L+G)$ gives the nominal stiffness B from Equation (11.2), without reading the chamber's actual gas. The local balance is

$$\begin{aligned} \dot{p}_m &= g_\omega g\omega + g_v v + BQ_b/V_c + \hat{d}, \\ \dot{\omega} &= (u_\omega - \omega)/3, \qquad \dot{v} = (u_v - v)/10, \\ g_\omega &= -B\eta V_{rev}/(60V_c), \qquad g_v = BA(L+G)/(60000V_c). \end{aligned} \tag{11.9}$$

Here g is the current gate state, Q_b is the applied injection in m^3/s, and $\hat{d}$ is a pressure-rate disturbance in bar/s. At the clay design point $B = 21.2933$ bar, $g_\omega = -0.00310527$ bar/s/rpm and $g_v = 0.000411525$ bar/s per mm/min. Torque is predicted as $a_f v + b_T$, with $a_f = a\exp(-f/0.5)$, and thrust as $a_p p_m + a_v v + b_F$, with $a_p = 8.29577$ percentage points/bar and $a_v = 0.2$ percentage points per mm/min. The biases b_T and b_F match the current measured outputs.

Bidirectional steps over 120 s, with 240 samples at 0.5 s, give pressure RMS residuals of 0.001549 and 0.001 763 bar for screw and advance. The independent 0.01 s reference integration includes the known 3 s transmitter lag omitted from the prediction. The descriptive screw IPDT fit is $k' = -0.00293016$ bar/s/rpm with 4.428 s effective delay. It summarizes lags, not transport delay.

Offset correction [20] subtracts the nominal rate from measured pressure increments every 1 s, using mean measured speeds, the previous gate and applied injection. It is clipped to ± 0.05 bar/s and filtered with $\alpha = 1 - \exp(-1/45)$ as $\hat{d}_k = (1-\alpha)\hat{d}_{k-1} + \alpha d_{res,k}$. The correction keeps updating under Layers and in Manual and starts again at Reset. No Kalman gain is used.

The model is discretized exactly every 10 s for 60 samples (600 s). Ten blocks of 1, 1, 2, 3, 4, 6, 8, 10, 12 and 13 samples give 20 command variables and three nonnegative horizon-wide slacks. With

$u_j = (u_{\omega,j}, u_{v,j})$, the fixed default objective, up to a reference-only constant, is

$$\begin{aligned} J = \sum_{j=1}^{60} \big[& 30(p_j - r_j)^2 + 0.04(u_{\omega,j} - \bar{u}_{\omega,j})^2 \\ & + 0.002(u_{v,j} - \bar{u}_{v,j})^2 \big] \\ & + \sum_{b=1}^{10} \big[(\Delta u_{\omega,b})^2 + 0.02(\Delta u_{v,b})^2 \big] \\ & + 10^4 s_p^2 + 3 s_T^2 + 3 s_F^2 . \end{aligned} \tag{11.10}$$

The weights are in bar, rpm, mm/min and percentage points. The terminal weight is zero and the Hessian carries a ridge of 10^{-7}. The pressure and move priorities multiply the tracking and slack terms and the move terms, and both stay at 1 in this chapter. Normally $\bar{u}_\omega = 11$ rpm and $\bar{u}_v$ is the allowed advance request.

Pressure $p_{SP} \pm 0.10$ bar, torque 85 % and thrust 90 % are soft, using one slack per output across all 60 samples. The pressure constraint remains centered on p_{SP} even when the objective reference r_j is boosted. Planned screw commands are hard-bounded at 0–18.7 rpm, and advance at zero to the live ceiling, also limited by the operator request when the supervisor is off. Block changes are at most 4 rpm and 10 mm/min, with exact holds inside each block. Initial measured speeds can exceed lowered command bounds while decaying. Sequence stops, interlocks, fallback and ceiling changes can reduce applied commands immediately, beyond planned move limits. The original gate, stop/restart sequence, 100 % torque interlock and capped 8 m^3/h hold retain priority.

The arrival preview divides the measured remaining stroke by the measured advance. The prediction sets both input references to zero after this estimated arrival, adds 0.03 bar to the pressure objective before arrival when it is less than 60 s away, and can pre-arm the existing hold in the last 20 s. Current gate and injection are otherwise held in the prediction. Setting the preview Off removes this forecast, and a declared timing error shifts only its time. Future restarts and events, hidden voids, the settlement and the live plant coefficients are not available to the plan. The layers have no forecast input, so each comparison below sets an MPC row with the preview Off beside the Layers row.

A plan is accepted only with an optimal status, correct shapes, a finite cost and iteration count, nonnegative finite slacks and finite predictions after conversion to pascals, and every sample must satisfy the command boxes, the block moves, the exact holds and the equality of the first move. A rejected plan holds the last applied commands for the first failed decision and then hands over to the tracked layered law within the hard bounds. A valid plan takes over again automatically. Thirteen injections of malformed results exercise these checks. Without noise, a failure or a manual transfer moves the first command by zero.

11.6 Results

The cases appear in Table 11.3. Each number below is a headless measurement on the default seed, taken in the page's own 0.5 s steps after any change a sentence names had been entered ahead of Run. Each case gives an identical run whether it is opened first, opened after some other case, or restarted with Reset. The exported file logs each instrument beside the quantity it measures, in a column whose name starts with `truth`.

11.6.1 The Inventory and the Loop

In case 1 the screw 1 rpm slow raised the axis cell from 2.397 to 2.565 bar in a minute and to 2.754 bar in two. Balancing the chamber set the screw to 5.22 rpm at the higher pressure, and the cell read 2.767 bar five minutes later, not 2.40 bar. A screw 1 rpm fast lowered it 0.190 bar in a minute. In case 2 the advance step of 50 % needed 3.37 rpm more screw (the feedforward computes 3.31 rpm). Without the feedforward the PI found it only as the pressure rose, which peaked 0.180 bar over

Table 11.3 The Cases, as Shipped

#	Start and settings	Teaching point
1	Manual, clay 8 m, 40 mm/min, screw 5.3 rpm	the inventory
2	Auto, clay 12 m, PI 15/90 s, screw feedforward Off, advance 50 to 75 mm/min at 15 min	loop and feedforward
3	Sand 55.5 m, 60 % foam, SP 2.15 bar, hold Off, 0.15 m to a ring	the forced stop
4	Clay 20.5 m, SP 2.05 bar under the building, pause at 8 h	budget and late survey
5	Sand 59 m, 60 % foam, SP 2.40 bar, river ahead, pause at 6 h	chainage feedforward
6	Sand 52 m, 25 % foam, SP 2.15 bar	plug, stiffness, infeasibility
7	Clay 47.5 m, 40 % foam, PI 80/20 s, chainage feedforward On	gain change at the boundary
8	Clay 34 m, SP 2.40 bar, void at 36–37 m, pause at 6 h	belt scale before survey
9	Auto, sand, 0.15 m stroke remaining, hold On, MPC pinned	arrival forecast comparison
10	Whole drive, screw held (gain 0), 50 mm/min, 25 % foam, no hold, no feedforward, boulder at 31 m	the whole problem

Table 11.4 Measured Results on the Cases. Settlement Is the Worst True Value at a Survey Point at the Instant Named

Case, instant	Setting	Settlement (mm)	Other
3, end of build	hold Off	1.76	min 1.516 bar
	hold On	0.00	min 2.056 bar
4, 8 h	SP 2.05 bar (shipped)	18.81	deficit 0.222 bar
	SP 2.10 bar	13.44	
	SP 2.16 bar	7.00	
	SP 2.20 bar	2.70	
	SP 2.23 bar	0.09	
	SP 2.40 bar	0.00	
5, 6 h	SP 2.40 bar (shipped)	100.5	blow-out 0
	chainage feedforward	0.0	blow-out 0
	SP 2.88 bar	11.8	blow-out 0
	SP 2.94 bar	1.2	blow-out 48.2 min
	SP 3.10 bar	0.0	blow-out 78.0 min
6, 3 h	foam 25 % (shipped)	60.6	spew 120 min
	foam 40 %	0.0	*B* 43.9 bar, torque 66 %
	foam 95 %	0.0	*B* 21.3 bar, torque 35 %
8, 6 h	SP 2.40 bar (shipped)	15.5	5.52 m^3 over
	SP 2.55 bar	5.2	1.87 m^3 over
	SP 2.65 bar	0.0	none over

setpoint 44 s after the step and stayed outside ± 0.02 bar for 227 s. With it the peak was 0.023 bar and the band was regained in 41 s, while retuning alone reached 0.074 bar at 40 rpm/bar and 40 s and 0.464 bar at 5 rpm/bar and 200 s. Acting on the cause instead of waiting for its effect is the same move as the coagulant feedforward of Volume 3, Chapter 6.

11.6.2 The Building, the River and the Boundary

Case 4 holds 2.05 bar, 0.18 bar under the clay's target, and at the 8 h pause the worst point had truly settled 18.81 mm at 27.5 m. The budget fixes a band, and raising the setpoint before the face arrives is the only remedy, as listed in Table 11.4. In case 5 the fixed 2.40 bar left the face 0.605 bar short under the river and settled the 65 m point 100.5 mm by 6 h, a trough already 60.2 mm deep at 4 h as shown in Figure 11.4b, while the chainage feedforward kept the chamber on or over its target and nothing settled. In case 7 the tight tuning held the clay within 0.028 bar with the screw between 5.0 and 8.6 rpm, and in the sand, whose chamber averaged 42.4 bar against 18.8 bar while mining, the screw swung from 0.3 to 17.0 rpm with the pressure 0.206 bar off. Halving the gain restored 0.060 bar in the sand, and the shipped 15/90 s is calm in both (0.039 and 0.076 bar).

🕮 **Pause.** Try Exercise 11.2 before reading on.

11.6.3 Foam and the Void

At 25 % foam in case 6 the muck spewed every minute the gate was open, 120 min in the first 3 h, the screw stopped at 0.0 rpm as the loop asked for less outflow, and the street settled 60.6 mm. At 40 % the plug held 3.17 bar against an inlet of 2.68 bar. More foam softened the chamber and needed more bentonite in the same three ring builds, 2.33, 3.26 and 4.67 m^3 at 40, 60 and 95 %. In case 8 the void began to run in at 62.0 min and the detector alarmed 228 s later, at 65.8 min with the face at 36.19 m, while the pressure loop simply took the extra muck out. The survey's first report of the damage came at the 35 m point at 140 min, 15.4 mm.

11.6.4 The Drive

The challenge asks for the 100 m in at most 48 h with every survey point at most 15 mm (true), no spewing, no blow-out risk and the torque at most 90 %, judged when the survey has reported every point. Most of the drive time is fixed by the lining. As shipped the drive took 57.56 h, of which 66 ring builds were 22.00 h, the restarts 2.73 h and the mining 31.97 h at 50 mm/min. The suggested configuration took 46.85 h, of which 65 builds were 21.67 h and the mining 24.24 h at a mean 68.1 mm/min, so the supervisor wins about a fifth of the time and only from the mining. Table 11.5 lists the configurations. The hold, the chainage feedforward and the supervisor are each necessary on wide sweeps. Without the hold the drive fails in all 15 combinations of a margin from 0.15 to 0.60 bar with a sand foam of 60, 80 or 100 %. Without the chainage feedforward it fails at all 31 fixed setpoints from 2.60 to 3.20 bar in steps of 0.02 bar, and without the supervisor at every fixed advance tried from 46 to 60 mm/min, on time at 46 to 47.9 mm/min (58.9 to 60.3 h) and on torque from 48 mm/min up. The screw feedforward is sufficient but not necessary. Without it the suggested 15/90 s loop fails on 3.1 min of blow-out risk, while four retuned loops pass, from 25/60 s at 4.0 mm to 40/20 s at 1.2 mm, and the two most aggressive (60/30 s and 80/20 s) fail on blow-out minutes again. Figure 11.3 sets the peak torque and the drive time of every configuration against the challenge's two limits.

> **TRY IT IN THE APP** **Challenge · The 100 m Drive, the Street Kept**
>
> Run the drive as shipped. It takes 57.6 h, the boulder takes the torque to 96.1 %, and the verdict is a fail. Press Reveal suggested tuning, press Reset and run again. With the advance supervisor slowing the advance at its torque target, the drive takes 46.9 h, the torque peaks at 86.1 %, and the verdict is a pass.

The mid-ranging limit is idle on the drive as shipped. With the suggested settings the supervisor's binding limit was the torque target for 37.7 % of the mining, its climb for 39.0 % and the ceiling for 23.3 %, and the screw never passed 17.8 rpm against its 18.7 rpm room limit. With a screw efficiency of 0.50 the limit bound for 31.6 % of the mining, held the chamber within 0.37 bar of its setpoint and slowed the drive to 48.5 h. Switched off by a test hook, the same screw sat at full speed for 347 min and the crown spent 48.4 min over its overburden. The challenge is built so that the reader must assemble the layers, and the full drive is shown at its building reach in Figure 11.4a from case 4.

> **DID YOU KNOW?** **No Fixed Advance Rate Passes the Drive**
>
> From the simulator: with the supervisor off, a fixed 48 mm/min took the boulder's torque to 90.04 %, over the 90 % limit, while 47.9 mm/min held it to 89.9 %. Both needed 58.9 h against the 48 h allowed (Table 11.5). Only a limit that moves with the ground wins both.

11.6.5 Predictive and Classical Comparisons

Table 11.6 compares the laws on matched seeds, instruments, events and settings, over windows of 6 h at the boundary, 2 h for each ring and 60 h for the challenge. The held ring is case 9, which starts

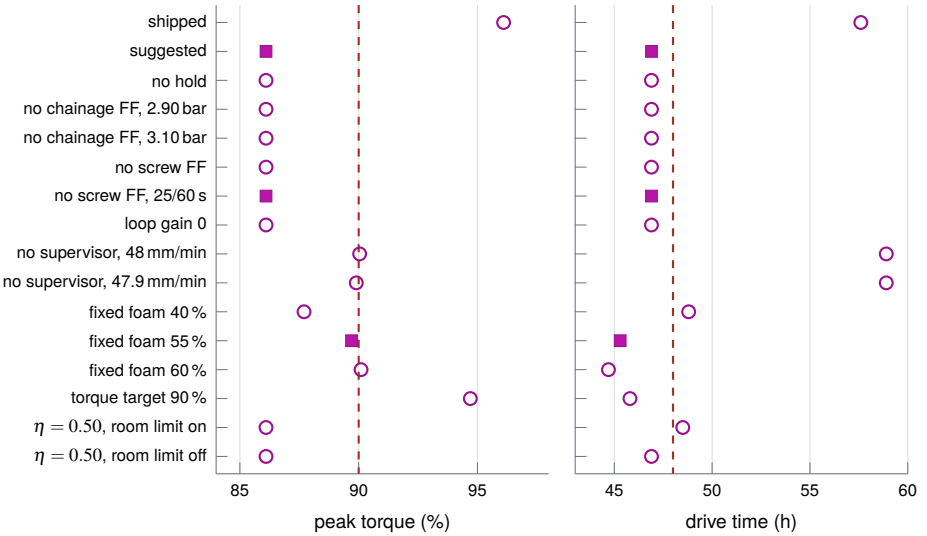

Figure 11.3 Peak torque and drive time of every configuration of Table 11.5, against the challenge's limits of 90 % and 48 h (dashed). Filled squares pass, open circles fail. Notice that the two fixed advances near 48 mm/min cannot meet both limits at once, while with the supervisor on the torque stays at or under 89.7 % unless its target is raised to 90 % or the foam is fixed at 60 %. Runs that fail inside both lines failed on settlement, spewing or blow-out minutes.

Table 11.5 The Drive. Worst True Settlement, Spewing and Blow-Out Minutes, Peak Torque and Drive Time

Configuration	S (mm)	Spew (min)	Blow (min)	T (%)	Drive (h)	Verdict
shipped	245.3	2079.1	1675.3	96.1	57.6	fails
suggested	0.0	0	0	86.1	46.9	passes
no hold	21.7	0	0	86.1	46.9	fails
no chainage FF, 2.90 bar	14.3	0	102.3	86.1	46.9	fails
no chainage FF, 3.10 bar	0.0	0	2302.7	86.1	46.9	fails
no screw FF	10.3	0	3.1	86.1	46.9	fails
no screw FF, 25/60 s	4.0	0	0	86.1	46.9	passes
loop gain 0	45.6	6.5	36.8	86.1	46.9	fails
no supervisor, 48 mm/min	0.0	0	0	90.04	58.9	fails
no supervisor, 47.9 mm/min	0.0	0	0	89.9	58.9	fails
fixed foam 40 %	69.6	371.1	0	87.7	48.8	fails
fixed foam 55 %	0.0	0	0	89.7	45.3	passes
fixed foam 60 %	0.0	0	0	90.1	44.7	fails
torque target 90 %	0.0	0	0	94.7	45.8	fails
$\eta = 0.50$, room limit on	1.6	0	0	86.1	48.5	fails (time)
$\eta = 0.50$, room limit off[a]	2.4	0.5	48.4	86.1	46.9	fails

[a] A test hook only, not a setting in the Control panel.

in sand with 0.15 m of stroke remaining and the hold On, and its Layers row selects Layers explicitly. The hold-Off ring is case 3, which keeps its hold Off.

The classical row applies the integrating SIMC rule [21] once, with $\lambda = \theta = 4.4\,\text{s}$ from the descriptive screw fit. It gives $K_c = 1/[|k'|(\lambda+\theta)] = 38.78159438$ rpm/bar and $\tau_I = 4(\lambda+\theta) = 35.2\,\text{s}$. The PI receives the pressure in bar. In the challenge the classical settings otherwise equal the suggested configuration. No tuning search or MPC adjustment was used.

Define control-band depth $d_c = \max(0, |p_m - p_{SP}| - 0.10)$ in bar and $I_c = \int d_c\,dt/60$ in bar min, with seconds as the integration variable. Physical depth instead uses the true target and blow-out limit, $d_g = \max(0, p_t - p, p - p_{blow})$, expressed in bar, with integral I_g. The two bands differ. The

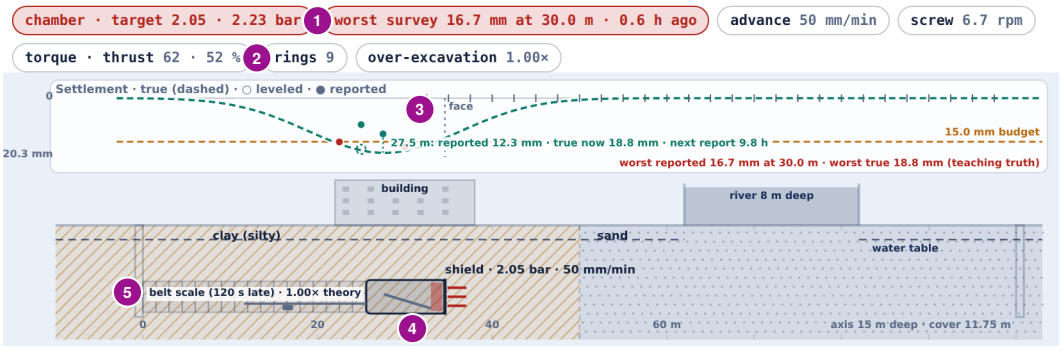

(a)

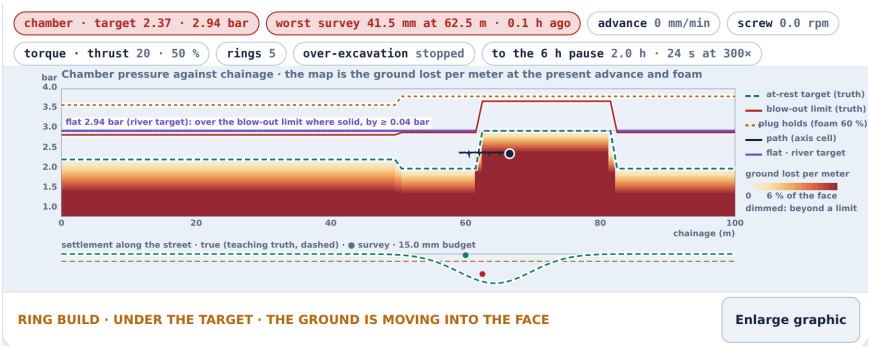

(b)

Figure 11.4 The late survey and the moving target. (a) Case 4 on the Section view at the 8 h pause, cropped to the readouts, the survey strip and the section, with the setpoint held at 2.05 bar against a 2.23 bar target (callout 1). The face had reached 34.6 m by the pause, and the rings readout counts nine builds. The screw turns at 6.7 rpm at 50 mm/min, and the torque and thrust, the supervisor's two machine limits, read 62 and 52 % (2). The survey's worst report is 16.7 mm at 30.0 m, 0.6 h old, and the dotted drop at 27.5 m (3) joins a held report of 12.3 mm to the true 18.8 mm, next reported at 9.8 h. The shield (4) carries the belt scale (5) behind it, which reads the discharge 120 s late. (b) Case 5 on the Envelope view at 4 h, during a ring build under the river with the fixed 2.40 bar setpoint. The axis cell reads 2.37 bar against a 2.94 bar target, the trough under the river has reached 60.2 mm (true) at 65 m and the survey's worst report is 41.5 mm at 62.5 m. The flat line at the river's target (violet in the app) is solid where it lies over the land's blow-out limit, by at least 0.04 bar at the present muck density of 1734 kg/m^3.

settlement is evaluation truth and never controller feedback.

On the held ring the MPC loses to the layers (0.26259 versus 0.04993 bar min), and with the hold Off it also loses (15.76819 versus 15.38332). The classical loop has the lowest integral in all four cases. With the preview On the MPC beats the preview Off on the held ring (0.26259 versus 5.45064 bar min), which supports the preview within this policy.

At 160 s the chamber under MPC On, Layers and MPC Off reads 2.19868, 2.08392 and 1.654 60 bar, as shown in Figure 11.5. The control changes when each run reaches the ring. A wrong preview can even help: an error of −120 s gives 0.02681 bar min on the held ring.

In the challenge the suggested and the classical layers pass at 46.8508 h of drive time. MPC On reaches the end at 44.7519 h but fails on 9 s of blow-out risk, MPC Off passes at 43.4329 h with 6.1221 mm of settlement, and the wrong preview fails on time at 50.9667 h. The survey is first complete at 48.0032, 45.5221, 44.2564 and 51.5776 h, in the same order. The 60 h integrals include

Table 11.6 Fixed-Law Comparisons at the Default 0.5 s Step. t_c Is Time Outside the Control Band, d_c Its Maximum Depth, S Worst True Settlement, X Advance and B_q Bentonite. All Solves Succeed.

Case	Law	I_c bar min	t_c min	d_c bar	I_g bar min	S mm	X m	B_q m^3	T_{max} %	F_{max} %	Solves
Boundary, 6 h	Layers	1.925	71.608	0.110	0.640	0.538	10.561	4.776	65.832	60.130	0
Boundary, 6 h	Classical	0.000	0.000	0.000	0.000	0.000	10.561	4.499	65.832	59.000	0
Boundary, 6 h	MPC On	0.302	15.508	0.072	0.011	0.004	10.524	5.119	65.832	58.883	2160
Boundary, 6 h	MPC Off	1.355	23.717	0.105	0.318	0.098	10.560	5.838	65.832	59.015	2160
Boundary, 6 h	MPC −120 s	0.050	3.158	0.066	0.002	0.001	9.724	4.261	65.832	58.697	2160
Hold Off, 2 h	Layers	15.383	53.192	0.539	12.740	4.158	3.176	0.000	50.722	57.923	0
Hold Off, 2 h	Classical	14.326	49.058	0.530	11.855	3.850	3.176	0.000	50.722	57.973	0
Hold Off, 2 h	MPC On	15.768	52.050	0.580	13.174	4.295	3.160	0.000	50.722	57.922	720
Hold Off, 2 h	MPC Off	20.323	56.633	0.616	17.400	5.643	3.176	0.000	50.722	57.921	720
Hold Off, 2 h	MPC −120 s	12.302	39.400	0.581	10.299	3.363	3.070	0.000	50.722	58.037	720
Held ring, 2 h	Layers	0.050	5.458	0.021	0.000	0.000	4.146	3.522	81.444	67.915	0
Held ring, 2 h	Classical	0.000	0.000	0.000	0.000	0.000	4.146	3.307	81.444	67.969	0
Held ring, 2 h	MPC On	0.263	9.600	0.046	0.000	0.000	4.543	3.599	75.962	66.073	720
Held ring, 2 h	MPC Off	5.451	25.833	0.402	4.253	1.362	4.712	5.110	76.086	66.154	720
Held ring, 2 h	MPC −120 s	0.027	0.908	0.058	0.001	0.000	3.162	3.212	75.936	66.056	720
Challenge, 60 h	Layers	6455.591	3568.808	2.913	5731.508	245.269	100.008	0.000	96.096	79.913	0
Challenge, 60 h	Suggested	0.601	48.867	0.070	0.001	0.000	100.017	41.987	86.127	75.812	0
Challenge, 60 h	Classical	0.001	0.333	0.008	0.000	0.000	100.017	37.423	86.127	75.855	0
Challenge, 60 h	MPC On	4.584	125.192	0.309	0.670	0.162	100.016	47.713	89.255	74.156	21600
Challenge, 60 h	MPC Off	94.380	862.375	0.634	56.027	6.122	100.010	77.305	89.753	74.192	21600
Challenge, 60 h	MPC −120 s	2.010	30.767	0.439	1.295	0.371	100.014	41.048	87.906	74.118	21600

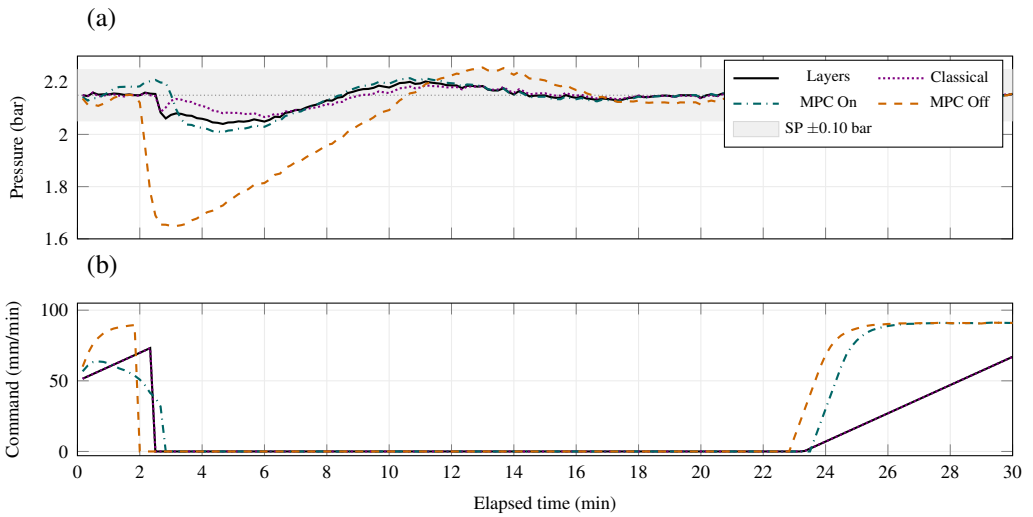

Figure 11.5 Measured ring-arrival transient over the first 30 min of the 2 h held-ring case, with Layers solid, the classical loop densely dotted, MPC On dash-dotted and MPC Off dashed. (a) Common pressure axes retain the Off minimum and show the setpoint band (shaded). (b) Applied advance commands on the same time axis. The full-run integrals in Table 11.6 include later ring cycles.

the holding after the drive has ended, 0.75607 bar min for MPC On. The advance averaged over the window differs from the productive rate, which uses the first arrival at 100 m. MPC On exceeds its soft 85 % torque target for 486.5 min, against 0.933 min for the suggested layers, both below the challenge's 90 % limit.

Faults injected unannounced at 300 s show further losses. With a water inflow, Layers, classical, MPC On and MPC Off settle 0.98982, 0.83311, 4.25629 and 5.970 91 mm. On spews for 1.792 min

and exceeds the blow-out limit for 6.500 min, against zero for Layers and classical. Off gives 0.383 and 1.525 min. The foam fault reaches 94.37 % torque under On against 83.15 % under Layers. Soft constraints allow these violations without a failed solve.

At 0.25 s, water-fault settlement changes from 4.256289 to 8.301 382 mm for On and 5.970910 to 7.586 344 mm for Off, reversing their order. The same seed draws a different noise path at the changed step. Matched noise-free runs give On 5.462907 to 5.488 265 mm and Off 6.615624 to 7.913 211 mm, so deterministic sensitivity also remains. Neither numerical convergence nor noise as the sole cause is established. The main pressure conclusions named above, MPC On against the layers and the classical loop, and the challenge grades survive the refinement. Other rankings need not: the wrong-forecast challenge integral changes from 2.009838 to 0.595714 bar min, slightly below suggested Layers' 0.598859 at the finer step, while still failing time. The fault settlement estimates are resolution-sensitive.

11.7 Verification, Agreement and What Is Not Claimed

The application's own suite of checks, expanded from 42 to 51, verifies the code. The chamber's mass balance closes to 2.0×10^{-13} of the mass entered over a disturbed two-hour run, a chamber lined out at its setpoint does not drift in ten minutes open loop, and a sealed chamber with seepage off changes neither pressure nor mass. A uniform 1 % loss superposed bowl by bowl reproduces $S_{max} = V_s/(\sqrt{2\pi}\, i)$ to every printed digit, 17.651 mm in the clay and 25.215 mm in the sand, with $S(i)/S(0) = 0.6065$ and $S(2i)/S(0) = 0.1353$, the Gaussian values. Halving the step moves case 4's worst settlement from 18.813 to 18.790 mm and changes no challenge number. For the original layered challenge, seeds 1 to 5 leave every number unchanged, because the noise enters only the cells, the belt and the survey and none of them decides a verdict. Playback at 60, 3600 and 7200× pauses at the same instants with the same verdicts, as Section 1.6.1 requires of the engine, and in a browser tab at 7200× the drive's verdict arrives in 24 to 29 s. The whole 48 h drive costs 0.8 s of headless computing, 2.3 µs per step.

Agreement with the literature is limited to what the cited records state. The trough factors 0.50 and 0.35 sit inside the published ranges [4]. The standstill loss in sand and not in clay has the direction field measurements report [8], and the model's states match the gas and liquid states of a published chamber model [9], but no magnitude was fitted to either. The suggested drive's 0.0 mm is stricter than any real drive, because the model has no shield-gap or tail-void loss. At the published median volume loss of 0.41 % of the face area for earth pressure balance drives [4], Peck's formula over this tunnel gives about 7.2 mm over the clay and 10.3 mm over the sand, and at the largest reported 0.83 %, 14.7 and 20.9 mm. The face-loss law is therefore no prediction of a real trough. It is the part of the loss that the face pressure controls. No educational evaluation has been made, and no learning gain is claimed.

Two plant fixes made in the last review of the simulator, a screw that starves once the chamber is under a quarter full and a spew whose time constant cannot fall under 1 s, moved no shipped or suggested run. They keep a screw run flat out by hand and a sand chamber at 0 % foam finite.

All ten cases run 52 h under MPC with finite outputs, bounded commands and zero solver failures. Tests reject 13 malformed plans and a stale-history mutant. Browser observations cover five decisions in 45 s at real time, with means of 1.16 and 1.56 ms for the solver and the full decision and 0.0222 % of a core for the whole control, and 8664 decisions in 12 s at 7200×, with 240 and 317 µs at near 60 frames a second. The 25 ms frame guard, which limits catch-up without skipping a decision that is due, was not active in either window. A forced replay with 4800 stops and replays at the browser's speeds preserve every value except the wall-clock times. Node and the browser differ in 113 rows, by at most 6.22×10^{-14} bar. The timings are host observations, not bounds.

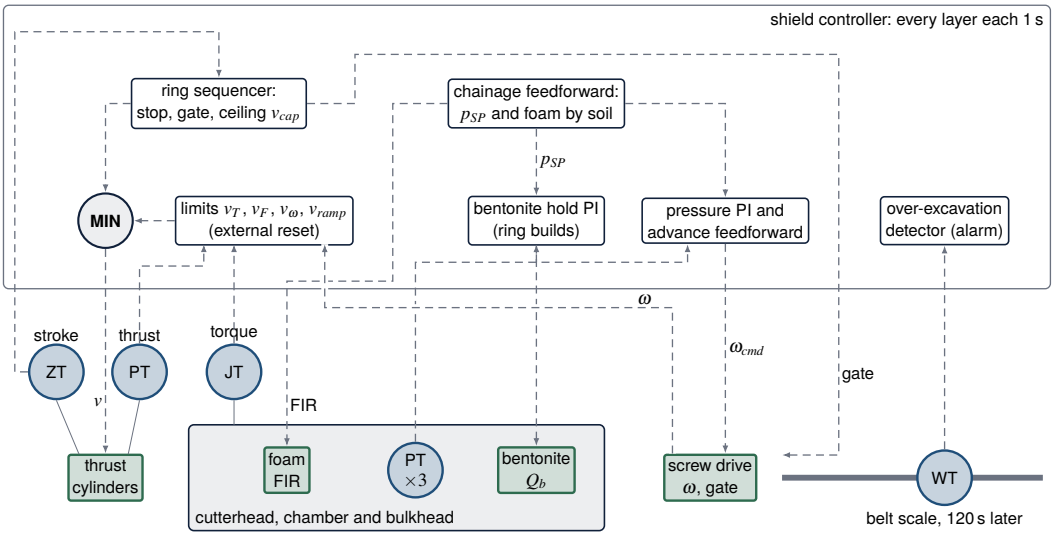

Figure 11.6 The face pressure's layers on real hardware. Three pressure cells (PT) on the bulkhead feed the pressure PI, which adds to the advance feedforward on the screw drive, and the bentonite hold, which injects slurry during ring builds. The chainage feedforward sets both setpoints and the foam. The thrust cylinders take the least of the ring sequencer's ceiling and the limits on torque (JT, from the cutterhead motors), thrust (PT, from the cylinder pressure) and screw speed. The belt scale (WT) reads the discharge 120 s late and only raises an alarm. Notice that every limit on the advance acts on the cylinders alone and leaves the screw loop free to hold the face.

11.8 Technology in Practice

A real shield reads and moves what the layers of Section 11.5 read and move. Figure 11.6 draws the screw loop, the bentonite hold and the advance supervisor's low selector with field hardware in place of the schematic's symbols.

What Measures It

In the earth pressure balance machine the pressure is transmitted to the face through the soil grains and is reduced by friction over the length of the screw conveyor. A slurry machine transmits it hydraulically and controls it with pressure gauges and control valves in a piping system, which gives a more precise and consistent control [2].

> **TECHNOLOGY Sensor: Pressure Cells on the Bulkhead**
>
> In an earth pressure balance machine the rotational speed of the screw is matched to the advance rate, and the pressure in the chamber is monitored with multiple sensors [2]. A process that carries solids tends to clog an instrument's impulse lines, as sludge does in wastewater service, and a remote seal's isolating diaphragm eliminates that problem. Its price is a time delay from the viscous fill fluid in the capillary, which makes the instrument slow, as if it carried a snubber [22]. No open source read gives a cell's range or accuracy. The model's three cells read through a 3 s lag with noise of 0.005 bar, and the loop uses the axis cell alone.

What Moves It

Conditioning agents are mixed with the soil through nozzles on the cutterhead, on the bulkhead and along the screw, to turn the muck into a plastic paste that keeps the screw full as a plug and lowers the torque of the cutterhead and the screw. Bentonite slurry is among those agents [23]. For a

rock machine, a high penetration rate does not by itself assure a high average advance rate, and the actual utilization typically runs in the range of 50 % [2]. The lining does the same to this drive: the suggested configuration spent 21.67 h of its 46.85 h building rings, so the supervisor can win time only from the mining.

> **TECHNOLOGY** **Final Element: Thrust Cylinders and the Screw Drive**
>
> The machine advances by shoving off the precast segments erected in its tail, and the force of its thrust cylinders passes from the pressure bulkhead onto the soil in the chamber. The amount of material removed is controlled by the speed of the screw and the opening of its discharge [2]. On one instrumented rock machine the thrust was read as hydraulic pressure in the cylinders, the cutterhead's load as the current of its motors, and the stroke by a device on the cylinder extension, which drew a sawtooth on the trend, one tooth for each stroke and retraction [24]. No open source read gives the speed of these drives. The model's screw and cylinders follow their commands through lags of 3 s and 10 s.

Where the Controller Runs

Manual or automatic operation of an earth pressure balance machine is possible through its integrated programmable logic controller and computer-control systems [2]. Supervisory and lower-level loops in such systems run continuously with cycle times from milliseconds to minutes [25]. Every layer of the application runs each 1 s, so the screw loop, the hold, the ring sequencer and the supervisor fit one controller scan, and the optional predictive coordinator plans on a 10 s grid in a supervisory layer above them.

What Is Modernizing It

Two developments bear on the face. The first is conditioning. New families of conditioning agents widen the range of both closed-face machine types and blur the line between them, and where the geology falls into more than one grading envelope a dual-mode open and closed machine, or a dual slurry and earth pressure balance system, may be developed [2]. The second is the machine's own data. On one instrumented rock machine a data acquisition system sent the pressures, currents, temperatures and the cylinder extension over a network to a laboratory's control system, where the mining cycle could be plotted from a distance [24], and a predictive controller built on a model learned from field data identified the screw speed and the advance speed as the major factors in the chamber pressure [1].

> **IN PRACTICE** **Face Pressure First, Then the Tail Void**
>
> Where ground is lost in the tunneling, the settlement can nearly always be mitigated. A highway tunnel manual names a closed, pressurized machine whose face pressure is at least the combined soil and groundwater pressure at tunnel level, and immediate, complete grouting of the space behind the segments, by automated systems that do not permit the machine to advance unless that void is grouted at the same time. Each percent of steering correction can translate to 1.5 % of ground loss. Good practice, with tight control of face pressure, is put at a volume loss of 0.5 % [2]. The grouting rule is one more interlock on the advance, like the torque interlock of Section 11.5.

11.9 Limitations

The drive is a teaching design, not a project. The condensed stiffness, the plug, the seepage and gas-loss coefficients, the face-loss closure and the torque law carry the directions of the cited work with magnitudes chosen so that each case teaches its point. The stiffness doubling across the boundary is mostly the two declared values of B_c. The street settles only through the face, so the absolute settlement is low against real drives and the comparison across configurations is the result. The trough is greenfield and transverse spreading is the only shape effect. The building does not interact with the ground, and no settlement develops with time once the face has stopped losing ground. The

chamber is always full, and a part-full chamber, face collapse, arching, steering, cutter wear and the slurry shield are outside the model. The declared site investigation equals the plant, so the chainage feedforward is exact except where a learner changes the ground. Rings are 1.509 m apart because the advance coasts about 8 mm through the cylinders' lag before each build, which gives 65 or 66 builds over the 100 m. The interlock and the flags are teaching devices rather than certified logic, and no part of the simulation is a safety case.

The MPC omits changing gas loss, conditioning, plug failure, voids and settlement. Residuals and output biases do not recover these dynamics. Its soft setpoint band differs from ground-support limits. Retained sequence and interlocks provide no actual-ground safety guarantee. The water-fault refinement limits quantitative settlement estimates. No retuning was used to force a ranking that holds in every case.

OPEN CHALLENGE **Choosing the Foam Without a Standard Test**

Foam conditioning is designed in the laboratory, yet no recognized standards are available for testing conditioned soil. The best foam injection ratio has to be found by specific tests, and pressure in the chamber lowers both the expansion and the injection ratio of the foam [23]. In this chapter the foam ratio sets the loop's gain as well as the plug's strength, so a standard test that links a laboratory mix to the stiffness and plug of a pressurized chamber would let a controller be tuned before the drive.

11.10 Conclusion

This chapter studied an earth pressure balance shield simulated as a full chamber whose pressure is its inventory, whose stiffness is set by the foam's gas, and whose outflow is a screw with a plug that can fail. One rpm of mismatch moves the clay chamber about 0.17 bar a minute and nothing returns it. Every ring build seals the chamber for twenty minutes, in which the sand drops it from 2.15 to 1.52 bar unless a bentonite hold keeps it over 2.06 bar. The target rises 0.95 bar under the river to the land's blow-out limit, so only a setpoint scheduled from the geology passes both reaches, and the survey's first report of a point under the building reads 12.3 mm while it is heading past 18.8 mm. Layered control drives the 100 m in 46.9 h against 57.6 h, with 21.7 and 22.0 h of them spent building rings, and the mid-ranging limit stays idle until the screw runs short. The settlement scale omits the losses no face pressure controls.

The optional MPC reduces the boundary band integral but fails the faster challenge on nine seconds of blow-out risk. A fuller model would add a tail-void and shield-gap loss fitted to published volume losses, a part-full chamber and a building that interacts with the trough, and improved predictive policies would sharpen the comparison. As it stands, the simulator is a plant for practicing inventory pressure control through forced stops against a product measured late.

KEY POINTS

- A shield holds up the ground with the muck in its chamber, and that pressure is an inventory that nothing returns once it drifts.
- Every ring build seals the chamber for twenty minutes, in which the sand drops it from 2.15 to 1.52 bar unless a bentonite hold keeps it over 2.06 bar.
- Only a setpoint scheduled from the geology passes both the land and the river.
- Layered control drives the 100 m in 46.9 h against 57.6 h.

Where else this applies. Chapter 7 caps a drill's thrust with pressure and torque overrides in the same kind of low selector, and Chapter 10 stacks three mid-ranging layers on one power block.

References

[1] X. Liu and K. Zhang, "Earth pressure balance control of shield tunneling machine based on nonlinear least squares support vector machine model predictive control," *Measurement and Control*, vol. 52, no. 1–2, pp. 3–10, 2019.

[2] C. J. Hung, J. Monsees, N. Munfah, and J. Wisniewski, "Technical manual for design and construction of road tunnels – civil elements," Federal Highway Administration, U.S. Department of Transportation, Washington, DC, Publication FHWA-NHI-10-034, 2009. [Online]. Available: https://www.fhwa.dot.gov/bridge/tunnel/pubs/nhi09010/tunnel_manual.pdf

[3] G. Anagnostou and K. Kovári, "Face stability conditions with earth-pressure-balanced shields," *Tunnelling and Underground Space Technology*, vol. 11, no. 2, pp. 165–173, 1996.

[4] J. Holzhäuser and M. Kuder, "Settlement predictions for TBM tunnelling in soft ground on the basis of empirical values," *tunnel*, no. 2, 2024. [Online]. Available: https://www.tunnel-online.info/en/artikel/settlement-predictions-for-tbm-tunnelling-in-soft-ground-on-the-basis-of-empirical-values-4092239.html

[5] M. D. Boscardin and E. J. Cording, "Building response to excavation-induced settlement," *Journal of Geotechnical Engineering*, vol. 115, no. 1, pp. 1–21, 1989.

[6] A. S. Merritt and R. J. Mair, "Mechanics of tunnelling machine screw conveyors: model tests," *Géotechnique*, vol. 56, no. 9, pp. 605–615, 2006.

[7] D. Peila, C. Oggeri, and R. Vinai, "Screw conveyor device for laboratory tests on conditioned soil for EPB tunneling operations," *Journal of Geotechnical and Geoenvironmental Engineering*, vol. 133, no. 12, pp. 1622–1625, 2007.

[8] D. Sebastiani, S. Miliziano, and A. Bezuijen, "Front-face pressure drop during the standstill phase for EPB mechanized tunnelling in coarse-grained soils," in *Geotechnical Aspects of Underground Construction in Soft Ground*. CRC Press, 2021, pp. 408–414.

[9] H. Yu, M. A. Mooney, and A. Bezuijen, "A simplified chamber pressure model for EPB TBM tunneling in granular soil," in *Geotechnical Aspects of Underground Construction in Soft Ground*. CRC Press, 2017, pp. 127–135.

[10] B. J. Allison and A. J. Isaksson, "Design and performance of mid-ranging controllers," *Journal of Process Control*, vol. 8, no. 5–6, pp. 469–474, 1998.

[11] T. Xu, A. Bezuijen, and M. Thewes, "Pressure infiltration characteristics of foam for EPB shield tunnelling in saturated sand – part 2: soil–foam mixture," *Géotechnique*, vol. 72, no. 4, pp. 295–308, 2022.

[12] C. Shao and D. Lan, "Optimal control of an earth pressure balance shield with tunnel face stability," *Automation in Construction*, vol. 46, pp. 22–29, 2014.

[13] B. Maidl *et al.*, *Mechanised Shield Tunnelling*. Wiley, 2012.

[14] S. Babendererde, "Tunnelling machines in soft ground: a comparison of slurry and EPB shield systems," *Tunnelling and Underground Space Technology*, vol. 6, no. 2, pp. 169–174, 1991.

[15] G. Lehmann and K. Bäppler, "Development tendencies in mechanised tunnelling – review and development tendencies for shield machines with slurry and earth pressure support," *Geomechanics and Tunnelling*, vol. 3, no. 3, pp. 245–255, 2010.

[16] N. Loganathan and H. G. Poulos, "Analytical prediction for tunneling-induced ground movements in clays," *Journal of Geotechnical and Geoenvironmental Engineering*, vol. 124, no. 9, pp. 846–856, 1998.

[17] R. J. Mair, "Tunnelling and geotechnics: new horizons," *Géotechnique*, vol. 58, no. 9, pp. 695–736, 2008.

[18] P. W. Mayne and F. H. Kulhawy, "K_o–OCR relationships in soil," *Journal of the Geotechnical Engineering Division*, vol. 108, no. 6, pp. 851–872, 1982.

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

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

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

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

[23] D. Peila, D. Martinelli, C. Todaro, and A. Luciani, "Soil conditioning in EPB shield tunnelling – an overview of laboratory tests," *Geomechanics and Tunnelling*, vol. 12, no. 5, pp. 491–498, 2019.

[24] J. Lach *et al.*, "Instrumentation of a reconditioned Robbins tunnel boring machine – 106th street tunnel," Fermi National Accelerator Laboratory, Batavia, IL, Technical Memo Fermilab-TM-2141, 2000.

[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.

11.11 Exercises

11.1 Concept check · 2 min

In case 1 the screw runs 1 rpm slower than the balance for two minutes, and the axis cell climbs from 2.397 to 2.754 bar. Balance the chamber then sets the screw to the speed that takes out what the advance brings in, and five minutes later the cell reads 2.767 bar. Why does the pressure not return to 2.40 bar?

(a) The screw drive's 3 s lag delays the return by several minutes

(b) The clay pushes water into the chamber, which holds the pressure up

(c) The pressure cells hold their old reading until the survey levels the face

(d) The pressure integrates inflow less outflow, so balanced flows stop it where it is

11.2 Concept check · 2 min

In case 7 a pressure loop tuned tight in the clay, 80 rpm/bar with a 20 s integral time, holds the axis cell within 0.028 bar. After the face crosses into the sand at the same 40 % foam, the screw swings between 0.3 and 17.0 rpm. What changed?

(a) The muck's stiffness about doubled, so each rpm moves the pressure about twice as fast

(b) The sand's seepage added a dead time between the screw and the pressure cells

(c) The screw's plug failed in the sand, so the screw speed no longer sets the outflow

(d) The chainage feedforward stepped the setpoint at the boundary and the loop overshot it

11.3 Concept check · 2 min

In case 5 a fixed setpoint of 2.40 bar settles the street 100.5 mm under the river. A fixed 2.88 bar still settles 11.8 mm, a fixed 2.94 bar spends 48.2 min at blow-out risk, and the chainage feedforward settles nothing. Why can no fixed setpoint pass both the land and the river?

(a) The plug cannot hold more than about 2.9 bar in sand, so a higher setpoint spews

(b) The river's target is level with the land's blow-out limit at the axis, so no flat line clears one and stays under the other

(c) The screw saturates under the river, so the loop cannot reach a higher setpoint

(d) The survey reports the river's settlement too late to correct a fixed setpoint

11.4 Theory and calculation · 6 min

The face's target and the blow-out limit, from Equation (11.1). The axis is 15 m deep, the water table 3 m below the street, the river 8 m deep, the shield's radius 3.25 m, and $\rho_w g = 9.81\,\mathrm{kN/m^3}$. The clay has $\gamma = 19.5\,\mathrm{kN/m^3}$ and $K_0 = 0.60$, the sand $\gamma = 20.0\,\mathrm{kN/m^3}$ and $K_0 = 0.45$.

(a) Find the vertical stress, the pore pressure and the target p_t at the axis in the clay on land. *Checkpoint:* 2.925 bar, 1.177 bar, 2.226 bar

(b) Find the target in the sand on land and under the river, and the rise between them. *Checkpoint:* 1.997 bar, 2.944 bar, 0.947 bar

(c) Find the crown's overburden $\sigma_{v,crown}$ in the sand on land. The blow-out flag compares the crown's chamber pressure $p - \rho_{mix} g R$ with it. At what axis pressure does the flag rise when the muck's density is 1734 kg/m^3, and how does that compare with the river's target? *Checkpoint:* 2.350 bar, 2.903 bar, 0.041 bar under it

11.5 Theory and calculation · 7 min

The screw as an integrator, from Section 11.4. One rpm of screw removes $\eta V_{rev} = 0.2625\,\mathrm{m^3}$ a minute from a chamber of $V_c = 30\,\mathrm{m^3}$.

(a) In case 1 the stiffness is $B = 21.3$ bar. Find the rate at which a 1 rpm mismatch moves the pressure, in bar a minute, and the time it takes to carry the axis cell from 2.40 to 2.91 bar. *Checkpoint:* 0.186 bar a minute, about 2.7 min

(b) Find the same rate in Pa/s per rpm in the clay at $B = 18.82$ bar and in the sand at $B = 37.72$ bar. By what factor does the loop gain of a PI with fixed tuning change at the boundary? *Checkpoint:* 274 and 550 Pa/s per rpm, a factor of 2.0

(c) A step test fits an integrating-plus-dead-time model with $k' = -0.002\,93$ bar/s per rpm and an effective delay of 4.4 s. Apply the SIMC rule for an integrating process with $\lambda = \theta$: $K_c = 1/[|k'|(\lambda + \theta)]$, $\tau_I = 4(\lambda + \theta)$. *Checkpoint:* 38.8 rpm/bar, 35.2 s

11.6 Theory and calculation · 6 min

The settlement trough, from Equation (11.6). A loss V_s per meter of tunnel spread evenly along the drive settles the axis by $S_{max} = V_s/(\sqrt{2\pi}\,i)$, with $i = Kz_0$, $z_0 = 15$ m, and the face area is $A = 33.18\,\text{m}^2$.

(a) Find S_{max} for a loss of 1 % of the face area in the clay ($K = 0.50$) and in the sand ($K = 0.35$). *Checkpoint:* 17.65 mm, 25.2 mm

(b) Find the settlement at a distance i and $2i$ from the trough's axis as fractions of S_{max}. How far from the axis is it half of S_{max} in the clay? *Checkpoint:* 0.607 and 0.135, about 8.8 m

(c) Repeat part (a) for the largest reported volume loss of 0.83 %. *Checkpoint:* 14.7 mm, 20.9 mm

11.7 App challenge · 8 min

Scenario 2, *The pressure loop on the screw*, default seed. Run it to 1500 s (25 min) at the case's 60× (about 25 s) and export the CSV file. The advance steps from 50 to 75 mm/min at 900 s.

(a) As shipped, with *Advance feedforward to the screw* Off, find the largest value of `p_mid_bar` less `p_sp_bar` after 900 s and its time, and the last row in which the two differ by more than 0.02 bar. *Checkpoint:* 0.177 bar at 940 s, 1120 s

(b) Set *Advance feedforward to the screw* to On, press Reset, run and export again. Find the largest deviation after 900 s, and `screw_rpm` at 890 and 1490 s. *Checkpoint:* 0.021 bar, 6.63 and 9.97 rpm

(c) From Equation (11.7), compute the change in ω_{ff} for the step, with $p_{mid} = 2.40$ bar, FIR 30 %, $B_c = 40$ bar, FER 10, $p_{atm} = 1.013$ bar, $A = 33.18\,\text{m}^2$ and $\eta V_{rev} = 0.2625\,\text{m}^3$. Compare it with `screw_ff_rpm` in the second file. *Checkpoint:* 3.31 rpm

Tolerance: 0.005 bar, 10 s and 0.03 rpm.

11.8 App challenge · 8 min

Scenario 3, *The ring build: twenty minutes with nothing coming in*, default seed. Run it to 2400 s (40 min) at the case's 60× (about 40 s) and export the CSV file.

(a) As shipped, with *Bentonite hold in the ring build* Off, find the rows in which `gate` is 0, the first row in which `truth_p_bar` is under `truth_target_bar`, and the lowest `p_mid_bar` while the gate is shut. *Checkpoint:* 230 to 1420 s, 370 s, 1.519 bar

(b) Read the scorecard tiles *Street · drive (true)* and *Face pressure deficit (true)*. *Checkpoint:* 1.8 mm at 57.5 m; 0.48 bar worst, 20 min under

(c) Set the hold to On, press Reset, run and export again. Find the lowest `truth_p_bar` while the gate is shut, the peak of `bentonite_m3h`, and the volume injected, the sum of `bentonite_m3h` times 10 s over 3600 s. *Checkpoint:* 2.056 bar, 5.78 m^3/h, 1.10 m^3

Tolerance: 10 s, 0.01 bar, 0.1 mm, 0.05 m^3/h and 0.05 m^3.

11.9 App challenge · 9 min

The challenge, *the 100 m drive, the street kept*, default seed, at 7200×. Each run pauses with the verdict when the survey has reported every point, after about 24 to 29 s.

(a) Run as shipped and read the verdict: worst settlement, minutes of spewing and of blow-out risk, peak torque and drive time. *Checkpoint:* 245.3 mm, 2079.1 min, 1675.3 min, 96.1 %, 57.6 h

(b) Press Reveal suggested tuning, then Reset and Run. Read the verdict, the clock when it arrives and the *Survey (reported)* tile. Why can the survey report more than the true settlement? *Checkpoint:* met, 0.0 mm, 46.9 h, 86.1 %; about 48.0 h; 1.0 mm at 87.5 m

(c) Set *Bentonite hold in the ring build* to Off, press Reset and Run. Read the verdict and the *Face pressure deficit (true)* tile. *Checkpoint:* not met, 21.7 mm at 97.5 m; 0.78 bar worst, 710 min under

Tolerance: 0.1 mm, 0.1 h and 0.1 min on the verdict values.

11.10 Application area · 4 min

The suggested drive settles the street 0.0 mm (true), and a project engineer wants to quote that number to the owner of the building over 22 to 38 m.

(a) Why is 0.0 mm no prediction of a real drive?

(b) Using Peck's formula of Exercise 11.6 and the published median volume loss of 0.41 % of the face area, what would a real drive add over the clay, and how much of the 15 mm budget would that leave for losses at the face? *Checkpoint:* 7.2 mm, 7.8 mm

(c) Where does a real settlement budget come from?

11.11 Application area · 4 min

In case 8 a loose pocket begins to run into the face at 62.0 min. The over-excavation detector alarms at 65.8 min, and the survey's first report of the damage, 15.4 mm at the 35 m point, comes at 140 min.

(a) What makes up most of the 228 s between the void and the alarm?

(b) How much earlier than the survey does the belt scale warn, and what fixes the survey's delay?

(c) The pressure loop holds its setpoint through the void. Why does it not see the over-excavation?

11.12 Review the agent · 5 min

An AI assistant reviews case 4 at its 8 h pause and writes: "The survey's first report of the 27.5 m point, the middle of the building, reads 12.3 mm, under the 15 mm budget. The 2.05 bar setpoint protects the building, and no change is needed." Find the flaw and say how you would check it in the application.

CHAPTER 12

Allocating Tires, Suction Fans and Cold-Gas Thrusters on One Nitrogen Tank

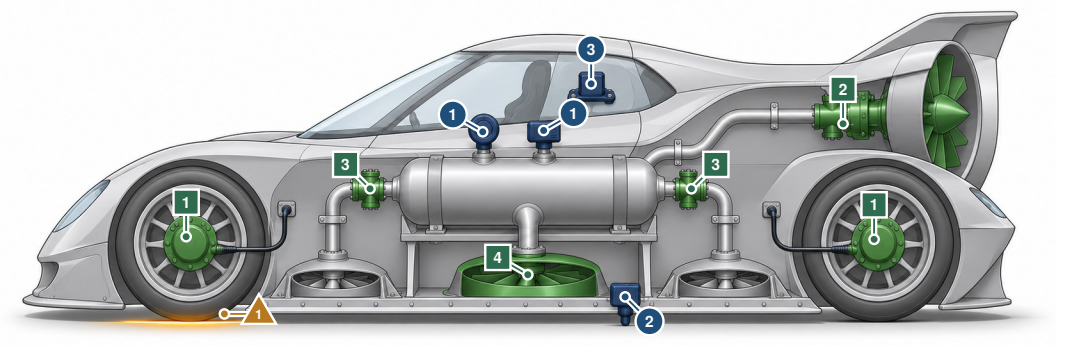

The controller moves 1 motors that drive the wheels 2 valves of the rear thrusters 3 valves of the corner lift thrusters 4 fan that sucks the floor down

The controller reads 1 tank pressure and temperature 2 height of the floor 3 tilt and turn of the body

What disturbs it 1 a wet patch on the strip

Notice that one tank feeds the rear and the corner thrusters, so the gas a fast launch spends is gone for the flight. Illustration, not to scale.

IN PLAIN TERMS **Three Ways to Push, One Tank of Gas**

The main idea. A test car pushes itself along with its tires, with fans that suck it onto the road for more grip, and with cold gas jets from one tank. The controller uses the free tires first and spends gas only on what they cannot give.

The challenge. The tank cannot be refilled, so gas spent on a fast start is gone for the flight. Escaping gas cools, its pressure falls faster than its amount, and the jets weaken while the gauge still shows gas left.

Why it matters. A machine that draws several actions from one store must keep its reserve in what runs out first.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact in the Introduction (Section 12.1), the history segment (Section 12.2), the try-it segment on the inverted ramp (Section 12.4), the try-it segment and the fact in the results (Section 12.6) and the key points
Study	50 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 12.6), Technology in Practice (Section 12.7), Limitations (Section 12.8) and the open challenge

The control question A car is pressed onto the road by suction fans and lifted by cold-gas thrusters, so free tires, cheap fans and costly nitrogen share one demand. What does each allocation choice buy and cost from a tank that cannot be refilled, and why does that tank run out by pressure before it runs out of mass?

Learning objectives

- Compute the suction of a sealed skirt and the tire force it makes available, and explain why the seal is a positive feedback.
- Allocate a launch to the free actuator first and the costly one on the remainder, and price a g target in nitrogen.
- Compute the gas cost of a hover and the runway of a tank, and explain why a runway taken from the gauge's mass is optimistic.
- Compare an integral alone with a mass feedforward on a vehicle that grows lighter as it flies.
- Recognize a demand the tank cannot meet, and set the reserve on the quantity that runs out first.

Prerequisites Sections 1.4.6 (split range), 1.4.8 (allocation), 1.4.9 (depleting actuators), 1.4.5 (an actuator that loses its gain) and 1.6.2 (velocity-form PID).

The app https://apmonitor.com/apps/performance-auto/

See also Volume 20, Chapter 4 for allocation among rotors and a wing whose authority moves with airspeed, Volume 20, Chapter 3 for a load feedforward on a multirotor, and Volume 20, Chapter 7 for actuators that store what they apply and must be emptied.

12.1 Introduction

The app of this chapter, Performance Auto, takes its premise from a 2026 public demonstration of a production-car prototype, a road car with ducted fans that hold it down and compressed nitrogen that lifts it. The generic car built here weighs 1723 kg with 423.0 kg of nitrogen in an 850 L tank at 700 bar. On its tires and fans it reaches 60 mph in 1.47 s. With both rear thrusters wide open it gets there in 0.60 s, at a peak of 4.83 g and for 40.5 kg of nitrogen, while thrust metered to 3.4 g takes 0.83 s and 29.1 kg. A 5 s hover on a full tank costs 114.1 kg, more than a quarter of the fill.

Problems in which several actuators of different speed, limit and cost share one demand are called control allocation, and Section 1.4.8 sets out the general problem. They began with aircraft that carry more control surfaces than axes to control [1] and now cover ships, road vehicles and spacecraft, with costs, limits and failures built into the allocation law [2]. A finite onboard consumable adds a budget to the allocation, the depleting actuator of Section 1.4.9. Cold-gas thrusters, the simplest rocket, are still chosen for small satellites because they are clean and reliable [3, 4], and every newton-second they deliver is drawn from a pressure vessel that empties as it cools. A process engineer meets the same shape in split-range control of a cheap and an expensive utility [5] (Section 1.4.6), and in a plant whose inventory must last a campaign.

✦ DID YOU KNOW? Cold-Gas Thrusters Start at Micronewtons

A 2025 survey of small-spacecraft propulsion lists cold-gas thrusters from 10 µN to 3.6 N, with specific impulses of 40 to 110 s [6]. Each rear thruster of this chapter's car is rated at 25 kN, at an assumed 73 s.

The car is built to correct two beliefs. The first is that the strongest actuator should be used whenever it helps. Opening the thrusters wide saves 0.23 s on the metered launch, spends 11.4 kg more nitrogen and passes the 3.5 g teaching limit. The second is that a tank gauge reports what is left to use. The gas cools as it expands, its pressure falls faster than its mass, and the thrusters are rated at the regulator's pressure, so a 12 s hover ends in a landing at 2.83 m/s with 22.6 % of the nitrogen still in the tank. A transfer function cannot show either mistake, because one is a choice among actuators and the other is a state of the actuator's supply.

Teaching material usually treats these pieces apart. Tire slip is taught on a braking corner with a fixed normal force [7], thruster modulation on a spacecraft in free space [8], and descent guidance on a vehicle whose propellant is not scored [9]. This chapter works with one simulated car that is a ground vehicle and an aircraft in the same run, and it quantifies what each allocation choice buys and what it costs from a single tank. The tire curve, the velocity-form proportional-integral-derivative (PID) loop, the pulse-width modulation of the valves and the descent schedule are established methods. The results are measurements on a single declared car, among them the findings that the tank runs out by pressure before mass, that a bang-bang roll entered at 5.1 m/s of climb ends 1.77 m above its entry rather than level in height, and that metering the launch thrust saves 28 % of its gas at a cost of 0.23 s. The model has 24 states and a tank balance that closes to 4×10^{-12} kg, its gas law reproduces the reference density at the fill to 0.12 %, and a strict split separates what the laws may read from what only the scorecard sees.

Section 12.2 traces the domain: downforce, cold-gas propulsion, allocation and the few control studies of vehicles that both drive and fly. Section 12.3 gives the car, from the tire and the skirt to the tank and the instruments. Section 12.4 draws from the model the three prices and the one budget that shape every choice, and Section 12.5 describes the laws as implemented. Section 12.6 measures ten cases, from a hand-flown launch to a scored one-tank flight. Section 12.7 sets the car's devices beside real hardware, Section 12.8 states the limitations and Section 12.9 concludes. The exercises at the end of the chapter run the cases on the app.

12.2 Suction, Cold Gas and the Allocation of Thrust

Downforce entered racing as a way to add normal force without adding mass. Inverted wings came first, then underbody shapes that accelerate the air between the floor and the track, and in some cars fans that evacuate a sealed space under the body so that the grip no longer depends on speed [10]. Fan suction raced briefly in the 1970s before the rules removed it, while underbody downforce remains the main source of grip in open-wheel racing [10]. The physics of the seal is older and comes from the air-cushion vehicle, whose fans pressurize a skirted cushion and whose fan flow is set by the leakage under the skirt hem [11]. The fan car runs that theory in reverse, with the cushion below ambient and the same leakage, so a hem that lifts a few centimeters lets the suction go.

Cold-gas propulsion is the simplest form of reaction control. A regulated gas is expanded through a nozzle without combustion, and nitrogen gives a specific impulse of the order of 70 s, low against chemical rockets but clean, cheap and safe to handle [12]. Small satellites still carry such systems for attitude control and fine maneuvers [3, 4]. Their tanks are blowdown vessels, and a vessel emptied in seconds behaves almost adiabatically, so the gas left inside cools by tens of kelvins and its pressure falls faster than its mass [13]. At 700 bar nitrogen is far from ideal, and the tables that a simulator can be checked against come from a reference equation of state [14]. On-off thruster valves cannot open for less than a minimum time, and pulse-width pulse-frequency modulation was developed to deliver a commanded average impulse from such valves [8]. The same quantization was flown with a pilot aboard, as the history segment after this paragraph tells.

> ⧖ HISTORY **1967: Rocket Pulses That Hunted in Flight**
>
> On-off thrusters were flown with a pilot aboard long before this car. A free-flight lunar-landing research vehicle carried a vertically mounted jet engine to create a pseudolunar gravity field, and its attitude rockets were fired by on-off solenoid valves. A switching amplifier opened them when the error between the pilot's rate command and the measured rate crossed a threshold. A describing-function analysis mapped the rate dead bands and control authorities that would give a limit cycle, and flight and ground tests in general agreed with it. Pilots found the cycling did not spoil control, but the noise of the rockets firing at about 2 Hz distracted them. In a fixed-base simulator the attitude propellant was least at one middle dead band and rose at smaller ones, partly through limit-cycle operation [15].

Allocation as a control problem dates from aircraft with redundant control surfaces, where the attainable moment set and its saturation limits were analyzed explicitly [1]. Later work added actuator rates, costs, priorities and failures, and today's allocation laws in marine, automotive and aerospace practice share that formulation [2]. The descent to a landing has its own lineage. The Apollo lunar module flew a guidance law that shaped its altitude and vertical speed toward a gentle touchdown [16], and modern powered descent poses the same problem as a convex trajectory optimization with propellant as the cost [9]. On the ground the friction available to a driven wheel peaks at a small slip and falls past it, the shape that traction control, launch control and anti-lock braking all regulate against [7], and the slip builds over a rolling distance, the tire's relaxation length [17]. Every limit the simulator scores is a teaching value. Real acceleration limits come from human-tolerance data [18], and real touchdown limits from the sink rate a landing gear is designed to absorb [19].

Vehicles that both drive and fly have a control literature of their own, though a short one. A winged flying-car prototype has flown under a single nonlinear position and attitude law whose forces and moments a real-time allocation spreads over distributed propellers and control surfaces [20]. A fly-drive vehicle has been modeled with its rotors, tires, chassis and suspension in one set of equations, and in hardware-in-the-loop simulation its rotors steady the suspension after touchdown through a predictive controller over an allocation layer [21]. A land-air simulation shows the wheel loads falling to zero through the switch from driving to flight and an oscillation once the car is clear [22], and a hybrid vehicle has tracked ground and air trajectories under one adaptive attitude law in experiments [23]. Where the allocation carries a price, a flight-tested law for an over-actuated transition aircraft minimizes total power in place of a norm of the commands [24], and Volume 20, Chapter 4 allocates among the rotors, wing and propeller of a lift-plus-cruise aircraft whose effectors gain and lose authority with the airspeed. In process systems engineering the subject is thin. A search of seven of its journals returned no paper on flying cars, roadable aircraft or fly-drive vehicles, and the nearest works are of two kinds. One is nonlinear model predictive control of a solar aircraft that stores energy in batteries and in altitude [25]. The other is split-range control, which extends the range of one controlled variable by moving through several inputs one at a time [5]. The car sits between them. Its launch is a split range in the sense of Section 1.4.6, tires before thrust, in which the costly input draws on a stock that nothing in the run can replenish, so every allocation choice is also an inventory decision.

12.3 The Car as Modeled

The simulated car is generic and every number is a declared teaching value, as shown in Figure 12.1. A coupe of 1300 kg dry mass rides on four corner springs at $\pm a$ along and $\pm t$ across. One lumped driven wheel of radius R stands for all four. A skirt seals a floor of area A_{floor} to the road, and fans evacuate it. Four corner lift thrusters of 7 kN and two rear thrusters of 25 kN each are fed from one nitrogen tank through a regulator and valves. The car's computer reads only its instruments: wheel speed, a ground-speed radar, a seat accelerometer, lasers for the height, an inertial unit, an optical position fix, the tank's pressure and temperature, and a fan tachometer.

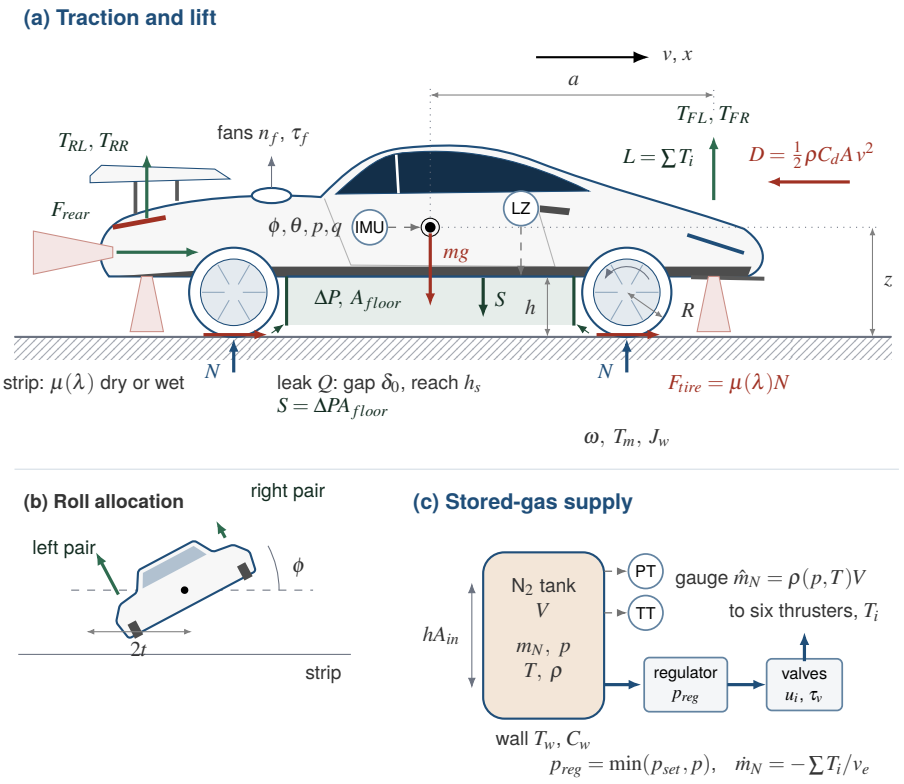

Figure 12.1 Fan car as modeled, in side elevation at the metered launch, with the nitrogen supply and a rear view for the roll. The weight mg, the suction S of the sealed floor and the corner lift L set the road's normal force N, which limits the tire force F_{tire}. The rear thrust F_{rear} and the drag D act along the strip, and z measures the center of mass above the strip. The tank holds m_N at p and T inside a wall at T_w, and the regulator feeds the valves at p_{reg}. LZ is the height laser, IMU the inertial unit, and PT and TT the tank transmitters that drive the gauge. Not to scale.

The variables are classified in Table 12.1. Twenty-four states are integrated: the position and velocity of the center of mass in three axes, the roll ϕ and pitch θ with their rates p and q, the wheel speed ω, the tire slip λ, the motor torque T_m, the fan speed n_f, five valve positions, the nitrogen mass m_N, the gas temperature T, the wall temperature T_w, and two ledgers that integrate the impulse delivered and any leak, so that the tank balance can be checked exactly. Units are seconds, kilograms, meters, newtons, bar absolute and kelvins, with forces quoted in kilonewtons.

12.3.1 The Wheel, the Tire and the Springs

The motors drive the lumped wheel through a first-order torque lag, and the tire pushes with the friction its slip allows on the normal force the springs carry, as shown in Equation (12.1).

$$\begin{aligned} J_w\dot{\omega} &= T_m - RF_x, \qquad F_x = \mu(\lambda)N \\ \sigma\dot{\lambda} &= (R\omega - v) - \max(|v|, v_{min})\,\lambda \\ \tau_m\dot{T}_m &= T_{cmd} - T_m \\ |T_{cmd}| &\le \min\big(F_{mot}R,\ P_{max}/\max(1,|\omega|)\big) \\ \mu(\lambda) &= c_1\big(1 - \mathrm{e}^{-c_2|\lambda|}\big) - c_3|\lambda|, \quad |\lambda| \le 1 \end{aligned} \tag{12.1}$$

Table 12.1 Variables by Role, With Their Values in the Metered Launch of Case 4 at $t = 1.4\,\text{s}$

Variable	Role	At 1.4 s	Range
λ	CV, tire slip	0.047	SP 0.02–0.5
h	CV, wheel height	—	SP 0.5–4 m
ϕ	CV, roll (a profile)	0°	—
T_{cmd}	MV, motor torque	100 %	± motor limit
u_{rear}	MV, rear thrust	25.5 kN	0–50 kN
u_i	MV, corner valves	0	0–1 each
n_{cmd}	MV, fan speed	1	0–1
surface	DV, wet patch	dry	10 m, peak 0.80
F_g	DV, side gust	0	0–6 kN
b, stuck, leak	DV, bump, valve, regulator	none	see text
v	state, speed	27.6 mph	—
N, S	output, road, suction	37.5, 21.1 kN	—
F_{tire}	output, tire force	31.6 kN	$\le \mu(\lambda)N$
m_N	state, nitrogen	409.4 kg	0–423.0 kg
p, T	output, state, tank	621 bar, 11 °C	to 5 bar
a_x	output, seat g (8 Hz)	3.38 g	limit 3.5 g

where $\sigma = 0.25\,\text{m}$ is the relaxation length and $v_{min} = 1\,\text{m/s}$ keeps the slip equation damped at a standstill. The Burckhardt coefficients are those of the anti-lock braking app of Volume 7, Chapter 4, $(c_1, c_2, c_3) = (1.2801, 23.99, 0.52)$ on dry asphalt and $(0.857, 33.822, 0.347)$ on wet, and both apps import the same functions from the framework's shared tire module. The dry peak is 1.170 at $\lambda = 0.170$ and the locked-wheel value 0.760. The slip is the traction form with the car's speed in the denominator, so μ is held at its value for $\lambda = 1$ beyond it, and λ is clipped to ± 1.5.

The four corner springs carry the weight, the suction and whatever the lift takes off, as shown in Equation (12.2).

$$\begin{aligned} \delta_i &= z_{free} - (z + x_i \sin\theta + y_i \sin\phi), \qquad N = \textstyle\sum_i F_i \\ F_i &= \max\!\left(0,\; k\delta_i + c\dot{\delta}_i\right) + 20k \max(0, \delta_i - 0.12) \end{aligned} \tag{12.2}$$

where $(x_i, y_i) = (\pm a, \pm t)$ and the last term is a bump stop 12 cm into the travel. A spring pushes only while compressed and only while the car is roughly upright ($\cos\phi > 0.5$ and $\cos\theta > 0.5$). At rest with a full tank the springs carry 16.90 kN and the center of mass sits at $z = 0.455\,\text{m}$. The tire's force acts at the road, so it pitches the body by $F_x z$, the squat of a launch, and a stiff saturating friction $-0.9N\,\text{sat}(v_y/0.05)$ holds the car laterally while it is in contact.

12.3.2 Suction Under the Skirt

The fans pull air from under the floor, and the outside air returns through the gap under the skirt. Equating the fan curve $\Delta P = \Delta P_{max} n_f^2 - k_Q Q^2$ to the orifice leak $Q = A_{gap}\sqrt{2\Delta P/\rho}$ gives the closed form of Equation (12.3).

$$\begin{aligned} \Delta P &= \frac{\Delta P_{max}\, n_f^2}{1 + 2k_Q A_{gap}^2/\rho}, \qquad S = \Delta P A_{floor} \\ A_{gap} &= L_s\left(\delta_0 + \max(0,\; h - h_s) + b\right) \\ \tau_f \dot{n}_f &= n_{cmd} - n_f \end{aligned} \tag{12.3}$$

where $h = z - d_{floor}$ is the floor's clearance, $h_s = 15\,\text{cm}$ the skirt's reach, $\delta_0 = 5\,\text{mm}$ the sealed leak around $L_s = 12\,\text{m}$ of skirt, and b the extra gap of a bump, 7 mm for 0.3 s. Sealed at full speed the pressure is 6.00 kPa and the suction 23.99 kN, 1.42 times the full-tank weight, for about 60 kW of fan power at an assumed efficiency of 0.6. The seal is positive feedback. A floor 2 cm above the skirt's reach keeps 42.8 % of the sealed suction and one at 5 cm keeps 13.0 %, so a car that starts to lift loses the force that held it. The suction acts at the floor's center and makes no moment.

12.3.3 Flight in Five Degrees of Freedom

Pitch is applied first and roll second, with no yaw, so the body's longitudinal and vertical axes in the world are $\mathbf{e}_x = (\cos\theta, 0, \sin\theta)$ and $\mathbf{e}_z = (-\sin\theta\cos\phi, -\sin\phi, \cos\theta\cos\phi)$. The translation and the two rotations are given by Equation (12.4).

$$\begin{aligned} m\ddot{\mathbf{r}} &= L\,\mathbf{e}_z + F_{rear}\,\mathbf{e}_x + (F_x - F_r)\mathbf{e}_X + F_y\mathbf{e}_Y \\ &\quad + (N - S)\,\mathbf{e}_Z - mg\,\mathbf{e}_Z - \mathbf{D} + F_g\,\mathbf{e}_Y \\ I_{xx}\ddot{\phi} &= t\,(T_{FL} + T_{RL} - T_{FR} - T_{RR}) + \textstyle\sum_i y_i F_i - 0.25 F_g \\ I_{yy}\ddot{\theta} &= a\,(T_{FL} + T_{FR} - T_{RL} - T_{RR}) + \textstyle\sum_i x_i F_i + F_x z \end{aligned} \tag{12.4}$$

where $L = \sum T_i$ over the corners, $F_r = f_r N$ is the rolling resistance, and the drag components are $\frac{1}{2}\rho A_j v_j |v_j|$ with $A_x = C_d A = 0.6\,\mathrm{m}^2$, $A_y = 3.6\,\mathrm{m}^2$ and $A_z = 8\,\mathrm{m}^2$. The gust F_g is a 2 s half sine that starts 1 s into the hover and acts 0.25 m above the center of mass. Roll and pitch are integrated as independent axes, which is exact for a roll with the pitch held at zero. The mass and inertias follow the tank, $m = m_{dry} + m_N$ and $I = I_{full}\, m/m_{full}$.

12.3.4 Thrusters, Valves and the Tank

Each thruster delivers its rating times its valve position, scaled by the regulator, and every newton costs mass at the exhaust speed, as shown in Equation (12.5).

$$\begin{aligned} T_i &= T_{rated}\,u_i\,\frac{p_{reg}}{p_{set}}, \qquad p_{reg} = \min(p_{set},\, p) \\ \tau_v \dot{u}_i &= u_{i,cmd} - u_i \\ \dot{m}_N &= -\frac{\sum T_i}{v_e} - \dot{m}_{leak}, \qquad v_e = I_{sp}\, g_0 \end{aligned} \tag{12.5}$$

where $p_{set} = 60\,\mathrm{bar}$ is the pressure at which the thrusters are rated, $v_e = 716.1\,\mathrm{m/s}$ for $I_{sp} = 73\,\mathrm{s}$, and τ_v is 20 ms for proportional valves and 5 ms for on-off valves. Below 5 bar the valves pass nothing and TANK EMPTY shuts every thruster. Holding the full-tank weight of 16.90 kN costs 23.6 kg/s, and 50 kN of rear thrust costs 69.8 kg/s. A yaw pair of up to 2 kN is assumed to hold the heading, and its gas is counted while a gust acts.

The tank's pressure comes from a real-gas law and its temperature from the energy balance of the gas left behind, as shown in Equation (12.6).

$$\begin{aligned} p &= \rho R_{N_2} T\, Z(\rho, T), \qquad \rho = m_N / V \\ Z &= 1 + \textstyle\sum_{k=1}^{5} r^k \left(a_{k0} + a_{k1}\tau + a_{k2}\tau^2\right) \\ m_N c_v \dot{T} &= hA_{in}(T_w - T) - \dot{m}_{out}\,\frac{T}{\rho}\left(\frac{\partial p}{\partial T}\right)_{\rho} \\ C_w \dot{T}_w &= hA_{in}(T - T_w) + hA_{out}(T_{air} - T_w) \end{aligned} \tag{12.6}$$

where $r = \rho/(1000\,\mathrm{kg/m^3})$, $\tau = 100\,\mathrm{K}/T$, and $c_v = e_0 + e_1 r + e_2 r\tau + e_3 r^2 + e_4 r^2\tau + e_5 r^3$ in J/(kg K). The fifteen a_{kj} and six e_j are listed in Table 12.2. They are least-squares fits to the reference tables for nitrogen [14] between 140 and 320 K and 5 and 705 bar, with worst errors of 0.9 % in pressure and 0.5 % in density. At 700 bar and 20 °C the fit gives 497.6 kg/m^3 against the tables' 498.2, so the tank holds 423.0 kg, of which 4.9 kg stays behind at 5 bar. For an ideal gas the last term of the energy balance reduces to $\dot{m}_{out} R_{N_2} T$, the familiar adiabatic cooling. With the wall switched off the blowdown follows the reference isentrope from the fill within 0.7 K down to 200 kg/m^3. The gauge inverts the same law at the measured pressure and temperature, $\hat{m}_N = \rho(\hat{p}, \hat{T})V$, by twelve fixed Newton iterations.

Table 12.2 Coefficients of the Nitrogen Fit of Equation (12.6)

k	a_{k0}	a_{k1}	a_{k2}
1	1.244544798	−2.619325373	−4.784349976
2	5.289904277	−20.10382874	27.53234519
3	−8.408539214	50.59369007	−58.79664928
4	10.45578251	−50.99223434	37.20049222
5	−0.337057974	23.03951548	−0.8679504764

c_v	e_0, e_1	e_2, e_3	e_4, e_5
	727.51089, 0.16531432	1698.7723, −930.60672	−2469.4212, 1991.6526

Table 12.3 Model Parameters. H Is Hidden and Fixed, A Is Adjustable in the Advanced Parameters Panel

Symbol	Meaning	Value	Unit	
m_{dry}	car, driver, empty tank	1300	kg	A
I_{xx}, I_{yy}	inertias, full tank	550, 2200	kg m^2	A
a, t	half wheelbase, half track	1.35, 0.90	m	H
z_{free}	CG height, springs unloaded	0.49	m	H
k, c	corner spring, damper	120, 9	kN/m, kN s/m	A
C_dA	drag area	0.6	m^2	A
ΔP_{max}, k_Q	fan rise, curve droop	6.35, 9.83	kPa, Pa s^2/m^6	A, H
A_{floor}, L_s	floor area, skirt perimeter	4.0, 12	m^2, m	H
δ_0, h_s	sealed gap, skirt reach	5, 150	mm	A, H
d_{floor}	floor below the CG	0.35	m	H
τ_f	fan spool lag	0.5	s	A
R, J_w	wheel radius, lumped inertia	0.34, 8	m, kg m^2	H, A
F_{mot}, P_{max}	tractive force, power limit	34, 1.0	kN, MW	A
τ_m, σ	motor lag, relaxation length	20, 0.25	ms, m	A
f_r	rolling resistance	0.012	—	A
T_c	corner thruster, each of four	7	kN	A
T_{rear}, T_{yaw}	rear pair, yaw pair	50, 2	kN	A
τ_v	valve lag, proportional, on-off	20, 5	ms	A
I_{sp}	specific impulse	73	s	A
V, p_{fill}	tank volume, fill pressure	850, 700	L, bar	A
T_{fill}	fill temperature	293.15	K	H
p_{set}, p_{empty}	regulator, empty	60, 5	bar	A, H
C_w, hA_{in}	wall capacity, gas to wall	60, 800	kJ/K, W/K	A
hA_{out}	wall to air	50	W/K	H
v_{hard}, v_{stop}	hard landing, stop	1.5, 3.0	m/s	A
a_{lim}	driver g limit (warning)	3.5	g	A
g, ρ, T_{air}	gravity, air, ambient	9.81, 1.2, 293.15	SI	H

12.3.5 Events, Instruments and Numerics

Hard events act between integration steps. The takeoff is the instant every corner unloads. The touchdown is located inside its step by interpolating to the first corner contact, with the vertical speed carried from the last free step's acceleration, and a touchdown faster than 1.5 m/s latches HARD LANDING, faster than 3 m/s stops the run. A touchdown faster than 1.5 m/s rolled past 25° or pitched past 15°, or any body corner reaching the road while the car is not wheels-down, latches CRASHED. On the inverted ramp the road is a ceiling, gravity's component along the road normal changes sign, and the car has fallen once it hangs 0.3 m below its unloaded ride height. Every instrument is the truth plus seeded colored noise, with standard deviations of 0.05 rad/s on the wheel speed, 0.03 m/s on the ground speed, 4 mm on the height, 0.02 m/s on the vertical speed, 2 cm on the position, 0.1° on the attitude, 0.3 bar and 0.3 K on the tank, and 0.5 % on the fan speed. The seat's acceleration is filtered with a 20 ms lag, close to 8 Hz, and that filtered value is the g the limit and the scorecard use. The engine of Section 1.6.1 integrates with fourth-order Runge-Kutta steps of $\Delta t = 1$ ms, the computer runs every 10 ms and the historian logs at 100 Hz. Halving the step leaves the metered launch, the hover, the roll and case 9 unchanged to the printed digit, and the four measured touchdowns agree to 0.001 m/s.

12.3.6 Assumptions

The parameters of Table 12.3 describe no real vehicle. One lumped wheel replaces four, so load transfer between axles never changes the grip, and the tire curve was fitted to braking slip and is read here in traction. The skirt is one gap around a flat floor, which is why the suction makes no pitch or roll moment, and the inverted ramp is a flat ceiling the car is already on. Roll and pitch are independent axes, which breaks down in a tumble. The specific impulse is held at 73 s, as if the gas were warmed to ambient between the regulator and the nozzles, while the tank itself cools. A nozzle fed at −150 °C would deliver nearer 47 s. Below 140 K the gas law is extrapolated, and at the end of case 9 (about 125 K and 27 bar) its density is within 1.1 % of the tables while its c_v is about 12 % low, so the last seconds' cooling is slightly overstated. Condensation is not modeled. The g limit, the landing speeds and the reserve are teaching values.

12.4 Three Prices and One Budget

12.4.1 A Normal Force on Command

With the fans off the springs carry the weight alone, and the dry tire can push with at most $1.17 \times 16.9 = 19.8$ kN, 58 % of the motors' 34 kN. Fully spooled the fans add 24.0 kN to the normal force, and the tire's limit rises to 47.8 kN, past the motors. The fans make the tire's limit a manipulated variable, as actively pitched wings do for a race car's tires [26]. They are slow, though, because the suction grows with the square of a speed that lags its command by 0.5 s, so 0.3 s after a start the fans turn at 45 % and give 20 % of their suction. On the ground the allocation is therefore a priority with a time dimension: spool the fans first, push through the tires, and add thrust only for the remainder. Chapter 1 calls this a split range in time (Section 1.4.6).

12.4.2 Positive Feedback at the Seal

The skirt seals until the floor is 15 cm up, 1 cm above the height at which the wheels unload. Past that the leak grows with the clearance and the suction collapses, as Equation (12.3) gives. On a ceiling the consequence is a fall. At 89 % fan speed the suction is 1.12 times the weight, and a bump that opens the skirt 7 mm for 0.3 s cuts it to 0.79 of sealed, 0.89 of the weight. The wheels unload, the floor moves off the ceiling, and the seal never recovers. At takeoff the same feedback works for the car, through the fall of the wheel loads to zero that other flying-car models also show [22]. Fully spooled the fans add 24.0 kN to the 16.9 kN weight, against 28 kN of lift, so the car cannot leave with them at full speed, and the sequencer spins them down first.

> **TRY IT IN THE APP** **5 · The Inverted Ramp**
>
> Load the scenario and press Run. The car drives on a ceiling held by its suction alone with the fans at 89 %, and the bump at 2 s drops it at 2.42 s. Press Reset, set Fan speed to 100 % and run again. The tires now keep at least 0.86 kN on the ceiling through the bump.

12.4.3 A Hover on a Lightening Car

A hover costs mg/v_e, 23.6 kg/s on a full tank, and the car loses a fifteenth of its all-up mass in a 5 s hover, 114.1 kg. An integral alone must find the weight and then chase it down, while a feedforward from the gauge carries it directly. The hover demand is 60 % of each corner thruster, which on on-off valves with a 10 ms minimum on and off time in a 20 ms period lies between the 50 and 100 % that a pulse can give. The minimum impulse bit is 70 N s per thruster.

12.4.4 The Tank Runs Out by Pressure

A blowdown of seconds is nearly adiabatic, so the gas cools and its pressure falls faster than its mass, as shown in Figure 12.2b for case 9. By the takeoff at 3.48 s the metered launch and the rollout have spent 7.1 % of the nitrogen, while the pressure has fallen 23.1 % and the gas has cooled to 1 °C. At

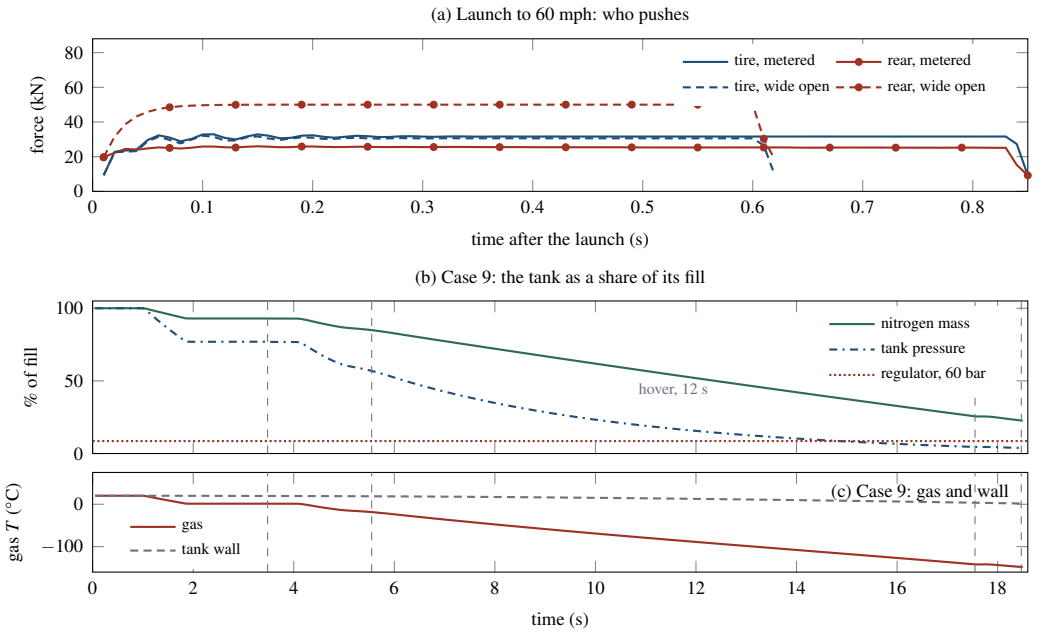

Figure 12.2 The launch allocation and the tank's two budgets, from the engine at the shipped settings. (a) Tire force and rear thrust after the launch, metered to 3.4 g (solid, case 4) and wide open (dashed, case 3), the rear thrust marked with dots. (b) Case 9 nitrogen mass (solid) and tank pressure (dash-dotted) as shares of their fill, with the regulator's 60 bar line (dotted), and thin dashed lines at the takeoff, the start and end of the 12 s hover and the touchdown. (c) The gas (solid) and wall (dashed) temperatures.

the start of the hover the tank holds 84.9 % of its mass at 56.9 % of its pressure, and by the end of a 12 s hover 25.6 % of the mass remains at 31.7 bar and −142 °C. The pressure met the 60 bar regulator at 14.86 s with 38.1 % of the mass left. From there the thrust falls with the pressure, and at the end of the hover full lift is 14.8 kN against a weight of 13.8 kN, 1.07 times it. The valves can still open fully, so the thrusters are the actuator of Section 1.4.5 whose gain falls without reaching its clamp. A runway computed from the gauge's mass with Equation (1.11) is therefore optimistic, since the thrust gives out before the mass does. The wall warms the gas by only about 2 to 6 K over a run, because the time constant $m_N c_v / h A_{in}$ is several minutes.

🕮 **Pause.** Try Exercise 12.1 before reading on.

12.4.5 The Price of a Roll

One corner pair gives 12.6 kN m about the roll axis. The roll profile asks for 90 % of it, $\alpha = 0.9 \cdot 2T_c t / \hat{I}_{xx}$, and a bang-bang turn of 2π takes $T_{roll} = 2\sqrt{2\pi/\alpha}$, 1.04 s in the challenge and 1.08 s in case 7, where the car is heavier. While the car is on its side its lift points sideways, and inverted it points at the road, so a roll from a hover at 2 m reaches the ground. The turn lingers near level at its start and end, where the firing pair still lifts, so the mean lift is not zero. Entered at 5.1 m/s of climb the roll ends 1.77 m above its entry in case 7, and the entry rate that ends the roll at its entry height is about 3.6 m/s.

12.5 The Controllers as Implemented

Every law runs every 10 ms on a measurement packet that holds only the instrument readings of Section 12.3 and the gauge, never the true mass, suction, tire force, surface, gust, stuck valve or

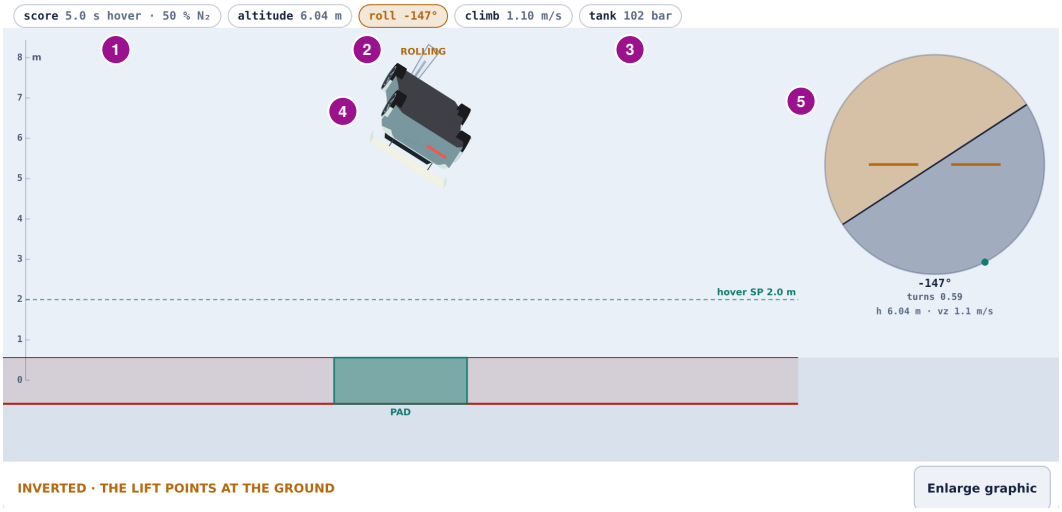

Figure 12.3 The one-tank challenge on its suggested settings, cropped to the Roll view and captured at $t = 11.70$ s, 0.57 s into the 1.04 s roll. The five badges above the scene give the score (1, 5.0 s of hover in band and 50 % of the nitrogen left), the altitude, the roll (2), the climb rate and the tank pressure (3, 102 bar). The car (4) is past inverted at −147°, which the dial (5) counts as 0.59 of a turn, 6.04 m up and still climbing at 1.10 m/s. At that instant the right pair was braking the turn, the collective stood at one pair's 14.0 kN against a weight of 14.8 kN, and the tank held 212 kg, 50 % of its fill.

leak. Truth feeds only the scorecard, the budget shading and the drawing. A stage sequencer from the shared framework carries the run as one active phase at a time: ARMED, SPOOL, LAUNCH, ROLLOUT, TAKEOFF, CLIMB, HOVER, LOB, ROLL, ARREST, LAND and DONE. SPOOL lasts the pre-spool lead, LAUNCH ends at 60 mph (26.82 m/s), ROLLOUT brakes to a stop, TAKEOFF ends at 5 cm of height, CLIMB ends within 0.15 m and 0.3 m/s of the hover, HOVER lasts its set time, LOB ends at the entry climb rate, ROLL lasts $T_{roll} + 0.05$ s, ARREST ends when the descent falls below 0.2 m/s, and LAND ends on the ground. The fans spin down at takeoff and play no part in flight.

12.5.1 The Slip Loop

The motor torque comes from the shared velocity-form PID of Section 1.6.2, Equations (1.12) to (1.15), with the error $e_k = \lambda_{SP} - \hat{\kappa}_k$ and the derivative term on the measured slip. Here $\hat{\kappa} = (R\hat{\omega} - \hat{v})/\max(|\hat{v}|, 3\,\mathrm{m/s})$ is the measured slip with its denominator held at 3 m/s at low speed. The running sum is never clamped, and the integral is corrected by back-calculation from the clamped output, with the tracking time constant equal to τ_I. The defaults are $\lambda_{SP} = 0.17$, $K_c = 40$ kN m per unit slip, $\tau_I = 50$ ms and no derivative. The output range is 0 to the motor limit in the launch and $-T_{lim}$ to 0 in the rollout, where the same loop brakes at $\lambda_{SP} = -0.10$ and a hold torque of 3 kN m opposes any motion slower than 0.3 m/s. At the launch the integral is preset to $0.9\,\mu_{peak}\hat{N}R$, the torque the tire can take on the normal force the gauge and the fan law predict. In Manual the throttle sets the torque and the loop follows it, which makes the switch back to Auto free of bumps.

12.5.2 Metered Rear Thrust

The rear thrusters are off, wide open, or metered to a g target a_t by supplying only what the tires do not, as shown in Equation (12.7).

$$\begin{aligned} F_{rear} &= \mathrm{clip}\big(\hat{m}a_t + \hat{D} - \hat{F}_{tire},\ 0,\ T_{rear}f\big) \\ \hat{F}_{tire} &= (T_{cmd} - J_w\hat{\dot{\omega}})/R, \qquad \hat{D} = \tfrac{1}{2}\rho C_d A\hat{v}^2 \end{aligned} \tag{12.7}$$

where $\hat{m} = m_{dry} + \hat{m}_N$, $f = p_{reg}/p_{set}$ from the measured pressure, and $\hat{\dot{\omega}}$ is the wheel's measured acceleration through a 10 ms lag. The motor command rather than the lagged torque is used because the rear valves lag the same 20 ms as the motors, so the two arrive together. The default target is 3.4 g.

Pause. Try Exercise 12.2 before reading on.

12.5.3 Altitude, Attitude and the Mixer

The altitude loop is a cascade. A ramped height setpoint h_{SP}, moving at 2.5 m/s, asks for a climb rate in proportion to the height error, and a PI on the climb-rate error adds lift to the mass feedforward, as shown in Equation (12.8).

$$\begin{aligned} v_{z,SP} &= \mathrm{clip}\big(K_h(h_{SP} - \hat{h}),\ \pm 2.5\,\mathrm{m/s}\big) \\ L &= \frac{\hat{m}g}{\cos\hat{\phi}\cos\hat{\theta}} + \mathrm{PID}\big(v_{z,SP} - \hat{v}_z\big) \end{aligned} \tag{12.8}$$

with $K_h = 2.5\,\mathrm{s}^{-1}$, a climb-rate gain of 10 kN s/m, $\tau_I = 2$ s and the cosine floored at 0.5. Without the feedforward the integral carries the whole weight. The rocket of Volume 3, Chapter 10 also carries a gravity feedforward while it burns its own propellant. In the lob and the arrest the collective is full lift, in the roll it is one pair's worth, and the loop tracks meanwhile so it takes over on the feedforward. The landing schedule is $v_{z,SP} = -\min\big(4,\ \max(0.5,\ 2\sqrt{\hat{h}})\big)$ m/s, a constant deceleration of $2\,\mathrm{m/s}^2$ that ends at a 0.5 m/s touchdown rate.

Roll and pitch are held by PD laws whose gains follow a natural frequency $\omega_a = 8\,\mathrm{rad/s}$ and damping $\zeta_a = 0.8$ on the gauge's inertias, with the profile's acceleration as feedforward in the roll, as shown in Equation (12.9).

$$\begin{aligned} \tau_\phi &= \hat{I}_{xx}\big(\alpha_{ff} + \omega_a^2 e_\phi + 2\zeta_a\omega_a(p_{ref} - \hat{p})\big) \\ \tau_\theta &= \hat{I}_{yy}\big(\omega_a^2(\theta_{SP} - \hat{\theta}) - 2\zeta_a\omega_a\hat{q}\big) \\ T_i &= \frac{L}{4} + s_{\phi,i}\frac{\tau_\phi}{4t} + s_{\theta,i}\frac{\tau_\theta}{4a}, \qquad 0 \le T_i \le T_c f \end{aligned} \tag{12.9}$$

where $s_{\phi,i}$ and $s_{\theta,i}$ are ± 1 by corner. The mixer gives attitude priority. If the torques alone need more spread than one thruster's range they are scaled down together, and the collective is then clipped to what they leave. Position loops lean the car, $\phi_{SP} = -\mathrm{clip}\big((0.8e_y - 1.4\hat{v}_y)/g,\ \pm 12^\circ\big)$ and the same in pitch within $\pm 8^\circ$, except in the lob, the roll and the arrest. The roll reference is the bang-bang profile, $\phi_{ref} = \frac{1}{2}\alpha s^2$ for the first half of the turn and $2\pi - \frac{1}{2}\alpha(T_{roll} - s)^2$ for the second, with its acceleration led by the valve lag.

12.5.4 Gas-Aware Abort and Pulsed Valves

The sequencer lands at once, never in the middle of a roll, when the gauge reads less than 8 % of the fill or, if the option is on, when the measured tank pressure reaches the regulator setpoint. On-off valves are pulsed every 20 ms and can neither open nor close for less than 10 ms, so a pulse is 0, 10 or 20 ms wide. The nearest-width law, which the app's Pulse law option labels Nearest width, opens each valve for the allowed width nearest its demand. The pulse-width law keeps an impulse account per valve, fires the widest allowed pulse the account holds, and banks at most one period, so the average delivered equals the average asked.

Pause. Try Exercise 12.3 before reading on.

Table 12.4 The Ten Cases and Their Lessons

#	Case as shipped	Teaching point
1	Manual, fans off, throttle 100 %	wheelspin past the friction peak
2	Fans and slip loop, 0.3 s pre-spool lead	suction raises the tire's limit, a slow actuator needs a lead
3	Rear thrusters wide open	fast, over the g limit, expensive
4	Rear thrusters metered to 3.4 g	the free actuator first, the costly one on the remainder
5	Inverted ramp at 12 m/s, fans 89 %, bump at 2 s	a seal's margin must cover a disturbance
6	On the pad, 5 s hover at 2 m, integral alone	the unstick, feedforward against integral
7	Hovering at 2 m, roll from the hover	a roll needs climb rate in hand
8	Hovering on on-off valves, nearest-width pulses	a minimum impulse bit sets a limit cycle
9	Metered launch, 12 s hover, abort off	a demand the tank cannot meet
10	Challenge: wide open, integral alone, on-off valves, roll from the hover, wet patch, 3 kN gust	the whole allocation on one tank



Figure 12.4 Two scene views at their telling instants. Left: the metered launch of case 4 at $t = 1.4$ s on the Forces view, 0.39 s after the launch at 27.6 mph. The road's 37.5 kN carries the 16.8 kN weight and 21.1 kN of suction, the tire gives 31.6 kN of its 43.9 kN limit at a slip of 0.05 with the motors saturated, and the rear thrusters add 25.5 kN toward the 57.1 kN the 3.4 g target asks. Right: case 9 as shipped at $t = 18.5$ s on the Budget view, just after the hard landing, with 96 kg of 423 kg left at 27 bar, −148 °C and 44 % thrust, the pressure under the regulator's line while the mass line still stands well above the reserve.

12.5.5 Three Models Kept Apart

The plant is everything in Section 12.3. The controllers' model is smaller: the declared gas law in the gauge, the fan law in the launch preset, the declared inertias in the roll profile and the PD gains, and the drag estimate in the metering law. No identified model is fitted inside the application, and every controller reads the plant only through its instruments.

12.6 Ten Cases, Measured

The ten cases are listed in Table 12.4. They share one reference configuration, the challenge's suggested settings less the choices each case makes for itself, and depart from it as their rows say. Each value below comes from a headless run at the default seed, advanced in the same fixed steps the browser takes (Section 1.6.1), with the change a row names applied before Run. Each case was also loaded three ways (first, from the list after another case, and by Reset), and the three runs were identical.

12.6.1 Launching on Three Actuators

Floored by hand with the fans off, the slip runs to 1.5, the wheels spin for 3.64 s and the car takes 3.68 s to 60 mph at a peak of 0.96 g. At 60 % throttle the slip peaks at 0.14, below the curve's peak,

Table 12.5 Launch to 60 mph on the Declared Car, Against the Design-Stage Estimate

Launch (fans on unless stated)	0–60 (s)	Peak (g)	N_2 (kg)	Estimate (s, g, kg)
Manual, floored, fans off (case 1)	3.68	0.96	0	—
Manual, 60 %, fans off (case 1)	2.43	1.15	0	—
Slip loop, fans off	2.49	1.15	0	2.3, 1.2, 0
Slip loop, 0.3 s lead (case 2)	1.76	1.93	0	—
Slip loop, 1 s lead	1.47	1.93	0	1.26, 2.2, 0
Rear wide open (case 3)	0.60	4.83	40.5	0.51, 5.5, 35
Metered to 3.0 g	0.94	3.00	24.3	—
Metered to 3.4 g (case 4)	0.83	3.39	29.1	0.81, 3.4, 21
Metered to 3.8 g	0.75	3.79	33.2	—
Metered, wet patch	0.83	3.41	29.4	—

and the time falls to 2.43 s at 1.15 g. At 61 % the wheels break loose and the time is back to 3.48 s, because past the peak the force falls as the slip grows. The launch results are collected in Table 12.5. With the fans spooled the motors rather than the tire set the pace, and a lead of 0.3 s costs 0.29 s against a lead of a second.

> **TRY IT IN THE APP** **1 · By Hand: Fans Off, Foot to the Floor**
>
> Load the scenario and press Run in Manual. Watch the slip badge and the WHEELSPIN warning: the slip runs to 1.5, the wheels spin for 3.64 s, and the car needs 3.68 s to 60 mph. Press Reset, set Throttle (motor torque) to 60 % and run again. The slip now peaks at 0.14 and the time falls to 2.43 s.

The metered launch is the allocation drawn in Figure 12.2a and on the Forces view of Figure 12.4. At the slowest playback speed, 0.1 ×, it lasts 8.3 s of wall time, slow enough to watch the allocation on the traction bar. Between 0.2 and 0.6 s after the launch the tire gives an average of 31.6 kN on 37.5 kN of normal force and the rear thrusters add 25.5 kN, so the seat reads 3.37 g. Wide open the tire gives about the same, 30.6 kN, and the thrusters 50.0 kN, which buys the last 0.23 s for 11.4 kg and 1.4 g over the target. On the wet patch the tire gives less and the metering law makes it up, for 0.3 kg more. The design-stage estimates were made for a lighter car, 1595 kg, with the motors' full force from the first instant. The engine's car weighs 1723 kg, its wheels and rotors add $J_w/R^2 = 69\,\text{kg}$ of effective mass, and the motors lag 20 ms, which together account for the differences. Figure 12.5 sets the ten launches of Table 12.5 side by side.

12.6.2 Suction, Takeoff and Hover

On the inverted ramp at 89 % the bump drops the car at 2.42 s, the suction collapsing to 5 % of sealed as the floor moves off the ceiling. At 94 % the car still falls, at 2.61 s, at 95 % it only just holds with the seal at 0.79, and at 100 % the tires keep at least 0.86 kN on the ceiling through the bump. The flight results are collected in Table 12.6. In case 6 the integral alone leaves the pad 0.97 s after Run, climbs at most at 1.58 m/s, hovers between 1.914 and 1.998 m and lands at 10.21 s. With the feedforward the car leaves at 0.75 s, hovers within 1.985 to 2.003 m and is down at 8.20 s. The 5 s hover costs about the same either way, and the 39.5 kg difference comes from the integral's slow climb and descent, which keep the car in the air 2.0 s longer. Lifting at once with the feedforward, while the fans are still spooling down, the car leaves at 0.27 s with 8.1 kN of suction still under the floor, which collapses as the skirt lifts.

12.6.3 The Roll and the Pulsed Valves

Rolled from the hover, the car is falling within a quarter turn and a body corner reaches the ground 0.73 s into the roll, at 75° of roll. With the lob it climbs at full lift to 5.12 m/s, enters the roll at 3.88 m, ends it 0.3° from level at 5.64 m, arrests the fall and lands at 1.23 m/s. A lob to 6 m/s ends the roll at 7.40 m, 2.80 m above its entry. On on-off valves the nearest-width law rounds the 60 % demand to

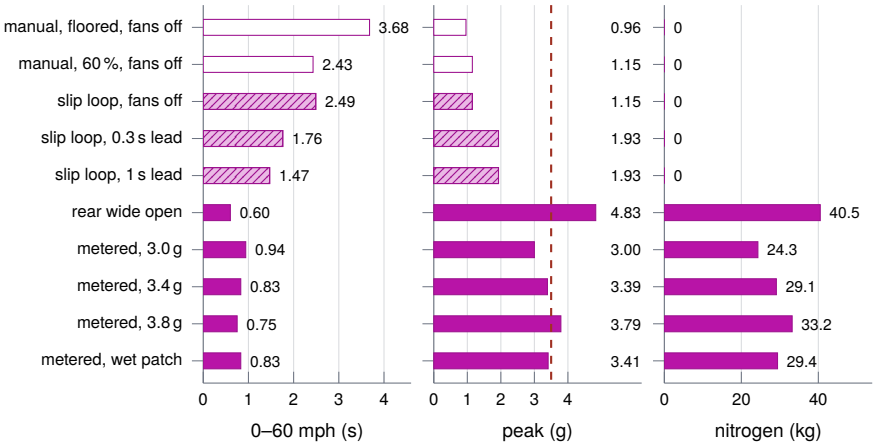

Figure 12.5 Every launch of Table 12.5: time to 60 mph, peak g and the nitrogen spent. Open bars are tires driven by hand, hatched bars the slip loop, and solid bars add the rear thrusters. The dashed line is the 3.5 g teaching limit. Notice that the fans and the slip loop take the launch from 3.68 to 1.47 s for no gas at all, and that past the slip loop each tenth of a second is bought with nitrogen, the last 0.23 s at the price of the g limit.

Table 12.6 Flight Cases. Hover Band After Its First Second, Worst Error Over the Whole Hover, Gas Spent Over the Whole Run

Case	Configuration	Hover (m)	Worst	N_2 (kg)
6	integral alone (shipped)	1.914–1.998	15.4 cm	207.6
	mass feedforward	1.985–2.003	3.5 cm	168.1
8	on-off, nearest width (shipped)	1.902–1.987	9.8 cm	204.2
	on-off, pulse-width law	1.992–2.008	1.0 cm	200.8
	proportional valves	1.995–2.006	0.6 cm	200.9
		Entry	Roll end	Touchdown
7	roll from the hover (shipped)	0.01 m/s, 2.00 m	CRASHED	—
	lob to 5.1 m/s	5.12 m/s, 3.88 m	0.3°, 5.64 m	1.23 m/s
	lob to 3.6 m/s	3.61 m/s, 2.92 m	0.3°, 2.95 m	1.34 m/s
	lob to 3 m/s	3.08 m/s, 2.69 m	0.1°, 2.15 m	1.31 m/s
		Hover	Left	Touchdown
9	12 s, abort off (shipped)	12.00 s	22.6 %	2.83 m/s
	pressure abort on	9.30 s	33.2 %	0.94 m/s
	10.8 s, abort off	10.80 s	—	1.47 m/s
	10.9 s, abort off	10.90 s	—	1.59 m/s

50, and the hover hunts between 1.902 and 1.987 m, the lift on each corner jumping between half and full as the loop's demand crosses the rounding point. The pulse-width law brings the worst error from 9.8 to 1.0 cm and spends 3.4 kg less.

12.6.4 An Honestly Infeasible Hover

Case 9 asks for 12 s of hover after the metered launch with the pressure abort off. The hover is flown, and the landing comes down at 2.83 m/s, a HARD LANDING, with 22.6 % of the nitrogen in the tank at 27 bar and −148 °C, the thrust down to 44 %, as the Budget view of Figure 12.4 shows. The gas has cooled 168 K from its fill. With the abort on, the car starts down when the tank reaches 60 bar after 9.30 s of hover and lands at 0.94 m/s. Without it, 10.8 s is the longest hover that still lands softly. A sweep of 864 combinations of the abort, the feedforward, the valves and pulse law, the climb-rate gain and integral time, the height gain and the reserve found none that hovers 12 s and lands under 1.5 m/s, the softest coming down at 2.38 m/s, and the best of 36 hand-flown descents lands at 2.49 m/s. The

Table 12.7 The One-Tank Challenge. Values at the Default Seed, Nitrogen Left at the End of the Run, Passes Over the Default Seed and Seeds 1–20

Configuration	0–60 (s, g)	Roll end	Touch (m/s)	Left (%)	Seeds passed
shipped	0.60, 4.82	crash	5.81	—	0/21
suggested	0.83, 3.41	0.1°	1.22	37.9	21/21
suggested, integral alone	0.83, 3.41	0.1°	0.50	25.9	21/21
suggested, on-off valves	0.83, 3.41	0.7°	1.21	37.6	21/21
suggested, wide open	0.60, 4.82	0.0°	1.26	35.7	0/21
suggested, roll from hover	0.83, 3.41	crash	—	—	0/21

limit belongs to the tank. Thrusters rated at 40 bar land the same hover at 1.08 m/s, and a gas-to-wall conductance of 10 kW/K lands it at 1.44 m/s.

12.6.5 The One-Tank Challenge

The challenge asks for 0–60 under 0.85 s at no more than 3.5 g, a roll ending within 10° of level, a 5 s hover within 0.2 m of 2 m, a touchdown under 1.5 m/s and at least 5 % of the nitrogen left, with the declared parameters. The run is judged when the car is down and settled, and a change to a model parameter, the hover height or a disturbance voids the attempt. The challenge configurations are compared in Table 12.7. As shipped the run fails twice. The wide-open launch reaches 60 mph in 0.60 s at 4.82 g, and the roll from the hover meets the ground at 5.81 m/s and 60° of roll, CRASHED. The suggested settings (the metered launch, the mass feedforward, proportional valves and the lob) pass at each of 31 seeds, at the native step and with it halved, and the shipped settings fail at all 62 of those runs. They spend 262.7 kg: 29.4 on the launch, 0.7 in the rollout, 33.9 to take off and climb, 110.4 in the hover, 24.7 on the lob, 21.7 on the roll, 5.4 in the arrest and 36.5 to land. The roll ends 0.1° from level and 2.04 m above its 3.60 m entry, and the car lands at 1.22 m/s with 37.9 % left, down and settled 14.3 s after Run. Two of the four changes are not needed on their own, as the table shows, and the scorecard prices them instead. The car rolling on the suggested settings is shown in Figure 12.3.

> **+ DID YOU KNOW? The Softest Landing Is Not the Cheapest**
>
> From the simulator: on the challenge's suggested settings, the integral alone in place of the mass feedforward gives the softest touchdown of Table 12.7, 0.50 m/s against 1.22 m/s, but leaves 25.9 % of the nitrogen against 37.9 %.

12.7 Technology in Practice

A real car of this kind would read and move what the laws of Section 12.5 read and move. Figure 12.6 draws the slip loop, the metered rear thrust and the tank's gauge with field hardware in place of the schematic's symbols.

What Measures It

The slip loop needs two speeds, the wheel's and the car's, and the model reads them from a wheel-speed sensor and a ground-speed radar with noise of 0.05 rad/s and 0.03 m/s. Its error is the measured slip, with the car's speed held at 3 m/s or more in the denominator. What no on-board instrument reads is the friction the road can give, the open challenge at the end of the Limitations. The tank carries a pressure and a temperature transmitter, and the gauge turns their readings into a mass.

TECHNOLOGY Sensor: Pressure and Temperature as a Tank Gauge

A gas tank has no level to read, so what it holds is computed from its pressure, its temperature and its known volume through an equation of state. Spacecraft tanks are gauged from the same three readings by the pressure-volume-temperature method, which rests on conservation of the gas's mass, so no venting or leakage is allowed. In one ground test its errors stayed within 1.5 %. It was most accurate in steady, uniform conditions, had more uncertainty at low tank pressures, and is not considered an accurate real-time technique. A pressure transducer failing slowly showed as readings that grew progressively different from those of a second transducer [27]. The model's gauge inverts the declared gas law from readings with 0.3 bar and 0.3 K of noise.

What Moves It

The gas reaches the valves through a pressure regulator. In a self-operated, spring-loaded regulator a spring pushes the plug off its seat while the downstream pressure acts on a diaphragm to push it back, so the less the downstream pressure, the wider the trim opens, and the spring sets the setpoint [28]. The model's regulator holds 60 bar until the tank falls to it and then passes the tank's pressure, $p_{reg} = \min(p_{set}, p)$ in Equation (12.5), which is why the thrusters lose gain below 60 bar while their valves still open fully. The motors and the fans are the other final elements, each modeled as a first-order lag on its command.

TECHNOLOGY Final Element: Solenoid Thruster Valves

A common on-off valve is the solenoid valve. An energized coil pulls a ferrous armature into its center, and the armature moves a small valve element, most often a poppet that covers a seat. A return spring holds the valve in its normal position, and the process pressure should act to help that spring, since enough pressure against it can keep the valve from returning when the coil is switched off. Fast action demands a coil force that greatly exceeds the spring's [28]. No open source read gives the opening time of a thruster valve. The model's on-off valves lag 5 ms and cannot open or close for less than 10 ms, and a stuck valve is one of its disturbances.

Where the Controller Runs

Control in a car is executed by electronic control units, microprocessors with input and output interfaces. Once computers entered the car they took on antilock braking, electronic braking and stability control, which use accelerometers and gyroscopes to brake the wheels one by one [29]. Every law of this chapter runs each 10 ms on one computer: the slip loop with its clamp on the motor torque, the metering law, the altitude cascade and the mixer that gives attitude priority. The stage sequencer and the gas-aware abort sit above them as a supervisory layer that chooses the active phase.

IN PRACTICE What a Stability Control Must Be Able to Do

In the United States, Federal Motor Vehicle Safety Standard No. 126 (49 CFR 571.126) defines electronic stability control by six attributes. Among them, it brakes the wheels one by one to induce a correcting yaw moment, uses a closed-loop algorithm to limit oversteer and understeer, and has an algorithm and a means to modify engine torque as necessary. It must work in every phase of driving except below 20 km/h, in reverse or during initialization, and stay capable of activation while antilock braking or traction control is active. An off control is allowed, but each new ignition cycle restores the default mode [30]. Torque taken away by a protective loop is the override of this volume.

What Is Modernizing It

Three developments bear on this car. The first is the allocation itself. A flight-tested law for an over-actuated transition aircraft minimizes total power in place of a norm of the commands [24], and a fly-drive vehicle has used a predictive controller over an allocation layer [21], where this chapter

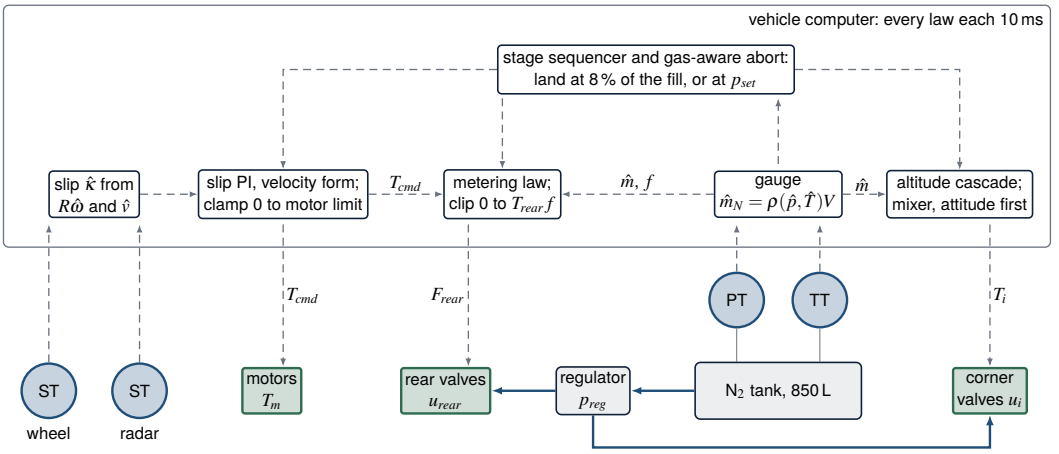

Figure 12.6 The allocation on field hardware. A wheel-speed sensor and a ground radar (ST) give the slip, and its PI loop drives the motors with the output clamped to the motor limit. The metering law asks the rear valves for only the force the tire does not give toward the g target, from the motor command and the gauge's mass. Pressure (PT) and temperature (TT) transmitters on the tank feed the gauge, whose mass also serves the altitude cascade and the gas-aware abort. Every thruster draws through one regulator. Notice that the tank pressure reaches the thrusters twice, as a reading through the gauge and as gas through the regulator, so a falling pressure cuts their thrust while the gauge still reads mass.

keeps a fixed priority. The second is the road. Intelligent tires that carry their own accelerometers are one sensing route to the friction coefficient, and neural networks trained on data are another [31]. The third is security. United Nations Regulation No. 155 makes cybersecurity a condition of a vehicle type's approval. Its list of threats includes sending arbitrary messages on the vehicle's network through an unprotected gateway and manipulating what a sensor collects, and it notes that data manipulation on sensors can be mitigated by correlating data from different sources [32].

12.8 Limitations

The car is generic and its numbers are teaching values, chosen so that the launch, the hover, the roll and the landing fit one tank. The model lumps four wheels into one, ignores load transfer and yaw, reads a braking tire curve in traction, and treats the skirt as one gap around a flat floor. The thrusters hold a constant specific impulse although the gas they draw reaches −148 °C, so the hover's gas cost is understated at the end of a long flight, and the gas law is extrapolated below 140 K. Condensation, nozzle icing and the regulator's own dynamics are outside the model. The tank was enlarged from a first estimate of 600 L to 850 L because on this model the smaller tank either loses thrust before the landing or ends at nitrogen's dew line, and the case 9 infeasibility belongs to that declared tank, not to control.

The controllers know the declared gas law, inertias and fan law, so model error reaches them only through the noise and the disturbances a case introduces. None of them is an optimization, and the allocation is a fixed priority rather than a solved program with costs such as a power-priced allocation [24] or a predictive controller over an allocation layer [21]. Chapter 13 solves a constrained least-squares allocation every period for the feet of a walking robot. The height-neutral entry of about 3.6 m/s is taught in case 7 but not flown in the challenge, which keeps an entry rate of 5.1 m/s, and the challenge's thresholds do not force the mass feedforward or proportional valves, which the scorecard prices instead. The g limit, the landing speeds and the reserve are teaching values, not certified limits.

OPEN CHALLENGE **Knowing the Grip Before the Tire Finds It**

The friction between tire and road sets the best slip of antilock braking and the yaw rate a stability control may ask for, yet no on-board sensor measures it. Estimators built on the vehicle's dynamics resist outside interference but need suitable excitation, added sensors that look at the road need none but are easily disturbed, and data-driven estimators depend on the completeness of their data. A 2022 review names fused estimators, friction shared among connected vehicles and the prediction of friction ahead as the ways forward [31]. The slip loop of this chapter keeps the dry curve's peak of 0.17 as its setpoint on the wet patch too, since the surface is hidden from the controllers.

12.9 Conclusion

A fan car with cold-gas thrusters has three actuators priced very differently, and this chapter measured what each allocation choice buys from one 850 L tank of nitrogen. The fans raise the tire's limit from 19.8 to 47.8 kN at no gas cost, metering the rear thrust to 3.4 g saves 28 % of the launch gas at a cost of 0.23 s, and a mass feedforward from the gauge saves 39.5 kg of a takeoff and hover against an integral alone. A roll on lift thrusters needs climb rate in hand, and a bang-bang turn entered at 5.1 m/s ends 1.77 m high. The tank runs out by pressure, so a 12 s hover lands at 2.83 m/s with 22.6 % of the nitrogen still inside, and no controller among 864 tried lands it softly.

The model is a teaching model with a constant specific impulse and a lumped wheel. A fuller model would add a nozzle temperature dependence, yaw and a lateral tire, a solved allocation with explicit costs [24] or a trajectory optimized against the stored gas as an energy budget is optimized for a solar aircraft [25], and a regulator with its own dynamics. Even as a teaching model, the car shows allocation against a budget that runs out in a quantity the gauge does not show.

KEY POINTS

- A car with tires, fans and gas thrusters has three ways to push, priced very differently, and all of its gas comes from one tank.
- The fans raise the tire's limit from 19.8 to 47.8 kN at no gas cost, and metering the rear thrust to 3.4 g saves 28 % of the launch gas for 0.23 s.
- A mass feedforward from the gauge saves 39.5 kg of a takeoff and hover against an integral alone.
- The tank runs out by pressure, so a 12 s hover lands at 2.83 m/s with 22.6 % of the nitrogen still inside.

Where else this applies. Volume 20, Chapter 4 allocates among rotors and a wing whose authority moves with the airspeed, and Volume 7, Chapter 4 holds the same tire curve near its friction peak while braking.

References

[1] W. C. Durham, "Constrained control allocation," *Journal of Guidance, Control, and Dynamics*, vol. 16, no. 4, pp. 717–725, 1993.

[2] T. A. Johansen and T. I. Fossen, "Control allocation—a survey," *Automatica*, vol. 49, no. 5, pp. 1087–1103, 2013.

[3] A. Anis, "Cold gas propulsion system – an ideal choice for remote sensing small satellites," in *Remote Sensing – Advanced Techniques and Platforms*. InTech, 2012.

[4] K. Lemmer, "Propulsion for CubeSats," *Acta Astronautica*, vol. 134, pp. 231–243, 2017.

[5] A. Reyes-Lúa and S. Skogestad, "Multi-input single-output control for extending the operating range: Generalized split range control using the baton strategy," *Journal of Process Control*, vol. 91, pp. 1–11, 2020.

[6] S. V. Weston, C. D. Burkhard, J. M. Stupl, R. L. Ticknor, B. D. Yost, R. A. Austin, P. Galchenko, L. K. Newman, and L. Santos Soto, "State-of-the-art small spacecraft technology," National Aeronautics and Space Administration, Ames Research Center, Moffett Field, CA, Technical Publication NASA/TP-20250000142, Jan. 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20250000142

[7] U. Kiencke and L. Nielsen, *Automotive Control Systems: For Engine, Driveline, and Vehicle*, 2nd ed. Berlin: Springer Berlin Heidelberg, 2005.

[8] B. Wie, *Space Vehicle Dynamics and Control*, 2nd ed. American Institute of Aeronautics and Astronautics, 2008.

[9] B. Açıkmeşe and S. R. Ploen, "Convex programming approach to powered descent guidance for mars landing," *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 5, pp. 1353–1366, 2007.

[10] J. Katz, "Aerodynamics of race cars," *Annual Review of Fluid Mechanics*, vol. 38, no. 1, pp. 27–63, 2006.

[11] L. Yun and A. Bliault, "Air cushion theory," in *Theory and Design of Air Cushion Craft*. Elsevier, 2000, pp. 48–83.

[12] G. P. Sutton and O. Biblarz, *Rocket Propulsion Elements*, 8th ed. Chichester, UK: John Wiley and Sons, 2010.

[13] A. Haque, S. Richardson, G. Saville, and G. Chamberlain, "Rapid depressurization of pressure vessels," *Journal of Loss Prevention in the Process Industries*, vol. 3, no. 1, pp. 4–7, 1990.

[14] R. Span, E. W. Lemmon, R. T. Jacobsen, W. Wagner, and A. Yokozeki, "A reference equation of state for the thermodynamic properties of nitrogen for temperatures from 63.151 to 1000 K and pressures to 2200 MPa," *Journal of Physical and Chemical Reference Data*, vol. 29, no. 6, pp. 1361–1433, 2000.

[15] C. R. Jarvis, "Flight-test evaluation of an on-off rate command attitude control system of a manned lunar-landing research vehicle," National Aeronautics and Space Administration, Washington, DC, NASA Technical Note D-3903, Apr. 1967.

[16] A. R. Klumpp, "Apollo lunar descent guidance," *Automatica*, vol. 10, no. 2, pp. 133–146, 1974.

[17] H. B. Pacejka, "Single-contact-point transient tire models," in *Tire and Vehicle Dynamics*. Elsevier, 2012, pp. 329–354.

[18] A. M. Eiband, "Human tolerance to rapidly applied accelerations: A summary of the literature," National Aeronautics and Space Administration, NASA Memorandum 5-19-59E, Jun. 1959.

[19] N. S. Currey, *Aircraft Landing Gear Design: Principles and Practices*. American Institute of Aeronautics and Astronautics, 1988.

[20] X. Shi, K. Kim, S. Rahili, and S.-J. Chung, "Nonlinear control of autonomous flying cars with wings and distributed electric propulsion," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 5326–5333.

[21] Q. Tan, X. Zhang, H. Liu, S. Jiao, M. Zhou, and J. Li, "Multimodal dynamics analysis and control for amphibious fly-drive vehicle," *IEEE/ASME Transactions on Mechatronics*, vol. 26, no. 2, pp. 621–632, 2021.

[22] T. Ai, B. Xu, C. Xiang, W. Fan, and Y. Zhang, "Modeling and multimode analysis of electrically driven flying car," in *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2020, pp. 1565–1571.

[23] J. Colmenares-Vázquez, P. Castillo, N. Marchand, and D. Huerta-García, "Nonlinear control for ground-air trajectory tracking by a hybrid vehicle: Theory and experiments," *IFAC-PapersOnLine*, vol. 52, no. 8, pp. 19–24, 2019.

[24] O. Pfeifle and W. Fichter, "Minimum power control allocation for incremental control of over-actuated transition aircraft," *Journal of Guidance, Control, and Dynamics*, vol. 46, no. 2, pp. 286–300, 2023.

[25] R. A. Martin, N. S. Gates, A. Ning, and J. D. Hedengren, "Dynamic optimization of high-altitude solar aircraft trajectories under station-keeping constraints," *Journal of Guidance, Control, and Dynamics*, vol. 42, no. 3, pp. 538–552, 2019, correction: vol. 43, no. 11, p. 2201, 2020, doi:10.2514/1.G003737.c1.

[26] F. Diba, A. Barari, and E. Esmailzadeh, "Handling and safety enhancement of race cars using active aerodynamic systems," *Vehicle System Dynamics*, vol. 52, no. 9, pp. 1171–1190, 2014.

[27] N. T. Van Dresar and G. A. Zimmerli, "Pressure-volume-temperature (PVT) gauging of an isothermal cryogenic propellant tank pressurized with gaseous helium," National Aeronautics and Space Administration, Glenn Research Center, Cleveland, OH, Technical Publication NASA/TP-2014-218083, Feb. 2014. [Online]. Available: https://ntrs.nasa.gov/citations/20140012695

[28] 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.

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

[30] National Highway Traffic Safety Administration, "Standard no. 126; electronic stability control systems for light vehicles, 49 CFR 571.126," Code of Federal Regulations, Title 49, Chapter V, edition of October 1, 2024, U.S. Government Publishing Office, 2024. [Online]. Available: https://www.govinfo.gov/content/pkg/CFR-2024-title49-vol6/pdf/CFR-2024-title49-vol6-sec571-126.pdf

[31] Y. Wang, J. Hu, F. Wang, H. Dong, Y. Yan, Y. Ren, C. Zhou, and G. Yin, "Tire road friction coefficient estimation: Review and research perspectives," *Chinese Journal of Mechanical Engineering*, vol. 35, no. 1, p. 6, 2022.

[32] United Nations Economic Commission for Europe, "UN regulation no 155: Uniform provisions concerning the approval of vehicles with regards to cybersecurity and cybersecurity management system," Official Journal of the European Union, L 82, 9 March 2021, pp. 30–59, 2021. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2021/387/oj/eng

12.10 Exercises

12.1 Concept check · 2 min

At the end of the 12 s hover of case 9, about a quarter of the nitrogen fill is still in the tank, yet the car comes down at 2.83 m/s, a hard landing. Which explanation is right?

(a) The gauge over-reads at low temperature, so far less gas is left than the tile shows

(b) The climb-rate integral wound up in the hover and overshoots the landing schedule

(c) The valves were pulsed, so the hover burned gas faster than mg/v_e

(d) The gas cooled as it expanded, so its pressure fell below the regulator's and the lift fell with it

12.2 Concept check · 1 min

The metered launch sets the rear thrust to $\mathrm{clip}(\hat{m}a_t + \hat{D} - \hat{F}_{tire}, 0, T_{rear}f)$, as in Equation (12.7). Which allocation does this law carry out?

(a) The tires give what they can for free, and the thrusters add only the remainder toward the g target

(b) The thrusters carry the g target, and the tires add whatever the slip loop allows on top of it

(c) The demand is split between the tires and the thrusters in proportion to their force ratings

(d) The thrusters fire only while the slip is above its setpoint, so that the tire is relieved

12.3 Concept check · 2 min

In case 8 the hover needs about 60 % of each corner thruster, and the on-off valves can stay open for 0, 10 or 20 ms of each 20 ms period. With the nearest-width pulse law (Nearest width in the app), what does the hover do?

(a) It holds 2 m, because the pulses average to the demand over many periods

(b) It hunts below 2 m, because the demand rounds to half and the corner lift jumps between half and full

(c) It climbs above 2 m, because the demand rounds up to a full pulse in every period

(d) It settles a few centimeters low and stays there, because the integral cannot see the rounding

12.4 Theory and calculation · 6 min

The skirt of Equation (12.3) has $\Delta P_{max} = 6.35\,\mathrm{kPa}$, $k_Q = 9.83\,\mathrm{Pa\,s^2/m^6}$, $\rho = 1.2\,\mathrm{kg/m^3}$, $L_s = 12\,\mathrm{m}$, $\delta_0 = 5\,\mathrm{mm}$ and $A_{floor} = 4.0\,\mathrm{m^2}$. The full-tank car weighs 16.90 kN, and the dry tire peaks at $\mu = 1.170$.

(a) Find the sealed suction at full fan speed and its ratio to the full-tank weight. *Checkpoint:* 23.99 kN, 1.42 times the weight.

(b) Find the largest tire force on the dry strip with the fans off and with them fully spooled, and compare both with the motors' 34 kN. *Checkpoint:* 19.8 kN (58 % of the motors) and 47.8 kN.

(c) On the inverted ramp a bump opens the gap by a further 7 mm. Find the share of the sealed suction left during the bump, and the lowest fan speed at which the bumped suction still equals the weight. *Checkpoint:* 0.79 of sealed, and a fan speed of about 94.4 %.

12.5 Theory and calculation · 6 min

The thrusters have $I_{sp} = 73\,\mathrm{s}$ with $g = 9.81\,\mathrm{m/s^2}$. The full-tank car has an all-up mass of 1723 kg, the rear pair is rated at 50 kN together, and a proportional valve follows its command with a first-order lag of $\tau_v = 20\,\mathrm{ms}$.

(a) Find the exhaust speed, the gas rate of a hover on a full tank and the gas rate of the rear pair wide open. *Checkpoint:* 716.1 m/s, 23.6 kg/s and 69.8 kg/s.

(b) Find the gas spent in a 5 s hover from a full tank, allowing for the mass that the hover itself removes. *Checkpoint:* 114.1 kg.

(c) In case 3 both rear valves are commanded wide open at the launch and the car reaches 60 mph 0.60 s later. Estimate the gas spent to that instant, allowing for the valve lag. *Checkpoint:* 40.5 kg.

12.6 Theory and calculation · 6 min

At the start of the hover of case 9 the tank holds 84.9 % of its 423.0 kg fill. The dry mass is 1300 kg, $v_e = 716.1\,\mathrm{m/s}$, $g = 9.81\,\mathrm{m/s^2}$, and the gauge reserve is 8 % of the fill.

(a) Find the gas rate of the hover at that instant. *Checkpoint:* 22.7 kg/s.

(b) Find the runway to the reserve with Equation (1.11) at that rate, and again allowing for the rate falling as the car lightens. *Checkpoint:* 14.3 s and 15.9 s.

(c) With the pressure abort on, the car starts down after 9.30 s of hover. Compare and explain. *Checkpoint:* The mass runway is 1.5 to 1.7 times the pressure runway.

12.7 App challenge · 7 min

Load scenario 4, Metered to 3.4 g: the tires first, at the default seed. Each run ends by itself when the car has braked to a stop, about 3.5 s after Run.

(a) Run it as shipped and read the 0–60 mph, Peak g (launch) and N_2 for the launch tiles. *Checkpoint:* 0.83 s, 3.39 g, 29.1 kg.

(b) Press Reset, set Metered g target to 3.00 g, and run again. *Checkpoint:* 0.94 s, 3.00 g, 24.3 kg.

(c) Load scenario 3, Rear thrusters wide open, and run it as shipped. From the three runs, find the nitrogen spent per second of 0–60 time saved, from 3.0 g to 3.4 g and from 3.4 g to wide open. *Checkpoint:* 0.60 s, 4.83 g, 40.5 kg. About 44 and 50 kg/s.

Read the tiles to their last digit, and accept a price within 3 kg/s.

12.8 App challenge · 9 min

Load scenario 9, A 12 s hover and a soft landing on one tank, at the default seed. Each run ends by itself when the car is down, 16 to 19 s after Run.

(a) Run it as shipped and read the Touchdown and N_2 left tiles. *Checkpoint:* 2.83 m/s (hard), 95 kg and 23 %.

(b) Press Reset, set Land when the tank nears the regulator On, and run again. Read how long the car hovered (Hover in band) and the Touchdown and N_2 left tiles. *Checkpoint:* 9.3 s, 0.94 m/s, 140 kg and 33 %.

(c) Press Reset, set Land when the tank nears the regulator Off and Hover time to 10.5 s, and run. Then press Reset, set Hover time to 11.0 s, and run again. Read the Touchdown tile each time. *Checkpoint:* 1.15 m/s (soft) and 1.72 m/s (hard).

Say what limits the hover and how much gas the soft landing leaves unused.

12.9 App challenge · 8 min

Load the challenge, One tank, at the default seed. Each run ends when the car is down and settled or has crashed, within about 17 s of Run.

(a) Run it as shipped. Read the Peak g (launch), Touchdown and Run tiles, and name the two clauses of the target that fail. *Checkpoint:* 4.82 g, 5.81 m/s, CRASHED.

(b) Press Reset, set Rear thrusters to Metered, Mass feedforward On, Thruster valves to Proportional and Aerial roll to Lob first, and run again. Read Peak g (launch), Roll ended, Touchdown, N_2 left and Challenge target. *Checkpoint:* 3.41 g, 0.1°, 1.22 m/s, 160 kg and 38 %, met.

(c) Press Reset, set Mass feedforward Off, and run again. Read Hover worst error, Touchdown, N_2 left and Challenge target. *Checkpoint:* 15.4 cm, 0.50 m/s, 109 kg and 26 %, met.

Say what the feedforward is worth on this flight.

12.10 Application area · 4 min

The model holds I_{sp} at 73 s, as if the gas were warmed to ambient, 293 K, between the regulator and the nozzles, while the tank gas cools to near −150 °C. For an ideal nozzle the exhaust speed scales with the square root of the absolute temperature of the gas at its inlet.

(a) Estimate I_{sp} for gas fed to the nozzles at −150 °C. *Checkpoint:* About 47.3 s.

(b) At the end of the hover of case 9 the car weighs 13.8 kN. Find the hover gas rate with each specific impulse, and say which results of the chapter would move and in which direction. *Checkpoint:* 19.3 kg/s and 29.7 kg/s.

12.11 Application area · 3 min

In case 9 as shipped the tank pressure meets the regulator's 60 bar after 14.86 s, with 38.1 % of the nitrogen fill still in the tank. The gauge reserve that triggers a landing is 8 % of the fill.

(a) Which limit is reached first in this flight, and what share of the fill would a reserve on mass need to protect full thrust here? *Checkpoint:* The pressure, and at least 38.1 %.

(b) Explain why no fixed share of mass serves as the reserve on every flight, and name the measurement the app's abort uses instead. *Checkpoint:* The tank pressure against the regulator setpoint.

12.12 Review the agent · 5 min

An AI assistant reviews case 6, Takeoff and a 5 s hover, and writes: "The mass feedforward only improves the hover's accuracy. A hover costs mg/v_e whichever loop holds it, so the run spends the same nitrogen with the feedforward on or off. Leave it off to keep the controller simple, since the integral will find the weight." Find the flaw in this claim and say how it could be checked.

CHAPTER 13

Quadruped Balance Simulator for Sharing Ground Forces Inside Friction Cones

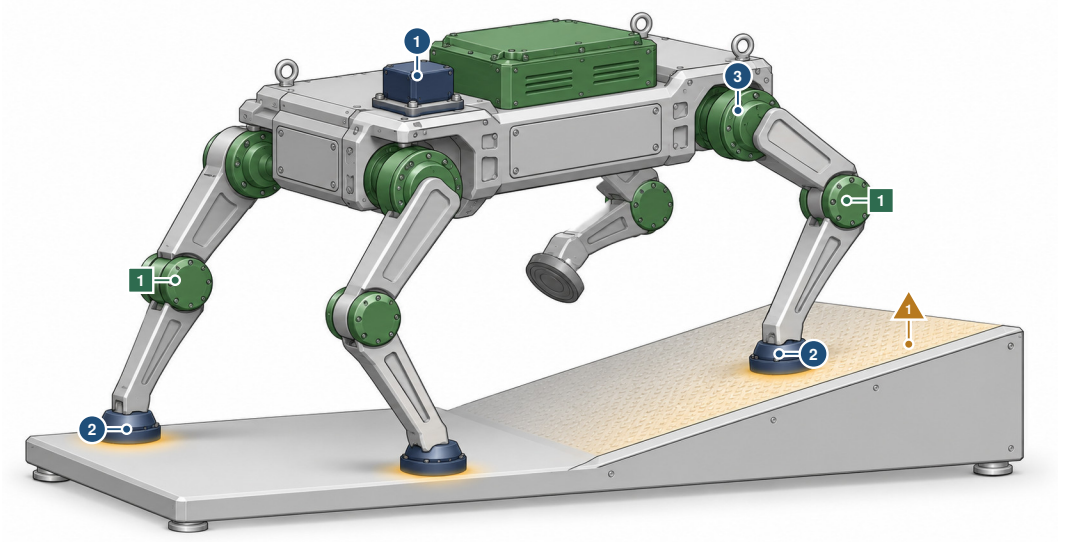

The controller moves 1 hip and knee motors that push each foot

The controller reads 1 tilt and turn sensor on the body 2 contact switches in the feet 3 joint angle encoders

What disturbs it 1 a ramp and a slippery patch

A robot balances only through the feet that are down, and each one can push but not pull, and only inside the grip of the ground under it. Illustration, not to scale.

IN PLAIN TERMS A Robot Dog That Can Only Push

The main idea. A four-legged robot keeps its body steady only by pushing on the ground with the feet that are down. Its computer decides how hard each foot pushes, and a foot can never pull or push sideways harder than the ground's grip allows.

The challenge. The grip, the feet on the ground and the strength of each leg change at every step. A computer that guesses wrong, say that a slippery floor grips well, asks for a push the ground cannot give, and the body tips.

Why it matters. Any machine that shares one job among several limited helpers must know what each can really give.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history segment (Section 13.2), the fact on leg odometry (Section 13.3), the fact and the two try-it segments in the results (Section 13.6) and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 13.6), Technology in Practice (Section 13.8), Limitations (Section 13.9) and the two open challenges

The control question A walking robot corrects its height, pitch and speed only by pushing on the ground, and each foot pushes but cannot pull, only inside its friction cone and only while it touches. Which limit binds in each situation, and what does a controller gain by knowing the true cone, the true contacts and the true leg limits, or by planning the forces ahead?

Learning objectives

- Reduce the feasible closed loop to three double integrators and place their roots from the gains.
- Derive the tipping and friction bounds on braking and tell which one caps a push.
- Explain how planning with the wrong friction or the scheduled feet turns a force request into an unintended moment.
- Compare a reactive allocation with a predictive force planner on a leg whose force limit switches with the gait.
- Configure friction and footholds so that the robot passes a mixed course, and judge the margins it leaves.

Prerequisites Sections 1.4.8 (control allocation), 1.4.5 (saturation and lost actuator effect), 1.6.2 (velocity form) and 1.6.4 (scenarios). Section 1.6.3 is needed only for the predictive planner.

The app https://apmonitor.com/apps/quadruped-balance/

See also Volume 20, Chapter 4 for an allocator whose effectors gain and lose authority, Volume 20, Chapter 6 for thrust allocation on a double-integrator hull and Volume 20, Chapter 13 for constraints that change at contact. Volume 9, Chapter 4 for offset-free predictive control with a disturbance estimate. In this volume, Chapter 12 for allocation by fixed priority.

13.1 Introduction

Legged robots move their bodies only by pushing on the ground. Lifting a sagging body, raising a dipping nose, braking after a shove and holding a posture on a slope all arrive as ground reaction forces at whichever feet are planted at that instant, and each of those forces is bounded three ways. A foot can push but cannot pull, its leg can press only as hard as its motors allow, and its component along the ground can be no larger than the friction coefficient times its normal component [1, 2]. A foot in the air carries nothing.

Control courses usually treat an actuator as a signal that a loop drives anywhere between two limits, so a loop that computes the right command is taken to have delivered it. A legged robot separates the command from the delivery. In the simulator of this chapter, a robot trotting across a patch of floor with a friction coefficient of 0.15 is shoved from behind, and an allocation that assumes a coefficient of 0.6 hands the feet 68.2 N of braking that it believes is feasible. The ground returns 14.8 N. The vertical forces had been arranged to cancel the nose-down moment of the braking that never arrived, so the body pitches up by 13.9° instead of braking level. What the loops request, what the allocation believes and what the ground returns are three different quantities, and the cases of the application are built around the gaps between them.

The problem is one member of a wider family. Several actuators share one requested total, each has hard limits of its own that change with the operating state, and an allocation must decide how much of the total each supplies and report what none of them can. The internal forces of a multifingered grasp chosen under friction and joint-torque limits [1], the deflections of redundant aircraft control surfaces chosen to produce a demanded moment [3], and the ground reaction forces that hold up a humanoid or a quadruped [2, 4] are one problem. Legs add a feature that the other members of the class rarely face, because the set of available actuators changes with the gait, as feet land, lift, slide and miss a foothold several times a second. Volume 20, Chapter 4 meets a milder form of it, effectors whose authority grows and fades with airspeed, and Volume 20, Chapter 13 a constraint set that changes at the instant of contact.

A block diagram of the system hides what matters. While every request is feasible, the allocation is transparent and the torso behaves as three decoupled double integrators, one per loop, which is the case that the linear analysis of Section 13.5 treats exactly. The cases of the application are the

moments when a limit binds. Tipping over the front feet caps braking before friction does, leg reach caps posture on a ramp, the true friction cone caps braking on a slippery floor, a foot over a hole contributes nothing, and a weak leg caps what its diagonal pair can hold. Reviews of control education place laboratory and software exercises at the center of good practice [5] and describe a long-standing distance between control theory and industrial practice [6]. A plant whose set of actuators changes at every step is a direct way to expose that distance.

This chapter determines, case by case, which limit binds in a quadruped balance simulator and what each controller choice costs, and it sets the simulator down in enough detail that it can be rebuilt without its source. The single rigid body, force allocation by quadratic programming, Raibert foot placement and predictive control of ground reaction forces each have an established literature [2, 7, 8]. The model is planar, with unilateral, friction-limited and force-limited contact, sliding feet, toe stubs, missed footholds and odometry that slip corrupts. It is controlled by an allocation that is feasible at every iteration or by a receding-horizon force planner. Closed-form checks against the running engine cover the three loops as double integrators while the allocation is feasible, a braking bound in which tipping over the front feet binds before friction does, and a copper-loss expression that accounts for the energy figure. Each case changes one quantity and reads the consequence from a scorecard that keeps the worst pitch, any fall and the worst push recovery.

Section 13.2 traces balance control on legs. Section 13.3 gives the torso, the legs, the contacts and the sensors, and Section 13.4 the three loops, the allocation, the trot and the planner. The allocation is the weighted least squares of Section 1.4.8 with three demands and a set of actuators that changes at every step. Section 13.5 derives what the feet can deliver, and Section 13.6 reports the nine measured cases. Section 13.7 lists the checks of the engine, Section 13.8 sets the robot's devices beside real hardware, and Section 13.9 states the limitations.

13.2 Balance on Legs, From Creeping Gaits to Force Planning

The first quantitative stability criterion for walking machines was static. A quadruped that lifts one leg at a time remains upright while the vertical projection of its center of mass lies inside the polygon of its supporting feet, and the shortest distance from that projection to the edge of the polygon was proposed as the stability margin of creeping gaits in 1968 [9]. The criterion suits slow machines. It survives in this application in two forms, the support interval drawn under the robot on its scene and the tipping limit on braking derived in Section 13.5.

Dynamic balance came from running machines. Simple algorithms that balance a machine on one leg were generalized to several legs, and the concept of a virtual leg mapped the trot, the pace and the bound, gaits that use the legs in pairs, onto one virtual leg in support at a time, with a quadruped controlled in trotting experiments [7]. The fixed-timing trot and the foot placement rule of this application descend from that work. Hydraulic machines later carried legged balance onto rough ground [10], and quadrupeds with torque-controlled joints, such as HyQ, which stands 1 m tall, weighs about 90 kg and carries 12 torque-controlled joints [11], made the force at each foot something a controller could command directly.

Force-based balance was worked out on humanoids. One method is told in the history segment after this paragraph. Another computes full-body joint torques from a desired motion of the center of mass and the contact forces that produce it [12]. On a 75 kg quadruped the same distribution, posed as a quadratic program with unilateral and no-slip constraints and minimum actuator effort, carried the robot walking between two walls inclined at 50° [2]. The mathematics had appeared earlier in grasping, where the internal forces of a multifingered hand were chosen by linear programming under friction and joint-torque limits [1]. Aircraft control allocation solves the same demanded-moment problem over redundant surfaces, with optimization methods that trade the error in the achieved moment against the control effort [3].

> ⧖ HISTORY **2007: Sharing a Humanoid's Weight Among Its Feet**
>
> Force-based balance was worked out on a hydraulic humanoid of 93.7 kg. Its controller set the ground force and the center of pressure that the body needed and turned them into joint torques for any number of contact points, without measuring the force at those points, and it kept the robot balanced against unknown external forces. The normal forces were shared by a pseudo-inverse that minimizes the sum of their squares, which avoids needless internal forces between the contacts. The horizontal forces followed in proportion to the normal ones, so the contacts that pressed hardest received the most friction. A weighting matrix could be added to respect the friction cone or to unload a limb with a broken joint, and every control process finished within 1 ms [4].

Prediction followed. Reducing the robot to a single rigid body driven by its ground reaction forces made force planning convex, and horizons of up to 0.5 s were solved to optimality in under 1 ms at 20 to 30 Hz on a quadruped that stood, trotted, bounded and galloped [8]. That formulation is the reference design for the predictive planner of this application, which keeps the single rigid body in the plane and solves it approximately with a fixed number of iterations.

The simplifications that made these controllers tractable concern the contacts. A rigid schedule assumes that a foot touches down when the gait says it will, while a contact estimate fused from kinematics, dynamics and the schedule detected touchdown with 99.3 % accuracy and a delay of 4 to 5 ms and let a controller respond leg by leg to early and late contacts [13]. The usual no-slip assumption corrupts state estimation and control alike on slippery ground, which led to separate treatment of slip detection and slip recovery [14]. Terrain templates and learned foothold selection carried the LittleDog quadruped over rough terrain [15], and a push larger than the standing feet can absorb has to be caught with a step, at a capture point [16]. The application reproduces each of these problems in its simplest form and deliberately keeps the fixed schedule, the no-slip odometry and the absence of stepping, so that their consequences can be measured.

13.3 Planar Torso, Legs and Contacts

The simulated robot is a 12 kg rigid torso moving in the sagittal plane on four massless legs, each made of two 0.20 m segments, with hips 0.19 m ahead of and behind the torso center. In the plane the two front legs share one hip and the two rear legs another, while every leg keeps its own foot, contact state and force, and the application declares this approximation on its scene and in its model panel. Figure 13.1 draws the robot with each symbol the equations use, and Table 13.1 gives the role of every signal.

13.3.1 Torso and Mass Properties

The torso is a rigid body driven only by the ground reactions of its planted feet, because a massless leg transmits exactly the force at its foot. The state describes the combined body of robot and backpack, whose translation and rotation obey Equation (13.1).

$$\begin{aligned} m_t \ddot{x} &= \sum_{i \in P} F_{x,i}, \qquad m_t \ddot{z} = \sum_{i \in P} F_{z,i} - m_t g \\ I_t \ddot{\varphi} &= \sum_{i \in P} \left(r_{x,i} F_{z,i} - r_{z,i} F_{x,i} \right), \qquad \mathbf{r}_i = \mathbf{p}_i - (x, z) \end{aligned} \tag{13.1}$$

where P is the set of feet in contact, $\mathbf{p}_i$ is the position of foot i in m, the forces are in N, and φ is in rad, positive nose-up. A backpack of mass m_b centered at the body-frame point $\mathbf{b} = (-0.10, 0.085)$ m, on the rear of the back, gives the mass properties of Equation (13.2).

$$m_t = m + m_b, \quad \mathbf{c} = \frac{m_b \mathbf{b}}{m_t}, \quad I_t = I + m \|\mathbf{c}\|^2 + m_b \|\mathbf{b} - \mathbf{c}\|^2 \tag{13.2}$$

where $\mathbf{c}$ is the offset of the combined center of mass from the torso center in the body frame. A 4 kg pack moves the center of mass 2.50 cm back and 2.13 cm up and raises I_t from 0.30 to 0.352 kg m^2.

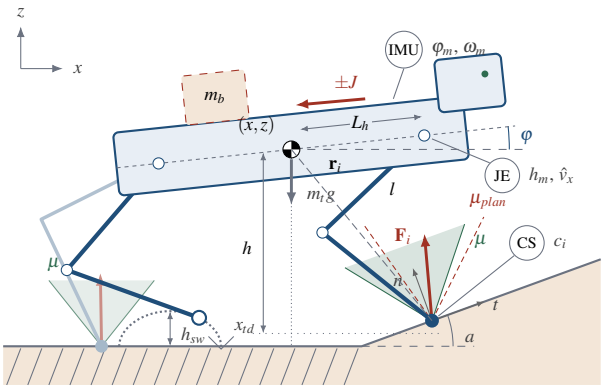

Figure 13.1 Planar robot dog on a floor that turns into a ramp of slope a. The torso, pitched by φ with its center of mass at (x, z) and weight $m_t g$, carries hips L_h ahead of and behind its center, and each leg of two segments of length l reaches a foot at moment arm $\mathbf{r}_i$. A planted foot receives the ground reaction $\mathbf{F}_i$, resolved along the ground t and its normal n, inside the true friction cone of half-angle arctan μ (green), while the allocation plans inside the dashed cone of half-angle arctan μ_{plan} about the normal it believes in. The near rear foot swings toward the touchdown point x_{td} with clearance h_{sw}, and the far rear leg, drawn lighter, stands on its own foot. The body height h runs from the planted feet to the hip line, the push $\pm J$ lands on the top of the back and a backpack m_b sits on its rear. The IMU, the joint encoders JE and the foot contact switches CS supply φ_m, ω_m, h_m, $\dot{v}_x$ and c_i.

Table 13.1 What Each Signal Is, With the Value It Holds While Standing

Symbol	Role	Nominal	Range
h	CV, body height	28 cm	SP 18 to 34
φ	CV, pitch, + nose-up	0°	SP ±15
v_x	CV, forward speed	0 m/s	SP −0.4 to 0.8
F_z	MV, total vertical force	117.7 N	0 to $\sum n_{max,i}$
M_y	MV, pitch moment	0 N m	±60
F_x	MV, total horizontal force	0 N	$\lvert F_x \rvert \leq 150$
t_i, n_i	allocated force per foot	0, 29.4 N	cone, F_{max}
x, z, φ	states, torso	0, 0.28 m, 0	—
v_x, v_z, ω	states, torso rates	0	—
$p_{x,i}$, $p_{z,i}$	states, four feet	±0.19 m, 0	reach 0.10 to 0.38
c_i	states, contact modes	stuck	air, stuck, sliding
—	state, fall flag	0	0 or 1
J	DV, push on the back	10 N s	2 to 16, events
m_b	DV, backpack	0 kg	0 to 6
—	DV, rear-left leg capacity	100 %	20 to 100
—	DV, ground ahead	flat	six courses
φ_m, ω_m	measured, IMU	—	noise, Table 13.2
h_m, $\dot{h}_m$, $\dot{v}_x$	measured, leg kinematics	—	noise, Table 13.2
c_i	measured, contact switches	—	ideal, every 2 ms

When the pack goes on or comes off during a run the torso does not move, and the state is relabeled to the new combined body with the velocity of the torso carried over.

13.3.2 Unilateral, Friction-Limited Contact

The ground is a profile with height $z_g(x)$, slope angle $a(x)$ and static friction coefficient $\mu_s(x)$, and a force at a foot is resolved into a component F_t along the ground and F_n along its normal, as shown in Equation (13.3).

$$F_x = F_t \cos a - F_n \sin a, \qquad F_z = F_t \sin a + F_n \cos a \tag{13.3}$$

A planted foot delivers the force (f_t, f_n) its leg applies whenever that force is physically possible, and the contact law of Equation (13.4) decides what reaches the body.

$$\begin{aligned} F_n &= \text{clamp}\,(f_n, 0, F_{max,i}) \\ (F_t, v_s) &= (f_t,\ 0) \qquad \text{if } |f_t| \le \mu F_n \\ (F_t, v_s) &= \left(\mu_k F_n \,\text{sgn} f_t,\ -\text{sgn}(f_t)\,\frac{|f_t| - \mu_k F_n}{b}\right) \end{aligned} \tag{13.4}$$

where the last line applies when the push leaves the cone, μ is μ_s for a stuck foot and $\mu_k = 0.85\,\mu_s$ for a sliding one, $F_{max,i}$ is the normal-force rating of leg i, v_s is the speed of the foot along the ground in m/s, and $b = 150\,\text{N s/m}$ is the back-drive damping of the leg, the one number a massless, force-controlled leg needs to say how fast it gives way. The tangential command is clipped to $\pm F_{max,i}$ first. A sliding foot moves opposite to the push it failed to hold and sticks again once $|f_t|$ falls back inside $\mu_k F_n$. The stick or slip decision is algebraic within a step, with the threshold chosen by the mode that the previous step left, so the right-hand side stays smooth inside each Runge-Kutta step. A foot in the air carries no force whatever its leg commands and moves with the velocity its swing controller gives it.

13.3.3 Legs, Swing, Terrain and Hard Events

Each leg reaches between $l_{min} = 0.5\,l$ and $l_{max} = 1.9\,l$ from its hip, 0.10 m to 0.38 m, and a swing foot moves no faster than $v_{sw} = 2.5\,\text{m/s}$. The knee of each two-segment leg points backward and is computed from the hip and the foot. Contact modes change only at hard events applied after each integration step. A swing foot touches down when it is at or below solid ground within its leg's reach and is not being driven upward, and a planted foot lifts off when the gait releases it. A planted foot farther than l_{max} from its hip loses the ground, which is counted as a reach loss, while an airborne foot at full reach is carried along by its hip. A swing foot moving forward below the top of an upward face is stopped at the face, a toe stub, and lands short. The robot falls when its pitch passes 40°, when a corner of the torso or a knee touches the ground, or when a planted leg folds below l_{min}, and the simulation then freezes.

Every course starts with at least 0.6 m of flat, dry floor with $\mu_s = 0.8$. The ramp runs from 0.7 to 2.7 m at an adjustable angle, the slippery patch from 0.75 to 1.95 m with $\mu_{patch} = 0.15$, the step rises 6 cm at 0.95 m, and the hole is 12 cm wide at 0.85 m, with a floor 1 m down, beyond any leg. The mixed course combines an 8° ramp from 0.6 to 1.35 m, a patch from 1.75 to 2.45 m, a 6 cm step at 2.95 m and a finish line at 3.8 m.

13.3.4 What the Controller Measures

The controller never reads the true state, as Equation (13.5) sets out.

$$\begin{aligned} \varphi_m &= \varphi + \sigma_\varphi \varepsilon_1, \qquad \omega_m = \omega + \sigma_\omega \varepsilon_2 \\ h_m &= \frac{1}{|P|} \sum_{i \in P} - \left(r'_{x,i} \sin\delta + r'_{z,i} \cos\delta\right) + \sigma_h \varepsilon_3 \\ \hat{v}_x &\leftarrow \text{lag}\left(\frac{1}{|P|} \sum_{i \in P} \left(v_{T,x} - \dot{p}_{x,i}\right) + \sigma_v \varepsilon_4,\ \tau_v\right) \end{aligned} \tag{13.5}$$

where ε_k are independent unit normal deviates drawn each control period from the seeded generator, $\delta = \varphi_m - \varphi$ is the IMU error, $\mathbf{r}'_i$ runs from the torso center to foot i, $v_{T,x}$ is the horizontal velocity

of the torso center and lag is the exact discrete first-order lag. The same odometry gives the vertical rate $\dot{h}_m$ used by the height loop, and the integral of $\hat{v}_x$ is the position estimate $\hat{x}$. The noise levels are $\sigma_\varphi = 0.1°$, $\sigma_\omega = 0.02\,\mathrm{rad/s}$, $\sigma_h = 1.5\,\mathrm{mm}$ and $\sigma_v = 0.02\,\mathrm{m/s}$, and the odometry is smoothed with $\tau_v = 15\,\mathrm{ms}$. Contact switches report the true contact of each foot, sampled every 2 ms. The odometry assumes that a planted foot does not move, so a foot sliding at v_s biases the estimate by the slip speed and the position estimate drifts.

> **DID YOU KNOW? Leg Odometry Drifts by About a Tenth**
>
> In a 60 s static walk of about a meter, a filter that fused leg kinematics with an inertial unit drifted in position by up to roughly 10 % of the distance traveled, from inaccurate leg kinematics and fault-prone contact detection [17]. The odometry of this chapter has no such filter.

13.3.5 Disturbances and Numerics

A push is a horizontal impulse J on the top of the back, 0.06 m above the hip line, applied as the velocity jumps of Equation (13.6).

$$\Delta v_x = \frac{J}{m_t}, \qquad \Delta\omega = -\frac{r_{z,J}\,J}{I_t} \tag{13.6}$$

in which $r_{z,J}$ is the height of the push point above the combined center of mass in m. The calibrated 10 N s push starts the unloaded robot at 0.833 m/s with a nose-down rate of 2.0 rad/s. The backpack is a step in m_b, and the rear-left leg can be derated to a fraction of F_{max}.

The equations are integrated by fourth-order Runge-Kutta at $\Delta t = 0.5\,\mathrm{ms}$, the controller runs every $\Delta t_c = 2\,\mathrm{ms}$ with the leg forces held between executions, and the historian logs at 50 Hz. The engine steps as Section 1.6.1 describes, with the hard events applied after each integration step. Halving the step to 0.25 ms moves the robot of case 7 with sensed contacts from 1.7123 to 1.7116 m after 4 s and its worst pitch from −13.28 to −13.42°, with the same 42 touchdowns in both runs. The noise comes from the seeded generator of the framework, so a run replays exactly. Every numerical value of the model is listed in Table 13.2.

13.3.6 Assumptions

The model is planar. There is no roll, no yaw and no sideways motion, and the left and right legs at each end share a hip, so the robot cannot balance itself sideways and the roll sway of a real trot is absent. The assumption holds for straight walking over ground that is the same across the width of the robot, and it fails for a ramp taken at an angle or a hole under one side only.

The legs are massless. Swinging a leg costs nothing and pushes nothing back on the body, and a touchdown has no impact. This is fair at a walk or a slow trot, where the legs carry a small fraction of the mass, and wrong for running.

Each leg delivers the force it is commanded after a 4 ms lag, up to a normal-force rating that does not depend on the configuration of the leg. Real limits are joint torques whose force equivalent changes with the knee angle.

The ground is rigid, with Coulomb friction and an assumed back-drive damping that sets how fast a foot slides. Real feet roll and deform, and real ground gives way.

The trot has fixed timing and no recovery step, so late contacts do not shift the schedule and a push beyond what standing feet can stop cannot be caught [13, 16].

The sensors carry white noise only, the contact switches are ideal, the terrain map is perfect when it is selected, and no estimator fuses the measurements, so a slip appears directly as position error [14].

13.4 Loops, Allocation, Gait and Planner

The controller runs every 2 ms. Three feedback loops turn the measured height, pitch and speed into a requested body wrench $\mathbf{w} = (F_x, F_z, M_y)$, an allocation turns the wrench into one force per planted

Table 13.2 Parameters of the Robot, Ground and Sensors. Class a Can Be Changed in the Parameter Panel, h Is Hidden From It, and c Is Set in the Source. The Controllers Keep the Values of m and I Listed Here

Symbol	Meaning	Value	Unit	Cl.
m	robot mass	12	kg	a
I	pitch inertia	0.30	kg m^2	a
L_h	hip to torso center	0.19	m	a
l	thigh and shank length	0.20	m	h
l_{min}, l_{max}	leg reach	0.10, 0.38	m	c
—	torso length, top, bottom	0.50, 0.05, −0.04	m	c
F_{max}	leg normal-force rating	150	N	a
v_{sw}	swing speed limit	2.5	m/s	a
τ_a	leg force lag	4	ms	a
b	back-drive damping	150	N s/m	h
μ_s, μ_{patch}	floor, patch friction	0.8, 0.15	—	a
μ_k/μ_s	kinetic to static ratio	0.85	—	a
—	ramp angle in case 4	25	°	a
—	step height, hole width	6, 12	cm	a
σ_φ, σ_ω	IMU noise	0.1, 0.02	°, rad/s	a
σ_h, σ_v	height, odometry noise	1.5, 0.02	mm, m/s	a
τ_v	odometry filter	15	ms	a
J	calibrated push	10	N s	a
—	push height above hip line	0.06	m	c
b	backpack position	−0.10, 0.085	m	c
—	pitch at which it falls	40	°	h
g	gravity	9.81	m/s^2	c
Δt, Δt_c	physics, control step	0.5, 2	ms	c
—	historian rate	50	Hz	c

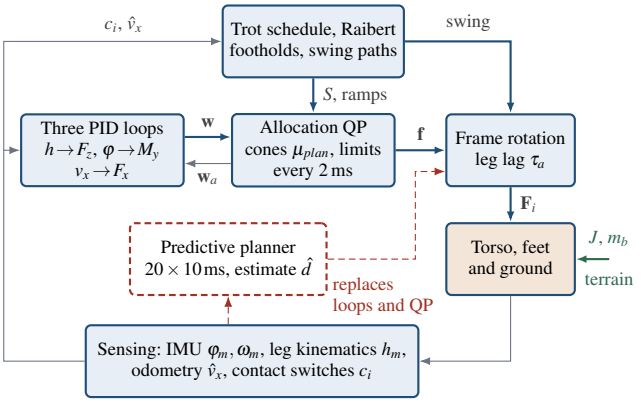

Figure 13.2 Control structure. Three loops request the wrench **w**, the allocation shares it among the feet in the set S that the trot schedule plans and the contact switches confirm, and the delivered wrench $\mathbf{w}_a$ returns to the loops for anti-windup. The forces **f** are rotated from the believed ground frame into the true one and reach the feet through the leg lag. The dashed path is the predictive planner, which replaces the loops and the allocation together when it is selected.

foot, and a fixed-timing trot decides which feet are planted. The structure is shown in Figure 13.2, and its settings are listed in Table 13.3.

13.4.1 Three Loops as Implemented

Each loop runs the velocity-form PID of Section 1.6.2, Equations (1.12) to (1.15), with $P_k = K_c e_k$ and one correction added to the integral, as shown in Equation (13.7).

$$\begin{aligned} u_k^u &= u_{b,k} + \sum_{j \le k} \left(P_j - P_{j-1} \right) + I_{k-1} \\ u_k &= \text{clamp}\left(u_k^u, u_{min}, u_{max} \right) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - u_k^u}{T_t} + \frac{w_{a,k} - u_k}{T_t} \right] \Delta t_c \end{aligned} \tag{13.7}$$

Here $u_{b,k}$ is a bias recomputed every period, u_k^u is the output the algorithm demands, built on the integral accumulated through scan $k-1$, $T_t = \tau_I$ is the tracking time and $w_{a,k}$ is the component of the wrench that the allocation actually delivered. The first correction is the back-calculation of the framework against the clamp. The second is added by the application so that an integral stops winding up against a force the feet could not supply, the case of an actuator that loses its effect without reaching its clamp in Section 1.4.5. Both act on the integral term only and only when the loop has integral action, and an integral time at its 5 s slider maximum removes that action. The framework PID can differentiate a filtered measurement, but all three loops set that derivative to zero and take their rate terms instead through the bias from measured rates, $u_b = mg - K_{c,h}\tau_{D,h}\dot{h}_m$ for the height loop and $u_b = -K_{c,\varphi}\tau_{D,\varphi}\,\omega_m$ for the pitch loop, so no derivative filter is needed and a setpoint step produces no derivative kick. Away from the limits, where neither correction acts, the recursion reduces to the positional laws of Equation (13.8).

$$\begin{aligned} F_z &= mg + K_{c,h}\left(e_h + \frac{1}{\tau_{I,h}} \int e_h \, dt \right) - K_{c,h}\tau_{D,h}\,\dot{h}_m \\ M_y &= K_{c,\varphi}\left(e_\varphi + \frac{1}{\tau_{I,\varphi}} \int e_\varphi \, dt \right) - K_{c,\varphi}\tau_{D,\varphi}\,\omega_m \\ F_x &= K_{c,v}\left(e_v + \frac{1}{\tau_{I,v}} \int e_v \, dt \right) \end{aligned} \tag{13.8}$$

with $e_h = h_{SP} - h_m$, $e_\varphi = \varphi_{SP} - \varphi_m$ and $e_v = v_{SP} - \hat{v}_x$. The bias mg is the nominal weight of the unloaded robot, 117.72 N, whatever mass the Parameter panel sets, so a backpack or a heavier robot has to be found by the integral. The speed setpoint is ramped at 0.8 m/s^2, and the integral of the speed error is the distance between a position reference that advances at the setpoint and the estimate $\hat{x}$, so a standing robot with $v_{SP} = 0$ holds its place. All three loops form the error as setpoint minus measurement with the same sign, so a body below its height setpoint receives more vertical force, a nose-down error a nose-up moment and a speed deficit more forward force, which makes each loop reverse acting in the ISA convention. The limits are $0 \le F_z \le \sum n_{max,i}$ over the planted feet, $|M_y| \le 60$ N m and $|F_x| \le \min(150,\ \mu_{plan} \sum n_{max,i})$ N, and no rate limit applies to any output. With the pitch measured from the ground, the setpoint becomes $\varphi_{SP} + \hat{a}$, where $\hat{a}$ is the estimated ground slope smoothed by a 0.25 s lag.

In manual mode the operator sets F_z, M_y and F_x directly, on sliders of 0 to 400 N, ±40 N m and ±120 N clamped to the same limits, and the allocation shares them among the planted feet. Each loop keeps running and tracks the delivered component by setting I_k so that u_k equals it, which makes the return to automatic bumpless. The same tracking runs while the planner is in charge.

Table 13.3 Controller Settings at Their Defaults, With Ranges. Settings Marked adv. Are Shown in the Advanced View, and the Final Block Is Set in the Source

Symbol	Meaning	Default (range)
h_{SP}	height setpoint	28 cm (18 to 34)
$K_{c,h}$, $\tau_{I,h}$, $\tau_{D,h}$	height loop	1500 N/m, 0.5 s, 0.12 s
φ_{SP}	pitch setpoint	0° (±15)
$K_{c,\varphi}$, $\tau_{I,\varphi}$, $\tau_{D,\varphi}$	pitch loop	60 N m/rad, 0.6 s, 0.10 s
v_{SP}	speed setpoint	0 or 0.4 m/s (−0.4 to 0.8)
$K_{c,v}$, $\tau_{I,v}$	speed loop	120 N s/m, 0.8 s
—	integral removed at	$\tau_I = 5$ s
μ_{plan}	friction assumed	0.6 (0.05 to 1.2)
—	forces go to	sensed contacts, or scheduled feet
—	leg limits planned with	nominal, or actual
—	force planning	reactive, or predictive
N	horizon, 10 ms blocks	20 (4 to 40)
—	planner iterations	40 (1 to 150), adv.
h_{sw}	swing clearance	5 cm (1 to 12)
k_v	foot placement gain	0.08 s (0 to 0.3)
—	footholds	blind, or terrain map
T, D	gait period, duty factor	0.36 s, 0.55, adv.
W	wrench weights	1, 1, 30 m^{-2}
λ_t, λ_n, β	allocation regularization	0.02, 0.002, 0.01
—	allocation iterations	40
R_t, R_n	planner force weights	2×10^{-3}, 2×10^{-4} N^{-2}
—	estimate limits, $\hat{\mathbf{d}}$	3 m/s^2, 15 rad/s^2

13.4.2 Allocation Inside the Cones

The allocation turns $\mathbf{w}$ into ground-frame forces $\mathbf{f} = (t_j, n_j)$ for the planted feet by the quadratic program of Equation (13.9).

$$\begin{aligned} \min_{\mathbf{f}} \; & (A\mathbf{f} - \mathbf{w})^\top W (A\mathbf{f} - \mathbf{w}) + \beta \|\mathbf{f} - \mathbf{f}^-\|^2 \\ & + \lambda_t \textstyle\sum_j t_j^2 + \lambda_n \sum_j n_j^2 \\ \text{s.t. } & 0 \le n_j \le n_{max,j}, \qquad |t_j| \le \mu_{plan}\, n_j \end{aligned} \tag{13.9}$$

where the column pair of A for foot j maps its force through the believed ground angle $\hat{a}_j$ and its moment arm to (F_x, F_z, M_y), the weight $W = \mathrm{diag}(1, 1, 30)$, with the last entry in m^{-2}, ranks the moment above the horizontal and vertical forces, which carry equal weight, $\lambda_t = 0.02$ keeps tangential pushes toward the middle of the cone, $\lambda_n = 0.002$ regularizes the normal forces, and $\beta = 0.01$ ties the answer to the previous period's $\mathbf{f}^-$. Each feasible set is a friction cone cut off by the leg's rating, a triangle in the (t, n) plane, and $n_{max,j}$ is the rating the controller plans with, nominal or actual, scaled by the loading ramp at the start and end of each stance. The program is solved by 40 iterations of accelerated projected gradient [18], with the step set from a power-iteration estimate of the largest eigenvalue of the Hessian and every foot's force projected back onto its triangle after each step. The answer is therefore feasible however few iterations run, and what it cannot deliver appears as a shortfall $\mathbf{w}_a \neq \mathbf{w}$ that the loops treat as saturation. Chapter 12 shares the force on one car body among tires, fans and thrusters by fixed priority instead of quadratic weights.

The allocation works in the frame the controller believes. Its moment arms run from the torso center rather than from the true center of mass and are rotated by the IMU error, and its ground angle is $\hat{a}_j$, which for blind footholds is the angle of the line through the planted front and rear feet and for the terrain map is the mapped slope at each foot. The command sent to leg i is converted into the true

ground frame at its foot by Equation (13.10).

$$\begin{bmatrix} f_t \\ f_n \end{bmatrix}_i = R(-a_i)\,R(-\delta)\,R(\hat{a}_i) \begin{bmatrix} t \\ n \end{bmatrix}_i \tag{13.10}$$

where $R(\cdot)$ is a planar rotation. This is where a wrong idea of the slope or of the pitch turns into a push the controller did not intend. Each leg then applies its command through the first-order lag τ_a, with zero force on a foot in the air, and the wrench is shared either among every foot the schedule calls planted or, by default, only among those whose contact switch also reports ground.

13.4.3 Trot and Foot Placement

The trot moves the diagonal pairs, front right with rear left and front left with rear right, half a period apart. With period $T = 0.36\,\mathrm{s}$ and duty factor $D = 0.55$ each pair stands for $0.198\,\mathrm{s}$ and swings for $0.162\,\mathrm{s}$, so both pairs are down together for $18\,\mathrm{ms}$ at each hand-over. The timing is fixed, and a foot that lands late or early does not move the schedule. The capacity the allocation may count on ramps up from 30 % over the first $20\,\mathrm{ms}$ of a stance and down over its last $30\,\mathrm{ms}$. Each swing foot aims at the touchdown point of Equation (13.11).

$$x_{td} = x_{hip}(t_{td}) + \frac{\hat{v}_x T_{st}}{2} + k_v\,(\hat{v}_x - v_{SP}) \tag{13.11}$$

where $x_{hip}(t_{td})$ is the hip position predicted at touchdown, $T_{st} = DT$ is the stance time and $k_v = 0.08\,\mathrm{s}$ moves the foot forward to brake when the body is too fast [7], with the offset from the hip limited to $\pm 0.12\,\mathrm{m}$. The path from lift-off to x_{td} follows a smoothstep in both coordinates plus a lift of $h_{sw}\sin \pi s$ at swing fraction s, at no more than v_{sw} and inside the reach of the leg. A blind foot aims at the ground height the planted feet stand on, extrapolated along the line through them, and a mapped foot aims at the true ground height with half of any rise added to its lift. A swing that ends in the air keeps reaching down at $0.4\,\mathrm{m/s}$, and a leg that reaches 97.5 % of l_{max} without finding ground records a missed foothold and waits for its next swing. Standing, a foot that has lost the ground is put back under its hip within $0.25\,\mathrm{s}$.

13.4.4 Predictive Force Planner

The planner replaces the loops and the allocation together. It is a receding-horizon controller of the kind outlined in Section 1.6.3, written for this plant. At every $10\,\mathrm{ms}$ mark of the clock, and again whenever a foot lands, lifts or misses its foothold, it chooses the ground-frame forces of every foot the schedule expects to be down over $N = 20$ intervals of $10\,\mathrm{ms}$, each held constant within its interval, on the planar single rigid body the controller believes in, as shown in Equation (13.12).

$$\begin{aligned} \min_{\mathbf{u}_0 \dots \mathbf{u}_{N-1}} \; & \sum_{k=1}^{N} (\mathbf{s}_k - \mathbf{r}_k)^{\top} Q (\mathbf{s}_k - \mathbf{r}_k) \\ & + \sum_{k=0}^{N-1} \sum_{j} \left[R_t t_{jk}^2 + R_n (n_{jk} - \bar{n}_k)^2 \right] \\ \text{s.t. } & \mathbf{s}_{k+1} = A_d \mathbf{s}_k + B_k \mathbf{u}_k + \mathbf{d}_g + \hat{\mathbf{d}}, \quad \mathbf{u}_k \in \mathscr{P}_k \end{aligned} \tag{13.12}$$

where $\mathbf{s} = (x, z, \varphi, v_x, v_z, \omega)$ starts from the measurements, A_d is an exact double integrator per axis over one interval, $\mathbf{d}_g$ is gravity, B_k maps the forces of the feet planned to be down in interval k, at their present positions or at their planned footholds, through the believed normals and torso-centered moment arms of the nominal mass and inertia, and $\bar{n}_k$ is an equal share of the weight among those feet. The diagonal of Q is 4×10^3, 3×10^5, 3×10^4, 1.5×10^3, 3×10^3 and 2×10^2 in SI units, with R_t and R_n in Table 13.3. The reference advances the position at the speed setpoint from the present estimated error, limited to $\pm 0.3\,\mathrm{m}$, and holds the height and pitch setpoints with zero vertical and

angular rates. Each $\mathscr{P}_k$ is the same set of triangles as the allocation uses, including the loading ramps. The problem has up to 160 force variables when all four feet stand and 89 on average in the trot of case 8, and it is solved by 40 iterations of the same accelerated projected gradient, with the gradient from one forward rollout and one adjoint sweep and the plan warm-started from the previous one shifted by an interval. Only the first interval is applied.

The integral action of the planner is the estimate $\hat{\mathbf{d}}$, a constant vertical and angular acceleration that its model fails to explain, updated at each solve from the error between the velocities it predicted and those it measures, the prediction being interpolated along the interval when a contact change brings the solve forward, with a time constant of 0.25 s and limits of ±3 m/s^2 and ±15 rad/s^2. The horizontal channel is left out because slip corrupts the odometry, and an estimator there would integrate a sliding foot's speed into more push. Volume 9, Chapter 4 obtains the same offset-free action in a linear MPC from a fuel bias carried by its Kalman estimator. The estimate is frozen whenever the support differs from the one the last plan assumed.

13.4.5 Simulation, Analysis and Controller Models

Equations (13.1) to (13.6) form the simulation model, which carries the true mass properties, the true ground at each foot and contact modes that switch at hard events. The analysis model of Section 13.5 is its reduction about standing with every request feasible, in which each loop acts on a double integrator. The internal model of the controller is smaller than either. The allocation knows the nominal geometry, the believed slope and the friction it assumes, and the planner knows a single rigid body of the nominal mass and inertia with a disturbance estimate, while neither knows the backpack, a mass or inertia set in the Parameter panel, the true friction of the patch, a leg's derating unless told, or the slip of a foot.

13.5 What the Feet Can Deliver

The statements of this section follow from Equations (13.1), (13.4), (13.8) and (13.9) by closed-form reasoning, and each is checked against the running engine.

13.5.1 Three Double Integrators While Every Request Is Feasible

When the allocation meets the wrench exactly and the leg lag is neglected, Equation (13.1) closed by Equation (13.8) splits into three independent loops around double integrators. The characteristic polynomial of the height loop is Equation (13.13), the pitch loop has the same form with I and the pitch gains, and the speed loop acting on the position error gives Equation (13.14).

$$ms^3 + K_{c,h}\tau_{D,h}s^2 + K_{c,h}s + K_{c,h}/\tau_{I,h} = 0 \tag{13.13}$$

$$ms^2 + K_{c,v}s + K_{c,v}/\tau_{I,v} = 0 \tag{13.14}$$

At the defaults the height loop has roots at -2.73 and $-6.13 \pm 7.34j$ 1/s, and without its integral it would be an oscillator of 11.2 rad/s with a damping ratio of 0.67. The pitch loop has roots at -2.04 and $-8.98 \pm 9.10j$ 1/s, or 14.1 rad/s with a damping ratio of 0.71 without its integral, and the speed loop has real roots at -1.46 and −8.54/s.

A crouch checks the reduction. Stepping the height setpoint from 28 to 24 cm at $t = 1$ s drives the allocated vertical force from 117.7 N down to 55 N and up to 138.3 N, and the body undershoots to 23.03 cm 0.35 s after the step, where Equation (13.13) with its setpoint zero predicts 22.87 cm at 0.38 s. The difference is the 4 ms leg lag and the noise that the linear model leaves out. The zero at $-1/\tau_{I,h} = -2/\text{s}$, slower than every closed-loop root, produces the undershoot, since with the proportional term acting on the measurement the same poles give none, and the body stays within 2 mm of the new setpoint only from 0.87 s after the step.

The regularization of Equation (13.9) leaves a small shortfall that manual mode exposes. With four feet sharing a vertical request F, the vertical part of Equation (13.9) gives $\sum n_j = F/(1 + \lambda_n/4)$, which for the 117.72 N weight is 0.059 N short. In manual mode nothing corrects it, and the engine

Table 13.4 The Calibrated Push at Three Speed Gains, Against the Linear Speed Loop of Equation (13.14), Which Assumes Any Braking Is Available

$K_{c,v}$ (N s/m)	60	120	300
roots (1/s)	−2.50, −2.50	−1.46, −8.54	−1.32, −23.7
initial braking request (N)	50	100	250
linear slide (cm)	12.3	6.8	3.0
engine slide (cm)	13.1	10.1	7.5
peak braking (N)	43.4	56.6	86.4
nose-down dip (°)	3.95	5.11	9.30
recovery (s)	1.02	0.25	0.33
largest push survived (N s)	11.0	11.5	12.0

measures a mean shortfall of 0.059 N and a sink of 6.15 cm in 5 s, against the 6.13 cm that a uniform 0.059 N on 12 kg gives. A wrench that balances the weight does not hold a pose.

13.5.2 Two Limits on Braking

A standing robot brakes a push with horizontal forces at its feet, a height h below the center of mass, and those forces pitch the body nose-down. The vertical forces must supply the balancing moment, and the rear feet, which cannot pull, reach zero first. With the rear feet unloaded, the moment balance about the center of mass gives the bound of Equation (13.15).

$$F_b \le \min\left(\mu_{plan} F_z,\ \frac{x_f F_z - M_y}{h}\right) \tag{13.15}$$

where F_b is the braking force, x_f is the distance the front feet stand ahead of the center of mass, F_z is the vertical force and M_y is the nose-up moment the pitch loop still requests. Standing still with no moment request, the tipping term is $L_h mg/h = 79.9$ N against a friction term of 70.6 N, so the planned cone would bind first. A 10 N s push changes that, because the body moves over its feet and the pitch loop starts to ask for a nose-up moment. At the braking peak, 30 ms after the push, the front feet stand only 16.6 cm ahead of the center of mass and the rear feet carry nothing. The loops ask for 82.6 N of braking, 120.6 N of vertical force and 9.82 N m nose-up, and the allocation, ranking the moment first, returns 56.6 N, 135.6 N and 6.78 N m, pressing the front feet harder than asked to buy moment. With those values Equation (13.15) gives 56.1 N against a friction term of 81.4 N. Tipping, not friction, caps the braking. Braking at a constant 56.6 N from the first instant would stop the body in $mv_0^2/(2F_b) = 7.37$ cm.

Table 13.4 compares the linear speed loop with the engine. At the softest gain the loop asks for 50 N at the start, below the tipping cap, and the slide of 13.1 cm is close to the linear 12.3 cm. At the default and the stiffest gains the requests of 100 and 250 N are cut to what tipping allows, and the slides are 10.1 and 7.5 cm instead of 6.8 and 3.0 cm. The stiffest loop brakes harder, up to 86.4 N, because at that instant the allocation lifts the vertical force to 159.8 N against the 121.5 N asked and settles for 4.56 N m of the 11.7 N m of moment asked. What a stiffer loop buys is a larger recoverable push. The largest push the standing robot survives, found on the 0.5 N s steps of the push setting, rises from 11.0 to 11.5 to 12.0 N s as the gain rises from 60 to 120 to 300, and a finer 0.25 N s grid finds the same edges. The slides at the edge are 16 to 18 cm, close to the 19 cm between the center of mass and the front feet. The price is the pitch dip, which at the calibrated push grows from 3.95 to 9.30°. A push beyond the edge cannot be stopped inside the support at all. At the default gain 12 N s carries the body 40.6 cm until its chest touches the ground at 1.79 s, and only a step could have caught it [16].

13.5.3 Offsets From an Unmodeled Load

The bias of the height loop is the nominal weight, so with its integral off the 39.24 N of a 4 kg pack must come from the proportional term, a sag of $m_b g / K_{c,h} = 2.62$ cm. The pack also sits 0.10 m behind the torso center about which the controller takes moments, and its weight tilts the body nose-up until the proportional pitch term balances it, as shown in Equation (13.16).

$$K_{c,\varphi}\, \varphi_{ss} = m_b g\, (0.10 \cos \varphi_{ss} + 0.085 \sin \varphi_{ss}) \tag{13.16}$$

Its solution is 3.96°, the second term being the extra lever that the tilt itself gives the raised pack. With both integrals off the engine settles at 25.37 cm, a sag of 2.63 cm, and at 3.96°. With integral times of 0.5 and 0.6 s both offsets vanish, the body is back within 2 mm and 0.2° of its setpoints 1.58 s after the pack goes on, and the feet settle at 44.5 N behind and 34.0 N ahead, against the 44.4 and 34.1 N that a zero moment about the shifted center of mass requires.

13.5.4 Braking on a Slippery Patch

On the patch, friction replaces tipping as the binding term, and no controller can beat it. A body moving 0.833 m/s faster than its reference and braked by the full kinetic friction $\mu_k mg$, with $\mu_k = 0.1275$, needs 0.67 s and 27.8 cm to return to the reference speed, or 23.6 cm at the static coefficient. The robot runs 28.2 cm ahead of its position reference with the allocation planning at 0.6 and 29.0 cm at 0.15, so the friction assumption changes almost nothing about the surge. What it changes is the moment. At 3.36 s, 60 ms after the push, the allocation planning at 0.6 commands 68.2 N of braking and the ground returns 14.8 N. The vertical forces were set to cancel the 19.1 N m nose-down moment that the planned braking would produce at the feet, only 4.2 N m of it arrives, and the body receives 12.3 N m nose-up where the controller commanded 2.95 N m nose-down. Planning at 0.15, the commanded braking of 17 to 21 N is close to the 15 to 18 N the ground returns, and once the 4 ms leg lag has passed the moment the body receives stays within about 1 N m of the moment commanded.

🕮 **Pause.** Try Exercise 13.2 before reading on.

13.5.5 What the Energy Figure Measures

The application reports an energy proxy per meter walked, the integral over the planted legs of the absolute mechanical power $|\mathbf{F}_i \cdot (\mathbf{v}_{hip,i} - \mathbf{v}_{foot,i})|$ plus a copper-loss term $k_{cu} \|\mathbf{F}_i\|^2$ with an assumed $k_{cu} = 0.01\,\mathrm{W/N^2}$ per leg, standing for the heat the motors dissipate while they hold a force, one of the principal losses of a legged robot [19]. With massless legs and feet that stay put the mechanical term is small, and in a trot two feet share the weight for most of the cycle, so the proxy reduces to Equation (13.17).

$$\frac{E}{d} \approx \frac{2 k_{cu} (mg/2)^2}{v} = \frac{k_{cu} (mg)^2}{2v}, \qquad \mathrm{CoT} \approx \frac{k_{cu} mg}{2v} \tag{13.17}$$

where E is the energy in J, d is the distance walked in m and $\mathrm{CoT} = E/(mgd)$ is the cost of transport. At 0.4 m/s it gives 173 J/m and a cost of transport of 1.47, against 178.3 J/m and 1.51 measured over 8 s of trotting in case 8 with the leg at full capacity, of which the copper term is 99 %. The figure therefore measures how evenly and for how long force is held rather than mechanical work, and swing clearance, which costs nothing when the legs are massless, leaves it at 178.31 J/m for every clearance from 3 to 12 cm.

13.6 Results Across the Cases

The nine cases are listed in Table 13.5. Their outcomes, measured at the default seed with the stated change and nothing else, are gathered in Table 13.6. Each change is made after the case is loaded and before the run starts, which gives the same run as the change carried through a reset. Runs last 4 to 17 s. Most results are peaks, times, distances or counts, and the few averages, the rms height and the energy per meter, are quoted with their windows. Figure 13.3 redraws the pitch ranges of cases 5, 8 and 9 from the table.

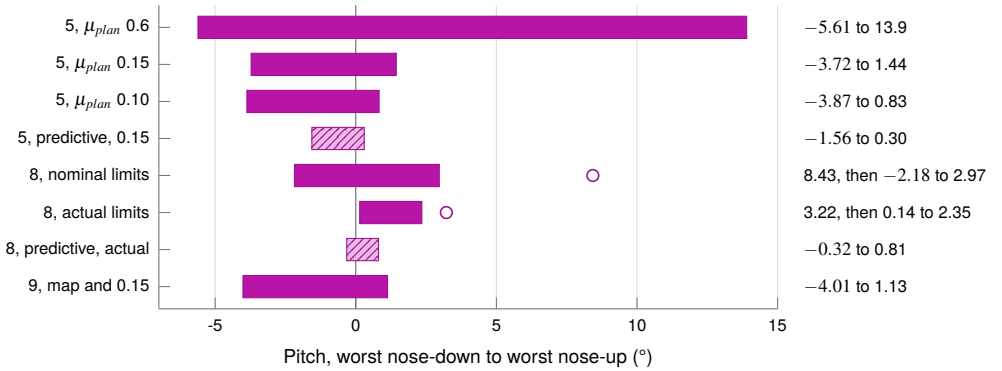

Figure 13.3 Pitch ranges of Table 13.6 for the slippery patch (case 5), the weak leg (case 8) and the graded course (case 9). Solid bars are the reactive allocation and hatched bars the predictive planner. In case 8 the open circle marks the first nose-up step and the bar the range that follows it. Notice that the widest range comes from planning at 0.6 on a patch of 0.15, and that the planner gives the narrowest range in both cases where it runs.

Table 13.5 The Nine Cases and What Each Isolates

#	Scenario and setup	Teaching point
1	Stand quietly	feedback holds a pose
2	Push recovery, 10 N s at 1 s	friction and tipping cap braking
3	Added backpack, 4 kg, integrals off	offsets and integral action
4	Ramp, trot at 0.4 m/s up 25°	leg reach and body posture
5	Slippery patch, push at 3.3 s	plan inside the true cone
6	Raised step, blind, $h_{sw} = 5$ cm	clearance or terrain knowledge
7	Missed contact, forces to scheduled feet	give force only to contacts
8	Force-limited leg at 30 %	reactive against predictive
9	Challenge, 3.8 m course, push at 5.5 s	everything at once

13.6.1 Standing, a Push and a Load

Standing on flat ground between 1 and 5 s, the four feet carry 29.42 to 29.43 N each, 117.71 N in all against a 117.72 N weight, averaged at the physics step, and the true height stays within 0.25 mm rms of 28 cm through 1.5 mm of height noise. The push of case 2 and the pack of case 3 follow the analysis of Section 13.5. The predictive planner answers the calibrated push differently from the reactive loops. It brakes harder, up to 91.2 N, while applying almost no net moment (–0.04 N m), which is about the full tipping capacity of Equation (13.15) with M_y near zero (90.5 N at that instant), and it holds the nose-down dip to 1.36° against 5.11°, at the cost of a slightly longer slide, 10.35 cm against 10.10 cm, and a recovery of 0.36 s.

13.6.2 Posture on a Ramp

Trotting at 0.4 m/s onto the 25° ramp with the pitch measured from level, the robot keeps its back horizontal, so the ground under the front hip is 17.7 cm higher than under the rear hip, and the rear legs must stretch while the front legs fold. The rear legs reach their full 38 cm. The rear left foot loses the ground at 3.05 s and, searching at full extension, cannot reach the ramp again, the rear right foot does the same at 3.21 and 3.51 s, and the front right knee touches the ramp at 3.82 s. The fall is certain but its timing is not, and over the default seed and seven others it comes between 3.54 and 3.84 s, by a knee or a leg folded to its limit, after two or three reach losses. Measured from the ground, the pitch setpoint follows the estimated slope, the body tilts up to 25.6°, the longest planted

Table 13.6 Measured Results at the Default Seed

Case and condition	Result
1, automatic, 1 to 5 s	29.42 to 29.43 N per foot, h 0.25 mm rms
1, crouch to 24 cm	low point 23.03 cm, within 2 mm after 0.87 s
1, manual, wrench equal to weight	0.059 N short, sinks 6.15 cm in 5 s
2, default	slide 10.1 cm, dip 5.11°, recovery 0.25 s
2, $J = 12$ N s	slides 40.6 cm, falls at 1.79 s
2, predictive planner	slide 10.3 cm, dip 1.36°, recovery 0.36 s
3, integrals off	25.37 cm, 3.96° nose-up
3, integrals on	within 2 mm and 0.2° after 1.58 s
4, level body	leg at its 38 cm reach, knee down at 3.82 s
4, pitch from the ground	longest leg 31.12 cm, climbs
4, level body, $h_{SP} = 25$ cm	longest leg 35.86 cm, climbs
4, level body, predictive	climbs to 2.80 m, knee down at 7.17 s
5, $\mu_{plan} = 0.6$	slip 15.6 cm, −5.61 to 13.9°, 1.42 s
5, $\mu_{plan} = 0.15$	slip 1.07 cm, −3.72 to 1.44°, 0.59 s
5, $\mu_{plan} = 0.10$	no slip, −3.87 to 0.83°, 0.77 s
5, predictive, $\mu_{plan} = 0.15$	slip 0.99 cm, −1.56 to 0.30°, 1.16 s
6, $h_{sw} = 5$ cm	six toe stubs, falls at 2.81 s
6, $h_{sw} = 7$, 8 cm	two stubs, none, both climb
6, terrain map, 5 cm	no stub, −0.15 to 0.05°
6, predictive, blind 5 cm	five stubs, falls at 2.81 s
7, scheduled feet	falls at 2.13 s, nose 27.9° down
7, sensed contacts	worst −13.28°, walks on
7, sensed, predictive	worst −3.13°, four misses, past at 6.22 s
7, scheduled, predictive	falls at 2.49 s
8, nominal limits	8.43°, then −2.18 to 2.97°, 0.40 m/s
8, actual limits	3.22°, then 0.14 to 2.35°, 0.317 m/s
8, predictive, actual limits	−0.32 to 0.81°, 0.40 m/s
9, as shipped	seven stubs, falls at 7.63 s
9, terrain map and $\mu_{plan} = 0.15$	finish 9.25 s, −4.01 to 1.13°, 0.55 s
9, $h_{sw} = 8$ cm, $\mu_{plan} = 0.15$	finish 9.57 s, recovery 2.54 s
9, terrain map only	finish 9.68 s, recovery 0.77 s, slip 4.58 cm

leg reaches only 31.12 cm, and the robot climbs to 4.66 m in 12 s. A level back at a 25 cm height also climbs, with its longest leg at 35.86 cm. Friction does not bind on this ramp, because the 0.8 of the floor and the 0.6 of the planning cone both exceed the $\tan 25° = 0.47$ that holding the weight requires. With a level body the predictive planner avoids the early reach losses and climbs to 2.80 m, but its front left knee touches the ground at 7.17 s. The outcome is marginal, since at two of the eight seeds the planner climbs the whole ramp to 4.70 m.

🕮 **Pause.** Try Exercise 13.1 before reading on.

13.6.3 The Slippery Patch

Walking steadily across the patch before the push, no planted foot carrying more than 1 N uses more than 67.8 % of the real cone, so the wrong assumption costs nothing until the push arrives at 3.3 s. A front foot touching down at 3.05 s with 0.6 N of load slides for an instant, no measurable distance, and the application counts and marks a slip only on a foot carrying more than 1 N, so the first slip marked in Figure 13.5 comes after the push. After the push, planning at 0.6 slides the feet 15.6 cm, holds the pitch between −5.61° and 13.9°, and takes 1.42 s to recover, and the odometry, fooled by the sliding feet, ends 8.3 cm behind the true position. Planning at the true 0.15, the feet slide 1.07 cm, the pitch stays between −3.72° and 1.44° and recovery takes 0.59 s. The predictive planner with the true cone narrows the pitch to −1.56° to 0.30° but takes 1.16 s to recover.

Sweeping the assumed coefficient shows that the choice is a trade, as shown in Figure 13.4. Slip and the nose-up peak rise with the excess of μ_{plan} over the true value up to about 0.6 and change little beyond it, reaching 16.3 cm and 14.5° at 0.8, while recovery slows from 0.59 to 1.44 s. Planning exactly at the true coefficient is not the best choice, because it puts every braking command on the edge of the real cone, where pitch noise and the leg lag tip it over. The setting moves in steps of 0.05,

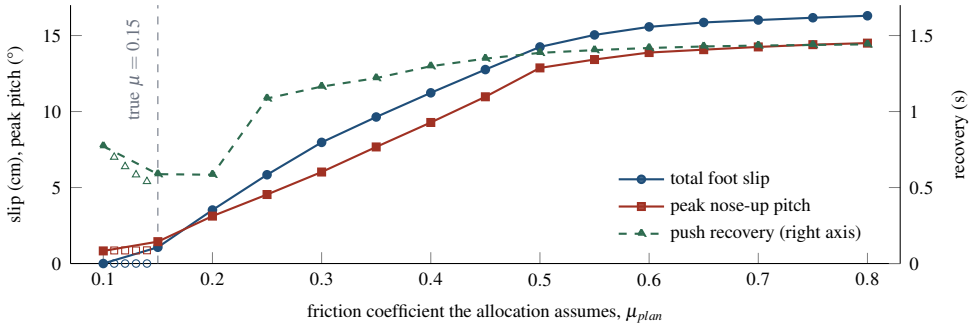

Figure 13.4 Case 5 with only the friction coefficient the allocation assumes changed. Foot slip and the nose-up peak grow with the excess over the true coefficient of 0.15. Filled markers lie on the 0.05 steps of the setting, and hollow markers from 0.11 to 0.14 are counterfactual runs between them, where the fastest recovery, 0.540 s with no slip, comes at 0.14.

and the next step down, 0.10, never slides a foot but takes 0.77 s to recover against 0.59 s at 0.15. Counterfactual runs between the two steps, which a learner cannot set, locate the optimum. At 0.14 the slip is exactly zero and recovery is fastest, 0.540 s, and further conservatism costs time without any gain, 0.64 s at 0.12.

> ✦ DID YOU KNOW? **The Steadier Planner Recovers Later**
>
> From the simulator: in Table 13.6 the predictive planner holds the pitch nearest to level in cases 2, 5 and 8, a dip of 1.36° against 5.11° after the push, yet it recovers later, in 0.36 s against 0.25 s after the push and in 1.16 s against 0.59 s on the patch at the true friction.

13.6.4 A Step and a Hole

Blind with 5 cm of clearance, the feet pass the 6 cm step low in their arc, strike its face six times, and the robot pitches onto its chest at 2.81 s. At 7 cm it strikes twice and climbs, and at 8 cm it climbs with no strike and a pitch between −0.25° and 0.11°. With the terrain map each foot aims at the top of the step and its path rises to meet it, so even 5 cm climbs without a strike, between −0.15° and 0.05°. The predictive planner does not rescue the blind 5 cm swing, which falls at 2.81 s after five strikes, because no choice of forces moves a foot that a face has stopped. Over the default seed and seven others the blind 5 cm swing strikes six or seven times and falls between 2.81 and 2.83 s, and 8 cm never strikes.

In case 7 the allocation first shares the weight among the feet that the schedule says are down. A front foot drops into the 12 cm hole, its share of the support goes nowhere, and the robot falls at 2.13 s, when its front left knee touches the ground with the nose 27.9° down at that physics step (the worst pitch on the scorecard, sampled at the control call before it, reads 27.1°). With force sent only to feet whose contact switch reports ground, the front left foot is stopped by the far wall of the hole at 2.06 s, reaches full extension without finding ground at 2.29 s, and the remaining feet carry the body through a worst nose-down pitch of 13.28°, after which the robot walks on. The predictive planner, which plans the landing of the next pair into the forces it applies now, holds the drop to 3.13°, although with forces sent to the scheduled feet it too falls, at 2.49 s. With sensed contacts it also brakes the lurch over the hole back to its position reference, so the rear feet land short, at the edge of the hole, and miss in turn. The robot rocks in place, its nose up to 8.21°, strikes the far wall six times in all and misses four footholds before all four feet are past the hole at 6.22 s. The position weight is the cause. In a counterfactual run with a quarter of the weight Q_x, the planner misses only the first foothold and is past the hole at 2.94 s, against 2.75 s for the reactive allocation.

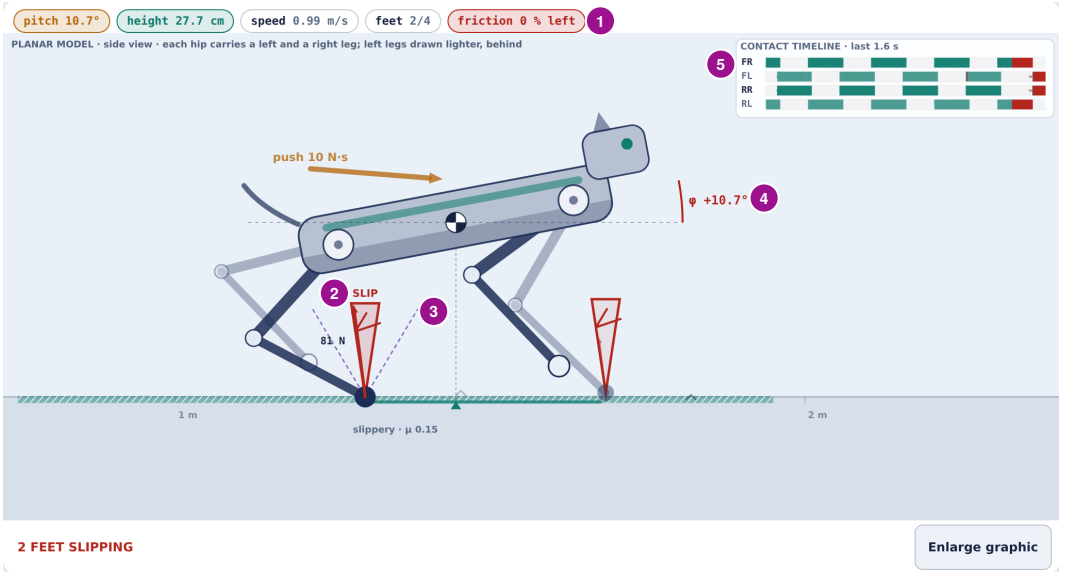

Figure 13.5 The scene panel of the quadruped application in scenario 5, Slippery patch, captured at $t = 3.50$ s, 0.20 s after the 10 N s push. The robot trots on the $\mu = 0.15$ patch with the allocation assuming 0.6, and both feet of the planted diagonal pair slide. (1) The friction badge reads 0 % left, since a sliding foot has no reserve, and the status line below the scene reads 2 feet slipping. (2) The near rear foot, carrying 81 N and labeled SLIP, has its true cone outlined in red, as has the far front foot. (3) The dashed planning cone at the near foot is visibly wider than the true one. (4) The body is pitched 10.7° nose-up under the fading push arrow, and the speed badge reads 0.99 m/s. (5) The contact timeline shows the pair that took the push sliding in red before it lifted and the pair now down sliding since it landed. At that instant the far front foot carried 49 N, the speed setpoint was 0.40 m/s and the odometry read 0.84 m/s, the allocated horizontal force was −65 N against −66 N asked and the pitch moment −18.1 N m, and the first loaded foot had slid 4.5 ms after the push.

> **TRY IT IN THE APP** **7 · Missed Contact**
>
> Load the scenario and press Run. Watch the front foot that steps into the hole and the contact timeline: with the force shared among the scheduled feet, the robot falls at 2.13 s. Press Reset, set Forces go to to Sensed contacts and run again. The remaining feet now carry the body through a worst nose-down pitch of 13.28°, and the robot walks on.

Pause. Try Exercise 13.3 before reading on.

13.6.5 A Weak Leg, Reactive and Predictive

With the rear-left leg at 30 % of its rating, 45 N, and the allocation planning with nominal limits, the allocation asks that leg for 138.8 N in the first weak stance and about 90 N in every later one, 82.8 to 93.1 N after the first second, and the leg delivers 45 N. The body pitches up 8.43° at 0.20 s and then rocks between −2.18° and 2.97° at 0.40 m/s, with the mean pitch 2.2° apart between the stance of the weak pair and that of the strong pair. The imbalance is structural. With only the weak pair down, a front foot and a capped rear foot must carry 117.7 N, the front foot must take 72.7 N, the vertical forces leave a nose-up moment of 5.3 N m, and canceling it with horizontal force at the feet would take 18.8 N of braking. Planning with the actual limit, the allocation stays feasible by giving up horizontal force rather than moment, which carries the largest weight in Equation (13.9). After a

first step of 3.22° the pitch stays between 0.14° and 2.35° while the speed falls to 0.317 m/s, averaged from 2 to 10 s. The predictive planner with the actual limit sees each weak stance coming inside its 0.2 s horizon and spreads the correction across the cycle. The pitch stays between −0.32° and 0.81° at 0.40 m/s, and the mean pitch of the two stances differs by 0.2°. Over 8 s the energy proxy charges the three configurations 213, 259 and 204 J/m against 178 J/m with a healthy leg, the slower walk and the uneven sharing of the second both adding held force per meter, as Section 13.5.5 explains.

> **TRY IT IN THE APP** **8 · Force-Limited Leg**
>
> Load the scenario and press Run. With the rear-left leg at 30 % of its rating, the body pitches up 8.43° and then rocks between −2.18° and 2.97°. Press Reset, set Leg force limits it plans with to Actual and run again. After a first step of 3.22° the pitch stays between 0.14 and 2.35°, but the walk slows to 0.317 m/s.

13.6.6 The Graded Course

The course runs over flat floor, an 8° ramp, a patch with $\mu = 0.15$ and a 6 cm step to a finish at 3.8 m, with the calibrated push at 5.5 s on the patch. The target is to finish without a fall, keep the pitch within ±10°, recover from the push within 2 s and let the feet slide less than 2 cm in total. The attempt is judged when the run stops itself, 3 s after the finish line or 0.35 s after a fall, on the case as given, so changing a model parameter or a disturbance voids it. As shipped, the push slides the feet 4.60 cm and is recovered in 0.77 s, but the blind 5 cm swing strikes the step seven times from 6.93 s on and the robot falls onto its chest at 7.63 s, with a worst pitch of 25.6°. With the terrain map and the allocation planning at 0.15 the robot crosses the finish at 9.25 s, the pitch never leaves −4.01° to 1.13°, the feet slide 0.46 cm and the push is recovered in 0.55 s, at all eight seeds tried. Against the target these are margins of 6.0° of pitch, 1.45 s of recovery and 1.54 cm of slip.

An 8 cm clearance with $\mu_{plan} = 0.15$ also finishes, at 9.57 s, but recovers only after 2.54 s and fails. Near the step the blind slope estimate tilts by 8.9 to 9.6°, more than the 8.5° half-angle of a cone planned at 0.15, so every force the allocation may use points slightly backward, the allocated forward force stays near −1 N against requests of up to 18.0 N, and the speed sags to 0.24 m/s before the robot is over the step. The failure is marginal. At the default seed the recovery from the push has not been held for its 0.3 s before the robot reaches the step, while at the seven other seeds it has, after 0.53 to 0.79 s, and the same settings pass. A 7.5 cm clearance with 0.15, or 8 cm with $\mu_{plan} = 0.20$, also passes at the default seed. The terrain map alone finishes at 9.68 s with a recovery of 0.77 s and a worst pitch of 5.24°, and an 8 cm clearance alone finishes at 9.61 s, but the push slides their feet 4.58 cm and 3.66 cm, past the 2 cm the target allows, at every seed tried. The step lesson of case 6 and the friction lesson of case 5 are both needed, and the target does not single out the terrain map among the step fixes.

13.7 Verification of the Engine

All 35 checks that the application carries pass on the version measured here. A standing robot with a quarter of the weight on each foot has derivatives below 10^{-12}, the closed-loop vertical forces at rest match the weight within 0.2 %, and no foot in the air carries force at any physics step of a 3 s trot. Neither the allocation nor the planner exceeds a friction or force limit at any control step of four whole runs or anywhere over the planner's horizon, the planner's iterations never raise its cost, and reloading the same scenario reproduces a run exactly. The results of this chapter use the default seed, and the orderings of the reactive cases survive other noise realizations. Over the default seed and seven others, planning at 0.10 on the patch never slides a foot and recovers in 0.772 to 0.780 s (0.540 to 0.546 s at the counterfactual 0.14), planning at 0.15 slides 0.93 to 1.09 cm, planning at 0.6 slides 15.52 to 15.78 cm, and the calibrated push of case 2 slides 10.10 to 10.28 cm. Two kinds of result are less robust. The 8 cm clearance with $\mu_{plan} = 0.15$ fails the challenge only at the default seed, as Section 13.6 reports. The predictive planner solves on the 10 ms marks of the clock and at every contact change, so the instant at which it is selected no longer decides its outcome. Selected at

0, 2, 4, 6, 8, 10 or 20 ms, it gives the same push response, the same recovery on the patch and the same crossing of the hole at the default seed. Its marginal results are the level-body ramp, which it climbs at two of eight seeds and falls on at six, and the hole with sensed contacts, which it crosses at seven of eight seeds and falls in at the eighth, at 5.64 s.

13.8 Technology in Practice

A real legged robot reads and moves what the loops of Section 13.4 read and move. Figure 13.6 draws the three loops, the allocation and the planner with field hardware in place of the schematic's symbols.

What Measures It

The controlled variables are the body's height, pitch and forward speed. The model takes the pitch and its rate from an inertial unit with 0.1° and 0.02 rad/s of noise, and the height and speed from leg odometry built on the joint angles of the planted legs, which a contact switch on each foot names. Real robots typically do not use such raw readings directly, but turn them into the pose and velocity of the body with a state estimator [20].

> **TECHNOLOGY Sensor: Inertial Unit and Leg Kinematics**
>
> An inertial measurement unit on the torso measures the body's acceleration and turn rate, and encoders on the joints give the joint angles, from which the known leg geometry places each foot relative to the body. Both carry noise, and the inertial readings also carry biases that wander as random walks. An extended Kalman filter can fuse the two, taking each planted foot as fixed and adding noise to its position to allow a little slip. With one foot or more on the ground, pitch, roll and velocity are observable, while position and heading are not, and faulty contact detection shows as a slow drift of the position estimate [17]. The model's unit and odometry carry white noise only.

What Moves It

Each foot force is made by the motors of its leg. A balance controller of this kind sets the ground force that the body needs and transforms it into joint torques [4], so each leg's own joint loops turn the force the allocation hands it into torques at the hip and knee. Real limits are those joint torques, whose force equivalent changes with the knee angle, as Section 13.3.6 notes, while the model rates every leg at one normal force.

> **TECHNOLOGY Final Element: Proprioceptive Leg Actuators**
>
> Hydraulic legs offer a large power-to-weight ratio but are expensive and need specialized upkeep. An electric motor behind a high-ratio gearbox gains torque, but the gearbox's static friction keeps the joint from being driven backward, so the torque cannot be set through the motor current, and the springs or torque sensors added after it limit bandwidth or suffer under impacts. A motor of high torque density with a low gear ratio sets the joint torque through its current alone, can be back-driven to soften an impact, and controls force at high bandwidth. Such proprioceptive actuators are now the dominant design for legged robots [20]. The model's leg reaches its commanded force after a 4 ms lag, up to a fixed 150 N rating that case 8 cuts to 30 %.

Where the Controller Runs

Reactive controllers of this kind take the state as known and constant over the control period, typically between 1 and 10 ms, and the set of feet in contact as given and fixed for each period. Posed as a quadratic program, with the friction cones approximated as pyramids, the problem is solved in hundreds of microseconds to a few milliseconds. A predictive layer on a simplified model commonly runs above it at a lower rate [21]. Learned controllers show the same split, with a policy updated at about 50 Hz and the joint torque still produced at 1 kHz [20]. In this chapter the loops and the

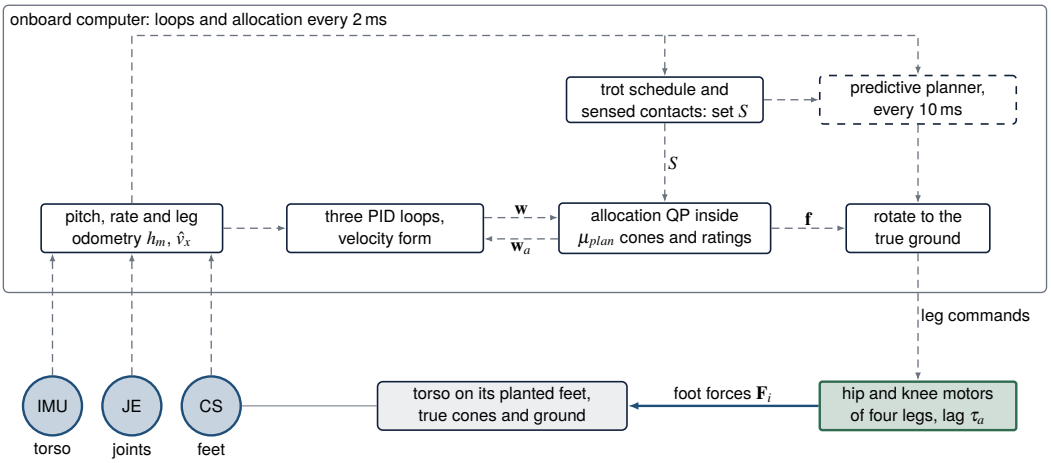

Figure 13.6 The balance loops on field hardware. An inertial unit (IMU) on the torso, encoders on the joints (JE) and a contact switch on each foot (CS) feed the pitch, the leg odometry and the set of planted feet. The three loops ask for a wrench $\mathbf{w}$, the allocation shares it among the feet that the trot schedule plans and the switches confirm, and the delivered wrench $\mathbf{w}_a$ returns to the loops. The forces are rotated into the true ground frame and sent to the hip and knee motors. The dashed planner replaces loops and allocation when selected. Notice that every limit the allocation respects is one it believes, set by μ_{plan} and the ratings it plans with, while the ground applies the true ones.

allocation form the reactive layer and run every 2 ms, inside that range, with each cone already a triangle in the plane, and the predictive planner, when selected, replaces both and solves every 10 ms and at each contact change.

What Is Modernizing It

Three developments bear on this robot. The first is the actuator. Proprioceptive actuators have become the dominant design, and open-source designs from universities and inexpensive hardware from industry have sped up experiments [20]. The second is learning. A controller trained in simulation with the robot's mass, the friction coefficient, the motor strength and the latency randomized transfers better to the real robot, much as a robust controller is designed for a set of parameters, and the approach is now widely adopted [20]. The third is the model inside predictive control. Recent work has used more complex models, some of them the full robot, made possible by better gradient-based methods and fast open-source dynamics libraries, and with Moore's law ending further gains may come from parallel computing [21]. The planner of this chapter keeps the single rigid body in the plane.

IN PRACTICE Soft or Strict Priorities When Tasks Conflict

The tasks given to a legged robot's balance controller conflict in practice, so the designer has to say how they trade. In the common quadratic-program approach, the friction cones and other inequalities stay hard constraints, and a weighted sum of the squared errors of the other tasks is minimized, a soft priority. Regularization terms keep the answer from jumping between consecutive control periods. Weights grow hard to tune as tasks multiply, and a strict hierarchy can rank the tasks instead, though the cost of its solvers has limited its use [21]. The allocation of this chapter is the soft kind, with the pitch moment weighted above the two forces, $W = \mathrm{diag}(1, 1, 30)$.

13.9 Limitations

Where each modeling assumption breaks down is stated in Section 13.3.6. Three of them bear on the results. Because the legs are massless, swing clearance and swing speed carry no dynamic or energetic cost, and a result that depends on either is optimistic. Leg ratings are constant forces rather than joint torques, and an assumed back-drive damping sets the speed of every slip.

The gait never adapts its timing, estimates its contacts beyond ideal switches, or steps to catch a push. The standing robot survives at most 11.0 to 12.0 N s, depending on the speed gain, and a larger push is a fall by construction.

The predictive planner is approximate, with 40 iterations per solve, and it does not rescue every case. It falls on the ramp with a level body, at the low step, and in the hole with forces sent to scheduled feet, and with sensed contacts it holds its position reference so tightly that it rocks at the edge of the hole for 3.5 s. Its disturbance estimate is limited to 3 m/s^2 vertically, which a 4 kg pack on the 12 kg model already exceeds, since it needs 3.27 m/s^2, and the planner then leaves a sag of 1.6 mm. At the 6 kg slider limit both limits are reached, leaving 6.7 mm and 0.35° where the reactive integrals leave 0.1 mm.

OPEN CHALLENGE A Safety Layer Over a Learned Walker

Controllers learned by reinforcement walk robustly outdoors yet struggle in safety-critical places such as stepping stones and confined spaces. One route to safety trains the policy under constraints. Another, for deployment, is a safety filter built from reachability analysis, control barrier functions or predictive methods. Such filters are well studied in control, but their use in locomotion has been limited, and more research on them could help certify legged robots for safety-critical service [20]. In the language of this volume, a safety filter is an override placed over a learned controller.

The sensors carry white noise only, the terrain map is perfect, and no estimator fuses the measurements. The energy figure is almost entirely an assumed copper loss. The numerical targets of the scenarios are design choices of a teaching simulator, not validated performance of a robot.

OPEN CHALLENGE Knowing the Friction and the Load

The friction under the feet and the payload on the back are the quantities this chapter's controllers lack. A recent review of learned locomotion calls measuring or estimating such factors accurately nearly impossible, yet needed for a controller to be robust or adaptive. Current methods build a belief from a history of measurements, or train a student policy, which sees only that history, to imitate a teacher policy that was given the friction and the terrain in simulation [20]. In this chapter the assumed friction is a setting, and case 5 measures what a wrong choice costs.

13.10 Conclusion

A legged robot balances only through the forces its feet can deliver, and the simulator of this chapter makes each limit on those forces measurable. A planar torso on four massless legs with unilateral, friction-limited and force-limited contact is held by three loops whose wrench an allocation shares

inside the friction cones, with a fixed-timing trot and an optional receding-horizon force planner. While every request is feasible the loops are three double integrators, and the cases are the moments when that stops being true. After a 10 N s push, tipping over the front feet caps the braking at 56.6 N while the planned cones would allow 81.4 N. A level body on a 25° ramp stretches its rear legs to their full 38 cm reach and falls, while a body laid along the ground climbs. On a slippery patch an allocation that assumes 0.6 slides the feet 15.6 cm and pitches the body up 13.9°, and one that assumes 0.10 slides them not at all. Forces sent to scheduled feet drop the robot into a hole, and only the predictive planner keeps both pitch and speed on a leg derated to 30 %. The model is planar, its legs are massless, its gait neither adapts its timing nor steps, and its energy figure is almost entirely an assumed copper loss. A fuller model would add roll, leg dynamics, contact estimation with event-based timing, recovery steps and friction estimation. The quadruped is an accessible plant, with nothing to install, for learning to allocate force against limits that change with every step.

KEY POINTS

- A walking robot moves its body only by pushing on the ground, and each foot pushes only inside its friction cone and only while it touches.
- While every request is feasible the loops act on three double integrators, and the cases are the moments when a limit binds.
- After a 10 N s push, tipping over the front feet caps the braking at 56.6 N, below the 81.4 N the planned cones allow.
- Assuming 0.6 on the patch slides the feet 15.6 cm, and forces sent to scheduled feet drop the robot into a hole.
- Only the predictive planner keeps both pitch and speed on a leg derated to 30 %.

Where else this applies. Volume 20, Chapter 6 shares a demanded force among a ship's thrusters, and Volume 20, Chapter 4 among effectors whose authority changes with the airspeed.

References

[1] J. Kerr and B. Roth, "Analysis of multifingered hands," *The International Journal of Robotics Research*, vol. 4, no. 4, pp. 3–17, 1986.

[2] M. Focchi, A. del Prete, I. Havoutis, R. Featherstone, D. G. Caldwell, and C. Semini, "High-slope terrain locomotion for torque-controlled quadruped robots," *Autonomous Robots*, vol. 41, no. 1, pp. 259–272, 2017.

[3] M. Bodson, "Evaluation of optimization methods for control allocation," *Journal of Guidance, Control, and Dynamics*, vol. 25, no. 4, pp. 703–711, 2002.

[4] S. Hyon, J. G. Hale, and G. Cheng, "Full-body compliant human–humanoid interaction: Balancing in the presence of unknown external forces," *IEEE Transactions on Robotics*, vol. 23, no. 5, pp. 884–898, 2007.

[5] 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.

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

[7] M. Raibert, M. Chepponis, and H. Brown, "Running on four legs as though they were one," *IEEE Journal on Robotics and Automation*, vol. 2, no. 2, pp. 70–82, 1986.

[8] J. Di Carlo, P. M. Wensing, B. Katz, G. Bledt, and S. Kim, "Dynamic locomotion in the MIT Cheetah 3 through convex model-predictive control," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 1–9.

[9] R. B. McGhee and A. A. Frank, "On the stability properties of quadruped creeping gaits," *Mathematical Biosciences*, vol. 3, pp. 331–351, 1968.

[10] M. Raibert, K. Blankespoor, G. Nelson, and R. Playter, "BigDog, the rough-terrain quadruped robot," *IFAC Proceedings Volumes*, vol. 41, no. 2, pp. 10 822–10 825, 2008.

[11] C. Semini, N. G. Tsagarakis, E. Guglielmino, M. Focchi, F. Cannella, and D. G. Caldwell, "Design of HyQ – a hydraulically and electrically actuated quadruped robot," *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, vol. 225, no. 6, pp. 831–849, 2011.

[12] B. J. Stephens and C. G. Atkeson, "Dynamic balance force control for compliant humanoid robots," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 1248–1255.

[13] G. Bledt, P. M. Wensing, S. Ingersoll, and S. Kim, "Contact model fusion for event-based locomotion in unstructured terrains," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 4399–4406.

[14] F. Jenelten, J. Hwangbo, F. Tresoldi, C. D. Bellicoso, and M. Hutter, "Dynamic locomotion on slippery ground," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4170–4176, 2019.

[15] M. Kalakrishnan, J. Buchli, P. Pastor, M. Mistry, and S. Schaal, "Fast, robust quadruped locomotion over challenging terrain," in *2010 IEEE International Conference on Robotics and Automation*, 2010, pp. 2665–2670.

[16] J. Pratt, J. Carff, S. Drakunov, and A. Goswami, "Capture point: A step toward humanoid push recovery," in *2006 6th IEEE-RAS International Conference on Humanoid Robots*, 2006, pp. 200–207.

[17] M. Bloesch, M. Hutter, M. A. Hoepflinger, S. Leutenegger, C. Gehring, C. D. Remy, and R. Siegwart, "State estimation for legged robots: Consistent fusion of leg kinematics and IMU," in *Robotics: Science and Systems VIII*, 2012.

[18] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.

[19] S. Seok, A. Wang, M. Y. Chuah, D. J. Hyun, J. Lee, D. M. Otten, J. H. Lang, and S. Kim, "Design principles for energy-efficient legged locomotion and implementation on the MIT Cheetah robot," *IEEE/ASME Transactions on Mechatronics*, vol. 20, no. 3, pp. 1117–1129, 2015.

[20] S. Ha, J. Lee, M. van de Panne, Z. Xie, W. Yu, and M. Khadiv, "Learning-based legged locomotion: State of the art and future perspectives," *The International Journal of Robotics Research*, vol. 44, no. 8, pp. 1396–1427, 2025.

[21] P. M. Wensing, M. Posa, Y. Hu, A. Escande, N. Mansard, and A. Del Prete, "Optimization-based control for dynamic legged robots," *IEEE Transactions on Robotics*, vol. 40, pp. 43–63, 2024.

13.11 Exercises

13.1 Concept check · 2 min

In scenario 4, Ramp, the robot trots up the 25° ramp with its pitch measured from level and falls when a knee touches the ramp. Which limit ends the run?

(a) The friction cone, because $\tan 25°$ exceeds the 0.6 the allocation plans with
(b) The leg rating, because the uphill legs must carry the weight at 150 N each
(c) The leg reach, because the rear legs stretch to their full 38 cm and lose the ground
(d) The pitch moment clamp, because holding the back level needs more than 60 N m

13.2 Concept check · 2 min

In scenario 5, Slippery patch, the allocation plans with a friction of 0.6 on a patch whose true coefficient is 0.15. After the push at 3.3 s the body pitches 13.9° nose-up instead of braking level. Why?

(a) The push lands above the center of mass and rotates the body nose-up
(b) The pitch integral winds up against the 60 N m moment clamp
(c) The IMU reads the slip as pitch, and the pitch loop follows the false reading
(d) The vertical forces cancel a nose-down braking moment that the ground never returns

13.3 Concept check · 2 min

In scenario 7, Missed contact, with Forces go to set to Scheduled feet, a front foot steps into the hole while the schedule says it is planted. What happens to the force the allocation assigns to that foot?

(a) It never reaches the body, and the allocation still believes it was delivered
(b) The contact law passes it to the diagonal partner, which then carries both shares
(c) The allocation reports it as a shortfall, and the loops treat it as saturation
(d) It reaches the body after the 4 ms leg lag, once the foot lands

13.4 Theory and calculation · 6 min

Use Equations (13.13) and (13.14) with the defaults of Table 13.3 and $m = 12\,\text{kg}$.

(a) Remove the integral from the height loop. Find the natural frequency and the damping ratio of $ms^2 + K_{c,h}\tau_{D,h}s + K_{c,h} = 0$. *Checkpoint:* 11.2 rad/s and 0.67.

(b) With the integral, $s = -2.73$ is a root of the cubic. Use the sum and the product of the roots to find the other two. *Checkpoint:* $-6.13 \pm 7.34j$ 1/s.

(c) Find the roots of the speed loop. *Checkpoint:* -1.46 and −8.54/s.

13.5 Theory and calculation · 6 min

Use Equation (13.15) with $L_h = 0.19\,\text{m}$, $h = 0.28\,\text{m}$, $mg = 117.72\,\text{N}$ and $\mu_{plan} = 0.6$.

(a) For the robot standing still with no moment request, $x_f = L_h$ and $F_z = mg$. Find the tipping term and the friction term. Which binds? *Checkpoint:* 79.9 N and 70.6 N, friction.

(b) At the braking peak after the 10 N s push, $x_f = 0.166\,\text{m}$, $F_z = 135.6\,\text{N}$ and $M_y = 6.78\,\text{N m}$. Find both terms again. *Checkpoint:* About 56 N and 81.4 N, tipping.

(c) The push starts the body at J/m. Braking at a constant 56.6 N from the first instant, how far does it travel before it stops? Compare with the engine slide of Table 13.4. *Checkpoint:* About 7.4 cm against 10.1 cm.

13.6 Theory and calculation · 5 min

In scenario 3, Added backpack, a 4 kg pack centered at $\mathbf{b} = (-0.10, 0.085)$ m goes on a 12 kg robot whose height and pitch integrals are off. Take $g = 9.81\,\text{m/s}^2$.

(a) Find the extra weight and the height sag it causes, given that the bias of the height loop stays at the nominal weight. *Checkpoint:* 39.24 N and 2.62 cm.

(b) Solve Equation (13.16) for the steady pitch. Start from the first-order value $m_b g\,0.10/K_{c,\varphi}$ and iterate once. *Checkpoint:* 3.75°, then 3.95°, and the solution is 3.96°.

(c) Find the offset $\mathbf{c}$ of the combined center of mass from Equation (13.2). *Checkpoint:* 2.50 cm back and 2.13 cm up.

13.7 App challenge · 8 min

Load scenario 2, Push recovery, at the default seed, and run it to $t = 4\,\text{s}$. Read the Worst pitch and Push recovery tiles of the scorecard.

(a) Report the nose-down dip and the recovery time at the default Speed gain of 120 N s/m. *Checkpoint:* 5.1° and 0.25 s, within 0.1° and 0.01 s.

(b) Press Reset, set Speed gain to 300 N s/m, and run to 4 s again. Then Reset with 60 N s/m. Report both pairs. *Checkpoint:* 9.3° and 0.33 s, then 3.9° and 1.02 s.

(c) Explain why the stiffest gain does not recover fastest, and why its dip is the largest. *Checkpoint:* Tipping caps braking, so a stiffer loop buys a harder pitch disturbance, not a proportional cut in the slide.

13.8 App challenge · 8 min

Load scenario 5, Slippery patch, at the default seed, and run it to $t = 8\,\text{s}$. Read the Foot slip, Worst pitch and Push recovery tiles.

(a) Report the three tiles with Friction the allocation assumes at its default of 0.6. *Checkpoint:* 15.6 cm, 13.9° (from −5.6) and 1.42 s.

(b) Press Reset, set the friction to 0.15 before pressing Run, and run to 8 s. Repeat with 0.10. Report the tiles for both. *Checkpoint:* 0.15: 1.1 cm, 3.7°, 0.59 s. 0.10: 0.0 cm, 3.9°, 0.77 s.

(c) The true coefficient is 0.15. Which setting would you ship for this patch, and why is the most conservative one not the best? *Checkpoint:* 0.15 recovers fastest of the three. Below it the allowed braking is smaller and recovery slows.

13.9 App challenge · 10 min

Load the challenge scenario, Challenge · short mixed course, at the default seed. Each run stops itself 3 s after the finish line or 0.35 s after a fall. Read the Robot, Course, Worst pitch, Push recovery and Foot slip tiles.

(a) Run it as shipped. When does it fall, and which part of the target had it already failed before the fall? *Checkpoint:* Falls at 7.63 s, after the push had slid the feet 4.6 cm.

(b) Press Reset, set Friction the allocation assumes to 0.15 and Footholds to Terrain map, and run again. Report the finish time, the worst pitch, the recovery and the slip, and the margin of each against the target. *Checkpoint:* 9.25 s, 4.0°, 0.55 s, 0.5 cm. Margins 6.0°, 1.45 s, 1.5 cm.

(c) Press the reload button, which restores the scenario's settings, and set only Footholds to Terrain map. Does the run pass? Explain with cases 5 and 6. *Checkpoint:* It finishes at 9.68 s but slides 4.6 cm and fails.

13.10 Application area · 4 min

The energy figure of Section 13.5.5 is a proxy with an assumed copper loss $k_{cu} = 0.01\,\mathrm{W/N^2}$ per leg.

(a) Evaluate Equation (13.17) for the trot at 0.4 m/s and compare with the measured 178.3 J/m. *Checkpoint:* 173 J/m and a cost of transport of 1.47, against 1.51 measured.

(b) The proxy reads the same 178.31 J/m for every swing clearance from 3 to 12 cm. Should a design team use it to choose the clearance of a real robot? *Checkpoint:* No. The legs are massless, so clearance costs nothing in the model.

13.11 Application area · 4 min

Friction cones in practice.

(a) Find the half-angle $\arctan\mu$ of the cone for the floor (0.8), the default plan (0.6) and the patch (0.15). *Checkpoint:* 38.7°, 31.0° and 8.5°.

(b) In the challenge with a blind 8 cm swing and $\mu_{plan} = 0.15$, the slope estimate near the step tilts by 8.9 to 9.6°. What does that do to the forward force the allocation may use, and what removes the problem? *Checkpoint:* Every allowed force points slightly backward. The terrain map supplies the true ground normal.

13.12 Review the agent · 5 min

A trainee asks an AI assistant how to make the standing robot of scenario 2 stop shorter after the 10 N s push. The assistant replies: *"The speed loop is a second-order system with roots set by Equation* (13.14)*. Raising the Speed gain from 120 to 300 N s/m moves the slow root and cuts the predicted slide from 6.8 to 3.0 cm, so the robot will stop in less than half the distance with no other effect."* Find the flaw and say how you would check it in the app.

CHAPTER 14

Defending a Currency Peg With a Finite Vault and Too Few Instruments

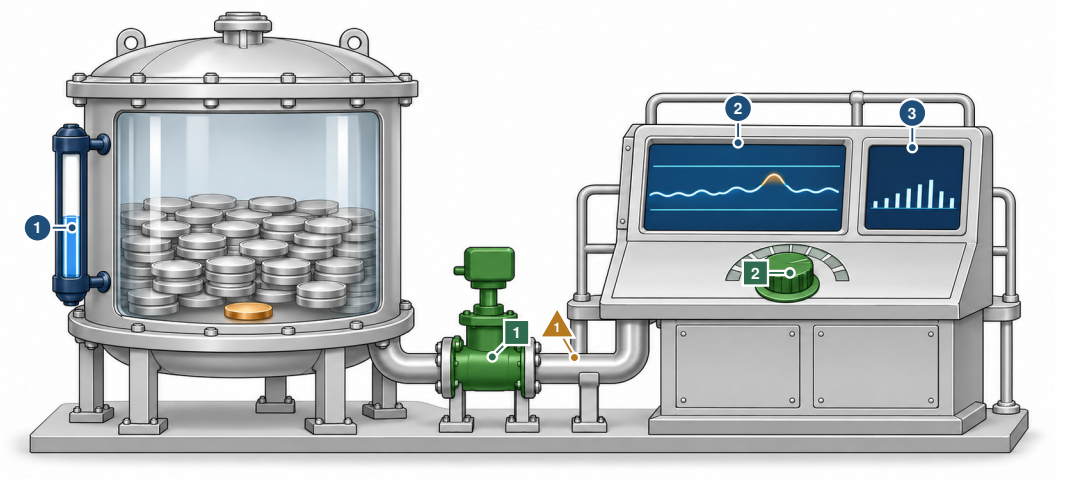

The controller moves 1 sales from the reserve vault 2 policy interest rate

The controller reads 1 reserves on the bank's books 2 exchange rate inside its band 3 quarterly output report

What disturbs it 1 money leaving the currency, an outflow or an attack

The vault pays for every sale through the green valve while the dial sets the interest rate, and the peg can end when the reserves fall to the floor marked by the amber disc. Illustration, not to scale.

IN PLAIN TERMS **Holding a Currency Steady With a Limited Vault**

The main idea. A central bank can promise a fixed price for its money in another currency. To keep the promise, it sells from a vault of foreign money when more people sell its money than buy it.

The challenge. The vault is limited. When it runs low, people expect the promise to break, and they pull money out even faster. A plan made once a week is too slow for a sudden attack, so a fast backup rule briefly limits money leaving the country.

Why it matters. Measuring each defense shows what it can buy. When the bank keeps lending to its government, no defense moves the day the promise breaks.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history segment and the fact on the foreign exchange market (Section 14.2), the two try-it segments and the fact on sterilization (Section 14.6) and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the measured cases and the challenge (Section 14.6), Technology in Practice (Section 14.7), Limitations (Section 14.8) and the two open challenges

The control question A central bank holds its exchange rate in a band with a vault of reserves that has a floor, and below a threshold the market's fear of a devaluation makes the vault drain faster. What can a currency board, a PI desk, a capped rate defense and an MPC that weighs defending against floating buy with two handles, and what does no defense change?

Learning objectives

- Explain why a reserve clock that divides the reserves by the outflow is optimistic once expectations feed on the vault.
- Derive the first-generation date of an attack from the shadow rate, and show why no intervention policy or policy rate moves it.
- Compare sterilized and unsterilized defense through the drain law and the overnight rate.
- Measure in the app what a capped rate defense, the MPC's float and a capital-controls override do against a speculative attack.
- Judge a defense by counting its limits against its instruments.

Prerequisites Sections 1.4.3 (selectors and override), 1.4.5 (saturation, rate limits and the calendar), 1.4.9 (actuators that deplete) and 1.6.2 (velocity form). Section 1.6.3 is needed only for the MPC, and Chapter 8 uses the same calendar.

The app https://apmonitor.com/apps/currency-peg/

See also Volume 6, Chapter 2 for two actuators against more variables worth holding and Volume 6, Chapter 11 for the same decision calendar on a chain of inventories. Volume 9, Chapter 7 for a store whose release is decided ahead of a forecast inflow and Volume 9, Chapter 11 for a discrete decision priced by optimization.



14.1 Introduction

A central bank that pegs its currency keeps a vault of foreign reserves and sells from it whenever the market wants more foreign currency than it offers. In the simulator of this chapter, capital that starts leaving at 6 % of gross domestic product (GDP) a year drains a vault that holds 15 % of GDP. A year in, the clock that divides the reserves above the floor by the outflow on the bank's books reads 9.8 months. The peg lasts eight, because below a threshold of 8 % of GDP the market starts to expect a devaluation, and the outflow grows to 9.8 % of GDP a year by year 1.6. When the bank instead lends the government 2 % of GDP a year without sterilizing, the peg ends at year 2.252 under every controller in the application, against a date of 2.25 years that follows from the model's equations alone, and the day before it the shipped intervention desk still holds 12.15 % of GDP.

The defense of a peg has been written as a problem of feedback and optimization for almost fifty years. A government that defends a price ceiling with sales from a stockpile must eventually meet a sudden attack by speculators [1], an argument made first for gold and then for a fixed exchange rate defended while domestic credit grows [2], after which a linear example gave the date of the attack in closed form [3]. An exchange rate held in a target zone by a policy that limits its variation is shaped by that expectation even while it sits inside the zone and nobody is defending it, a problem formally similar to option pricing [4]. Second-generation models center on the government's continuous comparison of the net benefits of devaluing and of defending [5], which is a switching decision, and the interest-rate defense has been analyzed as a policy that delays a crisis up to a point and hastens it beyond that point [6]. The trilemma, a stark tradeoff among exchange stability, monetary independence and open capital markets [7], goes back to stabilization policy analyzed under highly mobile capital [8], and in control terms it counts instruments against targets. Control is the natural language for the problem, since the domain already reasons in it. The band is a hard output constraint, the reserves are an inventory with a floor, the interest rate is a second actuator whose cost lands on another output, capital controls are a third actuator thrown as an override, and abandoning the peg is a decision to stop controlling.

The same shape recurs well outside currencies. A level held above a floor by an inventory that pays for every imbalance, a drain that grows as the level falls, a second actuator that draws the inflow

back at a cost charged to a different output, and a fast override that adds authority for a limited time describe a bank's liquid assets facing depositors, and speculation against a currency has been likened to a run on a bank [5]. They also describe a vessel kept off a low-level trip by a makeup stream whose demand rises as the level falls, and the nitrogen tank of Chapter 12 likewise gives out by its pressure while nitrogen is still in it. The misconception this chapter corrects is that a vault lasts as long as its reserves divided by the outflow, and that a defense tuned well enough can hold any peg its reserves can pay for. The clock of the first case is optimistic by 1.8 months, and in the deficit case no controller moves the attack by a single day, while the defenses that keep the most reserves hand the most to the speculators, 11.44 % of GDP under a hand that holds the policy rate at 8 % against 7.14 % under the shipped desk.

However, the crisis models compute the date of an attack from perfect foresight or from the conditions of an equilibrium, with the policy as a parameter of the analysis [5, 9], so a student seldom has a peg on which a hand, a currency board, a PI desk, a rate rule and an optimizer defend one vault through one calendar against flows that answer the money market. This chapter describes such a plant completely and measures on it what each defense buys and what no defense can change. The economics comes from the literature cited below: the shadow rate, the first-generation date, the second-generation choice and the trilemma. The plant is a continuous-time vault whose weirs conserve reserves against the integrated flows to 1.4×10^{-13} points of GDP in every case, a money market that raises a currency board's overnight rate by exactly $1/\alpha$ per point sold, and a shadow rate that reproduces a linear first-generation date to within a day at five settings. The bank adds an MPC that prices floating against defending at every decision and a capital-controls override, all on the decision-calendar module that the economics applications share. Nine cases and a graded challenge carry the measurements, and every number this chapter quotes from them was measured headless on the application's engine.

Section 14.2 traces the problem from a gold stockpile to a franc floor. Section 14.3 gives the economy, its vault, money market, capital flows and shadow rate, and Section 14.4 the reasons its vault drains faster than its clock and the date that no controller moves. Section 14.5 gives the bank's two handles, its calendar, the PI desk, the rate defense, the MPC that chooses between defending and floating, and the capital-controls override. Section 14.6 reports the nine cases and the challenge, Section 14.7 sets the bank's handles beside the instruments of real central banks, Section 14.8 states the limitations and Section 14.9 summarizes. The exercises in Section 14.10 run the cases.

14.2 From a Gold Stockpile to a Franc Floor

Speculative attacks on fixed rates were once taken for market pathologies that well-functioning markets would avoid, and the work of the late 1970s recast them as the competitive response of such markets to an inconsistent government policy [10]. The first model concerned a government that pegs the price of gold, or defends a ceiling on it, with sales from its stockpile, and it showed that the policy must end in a sudden attack by speculators [1]. The same reasoning was applied to a fixed exchange rate defended by a stock of reserves while domestic credit expands [2], and a linear version gave the date of the attack explicitly [3]. In those classical models a fixed rate contains inflationary pressures that end in a sudden balance-of-payments crisis, which frees the currency to depreciate [5]. A survey of 1991 set out the construction this application follows. The shadow floating rate is the rate that would prevail if the reserves were gone and the currency floated, the fixed rate is safe while the shadow rate stays on the strong side of it, and the regime collapses on the day the two meet, the earlier the larger the interest semi-elasticity of money demand [9].

Europe supplied the experience that tested the theory in the exchange rate mechanism (ERM) of the European monetary system, whose crisis of 1992 the history segment below recounts. In August 1993 most ERM bands were widened from ± 2.25 % to ± 15 % [5].

HISTORY 1992: Unlimited Intervention at the Edge of a Band

The exchange rate mechanism obliged both central banks to intervene when a bilateral rate reached the edge of its band, 2.25 % from the central parity in the normal band, and its founding act made those interventions unlimited. Capital controls had protected central banks' reserves against speculative attacks until the Single European Act mandated their elimination by July 1, 1990, while still allowing emergency controls for no more than six months. On September 16, 1992, the Bank of England reportedly spent as much as $20 billion, half its foreign exchange reserves, raised its discount rate from 10 to 12 % and announced a second increase to 15 %, and sterling left the mechanism at the end of the day. That month the Swedish central bank raised its marginal lending rate from 16 to 75 %, and overnight rates later reached an annualized 500 % [11].

Those attacks struck pegs that were not unsustainable in the sense of the classical models, and a second generation of models showed why. The discomfort a government suffers from speculation shapes the speculators' incentives and can admit more than one equilibrium, so even a sustainable peg may be attacked and broken [5]. Sterilization exposed a gap of a different kind. In practice a monetary authority insulates the monetary base from an attack by sterilizing the reserve outflow, whereas in the standard models a sterilized attack leaves a regime that either works indefinitely or collapses at inception, a contradiction that the Mexican crisis of December 1994 brought forward [10]. In Thailand the baht came under attack in late 1996 and twice more early in 1997, the government reportedly spent $16 billion in its defense, and the baht was cut loose on July 2, 1997. Creditors then withdrew capital from the region with little effort to distinguish among its economies, and by early September the currencies of Indonesia, Malaysia, the Philippines and Thailand had each fallen by 20 % or more [12].

Two regimes mark the ends of the design space. Hong Kong's linked exchange rate system has held 7.8 Hong Kong dollars to the US dollar since October 1983, and in May 2005 it added a convertibility undertaking on the strong side at 7.75 and moved the weak side to 7.85, a symmetric zone that made the system less discretionary and, by the evidence of option prices, more credible [13]. From September 2011 to January 2015 the Swiss central bank defended a floor of 1.20 francs per euro from the other side, by buying foreign currency without limit, and option prices put the floor's three-month survival probability below 75 % in its first months and at 95 % by August 2014 [14]. More than 130 years of interest rates largely bear out the trilemma's constraints [7], although a global financial cycle driven by the center country's monetary policy can bind even a floating rate, which leaves targeted capital controls among the remaining options [15]. The application's weirs are the ERM's and Hong Kong's unlimited interventions at the edges, its desk is intervention inside the band, and its override is a temporary flow measure.

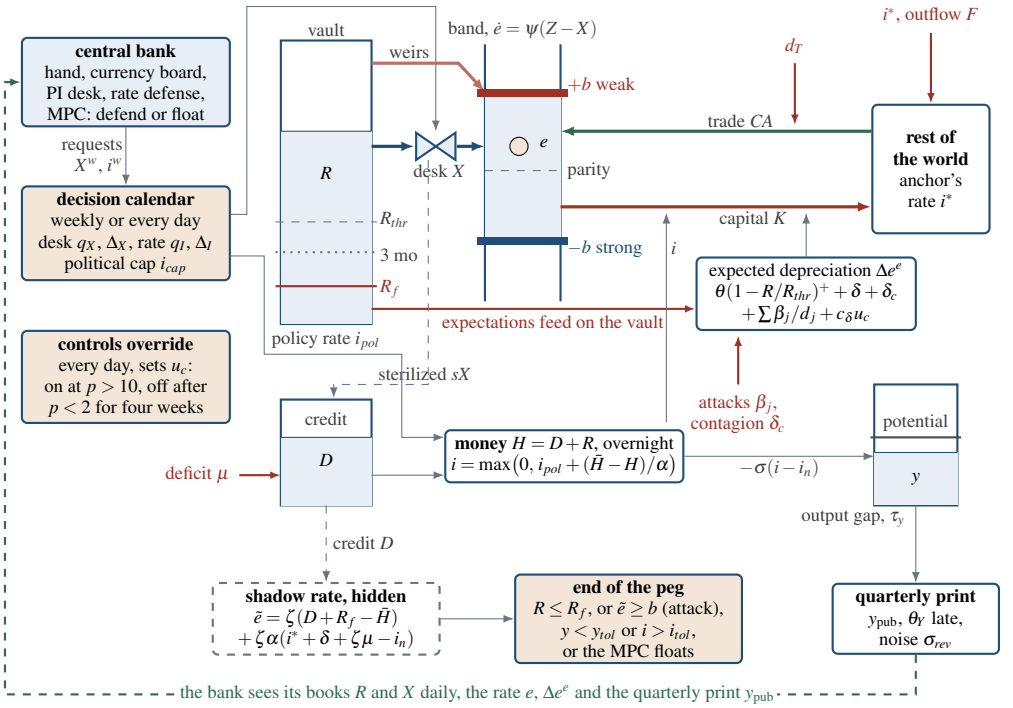

Figure 14.1 The peg as a vault, a band and a money market, not to scale. The vault R pays the desk's sales X and whatever the weirs at $\pm b$ sell or buy, and the float e moves at $\psi(Z - X)$ on the private demand $Z = -(CA + K)$. Domestic credit D grows with the monetized deficit μ and, when the bank sterilizes ($s = 1$), with every sale, money is $H = D + R$, and the overnight rate i rises above the policy rate i_{pol} when H falls below $\bar{H}$. Capital K answers i against the anchor's rate i^* and expected depreciation Δe^e, which grows below the threshold R_{thr}. Output y pays for i. The hidden shadow rate $\tilde{e}$ follows D, and the peg ends at the floor R_f, when $\tilde{e}$ reaches b, past the government's tolerance or by the MPC's choice. Thick arrows that enter from outside the plant (red) are disturbances, thin arrows (gray) carry decisions and signals, the dashed box marks the hidden shadow rate, and the long dashed path along the bottom and the left edge (green) is what the bank sees.

> ✦ **DID YOU KNOW? Seven and a Half Trillion Dollars a Day**
>
> By the triennial survey of central banks, foreign exchange trading averaged $7.5 trillion a day in April 2022, up 14 % from 2019 [16]. The 1992 defense in the history segment above spent as much as $20 billion, and that was half of one central bank's foreign exchange reserves [11].

14.3 A Vault, a Band and a Money Market

The simulated economy is a small open economy pegged to an anchor currency, in years. Stocks are in percent of GDP, flows in percent of GDP a year, interest rates in percent a year, and the exchange rate e in percent above the central parity, so that a positive e is a weaker currency. Figure 14.1 draws the plant as the application does, a vault, a channel with two weirs and arrows for the flows, and Table 14.1 classifies its signals.

Table 14.1 Signals of the Peg by Role, With Their Defaults

Symbol	Role	Default	Range
e	CV, exchange rate from parity, SP e_{sp}	0, SP 0	$\pm b$ pegged
R	CV and state, reserves, floor R_f	15 %	—
i	output, overnight rate, judged in the challenge	5 %	≥ 0
y	output and state, true gap, hidden from the bank	0	—
y_{pub}, $\hat{y}$	measured, quarterly print, and the bank's estimate	0, 0	—
X	MV, desk sales inside the band	0	−20 to 60
i_{pol}	MV, policy rate	5 %	0–25
u_c	MV, capital controls in force	0	off, on, override
F, i^*	DVs, persistent outflow, anchor's rate	0, 3 %	0–20, 0–10
μ, d_T	DVs, monetized deficit, trade shock	0, 0	0–5, −6–3
δ_c, β_j	DVs, contagion premium, attack bets	0	events
D, H	state and output, domestic credit, money	35, 50	—
$\tilde{e}$	output, shadow rate, hidden	−20	—
Δe^e	output, expected depreciation (market price)	2	—
q, I_y	states, integrated flows, integrated gap	0	—
C, P_s, Π	states, defense cost, private holdings, profit	0	—

Stocks in % of GDP, flows in % of GDP a year, rates and e in %.

14.3.1 The Vault, the Float and the Weirs

The bank's reserves R fall by what its desk sells, and the exchange rate floats on the private demand for foreign currency that the desk does not meet, as shown in Equation (14.1).

$$\dot{R} = -X, \qquad \dot{e} = \psi\,(Z - X), \qquad Z = -(CA + K) \tag{14.1}$$

where $\psi = 20$ percent of rate per point of GDP, CA is the trade balance and K the net capital flow. After each integration step the weirs act. A rate past the weak edge b is put back at b and the excess $(e - b)/\psi$ is sold from the vault, and a rate past $-b$ is put back at $-b$ and the shortfall bought into it. The weirs are marginal intervention and the desk is intramarginal intervention, and the vault pays for both. Because the weirs move R and e together, the vault always holds what the flows say it should, as shown in Equation (14.2).

$$R - \frac{e}{\psi} - q = \text{const}, \qquad \dot{q} = CA + K \tag{14.2}$$

The application's check finds the identity held to 1.4×10^{-13} points of GDP in every scenario while its peg stands. The default band is $b = 1\,\%$, a teaching value, with presets for the ERM's narrow band of 2.25 % and Hong Kong's undertakings at 7.75 and 7.85 to the dollar, $\pm 0.641\,\%$ of 7.80.

14.3.2 Money, Sterilization and the Overnight Rate

Money is domestic credit plus reserves, and the overnight rate clears a linear money market around the money held at the peg, as shown in Equation (14.3).

$$\begin{aligned} H &= D + R, \qquad \dot{D} = \mu + sX \\ i &= \max\left(0,\; i_{pol} + \frac{\bar{H} - H}{\alpha}\right) \end{aligned} \tag{14.3}$$

where $\bar{H} = 50\,\%$ of GDP, $\alpha = 1.5$ points of GDP per point of rate, μ is the monetized deficit and s is 1 when the bank sterilizes and 0 when it does not. A weir sale adds s times the sale to D in the same way. Unsterilized, a point of GDP sold takes a point of money out and raises the overnight rate $1/\alpha = 0.67$ points above the policy rate, the automatic adjustment of a currency board, which the application's check confirms on every day of case 2 to 3.6×10^{-15}. Sterilized, the bank buys domestic bonds with the money a sale brings in, so money and the overnight rate stay where they were and domestic credit rises by the sale instead.

14.3.3 Capital, Expectations and Attacks

Capital moves on uncovered interest parity with imperfect mobility, against an expected depreciation that grows as the vault falls below a threshold, as shown in Equation (14.4).

$$\begin{aligned} K &= \kappa_c\,(i - i^* - \Delta e^e) - (1 - c_K u_c)F \\ \Delta e^e &= \theta \max\!\left(0, \frac{R_{thr} - R}{R_{thr}}\right) + \delta + \delta_c(t) + \sum_{j\,\text{open}} \frac{\beta_j}{d_j} + c_\delta u_c \end{aligned} \tag{14.4}$$

where $\kappa_c = \kappa(1 - c_K u_c)$, $\kappa = 1.5$ points of GDP a year per point of return, $\theta = 10\,\%$ a year with the vault empty, $R_{thr} = 8\,\%$ of GDP, $\delta = 2$ points of fundamental premium, and u_c is 1 while capital controls are in force. Below the threshold each point of reserves lost adds $\theta / R_{thr} = 1.25$ points to Δe^e, which the check confirms to 10^{-12}, and the plant's feedback turns positive. A speculative attack j is a bet on a devaluation of $\beta_j = S_j / \kappa_0$ percent within d_j years, with $\kappa_0 = 1.5$, which at the shipped mobility moves S_j points of GDP out of the currency. The three attacks move 5 points in a quarter ($\beta = 3.33$, 20 % of GDP a year), 10 points in a quarter ($\beta = 6.67$, 40 % a year) and 15 points in six weeks ($\beta = 10$, 130 % a year). Contagion δ_c adds 5 points for half a year. Controls cut capital mobility and the persistent outflow F by $c_K = 80\,\%$ at a trade cost $c_T = 1$ point of GDP a year and a credibility cost $c_\delta = 1$ point of premium.

14.3.4 Output and the Trade Balance

The overnight rate, not the policy rate, is what output pays, and the trade balance answers the gap and the position of the rate in the band, as shown in Equation (14.5).

$$\begin{aligned} \tau_y \dot{y} &= -y - \sigma(i - i_n) + \phi\,(\eta e + d_T - c_T u_c) \\ CA &= -m_y y + \eta e + d_T - c_T u_c \end{aligned} \tag{14.5}$$

where $\tau_y = 0.5\,\text{yr}$, $\sigma = 0.7$ points of gap per point of rate, $i_n = 5\,\%$ is the neutral rate, $\phi = 0.5$, $\eta = 0.2$ points of GDP a year per percent of rate and $m_y = 0.3$. At the defaults $i_n = i^* + \delta$, so capital, trade and the vault are all at rest.

14.3.5 The Shadow Rate and the End of the Peg

The shadow rate is where the currency would float if speculators took every reserve above the floor and the bank stopped defending. Money would then be $D + R_f$, the price level would follow money, and money demand would fall with the float's interest rate, the anchor's rate plus the premium plus the trend depreciation $\zeta\mu$, as shown in Equation (14.6).

$$\tilde{e} = \zeta\,(D + R_f - \bar{H}) + \zeta\alpha\,(i^* + \delta + \zeta\mu - i_n), \qquad \zeta = 100/\bar{H} \tag{14.6}$$

where $\zeta = 2$ percent of rate per point of GDP and $R_f = 5\,\%$ of GDP. The shadow rate depends on domestic credit and not on the reserves in the vault, which the application checks. After each step the peg ends on the first of four conditions. When $\tilde{e} \geq b$ with $R > R_f$, the reserves above the floor are free money and are taken in one day, too fast to sterilize, which is the first-generation attack. When $R \leq R_f$, the bank gives up at its floor. When the government's tolerance is on and $y < y_{tol}$ or $i > i_{tol}$, with $y_{tol} = -3\,\%$ and $i_{tol} = 12\,\%$ (each moved halfway to neutral by a political shock), the defense has cost more than the government will bear. When the MPC of Section 14.5 asks to float, the peg ends by choice. At the exit the bank stops selling and sets its rate to i_n, the open attacks close, and the currency jumps to

$$\begin{aligned} e^+ &= \max\left(\tilde{e},\; b + J\right) \\ J &= \textstyle\sum_{j\,\text{open}} \beta_j + \theta \max\left(0, 1 - R/R_{thr}\right) + \delta_c \end{aligned} \tag{14.7}$$

with $J = 0$ at a first-generation attack, after which e drifts at $\zeta\mu$ and the vault is frozen.

Table 14.2 Teaching Values. M Marks a Model-Dialog Slider, D a Disturbance, A an Advanced Control and C a Source Constant, and None Is Hidden

Symbol	Meaning	Value	Range
M, model-dialog sliders			
$R(0)$, R_f	reserves at the start, floor (% of GDP)	15, 5	6–25, 0–10
R_{thr}, IM	threshold (% of GDP), imports (%/yr)	8, 30	4–14, 10–60
κ, θ	capital mobility, devaluation expected when empty	1.5, 10	0.25–5, 0–30
δ	fundamental premium (points)	2	0–6
α, $\bar{H}$	money demand slope, money at the peg	1.5, 50	0.25–5, 20–100
i_n	neutral domestic rate (%)	5	1–10
ψ	rate moved (%) per unmatched point of GDP	20	5–100
η, m_y	trade response to e and to y	0.2, 0.3	0–1, 0–1
σ, τ_y	output response to i, gap lag (yr)	0.7, 0.5	0.1–2, 0.1–2
ϕ	output response to trade	0.5	0–1.5
c_K, c_T, c_δ	controls: share stopped, trade cost, premium	0.8, 1, 1	0–1, 0–3, 0–3
y_{tol}, i_{tol}	government's tolerance (%)	−3, 12	−8–−0.5, 6–25
θ_Y, σ_{rev}	print delay (yr), print noise	1/12, 0.3	0–13 wk, 0–1.5
D, disturbances			
F, d_T	persistent outflow, trade shock	0, 0	0–20, −6–3
i^*, μ	anchor's rate (%), monetized deficit	3, 0	0–10, 0–5
—	political shock: tolerance factor	0.5	—
A, advanced control			
b	band half-width (%)	1	0.641, 2.25
C, source constants			
κ_0	mobility that sizes an attack	1.5	—
—	attack sizes, lengths (yr)	5, 10, 15 in 1/4, 1/4, 6/52	—
—	contagion premium, length	5, 0.5 yr	—
Δt	integration step, control period, log	1/365 yr	—
—	seed	20261005	—

14.3.6 What the Bank Can See

The bank reads its reserves, its domestic credit, its sales and the credit it extends to the government every day from its own books, and the exchange rate, the overnight rate, the anchor's rate and expected depreciation as market prices. It sees output only through a quarterly print of the quarter's average gap, published $\theta_Y = 1/12$ year after the quarter ends, as shown in Equation (14.8).

$$y_{\text{pub},j} = 4\int_{(j-1)/4}^{j/4} y\,\mathrm{d}t + \sigma_{rev}\,\nu_j \tag{14.8}$$

where $\sigma_{rev} = 0.3$ points and ν_j is a standard normal deviate from the engine's seeded generator. The shadow rate and the speculators' position are the model's and never the bank's.

14.3.7 Start, Integration and Bookkeeping

Unless a case declares otherwise, it starts at rest with $R = R(0) = 15\,\%$ of GDP, six months of the imports of 30 % of GDP a year, $D = \bar{H} - R(0) = 35$, $e = 0$, $y = 0$ and the policy rate at 5 %. The engine of Section 1.6.1 integrates the nine states by classical fourth-order Runge–Kutta on a step of a day, $\Delta t = 1/365$ year, the controller runs once a day ahead of each step, and every day is logged. The default equilibrium has derivatives of exactly zero and holds for two years with no drift, and at half the step the challenge keeps both verdicts with its lowest reserves within 0.0091 points of GDP. Four integrals keep the books, the flows q of Equation (14.2), the gap I_y for the print, the defense cost $C = \int [y^2 + ((i - i_n)^+)^2]\,\mathrm{d}t$ and the private sector's holdings P_s bought from the vault, with a teaching profit Π that collects the gain on the rate's moves less the carry given up. Noise comes from seed 20261005 by default. Table 14.2 gathers the constants.

14.3.8 Assumptions

Expectations are mechanical. Expected depreciation is a premium, a term that rises as the vault falls, the contagion premium and the scripted bets. Only the first-generation attack looks ahead, through

the shadow rate. A market that learned the bank's rule would change these relations, which is the critique of [17], so the model shows what a defense does to this economy and not what it would do to one that anticipates it.

Flows, not portfolios. A constant gap in returns gives a constant capital flow, as in the textbook flow models, with no portfolio stocks that fill up, and one mobility κ serves every kind of capital. A long attack is bounded only by its scripted length.

One money and one bank. Base money and the deposits it supports are one stock H, the money market is linear about $\bar{H}$, and sterilization is all or nothing. There are no banks with balance sheets, no foreign-currency debt and no default, which leaves out the banking side of crises such as Thailand's, where public money also propped up failed banking institutions [12].

Linear money demand about the peg. The shadow rate of Equation (14.6) turns a point of GDP of money into ζ percent of price level, which is accurate only near the peg.

An ending rule, not a float. After the exit the currency follows Equation (14.7) and a drift, which is a rule for the end of the story and not a model of a floating rate.

Teaching values. Every threshold, elasticity, attack size, tolerance and rule is chosen to expose one mechanism. The real ones come from the crisis literature [2, 3, 5, 6] and from the institutions themselves, such as the ERM's margins [11] and Hong Kong's undertakings [13]. The application uses three months of imports as its rule of thumb for a safe vault, and its floor of 5 % of GDP is two.

14.4 Why the Vault Drains Faster Than Its Clock

14.4.1 A Float That Crosses the Band in a Day

The exchange rate is the fastest state in the plant. During the 10-point attack the unmatched demand is 40 % of GDP a year, which moves the rate at $\psi Z = 800\,\%$ a year, across the two-point band in 0.9 day. Inside the band the trade balance pulls the rate back at $\psi\eta = 4$ per year, a quarter-year time constant, which is slow beside an attack. The float also stores almost nothing. By Equation (14.2) a full crossing of the band is worth $2b/\psi = 0.1$ point of GDP of reserves, so where in the band the bank chooses to sell changes little about how fast the vault empties. A desk that decides once a week cannot hold the float off the weak weir in an attack, and the weir sells what the desk misses. In a tuning sweep of case 4 without sterilization, the float spends 64 % of the time from year 0.1 to year 1 at an edge under the weekly desk at the shipped gain of 3, against 7 % for a desk deciding every day at a gain of 10.

14.4.2 The Drain Law

Once the desk and the weirs meet the whole private demand, the vault absorbs the whole private flow, $\dot{R} = CA + K$, and at a given gap that flow is a static function of the vault. Evaluated with the policy rate at neutral, the trade term at zero and $y = 0$, it is the drain law of Equation (14.9).

$$\dot{R} = \kappa\left[(1-s)\frac{R(0)-R}{\alpha} - \theta\max\left(0, 1-\frac{R}{R_{thr}}\right)\right] - F \tag{14.9}$$

where the first term is the money market drawing capital back as unsterilized sales drain money, since then $H = \bar{H} - (R(0) - R)$. The vault is a depleting actuator whose drain grows as the stock falls, the case of Section 1.4.9. Figure 14.2(a) plots the law from the application's own flow function. Sterilized, the vault is neutrally stable above the threshold and loses $\kappa\theta/R_{thr} = 1.875$ points of GDP a year for every point below it, a positive feedback with a doubling time of 0.37 year. Unsterilized, every point sold draws $\kappa/\alpha = 1$ point of GDP a year back, which is enough to make the drain law positive everywhere below $R(0)$ when nothing else is leaving, so a board refills its vault after an attack ends. A persistent outflow of 6 % of GDP a year shifts both curves down by 6. The sterilized vault then drains at a constant rate to the threshold and faster below it, while the unsterilized one balances at $R(0) - \alpha F/\kappa = 9\,\%$ of GDP with the overnight rate 4 points above neutral. With the float on the weak weir, the trade balance gains $\eta b = 0.2$ point of GDP a year, which moves the sterilized point of no return from 8 to 7.89 % of GDP.

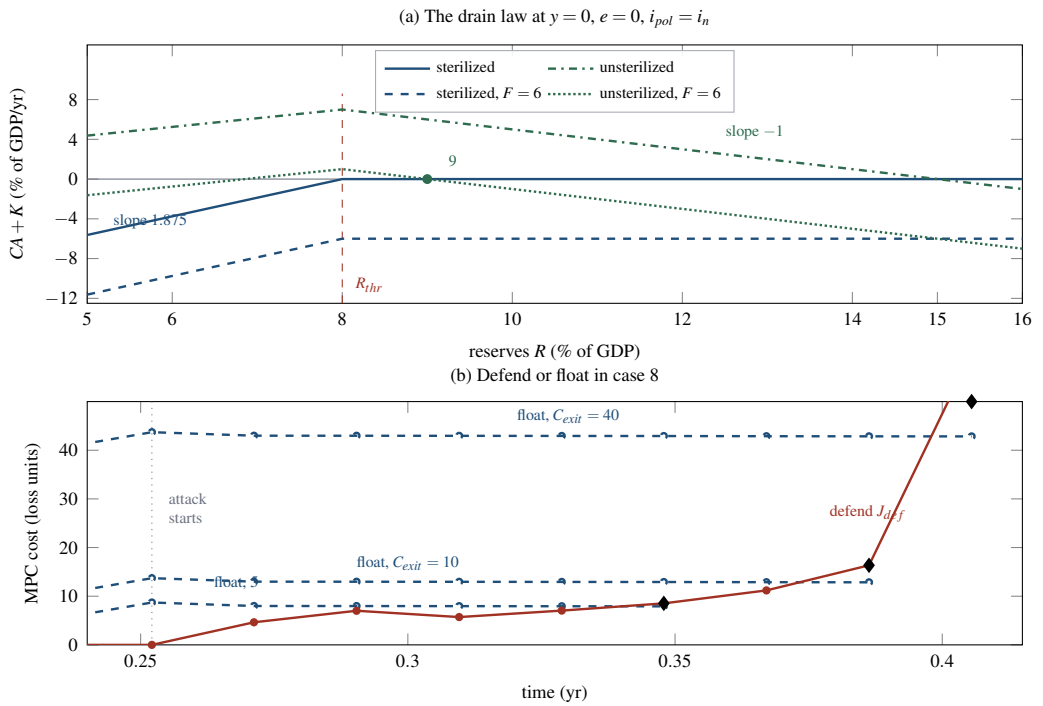

Figure 14.2 Two characteristics from the application's own functions and engine. (a) The drain law, the private inflow to the vault against the reserves held, at $y = 0$ with the float at parity and the policy rate at neutral, sterilized (solid with no outflow, dashed with 6 % of GDP a year leaving) and not (dash-dotted with no outflow, dotted with the outflow). (b) The MPC's cost of its best defense (solid, filled circles) and its cost of floating (dashed, open circles) at each weekly decision of case 8, for costs of floating of 5, 10 (shipped) and 40, with the attack starting at 0.25 year. A diamond marks the decision at which each run floats, the last one off the scale at 2231.

14.4.3 The Date No Controller Moves

With a monetized deficit and no sterilization, domestic credit grows as $D(t) = D(0) + \mu t$ whatever the bank does with its vault or its rate, and with $\bar{H} = D(0) + R(0)$ the shadow rate of Equation (14.6) reaches the weak edge b at the date of Equation (14.10).

$$T_K = \frac{R(0) - R_f}{\mu} + \frac{b}{\zeta\mu} - \frac{\alpha\,(i^* + \delta + \zeta\mu - i_n)}{\mu} \tag{14.10}$$

In case 3, $\mu = 2$ points of GDP a year, the shadow rate starts at −8 % and rises at $\zeta\mu = 4\,\%$ a year, and $T_K = 5 + 0.25 - 3 = 2.25$ years. No intervention policy and no policy rate appears in it. The application's check runs the simulation at five settings and finds the attack within a day of the formula each time, at 7.501 against 7.50 years for $\mu = 1$, 2.252 against 2.25 for $\mu = 2$, 0.501 against 0.50 for $\mu = 3$, 0.252 against 0.25 from a vault of 11 at $\mu = 2$, and 0.877 against 0.875 from 20 at $\mu = 4$. The day is the integration step, since the attack falls on the first step at which the shadow rate has reached the edge.

A survey restates the linear first-generation model, in the tradition of the linear examples of [3], with the date $t_c = \theta R_0/\mu - \alpha$ in logarithms, where money is a weighted sum of credit and reserves, θ is the ratio of the weights, reserves are exhausted at zero and α is the semi-elasticity of money demand [9]. Equation (14.10) is that date written in levels of GDP, and it differs in four places. The reserves count down to a floor R_f instead of zero. The attack comes when the shadow rate reaches the

Table 14.3 The Bank's Settings, Defaults and Ranges. A Marks an Advanced Control

Symbol	Setting	Default	Range
Who decides and when			
—	controller in AUTO	PI desk	board, PI, rate, MPC
—	calendar	weekly	weekly, every day
q_X, Δ_X	desk quantum, largest change (A)	0.5, 20	0–5, 1–80
q_I, Δ_I	rate quantum, largest move (A)	0.25, 3	0–1, 0.25–10
s	sterilization	on	off, on
u_c	capital controls	off	off, on, override
—	tolerance, band (A)	on, ±1 %	on, off, 3 bands
The hand			
X^w, i^w	desk and rate requests	0, 5 %	−20–60, 0–25
PI desk			
e_{sp}	band setpoint	0	±1
K_c, τ_I	gain, integral time	3, 0.03 yr	0–5, 0–0.5
Interest-rate defense			
k_r, i_{cap}	rate per point of outflow, political cap	0.25, 12 %	0–1, 5–25
MPC, defend or float			
w_R, w_y	price of reserves spent, weight on the gap	1, 1	0–10, 0–5
C_{exit}	cost of floating	10	0–50
H_p, τ_{att}	horizon, expected fading of an outflow (A)	2, 0.25 yr	0.5–3, 0.05–2

Desk in % of GDP a year, rates in %, k_r in points per % of GDP a year.

weak edge of a band rather than a single parity, which adds $b/(\zeta\mu) = 0.25$ year. The semi-elasticity term is $\alpha\zeta$, the same flight of money divided by the credit growth, and it reduces to the textbook's form only while $i^* + \delta = i_n$, as it does in every shipped case. The fourth difference is in the path to the date. In the textbook model the fixed rate holds money demand constant, so reserves run down in proportion to the growth of credit and the attack takes just the money that demand no longer wants. Here capital is imperfectly mobile and the overnight rate clears the money market, so the vault falls more slowly than credit grows and the attack takes whatever stands above the floor. In case 3 the first-generation count of the vault, $R(0) - \mu t$, stands at 10.6 % at year 2.2, against 12.25 % in the vault under the shipped desk, and the attack takes 7.14 points where the count puts only 5.5 above the floor. The difference is the excess money, 1.65 points above $\bar{H}$ that the market holds at an overnight rate of 3.90 %. The date survives unchanged because it depends on credit alone. Sterilized sales are credit as well and bring the day forward. The standard model cannot represent that channel, because there a sterilized attack leaves a regime that either works indefinitely or collapses at inception [10], and this model keeps it because the final attack is never sterilized.

14.5 Two Handles, a Calendar and an Override

The bank has two handles, the desk's sales X inside the band and the policy rate i_{pol}, and one more switch, capital controls. The controller executes every day. At each execution it updates the bank's estimate of the gap, reads its books and runs the controls override. It then checks the calendar, and only an execution that counts as a decision consults the selected policy. Table 14.3 lists the settings.

14.5.1 Weekly Instants, Quanta and Caps

Both the calendar and the shaping are those of `core/cadence.js`, a module of the framework that the economics applications share. The same shaping governs the cash decisions of Chapter 8 and the policy rate of Volume 2, Chapter 11, and Section 1.4.5 sets it beside the rate limits of other chapters. The weekly calendar places an instant every $1/52$ year from $t = 0$, and a decision happens at the first daily execution that reaches an instant not yet acted on, so decisions fall seven or eight days apart. The continuous calendar makes every daily execution a decision. At a decision each handle's request

is shaped as shown in Equation (14.11).

$$\begin{aligned} \tilde{u} &= q\,\mathrm{round}(u^w/q), \qquad \Delta_q = q\max\left(1, \lfloor \Delta/q \rfloor\right) \\ u_k &= \mathrm{clip}_{[u_{min},u_{max}]}\left(u_{k-1} + \mathrm{sat}_{\Delta_q}(\tilde{u} - u_{k-1})\right) \end{aligned} \tag{14.11}$$

where for the desk $q_X = 0.5$, $\Delta_X = 20$ and $[-20, 60]$ points of GDP a year, and for the rate $q_I = 0.25$, $\Delta_I = 3$ and $[0, 25]$ percent. Halves round away from zero, the bounds snap inward to the grid before the clip, a request that is not a number changes nothing, and the political cap i_{cap} trims the rate defense's and the MPC's requests before the shaping, though never the hand's. Between decisions both handles hold. The application's check finds every decision on its grid, inside its cap and on its calendar instant under the PI, the rate defense and the MPC on both calendars. In AUTO each handle's slider follows its controller's request before the quantum and the caps, and a switch of mode or controller holds both handles at the first decision after it, so every controller transfers without a jump.

14.5.2 The Hand and the Currency Board

In MANUAL the hand holds both handles, and its two sliders are requests that the next decision shapes exactly as it shapes a controller's. A helper sets the desk to the outflow the books showed over the last interval. The currency board holds both handles mechanically. Its desk sells nothing, so the weirs do all the selling, its policy rate stays at i_n, and its sales are never sterilized, whatever the switch says, so its money falls one for one with its vault.

14.5.3 The PI Desk

The desk's PI is the velocity-form proportional-integral-derivative (PID) controller of Section 1.6.2, with no derivative term, reverse acting on the rate's position in the band and updated only at a decision, as shown in Equation (14.12).

$$\begin{aligned} \varepsilon_k &= e_k - e_{sp}, \qquad P_k = K_c \varepsilon_k \\ S_k &= S_{k-1} + P_k - P_{k-1}, \qquad \tilde{X}_k = S_k + I_{k-1} \\ X_k &= \mathrm{clip}_{[-20,60]}(\tilde{X}_k) \\ I_k &= I_{k-1} + \left[K_c \varepsilon_k/\tau_I + (X_k - \tilde{X}_k)/\tau_I\right]\Delta t_k \end{aligned} \tag{14.12}$$

where $K_c = 3$ points of GDP a year per percent of rate, $\tau_I = 0.03\,\mathrm{yr}$, Δt_k is the time since the previous decision, and the symbols map onto Equations (1.13) to (1.15) with $\tilde{X}_k = v_k$, no bias and a tracking time $T_t = \tau_I$. The request is the unclamped output $\tilde{X}_k$, which the calendar then shapes, and the integral never learns of the quantum or the cap. A weaker rate asks for more sales. When the PI does not hold the desk it tracks the applied sales, and $\tau_I = 0$ removes the integral.

14.5.4 The Interest-Rate Defense

The rate defense is a proportional controller on the outflow that the books show between decisions, as shown in Equation (14.13).

$$i_k^w = \min\left(i_n + k_r \max\left(0, \frac{R_{k-1} - R_k}{\Delta t_k}\right),\ i_{cap}\right) \tag{14.13}$$

where $k_r = 0.25$ points per point of GDP a year and $i_{cap} = 12\,\%$. It holds the policy rate while the hand keeps the desk. Its gain is limited by politics rather than by stability, because the 130 % a year of the large attack would need a rate above 90 % to stop it at the shipped mobility.

14.5.5 Defend or Float as Model Predictive Control

At each decision the MPC plans the policy rate over $H_p = 2$ years in monthly steps on the bank's own model of the defended peg, in which every flow is met from the vault and every coefficient is the

declared value of Table 14.2, as shown in Equation (14.14).

$$\begin{aligned}\dot{R} &= \kappa_c\Big(u + \tfrac{\bar{H}-H}{\alpha} - i^* - \delta - c_\delta u_c - \rho\,\theta\big(1-\tfrac{R}{R_{thr}}\big)\Big)\\ &\quad - m_y\hat{y} + CA_x + d\\ \dot{H} &= \mu + (1-s)\dot{R}\\ \tau_y\dot{\hat{y}} &= -\hat{y} - \sigma\Big(u + \tfrac{\bar{H}-H}{\alpha} - i_n\Big) + \phi\,CA_x\end{aligned} \tag{14.14}$$

where u is the policy rate, $CA_x = \eta e + d_T - c_T u_c$ is frozen at the decision, ρ is 1 in the regime below the threshold and 0 above it, and d is a preview of the outflow the bank cannot explain, its smoothed reading $\hat{d}$ fading as $\hat{d}\,\mathrm{e}^{-(k+1/2)T_s/\tau_{att}}$ with $\tau_{att} = 0.25\,\mathrm{yr}$. The model is linear in each regime. It is solved in the regime the vault is in now and solved once more in the other regime when more than half of the planned months cross the threshold. The model is discretized exactly for a zero-order hold over $T_s = 1/12$ year, with the rate held over a step added as a state so that the overnight rate is an output, and the framework's linear MPC of Section 1.6.3, `core/linmpc.js`, poses the quadratic program (QP) of Equation (14.15). It keeps the move suppression and priced slacks of Equation (1.17) and replaces the tracking term with costs on the gap and the rate.

$$\begin{aligned}\min_{u,\varsigma}\ & \textstyle\sum_{k=1}^{N} T_s\big[w_y\hat{y}_k^2 + w_i(i_k - i_n)^2\big]\\ & + \nu\textstyle\sum_{j=0}^{9}(u_j - u_{j-1})^2 + \sum_{l=1}^{4}\big(r_l\varsigma_l^2 + \lambda_l\varsigma_l\big)\\ \text{s.t. }\ & 0 \le u_j \le i_{cap}, \quad R_k \ge R(t) - \varsigma_1\\ & R_k \ge R_f + 0.25 - \varsigma_2, \quad \hat{y}_k \ge y_{tol} + 0.25 - \varsigma_3\\ & i_k \le i_{tol} - 0.25 + \varsigma_4, \quad \varsigma \ge 0\end{aligned} \tag{14.15}$$

where $N = 24$, the ten move blocks are four of one month, two of two, and then 3, 4, 6 and 3 months, u_{-1} is the rate in force, $w_y = 1$, $w_i = 0.1$, $\nu = 0.05$, $r_1 = 0.02 w_R + 10^{-3}$ and $\lambda_1 = w_R$ price the deepest drawdown of the vault, and $r_l = 10^4$ and $\lambda_l = 10^3$ make the floor and the government's tolerance stiff. One slack serves each band over the whole horizon, so ς_1 is the plan's deepest drawdown. The program has 14 variables and 120 rows, a ridge of $10^{-9}(1 + d_{max})$ keeps it strictly convex, and the dual active-set method of [18] in `core/qp.js` solves it. In case 8 the MPC plans 21 times, at the start and at every weekly decision to the float, and five of the plans after the start are solved a second time in the other regime. The 25 QPs after the start take 13.1 active-set iterations on average and 22 at most, and none fails.

The first move becomes the rate request, and the plan then prices both choices, as shown in Equation (14.16).

$$\begin{aligned}J_{def} &= \textstyle\sum_k T_s\big[w_y\hat{y}_k^2 + w_i(i_k - i_n)^2\big] + \nu\sum_k \Delta u_k^2\\ &\quad + w_R \max_k\big(R(t) - R_k\big)^+ + 10^4\bar{\varsigma}^2 + 10^3\bar{\varsigma}\\ J_{flt} &= C_{exit} + w_E J_e^2\\ &\quad + \textstyle\sum_{k=1}^{N} T_s w_y\big[\hat{y}_\infty + (\hat{y} - \hat{y}_\infty)\mathrm{e}^{-kT_s/\tau_y}\big]^2\end{aligned} \tag{14.16}$$

where $\bar{\varsigma}$ is the plan's largest shortfall against the floor and the tolerance, $J_e = \max(0, b + J - e)$ is the jump the market prices by Equation (14.7), $w_E = 0.05$, $C_{exit} = 10$, and $\hat{y}_\infty = \phi(CA_x + \eta J_e)$ is the gap the cheaper currency would support with the rate back at neutral. The MPC floats when $J_{flt} < J_{def}$, which is the second-generation comparison [5] made with a number on each side. A discrete choice priced by two costs is also the core of the integer pump schedules of Volume 9, Chapter 11. If a solve fails, the rate in force stays and the failure is counted.

14.5.6 The Capital-Controls Override

Controls are off, on, or an override that decides by itself. The override reads the outflow the bank cannot explain with its own model, the attack, a portfolio shift or contagion, as shown in Equation (14.17).

$$\begin{aligned} g &= \frac{R_k - R_{k-1}}{\Delta t} + m_y \hat{y} - CA_x \\ &\quad - \kappa_c \left(i - i^* - \theta \left(1 - \tfrac{R}{R_{thr}}\right)^+ - \delta - c_\delta u_c \right) \\ \hat{d} &\leftarrow \hat{d} + (g - \hat{d}) \min(1, \Delta t / \tau_f) \\ p &= -\hat{d} / \max(0.05,\, 1 - c_K u_c) \end{aligned} \tag{14.17}$$

where $\tau_f = 1/52$ year and p is the pressure behind the outflow, with what the controls stop added back. Controls come on when $p > 10$ points of GDP a year and lift once $p < 2$ has held for four weeks. The override runs at every daily execution, outside the calendar, as a fast protective layer under the weekly policy, and the same reading $\hat{d}$ is the MPC's preview.

14.5.7 The Estimate, and the Three Models

The bank's estimate $\hat{y}$ is the first line of Equation (14.5) run forward every day with the overnight rate in force. At each print it is corrected by $0.6\,(y_{\text{pub}} - \bar{\hat{y}})\,\mathrm{e}^{-a/\tau_y}$, where $\bar{\hat{y}}$ is the bank's own average for the quarter and a the time since the quarter's middle, a teaching gain rather than a Kalman design. Equations (14.1) to (14.8), with the weirs, the four endings and the scripted events, are the simulated economy. The MPC's internal model is Equation (14.14), which sends every flow straight to the vault without a float or weirs, drops the overnight rate's floor at zero, foresees no attack the bank has not yet seen, and reads the gap only through its estimate. The models it stands for are the crisis models of [2, 3, 5] and the analysis of the interest-rate defense in [6], whose coefficients are not copied here. The PI, the board and the rate defense use no model, and nothing in the application is fitted to data. The bank's estimate, its reading $\hat{d}$ and the MPC keep the declared values of Table 14.2, and so does the neutral rate that the board and the rate defense set. A parameter changed by the learner therefore changes the economy and not the bank's models. The plan starts from the reserves and the credit on the bank's balance sheet, and the reserve floor and the government's tolerance are read as set.

14.6 Nine Defenses, Measured

Table 14.4 lists the cases. Each begins at the defaults of Tables 14.2 and 14.3, loaded through the dropdown, alters only what its row names, and is run over its listed length on the default seed. A variant changes one more setting before the run starts, and the application gives the same run whether that setting is chosen before the first Run or carried through Reset, as the briefings direct, so it is in force from the first day. Section 1.6.4 describes the briefings, the scorecard and the exported file from which these values are read. Table 14.5 collects the outcomes.

14.6.1 The Clock and the Equilibrium

In the reference the flows balance because $i - i^* - \delta = 5 - 3 - 2 = 0$, and the vault holds 15 % of GDP exactly for two years. Raised by hand to 6 %, the policy rate draws capital in at 1.5 % of GDP a year, the float goes to the strong weir, which buys it, and the vault is at 16.38 % a year later, while the gap is –0.32 % after a quarter and –0.69 % after a year. Selling 5 % of GDP a year from the desk into the balanced market does nothing lasting, because the strong weir buys back 4.74 of the 5 points within the year.

In case 1 capital leaves at 6 % of GDP a year from month 1 and the bank sterilizes. With the desk at zero the weak weir sells all ten points. The vault is 12.64 % at year 0.5 and 9.75 % at year 1, when the books show 5.8 points a year leaving and the clock reads 9.8 months. It crosses the threshold at year 1.304 and three months of imports at 1.384, after which the outflow is 7.9 % of GDP a year at year 1.5 and 9.8 at 1.6, and the peg falls at the floor at year 1.666, eight months after the clock's

Table 14.4 Cases 1 to 9, Challenge C and Reference R

#	Change from the defaults, run length	Teaching point
1	MANUAL, outflow of 6 from month 1, 2 yr	the clock is optimistic
2	currency board, 10-point attack at 0.25 yr, 1.5 yr	automatic adjustment
3	$\mu = 2$, unsterilized, 3 yr	a date no one moves
4	PI desk, 10-point attack at 0.25 yr, 1.5 yr	sterilization
5	rate defense, $i_{cap} = 15$, $\Delta_I = 10$, 15-point attack, 1 yr	a gain capped by politics
6	i^* from 3 to 6 % at 0.25 yr, 2 yr	importing the anchor's policy
7	$R(0) = 11$, contagion at 0.25 yr, 2 yr	a basin of attraction
8	MPC, unsterilized, 10-point attack at 0.25 yr, 1 yr	the choice to float
9	controls on, 10-point attack at 0.25 yr, 1.5 yr	the third instrument
C	unsterilized, 10-point attack at 0.5 yr, 1.5 yr	three limits, two handles
R	exact prints, nothing happens, 1 yr	the equilibrium

Table 14.5 How Each Peg Ended, the Lowest Vault, the Highest Overnight Rate and the Worst Gap Over Each Case's Own Run

Case and configuration	End (yr)	R_{low}	i_{max}	y_{min}
1, weirs only, as loaded	floor 1.666	5.00	5.00	0.00
1, desk sells 6	floor 1.608	5.00	5.00	−0.04
1, unsterilized	held	10.37	8.09	−1.77
1, policy rate 7 %	held	10.54	7.00	−1.27
2, currency board	held	6.22	10.85	−2.47
3, PI desk, as loaded	attack 2.252	12.14	5.00	0.00
3, policy rate 8 % by hand	attack 2.252	15.00	8.00	−0.95
3, sterilized	attack 1.321	13.14	5.00	0.00
4, sterilized, as loaded	floor 0.501	4.97	5.00	0.00
4, unsterilized	held	6.12	10.92	−2.44
5, cap 15 %, as loaded	tolerance 0.274	12.26	15.00	−0.02
5, cap 12 %	floor 0.337	4.70	12.00	−0.54
5, controls on	held	11.44	11.75	−1.16
6, policy rate 5 %, as loaded	attack 1.589	8.99	5.00	0.00
6, policy rate 8 %	held	15.00	8.00	−2.06
6, unsterilized	held	11.50	7.33	−1.37
6, rate defense	held	9.66	6.00	−0.40
7, $R(0) = 11$, as loaded	floor 1.441	5.00	5.00	−0.01
7, $R(0) = 12$	held	8.19	5.00	−0.01
7, unsterilized	held	8.05	6.97	−0.85
8, MPC, as loaded	float 0.389	10.32	11.17	−0.54
8, $C_{exit} = 5$	float 0.351	11.50	10.83	−0.27
8, $C_{exit} = 40$	float 0.408	9.73	11.17	−0.67
8, $\tau_{att} = 1$ yr	float 0.293	13.56	11.96	−0.11
8, 5-point attack	held	10.96	7.77	−0.67
9, controls on, as loaded	held	11.20	5.00	−0.48
9, override	held	12.25	5.00	−0.27
9, controls off	floor 0.501	4.97	5.00	0.00
R, reference	held	15.00	5.00	0.00

R_{low} in % of GDP, rates and gap in %, on the true states, as the scorecard reads them. R_{low} and i_{max} include the day that ends a peg, at an attack before the speculators' take. The five-point row replaces the scripted attack, and the reference case reproduces it with the MPC, sterilization off, print noise 0.3 and the small attack fired at 0.25 yr.

reading. The currency then falls 3.76 points, the devaluation the market had been pricing. Selling 6 % a year from the desk from the start holds the float just strong of parity, −0.19 % on average from year 0.2 to 1.2, and brings the floor forward to year 1.608, because a stronger rate costs a little trade balance. The same sales requested when the float first reaches the weak weir, at year 0.096, leave it there, and the floor comes at 1.630. Unsterilized, the same outflow leaves the vault at 10.37 % after two years with the overnight rate at 8.09 % and a gap of −1.77 %, the drain law's balance on its way, and a policy rate of 7 % holds it at 10.54 %.

TRY IT IN THE APP **1 · Sell Reserves by Hand and Read the Clock**

Load the scenario and press Run without touching the desk, so the weak weir does all the selling. Watch the clock under the vault: a year in it reads 9.8 months, yet the vault reaches the floor at year 1.666, eight months later, because the outflow grows once the vault is below the threshold.

Pause. Try Exercise 14.1 before reading on.

14.6.2 A Board, a Sterilizing Desk and the Trilemma's Center

In case 2 speculators bet on a devaluation of 6.67 % within three months and move ten points of GDP out, and the currency board meets them with its weirs alone. The overnight rate rises by itself to 10.85 % at year 0.504, the end of the attack, with the vault at 6.22 %, 2.5 months of imports, and below the threshold from year 0.449. The high rate draws capital back, the strong weir starts to buy, and by year 1.5 the vault is 11.78 % with the rate at 7.15 %. The bill is output, a gap of −2.47 % at year 1.10, long after the attack, and the speculators lose 0.62 % of GDP, the carry they paid for a devaluation that did not come. The Hong Kong band gives nearly the same run, 6.19 % and 10.87 %.

Case 4 meets the same attack with the PI desk. Sterilized, the overnight rate stays at 5 %, the outflow keeps running at 40 % of GDP a year, and the vault reaches the floor at year 0.501, the last day of the attack, after which the currency falls 10.46 points. The sterilized sales have also driven the shadow rate from −20 % to 0.07 %. Unsterilized, the overnight rate climbs to 10.92 % at year 0.507, the vault bottoms at 6.12 % and recovers to 11.89 % by year 1.5, and the gap reaches −2.44 % at year 1.06. At year 0.45 the sterilized vault holds 7.10 % against 7.86 %, and the check finds the sterilized vault lower at every instant before it ends. A desk deciding every day at a gain of 10 keeps the float off the weir and loses the peg one day sooner, at year 0.499. Sterilized intervention keeps a fixed rate, free capital and an independent rate together, the trilemma's empty center, and the vault pays for it until it cannot.

Pause. Try Exercise 14.3 before reading on.

14.6.3 The Date No Controller Moves

In case 3 the bank lends the government 2 % of GDP a year and its desk does not sterilize. At year 2.2 the vault still holds 12.25 % of GDP and loses only 1.5 % a year, and the clock reads 58 months, as shown in Figure 14.3. The shadow rate is then 0.80 % against an edge of 1. At year 2.252 the shadow rate reaches the edge and the speculators take the 7.14 points above the floor in a day. Every controller loses the peg on that day, and only the size of the last attack differs. The board leaves 7.30 points for it, the rate defense 7.72, or 8.53 at a gain of 1, the MPC, which raises the policy rate to 8 %, 10.03, and the hand holding 8 % 11.44, at the cost of a gap of −0.95 %. A rate that defends harder fills the vault with reserves that the attack then takes. Under the PI desk, sterilizing brings the day forward to year 1.321, because sterilized sales are credit too, while the board, which never sterilizes, keeps the attack at 2.252. This is the application's honestly infeasible case, and its check sweeps the hand at two policy rates, the board, two PI desks, two rate defenses and the MPC and finds every one of them ending within a day of T_K.

Pause. Try Exercise 14.2 before reading on.

DID YOU KNOW? **Sterilizing Brings the Attack Forward**

From the simulator: in case 3 the PI desk without sterilization loses the peg to the attack at year 2.252, and so does a hand that holds the policy rate at 8 %, whose vault never falls below 15.00. The same desk sterilizing loses it at year 1.321, because sterilized sales add to domestic credit (Table 14.5).

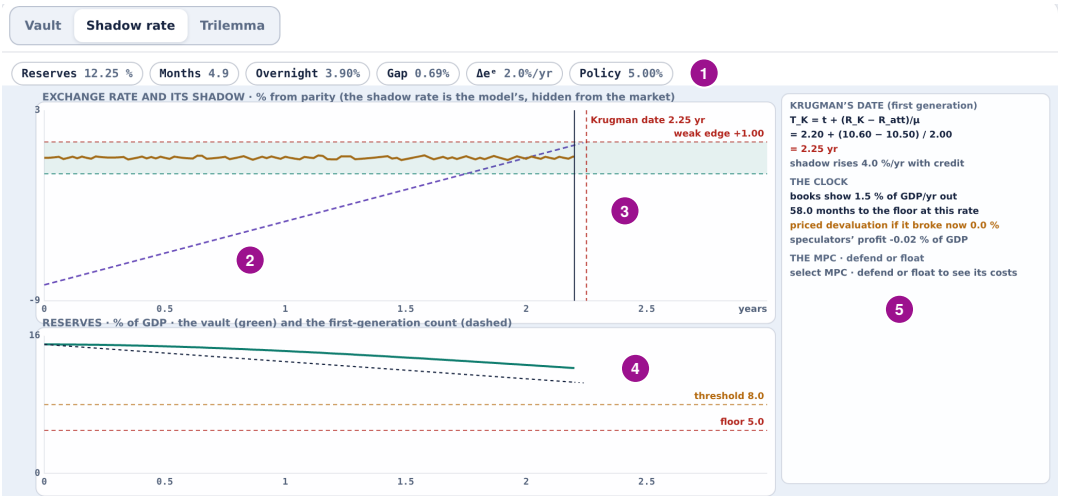

Figure 14.3 Case 3, a monetized deficit of 2 % of GDP a year, at $t = 2.20$ yr in the Shadow rate view, in AUTO as loaded under the PI desk, cropped to the view: the header, the control panel, the trends and the learning panel are cut. (1) The badges read reserves 12.25 %, 4.9 months of imports, an overnight rate of 3.90 % below the 5.00 % policy rate, a gap of 0.69 % and expected depreciation of 2.0 % a year. (2) The upper plot shows the exchange rate near parity and the hidden shadow rate (dashed) rising in a straight line from −8 % toward the weak edge at +1.00, and (3) Krugman's date is marked at 2.25 years. (4) The lower plot shows the vault against its dashed first-generation count. (5) The side panel computes $T_K = 2.20 + (10.60 - 10.50)/2.00 = 2.25$ years, says the shadow rate rises 4.0 % a year with credit, and gives the clock as 1.5 points of GDP a year out and 58.0 months to the floor. The attack comes 19 days later.

14.6.4 A Capped Rate and an Imported Tightening

In case 5 speculators bet on a 10 % devaluation within six weeks and move capital out at 130 % of GDP a year from year 0.25. The rate defense first acts at the weekly decision of year 0.271, a week into the attack, and asks for far more than its 15 % cap, which one decision of up to ten points reaches. The next day the overnight rate is past the government's 12 % and the peg ends with 12.26 % of GDP still in the vault, and the currency falls 10 %, the devaluation the speculators bet on. With the cap at 12 % the rate sits at the cap inside the tolerance, and seven points draw back only 10.5 % of GDP a year against 130, so the vault reaches the floor at year 0.337 and the currency falls 14.12 points. With controls on, the peg holds with the vault at 11.44 % at its lowest and a gap of −1.16 %. The scorecard counts the day that ends the first run, so its highest overnight rate reads the 15 % that broke the peg.

In case 6 the anchor raises its rate from 3 to 6 % at year 0.25 for reasons of its own. With its own rate at 5 % and sterilized sales, the bank loses 4.5 % of GDP a year, from 13.89 % at year 0.5 to 9.39 % at 1.5, and the sterilized sales carry the shadow rate to the band, where speculators take the last 3.99 points at year 1.589. Importing the tightening keeps the peg. At a policy rate of 8 % the flows balance, the vault rises to 17.08 % and the gap is −2.06 % after two years. Unsterilized at 5 %, the money market imports it for the bank, with the overnight rate drifting to 7.33 % and the gap to −1.37 %. The rate defense stops at 6 %, with a gap of −0.40 % and a vault still draining at 9.66 %.

14.6.5 Contagion and the Basin of a Sound Peg

Case 7 has no deficit and no overvaluation. Its vault starts at 11 % of GDP, and a crisis next door adds 5 points to the premium from year 0.25 to 0.75, which alone moves $\kappa \cdot 5 \cdot 0.5 = 3.75$ points of GDP out. The vault crosses the threshold at year 0.655 and holds 7.22 % when the contagion ends,

Table 14.6 The Challenge in Eighteen Configurations, as the Application Scores It

Configuration	Months	i_{max}	y_{min}	Controls	End
PI unsterilized, shipped	2.45	10.92	−2.44	0	held
PI sterilized	1.99	5.00	0.00	0	floor 0.751
policy rate 6 %	2.81	11.31	−2.59	0	held
policy rate 7 %	3.16	11.73	−2.82	0	held
policy rate 8 %	3.58	12.03	−1.81	0	tolerance 0.745
policy 10 %, sterilized	4.43	10.00	−3.00	0	tolerance 0.967
currency board	2.49	10.85	−2.47	0	held
rate defense	5.39	12.01	−0.14	0	tolerance 0.545
rate, $k_r = 0.5$, cap 10 %	4.77	12.05	−0.55	0	tolerance 0.597
MPC	4.13	11.17	−0.54	0	float 0.638
MPC sterilized	4.18	11.75	−0.66	0	float 0.638
MPC, $C_{exit} = 40$	3.89	11.17	−0.67	0	float 0.658
controls on throughout	4.79	7.02	−1.63	1.501	held
override, suggested	4.97	6.72	−0.88	0.384	**met**
override, sterilized	4.90	5.00	−0.27	0.386	**met**
override, MPC	4.95	5.25	−0.15	0.378	**met**
override, rate defense	5.10	8.48	−1.09	0.375	**met**
override, board	5.02	6.63	−0.88	0.378	**met**

Limits 3 months, 10 %, −3 %, 0.5 yr of controls. Lowest months of imports, rates and gap in %, controls in years, end in years. Each configuration is chosen before the run starts. The scorer counts the day that ends a peg.

and from there the drain law carries on alone, to 6.69 % at year 1.0, 5.89 % at 1.25 and the floor at 1.441, eight months after the cause has gone. A vault that starts at 12 % survives with 8.19 % at its lowest, one that starts at 11.75 dips to 7.87 %, just below the threshold, and keeps sliding, and one at 11.5 is at 5.25 % after two years. The basin's edge is therefore the threshold plus what the contagion drains, near 11.8 %. At 11 %, letting the money market tighten holds the peg at 8.05 % with a gap of −0.85 %, and a policy rate of 7 % holds it at 9.64 %.

14.6.6 Defend or Float

Case 8 repeats the ten-point attack, unsterilized, with the MPC deciding every week. It raises the policy rate to 8.5 %, and the drain takes the overnight rate to 11.17 %. Each week that the outflow keeps coming, the projected bill of defending grows, from 7.01 against 12.98 for floating at the decision of year 0.290 to 8.53 against 12.92 at 0.348 and 11.19 against 12.88 at 0.367, as shown in Figure 14.2(b). At the decision of year 0.386 defending would cost 16.35 and floating 12.86, and the peg ends by choice at year 0.389, seven weeks into the attack, with 10.32 % of GDP in the vault and a fall of 6.67 points, the devaluation priced. A cost of floating of 5 moves the date to year 0.351, and one of 40 defends one week longer and floats at 0.408, when no plan can avoid the floor or the tolerance and the cost of defending passes 2000. Expecting the outflow to last a year, $\tau_{att} = 1$, it floats two weeks in, at year 0.293. Against the five-point attack it defends throughout, and over attacks of 2 to 16 points from month 6 in the challenge's economy it defends 2, 4 and 6 and floats at years 0.715, 0.638, 0.581, 0.562 and 0.542 for 8, 10, 12, 14 and 16, earlier the larger the attack, as the application's check requires, and sterilizing leaves the date unchanged at year 0.389, with a policy rate that reaches 11.75 %.

14.6.7 The Third Instrument

In case 9 the ten-point attack meets the sterilizing desk with controls on, so only a fifth of it moves. Kept on, the controls hold the peg with the vault at 12.40 % at year 0.5, but they keep costing after the attack is over, and the vault slides to 11.20 % by year 1.5 with a gap of −0.48 %. The same controls in the reference case, with no attack, cost 1.8 points of reserves in 1.5 years and the same gap. With the switch on the override, the controls start lifted, come on at the execution of year 0.260, three days into the attack, and lift at 0.644, which holds the peg with the vault at 12.25 % at its lowest and a gap

of −0.27 %, controls in force for 0.384 year, while with controls off the vault reaches the floor at year 0.501.

> **TRY IT IN THE APP 9 · Capital Controls as an Override, With the Bill**
>
> Run the scenario on the Trilemma view with controls on and read the lowest reserves, 11.20 % of GDP. Set Capital controls to Override, press Reset and run again. The controls come on three days into the attack and lift at year 0.644, and the vault's lowest is 12.25 %.

The challenge fires the ten-point attack at month 6 against the PI desk unsterilized at a policy rate of 5 %, weekly. Throughout 1.5 years the vault must stay at or above three months of imports, 7.5 % of GDP, the overnight rate at or below 10 % and the gap at or above −3 %, with the peg standing and any controls in force for at most six months. The limit on the overnight rate is judged rather than on the policy rate, because output pays the overnight rate. Table 14.6 compares 18 configurations. As shipped, the vault falls below three months at year 0.712, eleven weeks into the attack, and bottoms at 6.12 %, 2.45 months, with the overnight rate at 10.92 %. No setting without controls passes. A higher policy rate trades the vault against the overnight rate, which passes 10 % from a policy rate of 6 % on and breaks the tolerance at 8 %. The rate defense and the MPC end the peg, sterilizing loses it at the floor, and a sterilized policy rate of 10 % holds 4.43 months until the gap reaches −3 % at year 0.967. Three limits and two handles leave no setting inside all three, the count of instruments against targets that Volume 6, Chapter 2 meets with two actuators and more variables worth holding.

The suggested override passes. It engages at the execution of year 0.510, three days into the attack, lifts at 0.893, and keeps the vault at 4.97 months at its lowest, the overnight rate at 6.72 % and the gap at −0.88 %, with controls in force 0.384 year (4.6 months). The override also rescues the sterilized desk, the MPC, the rate defense and the board, and Figure 14.4 draws every configuration against the three limits. Left on for the whole run, the controls hold the peg but cost 1.5 years of controls and a gap of −1.63 %. Over six seeds, the default and 1 to 5, the shipped run is identical to six decimals, because the seed reaches only the GDP prints and the desk's PI never reads them, and the suggested override passes on all six, with seed 5 alone lifting the controls a day earlier.

The override's speed is part of that verdict. It runs every day, outside the weekly calendar that binds the desk and the rate. In a counterfactual copy of the application's controller with the override checked only at the weekly decisions, it engages at year 0.521 and, because the weekly readings of the pressure at years 0.923 and 0.943 land above 2, it lifts only at 1.041, so controls stay 0.521 year and the suggested run fails the six-month limit. The same copy with a release level of 3 instead of 2 lifts at 0.904 and passes, with 4.84 months and controls for 0.384 year. A protective layer that acts faster than the policy it protects is how an override is normally built (Section 1.4.3), and the challenge's statement, the briefings and the Model dialog say that the override is checked every day, outside the calendar.

14.6.8 What the Calendar Changes

Deciding every day instead of every week moves the float more than the vault. The daily desk keeps the float off the weirs and changes the end of case 4 by one day. The MPC deciding every day floats at year 0.381 instead of 0.389, with 10.67 % of GDP in the vault, and in the challenge at 0.630 instead of 0.638. The rate defense deciding every day moves to its cap one day into the attack and breaks the tolerance by the end of that day, at year 0.258 instead of 0.274, while at the 12 % cap it reaches the floor on the same day, 0.337, on either calendar. A daily PI at a gain of 10 still fails the challenge with 2.42 months. The lag that matters here is not the calendar's but the vault's, which a week of decisions neither adds nor removes.

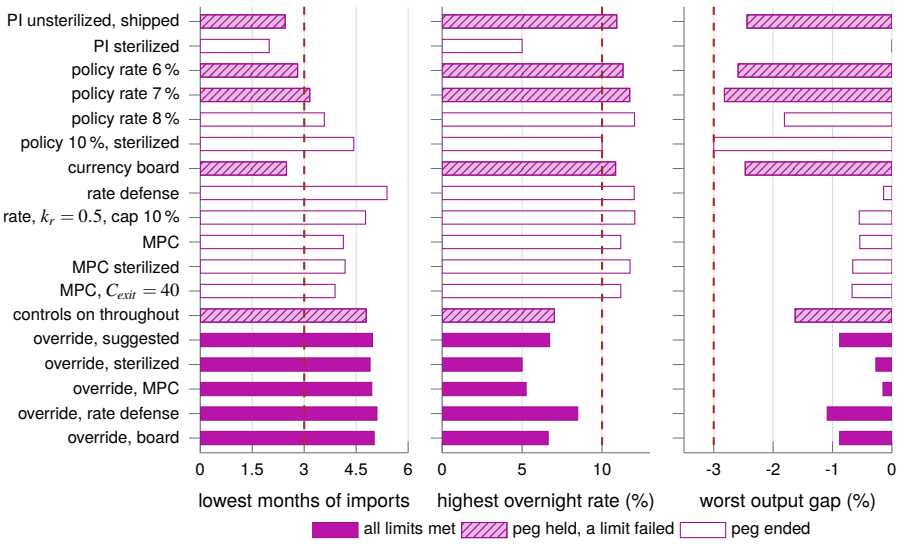

Figure 14.4 The challenge's three measures for the 18 configurations of Table 14.6, which also prints the values and the time in controls. Dashed lines are the limits: at least three months of imports, an overnight rate of at most 10 % and a gap of at least −3 %. Dark bars met every limit, hatched bars held the peg but failed a limit, and white bars lost the peg. The hatched bar of controls on throughout failed the limit on time in controls. Notice that every configuration without controls fails at least one limit, while each of the five with the override, which acts every day, clears all three.

14.7 Technology in Practice

A real central bank reads the same stocks and moves the same handles as the bank of Section 14.5, but its instruments are its books, the market's prices and its dealing room. Figure 14.5 draws the desk, the rate and the override with those instruments in place of the schematic's symbols.

What Measures It

The controlled variables are the exchange rate and the vault. Hong Kong's monetary authority publishes the details of its monetary base every day, and since March 1999 a monthly account that sets the backing assets beside the monetary base [19]. By a 2013 review, most central banks of emerging market economies had increased their market monitoring, particularly of cross-border portfolio flows [20]. The model's bank reads its reserves and its credit every day from its own books and the rate as a market price, and it sees output only through a late quarterly print, as Section 14.3 sets out.

> **TECHNOLOGY Sensor: Reserves on the Books and the Rate on the Screen**
>
> A central bank reads its reserves from its own books and the exchange rate from the market. Since Hong Kong's real-time gross settlement began in December 1996, interbank transactions settle across the books of its monetary authority, which publishes the monetary base every day and, in real time, the forecast change in the banks' aggregate balance caused by its foreign exchange transactions [19, 21]. Dealers post prices through brokers and electronic platforms, and some central banks survey dealers' open positions every day [22]. Readings can mislead: by June 1997 the Thai authorities reportedly had, beyond the reserves already spent, $23 billion in forward swaps outstanding, and the available reserves had fallen far below the short-term debts owed to international banks [12]. The model's bank reads R exactly every day, and its vault carries no forward commitments.

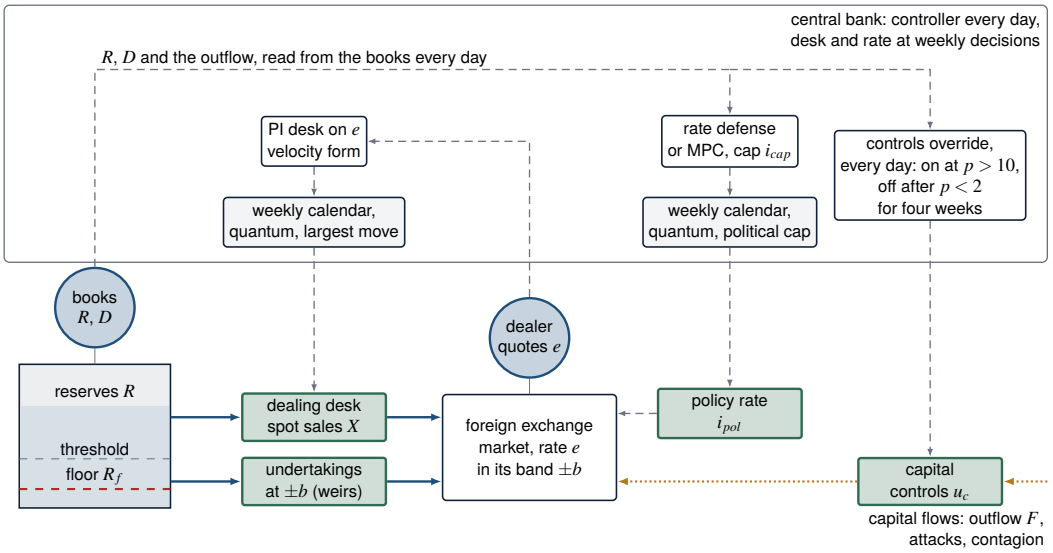

Figure 14.5 The desk, the rate and the override of Section 14.5 with a central bank's instruments. Green boxes are the final elements: spot sales from the dealing desk, the undertakings at the band edges, the policy rate and capital controls. Blue circles are the measurements: the bank's own books and the dealers' quotes. Solid arrows carry reserves, dashed arrows signals, and the dotted arrows the private capital flows that the controls slow. Notice that the desk and the rate wait for the weekly calendar, while the override and the undertakings act without it.

What Moves It

Beside the desk's spot sales, the interest rate is the second handle, and its cost lands on borrowers. Where mortgage payments are indexed to money market rates, a defense that holds those rates high for long periods can raise mortgage rates dramatically, and authorities that doubt a quick success may hesitate to use it [11]. Capital controls took a variety of forms, from taxes on holdings of foreign currency assets to restrictions on the ability of banks to lend abroad [11]. The model folds them into one switch that cuts capital mobility and the persistent outflow by c_K, and its political cap i_{cap} stands for the hesitation over rates.

> **TECHNOLOGY Final Element: Undertakings at the Edges and a Dealing Desk**
>
> Hong Kong's monetary authority undertakes to sell its currency to banks at 7.75 to the US dollar and to buy it at 7.85. At the weak side it buys, the banks' aggregate balance contracts and interest rates rise, an automatic adjustment that pushes the rate back inside the zone [23]. In a survey of emerging market central banks, most intervention was a spot transaction in the onshore wholesale market, mostly in normal business hours, when the market is deepest. It is dealt directly with market-makers, through voice or electronic brokers, on trading platforms or by auction [22]. In the model the undertakings are the weirs at $\pm b$, which act after every step, and the dealing desk is the desk X, which moves only at a decision.

Where the Controller Runs

In a survey of central banks, most intervene mainly in their own time zone and during normal business hours. A currency board is the exception: its intervention is passive and automatic, and Hong Kong supports its arrangement on a 24-hour basis through offices in other financial centers [22]. The application's controller executes every day. Its desk and rate move only at the decisions of the weekly calendar of `core/cadence.js`, while the weirs act after every integration step and the

capital-controls override runs at every daily execution, outside the calendar. That fast protective layer under a slower policy is the override structure of Section 1.4.3.

> **IN PRACTICE** **Reserve Adequacy by Three Ratios**
>
> Central banks judge a vault by ratios. The main indicator is reserves against short-term external debt, the debt maturing within one year. A country is typically considered prudent if its reserves cover that debt, a test known as the Greenspan–Guidotti rule. Reserves against imports give the months a country could go on importing if all other capital flows stopped, and a rule of thumb asks for three to four months. Reserves against broad money gauge how much capital residents could take out [24]. The application's safe vault is three months of imports, the challenge's limit, and its floor is two.

What Is Modernizing It

Three developments bear on this loop. The first is rules. Some central banks sell option contracts under well publicized auction rules, which transparently adds an element of automatic stabilization and gives the private sector a more complete set of hedging instruments [22]. The second is a wider set of handles. Under heavy pressure in 2013, several emerging market economies sold reserves, raised interest rates and provided the private sector with insurance against exchange rate risk, and targeted macroprudential and capital flow management measures are discussed as a further way to reduce their vulnerability [20]. The third is disclosure: a real-time forecast of the change in the banks' balances lets the market anticipate liquidity and adjust interbank rates more efficiently [19].

14.8 Limitations

The economy is a teaching model. Its expectations are mechanical and look backward, apart from the shadow rate's trigger, so the objection to evaluating policy with relations fitted under another policy [17] applies in full, and the multiple equilibria of self-fulfilling crises [5] appear only as the threshold of a fixed expectation term, not as a game among speculators. Capital flows on a level difference without portfolio stocks, one mobility serves all capital, money is one stock with a linear demand, and there are no banks, no foreign-currency debt and no float after the exit, only a rule. The first-generation date matches the linear textbook model only in its timing, as Section 14.4 sets out. Every coefficient, threshold, tolerance and attack size is a teaching value, and none of them describes a particular country.

> **OPEN CHALLENGE** **How Large a Vault Is Enough?**
>
> No single ratio settles the question: the most popular is not the only one, and their values depend on the database used [24]. Some central bankers note that the level of reserves perceived as adequate can vary considerably over time, particularly in crisis periods, and others that the financial costs of intervention can be relatively large as a percentage of GDP in emerging market economies [20]. The model fixes its threshold R_{thr}, where expectations turn, as a teaching value. A threshold read from the market as it moves, rather than declared in advance, would give the override a better limit.

The bank's estimate uses a fixed teaching gain, its reading of the unexplained outflow is smoothed over a week, and the MPC's model sends every flow straight to the vault and foresees no attack it has not yet seen, while its cost of floating is a declared number on which its date depends, from year 0.351 to 0.408 between costs of 5 and 40. With the economy's capital mobility doubled to 3 in case 8, the MPC still plans with 1.5, and the government's tolerance ends the peg at year 0.293. The override acts every day, outside the weekly calendar, and the suggested pass depends on that.

OPEN CHALLENGE Telling a Passing Attack From a Lasting Shift

The override of this chapter acts on the outflow that the bank cannot explain with its own model. Central banks find it difficult to know in advance why international investors move, or to identify the short-term component of capital flows. The effect of an intervention is hard to measure too, because nobody sees the exchange rate that would have prevailed without it. Senior central bankers meeting in 2013 reached no consensus on how effective intervention is [20]. An estimator that separated a passing attack from a lasting shift, and said how sure it was, would be progress.

14.9 Conclusion

This chapter wrote a currency peg as a band held by two weirs, a vault with a floor that pays every flow, a money market that tightens when sales go unsterilized, and an expectation term that turns the plant's feedback positive below a threshold. Its reserve clock reads 9.8 months a year into a steady outflow and the peg lasts eight, a currency board's overnight rate rises by itself to 10.85 % and holds the vault through a ten-point attack, and the same attack empties a sterilizing desk's vault on its last day. Without sterilization a monetized deficit sets the attack at year 2.252 against a derived 2.25 under every controller, which decide only how much of the vault the speculators take, from 7.14 to 11.44 points of GDP. A rate defense capped at 15 % ends the peg the day after it acts, and an MPC floats seven weeks into an attack with 10.32 % of GDP still in the vault. In the challenge twelve configurations without controls fail and a capital-controls override passes with 4.97 months of imports, an override that acts every day. The expectations are mechanical and every coefficient is a teaching value. A fuller model would add a market that learns the bank's rule, portfolio stocks and banks with foreign-currency debt, and a float after the exit. The vault makes the arithmetic of a peg, its date and its missing instrument something a student measures rather than accepts.

KEY POINTS

- A bank that fixes its exchange rate pays for every unmatched flow from a vault with a floor, and the vault drains faster once the market expects a devaluation.
- A year into a steady outflow the reserve clock reads 9.8 months, and the peg lasts eight.
- Without sterilization a monetized deficit sets the attack at year 2.252 against a derived 2.25 under every controller, which decide only how much of the vault the speculators take.
- In the challenge twelve configurations without controls fail, and a capital-controls override that acts every day, outside the weekly calendar, passes with 4.97 months of imports.

Where else this applies. Chapter 8 defends a cash floor with an override that gives way in a fixed order, and Volume 6, Chapter 2 meets the same count of two actuators against more variables worth holding.

References

[1] S. W. Salant and D. W. Henderson, "Market anticipations of government policies and the price of gold," *Journal of Political Economy*, vol. 86, no. 4, pp. 627–648, 1978.

[2] P. Krugman, "A model of balance-of-payments crises," *Journal of Money, Credit and Banking*, vol. 11, no. 3, pp. 311–325, 1979.

[3] R. P. Flood and P. M. Garber, "Collapsing exchange-rate regimes: Some linear examples," *Journal of International Economics*, vol. 17, no. 1-2, pp. 1–13, 1984.

[4] P. R. Krugman, "Target zones and exchange rate dynamics," *The Quarterly Journal of Economics*, vol. 106, no. 3, pp. 669–682, 1991.

[5] M. Obstfeld, "Models of currency crises with self-fulfilling features," *European Economic Review*, vol. 40, no. 3-5, pp. 1037–1047, 1996.

[6] A. Lahiri and C. A. Végh, "Delaying the inevitable: Interest rate defense and balance of payments crises," *Journal of Political Economy*, vol. 111, no. 2, pp. 404–424, 2003.

[7] M. Obstfeld, J. C. Shambaugh, and A. M. Taylor, "The trilemma in history: Tradeoffs among exchange rates, monetary policies, and capital mobility," *Review of Economics and Statistics*, vol. 87, no. 3, pp. 423–438, 2005.

[8] R. A. Mundell, "Capital mobility and stabilization policy under fixed and flexible exchange rates," *Canadian Journal of Economics and Political Science*, vol. 29, no. 4, pp. 475–485, 1963.

[9] P.-R. Agénor, J. Bhandari, and R. Flood, "Speculative attacks and models of balance-of-payments crises," National Bureau of Economic Research, Cambridge, MA, NBER Working Paper w3919, 1991.

[10] R. P. Flood, P. M. Garber, and C. Kramer, "Collapsing exchange rate regimes: Another linear example," *Journal of International Economics*, vol. 41, no. 3-4, pp. 223–234, 1996.

[11] B. Eichengreen and C. Wyplosz, "The unstable EMS," *Brookings Papers on Economic Activity*, vol. 1993, no. 1, pp. 51–144, 1993.

[12] S. Radelet and J. D. Sachs, "The East Asian financial crisis: Diagnosis, remedies, prospects," *Brookings Papers on Economic Activity*, vol. 1998, no. 1, pp. 1–90, 1998.

[13] H. Genberg and C.-H. Hui, "The credibility of Hong Kong's link from the perspective of modern financial theory," *Journal of Money, Credit and Banking*, vol. 43, no. 1, pp. 185–206, 2011.

[14] U. J. Jermann, "Financial markets' views about the Euro–Swiss franc floor," *Journal of Money, Credit and Banking*, vol. 49, no. 2-3, pp. 553–565, 2017.

[15] H. Rey, "Dilemma not trilemma: The global financial cycle and monetary policy independence," National Bureau of Economic Research, Cambridge, MA, NBER Working Paper w21162, 2015.

[16] Bank for International Settlements, "Triennial survey shows global foreign exchange trading averaged $7.5 trillion a day in april 2022, OTC interest rate derivatives $5.2 trillion," Press release, 27 October 2022, 2022. [Online]. Available: https://www.bis.org/press/p221027.htm

[17] R. E. Lucas, "Econometric policy evaluation: A critique," *Carnegie-Rochester Conference Series on Public Policy*, vol. 1, pp. 19–46, 1976.

[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] Hong Kong Monetary Authority, "Transparency of currency board operations," Web page, last revised 26 August 2019, 2019, accessed 2026-11-21. [Online]. Available: https://www.hkma.gov.hk/eng/key-functions/money/linked-exchange-rate-system/transparency-of-currency-board-operations/

[20] M. S. Mohanty, "Market volatility and foreign exchange intervention in EMEs: what has changed? An overview," in *Market Volatility and Foreign Exchange Intervention in EMEs: What Has Changed?*, ser. BIS Papers. Basel: Bank for International Settlements, 2013, no. 73, pp. 1–10. [Online]. Available: https://www.bis.org/publ/bppdf/bispap73.htm

[21] Hong Kong Monetary Authority, "Milestones of monetary reform," Web page, last revised 13 November 2023, 2023, accessed 2026-11-21. [Online]. Available: https://www.hkma.gov.hk/eng/key-functions/money/linked-exchange-rate-system/milestones-of-monetary-reform/

[22] D. Archer, "Foreign exchange market intervention: methods and tactics," in *Foreign Exchange Market Intervention in Emerging Markets: Motives, Techniques and Implications*, ser. BIS Papers. Basel: Bank for International Settlements, 2005, no. 24, pp. 40–55. [Online]. Available: https://www.bis.org/publ/bppdf/bispap24.htm

[23] Hong Kong Monetary Authority, "How does the LERS work?" Web page, last revised 22 October 2025, 2025, accessed 2026-11-21. [Online]. Available: https://www.hkma.gov.hk/eng/key-functions/money/linked-exchange-rate-system/how-does-the-lers-work/

[24] International Relations Committee Task Force, "The accumulation of foreign reserves," European Central Bank, Frankfurt am Main, ECB Occasional Paper 43, 2006.

14.10 Exercises

14.1 ○ Concept check · 2 min

In case 1 capital leaves at 6 % of GDP a year and the bank sterilizes. A year in, the reserves above the floor divided by the outflow on the books give a clock of 9.8 months, yet the peg ends eight months later. Why is the clock optimistic?

(a) It divides by the outflow of the first month, which is smaller than the outflow a year later

(b) Below the threshold each point of reserves lost raises expected depreciation, so the outflow grows

(c) The weak weir sells at a worse rate than the desk would, so each point of outflow costs more

(d) The weekly calendar delays the sales, and the weirs then sell the backlog at the edge of the band

14.2 Concept check · 1 min

In case 3 a monetized deficit of 2 % of GDP a year meets an unsterilizing bank. The hand at a policy rate of 8 %, the currency board, the rate defense and the MPC all lose the peg at year 2.252. Why does none of them move the date?

(a) All of them hold the vault at the same level, so it reaches the floor on the same day

(b) None of them raises the policy rate early enough to draw capital back before the attack

(c) The date is set by the shadow rate, which follows domestic credit, and none of them changes credit

(d) The government's tolerance ends every defense at the same overnight rate of 12 %

14.3 Concept check · 2 min

In case 4 the PI desk meets the ten-point attack. Sterilized, the vault reaches the floor at year 0.501. Unsterilized, the same desk keeps the peg with the vault at 6.12 % at its lowest. What makes the difference?

(a) Unsterilized, the desk holds the float off the weak weir, so the weirs sell nothing

(b) Unsterilized, the speculators see the rate rise and bet a smaller devaluation

(c) Unsterilized sales add to domestic credit, which pushes the shadow rate away from the band

(d) Each point sold takes a point of money out, raises the overnight rate and draws capital back

14.4 Theory and calculation · 6 min

The first-generation date of Equation (14.10) uses $\zeta = 100/\bar{H} = 2$, $\alpha = 1.5$, $i^* = 3\,\%$, $\delta = 2$, $i_n = 5\,\%$, $R_f = 5\,\%$ and $b = 1\,\%$.

(a) For case 3, $R(0) = 15$ and $\mu = 2$ with $D(0) = \bar{H} - R(0)$, find the shadow rate at the start from Equation (14.6), its rate of rise and T_K. *Checkpoint:* $-8\,\%$, 4 % a year and $T_K = 2.25\,\text{yr}$.

(b) Find T_K for $\mu = 1$ from $R(0) = 15$, and for $\mu = 2$ from $R(0) = 11$. *Checkpoint:* 7.5 yr and 0.25 yr.

(c) At year 2.2 of case 3 the first-generation count $R(0) - \mu t$ gives 10.6 %, the vault holds 12.25 % and the overnight rate is 3.90 % under a policy rate of 5 %. Use Equation (14.3) to account for the difference.

14.5 Theory and calculation · 6 min

Use the drain law of Equation (14.9) with $\kappa = 1.5$, $\theta = 10$, $R_{thr} = 8$, $\alpha = 1.5$ and $R(0) = 15$.

(a) Sterilized with no outflow, find the slope of the law below the threshold and the time in which a deviation below the threshold doubles. *Checkpoint:* 1.875 points of GDP a year per point, 0.37 yr.

(b) Unsterilized, with a persistent outflow $F = 6$, find the reserves at which the law balances and the overnight rate there, with the policy rate at neutral. *Checkpoint:* 9 % of GDP and 9 %, 4 points above neutral.

(c) In case 1 (sterilized, $F = 6$ from month 1) the float first reaches the weak weir at year 0.096, and with the float on the weir the trade balance returns $\eta b = 0.2$ point of GDP a year. Estimate when the vault crosses the threshold. *Checkpoint:* 1.30 yr; the run gives 1.304.

14.6 Theory and calculation · 5 min

An attack j is a bet of $\beta_j = S_j/\kappa_0$ percent within d_j years, with $\kappa_0 = \kappa = 1.5$. It enters Δe^e as β_j/d_j.

(a) Find the capital flow of the ten-point attack ($S = 10$, $d = 1/4$) and of the fifteen-point attack ($S = 15$, $d = 6/52$). *Checkpoint:* 40 % and 130 % of GDP a year.

(b) With the vault above the threshold, find the overnight rate at which the fifteen-point attack draws no net capital out. *Checkpoint:* 91.7 %.

(c) In case 5 at a 12 % cap the rate sits 7 points above neutral. Find the inflow those 7 points draw against the fifteen-point attack. Then find the time the ten-point attack takes to carry the float across the two-point band at $\psi = 20$ with nothing sold or drawn against it. *Checkpoint:* 10.5 % of GDP a year, and 0.9 day.

14.7 App challenge · 8 min

Load scenario 1, Sell reserves by hand and read the clock, at the default seed. Leave the Intervention request at zero, so the weirs do all the selling, and press Run. The run pauses when the peg ends, about 1.7 yr in, which takes about seven seconds at the default speed of 0.25 yr a second. Export the CSV file.

(a) In the row at year 1.000, read the reserves `R` and the outflow on the books `outflow_read`, and compute the clock in months. *Checkpoint:* 9.75 %, 5.80 and 9.8 months (±0.1).

(b) Read the Peg tile. How many months did the peg last after year 1, and how optimistic was the clock? *Checkpoint:* Abandoned at 1.67 yr: 8.0 months, so 1.8 months optimistic.

(c) Find in the file the year at which `R` first falls below 8 %, and the `outflow_read` in the rows at years 1.5 and 1.6. Explain the growth with Equation (14.4). *Checkpoint:* Year 1.304, outflow 7.94 and 9.84 (±0.05).

14.8 App challenge · 9 min

Load scenario 3, Krugman: a monetized deficit and its date, at the default seed, on the Shadow rate view. The PI desk does not sterilize. Each run lasts up to 3 yr and pauses at the attack, about ten seconds at the default speed of 0.25 yr a second.

(a) Press Run and read the Peg tile, the Lowest reserves tile and Krugman's date on the Shadow rate view. How many points of GDP does the attack take? *Checkpoint:* Abandoned at 2.25 yr, lowest reserves 12.14 %, Krugman's date 2.25 yr, so 7.14 points.

(b) Set Controller to Currency board, press Reset and run again. Read the same tiles. *Checkpoint:* Abandoned at 2.25 yr, lowest reserves 12.30 %, so 7.30 points.

(c) Press the reload button, set Sterilization to On and run. Read the Peg tile and explain why the date moves now. *Checkpoint:* Abandoned at 1.32 yr.

14.9 App challenge · 9 min

Load the challenge, Survive the attack, at the default seed. The ten-point attack starts at year 0.5 and the run lasts 1.5 yr, six seconds at the default speed of 0.25 yr a second. Imports are 30 % of GDP a year, so a month of imports is 2.5 % of GDP. After each run, export the CSV file.

(a) Run the challenge as loaded. Read Lowest reserves, Highest overnight rate, Worst output gap and the Challenge tile, and convert the lowest reserves to months of imports. *Checkpoint:* 6.12 %, 2.45 months; 10.92 %; −2.44 %; not met.

(b) Set Capital controls to On, press Reset and run. Read the same tiles, and count the rows of the `controls` column that read 1. *Checkpoint:* 11.98 %, 4.79 months; 7.02 %; −1.63 %; 549 rows, 1.5 yr; not met.

(c) Set Capital controls to Override, press Reset and run. Read the same tiles and the time the controls were in force. *Checkpoint:* 12.41 %, 4.96 months (±0.02); 6.72 %; −0.88 %; 140 rows, 0.384 yr; met.

14.10 Application area · 3 min

Central banks often judge a vault in months of imports.

(a) With imports of 30 % of GDP a year, express the starting vault of 15 %, the threshold of 8 %, the floor of 5 % and case 7's starting vault of 11 % in months. *Checkpoint:* 6, 3.2, 2 and 4.4 months.

(b) Case 7 starts at 4.4 months, above the common three-month rule, and a six-month contagion still ends the peg. What should a reserve target be set against, in the language of Section 1.4.9?

14.11 Application area · 4 min

Capital controls in this model cut capital mobility and the persistent outflow by 80 %, at a trade cost of 1 point of GDP a year and a premium of 1 point.

(a) In the reference case, with no attack, controls kept on for 1.5 years cost 1.8 points of reserves and a gap of −0.48 %. Name the two channels of that bill in Equations (14.4) and (14.5).

(b) In the challenge the override checks the outflow every day. A copy that checks it only at the weekly decisions keeps controls 0.521 year and fails. Why should a protective override run faster than the policy it protects?

14.12 Review the agent · 5 min

An AI assistant reviews case 5, where the rate defense capped at 15 % ends the peg at year 0.274, and writes: "The rate defense failed because its gain $k_r = 0.25$ is too timid. Raise k_r to 1 and the political cap to 25 %, so the rate can climb high enough to stop the outflow while the vault still holds 12 % of GDP." Find the flaw, and say how you would check it in the app.

CHAPTER 15

Summary and Outlook



Open gaps 1 actuator sizing 2 release of overrides 3 margins from prices 4 usable stock 5 unmeasured limits 6 selectors and optimizers

Notice that the gaps of release and of moving limits meet on the oven and the shield, while the gap of usable stock joins the office, the fan car and the vault. Illustration, not to scale.

IN PLAIN TERMS **What Thirteen Plants Taught About Limits**

The main idea. Every chapter of this book ran something close to a limit it must not cross, from a jet engine to a company. Some designs let a safety rule take over the controls as a limit came near. Others shared the work among several handles, the cheap ones first.

The challenge. A safety rule works only if its handle moves fast and far enough, and only if it lets go when the danger has passed. One that stays on costs output. One that comes late costs the machine.

Why it matters. Testing each safety rule away from the normal case, and pricing what it gives up, separated the designs that worked from those that failed.

Chapter 1 asked two questions of a plant at its limits. When several variables must stay inside a limit and one actuator moves them all, which variable gets the actuator? When several actuators move one variable, which one should move, and for how long? Thirteen chapters answered them on machines, process lines, vehicles, a company and a central bank. This chapter sets the answers side by side, draws a short guide for choosing a structure, looks at where practice is heading, reports how the volume was made with AI agents and what they got wrong, and lists the problems that remain open.

15.1 What the Chapters Showed

Table 1.4 listed what the selectors of each chapter protect, how its actuators are shared and which stock depletes. Tables 15.1 and 15.2 return to that table with the measured result, the first for one main actuator and several limits, the second for more actuators than controlled variables. Each row names the limit or the sharing problem, a structure that met it, one that did not, and the result with the table or section that reports it.

The columns of the two tables make five points.

A protective loop is bounded by its actuator before its tuning. On the compressor of Chapter 2 the controller acted correctly with a slow valve and the machine surged anyway, to −105.4 % of margin. With a valve too small it surged to −129.0 %, and no gain or integral time changed either outcome. The trigeneration plant of Chapter 9 emptied its uninterruptible supply in every one of 54 settings once demand exceeded the turbine and the grid, the fan car of Chapter 12 found no controller among 864 that landed a 12 s hover softly, and the CO_2 cycle of Chapter 10 made 1188 kW of a requested 1270 kW with every layer at its stop (Section 10.6.4). Each chapter said so, as Section 1.4.5

Table 15.1 What Each Selector Protected, Chapters 2 to 8

Chapter	Limit	What worked	What failed	Measured result
2, compressor	Surge line, guarded by one recycle valve against a fast throttle closure.	A PI override resting on its closed clamp, a fast valve, and the throttle signal as a minimum opening through a high select.	Any tuning on a valve too slow to stroke or too small to pass the flow.	Minimum margin −8.6 standard, +3.3 fast, −105.4 slow, −129.0 undersized, in percent (Table 2.5). High select +16.16 against −8.61 (Table 2.6).
3, turbofan	Six caps and two floors on one fuel valve. A slam outruns the spools.	An 8 % acceleration schedule with integrator tracking on every limiter.	No schedule. In the challenge, the schedule or tracking alone.	Challenge passes at 5.86 s, 9.4 % margin, 938 °C. Shipped 8.49 s and 996 °C, schedule alone 1019 °C (Table 3.5).
4, combine	Engine, rotor torque and loss limit one ground speed. Moisture is unseen.	Five demands in a low selector with external reset, a 90 % engine limit, preview, a loss supervisor and trims.	The overrides off at the shipped feed. A 95 % engine limit.	Against the rain 10.77 t, worst pass 1.71 %, 1.45 s overload, against 9.03 t, two plugs and 74.70 s. At 95 %, 31.50 s (Table 4.5).
5, extruder	Torque, die pressure and dryer air on a melt heated by its own shear. A wrong flour certificate.	Overrides as clamps, a density trim restarted at each lot, a throughput supervisor at 420 kg/h.	The supervisor alone. The trim carried from one lot into the next.	Accepted 290.2 against 186.3 kg shipped and 172.5 kg for the supervisor alone. No lot restart: torque 99.2 % (Table 5.8).
6, coater	A 25 % LEL cap that binds in whichever zone holds the drying front.	An override with three handles in order, exhaust then speed or zone heat, under a push to the first binding limit.	The push without the override. The push without the solids feedforward.	3076 m^2/h at 21.9 % LEL passes. Without override 3400 m^2/h at 28.1 % LEL, without feedforward 2075 m^2/h (Table 6.6).
7, drill	Torque, mud ratio or thrust holds the rate, by soil. The pressure limit falls with the cover.	A rate optimizer on the binding limit, a pressure override on the thrust as backstop, and guidance.	Fixed clay settings in sand. More mud flow alone on the exit climb.	Exit climb: breakout at 149.8, and at 154.8 m with more flow. Override 44.2, optimizer 76.6 m/h (Table 7.5). Crossing 0.60 m off in 2.13 h, no return (Table 7.6).
8, startup	A cash floor, and a credit line whose ceiling moves with trailing EBITDA.	A runway override split over three ranges ahead of the growth loops.	The board by hand. Higgins's rule. A cash-flow MPC.	Lowest cash $0.79 against −$2.03 million, insolvent at 1.88 yr (Table 8.5). Challenge 20.7 % a year, $272.5 million. Higgins fills 79.5 % (Table 8.7).

recommends: a saturated protective actuator is a sizing statement. Figure 15.1 sets one like pair of numbers from each row of the two tables side by side.

Anti-windup decides when the override takes over. The turbofan of Chapter 3 passed its challenge only with integrator tracking on every limiter. With the tightened schedule alone the exhaust limiter's integral had wound up while the engine was cold, and the indicated temperature reached 1019 °C against a 950 °C limit. The tunnel shield of Chapter 11 starts each of its five advance limits from the request actually selected, so none winds up while another holds the advance. The CO_2 cycle hands the bypass controller's unsaturated demand upward, because the clamped position reports the same error for a small and a large shortfall (Section 10.5.3). External reset is part of the protection, not a refinement of it (Section 1.4.4).

Table 15.2 How Spare or Limited Actuators Were Shared, Chapters 9 to 14

Chapter	Sharing problem	What worked	What failed	Measured result
9, trigeneration	Cheap absorption cooling answers in minutes, electric chillers in seconds, a store in between.	A fast electric trim and a VPC that hands the duty back to absorption, with a stronger store discharge.	Absorption alone. The store gain alone. Any tuning once capacity is lost.	IAE 951.41 against 5.16 K s, reserve regained at 806 s (Table 9.7). Challenge 0 against 600 s out of band (Table 9.8).
10, CO_2 cycle	Bypass fast and wasteful, firing slow and capped, compressor pressure the only headroom.	Three mid-ranging layers, each handing the unsaturated demand upward, with a 3 % resting bypass.	Firing alone, or bypass and firing. A resting bypass of 0 or 8 %.	Step up: in band at 6.8 min, IAE 453 against 16751 (Table 10.7). Day: 0 gives 116.3 min out of band, 8 discards 509 kW h (Table 10.10).
11, tunnel shield	The screw holds the chamber, ring builds stop it, and five limits bound the advance.	The least of five limits with external reset, a screw-room limit, a bentonite hold and a chainage feedforward.	A fixed advance. No hold. A fixed pressure setpoint under the river.	Drive 46.9 h and 0.0 mm against 57.6 h and 245.3 mm. Fixed advance: torque 90.04 %. No hold: 21.7 mm (Table 11.5).
12, fan car	Fans, tires and cold-gas thrusters at three prices. The tank runs out by pressure.	Fixed priority: fans, tires at the slip limit, thrust metered to 3.4 g. A pressure abort.	Thrust wide open. A 12 s hover without the abort.	Launch 0.83 s on 29.1 kg against 0.60 s on 40.5 kg (Table 12.5). A 12 s hover lands at 2.83 m/s with 22.6 % left, the abort at 0.94 m/s (Table 12.6).
13, quadruped	Feet that push but cannot pull, inside friction cones that change with each step.	Weighted least squares inside the true cone, the actual leg limits and the sensed contacts.	A cone planned at 0.6 on 0.15. Nominal limits for a weak leg. Scheduled feet.	Slip 15.6 against 1.07 cm. Weak leg 8.43° against 3.22°. Scheduled feet fall. Course finished at 9.25 s against a fall (Table 13.6).
14, currency peg	A vault with a floor, a drain that grows as it empties, two instruments for three limits.	A capital-controls override checked every day, outside the weekly calendar.	Every setting without controls. Any controller against a monetized deficit.	Override: 4.97 months of imports, 6.72 %, controls 0.384 yr. Shipped 2.45 months (Table 14.6). Deficit: attack at 2.252 yr under every controller (Table 14.5).

Every override has a measured price in throughput. The coater's override held every zone under the cap through a fan loss at about a fifth less output, a mean of 27.6 m/min (Table 6.4). On the drill's exit climb the override alone made 44.2 m/h where the optimizer, which predicts the pressure, made 76.6 m/h and stayed clear. The startup's 24-month runway target kept $1.50 million at the lowest and left $42.3 million of founders' value at year 5, against $74.0 million at 12 months (Table 8.5). Section 1.4.2 called this the back-off, and each chapter measured it.

A resting position or a priority is a purchase. The CO_2 cycle's resting bypass bought tracking with fuel, at about 3.8 kW h for each minute out of band recovered, and only 3 %, of the three positions tried, passed both clauses of the dispatch day. The fan car's metered launch spent 28 % less gas than the wide-open one for 0.23 s. The quadruped's assumed friction coefficient traded speed of recovery against slip, and the trigeneration plant held 0.5 MW of electric cooling as the reserve to which its valve-position controller returns the fast actuator. None of these numbers comes from a tuning rule. Each is a price.

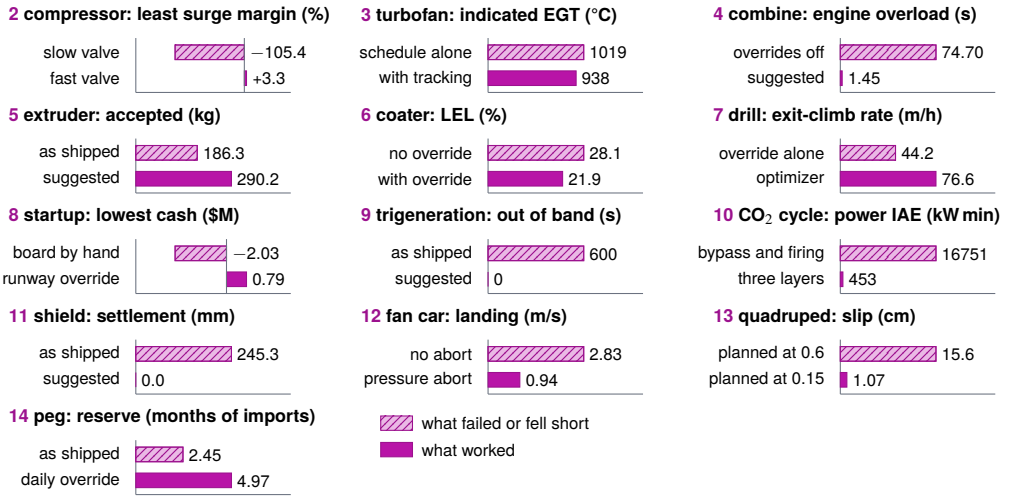

Figure 15.1 One like pair of numbers from each row, as the "Measured result" column of Tables 15.1 and 15.2 prints them. Each panel has its own scale, and the compressor and startup panels run below zero, with a line at zero and the negative value printed to its right. The hatched bar is the structure or setting that failed or fell short and the solid bar the one that worked; a value of zero has a label and no bar. On the drill both rates kept clear of the pressure limit, and the override alone was the slower. Notice that the closest pairs, on the turbofan and the coater, still straddle their limits, 950 °C of indicated exhaust temperature and 25 % LEL: a small difference in the number was the difference between passing and failing.

A stock that depletes turns time into a limit. The peg of Chapter 14 held only with a protective layer that ran every day. Checked at the weekly decisions instead, the same override lifted late and kept the controls in force for 0.521 year, past the six-month limit. Against a monetized deficit no controller moved the date of the attack, which credit growth set at 2.252 yr. The fan car's tank ran out by pressure with 22.6 % of its nitrogen still inside, so the abort that read the tank pressure landed the car softly while the reserve on the mass gauge never fired. Where a protection worked, it read the quantity that actually runs out, and acted faster than the policy it protected (Section 1.4.9).

15.2 Choosing a Structure

The chapters suggest an order of questions for a plant at its limits. Each names a chapter where its answer decided the result.

1. *Which limits are active at the best operating point, and what does each unit of back-off cost?* Price the margin first, in grain, kilograms or meters an hour (Chapter 4, Chapter 5).
2. *Can the protective actuator act fast enough and far enough?* Check the stroke time and capacity against the fastest credible upset before any tuning (Chapter 2).
3. *One actuator and several limits?* Put the demands in a low or high selector, floors after caps, with external reset from the selected output (Chapter 3, Chapter 4). Then test that each override lets go when its limit clears (Chapter 6).
4. *Several actuators for one variable?* Split range shares one demand among actuators of similar cost. A fast, costly actuator beside a slow, cheap one calls for mid-ranging on the unclamped demand, with the resting position set as a price (Chapter 9, Chapter 10). Many actuators with limits call for allocation, by priority where one actuator is free and another is spent (Chapter 12) and by weighted least squares where the limits change each period (Chapter 13).

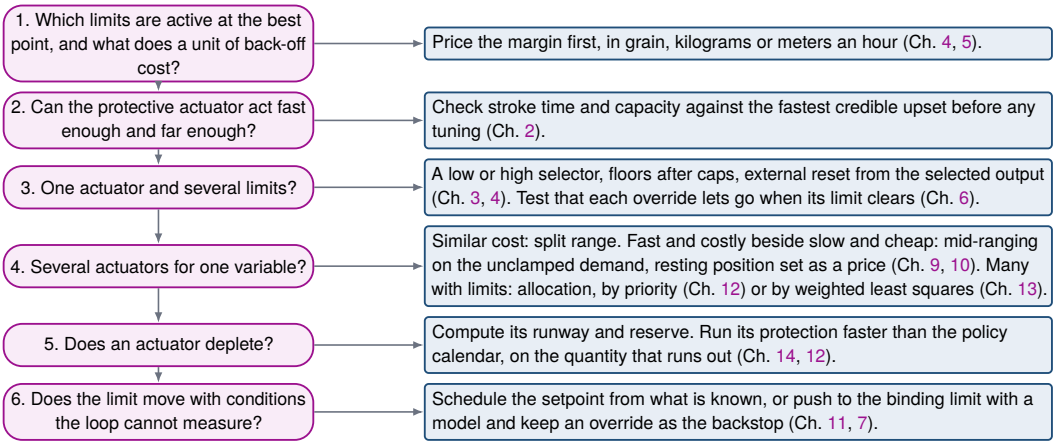

Figure 15.2 The six questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the price of the margin and the reach of the actuator come before any selector is drawn: a structure cannot protect a limit that its actuator cannot reach.

5. *Does an actuator deplete?* Compute its runway and reserve, and run its protection faster than the policy calendar, on the quantity that actually runs out (Chapter 14, Chapter 12).
6. *Does the limit move with conditions the loop cannot measure?* Schedule the setpoint from what is known, as the shield's chainage feedforward does from the geology, or push to the binding limit with a model and keep an override as the backstop (Chapter 11, Chapter 7).

Figure 15.2 draws the six questions in order.

A predictive controller handles limits and prices in one optimization, and it paid where it coordinated handles that the selectors moved in sequence. Over the CO_2 cycle's dispatch day it left 781 kW min of power error against 8425 for the shipped hierarchy and 1050 for a retuned one (Table 10.11). Through a wet flour lot it accepted 212.9 kg with nothing off specification, against 169.4 kg for the loops (Table 5.11). It did not always win. On the compressor the PI override given the throttle signal through a high select held +16.16 % of margin against the predictive controller's +15.00 % (Table 2.6). On the combine the MPC failed the challenge with 48.60 s of engine overload where the selector network passed with 1.45 s (Section 4.6.8). Much of the extruder's gain came from its estimate: the loops given the MPC's estimate accepted 186.0 kg against the MPC's own bias form at 177.6 kg (Table 5.9). Volume 9 treats that choice in full.

15.3 Where Practice Is Heading

Simple elements are being given a theory. A 2023 review collects the standard elements of advanced regulatory control, among them selectors, split range, cascade and feedforward, and challenges the view that predictive control can solve every control problem while these elements are ad hoc. It calls for better teaching of the elements and for research on their theory and design methods [1]. Split range itself has been generalized to several independent controllers with one setpoint, of which only one is active at a time, and the same work notes that predictive control needs a detailed dynamic model and does not use one input at a time [2]. Feedback optimizing control, which reaches the economic optimum by choosing what to control rather than by solving an optimization online, has been surveyed as a family of real-time optimization methods [3]. The design and performance of mid-ranging controllers were studied in the 1990s [4], and the resting position of Chapter 10 shows that the question is still economic as much as dynamic. In the editor's

view, the selector networks of this volume will increasingly be designed from the active constraints of the plant's optimum rather than drawn from experience, as Chapter 1 began to do.

Allocation is solved online. Control allocation, the sharing of a demanded force or moment among redundant actuators, has been surveyed as a field of its own [5]. On a quadruped robot, ground reaction forces have been planned by convex predictive control over horizons up to 0.5 s and solved in under 1 ms, with one set of gains and weights across many gaits [6]. The fixed priority of Chapter 12 and the weighted least squares of Chapter 13 are the two ends of that range. Industrial predictive control technology, the other route to the same problem, was surveyed in 2003 [7].

Constraint enforcement becomes a separate layer. Reference and command governors modify the setpoint given to a loop already designed, so that its constraints hold, and their theory and applications have been surveyed [8]. Control barrier functions express safety as the forward invariance of a set and join it with performance objectives in a quadratic program, demonstrated on adaptive cruise control and lane keeping [9]. A 2023 survey presents safety filters that keep state constraints satisfied under the physical limits of the system whatever the controller behind them asks, and observes that missing safety certificates still keep many new designs inside research environments [10]. In the editor's view, this is where the override of this volume is going. A controller written or tuned by an agent, or learned from data, will sit behind a layer that protects the limits, and that layer will have to be simpler, faster and better verified than the controller it protects, as the daily override of Chapter 14 was faster than the weekly policy.

> ⧖ HISTORY **1942 to 1997: From a Fuel Stop to an Estimated Limit**
>
> The timeline of Chapter 1 places a jet-engine governor in 1942. The first jet engine built in the United States that year had a hydromechanical governor, a minimum-flow stop on the fuel valve against flame-out and a maximum-flow schedule against over-temperature. Flight research later brought estimates into the control. A distortion-tolerant control, flight-tested in 1997, estimated the inlet distortion from engine-face pressures and flight-control estimates and trimmed the engine's actuator commands for stability. A Kalman filter that estimated engine parameters in flight informed a self-tuning onboard engine model implemented in a military engine's control unit [11]. In the editor's reading, the fixed caps of Chapter 3 stand at the start of that arc, and the estimated limits of Section 15.3 at its end.

Limits are estimated where no sensor reaches them. In an aircraft engine, thrust cannot be measured, and limited quantities such as turbine inlet temperature and minimum stall margin cannot be measured directly. An onboard engine model with a tracking filter lets a loop close on such unmeasured values, but the health parameters to be estimated generally outnumber the sensors, which can leave significant error in the unmeasured outputs. One such approach was not yet fully implemented in operating engine controls, because of the processing it needs and of certification issues with a model-based sensor estimate [12]. The limiters of Chapter 3 read an indicated exhaust temperature. In the editor's view, an override that acts on an estimated limit will need its back-off sized from the error of the estimate, which is gap 5 of Section 15.5.

Predictive control reaches the protective layer. For industrial centrifugal compressors, nonlinear predictive controllers and a sequential quadratic programming approximation have been formulated for process and anti-surge control together, tested in simulations that emulate a realistic test bench and timed on an industrial programmable logic controller. In those scenarios they outperformed standard industrial control on surge margin, tracking and actuation without significantly more computation [13]. A study of predictive control for a supercritical carbon dioxide power converter argues that one predictive regulator might outperform several PID loops that each need their own tuning and decoupling, because it coordinates the inputs [14]. On the compressor of Chapter 2, the PI

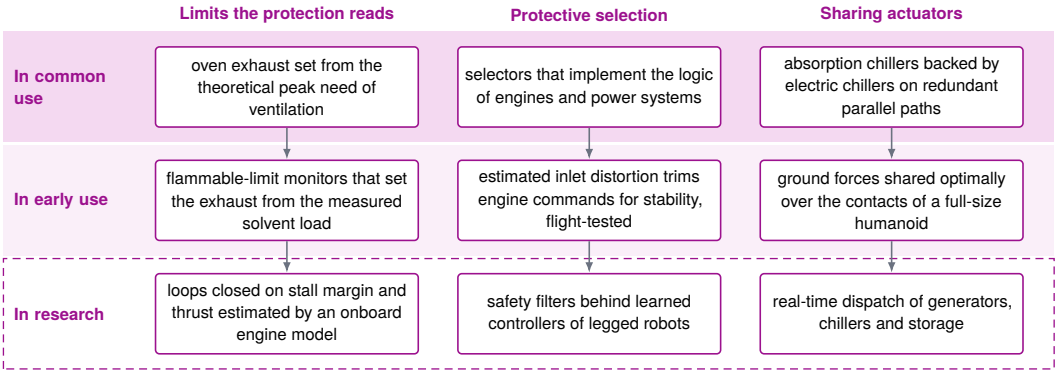

Figure 15.3 Developments in three threads. The top lane holds what the cited source calls standard, favored or common practice, the middle lane what it reports offered as equipment, flight-tested or shown on hardware, and the bottom lane what it reports from research, as a proposal or as a need, in the editor's reading. Sources by thread: limits the protection reads [12, 15], protective selection [11, 16, 17] and sharing actuators [18–20]. Notice that each thread moves from a fixed rule toward a computation that must itself be checked while the plant runs.

override given the throttle signal kept more margin than the predictive controller (Section 15.2). In the editor's view the predictive layer will take over the coordination, while a simple override keeps the last word on the limit.

Ventilation follows the measured solvent load. Ovens that drive off flammable solvents have set their exhaust from the theoretical peak need of ventilation, so that the highest vapor concentration the process can produce stays well below the lower flammable limit. Monitoring equipment that tracks the solvent load in real time and sets the ventilation from the real need keeps a safe ratio with less exhaust air, and linking it to an adjustable-speed drive on the exhaust fan improves the efficiency further [15]. The coater of Chapter 6 reads six zone monitors and trims the exhaust first when a zone nears its cap. In the editor's view a limit measured continuously is what lets the back-off of Section 15.1 shrink. Figure 15.3 sets these developments in three threads.

15.4 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.

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.

15.4.1 How This Volume Was Made

Each of the thirteen apps was built and tested by agents. An agent then wrote the paper on it with verification notes that trace each number to a run. Each paper received a referee report, and at least four of the thirteen are self-reviews by the agent that wrote the paper. Defects went into a log for each app. For this book one agent per chapter reran the numbers its exercises use. The editor set the scope of each round and what each agent could change, and statements that runs contradicted were left for the editor.

The defect logs. Twelve of the thirteen logs give each entry a severity, 108 entries in all. Of these, 35 are wrong behavior of an app, 56 are documentation, 15 are cosmetic or display, and 2 are of other kinds. Nine of the 35, in seven logs, are a controller or an operator aid that read the plant's true state or the run's hidden settings where a plant would offer only a measurement. The drill's log records no defect and says that every number its paper quotes was reproduced on the shipped app.

A limit with no override. The extruder's dryer air had a hard limit of 150 °C, while the moisture cascade could ask for up to 160 °C. With humid intake air and a low moisture setpoint the air ran at 160 to 161.1 °C for half an hour, and a heater driven by hand reached 175.0 °C. No shipped or suggested run had passed 147.4 °C, so the scenarios could not show it. An audit made weeks after

the paper, when a predictive option was added, found it in runs at the ends of the setting ranges, and a high-limit override that acts in every mode now holds the air under the limit.

An override that did not let go. On the coater a zone that joined the speed-limited set at a low speed was then held at its target inside the hold band and never left, so the limit stayed down for good. One run of case 3 rested at 28.1 m/min, and sweeps of the moment of the setpoint change found 10 of 81 and 22 of 82 runs resting under 31 m/min. A release rule brought both sweeps to none. A second rule fixed a deep cut in which the upstream zone took the speed handle as well and the line settled near 21.5 m/min instead of 27.6 in 3 of 8 seeds, and in 6 of 8 at half the step. A third candidate rule was dropped because it let the line run on cut heat at 44.7 m/min. Both defects showed away from the case as shipped, at another moment of the change or another seed, and neither rule changed the challenge's runs, so the graded case could not have shown them.

A clamp that ignored its limit. The CO_2 cycle's turbine inlet setpoint was clamped at a constant 618 °C, two degrees under the default metal limit. With the learner's limit lowered to 605 °C, scenario 8 spent 452.8 min of 500 above it, and the predictive controller's bound was the same constant. An audit that moved the limit away from its default found it. The clamp now follows the limit the run declares, and no run at the shipped limit moved, which is why no shipped scenario had shown it.

A supervisor that pushed the wrong way. The extruder's supervisor treated a light product like a dense one and lowered the feed for both. Less feed puts more work into each kilogram and makes the pieces lighter still, so on the light side the correction drove the violation it answered. The coordinator measured a one-sided fix in a scratch copy, found that it moved the challenge's headline result and doubled the off-specification product of the light-lot case, from 96.6 to 192.8 kg, and kept the condition as implemented, stated in the chapter.

A score that skipped the violation. The peg's scorecard recorded its extremes only while the peg stood. On the day a rate defense broke the peg at a 15 % overnight rate, the tile still read 5.00 %. Once fixed, the challenge scorer counted the exit day, and seven rows of its table moved. Each of these five errors lay in a protective element or in the measure that judges one, and most showed only away from the cases as shipped.

The book round. The thirteen chapter records list 179 rows of numeric verification. Most chapters showed no drift. The coater's case 3 starting values, once printed as 18.5 and 18.3 % LEL where runs give 18.25 and 18.57, and the shield's settlement after one stop, 1.76 mm (1.78 only once the second ring build has begun), were corrected in paper and chapter, and three chapters quoted counts of app checks that have since grown. The CO_2 cycle's Conclusion disagreed with its own tables in five figures and now uses the table values. The agent writing Chapter 1 flagged a sentence that ranked the quadruped's vertical force above its horizontal force, where the two weights are equal, and the chapter pass reworded it. At the end of that round the editor's list of defects held eleven entries for this volume.

Every error above was caught by a check that could fail: a setting driven to the end of its range, a sweep of timing and seed, a limit moved away from its default, a score compared with the logged states. For constraint control the record adds one rule to those of the shared section. Test every protective element away from the nominal case, and ask of each override when it lets go.

15.5 Gaps in Override and Constraint Control

Six gaps belong to the topic of this volume, and the gaps common to agent-era practice follow in Section 15.5.1. In the editor's view none is closed. Each item gives the chapters' evidence, then the

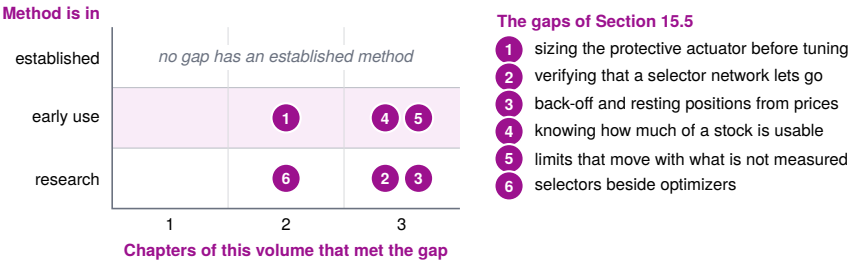

Figure 15.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: two for gap 1 (Chapters 2 and 9), three for gap 2 (6, 11 and 3), three for gap 3 (10, 4 and 13), three for gap 4 (12, 8 and 14), three for gap 5 (7, 11 and 6) and two for gap 6 (4 and 10). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that four gaps were met by three chapters each. Of those four, the two in early use, usable stock and moving limits, rest on a rule of thumb and on prediction models that their own sources call variable or inaccurate.

state of the literature with its sources. Why each gap is hard and what progress would look like are the editor's view. Figure 15.4 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. *Sizing the protective actuator before tuning it.* The compressor's slow valve and the trigeneration plant's lost capacity were found by simulation, case by case. A review of anti-surge control notes that sensor delay and slow valve response often limit the control algorithm, that faster actuators add cost and complexity, and that strategies shown on test benches or in simulation must still be validated on full-scale compressors with safety mechanisms [23]. The gap is hard because a saturated actuator looks on a trend like a slow tuning, as the compressor's did when no gain changed the outcome. A method that states the stroke time and capacity a protection needs against a stated upset is not standard.

2. *Verifying that a selector network lets go.* The coater's resting points were found only by sweeping the timing of a change over 81 runs, and the shield and the turbofan carry networks of five and eight limits. The Open challenge box after this list takes the gap further.

3. *Setting back-off and resting positions from prices.* The resting bypass, the engine limit of the combine and the planned friction of the quadruped were each chosen by a sweep against a scorecard. A design study of the same kind of cycle chose the throttle valve's pressure drop so that the turbine bypass valve stays around 10 % open during throttling [24], and a report on the dispatch of combined cooling, heating and power found savings of 8 to 25 % that depend on pricing, season, forecast accuracy and grid constraints [20]. The gap is hard because the price of a margin moves with markets and seasons while the margin is set once, at commissioning. Methods that set margins from prices, and move them as prices move, are thin.

4. *Knowing how much of a stock is usable.* The fan car's tank ran out by pressure with nitrogen inside, the startup's credit ceiling moved with its earnings, and the peg's date was set by credit the defense did not touch. For reserves, a central bank task force found the ratio of short-term debt the most popular adequacy indicator but not the only one, and gave three to four months of imports as a rule of thumb [25]. At a later meeting of central bank officials, some participants noted that the adequate level can vary considerably over time, particularly in crises [26]. The gap is hard because what can be used depends on the state in which it is used, a tank pressure, a trailing earnings figure or the confidence of a market. The usable runway is still modeled plant by plant.

5. *Limits that move with what is not measured.* The drill's pressure limit followed the cover and the soil, the shield's target rose under the river, and the coater's binding zone moved with the drying front. A thesis on drilling found that the commonly used Bingham plastic model overestimates the annular pressure and that the commonly used cavity expansion method significantly overestimates the allowable pressure, which can lead to hydraulic fracture [27]. The gap is hard because such a limit is often learned only when it is crossed. A limit estimate that states its own uncertainty, so the back-off can follow it, is not yet routine.

6. *Selectors beside optimizers.* The combine's MPC collected more grain in the challenge and failed it on overload, and the CO_2 cycle's MPC kept limits that depend on inner loops tuned as its model assumes. A review of learned legged locomotion notes that safety filters have been studied extensively in control but applied to locomotion only in a limited way, and that more research could help certify such controllers [17]. The gap is hard because the optimizer and the override see the same limit through different models, and either can be wrong. A rule for which layer owns which limit, and how an override backs up an optimizer whose model is wrong, is still written per plant.

> ⚠ **OPEN CHALLENGE** **An Override Network That Is Proved to Let Go**
>
> Gap 2 runs through the coater, the shield and the turbofan. A fieldbus standard gives the control selector a back-calculation signal, so a controller it de-selects halts integral action [28]. A unified framework treats an override as a substitution of the plant input, which wrongly updates the controllers not chosen and bumps the input, and shows every known linear anti-windup scheme as a special case [29]. In the editor's view the gap is hard because release depends on the order and timing of the limits, which no linear framework covers. Progress would be a check that lists each combination of active limits and shows its exit.

15.5.1 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 [30]. 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 [31]. 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 [32]. 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 [33]. 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 protective loop was bounded by its actuator before its tuning, and external reset decided when an override took over.
- Every override, resting position and priority had a measured price in throughput, fuel or gas.
- A stock that depletes made time a limit, and protections worked where they read what actually runs out.

References

[1] S. Skogestad, "Advanced control using decomposition and simple elements," *Annual Reviews in Control*, vol. 56, p. 100903, 2023.

[2] A. Reyes-Lúa and S. Skogestad, "Multi-input single-output control for extending the operating range: Generalized split range control using the baton strategy," *Journal of Process Control*, vol. 91, pp. 1–11, 2020.

[3] D. Krishnamoorthy and S. Skogestad, "Real-time optimization as a feedback control problem: A review," *Computers & Chemical Engineering*, vol. 161, p. 107723, 2022.

[4] B. J. Allison and A. J. Isaksson, "Design and performance of mid-ranging controllers," *Journal of Process Control*, vol. 8, no. 5–6, pp. 469–474, 1998.

[5] T. A. Johansen and T. I. Fossen, "Control allocation: A survey," *Automatica*, vol. 49, no. 5, pp. 1087–1103, 2013.

[6] J. Di Carlo, P. M. Wensing, B. Katz, G. Bledt, and S. Kim, "Dynamic locomotion in the MIT Cheetah 3 through convex model-predictive control," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 1–9.

[7] 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.

[8] E. Garone, S. Di Cairano, and I. Kolmanovsky, "Reference and command governors for systems with constraints: A survey on theory and applications," *Automatica*, vol. 75, pp. 306–328, 2017.

[9] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.

[10] K. P. Wabersich, A. J. Taylor, J. J. Choi, K. Sreenath, C. J. Tomlin, A. D. Ames, and M. N. Zeilinger, "Data-driven safety filters: Hamilton-Jacobi reachability, control barrier functions, and predictive methods for uncertain systems," *IEEE Control Systems*, vol. 43, no. 5, pp. 137–177, 2023.

[11] L. C. Jaw and S. Garg, "Propulsion control technology development in the United States: A historical perspective," NASA Glenn Research Center, Cleveland, OH, Technical Memorandum NASA/TM-2005-213978, 2005, nTRS 20050238468.

[12] S. Garg, "Aircraft turbine engine control research at NASA Glenn Research Center," NASA Glenn Research Center, Cleveland, OH, Technical Memorandum NASA/TM-2013-217821, 2013, nTRS 20130013439.

[13] G. Torrisi, S. Grammatico, A. Cortinovis, M. Mercangoz, 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.

[14] R. Ponciroli and R. Vilim, "Analyses of model-based predictive control for a S-CO_2 Brayton cycle power converter," Argonne National Laboratory, Argonne, IL, USA, Tech. Rep., 2016.

[15] U.S. Department of Energy, "Use lower flammable limit monitoring equipment to improve process oven efficiency," National Renewable Energy Laboratory for the U.S. Department of Energy, Industrial Technologies Program, Golden, CO, Process Heating Tip Sheet #11 DOE/GO-102007-2511, Oct. 2007.

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

[17] S. Ha, J. Lee, M. van de Panne, Z. Xie, W. Yu, and M. Khadiv, "Learning-based legged locomotion: State of the art and future perspectives," *The International Journal of Robotics Research*, vol. 44, no. 8, pp. 1396–1427, 2025.

[18] K. Darrow and B. Hedman, "Opportunities for combined heat and power in data centers," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2009/052, Mar. 2009, prepared by ICF International.

[19] S. Hyon, J. G. Hale, and G. Cheng, "Full-body compliant human–humanoid interaction: Balancing in the presence of unknown external forces," *IEEE Transactions on Robotics*, vol. 23, no. 5, pp. 884–898, 2007.

[20] S. Samuelsen and J. Brouwer, "Novel controls for time-dependent economic dispatch of combined cooling heating and power (CCHP)," University of California, Irvine, Irvine, CA, Final Report DOE-UCI-01113, 2013.

[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] J. Du, Y. Zhang, M. Martínez García, and A. Spencer, "Review on compressor surge monitoring, modeling, and anti-surge control approaches," *Journal of Dynamics, Monitoring and Diagnostics*, vol. 4, no. 4, pp. 292–307, 2025.

[24] A. Moisseytsev and J. J. Sienicki, "Development of the ANL plant dynamics code and control strategies for the supercritical carbon dioxide Brayton cycle and code validation with data from the Sandia small-scale supercritical carbon dioxide Brayton cycle test loop," Argonne National Laboratory, Argonne, IL, USA, Tech. Rep. ANL-ARC-218, 2011.

[25] International Relations Committee Task Force, "The accumulation of foreign reserves," European Central Bank, Frankfurt am Main, ECB Occasional Paper 43, 2006.

[26] M. S. Mohanty, "Market volatility and foreign exchange intervention in EMEs: what has changed? An overview," in *Market Volatility and Foreign Exchange Intervention in EMEs: What Has Changed?*, ser. BIS Papers. Basel: Bank for International Settlements, 2013, no. 73, pp. 1–10. [Online]. Available: https://www.bis.org/publ/bppdf/bispap73.htm

[27] A. Rostami, "Prediction and evaluation of annular pressure in horizontal directional drilling," Ph.D. dissertation, University of Alberta, Edmonton, AB, 2017.

[28] 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.

[29] M. V. Kothare, P. J. Campo, M. Morari, and C. N. Nett, "A unified framework for the study of anti-windup designs," *Automatica*, vol. 30, no. 12, pp. 1869–1883, 1994.

[30] 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.

[31] 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.

[32] 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.

[33] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

15.6 Exercises

15.1 Discussion · 8 min

After a surge on the slow-valve case of Chapter 2, an agent proposes raising the anti-surge controller's gain and shortening its integral time, because "the controller reacted too slowly". Using the chapter, explain why the proposal will not prevent the surge, list the checks you would ask for first, and name the changes that did hold the closure.

15.2 Discussion · 8 min

A design review lists integrator tracking as an optional refinement of a selector network. Using Chapter 3 and Chapter 11, argue that the anti-windup arrangement is part of the protection, and state what a reviewer should ask about it on a new selector.

15.3 Discussion · 8 min

A manager asks that the resting bypass of the CO_2 cycle in Chapter 10 be set to zero to save heat. Using that chapter and Chapter 9, answer the manager, and describe how you would set a resting position on a new plant.

15.4 Discussion · 8 min

Using Chapter 14, Chapter 12 and Chapter 8, write three requirements for a protective layer on an actuator that depletes, and support each with a measured result.

15.5 Discussion · 9 min

An agent delivers a new override network for a coating line with a vapor limit, and its run of the graded challenge passes. Using the cases of Section 15.4.1, list the tests you would require before the network is accepted, and say which defect each test would have caught.

15.6 Discussion · 8 min

A project proposes one model predictive controller in place of every selector and override. Using Chapter 2, Chapter 4, Chapter 5 and Chapter 10, argue where it should take over, where the selectors should stay, and how to arrange the two.

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) 17.4 kg/min, 0.4 kg/min more than at 85 °C. (b) 0.15 kg/min per minute.

1.3 23 kg/min at 83 °C, 20 kg/min at 84.5 °C. It takes over at 85 °C.

1.5 (c). A clamped output carries no information about how far past the clamp the fast controller wanted to go.

1.7 (a). A short tracking time holds the demand close to the applied output, so a fast approach is enough to pass it below.

1.9 (d). With $T_t = \tau_I$ the error terms cancel and the integral is a lag of the applied output.

Chapter 2. Compressor Surge and Protective Override

2.1 (c). The trace turns positive once $\psi_c' > 1/(B^2\psi_T')$, inside the peak, while $\det J$ stays positive at every throttle position.

2.3 (a). With $u_{cmd} = 0$ and $T_t = \tau_I$ the error terms of Equation (2.10d) cancel, and $I_k = (1 - \Delta t/\tau_I)\, I_{k-1}$.

2.5 (a) $\psi^* = 0.4822$. (b) $\psi_c' = 0.907$, $\psi_T' = 2.756$, $B_{crit} = 0.633$. The default B surges and $B = 0.5$ does not. (c) Nine times smaller.

2.7 (a) -135.1 %, 0.631, $\varphi_{min} = -0.175$, alarm at 11.22, period 7.51. (b) -43.8 %, 0.211, flow 0.3207 and no alarm. With $B = 1.5$, alarm at 34.32 and -137.9 %. At $B = 0.5$ each step rings and decays, because B_{crit} at the final point, 0.596, is above 0.5.

2.9 (a) $+15.0$ %, 2.46, 85.6. (b) $+4.7$ %, 2.27, 168. (c) -8.6 %, 2.11, 529. The signal supplies 10.3 of the 23.6 points over the PI. The plan adds less recycle and far less valve travel, not margin.

2.11 (a) 0.1865 and 0.2197, 32.4 % and 36.6 % of the compressor flow, 18 % more recycle for 5 more points of margin. (b) The true cost is higher: hot gas raises the inlet temperature and moves the characteristic, so a cooler is needed and the extra power is lost as heat for as long as the valve stays open.

Chapter 3. Turbofan Limiter Selection

3.1 (b). With the HPT nozzle choked, $W_4\sqrt{T_4}/P_4$ is fixed, so the gas heated at once by the fuel passes only at a higher pressure.

3.3 (d). The floors are applied after the caps, so a minimum fuel above what the schedule allows wins the valve, and the schedule cannot protect the margin.

3.5 (a) $\dot{I}_i = (W_f - I_i)/T_i$. (b) $e = \rho T_i/K_i$, with $\rho = 0.534$ kg/s per second, 89 K. (c) About 45 K. With tracking off the integral winds up toward 1.6 kg/s and the limiter takes over only after the EGT has passed its line.

3.7 (a) 5.37 s, 11.3 %. (b) 4.74 s, 9.5 %, EGT limiter. (c) 4.66 s, N_1 peaks at 98.0 % against 96.1 % with tracking on, because the governor's integrator wound up while the schedule held the valve.

3.9 (a) 8.49 s, 14.8 %, 996 °C; misses time and EGT. (b) 5.17 s with 1019 °C (fails on EGT); 8.55 s with 934 °C (fails on time). (c) Both pass at 5.86 s and 9.4 %, 936 °C fixed and 938 °C scheduled. The gain schedule is not what passes.

3.11 (a) 78 K, about 3.9 K per kelvin if linear. (b) About ISA + 5 K. Scenario 8 runs at ISA + 15 K, where the limiter trims the worn engine to 91.5 % of its rating.

Chapter 4. Combine Harvester Feed-Rate Overrides

4.1 (b). Moisture lowers the rotor's capacity and raises the loss without changing the feed, and the feed loop measures only the feed.

4.3 (d). It is conditional integration in the outer loop of the cascade: a raise is refused by tracking the previous value.

4.5 (a) 9.41 and 14.9 t/h per km/h, loop gain 0.742. (b) Phase -180.0 degrees, $|L| = 1.04$, so the loop is just past its stability limit and hunts with a period near $2\pi/0.4095 = 15.3$ s. (c) 1.93. A longer τ_I changes the phase too and moves the crossover.

4.7 (a) 10.48 t, 1.00 %, 1.84 t standing; 12.23 t, 1.73 %, nothing standing. (b) 0.106 and 0.215 t lost. The faster machine loses 0.11 t more and saves 1.75 t from the rain.

4.9 (a) 9.03 t, 10.86 %, two plugs, 74.7 s: fails all four targets. (b) 10.77 t, 1.71 %, none, 1.5 s: passes. (c) 10.25 t, 1.91 %, none, 0.0 s: misses 10.5 t. The trims let the rotor separate better, the supervisor spends that margin on feed, and the engine becomes the active constraint.

4.11 (a) About 55 t/h, against 55.1 t/h measured over 80 s to 120 s. (b) $n_e \approx 0.905$, a rotor near 860 rpm, against 89.9 and 854 rpm measured. The rotor loses capacity, and the loss rises from 1.00 to 1.64 %.

Chapter 5. Cereal Extrusion Overrides and Density Trim

5.1 (a). Equations (5.14) and (5.16) use the same certificate value, so the estimate agrees with the setpoint whatever the flour holds.

5.3 (c). The SME falls from 43.0 to 34.3 kWh/t, the melt to 150.2 °C and the density rises to 145.6 g/L.

5.5 (a) 77.5 %. (b) 87.5 %, a slack of -1.84, so the supervisor backs off rather than push. (c) About 433 kg/h and 439 kg/h. Case 4 settles at 437.2 kg/h, between them.

5.7 (a) 209 s and 331 s, 64 kg, all dense. (b) 24 kg, a trim of -2.67 points against the 2.61 points of the certificate error. (c) 18 kg.

5.9 (a) 186 kg accepted, 83 kg off-spec, 74.0 %, 93.7 kg short. (b) 290 kg, 34 kg, 87.4 %, passes. (c) 244 kg, 75 kg off-spec of which 44 kg light, 99.2 %, fails. Lot B's correction carries into lot C.

5.11 (a) $0.876 \times 540 = 473$ N m and 22.3 kW. On a 600 N m drive that is 78.8 %, the 78.9 % from the unrounded torque. (b) The extra water sits on the supervisor's quality margin and raising the override limit leaves little room to the hard limit, so the larger drive is the option that keeps every margin.

Chapter 6. Explosive-Limit Control of an Electrode Dryer

6.1 (d). A wet zone evaporates at the rate its air sets at any speed, so only the zone the drying front enters gains vapor.

6.3 (b). The transit time is the floor for both loops, so the predictor buys gain, not time.

6.5 (a) 0.120 kg/s. (b) 77.86 m/min. (c) 54.5 m/min, well above the 40.6 m/min that the push holds in the challenge after the fan loss.

6.7 (a) 30.3 % LEL in zone 4, crossing the cap about 62 s after the change. (b) 22.9 % LEL in zone 4 and a mean of 35.6 m/min. The speed hunts because zone 4 gains about 2 points per m/min and feels a speed change only about a minute later.

6.9 (a) 2520 m^2/h, missing the cap, the coat weight, the migration and the throughput. (b) 3076 m^2/h, all targets met. (c) 2075 m^2/h, missing the coat weight and the throughput.

6.11 (a) 60 min and 2520 m^2. (b) About 74.0 kAh per side. The exit web pyrometer moves with the residual solvent. Nothing on the line measures the migration.

Chapter 7. Horizontal Directional Drilling Rate Limits

7.1 (c). Pushed, the bit cut at 1.19 m/min against 2.43 m/min rotating, so every extra pushed rod is a slower rod.

7.3 (a). P_i grows with the drilled length s_i, and P_{lim} fell from 511 kPa at 160 m to 185 kPa at 170 m.

7.5 (a) $A_{ann} = 75.6\,\mathrm{cm}^2$, $v = 0.368\,\mathrm{m/s}$, $\eta = 0.891$, $c = 0.267$, about 5.87 m/min. (b) 3.20 m/min, so the shipped bit cut 2.6 times faster than its mud could carry. (c) 2.78 m/min, above the 2.65 m/min the torque ceiling allows.

7.7 (a) 126.0 m/h, one return, first at 114.6 m. (b) 108.4 m/h, no return, worst ratio 0.71, held by the mud. (c) 0.13 kWh/m as shipped and 0.24 kWh/m with the optimizer, mostly the pump at full flow and 8 MPa.

7.9 (a) Not met: exit 2.54 m off, one inadvertent return, torque 91.1 %, in 1.83 h. (b) Met: exit 0.60 m off, cover 4.38 m, no return, torque 85.8 %, 2.13 h. (c) The exit, 1.70 m off: the head is dead reckoned from the river to the exit.

7.11 (a) 71.6 s and 83.6 s per rod give 150.8 and 129.2 m/h, 17 % more for a bit 55 % faster, against 16 % measured. (b) The fast bit gains 26 %, the slow one 22 %: the rod change is the larger share of a fast rod.

Chapter 8. Growth Control Against a Cash Floor

8.1 (c). Fixed assets grow by \$1.82M and stock at book by \$1.17M, while equity barely moves.

8.3 (a). Above 1 the first band is saturated and the second draws the line. R&D is cut only above 2.

8.5 (a) \$0.81M and \$0.42M. (b) \$0.308M, with time constants of 6.4 and 10.7 weeks. (c) The ceiling is set by trailing EBITDA, which a payment delay leaves unchanged.

8.7 (a) 1.58 yr, with \$8.37M of revenue, \$1.11M of EBITDA and \$0.51M of net income. (b) $\Delta C = \Delta E + \Delta AP - \Delta AR - \Delta I^{\$} - \Delta W^{\$} - \Delta F = -0.01 + 0.09 - 0.01 - 1.17 - 0.10 - 1.82 = -3.02$. (c) \$2.20M, 69.7 % and \$6.60M: solvent, but smaller than at the start.

8.9 (a) Insolvent at 4.40 yr, in the first recession. (b) Passes with 20.7 % growth, 97.7 % filled and \$272.5M. (c) Misses the 95 % fill target at 79.5 %, because Higgins's rule caps the factory at the rate the equity finances while demand grows faster.

8.11 (a) $0.481 \times 21.9/4 = 2.63$ thousand units, about \$0.39M at book. (b) Obsolete stock and a second product are outside the model. In case 3 feedback alone fills 91.6 % of orders and overstocks after each peak, while the seasonal feedforward fills every order.

Chapter 9. Valve-Position Control in Trigeneration

9.1 (b). The bypass loop holds the hot water at 95 °C by rejecting the surplus exhaust heat, 2.73 MW more of it, so the absorption chiller never sees it.

9.3 (d). At the limit the selector passes the limiter's valve, and the turbine sits at its 7.95 MW capacity however the power loop is tuned.

9.5 (a) $k = 0.0763\,\mathrm{MW}$ per percent. (b) 7.9 points, to 83.8 %. (c) $E_h = \Delta v / K_I = 7.87/0.05 = 157\,\mathrm{MJ}$, about 44 kWh.

9.7 (a) 8.43 MW and 8.15 MW against a setpoint of 6.8 MW, and the limit is never reached. (b) 204 s, 8.71 MW, 491.6 °C. (c) About 651 s. Back-calculation held the power loop's integral at the limiter's ceiling.

9.9 (a) 600 s out of band, store 23.6 %, fails. (b) Trim only: 242 s, 61.3 %, fails. Gain only: 600 s, 19.6 %, fails twice. (c) 0 s, 58.1 %, passes, with 0.33 MWh imported against 0.18 MWh.

9.11 (a) $5.79/1.4 = 4.14\,\mathrm{MW}$ against 7.94 MW, and $1 + 1/1.4 = 1.71$ MW rejected per MW against 2.37. (b) One hot-water circuit keeps the teaching plant simple. A two-stage machine needs a hotter circuit or a steam circuit from the recovery unit.

Chapter 10. Supercritical CO_2 Brayton Mid-Ranging

10.1 (a). The model is a simple recuperated cycle with no recompression branch, so the flow split is not an actuator.

10.3 (c). The power transducer measures net power; the other three are computed from the states for display.

10.5 (a) $u = 0.955$, factor 1.782, 37.13 kJ/kg. (b) 4.40 kg/s, 172.0 kW, 51.3 kW more. (c) Turbine −213.0 kW, compressor −51.3 kW, pump −67.2 kW: 331.5 kW, the table's 929.0 to 597.4 kW.

10.7 (a) 22.0 min, two minutes after the supply warms; 29.80 °C, 597.4 kW, 35.33 %. (b) 4.40, 37.13, 593.0 and 20.58: everything but the compressor inlet holds, and the coolant valve has no travel left.

10.9 (a) It passes: 61.7, 12.6, 23.80 MPa, 280 kWh, inside the limits by 18.3, 7.4, 0.20 and 120. (b) The energy clause: 509 kWh against 400, with the best tracking, 17.6 min. (c) $377/98.7 = 3.8$ kWh per minute.

10.11 (a) $2298/930.4 = 2.47$ and $2635/930.0 = 2.83$; 337 kW. (b) $337 \times 8760 \times 0.02 \approx \$59\,000$; trend the heat-to-power ratio or the calculated recuperator effectiveness with an alarm.

Chapter 11. Tunnel Face Pressure Control

11.1 (d). A full chamber has no self-regulation: balancing the flows stops the pressure, and only a period of excess outflow brings it back.

11.3 (b). The target under the river is 2.944 bar, and the land's blow-out limit expressed at the axis cell is 2.91 to 2.95 bar.

11.5 (a) 0.186 bar a minute, about 2.7 min. (b) 274 and 550 Pa/s per rpm, a factor of 2.0. (c) $K_c = 38.8\,\mathrm{rpm/bar}$, $\tau_I = 35.2$ s.

11.7 (a) 0.177 bar at 940 s, last outside ±0.02 bar at 1120 s. (b) 0.021 bar; 6.63 and 9.97 rpm. (c) 3.31 rpm, as `screw_ff_rpm` rises from 6.63 to 9.94 rpm.

11.9 (a) 245.3 mm, 2079.1 min spewing, 1675.3 min of blow-out risk, 96.1 % torque, 57.6 h: every limit fails. (b) Met: 0.0 mm, 46.9 h, 86.1 %, the verdict at about 48.0 h. The survey reads 1.0 mm at 87.5 m because each report carries 0.5 mm of noise. (c) Not met: 21.7 mm, with 0.78 bar worst deficit and 710 min under the target.

11.11 (a) The conveyor's 120 s to the belt scale and the 60 s averages of both signals. (b) About 74 min earlier. A point is leveled 1.5 h after the face passes and reported 0.5 h later. (c) The extra ground arrives as extra muck, and the screw simply takes it out.

Chapter 12. Performance Auto: Allocation on One Gas Tank

12.1 (d). Blowdown cooling drops the pressure faster than the mass, and below 60 bar the thrusters deliver only p/p_{set} of their rating.

12.3 (b). A demand of 60 % rounds to a 10 ms pulse, so the car sinks until the loop's demand crosses the rounding point, and the height hunts between 1.902 and 1.987 m.

12.5 (a) 716.1 m/s, 23.6 kg/s, 69.8 kg/s. (b) 114.1 kg, against 118.0 kg at the starting rate. (c) $0.60 \times 69.8 - 0.02 \times 69.8 = 40.5$ kg.

12.7 (a) 0.83 s, 3.39 g, 29.1 kg. (b) 0.94 s, 3.00 g, 24.3 kg. (c) 0.60 s, 4.83 g, 40.5 kg. The second saved costs about 44 kg of nitrogen from 3.0 to 3.4 g and about 50 kg from 3.4 g to wide open, which also passes the 3.5 g limit.

12.9 (a) 4.82 g and a crash at 5.81 m/s: the launch passes the 3.5 g limit and the roll from the hover reaches the ground. (b) 3.41 g, 0.1°, 1.22 m/s, 160 kg (38 %), met. (c) 15.4 cm, 0.50 m/s, 109 kg (26 %), met. The feedforward is worth about 51 kg, 12 points of the fill.

12.11 (a) The pressure limit, at 38.1 % of the fill, long before the 8 % gauge reserve. (b) The mass left when the pressure meets the regulator depends on how fast and how long the gas was drawn. The abort reads the measured tank pressure against the regulator setpoint.

Chapter 13. Quadruped Balance Inside Friction Cones

13.1 (c). With a level back the ground under the front hip is 17.7 cm higher than under the rear hip, so the rear legs run out of reach and a foot loses the ground.

13.3 (a). A foot in the air carries no force, so its share is missing from the wrench, and the allocation, trusting the schedule, sees no shortfall.

13.5 (a) 79.9 N and 70.6 N, so friction binds. (b) About 56 N and 81.4 N, so tipping binds. (c) About 7.4 cm, against 10.1 cm measured.

13.7 (a) 5.1°, 0.25 s. (b) 9.3° and 0.33 s at 300, 3.9° and 1.02 s at 60. (c) Braking is capped by tipping, so the stiff loop mostly converts its larger request into a larger nose-down moment.

13.9 (a) It falls at 7.63 s, worst pitch 25.6°, and the 4.6 cm of slip had already failed the target. (b) 9.25 s, 4.0°, 0.55 s, 0.5 cm. (c) No. The map climbs the step, but the 0.6 cone still slides the feet on the patch.

13.11 (a) 38.7°, 31.0°, 8.5°. (b) The narrow cone, tilted by more than its own half-angle, admits no forward force. The allocated forward force stays near −1 N and the speed sags. A true ground normal, from the terrain map, removes it.

Chapter 14. Defending a Currency Peg With a Finite Vault

14.1 (b). Below $R_{thr} = 8\,\%$ each point lost adds 1.25 points to Δe^{e}, and the outflow grows from 5.8 to 9.8 % of GDP a year.

14.3 (d). Unsterilized, the overnight rate rises $1/\alpha = 0.67$ point per point sold, to 10.92 %, and κ times that rise draws capital back.

14.5 (a) 1.875 points of GDP a year per point, doubling in $\ln 2/1.875 = 0.37$ yr. (b) $15 - 1.5 \times 6/1.5 = 9\,\%$, with the overnight rate at 9 %. (c) $0.096 + 7/5.8 = 1.30$ yr.

14.7 (a) $12(9.75 - 5)/5.80 = 9.8$ months. (b) The peg ends at year 1.666, 8.0 months after year 1, so the clock was 1.8 months optimistic. (c) Below 8 % from year 1.304. The outflow grows to 7.94 and 9.84 as expected depreciation rises 1.25 points per point of reserves lost.

14.9 (a) 2.45 months, 10.92 %, −2.44 %: not met, on reserves and the rate. (b) 4.79 months, 7.02 %, −1.63 %, but controls for 1.5 years against a limit of 0.5: not met. (c) 4.96 months, 6.72 %, −0.88 %, controls for 0.384 year: met.

14.11 (a) The trade balance loses $c_T = 1$ point a year, which also lowers output through ϕ, and the premium c_δ draws capital out. (b) So that it engages and releases on the fast signal it protects against. Weekly readings caught the decaying pressure above its release level twice and kept the controls on past the limit.

Chapter 15. Summary and Outlook

15.1 A good answer notes that in the slow-valve case the controller's command left zero $1.0\hat{t}$ after the closure and reached 100 % at $\hat{t} = 11.0$ while the valve was only 34 % open, so the decision was right and the actuator was late, and the minimum margin was −105.4 %. With the undersized valve the command sat at 100 % for 56.2 % of the run and the margin reached −129.0 %. The aggressive tuning of scenario 10 bought margin at seven times the valve activity and left a control-induced oscillation. The checks: command against valve position, stroke time and rate limit against the time to surge, and the capacity needed at the lowest throttle position. What held: a fast valve (+3.3 %) and the throttle signal as a minimum opening through a high select (+16.16 %).

15.3 A good answer cites the dispatch day: at 3 % the plant was out of band 61.7 min and discarded 280 kW h, at zero 116.3 min out of band, failing the tracking clause, and at 8 % 17.6 min with 509 kW h, failing the energy clause. The resting position is a purchase of responsiveness with fuel, about 3.8 kW h per minute of tracking recovered, and zero removes the fast path's authority in the direction that adds power. The trigeneration plant's 0.5 MW electric reserve plays the same role. On a new plant, price the slow actuator's cost and the cost of time out of band, sweep the resting position on the plant's real disturbance pattern, and pick the cheapest position that meets the service requirement.

15.5 A good answer lists tests such as: sweep the moment of a setpoint change and the seed, which caught the coater's override that never let go (10 of 81 and 22 of 82 runs resting under 31 m/min) and the deep cut that settled near 21.5 m/min in 3 of 8 seeds. Move each limit away from its default and check that every clamp follows it, which caught the CO_2 cycle's constant 618 °C clamp. Drive every setting to the end of its range and check that every hard limit has a protection, which caught the extruder's dryer air at 175.0 °C. Start from settings changed before Run, which caught the extruder starting past its torque limit.

Compare the score with the logged states on runs that end in a violation, which caught the peg's scorer. Read the code for what each controller reads.

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 chiller A chiller driven by hot water instead of electricity, cheap to run but slow to change its cooling. *Chapter 9*.

active constraint A limit that the best operating point presses against, so the limited variable must be controlled at it. *Chapter 1*.

anti-surge valve A valve that returns gas from a compressor's outlet to its inlet, keeping enough flow through the machine to avoid surge. *Chapter 2*.

anti-windup Logic that stops a controller's integral term from growing while its output is not the one the actuator receives. *Chapter 1*.

back-calculation An anti-windup method that pulls a controller's integral toward the output actually applied after limits and selectors. *Chapter 1*.

back-off The margin between a setpoint and the limit it guards, set by the largest excursion the control leaves. *Chapter 1*.

bumpless transfer A switch of mode or of the controller in charge that does not jolt the actuator. *Chapter 1*.

capital controls Rules that limit how much money may leave or enter a country, used here as a fast override in a currency defense. *Chapter 14*.

cascade control Two loops in series in which an outer controller sets the setpoint of an inner, faster loop. *Chapter 4*.

cash runway How long a company can keep spending before its cash falls to a floor, in months. *Chapter 8*.

central bank The public bank that issues a country's money, sets its policy interest rate and holds its foreign reserves. *Chapter 14*.

cold store A tank of chilled water kept to supply cooling quickly, which empties as it is used. *Chapter 1*.

cold-gas thruster A jet that pushes by releasing stored compressed gas through a nozzle, without burning anything. *Chapter 12*.

combined heat and power Making electricity and useful heat from one fuel at one site, abbreviated CHP. *Chapter 9*.

compressor map A chart of a compressor's pressure ratio against its flow at several speeds, with the surge line as its boundary. *Chapter 3*.

constraint control Control whose main task is to keep variables inside limits, often by running the plant as close to them as is safe. *Chapter 1*.

control allocation Sharing one or more demands among several actuators by a rule or an optimization that respects each actuator's limits. *Chapter 1*.

corrected quantities Engine speeds and flows scaled by the inlet temperature and pressure, so one map holds at any altitude and on any day. *Chapter 3*.

critical point The temperature and pressure beyond which a fluid is neither liquid nor gas, near which its properties change sharply. *Chapter 10*.

currency board An arrangement that fixes an exchange rate by rule and buys or sells foreign money at that rate on demand. *Chapter 14*.

currency peg A promise by a central bank to hold its exchange rate at a fixed value, or inside a narrow band around it. *Chapter 14*.

decision calendar The schedule on which an economic controller may change its decisions, such as a weekly or a quarterly meeting. *Chapter 8*.

depleting actuator An actuator that draws on a finite stock, such as stored cooling, gas, cash or reserves, so it can act only for a while. *Chapter 1*.

die pressure The pressure of the dough just before the holes of an extruder's die, which rises with a stiffer or wetter dough. *Chapter 5*.

distributed control system A plant control system of many controllers and stations on a network under a supervisory level, abbreviated DCS. *Chapter 1*.

drying front The place along a dryer where a coating loses the last of its free solvent and its vapor release peaks. *Chapter 6*.

earth pressure balance shield A tunneling machine that holds up the ground ahead with its own excavated soil, packed in a closed chamber and let out by a screw. *Chapter 11*.

economic model predictive control Predictive control whose cost is money or another economic measure instead of the distance from setpoints. *Chapter 8*.

electric chiller A chiller driven by an electric compressor, which can change its cooling duty quickly. *Chapter 9*.

exhaust gas temperature The temperature of the gas

leaving a turbine, read by thermocouples and limited to protect the hot parts, abbreviated EGT. *Chapter 3*.

external reset Feeding a controller's integral term from the signal actually applied downstream, so a controller not selected does not wind up. *Chapter 1*.

fan car A car whose fans suck it toward the road, raising the grip its tires can give. *Chapter 12*.

feedforward control Control action computed from a measured disturbance or a known demand before it has moved the controlled variable. *Chapter 2*.

flameout The flame in an engine's combustor going out, which a minimum limit on the fuel flow prevents. *Chapter 3*.

foreign exchange reserves The foreign money and assets a central bank holds and can sell to support its own currency. *Chapter 14*.

friction cone The set of forces a foot can apply without sliding: the sideways part may be at most the grip times the push. *Chapter 13*.

full-authority digital engine control An engine computer that sets the fuel and other engine actuators on its own, abbreviated FADEC. *Chapter 3*.

gain scheduling Changing a controller's settings with the operating point, so the loop behaves alike where the plant's gain differs. *Chapter 3*.

grain loss Grain thrown out of the back of a combine with the straw and chaff instead of reaching the tank. *Chapter 4*.

Greitzer parameter A number comparing the gas stored in a compressor's plenum with the inertia of the gas in its duct, which decides surge or stall. *Chapter 2*.

ground reaction force The push of the ground on a foot, which balances the body's weight and moves it. *Chapter 13*.

high selector A function that passes the largest of several signals, so each of its inputs acts as a floor. *Chapter 1*.

horizontal directional drilling A trenchless method that drills a steerable path under an obstacle such as a river and pulls a pipe back through it. *Chapter 7*.

hydrofracture Splitting of the ground by drilling mud whose pressure exceeds what the soil can hold, which can send mud to the surface. *Chapter 7*.

hydrostatic drive A drive in which an engine-driven pump sends oil to a motor, so the speed can be varied smoothly. *Chapter 4*.

inertial measurement unit A sensor of rotation rates and accelerations from which a body's tilt and motion are computed, abbreviated IMU. *Chapter 13*.

integrated absolute error The area between the controlled variable and its setpoint over a run, abbreviated IAE, a common score of a loop. *Chapter 1*.

integrator windup Growth of a controller's integral term while its output has no effect on the actuator, which later delays or overshoots its action. *Chapter 1*.

Kalman filter An estimator that blends a model's prediction with noisy measurements, weighting each by its expected error. *Chapter 1*.

limit cycle A steady oscillation that a nonlinear system settles into on its own, such as the repeating flow reversal of surge. *Chapter 2*.

load following Changing a power plant's output to follow a demand that rises and falls. *Chapter 10*.

low selector A function that passes the smallest of several signals, so each of its inputs acts as a cap. *Chapter 1*.

lower explosive limit The smallest concentration of a vapor in air that can burn, abbreviated LEL, of which a dryer keeps only a fraction. *Chapter 6*.

mid-ranging control A structure in which a fast actuator holds the variable and a slow one moves until the fast one is back near mid-range. *Chapter 1*.

model predictive control Control that uses a model to predict the plant over a horizon and plans moves that respect limits, abbreviated MPC. *Chapter 1*.

moving horizon estimation Estimating a plant's states by fitting a model to a window of recent measurements, abbreviated MHE. *Chapter 1*.

N-methyl-2-pyrrolidone The solvent of the cathode paste in battery electrode coating, abbreviated NMP, whose vapor can burn. *Chapter 6*.

override control A structure in which a controller guarding a limit takes the actuator from the normal controller when that limit comes near. *Chapter 1*.

peel test A laboratory test that pulls a dried coating from its foil to measure how strongly it holds. *Chapter 6*.

PI controller A controller with proportional and integral action, the form of nearly every loop in this volume. *Chapter 1*.

plenum The volume of gas between a compressor and its throttle, whose pressure pushes back on the flow. *Chapter 2*.

policy rate The interest rate a central bank sets to steer the cost of borrowing in its currency. *Chapter 14*.

programmable logic controller A rugged industrial computer that reads inputs and drives outputs by a stored program, abbreviated PLC. *Chapter 1*.

quadratic programming Optimization with a

quadratic cost and linear limits, abbreviated QP, solved fast enough to run inside a control loop. *Chapter 13*.

rate limit A cap on how fast a signal or an actuator may change. *Chapter 1*.

rate of penetration The speed at which a drill bit advances along its path, abbreviated ROP. *Chapter 7*.

recuperator A heat exchanger that passes heat from a turbine's exhaust to the gas leaving the compressor. *Chapter 10*.

recycle line A pipe that returns gas from a compressor's outlet to its inlet through the anti-surge valve. *Chapter 2*.

resting position The opening at which a valve-position controller holds the fast actuator, leaving it room to move both ways. *Chapter 1*.

rotating stall A pattern of stalled blade passages that travels around a compressor and cuts its flow and pressure rise. *Chapter 2*.

saturation The state in which an actuator sits at a limit, so a larger command moves the process no further. *Chapter 1*.

selector A function that passes one of several signals by rule, such as the smallest or the largest. *Chapter 1*.

settlement trough The bowl-shaped sinking of the ground surface above a tunnel. *Chapter 11*.

soft constraint A limit that a controller may cross at a penalty, instead of one that must never be crossed. *Chapter 2*.

soft sensor An estimate of a hard-to-measure quantity computed from measurements that are easy to take. *Chapter 5*.

specific mechanical energy The work of an extruder's screws per kilogram of product, abbreviated SME, a measure of how hard the dough is worked. *Chapter 5*.

speculative attack A sudden sale of a currency by investors who bet that its fixed rate will be abandoned. *Chapter 14*.

split point The controller output at which a split-range controller hands over from its first actuator to its second. *Chapter 1*.

split-range control One controller output divided into ranges, each moving a different actuator in a fixed order. *Chapter 1*.

supercritical carbon dioxide cycle A closed power cycle that compresses carbon dioxide near its critical point, where it is dense and needs little work. *Chapter 10*.

supervisory control A layer that writes setpoints to the controllers below it instead of moving actuators itself. *Chapter 1*.

surge A violent cycle in which the flow through a compressor collapses and reverses, again and again. *Chapter 2*.

surge margin How far a compressor's operating point lies from its surge line, usually in percent. *Chapter 2*.

sustainable growth rate The yearly sales growth a company can finance from its own earnings without new equity or a change in its borrowing. *Chapter 8*.

tracking time The time constant with which back-calculation pulls a controller's integral toward the output actually applied. *Chapter 1*.

trigeneration Making electricity, heat and cooling from one fuel, the cooling from heat that would otherwise be wasted. *Chapter 9*.

trilemma The finding that a country cannot at once fix its exchange rate, allow free movement of capital and run its own interest-rate policy. *Chapter 14*.

turbine bypass A valve that sends gas around a turbine, the fastest but least efficient way to change a cycle's power. *Chapter 10*.

turbofan engine A jet engine whose core turns a large fan, so most of its air bypasses the core. *Chapter 3*.

uninterruptible power supply A battery system that carries critical loads through a loss of power, abbreviated UPS. *Chapter 9*.

valve-position control Slow integral action that moves a second actuator until a fast one returns to its resting position, abbreviated VPC. *Chapter 1*.

variable-speed drive An electronic drive that sets a motor's speed by the frequency of the power it supplies. *Chapter 1*.

velocity form A digital PID that computes only the change of its output at each execution. *Chapter 1*.

wheel slip control Control of the difference between wheel speed and vehicle speed, so a tire stays near its peak grip. *Chapter 12*.

working capital Money tied up in stock and unpaid customer invoices, less what a company owes its suppliers. *Chapter 8*.

Index of Control Topics

active constraint, 4, 95, 144, 174
  moving with the conditions, 164, 177
active surge control, 27, 46, 378
actuator sizing, 39, 373
advanced regulatory control, 377
AI agents
  engineering practice, 379
  open problems, 383
  reproduction, 381
  truth audit, 380
  verification record, 380
  while studying, 18
anchoring and adjustment, 199

back-calculation, 33, 60, 86, 145, 253, 328, 357
back-off, 4, 117, 147, 174, 335, 375
bumpless transfer, 17, 146, 199, 253, 307, 328, 357

capacity limit, 258
cascade control, 87, 117, 308
  inner-loop speed, 253
clamp
  protective limit, 381
code verification, 149
conditional integration, 88
constrained optimization, 174
constraint adaptation, 146
constraint control, 1, 104, 136, 164, 351
  open problems, 381
  pushing, 153
control allocation, 11, 227, 297, 320, 378
  fixed priority, 171, 305
  friction cones, 320
  metered remainder, 236, 308
  priority, 375
  weighted least squares, 329
control structure
  selection, 376
control-induced oscillation, 40

data export, 18
dead reckoning, 173
decision calendar, 198, 356

depleting actuator, 12, 171, 188, 226, 297, 346, 376
  reserve, 308
disturbance model, 36, 89, 119, 148, 255, 331
double integrator, 331
duty-cycle actuator, 165

external reset, 2, 7, 34, 60, 87, 117, 146, 174, 226, 281, 374

fault detection, 282
feedforward
  from a measured mass, 305
  measured load, 226, 281
  seasonal, 199
  setpoint schedule, 282
feedforward control, 146
  as a moving limit, 87
  model error, 116
  static, 41
friction cone, 325

gain margin, 86
gain scheduling, 54, 59
Greitzer model, 27
Greitzer parameter, 27, 32

Hopf bifurcation, 32

infeasible demand, 311
infeasible setpoint, 178, 280, 361
inferential control, 104, 106, 152
integral action, 333
integral of absolute error, 36
integrated absolute error, 227
integrating process, 272
  integrating-plus-dead-time model, 286
integrator tracking, 60
integrator windup, 34, 60, 199, 226
inventory control, 246, 271

Kalman filter, 17, 119
  bias estimation, 173
  extended, 36

limit cycle, 24, 32, 151, 300, 311
limit protection, 51, 70, 378

linearization, 31

measurement
  late and revised, 280
measurement bias, 92, 99, 146, 171
measurement filter, 40, 44
measurement noise, 40
mid-ranging control, 9, 215, 243, 271
model mismatch, 232
model predictive control, 34, 61, 88, 118, 147, 227, 254, 266, 282, 330
  against selectors, 377
  economic, 189, 200, 358
  in the apps, 17
moving horizon estimation, 17

offset, 333
override control, 2, 5, 24, 51, 76, 86, 104, 116, 136, 164, 188, 291, 346
  capital controls, 359, 368
  handle choice, 146
  pressure, 174, 182
  protective, 13, 25, 45, 291
  runway, 200, 209

pattern search, 147, 174
performance monitoring, 260
perturb-and-observe search, 88
phase margin, 86
PID control
  anti-windup, 3, 7, 16, 383
  velocity form, 16, 33, 86, 116, 145, 199, 252, 307, 328, 357
PID tuning
  SIMC rule, 286
positive feedback, 305, 352
  basin of attraction, 363
preview control, 79, 283
process gain
  varying, 85, 279
pulse-width pulse-frequency modulation, 299
pure pursuit, 173

quadratic programming, 329, 341
quantization, 198, 308, 356

rate limit, 8, 34, 60
rate of penetration, 165
ratio control, 116, 189
  feedback trim, 117, 145, 199
receding horizon, 330
reference governor, 378
residence-time distribution
  tanks in series, 114

safety filter, 341, 378
sampled analyzer, 117, 142, 353
saturation, 8, 39, 123
  capacity limit, 232
scenarios, 17
  completion rules, 18
scorecard, 18
selector, 5, 41, 51, 86
  min/max, 52, 146, 174, 226, 281
  realized as a clamp, 117
  takeover by rate, 63
sequencer, 173, 281, 307
setpoint from the plant state, 200
simulation engine, 15
simulator framework, 15
Smith predictor, 145
soft constraint, 35, 89, 120, 148, 256, 358
soft sensor, 115
  bias update, 119
  soil classification, 173
split-range control, 8, 110, 200
stability
  trace and determinant, 31
state estimation, 359, 378
steady-state gain, 115
step test, 118, 250
successive linearization, 200
supervisory control, 13, 87, 117, 174, 254, 358
system dynamics, 191

throughput control, 76, 147
throughput supervisor, 381
time-scale separation, 244
transport delay, 81, 140

underactuated control, 364

valve characteristic, 249
valve-position control, 9, 215, 243
  reserve setpoint, 226
  resting position, 260, 375
  unsaturated demand, 254

wheel slip control, 307

Index of Applications

absorption chiller, 218, 236
    capacity lag, 221
acceleration schedule, 59
accounts receivable, 194
acquisition, 204
amylose, 122
annular pressure, 170, 181
anti-surge valve, 27
    stroke time, 39, 45
at-rest earth pressure, 274

Bass diffusion model, 190
battery-electric drive, 167, 182
beacon, 167
belt dryer, 112
belt scale, 278
bentonite, 279
beta gauge, 139
binder migration, 139
blow-out, 275
blowdown tank, 303, 313
    adiabatic cooling, 305
bodie, 65
Brayton cycle, 243
breakfast cereal, 105
bulk density, 111, 131

capital controls, 352
capital mobility, 352
capture point, 323
cash conversion cycle, 191
cash flow, 188
cash runway, 195, 208
cavity expansion limit, 165, 167
central bank, 346
centrifugal compressor, 24
certification
    five-second rule, 54
chamber stiffness, 276
choked nozzle, 57
cleaning shoe, 82
cold-gas thruster, 297, 299
combine harvester, 76, 77
combined cycle, 217
combined heat and power, 216, 238
competitor launch, 204
compressor
    near-critical inlet, 246
compressor characteristic, 25
compressor inlet temperature, 252, 263
compressor map, 56
compressor stall, 57
compressor surge, 24, 57
contact estimation, 323
contagion, 352
cook index, 111
cooled reactor, 4
cooling tower, 222, 252
corrected quantities, 54
cost of transport, 333
Coulomb friction, 325
credit line, 194, 209
critical point, 244, 264
critical solvent content, 141
crop moisture, 81
currency board, 351, 366
currency crisis
    first-generation model, 352
    second-generation model, 358
currency peg, 346
cuttings bed, 169

data center, 215
data hall, 222
deceleration schedule, 59
deep surge, 32
die deposit, 110
die pressure, 110
dilution, 204
downforce, 299
drain law, 354
drying
    constant-rate period, 141
    falling-rate period, 141
drying front, 144

earth pressure balance shield, 271, 273
electric chiller, 218, 236
electrode coater, 136
electrode dryer, 139
electronic stability control, 313
engine control, 51, 69

engine deterioration, 58
engine lugging, 83
equity round, 194
exhaust bypass, 218
  hot-water loop, 226
exhaust fan, 151, 157
exhaust gas temperature, 57, 69
exit climb, 178
expansion index, 107, 111
expected depreciation, 353
extruder, twin-screw, 104, 129

face stability, 273
fan car, 297, 299
fan-speed governor, 59
feeder house, 83
flameout, 64
flat rating, 64
flying car, 300
foam injection ratio, 274
food extrusion, 104, 105
foreign exchange intervention, 349, 351
foreign exchange reserves, 347, 365
forward crop sensor, 80, 83
friction estimation, 341
full-authority digital engine control, 53

gas turbine, 215
  ambient derating, 218
grain damage, 82
grain loss, 77, 98
grid import limit, 223
ground reaction force, 321, 323
ground-effect suction, 302
growth trap, 203

heat-recovery unit, 218
Helmholtz frequency, 33
horizontal directional drilling, 164, 166
hung stall, 65
hydrofracture, 166
hydrostatic drive, 78, 98

inadvertent return, 165, 183
inertial measurement unit, 339
insolvency, 202
interest-rate defense, 357

lean blow-out, 57
learning curve, 190
leg odometry, 325, 339
legged robot, 320

lethality, 130
lithium-ion battery manufacturing, 136
load curtailment, 226
load following, 245
load rejection, 257
loan covenant, 194
lodged crop, 81
loss curve, 85
loss monitor, 83, 97
loss-in-weight feeder, 107
lower explosive limit, 138, 158, 379

market share
  logit, 191
migration index, 142
minimum impulse bit, 305
Moore-Greitzer characteristic, 29
mud ratio, 169

N-methyl-2-pyrrolidone (NMP), 137, 158
nitrogen
  real-gas law, 303

output gap, 353
over-excavation, 282
overnight rate, 351

pecking order theory, 191
peel test, 139
pilot bore, 166
plenum, 27, 57
power-law fluid, 109
preconditioner, 107
pressure regulator, 303
production planning, 199
proprioceptive actuator, 339
push recovery, 326

quadruped robot, 320

Raibert foot placement, 322
rate law, 169
recession, 206
recuperator, 246
  fouling, 259
recycle line, 27
research and development pipeline, 193
reserve adequacy, 354, 367
residual solvent, 152
ring build, 279
river crossing, 179
rod change, 171
rotating stall, 26
rotor plugging, 82

Rule of 40, 195

screw conveyor, 274
  plug, 277
screw wear, 110
seasonal demand, 193
separation loss, 82
settlement survey, 278
settlement trough, 273, 274
shadow exchange rate, 352
shore power, 177
skirt seal, 302
slant-face bit, 167
slot die, 139
slurry shield, 273
slurry solids, 152
soil conditioning, 274, 292
solenoid valve, 313
solvent oven, 138
specific energy, 166
specific impulse, 299
specific mechanical energy, 105, 129
speculative attack, 347
spewing, 277
spool dynamics, 57
stall recognition, 60
startup company, 188
starve feeding, 108
static stability margin, 322
sterilization, 351
supercritical carbon dioxide, 263
supercritical carbon dioxide cycle, 243, 245
surge control line, 27
surge detection, 30, 46
surge line, 29, 56
surge margin, 36, 56, 71
sustainable growth rate, 191
  on trailing books, 196

terrain map, 336
thermal storage
  cold store, 218
throttle chop, 64
throttle slam, 62
timeliness loss, 77, 78
tipping limit, 332
tire
  Burckhardt curve, 301
  friction estimation, 315
  relaxation length, 301
toe stub, 336
trailing twelve months, 194
trenchless construction, 164
trigeneration, 215
trilemma, 347
trot, 330
tunnel face pressure, 272, 290
tunneling, 271
turbine bypass, 244
turbofan engine, 51
  two-spool high-bypass, 54

uncovered interest parity, 352
uninterruptible power supply, 223
utility drill, 164

vapor monitor, 139, 157
venture capital, 191, 210
viscous dissipation, 105
volume loss, 273

walkover locator, 167
work in process, 193
working capital, 197

yield sensor, 83

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.