INTERACTIVE CONTROL ENGINEERING

Control of Energy Conversion and Storage

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 17

Control of Energy Conversion and Storage

Fast regulation and slow dispatch on a changing grid

John D. Hedengren
Editor

Advanced Process Solutions, LLC

Interactive Control Engineering: Principles and Applications
Volume 17: Control of Energy Conversion and Storage
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/v17-energy/

How to cite this volume. John D. Hedengren, Ed., *Control of Energy Conversion and Storage*, vol. 17 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 converting and storing energy on a grid whose supply and demand change by the second and by the day. Chapters 2 to 6 cover conversion: engines, fuel cells, wind, waves and concentrated solar heat. Chapters 7 to 9 cover the grid edge: an islanded microgrid, a distribution feeder with rooftop solar and a fast charger. Chapters 10 to 13 cover storage and dispatch, where the state of charge is an inventory and the decision is economic.

Each chapter names the technology volume that teaches its main technique. Volume 18 covers nuclear plants.

John D. Hedengren
2026

Contents

Preface v

Notation x

1 Introduction: From Inertia to Dispatch 1

1.1 The Question This Volume Answers . . . 1
1.2 A Short History . . . 2
1.3 Notation and Conventions . . . 4
1.4 From Inertia to Dispatch . . . 4
1.5 Instruments and Final Elements in This Volume . . . 12
1.6 The Simulator Framework . . . 13
1.7 Working With AI Agents While Studying . . . 17
1.8 Chapter Roadmap . . . 18
1.9 Exercises . . . 20

2 Gas Engine Cogeneration Where Power Settles in Seconds and Heat in Minutes 23

2.1 Introduction . . . 24
2.2 Recovered Heat, the Grid and the Fuel Valve . . . 25
2.3 Three Engine Modules, a Buffer and a Data Hall . . . 26
2.4 The Control Layers as Implemented . . . 31
2.5 Matched Experiments, a Capacity Limit and the Challenge . . . 35
2.6 Technology in Practice . . . 41
2.7 Limitations . . . 43
2.8 Conclusion . . . 44
2.9 Exercises . . . 46

3 Solid Oxide Fuel Cell Current That Waits for Its Fuel 49

3.1 Introduction . . . 50
3.2 From Ceramic Electrolytes to Fuel-Limited Load Following . . . 51
3.3 An Islanded Fuel Cell Plant and Its Data Hall . . . 52
3.4 Admitting Current: Governor, Estimator and Loops . . . 59
3.5 Experiments on One Plant . . . 63
3.6 Technology in Practice . . . 72
3.7 Limitations . . . 74
3.8 Conclusion . . . 75
3.9 Exercises . . . 77

4 Torque-to-Pitch Handover and Gain Scheduling in a Wind Turbine Simulator 80

4.1 Introduction . . . 81
4.2 Background . . . 82
4.3 Turbine Model . . . 83
4.4 Pitch Sensitivity and the Gain Schedule . . . 88
4.5 Controller Design . . . 89
4.6 Results . . . 94
4.7 Wind Measurement in the Region 2 Setpoint . . . 100
4.8 Technology in Practice . . . 102
4.9 Limitations . . . 104
4.10 Conclusion . . . 105

4.11 Exercises . . . 107

5 Energy-Maximizing Control of a Heaving Buoy on a Fading Wave Forecast 111
5.1 Introduction . . . 112
5.2 Point Absorbers, Phase Control and Short Forecasts . . . 113
5.3 The Buoy, the Sea and the Machine . . . 114
5.4 Why the Best Force Looks Ahead . . . 118
5.5 The Force Laws as Implemented . . . 119
5.6 Ten Cases on One Buoy . . . 121
5.7 Technology in Practice . . . 127
5.8 Limitations . . . 128
5.9 Conclusion . . . 129
5.10 Exercises . . . 131

6 A Parabolic-Trough Simulator for Choosing a Setpoint Rather Than Defending One 134
6.1 Introduction . . . 135
6.2 From SEGS to a Weather Forecast . . . 136
6.3 Process Model . . . 137
6.4 Open-Loop Character . . . 143
6.5 Regulatory Control . . . 143
6.6 The Supervisory Layer . . . 147
6.7 Scenarios and Measured Results . . . 148
6.8 Technology in Practice . . . 155
6.9 Limitations . . . 157
6.10 Conclusion . . . 158
6.11 Exercises . . . 159

7 Forming the Grid From a Battery When a Campus Microgrid Islands 163
7.1 Introduction . . . 164
7.2 How Small Power Systems Learned to Run Alone . . . 165
7.3 The Campus Behind One Breaker . . . 166
7.4 Five Relations Behind the Cases . . . 172
7.5 The Microgrid Controller, Layer by Layer . . . 174
7.6 Ten Cases and a Two-Hour Island . . . 176
7.7 An Optional One-Hour Energy Dispatch . . . 181
7.8 Technology in Practice . . . 183
7.9 Limitations . . . 185
7.10 Conclusion . . . 186
7.11 Exercises . . . 187

8 Hunting, Reverse Flow and an Unmeasured Customer on a Solar Distribution Feeder 191
8.1 Introduction . . . 192
8.2 Voltage Regulation on Radial Feeders . . . 193
8.3 The Feeder as Modeled . . . 194
8.4 Four Mechanisms Behind the Cases . . . 198
8.5 The Controllers as Implemented . . . 199
8.6 Nine Cases and a Challenge . . . 203
8.7 Technology in Practice . . . 210
8.8 Limitations . . . 212
8.9 Conclusion . . . 212
8.10 Exercises . . . 214

9 Electric Vehicle Fast Charge Simulator for Selector Control of an Unmeasured Constraint 218
9.1 Introduction . . . 219
9.2 From Constant Current to a Charge Current Limit Map . . . 220
9.3 Pack and Cell Model . . . 221
9.4 The Constraint Surface . . . 227
9.5 The Charge Controller . . . 228
9.6 Scenarios and Measured Results . . . 232
9.7 Technology in Practice . . . 239
9.8 Limitations . . . 241
9.9 Discussion . . . 242
9.10 Conclusion . . . 243
9.11 Exercises . . . 244

10 A Vanadium Flow Battery Simulator for Inventory-Constrained Dispatch on a Duck Curve 248
10.1 Introduction . . . 249
10.2 From the NASA Redox Program to a Feeder . . . 250
10.3 Process Model . . . 251
10.4 Open-Loop Character . . . 258
10.5 Regulatory Control . . . 259
10.6 The Dispatch Layer . . . 261
10.7 Scenarios and Measured Results . . . 262
10.8 Technology in Practice . . . 266
10.9 Limitations . . . 269
10.10 Conclusion . . . 270
10.11 Exercises . . . 271

11 A Salt-Cavern Air Storage Simulator for Inventory Policy Under a Capacity Commitment 275
11.1 Introduction . . . 276
11.2 From Huntorf to a Capacity Market . . . 277
11.3 Process Model . . . 278
11.4 Open-Loop Character . . . 284
11.5 Regulatory Control . . . 285
11.6 Three Inventory Policies . . . 287
11.7 Scenarios and Measured Results . . . 288
11.8 Technology in Practice . . . 292
11.9 Limitations . . . 295
11.10 Conclusion . . . 295
11.11 Exercises . . . 297

12 A Solar Hydrogen Simulator for Safety Interlocks and the Question of Keeping the Oxygen 300
12.1 Introduction . . . 301
12.2 From Utsira to a Firm Evening . . . 302
12.3 Process Model . . . 303
12.4 Open-Loop Character . . . 309
12.5 Regulatory Control . . . 309
12.6 The Supervisory Machine . . . 311
12.7 Scenarios and Measured Results . . . 312
12.8 Technology in Practice . . . 315
12.9 Limitations . . . 318

12.10 Conclusion . . . 318
12.11 Exercises . . . 320

13 Supervisory Economic Dispatch of Nuclear Heat, Storage and Hydrogen Production 323
13.1 Introduction . . . 324
13.2 Nuclear Cogeneration and Flexible Operation . . . 325
13.3 Plant Model . . . 327
13.4 What an Hour Is Worth . . . 333
13.5 Two Control Layers . . . 334
13.6 Results . . . 337
13.7 Technology in Practice . . . 343
13.8 Limitations . . . 346
13.9 Discussion . . . 346
13.10 Conclusion . . . 347
13.11 Exercises . . . 349

14 Summary and Outlook 352
14.1 What the Chapters Showed . . . 352
14.2 Choosing a Method . . . 356
14.3 Where Practice Is Heading . . . 356
14.4 Engineering With AI Agents . . . 359
14.5 Gaps in the Control of Energy Conversion and Storage . . . 362
14.6 Exercises . . . 366

Answers to Selected Exercises 367

Glossary 372

Index of Control Topics 375

Index of Applications 378

About the Editor 381

Volumes in This Series 382

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: From Inertia to Dispatch

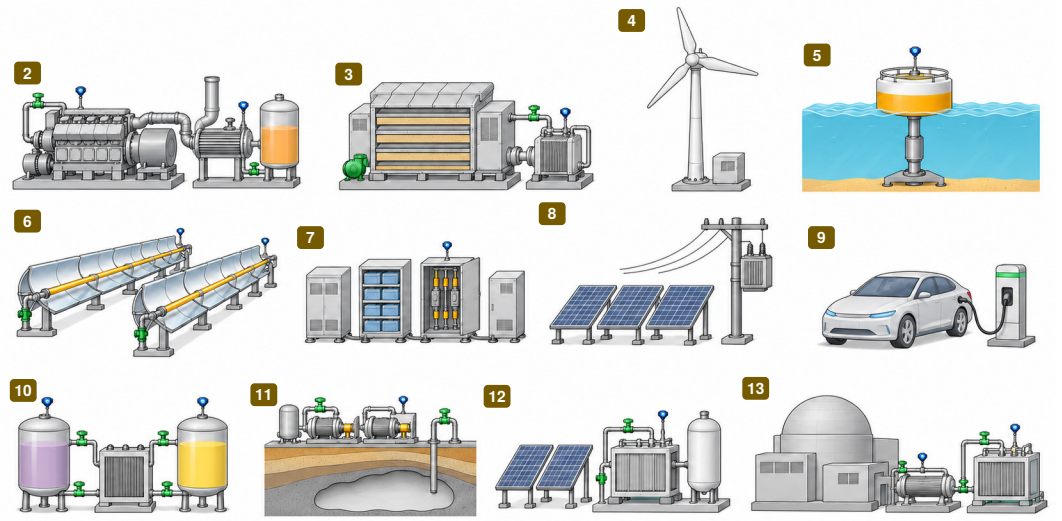

The chapters 2 gas engine cogeneration 3 fuel cell cogeneration 4 wind turbine 5 wave buoy 6 parabolic trough field 7 campus microgrid 8 solar feeder 9 fast charger 10 vanadium flow battery 11 salt-cavern air store 12 solar hydrogen plant 13 nuclear heat and hydrogen

Notice that the top row converts gas, wind and waves, the middle row collects sunlight and runs the networks and a charger, and the bottom row stores energy in electrolyte, air and hydrogen or turns reactor heat into hydrogen. Illustration, not to scale.

IN PLAIN TERMS **Fast Loops and Slow Plans for Energy**

The main idea. Engines, fuel cells, wind turbines, solar fields and batteries turn energy from one form into another or keep it for later. Fast controls act on each one within seconds, and slow plans look hours or days ahead.

The challenge. Electricity must match demand at every moment. Fuel, heat, sun and wind change far more slowly than the electronics that deliver the power. A store can cover the gap only while it is neither full nor empty, and the limit that matters is often not measured.

Why it matters. Good control keeps the lights on while it spends fuel and stored energy at the best times.

THREE WAYS TO READ THIS CHAPTER

Explore	7 min	This page, Section 1.2 with its timeline and two History boxes, Figure 1.2, the Did you know? boxes in Sections 1.4.2, 1.4.4 and 1.4.5, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, working each number of the example E-1, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.7 with Table 1.2, Section 1.5 with Figures 1.6 and 1.8, and Section 1.8.

1.1 The Question This Volume Answers

An energy system is controlled on clocks that span nine decades. A grid-forming inverter answers a load step within milliseconds. A gas engine settles its power in seconds and its recovered heat in about a minute. A solar field passes its fluid from inlet to outlet in minutes, a flow battery empties in hours, and a salt cavern is managed against a reserve that changes with the season. Every plant in this volume converts or stores energy, and every one of them is controlled in layers: fast loops that hold

a frequency, a voltage, a temperature or a current, and slow layers that decide how much to make, store or sell over the next hour, day or season. This volume answers one question: how is each layer designed, and how does a fast regulating layer share a plant with a slow dispatching one?

Seven control problems define the industry, and every chapter meets several of them.

- *Supply and demand must balance at every instant.* Electricity is not stored in the wires, so any imbalance changes the speed of rotating machines, the frequency or a bus voltage within seconds.
- *Fast actuators stand in front of slow supplies.* A power converter moves in milliseconds, while the fuel, heat or wind behind it moves in seconds to hours, and a store must cover the difference.
- *Stores are inventories with two limits.* A battery, a tank of salt, a cavern of air or a vessel of hydrogen integrates what goes in and out, and it is useless both full and empty.
- *One input makes two products.* A fuel valve makes electricity and heat at once, and the two answer on different clocks.
- *The value of energy moves with time.* Prices, loads, sun, wind and waves change through the day, so the best action now depends on a forecast of what comes later.
- *The binding limit is often unmeasured.* An anode potential, the voltage at the far end of a feeder or the fuel arriving at a cell is inferred, not read.
- *Protection outranks economics.* An interlock, an overtemperature override or a shed relay must win over any dispatch, however profitable.

Twelve chapters answer the question on twelve plants, each with an app that the reader runs in a browser. Chapters 2 and 3 make power and heat for data centers from natural gas, in reciprocating engines and in a solid oxide fuel cell. Chapters 4, 5 and 6 take energy from wind, waves and sunlight. Chapters 7 and 8 control the networks that carry it: a campus microgrid that islands from the utility and a distribution feeder with rooftop solar. Chapter 9 charges a vehicle pack. Chapters 10, 11 and 12 store energy in an electrolyte, in compressed air and in hydrogen, and Chapter 13 dispatches the constant heat of a nuclear reactor between a turbine, a salt store and an electrolyzer.

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

- Place each loop of the plant on the time-scale hierarchy and check that each layer is slow enough to treat the one below as settled.
- Compute the first rate of change of frequency after a load step and the settled deviation that droop leaves, and say which layer removes it.
- Treat a store as an inventory: its duration, its losses, the loop that keeps it in the middle and the limits that taper its power.
- Follow one product of a cogeneration plant and size the store that absorbs the mismatch of the other.
- Say what a forecast and a planning horizon are worth on a given day, and price the energy left in a store at the end of a plan.

This chapter develops each of these on one small example and names the chapter that uses it. PID forms and tuning are taken as known from Volume 2, feedforward and cascade from Volume 3, selectors, overrides and actuators that deplete from Volume 4, estimation and Kalman filtering from Volume 5, inventory control from Volume 6, gain scheduling from Volume 7, state machines and switching from Volume 8, and predictive and economic control from Volume 9. Volume 18 treats nuclear reactor physics.

1.2 A Short History

The milestones of Figure 1.1 follow the plants of this volume, from the first fuel cell to a definition of fast charging.

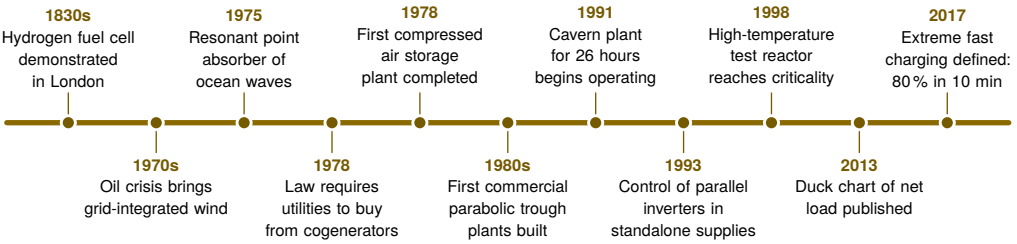

Figure 1.1 Milestones of energy conversion and storage and their control. The dates are from [1] (1830s), [2] (1970s), [3] (1975), [4] (1978, the law), [5] (1978, the air store, and 1991), [6] (1980s), [7] (1993), [8] (1998), [9] (2013) and [10] (2017). Notice that most of the converters came before the controls that join them to a grid. Events are evenly spaced, not to a time scale.

> ⧖ HISTORY **1830s to 1960s: A Fuel Cell Waits for a Use**
>
> A hydrogen fuel cell was demonstrated in London in the 1830s, and the technology then remained without a practical application for over a hundred years. Fuel cells returned to the laboratory in the 1950s, when the United States space program needed new power systems with low or no emissions [1]. In the 1960s the space program demonstrated a 1 kW polymer electrolyte cell in Gemini and a 1.5 kW alkaline cell in Apollo, both on pure reactant gases [11]. Today Chapter 3 runs a solid oxide cell on natural gas, and Chapter 12 runs a membrane cell on hydrogen it made itself.

The energy crises of the 1970s brought the other sources. The oil crisis of that decade rekindled interest in renewable energy and led to the commercial adoption of wind energy systems on the grid [2], the machines of Chapter 4. In 1975 a wave absorber much smaller than a wavelength was named a point absorber, a damped oscillator that passes energy from the wave to a load [3], the device of Chapter 5.

A United States law of November 1978 required electric utilities to offer to buy electric energy from qualifying cogeneration facilities and to sell to them [4], the kind of plant Chapter 2 controls. The first compressed air energy storage plant was completed near Bremen in the same year, and a plant in Alabama designed for 26 hours of generation at full power from a single salt cavern has run since 1991 [5], the store of Chapter 11. The first commercial parabolic trough plants were built in the Mojave Desert in the 1980s [6]. A test reactor built for high-temperature heat reached first criticality in 1998 [8], a step toward the reactor heat for hydrogen of Chapter 13.

In 2013 the California system operator published the duck chart of net load [9], the curve that the flow battery of Chapter 10 serves. In 2017 a gap assessment defined extreme fast charging as recharging up to 80 percent of a battery's capacity in ten minutes or less [10], the goal that Chapter 9 pursues against a limit no instrument reads.

> ⧖ HISTORY **1993 to 2018: Inverters Learn to Hold a Grid**
>
> Droop control for inverters was first proposed in the early 1990s. It trades frequency and voltage against real and reactive power much as a synchronous machine does, so parallel inverters can share a load without communicating [7, 12]. Such controls were first meant for microgrids and small islands [12]. On the grid itself, United States standards long kept inverters from supporting the voltage, until an amendment of May 2014 allowed advanced functions [13]. Chapter 7 forms its island from a battery inverter, and Chapter 8 runs Volt-VAR curves along a feeder.

1.3 Notation and Conventions

Electric power is in megawatts (MW) or kilowatts and energy in megawatt-hours. Where heat and electricity meet, MWth and MWe keep them apart. Frequency f is in hertz with nominal value $f_0 = 60\,\text{Hz}$, and a voltage v is in per unit (pu) of its nominal value. A machine or converter of rating S has an inertia constant H, the stored kinetic energy divided by the rating, in seconds. A droop R is the per-unit fall of frequency that drives a unit from no load to its rating. A store holds at most E megawatt-hours, its state of charge is the fraction SOC, and it charges at power P_c with efficiency η_c and discharges at P_d with efficiency η_d. A cogeneration unit converts fuel power P_f to electricity with efficiency η_e and to recovered heat Q_h with efficiency η_h.

The framework's h is its integration step and its a the controller action, and neither is used for anything else here. Its S_k is a running sum in the PID algorithm and its Q, R and N are the weights and the horizon of a predictive controller, unrelated to the rating, the droop and the heat of this chapter except that N is again a horizon. Loop action follows Section 1.6.2: $a = +1$ is reverse acting, for a process whose measurement rises with the manipulated variable, and $a = -1$ is direct acting. A frequency loop on a generator's setpoint is reverse acting, because more power raises the frequency. A loop that holds a state of charge by discharging is direct acting, because more discharge lowers the charge. A cooling loop on a temperature is direct acting. Three chapters, Chapter 6, Chapter 9 and Chapter 10, write the action flag as γ with the same meaning.

The chapters keep their own symbols, and several collide. H is an inertia constant in Chapter 7, the significant wave height H_s and the forecast horizon H_f in Chapter 5 and an exhaust enthalpy in Chapter 2. R is the governor droop in Chapter 7, the compensator resistance R_{set} in Chapter 8 and the rotor radius in Chapter 4. Q is reactive power in Chapter 8 and Chapter 7, electrolyte flow in Chapter 10, a pack capacity in ampere hours in Chapter 9 and a heat duty in Chapter 13 and Chapter 6. λ is the tip-speed ratio in Chapter 4, the membrane hydration in Chapter 12, the cavern leakage rate in Chapter 11 and a move-suppression weight in Chapter 6 and Chapter 13. φ is the anode potential in Chapter 9, an autoregressive coefficient in Chapter 5 and the idle fuel share φ_0 in Chapter 2, while Chapter 8 writes its Volt-VAR curve $\phi(v)$ and Chapter 13 a pressure ratio ϕ. δ is a machine angle in Chapter 7, the Volt-VAR deadband in Chapter 8, the electrolyte imbalance in Chapter 10 and a degradation state in Chapter 13. S is a rating in Chapter 7 and Chapter 8 and the cavern's state of charge in Chapter 11, and z is the state of charge in Chapter 9 but an overload state in Chapter 7. Chapter 7 also writes its dispatch slot length as h, which is not the framework's step.

1.4 From Inertia to Dispatch

1.4.1 The Flow Scheme

Figure 1.2 places each chapter's plant on the path from primary energy to the user. Gas engines and fuel cells convert natural gas to electricity and heat close to a large user such as a data center [14]. Wind turbines, wave converters and solar fields convert flows that nobody controls, and their output follows the weather. A nuclear reactor makes heat at a constant rate. Networks carry the electricity: a transmission grid, a distribution feeder that ends at houses with their own solar, or a microgrid that can run alone behind one breaker. Stores sit on the network or inside a plant. A flow battery and a compressed-air cavern trade electricity across hours and days, a hydrogen plant across a day

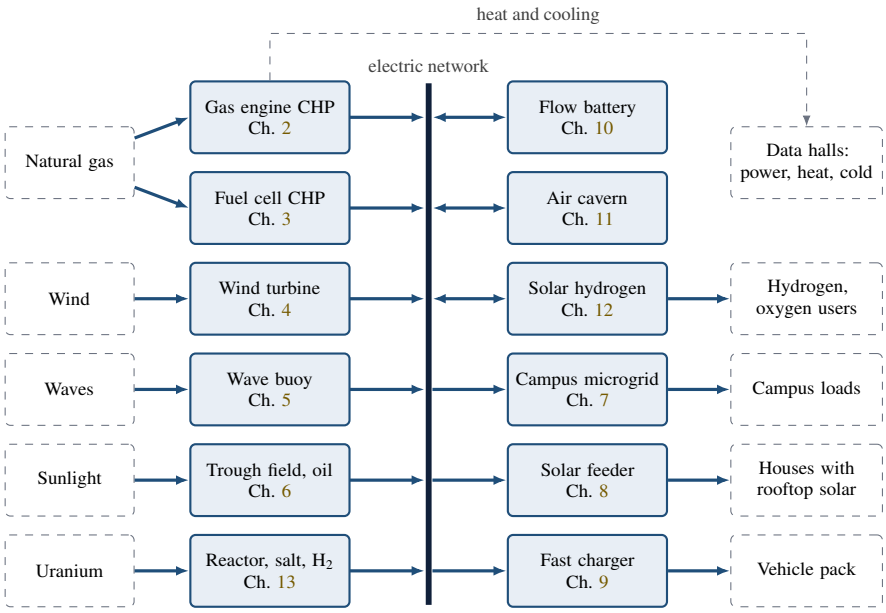

Figure 1.2 The Plants of the Volume on the Path From Primary Energy to the User. Filled Blocks Are the Plants of Chapters 2 to 13, Dashed Blocks Lie Outside the Volume. Converters on the Left Feed the Electric Network (Vertical Bar), Stores Trade With It in Both Directions, and Networks and Chargers on the Right Deliver It. The CHP Plants Also Serve Data Halls With Heat and Cooling (Dashed Arrow), and the Nuclear Plant Also Sells Hydrogen (Not Drawn).

and night, and tanks of hot water or molten salt store heat inside the plants that make it. Electric vehicles are users that are also stores. Planning such systems as a whole, with heat, electricity and fuels coupled, is the subject of multi-energy system studies [15], and the flexibility that storage, dispatchable plant and demand provide is what lets a grid take a large share of variable renewable energy [16, 17].

The techniques the chapters share are developed below on one running example, an island energy system called E-1. A gas engine-generator set rated 6 MW with recovered heat, a battery of 2 MW and 4 MWh, a 3 MW solar array and a hot-water store serve the electricity and heat of a small town. Every number of E-1 is declared here or computed from these declarations by the script kept with this chapter's record. The numbers are chosen for clear arithmetic and are not taken from any one chapter.

1.4.2 The Time-Scale Hierarchy

The rotating machines of a grid store kinetic energy, and that energy is the first buffer against any imbalance. If the mechanical power P_m delivered to the machines differs from the electrical power P_e drawn from them, the frequency changes as Equation (1.1) shows, which is the swing equation of a single machine or of a grid treated as one [18].

$$\frac{2HS}{f_0}\frac{\mathrm{d}f}{\mathrm{d}t} = P_m - P_e, \qquad \mathrm{ROCOF}(0^+) = -\frac{f_0\,\Delta P}{2HS} \tag{1.1}$$

The product HS is the stored kinetic energy. Right after a load step ΔP, before any governor moves, the frequency falls at the rate of change of frequency (ROCOF) on the right. E-1's engine-generator set has $H = 2.5$ s on its 6 MW rating, so it stores 15 MJ, and a 1 MW load step makes the frequency fall at 2.00 Hz/s. With nothing else acting, the island would lose 1 Hz in half a second. Inertia therefore buys time and nothing more. It sets how fast the frequency falls, while the layers above decide where

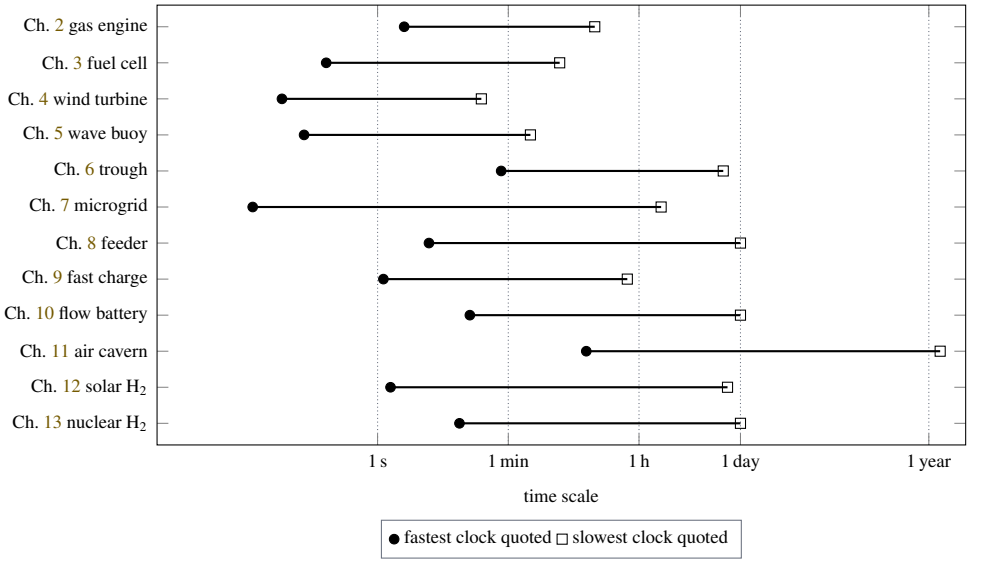

Figure 1.3 The Clocks of Each Chapter's Plant on a Logarithmic Axis. The Filled Circle Is the Fastest Response or Control Period the Chapter Quotes and the Open Square the Slowest Time Constant or Planning Horizon It Quotes, for Example a Module's Power 63 Percent Settled in 2.3 s and a Jacket Still Warming After 15 min (Ch. 2), or a Header Holding 13 s of Reactor Output per Bar and a 24 h Plan (Ch. 13).

the fall stops. Grids with a large share of converter-connected wind and solar lose this stored energy, which is why the first swing has become a design question of its own [17, 19].

Chapter 7 uses Equation (1.1) as its first relation. Its battery inverter forms the grid as a virtual synchronous machine with a virtual inertia of 2 s on 1 MW, and the bound $f_0 \Delta P/(2H_v S_i)$ gives 4.94 Hz/s for its 0.329 MW import, while a relay averaging over 0.1 s reads 2.48 Hz/s. The same balance governs a rotor that is not on a grid. The rotor of Chapter 4, 4.0×10^7 kg m^2 at its rated 1.267 rad/s, stores 32.1 MJ, which is 6.42 s of its 5.0 MW rating, and its open-loop time constant falls from 25.8 to 15.0 s across the range below rated wind. The steam header of Chapter 13 is the thermal equivalent: it stores 3.0 MWh per bar, about thirteen seconds of reactor output, so its pressure loop must close the heat balance continuously. Chapter 2 models its engines on a bus held by an uninterruptible power supply (UPS), so no swing equation appears there and the UPS takes each instantaneous imbalance instead.

> **+ DID YOU KNOW? How Fast a Frequency May Fall**
>
> A 2018 guidance of the European transmission operators asks generating units to stay connected while the rate of change of frequency stays within 2 Hz/s averaged over 0.5 s, 1.5 Hz/s over 1 s and 1.25 Hz/s over 2 s [20]. E-1's first fall, 2.00 Hz/s before any governor moves, is as steep as the limit of the shortest window.

Above inertia, control is arranged in layers that act on longer and longer clocks [21, 22]. Primary control, a droop in each governor or converter, stops the fall within seconds. Secondary control returns the frequency, or a voltage, to its nominal value within tens of seconds to minutes. Tertiary control and economic dispatch set each unit's share over minutes to hours, and commitment and storage policy look a day to a season ahead. Each layer can treat the layer below as settled only if it is several times slower, which is the same rule that sets an outer loop of a cascade several times slower than its inner loop in Volume 3. Figure 1.3 shows how widely the plants of this volume spread across that hierarchy.

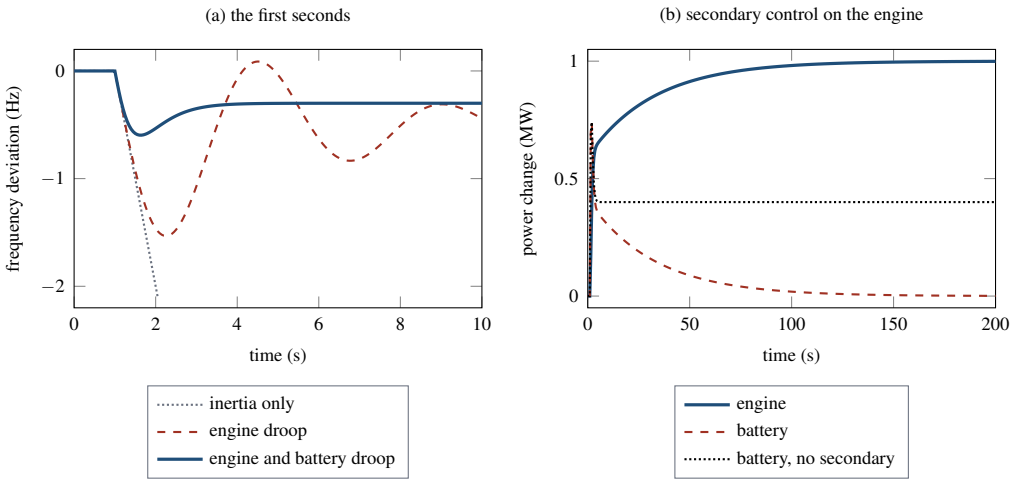

Figure 1.4 E-1 After a 1 MW Load Step at $t = 1$ s. (a) Frequency Deviation on Inertia Alone (Dotted, Leaving the Axis at 2.0 Hz/s), With the Engine's 5 % Droop (Dashed) and With the Battery's 2.5 % Droop Added (Solid). (b) With an Integral Secondary Loop on the Engine's Setpoint, the Engine (Solid) Takes Over the Step and the Battery (Dashed) Returns to Zero, While Without It the Battery Keeps 0.40 MW (Dotted).

1.4.3 Droop and Restoration

A droop makes each unit's power fall in proportion to the frequency's rise, as a proportional controller with no integral action. A unit of rating S and droop R changes its power by $\Delta P = -K\,\Delta f$ with the stiffness $K = S/(Rf_0)$ in megawatts per hertz. When several units share a load step on one frequency, each takes a share in proportion to its stiffness and the frequency settles where the shares add up to the step, as Equation (1.2) shows.

$$\Delta f = -\frac{\Delta P_L}{\sum_i K_i}, \qquad \Delta P_i = \frac{K_i}{\sum_j K_j}\,\Delta P_L, \qquad K_i = \frac{S_i}{R_i f_0} \tag{1.2}$$

E-1's engine governor has a droop of 5 % on 6 MW, a stiffness of 2.0 MW/Hz. Its battery inverter has a droop of 2.5 % on 2 MW, 1.33 MW/Hz. Chapter 7 quotes stiffness per unit of speed, which is f_0 times these numbers. After the 1 MW step the engine alone would settle 0.50 Hz low. With both, the frequency settles 0.30 Hz low and the engine carries 60 % of the step and the battery 40 %.

How deep the frequency falls first depends on how fast each unit answers. Figure 1.4a integrates Equation (1.1) with an engine whose governor and air path lag by 2 s and an inverter that lags by 0.1 s. On inertia alone the frequency falls in a straight line. With the engine's droop alone it falls to a nadir 1.53 Hz low 1.26 s after the step and rings before settling at 0.50 Hz low, because a 2 s lag on a stiff droop is an underdamped loop with a damping ratio of 0.18. With the battery's droop added, the nadir is 0.60 Hz low after 0.63 s, and at its peak the battery carries 0.78 MW, almost the whole step, while the engine catches up. The fast unit takes the transient and the slow unit the steady share.

Droop alone leaves the frequency off nominal and leaves the battery delivering its share for as long as the load stays. At 0.40 MW, a battery at half charge would be empty after 4.5 h. Secondary control removes both. An integral loop on the engine's setpoint, $\mathrm{d}P_{set}/\mathrm{d}t = -k_I\,\Delta f$ with $k_I = 0.1$ MW per hertz-second, is reverse acting and has a time constant of about $\sum K_i/k_I = 33$ s. Figure 1.4b shows the engine taking over the whole step and the battery handing it back. The frequency is 0.12 Hz low 30 s after the step and within 0.003 Hz 150 s after it. Over the whole event the battery delivers 3.70 kWh, less than a tenth of a percent of its capacity. Regulation costs a store power, not energy. Its energy is spent by the dispatch that decides what it does over hours.

Chapter 7 measures every part of this. Its genset and battery share a load in the ratio of their stiffnesses, 1.26 measured against 1.26 declared, its secondary returns the frequency to 60.000 Hz on the genset, and its falls after a step stay between 0.302 and 0.304 Hz from no virtual inertia to 3 s. The same division of labor, a fast actuator that carries a transient and a slow one that takes the steady load, is the structure of Chapter 2 and Chapter 3, where a UPS carries the bus while engines or a fuel processor catch up. Hierarchical control of this kind is the standard arrangement of microgrids and of the converters in them [23, 24].

Voltage has a droop of its own. On a radial feeder the voltage at a node differs from the substation by approximately the resistance times the active flow plus the reactance times the reactive flow, and a smart inverter absorbs reactive power in proportion to how far its voltage exceeds a deadband δ, with slope k. For one inverter at the end of a line of resistance r and reactance x, exporting active power P, Equation (1.3) gives the voltage rise and its settled value under the Volt-VAR droop.

$$\Delta v \approx rP - xq, \qquad q = k\,(\Delta v - \delta) \;\Rightarrow\; \Delta v = \frac{rP + xk\delta}{1 + xk} \tag{1.3}$$

At $r = 0.05$ pu, $x = 0.10$ pu and full export the rise is 0.050 pu without reactive power. With $k = 5$ and $\delta = 0.02$ pu the inverter absorbs 0.10 pu and the rise settles at 0.040 pu. Like the frequency droop, the Volt-VAR droop shares the correction and leaves a deviation. The slow layer that removes it is the tap changer at the substation, which moves the whole feeder in steps after a timer. Chapter 8 shows what happens when the slow layer misreads the fast one. Its tap changer's line-drop compensator reads the inverters' reactive swing as load, and with the Volt-VAR slope raised from 7.33 to 22 the tap changer makes 18 operations on a cloudy day instead of 2.

1.4.4 Storage State as an Inventory

A store is an inventory. Its energy integrates what is charged and what is discharged, with an efficiency on each side and a standing loss, as Equation (1.4) shows.

$$E\,\frac{\mathrm{dSOC}}{\mathrm{d}t} = \eta_c P_c - \frac{P_d}{\eta_d} - \sigma_s E\,\mathrm{SOC}, \qquad 0 \le \mathrm{SOC} \le 1 \tag{1.4}$$

A store has no self-regulation: with charge and discharge equal in power, the state of charge still drifts by the losses. Two numbers size it, the power rating and the energy, and their ratio is its duration at rating. E-1's battery has 2 MW and 4 MWh, a duration of 2 h, and an efficiency of 0.9 each way, a round trip of 0.81. Full to empty it delivers 3.60 MWh, and empty to full it takes 4.44 MWh. Storage technologies differ most in these two numbers and in how independently they can be bought [25, 26]. Chapter 10 buys them separately, 100 MW of stacks on 600.4 MWh of electrolyte, six hours from full to empty and 3.9 h of the working band at rated discharge. Chapter 11 stores 24.6 GWh in a cavern, which is 81.9 h at its expander rating, and Chapter 12 holds 15.7 h of hydrogen at its fuel cell rating.

> **+ DID YOU KNOW? Troughs Now Come With Storage**
>
> Parabolic trough plants can no longer compete with solar panels for daytime power, so the plants built since 2013 have included three hours or more of thermal storage. Their solar fields are oversized to run at full power by day and still charge the store [6].

The losses make even a balanced duty drain the store. If E-1's battery regulates with a square wave of ± 0.5 MW, equal time charging and discharging, its state of charge falls by 1.32 % an hour, about a third of the battery in a day. Something must therefore hold the state of charge, and that loop is an inventory loop of the kind that Volume 6 treats for tanks. A proportional state-of-charge manager adds $P = K\,(\mathrm{SOC} - \mathrm{SOC}_{ref})$ to the battery's discharge, which is direct acting. Against Equation (1.4) without losses it gives the first-order return of Equation (1.5).

$$\frac{E}{K}\,\frac{\mathrm{dSOC}}{\mathrm{d}t} = -(\mathrm{SOC} - \mathrm{SOC}_{ref}), \qquad \tau_{soc} = \frac{E}{K} \tag{1.5}$$

With $K = 2\,\text{MW}$ per unit of charge, 0.02 MW per percent, E-1's manager returns the charge with a time constant of 2 h, and the efficiencies make it 1.8 to 2.2 h depending on the direction. Like an averaging level loop, it should be slow: a fast manager would fight the regulation it exists to support, and a proportional one settles 2.6 points below its reference against the drift of the square wave, which is harmless in a band. The microgrid of Chapter 7 uses such a manager with a band from 40 to 80 % and 0.05 MW per percent outside it, and the UPS reserve loops of Chapter 2 and Chapter 3 hold a stored energy for the next trip or burst.

Near its limits a store loses power, not only energy. A battery management system tapers the allowed charge and discharge power over the last points of the band, and the controller above must receive those limits as the bounds of its output. Chapter 10 calls this cascade anti-windup and Chapter 7 tapers its battery's limits over 20 % of charge. Several plants hold two inventories on two clocks. Chapter 10 passes electrolyte through its stacks in 10 s and through its tanks in 11.1 min, so the charge measured at the stack is a biased estimate of the tank. Chapter 11 stores gas in the ground and heat in a packed bed, and finds that the heat more than the gas decides whether its winter commitment is kept. Chapter 9 fills an inventory as its product: the pack's state of charge is the integral of the current that four limits allow.

1.4.5 Heat and Power From One Fuel

A cogeneration unit makes electricity and heat from one fuel input, as Equation (1.6) shows [14].

$$P_e = \eta_e P_f, \qquad Q_h = \eta_h P_f = \psi P_e, \qquad \psi = \frac{\eta_h}{\eta_e} \tag{1.6}$$

The heat-to-power ratio ψ ties the two products together. E-1's engine converts 40 % of its fuel to electricity and recovers 50 % as hot water, so $\psi = 1.25$. One fuel command cannot hold an electrical balance and a heat balance at once. The plant follows one product and lets the other float: electrical following holds the power and makes heat as a by-product, and thermal following holds the heat and lets a grid or a store take the electrical difference. An island has no grid, so E-1 must follow its electricity. Its hot-water store then integrates the heat mismatch $Q_h - Q_L$ against the heat demand Q_L, and when it is full the surplus goes to a dump radiator, and when it is empty a boiler makes up the deficit.

> ✦ DID YOU KNOW? **Two Products in a Fixed Proportion**
>
> A 2017 federal catalog gives typical power-to-heat ratios of 0.5 to 1.2 for reciprocating engines and 1 to 2 for fuel cells. An engine starts in about 10 s, a fuel cell in 3 hours to 2 days [1]. E-1's engine, with $\psi = 1.25$, makes 0.8 units of electricity per unit of heat.

Figure 1.5 follows one day. The engine supplies the electric demand less the solar output, so at night it runs near 4 MW and recovers 5 MW of heat against a demand near 3.8 MW, and at noon the solar takes most of the load and the engine recovers only about 2 MW while the town still wants 3.8 MW. Over the day the heat recovered, 102.1 MWh, nearly matches the heat demanded, 99.1 MWh, yet a 6 MWh store started half full is full by 02:30 and empty by 12:22. It dumps 6.53 MWh at night and in the evening and buys 6.59 MWh from the boiler in the afternoon. The mismatch is in time, not in quantity, and the store that removes it is set by the range of the cumulative mismatch, as Equation (1.7) shows.

$$E_h \geq \max_t \int_0^t (Q_h - Q_L)\,\mathrm{d}t' - \min_t \int_0^t (Q_h - Q_L)\,\mathrm{d}t' \tag{1.7}$$

For E-1's day that range is 12.6 MWh. A store of that size started at 6.5 MWh neither dumps nor fires the boiler, while one of 12 MWh started half full still dumps 0.12 MWh and buys 0.59 MWh. The starting level matters as much as the size, which is a dispatch decision taken with a forecast of tomorrow's heat.

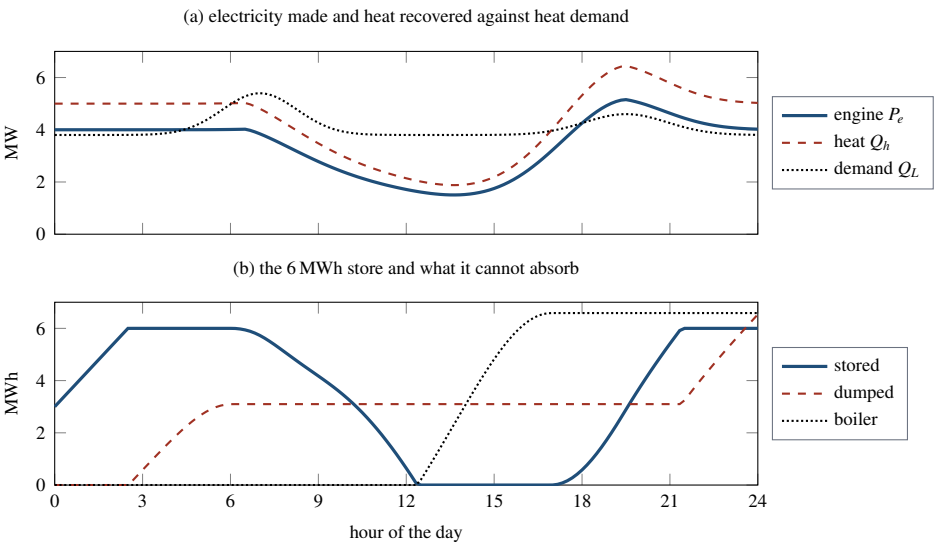

Figure 1.5 One Day of E-1 With the Engine Following the Electric Demand Less the Solar Output. (a) Engine Power (Solid), Recovered Heat (Dashed) and Heat Demand (Dotted). (b) Energy in a 6 MWh Hot-Water Store Started Half Full (Solid), With the Heat Dumped When It Is Full (Dashed) and Bought From a Boiler When It Is Empty (Dotted), Both Cumulative.

The two products also answer on different clocks. In Chapter 2 a fuel step moves a module's power 63 percent of the way in 2.3 s and its recovered heat in 68 s, and its jacket is still warming fifteen minutes later. When one of three modules trips on the island, the UPS carries the lost 1.0 MW for 12.5 s, and the hot-water buffer goes on losing heat for fourteen minutes. That chapter's plant offers base load, electrical following and thermal following as selectable authorities, and an overtemperature override that takes engine capacity when heat cannot be stored or rejected. Its fuel cell counterpart, Chapter 3, has the coupling in another form. Its current moves in fractions of a second, while a step of its fuel valve needs 17.3 s to deliver 63 percent of its hydrogen to the anode, so a governor must hold the current back and a UPS must carry the load meanwhile. Chapter 13 reverses the problem. Its reactor makes a constant 860 MWth, and all the flexibility lies downstream in a turbine, a salt store, an electrolyzer and a dump condenser whose duties must add up to that heat at every instant. Chapter 6 collects heat whose best temperature moves with the sun. Its receiver loses heat faster the hotter it runs while the power block converts heat better, and the outlet temperature that maximizes net output falls from the design value at full sun to the bottom of the search interval at a quarter of it.

1.4.6 Forecasts

A dispatch decision depends on what the load, the weather or the price will do, and the forecast that supplies it degrades with the lead time. The simplest model shows how. If a deviation from the expected value follows a first-order autoregression with coefficient φ per step, the best forecast k steps ahead is φ^k times the present deviation, and its skill, one minus the forecast's mean squared error over the variance, is given by Equation (1.8).

$$x_{k+1} = \varphi\, x_k + w_k, \qquad \hat{x}_{k+j} = \varphi^j x_k, \qquad \text{skill}(j) = \varphi^{2j} \tag{1.8}$$

With $\varphi = 0.9$ the skill is 0.81 one step ahead, 0.43 four steps ahead and 0.19 eight steps ahead. Real forecasts of sun, wind and load combine weather models, statistics and recent measurements, but their skill falls with lead time in the same way [27, 28]. Chapter 5 measures exactly this. Its autoregressive

Table 1.1 E-1's Battery, 2 MW and 4 MWh, Over Three Days of Hourly Prices: Margin of Each Policy and Its Share of the Best Possible, $1,812

Policy	Margin ($)	Share
Plan over 1 h, no value on the energy left	54	3.0 %
Plan over 1 h, energy left valued at yesterday's water value	734	40.5 %
Plan over 2 h, no value on the energy left	1,282	70.7 %
Plan over 3 h, no value on the energy left	1,740	96.0 %
Plan over 6 h or longer	1,812	100.0 %
Plan over 24 h, forecast blind to the $400 hour	1,131	62.4 %
Rule: charge below $25, discharge above $60, at rating	1,306	72.1 %

forecast of the wave force has a skill of 0.96, 0.85, 0.82 and 0.78 at leads of 1, 2, 3 and 4 s, and half of what a perfect four-second forecast adds to the energy captured comes from its first 0.5 s.

A dispatch layer plans over a horizon of N steps, applies the first move and plans again, the receding horizon of a predictive controller with an economic objective [29]. The horizon must reach far enough to see why a store should be filled or kept. E-1's battery is dispatched against three days of hourly prices with a perfect forecast, starting half full, by dynamic programming over a grid of 0.1 MWh. Dispatch against a price of this kind is the core problem of power system economics [30]. The prices follow a daily cycle from $12 to $95 per megawatt-hour, 10 % higher on the third day, and on the second day the hour from 15:00 is priced at $400. Table 1.1 gives the margin of each policy as a share of the best possible, $1,812 over the three days. A plan one hour long sees no reason to charge and earns 3 %. A plan of three hours earns 96 %, and from six hours on the plan earns everything, because six hours reach from the cheap midday hours to the dear evening ones.

A short plan treats energy left at its end as worthless and so never stores any. The remedy is a terminal value, the price at which a stored megawatt-hour would be sold later. The merit order gives it, as Equation (1.9) shows. A full store supplies n blocks of length Δ at its rating P_{max}, so the n-th highest price $p_{[n]}$ among the blocks ahead is the price at which its last megawatt-hour would be sold. A megawatt-hour in the store is then worth $p_w = \eta_d\, p_{[n]}$, the marginal value of stored energy, which is the reservoir water value of hydro scheduling [31].

$$n = \left\lceil \frac{\eta_d E}{P_{max}\,\Delta} \right\rceil, \qquad p_w = \eta_d\, p_{[n]}, \qquad \text{charge if } p < \eta_c\, p_w, \quad \text{discharge if } p > \frac{p_w}{\eta_d} \tag{1.9}$$

For E-1, $n = \lceil 0.9 \times 4/2 \rceil = 2$ one-hour blocks, and the second highest price of the first day is $90, so a stored megawatt-hour is worth $81. The battery should charge below $72.90 and discharge above $90. Valuing the energy left at the end of each plan at the previous day's water value lifts the one-hour plan from 3 % to 40.5 %, while at two and three hours it earns 63.9 and 92.9 %, slightly less than without it, because a value learned from yesterday's evening holds energy that a longer view would sell. A terminal value substitutes in part for a horizon that is too short. The last two rows of Table 1.1 show what a forecast cannot give. A plan blind to the $400 hour holds a full battery through it, waiting for the evening, and loses $681. A rule with fixed thresholds that reacts to the measured price sells in the $400 hour but earns only 72 % of the best, because its thresholds are wrong for ordinary hours.

The chapters measure the same trade. Chapter 10 finds its forecast worth nothing on an ordinary day, where feedback books $16 more, and $16,841 on a day with one scarce hour, and shows that a plan with no terminal condition sells its tank at the edge of the horizon. Chapter 13 sets the storage thresholds of its optimizing architecture from a merit order of its forecast block prices, finds that three price thresholds capture 78 % of what its economic optimizer adds over a fixed split on an ordinary day, and shows the two separating by $85,447 when the forecast misses an evening spike. Chapter 11, Chapter 6 and Chapter 12 find a threshold price for what they store by bisection over 30 daily blocks,

Table 1.2 For Each Chapter, the Fast Regulating Layer, the Slow Deciding Layer and What Joins Them

Ch.	Fast layer	Slow layer	What joins them
2	Module power loops, seconds	Authority selector, staging, MPC supervisor	UPS and hot-water buffer, overtemperature override
3	Current governor, voltage guard	Fuel feedforward and trim, reserve loop	UPS, estimate of arriving hydrogen
4	Torque and pitch loops	Region logic, gain schedule	Rotor inertia, wind preview
5	Force law every 0.1 s	Sea-state estimate, 120 s	Forecast of the wave force
6	Outlet loop on field flow	Setpoint optimizer, day schedule	Two-tank store, defocus override
7	Droops, virtual inertia, shedding	Secondary, SOC manager, genset start	Battery charge and its limits
8	Inverter Volt-VAR curves	Tap changer, central program	Compensator, look-ahead
9	Selector of current limits	Charge to target, predictive plan	Anode map, selector preload
10	Power loop, 18 s	24-hour dispatch program	Tanks, bounded net-load trim
11	Grid trim and heat loops	Inventory policy, seasonal reserve	Cavern, thermal store, capacity floor
12	Rectifier loops, 1.5 s	State machine, evening schedule	Gas inventories, interlocks
13	Six PI loops, header pressure	Economic plan every 15 min	Header identity, salt store

28 half-hour blocks and 32 half-hour blocks. Chapter 8 gives its predictive controller a preview of the clear-sky rise, which takes a cloudy day at 150 % solar from 15.9 to 0 minutes out of band. On the shipped challenge a program judged against the declared shapes 2 h ahead spends 10.2 minutes out against the controller's 11.1, and neither holds the band. The predictive controllers of Chapter 4 and Chapter 5 use a preview of the wind or a forecast of the wave force a few seconds ahead.

1.4.7 Fast Regulation and Slow Dispatch

Every chapter of the volume joins a fast regulating layer to a slow deciding one, and the join takes one of four forms. A store lets the two run at different rates: the fast layer draws on it and the slow layer refills it, as E-1's battery lends 3.70 kWh that its engine repays. A setpoint or reference written by the slow layer is held by the fast one, as a dispatch plan is trimmed by a net-load loop. A limit, selector or interlock lets protection take the actuator from either layer. A forecast lets the slow layer act before the fast layer sees anything. Table 1.2 lists the two layers and the join for each chapter. Chapter 14 adds the measured results and compares fast regulation with slow dispatch across the technologies.

1.5 Instruments and Final Elements in This Volume

Figure 1.6 lists what the controllers of each chapter read and what they move. Ten chapters read a power, a voltage or a current, and eight a temperature. Eight control a value that only a model computes, such as a state of charge or an anode potential, and seven act on a price or a forecast. A power converter is the final element in nine chapters.

Power and voltage reach the controller through instrument transformers. A potential transformer steps the line voltage down to a safe level, typically 120 V, and a current transformer steps the line current down, typically to 5 A at full load, for the meters and the protective relays to read. An energized current transformer must never have its secondary opened, because it may then develop an extremely high voltage [32]. Temperatures come mostly from thermocouples (Figure 1.7), whose most common failure is to burn out, after which the open input can read electrical noise as a wildly varying temperature [32].

A grid-following converter latches onto the grid voltage with a phase-locked loop and, for fixed power commands, acts as a constant source of real and reactive power. Unlike a synchronous machine, which can supply large currents for short periods, it has a hard current limit [12], the headroom

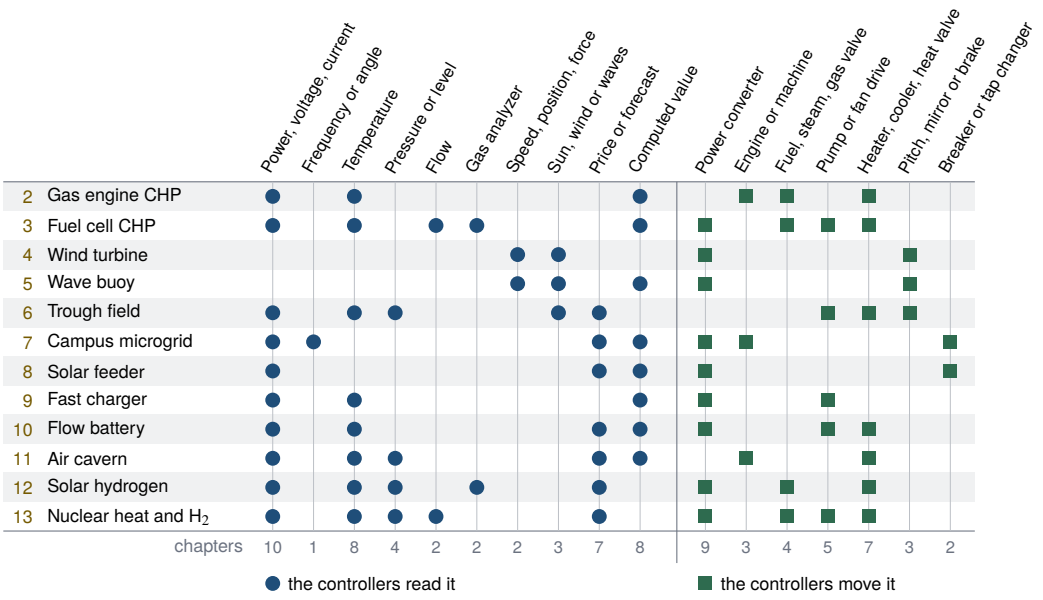

Figure 1.6 Sensors and final elements by chapter: a round mark where the chapter's controllers read the device and a square mark where they move it, as its variables table, control structure or technology section names them. A computed value is not a sensor. Notice that only the wind turbine and the wave buoy work without a power meter, and that the converter is the commonest handle.

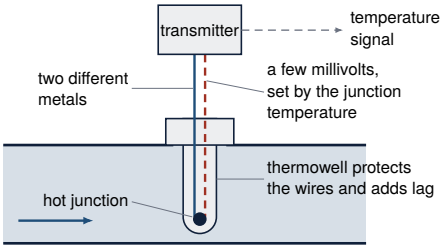

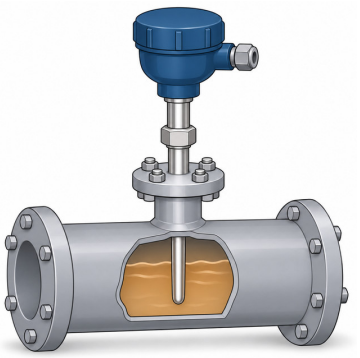

Figure 1.7 Principle of a thermocouple in a thermowell. Notice that the well that protects the junction also adds a lag [32], 2.2 s at the trough outlet of Chapter 6 and 5 s in the stack of Chapter 3.

that Chapter 7 must keep. Pumps and fans are moved by drives that set a motor's speed through the frequency of its supply [32], as for the electrolyte of Chapter 10.

Figure 1.8 places the volume's loops on their platforms. Supervisory and lower loops run with cycle times from milliseconds to minutes. A safety instrumented system is often independent of the control system, so that a failure of the control system does not weaken it [33]. The apps run their fastest laws every 0.02 s on the rotor of Chapter 4 and plan twice a day in the cavern of Chapter 11.

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

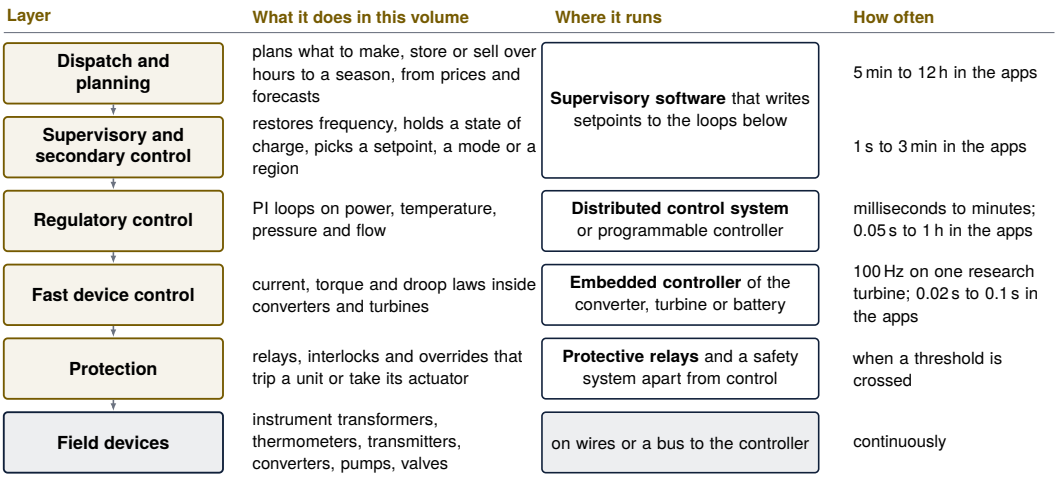

Figure 1.8 The control layers of this volume, the platform that hosts each and how often each acts. Loop periods are from [33], the 100 Hz from the research turbine that Chapter 4 cites and the app ranges from the chapters' technology sections. Notice that the dispatch layer acts once in minutes or hours, while the converter laws below it act many times a second.

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

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

way in every chapter. The exceptions are standalone studios, such as the cascade studio and the two Temperature Control Lab apps, which follow the same ideas in their own code. This section describes the framework once and states where the standalone apps differ. Table 1.3 lists what a reader meets.

1.6.1 The Simulation Engine

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

Three clocks run at separate rates. The plant advances every step h. The controller executes every control period Δt, equal to h unless the app declares a slower controller, and holds its outputs between executions. The log takes one sample at a fixed interval of simulated time for the trends, the scorecard and the exported file. The order of work in one step is shown in Figure 1.9. The controller also executes once when a scenario is loaded, so every run starts from an initialized controller.

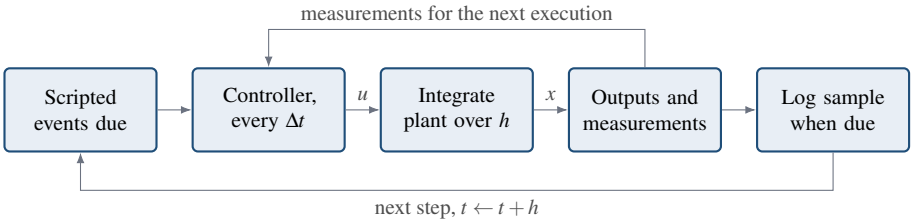

Figure 1.9 Work of the Engine in One Integration Step

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 (1.10) to (1.13) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.10}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.11}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.12}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.13}$$

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 (1.13) drives I toward the value that reproduces the output applied, which is back-calculation [34]. With no integral action, Equation (1.13) is skipped. Clamping the running sum would instead discard each increment that

crosses a limit and keep each one that returns, so symmetric noise would push the output off the limit even in a proportional-only loop.

The derivative acts on the measurement, so a setpoint step gives a proportional kick and no derivative kick. A change of K_c acts on the whole error at the next execution, as in the position form. The action $a = +1$ suits a process whose PV rises with the MV, a reverse-acting controller in the ISA convention, and $a = -1$ gives the direct-acting one.

In Manual the controller keeps executing and sets $I_k = u_{\mathrm{man}} - u_b - S_k$, so the output it would send equals the manual value u_{man}. The first automatic output differs from it by one period's change in P and D. In Auto the engine copies each controller output to the manual station, so a switch to Manual is also bumpless. The standalone studios with PID loops use velocity-form code of their own, with the limit on the output and anti-windup on the integral alone, in which a gain change applies only to later increments.

1.6.3 The Predictive Control and Estimator Options

More than 100 apps offer model predictive control (MPC) as an option beside the structure they open with, which stays the default. Volume 9 develops the theory. In most of them the controller plans with a small linear model in discrete time as shown in Equation (1.14).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.14}$$

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 (1.15) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q\,(y_k - r_k) + \sum_{j} \Delta u_j^{\top} R\,\Delta u_j + \sum_{i} \rho_i\, s_i^2 \tag{1.15}$$

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 [35]. 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 [36]. 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 [37].

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

1.6.4 Scenarios, Progress and Data

An app teaches through scenarios, the numbered cases at the top of the Control panel, nine to thirteen in most apps. Loading one sets the model parameters, the controller settings, the disturbances, the mode, the initial state and any scripted events. The information button beside the list opens a briefing in four parts: what the case is, why it matters, what to do and one detail for the practicing engineer.

The Learn card holds the lesson back, with Predict and Try It before the run, hints on request and Explain afterward.

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

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

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

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

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

1.7 Working With AI Agents While Studying

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

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

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

- *Units.* Every term of an equation carries the same units, and a process gain is in PV units per MV unit. A gain computed with the valve as a fraction is 100 times the gain of a loop whose valve is in percent.
- *Signs.* Decide from the physics which way the PV moves when the MV rises. Compare that with the sign of the fitted gain and with the controller action of Section 1.6.2.
- *Steady state.* Set the derivatives to zero. The model must return the operating point that Show Model reports, and a linearization must give zero deviation there.

Table 1.4 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Gas engine CHP	Power recovers in seconds and heat in minutes.	introductory	2.5 h
3	Fuel cell CHP	Current must wait for its fuel, paid from storage.	intermediate	2.5 h
4	Wind turbine	Two actuators hand over one speed, with a scheduled gain.	intermediate	3 h
5	Wave buoy	The best force depends on waves still to come.	advanced	2 h
6	Trough field	The best setpoint moves with the sun.	intermediate	3 h
7	Campus microgrid	Inertia sets the fall, droop and headroom stop it.	introductory	3 h
8	Solar feeder	Devices right alone can hunt together.	intermediate	2.5 h
9	Fast charger	The limit that ends a charge has no instrument.	advanced	3 h
10	Flow battery	The inventory, not the converter, sets the value.	introductory	2.5 h
11	Air cavern	Heat, stored or bought, defends the commitment.	intermediate	2.5 h
12	Solar hydrogen	Interlocks choose the mode, economics the setpoint.	intermediate	2.5 h
13	Nuclear heat and H_2	A horizon is insurance that looks dear on ordinary days.	advanced	2.5 h

- *A run on the app.* Load the named scenario at the default seed, enter the proposed numbers and run. The trend and the scorecard decide. A tuning that reads well on paper and holds the valve at a limit for half the run has failed.

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

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

1.8 Chapter Roadmap

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

Chapters 2, 7 and 10 form a first course: two products on two clocks, inertia, droop and restoration on one bus, and a store whose value is set by when it is spent. They rest on Sections 1.4.3 to 1.4.5. Chapters 3, 4, 6, 8, 11 and 12 add a slow supply behind a fast actuator, a scheduled gain, a moving optimum, interacting devices, a seasonal inventory and a supervisory state machine. Chapters 5, 9 and 13 combine prediction or inference with constraints and are best left until Section 1.4.6 is familiar. A reader who follows the energy can read the chapters in their order, which runs from conversion through networks to storage and dispatch.

KEY POINTS

- An energy system runs in layers: fast loops hold a frequency, a voltage, a temperature or a current, and slow layers decide how much to make, store or sell.
- Inertia sets how fast the frequency falls, while the layers above decide where the fall stops.

- Droop shares a step by stiffness and leaves an offset that secondary control removes, and regulation costs a store power, not energy.
- A store has no self-regulation, so a slow state-of-charge manager must hold it.
- A terminal value substitutes in part for a horizon that is too short.

Where else this applies. Droop and restoration on one machine are the subject of Volume 1, Chapter 8, inventory control returns in Volume 6 and economic predictive control in Volume 9.

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] P. Veers, K. Dykes, E. Lantz, S. Barth, C. L. Bottasso, O. Carlson, A. Clifton, J. Green, P. Green, H. Holttinen, D. Laird, V. Lehtomäki, J. K. Lundquist, J. Manwell, M. Marquis, C. Meneveau, P. Moriarty, X. Munduate, M. Muskulus, J. Naughton, L. Pao, J. Paquette, J. Peinke, A. Robertson, J. Sanz Rodrigo, A. M. Sempreviva, J. C. Smith, A. Tuohy, and R. Wiser, "Grand challenges in the science of wind energy," *Science*, vol. 366, no. 6464, p. eaau2027, 2019.

[3] K. Budal and J. Falnes, "A resonant point absorber of ocean-wave power," *Nature*, vol. 256, no. 5517, pp. 478–479, 1975.

[4] U.S. Congress, "Public utility regulatory policies act of 1978, Public Law 95-617, 92 Stat. 3117," United States Statutes at Large, vol. 92, govinfo.gov, Nov. 1978. [Online]. Available: https://www.govinfo.gov/content/pkg/STATUTE-92/pdf/STATUTE-92-Pg3117.pdf

[5] S. Succar and R. H. Williams, "Compressed air energy storage: Theory, resources, and applications for wind power," Princeton Environmental Institute, Princeton University, Princeton, NJ, Report, Apr. 2008.

[6] M. Mehos, H. Price, R. Cable, D. Kearney, B. Kelly, G. Kolb, and F. Morse, "Concentrating solar power best practices study," National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5500-75763, 2020.

[7] M. C. Chandorkar, D. M. Divan, and R. Adapa, "Control of parallel connected inverters in standalone AC supply systems," *IEEE Transactions on Industry Applications*, vol. 29, no. 1, pp. 136–143, 1993.

[8] International Atomic Energy Agency, *Opportunities for Cogeneration with Nuclear Energy*, ser. IAEA Nuclear Energy Series. Vienna: International Atomic Energy Agency, 2017, no. NP-T-4.1.

[9] P. Denholm, M. O'Connell, G. Brinkman, and J. Jorgenson, "Overgeneration from solar energy in california: A field guide to the duck chart," National Renewable Energy Laboratory, Golden, CO, USA, Tech. Rep., 2015.

[10] S. Ahmed, I. Bloom, A. N. Jansen, T. Tanim, E. J. Dufek, A. Pesaran, A. Burnham, R. B. Carlson, F. Dias, K. Hardy, M. Keyser, C. Kreuzer, A. Markel, A. Meintz, C. Michelbacher, M. Mohanpurkar, P. A. Nelson, D. C. Robertson, D. Scoffield, M. Shirk, T. Stephens, R. Vijayagopal, and J. Zhang, "Enabling fast charging – a battery technology gap assessment," *Journal of Power Sources*, vol. 367, pp. 250–262, 2017.

[11] U.S. Department of Energy, National Energy Technology Laboratory, "Fuel cell handbook," U.S. Department of Energy, Office of Fossil Energy, National Energy Technology Laboratory, Morgantown, WV, Report, Seventh Edition, Contract DE-AM26-99FT40575, Nov. 2004. [Online]. Available: https://netl.doe.gov/sites/default/files/netl-file/FCHandbook7.pdf

[12] Y. Lin, J. H. Eto, B. B. Johnson, J. D. Flicker, R. H. Lasseter, H. N. Villegas Pico, G.-S. Seo, B. J. Pierre, and A. Ellis, "Research roadmap on grid-forming inverters," National Renewable Energy Laboratory, Golden, CO, Technical Report NREL/TP-5D00-73476, 2020.

[13] National Renewable Energy Laboratory, "Advanced inverter functions to support high levels of distributed solar: Policy and regulatory considerations," Brochure NREL/BR-6A20-62612, Golden, CO, OSTI 1164879, 2014.

[14] D. W. Wu and R. Z. Wang, "Combined cooling, heating and power: A review," *Progress in Energy and Combustion Science*, vol. 32, no. 5–6, pp. 459–495, 2006.

[15] P. Mancarella, "MES (multi-energy systems): An overview of concepts and evaluation models," *Energy*, vol. 65, pp. 1–17, 2014.

[16] P. D. Lund, J. Lindgren, J. Mikkola, and J. Salpakari, "Review of energy system flexibility measures to enable high levels of variable renewable electricity," *Renewable and Sustainable Energy Reviews*, vol. 45, pp. 785–807, 2015.

[17] B. Kroposki, B. Johnson, Y. Zhang, V. Gevorgian, P. Denholm, B.-M. Hodge, and B. Hannegan, "Achieving a 100% renewable grid: Operating electric power systems with extremely high levels of variable renewable energy," *IEEE Power and Energy Magazine*, vol. 15, no. 2, pp. 61–73, 2017.

[18] A. Ulbig, T. S. Borsche, and G. Andersson, “Impact of low rotational inertia on power system stability and operation,” *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 7290–7297, 2014.

[19] F. Milano, F. Dörfler, G. Hug, D. J. Hill, and G. Verbič, “Foundations and challenges of low-inertia systems,” in *Proceedings of the 2018 Power Systems Computation Conference (PSCC)*, 2018, pp. 1–25.

[20] ENTSO-E, “Rate of change of frequency (RoCoF) withstand capability: ENTSO-E guidance document for national implementation for network codes on grid connection,” European Network of Transmission System Operators for Electricity, Brussels, Guidance document, Jan. 2018. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/Network%20codes%20documents/NC%20RfG/IGD_RoCoF_withstand_capability_final.pdf

[21] IEEE/CIGRE Joint Task Force on Stability Terms and Definitions, “Definition and classification of power system stability,” *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1387–1401, 2004.

[22] J. M. Guerrero, J. C. Vasquez, J. Matas, L. G. de Vicuña, and M. Castilla, “Hierarchical control of droop-controlled AC and DC microgrids: A general approach toward standardization,” *IEEE Transactions on Industrial Electronics*, vol. 58, no. 1, pp. 158–172, 2011.

[23] J. Rocabert, A. Luna, F. Blaabjerg, and P. Rodríguez, “Control of power converters in AC microgrids,” *IEEE Transactions on Power Electronics*, vol. 27, no. 11, pp. 4734–4749, 2012.

[24] D. E. Olivares, A. Mehrizi-Sani, A. H. Etemadi, C. A. Cañizares, R. Iravani, M. Kazerani, A. H. Hajimiragha, O. Gomis-Bellmunt, M. Saeedifard, R. Palma-Behnke, G. A. Jiménez-Estévez, and N. D. Hatziargyriou, “Trends in microgrid control,” *IEEE Transactions on Smart Grid*, vol. 5, no. 4, pp. 1905–1919, 2014.

[25] K. C. Divya and J. Østergaard, “Battery energy storage technology for power systems: An overview,” *Electric Power Systems Research*, vol. 79, no. 4, pp. 511–520, 2009.

[26] X. Luo, J. Wang, M. Dooner, and J. Clarke, “Overview of current development in electrical energy storage technologies and the application potential in power system operation,” *Applied Energy*, vol. 137, pp. 511–536, 2015.

[27] R. H. Inman, H. T. C. Pedro, and C. F. M. Coimbra, “Solar forecasting methods for renewable energy integration,” *Progress in Energy and Combustion Science*, vol. 39, no. 6, pp. 535–576, 2013.

[28] T. Hong and S. Fan, “Probabilistic electric load forecasting: A tutorial review,” *International Journal of Forecasting*, vol. 32, no. 3, pp. 914–938, 2016.

[29] J. B. Rawlings, D. Angeli, and C. N. Bates, “Fundamentals of economic model predictive control,” in *Proceedings of the 51st IEEE Conference on Decision and Control (CDC)*, 2012, pp. 3851–3861.

[30] D. Kirschen and G. Strbac, *Fundamentals of Power System Economics*. Wiley, 2004.

[31] J. W. Labadie, “Optimal operation of multireservoir systems: State-of-the-art review,” *Journal of Water Resources Planning and Management*, vol. 130, no. 2, pp. 93–111, 2004.

[32] 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.

[33] 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.

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

[35] 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.

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

[37] 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 · 3 min

An island grid has 50 MW of synchronous machines with an inertia constant of 4 s on that rating, at 60 Hz. A 2 MW unit trips.

(a) Find the stored kinetic energy and the rate of change of frequency right after the trip. *Checkpoint:* 200 MJ, a fall of 0.30 Hz/s.

(b) What is the smallest inertia constant on the same rating that keeps the first rate of change at or below 0.5 Hz/s? *Checkpoint:* 2.40 s.

1.2 Theory and calculation · 4 min

Two units share an island load on one frequency at 60 Hz: a 10 MW unit with a droop of 4 % and a 5 MW unit with a droop of 2.5 %. The load rises by 1.5 MW and only primary control acts.

(a) Find each unit's stiffness in megawatts per hertz. *Checkpoint:* 4.17 MW/Hz and 3.33 MW/Hz.

(b) Find the settled frequency deviation and each unit's share of the step. *Checkpoint:* 0.200 Hz low, 0.833 MW and 0.667 MW.

1.3 Concept check · 2 min

An island with droop-only primary control loses 1 MW of generation. What does doubling the inertia constant of its machines do?

(a) it halves the frequency deviation at which the island settles

(b) it halves the first rate of fall and leaves the settled deviation unchanged

(c) it halves both the first rate of fall and the settled deviation

(d) it leaves the first rate of fall unchanged and halves the settled deviation

1.4 Theory and calculation · 4 min

E-1's engine and battery droops have stiffnesses of 2.0 and 1.33 MW/Hz, and after a 1 MW load step the battery first takes its droop share of 0.40 MW. A secondary loop moves the engine's setpoint at $\mathrm{d}P_{set}/\mathrm{d}t = -k_I \Delta f$ with $k_I = 0.05$ MW per hertz-second.

(a) Estimate the time constant with which the frequency returns. *Checkpoint:* about 67 s.

(b) The battery's share decays with the same time constant. Estimate the energy the battery delivers over the event. *Checkpoint:* about 26.7 MJ, or 7.4 kWh.

1.5 Theory and calculation · 4 min

A solar inverter at the end of a feeder exports 1 pu through a line of resistance 0.04 pu and reactance 0.12 pu. Its Volt-VAR curve absorbs reactive power with a slope of 4 beyond a deadband of 0.01 pu.

(a) Find the voltage rise without reactive power and the reactive power that would cancel it entirely. *Checkpoint:* 0.040 pu, 0.333 pu.

(b) Find the settled rise and the reactive power absorbed under the Volt-VAR curve. *Checkpoint:* 0.0303 pu, 0.081 pu.

1.6 Theory and calculation · 6 min

A battery of 5 MW and 20 MWh has an efficiency of 0.92 each way.

(a) How much energy must the grid supply to charge it from 30 to 80 %, and how long does that take at rating? How much does it return discharging back to 30 %? *Checkpoint:* 10.87 MWh, 2.17 h, 9.20 MWh.

(b) A proportional manager adds 0.1 MW of discharge per percent of charge above its reference. Find its time constant and the time to remove 95 % of an offset, ignoring losses. *Checkpoint:* 2.0 h, 6.0 h.

(c) The battery regulates with a square wave of ± 1 MW, half the time each way. At what mean rate does its charge drift, and how far below its reference does the manager settle? *Checkpoint:* 0.0835 MW, 0.42 % an hour, 0.84 points.

1.7 Theory and calculation · 4 min

An engine converts 38 % of its fuel to electricity and recovers 46 % as heat. On an island it makes 3 MW of electricity while the heat users take 2.5 MW, and an empty 6 MWh hot-water store takes the surplus.

(a) Find the heat-to-power ratio, the fuel power and the heat recovered. *Checkpoint:* 1.211, 7.89 MW, 3.63 MW.

(b) How long until the store is full, and what happens after that? *Checkpoint:* 5.3 h, the surplus is dumped.

1.8 Concept check · 2 min

E-1's engine follows the island's electric demand less the solar output. At 03:00 its hot-water store is full and the recovered heat exceeds the heat demand. Which statement is right?

(a) the power loop lowers the fuel until the heat matches the demand

(b) the store goes on absorbing heat beyond its capacity until the morning

(c) the engine switches to thermal following and the frequency is held by the store

(d) the surplus heat is rejected in the dump radiator while the electricity is held

1.9 Concept check · 2 min

A battery is dispatched by a plan that applies only its first move each hour and is replanned, with a perfect price forecast and no value placed on the energy left at the end of the plan. Prices are lowest at midday and highest in the evening, six hours later. How long must the plan be?

(a) long enough to reach from the cheap hours to the dear ones

(b) one hour, because only the first move is ever applied

(c) as long as the whole run, or it will empty the battery each hour

(d) any length, because a perfect forecast makes the horizon irrelevant

1.10 Theory and calculation · 5 min

A forecast error follows a first-order autoregression with coefficient 0.9 per 15 min step. A battery of 2 MW and 4 MWh, with an efficiency of 0.9 each way, faces hourly prices of \$95, 90, 70, 70, 55 and 48 per megawatt-hour over the next six hours.

(a) Find the forecast's skill one hour ahead, and the lead at which it falls to one half. *Checkpoint:* 0.43, about 49 min.

(b) Find the number of hours a full battery can supply at rating, the water value of a stored megawatt-hour, and the charge and discharge thresholds. *Checkpoint:* 2 hours, \$81, charge below \$72.90, discharge above \$90.

CHAPTER 2

Gas Engine Cogeneration Where Power Settles in Seconds and Heat in Minutes

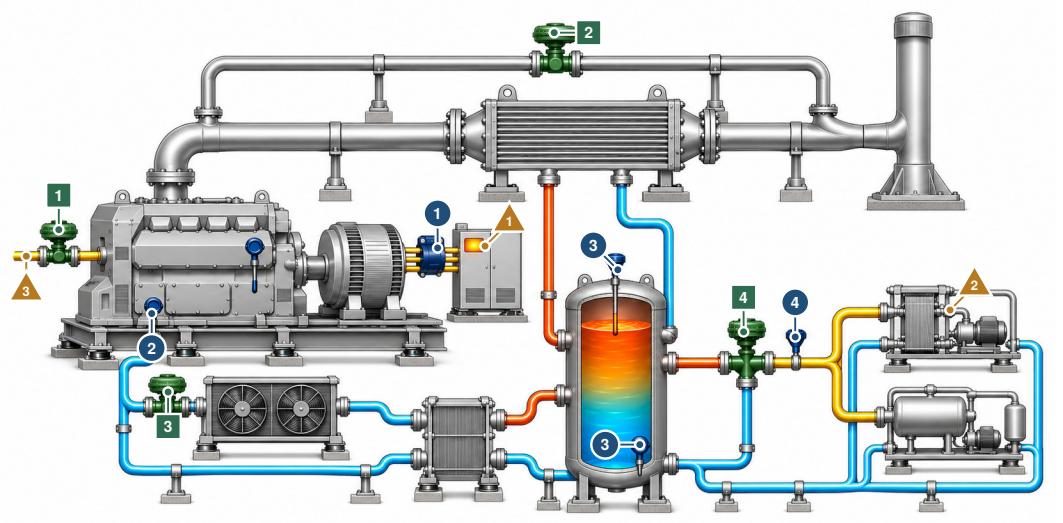

The controller moves 1 fuel valve of the engine 2 exhaust bypass around the heat exchanger 3 valve to the cooling radiator 4 mixing valve for the hot-water supply

The controller reads 1 electric power of the generator 2 cooling water leaving the engine 3 buffer tank, top and bottom 4 hot-water supply temperature

What disturbs it 1 electric load of the site 2 heat demand of the customer 3 energy content of the gas

Notice that the fuel valve changes the generator's power within seconds, while its heat must first warm the engine, the exchangers and the buffer tank. One of the plant's three engines is drawn. Illustration, not to scale.

IN PLAIN TERMS **One Fuel Valve, Two Clocks**

The main idea. A gas engine makes electricity, and its hot cooling water and exhaust heat water for a building and a chiller. More fuel raises the power in seconds, but the heat takes minutes, because the engine must warm up first.

The challenge. A tank of hot water soaks up the difference, so a slow cooling loop can look fine for minutes. When the tank is full and the engines run hot, a safety limit cuts their power, and a backup battery must fill the gap.

Why it matters. A plant that cannot store or shed its heat loses power when it needs it most.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the fact box of the Introduction, the history box of Section 2.2, the try-it box of Section 2.5, the try-it and fact boxes of Section 2.5.5, and the key points
Study	50 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 2.5, Section 2.6, Section 2.7 and its open challenges

The control question One fuel valve per engine sets both the electricity and the recovered heat, but the power settles in seconds and the heat in minutes. Which control layer of a three-engine plant at a data center holds which balance, and when does a shortfall stop being a tuning problem and become a lack of capacity?

Learning objectives

- Explain from the jacket and exchanger heat capacities why a fuel step moves a module's power long before its recovered heat.
- Compute the Willans-line efficiency, the buffer's stored energy and the UPS reserve from the plant parameters.
- Compare feedback with metered-load feedforward, and a long with a short anti-windup tracking time, on matched runs.
- Tell an overtemperature override or a capacity limit from a badly tuned loop by reading the allowed power, the jackets and the UPS.
- Tune the jacket and supply loops so the plant passes the islanded challenge.

Prerequisites Sections 1.4.5 (heat and power from one fuel), 1.4.4 (storage as an inventory), 1.6.2 (velocity-form PID and anti-windup) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive supervisor.

The app https://apmonitor.com/apps/gas-engine-chp/

See also Volume 4, Chapter 2 and Volume 4, Chapter 3 for protective overrides and limiters on one actuator, and Volume 4, Chapter 9 for a trigeneration plant at a data center. Volume 3, Chapter 6 for feedforward from a measured disturbance and Volume 3, Chapter 4 for limits that act on a fuel command.

2.1 Introduction

Reciprocating engines are one of the five prime-mover families of combined heat and power (CHP) that a catalog prepared for the U.S. Environmental Protection Agency describes. The catalog counts more than 2,000 active engine CHP installations in the United States with nearly 2.3 GW of capacity, many of the larger ones built from several units, and puts the heat in an engine's exhaust and cooling systems at 60 to 70 % of its fuel energy [1]. Recovered as hot water that heat can drive a single-stage absorption chiller, and the U.S. Department of Energy lists data centers among the sites with the steady year-round cooling load that such a chiller suits [2]. In the simulated plant three 2 MW modules serve a 2.4 MW data hall. Opening one module's fuel valve from 51.7 to 61.5 % moves its electrical output from 0.985 to 1.213 MW, 63 percent of the way in 2.3 s, while its recovered heat rises from 1.08 to 1.23 MW and needs 68 s for the same share, and its jacket is still warming fifteen minutes later.

> **+ DID YOU KNOW? A Fast-Growing Electric Load**
>
> Data centers in the United States used an estimated 192 TWh of electricity in 2024, 4.7 % of the country's total. The reference case of the same study puts their use in 2030 at 649 TWh, or 11.8 % of the forecast national total [3].

The simulator is built to correct one belief, that a cogeneration plant has recovered once its electrical output has. When one of three modules trips on an island, the UPS carries the lost 1.0 MW for 12.5 s and the other two modules take the load, yet the hot-water buffer goes on losing heat for the next fourteen minutes, from 45 to 34 % of its charge. The same plant shows a second, more expensive version of the belief. A slow engine-cooling loop looks adequate while the buffer still has cool water at its bottom, and in the challenge the buffer hides such a loop for eight minutes, after which the jackets pass 100 °C, the overtemperature override takes capacity for 408 s, and the UPS falls from 80 to 52 % of its charge.

The class of problem is general. One actuator sets two products whose responses are separated in time, a store absorbs the difference between them, and a protective limit takes authority when the store

is full. Section 1.4.5 sets out this coupling on the volume's running example. The boiler-turbine unit, in which fuel, steam and feedwater act on power and pressure together, is a standard multivariable problem of power generation [4], which Volume 9, Chapter 4 treats with predictive control. Combined cooling, heating and power extends the same coupling to a third product [5], and Volume 4, Chapter 9 follows it on a gas turbine at a data center. Whether a CHP plant should follow its electric or its thermal load is a known operating question with measured consequences [6], thermal storage changes how such a plant operates [7], and the controls that keep a saturated loop from winding up and that hand authority to a constraint are general tools of process control [8, 9]. In process terms the buffer is a surge volume on an energy inventory, the dump radiators are its overflow, and the override is a high selector on the fuel valve.

However, these pieces are seldom found on one plant. Operating strategies are compared by primary energy, cost and emissions [6], engine cogeneration models are written and calibrated for energy performance [10, 11], and earlier energy benchmarks pose dispatch as an optimization [12]. A student who wants to watch a cooling loop take engine capacity away needs a plant whose loops, limits and stores are settings with consequences in megawatts, kelvin and kilowatt-hours of reserve. This chapter establishes what each control layer of a three-engine CHP plant at a data center buys, and where control ends and capacity begins, on a simulator that runs in a browser.

Every method used is an established one: the Willans line, the effectiveness exchangers, the velocity-form PID with back-calculation, the high selector and hysteresis staging. The results are a set of matched comparisons measured on one plant at one seed, and one of their findings is that the reserve a slow cooling loop consumes depends on the buffer volume and the jacket heat capacity as much as on the loop. The plant has 33 physical states whose 38 ledgers close the first law to 5×10^{-16} of the fuel, a UPS that cannot pass its floor inside the integrator's stages, and controllers that read only a declared packet of meter readings. Section 2.2 traces the domain. Section 2.3 gives the plant equation by equation, and Section 2.4 the control layers, with the predictive supervisor in Section 2.4.4. Section 2.5 reports the measurements on the eleven cases, from a manual fuel step to the challenge: the predictive supervisor against the PI authorities, three matched comparisons, a demand beyond the engines' capacity, the remaining cases and the challenge, followed by the checks of the code in Section 2.5.6. Section 2.6 describes the hardware of real plants, Section 2.7 states what the model leaves out, and Section 2.8 draws the results together.

2.2 Recovered Heat, the Grid and the Fuel Valve

The straight line of fuel against output that every averaged engine map still uses carries the name of P. W. Willans, whose economy trials of a non-condensing steam engine were reported in 1888 [13]. For a gas engine the same line gives a fixed fuel flow at zero load and a constant marginal efficiency above it, and the EPA catalog puts the efficiency at half load 8 to 10 % below the full-load value [1]. Mean-value engine models of the kind used for engine control add the air path and the torque response as lumped lags [14].

> ⧖ HISTORY **Utilities Must Buy From Cogenerators (1978)**
>
> The Public Utility Regulatory Policies Act became law in November 1978. Its section on cogeneration and small power production told the federal regulator to write rules requiring electric utilities to sell electricity to qualifying cogeneration facilities and to buy electricity from them [15, 16]. The act defined a cogeneration facility as one that produces electricity together with steam or another useful form of energy, such as heat, used for industrial, commercial, heating or cooling purposes [15]. Its efficiency standard counts only half of a plant's useful heat, because electricity is the more valuable product [1]. A cogenerator running in parallel with the grid could now sell a surplus of electricity and buy a shortfall, so the grid could close its electrical balance while the plant chose which load to follow.

Heat recovery from engines settled into a standard arrangement. The jacket coolant carries up to 30 % of the fuel energy and can make 190 to 230 °F hot water, the most common recovery circuit is a closed loop with an excess heat exchanger that sends surplus heat to a cooling tower or a radiator, and most exhaust heat-recovery units are designed for a 250 to 350 °F outlet so that the exhaust does not condense in its piping [1]. Reviews of cogeneration technology place the engine beside microturbines, Stirling engines and fuel cells as a prime mover for buildings [17], and dynamic models of engine cogeneration have been written for residential plants and checked against measured transient and steady operation [10, 11]. Larger plants are often built from several engines, which the catalog notes raises a plant's capacity and availability [1], and which makes load sharing and staging part of the control problem.

Running in parallel with the utility made the control problem what it is now, as the history box earlier in this section tells. A grid-connected plant can follow either load and let the grid close the electrical balance. Over several climates, plants run following the thermal load reduced primary energy use in every city evaluated, while CHP run following the electric load always increased it [6]. Thermal storage loosens the tie between the two products and changes how a CHP unit operates and what it emits [7]. Absorption cooling adds a third product, and single-stage machines driven by 200 to 240 °F hot water are often paired with engine CHP [2, 5].

The control tools themselves came from process practice. Integrator windup against a saturated actuator and the schemes that avoid it were unified in 1989 [8], selectors and overrides that hand an actuator to a constraint are treated together with windup in the constrained-control literature [9], and measured-disturbance feedforward is standard textbook material [18]. The plant simulated here combines these at a data center, where the UPS that already protects the critical load becomes the electrical store that bridges an engine trip or a start delay.

2.3 Three Engine Modules, a Buffer and a Data Hall

The site in Figure 2.1 is an invented teaching design taken from no particular installation. Three identical lean-burn gas-engine generator modules of 2.0 MW gross each feed one site bus. Each module has its own heat recovery. Recovery water drawn from the bottom of the buffer passes the module's jacket-water plate exchanger (JKT) and then its exhaust exchanger (EXH), and returns to the top of the buffer. Jacket heat that the plate exchanger does not take goes to the module's dump radiator, and an exhaust diverter can send gas around the exchanger to the stack. A three-way mixing valve blends tank water with return water into a distribution loop that passes a single-effect absorption chiller and then the thermal customer. The site is a data hall with electric chillers, a cooling-water loop with a tower, and auxiliary loads, connected either to the grid or, islanded, to a 1.5 MW, 0.25 MWh UPS.

Table 2.1 sorts the variables into controlled, manipulated and disturbance variables and states. Per module, seven states are integrated, the fuel power P_f, the shaft power P_m, the jacket node T_j, the exhaust-exchanger metal T_x, the dump valve d, the diverter b and the recovery flow $\dot{m}_r$. The site adds four buffer nodes, the distribution return node T_u, the mixing valve x, the IT node, the chilled-water node, the electric chiller duty, the absorption capacity, the cooling-water node and the UPS energy,

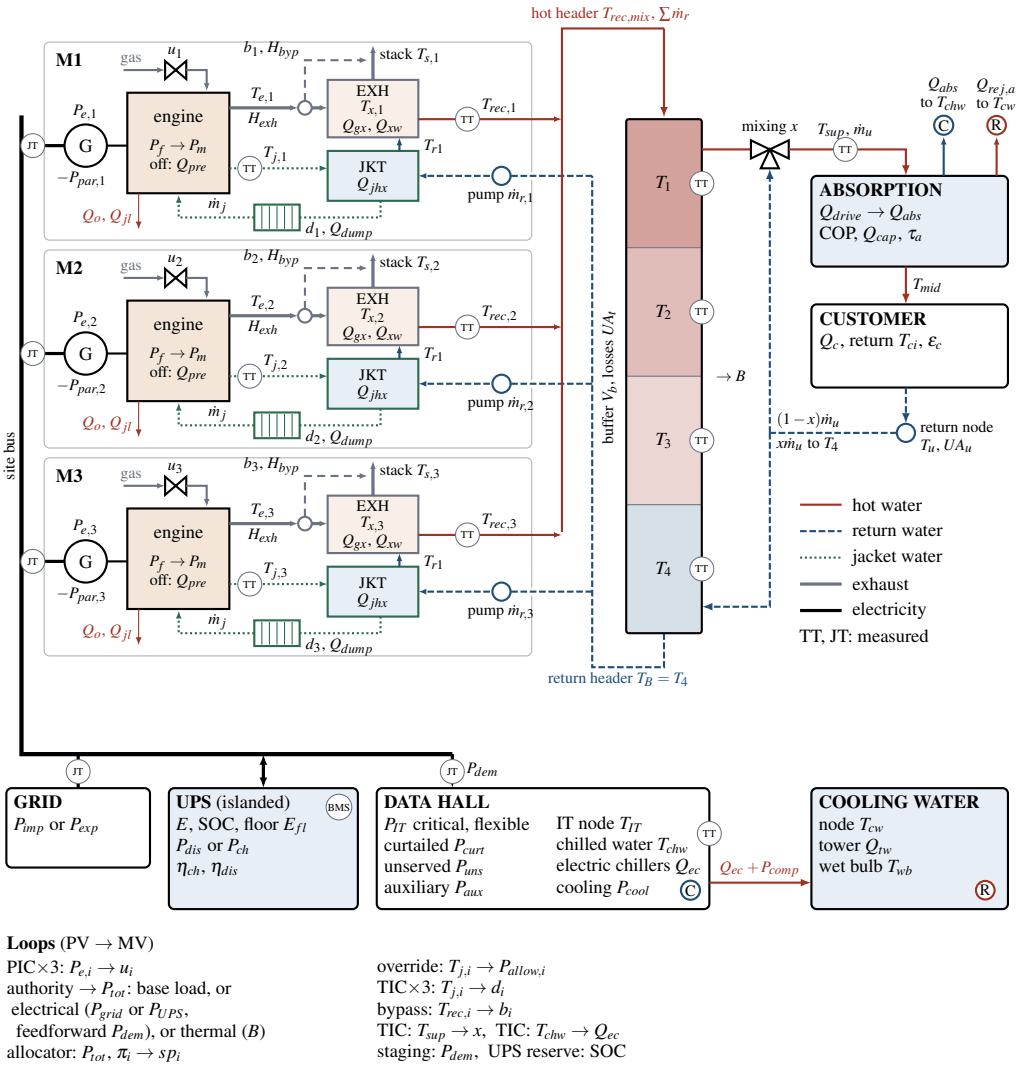

Loops (PV $\to$ MV)
PIC×3: $P_{e,i} \to u_i$
authority $\to P_{tot}$: base load, or
electrical (P_{grid} or P_{UPS},
feedforward P_{dem}), or thermal (B)
allocator: P_{tot}, $\pi_i \to sp_i$

override: $T_{j,i} \to P_{allow,i}$
TIC×3: $T_{j,i} \to d_i$
bypass: $T_{rec,i} \to b_i$
TIC: $T_{sup} \to x$, TIC: $T_{chw} \to Q_{ec}$
staging: P_{dem}, UPS reserve: SOC

Figure 2.1 The simulated plant as drawn on the application's Plant view. Each module i turns fuel $P_{f,i}$ through valve u_i into shaft power $P_{m,i}$ and electrical output $P_{e,i}$, jacket heat that leaves at $T_{j,i}$ through the plate exchanger (JKT, Q_{jhx}) and the dump radiator (d_i, Q_{dump}), and exhaust at $T_{e,i}$ that passes the exchanger metal $T_{x,i}$ or the bypass b_i to the stack. Recovery water at $\dot{m}_{r,i}$ leaves the buffer bottom at T_B, reaches T_{r1} and $T_{rec,i}$, and returns to the buffer top. The mixing valve x sets the supply T_{sup} to the absorption chiller and the customer, whose return node T_u feeds the valve and the buffer bottom. The grid, the UPS and the data hall hang on the site bus below the modules. Circles C and R carry the absorption cooling to the hall's chilled water and its rejection to the cooling water. Hot water is drawn solid, return water dashed, jacket water dotted, exhaust thick gray and electricity thick black. TT and JT mark the measured temperatures and powers. There is no rotor state. Process model only, not a piping diagram, not to scale.

33 physical states in all, and 38 ledger states integrate every boundary flux so that each balance can be checked independently. Inside the model the units are seconds, watts, joules and kelvin, and the interface shows MW, MWh and °C.

Table 2.1 Controlled, Manipulated and Disturbance Variables and States at the Start of Case 3, Grid-Connected Under Electrical Following. Ranges Use the Units of the Start Column

Variable	Role	Start	Range
P_{grid}	CV, grid import	0.200 MW	SP −0.5–1.5
$P_{e,i}$	CV, module power	0.985 MW	0.5–2.0
T_{sup}	CV, supply	90.0 °C	SP 80–95
$T_{j,i}$	CV, jacket outlet	83.1 °C	SP 95, trip 105
B	CV, buffer charge	43 %	SP 10–90
u_i	MV, fuel valve	51.7 %	0–100
x, d_i, b_i	MV, mix, dump, bypass	93.5, 0, 0 %	0–100
P_{tot}	MV of the authority	2.955 MW	module limits
P_{IT}	DV, IT load	2.40 MW	1.0–3.6
Q_{req}	DV, customer demand	3.20 MW	0–4.0
q	DV, gas heating value	100 %	80–105
T_a, T_{wb}	DV, ambient, wet bulb	25, 20 °C	5–45, 5–30
events	DV, scripted events	—	—
P_f, P_m	states, fuel, shaft	2.587, 1.016 MW	—
T_j, T_x	states, jacket, metal	83.1, 96.2 °C	—
$T_1..T_4$	states, buffer	91.5, 91.5, 68.1, 68.1	—
T_{IT}, T_{chw}, T_{cw}	states, hall, cooling	29.7, 15.0, 29.3	—
E	state, UPS energy	80 % of 0.25 MWh	floor 20

2.3.1 One Module: Fuel, Shaft and the Heat Split

The fuel valve sets the fuel power on the lower heating value (LHV) through a first-order lag, and the shaft follows the Willans line through an air-path lag that is slow on a rise and fast on a cut, as shown in Equation (2.1).

$$\begin{aligned} \tau_f \frac{\mathrm{d}P_f}{\mathrm{d}t} &= u\, q\, P_{f,max} - P_f \\ P_m^* &= \eta_m \left(P_f - \varphi_0 P_{f,r}\right) \\ \tau \frac{\mathrm{d}P_m}{\mathrm{d}t} &= P_m^* - P_m, \quad \tau = \begin{cases} \tau_m & P_m^* > P_m \\ \tau_{dn} & \text{otherwise} \end{cases} \\ P_f &= P_m + \kappa_j P_f + \kappa_o P_f + H_{exh} \end{aligned} \tag{2.1}$$

where q is the gas heating value as a fraction of nominal, $P_{f,max} = 1.05 P_{f,r}$, $P_{f,r}$ the rated fuel, η_m the marginal shaft efficiency, κ_j and κ_o the shares of fuel to the high-temperature jacket circuit and to the low-temperature circuit and radiation ($Q_{jg} = \kappa_j P_f$, $Q_o = \kappa_o P_f$), and H_{exh} the exhaust enthalpy above ambient. The exhaust is the residual of the fuel budget, so the balance closes at every instant and a shaft that lags its fuel heats the exhaust instead of losing energy. At rating the module burns 4.762 MW of fuel for 2.062 MW of shaft power, 0.952 MW to the jacket, 0.510 MW to the low-temperature circuit and 1.238 MW of exhaust at 400.9 °C, so the electrical efficiency is 42.0 %. With $\varphi_0 = 0.10$ the line gives $\eta_m = 0.4811$ and, relative to rated efficiency, 76.9, 90.9 and 96.8 % at a quarter, a half and three quarters of rated shaft power. The generator gives $P_e = \eta_{gen} P_m$, or P_m / η_{gen} while a module without fuel motors. The rating falls by 0.18 % per kelvin of ambient above 25 °C, the EPA catalog's 1 percent per 10 °F above 77 °F [1].

The engine is a grid-connected constant-speed reduction with no rotor state. Grid-connected, the grid holds the speed and the mechanical balance is the algebraic $P_e = \eta_{gen} P_m$. Islanded, the UPS-backed bus of Section 2.3.5 holds the voltage and takes the instantaneous imbalance, so the engines are power sources at constant speed and no swing equation, inertia or droop appears. One consequence is that at a trip the fuel burned during its 0.1 s lag heats the exhaust instead of a rotor, 0.174 MJ of exhaust enthalpy in the first second after module 3 trips in case 6.

2.3.2 Jacket, Exhaust Exchanger and Dump Radiator

The jacket node holds the block, the coolant and the high-temperature circuit. The plate exchanger is a counterflow exchanger between the jacket water and the recovery water, the dump radiator acts on the jacket water leaving it, and the exhaust exchanger is a gas-to-metal-to-water chain, as shown in Equation (2.2).

$$\begin{aligned} C_j \frac{\mathrm{d}T_j}{\mathrm{d}t} &= Q_{jg} + Q_{pre} - Q_{jhx} - Q_{dump} \\ &\quad - UA_{jl}(T_j - T_a) \\ Q_{jhx} &= \varepsilon_{cf}(UA_j, C_h, C_r)\, C_{min}(T_j - T_B) \\ Q_{dump} &= d\,\varepsilon_w(G_d, C_h)\, C_h\, (T_{jm} - T_a)^+ \\ C_x \frac{\mathrm{d}T_x}{\mathrm{d}t} &= Q_{gx} - Q_{xw} \\ Q_{gx} &= \min\big(\varepsilon_w(UA_g, C_{gx})(T_g - T_x), \\ &\qquad (T_g - T_{s,min})^+\big)\, C_{gx} \\ Q_{xw} &= \varepsilon_w(UA_w, C_r)\, C_r\, (T_x - T_{r1}) \end{aligned} \tag{2.2}$$

where $C_h = \dot{m}_j c_p$ and $C_r = \dot{m}_r c_p$ are the jacket and recovery heat-capacity flows, $C_{min} = \min(C_h, C_r)$, ε_{cf} the counterflow effectiveness, which becomes $\mathrm{NTU}/(1+\mathrm{NTU})$ for balanced flows, $\varepsilon_w(UA, C) = 1 - \mathrm{e}^{-UA/C}$ the effectiveness of a stream against an isothermal wall [19], $T_{jm} = T_j - Q_{jhx}/C_h$, $T_{r1} = T_B + Q_{jhx}/C_r$, $T_{rec} = T_{r1} + Q_{xw}/C_r$, and $(\cdot)^+ = \max(0, \cdot)$. The cap in Q_{gx} applies only when the exchanger takes heat from the gas. The exhaust flow is $\dot{m}_g = \dot{m}_{g,r}(0.2 + 0.8 P_f / P_{f,r})$ while the engine turns, the gas reaches the exchanger at $T_g = T_a + H_{exh}/(\dot{m}_g c_{p,g})$, and the diverter leaves $C_{gx} = (1-b)\dot{m}_g c_{p,g}$ on the exchanger and sends $H_{byp} = bH_{exh}$ around it, so the stack carries $H_{exh} - Q_{gx}$. A stopped module's preheater, $Q_{pre} = \mathrm{clip}_{[0,\,15\,\mathrm{kW}]}\big(K_{pre}(60\,^\circ\mathrm{C} - T_j)\big)$ with $K_{pre} = 15\,\mathrm{kW/K}$, holds its jacket near 60 °C. At the start of case 3 each module's recovery water rises from 68.1 to 79.0 °C in the plate exchanger, with the jacket at 83.1 °C and 0.500 MW to the water, and from 79.0 to 91.5 °C in the exhaust exchanger, whose metal sits at 96.2 °C while the gas falls from 396.8 to the 120.0 °C minimum and gives 0.579 MW. The jacket heat reaches the water at the lower temperature level and the exhaust lifts it.

2.3.3 Recovery Flow, Buffer, Users and Absorption

Each recovery pump follows its module's measured fuel $\bar{P}_f$ through a 120 s filter on the schedule

$$\dot{m}_r = \dot{m}_{r,max}\,\mathrm{clip}_{[0.15,1]}\big(0.15 + 0.85\bar{P}_f / P_{f,r}\big),$$

reached through a 5 s lag, an equipment schedule rather than a tunable loop. The buffer is four fully mixed nodes of mass M_k, with the recovery flow entering at the top at the mixed recovery temperature, the users' draw $x\dot{m}_u$ leaving at the top and returning to the bottom, and the net flow between nodes carried upwind, as shown in Equation (2.3).

$$\begin{aligned} M_k \frac{\mathrm{d}T_k}{\mathrm{d}t} &= \delta_{k1}\dot{m}_t(T_{rm} - T_1) + \delta_{k4} x\dot{m}_u(T_u - T_4) \\ &\quad + \dot{m}_d^+(T_{k-1} - T_k) + \dot{m}_d^-(T_{k+1} - T_k) \\ &\quad - (UA_t/c_p)(T_k - T_a) \\ T_{sup} &= xT_1 + (1-x)T_u \\ T_{mid} &= T_{sup} - Q_{drive}/C_u \\ Q_c &= \min\big(Q_{req},\ \varepsilon_c C_u (T_{mid} - T_{ci})^+\big) \\ M_u c_p \frac{\mathrm{d}T_u}{\mathrm{d}t} &= C_u\big(T_{mid} - T_u\big) - Q_c - UA_u(T_u - T_a) \end{aligned} \tag{2.3}$$

where $\dot{m}_t = \sum_i \dot{m}_{r,i}$ enters node 1 at the flow-weighted recovery temperature T_{rm}, $\dot{m}_d = \dot{m}_t - x\dot{m}_u$ is the net downward flow with $\dot{m}_d^+ = \max(0, \dot{m}_d)$ acting for $k > 1$ and $\dot{m}_d^- = \max(0, -\dot{m}_d)$ for $k < 4$, δ is the Kronecker delta, $C_u = \dot{m}_u c_p$, and the buffer charge B is the mean over the nodes of $(T_k - 70\,°\mathrm{C})/25\,\mathrm{K}$, clipped to the unit interval. Between 70 and 95 °C the 8 m^3 buffer stores 0.227 MWh, about four minutes of the 3.24 MW recovered at the design point.

The absorption chiller is the capacity-lag reduction with no internal energy storage, in which the cooling and its drive heat are chosen together, as shown in Equation (2.4).

$$\begin{aligned} \tau_a \frac{\mathrm{d}Q_{cap}}{\mathrm{d}t} &= Q_{tgt} - Q_{cap} \\ Q_{abs} &= \min\left(Q_{cap},\ \mathrm{COP}\, C_u (T_{sup} - T_{g,min})^+\right) \\ Q_{drive} &= Q_{abs}/\mathrm{COP} \\ Q_{rej,a} &= Q_{drive} + Q_{abs} + P_{aux,a} \end{aligned} \tag{2.4}$$

where COP is linear in T_{sup} through 0.70 at 87.8 °C and 0.74 at 97.8 °C, the two single-stage hot-water points of the DOE fact sheet [2], and bounded to 0.55–0.76, $T_{g,min} = 70\,°\mathrm{C}$, and the target Q_{tgt} is the chilled-water loop's request held between a 20 % turndown and a capacity map that is linear in T_{sup}, zero at 70 °C and 1.0 MW at 90 °C, capped at 1.15 times that rating, and falls 3 % per kelvin of cooling water above 29.4 °C. Without driving heat there is no cooling, whatever the lag state holds.

2.3.4 Data Hall and Cooling Water

The hall's IT node and the chilled-water and cooling-water loops are Equation (2.5).

$$\begin{aligned} C_{IT} \frac{\mathrm{d}T_{IT}}{\mathrm{d}t} &= P_{IT,del} - UA_{IT}(T_{IT} - T_{chw}) \\ C_{chw} \frac{\mathrm{d}T_{chw}}{\mathrm{d}t} &= UA_{IT}(T_{IT} - T_{chw}) + Q_{aux,IT} \\ &\quad - Q_{abs} - Q_{ec} \\ C_{cw} \frac{\mathrm{d}T_{cw}}{\mathrm{d}t} &= Q_{rej,a} + Q_{ec} + P_{comp} - Q_{tw} \end{aligned} \tag{2.5}$$

where $P_{IT,del}$ is the delivered and not the requested IT power, half of the auxiliary load's heat reaches the chilled water, the electric chillers follow their command through a 30 s lag at a COP of 5.5 that falls 2.5 % per kelvin of cooling water above 29 °C ($P_{comp} = Q_{ec}/\mathrm{COP}_{ec}$), and the tower's duty is $Q_{tw} = fUA_{tw}(T_{cw} - T_{wb})^+$ with a local proportional fan f across 28–30 °C. The cooling electricity P_{cool} is the compressor power plus pumps, tower fans and the absorption chiller's 20 kW of auxiliaries. At the start of case 3 the electric chillers carry 2.46 MW at a COP of 5.47 with the IT node at 29.7 °C, and T_{IT} is an effective equipment node, not a processor or rack-inlet temperature.

2.3.5 The Site Bus and the UPS Boundary Rule

Every electrical term is nonnegative and the balance is evaluated term by term at every step, as shown in Equation (2.6).

$$\begin{aligned} &P_{gen} + P_{imp} + P_{dis} + P_{uns} \\ &\quad = P_{IT} - P_{curt} + P_{cool} + P_{aux} + P_{par} + P_{ch} + P_{exp} \\ &\frac{\mathrm{d}E}{\mathrm{d}t} = \eta_{ch} P_{ch} - P_{dis}/\eta_{dis} \\ &P_{dis} \le \min\left(P_{dis,max},\ \frac{(E - E_{fl})\eta_{dis}}{\Delta t}\right) \end{aligned} \tag{2.6}$$

where P_{gen} is net of each module's auxiliaries, 30 kW plus its dump fans, and P_{par} carries the auxiliaries and preheaters of modules that are stopped or tripped. Grid-connected, the grid takes the imbalance after the UPS's own recharge request, and import and export are never both nonzero.

Islanded, the UPS takes it within the step's limits, and a deficit beyond them is unserved, allocated to flexible IT first, then the auxiliary load, the cooling and the critical IT last, while a surplus beyond the UPS's charge acceptance is refused and its shaft power returns to the exhaust. The step's discharge and charge limits are fixed from the stored energy at the start of the step and held through the integrator's stages. Because the stage weights of the fourth-order Runge-Kutta method are nonnegative and sum to one, the stored energy lands on its floor at most and never past it, and nothing is clipped after integration. A constant 1 MW from 80 % empties the UPS at $0.25 \times 0.60 \times 0.95 = 0.1425$ MWh, 513 s, matched by the model within 3.2×10^{-10} s. Power and energy are separate limits, so at the 1.5 MW rating the same reserve lasts 342 s and a larger deficit goes partly unserved at once.

2.3.6 Events, Instruments and Numerics

Every discrete event happens once, at a step boundary after integration, with a mark on the trends. A module trips when its jacket reaches 105 °C or when it motors below −3 % of rating for 2 s. A trip opens the breaker, removes the shaft at once and the fuel with a 0.1 s lag, and leaves the pumps running, so the heat retained in the jacket and the metal keeps reaching the water. A stopped module started by the staging supervisor or the operator burns idle fuel and closes its breaker 90 s later, and a tripped one may be reset once its jacket is below 90 °C.

The controllers read only a measurement packet. Each signal passes a first-order transmitter lag and then noise, so the noise setting is the noise a trend shows, and every scan draws the same number of random numbers whatever the settings. The module power transducers lag 0.1 s with 2 kW of noise, the jacket and supply RTDs 3 s and 2 s with 0.05 K, the recovery-outlet RTDs 3 s, the four tank sensors 5 s, the grid meter 0.2 s with 3 kW, the UPS power 0.1 s, the chilled water 5 s and the absorption heat meter 10 s. The site demand meter reports the average of each whole second at its end, so a load step reaches the feedforward 1 to 2 s late, the battery management system reports the charge once a second, and the fuel meters read volume at the nominal heating value, so they cannot see a change of gas.

Each case starts at a steady state of its own settings, solved for the data hall, the chillers and the tower by bisection, for each module from the Willans line, and for the recovery chain and the buffer by fixed-point iteration, with the buffer's starting charge as a parameter. The framework's fourth-order Runge-Kutta integration (Section 1.6.1) advances the states every 0.01 s, the controllers scan every 0.05 s, and the historian logs every 0.5 s. Parameters are listed in Table 2.2 with the class of their provenance, and the remaining symbols in Table 2.3.

2.3.7 Assumptions

The engine is an averaged map, a Willans line with fixed heat shares, so knock, emissions, gas composition beyond the heating value, a fully cold start and wear are absent. The value $\varphi_0 = 0.10$ reproduces the catalog's statement of 8 to 10 % less efficiency at half load as a 9.1 % loss [1]. The jacket and the exchanger metal are one node each and the buffer four nodes, so a sharp thermocline is smeared and warm water reaches the buffer bottom a little early. Because the jacket node is the engine outlet, a restricted jacket-water flow widens the rise across the engine rather than heating the outlet, and the hot spots that make low coolant flow dangerous are outside the model. The absorption chiller has no solution chemistry, crystallization or internal heat inventory, and the tower has no water balance, so no water-use claim is made. The island is a power interface that assumes a working voltage source, with no frequency, synchronization or protection coordination. The contracts are scored, not enforced, in base load and electrical following. Every size and threshold is a synthetic teaching value, and real limits come from the equipment manufacturer's data and the applicable codes.

2.4 The Control Layers as Implemented

One framework scan of $\Delta t_c = 0.05$ s runs the instruments, the module power loops, the authority, the override and the allocator. Integer counters inside it run the hot-water, jacket and bypass loops every

Table 2.2 Model Parameters. Class Gives the Provenance (L Literature-Based, F Fitted, T Synthetic Teaching Value) and Whether the Parameters Panel Shows (V) or Hides (H) the Value

Symbol	Meaning	Value	Unit	Class
Engine modules				
P_r	module gross rating	2.0	MW	T, H
η_e, η_{gen}	electrical (LHV), generator	42.0, 97	%	T, V
φ_0	zero-load fuel, share of rated	0.10	—	F, V
κ_j, κ_o	fuel to jacket, to LT circuit	0.20, 0.107	—	T, V
$\dot{m}_{g,r}$	exhaust flow at rating	3.05	kg/s	F, V
$c_{p,g}$	exhaust specific heat	1080	J/(kg K)	T, H
τ_f, τ_m, τ_{dn}	valve, air path, fuel cut	0.3, 2.0, 0.1	s	T, V/H
r, P_{min}	load ramp, minimum load	40, 0.5	kW/s, MW	T, V
P_{par0}, P_{fan}	auxiliaries, dump fans	30, 20	kW	T, V/H
t_{start}	warm start to breaker	90	s	T, V
—	derate above 25 °C	0.18	%/K	L
Heat recovery and protection				
C_j, C_x	jacket node, exchanger metal	8, 2	MJ/K	T, V
$\dot{m}_j$, UA_j	jacket flow, plate exchanger	20, 80	kg/s, kW/K	T, V
G_d, UA_{jl}	dump radiator, surface loss	18, 0.3	kW/K	T, V/H
UA_g, UA_w	exhaust exchanger, gas, water	6.25, 60	kW/K	T, V/H
$T_{s,min}$, $\dot{m}_{r,max}$	minimum stack, recovery flow	120, 18	°C, kg/s	L/T, V
T_{warn}, T_{ov}, T_{trip}	jacket warn, override, trip	98, 100, 105	°C	T, V
Buffer, users and absorption				
V_b, ρ	buffer volume, density	8, 975	m^3, kg/m^3	T, V/H
UA_t	standing loss per node	25	W/K	T, H
$\dot{m}_u$, M_u	distribution flow, holdup	35, 3000	kg/s, kg	T, V/H
T_{ci}, ε_c	customer return, effectiveness	55, 0.8	°C, —	T, V/H
COP	absorption at 87.8, 97.8 °C	0.70, 0.74	—	L
$Q_{abs,r}$, τ_a	absorption rating, lag	1.0, 180	MW, s	T, V
$P_{aux,a}$	absorption auxiliaries	20	kW	T
Data hall and cooling				
C_{IT}, UA_{IT}	IT node, transfer	60, 0.163	MJ/K, MW/K	T, H
C_{chw}, C_{cw}	chilled, cooling water	210, 126	MJ/K	T, H
$Q_{ec,max}$, COP_{ec}	electric chillers	4.0, 5.5	MW, —	T, V
UA_{tw}, P_{aux}	tower, site auxiliaries	0.5, 0.12	MW/K, MW	T, V
UPS and grid				
$P_{dis,max}$, $P_{ch,max}$	UPS discharge, charge	1.5, 0.75	MW	T, V/H
E_{cap}, η	UPS energy, efficiency	0.25, 95	MWh, %	T, V
E_{fl}, E_{ceil}	UPS floor, ceiling	20, 95	%	T, V/H
$P_{imp,max}$, $P_{exp,max}$	import, export contract	1.5, 0.5	MW	T, V

Table 2.3 Notation for the Flows, Temperatures and Bus Terms

Symbol	Meaning	Symbol	Meaning
Q_{jg}, Q_o	fuel heat to jacket, LT circuit	H_{exh}, H_{byp}	exhaust enthalpy, bypassed
Q_{jhx}	jacket heat to recovery water	Q_{dump}, Q_{jl}	dump radiator, surface loss
Q_{gx}, Q_{xw}	gas to metal, metal to water	Q_{pre}	preheater, stopped module
T_g, T_e	exhaust gas, reported value	T_{jm}	jacket water after JKT
T_B, T_{r1}	tank bottom, between exchangers	T_{rec}	recovery water outlet
C_h, C_r, C_u	heat-capacity flows	ε_{cf}, ε_w	exchanger effectiveness
T_k, M_k	buffer node, its mass	T_u	return node
T_{mid}, Q_c	customer inlet, heat	Q_{req}	customer demand
Q_{cap}, Q_{tgt}	absorption capacity, target	Q_{abs}, Q_{drive}	cooling, drive heat
$Q_{rej,a}$	absorption rejection	$T_{g,min}$	lowest generator outlet
$P_{IT,del}$, $Q_{aux,IT}$	delivered IT, auxiliary heat	Q_{ec}, P_{comp}	electric cooling, power
Q_{tw}, f	tower duty, fan	P_{gen}	net generation
P_{imp}, P_{exp}	grid import, export	P_{dis}, P_{ch}	UPS discharge, charge
P_{uns}, P_{curt}	unserved, curtailed	P_{cool}, P_{par}	cooling, stopped modules
P_{dem}	site demand meter	P_{tot}, sp_i	total command, setpoints
π_i, P_{allow}	participation, override cap	P_{rech}	UPS recharge request

0.5 s, thermal following, the chilled water and the UPS reserve every 1 s, and staging every 5 s, each

loop handed the time that actually passed since its own last execution. Every counter, filter, sequencer and PID state is created again on a reset. Table 2.4 lists the loops.

2.4.1 One PID for Every Loop

All of the loops use the framework's velocity-form PID of Section 1.6.2, Equations (1.10) to (1.13), with the action $a = \pm 1$ of Table 2.4, a bias u_b and the back-calculation tracking time T_t, which defaults to τ_I. No loop here uses derivative action or an output slew. In Manual, and for any loop that does not hold authority, `track()` solves the output equation for the integral so that the output equals the applied value, which makes every transfer bumpless, and no PID receives both an update and a track in one scan.

2.4.2 Module Power, Authority and Allocation

Each module's power loop moves its fuel valve on the measured terminal power, with the commissioned engine map as its bias, the valve in percent that the Willans line needs at nominal gas, $u_b = 100\left[\varphi_0 P_{f,r} + sp_i/(\eta_{gen}\eta_m)\right]/P_{f,max}$ clipped to 0–120 %, so the PI only trims. Its tracking time is 2 s by default and 200 s as case 2 ships, and its upper limit falls to the valve the map expects for the override's allowed power while that override derates.

One number, the total engine power P_{tot} in MW, comes from the selected authority. Base load holds the site power target. Electrical following holds the grid import at its target, or islanded the UPS power at minus its recharge request, with the bias $u_b = P_{dem} + P_{par} - SP$ when the metered-load feedforward is on and limits equal to the sum of the modules' limits. When the feedforward is switched on or off, its change of bias is taken out of the integral, so the output does not jump. Thermal following holds the buffer charge from the four tank sensors, with its output clipped to what the grid contract allows, $[P_{dem} + P_{par} - P_{imp,max},\ P_{dem} + P_{par} + P_{exp,max}]$ within the modules' limits. The two controllers that do not hold authority track P_{tot} in the same units, and in Manual both track the measured total of the module powers. An island forces electrical following. If islanded with the UPS at its floor or at 98 % of its rating, every running module goes to its limit and both controllers track, because the demand meter then reads only the load that is served.

The allocator divides P_{tot} among the running modules by their participation factors π_i, inside each module's minimum load and its allowed power, and spills what one cannot take to the others in up to four passes. Each setpoint then moves at most $r\Delta t_c = 2\,\mathrm{kW}$ per scan, and a module handing load to another gives it up only as fast as the other picks it up, so a transfer never takes power off the bus.

2.4.3 Override, Heat Loops, Staging and Reserve

The overtemperature override is protection, separate from the tunable regulation and effective in Manual, as shown in Equation (2.7).

$$P_{allow} = P_r - \mathrm{clip}_{[0,1]}\left(\frac{T_{j,m} - T_{ov}}{T_{trip} - 0.5 - T_{ov}}\right)\left(P_r - P_{min}\right) \tag{2.7}$$

where $T_{j,m}$ is the measured jacket temperature, $T_{ov} = 100\,^{\circ}\mathrm{C}$, $T_{trip} = 105\,^{\circ}\mathrm{C}$ and $P_{min} = 0.5\,\mathrm{MW}$, a slope of 0.333 MW/K. The allowed power falls at once and recovers at the load ramp, and the override flag clears only below 99 °C at full rating. Because it acts on the unfiltered RTD, it turns RTD noise into a slightly lower allowed power. A turbofan's fuel valve carries several such limiters at once, and Volume 4, Chapter 3 shows how a selector lets the most restrictive one win.

The jacket loops open their dump radiators to hold 95 °C at the jacket outlet. The exhaust bypass, also protective and automatic in Manual, opens above a 100 °C recovery outlet. The hot-water loop sets the mixing valve to hold the supply at 90 °C. The chilled-water loop sets a total cooling command that the absorption chiller takes first, with the electric chillers given the rest less the metered absorption duty. The UPS reserve loop, a proportional state-of-charge manager of the kind Section 1.4.4 describes, asks for a recharge of 0.05 MW per point below its target, at most 0.3 MW, and islanded the flexible IT may be curtailed below 35 % of charge and released above 50 %.

Table 2.4 Loops, Their Scan, Default Tuning and Direction. Action +1 Is Reverse Acting, −1 Direct

Loop	PV → MV	Scan (s)	K_c	τ_I	Act.
module power (×3)	$P_{e,i} \to u_i$	0.05	20 %/MW	2 s	+1
electrical following	$P_{grid} \to P_{tot}$	0.05	0.3 MW/MW	10 s	−1
thermal following	$B \to P_{tot}$	1	0.05 MW/%	600 s	+1
hot-water supply	$T_{sup} \to x$	0.5	4 %/K	20 s	+1
jacket (×3)	$T_{j,i} \to d_i$	0.5	10 %/K	60 s	−1
exhaust bypass (×3)	$T_{rec,i} \to b_i$	0.5	10 %/K	20 s	−1
chilled water	$T_{chw} \to Q_{cool}$	1	0.8 MW/K	120 s	−1
UPS reserve (P)	SOC → P_{rech}	1	0.05 MW/%	—	—

Curtailment opens the breakers of the whole flexible share, so the command reads no IT request. The staging supervisor runs every 5 s on the site demand meter. It starts a module when the demand exceeds a start share of the committed capacity and stops one when the others would carry it below a stop share, through the framework's stage sequencer with minimum run and off times and 30 s between moves, and its requests become plant events at the next step boundary. The defaults are shares of 85 and 70 % with 600 s of minimum run and 300 s of minimum off time, and case 8 ships 80 and 75 % with no minimum times.

The original loops keep the plant and their commissioned model apart. The plant is everything in Section 2.3. No model is identified from data. The controllers' original internal model is the commissioned Willans line in the power loops' bias, built from the defaults of Table 2.2 and fixed when the application loads, which assumes nominal gas, and the 70 °C to 95 °C scale of the buffer charge.

2.4.4 A Predictive Power and Dump Supervisor

The grid-connected MPC option, set up as Section 1.6.3 describes, uses an independent aggregate model, with state $x = [P_e, \overline{T}_j, \overline{T}_b]^T$ and commands $u = [P_{tot}, d]^T$. Power is in MW and temperatures in degrees Celsius. The first state follows the command with a 2.3 s surrogate lag. The thermal balances are

$$\begin{aligned} C_j \dot{\overline{T}}_j &= \kappa_j P_f - Q_{jhx} - Q_{dump} - Q_{jl}, \\ C_b \dot{\overline{T}}_b &= n(Q_{jhx} + Q_{xw}) - Q_{water} - Q_{bl}, \end{aligned} \tag{2.8}$$

where n is the number of running modules, $C_b = V_b \rho c_p$, and heat rates are in watts. A homogeneous module uses the declared Willans line and nominal exchanger equations to compute fuel P_f, jacket recovery Q_{jhx}, dump Q_{dump}, surface loss Q_{jl} and quasisteady exhaust recovery Q_{xw}. Tank loss is Q_{bl}. The measured supply-return difference times nominal distribution flow and heat capacity gives Q_{water}. The measured bottom-to-mean tank gradient is held, rather than forecasting a thermocline. Exhaust metal, valve and pump lags and module imbalance are omitted. All physical coefficients remain at their declared defaults. No customer schedule, true exchanger state or altered plant coefficient reaches this model.

Measured power and temperatures initialize all states every 10 s. The model is centrally linearized with its affine term, then exactly discretized. Sixty steps and twelve move blocks give a 600 s horizon in the shared quadratic optimizer, following the finite-horizon approach reviewed in [20]. Bounded thermal rate disturbances correct the previous raw one-step residual with gain 0.2 and limits 0.1 K/s for the jacket and 0.05 K/s for the tank. This is a bias correction, following [21], rather than physical parameter identification. The inner governors, allocation, staging, exhaust protection and trips remain. Islanded operation uses the original electrical authority. Rejected plans hold the handover command, mark the historian and return to tracked site and jacket PIs until recovery.

State tracking weights are 20, 0.03 and 0.15 under electrical or base following, and 0.1, 0.03 and 2 under thermal following. Targets are the selected power or import target, jacket target and $70 + 25B_{SP}$ for mean tank temperature. Move weights are 0.5 and 4, and a linear power price of 0.01 is a nominal

Table 2.5 The Eleven Cases and a Measured Outcome of Each

#	Case and events	Outcome as shipped
1	Manual, grid; learner opens a fuel valve ten points	power 63 % in 2.3 s, recovered heat in 68 s
2	Base load 3 → 6 → 3 MW at 60, 180 s, gas 88 %	valve pinned from 79 s at 1.83 MW
3	Electrical following, IT +0.6 MW at 60 s	import IAE 18.8 MW s (60–180 s)
4	Customer 3.2 → 0.8 MW at 60 s	first rejection 399 s, 2.72 MW at 900 s
5	Customer 3.2 → 2.2 MW at 60 s	electrical against thermal following
6	Island, absorption on, module 3 trips at 60 s	UPS 12.5 s, buffer 45 → 34 % by 900 s
7	Module 1 dump valve stuck at 15 %, 60–900 s	override from 255 s, down to 1.19 MW
8	Island, staging on, IT ±0.9 MW every 6 min	six starts in 30 min
9	Island, module 3 out, IT to 3.5 MW over 60–300 s	UPS floor at 951 s, 32.8 kWh unserved
10	Predictive power and dump, customer loss at 60 s	smaller fuel ledger, larger rejection and import
11	Challenge: customer lost at 30, trip at 120, IT +0.3 MW at 240 s	lowest UPS 52.0 %, fails

fuel proxy. Power commands respect available module and grid-contract limits, with module capacity taking priority if infeasible. Dump bounds 0–1 and planned block moves 0.4 MW and 0.15 are hard. Mean tank 70–98 °C and mean jacket 0.5 K below the override threshold are soft, priced at quadratic 100 and linear 10. Individual module temperatures and actual protective moves are not bounded by these aggregate constraints.

The exact power discretization agrees with its closed form within 1.2×10^{-13}, and a two-control optimizer check agrees with direct normal-equation inversion within 6.7×10^{-16}. In the new case the raw ten-second prediction residual RMS is 0.0322 MW, 0.0378 K and 0.0162 K, respectively. These are local residuals, not horizon-wide accuracy guarantees.

2.5 Matched Experiments, a Capacity Limit and the Challenge

The eleven cases appear in Table 2.5, and Section 1.6.4 describes how a case is loaded and scored. The results come from the shipped engine driven by a script at its default seed. The script loads each case through the same path as the case menu, makes any listed change before the first step, so that the engine treats it as an initial condition, and advances the plant at its own physics step without the browser's frame loop. Integrated absolute error (IAE) is taken on the simulator's truth at the physics step over the stated window, actuator travel is the total variation of the logged command over the same window, and recovery is the last instant the error is outside the stated band, after which it stays inside to the end of the run.

2.5.1 A Small Fuel Reduction With Other Costs

Case 10 repeats the customer loss with MPC and the original PI authorities, with the same seed 20261020, measurements and no event preview. At 900 s, MPC uses 1.9328 versus 1.9414 MWh fuel, a 0.45% reduction, but imports 0.06567 versus 0.050 00 MWh electricity and rejects 0.50885 versus 0.294 01 MWh. Closing tank heat above 0 °C is 0.78519 versus 0.892 63 MWh. This is not a net efficiency gain. Figure 2.2 shows the earlier rejection (dashed) that leaves the store colder. Both hold the supply band without a protective trip. Import IAE is 91.69 versus 2.13 MW s, so the economic/thermal compromise also worsens import tracking. Halving the physics step preserves these orderings.

Total power-command travel is 0.538 versus 18.069 MW, but this is partly the ten-second hold rejecting meter noise. With site-balance gain zero, simple metered-load feedforward achieves 0.006 83 MW travel, 2.135 MW s import IAE and 1.941 35 MWh fuel. Six thermal-PI settings and three fixed powers give further counterfactuals. The full measurement set of 33 rows includes the IT case with feedforward on for both controllers. Table 2.6 keeps the unfavorable comparisons. Under thermal following in case 5, MPC uses 1.661 versus 1.491 MWh fuel. In the stuck-valve case its peak jacket is 99.24 versus 102.30 °C, but import rises to 0.1539 from 0.0501 MWh. Islanded challenge runs are identical because MPC stands by, including the original suggested pass and shipped failure.

Table 2.6 Grid-Connected Comparisons at 900 s, PI / MPC

Case	Fuel MWh	Import MWh	Reject MWh	Plans / µs
New	1.941 / 1.933	.050 / .066	.294 / .509	90 / 1532
Following	1.941 / 1.914	.050 / .063	.011 / .187	90 / 638
IT step	2.263 / 2.261	.057 / .059	.000 / .082	90 / 434
Stuck valve	2.464 / 2.258	.050 / .154	.655 / .642	90 / 594

Plans and mean Node solve microseconds refer to MPC. No plan was rejected in these comparisons.

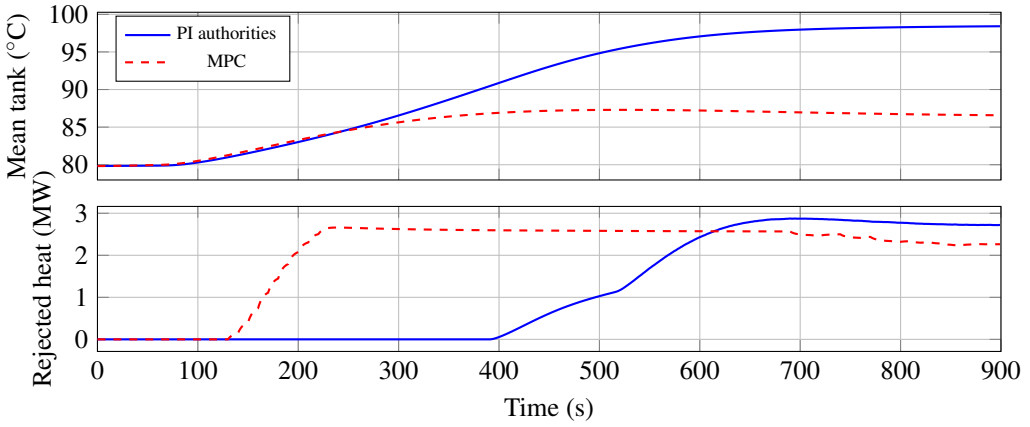

Figure 2.2 Case 10, customer demand reduced at 60 s, under the PI authorities (solid) and the MPC (dashed). The nominal MPC opens rejection earlier, leaving less tank heat. Its small fuel reduction comes with larger electrical import and rejected heat.

Table 2.7 Three Matched Comparisons on the Simulator's Truth. Energies Are Over the 300 s Run, Storage at Its Start and End. Grid-Connected, so No Load Is Unserved or Curtailed and the UPS Holds 80 % Throughout

Comparison, window, band	Setting	IAE (MW s)	Max dev. (MW)	Recovery (s)	Travel	Saturated (s)	Fuel (MWh)	Buffer (%)
A. Case 3, 60–180 s, ±20 kW	feedback only	18.82	0.600	144.4	0.67 MW	0	0.722	43.0 → 48.5
	with feedforward	4.14	0.600	31.4	1.19 MW	0	0.735	43.0 → 49.8
B. Case 2, 60–240 s, ±50 kW	$T_t = 200$ s	30.09	0.772	41.3	81 points	124.5	0.863	43.1 → 61.2
	$T_t = 2$ s	17.97	0.017	4.9	84 points	102.7	0.840	43.1 → 59.8
C. Case 3, 60–180 s, ±20 kW	gas 100 %	4.14	0.600	31.4	1.19 MW	0	0.735	43.0 → 49.8
	gas 88 %	4.40	0.600	39.0	1.18 MW	0	0.735	43.0 → 51.0

A and B compare the settings named, and C runs case 3 with the feedforward on. A and C regulate the true grid import to its 0.2 MW target and B module 1's true power to its setpoint, and recovery is the time to stay inside the band to the end of the run. Every run delivers 0.267 MWh of useful heat. Travel is of P_{tot} in A and C and of module 1's fuel valve in B. In B the maximum deviation is the overshoot above the setpoint after the 180 s drop, and recovery and the band are counted from that drop. Halving the step to 5 ms gives 18.82, 4.14 and 4.40 MW s, and in B overshoots of 0.772 and 0.017 MW with recovery times of 41.3 and 4.9 s.

Across six browser runs at different viewports and themes, sixty plans per run average 1.07–1.33 ms. Planner cost is 0.011–0.013% of one core at real time and 1.07–1.33% at 100-times playback. Measurement and state preparation are outside this timing.

2.5.2 Three Matched Comparisons

Each pair in Table 2.7 shares the initial state, the seed, the measurements, the limits and the disturbance, and differs in one setting. The trajectories are shown in Figure 2.3.

The first comparison is nominal disturbance rejection. In case 3 the IT load rises 0.6 MW at 60 s

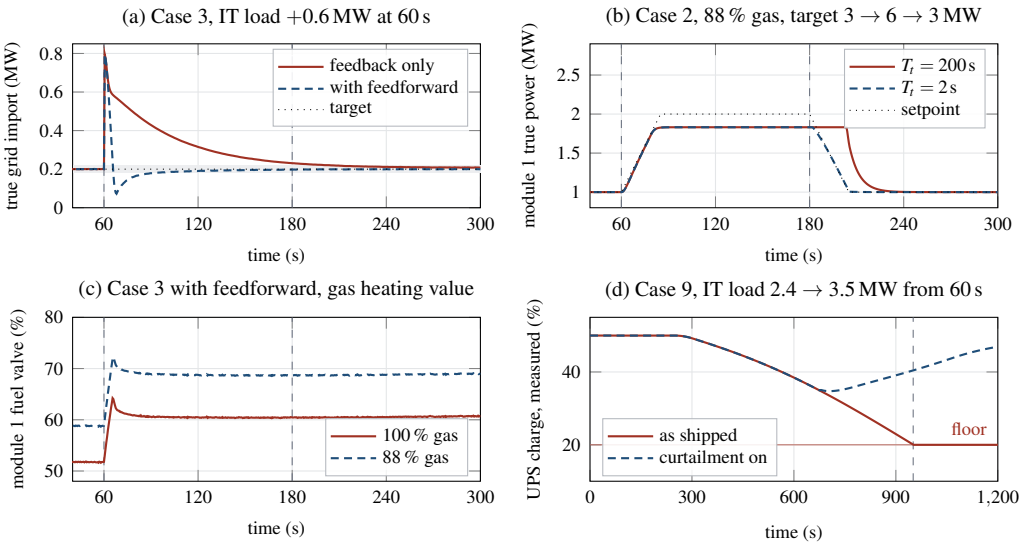

Figure 2.3 The three matched comparisons and the capacity limit, from the exported records. (a) Case 3, true grid import after the 0.6 MW IT step at 60 s with the ±20 kW band shaded, feedback only (solid) and with the metered-load feedforward (dashed). (b) Case 2, module 1's true power on 88 % gas against its ramped setpoint, tracking time 200 s (solid) and 2 s (dashed). (c) Case 3 with the feedforward, module 1's fuel valve at nominal (solid) and 88 % gas (dashed). (d) Case 9, measured UPS charge as shipped (solid) and with curtailment (dashed), the floor reached at 951 s. Dashed verticals mark the events and the IAE windows.

and the full 0.60 MW lands on the grid import in both runs, because nothing measures a load before it arrives. On feedback alone, with $K_c = 0.3$ and $\tau_I = 10$ s, the import stays more than 20 kW off its 0.2 MW target for 144 s. With the demand meter fed forward, the engines take the step once the meter reports it 1 to 2 s later, and the IAE falls from 18.82 to 4.14 MW s, 22 percent of the feedback-only value. Volume 3, Chapter 6 makes the same case for measuring a disturbance before it reaches the loop. The feedforward is not free. Its step and the proportional action both answer the same error, so the true import overshoots to 0.073 MW at 68 s, 0.127 MW below its target, and the total command travels 1.19 against 0.67 MW.

The second comparison is an actuator limit. On gas with 88 % of its heating value, a base-load target of 6.0 MW from 60 s pins module 1's valve at 100 % from 79 s, where the module makes 1.83 MW against a 2.0 MW setpoint. With the shipped tracking time of 200 s the integral holds 140 points of valve when the target falls back at 180 s, the valve stays open 23 s, and the module rides 0.77 MW above its falling setpoint and needs 41 s to come within 50 kW. With 2 s the integral holds 4.8 points, the valve leaves 100 % after 1.8 s, the overshoot is 0.017 MW and the recovery 4.9 s. The shortfall while pinned is 0.174 MW in both runs, because a shorter tracking time adds no capacity. Base load does not enforce the export contract, so the site exports up to 2.34 MW against a contract of 0.5 MW in both runs, which the scene marks.

The third comparison is a model mismatch. The power loops' map and the fuel meters assume nominal gas, so at 88 % gas each loop starts with 7.1 points of valve in its integral that the map does not supply. The site barely notices, with an import IAE of 4.40 against 4.14 MW s and a recovery of 39 against 31 s, because the integral absorbs the mismatch as long as the valve has room, as shown in Figure 2.3c. The same mismatch becomes windup only when the valve saturates, which is the second comparison.

2.5.3 Where Control Stops and Capacity Begins

Case 9 has module 3 out for maintenance, the UPS at 50 %, and the IT load rising from 2.4 to 3.5 MW over four minutes from 60 s. Two modules net 3.94 MW after their auxiliaries, and the UPS's usable energy from 50 % to its floor is $0.25 \times 0.30 \times 0.95 = 0.071$ MWh. Both modules reach their rating at 243 s, the UPS gives up to 0.49 MW and reaches its 20 % floor at 951 s, and 32.8 kWh of flexible IT goes unserved by 1200 s, with none of the critical IT. Two engines at full net capacity throughout would empty the UPS at 955 s, so the shipped settings lose 4 s to their own response. A sweep of the feedforward on and off, site-balance tunings of 0.1 with 100 s, 0.3 with 10 s and 3 with 2 s, and two participation splits reaches the floor between 491 and 951 s in all twelve runs, never later than that bound, with 32.8 to 79.6 kWh unserved, as shown in Figure 2.3d. Changing the gain once both modules are at their rating leaves the floor at 950.90 s exactly, because the loop's output is already pinned. Only less load keeps a reserve. With curtailment, 102 kWh of flexible work is curtailed, nothing goes unserved and the UPS keeps 34.7 %.

Pause. Try Exercise 2.2 before reading on.

2.5.4 The Remaining Cases

Case 1 separates the two clocks within one module. The exhaust part of the recovered heat covers 81 % of its change within a minute, because only the 2 MJ/K of exchanger metal must warm, while the jacket part covers 39 %, because 8 MJ/K of block and coolant must warm first. The buffer goes from 43 to 55 % in fifteen minutes and the grid import from 0.20 MW to an export of 0.03 MW within seconds. In case 4 the customer's cut from 3.2 to 0.8 MW fills the buffer from 43 to 81 % before the first heat is rejected at 399 s, through the exhaust bypasses first and the dump radiators from 519 s, and by 900 s 2.72 MW is rejected, 1.38 MW through the radiators and 1.34 MW around the exchangers, while the engines' 2.955 MW moves by at most 6 kW.

> **TRY IT IN THE APP** **1 · A Fuel Step: Which Arrives First, Power or Heat?**
>
> Load scenario 1 and press Run. After 30 s open module 1's fuel valve by about ten points and watch its power and its recovered heat on the trend. The power covers 63 percent of its change in 2.3 s, and the recovered heat needs 68 s for the same share. The buffer keeps climbing, from 43 to 55 % in fifteen minutes.

Pause. Try Exercise 2.1 before reading on.

Case 5 compares the two authorities after the customer drops to 2.2 MW. Over 1800 s electrical following holds the import at 0.20 MW, burns 3.88 MWh of fuel, imports 0.10 MWh, rejects 0.29 MWh and leaves the buffer at 89 %. Thermal following is held by the 1.50 MW import contract from 481 s, so the buffer stays at 51 to 55 %, above its 45 % target. It burns 2.74 MWh, imports 0.63 MWh and rejects nothing. Both deliver 1.12 MWh of useful heat, at CHP efficiencies on LHV of 65.7 and 73.6 %, but the two runs end with different storage, the buffer gaining 0.146 MWh under electrical following and 0.036 MWh under thermal following, and with 0.53 MWh more electricity from the grid in the second. The pair demonstrates who holds authority. It is not a savings estimate.

In case 6 the UPS takes the whole 1.0 MW at module 3's trip and stops discharging 12.5 s later after 1.9 kWh, and modules 1 and 2 settle at 1.51 MW each. The two recover 2.91 MW at 900 s against the 3.28 MW three recovered, the tripped module's retained heat fades from 0.76 MW at 90 s to 0.16 MW at 300 s, and the absorption chiller keeps drawing 1.4 MW of drive heat, so the buffer falls from 45 to 39 and 34 % at 600 and 900 s while the supply stays within 0.15 K of 90 °C. In case 7 module 1's jacket passes its warning at 151 s and the override threshold at 255 s, the override cuts its allowed power to as little as 1.19 MW and holds the jacket below 102.3 °C, modules 2 and 3 carry 1.35 MW each against module 1's 1.23 MW at 800 s, and the override clears 37 s after the valve is freed. In case 8 the start and stop thresholds of 3.2 MW and 3.0 MW sit 0.2 MW apart on a metered demand that swings 1.85 MW from peak to trough, so the reserve starts six times in 30 min and the

Table 2.8 The Challenge Over 900 s. Critical IT Unserved and Protective Trips Are Zero in Every Row

Configuration	UPS min (%)	Out (s)	Override (s)	T_j peak (°C)	UPS gave (kWh)	Verdict
shipped, 0.5/1200, 1/300	51.97	570	408	102.0	68.9	fails
suggested, 10/60, 4/20	79.19	0	0	96.8	2.2	passes
jacket loop tuned only	79.19	570	0	96.7	2.2	fails
supply loop tuned only	44.40	0	432	102.2	86.9	fails

Settings: jacket K_c in %/K over τ_I in s, then supply K_c over τ_I. Out is the time the supply spent outside ±3 K. T_j peak is the true jacket maximum of modules 1 and 2.

UPS carries up to 0.21 MW through each 90 s start. A 50 % stop share with a 600 s minimum run gives one start, the stop share alone one, and the minimum run alone three.

2.5.5 The Challenge

The challenge runs islanded with the absorption chiller on, three modules at 1.006 MW, the buffer at 67 % and the UPS at 80 %. The customer is lost at 30 s, module 3 trips at 120 s and the IT load rises 0.30 MW at 240 s. Over 900 s the target asks for no critical IT unserved, no protective trip, the supply within ±3 K of 90 °C outside a 90 s allowance after each event, and the UPS never below a 60 % reserve. After the trip and the rise the demand is 3.29 MW, and 3.47 MW even if the electric chillers took all of the absorption cooling, below the two remaining modules' 3.94 MW, so the target is feasible. Table 2.8 lists four configurations, and Figure 2.4 compares the shipped and the suggested runs. Figure 2.5 charts the table.

> ✦ **DID YOU KNOW?** **Tuning One Loop Can Make Things Worse**
>
> From the simulator: in the challenge, tuning only the hot-water supply loop keeps the supply in its band but lets the lowest UPS charge fall to 44.40 %, below the 51.97 % of the shipped tuning, with 432 s of override (Table 2.8). Only the two loops tuned together pass.

As shipped, the jacket loops run at 0.5 %/K with a 1200 s integral time and the supply loop at 1 %/K with 300 s. For eight minutes nothing looks wrong, because the recovery water still comes from cool water at the bottom of the buffer. Once the warm water reaches the bottom, the two remaining engines' jackets pass 100 °C at 492 s, the override takes capacity for 408 s, the UPS makes up 68.9 kWh and falls to 51.97 %, and the slow supply loop leaves the supply up to 5.7 K high for 570 s. The suggested tuning, 10 %/K with 60 s on the jackets and 4 %/K with 20 s on the supply, keeps the jackets at or below 96.8 °C with no override, the UPS at 79.19 % and the supply within 0.63 K. Each loop guards a different part of the target, so either one tuned alone fails the other part. The override's response to RTD noise moves the shipped result. With the RTD noise removed the lowest UPS charge is 52.91 %, at the shipped 0.05 K it is 51.97 % and at three times that 49.62 %, while the suggested run stays at 79.19 % in all three because it never reaches the override. At seeds 1 to 5 the shipped run reaches 51.97 to 52.00 % and fails, and the suggested one reaches 79.19 % and passes. The shipped run at its end is shown in Figure 2.6.

> **TRY IT IN THE APP** **11 · Challenge: A Customer Loss, a Trip and an IT Rise, Islanded**
>
> Load scenario 11 and run it as shipped at 100 times speed. Watch the jackets on Heat recovery. For eight minutes nothing looks wrong, then the jackets pass 100 °C at 492 s and the UPS falls to 51.97 %. Set the jacket loops to 10 %/K and 60 s and the supply loop to 4 %/K and 20 s, then Reset and run. The UPS holds 79.19 %.

Pause. Try Exercise 2.3 before reading on.

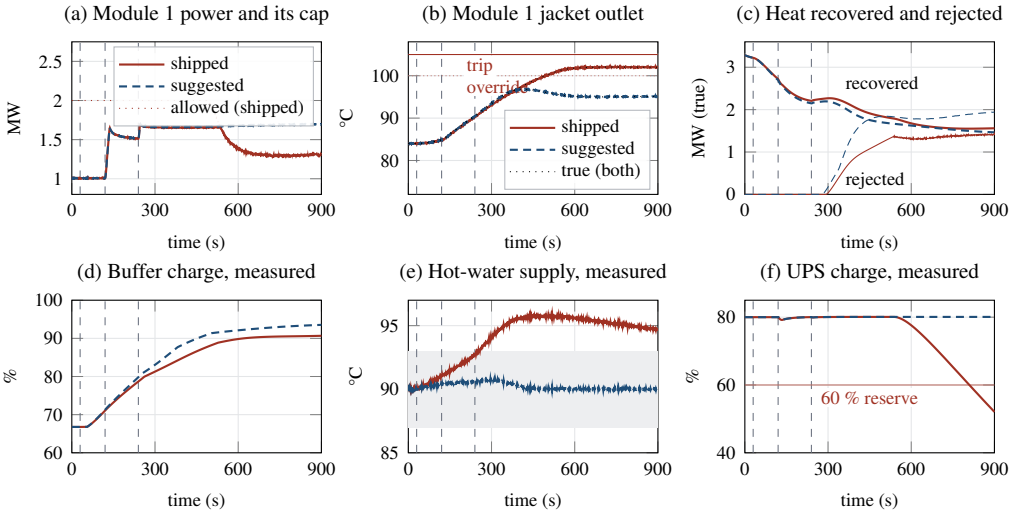

Figure 2.4 The challenge as shipped (solid) and with the suggested tuning (dashed), from the exported records. (a) Module 1's measured power and, dotted, the override's allowed power in the shipped run. (b) Module 1's jacket outlet, measured (solid and dashed) and true (dotted), against the override and trip thresholds. (c) True heat recovered and rejected. (d) Measured buffer charge. (e) Measured supply with the ±3 K band shaded. (f) Measured UPS charge against the 60 % reserve. Gray dashed lines mark the events at 30, 120 and 240 s.

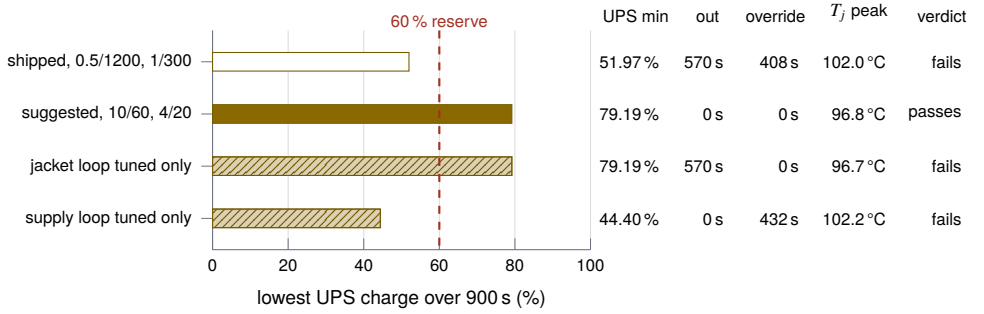

	UPS min	out	override	T_j peak	verdict
shipped, 0.5/1200, 1/300	51.97 %	570 s	408 s	102.0 °C	fails
suggested, 10/60, 4/20	79.19 %	0 s	0 s	96.8 °C	passes
jacket loop tuned only	79.19 %	570 s	0 s	96.7 °C	fails
supply loop tuned only	44.40 %	0 s	432 s	102.2 °C	fails

Figure 2.5 The four configurations of Table 2.8 as a chart: the lowest UPS charge over the 900 s (bars and first column), with the time the supply spent outside ±3 K, the time under override, the jacket peak and the verdict in the other columns. Settings are the jacket gain in %/K over its integral time in s, then the supply's. The white bar is the shipped tuning, the solid bar both loops tuned, hatched bars one loop tuned. Notice that tuning only the supply loop leaves less reserve than the shipped tuning, and that only the solid bar both clears the dashed reserve line and keeps the supply in its band.

2.5.6 Verification

These figures verify the code and are not control-performance results. Over a 20-minute base-load run of one module the integrated fuel equals electricity, losses, recovered and rejected heat and the change in stored heat to 5.0×10^{-16} of the fuel, and through case 6's trip with a customer loss the engines' and the hot water's ledgers close to 2.0×10^{-14} of their throughput, each from independent ledger states. The bus balances to 0 W at every one of the challenge's 90,000 steps, with import and export, and charge and discharge, never both nonzero. At case 3's start the recovered heat exceeds the users' heat by 0.25 % of rated fuel input, the 36 kW by which the default demand falls short of

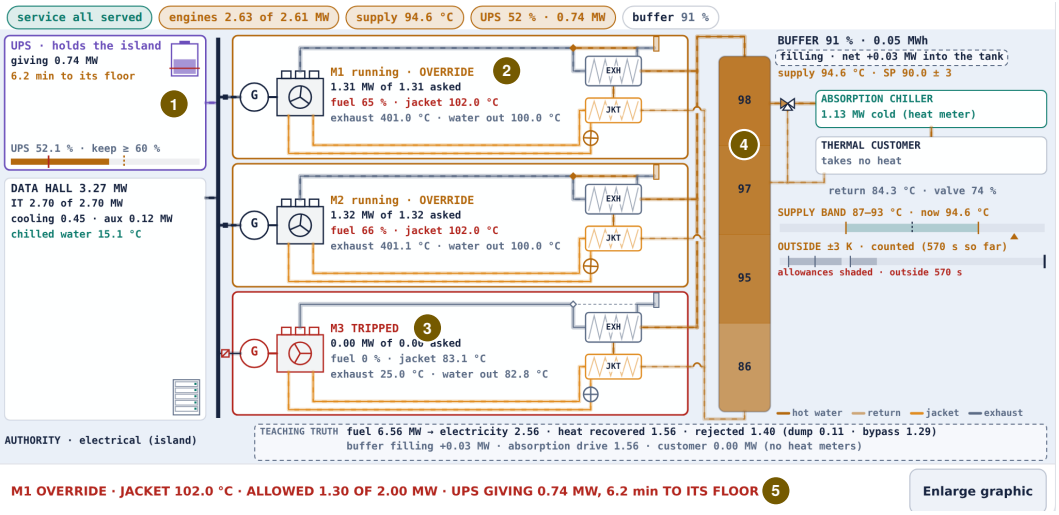

Figure 2.6 The challenge as shipped at $t = 900$ s on the Plant view, cropped to the plant. (1) The UPS stands at 52.1 %, giving 0.74 MW on its meter with 6.1 min to its floor. (2) Modules 1 and 2 run in OVERRIDE at 1.31 MW and 1.32 MW with their jackets at 102.0 °C, and (3) module 3 has tripped. (4) The buffer is 91 % charged with the customer taking no heat and 1.40 MW rejected, 1.29 MW of it through the exhaust bypasses, and the supply reads 94.6 °C against a ±3 K band. (5) The status line names module 1's override and its allowed 1.30 MW.

recovery. Halving the step to 5 ms changes no scored output of any case by more than 0.008 % of its span, a run replays bit for bit after a reset, and playing it at 1 or 100 times real time logs identical rows. A simulated hour costs 2.3 s of computing in a headless engine, and in a browser at the top speed of 100 times real time the challenge takes 9.0 s and the 1800 s staging case 18.0 s. All 72 of the application's checks pass. Agreement with reference data is a separate and narrower claim. The absorption map passes through the DOE fact sheet's two single-stage points exactly [2], and the Willans line reproduces the catalog's half-load statement as a 9.1 % loss [1], while no part of the plant has been compared with measurements of a real one.

2.6 Technology in Practice

The layers of Section 2.4 run on power meters, temperature sensors, fuel valves and water valves. Figure 2.7 draws the loops of module 1 and the site with that hardware in place of the schematic's symbols.

What Measures It

The module powers and the grid import are electrical measurements. Potential transformers step the line voltage down to a safe level, typically 120 V, and current transformers step the line current down, typically to 5 A at full load, for the wattmeter and the protective relays to read [22]. Standard accuracy classes of potential transformers are 0.3, 0.6 and 1.2, an uncertainty of ±0.3, ±0.6 and ±1.2 % of the rated turns ratio. A current transformer must never have its secondary opened while energized, because it may then develop an extremely high voltage [22]. The chapter's module power transducers JT-101 to JT-103 lag 0.1 s with 2 kW of noise, and the grid meter JT-100 lags 0.2 s with 3 kW.

plant controller: JIC-100, JY-100, JIC-101 and TY-111 every 0.05 s, TIC-111 and TIC-120 every 0.5 s. Protective relays trip the breakers apart from it.

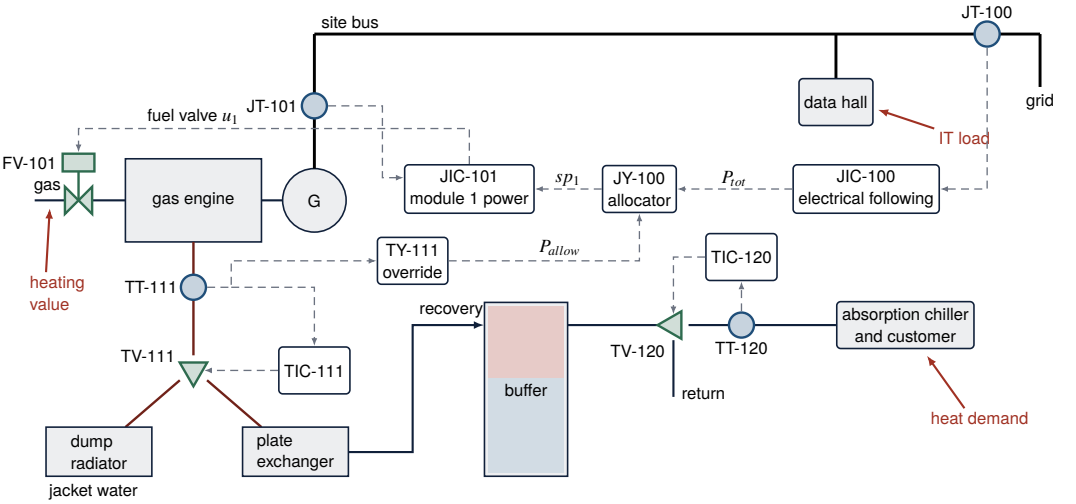

Figure 2.7 Module 1's loops and the site loop drawn with their hardware. The power transducers JT-100 at the grid and JT-101 at the generator read power through instrument transformers. The electrical-following controller JIC-100 sets the total power, the allocator JY-100 shares it, and JIC-101 moves the fuel valve FV-101. The jacket RTD TT-111 feeds both the jacket loop TIC-111, which opens the dump valve TV-111 toward the radiator, and the override TY-111, which caps the module's allowed power. TIC-120 holds the supply with the mixing valve TV-120. The exhaust exchanger and its bypass are left out. Notice that one jacket reading serves both a slow loop and a fast limit on the power.

> **TECHNOLOGY** **Sensor: A Resistance Thermometer in the Jacket Water**
>
> The jacket temperature that the override reads comes from a resistance temperature detector (RTD), a fine wire of pure metal, usually platinum, whose resistance always rises with temperature. It is relatively insensitive but very linear, and 100 Ω at 0 °C is a common reference resistance [22]. It sits in a thermowell, a pressure-tight sheath that adds a first-order lag, and an air gap at the bottom of the well adds a second [22]. A transmitter that drives its current loop below 3.6 mA or above 21 mA reports a failed sensor under the NAMUR signal levels [22]. The chapter's jacket RTDs lag 3 s with 0.05 K of noise.

What Moves It

The electrical side is moved by the fuel. The simplest gas engines draw air and fuel through a carburetor or another mixer, and high performance engines are turbocharged [1]. The chapter folds that fuel system into one valve FV-101 per module with a 0.3 s lag and an air path of 2 s. The heat side is moved by water valves on the jacket and supply circuits.

TECHNOLOGY Final Element: Three-Way Valves on the Water Circuits

A three-way valve has one stem and three ports. At either end of its travel the stem joins the common port to one of the others, and between the ends all three share the flow, so one body divides a stream or blends two, as a diverting or a mixing valve [22]. On an engine an excess heat exchanger sends the jacket heat that the users do not take to a cooling tower or a radiator, because an engine must be cooled even when its heat is not used [1]. The chapter's supply valve x is such a mixing valve, and each dump valve d_i sends jacket heat to its radiator.

Where the Controller Runs

Programmable logic controllers (PLCs) manage local processes by feedback and serve as local controllers under a supervisory scheme, and loops at both levels run at cycle times from milliseconds to minutes [23]. Beside them, protective relays, named by standard function numbers such as 50 for instantaneous overcurrent and 87 for differential current, trip the circuit breakers on their own [22]. Packaged plants that can island leave the grid through an automatic transfer switch and carry critical loads during an outage [1]. The chapter's plant controller scans every 0.05 s for the power loops, the authority, the allocator and the override TY-111, every 0.5 s for the jacket and supply loops, and every 5 s for staging. Its override is a protection layer kept apart from the tunable regulation, as relays are on a real site.

IN PRACTICE Keeping an Engine Plant in Service

An engine plant runs on its maintenance schedule. Oil, filters, coolant and spark plugs are inspected or replaced at short intervals, typically 500 to 2,000 hours, and an analysis of the oil tracks engine wear. A top-end overhaul, a rebuild of the cylinder heads and the turbocharger, falls between 8,000 and 30,000 hours, and a major overhaul with new pistons and liners after 30,000 to 72,000 hours [1]. Fuel quality is watched as well. Engine fuel is delivered by volume, so a gas with a lower heating value needs more volume for the same energy and forces a derating [1], the weaker gas of Section 2.5.2.

What Is Modernizing It

Three developments touch the loops of this chapter. Electronic engine control and better combustion chambers, including precombustion chambers, let engines run on leaner mixtures, and a federal development program set goals of 50 % brake thermal efficiency and more than 80 % CHP efficiency for the next engines [1]. Many service contracts now include remote monitoring of engine performance and allow predictive maintenance [1], so the jacket temperatures and powers of this chapter become data for a fleet. The same connectivity has a cost. Ethernet, Internet Protocol and wireless devices replacing older proprietary networks raise the likelihood of cybersecurity vulnerabilities and incidents in control systems [23]. The predictive supervisor of Section 2.4.4 is a fourth development, and the chapter measures both its gains and its costs.

2.7 Limitations

The sizes, maps and thresholds are synthetic teaching values set inside the published ranges of the EPA catalog and the DOE fact sheet [1, 2], not fitted to any engine, chiller or site. Only φ_0 and the rated exhaust flow are fitted, the first to a sentence of the catalog and the second scaled from its table of example systems. The jacket, the exchanger metal and the buffer are lumped, so the model gives lags of the right size rather than temperature profiles, and the challenge's lesson depends on the small buffer and jacket. With a 15 m^3 buffer and a 14 MJ/K jacket the shipped run still reaches the override, at 758 s for 142 s, but keeps 77.81 % of the UPS charge and fails on the supply alone, and with only one of the two enlarged it keeps 70.55 or 61.75 %, while the suggested tuning passes in all four combinations. There is no rotor, so the island shows power and energy limits rather than frequency, stability or a transfer without interruption. The swing equation of Section 1.4.2 and the island frequency of Chapter 7 are what such a model leaves out.

OPEN CHALLENGE **Hydrogen in the Pipeline Gas**

Gas utilities are beginning to blend hydrogen into their pipelines, and engines that burn pipeline gas will receive the blend [24, 25]. Hydrogen carries less energy per unit volume, so at a 20 percent blend the volume of fuel must rise by 15 percent to deliver the same energy. In tests on a stoichiometric engine whose controller had no warning of the change, a sudden step in hydrogen swung the air-fuel ratio and released a burst of nitrogen oxides, and better tuned controllers released less [24]. Fuels above about 5 percent hydrogen by volume may need special measures [1]. The chapter's power loops meet a weaker gas only as model mismatch. Progress would be fuel control that follows a changing gas without being told.

The original controllers read declared measurements and use the commissioned engine map, so weak gas or an edited engine parameter remains model mismatch. The added MPC has its own nominal aggregate thermal model, with measured states and bounded rate correction. It lacks exhaust-metal dynamics, evolving stratification and module imbalance, and its soft mean limits are not individual temperature guarantees. The protection layer is a teaching guardrail, not certified trip logic, and the load-shed priority is fixed rather than set by a site study. No study with learners has been made, so this chapter claims no measured improvement in learning. The CHP efficiency uses net electricity plus useful heat over LHV fuel at a teaching boundary, and a simulated pass is neither equipment qualification nor a savings or reliability estimate.

OPEN CHALLENGE **Data Center Load That Can Wait**

In the reference case of a national laboratory study, data centers take a third of the national growth in electric load from 2024 to 2030. The study names the work still to be done. The variation of a data center's demand over time should be measured, to learn how such loads can be served by the grid and by on-site generation, and ways to shift the timing of the load should be explored, so that a given set of energy resources serves more computing [3]. The chapter's flexible IT share, curtailed by a fixed rule on the UPS charge, is a first form. Progress would be a supervisor that schedules engines, heat store and flexible computing together against a demand forecast.

2.8 Conclusion

This chapter studied a gas-engine CHP plant at a data center as a laboratory in which one fuel valve per engine serves an electrical load that answers in seconds and a thermal load that answers in minutes. A ten-point fuel step moves a module's power 63 percent of the way in 2.3 s and its recovered heat in 68 s. Metered feedforward cuts the grid-import IAE from 18.82 to 4.14 MW s, a short tracking time cuts a windup overshoot from 0.772 to 0.017 MW, and a gas 12 % weaker than the map assumes costs only 4.40 against 4.14 MW s while the valve has room. When demand exceeds two engines, no tuning in a twelve-run sweep empties the UPS later than the 955 s that capacity allows. In the challenge a slow cooling loop that the buffer hides for eight minutes costs 408 s of override and 28 points of UPS reserve, and the suggested tuning keeps 79.19 %. The predictive supervisor, which plans from current measurements, shows why a smaller fuel ledger can come with more import, more rejection and less stored heat.

The model is lumped and its operating points are synthetic. A fuller model would add an averaged frequency mode on the island, a distributed buffer, a fitted part-load map, a cold start, and a supervisory layer that schedules the buffer against a known demand forecast, the kind of dispatch that Section 1.4.6 sets out. The simulator shows that heat a plant cannot store or reject comes back as lost electrical capacity.

KEY POINTS

- One fuel valve per engine sets both the electricity and the heat, and the electricity answers in seconds while the heat takes minutes.
- Metered-load feedforward cuts the grid-import IAE from 18.82 to 4.14 MW s, and a short tracking time cuts a windup overshoot from 0.772 to 0.017 MW.
- When demand exceeds two engines, no tuning in a twelve-run sweep empties the UPS later than the 955 s that capacity allows.
- In the challenge a slow cooling loop that the buffer hides for eight minutes costs 408 s of override and 28 points of UPS reserve.

Where else this applies. A store that carries the load while its source catches up also serves the fuel cell of Chapter 3, and a limiter that takes a fuel valve away from its loop governs Volume 4, Chapter 3.

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, 2017. [Online]. Available: https://www.energy.gov/sites/default/files/2017/06/f35/CHP-Absorption%20Chiller-compliant.pdf

[3] S. J. Smith, A. Hubbard, A. Newkirk *et al.*, "United states data center energy usage report: 2025 update," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-2001758, 2026.

[4] W. Tan, H. J. Marquez, T. Chen, and J. Liu, "Analysis and control of a nonlinear boiler-turbine unit," *Journal of Process Control*, vol. 15, no. 8, pp. 883–891, 2005.

[5] D. W. Wu and R. Z. Wang, "Combined cooling, heating and power: A review," *Progress in Energy and Combustion Science*, vol. 32, no. 5–6, pp. 459–495, 2006.

[6] P. J. Mago, N. Fumo, and L. M. Chamra, "Performance analysis of CCHP and CHP systems operating following the thermal and electric load," *International Journal of Energy Research*, vol. 33, no. 9, pp. 852–864, 2009.

[7] D. Haeseldonckx, L. Peeters, L. Helsen, and W. D'haeseleer, "The impact of thermal storage on the operational behaviour of residential CHP facilities and the overall CO2 emissions," *Renewable and Sustainable Energy Reviews*, vol. 11, no. 6, pp. 1227–1243, 2007.

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

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

[10] H. I. Onovwiona, V. I. Ugursal, and A. S. Fung, "Modeling of internal combustion engine based cogeneration systems for residential applications," *Applied Thermal Engineering*, vol. 27, no. 5–6, pp. 848–861, 2007.

[11] A. Rosato and S. Sibilio, "Energy performance of a micro-cogeneration device during transient and steady-state operation: Experiments and simulations," *Applied Thermal Engineering*, vol. 52, no. 2, pp. 478–491, 2013.

[12] 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.

[13] P. W. Willans, "Economy trials of a non-condensing steam-engine: simple, compound and triple," *Minutes of the Proceedings of the Institution of Civil Engineers*, vol. 93, pp. 128–188, 1888.

[14] L. Guzzella and C. H. Onder, *Introduction to Modeling and Control of Internal Combustion Engine Systems*. Berlin, Heidelberg: Springer, 2010.

[15] U.S. Congress, "Public utility regulatory policies act of 1978, Public Law 95-617, 92 Stat. 3117," United States Statutes at Large, vol. 92, govinfo.gov, Nov. 1978. [Online]. Available: https://www.govinfo.gov/content/pkg/STATUTE-92/pdf/STATUTE-92-Pg3117.pdf

[16] R. F. Hirsh, *Power Loss: The Origins of Deregulation and Restructuring in the American Electric Utility System*. Cambridge, MA: MIT Press, 1999.

[17] H. I. Onovwiona and V. I. Ugursal, “Residential cogeneration systems: review of the current technology,” *Renewable and Sustainable Energy Reviews*, vol. 10, no. 5, pp. 389–431, 2006.

[18] 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.

[19] R. K. Shah and D. P. Sekulić, *Fundamentals of Heat Exchanger Design*. Hoboken, NJ: Wiley, 2003.

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

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

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

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

[24] N. Katsampes, D. Montgomery, G. Arney, and D. B. Olsen, “Hydrogen-natural gas fuel blending in a “rich burn” engine with 3-way catalyst,” *Frontiers in Fuels*, vol. 2, p. 1416716, 2024.

[25] G. Vieira, R. Lorenzen, M. Patterson, and D. Olsen, “Methane emission reduction through hydrogen blending in a large bore 2-stroke lean-burn natural gas compressor engine,” *Frontiers in Fuels*, vol. 2, p. 1404367, 2024.

2.9 Exercises

2.1 Concept check · 2 min

In case 1 a ten-point step of module 1’s fuel valve moves its electrical output 63 percent of the way in 2.3 s, while its recovered heat needs 68 s for the same share. What sets the slower clock?

(a) The jacket node and the exchanger metal, which must warm before more heat passes to the water

(b) The recovery pump, which cuts its flow while the fuel valve is moving

(c) The air-path lag, which delays the exhaust enthalpy as long as the heat takes

(d) The four buffer nodes, whose mixing smears the heat before it is recovered

2.2 Concept check · 1 min

In case 9 both remaining modules reach their rating at 243 s, and the UPS supplies the rest until it reaches its floor at 951 s. The site-balance gain is raised once both modules are at their rating. What happens?

(a) The UPS reaches its floor later, because a faster loop returns the deficit to the engines sooner

(b) The UPS reaches its floor at the same instant, because the loop’s output is already pinned

(c) The UPS reaches its floor sooner, because the higher gain makes the module power cycle

(d) The UPS never reaches its floor, because the integral recovers the energy lost earlier

2.3 Concept check · 2 min

In case 4 the customer’s heat demand falls from 3.2 to 0.8 MW at 60 s under electrical following. Why is no heat rejected for more than five minutes after the cut?

(a) The jacket loops are slow, so the dump radiators open only after their integrals have grown

(b) The exhaust bypasses are held shut until the buffer charge passes 85 %

(c) The recovery water still comes from cool water at the buffer bottom, so the heat goes to the buffer

(d) The engines cut their fuel under electrical following, so less heat is recovered for a while

2.4 Theory and calculation · 6 min

At rating a module burns 4.762 MW of fuel (LHV) for 2.062 MW of shaft power. Its zero-load fuel is $\varphi_0 = 0.10$ of the rated fuel and its generator efficiency is 97 %.

(a) Find the marginal shaft efficiency η_m of the Willans line in Equation (2.1). *Checkpoint:* $\eta_m = 0.4811$.

(b) Find the electrical efficiency at rating. *Checkpoint:* 42.0 %.

(c) Show that at a share s of rated shaft power the efficiency relative to the rated value is $s/\big(\varphi_0 + s(1-\varphi_0)\big)$, and evaluate it at half load. *Checkpoint:* 90.9 %, a loss of 9.1 %.

2.5 Theory and calculation · 6 min

The buffer holds 8 m^3 of water at 975 $\mathrm{kg/m}^3$ with $c_p = 4186\,\mathrm{J/(kg\,K)}$, and its charge counts the band from 70 to 95 °C.

(a) Find the energy the buffer stores across the band. *Checkpoint:* 816 MJ, or 0.227 MWh.

(b) The modules recover 3.24 MW at the design point. How long could the buffer absorb that heat with no user taking any? *Checkpoint:* 252 s, about four minutes.

(c) At the start of case 3 the four nodes read 91.5, 91.5, 68.1 and 68.1 °C. Find the buffer charge B with each node's share clipped to the unit interval, and without the clip. *Checkpoint:* 43 % clipped, 39.2 % without.

2.6 Theory and calculation · 8 min

The UPS stores 0.25 MWh, discharges with an efficiency of 95 % and stops at a 20 % floor. The override follows Equation (2.7) with $P_r = 2.0$ MW, $P_{min} = 0.5$ MW, $T_{ov} = 100\,^{\circ}\mathrm{C}$ and $T_{trip} = 105\,^{\circ}\mathrm{C}$.

(a) Starting from 80 %, find the energy the UPS can deliver before its floor and how long it lasts at a constant 1 MW. *Checkpoint:* 0.1425 MWh and 513 s.

(b) How long does the same reserve last at the 1.5 MW rating? *Checkpoint:* 342 s.

(c) Find the slope of the override and the allowed power at a measured jacket of 102.0 °C. In the challenge at 900 s both running modules read about 102.0 °C while their allowed power is near 1.30 MW. Explain the difference. *Checkpoint:* 0.333 MW/K and 1.333 MW.

2.7 App challenge · 8 min

Load scenario 3, An IT load step: feedback alone, or metered feedforward, which ships with Metered-load feedforward Off. Run to 300 s at 20x with the default seed and export the CSV file. Then switch Metered-load feedforward On, press Reset, run to 300 s again and export a second file. Use the column `truth_grid_import_MW` and the rows up to 300 s.

(a) In each file find the last row after the 60 s step at which the true import is more than 20 kW from its 0.2 MW target. Give the recovery time after the step, within 1 s. *Checkpoint:* 144 s on feedback alone and 31 s with the feedforward.

(b) In the feedforward run, find the lowest true import after the step and the time it occurs, within 0.005 MW and 1 s. *Checkpoint:* 0.073 MW at 68 s.

(c) Both runs show the import near 0.8 MW in the row at 60.5 s. Explain why the feedforward cannot prevent this, and why the meter's delay of 1 to 2 s costs little. *Checkpoint:* Nothing measures the load before it lands; the loop alone needs minutes.

2.8 App challenge · 8 min

Load scenario 2, Tune a module power loop: a pinned valve and windup. It ships on gas with 88 % of its heating value, a base-load target that steps from 3.0 to 6.0 MW at 60 s and back at 180 s, and an anti-windup tracking time of 200 s. Run to 240 s at 5x with the default seed and export the CSV file.

(a) In the row at 150 s, read `fuel1_cmd_pct`, `truth_P1_MW` and `P1_sp_MW`. *Checkpoint:* 100 %, 1.83 MW and 2.00 MW.

(b) After 180 s, find the largest value of `truth_P1_MW` minus `P1_sp_MW`, within 0.01 MW, and the time after 180 s of the last row at which the difference is more than 0.05 MW in size, within 1 s. *Checkpoint:* 0.77 MW and 41 s.

(c) Set Anti-windup tracking time to 2 s, press Reset, run to 240 s and repeat (a) and (b). Give the time within 0.5 s. *Checkpoint:* The same row at 150 s; 0.02 MW and 4.5 s.

2.9 App challenge · 9 min

Load the challenge, case 11 of Table 2.5, which the scenario menu lists as "11 · Challenge · a customer loss, a trip and an IT rise, islanded". Use the default seed and 100x, so each 900 s run takes about 9 s. After each run read the scorecard tiles Lowest UPS charge, Supply outside ±3 K, Override · starts and Challenge target.

(a) Run to 900 s as shipped. *Checkpoint:* 52.0 %, 570 s, 408 s, not met.

(b) Set Jacket gain to 10 %/K and Jacket integral time to 60 s, press Reset and run to 900 s. *Checkpoint:* 79.2 %, 570 s, 0 s, not met.

(c) Also set Supply temperature gain to 4 %/K and Supply temperature integral time to 20 s, press Reset and run to 900 s. Say which part of the target each loop guards. *Checkpoint:* 79.2 %, 0 s, 0 s, met.

2.10 Application area · 3 min

In case 5 the customer's heat demand falls to 2.2 MW. Over 1800 s electrical following reaches a CHP efficiency of 65.7 % and thermal following 73.6 %, both on LHV.

(a) What holds thermal following back from its 45 % buffer target? *Checkpoint:* The 1.50 MW import contract, from 481 s.

(b) A site manager reads the two efficiencies as a saving of about eight points from switching authority. Give two reasons the pair is not a savings estimate. *Checkpoint:* Different stored heat at the end, and more grid electricity.

2.11 Application area · 4 min

The model takes the exhaust only down to a 120 °C stack minimum, and its jacket node is the engine outlet.

(a) Exhaust heat-recovery units are usually designed for a 250 to 350 °F outlet. Convert the lower bound to degrees Celsius and say why the recovery stops there. *Checkpoint:* 121 °C; the exhaust must not condense in its piping.

(b) In the model a restricted jacket-water flow widens the rise across the engine without heating the outlet. Why is the same fault more dangerous on a real engine? *Checkpoint:* Low coolant flow makes local hot spots, which one node cannot show.

2.12 Review the agent · 5 min

An AI assistant summarizes case 10 as follows: "At 900 s the MPC uses 1.9328 against 1.9414 MWh of fuel, a 0.45 % reduction. The predictive supervisor is therefore the more efficient way to run the plant, and it should replace the PI authorities." Find the flaw and say how you would check the claim on the app.

CHAPTER 3

Solid Oxide Fuel Cell Current That Waits for Its Fuel

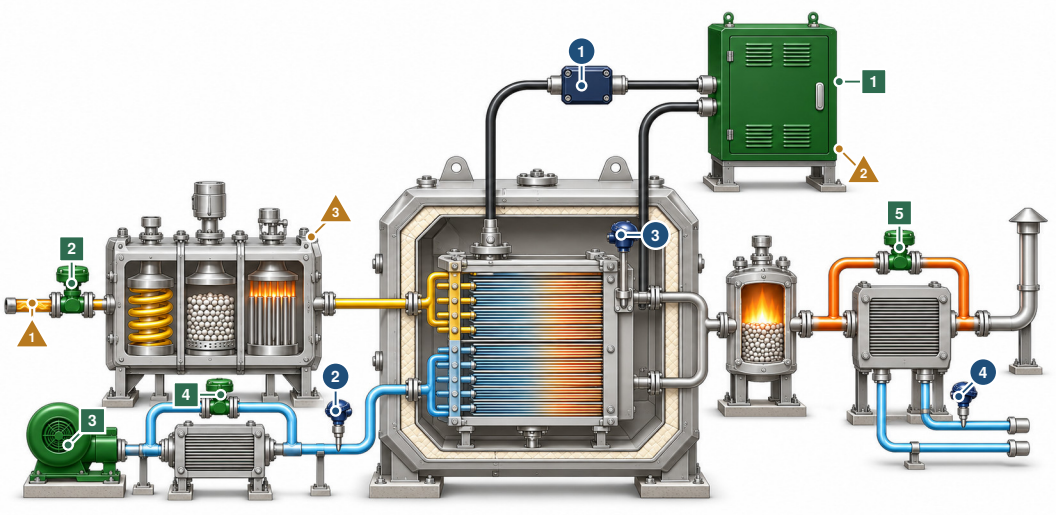

The controller moves 1 inverter that sets the stack current 2 fuel valve 3 air blower 4 valve that lets air skip the preheater 5 valve that lets exhaust skip the water heater

The controller reads 1 stack current and voltage 2 air entering the stack 3 hot end of the stack 4 hot water to the customer

What disturbs it 1 quality of the fuel gas 2 power the data center asks for 3 a fuel processor that slows down

Notice that the inverter can change the stack current at once, while the fuel must first pass the three compartments of the fuel processor before it reaches the stack. The data center and its battery are not drawn. Illustration, not to scale.

IN PLAIN TERMS A Current That Must Wait for Its Fuel

The main idea. A solid oxide fuel cell makes power for a data center from gas without burning it. Its current can change in a blink, but the hydrogen it uses comes later, after a slow trip through a fuel processor.

The challenge. If the current runs ahead of its fuel, the cell starves and its voltage sags, so the simple control asks for more current still. A limit called a governor holds the current back until the fuel is there, and a battery carries the load while it waits.

Why it matters. A starved cell can be hurt, and a data center cannot pause while its power catches up.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the fact box of the Introduction, the history box of Section 3.2, the try-it box of Section 3.5, the fact and try-it boxes of Section 3.5.3, and the key points
Study	55 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 3.5, Section 3.6, Section 3.7 and its open challenges

The control question A solid oxide fuel cell can change its current within a fraction of a second, but the hydrogen that current consumes reaches the anode only after tens of seconds in a fuel processor. How does a controller make the current wait for its fuel, and what carries the load while it waits?

Learning objectives

- Compute the hydrogen a current reacts, its fuel utilization and the current a governor admits.
- Explain the positive feedback by which load-to-current control runs the anode short of hydrogen.
- Predict the delay of the hydrogen through a fuel processor held as stages in series.
- Tell a current limited by arriving fuel from one limited by heat, by the inverter rating or by missing capacity.
- Evaluate how a wrong estimate of gas quality or holdup erodes the governor's margin.

Prerequisites Sections 1.4.2 (time scales), 1.4.4 (a store as an inventory), 1.4.5 (heat and power from one fuel), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios), with Chapter 2 for the shared bus and UPS. Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/sofc-chp/

See also Volume 9, Chapter 4 for an offset-free controller that carries a fuel bias and Volume 9, Chapter 8 for the data hall. Volume 5, Chapter 2 for fuel that acts after a transport delay and Volume 5, Chapter 3 for a slow analyzer.

3.1 Introduction

Solid oxide fuel cells are one of four fuel cell types used for stationary combined heat and power (CHP), and a catalog prepared for the U.S. Environmental Protection Agency gives them projected electric efficiencies of 45 to 60 %, with smaller units typically near 50 % [1]. In the simulated plant, 1000 equivalent cells of 1.2 m^2 at 790 °C carry an islanded data center of 2.8 MW of IT load and 0.635 MW of cooling and auxiliaries at 3.435 MW net, on 7.22 MW of methane on its lower heating value (LHV), and return 1.10 MW of heat as hot water. With every actuator held, a five percent step of the stack current moves the cell voltage from 0.7479 to 0.7397 V within 0.2 s, the anode's fuel margin completes 63 percent of its fall in 1.38 s, and the hydrogen arriving does not move, while a step of the fuel valve needs 17.3 s to deliver 63 percent of its hydrogen to the anode.

+ DID YOU KNOW? A Small Installed Fleet

The EPA catalog of 2017 counted about 83.6 MW of fuel cell CHP installed in the United States, with California first at just under 45 MW. A typical stack joins several hundred cells in electrical series [1].

The simulator is built against a common expectation, that a fuel cell follows its load as quickly as its power electronics do. A load-to-current law that divides the power the site needs by the stack voltage takes the fuel utilization to 0.993, 3.2 s after a 0.35 MW workload step, and the anode's fuel margin to 0.157 against a limit of 0.25. The thinning hydrogen lowers the voltage, the law then asks for more current, and only the inverter's floor on its input voltage stops the runaway. The current can move before the fuel is ready, so something has to make it wait and something else has to carry the load while it does.

The control problem belongs to a general class. A fast actuator stands in front of a slow supply, and a constraint on an internal state (here the hydrogen in the anode) must hold at every instant while the reference changes. The reference governor answers that class by modifying the command so that the constrained response stays admissible, a theory that grew from maximal output admissible sets [2] into a body of work with applications across many plants [3]. Model predictive control holds such constraints over a receding horizon [4], and for tubular solid oxide stacks it has been studied with multivariable models [5] and with constraints against starvation and thermal cracking in load following [6]. The same shape appears when a proton exchange membrane stack must not draw

current faster than its air compressor supplies oxygen [7], and at a fuel cell powered data center whose fuel delivery cannot follow a workload surge, where storage or power capping covers the shortfall [8].

However, the published material treats these pieces apart. Dynamic stack and plant models were written for stack design and for power system simulation [9–11], system studies showed that fuel flow set in proportion to current is effective and that fuel delivery delay limits the transient response [12], control studies chose the controlled and manipulated variables for load following [13, 14], and the data center work sizes storage at rack scale [8]. A student who wants to switch a governor off and watch the anode run short of hydrogen needs a plant in which the current, the processor's holdup, the anode inventory, the stack temperatures and a finite UPS are settings with consequences in amperes, moles and kilowatt-hours. This chapter weighs what a reference governor, a current-to-fuel feedforward and a reserve loop each add to naive load-to-current control on a simulated islanded solid oxide fuel cell plant at a data center, where the estimator's commissioning assumptions erode the governor's margin, and where no control setting replaces missing capacity.

The methods are established ones: the Nernst and Butler-Volmer relations, a holdup modeled as stages in series, the velocity-form proportional-integral-derivative (PID) controller with back-calculation and a low selector of constraint ceilings. The results are matched comparisons on one plant at one seed, and two of them correct the app's first account of its own cases. The current held back after a burst comes from the utilization trim reading the fuel still in the processor as a surplus, not from the UPS recharge, and 17.7 of the 26.1 kWh the UPS gives in the challenge are demand above the inverter's rating, which no governor setting moves. The plant has 27 physical states and 50 ledger states whose energy and atom balances close to 5×10^{-7} of their throughput. It evaluates no logarithm outside its domain, and its controllers are confined to a declared packet of measurements. A predictive controller offered as an option is compared with the governor by ablation, and ten graded cases carry the lessons into the app.

Section 3.2 traces the domain from ceramic electrolytes to fuel-limited load following. Section 3.3 gives the plant equation by equation, from the fuel processor's holdup through the cell voltage and the two stack halves to the bus, the UPS and the data hall, and says where its simplifications stop holding. Section 3.4 gives the controllers: the load-to-current law, the reference governor and its ceilings, the estimator of the arriving hydrogen, the fuel, air and heat loops, the reserve loop, the workload cap and the optional predictive controller of Section 3.4.6. Section 3.5 reports the measurements: three matched comparisons (Section 3.5.1), the recharge that was not the cause (Section 3.5.2), the combined challenge (Section 3.5.3), the predictive controller against the governor (Section 3.5.4) and the checks of the implementation (Section 3.5.5). Section 3.6 describes the hardware of real plants, Section 3.7 states the limitations, Section 3.8 draws the conclusions and Section 3.9 holds the exercises.

3.2 From Ceramic Electrolytes to Fuel-Limited Load Following

The fuel cell is older than the engines it now competes with, as the history box in this section tells. The solid oxide cell's electrolyte is an oxide-ion conductor, usually zirconia, with a nickel and zirconia cermet anode and a lanthanum manganite cathode, and it works above 600 °C [15]. The technology's progress was reviewed at the turn of the century [16].

High temperature is what makes the cell a CHP prime mover and also what slows it. The waste heat is hot enough for cogeneration or a bottoming cycle [17], and the catalog warns that long heat-up and cool-down periods restrict the cells' use in cyclic service [1]. The fuel side sets a second limit. Because the reactant activities fall as utilization rises and the cell voltage cannot exceed the lowest local Nernst potential, which in co-flow sits at the outlet, a utilization of 90 % at conditions typical of these cells costs more than 200 mV of Nernst voltage [17]. Running an anode short of hydrogen is also a risk to the cell, since a nickel anode that oxidizes and is reduced again expands, by the order of

1 % on its first complete oxidation in several studies, and the strain accumulates until it can disrupt the anode or the cell [18].

> ⧖ HISTORY **From a London Demonstration to Ceramic Stacks (1830s to 2007)**
>
> William Grove demonstrated a hydrogen fuel cell in London in the 1830s, and the idea waited more than a century for a practical use. Fuel cells returned to the laboratory in the 1950s, when the United States space program needed power systems with low or no air emissions [1]. The solid oxide cell has the longest continuous development of any type, from the late 1950s, and its solid electrolyte can be cast as tubes, flat plates or monoliths [17]. Early electrolytes conducted well enough only near 1000 °C. Planar cells for 700 to 850 °C instead of 900 to 1000 °C followed, and the lower temperature admits less expensive materials of construction [17]. Test units ran for more than ten years, yet in 2007 these cells were not yet shipping commercially [1].

Dynamic models arrived in two lines. Stack designers simulated planar stacks in three dimensions and in time [9], and power engineers built integrated plant models for system simulation [10] and for the load-following performance of fuel cells beside microturbines [11]. A dynamic model of a 25 kW tubular system put the voltage's response to a change of anode hydrogen on a scale of seconds and the temperature's on a scale of hours, showed that changing the fuel flow in proportion to the current is highly effective, and showed that a fuel flow delay imposes major limitations on the transient response [12]. Model-based control studies of internal-reforming cells followed [13], and a study of four load-following schemes took the DC power, the operating temperature, the fuel utilization and the steam-to-carbon ratio as controlled variables, the current density and the air, fuel and steam flows as manipulated ones, and the cathode outlet air temperature as the measured stand-in for a cell temperature that is hard to measure directly [14].

The answer that limits the current itself came from the low-temperature cell. For a proton exchange membrane stack a linear filter on the current demand had been proposed against oxygen starvation, and a load governor published in 2005 regulated the current drawn so that pointwise-in-time constraints on the oxygen excess ratio and the cathode's oxygen were strictly enforced, and compared its result with that filter's [7]. Data center designers then met the fuel-delivery version of the problem. In a rack-level study the fuel cell matched a power ramp 1.5 minutes after it began, drawing first on the fuel already in the stack, and an energy storage device or server power capping covered the shortfall in between, the storage recharging as a constant load on the fuel cell [8]. Published grid benchmarks pose load following with storage as dynamic optimization with perfect foresight [19], and the plant of this chapter turns that coupling into causal feedback.

3.3 An Islanded Fuel Cell Plant and Its Data Hall

The plant in Figure 3.1 is a teaching design with synthetic values throughout, drawn as the application's Plant view draws it. Fuel gas, methane with an inert share, passes a valve and a steam ratio station into a fuel processor whose external reformer and shift turn each methane molecule and two of steam into four of hydrogen and one of carbon dioxide. The reformate feeds the anodes of a stack of N_s series cells, preheated air from a blower feeds the cathodes, and both off-gases burn in a burner whose exhaust heats the reformer, the air, the steam and finally a hot-water loop. An inverter delivers the stack's power to an islanded bus that also carries the data hall, its chillers and a UPS. There are 27 physical states and 50 ledger states that integrate every boundary flux, and the variables are classified in Table 3.1. Internally the model works in seconds, watts, joules, moles, kelvin and pascals.

3.3.1 Fuel Valve, Steam Ratio and Processor Holdup

Every species enthalpy is a formation enthalpy plus one constant mean heat capacity, $h_i(T) = \Delta h_{f,i} + c_{p,i}(T - T^\circ)$ with $T^\circ = 298.15\,\mathrm{K}$, so the processor, the stack, the burner and the exchangers close on one basis and no heating value or reaction heat is ever added on top. The valve follows its command through a 1 s lag up to 12 mol/s, the ratio station meters steam at $S/C = 2.5$ on the metered

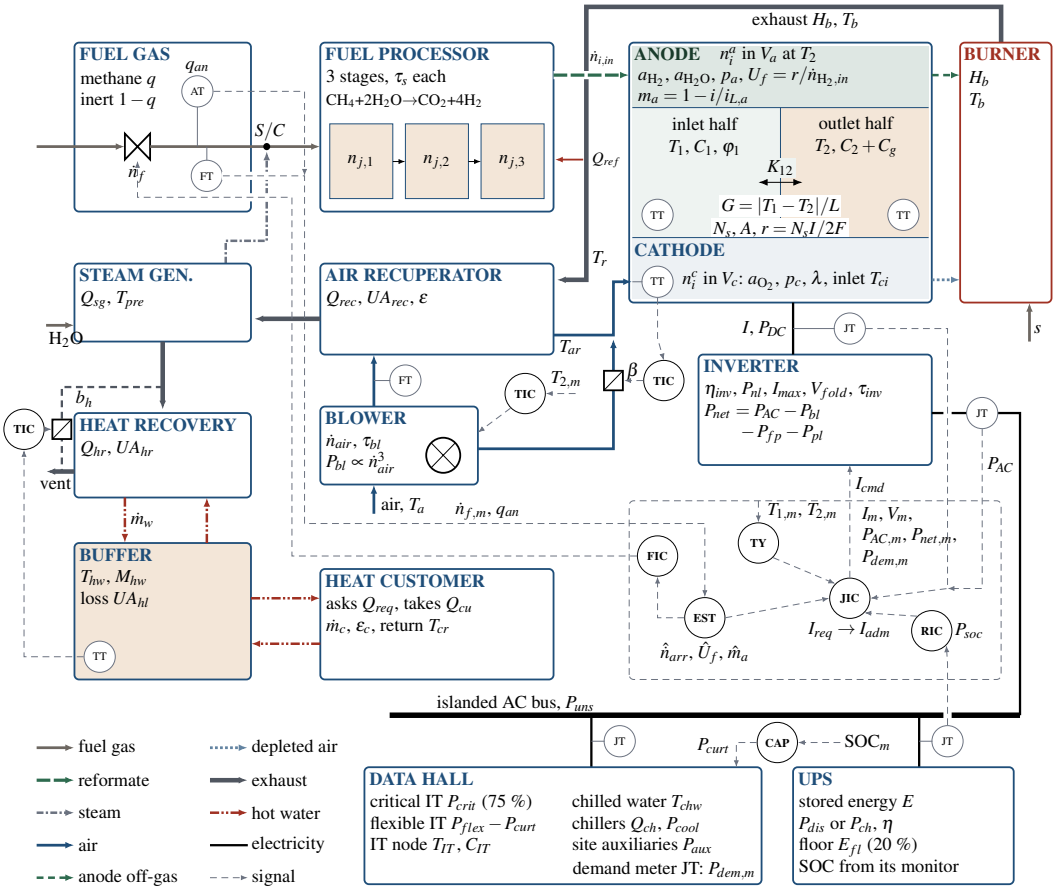

Figure 3.1 The simulated plant as the Plant view draws it. Fuel gas $\dot{n}_f$ of methane fraction q passes the valve, joins steam at the ratio S/C and travels through three processor stages $n_{j,k}$ to the anode, whose inventory n_i^a sets a_{H_2} and a_{H_2O}. The blower's air $\dot{n}_{air}$ leaves the recuperator, mixes with the bypassed share β at T_{ci} and feeds the cathode inventory n_i^c. The two stack halves at T_1 and T_2 exchange $K_{12}(T_2 - T_1)$, the current I makes P_{DC}, and the off-gases burn to H_b with supplementary fuel s. The exhaust gives Q_{ref}, Q_{rec}, Q_{sg} and Q_{hr} in that order before the vent, and the bypass b_h sends part of it around the heat-recovery exchanger. Dashed lines are measurements and commands: JIC (load-to-current and the reference governor) sets the inverter from the meters, the estimator EST, the thermal limiter TY and the recharge request P_{soc} of RIC, FIC sets the valve, TIC loops set the blower, β and b_h, and CAP curtails flexible IT. A process model, not a piping diagram, and not to scale.

gas as if it were pure methane, and the processor holds the feed in three equal stages, as shown in

Table 3.1 Variable Roles and Values at the Design Point (25 °C, 2.8 MW of IT, 1.1 MW of Heat Demand)

Variable	Role	Design	Range
P_{net}	CV, net power	3.435 MW	SP
U_f, m_a	CV, utilization, margin	0.750, 0.733	< 0.85, > 0.25
T_2, T_{ci}	CV, outlet node, air inlet	790, 620 °C	SP
T_{hw}, SOC	CV, hot water, UPS charge	80 °C, 70 %	SP
I	MV, stack current	4996 A	0–7200
$\dot{n}_f$	MV, fuel gas	8.629 mol/s	2–12
$\dot{n}_{air}$	MV, air	457.6 mol/s	150–650
β, b_h	MV, two bypasses	9.6, 22.9 %	0–60, 0–100
P_{IT}	DV, IT load, 75 % critical	2.80 MW	1.5–3.6
Q_{req}, q	DV, heat demand, methane	1.10 MW, 1.00	0–2, 0.7–1
$\Delta\tau_{fp}$, f_{air}	DV, slowdown, blower	0 s, 100 %	0–90, 30–100
$n_{j,k}$	9 states, processor stages	15 s of feed	—
n_i^a, n_i^c	4 + 2 states, electrodes	73.6, 254 mol	—
T_1, T_2	states, stack halves	724.8, 790 °C	—
T_{hw}, T_{IT}, T_{chw}	states, water, hall	80, 23.3, 14	—
Q_{ch}	state, chiller duty	2.85 MW	—
E	state, UPS energy	0.175 MWh	20–95 %
5 states	actuator lags	0.05–5 s	—

Table 3.2 Symbols of the Equations Not Defined in Tables 3.1 and 3.3

Symbol	Meaning	Symbol	Meaning
r, i	H_2 reacted, I/A	E°, E_{rev}	standard, Nernst
a_i, x_i	activity, mole fraction	p_a, N_a	anode pressure, moles
$\dot{n}_{a,out}$	anode outflow	ρ, ρ_{eff}	resistance (ASR)
$i_{L,a}$, $i_{L,c}$	limiting currents	φ_1, φ_2	current shares
C_g	gas heat capacity	G_i, B_i	streams leaving node 1
G	gradient proxy	$Q_{l,j}$	enclosure loss
H_b, s	burner, extra fuel	T_r, T_{x2}	gas temperatures
Q_{ref}, Q_{sg}	reformer, steam	Q_{rec}, Q_{hr}	recuperator, recovery
T_{ar}	recuperated air	Q_{cu}, H_{vent}	heat taken, vent
P_{DC}, P_{AC}	stack, inverter output	P_{bl}, P_{fp}, P_{pl}	blower, processor, plant
P_{gen}, P_{par}	generation, parasitic	P_{dis}, P_{ch}	UPS out, in
P_{uns}, P_{curt}	unserved, curtailed	P_{dem}, P_{soc}	demand, recharge
I_{req}, I_{adm}	request, admitted	I_{ss}, I_{tr}	fuel ceilings
I_V, I_{th}	guard, thermal	I_{air}, I_{inv}	air, inverter
$\hat{n}_{feed}$, $\hat{n}_{arr}$	H_2 fed, arriving	$\hat{U}_f$, $\hat{m}_a$	estimates
I_{ff}, $\dot{n}_{ff}$	feedforward	$\hat{U}_{ff}$, x_{th}	trim PV, excess

Equation (3.1).

$$
\begin{aligned}
\frac{\mathrm{d}n_{j,k}}{\mathrm{d}t} &= \dot{n}_{j,k-1} - \frac{n_{j,k}}{\tau_s}, \qquad k = 1,2,3 \\
\tau_s &= (\tau_{fp} + \Delta\tau_{fp})/3 \\
\dot{n}_{j,0} &= \left(q\dot{n}_f,\ (S/C)\,\dot{n}_f,\ (1-q)\dot{n}_f\right) \\
\dot{n}_{\mathrm{H_2},in} &= 4n_{\mathrm{CH_4},3}/\tau_s \\
\dot{n}_{\mathrm{H_2O},in} &= (n_{\mathrm{H_2O},3} - 2n_{\mathrm{CH_4},3})/\tau_s \\
Q_{sg} &= H_{feed}(T_{pre}) - H_{raw}(T_a) \\
Q_{ref} &= H_{a,in}(T_{fp}) - H_3(T_{pre})
\end{aligned}
\tag{3.1}
$$

where j runs over methane, steam and inert, the stages leave at their contents over τ_s (at least 0.5 s), the outlet also carries $n_{\mathrm{CH_4},3}/\tau_s$ of carbon dioxide and $n_{inert,3}/\tau_s$ of inert, Q_{sg} raises and preheats the liquid water and the feed to $T_{pre} = 150\,^\circ\mathrm{C}$, and Q_{ref} takes the products to $T_{fp} = 750\,^\circ\mathrm{C}$. Reforming and shift are complete at the outlet, so there is no carbon monoxide and its electro-oxidation, internal reforming and anode recycle are absent. The steam ratio per carbon is $2.5/q$, 2.94 at 85 % methane, never below the 2 that complete reforming needs. A reformer with its own furnace, equilibrium

and steam supply is the plant of Volume 11, Chapter 12, where the steam must likewise never fall behind the carbon fed. At design $Q_{ref} = 2.421\,\mathrm{MW}$ and $Q_{sg} = 1.114\,\mathrm{MW}$. Lengthening the holdup also delays the feed already in transit, so a slower processor briefly delivers less.

3.3.2 Electrode Inventories and Faraday's Law

The anode holds hydrogen, steam, carbon dioxide and inert in one mixed volume V_a, the cathode oxygen and nitrogen in V_c, both at the outlet node's temperature, and each drains through a linear conductance into the burner's back pressure p_b, as shown in Equation (3.2).

$$\begin{aligned} r &= \frac{N_s I}{2F}, \qquad \frac{\mathrm{d}n_i^a}{\mathrm{d}t} = \dot{n}_{i,in} - x_i^a \dot{n}_{a,out} + \nu_i r \\ \frac{\mathrm{d}n_{\mathrm{O_2}}^c}{\mathrm{d}t} &= x_{\mathrm{O_2}} \dot{n}_{air} - x_{\mathrm{O_2}}^c \dot{n}_{c,out} - \tfrac{1}{2} r \\ \dot{n}_{a,out} &= k_a \big(p_a - p_b\big)^+, \qquad p_a = N_a R T_2 / V_a \\ a_i &= n_i R T_2 / (V p^\circ), \qquad U_f = r / \dot{n}_{\mathrm{H_2},in} \end{aligned} \tag{3.2}$$

where $\nu = (-1, +1, 0, 0)$ for hydrogen, steam, carbon dioxide and inert, N_a is the anode's total moles, $x_i^a = n_i^a / N_a$, $p^\circ = 1\,\mathrm{atm}$, $p_b = 104.3\,\mathrm{kPa}$ and $x_{\mathrm{O_2}} = 0.21$. The nitrogen balance of the cathode has the same form without the reaction. The utilization is undefined below 1×10^{-3} mol/s of hydrogen. The anode holds 73.64 mol at 108.5 kPa, about 1.5 s of its throughput, which is why its margin answers a current step on a clock of about a second. Because one equivalent string carries the whole array, fuel and electrical quantities scale together.

Pause. Try Exercise 3.2 before reading on.

3.3.3 Cell Voltage and the Domain of Each Logarithm

The voltage is the Nernst potential at the outlet composition less activation, ohmic and concentration losses, with the standard potential computed from the same heat capacities as the energy balance, as shown in Equation (3.3).

$$\begin{aligned} E^\circ(T) &= -\big[\Delta h_r(T) - T \Delta s_r(T)\big] / (2F) \\ E_{rev} &= E^\circ(T_2) + \frac{R T_2}{2F} \ln \frac{a_{\mathrm{H_2}} \sqrt{a_{\mathrm{O_2}}}}{a_{\mathrm{H_2O}}} \\ V &= E_{rev} - \sum_{e=a,c} \frac{R \bar{T}}{F} \sinh^{-1} \frac{i}{2 i_{0,e}(\bar{T})} - i \rho_{eff} \\ &\quad + \frac{R T_2}{2F} \ln\left(1 - \frac{i}{i_{L,a}}\right) + \frac{R T_2}{4F} \ln\left(1 - \frac{i}{i_{L,c}}\right) \\ \rho_{eff} &= \frac{2}{1/\rho(T_1) + 1/\rho(T_2)} \\ \rho(T) &= \rho_0 \exp\big[(E_\Omega / R)(1/T - 1/T_{800})\big] \\ i_{L,a} &= k_{L,a} a_{\mathrm{H_2}}, \qquad i_{L,c} = k_{L,c} a_{\mathrm{O_2}} \\ m_a &= 1 - i / i_{L,a} \end{aligned} \tag{3.3}$$

where $i = I/A$ in $\mathrm{A/cm^2}$, Δh_r and Δs_r carry the constant Δc_p of the reaction, $\bar{T} = (T_1 + T_2)/2$, the exchange current densities are Arrhenius in $\bar{T}$ with a symmetric transfer coefficient of one half, ρ is the area-specific resistance (ASR) with $T_{800} = 1073.15\,\mathrm{K}$, the two halves conduct in parallel and node 1 carries the share $\varphi_1 = \rho(T_2)/[\rho(T_1) + \rho(T_2)]$ of the current, and m_a is the anode fuel margin, the share of the anode's limiting current still unused. At design $E_{rev} = 0.8916\,\mathrm{V}$ less 0.0645, 0.0608 and 0.0184 V gives 0.7479 V at $0.416\,\mathrm{A/cm^2}$, and $\varphi_1 = 0.337$.

Every logarithm has a physical domain, positive activities and $i < i_L$. The functions return an undefined value outside it and never clip an argument. The inverter keeps the plant inside the domain, because it cannot hold its DC input below $V_{fold} = 0.60\,\mathrm{V}$ per cell. Its current reference follows the command through a 50 ms lag, and when the voltage at that reference would fall below the floor, or the bus could not take the power, it draws instead the largest current that satisfies both, found by 40 bisections on the monotone voltage curve with an undefined voltage counted as infeasible. The current that flows is therefore always one at which the cell equation is defined, and refused generation is never created, because less current flows.

3.3.4 Two Stack Halves and Their Heat

Both gas streams enter node 1 and leave node 2, each node's gas leaves at its temperature, and the conduction between the halves is equal and opposite, as shown in Equation (3.4).

$$\begin{aligned} C_1 \frac{\mathrm{d}T_1}{\mathrm{d}t} &= \textstyle\sum_i \dot{n}_{i,in} c_{p,i} (T_{fp} - T_1) \\ &\quad + \textstyle\sum_i \dot{n}_i^{air} c_{p,i} (T_{ci} - T_1) + K_{12}(T_2 - T_1) \\ &\quad - \varphi_1 \left[r \Delta h_r(T_1) + P_{DC} \right] - Q_{l,1} \\ (C_2 + C_g) \frac{\mathrm{d}T_2}{\mathrm{d}t} &= \textstyle\sum_i (G_i + B_i) c_{p,i} (T_1 - T_2) \\ &\quad + R T_2 \frac{\mathrm{d}N}{\mathrm{d}t} - \varphi_2 \left[r \Delta h_r(T_2) + P_{DC} \right] \\ &\quad - K_{12}(T_2 - T_1) - Q_{l,2} \end{aligned} \tag{3.4}$$

where G_i and B_i are the anode and cathode streams leaving node 1 after its share of the reaction, $C_g = \sum n_i (c_{p,i} - R)$ is the heat capacity of the electrode gas inventories, $RT_2\,\mathrm{d}N/\mathrm{d}t$ their flow work, $Q_{l,j} = \frac{1}{2} U A_l (T_j - T_a)$ the enclosure loss and $\varphi_2 = 1 - \varphi_1$. The gradient proxy is $G = |T_1 - T_2|/L$ with $L = 0.10\,\mathrm{m}$, 651.8 K/m at design, a coarse stress indicator and not a prediction of a hot spot or of lifetime. The stack makes 2.67 MW of heat for its 3.736 MW of DC power, 0.715 watts per watt, and the cathode air carries most of it away at $\lambda = 7.42$.

3.3.5 Burner, Recuperator and the Heat Customer

The off-gas burns completely, and supplementary fuel gas s is fired only when the reformer's heating gas would otherwise leave below $T_{fp} + \Delta T_{ref} = 790\,^{\circ}\mathrm{C}$, as shown in Equation (3.5).

$$\begin{aligned} H_b &= H_{a,out}(T_2) + H_{c,out}(T_2) + s\, h_{fuel}(T_a) \\ s &= \frac{\left[Q_{ref} + H_b'(T_{fp} + \Delta T_{ref}) - H_{b,0} \right]^+}{h_{fuel} - \Delta h_p} \\ Q_{rec} &= \varepsilon C_{min} (T_r - T_a) \\ T_{ci} &= (1 - \beta) T_{ar} + \beta T_a \\ Q_{hr} &= \varepsilon' C'_{min} (T_{x2} - T_{hw}) \\ H_{vent} &= H_b - Q_{ref} - Q_{rec} - Q_{sg} - Q_{hr} \end{aligned} \tag{3.5}$$

where $H_{b,0}$ is the off-gas enthalpy, H_b' the products' enthalpy flow, Δh_p the enthalpy of the products of one mole of fuel gas at 790 °C, and every temperature is read from the enthalpy flow on the products' own composition, $T = T^{\circ} + (H - \sum \dot{n}_i \Delta h_{f,i}) / \sum \dot{n}_i c_{p,i}$. The recuperator heats the unbypassed air $(1 - \beta)$ against the gas leaving the reformer at T_r with the counterflow effectiveness ε of $UA_{rec} = 55\,\mathrm{kW/K}$ and delivers it at T_{ar}, the steam generator takes Q_{sg}, and the heat-recovery exchanger heats the water loop with the gas not bypassed, $(1 - b_h)$, which arrives at T_{x2}. The hot water and the

customer close as

$$\begin{aligned} M_{hw}c_w\frac{\mathrm{d}T_{hw}}{\mathrm{d}t} &= Q_{hr} - Q_{cu} - UA_{hl}(T_{hw} - T_a) \\ Q_{cu} &= \min\!\left(Q_{req},\ \varepsilon_c \dot{m}_c c_w (T_{hw} - T_{cr})^{+}\right) \end{aligned} \tag{3.6}$$

At design the burner reaches 942.6 °C, the supplementary fuel is 0.374 mol/s, and the train gives 2.421 MW to the reformer, 8.631 MW to the air, 1.114 MW to the steam and 1.122 MW to the water, venting at 105.0 °C. Burned off-gas is fuel the stack did not consume, so it is never counted as stack fuel, and the efficiencies on one LHV boundary are 47.55 % electric and 62.78 % CHP.

3.3.6 Inverter, Bus and the Stored-Energy Rule

The bus balance and the UPS come from the framework's shared site module `core/cogen.js`, written for the gas engine plant of Chapter 2, moved to the framework with the turbine trigeneration plant of Volume 4, Chapter 9 and imported here unchanged, as shown in Equation (3.7).

$$\begin{aligned} P_{net} &= \eta_{inv} N_s V I - P_{nl} - P_{bl} - P_{fp} - P_{pl} \\ P_{gen} + P_{dis} + P_{uns} &= P_{dem} + P_{ch} \\ P_{dem} &= P_{crit} + P_{flex} - P_{curt} + P_{cool} + P_{aux} + P_{par} \\ \frac{\mathrm{d}E}{\mathrm{d}t} &= \eta_{ch} P_{ch} - P_{dis}/\eta_{dis} \\ P_{dis} &\le \min\!\left(P_{dis,max},\ (E - E_{fl})\eta_{dis}/\Delta t\right) \end{aligned} \tag{3.7}$$

where $P_{bl} = 110\,\mathrm{kW}\,(\dot{n}_{air}/450\,\mathrm{mol/s})^3$, $P_{fp} = 15\,\mathrm{kW} + 1\,\mathrm{kW}$ per mol/s of fuel gas, $P_{pl} = 40\,\mathrm{kW}$, $P_{gen} = P_{net}^{+}$, and a negative net output appears as P_{par} on the demand side. All terms are nonnegative, and each step checks the balance one term at a time. Islanded, the UPS takes the imbalance inside the step's limits, which the energy held when the step begins sets and all four Runge-Kutta stages keep. Those stages carry nonnegative weights, so the integrated energy can at most land on the floor, nothing is clipped afterward and no step needs splitting. Charge is limited the same way toward the 95 % ceiling and by the 0.5 MW rating, and the inverter's current limit includes the charge limit, so a surplus the bus cannot take is never generated. A deficit beyond the UPS is unserved, and this application replaces the module's priority order with a proportional one, so an unplanned shortfall falls on every feeder in the share P_{uns}/P_{dem} and the workload cap is the only planned shedding of the flexible 25 % of the IT load. From its 70 % start the 1 MW, 0.25 MWh UPS holds $0.25 \times 0.50 \times 0.95 = 0.118\,75\,\mathrm{MWh}$ above its 20 % floor, 7.1 min at its rating.

3.3.7 Data Hall, Instruments, Events and Numerics

The data hall is an equipment node and a chilled-water node, as shown in Equation (3.8).

$$\begin{aligned} C_{IT}\frac{\mathrm{d}T_{IT}}{\mathrm{d}t} &= P_{IT,del} - UA_{IT}(T_{IT} - T_{chw}) \\ C_{chw}\frac{\mathrm{d}T_{chw}}{\mathrm{d}t} &= UA_{IT}(T_{IT} - T_{chw}) + Q_x - s_c Q_{ch} \\ \tau_{ch}\frac{\mathrm{d}Q_{ch}}{\mathrm{d}t} &= \mathrm{clip}_{[0,\,Q_{ch,max}]}(Q_{cmd}) - Q_{ch} \\ Q_{cmd} &= UA_{IT}(T_{IT} - T_{chw}) + Q_x \\ &\quad + K_{ch}(T_{chw} - T_{chw,SP}) \end{aligned} \tag{3.8}$$

where $P_{IT,del}$ is the IT power actually served, Q_{cmd} the chiller plant's own duty command, Q_x the half of the site auxiliary load that ends in the chilled water, s_c the served share of the cooling feeder, and the chillers draw $P_{cool} = Q_{ch}/\mathrm{COP} + P_{pump}$ with a COP of 6 and 60 kW of pumps and fans. T_{IT} is a lumped thermal state of the IT equipment, not a measured chip or inlet temperature.

Table 3.3 Model Parameters. Class Is the Source of Each Value (L Checked Against the Literature, F Fitted, T Synthetic Teaching Value) With Its Visibility in the Parameters Panel (V Shown, H Hidden)

Symbol	Meaning	Value Unit	Class	Symbol	Meaning	Value Unit	Class
Stack and electrochemistry				*Air, burner and heat recovery*			
N_s, A	cells, area each	1000, 1.2 —, m^2	T, H	$\dot{n}_{air,max}$, $\dot{n}_{air,min}$	blower range	650, 150 mol/s	T, V/H
$E^\circ(800\,^\circ\mathrm{C})$	from Δh, Δs, Δc_p	0.982 V	L	τ_{bl}, P_{bl}	lag, power at 450	2, 110 s, kW	T, H
ρ_0, E_Ω	ASR, its activation	0.10, 91.4 $\Omega\,cm^2$, kJ/mol	T, F	UA_{rec}, β_{max}	recuperator, bypass	55, 60 kW/K, %	T, V/H
$i_{0,a}$, $E_{a,a}$	anode exchange current	1.5, 80 A/cm^2, kJ/mol	T, H	τ_β, τ_{b_h}	damper lags	3, 5 s	T, H
$i_{0,c}$, $E_{a,c}$	cathode exchange current	0.6, 100 A/cm^2, kJ/mol	T, H	ΔT_{ref}	reformer approach	40 K	T, H
$k_{L,a}$, $k_{L,c}$	limiting currents	8, 12 A/cm^2	T, H	UA_{hr}, $\dot{m}_w$	recovery, water	45, 15 kW/K, kg/s	T, V/H
V_a, V_c	gas volumes	6, 20 m^3	T, V/H	M_{hw}, UA_{hl}	buffer, loss	20, 0.4 t, kW/K	T, V/H
k_a, k_c	outlet conductances	11.4, 56 mol/(s MPa)	T, H	$\dot{m}_c$, T_{cr}, ε_c	customer loop	20, 55, 0.8 kg/s, °C, —	T, H
C_1, C_2	stack halves	3, 3 MJ/K	T, V	*Power, site and UPS*			
K_{12}, UA_l, L	conduction, loss, distance	10, 0.2, 0.10 kW/K, kW/K, m	T, H	η_{inv}, P_{nl}	efficiency, no-load	0.97, 10 —, kW	T, V/H
Fuel processor				I_{max}, $P_{AC,max}$, V_{fold}	limits, foldback	7200, 4.3, 0.60 A, MW, V	T, V/H
τ_{fp}, stages	processor holdup	15, 3 s, —	T, V	τ_{inv}	current lag	50 ms	T, H
S/C, q_{nom}	steam ratio, assumed q	2.5, 1 —	T, V/H	P_{crit}/P_{IT}	critical share	75 %	T, V
T_{pre}, T_{fp}	feed, reformer outlet	150, 750 °C	T, H	C_{IT}, UA_{IT}	IT node	20, 0.3 MJ/K, MW/K	T, H
$\dot{n}_{f,max}$, τ_{fv}	valve open, lag	12, 1 mol/s, s	T, V/H	C_{chw}, τ_{ch}, K_{ch}	chilled water, loading	150, 60, 1 MJ/K, s, MW/K	T, H
Protection (teaching thresholds)				$Q_{ch,max}$, COP, P_{pump}	chillers, pumps	4, 6, 60 MW, —, kW	T, H
$U_{f,max}$	utilization ceiling	0.85 —	T, V	P_{aux}	site auxiliaries	0.10 MW	T, H
m_{lim}, m_{trip}	margin limit, trip	0.25, 0.10 —	T, H/V	$P_{dis,max}$, $P_{ch,max}$	UPS ratings	1.0, 0.5 MW	T, V/H
T_{lo}, T_{hi}, T_{trip}	node limits, trip	650, 830, 860 °C	T, H/V/H	E_{cap}, η	energy, efficiency	0.25, 0.95 MWh, —	T, V/H
G_{max}, R_{max}	gradient, node rate	1000, 5 K/m, K/min	T, V/H	floor, ceiling, SOC_0	UPS charge	20, 95, 70 %	T, V/H/V

Controllers read only a measurement packet. Each analog reading is the truth through a first-order transmitter lag, discretized at the scan with the value at the scan instant, with noise added after the lag: current and voltage (0.02 s, 1 A and 0.2 V on about 750 V), the power meters (0.1 s, 2 kW), the flow meters (0.5 s), the stack thermocouples (5 s, 0.3 K), the cathode inlet and hot water (3 s) and the UPS output (0.05 s), with the battery monitor reporting the charge once a second. A sampled analyzer reports the methane fraction from a sample drawn every 60 s through a 30 s line and analyzed for 60 s, so a reading is 90 to 150 s old. No instrument in the control system reads the anode's hydrogen. The site demand meter reads the load request, so shed load still registers. Every scan draws the same count of random numbers whatever the settings.

Events happen once, at a step boundary, with a mark on the trends. The stack trips and latches at a true anode margin of 0.10, read from an anode-outlet analyzer in the protection chain (a declared teaching simplification), or when a measured node reaches 860 °C, and a trip zeroes the current and puts the fuel to its 2.0 mol/s hot standby in any mode. Each case starts in equilibrium for its own settings, solved by bisection on the current for the site's demand, Newton iterations on the two node balances with the air, and bisection on both bypasses, and its derivatives there are below 2.4×10^{-11}. A fixed step of 0.02 s of the classical fourth-order Runge-Kutta scheme carries the states, the fastest retained modes being the cathode's 35 ms and the anode's 57 ms pressure relaxation and the inverter's 50 ms lag, and the historian logs one row a second. The integration step, the scan and the log are the three clocks of the engine of Section 1.6.1. Table 3.3 lists the parameters with the class of each value.

3.3.8 Where the Simplifications Stop Holding

The plant is hot and operating, so cold start, thermal cycling and degradation chemistry are outside it. One mixed volume per electrode puts the whole anode at its outlet composition, the operative one for the voltage in co-flow [17], but it cannot show a local depletion upstream. The design point, 0.31 W/cm^2 at 0.748 V, sits inside the handbook's 300 to 400 mW/cm^2 for planar stacks on reformate [17] and its 47.55 % inside the catalog's range [1], but no polarization constant was fitted to a cell. The constant mean Δc_p puts E° 0.006 to 0.014 V below the handbook's two-digit values above 650 °C, 0.982 V against 0.99 V at 800 °C. The 3 MJ/K per half lets the stack's temperatures move within a fifteen-minute case, while 1200 m^2 of planar cells with metal interconnects would hold roughly 15 to 50 MJ/K, so the thermal lessons come early (Section 3.5.5). The proportional shortfall is an averaged energy account, since a real islanded bus without shedding would drop whole

Table 3.4 Loops, Scans, Default Tuning and Action (+1 Raises the Output When the Measurement Is Below Its Setpoint), All PI

Loop	PV → MV	Scan	K_c	τ_I	Act.	Output range
trim	$\hat{U}_{ff} \to \dot{n}_f$	1 s	5 mol/s	15 s	−1	$[2, 12] - \dot{n}_{ff}$
air	$T_{2,m} \to \dot{n}_{air}$	1 s	6 mol/(s K)	120 s	−1	$[150,\ \dot{n}_{air,m} + 52]$
preheat	$T_{ci,m} \to \beta$	1 s	1 %/K	30 s	−1	0–60 %
hot water	$T_{hw,m} \to b_h$	1 s	30 %/K	90 s	−1	0–100 %
thermal	$x_{th} \to I_{th}$	1 s	60 A/K	20 s	+1[a]	$[1080,\ \bar{I}]$
guard	$V_m/N_s \to I_V$	0.1 s	20 kA/V	2 s	−1	$[1080,\ \bar{I}]$
reserve	$\mathrm{SOC}_m \to P_{soc}$	5 s	3 MW	300 s	+1	0–0.3 MW

[a] On the negative excess, so an excess lowers the ceiling. $\bar{I} = \min(I_{max},\ \max(I_{req}, I_{adm}) + 300\,\mathrm{A})$. Setpoints $U_{f,SP} = 0.75$, $T_{2,SP} = 790\,^\circ\mathrm{C}$, $T_{ci,SP} = 620\,^\circ\mathrm{C}$, $T_{hw,SP} = 80\,^\circ\mathrm{C}$, $V_g = 0.68\,\mathrm{V}$, $\mathrm{SOC}_{SP} = 70\,\%$.

feeders or collapse. Real limits come from the fuel cell, inverter and UPS manufacturers' data and the applicable codes.

3.4 Admitting Current: Governor, Estimator and Loops

One framework scan of $\Delta t_c = 0.1\,\mathrm{s}$ runs the instruments, the estimator, the load-to-current law, the voltage guard, the governor, the inverter command and the workload cap. Integer counters inside it run the fuel, air, bypass, hot-water and thermal-limiter loops every 1 s and the reserve loop every 5 s, each integrating over the real interval between its own runs, and counters, filters, PID states and the estimator are created again at every reset. Table 3.4 lists the loops. Their action follows Section 1.3: +1 is reverse acting, as in the reserve loop, where more recharge raises the charge, and −1 is direct acting, as in the air loop, where more air cools the outlet node.

3.4.1 One PID Code for Every Loop

Every loop runs the framework's velocity-form PID of Section 1.6.2, Equations (1.10) to (1.13), with the back-calculation time $T_t = \tau_I$ and the loop's own elapsed interval Δt_k in place of the control period. Every loop here is PI. A loop that does not own its output calls `track()`, which back-solves the integral for the value actually applied, and no PID is updated and tracked in one scan, so every transfer is bumpless. The plant signs were measured from the design point: 100 A more lowers the voltage from 0.7479 to 0.7447 V in 0.2 s, 0.4 mol/s more fuel lowers the utilization from 0.750 to 0.719 after 30 s, and 20 mol/s more air cools the outlet node by 1.06 K in 60 s.

3.4.2 Load-to-Current and the Reference Governor

The current request is the power asked divided by the measured voltage, and the governor admits the smallest of six ceilings, as shown in Equation (3.9). A low selector over limits on one actuator is also the structure of the jet engine fuel control of Volume 4, Chapter 3.

$$
\begin{aligned}
I_{req} &= \frac{P_{dem,m} + P_{soc} + (P_{AC,m} - P_{net,m})^{+} + P_{nl}}{\eta_{inv} V_m} \\
I_{ss} &= \frac{2 F U_{gov} \hat{\dot{n}}_{feed}}{N_s}, \qquad I_{tr} = \frac{2 F U_{gov}}{N_s} \min_{[t,\,t+H]} \hat{\dot{n}}_{arr} \\
I_{air} &= \frac{4 F x_{\mathrm{O}_2} \dot{n}_{air,m}}{N_s \lambda_{min}} \\
I_{inv} &= \min\left(I_{max},\ (P_{AC,max} + P_{nl})/(\eta_{inv} V_m)\right) \\
I_{adm} &= \min(I_{req}, I_{ss}, I_{tr}, I_V, I_{air}, I_{th}, I_{inv})
\end{aligned}
\tag{3.9}
$$

where $P_{AC,m} - P_{net,m}$ is the plant's own consumption from two meters, $U_{gov} = 0.80$ is the declared margin below the 0.85 ceiling, $\lambda_{min} = 2$, and $H = 5\,\mathrm{s}$. I_{ss} is the plan's steady ceiling on what the metered feed will make, 5329 A at design, and I_{tr} the transient anode constraint on the smallest

arrival the estimator predicts. The voltage guard I_V is a PI on the measured cell voltage with a 0.68 V setpoint, the backstop for a shortfall the estimator cannot see. The thermal limiter I_{th} is a PI on the excess

$$x_{th} = \max\big(T_{max,m} - T_{lim},\ (G_m - G_{lim})L,\ (\dot{T}_m - R_{max}) \cdot 60\,\mathrm{s}\big) \tag{3.10}$$

with $T_{lim} = 820\,^{\circ}\mathrm{C}$, $G_{lim} = 900\,\mathrm{K/m}$ and the heating rate taken from thermocouples filtered over 10 s and differenced over 60 s, so their noise reads as about 0.1 K/min against the 5 K/min limit. The guard and the limiter wait at $\bar{I}$, 300 A above the larger of the requested and the admitted current, tracking that ceiling while they have nothing to correct, and report no ceiling while they wait, so the gap never acts as a ramp limit and neither can wind up. A switch to Manual starts the operator's current at I_{adm}, where the fuel, air and inverter ceilings still cap it while the guard acts only in Auto, and with the governor off only I_{th} and I_{inv} remain, both protective in either mode. The inverter command moves toward I_{adm} at most 3000 A/s.

3.4.3 An Estimate of the Arriving Hydrogen

The governor and the trim need the anode's hydrogen, which no instrument reads, so an estimator keeps a declared model driven by the packet, as shown in Equation (3.11).

$$\begin{aligned}
\frac{\mathrm{d}\hat{c}_k}{\mathrm{d}t} &= (\hat{c}_{k-1} - \hat{c}_k)/\hat{\tau}_s, \qquad \hat{c}_0/\hat{\tau}_s = \hat{q}\dot{n}_{f,m} \\
\hat{n}_{arr} &= 4\hat{c}_3/\hat{\tau}_s, \quad \hat{n}_{feed} = 4\hat{q}\dot{n}_{f,m}, \quad \hat{\tau}_s = \hat{\tau}_{fp}/3 \\
\frac{\mathrm{d}\hat{n}_{\mathrm{H_2}}}{\mathrm{d}t} &= \hat{n}_{arr} - \frac{\hat{n}_{\mathrm{H_2}}}{\hat{N}_a}\hat{n}_{out} - \frac{N_s I_m}{2F} \\
\hat{N}_a &= \frac{p_{nom} V_a}{R T_{2,m}}, \qquad \hat{U}_f = \frac{N_s I_m}{2F\hat{n}_{arr}} \\
\hat{m}_a &= 1 - \frac{I_m/A}{k_{L,a}\hat{a}_{\mathrm{H_2}}}
\end{aligned} \tag{3.11}$$

where the inputs are the metered fuel flow $\dot{n}_{f,m}$ (0.5 s lag), the methane fraction $\hat{q}$ (the operator's entry, or the analyzer's report 90 to 150 s old), the measured current I_m and the outlet thermocouple $T_{2,m}$. The stages use the commissioning holdup $\hat{\tau}_{fp} = 15\,\mathrm{s}$ and are integrated by explicit Euler at 0.05 s substeps, the anode is one volume at its nominal pressure $p_{nom} = p_b + 4\,\mathrm{kPa}$, its outflow is the arriving gas per methane, $4 + (2.5/\hat{q} - 2)^+ + 1 + (1-\hat{q})/\hat{q}$ moles, and $\hat{a}_{\mathrm{H_2}} = (\hat{n}_{\mathrm{H_2}}/\hat{N}_a)(p_{nom}/p^{\circ})$. The horizon of I_{tr} moves the valve toward its present command through its 1 s lag and the stages forward in 0.25 s steps for 5 s, assuming the measured inputs stay constant and knowing no scheduled event. The utilization estimate is undefined below 0.5 mol/s of arriving hydrogen, and the margin estimate is set to $-\infty$ rather than evaluated when $\hat{a}_{\mathrm{H_2}} \le 0$, so the estimator takes no logarithm. At the priming call it starts at its own steady state for the first packet. Its bias is $\hat{q} - q$ and $\hat{\tau}_{fp} - \tau_{fp}$, the only two quantities it can get wrong, because its anode model takes the plant's declared volume, area, limiting current and steam ratio as commissioning values. Neither the estimator nor the governor reads simulator truth.

3.4.4 Fuel, Air, Preheat and Hot Water

The feedforward sends fuel for the current at the utilization setpoint as soon as it is asked for, and the trim corrects the utilization of that same current, as shown in Equation (3.12).

$$\begin{aligned}
I_{ff} &= \min(I_{req},\ I_{air},\ I_{th},\ I_{inv}) \\
\dot{n}_{ff} &= \frac{N_s I_{ff}}{8F\hat{q}U_{f,SP}}, \qquad \hat{U}_{ff} = \frac{N_s I_{ff}}{2F\hat{n}_{arr}} \\
\dot{n}_{f,cmd} &= \mathrm{clip}_{[2,12]}\big(\dot{n}_{ff} + \mathrm{PI}_u(U_{f,SP} - \hat{U}_{ff})\big)
\end{aligned} \tag{3.12}$$

The fuel ceilings are left out of I_{ff} on purpose, so the fuel starts moving at once and the admitted current climbs as the hydrogen arrives. A trim on the utilization of the admitted current would read the held current's low utilization as a fuel surplus, cut the feed and deadlock with the governor, which the first build measured as an emptied UPS in the governed challenge. Switching the feedforward moves its old contribution into the trim's integral, so the valve does not jump. The air loop moves the blower to hold the outlet node, with its ceiling following the metered flow so that a restricted blower cannot wind it up, the preheat loop holds the cathode inlet with the bypass, and the hot-water loop opens the heat-recovery bypass when the buffer runs hot.

3.4.5 Reserve Loop and Workload Cap

The reserve loop asks the fuel cell for extra power P_{soc} to restore the UPS toward 70 %. While the governor holds the current it does not integrate but tracks what the admitted current actually leaves above the demand, so it cannot wind up and the governor keeps priority. The workload cap, when armed, sheds the whole flexible share while the reported charge is below 40 % or the UPS gives more than 0.9 MW, and releases it once the charge is above 45 %, the UPS below 0.5 MW, the current unheld and 30 s have passed. Three models are kept apart. The plant is Section 3.3, nothing is fitted to data, and the internal model of the governor, the feedforward and the trim is the estimator of Equation (3.11) with its two commissioning values, while Section 3.4.6 gives the predictive controller's own.

3.4.6 Model Predictive Control of Current and Fuel

A structure switch offers a model predictive controller (MPC) in place of the governor's two fuel ceilings and the feedforward. It is never the default, and case 10 alone loads it. Every $T_s = 1$ s, or ten scans, a quadratic program plans the current I and the fuel command $\dot{n}_f$ over $N = 60$ steps and its first step is applied, the receding-horizon scheme of industrial predictive control [4] that Section 1.6.3 describes for every app. Seen from this controller, the governor of Equation (3.9) is a one-constraint relative that holds the utilization on the hydrogen it expects by cutting the current alone, while the feedforward sends the steady fuel for the current asked now. The plan holds that limit, the margin and the cell voltage over a minute with the valve as a second handle, so fuel can be sent ahead of a current that waits for it.

The model is written from Section 3.3 with the declared defaults of Table 3.3, never with the run's parameters, so a change made to the plant is model mismatch. Its states are the valve, the estimator's three stages and the anode's hydrogen, as shown in Equation (3.13).

$$
\begin{aligned}
\frac{\mathrm{d}v}{\mathrm{d}t} &= \frac{\dot{n}_f - v}{\tau_{fv}}, \qquad \frac{\mathrm{d}c_1}{\mathrm{d}t} = \hat{q}v - \frac{c_1}{\hat{\tau}_s} \\
\frac{\mathrm{d}c_k}{\mathrm{d}t} &= \frac{c_{k-1} - c_k}{\hat{\tau}_s}, \qquad k = 2,3 \\
\frac{\mathrm{d}n}{\mathrm{d}t} &= \frac{4\hat{\theta}c_3}{\hat{\tau}_s} - \frac{n}{N_a}\dot{n}_{out} - \frac{N_s I}{2F} \\
\dot{n}_{out} &= \frac{c_3}{\hat{\tau}_s}\,\frac{2\hat{\theta}\hat{q} + (S/C)q_{nom} + 1}{\hat{q}} \\
N_a &= \frac{p_a V_a}{R T_{2,m}}, \qquad p_a = p_b + \dot{n}_{out}/k_a
\end{aligned}
\tag{3.13}
$$

Here v is the delivered flow, c_k are stage contents in moles of methane as in Equation (3.11) with the estimator's commissioning $\hat{\tau}_s = \hat{\tau}_{fp}/3$ and $\hat{q}$, n is the anode's hydrogen, $\hat{\theta}$ is the quality ratio below, $q_{nom} = 1$, and N_a is the anode's total moles at the measured outlet temperature. The cell voltage and the margin follow from Equation (3.3) with the declared cell constants and the measured temperatures and air held over the horizon. The free response is integrated by RK4 with both inputs held, and the

moves act through a linearization at each decision, sampled with an exact zero-order hold and joined to the free response by an affine term per step. Run open loop for 60 s after a step of fuel or current on the lined-out plant with the noise off, the model's RMS errors are at most 0.0021 in utilization, 0.0001 in margin and 0.0010 V in cell voltage, and its Jacobian matches one derived by hand to 2.8×10^{-7}.

The quality ratio $\hat{\theta}$, the hydrogen really arriving over the hydrogen the estimator credits, makes the plan offset free as an input disturbance [20], the disturbance state of Section 1.6.3. The offset-free predictive controller of Volume 9, Chapter 4 carries a fuel bias of the same kind. Together with an anode inventory $\hat{n}$ it is predicted every 0.1 s scan with $\hat{\theta}$ held and corrected by the measured cell voltage, as shown in Equation (3.14).

$$\begin{bmatrix} \hat{n} \\ \hat{\theta} \end{bmatrix} = \begin{bmatrix} \hat{n}^{-} \\ \hat{\theta}^{-} \end{bmatrix} + L \left[\frac{V_m}{N_s} - V\left(I_m, \hat{n}^{-}, \hat{\theta}^{-}\right) \right] \tag{3.14}$$

The minus marks the prediction, $V(\cdot)$ is the model's cell voltage, and L is the steady-state Kalman gain at the design point, with the entries 53.5536 mol/V for $\hat{n}$ and 1.881 78 V^{-1} for $\hat{\theta}$. It is computed for process variances of 1×10^{-3} mol^2 and 2×10^{-6} per scan, a random walk on $\hat{\theta}$, against 4×10^{-7} V^2 on the voltage. It matches an independent Riccati iteration to 4.4×10^{-11}, the error decays with time constants of 0.65 and 0.39 s, and lined out with the instruments' noise $\hat{\theta}$ has a standard deviation of 0.0007 within its bounds of 0.5 and $1.01/\hat{q}$. Leaner gas shows once it reaches the anode, within 27.6 s of the change in case 10 against the analyzer's 90 to 150 s, but a late processor or a drifted stack is read as a change of gas too (Section 3.7).

The plan minimizes Equation (3.15) subject to Equation (3.16).

$$\begin{aligned} J = {} & q_I \sum_{k=0}^{N-1} \left(\frac{I_k - I_{tgt}}{I_{ref}} \right)^2 + r_f \sum_{k=0}^{N-1} \left(\dot{n}_{f,k} - \dot{n}_{f,ref} \right)^2 \\ & + \sum_{j=1}^{12} \left[s_I \left(\frac{\Delta I_j}{I_{ref}} \right)^2 + s_f \, \Delta \dot{n}_{f,j}^2 \right] + \sum_{r=1}^{3} \left(\rho_r \sigma_r^2 + \ell_r \sigma_r \right) \end{aligned} \tag{3.15}$$

$$\begin{aligned} & 0 \le I_k \le I_{box}, \qquad 2 \le \dot{n}_{f,k} \le \dot{n}_{f,max} \\ & \frac{N_s I_k}{2F} \le \frac{4 U_{gov} \hat{\theta} c_{3,k}}{\hat{\tau}_s} + \sigma_1 \\ & m_{a,k+1} \ge m_{lim} - \sigma_2, \qquad V_{k+1} + b_V \ge V_g - \sigma_3 \end{aligned} \tag{3.16}$$

Here $I_{ref} = 1000$ A, $q_I = 100$, $s_I = 0.01$, $r_f = s_f = 0.05$ per $(\mathrm{mol/s})^2$, and Δ is the change of a block from the one before, the first from the commands applied now, over twelve blocks (four of one step, two of two, then 3, 4, 6 and 9 steps and two of 15), which with the three slacks σ_r gives 27 decision variables and 231 rows. The box $I_{box} = \min(I_{max}, I_{inv}, I_{air}, I_{th})$ and the valve's range are hard. The soft rows are the governor's own limit on the hydrogen arriving when each step's current applies, with $c_{3,k}$ the third stage at the step's start, the warning margin $m_{lim} = 0.25$, and the guard's setpoint V_g with the bias b_V of the measured voltage over the model's. Each row carries a single slack shared by all 60 steps, priced at $\rho_r = 10^4$, 10^3 and 10^2 and $\ell_r = 10^3$, 10^2 and 10^2 in mol/s of hydrogen, tenths of margin and hundredths of a volt. The target I_{tgt} is the request capped by the box, by the current that full fuel carries at the limit and by the current at which the steady voltage at full fuel meets V_g, and $\dot{n}_{f,ref}$ is the feedforward of Equation (3.12) corrected by $\hat{\theta}$. The weights are design values kept on every case.

Between plans the current follows the request inside the first step's allowance, the largest current that keeps the first step's rows on their bounds, because holding the planned move left the current below the noisy request for 495.2 s of case 3 instead of 10.2 s. The admitted current never passes a protective ceiling and never takes the utilization row above U_{gov}, and the voltage guard, the air ceiling, the thermal limiter and the inverter still low-select it as in Equation (3.9). A plan is accepted

only on an optimal status, the expected array shapes, every planned command inside its box and finite costs, slacks and predictions, and a test injects four faults that must each be rejected. A rejected plan returns the current and the fuel to the governor, feedforward and trim, crediting hydrogen scaled by $\hat{\theta}$. Forced rejections at twelve instants of case 10, with the noise off, admitted exactly the current the MPC would have admitted, and thirty in a row after the gas change kept the true utilization at 0.802 or below. The MPC reads only the measurement packet, the estimator's stages, the applied commands, the settings and the run's limits, and fed one run's packets over a different hidden plant it issued identical commands for 8999 scans. A plan cost 0.38 ms on average in one browser run (largest 4.5 ms), 3.8 % of one core at 100 times real time.

3.5 Experiments on One Plant

A script loads each case as the case menu does, applies any setting named in the text before time starts (so it is a starting condition), steps the shipped application at its 0.02 s step on the default seed, and exports the CSV record from which the figures are drawn, as Section 1.6.4 describes for every app. On the simulator's truth the power error is $e = P_{dem} + P_{soc} - P_{net}$, the power loop's own setpoint less its output. IAE is $\int |e|\,\mathrm{d}t$ over the stated window in kilowatt-hours, the largest error and the UPS peak are taken from 0.2 s after the event (inside the first scan a step's jump is caught at the end of one physics step, which depends on the step size), recovery is the last moment in the window with $|e| > 20\,\mathrm{kW}$, less the event time, with the rest of the window as its dwell, held time counts the seconds with the admitted current below the request, and travel is the total variation of the fuel valve state. The ten cases are listed in Table 3.5, where case 1 shows the three clocks of the Introduction before any loop is closed.

> **TRY IT IN THE APP** **1 · A Manual Current Step: What Moves, and When**
>
> Load scenario 1 on the Fuel and current view and press Run in Manual. After 30 s type 5250 into Stack current and watch for two minutes. The cell voltage falls from 0.7479 to 0.7397 V within 0.2 s, the anode fuel margin falls from 0.733 to 0.671 by 5 s, and the fuel stays at 8.63 mol/s.

3.5.1 Three Matched Comparisons

Each group of Table 3.6 starts from one initial state with the same seed, meters, limits and timed disturbance and changes only the settings named, and Figure 3.2 shows the trajectories.

Group A is nominal disturbance rejection on the identical step. Naive control asks for the current at once, the utilization reaches 0.993 and the margin 0.157, and for 20.0 s the inverter's foldback, not a controller, stops the current, while the UPS gives up to 0.465 MW. The twenty seconds of high current at low voltage heat the outlet node at up to 8.0 K/min on the filtered thermocouples (3.7 K/min governed) against the limiter's 5 K/min, so the limiter holds the current for 48.9 s and naive control recovers last, at 94.1 s. The feedforward alone moves the fuel at once but cannot make it arrive sooner, so the margin still falls to 0.158 and the foldback lasts 11.2 s. The governor alone keeps the margin at 0.614 and the utilization at 0.800, at the price of 30.5 s of held current and 1.16 kWh from the UPS, and with both the current waits only 13.5 s, the UPS gives 0.65 kWh, the IAE falls from 2.638 to 0.797 kWh and the margin stays at 0.623 or above, as shown in Figure 3.2a and b. Every run ends within 0.7 points of its 70 % start, so no hidden discharge of the store is rewarded, and across seeds 1 to 5 the naive margin is 0.157 and the governed 0.622 to 0.623.

Pause. Try Exercise 3.1 before reading on.

Group B shows a capacity shortfall that more control cannot remove. A 1.2 MW burst for ten minutes takes the demand to 4.83 MW, the fuel cell reaches its inverter's AC rating at 4.10 MW of net output, and the UPS, from 30 %, holds only $0.25 \times 0.10 \times 0.95 = 0.024\,\mathrm{MWh}$ above its floor. It reaches the floor at 201.3 s, and from then on the shortfall falls on every feeder, so 68.0 kWh of critical IT is lost among 109.2 kWh unserved. A 0.5 MW recharge limit still leaves 68.0 kWh of

Table 3.5 The Ten Cases With One Measured Result as Shipped and the Completion Rule of Each

#	Case and events	Outcome as shipped	Required to complete (kind)
1	Manual, current 4996 → 5250 A after 30 s	voltage 0.7479 → 0.7397 V in 0.2 s, margin 0.733 → 0.671 by 5 s, fuel unchanged at 8.63 mol/s, U_f 0.750 → 0.788	step to ≥ 5200 A after 30 s, run 120 s, confirm (experiment)
2	Governor and feedforward off, IT +0.35 MW at 60 s	U_f 0.993 at 3.2 s, margin 0.157 at 5.4 s, foldback 20.0 s, UPS up to 0.465 MW, thermal limiter 107–157 s	see the margin below 0.25, confirm (demonstration)
3	Case 2 with feedforward and governor	carries the site by itself 13.2 s after the step, UPS 0.64 kWh, margin ≥ 0.623, U_f ≤ 0.800	see the fuel limit, then the site carried, confirm (experiment)
4	UPS at 35 %, recharge 0.1 MW, IT +0.9 MW 360–720 s	38.8 % at 360 s, floor at 705 s, 2.6 kWh unserved, 1.6 kWh critical	see the floor, recharge limit ≥ 0.3 MW, Reset, run 900 s, confirm (experiment)
5	Processor 55 s, estimate 15 s, IT +0.5 MW at 60, methane 85 % at 420 s	at 600 s $\hat{U}_f$ 0.685 against 0.805, guard holds 5974 of 6614 A, UPS 0.42 MW, 46.8 % at 900 s	see the guard bind, type 85 %, run 240 s, confirm (demonstration)
6	Blower to 50 % at 60, IT +0.4 MW at 420 s	outlet node 813.7 °C at 420 s, thermal override from 503 s, peak 825.0 °C, UPS up to 1.0 MW, 41.6 % at 900 s	see the override after 420 s, confirm (demonstration)
7	Heat demand 1.1 → 0.2 MW at 60 s	hot water peaks 80.97 °C, recovered heat 1.12 → 0.22 MW, nodes move under 0.1 K	bypass over 70 %, run to 600 s, confirm (exploration)
8	UPS at 30 %, IT +1.2 MW 60–660 s, cap off	fuel cell at most 4.10 MW of 4.83 MW, floor at 201 s, 68 kWh critical of 109 kWh unserved	see critical loss, arm the cap, Reset, run 720 s, confirm (demonstration)
9	Challenge, Table 3.7	utilization 1.048 and margin 0.158 at the first burst, fails	pass the shared verdict at 1200 s (challenge)
10	MPC, IT +0.5 MW 60–300 and 660–900 s, methane 0.90 at 420 s with the estimate at 1.00	peak U_f 0.807 and margin ≥ 0.568 with 5.58 kWh from the UPS, against 0.877, 53.2 s above 0.85 and 19.14 kWh under the governor	run 1200 s on the MPC, switch to the governor, run 1200 s, confirm (experiment)

critical IT unserved, and raising the governor's utilization limit to 0.85 leaves 67.5 kWh, because neither adds a watt of capacity. Only less load helps. With the cap armed the flexible quarter is shed at once by plan, the demand never passes 3.68 MW, nothing goes unserved, the UPS keeps 28.1 % or more and ends at 43.3 %, and 225.0 kWh of flexible work is deferred instead, as shown in Figure 3.2c and d. The capped run serves less and burns 1.896 MWh against 2.194 MWh, which is not a saving.

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

Group C is a model mismatch. The processor holds 55 s of feed against the estimator's 15 s, and at 420 s the gas falls to 85 % methane while the estimator assumes 100 %. At 600 s the estimate reads a utilization of 0.685 while the truth is 0.805, and only the voltage guard holds the stack, at 0.680 V per cell and 5974 A against 6614 A requested, so the UPS gives about 0.4 MW to the end and falls to 46.8 %. With 85 % typed at 480 s the estimate meets the truth (0.7885 against 0.7902 at 600 s), the feedforward opens the valve to its 12 mol/s limit and the current rises to 6209 A, and the UPS still gives 0.279 MW because twelve moles a second of this gas cannot carry the site at a safe utilization. The analyzer, used from the start, reaches the same state a couple of minutes later (0.794 against 0.776 at 600 s), and the UPS energy in the window falls from 48.37 to 34.65 and 36.50 kWh. Correcting only the processor estimate to 55 s helps only before the gas change, where the true utilization peaks at 0.816 instead of 0.887.

3.5.2 The Recharge That Was Not the Cause

Case 4 was built to show that restoring the reserve fast makes a second load step. It does, but a different mechanism dominates. With the suggested recharge (0.4 MW, 6 MW per unit) the UPS is at 49.1 % when the burst arrives, falls no lower than 28.4 %, and after the burst ends at 720 s the current waits for fuel for 28.9 s of the next three minutes. Setting the recharge limit to zero at 720 s leaves

Table 3.6 Three Matched Comparisons on the Simulator's Truth. No Run Tripped or Exceeded the Thermal Limits

Group, window	Setting	IAE (kWh)	Max e (MW)	Rec. (s)	Low m_a	Peak U_f	Held (s)	Fold (s)	Travel (mol/s)	UPS (MW)	UPS (kWh)	Unserved, crit., curt. (kWh)	SOC (%)
A. Rejection, case 2,	naive (shipped)	2.638	0.466	94.1	0.157	0.993	48.9	20.0	4.79	0.465	2.42	0, 0, 0	70 → 69.3
0–400 s, IT +0.35 MW	feedforward only	1.115	0.457	17.0	0.158	0.992	0	11.2	13.17	0.457	0.91	0, 0, 0	70 → 70.0
at 60 s	governor only	1.306	0.241	27.6	0.614	0.800	30.5	0	2.02	0.241	1.16	0, 0, 0	70 → 69.7
	both (suggested)	0.797	0.237	13.2	0.623	0.800	13.5	0	5.96	0.237	0.65	0, 0, 0	70 → 69.8
B. Capacity, case 8,	as shipped	136.5	1.081	695.1	0.584	0.801	632.7	0	16.89	0.798	26.52	109.2, 68.0, 0	30 → 25.4
0–900 s, IT +1.2 MW	recharge 0.5 MW	136.8	1.079	725.1	0.584	0.801	634.9	0	19.30	0.792	26.58	109.1, 68.0, 0	30 → 27.7
60–660 s	$U_{gov} = 0.85$	136.3	1.058	695.1	0.440	0.851	607.6	0	21.73	0.804	27.23	108.3, 67.5, 0	30 → 25.6
	workload cap	12.7	0.956	647.6	0.656	0.800	302.8	0	31.56	0.956	11.51	0, 0, 225.0	30 → 43.3
C. Mismatch, case 5,	$\hat{q} = 1.0$ (shipped)	48.37	0.426	n.r.	0.529	0.813	480.0	0	2.27	0.426	48.37	0, 0, 0	67.2 → 46.8
420–900 s, methane	$\hat{q} = 0.85$ typed at 480 s	34.65	0.402	n.r.	0.532	0.814	480.0	0	1.63	0.402	34.65	0, 0, 0	67.2 → 52.6
1.00 → 0.85 at 420 s	analyzer from the start	36.50	0.427	n.r.	0.542	0.811	480.0	0	1.89	0.427	36.50	0, 0, 0	67.2 → 51.8

Rec. is the recovery time from the event with a 20 kW band, its dwell the rest of the window (A: 245.9, 323.0, 312.4, 326.8 s, B: 144.9, 114.9, 144.9, 192.4 s), and n.r. means not inside the band at the window's end. Held counts the seconds with the admitted current below the request (A naive: the thermal limiter, A governed: the arriving-fuel ceiling, B: mostly the AC rating, C: the voltage guard), Fold the inverter's foldback, Travel the fuel valve state's. The current's own travel is mostly the noise of the measured voltage the law divides by (33 A/s before the step, 13.2 kA of the 14.3 to 18.6 kA in A). Fuel on LHV over the window: 0.9217, 0.9205, 0.9117, 0.9060 MWh (A), 2.194, 2.213, 2.204, 1.896 MWh (B), 1.163, 1.205, 1.200 MWh (C), with the same heat delivered in every run of a group. At 0.01 s no entry moves by more than one unit beyond the last digit shown except B's $U_{gov} = 0.85$ maximum error (1.027).

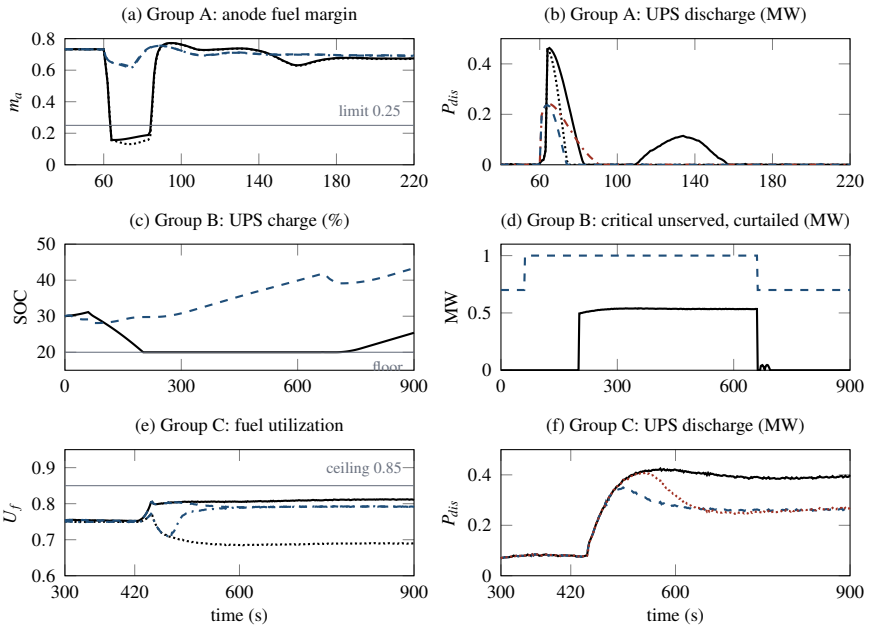

Figure 3.2 Matched comparisons A to C, drawn from the exported CSV records. (a) Case 2 after the 0.35 MW step at 60 s: true anode margin naive (solid) and with feedforward and governor (dashed), with the estimator's margin of each (dotted, dash-dotted), against the 0.25 limit. (b) UPS discharge of the four runs of group A: naive (solid), feedforward only (dotted), governor only (dash-dotted), both (dashed). (c) and (d) Case 8, as shipped (solid) and with the workload cap (dashed): true UPS charge against the 20 % floor, and the critical IT unserved as shipped with the flexible IT curtailed by the cap (dashed). (e) and (f) Case 5 after the gas changes at 420 s: true (solid) and estimated (dotted) utilization as shipped, the same pair with 85 % typed at 480 s (dashed, dash-dotted), and the UPS discharge as shipped (solid), typed (dashed) and with the analyzer (dotted).

the same 28.9 s, while setting the trim's gain to zero instead leaves 5.0 s. As the load falls, the trim's measurement, the utilization of the requested current over the estimated arriving hydrogen, reads the fuel still in the processor as a surplus and cuts the feed from 10.96 to 8.43 mol/s (9.27 mol/s with the trim's gain at zero), and the stack then waits for fuel to come back. The recharge's own cost appears with an aggressive setting (0.5 MW, 20 MW per unit, 60 s), which waits 41.6 s and lowers the margin to 0.575 against 0.621, and keeps the fuel cell at 3.93 MW against a 3.56 MW demand at 780 s. This is a property of any trim on a transport-delayed estimate, and the app's texts now state it.

3.5.3 The Combined Challenge

The challenge runs 1200 s from the design point with the UPS at 70 %. A 0.6 MW burst lasts four minutes from 120 s, the processor slows by 15 s over five minutes from 300 s while the estimator keeps 15 s, a 0.8 MW burst lasts four minutes from 600 s, and the heat demand falls to 0.2 MW at 840 s. One verdict function, shared by the scorecard, the completion rule and the tests, asks over the whole window for a true utilization below 0.85, an anode margin above 0.25, both nodes within 650 to 830 °C with the gradient within 1000 K/m, no trip, no critical IT unserved, the UPS never at its floor (lowest charge above 20.5 %) and at least 50 % at the end. Table 3.7 lists five policies run from the same initial state with the same meters and events, Figure 3.3 draws them as a chart, and Figure 3.4 compares three of them.

As shipped the first burst runs the anode short at once, the utilization reaches 1.048 at 122 s and the margin 0.158, and between the bursts the naive loop cycles on its own as the processor slows, its

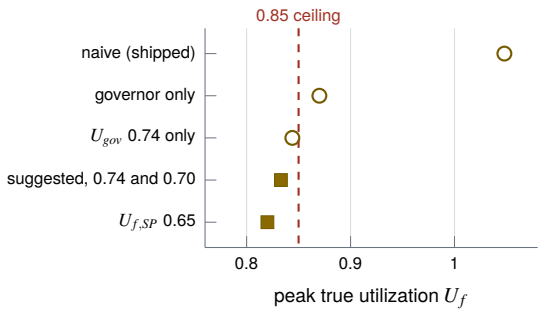

U_f	low m_a	UPS (kWh)	SOC low, end	verdict
1.048	0.158	43.4	53.1, 60.3 %	fails
0.870	0.412	31.2	57.4, 61.9 %	fails
0.844	0.466	55.6	46.6, 46.6 %	fails
0.833	0.492	26.1	59.4, 63.2 %	passes
0.820	0.517	20.7	61.6, 65.2 %	passes

Figure 3.3 The five policies of Table 3.7 as a chart: the peak true utilization over the 1200 s (marks and first column), with the lowest anode margin, the energy the UPS gave, its lowest and final charge and the verdict in the other columns. The suggested policy is the governor on with $U_{gov} = 0.74$ and $U_{f,SP} = 0.70$, the last row the same with $U_{f,SP} = 0.65$. Open circles fail, filled squares pass. Notice that the governor alone still crosses the dashed ceiling, and that lowering its limit without the setpoint stays under the ceiling but takes the most energy from the UPS.

Table 3.7 The Challenge Over 1200 s Under Five Policies. No Run Tripped, Left Critical IT Unserved or Left the Node and Gradient Limits

Policy	Peak U_f	Low m_a	UPS (MW)	UPS (kWh)	SOC low, end	Fuel (MWh)	Verdict
naive (shipped)	1.048	0.158	0.847	43.4	53.1, 60.3	2.863	fails
governor only	0.870	0.412	0.759	31.2	57.4, 61.9	2.778	fails
U_{gov} 0.74 only	0.844	0.466	0.873	55.6	46.6, 46.6	2.636	fails
suggested	0.833	0.492	0.772	26.1	59.4, 63.2	2.745	passes
$U_{f,SP}$ 0.65	0.820	0.517	0.784	20.7	61.6, 65.2	2.846	passes

The suggested policy is the governor on with $U_{gov} = 0.74$ and $U_{f,SP} = 0.70$, the last row the same with $U_{f,SP} = 0.65$. Nodes 704.7 to 801.5 °C and gradients 738 to 859 K/m over all five. The UPS starts at 70 % in every run and the heat delivered is 0.277 MWh in each.

margin falling below 0.25 at 387, 436, 490 and 550 s with the cell at the inverter's floor before the late fuel arrives. The governor alone fixes the bursts but not the slower processor, since at the second burst the estimator still assumes 15 s and admits current a little before the hydrogen is there, and the true utilization reaches 0.870 while the estimate reads 0.799. Lowering only the governor's limit to 0.74 leaves the trim's 0.75 setpoint above it, so the governor holds the current below the request for 1121 s of the 1200 s and the UPS ends at 46.6 %.

> **✦ DID YOU KNOW? A Safer Limit Alone Drained the UPS**
>
> From the simulator: in the challenge, lowering only the governor's limit to 0.74 drew 55.6 kWh from the UPS, more than the 43.4 kWh of naive control, and left it at 46.6 %, so the run failed (Table 3.7). The limit and the trim's setpoint have to move together.

Lowering the limit to 0.74 and the setpoint to 0.70 together passes, with the utilization at 0.833 at worst, the margin above 0.49 and the UPS ending at 63.2 %. The suggested run uses less UPS energy than the governor alone, 26.1 against 31.2 kWh, and 20.0 of its 26.1 kWh come in the second burst, whose 4.37 MW demand at its peak exceeds what the inverter's 4.3 MW AC rating delivers after the plant's own consumption. Of that energy 17.7 kWh lie above the rating, against 20.1 with the governor alone and 24.5 as shipped, and no governor setting can move it to the fuel cell. Over 31 seeds, the default and seeds 0 to 29, the shipped policy passes 0 times and the suggested one 31 times, at both 0.02 s and 0.01 s, with its peak utilization 0.832 to 0.833. Its state at 611 s is shown in Figure 3.5.

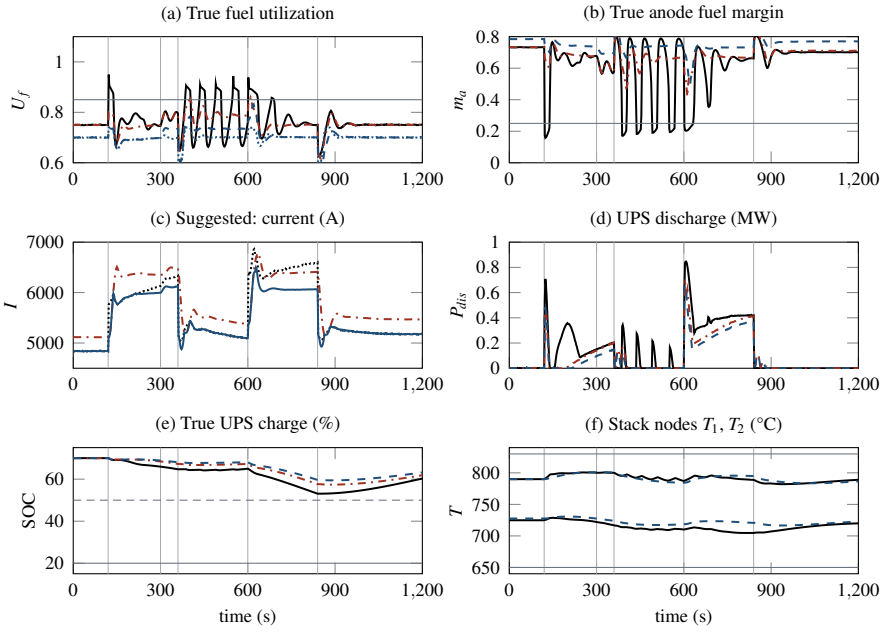

Figure 3.4 The challenge, from the exported records: naive as shipped (solid), the governor alone (dash-dotted) and the suggested policy (dashed), with the events at 120, 300, 360, 600 and 840 s marked (bursts begin, the processor slows, the first burst ends, the second begins, the second ends as the heat demand falls). (a) True utilization against the 0.85 ceiling, with the suggested run's estimate (dotted). (b) True anode margin against the 0.25 limit. (c) Suggested run: requested current (dotted), the arriving-fuel ceiling (dash-dotted) and the flowing current (solid). (d) UPS discharge. (e) True UPS charge against the 20 % floor and the 50 % end target. (f) Stack nodes T_1 and T_2 of the shipped (solid) and suggested (dashed) runs, inside the 650 to 830 °C limits.

> **TRY IT IN THE APP** **9 · Challenge: Workload Bursts, a Slower Fuel Processor and a Heat-Demand Drop**
>
> Load scenario 9 and run it as shipped at 100 times speed. At the first burst the utilization reaches 1.048 and the margin 0.158, and the verdict fails. Turn the Reference governor on, set the Governor utilization limit to 0.74 and the Fuel utilization setpoint to 0.70, then Reset and run the 1200 s. The utilization peaks at 0.833 and the verdict passes.

The margin to the ceiling is thin, 0.0175 of utilization. One factor at a time, with no retuning, the suggested setting fails for a true processor holdup above 20.75 s (the estimator keeping 15 s), a challenge slowdown above 20.05 s against its 15 s, or methane below 0.9626 with the estimate at 1.0, the direction in which an estimate errs dangerously, since gas leaner than assumed makes less hydrogen than the governor credits. A utilization setpoint of 0.65 under the same limit passes with 0.030 of margin for 3.7 % more fuel, 2.846 against 2.745 MWh. Two plan targets are read by the verdict as the plant allows. The plan asks for a positive anode fuel margin, but zero is unreachable, because the inverter folds back at 0.60 V per cell, naive control bottoms out at 0.157 and the stack trips at 0.10, so the verdict asks for 0.25, which lies beyond the 0.85 ceiling at steady state (0.49 there). A charge below 20 % is impossible by construction, so the verdict asks that the UPS never reach its floor.

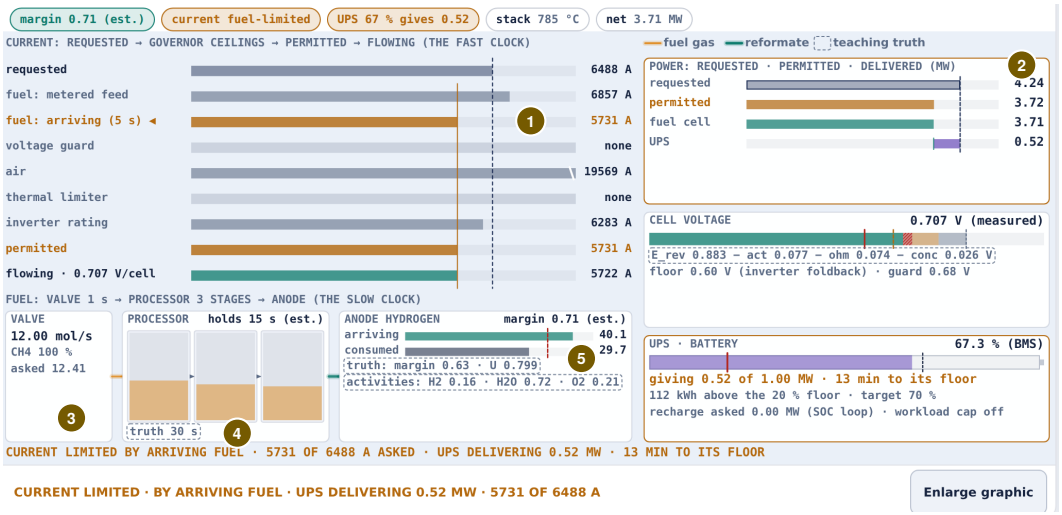

Figure 3.5 Suggested tuning in the challenge at $t = 611$ s on the Fuel and current view, eleven seconds into the second burst, cropped to the plant. (1) The current ladder shows 6488 A requested against ceilings of 6857 A on the metered feed, 5731 A on the arriving fuel, which binds, and 6283 A on the inverter's rating, with 5722 A flowing at 0.707 V per cell. (2) The fuel cell delivers 3.71 MW of the 4.24 MW asked and the UPS 0.52 MW from 67.3 %, 13 min from its floor. (3) The valve is at its 12.00 mol/s limit, (4) the processor's true holdup is 30 s against the estimator's 15 s, the estimator credits 40.1 mol/s of arriving hydrogen against 29.7 mol/s consumed, and (5) the anode's teaching truth reads a margin of 0.63 and a utilization of 0.799 against the estimate's 0.71 and 0.739.

3.5.4 The Predictive Controller Against the Governor

Table 3.8 compares the two structures on cases 3 to 6, 8 and 10 and the challenge over 1200 s, which holds every case's events, on the default seed, with identical instruments, noise and events and no preview for either structure. The governor keeps its shipped tuning ($U_{gov} = 0.80$, $U_{f,SP} = 0.75$), the challenge also runs with the governor on and with its suggested setting, and the MPC keeps case 10's settings everywhere. Ablations divide the credit between the MPC's information and its planning. They are the governor given $\hat{\theta}$ (a counterfactual built outside the application), the MPC with $\hat{\theta}$ off, and the MPC with its fuel capped at the larger of its steady reference and the fuel applied, so that none is sent ahead. The hindsight governor is the best of 96 fixed settings on each case (U_{gov} from 0.74 to 0.83, $U_{f,SP}$ from 0.65 to 0.75, two transient horizons, the feedforward on or off and two recharge loops) by the least UPS plus unserved energy with $U_f < 0.85$, $m_a > 0.25$ and no trip. Every MPC run solved 1200 plans with none rejected, and only the runs with $\hat{\theta}$ off paid a slack (11, 12 and 20 plans in cases 5, 10 and the challenge).

In case 10, two 0.5 MW bursts run from 60 to 300 and from 660 to 900 s, and at 420 s the gas falls to 90 % methane while the estimate stays at 100 %. At the first burst the plan holds the valve at 12 mol/s for 7.0 s, and the current reaches the request in 13.0 against 16.0 s under the governor, with 1.85 against 2.73 kWh from the UPS between 60 and 360 s. After the gas change $\hat{\theta}$ falls below 0.91 within 27.6 s and settles at 0.900, and the true utilization stays at 0.807 or below. The governor still credits pure methane, the thinner anode lowers the voltage, and the same power draws more current, as shown in Figure 3.6, the positive feedback of the Introduction opposed only by the voltage guard. The governor's utilization spends 53.2 s above 0.85, all after the gas change, and in the second burst the guard holds the current for 280.3 s while the UPS gives 14.04 against 3.58 kWh. Over the case the UPS gives 19.14 against 5.58 kWh on 2.7716 against 2.7010 MWh of fuel.

Table 3.8 The Governor Against the Predictive Controller Over 1200 s on the Default Seed, With the Rows That Separate the Estimate From the Plan. No Run Tripped, and Only the Naive Challenge Took the True Margin Below 0.25

Case	Structure	Peak U_f	Above 0.85 (s)	$\int (U_f - 0.85)^+ dt$ (s)	Low m_a	UPS (kWh)	Unserved (kWh)	Critical (kWh)	Fuel (MWh)
10 · methane 0.90	governor	0.877	53.2	0.384	0.421	19.14	0	0	2.7716
	MPC	0.807	0	0	0.568	5.58	0	0	2.7010
	MPC, $\hat{\theta}$ off	0.880	36.9	0.219	0.403	10.91	0	0	2.7684
	MPC, no fuel sent ahead	0.807	0	0	0.564	6.63	0	0	2.7014
	governor + $\hat{\theta}$	0.805	0	0	0.559	6.45	0	0	2.6996
	governor, $\hat{q} = 0.90$ at 425 s	0.800	0	0	0.568	6.32	0	0	2.7004
	governor, lead-lag 15/8 s	0.888	80.1	0.842	0.365	23.87	0	0	2.7588
	hindsight / MPC (0.74, 0.65)	0.822 / 0.744	0	0	0.513 / 0.664	2.52 / 2.66	0	0	2.7102 / 2.7546
5 · processor 55 s	governor	0.887	10.1	0.113	0.382	88.51	0	0	2.9158
	MPC	0.825	0	0	0.549	62.21	0	0	2.9886
	MPC, $\hat{\theta}$ off	0.870	11.8	0.103	0.430	62.90	0	0	2.9900
	governor + $\hat{\theta}$	0.825	0	0	0.538	64.64	0	0	2.9899
	hindsight / MPC (0.74, 0.65)	0.841 / 0.763	0	0	0.488 / 0.661	51.71 / 54.14	0	0	3.0230 / 2.9173
9 · challenge	naive (shipped)	1.048	146.5	5.651	0.158	43.41	0	0	2.8630
	governor on	0.870	7.6	0.047	0.412	31.22	0	0	2.7776
	suggested (0.74, 0.70)	0.833	0	0	0.492	26.10	0	0	2.7447
	MPC	0.814	0	0	0.559	28.16	0	0	2.7680
	MPC, $\hat{\theta}$ off	0.854	1.6	0.003	0.474	29.55	0	0	2.7835
	governor on + $\hat{\theta}$	0.815	0	0	0.537	30.04	0	0	2.7594
	hindsight / MPC (0.77, 0.65)	0.844 / 0.783	0	0	0.473 / 0.617	20.55 / 19.41	0	0	2.8470 / 2.8696
3 · step	governor	0.800	0	0	0.623	0.70	0	0	2.8323
	MPC	0.801	0	0	0.629	0.59	0	0	2.8344
	MPC, no fuel sent ahead	0.801	0	0	0.618	0.84	0	0	2.8364
	governor, lead-lag 15/2 s	0.800	0	0	0.626	0.54	0	0	2.8325
	hindsight / MPC (0.77, 0.65)	0.761 / 0.758	0	0	0.694 / 0.699	0.26 / 0.26	0	0	2.9103 / 2.9163
4 · reserve	governor	0.800	0	0	0.584	44.69	2.62	1.60	2.6794
	MPC	0.801	0	0	0.589	44.66	2.54	1.56	2.6815
	hindsight / MPC (0.77, 0.65)	0.770 / 0.774	0	0	0.646 / 0.646	40.02 / 39.96	0	0	2.7956 / 2.8014
6 · blower 50 %	governor	0.800	0	0	0.637	118.87	0.39	0.24	2.1875
	MPC	0.802	0	0	0.644	117.50	0	0	2.1956
	hindsight / MPC (0.74, 0.65)	0.740 / 0.744	0	0	0.724 / 0.726	53.08 / 52.80	0	0	2.5674 / 2.5697
8 · low reserve	governor	0.801	0	0	0.584	26.52	109.20	68.00	2.8844
	MPC	0.802	0	0	0.587	26.82	108.93	67.85	2.8959
	hindsight / MPC (0.77, 0.65)	0.768 / 0.770	0	0	0.649 / 0.644	28.25 / 28.25	97.08 / 97.27	60.41 / 60.54	3.0157 / 3.0169

Pairs in parentheses are $(U_{gov}, U_{f,SP})$, and a slash separates the hindsight governor from the MPC run at the same pair, a single entry applying to both. The hindsight governors keep the recharge loop as shipped and a transient horizon of 5 s, with the feedforward off in cases 3 and 6. The lead-lag is a 15 s lead over the best of 2, 3, 5 and 8 s lags, and $\hat{q}$ typed is the operator's estimate set 5 s after the gas change. Utilization and margin are the simulator's truth at every physics step, and the energies come from the plant's ledgers.

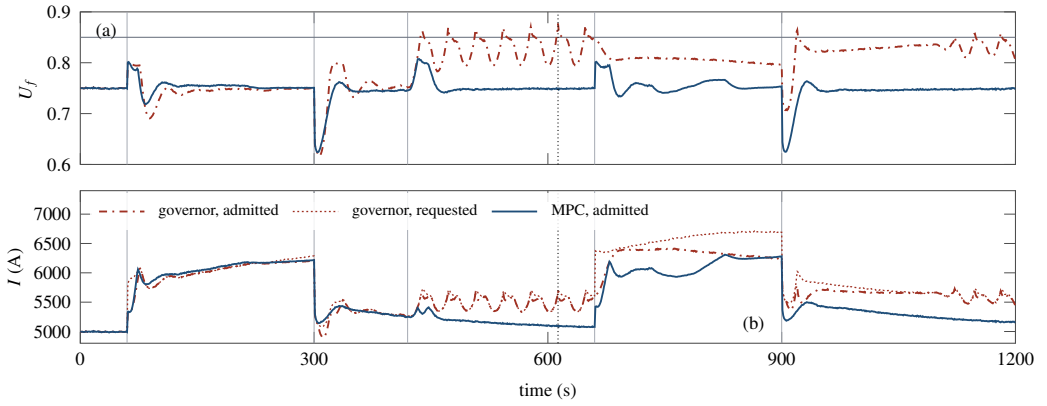

Figure 3.6 Case 10 under the governor (dash-dotted) and the predictive controller (solid), from the exported records. (a) True fuel utilization against the 0.85 ceiling. (b) Admitted current, with the governor's request dotted where its voltage guard holds the current below it. Thin lines mark the bursts at 60 and 660 s, their ends at 300 and 900 s and the gas change at 420 s. At the dotted line, 613 s, both serve the same 3.43 MW demand, the governor with 5689 A at 0.670 V per cell and a utilization of 0.877, the MPC with 5097 A at 0.736 V and 0.749.

The ablations assign the result. With $\hat{\theta}$ off the MPC reaches 0.880, while the governor given $\hat{\theta}$ holds 0.805, and 0.800 with its estimate typed to 90 % after the change. The utilization result belongs to the estimate read from the cell voltage, which the governor could read as well. The plan's own gain is the fuel sent ahead, 5.58 kWh against 6.63 kWh without it and 6.45 kWh for the governor given $\hat{\theta}$. A lead-lag feedforward also sends fuel ahead and beats the plan on case 3's clean step, 0.54 against 0.59 kWh, but without $\hat{\theta}$ it leaves case 10 at 0.888. In case 5, whose processor holds 55 s against 15 s, the governor peaks at 0.887 in the first burst and the MPC at 0.825, because $\hat{\theta}$ reads the late hydrogen as leaner gas, and in the challenge with the governor on the peaks are 0.870 and 0.814. Where the plant matches the estimate (cases 3, 4, 6 and 8) the two structures peak within 0.002 of each other and use UPS energy within 1.4 kWh. The rankings hold for seeds 11 to 13 and a 0.01 s step, the MPC peaking at 0.807 to 0.808, 0.825 and 0.814 against 0.877 to 0.878, 0.887 to 0.888 and 0.870 to 0.871 in cases 10, 5 and the challenge.

The challenge's rule is unchanged, and the MPC at its shipped settings meets it, with a peak of 0.814, a margin of at least 0.559 and 63.6 % left in the UPS. Two limits the verdict does not score are passed under it as under the suggested governor. The node heating rate on the limiter's filtered reading reaches 7.25 K/min (7.22) against a 5 K/min alarm, and the AC output reaches 4.3048 MW (4.3051 MW) against the 4.3 MW rating, which the inverter applies on the measured voltage and which holds the MPC's current for 96.3 and 133.3 s in the two bursts.

On stored energy the MPC loses to a tuned governor. The hindsight governor needs less UPS plus unserved energy than the MPC in every case, 0.26 against 0.59 (case 3), 40.02 against 47.20 (4), 51.71 against 62.21 (5), 53.08 against 117.50 (6), 125.33 against 135.75 (8), 2.52 against 5.58 (10) and 20.55 against 28.16 kWh (the challenge), in all seven with a utilization setpoint of 0.65 and 0.3 to 16.9 % more fuel. Held to no more fuel than the MPC, it still wins cases 3, 4, 5 and 8 and the challenge (0.27, 43.07, 61.19, 130.64 and 26.10 kWh) and loses cases 6 and 10 (118.41 and 34.56 kWh). Run at the hindsight limits, the MPC comes within 6 % of that governor, losing cases 5 and 10 (54.14 against 51.71 and 2.66 against 2.52 kWh) and winning the challenge (19.41 against 20.55 kWh), while its true peak stays lower wherever the plant departs from the estimate (0.763, 0.744 and 0.783 against 0.841, 0.822 and 0.844). The operating limits decide the stored energy more than the structure does.

3.5.5 Verification, Reference Data and Cost

These figures check the implementation and claim nothing about control performance. A script written apart from the application recomputes every boundary flux at each step's start and end and integrates it by the trapezoid rule over nine disturbed runs. The plant's energy closes to 4.5×10^{-7} of the fuel's LHV, the atoms to 3.2×10^{-7} of their throughput, the hot water and the data hall to 3.5×10^{-7} and 8.5×10^{-8} and the UPS's stored energy to 4×10^{-5}, all trapezoid error, and the bus residual is 0 W at every step with no negative term and no simultaneous charge and discharge. No step evaluated an undefined voltage, and the lowest hydrogen activity was 0.0746. With the stack tripped at 120 s the UPS at its 1 MW rating reaches its floor at 545.32 s against the analytic $120 + (E(120) - 0.2E_{cap})\eta_{dis}/1\,\mathrm{MW} = 545.31\,\mathrm{s}$. The design point recomputed by hand agrees with the application to 2.1×10^{-12}, halving the step changes no challenge verdict, and all 89 of the application's checks pass, 22 of them added with the predictive controller. One 1200 s challenge takes 0.46 s headless and 12.0 s at the browser's top speed of 100 times real time.

Two sensitivities bound the lessons. Case 6's thermal override comes at 503.1 s with 3 MJ/K per half, at 695.1 s with 6 MJ/K and not within 900 s with 15 MJ/K, and case 2's heating-rate hold of 48.9 s disappears from 6 MJ/K, while the three challenge verdicts are the same from 1.5 to 15 MJ/K. With $\rho_0 = 0.112\,\Omega\,\mathrm{cm}^2$, an aged stack, the voltage guard binds early and the governor alone passes, and with $0.08\,\Omega\,\mathrm{cm}^2$ the suggested setting keeps only 0.0066 of margin. Agreement with reference data is the weaker claim of the E°, power-density and efficiency checks above, the model has not met measurements of a stack, and no learner study has been made.

3.6 Technology in Practice

The governor and the loops of Section 3.4 act through an inverter, a fuel valve, a blower and two dampers, and they read meters, thermocouples and an analyzer. Figure 3.7 draws the current and fuel loops with that hardware in place of the schematic's symbols.

What Measures It

No instrument reads the anode's hydrogen, so the loops read what surrounds it: the stack current and voltage, the metered fuel, the gas quality and the stack temperatures. The temperatures come from thermocouples. The most common failure of a thermocouple is to fail open, or burn out, and an open input can pick up electrical noise that reads as a wildly varying temperature, so the instrument is set to read high or low on a burnout, whichever is the safer side. Where too much heat is the danger, a high reading drives the controller toward the safe state [21]. The chapter's stack thermocouples lag 5 s with 0.3 K of noise, and the gas quality comes from a sampled analyzer, whose principle Figure 3.8 shows.

> **TECHNOLOGY Sensor: A Process Gas Chromatograph for the Gas Quality**
>
> The methane fraction that the estimator needs comes from an online analyzer, and the most common one in continuous process analysis is the gas chromatograph. A sample valve injects a precise quantity of gas into a long, narrow column, a carrier gas washes it through, and each species leaves the column at its own time, so a simple detector sees a train of peaks [21]. The analyzer's computer can turn the concentrations into the heating value of natural gas. Time is the price. Cycle times of minutes are common even for online chromatographs, and enough dead time makes a feedback loop oscillate [21]. The chapter's analyzer draws a sample every 60 s through a 30 s line and analyzes it for 60 s, so each reading is 90 to 150 s old.

What Moves It

The inverter sets the current, the fuel valve sets the gas to the processor, and the blower and the two dampers set the air and the heat. Between the valve and the anode the processor is itself a chain of lags, and enough lags in series make a response that looks dead at first [21], which is why the chapter's three stages need 17.3 s to deliver 63 percent of a fuel step's hydrogen.

plant controller: EST-312, JIC-300 and JY-300 every 0.1 s, FIC-310 and TY-320 every 1 s. The inverter follows I_{adm} through a 50 ms lag and folds back at 0.60 V per cell.

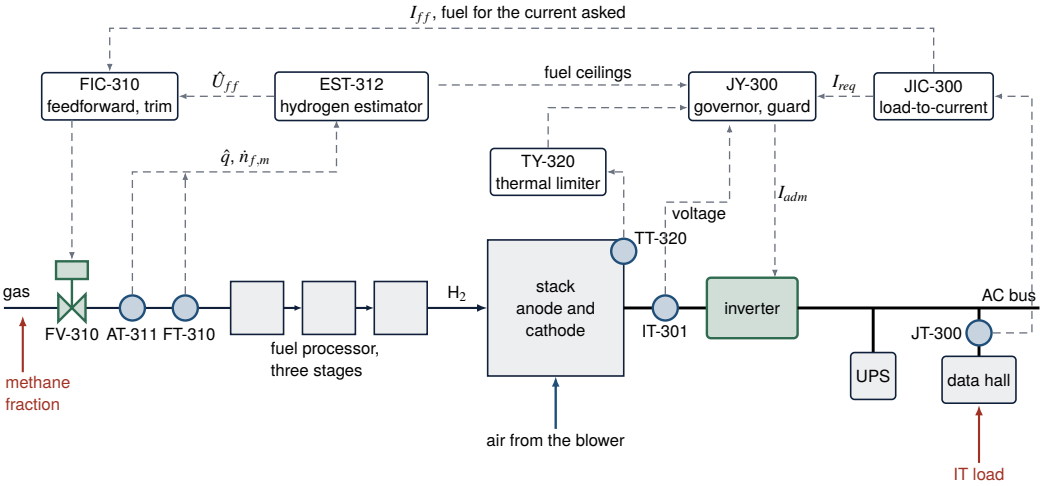

Figure 3.7 The current and fuel loops drawn with their hardware. The site demand meter JT-300 feeds the load-to-current law JIC-300, which asks for I_{req}. The estimator EST-312 turns the metered fuel FT-310 and the gas quality of the analyzer AT-311 into the hydrogen arriving, and the governor JY-300 admits the least of its fuel ceilings, the voltage guard on the stack meter IT-301, the thermal limiter TY-320 on the thermocouple TT-320 and the request. The inverter draws the admitted current. FIC-310 opens the fuel valve FV-310 for the current asked and trims it on the estimated utilization. The air, preheat and hot-water loops, the reserve loop and the workload cap are left out. Notice that the inverter acts on the fast clock while every fuel ceiling waits for the three processor stages.

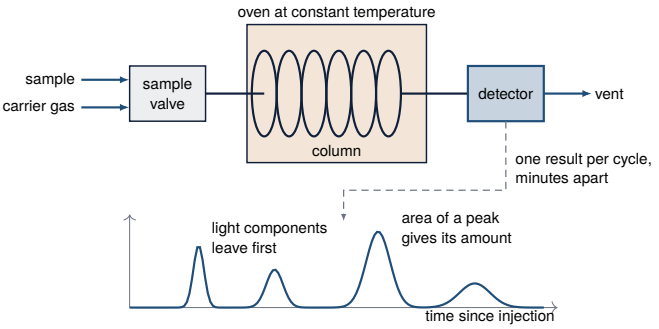

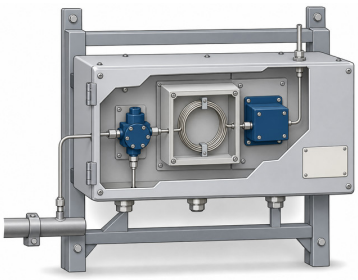

Figure 3.8 Principle of a process gas chromatograph, one kind of online analyzer for the gas quality. A valve injects a sample into carrier gas, the column holds back each species for its own time, and the detector sees them leave one by one. Notice that a result exists only at the end of each cycle, which is why the chapter's analyzer reports a methane fraction 90 to 150 s old.

TECHNOLOGY **Final Element: The Inverter That Sets the Stack Current**

A fuel cell stack is an unregulated voltage source whose voltage falls as its load rises [1]. A direct-current converter boosts that voltage and is the device that actually draws the current from the stack, and an inverter then makes alternating current by pulse width modulation at high frequency [1, 17]. The conversion is typically 92 to 96 % efficient. A specification for a small stand-alone unit asks the inverter to shut down when its input voltage dips below a minimum and names IEEE Standard 929 [17]. The chapter's inverter follows its command through a 50 ms lag and folds back at 0.60 V per cell.

Where the Controller Runs

A fuel cell system has three primary subsystems: the stack, the fuel processor and the power conditioner [1]. Programmable logic controllers (PLCs) manage local processes by feedback under a supervisory scheme, and loops at both levels run at cycle times from milliseconds to minutes [22]. The chapter's controller runs the estimator, the load-to-current law and the governor every 0.1 s, the other loops every 1 s and the reserve loop every 5 s. Its governor is a low selector, a function that passes the least of several signals, and a control selector block can report its selection back so that a de-selected controller stops integrating [21], the role of `track()` in Section 3.4.

> **IN PRACTICE** **Keeping a Fuel Cell Plant in Service**
>
> A fuel cell plant has few moving parts, but its catalysts and its stack wear out on a schedule. Major overhauls replace the shift catalyst every 3 to 5 years, the reformer catalyst every 5 years and the stack every 5 to 10 years, and filters are changed every 2,000 to 4,000 hours [1]. Early commercial units were watched by remote monitoring with an allowance for predictive maintenance. Fuel quality is watched as well. A fuel of lower energy content gives a less concentrated hydrogen-rich gas at the anode and some loss of stack performance, and sulfur and chlorides have to be removed before the processor's catalyst [1], the leaner gas of Section 3.5.1.

What Is Modernizing It

Three developments touch the loops of this chapter. Internally reforming cells, which join the endothermic reformer and the exothermic cell, are harder to design and control but are expected to take a significant market share [1]. The chapter's plant reforms outside the stack, and its Conclusion names internal reforming for a fuller model. Process chromatographs send several concentrations over a digital bus, and temperature programming of the column shortens their dead time [21]. Ethernet, Internet Protocol and wireless devices replacing older proprietary networks raise the likelihood of cybersecurity incidents in control systems [22]. The predictive controller of Section 3.4.6, which reads the gas quality from the cell voltage within 27.6 s where the analyzer needs 90 to 150 s, is a fourth development, and the chapter measures both its gains and its costs.

3.7 Limitations

All sizes and thresholds are synthetic teaching values set against the NETL handbook and the EPA catalog [1, 17] and fitted to no stack beyond the two conductivities behind E_{Ω}. The processor has no heat balance of its own, the burner has no temperature limit (after a trip it briefly burns the feed still in transit, to about 1340 °C on the truth), the two-node stack shows a gradient but no hot spot, and the light stack brings case 6's override and case 2's heating-rate hold earlier than a real stack would. The protection reads the true margin while the controllers do not, the estimator knows the plant's declared anode parameters, and the trim's reading of processor holdup as surplus remains in the shipped tuning. No forecast is present, the predictive controller's quadratic program is the only optimizer, no case outcome qualifies equipment, and none is a savings or reliability estimate.

> **OPEN CHALLENGE** **Stack Life Against Load Following**
>
> The high temperature that lets a solid oxide cell reform its own fuel also adds materials and mechanical design difficulties that shorten stack life and raise cost. The recommended stack rebuild every 5 to 10 years is expensive [1]. Mismatched thermal expansion and hard sealing between flat cells limit thermal cycling and stack life [17], and long heat-up and cool-down periods keep these cells out of many cyclic duties [1]. The chapter's gradient proxy is a coarse stress indicator, not a prediction of lifetime. Progress would be a controller that follows the load while it counts the wear each transient costs the stack.

The predictive controller's model takes the processor's holdup and the gas quality from the estimator's commissioning values and corrects them only through $\hat{\theta}$, which reads every unexplained

voltage deficit as fuel. In case 10, a stack whose area-specific resistance is 0.12 or 0.14 $\Omega\,\mathrm{cm}^2$ instead of 0.10 reads as leaner gas, on the safe side, with $\hat{\theta}$ up to 0.070 and 0.135 below the true ratio and peaks of 0.773 and 0.751 on 2.7690 and 2.9115 MWh of fuel against 2.7010. A stack at 0.085 or 0.07 $\Omega\,\mathrm{cm}^2$ reads as richer gas, with $\hat{\theta}$ up to 0.048 and 0.097 high, and the MPC peaks at 0.830 and 0.850, the second above the ceiling for 0.6 s, against 0.888 and 0.889 under the governor. A processor 20 or 40 s slower than commissioned is absorbed in part, with peaks of 0.818 and 0.825 against 0.878 and 0.888. The model also holds the stack temperatures and the air over the horizon, leaves out the UPS and the anode's pressure dynamics, and shares one slack per soft row across all 60 steps.

> ⚠ **OPEN CHALLENGE** **Storage Sized for the Worst Surge**
>
> Fuel delivery keeps fuel cells slow to follow a sudden rise in data center demand, and the gap becomes a temporary power shortfall. Both answers cost something. Capping the servers' power limits their performance and can break service agreements, while storage alone has to be oversized for the largest surge, which greatly raises the total cost of ownership. For a fuel cell it is the ramp rate that needs capping, not the peak, since the cell eventually meets the peak demand [8]. The chapter's workload cap acts on fixed thresholds of the UPS charge and output. Progress would be one supervisor that sets the store, the cap and the fuel limits together against the measured surges of a real hall.

3.8 Conclusion

This chapter has described a solid oxide fuel cell plant for an islanded data center whose current moves on the inverter's 50 ms, whose anode answers in 1.38 s and whose fuel arrives in 17.3 s. On an identical 0.35 MW step, naive load-to-current control takes the margin to 0.157 and leaves the inverter's foldback to stop the current, the feedforward alone changes nothing at the anode, and the governor with the feedforward keeps the margin at 0.623 and cuts the power IAE over 400 s from 2.638 to 0.797 kWh. When the estimator believes the gas is pure methane, only the voltage guard holds the stack, and when a burst exceeds the inverter's rating from a low reserve, only shedding load protects the critical IT. An added predictive controller keeps case 10's utilization at 0.807 where the governor reaches 0.877, through an estimate of gas quality read from the cell voltage that holds the governor at 0.805 as well, and its plan adds the fuel sent ahead, 5.58 against 6.63 kWh from the UPS without it. In the challenge, a governor limit of 0.74 with a setpoint of 0.70 passes on 31 of 31 seeds with 0.0175 of utilization to spare. Its stack is light, its processor has no heat balance of its own and its thresholds are synthetic. A fuller model would add internal reforming, a processor with its own heat balance, a distributed anode that can deplete locally, a heavier stack, and a predictive governor that knows the processor's holdup. The simulator shows that the current of a fuel cell has to wait for its fuel, and that the wait has to be paid for from storage.

KEY POINTS

- A fuel cell's current can change in a fraction of a second, but its fuel arrives many seconds later, so the current has to wait and storage carries the load.
- On an identical 0.35 MW step, the governor with the feedforward keeps the anode margin at 0.623 where naive control takes it to 0.157.
- When the estimator believes the gas is pure methane, only the voltage guard holds the stack, and when a burst exceeds the inverter's rating from a low reserve, only shedding load protects the critical IT.
- In the challenge, a governor limit of 0.74 with a setpoint of 0.70 passes on 31 of 31 seeds.

Where else this applies. A low selector of limits on one actuator also guards the fuel of the jet engine of Volume 4, Chapter 3, and the same site bus and UPS carry the gas engines of Chapter 2.

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] E. G. Gilbert and K. T. Tan, "Linear systems with state and control constraints: the theory and application of maximal output admissible sets," *IEEE Transactions on Automatic Control*, vol. 36, no. 9, pp. 1008–1020, 1991.

[3] 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.

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

[5] B. J. Spivey and T. F. Edgar, "Dynamic modeling, simulation, and MIMO predictive control of a tubular solid oxide fuel cell," *Journal of Process Control*, vol. 22, no. 8, pp. 1502–1520, 2012.

[6] B. J. Spivey, J. D. Hedengren, and T. F. Edgar, "Constrained control and optimization of tubular solid oxide fuel cells for extending cell lifetime," in *2012 American Control Conference (ACC)*. IEEE, 2012, pp. 1356–1361.

[7] J. Sun and I. V. Kolmanovsky, "Load governor for fuel cell oxygen starvation protection: a robust nonlinear reference governor approach," *IEEE Transactions on Control Systems Technology*, vol. 13, no. 6, pp. 911–920, 2005.

[8] Y. Li, D. Wang, S. Ghose, J. Liu, S. Govindan, S. James, E. Peterson, J. Siegler, R. Ausavarungnirun, and O. Mutlu, "Techniques for efficiently handling power surges in fuel cell powered data centers: Modeling, analysis, results," arXiv preprint arXiv:1806.04506, 2018. [Online]. Available: https://arxiv.org/abs/1806.04506

[9] E. Achenbach, "Three-dimensional and time-dependent simulation of a planar solid oxide fuel cell stack," *Journal of Power Sources*, vol. 49, no. 1–3, pp. 333–348, 1994.

[10] J. Padullés, G. W. Ault, and J. R. McDonald, "An integrated SOFC plant dynamic model for power systems simulation," *Journal of Power Sources*, vol. 86, no. 1–2, pp. 495–500, 2000.

[11] Y. Zhu and K. Tomsovic, "Development of models for analyzing the load-following performance of microturbines and fuel cells," *Electric Power Systems Research*, vol. 62, no. 1, pp. 1–11, 2002.

[12] F. Mueller, J. Brouwer, F. Jabbari, and S. Samuelsen, "Dynamic simulation of an integrated solid oxide fuel cell system including current-based fuel flow control," *Journal of Fuel Cell Science and Technology*, vol. 3, no. 2, pp. 144–154, 2006.

[13] P. Aguiar, C. S. Adjiman, and N. P. Brandon, "Anode-supported intermediate-temperature direct internal reforming solid oxide fuel cell. II. Model-based dynamic performance and control," *Journal of Power Sources*, vol. 147, no. 1–2, pp. 136–147, 2005.

[14] Y. Komatsu, S. Kimijima, and J. S. Szmyd, "Numerical analysis on dynamic behavior of solid oxide fuel cell with power output control scheme," *Journal of Power Sources*, vol. 223, pp. 232–245, 2013.

[15] N. Q. Minh, "Ceramic fuel cells," *Journal of the American Ceramic Society*, vol. 76, no. 3, pp. 563–588, 1993.

[16] S. Singhal, "Advances in solid oxide fuel cell technology," *Solid State Ionics*, vol. 135, no. 1–4, pp. 305–313, 2000.

[17] U.S. Department of Energy, National Energy Technology Laboratory, "Fuel cell handbook," U.S. Department of Energy, Office of Fossil Energy, National Energy Technology Laboratory, Morgantown, WV, Report, Seventh Edition, Contract DE-AM26-99FT40575, Nov. 2004. [Online]. Available: https://netl.doe.gov/sites/default/files/netl-file/FCHandbook7.pdf

[18] D. Sarantaridis and A. Atkinson, “Redox cycling of Ni-based solid oxide fuel cell anodes: A review,” *Fuel Cells*, vol. 7, no. 3, pp. 246–258, 2007.

[19] 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.

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

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

3.9 Exercises

3.1 Concept check · 2 min

Under naive load-to-current control, with the governor and the feedforward off, the IT load steps up by 0.35 MW. The current rises at once, and the anode fuel margin then goes on falling for several seconds. Which mechanism drives the continued fall?

(a) The utilization trim cuts the fuel valve as the utilization rises, so less hydrogen is fed

(b) The stack cools under the higher current, so its resistance rises and the voltage falls

(c) The UPS takes the load, and the inverter lowers the current to leave room for its recharge

(d) The thinner hydrogen lowers the cell voltage, and the law then asks for still more current

3.2 Concept check · 1 min

Case 1 runs in Manual with the fuel valve and the air held. The stack current is stepped by about 5 %. In what order do the cell voltage, the anode fuel margin and the hydrogen arriving at the anode respond?

(a) The voltage within 0.2 s, the margin within seconds, the arriving hydrogen not at all

(b) The arriving hydrogen first, then the margin, and the voltage last of the three

(c) The margin within 0.2 s, the voltage within seconds, the arriving hydrogen after 17 s

(d) All three together, on the processor’s clock of about 17 s to 63 percent

3.3 Concept check · 2 min

In case 8 the UPS starts at 30 % and a 1.2 MW burst takes the demand above what the inverter can deliver for ten minutes. Which change keeps all of the critical IT load served?

(a) Raising the recharge power limit to 0.5 MW, so the UPS refills faster

(b) Arming the workload cap, so the flexible quarter of the IT load is shed

(c) Raising the governor’s utilization limit to 0.85, so more current is admitted

(d) Switching to the predictive controller, so fuel is sent ahead of the burst

3.4 Theory and calculation · 5 min

The stack has $N_s = 1000$ cells in series and carries $I = 4996\,\mathrm{A}$ at the design point. The valve delivers 8.63 mol/s of pure methane, and the processor makes four moles of hydrogen of each. Take $F = 96\,485\,\mathrm{C/mol}$.

(a) Find the hydrogen the current reacts and the fuel utilization. *Checkpoint:* 25.89 mol/s and 0.750.

(b) Find the steady ceiling I_{ss} of Equation (3.9) at $U_{gov} = 0.80$, and the utilization if the current were stepped to 5250 A with the fuel held, as in case 1. *Checkpoint:* 5329 A and 0.788.

(c) Find the fuel the feedforward of Equation (3.12) sends for 4996 A at $U_{f,SP} = 0.75$. *Checkpoint:* 8.63 mol/s.

3.5 Theory and calculation · 7 min

The valve follows its command through a 1 s first-order lag, and the processor holds the feed in three equal stages of $\tau_s = \tau_{fp}/3$ each. The hydrogen arriving at the anode therefore answers a valve step through four first-order lags in series.

(a) Find the mean delay of the chain at the design holdup $\tau_{fp} = 15\,\text{s}$. *Checkpoint:* 16 s.

(b) For the three stages alone, the share of a step delivered at time t is $1 - e^{-x}(1 + x + x^2/2)$ with $x = t/\tau_s$. Find the x at which it reaches 63.2 %, the time that gives at $\tau_s = 5\,\text{s}$, and an estimate for the whole chain made by adding the valve's 1 s. *Checkpoint:* $x \simeq 3.26$, 16.3 s, about 17.3 s.

(c) In case 5 the processor holds 55 s of feed while the estimator keeps 15 s. Estimate the mean delay and the 63 percent time the same way. *Checkpoint:* 56 s and about 61 s.

3.6 Theory and calculation · 5 min

The UPS stores 0.25 MWh, keeps a floor at 20 % of it, discharges at up to 1 MW and delivers with an efficiency of 0.95.

(a) From its 70 % start, find the energy it can deliver above the floor and how long that lasts at its rating. *Checkpoint:* 0.119 MWh and 7.1 min.

(b) In case 8 it starts at 30 %. Find the energy it can deliver, and how long that covers the gap between the 4.83 MW demand and the 4.10 MW the fuel cell delivers at its inverter's rating. *Checkpoint:* 0.0238 MWh and about 117 s.

(c) The run of case 8 reaches the floor at 201.3 s, 141 s after the burst begins at 60 s. Give one reason it outlasts the estimate of part (b). *Checkpoint:* The UPS charges before the burst, and the demand builds up.

3.7 App challenge · 8 min

Load scenario 2, Naive load-to-current: a workload step and the fuel margin, and change no setting. Run to 400 s at a speed of 20 with the default seed, export the CSV file, and read the scorecard tiles Lowest anode fuel margin and Fuel utilization peak.

(a) Report both tiles, and the column `metric_ups_kWh` in the row at 400 s. *Checkpoint:* 0.16, 0.993 and 2.42 kWh, within 0.02 kWh.

(b) Press Reveal suggested tuning, which turns the feedforward and the governor on, press Reset, run to 400 s again and report the same three values. *Checkpoint:* 0.62, 0.800 and 0.65 kWh, within 0.02 kWh.

(c) The naive run puts no limit of its own on the current. Explain why it nonetheless draws more energy from the UPS than the governed run. *Checkpoint:* The foldback and the thermal limiter hold its current.

3.8 App challenge · 9 min

Load scenario 5, A slow fuel processor and lower-quality gas: current limitation and UPS energy, and change no setting. The processor holds 55 s of feed while the estimator keeps 15 s, and the gas falls to 85 % methane at 420 s. Run to 900 s at a speed of 20 with the default seed and export the CSV file.

(a) In the row at 600 s, read `Uf_est` and `truth_Uf`, and `I_req_A` and `I_meas_A`. Divide the estimate by the truth. *Checkpoint:* 0.685 against 0.805, ratio 0.85; 6614 A against 5973 A, within 10 A.

(b) Name the ceiling that holds the current, from the current card or the column `limit`, and explain the ratio of part (a). *Checkpoint:* The voltage guard. The estimator credits pure methane.

(c) Press the reload button, run to a time between 470 and 490 s, type 85 % into Fuel quality estimate (methane) and run on to 900 s. Read the tile UPS lowest · now and compare it with the run as shipped. *Checkpoint:* 52.6 %, within 0.3 points, against 46.8 %.

3.9 App challenge · 9 min

Load scenario 9, the challenge, which runs 1200 s with two workload bursts, a slowing processor and a drop of the heat demand. Run each part to 1200 s at a speed of 100 with the default seed, pressing the reload button before each part, and read the tiles Fuel utilization peak, UPS lowest · now and Challenge target.

(a) Run as shipped, with the reference governor off. *Checkpoint:* 1.048, 53.1 and 60.3 %, not met.

(b) Switch Reference governor On and set Governor utilization limit to 0.74, leaving Fuel utilization setpoint at 0.75. *Checkpoint:* 0.844, 46.6 and 46.6 %, not met.

(c) Switch the governor On with the limit at 0.74 and the setpoint at 0.70. Explain why part (b) fails while this run passes. *Checkpoint:* 0.833, 59.4 and 63.2 %, met.

3.10 Application area · 3 min

The protection trips the stack at a true anode fuel margin of 0.10, and the challenge verdict asks for a margin above 0.25.

(a) Why must a plant keep a solid oxide anode from running short of hydrogen, apart from the power it loses? *Checkpoint:* The nickel anode oxidizes and expands.

(b) Which reading does the simulated trip use, and which readings must a real control system work from instead? *Checkpoint:* The true margin; the voltage, the current and the fuel flow.

3.11 Application area · 4 min

At the design point the plant takes 7.22 MW of methane on its lower heating value, delivers 3.435 MW net and returns 1.10 MW of heat as hot water.

(a) Compute the electric and the combined heat and power efficiencies. *Checkpoint:* 47.6 % and 62.8 %.

(b) The model's stack has 3 MJ/K per half, while a real stack of this area would hold roughly 15 to 50 MJ/K. Which lessons of the cases would come later on a real stack, and which would not change? *Checkpoint:* Case 6's override and case 2's heating-rate hold come later; the challenge verdicts do not change.

3.12 Review the agent · 5 min

In case 4, with the suggested recharge (0.4 MW, 6 MW per unit), the current waits for fuel for about half a minute after the burst ends at 720 s. An AI assistant explains: "The recharge loop causes this. As soon as the burst ends it asks the fuel cell for up to 0.4 MW to refill the UPS, and that new load outruns the fuel. Set Recharge power limit to zero at 720 s and the wait disappears." Find the flaw and say how you would check it in the app.

CHAPTER 4

Torque-to-Pitch Handover and Gain Scheduling in a Wind Turbine Simulator

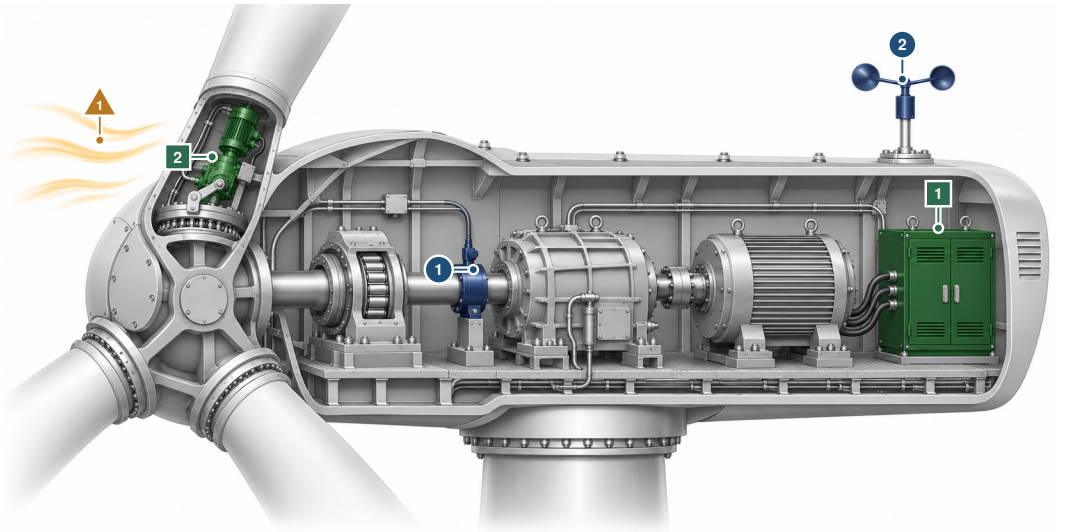

The controller moves 1 converter that sets the generator torque 2 drive that turns a blade, one of three

The controller reads 1 rotor speed 2 wind speed

What disturbs it 1 the wind, with its ramps, gusts and turbulence

Notice that the converter in the nacelle and the pitch drive in the hub act on the same rotor speed: torque holds it in light wind and pitch in strong wind. The chapter hands its controller the true wind at the rotor, where a real anemometer on the roof sits in the rotor's wake. Illustration, not to scale.

IN PLAIN TERMS One Rotor Speed, Two Ways to Hold It

The main idea. A large wind turbine must keep its blades turning at a safe speed while the wind keeps changing. In light wind the pull of the generator sets the speed. In strong wind the blades twist to spill the extra wind.

The challenge. The switch between the two must be smooth, or the gearbox takes a jolt each time the wind crosses over. The blades also grow about five times more sensitive as the wind rises, so a fixed setting moves them far more than needed in a gale.

Why it matters. Blade movement wears the bearings, and they must last twenty years.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	11 min	The first page, the fact box of the Introduction, the history box of Section 4.2, the two try-it boxes and the fact box of Section 4.6, and the key points
🕮	**Study**	65 min	The whole chapter with its pause points and the twelve exercises
⛰	**Practice**	20 min	The first page, Section 4.6, Section 4.8, Section 4.9 and its open challenges

The control question A large wind turbine regulates one rotor speed with two actuators, generator torque below rated wind and collective pitch above it, and the pitch sensitivity of its rotor grows by a factor of 4.9 across the range above rated. How should the two actuators hand over, and what do a gain schedule and a predictive controller told the wind ahead buy and cost?

Learning objectives

- Derive the Region 2 speed setpoint and the torque-law constant from the power coefficient at its peak.
- Explain how a setpoint offset, the region index with hysteresis and dwell, and back-calculation hand the rotor from torque to pitch.
- Place the pitch loop by pole assignment and compute its loop gain with and without the gain schedule across Region 3.
- Judge a pitch tuning on blade travel and peak pitch rate as well as on speed error.
- Separate what a wind preview gives a predictive controller from what a tighter speed objective gives, and state its energy cost in turbulence.

Prerequisites Sections 1.6.2 (velocity form and back-calculation), 1.6.4 (scenarios, scorecard and data) and 1.4.2 (inertia as stored energy). Sections 1.6.3 and 1.4.6 are needed only for the predictive controller and its preview.

The app https://apmonitor.com/apps/wind-turbine/

See also Volume 7, Chapter 9 and Volume 7, Chapter 3 for loops scheduled on an operating variable. Volume 4, Chapter 10 for several actuators on one variable, Volume 4, Chapter 2 for a protective loop that rests on its clamp, and Volume 4, Chapter 3 for limiter selection around a scheduled governor.

4.1 Introduction

Variable-speed pitch-regulated wind turbines present a control problem whose difficulty is the objective rather than the plant. There is one measurement worth regulating, rotor speed, and there are two actuators that change it. Generator torque removes energy from the shaft through the converter. Collective blade pitch prevents the wind from putting energy into the shaft in the first place. Both slow the rotor, both are available at every instant, and driving them from the same target produces two integral controllers that settle into any combination satisfying their common setpoint. The supercritical carbon dioxide cycle of Volume 4, Chapter 10 meets the same arrangement with three actuators for one variable.

What separates them is the objective, and the objective changes with the wind. Below the rated wind speed there is less power in the flow than the machine is built to take, so nothing has to be limited and the goal is to capture as much as possible. Capture is maximized at one particular tip-speed ratio, so the rotor speed setpoint tracks the wind rather than staying constant, and generator torque is the actuator that makes the rotor follow it. Above rated wind there is more power available than the drivetrain and the generator can accept, and the goal reverses. The rotor speed setpoint is fixed at the speed rating, generator torque is held at the torque rating because that is what a rating means, and the only remaining way to keep the rotor inside its limits is to feather the blades.

The handover between those two arrangements carries three requirements. It has to be bumpless, because a step in generator torque is a fatigue load on the gearbox. It has to be free of chatter, because in turbulent wind the boundary is crossed many times a minute. The loop that takes over above rated also works against a plant whose gain is not constant. The derivative of aerodynamic torque with respect to pitch grows from -8.9 MN m/rad at fine pitch to -46.6 MN m/rad at 25 m/s in the model used here, so a pitch controller tuned where it first takes over runs at roughly five times its design loop gain at the top of the envelope. Gain scheduling exists to remove that variation, and what it costs when it is switched off is measured in Section 4.6.2.

The physics behind the model is a pair of conservation statements. Aerodynamic power is the kinetic energy flux through the swept area multiplied by a dimensionless power coefficient, and the rotor speed follows an angular momentum balance on a single lumped inertia, so the plant is a first-

order lag whose aerodynamic damping is itself a function of the operating point. The distinguishing dynamic property is a process gain that varies by almost a factor of five over one operating region while its sign never changes, which is the textbook case for gain scheduling [1, 2]. The experiments of this chapter cover split-range and sequenced actuators on one measurement, setpoint tracking against a computed setpoint, feedforward from a steady-state inverse, anti-windup on a loop that rests against a clamp for most of its life, actuator rate limits as a bandwidth ceiling no tuning can raise, hysteresis and dwell on discrete logic, and gain scheduling against a measured nonlinearity [3]. The behavior matters industrially because pitch travel is blade-bearing duty on a component that must last twenty years, and because a turbine that loses speed regulation above rated is a protection event rather than a production loss.

> **DID YOU KNOW? A Rotor Wider Than Three Airliners**
>
> By 2019 blades approached 80 m and tip heights often passed 200 m, so three of the largest passenger aircraft, each 80 m across the wings, could fit inside the swept area of one rotor. Wind then supplied about 5 % of the world's electricity [4]. The chapter's reference rotor has a radius of 63 m.

This chapter works with a wind turbine simulator whose model, region logic and measured numbers can be reproduced from the text alone. It separates three mechanisms that are usually taught together: the sequencing of two actuators onto one measurement, the region logic that selects torque operation and labels which actuator is in charge, and the gain schedule that keeps the second loop where it was designed. Their roles differ. A 0.3 % setpoint offset lets the continuously active pitch loop begin regulating before a region declaration, the region index selects rated torque once Region 3 is declared, and switching the gain schedule off buys speed regulation at more than twice the blade travel. Section 4.2 describes how turbine control arrived at this two-region form. Section 4.3 gives the five-state model, with a fitted power coefficient surface whose departures from the textbook correlation are stated and justified, two rate-limited actuators and the wind. Section 4.4 measures the gain nonlinearity of the pitch loop and the schedule that removes it. Section 4.5 describes the two loops, the region logic with its hysteresis and dwell timers, and two settings of a model predictive controller: the controller of Section 4.5.6, which separates what a wind preview contributes from what a tighter speed objective contributes and costs, and the scenario tree of Section 4.5.7, which plans on three winds inside the preview's declared error and measures what hedging against that error is worth. Section 4.6 reports the thirteen scenarios, Section 4.7 discusses the wind measurement the Region 2 setpoint relies on, Section 4.8 describes the sensors, actuators and controllers of real turbines, and Sections 4.9 and 4.10 state the limitations and the conclusions.

4.2 Background

The first generation of commercial machines avoided the problem described above with a fixed rotor speed and blades that stall, as the history box in this section tells, and paid for it with a tip-speed ratio that is right at one wind speed only and a loose power cap.

Variable speed and blade pitch were introduced to fix each of those in turn. Allowing the rotor to change speed lets the tip-speed ratio be held at its optimum through the whole of the low-wind range, which is worth several percent of annual energy on a typical site. Pitching the blades replaces stall with a commanded, repeatable limit on captured power. The two together are the configuration of essentially every megawatt-class machine built since the 1990s, and they are also what creates the two-actuator problem, because the converter and the pitch system are both authoritative over rotor speed. The objectives, the actuator assignments and the region boundaries were laid out early in a form that is still recognizable [5], and the pitch-regulated variable-speed configuration and its power limiting were described in detail shortly after [6].

> ⧖ HISTORY **From Stalling Blades to Pitch and Variable Speed (1970s to 2023)**
>
> The oil crisis of the 1970s rekindled interest in renewable energy and led to the commercial adoption of grid-connected wind turbines [4]. Most early turbines ran at constant speed. On a fixed-frequency grid an induction generator turns only a few percent above synchronous speed while the wind varies across a wide range [7]. In a fixed-pitch, stall-controlled machine the blades are shaped so that stall limits the power in high wind, no pitch mechanism is needed, and control means little more than starting and stopping. A fixed speed puts the power coefficient at its peak at a single wind speed only [8]. By 1999 more and more variable-speed turbines were going into wind farms [7]. By 2023 the industry had nearly 40 years of field experience, 20 of them with multi-megawatt turbines that had converged on three blades, pitch control and variable speed [9].

Two lines of work then developed in parallel. Below rated, the question was how to hold the optimum without measuring the wind. The law that became standard is a static one in which generator torque is made proportional to the square of the measured rotor speed, with the constant chosen so that the rotor is in equilibrium exactly at the optimal tip-speed ratio, and the practical questions became how much the constant is in error and how slowly the rotor converges under it [10]. Above rated, the question was how to place a pitch loop whose plant gain moves with operating point. The standard answer, from the design of commercial controllers, is a proportional-integral loop on rotor speed with both gains divided by a function of the pitch angle [11], later extended to use the same actuator for tower and drivetrain load reduction [12]. Surveys of the field [13] and the open reference designs published for research use [8, 14] both present the machine as a set of regions with different actuators, and the transition between them as a named piece of logic rather than an afterthought.

The wider setting is that wind now supplies a large enough share of generation that turbine controllers are grid assets as well as machine protection, and the open problems in the field are increasingly about loads, atmospheric inflow and system integration rather than about the loops themselves [4]. The loops themselves therefore remain the entry point to the subject, because a pitch controller that trades speed regulation against blade-bearing duty is the smallest complete example of the trade every later turbine question turns on.

4.3 Turbine Model

The application simulates a three-bladed upwind horizontal-axis turbine with the geometry and ratings of the NREL 5 MW offshore reference turbine [14], which is the standard open baseline for control studies. Figure 4.1 shows the physical arrangement and names every symbol that appears in the equations. Everything in the model lives on the low-speed shaft, so the 97:1 gearbox appears only in the conversion of the published rated generator torque of 43 093.55 N m at the high-speed shaft into the 4.180 074 MN m used here.

4.3.1 Variable Classification

The single controlled variable and the two manipulated variables are listed in Table 4.1, together with the disturbances, the five states and the principal derived outputs. Only one loop is scored, because there is only one process value being regulated.

4.3.2 Rotor Aerodynamics

Aerodynamic power is the product of the kinetic energy flux through the swept area and a dimensionless power coefficient that depends on the tip-speed ratio and the collective pitch angle, as shown in Equation (4.1).

$$P_{aero} = \tfrac{1}{2}\rho A V^3 C_P(\lambda, \beta), \qquad \lambda = \frac{\omega R}{\max(V, 0.3)} \tag{4.1}$$

where ρ is air density, $A = \pi R^2$ is the swept area, V is the wind speed reaching the rotor and λ is the ratio of blade tip speed to wind speed. Other chapters of this volume use λ for other quantities, as

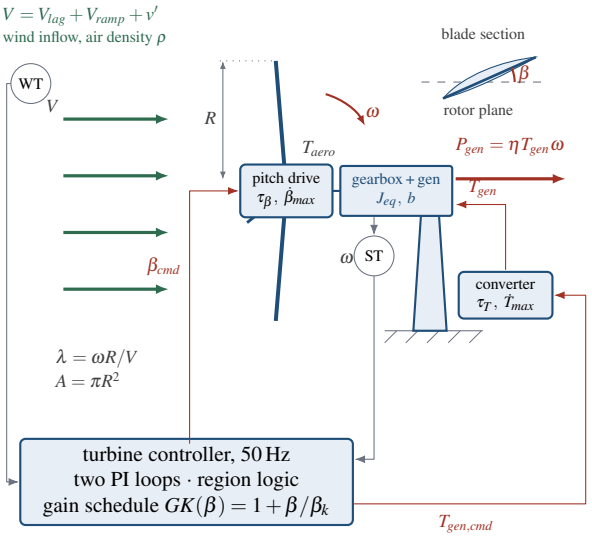

Figure 4.1 Rotor, drivetrain, actuators and instruments of the simulated machine. Rotor speed ω is measured exactly, and the same controller is given the true instantaneous wind speed V that reaches the rotor.

Table 4.1 Variables, Their Roles and Their Ranges

Symbol	Role	Range	Unit
ω	CV, state	0–12.1 nominal	rad/s (rpm)
ω_{SP}	setpoint, computed	0.45–1.2633	rad/s
T_{gen}	MV 1, state	0–4.180074	MN·m
β	MV 2, state	0–90 (0–30 in AUTO)	deg
λ^*	operator setting	4–10	—
V_{lag}	state, wind lag	3–28	m/s
V_{ramp}	state, wind ramp integral	unbounded	m/s
$\bar{V}$	DV, mean wind setting	3–28	m/s
$\dot{V}$	DV, wind ramp rate	−0.4–+0.4	m/s per s
I_{turb}	DV, turbulence intensity	0–20	%
V	inflow seen by rotor	0.2–45	m/s
λ	output	$\omega R / \max(V, 0.3)$	—
C_P	output	0–0.482	—
T_{aero}	output	—	N·m
P_{gen}	output	0–5.0	MW
region	output, annunciation	2, 2.5, 3	—

Section 1.3 lists. The floor of 0.3 m/s on the denominator prevents a division by zero in a calm and has no effect anywhere the machine operates. The upper bound on C_P is the Betz limit of $16/27 \approx 0.593$, which follows from a momentum balance on the stream tube and holds for any rotor whatever [15]. The surface used here peaks at 0.482, which is 81 % of that ceiling.

Aerodynamic torque follows from power and speed, as shown in Equation (4.2).

$$T_{aero} = \frac{P_{aero}}{\max(\omega,\, \omega_\varepsilon)}, \qquad \omega_\varepsilon = 0.05\,\mathrm{rad/s} \tag{4.2}$$

The speed floor ω_ε keeps the quotient finite at standstill. Dimensionally the expression is W divided by rad/s, which is N m as required.

The power coefficient is a fitted analytic surface of the family used for power system studies [16], given in Equation (4.3), with the pitch angle in degrees inside this expression only.

$$\lambda_t = s\lambda, \qquad \frac{1}{\lambda_i} = \frac{1}{\lambda_t} - c_8$$
$$C_P = k\left[c_1\left(\frac{c_2}{\lambda_i} - c_3(\beta + q\beta^2) - c_4\right)\mathrm{e}^{-c_5/\lambda_i} + c_6\lambda_t\right] \tag{4.3}$$

where $c_1 = 0.5176$, $c_2 = 116$, $c_3 = 0.40$, $c_4 = 6.5$, $c_5 = 12.25$, $c_6 = 0.0068$, $c_8 = 0.005$, $q = 0.08$, $s = 0.944705354917025$ and $k = 0.504333798658105$. Three departures from the textbook coefficient set are deliberate and each fixes a defect that would otherwise corrupt the lesson. The reciprocal $1/\lambda_i$ normally carries an additional $-0.035/(\beta^3+1)$ term and a $+0.08\beta$ term, which together make C_P non-monotonic in β at low tip-speed ratio, so that feathering a degree or two raises the modeled power and inverts the sign of the pitch loop. Both are dropped. The pitch loss is written $-c_3(\beta+q\beta^2)$ rather than $-c_3\beta$, because a linear term makes $|\partial C_P/\partial\beta|$ independent of pitch and would remove the nonlinearity the application exists to teach. The two scalings s and k place the optimum exactly at $\lambda^* = 7.55$ with $C_{P,max} = 0.482$, because the unscaled expression peaks at $\lambda = 7.13252543$ with a value of 0.95571624. The remaining coefficients were fitted so that the resulting steady pitch schedule reproduces the published fine-pitch schedule for this machine to 0.6° rms between 12 and 25 m/s and so that C_P falls to zero near $\lambda = 18.5$.

One term is added below the fitted surface. A stopped rotor in a gale is not torque-free, but every fit of this family goes to zero at low λ, which would leave a stopped rotor permanently stopped. A starting-torque floor $C_{P,start} = c_{q0}\lambda\,\mathrm{e}^{-(\lambda/\lambda_s)^2}$ with $c_{q0} = 0.015$ and $\lambda_s = 1.5$ is applied as a lower bound on C_P. At $\lambda = 3$ the floor is 0.0005 against a fitted 0.13, and at the optimum it is below 10^{-9}.

The resulting surface is shown in Figure 4.2. The two features that matter for control are both visible. There is a single maximum in λ at fixed β, which is what makes an optimal tip-speed ratio exist and what makes the Region 2 equilibrium stable. And the curves collapse toward zero as β grows, faster at each successive five degrees, which is the quadratic pitch term and the reason the pitch loop gain has to be scheduled.

4.3.3 Drivetrain

Blades, hub, shaft, gearbox and generator rotor are lumped into a single inertia referred to the low-speed shaft, giving the one rotor state shown in Equation (4.4).

$$J_{eq}\,\dot{\omega} = T_{aero}(\omega, V, \beta) - T_{gen} - b\,\omega \tag{4.4}$$

where $J_{eq} = 4.0 \times 10^7\,\mathrm{kg\,m^2}$ and $b = 1.0 \times 10^4\,\mathrm{N\,m\,s}$ is a viscous drag coefficient standing in for bearing and gear losses. Electrical power is the shaft power taken by the generator times a constant efficiency, $P_{gen} = \eta\,T_{gen}\,\omega$ with $\eta = 0.944$.

The open-loop rotor is stable, and its time constant is a strong function of where the machine is operating. Linearizing Equation (4.4) about a steady point gives $\tau_{OL} = J_{eq}/(|\partial T_{aero}/\partial\omega| + b)$, and $\partial T_{aero}/\partial\omega$ at the optimum is negative because past λ^* the power coefficient falls faster than the speed rises. Evaluated on the fitted surface, the aerodynamic damping is -1.54 MN m s at 6 m/s and -2.65 MN m s at rated speed and rated wind, so the open-loop time constant falls from 25.8 s to 15.0 s across Region 2. The plant is slow on the scale of a control loop, which is why the application runs with a 60 s default trend window while the magnetic suspension of Volume 7, Chapter 6 needs milliseconds. The energy this inertia stores at rated speed, in seconds of rated power, is the inertia constant of Section 1.4.2.

4.3.4 Actuators

Both actuators are first-order lags with a hard rate limit and end stops, as shown in Equations (4.5) and (4.6).

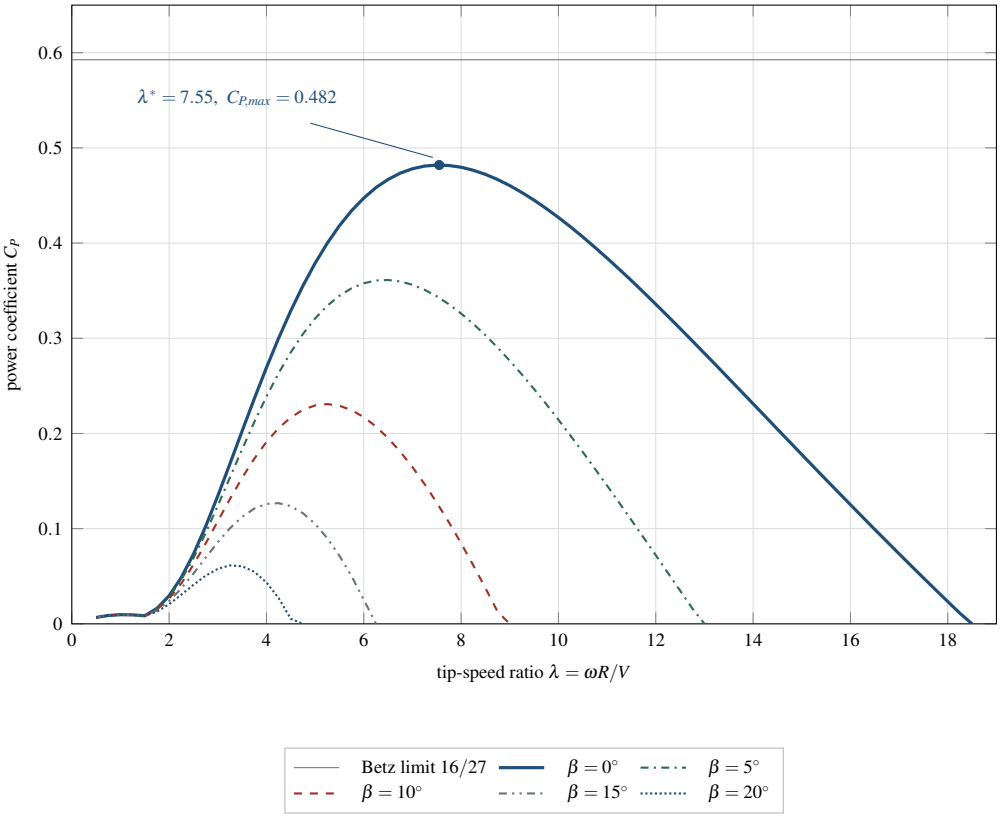

Figure 4.2 Power coefficient surface as implemented, at five collective pitch angles: 0° very thick solid, 5° dash-dotted, 10° dashed, 15° dash-dot-dotted and 20° densely dotted, under the thin horizontal Betz limit. The curves are evaluated from Equation (4.3) with the fitted coefficients, and the marker is the operating point the Region 2 controller holds.

$$\dot{\beta} = \mathrm{clamp}\left(\frac{\beta_{cmd} - \beta}{\tau_\beta},\ -\dot{\beta}_{max},\ +\dot{\beta}_{max}\right) \tag{4.5}$$

$$\dot{T}_{gen} = \mathrm{clamp}\left(\frac{T_{gen,cmd} - T_{gen}}{\tau_T},\ -\dot{T}_{max},\ +\dot{T}_{max}\right) \tag{4.6}$$

with $\tau_\beta = 0.25\,\mathrm{s}$, $\dot{\beta}_{max} = 8\,^\circ/\mathrm{s}$, $\tau_T = 0.05\,\mathrm{s}$ and $\dot{T}_{max} = 1.5\,\mathrm{MN\,m/s}$. The derivative is additionally set to zero when the state is against an end stop and the demand pushes further into it, so the right-hand side stays a pure function of the state and the fourth-order Runge-Kutta integrator sees no discontinuous jump inside a step. Pitch is bounded to $[0, 90^\circ]$ physically and generator torque to $[0, T_{rated}]$, and both bounds are re-imposed after each step.

The asymmetry between the two actuators is what makes several of the scenarios behave the way they do. Generator torque reaches a new value in about a tenth of a second, while pitch needs nearly three seconds to travel the 22° that 25 m/s requires, and pitch is the actuator that has to reject the fastest disturbances.

4.3.5 Wind

The wind reaching the rotor is a single scalar with three parts and two states, as shown in Equation (4.7).

Table 4.2 Machine Parameters and Where the Application Shows Them

Symbol	Meaning	Value	Unit
Rotor			
R	rotor radius	63	m
A	swept area, πR^2	12 469	m^2
ρ	air density	1.225	kg/m^3
J_{eq}	drivetrain inertia (LSS)	4.0×10^7	$kg{\cdot}m^2$
b	drivetrain damping[a]	1.0×10^4	N·m·s
λ^*	optimal tip-speed ratio[h]	7.55	—
$C_{P,max}$	peak power coefficient[h]	0.482	—
Ratings			
ω_{rated}	rated rotor speed	1.267109 (12.10)	rad/s (rpm)
ω_{min}	minimum rotor speed[a]	0.45 (4.30)	rad/s (rpm)
$0.997\,\omega_{rated}$	torque-loop setpoint cap	1.263308 (12.06)	rad/s (rpm)
T_{rated}	rated generator torque	4.180074	MN·m
η	generator efficiency[a]	0.944	—
V_{rated}	rated wind speed[h]	11.4	m/s
K	torque-law constant	2.1388×10^6	$N{\cdot}m{\cdot}s^2$
Actuators			
τ_β	pitch actuator lag	0.25	s
$\dot{\beta}_{max}$	pitch rate limit	8	deg/s
β_{max}	feather limit[a]	90	deg
β_{ctl}	pitch limit in AUTO[a]	30	deg
τ_T	torque actuator lag[a]	0.05	s
$\dot{T}_{max}$	torque rate limit[a]	1.5	MN·m/s
Limits and wind			
ω_{trip}	overspeed alarm[a]	1.20	× rated
τ_V	mean-wind setting lag[h]	2.0	s
ω_ε	torque speed floor	0.05	rad/s

[a] adjustable only in advanced mode. [h] not exposed for editing.

$$\dot{V}_{lag} = \frac{\bar{V} - V_{lag}}{\tau_V}, \quad \dot{V}_{ramp} = \dot{V}, \quad V = V_{lag} + V_{ramp} + v' \tag{4.7}$$

where $\bar{V}$ is the operator-set mean wind, $\tau_V = 2.0\,\text{s}$ so that a change in that setting is not a physical step in the wind, $\dot{V}$ is the operator-set ramp rate integrated independently, and v' is the turbulent component. The ramp integrator is frozen when the sum $V_{lag} + V_{ramp}$ leaves $[0.5, 40]$ m/s. Turbulence is a first-order filtered Gaussian process of time constant 6 s, scaled so that its stationary standard deviation equals the turbulence intensity times $\max(V_{lag} + V_{ramp}, 1)$, and it is generated once per controller step rather than inside the derivative function, because the derivative function is called four times per integration step and must be free of side effects. The random stream is a seeded mulberry32 generator, so a scenario with turbulence replays identically.

4.3.6 Parameters, Initialization and Numerics

Table 4.2 lists every parameter, its value and whether the application exposes it. Hidden parameters are used by the model but are not exposed for editing, which is what makes the low-wind scenarios a task for the reader rather than a demonstration, because a reader who does not already know λ^* has to find it from the reported C_P.

Every scenario begins on its own steady state. The initial speed and the initial generator torque are

$$\omega_0 = \text{clamp}\big(\lambda^*_{model} V_0 / R,\ \omega_{min},\ \omega_{rated}\big),$$
$$T_{gen,0} = \text{clamp}\big(T_{aero}(\omega_0, V_0, \beta_0) - b\,\omega_0,\ 0,\ T_{rated}\big),$$

and the initial pitch β_0 is the angle that makes the aerodynamic torque equal $T_{rated} + b\,\omega_0$, found by 60 bisection steps on the monotone C_P surface. The initialization uses the model's optimal ratio, not a subsequently edited controller setpoint. Before control commands are applied it gives $\dot{\omega} = 0$ at $t = 0$ in every region, and a verification test in the application confirms it to 10^{-12} rad/s^2 at 8 m/s.

Integration uses fourth-order Runge-Kutta at a fixed step of $dt = 0.01$ s. The controller executes every 0.02 s, which is the 50 Hz cadence of a real turbine controller, and the historian logs at 10 Hz. The three clocks, the order of work within a step and the independence of the results from the playback speed are those of Section 1.6.1. Rotor speed is clipped to non-negative values after each step and to eight times rated as a numerical guard, and the overspeed annunciator latches at $1.20\,\omega_{rated}$ and clears below $1.15\,\omega_{rated}$. The alarm is annunciation only, because the manual scenario exists to show where an unprotected rotor comes to rest.

4.3.7 Assumptions

The model is reduced deliberately, and each reduction removes something that matters somewhere.

One power coefficient surface stands for the entire rotor. There is no blade element momentum theory, no dynamic inflow, no stall hysteresis and no unsteady aerodynamics, so anything that depends on how fast the flow field itself rearranges is outside what the model can say. The first second after a large pitch step and deep stall at very low λ are both in that category.

The surface is fitted, not derived. It reproduces the reference machine over λ from 3 to 12 and β from 0 to about 25°, which covers Regions 2 and 3, and is an extrapolation outside that box. At rated wind and rated speed with $\beta = 0$ the surface gives 5.410 MW of aerodynamic power against 5.297 MW of rated mechanical power, so the modeled machine reaches its rating 2.1 % early. That residual is the price of one smooth analytic expression fitted to a real rotor.

One rotor-speed state. The drivetrain torsional mode near 1.7 Hz on the real machine is absent, which matters because it is exactly the mode an over-bandwidth pitch loop excites in practice [17].

No tower. Fast pitch action drives the tower fore-aft mode on a real turbine, and that coupling is the main reason commercial pitch loops are scheduled conservatively [11]. It is not represented, so the cost of an unscheduled loop is measured here as blade travel. A margin against tower motion cannot be quantified without the omitted dynamics.

The wind is one scalar at the hub. There is no shear, no tower shadow, no rotational sampling and no yaw error, and the turbulence is a band-limited first-order process rather than a Kaimal spectrum.

The electrical side is a constant. Generator power is shaft power times 0.944, so losses do not vary with load and there are no converter or grid dynamics of any kind.

Air density and inertia are constant, and the pitch limit of 30° in AUTO is a controller setting rather than a physical stop. It exists to keep the integrator out of the region where the fitted C_P has already reached zero and the pitch loop would have no gain at all. In MANUAL the blades still reach 90°.

4.4 Pitch Sensitivity and the Gain Schedule

The reason this machine needs a gain schedule is visible in one derivative. Table 4.3 lists the steady pitch angle that holds rated speed at rated torque, and the pitch sensitivity of aerodynamic torque and of aerodynamic power at that point, computed from Equations (4.1) to (4.3) at $\omega = \omega_{rated}$.

The sensitivity grows by a factor of 4.9 between the bottom and the top of Region 3. A proportional-integral loop whose gain $K_{c,\beta}$ is fixed therefore has a loop gain product $K_{c,\beta}\,|\partial T_{aero}/\partial\beta|$ that grows with it, from 33.6 where the tuning was placed, at the fine-pitch sensitivity of 8.92 MN m/rad, to 175.6 at 25 m/s with the default tuning, a factor of 5.2. The standard correction [11] divides both the proportional and the integral gain by a linear function of the pitch angle itself, as shown in Equation (4.8).

$$GK(\beta) = 1 + \frac{\beta}{\beta_k}, \qquad K_{c,\beta}^{\mathrm{eff}} = \frac{K_{c,\beta}}{GK(\beta)}, \qquad \tau_{I,\beta}^{\mathrm{eff}} = \tau_{I,\beta} \tag{4.8}$$

Table 4.3 Steady Pitch and Pitch Sensitivity Across Region 3

V (m/s)	β (deg)	$\partial T_{aero}/\partial\beta$ (MN·m/rad)	$\partial P_{aero}/\partial\beta$ (MW/rad)
11.4	0.48	−9.59	−12.16
15.0	10.95	−31.13	−39.45
18.0	15.27	−41.36	−52.41
21.0	18.51	−46.42	−58.81
25.0	21.97	−46.61	−59.05

where $\beta_k = 4.98°$ was fitted to this model's own sensitivity curve, close to the 6.30° published for the real 5 MW machine. Dividing K_c and K_c/τ_I by the same factor leaves the effective integral time untouched, so the schedule moves the bandwidth of the loop and not the shape of its response. Both effective values are logged as separate columns, and a 120 s run of scenario 7 ends with $GK = 5.266$, $K_{c,\beta}^{\text{eff}} = 0.7155\,\text{s}$ and $\tau_{I,\beta}^{\text{eff}} = 2.3333\,\text{s}$ against a set value of 2.3333 s, which confirms the claim from the data rather than from the code. The scheduled loop gain product at 25 m/s is 32.4 against 33.6 at fine pitch, so the schedule holds the loop within 4 % of where it was placed.

Scheduling a gain against a measured operating point is a well-studied device with well-known failure modes, chiefly that the schedule variable must move slowly compared with the loop it is scheduling [18]. Pitch angle satisfies that here because the blades move at most 8 °/s against a loop placed at $\omega_n = 0.6\,\text{rad/s}$. The riderless bicycle of Volume 7, Chapter 9 schedules its loop on forward speed and the turbofan governor of Volume 4, Chapter 3 on corrected core speed, two other schedule variables chosen because they move more slowly than the loop.

4.5 Controller Design

4.5.1 Velocity Form PID Algorithm

Both loops use the shared velocity-form proportional-integral controller of Section 1.6.2 on an exact 0.02 s cycle [19]. For the scheduled pitch loop the gain moves during a run, so the increment must difference the fully formed proportional terms, as in Equation (4.9).

$$\begin{aligned} e_k &= a(\text{SP}_k - \text{PV}_k), \qquad a = -1 \\ P_k &= K_{c,k} e_k, \qquad \Delta P_k = P_k - P_{k-1} \\ \Delta I_k &= (K_{c,k}/\tau_I) e_k \Delta t \end{aligned} \tag{4.9}$$

For constant gain this reduces to $K_c(e_k - e_{k-1})$. With scheduling, the term $K_{c,k}e_k - K_{c,k-1}e_{k-1}$ also carries the gain change. The implementation applies the command before advancing its integral state, as in Equation (4.10).

$$\begin{aligned} s_k &= s_{k-1} + \Delta P_k \\ c_k &= u_{bias,k} + s_k + I_{k-1} \\ u_k &= \text{clamp}(c_k, u_{min}, u_{max}) \\ I_k &= I_{k-1} + \Delta I_k + (u_k - c_k)\Delta t / T_t \end{aligned} \tag{4.10}$$

These are Equations (1.11) to (1.13) of Chapter 1 with $s_k = S_k$ and $c_k = v_k$, and $T_t = \tau_I$ by default. The running sum is never clipped, and back-calculation acts on the integral alone, as Section 1.6.2 explains. The pitch and torque tracking times are 2.3333 and 12 s, respectively, at default tuning. Actuator rate limits act in the plant dynamics, separately from the controller output clamp. Both derivative gains are zero. In manual, and for torque while Region 3 fixes its command, the tracking path sets $I = u_{applied} - u_{bias} - s$ using the current error. This aligns the controller output with the applied command for transfer back to automatic regulation, and later error and integral evolution may move it.

The direction flag is the one place a sign error would propagate everywhere. More generator torque slows the rotor and more pitch slows the rotor, so the process gain of both loops is negative. Both loops are therefore direct acting, their output rising as the measured speed rises, and both carry $a = -1$ in the convention of Section 1.3. Setting that flag once on the shared controller is the entire sign handling in the application.

4.5.2 Region 2 Torque Control

Below rated the rotor speed setpoint is the speed that puts the machine at the peak of the C_P curve, capped just below the speed rating, as shown in Equation (4.11).

$$\omega_{SP} = \text{clamp}\left(\frac{\lambda^* V}{R},\ \omega_{min},\ 0.997\,\omega_{rated}\right) \tag{4.11}$$

Generator torque is then a feedforward plus proportional-integral trim, as shown in Equation (4.12).

$$T_{gen,cmd} = K\omega_{SP}^2 + K_{c,T}(\omega - \omega_{SP}) + \frac{K_{c,T}}{\tau_{I,T}}\int(\omega - \omega_{SP})\,dt \tag{4.12}$$

where the bias term is the exact steady-state aerodynamic torque at the optimum. Substituting $\lambda = \lambda^*$ into Equations (4.1) and (4.2) gives $T_{aero} = K\omega^2$ with $K = \rho\pi R^5 C_{P,max}/(2\lambda^{*3}) = 2.1388 \times 10^6\,\text{N m s}^2$, and a verification test in the application confirms the identity to a relative 10^{-9} at four rotor speeds.

The same identity is the basis of the alternative Region 2 law offered as an advanced option, $T_{gen} = K\omega^2$ applied directly [10]. That law carries no speed setpoint and no wind measurement, which is the reason it is the baseline in the reference designs, and its two known weaknesses are that convergence to the optimum is limited by the rotor time constant and that the result is only as good as K. Both laws take K, and Equation (4.11) takes R, from the declared values of Table 4.2 once when the application loads, so an air density or rotor radius edited in the application detunes them as it would a commissioned controller, and with a 70 m rotor at 6 m/s the $K\omega^2$ law settles at $\lambda = 8.87$ and the proportional-integral law at $\lambda = 8.39$. The law does use the measured speed ω. What it does without is the wind and a speed setpoint.

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

4.5.3 Region 3 and the Handover

Above rated, generator torque is held at T_{rated} and the pitch loop regulates speed against the rating, as shown in Equation (4.13).

$$\begin{aligned} \beta_{cmd,k} &= \text{clamp}\left(K_{c,\beta}^{\text{eff}} e_{\omega,k} + I_{k-1},\ 0,\ \beta_{ctl,max}\right) \\ e_{\omega,k} &= \omega_k - \omega_{rated} \end{aligned} \tag{4.13}$$

Pitch control remains active in every region with a setpoint at ω_{rated}, while the torque setpoint is capped at $0.997\omega_{rated}$. This offset prevents both loops from seeking the same zero-error condition. Below rated, pitch rests at its lower limit and its integral follows the tracking law of Equation (4.10), as the protective override of Volume 4, Chapter 2 rests on its lower clamp through healthy operation. The tracking state need not lie exactly at the boundary on every scan. As speed rises, pitch can begin feathering before Region 3 is declared. Once that declaration occurs, an explicit selector commands rated generator torque and tracks the torque controller to that command. The onset of pitch regulation and the torque selector are distinct parts of the handover.

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

4.5.4 Region Index, Hysteresis and Dwell

The region index selects rated torque in Region 3. Outside Region 3 the selected PI or $K\omega^2$ law controls torque. The index also drives annunciation. Pitch remains active throughout, and its gain schedule depends on blade angle and the schedule setting, rather than on the index. Region logic carries hysteresis on both edges and dwell timers on all but one transition. Region 3 is entered when $\beta_{cmd} \geq 0.5^\circ$, or when $T_{gen,cmd} \geq 0.999\,T_{rated}$ and $\omega \geq 1.002\,\omega_{rated}$ together. It is left only when $\beta_{cmd} \leq 0.1^\circ$ and $\omega \leq 0.990\,\omega_{rated}$ together. Those two conditions cannot both hold, so the boundary is a dead band and not a threshold. The 2 against 2.5 label is decided by rotor speed rather than by the setpoint, with 2.5 declared above $0.96\,\omega_{rated}$ and 2 below $0.90\,\omega_{rated}$, because the setpoint follows the wind and in 14 % turbulence it crosses the cap several times a second while the rotor speed is filtered by 4×10^7 kg m^2 of inertia.

Dwell timers sit on top of the hysteresis. Leaving Region 3 requires 2.0 s since the last change, and any transition between 2 and 2.5 requires 6.0 s. Entry into Region 3 is never delayed, because that transition selects rated torque while pitch is already regulating speed. A verification test runs the wind ramp from 7 to 22 m/s and requires between one and four region changes over 130 s, and the measured value is two.

4.5.5 Tuning and Defaults

The pitch loop is placed by pole assignment on the integrating rotor. Treating $J_{eq}\dot{\omega} = (\partial T_{aero}/\partial\beta)\beta$ as the plant and requiring a closed-loop damping ratio $\zeta = 0.7$ at $\omega_n = 0.6$ rad/s gives $K_{c,\beta} = 2\zeta\omega_n J_{eq}/|\partial T_{aero}/\partial\beta| = 3.775$ s and $\tau_{I,\beta} = 2\zeta/\omega_n = 2.333$ s at the fine-pitch sensitivity of 8.9 MN m/rad. The shipped defaults are $K_{c,\beta} = 3.7673$ s, displayed as 22.6 °/rpm, and $\tau_{I,\beta} = 2.3333$ s. The torque loop is tuned for a closed-loop rotor time constant near 8 s, with $K_{c,T} = 3.0 \times 10^6$ N m s/rad displayed as 314 kNm/rpm and $\tau_{I,T} = 12$ s. The gain schedule is on by default and the Region 2 law defaults to the proportional-integral form.

4.5.6 Model Predictive Controller

A second controller structure can take the place of both loops and of the torque selector of the region logic. The model predictive controller (MPC) acts only in AUTO and only when selected, and the scheduled loops remain the default of the application and of every scenario except scenarios 11 and 12. Every $T_s = 0.1$ s, or five controller cycles, it solves one quadratic program for the pitch and torque commands over $N_p = 50$ steps, or 5 s, in 12 move blocks that grow from four single steps to one of 15 steps, with the last block taking the remaining 5. The first move is applied and held until the next plan, as Section 1.6.3 describes [20].

The model is written rather than identified. It is Equation (4.14), the application's own equations evaluated with the declared values of Table 4.2, which are copied when the application loads, so a parameter that a reader edits is model mismatch.

$$\begin{aligned} J_{eq}\dot{\omega} &= T_{aero}(\omega, V, \beta) - T_{gen} - b\,\omega + d \\ \tau_\beta\dot{\beta} &= \beta_{cmd} - \beta, \qquad \tau_T\dot{T}_{gen} = T_{gen,cmd} - T_{gen} \end{aligned} \tag{4.14}$$

Here d is an unmeasured torque disturbance, and the rate limits and end stops of Equations (4.5) and (4.6) become constraints. At every decision the model is linearized at the measured speed ω_0 and pitch β_0 and the first previewed wind V_0, with $a_\omega = \partial T_{aero}/\partial\omega$ and $a_\beta = \partial T_{aero}/\partial\beta$ from finite differences, as Equation (4.15) states.

$$\begin{aligned} \dot{x} &= \begin{bmatrix} \frac{a_\omega - b}{J_{eq}} & \frac{a_\beta}{J_{eq}} & -\frac{1}{J_{eq}} \\ 0 & -\frac{1}{\tau_\beta} & 0 \\ 0 & 0 & -\frac{1}{\tau_T} \end{bmatrix} x + \begin{bmatrix} 0 & 0 \\ \frac{1}{\tau_\beta} & 0 \\ 0 & \frac{1}{\tau_T} \end{bmatrix} u + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} f_k \\ f_k &= \frac{T_{aero}(\omega_0, V_k, \beta_0) - \hat{T}_{gen} - b\,\omega_0 + \hat{d}}{J_{eq}} \end{aligned} \tag{4.15}$$

The states $x = (\delta\omega,\ \delta\beta,\ \delta T_{gen})$ are deviations from the measured speed and pitch and the estimated torque $\hat{T}_{gen}$, the inputs $u = (\beta_{cmd} - \beta_0,\ T_{gen,cmd} - \hat{T}_{gen})$ are measured from where the actuators are, and the system is sampled at T_s with an exact zero-order hold. The previewed wind V_k of period k enters through the forcing f_k, evaluated before the solve, so the program stays linear. Electrical power is linearized as $\hat{P} = \eta(\hat{T}_{gen}\omega_0 + \hat{T}_{gen}\,\delta\omega + \omega_0\,\delta T_{gen})$. At rated speed the linearization gives $a_\beta = -9.60$, -31.14 and -46.61 MN m/rad at 11.4, 15 and 25 m/s, as in Table 4.3. Open loop from 8 to 22 m/s, its 5 s speed response to pitch and torque steps is within 6.89 % of the plant's, but it predicts a 14 to 20 m/s wind step 12.5 % low, 3.52 against 4.03 rpm, because it keeps the sensitivities of the moment.

The plan tracks the operating point that each previewed wind calls for and has no regions. The speed target is $\omega_{t,k} = \mathrm{clamp}(\lambda^* V_k/R,\ \omega_{min},\ \omega_{rated})$, the formula of Equation (4.11) capped at the rating itself, with fine pitch and the balancing torque as the input targets $\beta_{t,k}$ and $T_{t,k}$. Where that torque would exceed T_{rated}, the speed target rises toward rated speed at fine pitch, and at rated speed the pitch target is the angle that sheds the excess, found by bisection on the declared C_P surface with $\hat{d}$ in the balance. Each plan minimizes Equation (4.16), in which β_c and T_c are the pitch and torque commands and j counts the move blocks.

$$
\begin{aligned}
\mathscr{J} = &\sum_{k=1}^{N_p} w_\omega(\hat{\omega}_k - \omega_{t,k})^2 + \sum_{o=1}^{2}\left(\kappa_o\varepsilon_o^2 + \nu_o\varepsilon_o\right) \\
&+ \sum_{k=0}^{N_p-1}\left[w_\beta(\beta_{c,k} - \beta_{t,k})^2 + w_T(T_{c,k} - T_{t,k})^2\right] \\
&+ \mu\sum_{j=1}^{12}\left[\left(\frac{\Delta\beta_{c,j}}{m_\beta}\right)^2 + \left(\frac{\Delta T_{c,j}}{m_T}\right)^2\right] \\
\text{s.t. } &\hat{\omega}_k \le \omega_{max} + \varepsilon_1, \quad \hat{P}_k \le P_{max} + \varepsilon_2, \quad \varepsilon_o \ge 0 \\
&0 \le \beta_{c,j} \le \min(\beta_{ctl}, \beta_{max}), \quad 0 \le T_{c,j} \le T_{rated} \\
&|\Delta\beta_{c,j}| \le \Delta\beta_{lim}, \quad |\Delta T_{c,j}| \le \dot{T}_{max} T_s
\end{aligned}
\tag{4.16}
$$

Speed enters in rpm with $w_\omega = 4$ per rpm^2, the inverse square of the torque loop's 0.5 rpm band, and $w_\beta = 0.04$ per deg^2, $w_T = 0.33$ per (MN·m)^2, $m_\beta = 1°$, $m_T = 0.35$ MN m and $\mu = 3$. The soft ceilings are $\omega_{max} = 1.15\,\omega_{rated}$, the overspeed alarm's reset level read from its live setting, and $P_{max} = \eta T_{rated}\omega_{rated} = 5.0$ MW, with slack prices $\kappa_1 = 10^3$ per rpm^2, $\nu_1 = 10^2$ per rpm, $\kappa_2 = 10^2$ per MW^2 and $\nu_2 = 10$ per MW, and a ceiling already crossed is re-entered along a 1 s funnel. Neither decides a result, because scenario 11 peaks at 100.35 % of rated speed without the power ceiling and at 100.29 % without the speed ceiling or with both. The command boxes are hard and read from the live settings, as the loops' clamps read them. The torque bound is hard rather than soft because the rating is the loops' clamp and the plant's end stop. The cap $\Delta\beta_{lim} = 0.95\,\dot{\beta}_{max}\tau_\beta(1 - \mathrm{e}^{-T_s/\tau_\beta}) = 0.626°$ keeps a held step through the lag below 7.6 °/s, 95 % of the limit, and a long move made of such steps averages 6.26 °/s. The torque cap is the declared rate, and there is no terminal weight. A dual active-set method [21] solves each program, and any plan without a reported optimum, with a value that is not finite or with a command outside its box is discarded.

A steady-state Kalman filter on the speed and a random-walk disturbance d makes the plan offset-free [22], the arrangement of Section 1.6.3. With $G = (60/2\pi)\,10^6\,\Delta t/J_{eq} = 4.77 \times 10^{-3}$ rpm per MN·m per 0.02 s step, the transition matrix $\left[\begin{smallmatrix}1 & G\\ 0 & 1\end{smallmatrix}\right]$, a measurement variance $r = (0.01\,\mathrm{rpm})^2$ and a disturbance intensity $\gamma r/G^2$ with $\gamma = 1$, its gain is $L = (0.769,\ 100.6)$. The prediction is a Heun step of the declared rotor equation driven by the measured pitch, the hub wind and the torque that the commands produce through the declared lag, because the loops read no torque measurement. It runs under both structures, so a switch starts from a current estimate.

The wind over the horizon comes from an instrument with three settings. With the preview off the plan holds the hub wind measured now, the number the Region 2 setpoint already reads. The perfect

preview, a teaching bound, is the wind that will arrive, projected from the wind states, the settings and any scripted change within 5 s through the declared 2 s lag, with the turbulence drawn ahead from the plant's own generator and queued, so the plant's wind is bit-identical under both structures. The LIDAR setting, the default, adds a declared error $e_k = (\ell_k/5\,\mathrm{s})\,\sigma_L\,n$, with $\ell_k = \min\big((k+\frac{1}{2})T_s,\ 5\,\mathrm{s}\big)$, $\sigma_L = \sqrt{\sigma_u^2 + (0.01\,V)^2}$, σ_u the turbulence standard deviation of Section 4.3, and n one draw per plan of a unit-variance noise with a 2 s correlation time from the instrument's own seeded stream. The error is zero at the rotor, and at 5 s it combines the evolution of the air on its way in with a line-of-sight measurement that is not the rotor-effective wind, as in LIDAR-assisted predictive control [23]. The wave buoy of Chapter 5 plans on a forecast of its wave force in the same way, and Section 1.4.6 discusses what such a forecast is worth.

The move price and γ were tuned once, on scenario 11 with the LIDAR, over nine values of μ and five of γ, scoring the IAE relative to the loops'. No setting, and no preview, met the bound of no more travel than the loops in move-scale units (least 22.1 against 17.5), because the loops hold torque still above rated, so the bound is reported rather than imposed. The largest μ within 10 % of the best score, 0.4286 at $\mu = 1$, is $\mu = 3$ at 0.4314. Two earlier definitions of the LIDAR error, a fixed 1.0 m/s and 5 % of the wind, gave scenario 1 106.8 and 33.3 MN m of torque travel in 180 s, against 6.65 with the shipped error. Under those errors scenario 11 peaks at 100.67 and 100.51 %, so the shipped definition removes laminar jitter and does not buy the headline result.

In MANUAL the MPC plans nothing, and the manual overspeed case is identical under both structures. A rejected plan hands control to the loops, which track the applied commands every cycle, marks the historian, and returns control with the next valid plan, and the first plan is solved at $t = 0$. A switch to the MPC moves the commands by at most one capped step, and a switch back, or a fallback, moves the pitch command by at most 0.09°, but above rated the loops' constant-torque law commands the rating at once, a step of up to 0.667 MN m at 120 s of scenario 4 against 0.062 MN m for the MPC's own largest step in the 5 s before, which the actuator follows at its rate. The application ships no trip and no feather action, and the overspeed alarm, the actuator limits, the AUTO pitch limit and the torque rating act under the MPC exactly as under the loops. With the hydraulics slowed to 0.5 °/s in scenario 11 the alarm latches under both, at 33.98 s under the loops and 35.18 s under the MPC.

The plant integrates the model with the run's parameters, the instruments are the exact speed, pitch and hub wind and the preview, and the controller reads only what the loops read, the declared table, its own state and the preview. A test that gives a second plant eleven other parameters and overwrites its states with the first plant's every step finds the commands bit-identical over 6000 steps, and 24 of 24 mutated copies that read a true state, a run parameter or extra future wind fail the tests. A plan takes about 0.2 ms in Chrome, 172 to 210 µs, and the whole controller about 3 % of a core at ten times real time, at 59.8 to 60.0 frames per second.

4.5.7 Scenario Tree on the Declared Preview Error

The MPC takes the previewed wind as exact, although the LIDAR declares an error with it. A secondary setting of the MPC, its wind scenarios, keeps that plan as the default (One) and adds a scenario tree (Tree of 3), the multi-stage form of MPC [24, 25], that plans on the forecast and on the forecast plus and minus its declared error. Each wind keeps a plan of its own after the first move, because the next decision will see which wind arrived, and only the move applied before that decision is shared, the condition known as non-anticipativity. The objective is the expected cost over equally likely winds [26], not the worst case of min-max feedback MPC. Every scenario except scenario 12 ships the setting at One.

The winds are $V_k^{(0)} = V_k$ and $V_k^{(1,2)} = V_k \pm \bar{e}_k$, where $\bar{e}_k = (\ell_k/5\,\mathrm{s})\,\sigma_L$ is the one-standard-deviation envelope of the LIDAR error $e_k = \bar{e}_k n$, computed from the channels the preview reads and never from the error it draws. Five seconds ahead $\bar{e}$ is 1.95 m/s at 14 m/s in 14 % turbulence, and over ten seeds of scenario 12 the error the LIDAR makes five seconds ahead, divided by $\bar{e}$, has a standard deviation

of 1.01. With the preview off the envelope surrounds the held hub wind, and with the perfect preview it is zero. Three equally weighted winds one standard deviation apart carry two thirds of the declared variance, a spread fixed before any run.

Each wind has its own copy of Equation (4.15) at the same linearization point, with its own forcing and targets, which gives nine states and six inputs over the same 50 steps and 12 move blocks, and the plan solves Equation (4.17).

$$\min \frac{1}{3}\sum_{i=0}^{2} \mathscr{J}^{(i)} \quad \text{s.t.} \quad \beta_{c,1}^{(i)} = \beta_{c,1}^{(0)}, \;\; T_{c,1}^{(i)} = T_{c,1}^{(0)}, \;\; i = 1,2 \tag{4.17}$$

Here $\mathscr{J}^{(i)}$ is Equation (4.16) on the wind $V^{(i)}$, with slacks of its own and every weight, price, cap and box unchanged, and each equality enters the solver as two inequality rows. The program has 78 variables and 602 rows against the one-wind plan's 26 and 198. With no declared error the branches coincide, and the move equals the one-wind plan's within 3.1×10^{-9} degrees over 7395 decisions in four scenarios. A plan is also rejected when a command of any branch leaves its box or the three first moves differ by more than 10^{-6}.

The hedge is small by construction. When every branch carries the same model and the same weights and no constraint is active, minimizing over each branch's later moves leaves one quadratic in the shared move whose linear term is the mean of the branches' terms, so the shared move is the mean of the first moves of three one-wind plans made on $V^{(0)}$, $V^{(1)}$ and $V^{(2)}$. The plans on the higher and the lower wind nearly cancel, and the remainder comes from the curvature of aerodynamic torque in the wind and from the fine-pitch bound. Although a soft ceiling or a box is active at every decision of scenario 12, the tree's first pitch move lies within 0.01° of that mean at 1136 of the 1199 decisions, and it exceeds the one-wind plan's by 0.027° on average, where a one-wind plan told $V + \bar{e}$ moves 0.144° more.

A tree plan takes 2.9 ms in Chrome against 0.21 ms for one wind, about fourteen times as long, and the controller uses 3.2 % of a core at 1× and 32 % at the 10× top speed, with 60 frames per second at both on a host carrying other work. Because the plan exceeds the 2 ms threshold the apps set for a plan, a frame guard stops stepping a frame after 25 ms of wall time while the tree plans, so a slower computer runs below the requested speed with the same steps and the same logged rows. A switch between one wind and the tree in the middle of a gust, with the turbulence and the preview off, moves the pitch command by exactly the 0.061° step that the plan takes without the switch.

4.6 Results

Table 4.4 lists the thirteen set-ups. Each pins only the quantities it is about and leaves everything else at the defaults of Table 4.2, so a scenario differs from the base machine in exactly the settings its row names.

Table 4.5 reports the scorecard for the declared horizons. Runs use the shipped engine with its 0.01 s physics step, exact control timing and the default seed, which is the app's 20260824 unless the scenario declares its own. Scenario selection restores its settings, and comparisons change only the stated setting before the first step. IAE integrates absolute rotor-speed error against the logged torque-loop setpoint, in radians, and travel is the sampled total variation of actual generator torque at 10 Hz. Saturation uses the logged controller flag, matching the torque trend shading. Deliberately holding rated torque in Region 3 is excluded because pitch regulates speed there. The rating remains visible on the torque strip. Physical torque-bound occupancy can be computed separately from the CSV.

Rows 11 and 12 run under the MPC of Section 4.5.6 with the LIDAR preview, as the scenarios load (row 12 with the scenario tree of Section 4.5.7), and every other row runs under the scheduled loops or, in scenario 10, in MANUAL.

Table 4.4 Scenarios and What Each One Changes

#	Name	Setting
1	Steady 6 m/s	Region 2 at the C_P peak, no turbulence
2	Wind ramp through rated	7 to 22 m/s at 0.15 m/s per s, $t = 10$ to 110 s
3	Step gust	9 m/s, +6 m/s at $t = 25$ s, back at 70 s
4	Turbulent wind	11 m/s mean, 14% intensity, seeded
5	Pitch rate limit	$\dot{\beta}_{max} = 2$ deg/s, 13 to 19 m/s step
6	Generator torque limit	$T_{rated} = 2.8$ MN·m at 10.5 m/s
7	Gain schedule off	21 m/s, 14% turbulence, +4 m/s at $t = 30$ s
8	Doubled rotor inertia	$J_{eq} = 8.0 \times 10^7$, 8 to 14 m/s step
9	Region-transition hunting	10.9 m/s mean, 9% turbulence, three mean shifts
10	Manual overspeed	MANUAL, 16 m/s, $T_{gen} = 0.4$ MN·m, $\beta = 0$
11	MPC or PI loops in a gust	MPC with LIDAR preview, 14 m/s, +6 m/s at $t = 30$ s, back at 45 s
12	Plan on one wind or three	MPC tree with LIDAR, the gust of 11 in 14% turbulence, seeded
13	Challenge	weak gains, schedule off, ramp to 18 m/s

Table 4.5 Measured Scorecard Values at the End of Each Run

#	t (s)	IAE (rad)	MV travel (MN·m)	Sat. (%)	Peak ω (% rated)	Peak $\dot{\beta}$ (deg/s)	P (MW)
1	60	0.000	0.00	0.0	57	0.0	0.75
2	130	2.911	2.68	0.0	101	3.0	5.00
3	120	4.473	3.44	0.0	106	8.0	2.42
4	120	4.819	14.72	2.2	105	7.1	4.90
5	90	1.292	0.00	0.0	111	2.0	5.00
6	90	0.789	0.00	0.0	100	4.6	3.35
7	120	0.865	0.00	0.0	102	8.0	5.04
8	120	1.838	2.22	0.0	104	6.0	5.00
9	180	5.652	11.58	0.3	104	7.2	4.97
10	200	566.646	3.63	0.0	345	0.0	1.65
11	120	0.447	1.07	0.0	100	2.5	4.97
12	120	0.913	6.21	0.0	102	6.1	4.53
13	200	25.981	9.77	0.0	126	1.8	5.13

Above the torque setpoint cap, a rotor held at rating still differs from the logged target by 0.0363 rpm. That offset contributes to IAE, but it cannot in general be subtracted as a common floor because the metric takes absolute error. Pitch regulation should also be assessed against rated rotor speed.

4.6.1 Handover Through Rated Wind

Scenario 2 ramps wind from 7 to 22 m/s at 0.15 m/s per second. The target reaches its cap at 10.54 m/s. Region 2.5 is declared at 34.76 s, with wind 10.714 m/s, and Region 3 at 37.06 s, with wind 11.059 m/s. At the latter change actual generator torque is about 2.96 MN·m. The selector now

commands 4.18 MN·m and the rate-limited actuator catches up. Thus the region change can occur before actual torque reaches its rating. Pitch was already active and supplies the entry condition. At 130 s the blades have feathered to about 19.4 degrees, power is 5.00 MW, and speed is within 0.001 rpm of rating. The logged 0.0363 rpm residual is the setpoint offset.

> **TRY IT IN THE APP** **2 · Wind Ramp Through Rated**
>
> Load scenario 2 and press Run in Auto, at 5 times speed if you like. Watch the region badge. It turns to 2.5 at 34.76 s and to 3 at 37.06 s, while the generator torque is still about 2.96 MN m. By 130 s the blades have feathered to about 19.4° and the power holds 5.00 MW.

At steady 6 m/s, scenario 1 holds 6.87 rpm and $C_P = 0.482$, producing 0.746 MW electrical from about 1.65 MW available wind power. Changing the controller target ratio to 5.5 and running 240 s gives $\lambda = 5.500$, $C_P = 0.4183$ and 0.649 MW, a 13.0% loss for the same wind and machine.

4.6.2 Cost of Disabling the Gain Schedule

Scenario 7 starts at 21 m/s with 14% turbulence and adds 4 m/s at 30 s. The same seeded 120 s input is used with scheduling off and on. Without the schedule, IAE is 0.865 rad and peak speed 102% of rating. With it, IAE is 3.058 rad and peak speed 107%. Terminal logged errors are 0.14 and 0.26 rpm, respectively. This endpoint is one turbulence realization, not a steady error. For the ideal integrating plant, increasing both PI gains at fixed integral time raises both natural frequency and damping ratio with the square root of gain. A faster trace therefore need not ring in this reduced model.

Actual pitch travel is 255.2 degrees without scheduling versus 108.4 degrees with it, a factor of 2.35. Command travel is 350.9 versus 118.4 degrees. Peak pitch rate reaches the 8 deg/s constraint without scheduling, but is 5.2 deg/s with it. This measures actuator duty, not a quantified fatigue life or a stability margin against the omitted tower and drivetrain modes.

> **TRY IT IN THE APP** **7 · Gain Schedule Off, High Wind**
>
> Load scenario 7 and run it for 120 s as it loads, with the Pitch gain schedule off, and note that the Peak pitch rate reaches the 8 °/s limit. Reset, switch Pitch gain schedule to On, change nothing else and run again. The speed IAE rises from 0.865 to 3.058 rad, but the blade travel falls from 255.2 to 108.4°.

Pause. Try Exercise 4.2 before reading on.

4.6.3 Actuator and Rating Limits

Scenario 5 reduces the pitch rate limit from 8 to 2 deg/s, one quarter of the default, and steps wind from 13 to 19 m/s. Peak speed is 111% of rating. Raising pitch gain from 3.7673 to 6 with the same 2 deg/s actuator changes the response: IAE rises from 1.292 to 7.452 rad and peak speed from 111 to 116%. It cannot increase actuator speed. Restoring 8 deg/s at default tuning reduces IAE to 0.749 rad and peak speed to 106%, with peak actual pitch rate 3.6 deg/s. Gain therefore still affects command history under a rate constraint, and "more gain changes nothing" would be an incorrect description.

Scenario 6 derates torque to 2.8 MN·m at 10.5 m/s. It starts in Region 3, settles at 3.13 degrees of pitch and 12.10 rpm, and produces 3.35 MW. Actual torque is at its bound throughout, but both the scorecard saturation value and its controller flag remain zero, because pitch has control authority.

Scenario 8 doubles inertia to 8×10^7 kg m^2. It peaks at 104% rated speed with a 6.0 deg/s peak pitch rate, and the scorecard settling value is 10.7 s. Doubling inertia halves acceleration for a fixed imbalance. At fixed PI gains, the ideal integrating-plant natural frequency scales as $J^{-1/2}$, so the poles do not simply halve their frequency. Restoring a chosen pole placement requires gains proportional to inertia.

Table 4.6 Scenario 11 Over 0 to 120 s, With the Ten-Seed Peak Speed and Energy of Scenarios 4 and 9 Over 0 to 180 s

	PI loops	PI, sched. off	MPC, preview off	MPC, LIDAR	MPC, perfect
Peak ω (% rated)	105.59	101.56	102.39	100.29	100.34
Speed IAE (rad)	1.036	0.558	0.595	0.447	0.468
Pitch travel (deg)	17.5	17.0	16.9	20.5	17.4
Torque travel (MN·m)	0.00	0.00	0.72	1.07	0.73
Peak P (MW)	5.28	5.08	5.01	5.00	5.00
Energy (kWh)	166.8	166.7	166.1	165.8	166.1
Mean solve, headless (µs)	—	—	161	179	158
Ten turbulence seeds, 0 to 180 s					
Sc. 4 peak ω (%)	104.72	—	102.42	101.32	100.97
Sc. 4 energy (kWh)	216.50	—	209.44	205.26	209.28
Sc. 9 peak ω (%)	103.65	—	101.74	101.36	100.83
Sc. 9 energy (kWh)	225.83	—	221.20	219.93	221.98

4.6.4 Region Transitions in Turbulence

Scenario 4 produces twelve region changes in 120 s, and scenario 9 produces fourteen in 180 s and IAE 5.652 rad. The pitch loop stays active, but crossing Region 3 changes the torque selector and its tracking mode. Actual torque responds through its lag and rate constraint, so a region change should not be described as a purely cosmetic label. The hysteresis and dwell limit repeated selection near the boundary. These counts alone do not quantify fatigue or predict operation without that logic.

4.6.5 Predictive Control Against the Scheduled Loops

Scenario 11 holds the machine at rated in a laminar 14 m/s, raises the mean wind to 20 m/s at 30 s and returns it at 45 s, and it loads the MPC with the LIDAR preview. Table 4.6 compares it over 120 s with the shipped loops, whose tuning is the challenge's suggested tuning, the loops with the schedule off and the MPC under the other two previews, in the same wind, with no case retuned and 1200 plans per run, none rejected. The MPC with the LIDAR is told the wind ahead and the loops are not, so the like-for-like comparison is the preview-off column, given only the hub wind that the loops also read.

Under the loops the wind at the rotor starts to rise at 30.03 s, and the speed peaks at 105.59 % of rating at 33.12 s while the power reaches 5.28 MW, as shown in Figure 4.3. Told of the gust, the MPC raises its pitch command 0.5° above its line-out at 28.0 s, holds 11.4° at 30 s against the loops' 9.0°, and keeps the speed at or below 100.29 %, a value reached at 45.65 s as the gust passes, with the power at its 5.00 MW rating. With the preview off it peaks at 102.39 % at 31.93 s. The preview is therefore worth 2.1 of the 5.3 points, 40 % of the gain over the loops, and its share stays between 29 and 41 % for gusts of 3, 6 and 9 m/s, a gust from 12 m/s, a gust lasting 3 s, and 5 or 9 % turbulence over ten seeds, where the LIDAR peaks below the preview off at every seed.

The remaining 3.2 points come from how tightly the MPC's weights hold the speed, not from its model reading the hub wind. Given no wind reading at all, with a 30 s average of the hub wind in its model and filter, the MPC with the preview off peaks at 102.50 %, and with its torque held at the rating at 102.21 %, while its move price moves the peak to 101.26 % at $\mu = 0.3$ and 104.16 % at $\mu = 30$. A shipped option of the loops that is as tight does as well on this gust. With the schedule off the loops peak at 101.56 % with an IAE of 0.558 rad and 17.0° of pitch travel. The tight loops pay elsewhere, because on the gust from below rated of scenario 3 and in the turbulence of scenarios 4 and 9 the schedule off and the doubled pitch gains ride the 8 °/s limit for 0.90 to 5.56 s and still reach 102.59 to 104.74 %, where the MPC with the preview off reaches 100.61, 103.02 and 101.41 % without touching the limit. In turbulence the hub-wind reading does matter, and on scenario 9 the MPC without it peaks at 102.48 % with an IAE of 2.60 rad, against 101.41 % and 2.03 rad.

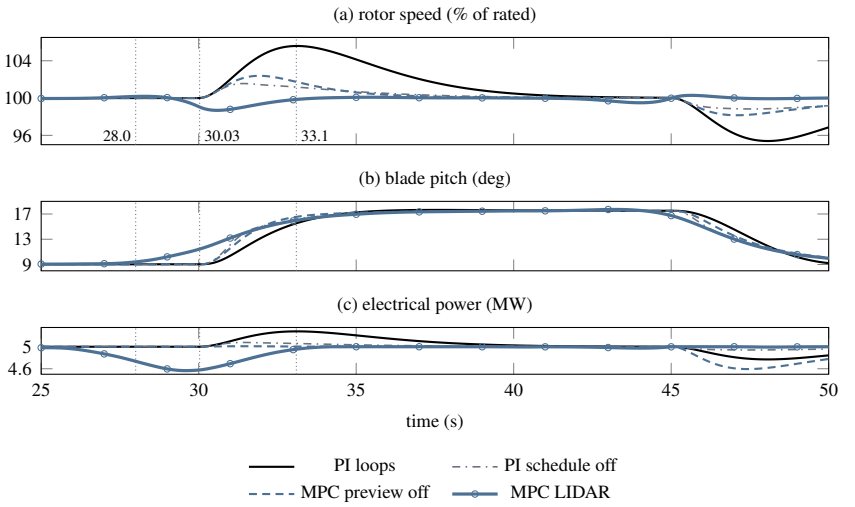

Figure 4.3 Scenario 11 from 25 to 50 s under the scheduled loops (black, solid), the loops with the schedule off (gray, dash-dotted), the MPC with the preview off (blue, densely dashed) and the MPC with the LIDAR preview (blue, very thick, with circles). The thin dotted vertical marks are the LIDAR plan's first pitch move of 0.5° at 28.0 s, the wind reaching the rotor at 30.03 s and the loops' peak at 33.1 s. (a) Rotor speed. (b) Blade pitch. (c) Electrical power.

The preview has a cost before the gust arrives. The plan feathers early and lowers its torque command to 3.83 MN m to keep the speed on target, so the power falls to 4.57 MW at 29.6 s while the loops hold 5.00 MW, as Figure 4.3c shows. The LIDAR's small error keeps the plan moving, so the MPC travels 20.5° of pitch against 17.5 and 1.07 MN m of torque where the loops hold rated torque still, and it makes 0.6 % less energy, 165.8 against 166.8 kWh, of which 0.41 kWh is the loops' power above the rating. A gust a second early or two seconds late still peaks at 100.19 to 100.49 %, and one announced that never comes costs 9.3° of needless pitch travel.

The energy losses in turbulence are larger, as the lower block of Table 4.6 shows. Over ten seeds the MPC with the LIDAR makes 5.19 % less energy than the loops in scenario 4, 3.26 % less with the preview off and 3.34 % with the perfect preview, and 2.61 % less in scenario 9. At the scenario's own seed, as it loads, the loss in scenario 4 is 5.8 %. In scenario 4, 10.07 of the 11.24 kWh are lost while the hub wind is below rated, where the MPC holds its speed target and keeps the blades pitched beyond 0.5° for 59.75 s of each run against the loops' 8.72 s, capturing 217.0 against 228.2 kWh from the wind, while the loops let the rotor speed up in the gusts and give that energy back at fine pitch and rated torque. The loops' power above the rating is only 0.89 kWh of the difference, and $\mu = 30$ does not recover it. The MPC was tuned on the laminar scenario 11 with energy outside the score, and this loss appeared afterward. The MPC therefore trades energy for speed regulation and is not the better controller overall. It also moves the torque 2.4 and 1.8 times as far in scenarios 3 and 9, raises scenario 9's power standard deviation from 0.42 to 0.52 MW, travels 6.65 MN m of torque in the steady wind of scenario 1, and has an IAE of 1.624 against 1.272 rad in scenario 7 at a third of the unscheduled loops' pitch travel. Figure 4.4 draws the peak speeds and the energy of Table 4.6 as bars.

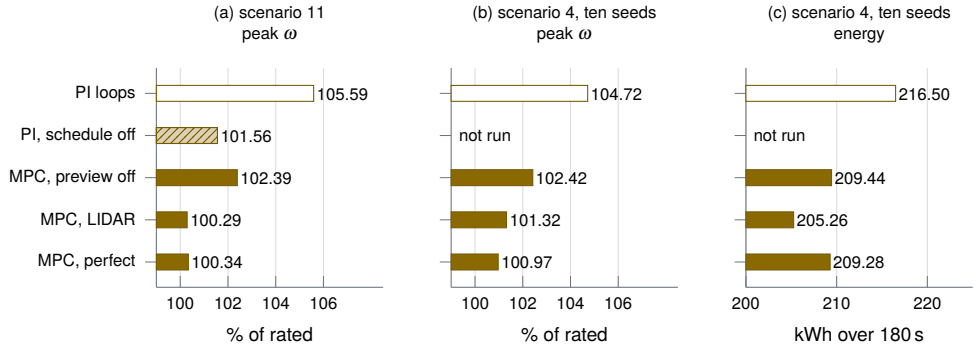

Figure 4.4 Table 4.6 as a chart: (a) the peak rotor speed in the scripted gust of scenario 11, (b) the peak speed and (c) the energy over ten turbulence seeds of scenario 4. The white bar is the shipped loops, the hatched bar the loops with the schedule off, solid bars the predictive controller under its three previews. The axes of (a) and (b) start at 99 % of rated and that of (c) at 200 kWh, so bar lengths show differences, not ratios. Notice that every predictive bar is shorter than the loops' in (b) and in (c): the controller that holds the speed tighter in turbulence also makes less energy.

> ✦ **DID YOU KNOW?** **Tighter Speed, Less Energy**
>
> From the simulator: over ten turbulence seeds of scenario 4, the predictive controller with the LIDAR preview held the peak rotor speed to 101.32 % of rated against the loops' 104.72 %, yet made 205.26 kWh against 216.50 kWh (Table 4.6). Better speed regulation did not bring more energy.

The same settings help where the loops lack authority. Over 180 s the MPC peaks at 100.32 % against 105.92 % on the gust of scenario 3, at 100.11 % against 111.27 % with the 2 °/s actuator of scenario 5, which its model still takes to be 8 °/s, and at 100.72 % against 104.17 % with the doubled inertia of scenario 8. After a Reset into AUTO from scenario 10, with the blades flat at 16 m/s, it overshoots to 112.68 % against 109.27 %, and the pitch move cap accounts for all of it, because it holds the command to 6.26 °/s while the loops run on the limit. With the cap at the full 8 °/s average the MPC peaks at 110.60 %, and with no cap at 108.75 %. Across the eleven scenarios that ran in AUTO before scenario 12 was added and the three previews, 66 000 plans were solved in runs of 200 s, none was rejected, and no command left its live bounds.

4.6.6 Planning on Three Winds

Scenario 12 repeats the gust of scenario 11 in 14 % turbulence and loads the MPC with the LIDAR and the tree. Table 4.7 compares seven structures over 0 to 120 s on the same ten turbulence seeds, with the MPC rows at the scenario's settings and nothing retuned. The hub wind hashes identical under the loops and the three LIDAR plans, and so do those plans' previews. Two rows are rivals on the same preview, the one-wind plan given the forecast plus its declared error, which is the cheapest use of the tree's information and is not offered in the application, and the loops with the LIDAR as a collective pitch feedforward 0.3 s ahead, a lead chosen on this scenario. The application has no tower and shows no load, so the thrust column is an estimate made outside it by ideal momentum theory, with the run's $C_P = 4a(1-a)^2$ solved for the axial induction $a \le 1/3$ and $F = \frac{1}{2}\rho\pi R^2 V^2\, 4a(1-a)$, a stand-in for tower load and nothing more.

Paired by seed, the tree peaks 0.10 points below the one-wind plan (by 0.02 to 0.47) with 2.7 % less estimated thrust, 1.0 % less pitch travel and 0.6 % less energy, each at all ten seeds. At the scenario's own seed the difference is 0.05 points, 101.99 against 102.04 % at 40.24 s, which the scene cannot show and the trends can, as shown in Figure 4.5. The one-wind plan on $V + \bar{e}$ peaks 0.52 points below the tree at all ten seeds, with 7.1 % less pitch travel and 1.4 % less estimated thrust (at

Table 4.7 Scenario 12 Over 0 to 120 s, Mean of Ten Turbulence Seeds

	Peak ω (% rated)	Thrust est. (MN)	Pitch (deg)	Torque (MN·m)	Energy (kWh)
PI loops	107.85	0.755	169.5	1.50	166.6
PI, LIDAR feedforward	101.76	0.671	270.1	1.37	166.3
MPC, one wind	101.76	0.589	124.0	9.50	160.8
MPC, tree	101.66	0.573	122.7	10.08	159.9
MPC, one wind on $V + \bar{e}$	101.14	0.565	113.9	11.33	154.8
MPC, one wind, preview off	104.00	0.674	211.0	16.02	161.9
MPC, tree, preview off	103.51	0.654	205.2	14.60	161.7

Every MPC row solves 1200 plans a run and rejects none. The thrust is an estimate by momentum theory that the application does not compute.

six seeds), but it makes 3.2 % less energy than the tree, a stronger hedge rather than a better one. At equal energy neither rule dominates, because told a quarter of the error the one-wind plan makes 159.7 kWh, peaks at 101.63 % and travels 121.0° against the tree's 159.9 kWh, 101.66 % and 122.7°, but carries 0.583 against 0.573 MN of estimated thrust. The feedforward loops match the one-wind plan's 101.76 % with 2.2 times its pitch travel and 3.39 s a run on the 8 °/s rate limit, which no MPC row touches. With the preview off the tree peaks 0.49 points below the one-wind plan at all ten seeds. Halving the integration step to 0.005 s leaves every paired verdict unchanged.

The declared error describes turbulence and not the timing or the size of a gust. With the plant's gust a second early, one or two seconds late, 2 m/s stronger or absent while the LIDAR reports the scripted one, the tree stays 0.08 to 0.15 points below the one-wind plan, as under the announced gust, and in the laminar wind of scenario 11 it equals the one-wind plan to three decimals under all six forecasts. In twelve worst cases, among them extreme wind, slowed hydraulics, a derated generator, a lowered alarm and model mismatch, 27 600 tree plans were solved with none rejected and no command outside its box.

4.6.7 Manual Operation and the Graded Challenge

Scenario 10 commands 0.4 MN·m and zero pitch at 16 m/s. From its balanced initial plant state, actuator motion creates a torque imbalance and the rotor accelerates to 41.77 rpm, or 345% rated, at $\lambda = 17.22$. The alarm latches at 1.60 s and the simulation continues. It is an annunciator, not a protective shutdown model. A real protection system is outside the modeled control layer.

The challenge checks the logged speed error from 132 s onward, Region 3 at the final sample, peak speed below 120% of rating, and peak pitch rate strictly below the configured actuator limit. The challenge inherits weak torque and pitch gains with scheduling off and a ramp from 9 to about 18 m/s under 8% turbulence. At 200 s this tuning has IAE 25.981 rad, 9.77 MN·m torque travel and peak speed 126% rated. Its worst logged speed error from 132 s onward is 1.086 rpm, so it fails both speed and overspeed requirements. Suggested tuning, applied before running, gives IAE 4.621 rad, 5.26 MN·m travel, peak speed 104%, peak pitch rate 6.4 deg/s and worst error 0.336 rpm in that window. It passes the challenge predicate, with 82.2% less IAE and 46.2% less torque travel.

With only the controller switched to the MPC the challenge also passes, with a worst error of 0.22 rpm from 132 s, a 102 % peak, 2.7 °/s of peak pitch rate and an IAE of 2.288 rad. The predicate scores no energy, of which the MPC makes 1.5 % less than the suggested loops, 0.2576 against 0.2614 MWh, with 19.42 against 5.26 MN m of torque travel. The challenge and its rules are unchanged.

4.7 Wind Measurement in the Region 2 Setpoint

Equation (4.11) needs the wind speed, and the application supplies the true instantaneous value including the turbulent component, sampled at the same 50 Hz the controller runs at. That is a perfect rotor-effective anemometer, and it is precisely the measurement a real turbine cannot make. A nacelle

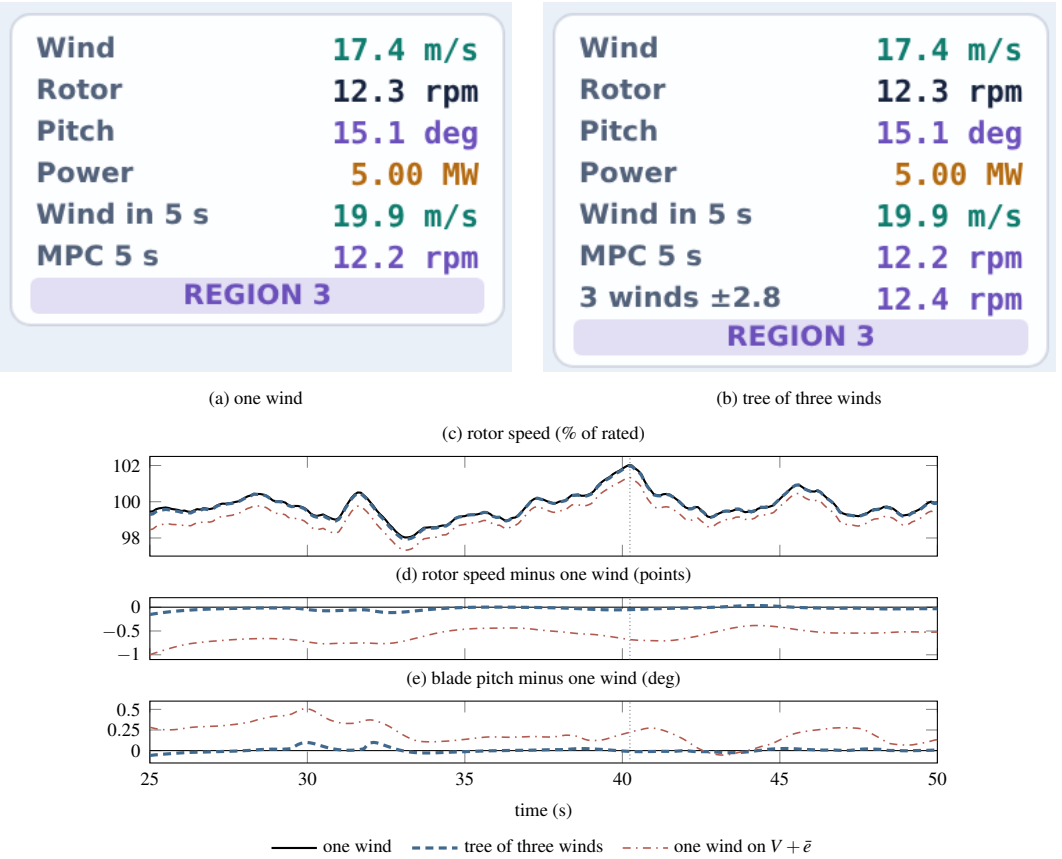

Figure 4.5 Scenario 12 at its own seed under the MPC with the LIDAR. (a), (b) The operating box at 40.3 s, just after both plans peak, where the tree adds the branches' spread 5 s ahead and the highest speed any branch predicts. (c) to (e) Rotor speed and the speed and pitch differences from the one-wind plan (black, solid) for the tree (blue, very thick, densely dashed) and the one-wind plan on $V+\bar{e}$ (red, dash-dotted), with the thin dotted line at the 40.24 s peak.

anemometer sits in the wake of the rotor it is supposed to be measuring, it reports one point rather than an average over a 12 469 m^2 disc, and the quantity the setpoint requires is the rotor-effective wind speed, which is a spatial average weighted by the aerodynamic loading. Lidar instruments measure the inflow ahead of the rotor and are an active research topic for exactly this reason [13].

The consequence has to be stated with the results. The Region 2 loop as written is a setpoint-tracking loop given information no machine has, and it converges faster than any real implementation would. The advanced option that replaces it, $T_{gen} = K\omega^2$, is the wind-free alternative and is what the reference designs use below rated [10, 14]. Running the two against each other in the same wind is the cleanest way in the application to show what a measurement is worth, and it is also the cleanest way to show what a feedforward term costs when its constant is wrong, because K depends on $C_{P,max}$ and λ^{*3} and both are properties of a rotor that erodes, ices and accumulates dirt over its life.

The MPC of Section 4.5.6 makes the information question explicit, because with the LIDAR preview it is told the next five seconds of wind and the loops are not, so every comparison in Section 4.6.5 carries the preview-off column, given the same hub wind that Equation (4.11) reads. In turbulence the preview's advantage rests on the instrument model, which hands the LIDAR the turbulence that will arrive with an error that is zero at the rotor. Told only the scripted mean wind,

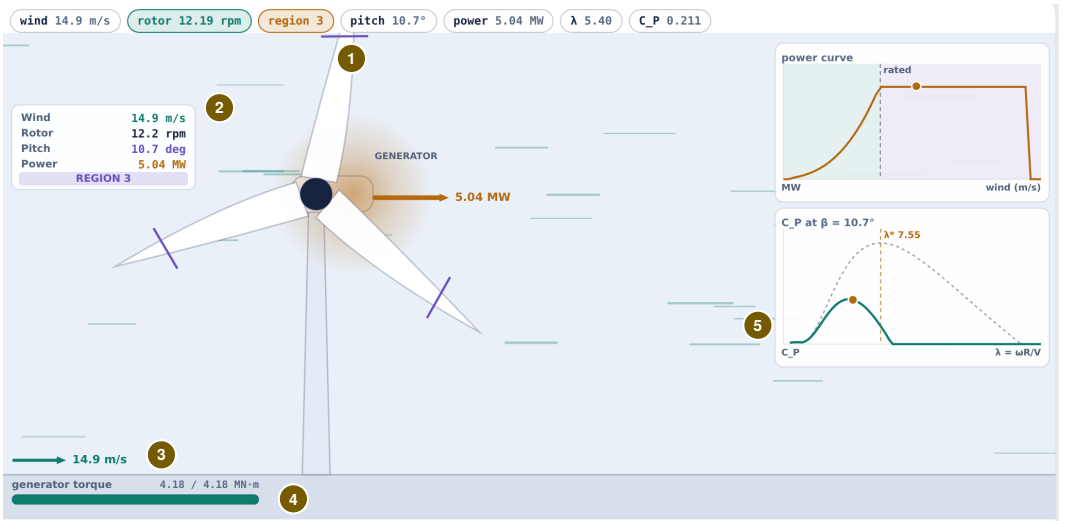

Figure 4.6 Wind ramp through rated, paused at $t = 62.7$ s, cropped to the plant view. The region changes came at 34.8 and 37.1 s, after which generator torque holds 4.18 MN m and the blades feather to 10.7° while rotor speed holds 12.19 rpm against its 12.10 rpm rating and power holds 5.04 MW. Callouts: (1) the badge row with the rotor speed and the region, (2) the readout box of wind, rotor speed, pitch and power, (3) the wind arrow, (4) the generator torque bar at its rating, and (5) the power curve and the power coefficient curve at the present pitch, with λ^* marked.

with the turbulence ahead decaying from its present value and a 1 % error, the MPC does worse than with the preview off, peaking over ten seeds at 102.15 against 101.74 % on scenario 9 and 103.21 against 102.42 % on scenario 4, and an instrument that sees the coming turbulence only through a 6 s moving average also loses, at 102.70 and 103.84 %. At three times the declared turbulence error plus 5 % of the wind, scenario 9's speed IAE is 5.686 rad against the loops' 5.652, with 2.6 times their pitch travel. In laminar wind only the 1 % part acts, which is why the scripted gust of scenario 11 is where the preview's contribution is robust.

4.8 Technology in Practice

The loops of Section 4.5 read one speed and one wind and move two actuators, the converter that sets the generator torque and the pitch system that turns the blades. Figure 4.7 draws them with that hardware in place of the schematic's symbols.

What Measures It

Rotor speed is the one controlled variable. In the reference design it is measured on the generator, after the gearbox, and filtered before either controller sees it [14]. A classic speed sensor is the tachogenerator, a small direct-current generator whose voltage is proportional to its shaft speed [27]. The wind is harder to read. A nacelle anemometer sits in the wake of its own rotor, as Section 4.7 explains, which is one reason the reference designs use the wind-free torque law below rated. The chapter hands its controller the exact rotor speed and the true rotor wind, so a filter's lag and an anemometer's error are both absent from its results.

turbine controller: SIC-101, SIC-102 and SY-103 every 0.02 s (50 Hz). The converter follows its command through a 0.05 s lag and 1.5 MN m/s; the pitch drives through 0.25 s and 8 °/s.

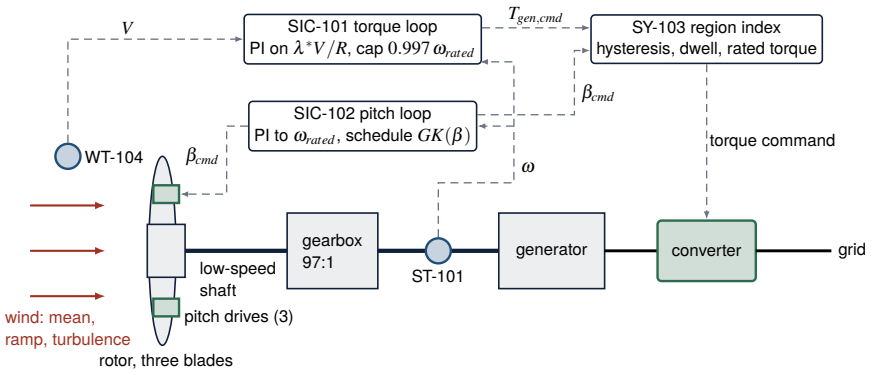

Figure 4.7 The torque and pitch loops drawn with their hardware. The speed sensor ST-101 on the generator shaft feeds both loops. The torque loop SIC-101 tracks the speed that the rotor wind WT-104 calls for, capped just below rated, and the pitch loop SIC-102 regulates to rated speed through the three pitch drives, with its gain divided by the schedule $GK(\beta)$. The region index SY-103 reads the pitch command and passes the torque loop's command below Region 3 and rated torque in it. The predictive controller, which can replace both loops and the selector, is left out. Notice that both loops act on the same speed and that only the selector decides which one sets the torque.

> **TECHNOLOGY** **Sensor: Generator Speed for Both Loops**
>
> In a utility-scale turbine both the torque and the pitch controller use the generator speed as their sole feedback input. The reference design filters it with a single-pole low-pass filter whose corner, 0.25 Hz, is about a quarter of the blade's first edgewise frequency, so that blade vibration does not excite the loops [14]. A research turbine reads its shaft and pitch positions with absolute position encoders sampled at 100 Hz [8]. Losing the speed signal is the dangerous failure. A drive's speed loop then reads a stopped shaft and runs the motor away to full speed, and a large propulsion drive carries two tachogenerators for redundancy [27]. On a trend the reading falls to zero while the true speed rises.

What Moves It

Below rated the converter moves the rotor. It sets the generator torque electronically, and on a research turbine the torque loop has a rated bandwidth of 500 rad/s [8], so torque is the fast actuator. The chapter's converter follows its command through a 0.05 s lag and a 1.5 MN m/s rate limit. Above rated the pitch system takes over, and its rate limit, not its gain, sets how fast it can meet a gust, as scenario 5 of Section 4.6 shows.

> **TECHNOLOGY** **Final Element: The Blade Pitch System**
>
> Each blade turns on a large grease-lubricated rolling bearing, most often a four-point contact ball bearing, driven through spur gears by geared motors or by hydraulic cylinders [28]. On a research turbine each blade has its own electromechanical servo that reaches 18 °/s, and the servo follows a pitch rate command with a fitted time constant of 1/60 of a second [8]. The reference design limits pitch to 8 °/s between 0 and 90° [14], the limits the chapter adopts with a 0.25 s lag. A slow pitch system shows on a trend as a pitch rate pinned at its limit and a rotor speed that overshoots, as in scenario 5.

Where the Controller Runs

A research turbine's control computer collects its measurements and closes its loops at a cycle rate of 100 Hz [8]. In the reference design the generator torque is a tabulated function of the filtered

generator speed across five regions, and the pitch loop is a gain-scheduled proportional-integral controller whose integral is saturated at the pitch limits for a fast transition between Regions 2 and 3 [14]. In process plants, programmable logic controllers (PLCs) close local feedback loops under a supervisory system, at cycle times from milliseconds to minutes [29]. The chapter's controller runs every 0.02 s, and it holds the handover in three places: the setpoint offset, the region index with its hysteresis and dwell, and back-calculation on the clamped loop.

> **IN PRACTICE** **Pitch Bearings Pay for Pitch Travel**
>
> Pitch bearing loads come from aeroelastic simulation of the design load cases of IEC 61400-1. The damage with the most critical consequence is a cracked bearing ring, found at bolt holes or stiffness changes, since it risks the blade detaching, and replacing a pitch bearing is costly because the blade must come off the hub [28]. Small oscillations wear the raceway, as standstill marks or false brinelling. Protection runs of a few degrees, and moves of 30° or more once a day, spread the grease again [28]. The pitch travel the chapter reports is the duty these bearings carry.

What Is Modernizing It

Three developments touch the loops of this chapter. Individual pitch control turns each blade on its own rather than all three together, to reduce fatigue loads, at the possible cost of slightly less power and more cyclic loading of the pitch bearing raceways [28]. Turbine controllers are moving toward feedforward that watches the incoming flow and flies the rotor in response [9], the role the LIDAR preview plays in Section 4.5.6. And because every wind turbine is digitally controlled, its operation has to be secure against cyber-physical attack [9], as Ethernet, Internet Protocol and wireless devices raise the likelihood of incidents across control systems [29].

4.9 Limitations

The model omits the two dynamics that set the real bandwidth limit on a Region 3 pitch loop. There is no drivetrain torsion and no tower motion, so the argument that an unscheduled loop is dangerous rather than merely expensive is made in prose here and cannot be demonstrated in the application. What the application can demonstrate is the actuator duty and the time spent against the rate limit. The omitted modes are close enough to the loop to matter. The first tower fore-aft mode of the reference machine is at 0.32 Hz [14] and the first drivetrain torsional mode is near 1.7 Hz, against a pitch loop placed at $\omega_n = 0.6\,\mathrm{rad/s}$, or 0.095 Hz. The scheduled tuning therefore sits a factor of 3.4 below the tower mode, which is the separation a commercial design aims for. On an integrating plant under proportional-integral control ω_n scales as the square root of the gain, so the unscheduled loop at five times the gain moves to 0.21 Hz and the separation falls to 1.5. That is where a real machine would begin to excite the tower, and the application cannot show it.

> **OPEN CHALLENGE** **Larger, Softer Rotors**
>
> Wind turbines are the largest flexible rotating machines in the world, and they must run for 20 years or more under constantly changing loads. As blades grow they flex into and out of the wind, and the rotor meets its own shed vorticity, which calls design assumptions into question. Aeroelastic stability has rarely driven blade design so far, but that may change for future large, flexible rotors [4]. The chapter's rigid rotor, without tower or drivetrain modes, cannot show this. Progress would be pitch controllers designed and verified with the flexible modes in the loop.

The power coefficient surface is an approximation with three stated departures from the textbook coefficient set, and it should not be presented to readers as the standard correlation. Its 2.1 % error at the rated point is small enough that the rating lands where a reader expects and large enough to be worth stating.

The app's Learn texts state the same 0.1-degree exit threshold, 2 s Region 3 exit dwell and 6 s Region 2/2.5 dwell as the controller, and the derated-torque scenario distinguishes operation at the rated bound from the controller saturation flag.

The MPC's model keeps the sensitivities of the moment over its horizon, although a gust from 9 to 15 m/s more than triples $|\partial T_{aero}/\partial\beta|$, omits the rate limits within a step, and inherits every omission of Section 4.3, so its tighter speed regulation is not shown to be acceptable to the tower or the drivetrain. Its LIDAR is a declared instrument, handed the coming turbulence with an error that vanishes at the rotor, which is kinder than a real nacelle LIDAR, and its turbulent results depend on that, as Section 4.7 shows. Its pitch move cap is conservative and is the whole cause of its larger overshoot after a Reset into AUTO in scenario 10, a fallback above rated in turbulence steps the torque command by up to 0.667 MN m, and no preview setting met the travel bound of the tuning pass. Most of all, it captures 2.6 to 5.2 % less energy than the loops over ten seeds of turbulence near rated, so it is a speed regulator and not an energy optimizer, and an economic objective that prices energy against speed error and actuator duty is the extension this comparison points to.

> **OPEN CHALLENGE** **Predictive Control on a Real Turbine**
>
> A model predictive controller can respect pitch limits, structural load limits and generator overspeed limits at once, and a lidar in the hub or on the nacelle can preview the wind and further reduce structural loads. Such controllers are computationally demanding, however, and by 2023 they had been demonstrated on real turbines only in very limited experiments [9]. The chapter's predictive controller plans on a declared model with a declared lidar. Progress would be field tests that measure its energy, pitch travel and loads against the scheduled loops on one machine.

The scenario tree plans on three winds one declared standard deviation apart with equal weights, which carry two thirds of the declared variance and nothing about a gust's timing or size, and it prices its soft limits by their expected value rather than holding them on every branch. Its thrust advantage is measured on a momentum-theory estimate of a machine without a tower, its cheapest rival is a stronger hedge that gives up a further 3.2 % of the energy, and each of its plans costs about fourteen times the one-wind plan's solve.

Finally, the loop is given a measurement no turbine has, as Section 4.7 describes, and rotor speed is read with no noise, no lag and no quantization. Wind is the only source of randomness in the plant, and the LIDAR preview of Section 4.5.6 draws its error from a separate seeded stream.

4.10 Conclusion

This chapter studied a wind turbine control simulator built around the handover between two actuators that reach the same measurement. The model is five states, a rotor speed on a lumped 4×10^7 kg m^2 inertia, two rate-limited actuator states and two wind states, driven by a fitted power coefficient surface whose peak is placed at $\lambda^* = 7.55$ with $C_{P,max} = 0.482$ and whose three departures from the textbook coefficient set are stated. Below rated a proportional-integral loop on generator torque holds the rotor at $\lambda^* V/R$ and captures 0.482 of the available power, and detuning the target tip-speed ratio to 5.5 costs 13.0 % of the electrical output at 6 m/s. Above rated the torque is frozen at 4.18 MN m and collective pitch regulates speed, with the pitch onset assisted by a 0.3 % setpoint offset and rated torque selected explicitly by the region index. The wind ramp crosses both boundaries with a terminal speed error of 0.04 rpm against the logged setpoint, and actual actuator motion remains continuous through its physical dynamics. The pitch sensitivity of aerodynamic torque grows by a factor of 4.9 across Region 3, and switching the schedule off improves IAE from 3.058 to 0.865 while raising blade travel from 108.4° to 255.2° and putting the actuator on its rate limit, so the fixed high gain buys speed regulation at the price of pitch travel. Margin against the omitted structural modes is an engineering concern beyond these simulations.

The model holds air density and inertia constant, and it contains no drivetrain torsion, no tower motion and no blade flexibility, so the loads that set the real bandwidth limit on a pitch loop are outside its reach. Adding a two-mass drivetrain, a fore-aft tower mode, individual pitch and a rotor-effective wind estimator would each bring the simulator closer to a commercial controller. Those additions would change what the application can demonstrate, not what it teaches. The predictive controller added as an option shows that a five-second preview supplies about two fifths of its gust advantage, that the rest is a tighter speed objective, which the loops match on that gust with the schedule off but not in turbulence, where the same objective costs 2.6 to 5.2 % of the energy over ten seeds. Its scenario-tree setting lowers a turbulent gust's peak by 0.10 points and the estimated thrust by 2.7 % for 0.6 % of the energy, because with one model in every branch its shared move stays close to the mean of three one-wind moves, whose outer branches nearly cancel. The wind turbine simulator puts operating regions, sequenced actuators and gain scheduling on a machine whose nonlinearity can be measured.

KEY POINTS

- A large wind turbine holds its rotor speed with two controls, the generator's torque in light wind and the angle of the blades in strong wind.
- A 0.3 % setpoint offset lets the pitch loop begin regulating before Region 3 is declared, and the region index then selects rated torque.
- The pitch sensitivity grows by a factor of 4.9 across Region 3, and switching the schedule off buys speed regulation at more than twice the blade travel.
- A five-second wind preview supplies about two fifths of the predictive controller's gust advantage, and in turbulence the controller captures 2.6 to 5.2 % less energy than the loops.

Where else this applies. A gain scheduled on a slowly moving variable also keeps the loops of Volume 7, Chapter 9 and Volume 4, Chapter 3 at their design point, and Volume 4, Chapter 10 shares one variable among three actuators.

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

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

[4] P. Veers, K. Dykes, E. Lantz, S. Barth, C. L. Bottasso, O. Carlson, A. Clifton, J. Green, P. Green, H. Holttinen, D. Laird, V. Lehtomäki, J. K. Lundquist, J. Manwell, M. Marquis, C. Meneveau, P. Moriarty, X. Munduate, M. Muskulus, J. Naughton, L. Pao, J. Paquette, J. Peinke, A. Robertson, J. Sanz Rodrigo, A. M. Sempreviva, J. C. Smith, A. Tuohy, and R. Wiser, "Grand challenges in the science of wind energy," *Science*, vol. 366, no. 6464, p. eaau2027, 2019.

[5] W. E. Leithead, S. de la Salle, and D. Reardon, "Role and objectives of control for wind turbines," *IEE Proceedings C Generation, Transmission and Distribution*, vol. 138, no. 2, pp. 135–148, 1991.

[6] E. Muljadi and C. P. Butterfield, "Pitch-controlled variable-speed wind turbine generation," *IEEE Transactions on Industry Applications*, vol. 37, no. 1, pp. 240–246, 2001.

[7] ——, "Pitch-controlled variable-speed wind turbine generation," National Renewable Energy Laboratory, Golden, CO, Conference Paper NREL/CP-500-27143, 2000, presented at the 1999 IEEE Industry Applications Society Annual Meeting. OSTI 753780.

[8] A. D. Wright and L. J. Fingersh, "Advanced control design for wind turbines; part I: Control design, implementation, and initial tests," National Renewable Energy Laboratory, Tech. Rep. NREL/TP-500-42437, 2008.

[9] P. Veers, C. L. Bottasso, L. Manuel, J. Naughton, L. Pao, J. Paquette, A. Robertson, M. Robinson, S. Ananthan, T. Barlas *et al.*, "Grand challenges in the design, manufacture, and operation of future wind turbine systems," *Wind Energy Science*, vol. 8, no. 7, pp. 1071–1131, 2023.

[10] K. E. Johnson, L. Y. Pao, M. J. Balas, and L. J. Fingersh, "Control of variable-speed wind turbines: standard and adaptive techniques for maximizing energy capture," *IEEE Control Systems*, vol. 26, no. 3, pp. 70–81, 2006.

[11] E. A. Bossanyi, “The design of closed loop controllers for wind turbines,” *Wind Energy*, vol. 3, no. 3, pp. 149–163, 2000.

[12] F. D. Bianchi, R. J. Mantz, and H. De Battista, *Wind Turbine Control Systems*, ser. Advances in Industrial Control. Springer London, 2007.

[13] L. Y. Pao and K. E. Johnson, “Control of wind turbines,” *IEEE Control Systems*, vol. 31, no. 2, pp. 44–62, 2011.

[14] J. Jonkman, S. Butterfield, W. Musial, and G. Scott, “Definition of a 5-MW reference wind turbine for offshore system development,” National Renewable Energy Laboratory, Tech. Rep. NREL/TP-500-38060, 2009.

[15] A. Betz, “The maximum of the theoretically possible exploitation of wind by means of a wind motor,” *Wind Engineering*, vol. 37, no. 4, pp. 441–446, 2013.

[16] J. G. Slootweg, S. W. H. de Haan, H. Polinder, and W. L. Kling, “General model for representing variable speed wind turbines in power system dynamics simulations,” *IEEE Transactions on Power Systems*, vol. 18, no. 1, pp. 144–151, 2003.

[17] B. Boukhezzar and H. Siguerdidjane, “Nonlinear control of a variable-speed wind turbine using a two-mass model,” *IEEE Transactions on Energy Conversion*, vol. 26, no. 1, pp. 149–162, 2011.

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

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

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

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

[23] D. Schlipf, D. J. Schlipf, and M. Kühn, “Nonlinear model predictive control of wind turbines using LIDAR,” *Wind Energy*, vol. 16, no. 7, pp. 1107–1129, 2013.

[24] P. O. M. Scokaert and D. Q. Mayne, “Min-max feedback model predictive control for constrained linear systems,” *IEEE Transactions on Automatic Control*, vol. 43, no. 8, pp. 1136–1142, 1998.

[25] S. Lucia, T. Finkler, and S. Engell, “Multi-stage nonlinear model predictive control applied to a semi-batch polymerization reactor under uncertainty,” *Journal of Process Control*, vol. 23, no. 9, pp. 1306–1319, 2013.

[26] D. Bernardini and A. Bemporad, “Stabilizing model predictive control of stochastic constrained linear systems,” *IEEE Transactions on Automatic Control*, vol. 57, no. 6, pp. 1468–1480, 2012.

[27] 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.

[28] M. Stammler, O. Menck, Y. Guo, and J. Keller, “Wind turbine design guideline DG03: Yaw and pitch bearings,” National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5000-89161, 2024.

[29] 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.

4.11 Exercises

4.1 ○ Concept check · 2 min

In the shipped loops the torque loop’s speed setpoint is capped at $0.997\,\omega_{rated}$, while the pitch loop regulates to ω_{rated} itself. What does this 0.3 % offset do?

(a) It leaves the overspeed alarm a margin, so the alarm does not latch during the handover

(b) It makes up for the generator efficiency, so the electrical power settles at exactly its rating

(c) It keeps the logged speed error inside the scorecard band while the wind sits near rated

(d) It keeps the two integrators off one zero-error point, so pitch can start before Region 3

4.2 Concept check · 2 min

Scenario 7 is run for 120 s with the pitch gain schedule off and then, after Reset, with it on, on the same seeded wind. Off: speed IAE 0.865 rad, blade travel 255.2°, peak pitch rate 8.0 °/s. On: 3.058 rad, 108.4° and 5.2 °/s. Which conclusion does the chapter support?

(a) The fixed gain buys tighter speed with more than twice the blade travel and a pitch rate on its limit

(b) The fixed gain is the better tuning above rated, because a lower speed IAE means less fatigue for the rotor

(c) The scheduled loop is close to instability at high wind, which is why its speed IAE is three times higher

(d) The two runs cannot be compared, because the turbulence differs between the first run and the Reset

4.3 Concept check · 1 min

The advanced Region 2 option replaces the proportional-integral law of Equation (4.12) with $T_{gen} = K\omega^2$. What does that law do without?

(a) A rotor speed measurement, because the constant K already holds the optimal ratio

(b) A wind measurement and a speed setpoint, because the aerodynamics settle the rotor at λ^*

(c) Slow convergence to the optimum, because it skips the integral action of the proportional-integral law

(d) A correct constant, because the rotor finds its optimum by itself as the blades erode over time

4.4 Theory and calculation · 6 min

The default machine has $\rho = 1.225\,\mathrm{kg/m^3}$, $R = 63\,\mathrm{m}$, $\lambda^* = 7.55$, $C_{P,max} = 0.482$, $\eta = 0.944$, $b = 1.0 \times 10^4\,\mathrm{N\,m\,s}$ and $\omega_{rated} = 1.267\,109\,\mathrm{rad/s}$. The wind is a steady 6 m/s.

(a) Find the Region 2 speed setpoint of Equation (4.11) in rad/s and in rpm. *Checkpoint:* 0.719 rad/s, 6.87 rpm.

(b) Find the torque-law constant K, the aerodynamic torque $K\omega^2$ at the setpoint, and the electrical power $\eta\,(K\omega^2 - b\,\omega)\,\omega$. *Checkpoint:* $2.139 \times 10^6\,\mathrm{N\,m\,s^2}$, 1.106 MN m, 0.746 MW.

(c) Find the wind speed at which the setpoint reaches its cap $0.997\,\omega_{rated}$. *Checkpoint:* 10.54 m/s.

4.5 Theory and calculation · 8 min

The pitch loop sees the integrating plant $J_{eq}\dot{\omega} = (\partial T_{aero}/\partial\beta)\,\beta$ with $J_{eq} = 4.0 \times 10^7\,\mathrm{kg\,m^2}$. The sensitivity $|\partial T_{aero}/\partial\beta|$ is 8.915 MN m/rad at fine pitch and 46.61 MN m/rad at 25 m/s, where the steady pitch is 21.97°. The schedule knee is $\beta_k = 4.98°$ and the shipped tuning is $K_{c,\beta} = 3.7673\,\mathrm{s}$, $\tau_{I,\beta} = 2.3333\,\mathrm{s}$.

(a) Place the loop at $\zeta = 0.7$ and $\omega_n = 0.6\,\mathrm{rad/s}$ at fine pitch: find $K_{c,\beta}$ and $\tau_{I,\beta}$. *Checkpoint:* 3.769 s and 2.333 s.

(b) With the shipped gain, find the loop gain product $K_{c,\beta}|\partial T_{aero}/\partial\beta|$, in MN m s/rad, at fine pitch and at 25 m/s, then the divisor GK of Equation (4.8) at 21.97° and the scheduled product there. *Checkpoint:* 33.6 and 175.6, $GK = 5.41$, 32.4.

(c) For the unscheduled loop at 25 m/s, find the natural frequency, in rad/s and Hz, and the damping ratio. *Checkpoint:* 1.37 rad/s or 0.218 Hz, $\zeta = 1.60$.

4.6 Theory and calculation · 6 min

The rotor of the default machine has $J_{eq} = 4.0 \times 10^7\,\mathrm{kg\,m^2}$, $\omega_{rated} = 1.267\,109\,\mathrm{rad/s}$, a rating of 5.0 MW and $b = 1.0 \times 10^4\,\mathrm{N\,m\,s}$.

(a) Find the kinetic energy at rated speed and express it in seconds of rated power. *Checkpoint:* 32.1 MJ, 6.42 s.

(b) The aerodynamic damping at rated speed and rated wind is −2.65 MN m s. Find the open-loop time constant. *Checkpoint:* 15.0 s.

(c) Scenario 8 doubles the inertia. With the shipped pitch tuning ($K_{c,\beta} = 3.7673$ s, $\tau_{I,\beta} = 2.3333$ s, $k = 8.915$ MN m/rad), find the natural frequency and damping ratio of the pitch loop, and the gain that restores the design, in seconds and as the app displays it in °/rpm, which is six times the value in seconds. *Checkpoint:* 0.424 rad/s, $\zeta = 0.495$; 7.535 s or 45.2 °/rpm.

4.7 App challenge · 6 min

Load scenario 1, Steady 6 m/s (Region 2), with the default seed. Run 60 s at 5× and export the CSV file. Then press the reload button, set Target tip-speed ratio to 5.50 before Run, run 240 s at 5× and export again. Read the last row of each file.

(a) Read `generator_power_W` at 60 s of the first run. *Checkpoint:* 0.746 MW, within 0.002 MW.

(b) Read `tip_speed_ratio`, `Cp` and `generator_power_W` at 240 s of the second run. *Checkpoint:* 5.50, 0.418 and 0.649 MW, within 0.002 MW.

(c) Express the power lost by the detuned ratio as a percentage of the power at $\lambda^* = 7.55$. *Checkpoint:* 13.0 %, within 0.3 points.

4.8 App challenge · 8 min

Load scenario 7, Gain schedule off, high wind, with the default seed. Run 120 s at 10× with Pitch gain schedule Off as loaded. Read the speed IAE on the scorecard and the Peak pitch rate tile, and export the CSV file. Then press Reset, switch Pitch gain schedule to On before Run, change nothing else, run 120 s again and repeat the readings.

(a) Give the speed IAE and the peak pitch rate of each run. *Checkpoint:* Off: 0.865 rad and 8.0 °/s. On: 3.058 rad and 5.2 °/s. IAE within 0.005 rad.

(b) From each file, sum the absolute changes of `pitch_deg` from row to row to get the blade travel, and find the ratio of the two. *Checkpoint:* 255.2° and 108.4°, within 0.5°, a ratio of 2.35.

(c) Read `ctrl_gain_divisor` and `ctrl_tauI_pitch_eff_s` in the last row of the second file. *Checkpoint:* $GK = 5.27$, within 0.01, and 2.3333 s.

4.9 App challenge · 8 min

Load scenario 11, MPC or PI loops in a gust, with the default seed. The mean wind steps from 14 to 20 m/s at 30 s and back at 45 s. Run each case to 60 s at 2×, export the CSV file, skip its comment lines that start with #, and find the largest value of `overspeed_pct`. The Peak speed tile rounds to a whole percent.

(a) Run the scenario as loaded, under the MPC with the LIDAR preview. *Checkpoint:* 100.29 %, within 0.02 points.

(b) Press Reset, switch Controller to PI loops before Run, and run again. *Checkpoint:* 105.59 %, within 0.02 points.

(c) Press the reload button, set MPC wind preview to Off before Run, and run again. Then find the share of the gain of (a) over (b) that the preview supplies. *Checkpoint:* 102.39 %, within 0.02 points, and a share of about 40 %.

4.10 Application area · 3 min

The first tower fore-aft mode of the reference machine is at 0.32 Hz and the first drivetrain torsional mode near 1.7 Hz. The scheduled pitch loop is placed at $\omega_n = 0.6\,\mathrm{rad/s}$, and the unscheduled loop at the top of Region 3 moves to about 0.21 Hz.

(a) Find the separation between the tower mode and the loop, as a ratio of frequencies, for the scheduled and for the unscheduled loop. *Checkpoint:* 3.4 and 1.5.

(b) Say why the app cannot show what the smaller separation costs, and what it measures instead. *Checkpoint:* No tower or drivetrain states. It measures blade travel and time on the rate limit.

4.11 Application area · 4 min

The Region 2 law of Equation (4.11) computes its setpoint from the wind speed, and the app hands it the true wind at the rotor. The reference designs use $T_{gen} = K\omega^2$ with $K = \rho\pi R^5 C_{P,max}/(2\lambda^{*3})$ instead.

(a) Give two reasons a nacelle anemometer cannot supply the wind that the setpoint needs. *Checkpoint:* It sits in the rotor's wake, and it reads one point, not the average over the disc.

(b) Over several years the blades erode, $C_{P,max}$ falls by 4 % and λ^* rises from 7.55 to 7.8. By what factor does the correct K change, and does a $K\omega^2$ law that keeps the commissioned K run the rotor above or below its new optimal ratio? *Checkpoint:* A factor of 0.871, about 13 % lower. Below.

4.12 Review the agent · 5 min

An AI assistant reviews the chapter's scenario 11 and writes: "The MPC holds the gust peak to 100.29 % of rated speed against the PI loops' 105.59 %, so the MPC is the better controller for this turbine and should replace the loops." Find the flaw in the reasoning and say how you would check the claim in the app.

CHAPTER 5

Energy-Maximizing Control of a Heaving Buoy on a Fading Wave Forecast

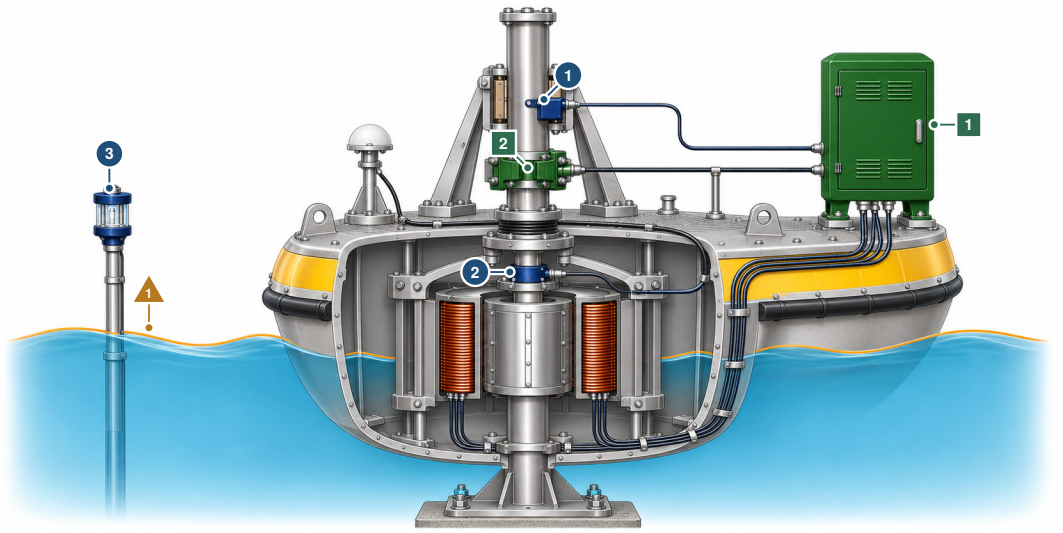

The controller moves 1 drive that sets the generator's force 2 brake that holds the buoy still

The controller reads 1 position and speed of the buoy on the spar 2 force between buoy and generator 3 wave probe upstream of the buoy

What disturbs it 1 the waves and the force they exert

Notice that the controller reads the buoy's motion, the force on its generator and the sea at a probe upstream, but never the wave force on the buoy itself, which it must estimate and forecast. Illustration, not to scale.

IN PLAIN TERMS A Buoy That Must Guess the Next Wave

The main idea. A floating buoy rides up and down on the waves. A generator between the buoy and a post fixed to the seabed turns that motion into electricity. How hard it pushes and pulls decides how much energy it catches.

The challenge. The best push depends on waves that have not arrived yet, so the controller must forecast them from the buoy's own motion. In big seas the generator can push only so hard, and the buoy must never hit the stops at the ends of its travel.

Why it matters. A good forecast of the next few seconds can triple the energy the same buoy brings home.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the fact box of Section 5.1, the history box of Section 5.2, the two try-it boxes, the fact box of Section 5.6, and the key points
Study	45 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 5.6, Section 5.7, Section 5.8 and the open challenges

The control question The force that takes the most energy from a wave depends on waves still to come. How much of that anticausal optimum can a controller reach on a forecast built from its own motion, inside a stroke and a force rating that bind in the most energetic seas?

Learning objectives

- Compute the natural period, the intrinsic impedance and the complex-conjugate tuning of a heaving buoy.
- Explain why the optimal velocity draws on the future wave force.
- Describe how a Kalman filter and an autoregressive model estimate and forecast a force no sensor reads.
- Show why the energy-maximizing quadratic program is convex and never credits more than the meter.
- Measure what a forecast horizon is worth, and set a survival threshold for a storm.

Prerequisites Sections 1.4.6 (forecast skill), 1.6.3 (predictive control, Kalman filter) and 1.6.4 (scenarios and data).

The app https://apmonitor.com/apps/wave-energy/

See also Volume 9, Chapter 7 and Volume 9, Chapter 5 for acting on a forecast or ahead of a demand. Volume 5, Chapter 11 and Volume 5, Chapter 13 for estimating what no instrument measures.

5.1 Introduction

In a sea of 1.5 m significant height and 9 s peak period, the simulated buoy, a cylinder 5 m across whose own period is 4.41 s, captures 0.767 kWh in five minutes through a generator set up as a damper tuned to that sea. A model predictive controller (MPC) handed the true wave force over the next 4 s captures 2.713 kWh from the same waves, and 0.862 kWh when it is handed nothing of the future. A forecast that the controller builds from the buoy's own motion brings 2.522 kWh at a 4 s horizon and 1.818 kWh at 1 s. The force that captures the most energy now depends on waves that have not yet reached the buoy.

The objective sets this plant apart from most of the series. A converter is paid for the energy it removes from its disturbance, so its controller optimizes an economic quantity rather than a distance from a setpoint, which is the defining trait of economic model predictive control [1]. The disturbance is also narrowband and partly predictable, so the optimum uses a prediction of it, and any economic controller whose prices or loads come from a forecast has the same shape (Section 1.4.6). The reservoir of Volume 9, Chapter 7 is one, deciding a release days ahead on an ensemble forecast whose error grows with lead time. Reviews of the field treat control as the technology that makes a converter's operation optimal [2], and they survey the optimal, predictive and MPC-like algorithms written for it [3].

✦ DID YOU KNOW? A Coastline's Worth of Waves

The wave energy available along the edge of the United States continental shelf has been estimated at 2640 TWh a year. With arrays packed at 15 MW per kilometer of coastline, 1170 TWh a year of it would be recoverable [4]. The chapter's buoy, tuned as a damper, captures 0.767 kWh in five minutes.

Three beliefs about harvesting waves are contradicted by the cases. The first is that a generator only ever takes power. Tuned as a damper and a spring, the same machine pushes 0.088 kWh back into the sea in five minutes and keeps 1.330 kWh, 1.73 times the damper's, at a peak force of 215 kN against 123 kN. The second is that a controller acts on the wave in front of it. The forecasts above are worth up to 3.1 times the energy of none, because the buoy must already be moving the right way when a crest arrives. The third is that a bigger sea should simply be harvested harder. In a 3.5 m sea a damper tuned for power drives the buoy onto its end-stops twice in five minutes, while the MPC keeps the stroke to 1.22 m and captures 7.011 kWh, and in the graded storm a controller on a real forecast keeps the machine whole only by locking it.

The parts of this problem usually reach a student separately. The hydrodynamic theory of absorption by oscillating bodies [5, 6], the constrained and predictive controllers built on it [3, 7] and short-term wave forecasting [8] live in separate papers, and none of them lets the controllers be run against one sea. This chapter reports, for one declared buoy in a browser simulator, what a damper, reactive control, latching and an energy-maximizing MPC each capture, what force each needs and what forecast each depends on, when all of them act through the same estimate of the wave force. Every method in it is established: the radiation state space [9, 10], the complex-conjugate optimum [6], latching [11] and constrained predictive control of a heaving buoy [7]. The measurements on one buoy and its seeded seas show, among other things, that most of a forecast's value lies within a quarter of the peak period and that the graded storm is passed by a survival lock rather than by harvesting more. The simulator behind them has an eleven-state model whose energy ledger closes to 1.4×10^{-6} through end-stop impacts, a quadratic program made convex by the passivity of the radiation model, a charge for reactive power that is a certified lower bound of the meter, and instruments that are the only route by which any controller learns of the sea. Each of the ten cases is credited by a completion rule of its own.

The chapter is organized as follows. Section 5.2 traces the domain from the first point absorbers to short-term wave forecasting. Section 5.3 states the buoy, the sea, the generator, the meter and the instruments, with every constant. Section 5.4 shows why the force that captures the most must look ahead and why the limits bind in the most energetic seas. Section 5.5 gives the force laws as implemented: the damper, reactive control, latching, the estimate and forecast of the wave force, the energy-maximizing quadratic program and the survival supervisor. Section 5.6 reports the ten cases as measured, the worth of a forecast, the graded storm, what the teaching bound bounds and the verification of the code. Section 5.7 describes the sensors, drives and controllers of real converters, Section 5.8 states the limitations, Section 5.9 summarizes, and the exercises follow in Section 5.10.

5.2 Point Absorbers, Phase Control and Short Forecasts

Modern research on wave energy conversion began in the early 1970s, and the history box at the end of this section tells how its first theory of point absorbers took shape.

The equation of motion that the simulator integrates descends from ship hydrodynamics. A ship-hydrodynamics report of 1962 described ship motions through an impulse response function [12], so that the radiation force becomes an inertia term, the added mass at infinite frequency, plus a convolution of the velocity with a memory kernel. States suit a controller better than a convolution does. A vertical cylinder in heave was given a state-space radiation model in 1995 [9], and identification of such parametric radiation models in the time and the frequency domain was compared for marine structures at zero speed in 2008 [10]. The added mass and damping of a vertical cylinder were computed in 1981 [13], and that matched-eigenfunction approach is the yardstick used in Section 5.6.6

for the simulator's teaching values. The irregular sea uses the spectrum from the measurements of wind-wave growth and swell decay of the Joint North Sea Wave Project (JONSWAP) [14].

Control entered as a question of phase. Latching holds the body still at its turning points, so that a body whose own period is short can keep time with longer waves, and its optimal timing was worked out in regular and irregular waves in 2006 [11]. Limits on motion and force move the optimum into the time domain. A study of a heaving sphere in 2011 optimized the machinery force under such limits, predicted the excitation force with an augmented Kalman filter on a damped harmonic oscillator, and reported a receding-horizon controller that absorbed 90 % or more of the ideal constrained optimum for a chosen range of realistic sea states [7]. Every such controller needs the future of the excitation, as the predictive controller of Chapter 4 needs a preview of the wind. The span of future information that acausal optimal control requires was examined in 2009 [15], short-term wave forecasting for real-time control in 2010 [8], and reviews in 2014 and 2017 set out energy-maximizing control and its predictive algorithms [2, 3]. A generator that loses energy in both directions changes the arithmetic of reactive control, and the effect of a non-ideal power take-off on a reactively controlled converter was quantified in 2014 [16].

> ⧖ HISTORY **From Wave Power to Point Absorbers (1974 to 2002)**
>
> Modern research on wave energy conversion began in the early 1970s, and criteria for the most a converter can absorb soon followed from analysis in the frequency domain [7]. A short paper in *Nature* in 1974 proposed the power of waves at sea as an income to live on, beside solar energy [17]. A second paper, in 1975, called a system a point absorber when its horizontal extent is much smaller than one wavelength: a damped oscillator through which energy passes from the wave to a load [18]. A theory of 1976 gave the efficiency of a damped oscillating body in one or two modes of motion and showed that complete absorption is possible in some cases [5]. A monograph of 2002 gathered the linear theory of waves, oscillating bodies and the absorption of wave energy [6]. Studies of those decades showed that control can raise a converter's absorption by up to 300 % [19].

5.3 The Buoy, the Sea and the Machine

The simulated converter is a vertical cylinder of radius $a = 2.5\,\mathrm{m}$ floating at a draft $d = 4\,\mathrm{m}$ in deep water, free to heave on a spar fixed to the seabed, as shown in Figure 5.1. A linear generator, the power take-off (PTO), has its stator in the spar and its translator carried by the buoy. A brake on the translator serves latching and the survival lock, collars on the spar stop the translator at $\pm x_{max} = \pm 1.5\,\mathrm{m}$ of heave, and a wave probe stands $L = 30\,\mathrm{m}$ upstream. Every number of the buoy, the machine and its limits is a teaching value. A real converter's mass, added mass and damping come from a radiation and diffraction solution for its own hull checked in a wave tank, and its force rating and stroke from the design of its generator and structure and the design sea states of its site. The variables are classified in Table 5.1.

5.3.1 Heave With a Radiation Memory

Linear potential flow gives the heave of a floating body as a balance of the excitation force, the radiation force, the hydrostatic spring and the machine, as shown in Equation (5.1).

$$(m + m_\infty)\,\ddot{x} = f_e(t) - f_r - k_h x - f_{pto} - f_{brk} - f_{end} \tag{5.1}$$

where $m = \rho \pi a^2 d = 80.5\,\mathrm{t}$ is the displaced mass, m_∞ the added mass at infinite frequency, $k_h = \rho g A_{wp} = 197.4\,\mathrm{kN/m}$ the stiffness of the waterplane $A_{wp} = \pi a^2$, f_e the excitation force, f_r the radiation force and f_{pto}, f_{brk} and f_{end} the generator, brake and end-stop forces, positive downward on the buoy. The radiation force is a convolution $\int_0^t K(t-\tau)\dot{x}(\tau)\,d\tau$ in the Cummins form, which the simulator replaces by two states, as shown in Equation (5.2).

$$\dot{z}_1 = z_2, \quad \dot{z}_2 = -a_0 z_1 - a_1 z_2 + \dot{x}, \quad f_r = b_1 z_2 \tag{5.2}$$

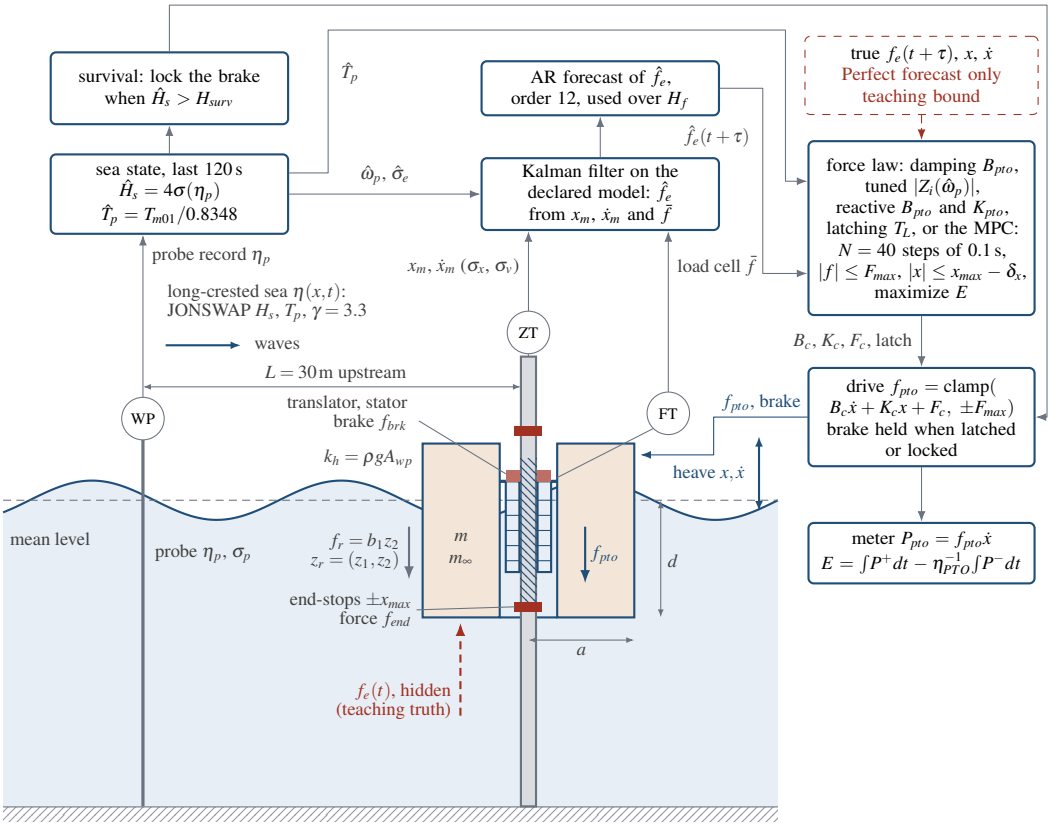

Figure 5.1 The heaving buoy in its sea and in its control loop. The hull is cut through its well to show the spar, the generator's stator and translator, the brake and the end-stop collars. The probe's record feeds the sea-state estimate, the position, velocity and load cell feed the Kalman estimate of the excitation force $\hat{f}_e$, and the true future excitation reaches only the labeled perfect-forecast teaching bound (dashed box and arrow). The estimators sit above the sea, and the force law, the drive and the meter stand beside the buoy.

The transfer function from velocity to force is $\hat{K}(s) = b_1 s/(s^2 + a_1 s + a_0)$, with a zero at the origin, relative degree one and stable poles, the structure a parametric radiation model must keep [10]. Its damping and added mass at frequency ω are given by Equation (5.3).

$$B(\omega) = \operatorname{Re}\hat{K}(i\omega) = \frac{b_1 a_1 \omega^2}{(a_0 - \omega^2)^2 + a_1^2 \omega^2}$$
$$a(\omega) = m_\infty + \operatorname{Im}\hat{K}(i\omega)/\omega \tag{5.3}$$

$B(\omega) \ge 0$ at every frequency because $b_1 a_1 > 0$, so the model is passive, and exactly so in time. The work the buoy does on the water splits into a stored part and a dissipated part, $\int f_r \dot{x}\, dt = \Delta[\frac{1}{2} b_1 (z_2^2 + a_0 z_1^2)] + b_1 a_1 \int z_2^2\, dt$. That identity is what the MPC of Section 5.5 is built on. Both radiation states start at zero, and the buoy starts at rest at $x = 0$ except in case 1, which releases it from $x(0) = 0.5$ m.

5.3.2 A Declared Kernel and Its Two-State Fit

The radiation coefficients are not taken from a boundary-element table. They are fitted to a declared kernel built from the Haskind relation for an axisymmetric body heaving in deep water,

Table 5.1 Variable Classification

Role	Symbol	Meaning	Range or unit
Objective	E	metered energy, maximized	kWh
Constrained	x	heave from rest	$\lvert x\rvert \le 1.5$ m
	f_{pto}	generator force	$\lvert f_{pto}\rvert \le 400$ kN
MV	B_c, K_c, F_c	drive's law and force	per strategy
	brake	latch hold, lock	$T_L \le 6$ s
	F_{man}	force by hand (Manual)	± 400 kN
DV	f_e	excitation force	hidden
	H_s, T_p	sea state	0.1 to 5 m, 4 to 14 s
	events	sea up, rogue group, PTO fault, probe outage	
States	$x, \dot{x}$	heave, velocity	m, m/s
	z_1, z_2	radiation states	m·s, m
	7 integrals	$E^{\pm}$, W_e, E_{rad}, E_{brk}, W_{end}, machine impulse	J, N·s
Measured	$x_m, \dot{x}_m$	encoder, velocity	σ_x, σ_v
	$\bar{f}$	load cell, period mean	exact
	η_p	probe at $-L$	σ_p
Estimated	$\hat{f}_e$	excitation force	kN
	$\hat{H}_s, \hat{T}_p$	sea state from the probe	m, s

$B = \omega k \Gamma^2/(2\rho g^2)$ with $k = \omega^2/g$, applied to the Froude–Krylov force on the flat bottom, which gives the excitation coefficient and the kernel of Equation (5.4).

$$\Gamma(\omega) = \rho g A_{wp}\, e^{-kd}, \qquad B_K(\omega) = \frac{\rho A_{wp}^2}{2g}\, \omega^3 e^{-2d\omega^2/g} \tag{5.4}$$

where Γ is the excitation force per meter of wave amplitude, in phase with the elevation at the buoy. Because B_K and Γ come from one relation, the largest power the model can take from a regular wave, $|\Gamma A|^2/(8B)$, equals the point-absorber limit J/k, the wave power in a crest width of one wavelength over 2π [5, 6, 18]. The three coefficients minimize $\sum_j \ln^2[B(\omega_j)/B_K(\omega_j)]$ on a 0.01 rad/s grid from 0.5 to 2.0 rad/s (wave periods from 3.1 to 12.6 s), which gives $b_1 = 15\,908.67$ N/m, $a_1 = 1.408\,497\,\mathrm{s}^{-1}$ and $a_0 = 1.699\,333\,\mathrm{s}^{-2}$, within 9.8 % of the kernel over the band and 2.8 % rms. The kernel's own initial value $K(0) = (2/\pi)\int B_K\, d\omega$ is 9640 N/m against the model's b_1, so the fit keeps the damping and gives up the shape of the memory at short times. The added mass at infinite frequency is a teaching value, a hemisphere's $\rho\pi a^3/3 = 16.8$ t, not a computed one.

5.3.3 The Sea and the Wave Force

The sea is long-crested and travels toward the buoy along $+x$. Its spectrum is the JONSWAP form with peak enhancement $\gamma = 3.3$, as shown in Equation (5.5).

$$\begin{aligned} S(\omega) &= A_\gamma \tfrac{5}{16} H_s^2 \omega_p^4\, \omega^{-5} \exp\!\big[-\tfrac{5}{4}(\omega_p/\omega)^4\big]\, \gamma^r \\ r &= \exp\!\big[-(\omega-\omega_p)^2/(2\sigma^2\omega_p^2)\big] \end{aligned} \tag{5.5}$$

with $\omega_p = 2\pi/T_p$, $\sigma = 0.07$ below the peak and 0.09 above it, and $A_\gamma = 1 - 0.287 \ln\gamma$. One component is drawn in each of 120 equal bins from 0.30 to 3.00 rad/s, its frequency uniform within the bin and its phase uniform on $[0, 2\pi)$, both from the case's seed, with amplitude $a_k = \sqrt{2S(\omega_k)\Delta\omega}$. The elevation, the excitation and the incident power per meter of crest follow in Equation (5.6).

$$\begin{aligned} \eta(x,t) &= \textstyle\sum_k a_k \cos(\omega_k t - k_k x + \varphi_k) \\ f_e(t) &= \textstyle\sum_k a_k\, \Gamma(\omega_k) \cos(\omega_k t + \varphi_k) \\ J(t) &= \tfrac{1}{2}\rho g \textstyle\sum_k a_k^2\, c_{g,k}, \qquad c_{g,k} = g/(2\omega_k) \end{aligned} \tag{5.6}$$

A change of sea state keeps the frequencies and phases and moves every amplitude linearly to the new spectrum over 60 s. A rogue group adds a narrow-band packet focused on the buoy 30 s after it is fired, with a crest of $1.2H_s$ at ω_p and an envelope of half-width $1.2T_p$, seen first at the probe.

5.3.4 Generator, Brake and End-Stops

Every 0.1 s the controller hands the drive an impedance law and a force, which the drive's own current loop applies continuously, clipped at its rating, as shown in Equation (5.7).

$$\begin{aligned} f_{pto} &= \mathrm{clamp}\big(B_c\dot{x} + K_c x + F_c,\ \pm F_{lim}\big) \\ f_{brk} &= \mathrm{clamp}\big(B_{lt}\dot{x},\ \pm F_{lock}\big) \text{ when engaged} \\ f_{end} &= \mathrm{sgn}(x)\, k_{end}(|x| - x_{max}) + c_{end}\dot{x} \end{aligned} \tag{5.7}$$

where f_{end} acts only while $|x| > x_{max}$, $F_{lim} = F_{max} = 400\,\mathrm{kN}$, halved by a PTO fault, $B_{lt} = 5\,\mathrm{MN\,s/m}$, $F_{lock} = 1.5\,\mathrm{MN}$, $k_{end} = 20\,\mathrm{MN/m}$ and $c_{end} = 570\,\mathrm{kN\,s/m}$. A damper sets B_c, a reactive law B_c and K_c, and the MPC F_c. The brake engages at the first velocity zero crossing after it is armed (and at least 0.2 s after its last release) and releases when its hold is over. Its work heats the brake and is never credited.

5.3.5 The Meter and the Energy Ledger

The meter counts what reaches the grid. Absorbed power $P = f_{pto}\dot{x} > 0$ is credited in full and power pushed back into the sea is charged at $1/\eta_{PTO}$ with $\eta_{PTO} = 0.8$, as shown in Equation (5.8).

$$E = \int_0^t \max(P,0)\,dt - \frac{1}{\eta_{PTO}} \int_0^t \max(-P,0)\,dt \tag{5.8}$$

The capture width ratio divides E by the incident energy across the diameter, $\int J\,2a\,dt$. Seven integrals ride in the state so that the integrator treats them as accurately as the motion: the two parts of the meter $E^{\pm}$, the excitation's work $W_e = \int f_e\dot{x}\,dt$, the radiated energy $E_{rad} = b_1 a_1 \int z_2^2\,dt$, the brake's and the end-stops' work and the impulse of every machine force. They close the ledger $W_e = (E^+ - E^-) + E_{rad} + \Delta E_s + E_{brk} + W_{end}$, with the stored energy $E_s = \frac{1}{2}(m+m_\infty)\dot{x}^2 + \frac{1}{2}k_h x^2 + \frac{1}{2}b_1(z_2^2 + a_0 z_1^2)$.

5.3.6 What the Controllers Measure

The controllers read a packet built once per control period and nothing else, as shown in Equation (5.9).

$$\begin{aligned} x_m &= x + v_x, \qquad \dot{x}_m = \dot{x} + v_v \\ \eta_p &= \eta(-L,t) + v_p \\ \bar{f}_j &= \tfrac{1}{\Delta t}\textstyle\int_{t_{j-1}}^{t_j} (f_{pto} + f_{brk} + f_{end})\,dt \end{aligned} \tag{5.9}$$

where v_x, v_v and v_p are seeded white noise of standard deviation $\sigma_x = 5\,\mathrm{mm}$, $\sigma_v = 0.01\,\mathrm{m/s}$ and $\sigma_p = 0.02\,\mathrm{m}$, and $\bar{f}_j$ is the mean force of a load cell between the buoy and the translator, on which the end-stop collars also act. The packet also carries the rating the drive can deliver and the brake's state. The probe's record of the 120 s before $t = 0$ is handed over at a reset, so the sea-state estimate exists from the start. The excitation force, the elevation at the buoy, the true state and the future are kept from the controllers, exported only in `truth_` columns and drawn dashed on the views. Only the perfect-forecast setting of the MPC adds the true state and the true future excitation to the packet, as a labeled teaching bound.

5.3.7 Constants, Initial State and Integration

Table 5.2 gives every parameter. The engine of Section 1.6.1 integrates Equations (5.1) to (5.8) with fourth-order Runge–Kutta at $h = 10\,\mathrm{ms}$. The stiffest dynamics are the end-stop spring, $\sqrt{k_{end}/(m+m_\infty)} = 14\,\mathrm{rad/s}$ or 44 steps a period, and the engaged brake, $B_{lt}/(m+m_\infty) = 51\,\mathrm{s^{-1}}$ or

Table 5.2 Parameters, Defaults and Their Standing

Symbol	Meaning	Value	Unit, standing
ρ	sea water density	1025	kg/m^3
a, d	radius, draft	2.5, 4.0	m, T
m	mass, $\rho\pi a^2 d$	80.5	t, A
m_∞	added mass, $\omega \to \infty$	16.8	t, T, A
k_h	hydrostatic stiffness	197.4	kN/m, A
b_1	radiation gain	15 908.67	N/m, F, A
a_1	radiation coefficient	1.408497	s^{-1}, F, A
a_0	radiation coefficient	1.699333	s^{-2}, F, A
F_{max}	PTO force rating	400	kN, T, A
x_{max}	end-stops	$\pm$1.5	m, T, A
η_{PTO}	reactive power charge	0.8	T, A
F_{lock}	brake rating	1500	kN, T, A
B_{lt}	brake damping	5000	kN·s/m, T, A
k_{end}	end-stop spring	20	MN/m, H
c_{end}	end-stop damping	570	kN·s/m, H
σ_x, σ_v	encoder noise	5, 0.01	mm, m/s, A
σ_p	probe noise	0.02	m, A
L	probe upstream	30	m, T
γ, N_W	JONSWAP, components	3.3, 120	F
Δt, N	control period, steps	0.1, 40	s
δ_x	MPC stroke margin	0.3	m, C
α	terminal value	0.5	C
p, T_w	AR order, window	12, 60	s, C
H_f	forecast horizon	4	s, C
H_{surv}	survival threshold	0 (off)	m, C

T teaching value, F fitted, A adjustable in the Model dialog, H hidden from the student, C a control on the faceplates. Case settings differ (Table 5.3).

$h\lambda = 0.51$, inside the method's stability region. The controller runs at $\Delta t = 0.1$ s and the historian logs at 10 Hz. Contacts, brake engagements, sea-state changes, the onset of a force cut and the probe's return are events, each processed a single time between two integration steps and marked in the log, and the derivative function reads only the state, the commands the drive holds and the sea's record, so a Reset replays a run bit for bit.

5.3.8 Assumptions

The model assumes linear potential flow in deep water and heave alone, so pitch, surge, viscous drag, moorings and the nonlinear Froude–Krylov force at large motion are absent. These break down for motions comparable to the draft and in steep or breaking seas. The excitation is in phase with the elevation at the buoy, with no diffraction correction. The sea is long-crested and stationary within each state, and its 120 components never repeat within a run. The meter charges only reactive power, so a real generator's losses on absorbed power are missing, and the plant's end-stop is a stiff spring with damping rather than a structural model of an impact.

5.4 Why the Best Force Looks Ahead

5.4.1 A Buoy Tuned to Short Waves

Released from 0.5 m in a 0.2 m sea, the buoy rings at its natural period. The characteristic quartic of Equations (5.1) and (5.2), $(Ms^2 + k_h)(s^2 + a_1 s + a_0) + b_1 s^2 = 0$ with $M = m + m_\infty$, has a dominant pair at $-0.0612 \pm 1.4353i\,\mathrm{s}^{-1}$, a damped period of 4.3777 s against $2\pi\sqrt{M/k_h} = 4.4103$ s, and a damping ratio of 0.043, so each swing is 0.76 of the one before. The seas of the cases have peak periods of 8 to 11 s, about twice the buoy's, and there the intrinsic impedance of Equation (5.10) is

dominated by its reactance.

$$Z_i(\omega) = B(\omega) + iX_i(\omega)$$
$$X_i(\omega) = \omega\big(m + a(\omega)\big) - k_h/\omega \tag{5.10}$$

At 9 s the reactance is −209 kN s/m against a radiation resistance of 4.5 kN s/m, so most of the wave's force goes into accelerating the buoy and the water it carries, not into moving it against a generator.

> **TRY IT IN THE APP** **1 · Free Float: The Natural Period (Start Here)**
>
> Load scenario 1 and press Run. Watch the Heave trend in the first 20 s. The buoy, let go 0.5 m above its rest position, swings at its own period, each swing a little smaller than the one before. The first three swings take 4.38 s each, about half the period of the 8 s waves around it.

5.4.2 The Complex Conjugate and Its Kernel in Time

A linear PTO law $f_{pto} = B_{pto}\dot{x} + K_{pto}x$ in a regular wave of excitation amplitude $F = \Gamma A$ gives the velocity amplitude and mean power of Equation (5.11).

$$\hat{u} = \frac{F}{\left|Z_i + B_{pto} - iK_{pto}/\omega\right|}, \qquad \bar{P} = \tfrac{1}{2}B_{pto}\hat{u}^2 \tag{5.11}$$

A damper alone takes the most at $B_{pto} = |Z_i(\omega)|$. A damper and a spring together take the most at the complex conjugate, $B_{pto} = B(\omega)$ and $K_{pto} = \omega X_i(\omega)$, where the spring cancels the reactance and the power reaches $|F|^2/(8B)$ [6]. Written as an operator on the excitation, the optimum velocity is $\hat{u}_{opt}(\omega) = \hat{F}_e(\omega)/[2B(\omega)]$. Because $1/[2B(\omega)]$ is real and even in ω, its impulse response is even in time, so the optimal velocity at t draws on the excitation at $t + \tau$ as much as at $t - \tau$. The optimum is anticausal, and the question every predictive controller for waves must answer is how much of that future it can see [15]. The reactive spring also costs force. At the tuning of case 3 the peak force rises from 123 kN to 215 kN, and the generator must run as a motor for part of each cycle.

Pause. Try Exercise 5.1 before reading on.

5.4.3 Limits That Bind When the Energy Is Largest

Tuned to the complex conjugate at 9 s, $B_{pto} = 4.5\,\text{kN s/m}$ and $K_{pto} = -146\,\text{kN/m}$, the buoy's predicted heave in case 2's 1.5 m sea has three standard deviations of 17.7 m, against a stroke of ±1.5 m. The force rating binds in the same seas, because reactive control's force grows with the motion it creates. Constraints therefore decide the controller in exactly the seas that carry the most energy, and the time-domain problem replaces the frequency-domain optimum [7]. A hit on an end-stop is damage rather than a lost kilowatt-hour, which is why the challenge counts contacts and not only energy.

5.5 The Force Laws as Implemented

The strategies act through Equation (5.7). Free float sets $B_c = K_c = F_c = 0$. Manual mode applies $F_c = F_{man}$ from a hand station, which in Auto follows the force the strategy asks for, so a switch to Manual holds it. Every law designs with the declared model parameters, the plant's own.

5.5.1 Damping, Reactive Tuning and Latching

Damping applies $B_c = B_{pto}$ from its dial. Tuned damping applies $B_c = |Z_i(\hat{\omega}_p)|$ at the probe's estimated peak frequency $\hat{\omega}_p = 2\pi/\hat{T}_p$, re-tuned whenever the estimate moves by a twentieth of a second or of a meter. Reactive control applies B_{pto} and K_{pto} from their dials. Its tuning helper sets $K_{pto} = \hat{\omega}_p X_i(\hat{\omega}_p)$ and the smallest B_{pto} above $B(\hat{\omega}_p)$, rounded to 5 kN/m and 5 kN s/m, for which three standard deviations of the predicted stroke stay inside $0.75\,x_{max}$, with $\sigma_x^2 = \sum_j S(\omega_j)\Delta\omega\,\Gamma^2(\omega_j)/[\omega_j^2|Z_i + B_{pto} - iK_{pto}/\omega_j|^2]$ on 120 frequencies from 0.3 to 3 rad/s, found by bisection. Latching damps at B_{pto}

while the buoy is free and arms the brake with the hold T_L. Its tuning sets $T_L = \max[0, (\hat{T}_{m01} - T_n)/2]$, at most 6 s, from the spectral mean period $\hat{T}_{m01} = 0.8348\,\hat{T}_p$ of a JONSWAP sea with $\gamma = 3.3$ and the natural period $T_n = 2\pi\sqrt{(m+m_\infty)/k_h}$, with $B_{pto} = |Z_i(\hat{\omega}_p)|/2$. In an irregular sea the peak period would hold the buoy too long. On case 4, $T_L = 2.75$ s from $\hat{T}_p$ captures 0.855 kWh against 1.236 kWh with 1.95 s from $\hat{T}_{m01}$.

5.5.2 Estimating and Forecasting a Force No Sensor Reads

The excitation force is what the waves would exert on a buoy held still, and nothing measures it. It is a variable a controller must infer, like the core temperature of the kiln in Volume 5, Chapter 11. A steady-state Kalman filter of the kind described in Section 1.6.3 carries it as a state of a lightly damped resonator beside the buoy's four states, as shown in Equation (5.12).

$$\begin{aligned} \ddot{e} + 2\zeta\hat{\omega}_p\dot{e} + \hat{\omega}_p^2 e = w(t), \qquad q = 4\zeta\hat{\omega}_p^3\hat{\sigma}_e^2 \\ \hat{\sigma}_e = \max\left[20\,\text{kN},\ (\hat{H}_s/4)\,\Gamma(\hat{\omega}_p)\right] \end{aligned} \tag{5.12}$$

with $\zeta = 0.2$ and white noise w of intensity q, so that the stationary variance of e is $\hat{\sigma}_e^2$, the approach of the 2011 heaving-sphere study [7]. The six states $[x, \dot{x}, z_1, z_2, e, \dot{e}]$ evolve by Equations (5.1) and (5.2) with $f_e = e$ and the load cell's $\bar{f}$ as a known input, discretized exactly at 0.1 s together with the process covariance by Van Loan's construction, with a small process noise on the velocity of intensity $[0.02\hat{\sigma}_e/(m+m_\infty)]^2$. The filter measures x_m and $\dot{x}_m$ with $R = \text{diag}(\sigma_x^2, \sigma_v^2)$, and its steady gain is redesigned when $\hat{\omega}_p$ moves by more than 3 % or $\hat{\sigma}_e$ by more than 20 %, at most once every 10 s. Because the end-stop collars act on the translator, the load cell reports their force too, and a contact does not appear to the filter as a burst of wave force.

The sea state comes from the probe's last 120 s. The estimate $\hat{H}_s = 4\sigma(\eta_p)$ follows from the record's variance and $\hat{T}_p = \hat{T}_{m01}/0.8348$ from its spectral mean period $2\pi m_0/m_1$, with the moments summed over a discrete Fourier transform from 0.3 to 3 rad/s at the record's own resolution, every 1 s. While the probe is out, the last estimate holds.

An autoregressive model of order $p = 12$ on samples of $\hat{f}_e$ every 0.5 s, $\hat{e}_k = \sum_{i=1}^{p} \varphi_i \hat{e}_{k-i}$, is fitted by least squares with a ridge of 10^{-6} of the largest diagonal entry to the last 60 s of the estimate, once a second. If the largest root of the fitted polynomial has a modulus $\rho > 0.98$, every coefficient is scaled, $\varphi_i \leftarrow \varphi_i(0.98/\rho)^i$, so an iterated forecast can only decay. The forecast is iterated from the latest p samples, read on the 0.1 s grid by linear interpolation, and set to zero beyond the forecast horizon H_f. Its skill, $1 - \text{MSE}/\overline{\hat{f}_e^2}$ at each lead, is scored against the estimate that later arrives, with a forgetting time of 60 s, which is all the controller itself can check. Section 1.4.6 shows with a first-order model, Equation (1.8), why such a skill falls with the lead.

5.5.3 An Energy-Maximizing Quadratic Program

The MPC chooses the force f_k, held over each 0.1 s step, for $N = 40$ steps (4 s), in 30 move blocks (20 single steps, then pairs). Its model is the plant's heave and radiation, Equations (5.1) and (5.2) without the brake and end-stops, discretized by an exact zero-order hold with the preview w_k of the excitation, the forecast or on the perfect setting the truth at each step's midpoint. The energy the PTO absorbs, $\sum_k f_k \Delta x_k$, is bilinear in force and motion, but the energy balance of the model rewrites it as the excitation's work less the radiated energy and the rise in stored energy. The controller solves the problem of Equation (5.13) every 0.1 s.

$$\begin{aligned} \max_{f,s}\ & \textstyle\sum_k w_k \Delta x_k - b_1 a_1 \Delta t \sum_k c_k z_{2,k}^2 \\ & - c\,\Delta t \textstyle\sum_k \left(f_k - \hat{B}\,\Delta x_k/\Delta t\right)^2 \\ & - (1-\alpha)E_s(N) - \rho_2 s^2 - \rho_1 s \\ \text{s.t. } & |f_k| \le F_{lim}, \quad |x_k| \le x_{max} - \delta_x + s \end{aligned} \tag{5.13}$$

where the sums over the force and the motion run over the N steps, $\Delta x_k = x_{k+1} - x_k$, the slack $s \geq 0$, $c_k = 1$ for $k = 1$ to $N-1$ and $c_N = \frac{1}{2}$ (the trapezoid on the radiated energy), $E_s(N)$ is the stored energy at the horizon's end, valued at the fraction α, and $c = (1/\eta_{PTO} - 1)/(4\hat{B})$ with $\hat{B} = |Z_i(\hat{\omega}_p)|$. Every quadratic weight is non-negative because $b_1 a_1 > 0$ and $\alpha \leq 1$, so the Hessian is positive semidefinite by construction, and the charge on reactive power makes it definite. The charge is a certified lower bound of the meter. For any f, v and $\hat{B} > 0$, $(f - \hat{B}v)^2 + 4\hat{B}fv = (f + \hat{B}v)^2 \geq 0$, so $(f - \hat{B}v)^2/(4\hat{B}) \geq \max(-fv, 0)$, and the objective never exceeds what the meter credits a plan evaluated step by step. The stroke band is soft, with one slack s priced at $\rho_2 = 10^{10}\,\mathrm{J/m^2}$ and $\rho_1 = 10^8\,\mathrm{J/m}$, far above any energy on offer, and the force bound is hard at the rating the drive reports. The quadratic program is solved by the framework's condensed linear MPC on its dense dual active-set solver (Section 1.6.3), capped at 300 iterations. A failed solve falls back to the passive optimum $\hat{B}$ for that period and is counted on the faceplate.

Pause. Try Exercise 5.3 before reading on.

5.5.4 Survival, the Teaching Bound and Three Models

The survival supervisor works on the probe's estimate alone. When $\hat{H}_s$ exceeds H_{surv}, the brake locks the buoy at its next velocity zero crossing while the strategy keeps control until the brake holds, after which the PTO idles. It unlocks when $\hat{H}_s$ has stayed 0.5 m below the threshold for 60 s, and a contact during survival stops the run. The perfect-forecast teaching bound of each case is the MPC of Equation (5.13) on the true state and the true future, at $H_f = N\Delta t = 4\,\mathrm{s}$, a margin of 0.01 m and $\alpha = 1$, with no survival lock. Three models are kept apart. The plant is the whole of Section 5.3. The filter and the MPC carry the plant's declared heave and radiation, with the excitation a resonator in one and a forecast preview in the other. The tunings use the frequency-domain impedance of Equation (5.10) at one estimated frequency.

5.6 Ten Cases on One Buoy

The ten cases are listed in Table 5.3. Every number below was measured on the app's own engine, headless, at the case's seed (20261030, and 3 for the challenge), over the case's own run, with the settings named entered before Run. The measured outcomes are collected in Table 5.4, and Figure 5.2 draws them as a chart. Each case carries a briefing and a completion rule (Section 1.6.4), and the rule credits the decision the case asks for rather than a state the plant happens to reach. Case 5, for example, is earned by finishing a run with the perfect forecast at a horizon of 0.5 s or less, then setting it to 3 s or more and capturing at least 2.0 kWh in a second run. A scripted walk earns every rule, and a run on the shipped settings earns none.

5.6.1 Resonance, Damping and a Two-Way Machine

Case 1 releases the buoy 0.5 m up in a 0.2 m sea at 8 s. Its first three swings take 4.38 s each, within 1 % of the 4.41 s formula. Over 120 to 180 s the heave's rms is 1.5 times the sea's at 8 s and 2.37 times with the peak period set to 4.5 s, near resonance. In case 2's 9 s sea the shipped damper of 40 kN s/m captures 0.474 kWh in five minutes (5.7 kW, a capture width ratio of 13 %). Tuned to the sea, 215 kN s/m, it captures 0.767 kWh (9.2 kW, 21 %), and swept by hand the best damping, 175 kN s/m, captures 0.785 kWh, so the optimum is flat. The tuned damper resists hard and lets the buoy move less, from 1.31 m to 0.79 m of stroke, and its velocity runs about 0.7 s ahead of the excitation. Reactive control in case 3, $K_{pto} = -145\,\mathrm{kN/m}$ and $B_{pto} = 210\,\mathrm{kN\,s/m}$, captures 1.330 kWh while the power dips to −22.8 kW in the worst wave. Latching in case 4's 10 s sea captures 1.236 kWh against the tuned damper's 0.747 kWh, with the buoy held 43 % of the time. A controller that spends almost half of each wave doing nothing captures 1.65 times as much, because the buoy then moves while the wave force is with it.

Pause. Try Exercise 5.2 before reading on.

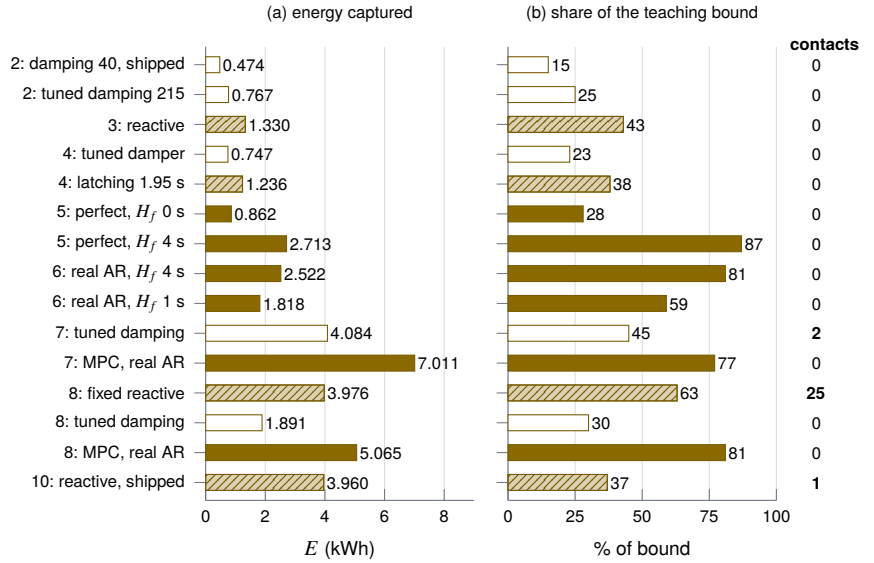

Figure 5.2 Table 5.4 as a chart: (a) the energy each controller captured and (b) its share of the case's perfect-forecast teaching bound, with the end-stop contacts of each run at the right. White bars are damping, hatched bars reactive control and latching, solid bars the MPC. Notice that in cases 7 and 8 the MPC captures the most and touches no stop, while the fixed reactive law of case 8 comes second while striking the stops 25 times.

Table 5.3 Scenarios

Case	Sea H_s, T_p	Shipped settings and lesson	Run
1	0.2 m, 8 s	free float from 0.5 m, natural period	180 s
2	1.5 m, 9 s	damping 40 kN·s/m, tune to the sea	300 s
3	1.5 m, 9 s	tuned damper, then reactive	300 s
4	1.5 m, 10 s	tuned damper, then latching	300 s
5	1.5 m, 9 s	MPC, Perfect forecast, H_f 0 s	300 s
6	1.5 m, 9 s	MPC, Real AR forecast, H_f 4 s	300 s
7	3.5 m, 9 s	tuned damping, then MPC	300 s
8	1.5 m, 8 s	reactive B 165, K −135, then 2.5 m, 11 s from 120 s	420 s
9	three seas	challenge: damping 170 kN·s/m	1200 s
10	4.0 m, 10 s	reactive B 630, K −155 (infeasible)	300 s

B in kN·s/m, K in kN/m. The challenge's seas are 1.5 m at 8 s, 2.5 m at 9 s from 340 s and 4 m at 10 s from 740 s.

5.6.2 What a Forecast Is Worth

Cases 5 and 6 run the MPC in the same 9 s sea, with the true future and with the autoregressive forecast of the controller's own estimate, as shown in Figure 5.3. With no forecast the plan always ends with the buoy at rest, and the MPC captures 0.862 kWh, 12 % more than the tuned damper's 0.767 kWh. A perfect forecast adds energy at every horizon: 1.790 kWh at 0.5 s, 1.847 at 1 s, 2.213 at 1.5 s, 2.500 at 2 s, 2.673 at 3 s and 2.713 at 4 s, with a peak force of 294 kN and 0.509 kWh pushed back into the sea. The real forecast tracks it closely up to 1 s (1.818 against 1.847 kWh) and falls behind as the lead grows, to 2.522 against 2.713 kWh at 4 s, 93 % of the perfect forecast's energy.

Table 5.4 Measured Results of Cases 2 to 8 and 10

Case	Controller	E (kWh)	Bound	Peak (kN)	Contacts
2	damping 40, shipped	0.474	15 %	44	0
2	tuned damping 215	0.767	25 %	123	0
3	reactive K −145, B 210	1.330	43 %	215	0
4	tuned damper	0.747	23 %	136	0
4	latching 1.95 s, B 120	1.236	38 %	166	0
5	Perfect, H_f 0 s	0.862	28 %	138	0
5	Perfect, H_f 4 s	2.713	87 %	294	0
6	Real AR, H_f 4 s	2.522	81 %	312	0
6	Real AR, H_f 1 s	1.818	59 %	237	0
7	tuned damping	4.084	45 %	291	2
7	MPC, Real AR	7.011	77 %	400	0
8	fixed reactive	3.976	63 %	313	25
8	tuned damping	1.891	30 %	238	0
8	MPC, Real AR	5.065	81 %	400	0
10	reactive, shipped	3.960	37 %	400[a]	1

Bound: the share of the case's perfect-forecast teaching bound (3.1032 kWh on cases 2, 3, 5 and 6, 3.2601 on case 4, 9.0726 on case 7, 6.2665 on case 8, 10.7592 on case 10). Peak: the largest force applied. [a] asked 897 kN before the rating cut it.

The forecast's own skill explains the gap. Averaged uniformly over 237 forecasts made every second from 60 to 296 s, it is 0.96, 0.85, 0.82 and 0.78 at leads of 1, 2, 3 and 4 s against the estimate that later arrived, and 0.95, 0.85, 0.83 and 0.78 against the true excitation, so the skill the controller can check is the skill it has. The forecast's worth arrives early. Half of the gain from no forecast to a 4 s perfect forecast comes from the first 0.5 s, and 88.5 % of it from the first 2 s, less than a quarter of the 9 s peak period.

> **TRY IT IN THE APP** **5 · The Anticausal Optimum: A Perfect Forecast Against None**
>
> Load scenario 5 and run the five minutes. The predictive controller has a perfect forecast but a Forecast horizon of 0 s, and it captures 0.862 kWh. Set the Forecast horizon to 4 s, press Reset and run again. The same waves now give 2.713 kWh, and the peak force rises from 138 kN to 294 kN.

The estimate behind the forecast is accurate. Sampled at the control instants after the first 30 s, its error is 6.0 % rms of the true force on case 6 and 4.0 % on case 7 under the MPC. After an end-stop contact, the largest error within 3 s is 40.0 kN on case 7's tuned damper and at most 51.6 kN in the challenge, because the load cell reports the collar's force.

5.6.3 Strokes, Ratings and Changing Seas

In case 7's 3.5 m sea the tuned damper reaches 1.54 m of stroke and hits the end-stops twice for 4.084 kWh. The MPC on the real forecast keeps the stroke to 1.22 m with no contact, at its 400 kN rating, and captures 7.011 kWh, and on the perfect forecast 7.506 kWh. A damper can only resist a motion that is already there, while the MPC sees the big wave in its forecast and starts braking the swing before it. Its plan carried slack three times, for 1.6 s in all, when even the full rating could not hold the plan 0.3 m inside the stops, and the margin absorbed it. In case 8 the sea grows from 1.5 m at 8 s to 2.5 m at 11 s at 120 s. The reactive settings tuned for the first sea hit the stops 25 times from 198 s on. Tuned damping follows the probe's estimate of the peak period, from 8.3 s at 120 s to 11.5 s at 300 s, raising its damping from 182 to 305 kN s/m, and never touches a stop, for 1.891 kWh. The MPC captures 5.065 kWh with no contact. With the probe out from 60 s, the estimate holds at 8.1 s at 240 s and the tuned damper still captures 1.869 kWh, because a passive damper is forgiving where a reactive tuning is not.

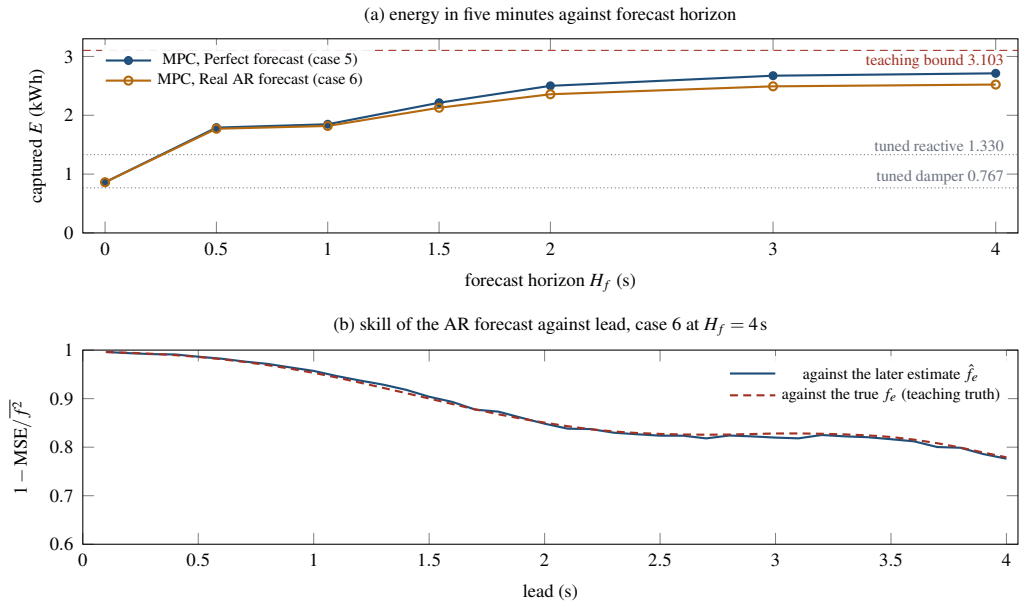

Figure 5.3 The worth of a forecast in the 1.5 m, 9 s sea of cases 5 and 6. (a) Energy captured in five minutes against the forecast horizon on the perfect forecast (solid line, filled circles) and on the real forecast (solid line, open circles), beside the tuned damper and tuned reactive control (dotted lines) and the case's teaching bound (dashed line). (b) The AR forecast's skill against lead, scored against the later estimate (solid line) and against the true excitation (dashed line).

Case 10 is the deliberately infeasible variant. Reactive control tuned for its 4 m, 10 s sea, $K_{pto} = -155\,\mathrm{kN/m}$ and $B_{pto} = 630\,\mathrm{kN\,s/m}$, asks the drive for 897 kN and still touches a stop once. Swept from 100 to 1000 kN s/m at that stiffness, the asked force runs from 395 kN with 132 contacts to 1007 kN with none, so every damping breaks one limit or the other, as does every softening stiffness with damping alone. The claim is scoped to the softening, reactive tuning. A stiffening spring of +50 kN/m or more, which detunes the buoy from these waves, keeps a linear law inside both limits at 3.8 kWh at best, and the MPC keeps both by planning its force, at 8.33 kWh on the real forecast.

> **+ DID YOU KNOW? More Energy and Fewer Hits**
>
> From the simulator: in the 3.5 m sea of case 7, the tuned damper captured 4.084 kWh and struck the end-stops twice. The predictive controller on its own forecast captured 7.011 kWh with no contact (Table 5.4). The controller that took more energy was also the one that spared the stops.

5.6.4 Twenty Minutes Through a Storm

The challenge runs one recorded sea, the same on every run, through 1.5 m, 2.5 m and a 4 m storm. It passes at 9.33 kWh or more, 40 % of the case's teaching bound of 23.3248 kWh, with no end-stop contact and no force asked of the drive beyond 400 kN, on the real forecast, as shown in Table 5.5. The shipped damper fails on contacts and energy. The MPC without survival captures 18.585 kWh, 79.7 % of the bound, but hits the stops twice in the storm. With the threshold at 3.3 m, the probe's $\hat{H}_s$ of 3.37 m trips the supervisor at 872.0 s, the brake engages at the buoy's next turning point at 872.4 s and holds it for the last 328 s, and the run keeps the 2.606 kWh and 5.927 kWh harvested in the first two seas. While locked, the brake carries at most 703 kN of its 1500 kN against an excitation of at most 537 kN, and the buoy creeps at no more than 0.141 m/s within 1.223 m. The state at 873.0 s is shown in Figure 5.4. Every threshold from 3.2 to 3.5 m passes, 3.0 and 3.1 m lock in the 2.5 m sea,

Table 5.5 The Challenge on Its Recorded Sea (Seed 3)

Policy	E (kWh)	Bound	Contacts	Verdict
shipped damping 170	8.909	38.2 %	23	fail
MPC real, no survival	18.585	79.7 %	2	fail
MPC real, H_{surv} 3.0 m	8.022	34.4 %	0	fail
MPC real, H_{surv} 3.3 m	**10.889**	**46.7 %**	**0**	**pass**
MPC real, H_{surv} 3.6 m	11.171	47.9 %	1	fail
tuned damping, H_{surv} 3.3 m	4.472	19.2 %	3	fail
perfect bound (teaching)	23.325	100 %	0	not graded

On six seeds the suggested setting (bold) never touches a stop and reaches 40 % of each seed's own bound on four (51.5, 46.9, 46.7 and 43.4 %, against 38.1 and 37.8 % on seeds 5 and 6). The MPC without survival touches the stops on seeds 3 and 5.

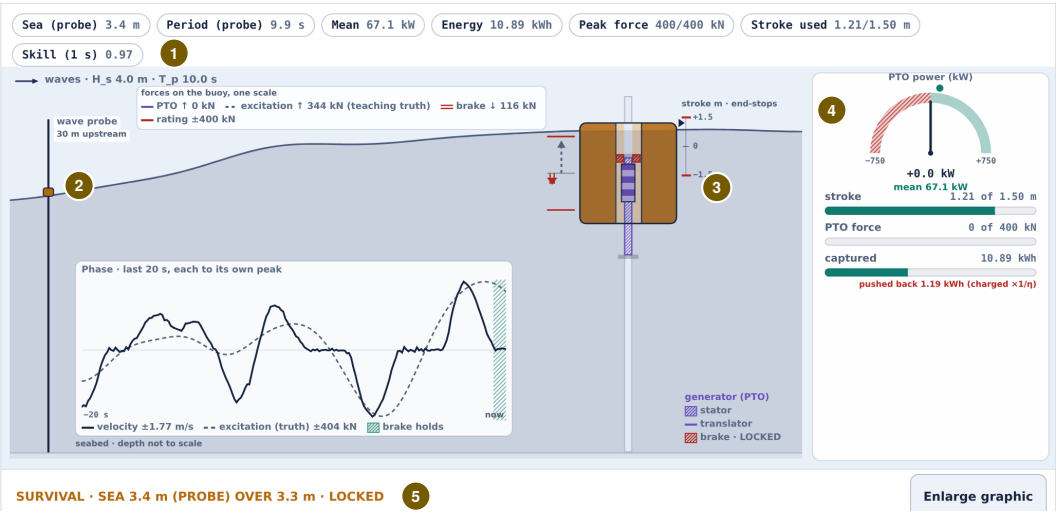

Figure 5.4 The challenge with the suggested settings (MPC on the real forecast, horizon 4 s, survival threshold 3.3 m) at $t = 873.0$ s, cropped to the Sea view. The probe reads 3.4 m and 9.9 s in the 4 m, 10 s storm, the brake has held the buoy at 1.21 m of heave since 872.4 s, and the generator applies 0 kN while the brake carries 116 kN against an excitation of 344 kN (teaching truth). The meter reads 10.89 kWh, and 1.19 kWh has been pushed back into the sea and charged at $1/\eta_{PTO}$. Callouts: (1) the badge row with the probe's sea state, (2) the wave probe upstream, (3) the buoy held by its brake on the spar, (4) the PTO power meter with the stroke, force and energy bars, and (5) the status line of the survival lock.

where $\hat{H}_s$ reaches 3.11 m, and 3.6 m touches a stop while locking. The lock costs 7.7 kWh against the MPC without survival, the price of not seeing the storm's biggest group coming.

The verdict belongs to one sea. On six seeds, each against its own bound, the shipped damper captures 37.4 to 39.3 % with 11 to 39 contacts every time, the MPC without survival rides out four storms and touches the stops on two, and the suggested setting stays off the stops on all six while reaching the target on four. A controller cannot know which sea it will meet, so the lock is insurance. The force term of the target never decides a verdict. Among nine student-like policies with survival and the 597 graded challenge runs of the search below, every one that asked for more than 400 kN also hit the stops or fell short of the energy, so the term states the rating and catches a law that leans on the drive's cut. A run that used the perfect forecast for even 1 s is not graded.

5.6.5 What the Bound Bounds

Each case's bound is the MPC of Equation (5.13) on the true future, planning to within 0.01 m of the stops where the student's controllers keep 0.3 m, and valuing the energy still stored at the horizon's end in full. A search of 6620 runs over every strategy and every setting a student can reach, 662 on each case, from damping and reactive laws over their whole dial ranges to latching, the MPC's horizon, margin, terminal value and the AR order and window on both forecasts, found none above it. The best perfect-forecast setting a student can reach (margin 0.05 m, $\alpha = 1$) comes to 97.6 to 98.5 % of it on cases 2 to 10, and ties it in case 1's small sea, where no margin binds. The best real forecast comes to 91.2 to 94.3 %. The word bound is scoped to what the app offers. It is not the most a buoy could capture. A plan solved offline over 60 s windows on the same discrete model, with the whole window's excitation known, stays within −0.51 to +0.22 % of the receding 4 s MPC over the first 50 s of each window, and gains 0.9 to 6.7 % only at the window's end by spending the energy stored in the buoy. The certified charge on reactive power is conservative. On case 6's sea, the perfect-forecast MPC at the bound's settings captures 3.1032 kWh with the certified charge and 3.3534 kWh with half of it, 8 % more, though the halved charge is no longer a lower bound of the meter. Every MPC in the app and its bound pay the same charge, so no comparison among them is unfair.

5.6.6 Code Verification and Agreement With Published Theory

Code verification tests the program against closed-form results and against the conservation laws it should obey. The hydrostatics computed by hand equal the source to the last digit printed. The free decay's period of 4.3777 s from the characteristic quartic matches the 4.38 s swings of the simulated decay, and the storage identity of the radiation model holds on a random velocity to 3.5×10^{-8} relative. In a regular wave at the model's resonance (4.3795 s, $A = 0.25$ m, $B_{pto} = B(\omega_0)$), the time-domain capture over 40 whole periods equals $|F|^2/(8B)$, 5.1911 kW, to 10^{-5}, and off resonance at 8 and 10 s the time and frequency domains agree to four digits. The energy ledger closes to 1.44×10^{-6} of the excitation's work over 900 s of the challenge with eight end-stop impacts, to 7.9×10^{-9} with the survival lock, to 9.2×10^{-8} under latching and to 4.4×10^{-8} with the force cut in case 10. Halving the step to 5 ms moves the challenge's energy by at most 0.0016 % with the same contacts. In the MPC, the balance form of the objective agrees with the exact $\sum f_k \Delta x_k$ to 0.0138 % on 200 random force sequences. The Hessian's smallest eigenvalue, with force in newtons and the app's 30 blocks, is 4.2×10^{-13} for the exact bilinear form, 2.1×10^{-14} for the balance form and 6.0×10^{-8} with the charge. The plant's continuous meter differs from the plan's step meter by at most 0.077 % of a run. At every one of 312 plans that carried slack, a feasibility linear program on the same state, preview, rating and blocks found that no force within the rating could avoid it, so the penalty is exact. In Node a solve takes 0.267 to 0.290 ms on average, 3.65 ms at most, with no failure. In Chrome at 30 times real time the challenge's twenty minutes take 40.0 s of wall time, with mean solves of 0.287 and 0.313 ms for the suggested setting and the bound and 6.30 ms at most.

Agreement with published theory is a separate matter, and here the model departs from the hull it draws. A matched-eigenfunction solution of this cylinder, the method of [13], gives an added mass at infinite frequency of 30.4 t against the teaching 16.8 t and a natural period of 4.67 s against 4.38 s. Against it the teaching model sits 28 to 31 % low on the added mass over 0.4 to 1 rad/s, and 14 to 27 % high on the radiation damping and 3 to 14 % high on the excitation over 0.4 to 0.8 rad/s. The lessons rest on the model's own consistency, its point-absorber limit, its passivity and its ledger, and not on these values. The JONSWAP form of Equation (5.5) integrates to $4\sqrt{m_0} = 1.0012 H_s$ and a ratio T_{m01}/T_p of 0.8344 against the 0.8348 the estimator uses. The probe leads the buoy by 2.55 s at the peak of the elevation's cross-correlation in the 9 s sea, between the deep-water phase delay of 2.13 s and group delay of 4.27 s. At 120 s the probe's $\hat{H}_s$ falls within 4.6 % of the sea as set and $\hat{T}_p$ within 3.9 % on four seas. The app has had no educational evaluation.

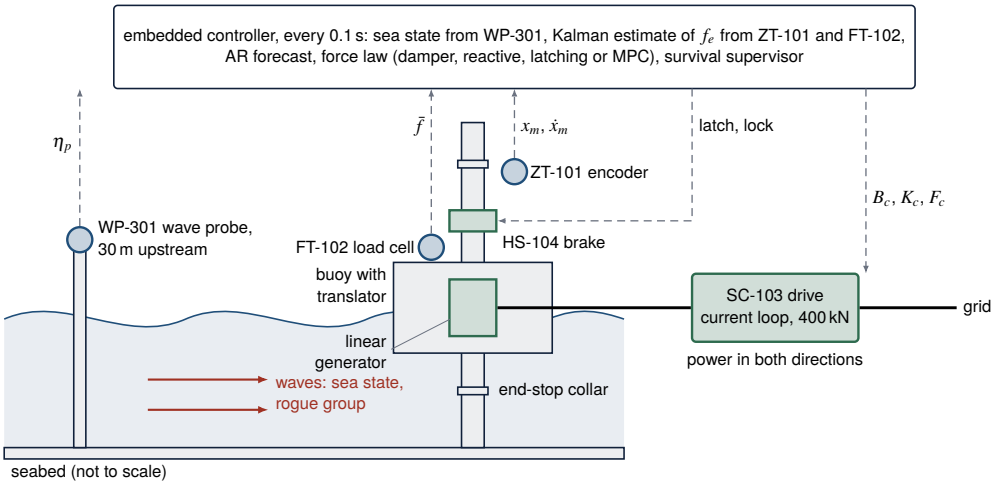

Figure 5.5 The force loop drawn with its hardware. The wave probe WP-301 reads the sea 30 m upstream, the encoder ZT-101 reads the heave and its velocity, and the load cell FT-102 reads the mean force between buoy and translator, collars included. The embedded controller turns them into a force law every 0.1 s, the drive SC-103 applies it continuously through the linear generator within 400 kN, and the brake HS-104 holds the buoy for latching and the survival lock. Notice that no instrument reads the wave force f_e that the optimum needs: the controller estimates it.

5.7 Technology in Practice

The force laws of Section 5.5 read three instruments and move two final elements, the drive of the linear generator and the brake. Figure 5.5 draws the loop with that hardware in place of the schematic's symbols.

What Measures It

The heave and its velocity come from a position sensor between buoy and spar. An encoder reads a position directly, and a speed is derived from the change of position with time [20]. The force comes from a load cell, strain gauges bonded to a metal element that stretches and compresses with the applied force as though it were an extremely stiff spring [21]. Stray electric currents through a load cell are read as changes of force [21]. The sea is read upstream, at sea by a measurement buoy. The chapter's encoder carries 5 mm and 0.01 m/s of noise, and its load cell reports the exact mean force of each period with the collars' force included, which is why an end-stop contact does not look like a burst of wave force to the filter.

> **TECHNOLOGY Sensor: A Wave Buoy That Feels the Sea**
>
> A wave buoy measures the sea through its own motion. An accelerometer with its axis perpendicular to the deck records the vertical acceleration, and the acceleration spectrum becomes a displacement spectrum when it is divided by the fourth power of the radian frequency. One buoy system analyzes 20 min records sampled at 1.5 Hz, and the significant wave height comes from the variance of the elevation, with 90 % confidence limits of about −10 % and +15 % [22]. The division amplifies electronic noise at low frequency, so a poorly corrected record shows low-frequency swell that is not real [22]. The chapter's probe reports the elevation every control period with 0.02 m of noise, and its sea-state estimate uses the last 120 s.

What Moves It

A power electronic converter sets the generator's current and so its force. A power take-off that pushes as well as pulls must pass power in both directions, a capability that many published designs

neglect [23]. The chapter's drive follows its law continuously between control instants, clipped at its 400 kN rating and at half of it after a fault. The brake is the second final element. It clamps the translator for latching and for the survival lock, with a rating of 1.5 MN, and its work is never credited.

> **TECHNOLOGY** **Final Element: A Direct-Drive Linear Generator**
>
> Waves deliver their energy at large forces and low velocities. Hydraulics or turbines can turn the slow motion into fast rotation, and a direct-drive linear generator takes it with fewer moving parts [24]. In one sea-trial converter a translator carrying permanent magnets moves inside the stator windings, end stops limit its stroke in the most powerful seas, and the generator is rated 10 kW at 0.67 m/s with 1.8 m of free stroke. Its voltage varies in amplitude and frequency, so its current is rectified and inverted again before it reaches the grid. A contact with an end stop shows on a trend as a sharp peak in the measured line force [24]. The chapter's drive applies its force law through a current loop clipped at 400 kN.

Where the Controller Runs

On a test device the most common arrangement is an integrated one: a dedicated controller packaged with the sampling hardware and embedded in the device, which acquires, scales and checks the measurements and also performs the control logic [20]. Supervisory and lower-level control loops run with cycle times from milliseconds to minutes [25]. The chapter's controller runs every 0.1 s and solves its quadratic program in at most 3.65 ms in Node. A practical controller often stores one set of gains for each bin of significant wave height and period and updates them when the sea changes, every several minutes [19], much as the tuned damper of this chapter re-tunes from the probe's estimate.

> **IN PRACTICE** **Check the Data Before the Power**
>
> The power performance of a wave energy converter is assessed against the technical specification IEC TS 62600-100, and accredited testing to it adds requirements and cost to sensor selection and data acquisition [20]. Each sensor carries its own calibration, because two identical load cells can read 3.10 V and 3.15 V under the same reference weight. In service, automated checks run in near real time: a value outside its expected range, a change too large or no change at all between samples, and a lost correlation between two signals that should agree, such as generator shaft speed and generator voltage [20].

What Is Modernizing It

Three developments touch this loop. Control co-design sizes the power take-off and its controller together, because a device designed for the most mechanical power may carry a needlessly costly specification [19]. All-electric power take-offs built on modular multilevel converters have been tested with controllers on field programmable gate arrays, although marine energy remains at an early stage of technology readiness without a commercial breakthrough [23]. Learning-based controllers have adjusted the damping of a marine current turbine as marine growth changed it, without measuring the growth [26], a remedy for the model drift that the estimator and the MPC of this chapter do not face.

5.8 Limitations

The hydrodynamics are a teaching model. The kernel is declared, the two-state fit is within 9.8 % of it, and against the cylinder's own coefficients the added mass is about 30 % low and the damping and excitation up to 27 % and 14 % high over the band of interest, so the natural period is 4.38 s where the hull would ring at 4.67 s. Numbers for a real buoy would move, while the ranking of the controllers rests on the model's own consistency. The estimator and the MPC use the plant's exact model, so model error, which a real controller's margin would also have to cover, is absent. The sea has 120 components, and the challenge's record shows a normalized autocorrelation of 0.49 at a lag of 478 s, so its wave groups resemble each other more than a real record's. The probe's peak period is within 4 %, not exact, and the latch rule inherits that error.

OPEN CHALLENGE **Marine Growth Moves the Optimum**

In simulations of a reference point absorber in a 2 m sea at its resonant 12 s period, heavy marine growth on the hull cut the power by 15 %, and re-tuning the damping recovered only about 6 % of the loss. A failure study of a 100 kW reference converter ranked marine growth, with up to 45 % less power, and power take-off failures among the costliest risks [26]. The chapter's hull never changes. Progress would be adaptive controllers and models that follow a hull as it degrades in service [26].

The reactive charge is a certified lower bound of the step meter, which the plant's continuous meter matches within 0.08 %, and it costs about 8 % of the perfect forecast's energy on case 6's sea. The shipped terminal value $\alpha = 0.5$ is not the best setting. At the shipped margin, $\alpha = 1$ changes the energy by -0.02 to $+0.48$ % on cases 2 to 10 and by 1.5 to 1.6 % in case 1. The challenge's verdict belongs to one recorded sea, and the suggested setting misses the target on two of six seeds. The 400 kN term of the target never binds alone. The force rating, the stroke, the brake, the efficiency and the survival threshold are teaching values, and passing the challenge in simulation certifies no converter and checks no design.

OPEN CHALLENGE **Optimal Control on a Real Machine**

Two issues limit the classical theory of complex-conjugate control on a real converter: the optimum is noncausal, and the interaction of body and water and the power take-off itself are nonlinear [19]. The excitation force cannot be measured directly, and a controller that maximizes mechanical power may absorb a net negative amount of electricity once the losses of the power take-off are counted [19]. The chapter's controllers use the plant's exact model and a meter that charges only reactive power, so neither problem appears in full. Progress would be controllers designed together with the machine and judged by the electricity they deliver.

5.9 Conclusion

A heaving buoy whose own period is half the waves' captures little through a damper and much more through a controller that knows the coming wave force, and this chapter has measured that difference on one declared model. The tuned damper takes 0.767 kWh in five minutes, reactive control 1.330 kWh at nearly twice the peak force, latching 1.236 kWh in a longer sea, and the MPC 2.713 kWh on a perfect 4 s forecast and 2.522 kWh on a forecast built from its own Kalman estimate. In a 3.5 m sea the MPC keeps the stroke to 1.22 m where a tuned damper hits the stops, and in the twenty-minute storm a survival lock on the probe's estimate passes the target at 10.889 kWh with no contact. The hydrodynamics are a teaching kernel that departs from the real cylinder's by up to 31 %. A fuller model would add pitch and surge, a hull computed by a boundary-element method, model error in the controller, a generator that loses energy both ways, and an array of buoys sharing one sea. The buoy makes anticausal optimal control something a reader can run, measure and break.

KEY POINTS

- A heaving buoy turns wave motion into electricity through a generator whose force a controller sets, and the force that captures the most depends on waves still to come.
- The tuned damper takes 0.767 kWh in five minutes, reactive control 1.330 kWh, and the MPC 2.713 kWh on a perfect 4 s forecast and 2.522 kWh on its own forecast.
- Most of a forecast's value lies within a quarter of the peak period.
- In a 3.5 m sea the MPC keeps the stroke to 1.22 m where a tuned damper hits the stops, and the graded storm is passed by a survival lock rather than by harvesting more.

Where else this applies. An economic controller acting on a forecast also decides the reservoir release of Volume 9, Chapter 7, and a preview of the disturbance feeds the predictive controller of Chapter 4.

References

[1] M. Ellis, H. Durand, and P. D. Christofides, "A tutorial review of economic model predictive control methods," *Journal of Process Control*, vol. 24, no. 8, pp. 1156–1178, 2014.

[2] J. V. Ringwood, G. Bacelli, and F. Fusco, "Energy-maximizing control of wave-energy converters: The development of control system technology to optimize their operation," *IEEE Control Systems*, vol. 34, no. 5, pp. 30–55, 2014.

[3] N. Faedo, S. Olaya, and J. V. Ringwood, "Optimal control, MPC and MPC-like algorithms for wave energy systems: An overview," *IFAC Journal of Systems and Control*, vol. 1, pp. 37–56, 2017.

[4] G. Hagerman, G. Scott, and P. Jacobson, "Mapping and assessment of the United States ocean wave energy resource," Electric Power Research Institute, Palo Alto, CA, Tech. Rep. 1024637, 2011.

[5] D. V. Evans, "A theory for wave-power absorption by oscillating bodies," *Journal of Fluid Mechanics*, vol. 77, no. 1, pp. 1–25, 1976.

[6] J. Falnes, *Ocean Waves and Oscillating Systems: Linear Interactions Including Wave-Energy Extraction*. Cambridge University Press, 2002.

[7] J. Hals, J. Falnes, and T. Moan, "Constrained optimal control of a heaving buoy wave-energy converter," *Journal of Offshore Mechanics and Arctic Engineering*, vol. 133, no. 1, p. 011401, 2011.

[8] F. Fusco and J. V. Ringwood, "Short-term wave forecasting for real-time control of wave energy converters," *IEEE Transactions on Sustainable Energy*, vol. 1, no. 2, pp. 99–106, 2010.

[9] Z. Yu and J. Falnes, "State-space modelling of a vertical cylinder in heave," *Applied Ocean Research*, vol. 17, no. 5, pp. 265–275, 1995.

[10] T. Perez and T. I. Fossen, "Time- vs. frequency-domain identification of parametric radiation force models for marine structures at zero speed," *Modeling, Identification and Control*, vol. 29, no. 1, pp. 1–19, 2008.

[11] A. Babarit and A. H. Clément, "Optimal latching control of a wave energy device in regular and irregular waves," *Applied Ocean Research*, vol. 28, no. 2, pp. 77–91, 2006.

[12] W. E. Cummins, "The impulse response function and ship motions," David Taylor Model Basin, Report 1661, Oct. 1962.

[13] R. W. Yeung, "Added mass and damping of a vertical cylinder in finite-depth waters," *Applied Ocean Research*, vol. 3, no. 3, pp. 119–133, 1981.

[14] K. Hasselmann *et al.*, "Measurements of wind-wave growth and swell decay during the Joint North Sea Wave Project (JONSWAP)," *Ergänzung zur Deutschen Hydrographischen Zeitschrift, Reihe A (8)*, no. 12, pp. 1–95, 1973.

[15] A. A. E. Price, D. I. M. Forehand, and A. R. Wallace, "Time-span of future information necessary for theoretical acausal optimal control of wave energy converters," in *2009 European Control Conference (ECC)*, 2009, pp. 3761–3766.

[16] R. Genest, F. Bonnefoy, A. H. Clément, and A. Babarit, "Effect of non-ideal power take-off on the energy absorption of a reactively controlled one degree of freedom wave energy converter," *Applied Ocean Research*, vol. 48, pp. 236–243, 2014.

[17] S. H. Salter, "Wave power," *Nature*, vol. 249, no. 5459, pp. 720–724, 1974.

[18] K. Budal and J. Falnes, "A resonant point absorber of ocean-wave power," *Nature*, vol. 256, no. 5517, pp. 478–479, 1975.

[19] G. Bacelli and R. G. Coe, "Comments on control of wave energy converters," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 1, pp. 478–481, 2021.

[20] R. Raye, C. Nichols, A. Bharath, C. Candon, and E. Baca, "Data acquisition and control for marine energy devices: Cost considerations," National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5700-90817, 2024.

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

[22] M. D. Earle, "Nondirectional and directional wave data analysis procedures," National Data Buoy Center, National Oceanic and Atmospheric Administration, Stennis Space Center, MS, NDBC Technical Document 03-01, 2003. [Online]. Available: https://www.ndbc.noaa.gov/wavemeas.pdf

[23] H. Nademi, B. J. Galindez, M. Ross, and M. Lopez, "Experimental validation of a modular all-electric power take-off topology for wave energy converter enabling marine renewable energy interconnection," *Journal of Marine Science and Engineering*, vol. 12, no. 8, p. 1323, 2024.

[24] M. Leijon, C. Boström, O. Danielsson, S. Gustafsson, K. Haikonen, O. Langhamer, E. Strömstedt, M. Stålberg, J. Sundberg, O. Svensson, S. Tyrberg, and R. Waters, "Wave energy from the North Sea: Experiences from the Lysekil research site," *Surveys in Geophysics*, vol. 29, no. 3, pp. 221–240, 2008.

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

[26] T. Brekken, "Applied research and development to support open-water testing at PacWave, task 7: WEC system control for marine growth and electrical faults," Oregon State University, Corvallis, OR, Final Technical Report DOE-OSU-EE0009969-7, 2025.

5.10 Exercises

5.1 Concept check · 2 min

The chapter calls the force that captures the most energy from an irregular sea anticausal. What makes it so?

(a) The probe stands 30 m upstream, so its record reaches the controller late

(b) The drive's current loop lags the command, so the force must be asked for early

(c) The optimal velocity is the excitation through a filter whose response is even in time

(d) The end-stops bind only in waves that have not yet reached the buoy

5.2 Concept check · 1 min

In case 4 latching captures 1.236 kWh against the tuned damper's 0.747 kWh, although the brake holds the buoy still 43 % of the time. Why does it capture more?

(a) The meter credits the work the brake absorbs during each hold

(b) Holding the buoy stops it radiating for the rest of each wave period

(c) The brake raises the force rating of the generator while it holds

(d) Released at the right moment, the buoy moves while the wave force is with it

5.3 Concept check · 2 min

The energy the generator absorbs over a plan, $\sum_k f_k \Delta x_k$, is bilinear in force and motion, yet the MPC of Equation (5.13) is a convex quadratic program. Why?

(a) The model's energy balance turns it into terms with non-negative weights, because $b_1 a_1 > 0$

(b) The hard force bound $|f_k| \le F_{lim}$ removes the negative curvature of the bilinear term

(c) The slack price $\rho_2 = 10^{10}\,\mathrm{J/m^2}$ outweighs the bilinear term in the Hessian

(d) The 30 move blocks reduce the decision variables until no indefinite direction is left

5.4 Theory and calculation · 6 min

The buoy is a vertical cylinder of radius $a = 2.5\,\mathrm{m}$ and draft $d = 4\,\mathrm{m}$ in sea water of density $\rho = 1025\,\mathrm{kg/m^3}$, with $g = 9.81\,\mathrm{m/s^2}$. Its added mass at infinite frequency is the hemisphere's, $m_\infty = \rho\pi a^3/3$.

(a) Compute the displaced mass m, the hydrostatic stiffness k_h, the added mass m_∞ and the undamped natural period $T_n = 2\pi\sqrt{(m+m_\infty)/k_h}$. *Checkpoint:* 80.5 t, 197.4 kN/m, 16.8 t and 4.41 s.

(b) The dominant roots of the characteristic quartic of Section 5.4 are $-0.0612 \pm 1.4353i\,\mathrm{s^{-1}}$. Compute the damped period, the damping ratio and the ratio of one swing's amplitude to the one before. *Checkpoint:* 4.378 s, 0.043 and 0.76.

5.5 Theory and calculation · 8 min

Use the two-state radiation model of Equations (5.2) and (5.3) with $b_1 = 15\,908.67\,\mathrm{N/m}$, $a_1 = 1.408\,497\,\mathrm{s^{-1}}$ and $a_0 = 1.699\,333\,\mathrm{s^{-2}}$, and the masses of the previous exercise, at the wave period of 9 s.

(a) Compute the radiation damping $B(\omega)$, the added mass $a(\omega)$ and the reactance $X_i(\omega)$ of Equation (5.10). *Checkpoint:* 4.48 kN s/m, 24.7 t and −209.4 kN s/m.

(b) Compute the damping $|Z_i|$ that a damper alone should apply and the stiffness $K_{pto} = \omega X_i$ of the complex conjugate. *Checkpoint:* 209.4 kN s/m and −146.2 kN/m.

(c) In a regular 9 s wave, what fraction of the complex-conjugate power $|F|^2/(8B)$ does the best damper alone take? Use Equation (5.11). *Checkpoint:* 0.042.

5.6 Theory and calculation · 5 min

The MPC charges reactive power through the term $c\,(f-\hat{B}v)^2$ with $c = (1/\eta_{PTO}-1)/(4\hat{B})$, where v is the velocity over the step.

(a) Show that $(f-\hat{B}v)^2 + 4\hat{B}fv = (f+\hat{B}v)^2$, and conclude that $(f-\hat{B}v)^2/(4\hat{B}) \geq \max(-fv, 0)$ for any f, v and $\hat{B} > 0$. *Checkpoint:* The inequality holds with equality when $f = -\hat{B}v$.

(b) Take $\hat{B} = 215\,\mathrm{kN\,s/m}$ and $\eta_{PTO} = 0.8$. For $f = -100\,\mathrm{kN}$ and $v = 0.5\,\mathrm{m/s}$, compare the MPC's charge $(1/\eta_{PTO}-1)(f-\hat{B}v)^2/(4\hat{B})$ with the meter's extra charge $(1/\eta_{PTO}-1)\max(-fv, 0)$. Repeat for $f = 100\,\mathrm{kN}$ at the same velocity. *Checkpoint:* 12.52 kW against 12.50 kW, and 0.016 kW against zero.

5.7 App challenge · 8 min

Load scenario 2, Passive damping: tuned to the sea. The damper is set to 40 kN s/m in a 1.5 m, 9 s sea. Use the default seed, run at 10× (30 s of wall time) to the pause at 300 s.

(a) Run as shipped. Read the Captured energy tile (energy and capture width ratio), the Peak PTO force tile and the stroke used on the End-stop contacts tile. *Checkpoint:* 0.47 kWh, 13 %, 44 kN and 1.31 m.

(b) Press Reset, press Tune to the sea and run again. Read the PTO damping it set and the same tiles. *Checkpoint:* 215 kN s/m, 0.77 kWh, 21 %, 123 kN and 0.79 m.

(c) Exercise 5.5 gives $|Z_i| = 209.4\,\mathrm{kN\,s/m}$ at 9 s. Read the note the helper writes and explain why it set 215 kN s/m. Then explain why the damper that takes more resists harder and lets the buoy move less. *Checkpoint:* The note reads $\hat{T}_p$ 9.1 s and $|Z_i(\hat{\omega}_p)| = 213\,\mathrm{kN\,s/m}$, rounded to the dial's 5 kN s/m.

5.8 App challenge · 9 min

Use the default seed, the speed 10×, and runs of 300 s to the pause. Read the Captured energy tile after each run.

(a) Load scenario 5, The anticausal optimum: a perfect forecast against none. Run it as shipped, with the Forecast horizon at 0 s. Set the Forecast horizon to 4.0 s, press Reset and run again. Also read the Peak PTO force and Reactive energy tiles of the second run. *Checkpoint:* 0.86 kWh, then 2.71 kWh at 294 kN with 0.51 kWh pushed back.

(b) Load scenario 6, The real forecast: energy against horizon. Run it as shipped, with the real forecast at 4.0 s. Set the Forecast horizon to 1.0 s, press Reset and run again. *Checkpoint:* 2.52 kWh, then 1.82 kWh.

(c) Compute the ratio of the perfect forecast's energy at 4 s to none, the share of the perfect forecast's energy that the real forecast keeps at 4 s, and explain why the real forecast loses more at long leads than at short ones. *Checkpoint:* 3.1 and 93 %, within one in the last digit.

5.9 App challenge · 8 min

Load the challenge, *Challenge · twenty minutes through three seas*. It uses one recorded sea at seed 3. Set the speed to 30× (40 s of wall time per run) and run each policy to the end at 1200 s. Read the Captured energy, End-stop contacts and Challenge target tiles.

(a) Run the shipped damper of 170 kN s/m. *Checkpoint:* 8.91 kWh, 23 contacts, not met.

(b) Set the Strategy to Predictive (MPC), keep the Real (AR) forecast at 4.0 s, set the Survival threshold to 3.3 m, press Reset and run. Also read the Reactive energy tile. *Checkpoint:* 10.89 kWh, no contact, met, with 1.19 kWh pushed back.

(c) Lower the Survival threshold to 3.0 m, press Reset and run again. Explain the result, and say what the 18.585 kWh that the MPC captures without survival, with two contacts, tells an engineer. *Checkpoint:* 8.02 kWh, no contact, not met.

5.10 Application area · 3 min

Reactive control in case 3 captures 1.330 kWh at a peak force of 215 kN, and the tuned damper 0.767 kWh at 123 kN. The reactive run pushes 0.088 kWh back into the sea.

(a) A linear generator is sized, and largely priced, by the largest force it must make. Compute the energy captured per kilonewton of peak force for each controller. *Checkpoint:* 6.24 Wh per kilonewton for the damper and 6.19 Wh for reactive control.

(b) The meter credits absorbed power in full and charges reactive power at $1/\eta_{PTO}$ with $\eta_{PTO} = 0.8$. How much did the reactive run absorb in all, before the charge? *Checkpoint:* 1.440 kWh.

5.11 Application area · 4 min

The chapter's hydrodynamics are teaching values. A matched-eigenfunction solution of the same cylinder gives an added mass at infinite frequency of 30.4 t against 16.8 t and a natural period of 4.67 s against 4.38 s.

(a) The latching rule of Section 5.5 sets $T_L = (\hat{T}_{m01} - T_n)/2$ with $\hat{T}_{m01} = 0.8348\,\hat{T}_p$ and $T_n = 4.41$ s. For a peak period of 10 s, compute the hold with the app's T_n and with the cylinder's 4.67 s. *Checkpoint:* 1.97 s and 1.84 s.

(b) In which direction would the complex-conjugate stiffness K_{pto} at 9 s move on the real cylinder, and why does the chapter say that its ranking of the controllers still holds? *Checkpoint:* Toward zero, a smaller negative spring.

5.12 Review the agent · 5 min

An AI assistant is asked how to set the damper of scenario 2 and replies as follows. "At a period of 9 s the radiation resistance of this buoy is $B(\omega) = 4.5\,\mathrm{kN\,s/m}$. By the maximum power transfer theorem the load must match the source resistance, so set the PTO damping to 5 kN s/m, the nearest dial value. This is the textbook optimum for a point absorber, and it will capture far more than the 215 kN s/m that Tune to the sea sets." Find the flaw and say how you would check it.

CHAPTER 6

A Parabolic-Trough Simulator for Choosing a Setpoint Rather Than Defending One

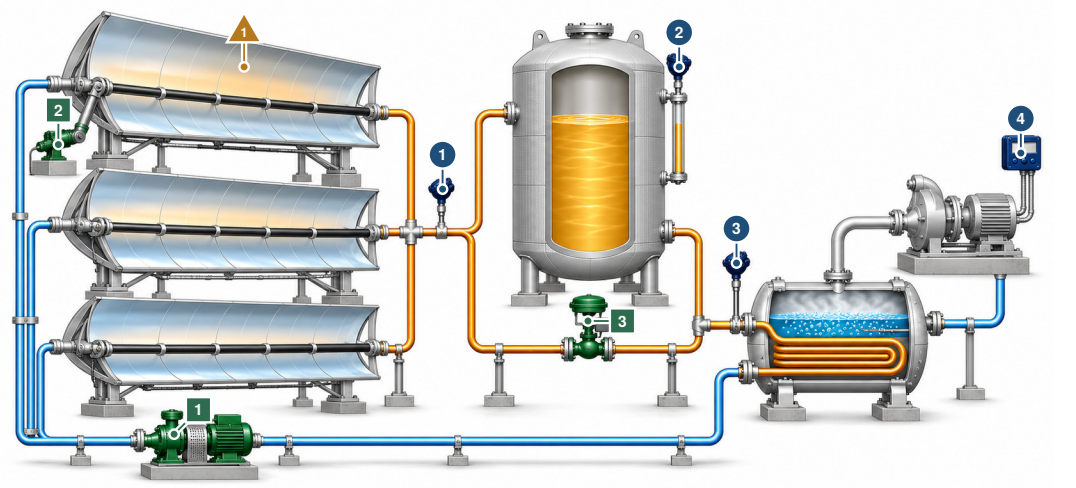

The controller moves 1 field pump speed 2 mirror drive that turns the troughs off the sun 3 bypass valve around the hot tank

The controller reads 1 oil temperature leaving the mirrors 2 hot tank level 3 oil temperature entering the steam generator 4 electric power out

What disturbs it 1 clouds that change the sunlight

Notice that the oil can reach the steam generator through the hot tank or around it, while the mirrors heat it only as fast as the sky allows. Illustration, not to scale.

IN PLAIN TERMS A Best Temperature That Moves With the Sky

The main idea. Long curved mirrors focus sunlight on a pipe, and oil in the pipe carries the heat to a storage tank and a power plant. A controller changes the pump speed to hold the hot oil at a chosen temperature.

The challenge. The pipe loses heat to the sky all the time, and more when it runs hotter, sunny or not. Under clouds, holding the usual high temperature wastes nearly a third of the heat, so the target must change with the sky.

Why it matters. On a broken day, a well chosen temperature makes about a third more electricity.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the history boxes of Section 6.2, the fact box of Section 6.3, the try-it boxes and the fact box of Section 6.7, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 6.7, Section 6.8, Section 6.9 and its open challenges

The control question A trough receiver loses heat as the fourth power of its temperature whether or not the sun is out, while the power block converts hotter oil more efficiently. When is the outlet temperature setpoint itself the decision worth automating, and when is the schedule of the power block worth more than either the setpoint or the tuning?

Learning objectives

- Compute the receiver loss, the field transit time and the outlet loop gain from the model, and explain how each moves with flow and irradiance.
- Explain why the optimal outlet temperature falls under cloud and why a fixed setpoint can be held perfectly and still be wrong.
- Describe how the level override, the high-temperature protection and the bypass interlock act on top of the four regulatory loops.
- Compare fixed and variable outlet temperature on a clear and a broken day, and the predictive controller against the outlet loop, from scenario runs.
- Separate what the tuning recovers from what the schedule earns in the graded challenge.

Prerequisites Sections 1.4.5 (a moving optimum), 1.4.6 (threshold and water value), 1.6.2 (velocity form), 1.6.3 (MPC option) and 1.6.4 (scenarios and scorecard).

The app https://apmonitor.com/apps/parabolic-trough/

See also Volume 9, Chapter 8 and Volume 9, Chapter 10 for supervisory setpoints that move with the weather or the load, and Volume 9, Chapter 6 for a delay set by the manipulated flow. Volume 5, Chapter 2 and Volume 5, Chapter 7 for control through a transport delay and a residence time.



6.1 Introduction

Ask a control engineer what to do about a collector outlet temperature and the answer arrives before the question finishes. Choose the design temperature, put a proportional-integral controller on the field flow, and defend it. The loop will do exactly that. On a cloudy morning it will defend 393 °C with complete success and throw away nearly a third of the hour doing it, and nothing about that failure is a tuning failure. The setpoint was wrong, and the plant has no mechanism for saying so.

That is the misconception this chapter corrects. A control course teaches regulation, and regulation assumes somebody upstream chose the right number to regulate to. In a solar field nobody did, because the right number moves. The receiver radiates as the fourth power of its metal temperature and does not care whether the sun is out, while the heat available to collect changes by a factor of five across a broken day. At full irradiance the conversion gain from running hot outweighs the radiation and the design temperature is roughly right. At a quarter of full irradiance, reaching the same temperature means running the field at a sixth of its flow, so the fluid spends six times as long under a receiver that is still losing 3.7 MW, and nearly a third of what is absorbed goes back to the sky.

The class of problem is steady-state optimization of an operating point that moves with a measurable disturbance, implemented as a supervisory layer that writes a setpoint into an otherwise conventional regulatory scheme. A fired heater with a variable draft, a solid oxide fuel cell trading polarization against thermal loss, a refrigeration plant choosing a condensing pressure against ambient, and a distillation column choosing a reflux ratio against feed composition all have this shape. In each, a loss term and a gain term depend on the same variable with opposite sign, the balance between them depends on the load, and holding the operating point fixed is a decision that somebody made once and nobody revisits. The formal treatment is real-time optimization above regulatory control, and the solar field is the clearest instance in this volume because the disturbance is visible in the sky. A data hall whose chilled water setpoint follows the outdoor wet bulb, in Volume 9, Chapter 8, has the same structure.

The chapter replicates on one plant a result already in the literature for this configuration [1, 2]. It reports the clear day and the broken day at the same length and finds that the variable temperature

strategy is worth 0.3 % of net generation on the first and 32 % on the second. The plant is a thirty-state distributed receiver model with a supervisory optimizer that runs twenty algebraic evaluations every three minutes, so it executes at interactive speed in a browser. Twelve scenarios separate what the regulatory layer can do from what only the supervisory layer can do, and the most valuable decision they expose is a setpoint rather than a gain.

Section 6.2 traces trough control from the SEGS plants to forecast-based dispatch. Section 6.3 gives the model, its parameters and its assumptions. Section 6.4 gives the open-loop character and the optimum outlet temperature against sky clarity. Section 6.5 gives the four regulatory loops, the override and the out-of-service logic, and Section 6.5.4 the optional predictive outlet controller. Section 6.6 gives the supervisory optimizer and the schedule. Section 6.7 reports the scenarios, the clear and broken days, the schedule, the override, the graded challenge and, in Section 6.7.6, the predictive controller against the outlet loop. Section 6.8 describes the hardware of a real field, Section 6.9 states the limitations and Section 6.10 concludes.

6.2 From SEGS to a Weather Forecast

The first commercial plants, told in the history box below, set the pattern the industry still follows: the collector loop outlet temperature is the controlled variable, the field circulation pump is the manipulated variable, and a fixed setpoint a few degrees below the stability limit of a synthetic aromatic fluid, about 400 °C, is defended through whatever the weather does [3, 4].

> ⧖ HISTORY **Nine Plants in the Mojave (1984 to 1991)**
>
> Nine commercial parabolic trough plants were built in the California Mojave Desert between 1984 and 1991, ranging from 14 to 80 MW of net output [5]. Falling energy prices and shrinking incentives stopped the building, and the developer filed for bankruptcy in 1991 when it could not finance a tenth plant [5]. The plants kept running. Many technical problems of the early units were found and corrected in the later ones, and together they gave more than 30 years of operating experience on the lifetime and performance of trough technology [5]. A loop in a modern field still takes oil in at about 292 °C and returns it at about 393 °C [5], the outlet temperature this chapter's fixed setpoint defends.

Control research on distributed collector fields is older than the commercial plants and was driven mostly by the ACUREX field at the Plataforma Solar de Almeria. Two decades of that work were collected in a survey whose first part covers the modeling and the classical schemes and whose second part covers adaptive, predictive and nonlinear approaches [6, 7], and in a book covering the same ground in depth [8]. The recurring difficulty in all of it is that the field is a distributed system with a transport delay that changes with the manipulated variable itself, so a controller tuned at design flow is wrong at a sixth of it. Feedforward from the irradiance, gain scheduling on the flow and series compensation of the delay are all standard answers, and all of them still take the setpoint as given.

Storage changed the question. Two-tank indirect storage decouples collection from generation, so a plant with a solar multiple above one can move its generation into the hours that pay for it [9, 10]. Once generation is schedulable the plant has an economic optimization to solve, not just a regulation problem, and the operating temperature becomes a decision variable rather than a constraint. Molten salt was proposed as the field fluid to raise that ceiling, and the engineering assessment of what it would cost in freeze protection is the clearest statement of why the synthetic aromatic survived [11].

> ⧖ HISTORY **Storage Comes to the Trough (2007)**
>
> Two-tank molten-salt storage was first demonstrated in a tower pilot plant in the 1990s. In the early 2000s a joint United States and European team adapted it to troughs by adding an intermediate heat exchanger between the field oil and the salt, and the concept was first used commercially in 2007 [5]. By 2020 about 31 trough plants around the world had between 3 and 10 hours of storage [5]. The chapter's two tanks hold the field oil itself, with no exchanger, and store 3.6 h at the power block rating.

The dynamic optimization of this exact plant configuration was published in a sequence of papers that the present simulator follows. A dynamic model of a trough plant with two-tank storage and a control scheme for it [12] was followed by optimization over a twenty-four hour horizon using weather forecasts [1] and then by the hybrid solar and fossil case [2]. The receiver loss correlations that make the optimum move are from the NREL receiver heat transfer model [13] and the heat-loss tests that calibrated it [14].

6.3 Process Model

The simulated plant has 73 800 m^2 of mirror aperture feeding 2815 m^2 of receiver tube, two tanks holding 1.8×10^6 kg of synthetic aromatic heat transfer fluid, and a power block of 27.8 MW_{th} rating at 36 % design efficiency. At clear noon the mirrors intercept 70.1 MW and the receivers absorb 50.8 MW, which is a solar multiple of 1.83 on the power block. Filled to its upper working level the hot tank holds 106.3 MWh above the cold tank, and 99.5 MWh of that lies between the two level limits, which is 3.6 h at the power block thermal rating. Figure 6.1 shows the arrangement and every symbol used below.

> ✦ DID YOU KNOW? **Two Million Square Meters of Mirrors**
>
> The nine first commercial trough plants held 354 MW of generating capacity and more than 2 000 000 m^2 of collectors. By the end of 2001 they had accumulated 127 years of operating experience [3]. The field of this chapter has 73 800 m^2 of mirror aperture.

The clock is in hours, temperatures in degrees Celsius, flows in kilograms per second, power in megawatts and tank level as a fraction of the total fluid inventory. Thirty states are integrated. Ten are receiver fluid node temperatures and ten are receiver metal node temperatures, which together carry the transport delay that dominates the outlet loop. The rest are the hot tank mass and temperature, the cold tank temperature, a steam generator lag, five actuator lags and a filtered sky state. The outlet and power loops read their measurements through first-order instrument lags of 2.2 s and 1.8 s, and the other two loops read the state directly.

6.3.1 Variables

Table 6.1 classifies what the application exposes. Four controlled variables pair with four manipulated variables, and one of those pairings is an override rather than a regulator.

6.3.2 The Sun on the Mirrors

Clear-sky direct normal irradiance and the optical efficiency are both written against the fraction of the solar day $x = (h - h_{up})/(h_{dn} - h_{up})$, zero at sunrise and one at sunset,

$$\mathrm{DNI} = \sigma\, \mathrm{DNI}_{max} \sin^{0.6}(\pi x), \qquad 0 < x < 1 \tag{6.1}$$

$$\eta_{opt} = \eta_{opt,0}\, \kappa \sin^{n_i}(\pi x)\, \Pi\!\left(\frac{x - 0.02}{0.09}\right) \Pi\!\left(\frac{0.98 - x}{0.09}\right) \tag{6.2}$$

where $\Pi(z) = \min(\max(z, 0), 1)$ ramps the row-end losses that make the first and last hour of the day nearly worthless. The absorbed power per receiver node follows,

$$q_a = \frac{A_{ap}}{10}\, \mathrm{DNI}\, \eta_{opt}\, (1 - f_d^a). \tag{6.3}$$

With the default parameters and 96 % mirror cleanliness, clear noon gives 950 W/m^2, 70.1 MW intercepted and 50.8 MW absorbed. The electricity price is a gaussian evening peak on a base with a midday depression, and the peak is at 19:30, after sunset. That offset is the entire commercial argument for thermal storage.

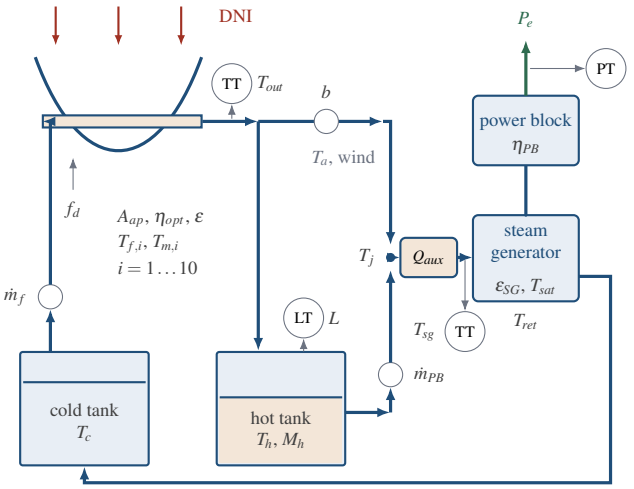

Figure 6.1 Plant arrangement and the symbols used in the equations. A single collector row stands for 73 800 m^2 of aperture, modeled as ten fluid nodes and ten metal nodes in series. Field fluid leaving the receiver either fills the hot tank or passes directly to the junction through the bypass, and it returns to the cold tank (not drawn) while it is less than 12 °C hotter than that tank. TT on the collector outlet supplies the controlled variable for the field flow loop, LT the hot tank level that drives the defocus override, TT after the heater the steam generator inlet T_{sg} that the bypass loop holds, and PT the net electrical output.

6.3.3 The Receiver

Each of the ten receiver segments carries a fluid node and a metal node. The metal node absorbs the sunlight, loses heat to the sky and conducts to the fluid, and the fluid node carries that heat downstream,

$$m_f c_p \frac{\mathrm{d}T_{f,i}}{\mathrm{d}t} = 3600\left[\dot{m}_f c_p (T_{f,i-1} - T_{f,i}) + hA(T_{m,i} - T_{f,i})\right] \tag{6.4}$$

$$m_m c_{p,m} \frac{\mathrm{d}T_{m,i}}{\mathrm{d}t} = 3600\left[q_a - hA(T_{m,i} - T_{f,i}) - q_{loss,i}\right] \tag{6.5}$$

with $m_f = M_f/10$, $m_m = M_m/10$ and $T_{f,0} = T_c$, the cold tank temperature. The loss from one node is radiation to the sky plus a convective envelope term that grows with wind,

$$q_{loss,i} = \varepsilon\sigma_{SB}\frac{A_{rec}}{10}\left[(T_{m,i}+273.15)^4 - (T_a+273.15)^4\right] + h_e\frac{A_{rec}}{10}(T_{m,i} - T_a) \tag{6.6}$$

$$h_e = h_{conv} + h_{wind}\, v_w^{0.6}. \tag{6.7}$$

Equation (6.6) is the origin of everything the application teaches. Summed over the field at an ambient of 25 °C and a 3 m/s wind, the loss is 0.12 MW at a metal temperature of 50 °C, 2.56 MW at 300 °C, 4.48 MW at 393 °C and 5.48 MW at 430 °C. None of those numbers depends on whether the sun is out.

Fluid properties are temperature dependent. Heat capacity is linear over the range the plant occupies, $c_p = c_{p,0} + c_{p,1}T$ in kJ/(kg °C) (entered in J/(kg °C) in the balances below, where it meets

Table 6.1 Variable classification

Symbol	Meaning	Unit	Range
Controlled variables			
T_{out}	collector outlet temperature	°C	312 to 400
P_e	net electrical output	MW	0 to 12
T_{sg}	steam generator inlet	°C	310 to 398
L	hot tank level (override)	%	0 to 100
Manipulated variables			
$\dot{m}_f$	field flow	kg/s	8 to 520
$\dot{m}_{PB}$	power block flow	kg/s	0 to 190
b	hot tank bypass	%	0 to 100
f_d	mirror defocus	%	0 to 100
Q_{aux}	supplemental heater	MW	0 to 12
Disturbances			
σ_∞	sky clarity	× clear	0.05 to 1.05
T_a	ambient temperature	°C	−10 to 48
v_w	wind speed	m/s	0 to 18
κ	mirror cleanliness	×	0.70 to 1.00
ε_{fc}	forecast error scale	×	0 to 1.5
σ_T	thermocouple noise, 1σ	°C	0 to 4
States			
$T_{f,i}$	receiver fluid node, $i = 1\ldots10$	°C	−30 to 520
$T_{m,i}$	receiver metal node, $i = 1\ldots10$	°C	−30 to 600
M_h	hot tank mass	kg	$0.001M$ to $0.998M$
T_h	hot tank temperature	°C	−30 to 500
T_c	cold tank temperature	°C	−30 to 500
T_{sg}	steam generator inlet lag	°C	−30 to 500
$\dot{m}_f^a$	applied field flow	kg/s	8 to 520
$\dot{m}_{PB}^a$	applied block flow	kg/s	0 to 190
b^a	applied bypass	—	0 to 1
f_d^a	applied defocus	—	0 to 1
Q_{aux}^a	applied heater duty	MW	0 to 12
σ	filtered sky clarity	—	0 to 1.05

conductances in watts), density is quadratic, and viscosity follows a Vogel form,

$$\mu(T) = \exp\left(\frac{544.149}{T + 114.43} - 2.59578\right) \quad [\text{mPa s}]. \tag{6.8}$$

Viscosity is in the model only because of the pump. At the design mean temperature of 346 °C the fluid is 0.243 mPa s and at 30 °C it is 3.23 mPa s. At design flow a field at 30 °C therefore needs 32.7 bar rather than 19.5 bar, although the pump power rises only 21 % because the cold oil is denser and occupies less volume. The model never limits the flow on that account.

6.3.4 Tanks, Junction and Power Block

Field fluid leaving the receiver splits three ways. A fraction b^a goes straight to the junction as bypass, as much of that as the power block will take is used directly, and everything else goes to the tanks,

$$\begin{aligned} \dot{m}_{byp} &= b^a \dot{m}_f, & \dot{m}_{dir} &= \min(\dot{m}_{byp}, \dot{m}_{PB}), \\ \dot{m}_{in} &= \dot{m}_f - \dot{m}_{dir}. \end{aligned} \tag{6.9}$$

What reaches the tanks is not automatically stored. Sending 200 °C oil into a 390 °C tank destroys the only thing the tank was holding, so the field return is recirculated to the cold tank until it is worth

more than a margin above it. The changeover is smoothed over a few degrees so that the right-hand side stays continuous for a Runge-Kutta step,

$$w_h = \Pi\left(\frac{T_{f,10} - (T_c + \Delta T_{rec})}{W_{rec}}\right), \quad \dot{m}_{hot} = w_h \dot{m}_{in}. \tag{6.10}$$

The tank balances follow, with $\dot{m}_{th} = \dot{m}_{PB} - \dot{m}_{dir}$ the draw from the hot tank and $M_c = M - M_h$. The hot pumps lose suction near the bottom of the tank. Below $M_h = 0.002M$ the part of their draw above $\dot{m}_{hot}$ fades linearly to zero at $0.001M$, and $\dot{m}_{PB}$ here and in the block equations is the flow they deliver, so an empty tank passes on only what flows into it and no state is clamped afterward,

$$\frac{\mathrm{d}M_h}{\mathrm{d}t} = 3600(\dot{m}_{hot} - \dot{m}_{th}) \tag{6.11}$$

$$M_h \frac{\mathrm{d}T_h}{\mathrm{d}t} = 3600\left[\dot{m}_{hot}(T_{f,10} - T_h) - \frac{UA_t(T_h - T_a)}{c_p}\right] \tag{6.12}$$

$$M_c \frac{\mathrm{d}T_c}{\mathrm{d}t} = 3600\left[\dot{m}_{PB}(T_{ret} - T_c) + (1 - w_h)\dot{m}_{in}(T_{f,10} - T_c) - \frac{UA_t(T_c - T_a)}{c_p}\right]. \tag{6.13}$$

The junction mixes bypass fluid with hot tank fluid, the supplemental heater raises it further, and a first-order lag of 9 s represents the steam generator metal,

$$T_j = \frac{\dot{m}_{dir} T_{f,10} + \dot{m}_{th} T_h}{\dot{m}_{PB}}, \qquad \tau_{SG} \frac{\mathrm{d}T_{sg}}{\mathrm{d}t} = T_j + \frac{10^6 Q^a_{aux}}{\dot{m}_{PB} c_p} - T_{sg}. \tag{6.14}$$

The power block is a counterflow steam generator of fixed effectiveness feeding a turbine with a part-load efficiency curve and a linear gain with inlet temperature,

$$\Delta T = \varepsilon_{SG}(T_{sg} - T_{sat}), \qquad Q = \frac{\dot{m}_{PB} c_p \Delta T}{10^6} \tag{6.15}$$

$$\eta = \eta_{PB}\, g(Q/Q_{PB})\,[1 + k_T(T_{sg} - 393)] \tag{6.16}$$

$$g(u) = \Pi^{1.05}_{0.2}\left(0.53 + 0.62u - 0.15u^2\right) \tag{6.17}$$

with the block producing nothing at all when $T_{sg} < T_{j,min}$ or when the load fraction falls below x_{min}, which is a genuine discontinuity and is what a real turbine does. At an inlet of 393 °C and 119 kg/s the block takes 26.4 $\mathrm{MW_{th}}$ and produces 9.33 MW gross at 35.4 %. At 340 °C the same flow gives 13.1 $\mathrm{MW_{th}}$ and 3.45 MW at 26.4 %, and at 310 °C the block is off.

Field pumping is quadratic in flow on a fixed hydraulic system with a viscosity correction,

$$\Delta p = \Delta p_0 \left(\frac{\mu(\bar{T})}{\mu_{ref}}\right)^{0.2} \left(\frac{\dot{m}_f}{\dot{m}_0}\right)^2, \quad P_{pump} = \frac{\Delta p\, \dot{m}_f}{\rho\, \eta_p\, 10^6}. \tag{6.18}$$

At the design flow of 222 kg/s the pressure drop is 19.5 bar and the pump takes 0.76 MW. At 350 kg/s it takes 2.97 MW and at the 520 kg/s limit it takes 9.75 MW, which is most of the plant output. Net electrical output subtracts pumping, heat tracing and a fixed auxiliary load from the gross, where P_{pump} is the field pump of Equation (6.18) plus 0.35 times the same law applied to the power block flow,

$$P_e = \max\left(0,\ \eta Q - P_{pump} - P_{fix} - P_{tr}\right). \tag{6.19}$$

6.3.5 Parameters

Table 6.2 gives every parameter with its default. The V column marks A for adjustable and shown, H for hidden from the beginner view, and F for a fixed physical constant.

6.3.6 Initialization and Numerics

Each scenario supplies a starting hot tank temperature, cold tank temperature and level. Before sunrise the field is not hot, so the receiver fluid nodes are initialized to the cold tank temperature with no profile along the row, and the metal nodes 2 °C above it. Giving a field at half past five in the morning a full temperature profile would credit the plant with a loop of collected heat it never collected, and that shows as a spurious couple of megawatt-hours in the harvest metric. When the clear-sky irradiance already exceeds 60 W/m^2, the nodes are initialized on a linear profile from the cold tank to the hot tank temperature, with the metal nodes 12 °C above the fluid.

Integration is classical fourth-order Runge-Kutta on a fixed step of 0.001 h, which is 3.6 s and about one time constant of the fastest actuator. The regulatory loops execute every 0.002 h, which is 7.2 s, the supervisory layer every 0.05 h, which is three minutes. The historian logs one sample every ninth physics step, 32.4 s. A post-step clip holds every temperature state, the tank mass, the two fractional actuators and the sky state inside their physical ranges.

6.3.7 Assumptions

1. **One collector row stands for the whole field.** All loops see the same irradiance, the same flow and the same inlet temperature. A real field has hundreds of parallel loops with a flow distribution, and the loop that runs hottest is the one that sets the defocus decision. The assumption understates the risk of a local overtemperature.

2. **The optical model is a smooth analytic surrogate.** The incidence angle modifier in Equation (6.2) is a power of the solar-day sine rather than a computed sun position, and the row-end term is a linear ramp. Neither is traceable to a collector test [14], and the shape is right while the numbers are representative rather than measured.

3. **Receiver loss is a single emissivity at all temperatures.** The NREL receiver model makes emissivity a function of the metal temperature [13], which steepens the loss curve above 350 °C. Using a constant understates the loss at the top of the range, so the value of lowering the setpoint is understated rather than exaggerated.

4. **Each tank is perfectly mixed and fully insulated to one** UA**.** Thermocline behavior, tank stratification and the wall thermal mass are all absent, so the tanks respond faster than real ones and never stratify on a partial charge [10].

5. **The power block is a steady-state map.** Equation (6.16) gives efficiency at once from inlet temperature and duty, with only a 9 s lag on the inlet. A real steam cycle has drum level, superheater metal and turbine thermal stress dynamics measured in tens of minutes, and the start and stop transients that carry most of the cost are missing.

6. **Fluid degradation is not modeled.** The stability limit at 400 °C is an alarm and a defocus trigger, not a mechanism. Exceeding it in the simulator costs nothing, where on a real plant it produces hydrogen that permeates into the receiver annulus and destroys the vacuum permanently.

7. **The price curve and the weather are synthetic.** The irradiance, the ambient and the price are smooth analytic profiles with a seeded forecast error. They are teaching data and are never labeled as measured.

Table 6.2 Model parameters and default values

Symbol	Meaning	Value	Unit	V
A_{ap}	mirror aperture area	73 800	m^2	A
A_{rec}	receiver tube surface area	2815	m^2	A
$\eta_{opt,0}$	peak optical efficiency	0.755	—	A
ε	receiver emissivity	0.12	—	A
h_{conv}	envelope loss, still air	0.35	W/(m^2 °C)	A
h_{wind}	wind term on envelope loss	0.25	W/(m^2 °C)	A
hA	metal to fluid conductance, per node	2.0×10^5	W/°C	A
M_f	fluid held in the field	60 000	kg	A
M_m	receiver tube mass	64 000	kg	A
$c_{p,m}$	tube heat capacity	490	J/(kg °C)	A
n_i	incidence-angle exponent	0.30	—	H
$c_{p,0}$	fluid c_p, 0 °C	1.60	kJ/(kg °C)	A
$c_{p,1}$	fluid c_p slope	0.0024	kJ/kg $°C^2$	A
ρ_0	fluid density, 0 °C	1074	kg/m^3	H
ρ_1, ρ_2	density slope, curvature	0.6367, 7.81×10^{-4}	kg m^{-3} $°C^{-1}$, $^{-2}$	H
T_{max}	fluid stability limit	400	°C	A
T_{freeze}	fluid freezing point	12	°C	A
$T_{f,min}$	freeze-protection setpoint	60	°C	A
M	total fluid inventory	1.8×10^6	kg	A
UA_t	tank heat loss, each	1500	W/°C	A
ΔT_{rec}	recirculation margin	12	°C	A
W_{rec}	recirculation changeover width	10	°C	H
L_{min}, L_{max}	working level limits	0.06, 0.94	—	A
Q_{PB}	power block thermal rating	27.8	MW	A
η_{PB}	block efficiency at rating	0.36	—	A
T_{sat}	evaporator saturation	285	°C	A
ε_{SG}	steam generator effectiveness	0.85	—	A
$T_{j,min}$	minimum inlet temperature	305	°C	A
x_{min}	minimum stable load	0.25	—	A
k_T	efficiency gain per degree	0.0013	1/°C	A
P_{fix}	fixed auxiliary load	0.32	MW	A
τ_{SG}	steam generator lag	9	s	A
Δp_0	field drop at design flow	2.0×10^6	Pa	A
$\dot{m}_0$	design field flow	222	kg/s	A
η_p	pump and drive efficiency	0.75	—	A
μ_{ref}	viscosity, design mean	0.277	mPa s	H
DNI_{max}	clear-sky noon irradiance	950	W/m^2	A
h_{up}, h_{dn}	sunrise, sunset	6.2, 19.3	h	A
τ_{sky}	cloud edge crossing time	0.04	h	A
t_0	clock time at the start of a run	4.0 (per case)	h	A
p_{base}	off-peak price	42	$/MWh	A
p_{peak}	evening price adder	165	$/MWh	A
c_{aux}	cost of supplemental heat	58	$/MWh	A
c_{st}	cost of a block start (not charged)	900	$	A
$T_{m,max}$	receiver metal limit	430	°C	A
$Q_{aux,max}$	heater rating	12	MW	A
ΔP_b	power tolerance band	0.4	MW	A
$\tau_f, \tau_{pb}, \tau_b, \tau_d$	actuator lags	3, 5, 5, 40	s	A
τ_{Tm}, τ_{Pm}	thermocouple, power transducer lags	2.2, 1.8	s	A
σ_{SB}	Stefan-Boltzmann constant	5.67×10^{-8}	W $m^{-2}K^{-4}$	F

Table 6.3 Optimal outlet temperature at solar noon against sky clarity

Sky × clear	DNI (W/m^2)	Q_{abs} (MW)	T_{opt} (°C)	Loss (MW)	$\dot{m}_f$ (kg/s)	Net gain over 393
1.00	950	50.8	393	3.98	207	0.0 %
0.70	665	35.6	393	3.81	140	0.0 %
0.50	475	25.4	340	3.15	238	8.8 %
0.35	333	17.8	327	2.96	232	21.4 %
0.25	238	12.7	325	2.89	167	36.0 %
0.15	143	7.6	325	2.84	81	–

6.4 Open-Loop Character

The field holds 60 000 kg of fluid and the design flow is 222 kg/s, so the transit from the cold tank inlet to the collector outlet takes 4.5 min. That is the dominant dynamic in the outlet loop and it is not a fixed dead time, because the flow is the manipulated variable. Cutting the flow to a sixth of design to defend a temperature under cloud stretches the transit to 27 min. A loop tuned at design flow is a different loop at a sixth of it, which is the standard difficulty in distributed collector fields [6]. The same delay set by the manipulated flow governs the buried network of Volume 9, Chapter 6.

Two thermal masses set the rest. The receiver metal is 31.4 MJ/°C and the fluid in the field is 146 MJ/°C, so the cloud of the third scenario, which removes 37 MW of absorbed power, moves the metal quickly and the fluid behind it. The tanks are three orders of magnitude larger, holding 99.5 MWh between the level limits, which is 3.6 h at the power block thermal rating.

Manual mode is where the inventory and the transit become visible. Starting at half past five in the morning with the whole field at the 295 °C of the cold tank, the operator sets a field flow and watches the outlet climb node by node. Setting it too high keeps the field cold and delays the first useful heat. Leaving it at the 40 kg/s the scenario starts with carries the outlet past the 400 °C stability limit at 07:42 and the receiver metal past its 430 °C limit at 07:50, while the inlet node is still near 310 °C. Raising the flow to 200 kg/s as the outlet passes 393 °C moves the outlet reading by 2 °C within 0.8 min, because a flow change acts on every node at once, but a fall of 20 °C takes 3.1 min. There is no setpoint anywhere in that run, and the reader has already met the central trade-off.

The trade-off is worth stating as a table. Table 6.3 reports the outlet temperature that maximizes net electrical output at solar noon for a range of sky clarity, computed by the application's own golden-section search over the supervisor's interval (325 to 393 °C for a 300 °C cold tank), against what the design setpoint of 393 °C would deliver at the same conditions.

Above half the clear-sky irradiance the optimum is the top of that interval, which is the design value, because the field is collecting enough that the flow needed to reach 393 °C is comfortable and the conversion gain from running hot outweighs the extra radiation. Below that the optimum falls sharply and levels out at 325 °C, which is the bottom of the interval, 25 °C above the cold tank, rather than a property of the field. With the interval opened down to 312 °C the quarter-sky optimum is 319 °C. At the bottom row the comparison stops being a percentage. Holding 393 °C at 143 W/m^2 produces exactly zero net electrical output, because the 3.51 MW of receiver loss leaves 4.1 MW to be carried at 18 kg/s, which is below the power block's minimum stable load, while the optimizer still extracts 0.89 MW.

🕮 **Pause.** Try Exercise 6.1 before reading on.

6.5 Regulatory Control

6.5.1 The Discrete Algorithm

All four loops use the velocity-form PID of Section 1.6.2, Equations (1.10) to (1.13). The chapter writes the action a as γ, $+1$ for a reverse-acting loop and -1 for a direct-acting one, as in Chapter 1.

Table 6.4 Regulatory loops as implemented

CV	MV	K_c	τ_I (h)	Character
T_{out}	$\dot{m}_f$	2.5	0.05	direct, delay dominant
P_e	$\dot{m}_{PB}$	10	0.01	reverse, fast
T_{sg}	b	1.2	0.03	reverse, interlocked
L	f_d	8	0.30	direct, override

The tracking time is $T_t = \tau_I$ on all four loops and no output has a rate limit. In manual mode, and while a loop is out of service, the controller holds its integral so that the output equals the value applied, and in automatic the manual stations follow the applied commands at every scan, so both transfers are bumpless. Every loop in this application runs with $\tau_D = 0$, so there is no derivative action anywhere and no derivative filter to choose. The regulatory sample period is $\Delta t_c = 7.2\,\mathrm{s}$.

Two limits act after a controller output and are invisible to its integral. The power block command is cut back toward the field flow that reaches the block or the hot tank, none of it while the return is recirculated, as the hot tank approaches its lower level limit, and the level loop's defocus demand passes the high selector of Equation (6.20). In both cases the integral is back-calculated against the controller's own clipped output rather than the value applied. At 15:05 on the broken day of Figure 6.3, with the tank at its lowest working level, the power controller sits at its 190 kg/s limit while the plant receives 140 kg/s.

6.5.2 Four Loops, One of Them an Override

Table 6.4 gives the pairings and default tunings. Two loops are direct acting. More field flow makes the outlet colder, and more defocus makes the tank level fall, so both process gains are negative.

The outlet loop has a process gain of about 0.38 °C/(kg s) at the design point and about four and a half minutes of transport delay, and the gain rises as the irradiance falls. The conventional rule for a plant whose gain varies with the load is to tune it at the weak condition. The default tuning is not set that way. With it the second scenario's step from 370 back to 393 °C settles within 1 °C in 12 min without overshoot under a clear sky. With the sky set to 0.25 before the run, the same gains are still swinging between 367 and 374 °C in the 18 min before that step, and the step back takes 69 min to settle, its 4.3 °C overshoot clipped by the high-temperature protection described below. Tripling the outlet loop gain to 7.5 settles the same step in 26.5 min, and no tested gain or hour of the day makes the loop unstable.

The bypass loop is the interesting structural case. It holds the steam generator inlet T_{sg}, read after the heater and the steam generator lag and shown as the loop's trend, by sending field fluid straight past the hot tank, and its process gain is proportional to how much hotter the field is than the tank. When the field is colder, opening the bypass makes the junction colder rather than hotter, so the gain changes sign. No fixed-action controller survives a sign change, so the loop is interlocked out whenever $T_{f,10} < T_h + 3\,^{\circ}\mathrm{C}$ or $\dot{m}_f < 15\,\mathrm{kg/s}$. While it is out the valve closes and the integral tracks the closed position, so a field cooling after sunset does not go on feeding the junction. A real plant interlocks it for exactly the same reason.

The level loop is an override rather than a regulator. It sits against zero for most of the day, because the tank level is below its setpoint and the direct-acting controller therefore demands negative defocus, which the output limit refuses. It wakes when the tank fills. A second demand comes from a high-temperature protection that ramps defocus linearly once the outlet passes $T_{max} - 4\,^{\circ}\mathrm{C}$, and the two are combined by a high selector,

$$f_d = \max\left(f_{d,L},\ \Pi_0^{100}\left[40\left(T_o - (T_{max} - 4)\right)\right]\right), \tag{6.20}$$

where T_o is the larger of $T_{f,10}$ and the thermocouple reading. The protection threshold sits deliberately above any setpoint the supervisor is allowed to choose. Placing it lower makes the two fight, because

the optimizer asks for the top of its range, the protection reads that as an excursion and defocuses part of the field, and the plant then collects less than it should while every loop reports success.

Pause. Try Exercise 6.2 before reading on.

6.5.3 Taking a Loop Out of Service

When the absorbed power falls below 2.5 MW the outlet temperature controller has no authority and no purpose. The application detects that condition, parks the setpoint on the measurement, and walks the field flow down to its floor at a rate limit rather than stepping it. Parking the setpoint matters for a reason that is about instrumentation rather than control. A setpoint left at 393 °C over a field that is at 180 °C after sunset puts a two hundred degree error on the trend and in the scorecard, and no action exists that would reduce it. The operator's own setpoint is kept behind the parked one, so a fixed 393 °C returns when the sun does. The same reasoning applies to the steam generator inlet setpoint, which is parked whenever the power block is not running.

Freeze protection is the other out-of-service mode and it runs in the opposite direction. When the coldest fluid node or the cold tank falls below 60 °C, electric heat tracing is energized until it is 8 °C above that, and the field flow floor is raised to 12 % of design so that the fluid keeps moving. The hot tank has its own tracing on the same rule, whose heat goes into whatever is left at its bottom. Both cost energy. The fluid freezes at 12 °C and a frozen header is an excavation rather than a restart, which is the assessment that kept synthetic aromatic in trough fields after molten salt was evaluated [11].

6.5.4 Model Predictive Control

A second outlet controller can take the field flow in place of the outlet loop of Table 6.4, within the general option of Section 1.6.3. It is a receding-horizon optimization [15] of the family developed for distributed collector fields at Almeria [7, 16], and it is an option rather than the default, since every scenario except the eleventh loads the proportional-integral loop. Nothing else changes. The supervisory setting still writes the outlet setpoint T_{SP} by the rules of Section 6.6, the division of labor proposed for trough plants in which an operating-point optimization sits above a predictive field controller [17], and the power, bypass and level loops, the power derate near an empty tank, the bypass interlock and the high selector of Equation (6.20) all stay in force over the plan just as over the loop. While the predictive controller is in charge the outlet loop tracks the applied flow at every scan with no preload, which makes switching between the two bumpless both ways.

Controller model. The plan predicts with the receiver lumped into five cells, each a fluid node and a metal node, written from Equations (6.4) to (6.7) with two plant nodes per cell, the declared defaults of Table 6.2 and the design weather (25 °C, 3 m/s, mirrors at 0.96), never the parameters of the running plant,

$$\frac{M_f}{5} c_p \frac{\mathrm{d}\hat{T}_{f,j}}{\mathrm{d}t} = 3600\big[\dot{m}_f c_p(\hat{T}_{f,j-1} - \hat{T}_{f,j}) + 2hA(\hat{T}_{m,j} - \hat{T}_{f,j})\big] \tag{6.21}$$

$$\frac{M_m}{5} c_{p,m} \frac{\mathrm{d}\hat{T}_{m,j}}{\mathrm{d}t} = 3600\Big[\frac{\hat{Q}_a}{5} - 2hA(\hat{T}_{m,j} - \hat{T}_{f,j}) - 2q_{loss}(\hat{T}_{m,j})\Big] \tag{6.22}$$

for $j = 1 \ldots 5$, with $\hat{T}_{f,0} = T_c$ from the inlet thermocouple and $\hat{Q}_a = A_{ap}\,\mathrm{DNI}\,\eta_{opt}\,(1 - \hat{f}_d)$. The irradiance is the pyrheliometer reading that the out-of-service test and the supervisor already use, and $\hat{f}_d$ is the defocus the high selector last commanded, passed through the declared 40 s drive lag. A cell holds 12 t of oil and 12.8 t of tube, so the transport delay is the declared inventory over the flow, $M_f/\dot{m}_f$, which is 4.5 min at 222 kg/s and 18 min at 55 kg/s. The delay therefore stretches with the manipulated variable by construction rather than by a schedule. Flow steps of $\pm 20\,\mathrm{kg/s}$ about 200 kg/s at clear noon give a gain of -0.436 on the plant against -0.423 on the cells, and steps

of $\pm 6\,\mathrm{kg/s}$ about 50 kg/s under a 285 W/m^2 noon give -1.644 against -1.589, all in °C per kg/s, so the cells read 2.9 and 3.3 % low. Plant and cells pass half of the response after about 37 000 kg of oil at both flows (37 440 and 36 720 kg on the plant), the signature of a delay that lives in fluid inventory rather than in seconds. The two step responses differ by 0.23 and 0.36 °C RMS at the two points, and over a sequence of flow steps and a cloud in manual mode the cells alone follow the outlet thermocouple to 1.35 °C RMS (4.30 °C at worst).

At each decision the cells are linearized numerically, by central differences, at the present estimate and at the flow at which the algebraic field of Equation (6.26) would hold $T_{SP} - b$ at the present irradiance, not about the applied flow, because the field's gain scales as $1/\dot{m}_f^2$ and its time scale as $1/\dot{m}_f$, so a model taken at a flow the plan is about to leave mispredicts the move it plans. The linearization is discretized with an exact zero-order hold, and its states are the ten cell temperatures and the previous planned flow.

Offset. An output bias removes steady error [18],

$$b_k = b_{k-1} + \left(1 - e^{-\Delta t_c/\tau_b}\right)\left(T_{out,k} - \hat{T}_{f,5} - b_{k-1}\right), \tag{6.23}$$

with $\tau_b = 72\,\mathrm{s}$, updated at every scan whichever controller acts, so a transfer to the predictive controller starts from a warm estimate, and held constant over the horizon. It absorbs what the cells omit, namely the true mirror cleanliness, the wind, the ambient temperature and the dispersion of ten nodes, and with it the open-loop gap quoted above falls to 0.28 °C RMS (1.01 °C at worst). The offset-free controller of Volume 9, Chapter 4 carries its offset in an estimated fuel bias instead.

Horizon and objective. A decision is taken every $T_s = 36\,\mathrm{s}$, five regulatory scans, over $N = 33$ decisions (20 min) grouped into twelve move blocks, four of one decision, three of two, two of three and one each of four, five and eight. The first move alone reaches the pump, and it stays there until the next decision. The plan minimizes

$$J = \sum_{k=1}^{N}\left(\hat{T}_k + b - T_{SP}\right)^2 + P\left(\hat{T}_{ss} + b - T_{SP}\right)^2 + \lambda K_u^2 \sum_{j=1}^{12} \Delta\dot{m}_j^2 + \rho s^2 + \sigma s \tag{6.24}$$

where $\hat{T}_k$ is the predicted outlet cell temperature, $\hat{T}_{ss}$ the outlet at which the linear model settles with the last planned flow held, K_u that model's steady gain from flow to outlet and $\Delta\dot{m}_j$ the change of flow at block j. A move is thus priced at λ times the square of the steady outlet shift it would cause, against one decision of squared error. The weights are $P = 10$, $\lambda = 0.02$, $\rho = 10^4$ per °C^2 and $\sigma = 10^3$ per °C. They were set once on the eleventh scenario and frozen: λ is the smallest value on a grid from 0 to 0.2 at which 1 °C of thermocouple noise raises the flow travel by no more than half (1.49 times at 0.02 against 1.67 at 0.01), and horizons of 10 to 40 min give 46.5 to 47.0 °C min of outlet error on that scenario, so the move price, not the horizon, limits the plan.

Constraints. The outlet is held under the threshold at which the protection of Equation (6.20) begins to defocus,

$$\hat{T}_k + b \le T_{max} - 4 + s, \qquad s \ge 0, \tag{6.25}$$

a soft constraint with one slack s for the whole horizon, on the run's own T_{max} as the protection reads it. The flow is held hard between the clamps of the outlet loop, 8 kg/s, or 26.6 kg/s (12 % of design) while freeze tracing is energized, and 520 kg/s. The defocus the level loop demands is held over the horizon, while the protection's share is not, because it releases as soon as the outlet returns under the threshold, which is what the plan is for. While the outlet reads less than the inlet plus ΔT_{rec} the field is still recirculating to the cold tank, and the flow is held on its floor, which is also where the loop leaves it. After that hold, until the outlet is within 2 °C of T_{SP} and one declared field inventory of 60 t

has been pumped, every planned move is also capped at the pump ramp the application declares for taking the field out of service, 600 kg/s per hour or 6 kg/s a decision. Uncapped, the first plan on every case that starts before sunrise drove the pump from 8 to 520 kg/s and back within a minute, pushing the oil warmed near the inlet past the oil the night had cooled at the outlet end. Capped, the flow stays at or below 110 kg/s in the first ten minutes after the hold on all seven such runs. Applied to every plan, the same cap raises the outlet error on the eleventh scenario from 46.8 to 1224.1 °C min, so it is confined to the start-up. Below 2.5 MW absorbed the plan stops and the field is taken out of service as described above.

Preview. The irradiance over the horizon enters the prediction as a known input. With the preview on it is the sky fraction the pyrheliometer reads now, persisted over the clear-sky profile of Equation (6.1) for the next 20 min, which is the rule the Forecast setting applies to its first hour. With the preview off the present irradiance is held. Neither carries a cloud edge, so every cloud in every scenario arrives unannounced, which is all a measurement at the field can offer.

Acceptance and fallback. The plan takes effect only if the solver returns an optimal status, every planned flow lies inside its bounds and inside the start-up cap where it applies, the first planned flow is the one applied, and its cost, slack and predicted outlets are all finite numbers. Otherwise the historian is marked, the outlet loop continues from its tracked state, and the next valid plan takes the flow back. Eight forced failures at 10:30 on the eleventh scenario, with the noise off, each changed the flow at the hand-over by 0.001 kg/s. The first plan is solved at $t = 0$, so a run started under the predictive controller is under it from the first logged row. In Chrome a plan costs a mean of 0.37 to 0.39 ms over the 350 decisions of the eleventh scenario (4.7 to 5.2 ms for the first, before the browser compiles it), which is 3.7 to 3.9 % of one core at one simulated hour per second and 18 to 19 % at five.

Three models kept apart. The plant is the thirty-state model of Section 6.3 at the run's parameters and weather. The identified model is a first-order-plus-dead-time fit to flow steps on that plant at 200 kg/s, a gain of −0.437 °C per kg/s, a dead time of 0.61 min and a lag of 3.19 min with a residual of 2.0 % of the response, which fixes only the retuned loop compared in Section 6.7.6. The controller's model is Equations (6.21) and (6.22) at the declared defaults, relinearized at every decision. A learner's change to any parameter is therefore a mismatch the bias must absorb, not knowledge the plan receives.

6.6 The Supervisory Layer

Three supervisory settings are available and the application exists for the difference between them. The first holds the operator setpoint. The second chooses the outlet temperature that maximizes net electrical output at the irradiance measured now. The third chooses it from the mean of the sky measured now and the sky forecast ten and twenty minutes ahead, and adds a schedule for the power block.

The objective is written once as algebra so that it can be evaluated without integrating the field. For a candidate outlet temperature T at inlet T_c, the absorbed power per node is known from Equation (6.3), the receiver metal sits above the fluid mean by whatever the conductance requires, and the flow follows from the energy that the ten nodes have to deliver,

$$\bar{T} = \tfrac{1}{2}(T_c + T), \quad T_m = \bar{T} + \frac{q_a}{hA}, \quad \dot{m} = \frac{10\,[q_a - q_{loss}(T_m)]}{c_p (T - T_c)}, \tag{6.26}$$

with q_{loss} the node loss of Equation (6.6). That flow is then passed through the power block map and the pump, and the objective is the net electrical output,

$$J(T) = P_e^{pb}\big(T, \min(\dot{m}, \dot{m}_{duty})\big) - P_{pump}(\dot{m}). \tag{6.27}$$

A golden-section search maximizes J over an interval in eighteen iterations, with no convergence test and no iteration count that depends on the data. The total cost of the supervisory layer is therefore

twenty evaluations of Equation (6.26) every three minutes, two to start the search and one per iteration. Any candidate that needs more than 2.2 times the design flow scores −99 MW, which marks it infeasible.

Two details of that search matter. The interval runs from 25 °C above the cold tank, or 312 °C if that is higher, to 393 °C, seven degrees below the fluid limit. A fixed lower end of 312 °C would ask for a rise of only 12 °C over a 300 °C cold tank, which at clear noon needs 1700 kg/s, more than three times the pump limit. Such probes score as infeasible and the search moves away from them, so in the present model a fixed lower end gives the same or a better answer at every condition tested (cold tank 295 to 310 °C, sky 0.25 to 1.0), and the margin binds only below about a third of clear sky, as Table 6.3 shows. The other detail is that the objective must be continuous. The plant has a genuine minimum stable load below which the power block produces nothing, and a cliff inside a golden section is not a search problem but a posing problem, because the objective stops being unimodal and the answer depends on where the first two probes landed. The optimizer is therefore handed a smoothed version of the same curve, gated over a band around the minimum load, while the plant keeps the discontinuity it actually has.

The chosen setpoint is rate limited before it reaches the loop, at a default of 500 °C/h. A supervisory setpoint that steps is a disturbance to the loop underneath it, and the loop underneath this one has four and a half minutes of transport delay.

The third setting adds a schedule. Twenty-eight half-hour blocks cover the remaining day, the net field power in each block is estimated from Equation (6.26) at the supervisor's current outlet temperature and the forecast sky (the sky measured now for the first hour), and the price is known. A merit order alone is not enough, because the tank holds 106.3 MWh and a clear day brings 370 to 380 MWh, so a schedule that saves everything for the evening fills the tank by early afternoon and then defocuses the rest off the mirrors. The schedule is therefore a causal forward pass that charges the tank block by block, runs the power block wherever the price clears a trial threshold, and counts whatever overflows,

$$\begin{aligned} \tilde{E}_k &= E_k + q_k \Delta \\ D_k &= \mathbb{1}[\pi_k \geq \theta]\, \min\!\left(Q_{PB}\Delta,\ \max(0, \tilde{E}_k)\right) \\ E_{k+1} &= \min\!\left(E_{cap},\ \tilde{E}_k - D_k\right) \end{aligned} \tag{6.28}$$

with the spill counted whenever the cap binds. Spill rises monotonically with the threshold θ, so fourteen bisections find the highest threshold that spills no more than 0.1 MWh, which is the dearest set of hours the tank can actually hold the energy for. The threshold plays the part of the water value of Section 1.4.6, found by forward simulation rather than by merit order. Nobody has to choose a threshold. The block runs whenever the price clears it, and also while a must-run rule holds, which latches when the hot tank comes within eight points of the level override setpoint and releases twelve points below it, so the turbine takes heat before the mirrors have to shed it. Measured on the evening scenario the threshold rises from $28.7 per megawatt-hour at noon to $205 at six in the evening, holding the block off in the cheap early afternoon, and then falls once the evening peak is being served.

6.7 Scenarios and Measured Results

Table 6.5 lists the twelve scenarios in selector order and gives the scorecard at the end of each window with the settings each case loads. Credit, Reset and reloading work as Section 1.6.4 describes, and the challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt.

6.7.1 A Cloud Against a Fixed Setpoint

The third scenario runs at eleven in the morning and drops the irradiance to a quarter at 11:48. The temperature loop does what it is told. It cuts the field flow from 199 kg/s to a mean of 33 kg/s between

Table 6.5 Scorecard at the end of each scenario window, as loaded

#	Scenario	Start	Window (h)	Generated (MWh)	Harvest (MWh)	Loss (MWh)	Defocus (MWh)	Margin ($)
1	Sunrise by hand	05:30	5	0.0	62.3	22.7	0.0	0
2	Outlet loop and its lag	11:00	3	29.1	136.6	11.9	0.0	847
3	Cloud on a fixed setpoint	11:00	4	40.0	135.3	15.9	13.8	1 205
4	Clear day	06:00	14	119.3	370.6	51.4	3.4	7 869
5	Broken day	06:00	14	25.4	111.0	48.2	9.4	1 168
6	Bypass and mixing	09:00	6	59.5	246.2	24.1	5.0	1 885
7	Full tank and defocus	09:00	8	32.0	135.9	30.1	184.9	1 138
8	Evening peak	06:00	18	80.6	261.6	53.0	0.1	6 413
9	Forecast error	08:00	14	61.6	186.2	43.5	3.5	5 680
10	Cold night	18:00	14	8.3	7.7	31.8	3.1	1 357
11	PI against MPC, passing clouds	09:30	3.5	35.0	111.3	13.8	0.0	1 086
12	Challenge	06:00	16	3.2	49.5	55.9	188.2	210

Table 6.6 Fixed against variable outlet temperature

Day	Strategy	Gen. (MWh)	Harvest (MWh)	Loss (MWh)	Pump (MWh)	Collect (%)
clear, 16 h	fixed	120.8	370.6	55.1	4.01	87.9
clear, 16 h	variable	121.1	377.8	51.8	6.25	88.9
broken, 14 h	fixed	25.4	111.0	48.2	0.39	70.5
broken, 14 h	variable	33.5	129.9	39.3	2.71	78.1

The strategy gives the same run whether it is selected before Run or carried through Reset.

12:00 and 12:36, resting on its 8 kg/s floor for nine minutes, and after a 25 °C dip at the cloud edge the outlet comes back within 4 °C of 393 °C 18 min into the cloud and stays within 4.3 °C of it until the clearing. Meanwhile the receiver loses 3.75 MW of the 11.9 MW it absorbs, nearly a third. The clearing at 12:36 is harder than the cloud, because the integral has to carry the flow from about 30 back toward 200 kg/s. The outlet peaks at 399.95 °C and rides between 395 and 400 °C until 13:26 while the high-temperature protection cycles the mirrors in and out. That protection accounts for all 13.8 MWh defocused, because the tank never passes 58 %. Over the four hours the scorecard shows 15.9 MWh of receiver loss against 135.3 MWh harvested, and the collection efficiency, which is heat into the fluid over heat absorbed by the receivers, at 89.1 %. The instructive part is where to look. The temperature trend says the controller succeeded and the energy trend says the plant did not.

> **TRY IT IN THE APP** **3 · What a Cloud Does to a Plant Defending a Temperature**
>
> Load scenario 3 and press Run in Auto on Fixed T. At 11:48 the sky falls to a quarter. Watch Field flow on the outlet strip fall from about 199 kg/s and rest on its 8 kg/s floor, while the outlet returns to within 4 °C of 393 °C. Then compare receiver loss with absorbed on the Where the sunlight goes strip: about 3.75 of 11.9 MW.

6.7.2 Clear Day Against Broken Day

The comparison the application is built around runs the identical day twice, once holding 393 °C and once letting the supervisor choose. Table 6.6 gives both days, and Figure 6.2 draws the same numbers as bars.

On a clear day the supervisor is worth nothing. It harvests 1.9 % more heat and loses 5.9 % less from the receivers, and it spends almost the entire difference on pumping, because the lower

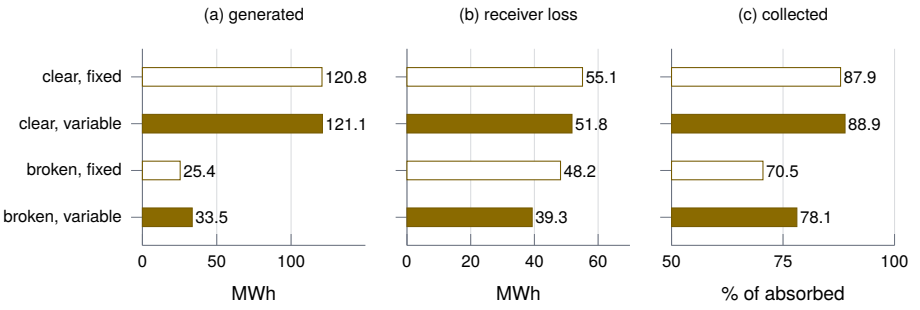

Figure 6.2 Table 6.6 as a chart: (a) net generation, (b) receiver loss and (c) the share of absorbed heat that reaches the fluid, for the fixed setpoint (white bars) and the variable setpoint (solid bars) on the clear day of 16 h and the broken day of 14 h. The axis of (c) starts at 50 %, so its bars show differences, not ratios. Notice that the variable setpoint cuts the loss on both days but raises generation only on the broken day.

outlet temperature it selects before 10:00 and after 16:00 (as low as 321 °C, while it holds 393 °C in between) requires more flow and pumping power rises as the cube of the flow. Net generation over sixteen hours is 121.07 MWh against 120.77 MWh, a difference of 0.3 % that is inside any reasonable uncertainty on the model. That is the correct result and it is worth reporting. A supervisory optimizer that reproduces the operator setpoint when the operator setpoint is right is a working optimizer, and a study that reported only the cloudy-day number would be reporting half a result.

> **DID YOU KNOW? More Heat, Almost No More Power**
>
> From the simulator: on the clear day the variable setpoint harvested 377.8 MWh against 370.6 MWh for the fixed one, but pumping took 6.25 MWh against 4.01 MWh, and net generation came to 121.1 MWh against 120.8 MWh (Table 6.6).

On a broken day the same optimizer is worth 32 %. It drops the outlet setpoint to about 318 °C during the cloudy stretches, the bottom of its interval, which keeps the flow near 134 kg/s rather than 30 kg/s and the receiver cooler, and the receiver loss falls from 48.2 MWh to 39.3 MWh at exactly the time the field can least afford it. Heat into the fluid rises by 17.0 % and net generation by 32.1 %, because the generation gain compounds the harvest gain through the part-load efficiency curve. The fraction of absorbed heat reaching the fluid goes from 70.5 % to 78.1 %. Published studies of this configuration report the same shape of result for the same reason [1, 2].

Two observations about the mechanism are worth making explicit. The optimizer is not clever, being a golden section on one scalar over an algebraic model. And it never touches an actuator. The ordinary temperature controller does all the work underneath, and it would have done exactly the same work for a wrong setpoint.

6.7.3 The Schedule Is Worth More Than the Setpoint

The eighth scenario runs a hazy day (sky 0.72) for eighteen hours from six in the morning, with the evening price peak at 19:30, after sunset. Under a fixed setpoint and a constant power demand the plant generates 78.1 MWh and books $2 875. Under the variable-temperature supervisor it generates 79.6 MWh and books $3 081. Under the schedule it generates 80.6 MWh, about the same, and books $6 413, which is 2.2 times as much. The schedule holds the power block off from 11:58 to 14:31, runs at rating through the late afternoon under the must-run rule (stopping only from 15:44 to 15:55) while the tank fills to 85 % by 17:14, lets the rule release at 18:21, and then runs into the evening peak from 19:06, when the price is $200 per megawatt-hour, until 21:11. Of its generation, 39.2 MWh

comes after 16:00, against 8.7 MWh under the fixed setpoint. On a clear sky the same schedule books only 1 % more than constant power ($8 208 against $8 112), because a field with a solar multiple of 1.83 and 3.6 h of storage has to run the block most of the day or overflow.

That ordering is the practical lesson of the application. Choosing the temperature is worth a few percent of energy. Choosing when to convert the energy is worth most of the revenue, and the two decisions are independent enough that a plant can do either without the other.

6.7.4 Defocus as an Override

The seventh scenario starts at nine in the morning with the tank nearly full and the power block asked for only 4 MW. Over eight hours the plant throws away 184.9 MWh, which is more than the 135.9 MWh it harvests. Every controller is holding its setpoint. There is simply nowhere to put the heat, so the level override takes the mirrors off the sun, and defocusing at noon on a plant sized for a whole day is normal operation rather than a fault. The setpoint that decides how much gets thrown away is the level override setpoint rather than any tuning constant. Raising it from 88 % to 96 % reduces the defocused energy from 189.5 MWh to 179.7 MWh over ten hours by giving the tank more room, and it leaves the drive less room to act before the tank is genuinely full. The power setpoint matters more. Asking the block for 10 MW instead of 4 MW over the same eight hours cuts the defocused energy to 69.4 MWh and raises generation from 32.0 to 80.0 MWh.

> **TRY IT IN THE APP** **7 · A Full Hot Tank, and Defocus as an Override**
>
> Load scenario 7 and run its eight hours in Auto with the Electrical output setpoint at 4 MW, then read Defocused: 184.9 MWh. Reset, set the setpoint to 10 MW before pressing Run, and run again. Defocused falls to 69.4 MWh and Generated rises from 32.0 to 80.0 MWh.

6.7.5 The Graded Challenge

The twelfth scenario opens sixteen hours from six in the morning with weak tuning, a fixed setpoint, mirrors ten percent dirty, the bypass in manual, the level override setpoint at 96 % and the power block asked for 10 MW from the start. To pass, one 16 h Auto run with the case's parameters and disturbances must generate more than 60 MWh, deliver more than 78 % of the absorbed heat into the fluid, defocus less than 20 MWh, book a margin above $4 500, and keep T_{sg} within 60 °C of its setpoint over the last third of the day, from 16:34. The default configuration generates 3.2 MWh, collects 49.8 %, defocuses 188.2 MWh, books $210 and holds T_{sg} within 5 °C of its setpoint, failing four of the five parts. The suggested answer is two changes, the schedule and the tuned gains, and with them the day generates 72.8 MWh, collects 84.1 %, defocuses 5.1 MWh, books $6 780 and holds T_{sg} within 29 °C. It passes every part, the narrowest being collection at 6.1 points above its threshold, and Reveal followed by Reset gives the same run. Together they are worth 69.6 MWh and $6 570, but they are worth different things. The tuned gains alone, inside the default structure, generate 68.7 MWh with 81.9 % collection and 11.2 MWh defocused, yet book only $2 663. The schedule with the weak gains books $4 751 but generates only 41.1 MWh and defocuses 75.0 MWh. The gains recover the energy and the schedule earns the money. Under the schedule the power loop tuning matters as much as the outlet loop tuning, since retuning the outlet loop alone gives 50.3 MWh and fails. The bypass in automatic, a 90 % level setpoint and a 388 °C inlet setpoint do not raise the score, and the lower level setpoint costs about $300, because the must-run rule then sells more of the day's heat in the afternoon (33.9 MWh against 31.5 MWh) and less after 18:00 (22.0 MWh against 24.7 MWh).

6.7.6 Predictive Control Against the Outlet Loop

Table 6.7 runs nine cases under three outlet controllers, each from $t = 0$ with the case's own settings, seed and duration, the controller chosen before the run and the supervisory setting the case loads (scenarios 4 and 5 also on Variable T, as their second runs ask). The shipped loop keeps its shipped tuning, and the challenge is run both as shipped and with its suggested settings. The predictive

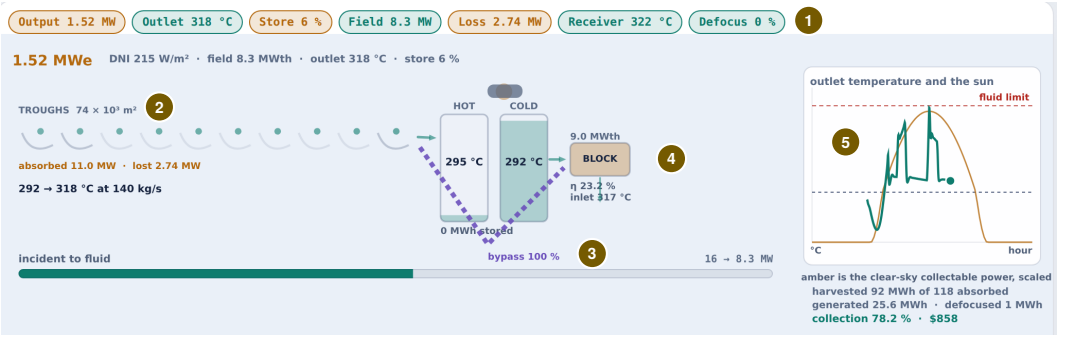

Figure 6.3 The application at 15:05 on the fifth scenario, nine hours into the broken day, with Variable T selected before the run. At 215 W/m^2 of direct normal irradiance the optimizer has selected 318 °C rather than 393 °C, and the field flow is holding at 140 kg/s rather than collapsing to defend a temperature the afternoon cannot support. Receiver loss is 2.74 MW against 11.0 MW absorbed, and the run has delivered 92 MWh of the 118 MWh absorbed into the fluid, a collection efficiency of 78.2 %. The hot tank is held at its lowest working level of 6 %, the consequence of an afternoon at a quarter of clear-sky irradiance. The view is cropped to the plant. Callouts: (1) the badge row with output, outlet, store, receiver loss and defocus, (2) the trough row with the absorbed power and the receiver loss, (3) the hot and cold tanks with the bypass carrying the whole block flow, (4) the power block with its efficiency and inlet temperature, and (5) the outlet temperature against the clear-sky collectable power, with the fluid limit marked.

controller runs with the eleventh scenario's settings on every case, with no retuning. The third row is the stronger of two classical structures given the same information, whichever has the lower outlet error on the case. One is the shipped loop retuned once by the SIMC rule on the identified model ($K_c = 6.0$, $\tau_I = 0.053$ h). The other is the shipped loop trimming a static feedforward, the flow that Equation (6.26) says holds the setpoint at the present pyrheliometer reading with the level loop's defocus held, started on the feedforward's own flow as the predictive controller starts on its own plan. Neither is in the application. Outlet error is the scorecard's IAE, the excursion columns count time and integrated degrees above the protection threshold $T_{max} - 4$ °C, and flow travel sums the absolute changes of the flow command. Defocused energy is reported rather than a count of defocus episodes, because the protection's command chatters across its first percent and a count depends on the counter.

The eleventh scenario is the case built for the comparison. It opens at 09:30 under a clear sky on Fixed T at 393 °C, three degrees under the protection threshold, and two clouds cross the field, a quarter of the sun from 10:00 to 10:48 and 0.45 of it from 11:36 to 12:00. The shipped loop answers the first cloud only as its error grows, so the outlet falls to 371.8 °C, and the integral then carries the flow down to its 8 kg/s floor after the outlet has already turned. The outlet climbs back to 396.8 °C inside the cloud, and when the cloud clears the flow is far too low, as shown in Figure 6.4. At 10:51 the loop's outlet peaks at 399.6 °C on 52.8 kg/s, while the predictive controller holds 394.9 °C on 142.7 kg/s, and 36 s later the protection has defocused 88 % of the loop's mirrors while the other run's mirrors are fully focused. Over the morning the loop spends 62.7 min above the threshold and defocuses 17.85 MWh, while the predictive controller keeps the outlet between 391.3 and 395.1 °C from the first cloud on. Net electricity is the same in both runs (34.99 against 34.98 MWh), because the block runs at its setpoint from the tank, and the difference is heat kept, with 37.8 MWh stored at 13:00 against 21.0 MWh.

Against the shipped loop the predictive controller lowers the outlet error on every case, by 11 % on the clear day at a fixed setpoint (4140.8 against 4666.4 °C min) and by 95 % on the eleventh scenario, and it spends no minute above the threshold on any case but the broken day on Variable T, where

Table 6.7 Outlet controllers compared on nine cases

Case (window)	Outlet controller	IAE (°C min)	Above 396 (min)	Above 396 (°C min)	Defocused (MWh)	Net (MWh)	Flow travel (kg/s)	Plans, rejected
11, passing clouds, Fixed T (3.5 h)	shipped PI	882.8	62.7	60.9	17.85	34.99	865	—
	MPC	46.8	0	0	0	34.98	622	350, 0
	MPC, preview off	47.1	0	0	0	34.98	622	350, 0
	PI with feedforward	**45.7**	0	0	0	34.98	574	—
3, one cloud, Fixed T (4 h)	shipped PI	624.8	44.0	49.1	13.84	39.99	616	—
	MPC	33.8	0	0	0	39.99	417	400, 0
	PI with feedforward	**31.0**	0	0	0	39.99	387	—
5, broken day, Fixed T (14 h)	shipped PI	8312.3	86.3	58.9	9.41	25.39	1050	—
	MPC	**7222.5**	0	0	0	**28.74**	999	1073, 0
	PI with feedforward	7461.9	11.4	11.6	0.89	28.51	905	—
5, broken day, Variable T (14 h)	shipped PI	5864.7	1.8	3.2	0.58	33.52	1870	—
	MPC	**5244.5**	1.4	1.0	0.19	32.40	6153	1073, 0
	PI retuned (SIMC)	5370.5	0	0	0	**33.61**	3277	—
4, clear day, Fixed T (14 h)	shipped PI	4666.4	18.2	29.2	3.43	119.34	625	—
	MPC	**4140.8**	0	0	0	**120.27**	737	1144, 0
	PI with feedforward	4304.9	14.9	25.9	2.87	119.57	654	—
4, clear day, Variable T (14 h)	shipped PI	3268.9	1.4	0.4	0.09	118.28	1252	—
	MPC	2805.9	0	0	0	**118.65**	4186	1144, 0
	PI with feedforward	**2701.7**	1.7	1.3	0.30	118.29	2027	—
9, forecast error, Forecast (14 h)	shipped PI	1826.6	31.9	18.9	3.55	61.58	1297	—
	MPC	**1020.0**	0	0	0	**63.13**	3935	1062, 0
	PI retuned (SIMC)	1260.9	10.0	5.1	0.96	62.85	2069	—
12, challenge as shipped (16 h)	shipped PI, weak tuning	7417.4	628	995.4	188.22	3.21	302	—
	MPC	**4326.7**	0	0	0	**72.78**	991	1136, 0
	PI retuned (SIMC)	4726.7	23.7	27.5	3.70	71.49	1130	—
12, challenge, suggested (16 h)	shipped PI, tuned	3593.3	28.6	19.0	5.11	72.85	1676	—
	MPC	**2839.7**	0	0	0.42	**74.63**	4439	1136, 0
	PI with feedforward	2896.1	10.7	9.2	2.77	73.55	2675	—

Bold marks the lowest outlet error and the highest net electricity of each case where the rows differ. The controller is selected before Run, and selecting it later and pressing Reset reproduces the same run. The defocused energy of the predictive controller on the suggested challenge is the level loop's, with no minute above the threshold.

it spends 1.4 min there against 1.8 min for the shipped loop. None of the 8 868 plans in the table was rejected. The comparison with the classical row says where that lead comes from. It is not the preview, because switching the preview off changes the outlet error by at most 0.64 % (on the eleventh scenario) and net electricity by at most 0.04 MWh, and it is not the receding horizon as such, because the loop trimming a static feedforward from the same pyrheliometer, started on its own flow, beats the predictive controller over the whole run on scenarios 11 and 3 (45.7 against 46.8 and 31.0 against

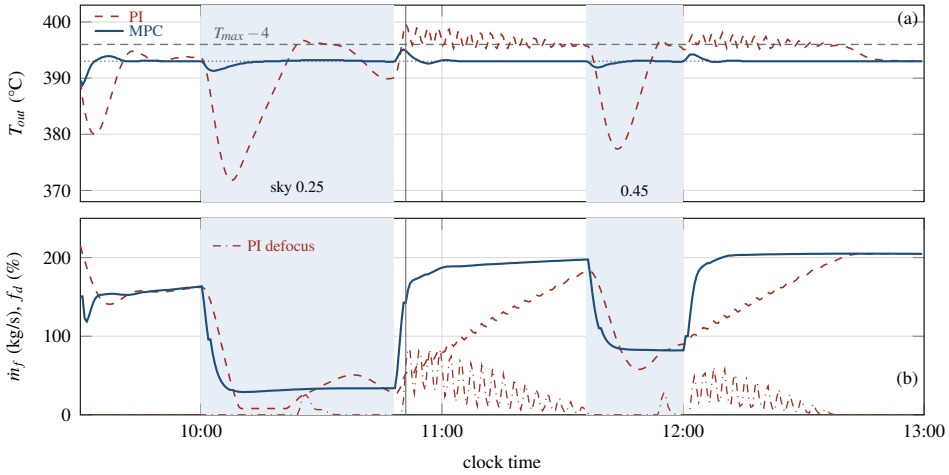

Figure 6.4 The eleventh scenario from 09:30 to 13:00 under the shipped outlet loop (red, dashed) and under the predictive controller (blue, solid), with the two clouds shaded and the instant 10:51 marked. (a) Outlet temperature against the 393 °C setpoint (gray, dotted) and the 396 °C protection threshold (gray, densely dashed). (b) Field flow under both controllers and the defocus the protection applies under the loop (red, dash-dotted), which reaches 88 % just after the first cloud clears.

33.8 °C min) and inside the clouds (22.3 against 34.9 and 15.0 against 22.8 °C min). The measured sun and a model of the field carry most of the result, whichever structure uses them. The predictive controller wins in two places. The first is the end of the morning approach, where on the clear day it begins raising the flow while the outlet is still about 80 °C below the setpoint and arrives without crossing the threshold, while the loop waits on its floor for the setpoint and overshoots to 400.0 °C, the fluid limit. The second is the cloud windows of scenarios 5, 9 and the challenge, where the outlet runs close to the threshold and the feedforward with its trim overshoots it for 10.1 to 11.4 min on the broken day at a fixed setpoint, on scenario 9 and on the suggested challenge, while the predictive controller holds the ceiling of Equation (6.25). Inside the clouds its error is 99.7 against 137.7 on the broken day, 919.2 against 1173.1 for the retuned loop on scenario 9, 859.7 against 1083.1 on the suggested challenge and 98.1 against 375.9 °C min for the retuned loop on the challenge as shipped.

🕮 **Pause.** Try Exercise 6.3 before reading on.

The predictive controller loses three cases on outlet error to the loop with feedforward, scenarios 11 and 3 as above and the clear day on Variable T (2805.9 against 2701.7 °C min), where the start-up cap slows the approach to the low early setpoints. It loses net electricity on the broken day on Variable T, 32.40 MWh against 33.52 MWh for the shipped loop and 33.61 MWh for the retuned one, with the same heat into the fluid (129.95 against 129.88 MWh). Of the 1.12 MWh difference, 1.05 MWh is lower gross output and 0.06 MWh is extra pumping (0.08 MWh more in the field and 0.02 MWh less at the block). The gross loss is thermal duty through the block (132.92 against 135.91 MWh at a conversion efficiency of 0.287 against 0.288), and the duty went into the hot tank. From the morning approach on, heat collects in the tank, and a tank kept warm leaves the field within 3 °C of it, so the interlock holds the bypass out, while under the loop the tank stays cooler and trips it less often. The bypass carries 40.1 % of the block's flow under the predictive controller against 88.4 % under the loop, 52.7 MWh enters the hot tank against 10.1 MWh, 0.30 MWh more leaks through its wall, and 2.79 MWh of the 2.99 MWh is still in the tanks when the run ends. Wherever the setpoint moves the predictive controller also moves the flow 2.6 to 3.3 times as far as the shipped loop. On the clear day on Variable T the loop with its feedforward read one transport delay ahead from the same preview makes 118.55 MWh, less than the predictive controller's 118.65 MWh.

plant control system: the four loops every 7.2 s, the supervisor every 3 min. Field pump lag 3 s, block pump and bypass 5 s, mirror drive 40 s. Thermocouple 2.2 s, power transducer 1.8 s.

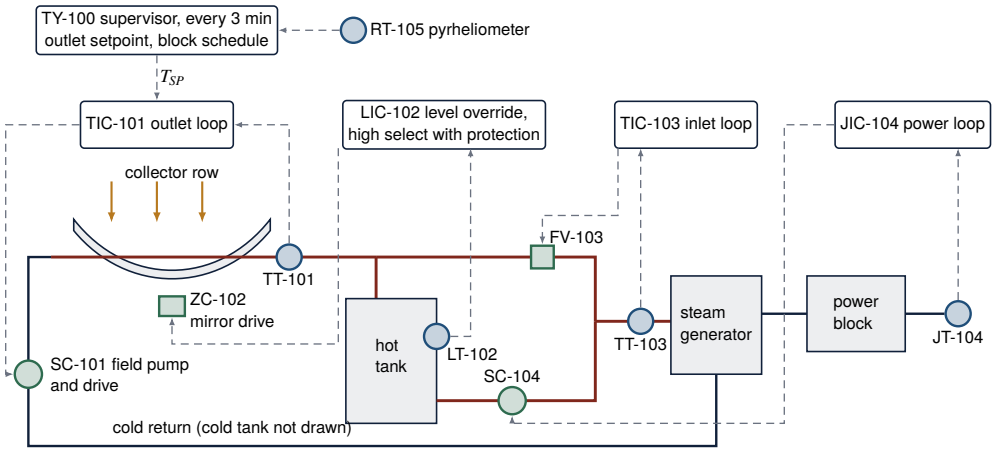

Figure 6.5 The four loops and the supervisor drawn with their hardware. The outlet thermocouple TT-101 feeds the outlet loop TIC-101, which sets the speed of the field pump SC-101, and the supervisor TY-100 writes that loop's setpoint from the pyrheliometer RT-105. The level LT-102 drives the mirror drive ZC-102 through the override LIC-102, high-selected with the outlet protection. TT-103 holds the steam generator inlet through the bypass valve FV-103, and the power transducer JT-104 sets the block pump SC-104. Notice that the supervisor moves no actuator: it writes one number into an ordinary loop.

The challenge's twentyfold difference, 72.78 against 3.21 MWh, belongs to the challenge's weak tuning, not to the predictive controller. With the weak gains the protection pins the outlet at 397 to 399 °C all day, the loop sees an error of only 4 to 6 °C and raises the flow by about 2 kg/s an hour, and 188.22 MWh is defocused. The shipped loop with only its outlet gains retuned makes 68.64 MWh, and the SIMC retune makes 71.49 MWh.

6.8 Technology in Practice

The loops of Section 6.5 read a temperature, a level and a power and move a pump, a valve and the mirrors, and the supervisor of Section 6.6 reads the sun. Figure 6.5 draws them with that hardware in place of the schematic's symbols.

What Measures It

The collector outlet temperature is the variable the whole chapter is about. A modern loop takes oil in at about 292 °C and returns it at about 393 °C, and some collectors in the loop are defocused when the outlet would otherwise pass about 395 °C [5]. The outlet is read by a temperature element in a thermowell. A pyrheliometer measures the direct normal irradiance, and a plant should have at least two or three instruments measuring it, checked often for alignment and cleanliness [5]. The chapter's thermocouple is a 2.2 s lag on the last fluid node, and its pyrheliometer reads the true irradiance with no error.

> **TECHNOLOGY** **Sensor: A Thermocouple at the Loop Outlet**
>
> A thermocouple is a junction of two dissimilar metal wires that produces a voltage approximately proportional to its temperature [19]. The common type K is listed for 32 to 2300 degrees Fahrenheit. In a pipe the element sits in a thermowell, which protects it but adds a first-order lag, and a lag severe enough can compromise the stability of feedback control [19]. The most common failure is to fail open, or burn out. An open input

can pick up electrical noise that reads as a wildly varying temperature, so the transmitter is set to read high or low when the element burns out [19], and the trend jumps to that end of its scale. The chapter's outlet thermocouple is a 2.2 s first-order lag with a noise setting of 0 to 4 °C.

What Moves It

The pump speed sets the field flow, which is also hard to measure on a hot oil loop. Plants have estimated it from the drive speed, the motor current or the pressure drop when their flow meters proved unreliable, at an accuracy typically not nearly as good as a working meter [5]. The mirrors are the second final element. The collector drive turns a trough partly or fully off the sun, and partial defocus of every collector in a loop by the same amount has worked well for outlet temperature control at a number of plants [5]. The chapter's defocus drive follows its command through a 40 s lag, and its field flow is known exactly.

TECHNOLOGY **Final Element: The Field Pump and Its Drive**

The field flow is set by the speed of the oil pumps. A variable-frequency drive sets the speed of an alternating-current motor by varying the frequency of the power that feeds it [19]. A study of operating plants recommends this capability on every oil pump, after one plant without it damaged its flow-control valves and shook its piping at each pump start [5]. Pump makers set limits on minimum flow, maximum flow and minimum speed, very low flow brings vibration, bearing wear and flow oscillations, and the failure reported most often is the pump seal [5]. The chapter's pump follows its command through a 3 s lag between 8 and 520 kg/s.

Where the Controller Runs

A trough plant is run from a distributed control system [5]. In that architecture a supervisory controller sends setpoints to the field controllers and requests data from them [20], and supervisory and lower-level loops run with cycle times from milliseconds to minutes [20]. The chapter's four loops execute every 7.2 s, and its supervisor every three minutes. The supervisor holds the chapter's lesson. It writes the outlet setpoint into an ordinary outlet loop, which does the same work for a right setpoint as for a wrong one. Tuning such a system takes time. Automation of oil flow control deserves higher priority, and tuning time should extend beyond commissioning to capture the plant's starts, stops and transients [5].

IN PRACTICE **Keep the Flow Up Under Cloud**

Operators report that loops at low flow under concentrated sun are where receivers break. Partial defocus used to hold the outlet temperature may put concentrated light on a receiver's glass-to-metal seals, and on a partly cloudy day a field whose flow is cut to drive up the outlet temperature may run some loops at low flow in full sun [5]. The recommended practice is a higher minimum flow, well into the turbulent region, also on partly cloudy days. One plant that holds about 50 % of design flow lost only 0.3 % of its receivers in its first ten years [5]. The chapter's outlet loop rests on its 8 kg/s floor in the cloud of scenario 3.

What Is Modernizing It

Three developments touch this chapter. Flow estimated from pump speed, current and pressure drop is a soft sensor that some plants already use in place of a failing meter, and two plants with newer automatic flow control schemes were pleased with them [5]. Performance models are asked to move from hourly to sub-hourly irradiance data, because hourly data do not capture the short cloud transients that a field feels [5], the transients the supervisor of this chapter answers. And as Ethernet, Internet Protocol and wireless devices replace older proprietary networks, the likelihood of cybersecurity vulnerabilities and incidents rises [20], which matters for a supervisor that writes setpoints into the field.

6.9 Limitations

The optical and receiver models are analytic surrogates rather than fitted correlations. The shape of the loss curve in Equation (6.6) drives every conclusion in this chapter, and a constant emissivity understates the loss at the top of the temperature range relative to the NREL receiver model [13]. The direction of that error is favorable to the fixed setpoint, so the reported advantage of variable-temperature operation is conservative, but no claim is made that the percentages would survive recalibration against a measured collector.

> **OPEN CHALLENGE** **Hydrogen in the Receiver Vacuum**
>
> A receiver tube sits inside an evacuated glass envelope. The oil slowly breaks down and gives off hydrogen, which permeates the steel into the vacuum. At a partial pressure of about 1 mbar the receiver's heat loss rises by a factor of 3 to 4, and the damage starts at the hottest receivers and works back along the loop over years [5]. Scrubbing hydrogen from the oil, injecting argon into the receivers and lower maximum operating temperatures are being considered, and the hydrogen in the oil is very difficult to measure [5]. The chapter holds one emissivity at all temperatures. Progress would be a measured receiver loss that the supervisor could weigh against the gain of running hot.

One collector row stands for the entire field. A real field has hundreds of parallel loops, a flow distribution across them, and a hottest loop that sets the defocus decision for all of them. Nothing in this model can pose the loop balancing problem, which is a substantial part of real trough plant operation.

The power block is a steady-state map with a nine-second inlet lag. Start and stop transients, drum level, superheater metal and turbine thermal stress are all absent, so the model understates the cost of the cycling that the schedule in Section 6.6 produces. The application counts a power block start in its exported data only after at least fifteen minutes off, and does not charge for starts in the margin. The eighth scenario records 3 starts in eighteen hours, while the block crosses its minimum stable load 22 times as the power loop works near it.

The schedule optimizes a threshold rather than a trajectory. It answers whether to run the block at rating in a given block, not at what load, so a plant that would be better served by running at part load for longer is not represented. The scheduler enforces tank capacity but makes no terminal-inventory guarantee. The application has no reserve setting, because holding energy back in trial draws does not impose a floor on actual discharge. A guaranteed reserve would require an operating constraint and a different scheduling objective.

The weather and the price are synthetic. The improvement numbers should be read as the shape of a result rather than as its magnitude, and any statement about what dynamic optimization is worth on a particular site requires that site's irradiance record and price curve.

The predictive outlet controller predicts with five cells at the declared design values, where the plant has ten nodes and whatever weather and parameters the run carries. It cannot see the true mirror cleanliness, the wind, the ambient temperature or a parameter the learner changes, it disperses the transit more than ten nodes do, and the output bias of Equation (6.23) carries all of that as one slowly varying offset. Its start-up cap is a rule rather than a model of the pump, borrowed from the ramp at which the field is taken out of service, and on the clear day at a variable setpoint it costs enough outlet error to lose that case to a feedback loop. Finally, the shipped loop is not the honest competitor. A proportional-integral loop trimming a static feedforward from the same pyrheliometer matches or beats the predictive controller through passing clouds, so a plant weighing the two should compare against that loop rather than against plain feedback.

OPEN CHALLENGE A Plant That Runs Itself

Trough plants are not yet automated end to end. A 2020 study found that their operating modes had not been refined to the point where the control system can run a plant automatically on a routine basis, so operators often write procedures as they go when equipment misbehaves [5]. Field flow is hard to automate because flow meters on the hot oil have been unreliable, and estimates from pump speed, motor current and pressure drop are less accurate than a working meter [5]. The chapter's supervisor trusts every measurement it reads. Progress would be a supervisory layer that keeps choosing a sound setpoint when a flow or temperature signal is lost.

6.10 Conclusion

A parabolic-trough field is a plant whose best operating point moves with a disturbance that anybody can see. Receiver loss depends on the fourth power of a temperature the operator chooses and not at all on the sunshine, while the heat available to collect changes by a factor of five across a broken day. The consequence is that the outlet temperature setpoint is a decision that has to be made repeatedly, and a control system with no mechanism for making it will defend a wrong number perfectly.

The measured comparison has two halves and both matter. Holding 393 °C on a clear day costs 0.3 % of net generation against an optimizer, because the extra harvest is spent on pumping. Holding it on a broken day costs 32 %, because the receiver goes on losing 3.7 MW while the field absorbs a fifth of what it does at clear noon. Above both sits a scheduling decision worth more than either. On a hazy day the schedule generates about what constant power does (80.6 against 78.1 MWh) and books 2.2 times the margin, by moving 39 MWh of generation past four in the afternoon. The regulatory loops underneath are ordinary, but their tuning is not a detail, because in the graded day the gains recover most of the energy while the schedule earns the money. An optional predictive outlet controller on a five-cell receiver model removes all but 1.4 min above the protection threshold on the nine measured cases, while a loop given a feedforward from the same pyrheliometer does as well through passing clouds, so the gain belongs to the measured sun and the model more than to the receding horizon.

KEY POINTS

- The best temperature for the oil leaving a solar field moves with the sky, so a controller can hold a fixed target perfectly and still be wrong.
- Holding 393 °C costs 0.3 % of net generation on a clear day and 32 % on a broken day.
- On a hazy day the schedule books 2.2 times the margin of constant power for about the same energy.
- A loop with feedforward from the same pyrheliometer does as well as the predictive controller through passing clouds.

Where else this applies. A supervisory setpoint that follows the weather also sets the chilled water of Volume 9, Chapter 8 and the aeration of Volume 9, Chapter 10.

References

[1] K. M. Powell, J. D. Hedengren, and T. F. Edgar, "Dynamic optimization of a solar thermal energy storage system over a 24 hour period using weather forecasts," in *2013 American Control Conference*, 2013, pp. 2946–2951.

[2] ——, "Dynamic optimization of a hybrid solar thermal and fossil fuel system," *Solar Energy*, vol. 108, pp. 210–218, 2014.

[3] H. Price, E. Lüpfert, D. Kearney, E. Zarza, G. Cohen, R. Gee, and R. Mahoney, "Advances in parabolic trough solar power technology," *Journal of Solar Energy Engineering*, vol. 124, no. 2, pp. 109–125, 2002.

[4] A. Fernández-García, E. Zarza, L. Valenzuela, and M. Pérez, "Parabolic-trough solar collectors and their applications," *Renewable and Sustainable Energy Reviews*, vol. 14, no. 7, pp. 1695–1721, 2010.

[5] M. Mehos, H. Price, R. Cable, D. Kearney, B. Kelly, G. Kolb, and F. Morse, "Concentrating solar power best practices study," National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5500-75763, 2020.

[6] E. F. Camacho, F. R. Rubio, M. Berenguel, and L. Valenzuela, "A survey on control schemes for distributed solar collector fields. Part I: Modeling and basic control approaches," *Solar Energy*, vol. 81, no. 10, pp. 1240–1251, 2007.

[7] ——, "A survey on control schemes for distributed solar collector fields. Part II: Advanced control approaches," *Solar Energy*, vol. 81, no. 10, pp. 1252–1272, 2007.

[8] E. F. Camacho, M. Berenguel, and F. R. Rubio, *Advanced Control of Solar Plants*, ser. Advances in Industrial Control. London: Springer London, 1997.

[9] D. Kearney, U. Herrmann, P. Nava, B. Kelly, R. Mahoney, J. Pacheco, R. Cable, N. Potrovitza, D. Blake, and H. Price, "Assessment of a molten salt heat transfer fluid in a parabolic trough solar field," *Journal of Solar Energy Engineering*, vol. 125, no. 2, pp. 170–176, 2003.

[10] D. Cocco and F. Serra, "Performance comparison of two-tank direct and thermocline thermal energy storage systems for 1 MWe class concentrating solar power plants," *Energy*, vol. 81, pp. 526–536, 2015.

[11] D. Kearney, B. Kelly, U. Herrmann, R. Cable, J. Pacheco, R. Mahoney, H. Price, D. Blake, P. Nava, and N. Potrovitza, "Engineering aspects of a molten salt heat transfer fluid in a trough solar field," *Energy*, vol. 29, no. 5-6, pp. 861–870, 2004.

[12] K. M. Powell and T. F. Edgar, "Modeling and control of a solar thermal power plant with thermal energy storage," *Chemical Engineering Science*, vol. 71, pp. 138–145, 2012.

[13] R. Forristall, "Heat transfer analysis and modeling of a parabolic trough solar receiver implemented in engineering equation solver," National Renewable Energy Laboratory, Golden, CO, USA, Tech. Rep., 2003.

[14] F. Burkholder and C. Kutscher, "Heat loss testing of Schott's 2008 PTR70 parabolic trough receiver," National Renewable Energy Laboratory, Golden, CO, USA, Tech. Rep., 2009.

[15] E. F. Camacho and C. Bordons, *Model Predictive Control*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2007.

[16] A. J. Gallego and E. F. Camacho, "Adaptative state-space model predictive control of a parabolic-trough field," *Control Engineering Practice*, vol. 20, no. 9, pp. 904–911, 2012.

[17] E. F. Camacho and A. J. Gallego, "Optimal operation in solar trough plants: A case study," *Solar Energy*, vol. 95, pp. 106–117, 2013.

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

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

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

6.11 Exercises

6.1 ○ Concept check · 2 min

At solar noon with a 300 °C cold tank, Table 6.3 puts the outlet temperature that maximizes net electrical output at 393 °C under a clear sky and at 325 °C at a quarter of it. Which explanation is right?

(a) Pumping power grows as the cube of the flow, so under cloud the optimizer runs cold to save pumping

(b) Receiver loss depends on the metal temperature and not on the sun, so under cloud it is a larger share of what is absorbed

(c) The fluid stability limit falls with the irradiance, so under cloud the optimizer must stay further below it

(d) The protection defocuses more of the field under cloud, so the optimizer moves away from its threshold

6.2 ○ Concept check · 1 min

For most of a clear day in Auto, the level loop's defocus output sits at zero. Why?

(a) The loop is interlocked out until the field is hotter than the hot tank

(b) Its integral time of 0.30 h is too long for it to act within a day

(c) The level is below its setpoint, so the direct-acting loop asks for negative defocus, which the limit refuses

(d) The high selector passes the protection's demand instead, which is larger all day

6.3 Concept check · 2 min

On scenario 11 the shipped outlet loop has an outlet error of 882.8 °C min. The predictive controller has 46.8 with its preview and 47.1 without it, and a PI loop trimming a static feedforward from the pyrheliometer, started on its own flow, has 45.7. What carries most of the predictive controller's lead over the shipped loop?

(a) The preview of the irradiance over the next twenty minutes, which shows each cloud edge in advance

(b) The receding horizon, because a plan over twenty minutes looks past the transport delay

(c) The soft ceiling at $T_{max} - 4\,^{\circ}\mathrm{C}$, which keeps the outlet out of the protection

(d) The measured irradiance and a model of the field, which any structure given both can use

6.4 Theory and calculation · 6 min

Use Equations (6.6) and (6.7) summed over the whole field, with $A_{rec} = 2815\,\mathrm{m}^2$, $\varepsilon = 0.12$, $h_{conv} = 0.35\,\mathrm{W/(m^2\,K)}$, $h_{wind} = 0.25\,\mathrm{W/(m^2\,K)}$, $\sigma_{SB} = 5.67 \times 10^{-8}\,\mathrm{W/(m^2\,K^4)}$, an ambient of 25 °C and a wind of 3 m/s. Take every metal node at the same temperature.

(a) Find h_e and the field loss with the metal at 393 °C, split into radiation and convection. *Checkpoint:* $h_e = 0.833\,\mathrm{W/(m^2\,K)}$, 3.62 MW plus 0.86 MW, 4.48 MW in all.

(b) Repeat at 300 °C, and give the ratio of the two radiative parts. *Checkpoint:* 2.56 MW, ratio 1.89.

(c) Express the loss at 393 °C as a share of the 50.8 MW absorbed at clear noon and of the 12.7 MW absorbed at a quarter of the sun. *Checkpoint:* 8.8 % and 35.3 %.

6.5 Theory and calculation · 6 min

The field holds $M_f = 60\,000\,\mathrm{kg}$ of oil with $c_p = 1.60 + 0.0024\,T$ in kJ/(kg K) (T in °C), and the receiver tube is $M_m = 64\,000\,\mathrm{kg}$ of steel with $c_{p,m} = 490\,\mathrm{J/(kg\,K)}$.

(a) Find the transit time of the oil through the field at the design flow of 222 kg/s and at a sixth of it. *Checkpoint:* 4.5 min and 27 min.

(b) Find the heat capacities of the tube metal and of the oil in the field, with c_p at the design mean of 346 °C. *Checkpoint:* 31.4 MJ/K and 146 MJ/K.

(c) The steady outlet rise over the inlet is $Q_{net}/(\dot{m}_f c_p)$. Find $\partial T_{out}/\partial \dot{m}_f$ at 222 kg/s with $Q_{net} = 46\,\mathrm{MW}$, and the factor by which it changes when both the heat and the flow fall to a quarter. *Checkpoint:* –0.384 °C per kg/s, four times larger.

6.6 Theory and calculation · 7 min

Use Equations (6.15) to (6.17) with $\varepsilon_{SG} = 0.85$, $T_{sat} = 285\,^{\circ}\mathrm{C}$, $Q_{PB} = 27.8\,\mathrm{MW}$, $\eta_{PB} = 0.36$, $k_T = 0.0013/^{\circ}\mathrm{C}$, a minimum stable load $x_{min} = 0.25$ and $c_p = 1.60 + 0.0024\,T$ in kJ/(kg K) evaluated at the mean of T_{sg} and T_{sat}. The block flow is 119 kg/s.

(a) At $T_{sg} = 393\,^{\circ}\mathrm{C}$, find the thermal duty, the efficiency and the gross output. *Checkpoint:* 26.4 $\mathrm{MW_{th}}$, 35.4 %, 9.33 MW.

(b) Repeat at $T_{sg} = 340\,^{\circ}\mathrm{C}$. *Checkpoint:* 13.1 $\mathrm{MW_{th}}$, 26.4 %, 3.45 MW.

(c) Show that the block is off at $T_{sg} = 310\,^{\circ}\mathrm{C}$. *Checkpoint:* load fraction 0.21, below 0.25.

6.7 App challenge · 8 min

Load scenario 3, *What a cloud does to a plant defending a temperature*, and change no setting. Run it at speed 0.5 with the default seed and pause between 15:00 and 15:10, four hours after the 11:00 start. Export the CSV file.

(a) Read the Defocused tile, and the highest value of the column `level_pct` in the file. *Checkpoint:* 13.8 MWh, within 0.1 MWh, and the level stays below 58 %.

(b) Average the columns `Qloss_MW` and `Qabs_MW` over the rows with `clock_h` from 12.00 to 12.60, inside the cloud, and give the loss as a share of the heat absorbed. *Checkpoint:* 3.75 MW of 11.9 MW, 31.5 % within one point.

(c) Say what the defocused energy was spent on, and which trend shows the controller succeeding while the plant fails. *Checkpoint:* The high-temperature protection after the clearing. The outlet strip against the "Where the sunlight goes" strip.

6.8 App challenge · 8 min

Load scenario 5, *The same comparison on an intermittent day*, which opens on Fixed T. Run it at speed 2 with the default seed and pause between 20:00 and 21:00. Then press Reset, set Supervisory strategy to Variable T before Run, and run again to the same window.

(a) Read Generated and Collection for the Fixed T run. *Checkpoint:* 25.4 MWh and 70.5 %, within one unit of the last digit.

(b) Read the same two tiles for the Variable T run. *Checkpoint:* 33.5 MWh and 78.1 %.

(c) Give the gain in net generation as a percentage of the Fixed T value, and explain why it is larger than the gain in collection. *Checkpoint:* About 32 %.

6.9 App challenge · 8 min

Load scenario 11, *PI against MPC on the outlet loop, through passing clouds*, which opens with Outlet loop on MPC. Use speed 0.5 and the default seed, and pause each run between 13:00 and 13:15. Run it as loaded. Then press Reset, set Outlet loop to PI and run again. Then set Outlet loop to MPC, set Preview to Off under Advanced, press Reset and run a third time.

(a) For each run read the IAE tile of the collector outlet temperature loop, in °C h, and convert it to °C min. *Checkpoint:* 46.8, 883 and 47.1 °C min, within 2 °C min for the PI and 0.2 °C min for the others.

(b) Read the Defocused tile for each run. *Checkpoint:* 0.0, 17.9 and 0.0 MWh.

(c) Give the reduction in outlet error from the PI to the MPC and the change from switching the preview off, both as percentages. *Checkpoint:* 95 % lower, about 0.6 % higher.

6.10 Application area · 3 min

The outlet setpoint of a trough field sits at 393 °C and the protection starts at $T_{max} - 4\,^{\circ}\mathrm{C}$. Say what sets T_{max}, what crossing it costs on a real plant and in this simulator, and why the industry kept the fluid that imposes it.

6.11 Application area · 4 min

A plant engineer proposes to use this simulator's minutes above 396 °C and its defocused energy to size the defocus margin of a real field with hundreds of parallel collector loops. What does the model hide that matters for that decision, in which direction does it bias the two numbers, and what plant data would the engineer need?

6.12 Review the agent · 5 min

An AI assistant reviews the graded challenge, scenario 12, and writes: “The chapter says the schedule earns the money. Switch Supervisory strategy to Forecast, leave everything else as shipped, and all five targets pass, since the weak gains only cost a little outlet error.” Find the flaw, and say how you would check it in the app.

CHAPTER 7

Forming the Grid From a Battery When a Campus Microgrid Islands

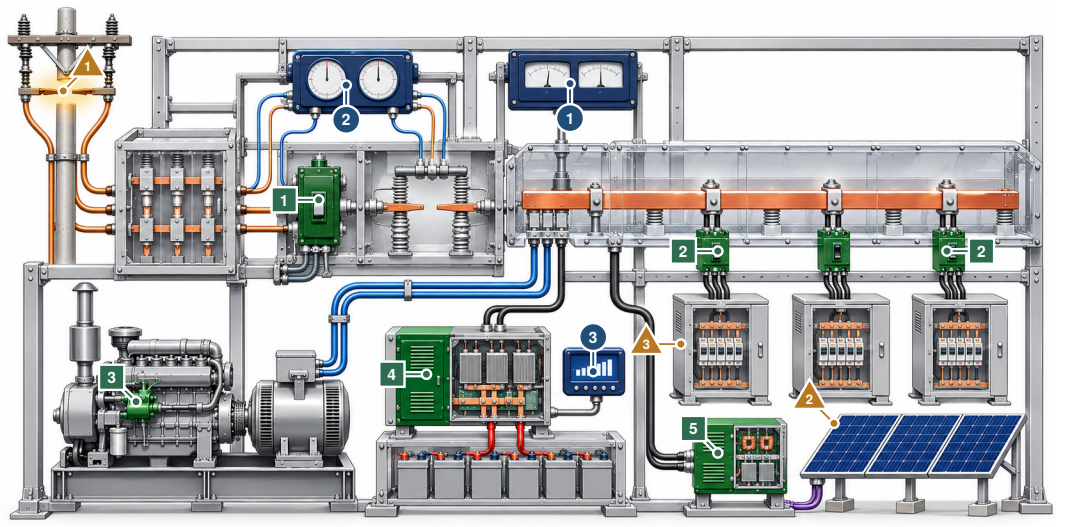

The controller moves 1 utility breaker, closed only through the sync check
2 breakers of the three load blocks that may be shed 3 power setting of the diesel genset, through its governor
4 battery inverter that makes the island's voltage 5 power cap of the solar inverter

The controller reads 1 frequency and voltage of the campus bus 2 synchroscope across the open utility breaker
3 battery charge, from its management system

What disturbs it 1 utility fault that opens the breaker 2 cloud passing over the solar
3 chiller starting on a load block

Notice the open gap of the utility breaker: from that moment the battery inverter alone makes the voltage and frequency that the solar and the loads follow, until the genset starts, and the critical feeder that is never shed is not drawn. Illustration, not to scale.

IN PLAIN TERMS **A Campus That Must Make Its Own Power**

The main idea. When the power line to a campus fails, the campus can run on its own battery, diesel generator and solar panels. One of them must set the beat of the power, its voltage and its frequency, and the others follow it.

The challenge. A battery that only follows cannot hold the campus, and the power dies in a tenth of a second. A battery that sets the beat must slow down gently under a sudden load and warn the safety switches when pushed past its limit.

Why it matters. A hospital wing, a data center and labs stay lit only if these settings are right.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	11 min	The first page, the history box of Section 7.2, the fact box of Section 7.4, the try-it and fact boxes of Section 7.6, and the key points
🕮	**Study**	55 min	The whole chapter with its pause points and the twelve exercises
⏢	**Practice**	20 min	The first page, Section 7.6, Section 7.8, Section 7.9 and its open challenges

The control question When the utility's breaker opens, one of the island's own sources has to make the voltage that the others follow. What does each layer above that source, from virtual inertia and droop to load shedding, charge management and the sync check, buy a campus in hertz, megawatts and critical load kept on?

Learning objectives

- Explain why a grid-following inverter cannot hold an island and a grid-forming one can.
- Compute the first-swing bound, the settling time constant and the droop deviation of a virtual synchronous machine.
- Predict how two droops share a load step from their stiffnesses, and what the secondary loop changes.
- Explain how a power-limit droop lets under-frequency relays hear a battery overload.
- Price an islanded hour in diesel and in stored energy, and judge a dispatch by both.

Prerequisites Sections 1.4.2 (swing equation, ROCOF), 1.4.3 (droop and restoration), 1.4.4 (state of charge), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios and data), with Section 1.6.3 for the energy dispatch only.

The app https://apmonitor.com/apps/islanded-microgrid/

See also Volume 1, Chapter 8 and Volume 1, Chapter 7 for inertia, droop offset, integral restoration and a governor on an islanded load. Volume 7, Chapter 2 and Volume 7, Chapter 6 for an outcome that changes at a threshold and a loop that creates its own operating point.

7.1 Introduction

A microgrid is a cluster of loads and small sources run as one controllable system, which the utility can treat as a single controlled cell and which can give its own customers local reliability when the grid cannot [1]. The campus simulated here is one built for that second role. While the breaker at the point of common coupling (PCC) is closed, the utility makes the voltage, holds the frequency and absorbs every mismatch, so every inverter on site can be a current source that follows a phase-locked loop (PLL). When the breaker opens, one of the island's own sources has to make the voltage that the others follow, and in the first tenth of a second nothing slower than that source can help. Microgrid control is organized as a hierarchy of primary, secondary and tertiary levels [2, 3], and the converters in it are classed by whether they form the grid, feed it or support it [4]. Section 1.4.2 places these levels on the clocks of the whole volume.

A battery on site is often taken to make a campus resilient, and the simulator shows why it does not by itself. A battery whose inverter follows the grid keeps pushing its current after the 10:00 fault opens the breaker, the bus voltage becomes that current through the load, and 40 ms later the bus is at 0.42 pu with both PLLs reading 60.92 Hz. The bus is dead at 60 ms, and the critical feeder loses 11.75 kWh before the genset can restore it. Once the battery forms the grid, two further misreadings remain, and both come from treating the inverter as a source rather than as a control law. Virtual inertia is read as depth, yet on a lone battery three 0.2 MW load steps at 0, 1.5 and 3 s of inertia fall at 2.61, 1.77 and 1.21 Hz/s and bottom at 59.697, 59.728 and 59.729 Hz, so the inertia sets the rate and the droop the depth. Inertia still decides the fault, because the solar rides through 3 Hz/s and no more, and with none the first swing reads 11.18 Hz/s, the solar trips and all three sheddable blocks are lost. A steady frequency is read as a source that copes, yet a grid-forming inverter holds its frequency while it overdraws its battery. With the inverter's power-limit droop turned down from 0.40 to 0.02, the chiller of case 7 keeps the battery 11.7 s over its continuous limit at 59.509 Hz, above the first shedding stage, until the battery's management system trips it and the critical feeder loses 17.57 kWh. At 0.40 the first stage sheds one block 7.1 s after the chiller starts and the critical feeder loses nothing.

The class of problem is general. A fast source with a small store beside a slow one with a large store, sharing load by local feedback without communication, appears wherever inverters run in parallel, and the parallel-inverter literature names isolated ac systems, distributed uninterruptible power supplies and photovoltaic systems as one control problem [5]. The islanding transition has been

studied with the campus's own mix, a rotating synchronous machine beside a converter-interfaced unit [6], and the difference between forming and following inverters is now argued to matter for any power system that inverters dominate [7]. In process terms the battery is a surge volume that answers every disturbance first and must then be refilled, the genset a slow supply that carries the average, and the state-of-charge band an averaging level controller on an energy inventory, as in Section 1.4.4. A block diagram of the droops has no place for the breaker, the trips or the limits that decide how the island ends.

These controls are usually evaluated in electromagnetic transient simulation or on laboratory hardware [6, 8], and the standards that govern islanding and reconnection state requirements rather than consequences [9]. Neither gives a learner a plant where the inertia, the droop, the band and the sync-check window are settings whose effects read in hertz, megawatts and kilowatt-hours of critical load left unserved. This chapter tallies what each layer of an island's control buys, from the inverter that forms the grid to the relay that lets the utility back, on a campus microgrid simulator described closely enough to be rebuilt. Every method in it is established. The virtual synchronous machine, the droop sharing, the shedding stages and the sync check follow the literature and standards cited. The simulator solves a one-bus phasor model whose power balance closes to round-off, switches a battery inverter between following and forming without a bump, pairs a virtual-impedance current limit with a power-limit droop that lets the relays hear an overload, and gives its controller only the meters a microgrid controller has. Its cases and two-hour challenge take one layer at a time and include an islanding that no setting survives whole.

Section 7.2 recounts how small power systems learned to run alone. Section 7.3 builds the campus: the bus, the genset, the battery inverter in both modes, the solar, the loads, the meters and the protection. Section 7.4 derives five relations behind the cases, from the first swing to the voltage across an open breaker. Section 7.5 programs the microgrid controller layer by layer. Section 7.6 measures ten cases and the two-hour challenge, and Section 7.7 adds an optional one-hour energy dispatch and prices the stored energy it leaves behind. Section 7.8 describes the hardware that real microgrids run on, Section 7.9 states the limitations and Section 7.10 draws the conclusions. The exercises of Section 7.11 work through the cases on the app.

7.2 How Small Power Systems Learned to Run Alone

A generator that feeds an isolated load alone can hold its speed exactly. The classical treatment of speed governing separates the isochronous governor, which integrates the speed error away and suits such a unit, from the droop governor, whose speed falls with output so that units in parallel divide a change of load in a fixed ratio [10]. The genset on the simulated campus has both, a droop by default and an isochronous option for a genset left alone on the island. The hydro unit of Volume 1, Chapter 7 also either supplies an islanded town or follows a stiff grid, and shares load through a permanent droop. On an interconnected system an area left short of generation is an emergency, in which an automatic load-shedding program can prevent the collapse of the area, and by 1968 the factors in applying underfrequency relays to such a program, with a method for calculating their settings, were in the transmission literature [11]. An IEEE guide now compiles the practice of shedding and of restoring the load afterward [12].

Power electronics removed the rotating mass but not the sharing problem, and the history box below tells how inverters came to share load by droop.

The microgrid concept gathered such sources into one system. A 2002 overview described a cluster of loads and microsources, each under 100 kW, operated as a single controllable system that the utility sees as one cell and that can give its customers local reliability and uninterruptible supply [1]. The islanding transition was then simulated for a micro-grid with one rotating synchronous machine and one converter-interfaced unit, after both a planned switching event and a fault, and a suitable control of the converter unit kept the island stable through it [6]. A hierarchy of primary droop with a virtual output impedance, a secondary level that restores the deviations the droop leaves, and a

tertiary level that manages the exchange with the distribution system was proposed as a general approach toward standardization [2], and a review of converter control sorted microgrid converters into grid-forming, grid-feeding and grid-supporting configurations [4].

> ⧖ HISTORY **Inverters Learn to Share Load (Early 1990s)**
>
> Droop control, the most established way for an inverter to form the grid, was first proposed in the early 1990s [5, 13]. Each inverter trades its frequency against its real power and its voltage against its reactive power, much as a synchronous machine does in steady state. All the units then settle at one frequency, and each supplies power in proportion to its droop slope, with no communication between them [13]. These controls were first designed for systems with a small footprint, microgrids and small islands. A national laboratory roadmap of 2020 found early grid-forming inverters serving islanded microgrids with a high share of renewable generation, and saw them as the first stage before grid-forming control reaches the bulk power system [13].

The inverter that forms an island was next asked to behave like the machine it replaces. A virtual synchronous machine was proposed in 2007 to make inverter-connected generation appear to the grid as an electromechanical synchronous machine, with the electrical effect of a rotating mass and a share of primary reserve [14]. Synchronverters, inverters operated to mimic synchronous generators, shared real and reactive power through frequency and voltage droop and could run in island mode [8]. A 2014 letter showed that a virtual synchronous machine and a frequency droop are equivalent under certain conditions [15], and the simulator implements the machine in that form. A 2020 review argued that a grid-forming inverter acts on frequency swings as a source of damping rather than of inertia while a grid-following one can aggravate them, and that protection against overload and a reserve of stored energy are what make the damping usable [7].

Interconnection standards came to the island in stages. The first standard of the IEEE 1547 series, approved in 2003, set requirements for interconnecting distributed resources, among them their response to abnormal conditions and to islanding [16], and the application quotes from the series the requirement that a distributed resource detect an unintended island and cease to energize it within 2 s. A 2011 guide addressed island systems that separate from and reconnect to part of the area power system while they serve their own loads [17]. The 2018 revision, whose Category III ride-through and trip defaults the simulated solar carries, treats interconnection and interoperability together [9], and Chapter 8 puts the same revision's Volt-VAR and Volt-Watt curves on a residential feeder. Reconnection rests on an older practice. A synchronizing panel tells the operator what to adjust on the governor and the exciter and when the breaker may close, and an automatic synchronizer with a manual backup often makes the decision [18]. The sync-check relay of the simulator, which the application labels 25, is that decision in software.

7.3 The Campus Behind One Breaker

Figure 7.1 shows the campus as a one-line diagram. Four sources stand above one bus, each on its own breaker: the utility behind the PCC, a 2 MW diesel genset, a 1 MW, 2 MWh battery inverter and 1.5 MW of rooftop solar. Four feeders hang below it. The critical feeder serves a hospital wing, a data center and laboratories and is never shed, and three sheddable blocks serve HVAC with the chiller plant, the dormitories, and the classrooms and offices. Every power is in per unit of a 1 MVA base, so a per-unit power is a megawatt, and time runs in seconds on a clock that starts each case at its own time of day. Table 7.1 gives each signal its role, unit and range.

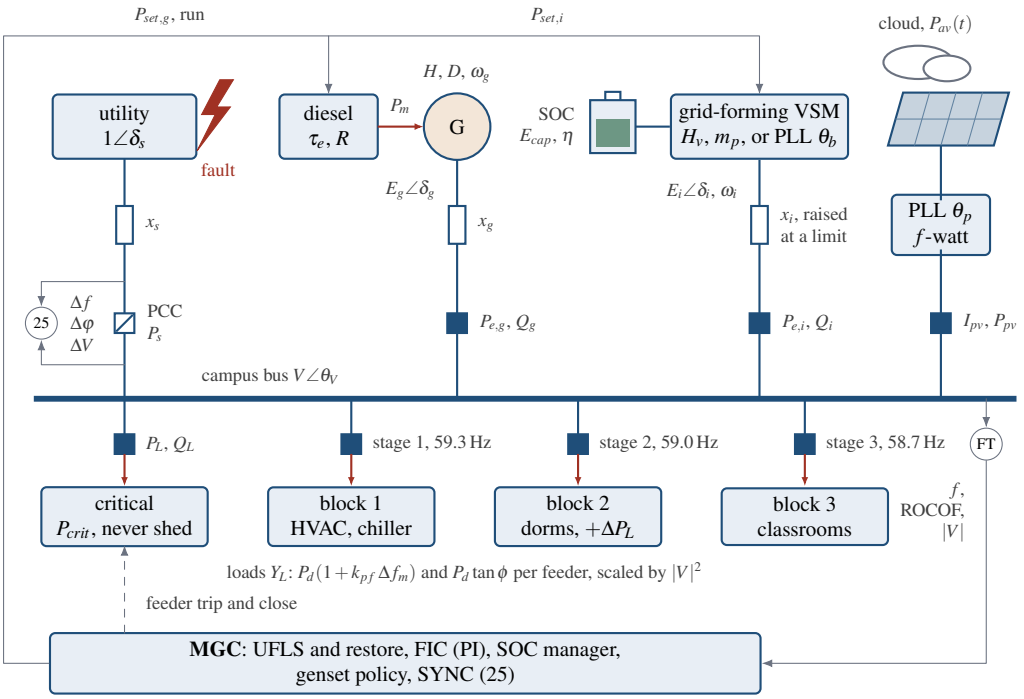

Figure 7.1 One-line diagram of the simulated campus. The utility, the genset and the battery inverter while it forms the grid are voltage sources $E_k\angle\delta_k$ behind the reactances x_s, x_g and x_i, and the solar, with the battery while it follows the grid, injects a current on its PLL angle θ_p or θ_b. The engine turns the governor's command into P_m through τ_e, and the rotor carries H, D and the speed deviation ω_g. The feeders are constant-impedance loads Y_L that depend on the measured frequency deviation Δf_m, and blocks 1 to 3 open at the shedding stages marked. The meter FT gives f, ROCOF and $|V|$ to the microgrid controller (MGC), which sets $P_{set,g}$ and $P_{set,i}$, starts the genset, trips and closes the feeders, and closes the PCC through the synchroscope (25) that compares its two sides. The PCC is drawn open, as it stands after the fault.

Table 7.1 The Campus's Signals by Role, Unit and Range

Symbol	Role	Unit	Range
f	CV of the secondary, measured	Hz	56.5–63
SOC	CV of the SOC manager, measured	%	5–95
$\Delta\varphi$	CV of the synchronizer, measured	°	±180
$P_{set,g}$, $P_{set,i}$	MVs, genset and battery setpoints	MW	0–2, ±1
B_1, B_2, B_3	MVs, block feeder breakers	0/1	—
—	MVs, genset run, PCC close	0/1	—
P_{cap}	MV, solar cap	MW	0–1.5
—	DVs, grid fault and return, chiller, genset or solar trip	event	—
P_{av}, ΔP_L	DVs, sun, cloud and haze, extra load on block 2	MW	0–1.5, 0–1
Δf_{grid}	DV, the utility's frequency offset	Hz	±0.1
δ_g, ω_g, P_m, E_g, ξ_g	states, genset rotor, engine, AVR, isochronous integral	rad, pu, MW	—
δ_i, ω_i, E_i, z	states, grid-forming inverter	rad, pu, MW	—
θ, x, V_m	states, each PLL and its voltage meter	rad, pu	—
P_{pv}, SOC, δ_s	states, solar, charge, utility angle	MW, —, rad	—
—	states, seven energy and fuel ledgers	MWs, L	—
$\lvert V\rvert$, P, Q, Δf, ΔV	measured, bus, every meter, across the PCC	pu, MW, Hz	—

7.3.1 The Bus

The bus has no state of its own. At every evaluation its complex voltage $V = |V|e^{j\theta_V}$ is the solution of one nodal equation, as shown in Equation (7.1).

$$V\left(Y_L + \sum_{k\in\mathscr{V}} \frac{1}{jx_k}\right) = \sum_{k\in\mathscr{V}} \frac{E_k e^{j\delta_k}}{jx_k} + \sum_{m\in\mathscr{C}} I_m \tag{7.1}$$

where $\mathscr{V}$ holds the voltage sources on line, the genset g, the battery inverter i while it forms the grid and the utility s while the PCC is closed, each an EMF E_k at angle δ_k behind a reactance x_k on the 1 MVA base, and $\mathscr{C}$ holds the grid-following currents of the solar p and of the battery b while it follows. Every angle is measured in a frame turning at $\omega_0 = 2\pi f_0$ with $f_0 = 60\,\text{Hz}$. The load admittance is $Y_L = G_{sh} - jB_{sh} + \sum_k [P_{d,k}(1 + k_{pf}\Delta f_m) - jQ_{d,k}]$ over the closed feeders, where $P_{d,k}$ and $Q_{d,k} = P_{d,k}\tan\phi$ are a feeder's demand at rated voltage with the power factor $\cos\phi = 0.97$, Δf_m is the per-unit deviation of the measured frequency held from the previous step, and $G_{sh} = B_{sh} = 0.004\,\text{pu}$ stand for the transformers' no-load losses and magnetizing. A voltage source delivers $S_k = V\,\overline{(E_k e^{j\delta_k} - V)/(jx_k)}$ and the loads take $P_L = |V|^2 \sum_k P_{d,k}(1 + k_{pf}\Delta f_m)$. One complex division gives V, so the power into the bus equals the power out of it to round-off, 3.2×10^{-15} MW at worst over the islanding of case 3.

7.3.2 The Genset

The genset is a classical machine, an EMF behind its transient, transformer and cable reactance $x_g = X_g/S_g$, driven by a diesel engine through a first-order lag and a droop governor, as shown in Equation (7.2).

$$\begin{aligned} 2HS_g \frac{\mathrm{d}\omega_g}{\mathrm{d}t} &= P_m - P_{e,g} - DS_g\,\omega_g \\ \tau_e \frac{\mathrm{d}P_m}{\mathrm{d}t} &= \mathrm{clip}_{[0,S_g]}\left(P_{set,g} - \frac{S_g}{R}\,\omega_g + \xi_g\right) - P_m \\ \frac{\mathrm{d}\delta_g}{\mathrm{d}t} &= \omega_0\,\omega_g \\ \frac{\mathrm{d}E_g}{\mathrm{d}t} &= K_{avr}\left(V_{set} - n_{q,g}\,Q_g/S_g - |V|\right) \end{aligned} \tag{7.2}$$

where ω_g is the rotor's per-unit speed deviation, P_m the mechanical power, $P_{e,g}$ and Q_g the electrical output, $\mathrm{clip}_{[a,b]}$ limits its argument to $[a,b]$, V_{set} is the voltage setpoint plus the synchronizer's trim, and the AVR's EMF stays between 0.2 and 1.6 pu. The isochronous option adds ξ_g with $\mathrm{d}\xi_g/\mathrm{d}t = -(S_g/R)\,\omega_g/T_{iso}$ and $T_{iso} = 4\,\text{s}$, and $\xi_g = 0$ otherwise. The engine burns $q_0 + q_1 \max(0, P_m)$ of diesel, with $q_0 = 110\,\text{L/h}$ running unloaded and $q_1 = 0.205\,\text{L/kWh}$ above it, and q_0 alone while it cranks and warms. A start takes $t_{start} = 45\,\text{s}$ from the crank to the breaker, a timer in the model, after which the breaker closes onto a dead bus or at the live bus's angle, speed and voltage, so the genset's synchronizer is perfect. While it is off its states follow the bus, so a restart is bumpless.

7.3.3 The Battery Inverter, Forming or Following

Forming the grid, the battery inverter is a virtual synchronous machine in the form whose droop is its damping [15], an EMF $E_i\angle\delta_i$ behind its filter, transformer and cable reactance $x_i = X_i/S_i$, as shown

in Equation (7.3).

$$\begin{aligned} 2H_v S_i \frac{\mathrm{d}\omega_i}{\mathrm{d}t} &= P_{set,i} - P_{e,i} - \frac{S_i\,\omega_i + m_{ov}\,z}{m_p} \\ \frac{\mathrm{d}\delta_i}{\mathrm{d}t} &= \omega_0\,\omega_i \\ \tau_{ov}\frac{\mathrm{d}z}{\mathrm{d}t} &= P_{over} - z, \qquad \left|\frac{\mathrm{d}z}{\mathrm{d}t}\right| \le \frac{\sigma_{ov} S_i}{f_0\, m_{ov}} \\ \frac{\mathrm{d}E_i}{\mathrm{d}t} &= K_q\big(V_{set} - n_q(Q_i/S_i - Q_{set}) - |V|\big) \end{aligned} \tag{7.3}$$

where H_v is the virtual inertia, m_p the frequency droop, ω_i the virtual speed deviation, z the lagged, slew-limited overload of Equation (7.4), and $Q_{set} = 0$ the reactive setpoint. With $H_v = 0$ the first line becomes the pure droop $\omega_i = (m_p/S_i)(P_{set,i} - P_{e,i}) - (m_{ov}/S_i)z$. The voltage loop integrates the bus voltage's error on a reactive droop, stays between 0.2 and 1.4 pu, and stops winding up while the inverter is limited.

Two limits make the voltage source a physical one. The inverter presents a larger impedance whenever its current would pass $I_{max} = 1.4$ pu or its power would pass the short-term allowance of its battery, $P_{hard} = \min(1.2S_i,\ P_{dis} + \Delta P_s)$ with $\Delta P_s = 0.5$ MW, multiplying x_i by $\max\left(|I|/I_{max},\ (P/P_{hard})^{0.8}\right)$ until both hold, so the bus voltage sags instead of the current rising. The power the island asks beyond the battery's continuous capability is then fed back, as shown in Equation (7.4).

$$P_{over} = [P_{e,i} - P_{dis}]^+ - [-P_{e,i} - P_{ch}]^+ + \lambda\,(P_{raw} - P_{e,i}) \tag{7.4}$$

where $[x]^+ = \max(0, x)$, P_{dis} and P_{ch} are the continuous limits of Equation (7.7), λ is 1 while the impedance is raised and 0 otherwise, and P_{raw} is the power the inverter would carry at its own reactance. This power-limit droop lowers the frequency in proportion to the shortfall, never faster than $\sigma_{ov} = 1$ Hz/s, so that the shedding relays hear an overload that a voltage source would otherwise hide.

Following the grid, the same inverter injects a current on the angle of its PLL, as shown in Equation (7.5).

$$\begin{aligned} I_b &= \frac{P_{ref} - jQ_{ref}}{V_m}\,e^{j\theta_b}, \qquad |I_b| \le S_i \\ \frac{\mathrm{d}\theta_b}{\mathrm{d}t} &= \omega_0\,(K_p v_q + x_b), \qquad \frac{\mathrm{d}x_b}{\mathrm{d}t} = K_i v_q \\ v_q &= \frac{\mathrm{Im}(V e^{-j\theta_b})}{|V|}, \qquad \tau_m \frac{\mathrm{d}V_m}{\mathrm{d}t} = |V| - V_m \end{aligned} \tag{7.5}$$

where P_{ref} is the battery's power setpoint, $Q_{ref} = 0$, $K_p = 2\zeta\,\omega_n/\omega_0$ and $K_i = \omega_n^2/\omega_0$ with $\omega_n = 2\pi f_{PLL}$, $f_{PLL} = 2$ Hz and $\zeta = 0.7$, $v_q = \sin(\theta_V - \theta_b)$ is the phase error, V_m the measured voltage behind a $\tau_m = 20$ ms lag, and the PLL reads the frequency $f_0(1 + K_p v_q + x_b)$. The current ceases below $V_m = 0.5$ pu, the momentary cessation of the Category III settings. On a grid θ_V belongs to the utility. On an island with no voltage source $V = \sum_m I_m/Y_L$, so θ_V is the PLL's own angle plus the angle of the loads' impedance, and the loop chases a reference that it turns itself. The PLL runs in both modes and in forming mode is only a meter. A change of mode solves the new mode's states from the present current, so the switch is bumpless.

7.3.4 Solar, Loads and the Battery's Store

The solar is a second grid-following current, $I_p = \min(P_{pv}, P_{av})\,e^{j\theta_p}/V_{m,p}$ on its own PLL with $|I_p|$ capped at $1.1P_{pk}$, so a cloud's edge reaches the bus at once and only a recovery waits for the power

loop. That loop follows the available power through the frequency-droop response of Equation (7.6).

$$\begin{aligned} \tau_{pv}\frac{\mathrm{d}P_{pv}}{\mathrm{d}t} &= P^{*}_{pv} - P_{pv} \\ P^{*}_{pv} &= \mathrm{clip}_{[0,P_{av}]}\big(r\min(P_{av},P_{cap}) - \Delta P_{fw}\big) \\ \Delta P_{fw} &= P_{pk}\frac{f_{pll,p} - f_0 - d_b}{f_0\,k_{OF}} \quad \text{for } f_{pll,p} > f_0 + d_b \end{aligned} \tag{7.6}$$

where P_{cap} is the operator's cap, r ramps from 0 to 1 over 300 s after a re-entry, $d_b = 0.036$ Hz and $k_{OF} = 5\,\%$, and $\tau_{pv} = 2.2$ s gives the 5 s response time that the application takes as the default. The available power is $P_{av} = P_{pk}\,c(h)\prod_n[1 - d_n(a_n - b_n)](1-\varepsilon)$ with $c(h) = [\max(0, \sin(\pi(h - 6.5)/13))]^{1.3}$ at clock hour h, which peaks at 13:00. Each cloud n of depth $d_n = 0.5$ has an entering edge a_n and a leaving edge b_n ten minutes later, each rising from 0 to 1 over 10 s, and ε is the haze.

The critical feeder demands $P_{crit} = 0.90$ MW of its 1 MW connected. Each block demands $P_{blk}\,o(h)$ with $P_{blk} = 0.5$ MW and an occupancy $o(h) = 0.55 + 0.30\exp[-((h-14)/3.5)^2]$, a chiller start adds 0.4 MW to block 1 through a 10 s soft start, and the extra-load disturbance adds to block 2. The power factor is 0.97 and $k_{pf} = 1.5$, so a 1 % fall in frequency takes 1.5 % off the load.

The battery's state of charge integrates its DC power, and its management system sets the continuous limits, as shown in Equation (7.7).

$$\begin{aligned} \frac{\mathrm{d}\,\mathrm{SOC}}{\mathrm{d}t} &= -\frac{P_{dc}}{3600\,E_{cap}}, \qquad P_{dc} = \begin{cases} P_{e,i}/\eta & P_{e,i} \geq 0 \\ \eta\,P_{e,i} & P_{e,i} < 0 \end{cases} \\ P_{dis} &= S_i\,\mathrm{clip}_{[0,1]}\Big(\frac{\mathrm{SOC} - \mathrm{SOC}_{min}}{\Delta}\Big) \\ P_{ch} &= S_i\,\mathrm{clip}_{[0,1]}\Big(\frac{\mathrm{SOC}_{max} - \mathrm{SOC}}{\Delta}\Big) \end{aligned} \tag{7.7}$$

where $P_{e,i}$ is the inverter's output in either mode, $\eta^2 = 92\,\%$ is the round trip and $E_{cap} = 2$ MWh, and the taper $\Delta = 20\,\%$ allows full power from 25 to 75 % and none below 5 % or above 95 %. Seven ledgers integrate each source's energy, the loads', the battery's DC energy and the diesel, and the energy ledgers close to 3.7×10^{-15} of the energy over case 3 and 1.2×10^{-14} over the two-hour challenge.

7.3.5 Meters, Protection and Scores

The frequency meter reads the bus voltage's angle through each step of length Δt, as shown in Equation (7.8).

$$\begin{aligned} f_k &= f_0 + \frac{\big[\theta_V(t_k) - \theta_V^{+}(t_{k-1})\big]_{2\pi}}{2\pi\,\Delta t} \\ \mathrm{ROCOF}_k &= \frac{f_k - f_{k-n}}{n\,\Delta t} \end{aligned} \tag{7.8}$$

where $\theta_V^{+}(t_{k-1})$ is the angle at the start of the step solved with that step's breakers, so the jump a breaker makes in the angle is never read as a frequency, $[\cdot]_{2\pi}$ wraps a difference into $\pm\pi$, and $n\,\Delta t = 0.1$ s is the averaging window the application takes from IEEE 1547 for ride-through. The devices' own protection is listed in Table 7.2. The grid-following inverters carry the Category III default trips that the application quotes from IEEE 1547-2018, cease to energize an island that has no voltage source after 2 s, and re-enter 300 s after the bus returns to 0.917–1.05 pu and 59.5–60.1 Hz. The solar also trips on a ROCOF beyond 3.0 Hz/s, which the application takes as the Category III ride-through and uses as a trip. The run is scored on the truth. The critical feeder counts as served only while the bus is live, its breaker closed and $|V| \geq 0.88$ pu, the frequency counts as out of band outside 59.5–60.5 Hz or with the bus dead, and the worst ROCOF counts from 0.2 s into a run.

Table 7.2 Protection Inside the Devices, as Implemented

Device	Condition	Delay
genset	$f_g < 57.0\,\text{Hz}$ or reverse power over 5 %	3 s
	$f_g > 63.0\,\text{Hz}$	0.5 s
forming inverter	$f_i < 56.5\,\text{Hz}$ or $f_i > 62.0\,\text{Hz}$	0.16 s
	on the current limit	10 s
	$P_{e,i} > 1.05\,P_{dis} + 0.01\,\text{MW}$ (BMS)	10 s
	SOC $\leq 2\,\%$	at once
following inverters	$V_m < 0.50\,\text{pu}$, $V_m < 0.88\,\text{pu}$	2, 21 s
	$V_m > 1.20\,\text{pu}$, $V_m > 1.10\,\text{pu}$	0.16, 13 s
	$f < 56.5\,\text{Hz}$, $f < 58.5\,\text{Hz}$	0.16, 300 s
	$f > 62.0\,\text{Hz}$, $f > 61.2\,\text{Hz}$	0.16, 300 s
	no voltage source on the bus	2 s
solar	$\lvert\text{ROCOF}\rvert > 3.0\,\text{Hz/s}$	at once

A genset trip locks the engine out for 60 s. The following inverters' frequency is their own PLL's reading.

7.3.6 Parameters, Initial State and Numerics

Table 7.3 gives the constants. Every case starts at a steady state for its clock, found by Newton steps on the angles and EMFs until every derivative except those of the state of charge and the ledgers is zero, within 2.7×10^{-12} at worst. The exception is the measured voltage V_m of the following battery in case 2, whose current depends on that reading, and whose derivative starts at 2.1×10^{-7} pu/s. Grid-connected, the battery sits at its schedule and the utility carries the rest. Islanded with the genset online, the genset is the angle reference and the battery's setpoint is iterated onto the dispatch rule of Section 7.5, so a case that starts islanded starts where the controller would hold it. The derivative function is pure. Each 20 ms step applies the controller's commands and any breaker change, solves the bus for the step's topology, holds the demand, the available solar and Δf_m, and advances four fourth-order Runge-Kutta sub-steps of 5 ms. The sub-steps are there because a grid-forming inverter with little inertia has a mode of time constant $2H_v m_p$, 6 ms at 0.1 s and 3 %. The engine and its clocks are those of Section 1.6.1. The meters, the protection and the device sequences act after the step. The controller scans every 20 ms and its slower layers act every second. A row is logged every 0.5 s with the largest ROCOF since the previous row, so a first swing that lasts a tenth of a second survives the log, and the Frequency view keeps every 20 ms step for 30 s. With no random input in the model, a case reloaded, or reset with the same settings, reproduces exactly, and the application passes its 33 verification checks.

7.3.7 Assumptions

One bus, solved as phasors. The campus is one bus at 60 Hz, with no electromagnetic transients, no unbalance, no harmonics, no impedance along the feeders and no fault inside the campus, only the utility's. The first cycle after a breaker operates, which the inverters' inner current loops govern, is outside the model.

A classical genset on a timer. The genset is an EMF behind a reactance with a first-order engine, its start is a timer, its synchronizer is perfect, and its fuel is a straight line. A real engine accepts load in limited steps, which the model represents only through the dispatch's ramp of 0.1 MW/s.

One grid-forming law. The inverter is a virtual synchronous machine of one form with a virtual-impedance current limit and a power-limit droop. Grid-forming controls differ widely in how they limit current, and the application's own assumptions warn that some lose synchronism there.

Ideal following inverters. Their inner current loops are ideal, and their protection is the default set the application quotes plus the 2 s island detection, applied only when no voltage source is on the bus, so the solar never trips for islanding on an island the battery has formed.

Table 7.3 Constants of the Campus Model by Class: V Adjustable in Show Model, H Hidden, C Set by the Code

Symbol	Meaning	Value	Class
Genset			
S_g, X_g	rating, reactance on its rating	2 MW, 0.30 pu	H, V
H, D	inertia constant, damping	1.5 s, 1.0 pu/pu	V
—	minimum load under the dispatch	30 %	V
$n_{q,g}$, K_{avr}	reactive droop, AVR gain	4 %, 6 s^{-1}	V, H
q_0, q_1	fuel, unloaded and per kWh	110 L/h, 0.205 L	V
T_{iso}	isochronous reset time	4 s	C
Battery inverter			
S_i, X_i	rating, reactance on its rating	1 MW, 0.15 pu	H, V
E_{cap}, η^2	energy, round trip	2 MWh, 92 %	V
I_{max}	hardware current limit	1.4 pu	V
K_q, n_q	voltage loop gain, reactive droop	20 s^{-1}, 5 %	H, V
m_{ov}, τ_{ov}, σ_{ov}	power-limit droop, lag, slew	0.40, 1 s, 1 Hz/s	V
ΔP_s	short-term allowance, 10 s	0.5 MW	V
SOC_{min}, SOC_{max}, Δ	BMS floor, ceiling, taper	5, 95, 20 %	V
τ_m, ζ	PLL voltage lag, damping ratio	20 ms, 0.7	C
Solar, utility and loads			
P_{pk}, τ_{pv}	solar rating, power loop lag	1.5 MW, 2.2 s	V
d_b, k_{OF}	frequency-droop dead band, slope	0.036 Hz, 5 %	V
—	ROCOF ride-through, then trip	3.0 Hz/s	V
$X_s = x_s$	utility reactance, 20 MVA short circuit	0.05 pu	V
P_{crit}, P_{blk}	critical demand, each block	0.90, 0.5 MW	V
pf, k_{pf}	power factor, frequency dependence	0.97, 1.5 pu/pu	V
—	chiller, soft start	0.4 MW, 10 s	V
G_{sh}, B_{sh}	transformer shunt	0.004 pu	C

Static loads. The loads are constant impedances with a frequency dependence, with no motor dynamics and no inrush, and the chiller arrives as a 10 s ramp, so a real compressor's starting current would dip the bus more than the model does.

An ideal meter. The relays act on the frequency of Equation (7.8) with no filtering of their own and no breaker time. Section 7.6 shows that the meter's step moves the reading of a first swing.

Thresholds set for teaching. The shedding stages, the restoration rule, the sync-check window, the battery's taper and the genset's trips are all chosen for teaching, and Section 7.9 says where real settings come from.

7.4 Five Relations Behind the Cases

Five relations from the model organize every measured case, and none depends on the microgrid controller above the devices.

7.4.1 The First Swing

At the instant the PCC opens, the virtual speed is still zero, and the inverter's electrical power rises by the import ΔP it must take over. The first line of Equation (7.3) then gives the initial rate of change of frequency of Equation (7.9), the first rate of Equation (1.1) written for the inverter.

$$\mathrm{ROCOF}(0^+) = -\frac{f_0\,\Delta P}{2H_v S_i} \tag{7.9}$$

For the 0.329 MW import of case 3 it is 4.94 Hz/s at 2 s, 6.59 at 1.5 s and 9.88 at 1.0 s, and the application's check, which steps the setpoint rather than the load, confirms the slope to 10^{-9} at 1, 2 and 3 s. The ROCOF a relay reads over 0.1 s is smaller, 2.48 Hz/s at 2 s, because the bus voltage dips from 0.988 to 0.958 pu in the first step and the constant-impedance loads draw 0.110 MW less, and because the window spans the bend of the swing. Equation (7.9) is therefore an upper bound. Kept under the 3 Hz/s ride-through, the bound would ask for 3.3 s of inertia, past the slider's range, while the reading itself clears the ride-through from 1.3 s at the shipped step (2.97 Hz/s).

+ DID YOU KNOW? A Continent Against a Campus

A 2018 report of the European transmission operators shows phasor measurement units recording the forced outage of 1000 MW of generation in Switzerland: the continental frequency fell by only 40 mHz, at 10 to 15 mHz/s [19]. On this campus the loss of a 0.329 MW import starts the frequency falling at 2.48 Hz/s.

7.4.2 Settling Without Inertia

With the battery alone on the bus, $P_{e,i}$ is set by the load and the virtual speed obeys the first-order law of Equation (7.10).

$$2H_v m_p \frac{\mathrm{d}\omega_i}{\mathrm{d}t} = -\omega_i + \frac{m_p}{S_i}\left(P_{set,i} - P_{e,i}\right) \tag{7.10}$$

The time constant $T_v = 2H_v m_p$ is 120 ms at the defaults, and the settled deviation $m_p f_0 \Delta P / S_i$ does not contain H_v. On its droop alone a lone battery would settle 0.36 Hz low after a 0.2 MW step. With the loads' frequency and voltage relief and the secondary already acting, the falls measured in case 4 from the frequency before each step are 0.302 to 0.304 Hz, within 2 mHz of one another from no inertia to 3 s. Inertia buys time, and the droop and the headroom decide where the fall stops. Volume 1, Chapter 8 draws the same division on a single-area grid, where inertia sets the initial rate of change of frequency and the stiffness of droop and load damping sets the arrested offset.

7.4.3 Sharing by Stiffness

With the genset online, both droops act on one frequency. The genset's governor and rotor damping give it the stiffness $K_g = S_g/R + DS_g = 42$ MW per unit of speed and the inverter has $K_i = S_i/m_p = 33.3$ MW per unit, so a load change divides between them and moves the frequency as shown in Equation (7.11).

$$\frac{\Delta P_{e,g}}{\Delta P_{e,i}} = \frac{K_g}{K_i} = 1.26, \qquad \Delta f = -\frac{f_0\left(\Delta P_{e,g} + \Delta P_{e,i}\right)}{K_g + K_i} \tag{7.11}$$

The ratio holds to 2 % in the application's check at governor droops of 5 and 3 %, where it becomes 2.06. Equation (1.2) writes the same division with each stiffness in megawatts per hertz and without the rotor damping, and this chapter quotes stiffness per unit of speed. The frequency stays where the droops leave it until an integral loop moves a setpoint.

7.4.4 The Overload a Voltage Source Hides

A voltage source on its limit holds its frequency, because nothing in its swing equation knows that the battery is overdrawn. At steady state the last term of Equation (7.3) lowers the frequency by $f_0 m_{ov}/S_i = 24$ Hz for each megawatt of shortfall, as shown in Equation (7.12).

$$\Delta f_{ov} = -\frac{f_0 m_{ov}}{S_i} z, \qquad \left|\frac{\mathrm{d}\Delta f_{ov}}{\mathrm{d}t}\right| \le \sigma_{ov} \tag{7.12}$$

A battery 0.05 MW short is driven toward a frequency 1.2 Hz low, through the first stage at 59.3 Hz, at no more than 1 Hz/s, which the solar rides through. The secondary loop cannot cancel the term while the battery carries the island alone, because its total is limited to the battery's continuous capability.

7.4.5 Closing Across the PCC

Two voltages of 1 pu that differ in phase by $\Delta\varphi$ put $2\sin(\Delta\varphi/2)$ per unit across the open breaker, which is 1.95 pu at the 153° of case 10 and 0.04 pu at the 2.4° at which the synchronizer closes. Closing at the wrong moment drives that difference through the utility's 0.05 pu and the island's sources at once, which is why the sync check of Section 7.5 stands between the operator and the breaker.

Table 7.4 Microgrid Controller Settings, Their Defaults and Allowed Ranges

Symbol	Setting	Default	Range
—	inverter mode	forming	following
H_v, m_p	virtual inertia, droop	2.0 s, 3 %	0–3, 2–5
f_{PLL}	PLL bandwidth	2 Hz	0.5–10
R	governor droop	5 %	2–8
τ_e, t_{start}	engine lag, start time	0.3, 45 s	to 1.5, 180
—	genset loading ramp	0.1 MW/s	0.02–1
—	genset start policy	on islanding	or manual, SOC
K_c, τ_I	secondary gain, integral time	0.25 MW/Hz, 2 s	0–2, 0.5–60
SOC_{lo}, SOC_{hi}	SOC band	40, 80 %	10–95
K_{soc}, r_{soc}	SOC gain, largest request	0.05 MW/%, 0.5 MW	to 0.5, 1
P_{gc}	battery schedule on the grid	0.45 MW	±1
f_1, f_2, f_3	shedding stages	59.3, 59.0, 58.7	57.5–59.8
t_{uf}	stage delay	0.2 s	0.06–2
f_{rest}, t_{rest}	restoration frequency, delay	59.9 Hz, 30 s	to 60, 300
Δf_{slip}	synchronizer slip	0.05 Hz	0.01–0.2
—	sync-check window	0.1 Hz, 10°, 5 %	to 0.5, 60, 20
t_{grid}	grid healthy first	60 s	0–300

7.5 The Microgrid Controller, Layer by Layer

The microgrid controller reads what a real one has: the bus frequency meter of Equation (7.8) and its ROCOF, the bus voltage, a power meter on every source, feeder and the PCC, the battery management system's state of charge and limits, the breaker contacts, and the two voltage transformers either side of the PCC. It never reads an EMF or an angle inside a machine. Primary control lives in the devices, the governor's droop, the inverter's virtual inertia and droop, the solar's frequency-watt response and the reactive droops of both voltage regulators. The controller adds five layers on two clocks, listed with their settings in Table 7.4. Every 20 ms it runs the shedding relays and the restoration sequence, the dispatch of the setpoints and the synchronizer, and every second the frequency restoration, the state-of-charge manager and the genset's start policy. Section 1.4.3 sets out the primary and secondary layers in general.

7.5.1 Shedding and Restoration

Stage k opens block k once the measured frequency has read under f_k for $t_{uf} = 0.2\,\mathrm{s}$, ten consecutive scans, with the bus live. The critical feeder is never shed by frequency. A dead bus opens every feeder after 0.5 s, and the critical feeder closes 2 s after the bus is live again. The blocks come back one at a time through a sequencer that the framework provides, shared/src/framework/core/sequencer.js, whose forced stop lets a relay take one stage off at once and whose minimum rest and spacing are both t_{rest}. A block returns only when the frequency has stayed above f_{rest} for t_{rest}, the forming sources have 1.2 times its load plus 0.05 MW of headroom, and t_{rest} has passed since a source last closed. The sequencer starts the lowest priority first, so block 3, the last shed, returns first. Every feeder the controller switches is fed forward. Its load is added to or taken from the secondary's total, its integral and the setpoint carrying it, so a restoration is not first seen as a frequency error, and on a shed the part of the block that a limited battery was not delivering is left out.

7.5.2 Frequency Restoration on the Shared Controller

The secondary is the velocity-form PID of Section 1.6.2, the one implementation every framework app shares, executed once a second with no derivative term and no bias. In the symbols of Equations (1.10) to (1.13) its error is $e_k = f_{SP} - \bar{f}_k$ with $P_k = K_c e_k$, where $\bar{f}_k$ is the meter's average over the period $T = 1\,\mathrm{s}$ and f_{SP} is 60 Hz or, while synchronizing, the utility's frequency plus the slip. The running sum S_k of proportional terms is never clamped, and the back-calculation with $T_t = \tau_I$ corrects the integral alone, so the clip acts only on the emitted total u_k. The limits are the forming sources' range, $[-P_{ch}, P_{dis}]$ for a battery whose SOC loop is in Auto, or its fixed manual setpoint, plus $[0.3S_g, S_g]$

while the genset is online. The tapered limits of Equation (7.7) thus reach the loop as the bounds of its output, the arrangement of Section 1.4.4 that Chapter 10 calls cascade anti-windup. A frequency below its setpoint raises the total, which is reverse action in the ISA sense and the action $a = +1$ of Section 1.6.2, as the application's direction check confirms. The total ramps to each new value over the next second rather than stepping. That ramp and the genset's 0.1 MW/s loading ramp act after the controller's output, so the integral does not see them, and while the genset is still ramping with the battery at its limit the total delivered lags u_k. The loop acts only while the campus is islanded, the bus is live, the loop is in Auto and it has an actuator, the online genset or a forming battery whose SOC loop is in Auto. Otherwise it tracks the dispatched total, $I_k = u_{now} - S_k$, so each hand-over is bumpless.

7.5.3 State of Charge and Dispatch

The SOC manager sets the battery's own share while the genset runs, as shown in Equation (7.13), a banded form of the proportional manager of Equation (1.5).

$$P_{soc} = \begin{cases} -\min\left(r_{soc},\ K_{soc}(\mathrm{SOC}_{lo} - \mathrm{SOC})\right) & \text{below} \\ 0 & \text{inside} \\ \min\left(r_{soc},\ K_{soc}(\mathrm{SOC} - \mathrm{SOC}_{hi})\right) & \text{above} \end{cases} \tag{7.13}$$

for a state of charge below, inside or above the band, with the SOC in percent, r_{soc} the largest request, and the result limited to $[-P_{ch}, P_{dis}]$. Under the by-SOC start policy the band's floor rises to its top while the genset runs, so the genset charges the battery up and can then stop. Without the manager the battery's share is its droop share of the total, $s\,u$ with $s = (S_i/m_p)/(S_i/m_p + S_g/R) = 0.455$. Every scan the genset's setpoint ramps at 0.1 MW/s toward u less the battery's share, between $0.3S_g$ and S_g, and the battery's ramps at 2 MW/s toward whatever the genset has not yet picked up, $u - P_{set,g}$, or toward all of u while it is alone. The battery therefore carries each transient and, once the genset has caught up, sits at its own share, the division of labor of Chapter 2, where a UPS carries the bus while the engines catch up. On the grid the battery follows its schedule P_{gc} at 0.2 MW/s.

7.5.4 Genset Start Policies

On islanding, the genset is cranked at the next execution of the 1 s layer after the PCC opens. Every scripted fault falls on one, so the crank follows the fault by one 20 ms step, and the genset is stopped a minute after the reconnection. By SOC, it starts when the state of charge falls below the band, when the battery has carried more than 90 % of what it may give for 5 s, or when the island has lost its grid former, and it stops above the band once the net load is under 80 % of what the battery may give. A one-stage instance of the same sequencer holds it to a minimum run of 10 min and rest of 5 min. The manual policy starts and stops it on the operator's command.

7.5.5 Synchronizer and Sync Check

Once the utility side has been healthy for t_{grid}, the synchronizer aims the secondary at the utility's frequency plus $\Delta f_{slip} = 0.05$ Hz, so the synchroscope's needle turns once in 20 s, and trims both voltage setpoints by the integral of $2(V_{grid} - |V|)$ per second within ±0.05 pu. The sync-check relay permits a close only while the frequency, phase and voltage differences are within 0.1 Hz, 10° and 5 %, and the synchronizer closes inside a quarter of the phase window. A close asked by hand outside the window is refused with its reason, a dead bus is closed onto the healthy utility without synchronizing, and a disturbance closes the breaker with no check at all, which is what the check exists to prevent. A planned island trims the battery's schedule by $5P_s$ per second until the PCC exchange P_s has stayed under 0.03 MW for a second, then opens the breaker, and abandons the attempt after 90 s.

7.5.6 Manual Mode and the Three Models

Each of the five stations, the frequency loop, the SOC loop, the feeder restoration, the synchronizer and the solar cap, has its own Manual and Auto, and the header's switch throws them all, copying

Table 7.5 Cases Run From the Reference Configuration

#	Setup	Lesson
1	grid, 09:30, 100 s, battery stepped by hand to 1.0, −0.5, 0 and 0.45 MW at 20, 40, 60 and 80 s	power moves, frequency does not
2	inverter following, 09:59, 480 s, fault at 60 s	nothing makes the voltage
3	09:59, 180 s, fault at 60 s	the first swing
4	islanded at 11:00 on the battery, genset manual, 210 s, block 2 +0.2 MW for 25 s at 10, 50, 90 and 160 s with H_v = 0, 1.5, 3 and 3 s set 5 s before, +0.4 MW from 120 s	inertia against headroom
5	09:59, 240 s, frequency loop in Manual at 0.3 MW, SOC loop in Manual at 0.45 MW, fault at 60 s, chiller at 180 s	sharing by stiffness
6	islanded at 11:59:30 with the genset, 150 s, cloud at 30 s	who answers first
7	islanded at 11:00 on the battery, genset manual, 20 % haze, 240 s, chiller at 20 s, genset started at 60 s	the last layer
8	islanded at 10:30 with the genset, SOC 50 %, 50 % haze, SOC manager off, 5400 s	the store and the average
9	09:59 at SOC 15 %, 240 s, fault at 60 s	honestly infeasible
10	islanded at 11:58:30 with the genset, 240 s, utility back at 30 s, 153.9° out of phase	the sync check
ch.	challenge, 09:59 to 12:10, fault 10:00, cloud 10:40, chiller 11:10, utility back 12:00, shipped with $H_v = 0$, SOC manager off, genset manual	a two-hour island

each controller's output into its station first, as Section 1.6.2 describes. In Manual the operator holds the genset's and the battery's setpoints, the three block breakers and the solar cap, and closes the PCC with a button that still passes through the sync check, while the droops in the devices keep acting. The simulation model is Equations (7.1) to (7.8). The controller identifies nothing from data, and the models inside it are small: the declared droops in the share s, the battery's limits and the genset's rating in the headroom test, and the feeder meters in the switched-load feedforward.

7.6 Ten Cases and a Two-Hour Island

Every case in Table 7.5 starts from one configuration, the battery forming the grid at 2 s of virtual inertia and a 3 % droop with 60 % charge, the genset cold on the on-islanding policy, the SOC manager on, the stages armed and the synchronizer automatic, and alters only the settings its row lists, with event times in seconds into the run. The values quoted below come from headless runs of the engine at the shipped 20 ms step, in simulated time, with each changed setting made at $t = 0$ before the run, and Table 7.6 gathers them. The table marks the challenge ch., since the app lists it apart from its numbered cases.

7.6.1 On the Grid, by Hand

Case 1 runs the campus on the grid at 09:30 with the SOC loop in Manual and moves the battery's setpoint every 20 s. At 1.0 MW the campus exported 0.107 MW and charging at 0.5 MW it imported 1.387 MW, so the battery's swing of 1.5 MW reappeared within 6 kW at the PCC while the meter stayed within 59.975 to 60.025 Hz. On the grid an inverter's setpoint decides only where the power comes from.

7.6.2 Following Against Forming

In case 2 the battery follows the grid. The 0.288 MW import vanished with the utility at 10:00, and with no voltage source left the bus voltage became the inverters' currents through the loads. The solar tripped at 20 ms on a ROCOF of 7.5 Hz/s, the bus was at 0.42 pu at 40 ms with both PLLs at 60.92 Hz, and it was dead at 60 ms. The battery's island detection tripped it at 2 s, the genset closed onto the dead bus at 10:00:45, and the critical feeder followed 2 s later and the blocks from 10:01:23,

Table 7.6 What the Engine Measured, Case by Case

Case and configuration	Result
1, battery at 1.0 and −0.5 MW	PCC exchange −0.107 and 1.387 MW, f within 59.975–60.025 Hz
2, following	solar trip at 20 ms, bus dead at 60 ms, 11.75 kWh critical lost
2, PLL at 0.5 Hz	0.84 pu and 60.70 Hz at 1.4 s, both trip at 2 s, 11.75 kWh lost
3, $H_v = 2$ s	2.48 Hz/s, nadir 59.516 Hz at 0.96 s, no shed
3, $H_v = 3$ and 1.5 s	1.97 and 2.82 Hz/s, nadirs 59.518 and 59.515 Hz
3, $H_v = 1.25$ and 1.0 s	5.44 and 5.83 Hz/s, solar trip, 3 sheds
3, $H_v = 0$	11.18 Hz/s, nadir 58.324 Hz, 3 sheds
4, $H_v = 0$, 1.5 and 3 s	2.61, 1.77 and 1.21 Hz/s, nadirs 59.697, 59.728 and 59.729 Hz
4, 0.07 MW headroom, 3 s	1.41 Hz/s, battery 1.085 MW, block 1 shed, nadir 59.088 Hz
4, $m_p = 2$ and 4 %	nadirs 59.793–59.827 Hz, and at 4 % 7.94 Hz/s with 3 sheds
5, $R = 5$ %	split 0.205 to 0.163 MW (1.26), $\Delta f = -0.294$ Hz
5, $R = 3$ %	split 0.251 to 0.122 MW (2.06), $\Delta f = -0.219$ Hz
5, frequency loop in Auto	60.000 Hz at 240 s, genset 0.726 MW
6, $m_p = 3$, 2 and 5 %	battery peaks 0.203, 0.269 and 0.135 MW, nadirs 59.637, 59.679 and 59.597 Hz
6, frequency loop in Manual	59.461 Hz at 150 s, battery 0.299 MW
7, shedding armed	block 1 shed at 59.23 Hz, nadir 59.182 Hz, 22.4 kWh shed
7, shedding off	BMS trip at 11:00:38, 17.23 kWh critical lost
7, $m_{ov} = 0.02$	59.509 Hz for 11.7 s over the limit, 17.57 kWh lost
8, manager off, on, floor 60 %	SOC to 9.1, 50.0 and 59.7 %, 399, 560 and 601 L
9, shedding armed	block 1 shed 0.58 s after the fault, nadir 59.024 Hz, nothing lost
9, shedding off	BMS trip 10 s after the fault, 9.26 kWh critical lost
10, synchronizer	close at −2.4° and 0.040 Hz, surge 0.327 MVA
10, no check at −153°	8.62 MVA, genset swing 8.42 MW, 57.85–61.92 Hz

about 30 s apart, after 11.75 kWh of critical load had gone unserved. A 0.5 Hz PLL only slowed the drift, to 0.84 pu and 60.70 Hz after 1.4 s, before both inverters tripped at 2 s with the same loss.

Pause. Try Exercise 7.1 before reading on.

Case 3 repeats the fault with the battery forming the grid. In the first 20 ms its output rose from 0.450 to 0.679 MW while the bus dipped to 0.958 pu, the ROCOF peaked at 2.48 Hz/s, the frequency bottomed at 59.516 Hz 0.96 s after the breaker opened, and the secondary brought it within 0.05 Hz of 60 Hz by 10.9 s, as shown in Figure 7.3. By 180 s the genset carried 0.791 MW with the battery at zero and the bus at 60.000 Hz. Figure 7.2 sets five inertias side by side. At 3, 2 and 1.5 s the worst ROCOF was 1.97, 2.48 and 2.82 Hz/s, the solar rode through, and the nadir stayed within 59.515 to 59.518 Hz. At 1.25 and 1.0 s it was 5.44 and 5.83 Hz/s. At 1.0 s the solar tripped 0.10 s after the breaker opened, the battery was asked for the solar's 1.03 MW besides the import, all three stages shed with a nadir of 58.226 Hz, and the inverter's current limit held the bus under 0.88 pu for 0.26 s, which cost even the critical feeder 0.06 kWh. With no inertia the swing read 11.18 Hz/s and the nadir was 58.324 Hz. Set too low, virtual inertia turns a fault the battery could carry into one it cannot.

TRY IT IN THE APP 3 · The Fault With a Grid-Forming Inverter

Load scenario 3 and press Run. Watch the scope of the Frequency view at 10:00, when the utility breaker opens. The bus frequency falls to 59.516 Hz about a second later and climbs back toward 60 Hz, the ROCOF trace stays inside the 3 Hz/s ride-through lines, and no load is shed.

7.6.3 Inertia Against Headroom

Case 4 islands the campus at 11:00 on the battery alone, which carries 0.574 MW of the 1 MW it may give, and steps block 2 by 0.2 MW with the inertia at 0, 1.5 and 3 s in turn. The ROCOF fell from 2.61 to 1.77 to 1.21 Hz/s while the nadir moved only from 59.697 Hz to 59.728 and 59.729 Hz, as Equation (7.10) predicts. With 0.4 MW more on block 2 the battery carried 0.930 MW with 0.070 MW of headroom, and the same step at 3 s fell at only 1.41 Hz/s but did not stop. The battery passed its continuous capability to 1.085 MW, the power-limit droop pulled the frequency down,

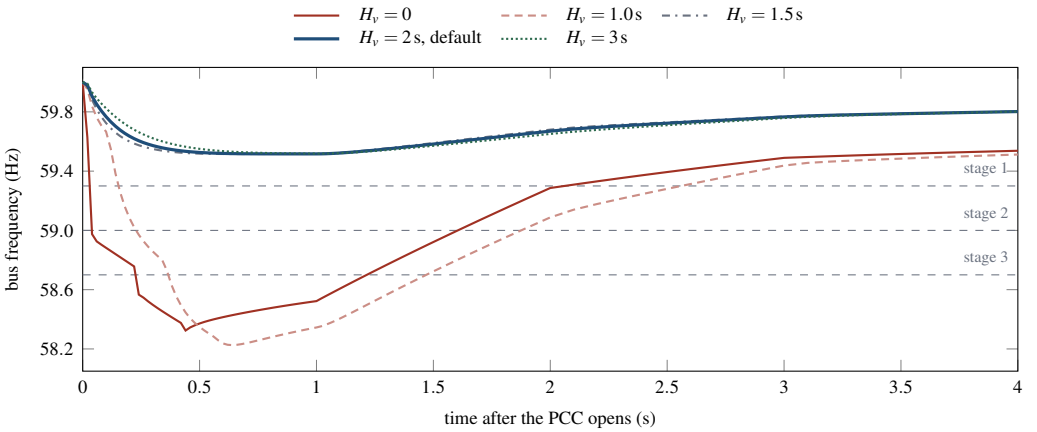

Figure 7.2 The first swing of case 3 at five virtual inertias, from the engine at its shipped 20 ms step, every step to 1.5 s and every 0.1 s after. At 1.5, 2 and 3 s the solar rides through and the curves meet at a nadir of 59.515 to 59.518 Hz. At 1.0 s and with none the swing passes 3 Hz/s, the solar trips and all three stages shed. The dashed lines are the stage thresholds.



Figure 7.3 The campus in case 3, The fault with a grid-forming inverter, at $t = 1.25$ min (10:00:15 on its clock), 15 s after the PCC opened, cropped to the Frequency view with no setting changed. (1) The badges read f 59.979 Hz, SOC 59.4 %, genset starting 30 s, and no load shed or critical load unserved. (2) The scope's last 30 s at every 20 ms show the bus falling from 60.0 Hz to the marked nadir of 59.516 Hz and recovering, and (3) the ROCOF trace under it dips toward the ± 3.0 Hz/s ride-through lines without reaching them. (4) The first-swing panel reads a virtual inertia of 2.0 s and a droop of 3.0 %, a 0.1 MW step's initial ROCOF of 1.50 Hz/s with τ 120 ms, the nadir 59.516 Hz with a worst ROCOF of 2.48 Hz/s, and the battery at +0.74 MW of the 1.00 MW it may give, a headroom of 0.26, with all three stages armed. (5) The status line reads GENSET STARTING · ONLINE IN 30 s.

stage 1 shed block 1 (0.35 MW) at 59.16 Hz, and the nadir was 59.088 Hz. A 2 % droop held the three nadirs within 59.793 to 59.827 Hz, while 4 % with no inertia stepped the frequency so far at

once that the meter read 7.94 Hz/s and all three blocks were shed. Inertia is a setting and headroom is megawatts left unused, and an island that needs a slow swing and a shallow one has to buy both.

🕮 **Pause.** Try Exercise 7.2 before reading on.

7.6.4 Two Droops and a Secondary

Case 5 takes the secondary out of the loop, with the genset held at 0.3 MW once it is online and the battery at 0.45 MW. By the end of the 45 s start the battery alone carried 0.716 MW at 59.521 Hz, and with the genset online the island sat at 59.969 Hz. A minute after the 10:02 chiller the genset had taken 0.205 MW more and the battery 0.163 MW, a ratio of 1.26, and the frequency was 0.294 Hz lower, the fall that Equation (7.11) gives for those two changes. A 3 % governor droop gave 0.251 and 0.122 MW, a ratio of 2.06, and a fall of 0.219 Hz. With the frequency loop in Auto the island sat at 60.154 Hz with the genset at its 0.6 MW minimum until the chiller, and the secondary then returned it to 60.000 Hz on the genset alone, at 0.726 MW by 240 s.

Case 6 runs the island at noon with the genset at its minimum and the battery idle in its band, and a cloud takes half of the 1.44 MW of solar over 10 s. At the bottom of the edge the genset had added 0.48 MW and the battery 0.22 MW, and the nadir was 59.637 Hz with the battery at its peak of 0.203 MW. The secondary then handed the loss to the genset, the battery falling to 0.002 MW by 60 s while the genset carried 1.30 MW, and the frequency was 60.000 Hz by 90 s. A 2 % battery droop raised the battery's peak to 0.269 MW and the nadir to 59.679 Hz, and 5 % lowered them to 0.135 MW and 59.597 Hz. With the frequency loop in Manual nothing hands the loss on, and the island stayed at 59.461 Hz with the battery giving 0.299 MW. The battery covers the first seconds, the engine the minutes, and the secondary decides when one hands over to the other.

7.6.5 The Last Layer

Case 7 islands the campus at 11:00 under 20 % haze on the battery alone, which carries 0.830 MW of its 1 MW, and starts a 0.4 MW chiller on block 1 at 11:00:20 while the operator starts the genset at 11:01. The chiller pushed the battery to 1.044 MW, past its continuous capability, and the power-limit droop pulled the frequency down until stage 1 opened block 1, chiller and all (0.63 MW), at 59.23 Hz 7.1 s after the chiller started. The nadir was 59.182 Hz, the island spent 1.6 s out of band and shed 22.4 kWh, and block 1 returned 30 s after the genset closed. With the shedding off the battery stayed over its limit until its management system tripped it at 11:00:38, and the critical feeder lost 17.23 kWh before the genset closed onto the dead bus. With the shedding armed but the power-limit droop at 0.02, the frequency held at 59.509 Hz while the battery spent 11.7 s over its limit, no stage saw its threshold, and the trip at 11:00:36 cost the critical feeder 17.57 kWh. The relays are only as good as the signal the source gives them.

🕮 **Pause.** Try Exercise 7.3 before reading on.

Case 9 faults the grid with the battery run down to 15 %, 14.6 % by the fault, where its management system allows 0.480 MW continuously against the 0.450 MW it was giving and the 0.329 MW import it must take over. With the genset 45 s away no setting keeps every feeder on. With the shedding armed the battery held the bus on its short-term allowance, stage 1 shed block 1 (0.32 MW) 0.58 s after the breaker opened at 59.10 Hz, the nadir was 59.024 Hz, 24 mHz above stage 2, and the battery then carried 0.445 to 0.448 MW while what it may give fell with its charge from 0.479 to 0.466 MW, until the genset closed at 10:00:45, with nothing lost on the critical feeder. With the shedding off the management system tripped the battery at 10:00:10 and the critical feeder lost 9.26 kWh. Shedding is a controlled failure, the HVAC given up so that the hospital wing keeps its lights.

7.6.6 Diesel Now or Battery Later

Case 8 runs the island for 90 min from 10:30 under 50 % haze with the genset online, the battery at 50 % and the SOC manager off, so the battery carries its droop share of the average. It gave about 0.58 MW for an hour and fell to 19.7 % by 11:30, after which the taper let it give less and less,

0.205 MW at 9.1 % by noon, while the genset took up the difference to 1.10 MW. The run burned 399 L. With the manager on, the battery stayed at 50.0 % carrying nothing while the genset carried 1.28 to 1.30 MW for 560 L, and the 161 L between the runs is the energy the first one borrowed from the battery. With the band's floor at 60 % the genset also recharged it, to 59.7 % by noon, for 601 L.

7.6.7 Back Through the Synchroscope

Case 10 begins islanded at 11:58:30 with the genset online, and at 11:59 the utility side returns 153.9° out of phase. The synchronizer waited 60 s for a healthy grid, aimed the island at 60.05 Hz from 12:00, and closed the PCC at 12:00:20 with the phase difference at −2.4°, the frequencies 0.040 Hz apart and the voltages matched, for a surge of 0.327 MVA, and the genset unloaded and stopped at 12:01:21. Asked to close by hand at −153°, the relay refused. Closed with no check at −153.4°, the PCC carried 8.62 MVA, the genset swung through 8.42 MW, four times its rating, and the frequency ran between 57.85 and 61.92 Hz.

> **TRY IT IN THE APP** **10 · Resynchronization Through the Synchroscope**
>
> Load scenario 10 and press Run on the Synchroscope view. The synchronizer waits 60 s for a healthy grid, slips the island and closes the utility breaker at 12:00:20 with the phases 2.4° apart, for a surge of 0.327 MVA. Asked to close by hand at −153°, the relay refuses. Closed there without the check, the breaker carried 8.62 MVA.

7.6.8 The Two-Hour Island

The challenge runs from 09:59 to 12:10, with the fault at 10:00, a cloud at 10:40, a chiller at 11:10 and the utility's return at 12:00. Its target, in teaching thresholds calibrated on this model, asks for no critical load unserved, at most 10 s outside 59.5 to 60.5 Hz, a ROCOF of at most 3.0 Hz/s, the battery within 20 to 90 % and at most 650 L of diesel, and Table 7.7 lists the configurations tried. As shipped, the fault trips the solar and sheds all three blocks, and block 1 is shed again at the cloud and at the chiller. After 12:01 the synchronizer's voltage trim raises the load past what the drained battery may give, so its management system trips it at 12:02:52 and the bus is dark for one 20 ms step before the PCC closes onto it. Removing the suggested values one at a time shows what each shipped setting costs. With the inertia back at zero the fault still trips the solar, costs 0.050 kWh and reads 11.18 Hz/s, and with the SOC manager off or the genset on manual start the battery falls to 13.0 or 8.2 %. With 2 s of inertia, the SOC manager and the genset started by SOC, the island meets all of it. The genset starts at 10:32:54 as the battery reaches 40 %, the SOC manager then asks up to 0.5 MW of charge, and the genset stops at 12:02:26 after 1.48 h online with the battery at 70.3 %. Started on islanding it also passes, with 583 L against 585 L but 2.03 h online and the battery at 55.7 %, so starting by SOC buys fewer engine hours and a fuller battery for the same diesel. As shipped the run fails four clauses of the target. With the suggested settings it meets every one, with 0.52 Hz/s of ROCOF and 65 L of diesel inside the limits. Revealing the suggested settings and then pressing Reset gives the same run. The challenge is scored on the day as scripted, so changing a model parameter or a disturbance voids the attempt. Figure 7.4 draws the seven configurations of Table 7.7 as a chart.

> **DID YOU KNOW?** **Less Diesel Is Not a Better Island**
>
> From the simulator: with the suggested settings but the SOC manager off, the two-hour challenge burns 347 L of diesel instead of 585 L and keeps the whole critical load on, yet it fails the target, because the battery falls to 13.0 % (Table 7.7).

7.6.9 The Step and the Meter

Halving the step to 10 ms keeps every conclusion that the application tests. Case 2 still collapses, dead at 50 ms instead of 60 ms, case 3 still sheds nothing with its nadir at 59.516 Hz, and case 7 sheds block 1 at 27.14 s at both steps. The reading of a first swing is another matter, because the meter

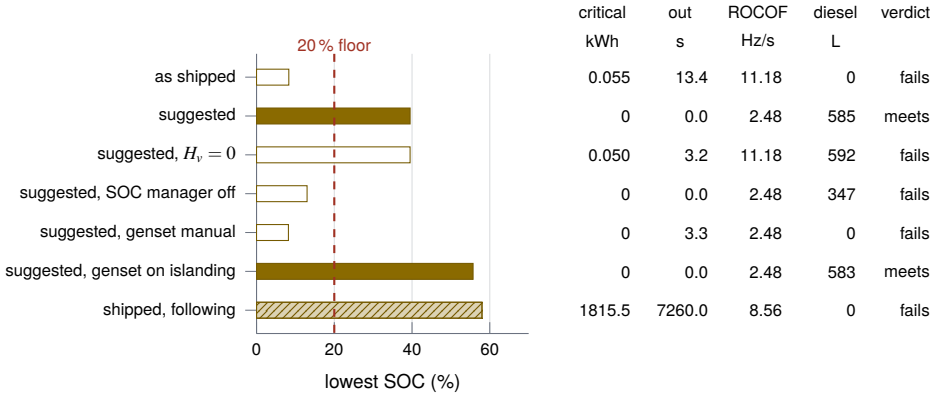

Figure 7.4 The seven configurations of Table 7.7 as a chart: the lowest state of charge over the two hours (bars), with the critical energy unserved, the time outside 59.5 to 60.5 Hz, the worst ROCOF, the diesel burned and the verdict against the target in the columns. Solid bars meet the target, white bars fail it, and the hatched bar is the shipped run with a following inverter. Notice that the shipped run, the run with the genset on manual start and the run with the SOC manager off all end below the dashed 20 % floor, and that only the two solid bars clear every clause.

Table 7.7 The Two-Hour Challenge in Seven Configurations

Configuration	Critical kWh	Out s	ROCOF Hz/s	SOC %	Diesel L
as shipped	0.055	13.4	11.18	8.3–60.0	0
suggested	0	0.0	2.48	39.5–73.7	585
with $H_v = 0$	0.050	3.2	11.18	39.5–74.0	592
SOC manager off	0	0.0	2.48	13.0–60.0	347
genset manual	0	3.3	2.48	8.2–60.0	0
genset on islanding	0	0.0	2.48	55.7–60.0	583
shipped, following	1815.5	7260.0	8.56	58.0–60.0	0

Only the suggested line and the genset on islanding meet the target.

of Equation (7.8) averages the frequency over one step, and at 2.5 ms it shows the frequency rising 24 mHz in the first 5 ms as the voltage dips. The default swing of case 3 reads 2.48, 2.62, 2.70 and 2.75 Hz/s at steps of 20, 10, 5 and 2.5 ms, converging under the 3.0 Hz/s ride-through, but at 1.5 s of inertia the readings are 2.82, 3.00, 5.20 and 5.25 Hz/s, and from 5 ms down the solar trips and all three blocks shed. The cliff that lies between 1.25 and 1.3 s at the shipped step lies between 1.6 and 1.7 s once the meter is resolved, while the 2 s default keeps its margin at every step. A ROCOF threshold is a statement about a meter as much as about the plant.

7.7 An Optional One-Hour Energy Dispatch

The Energy dispatch choice leaves Rules as the default and places a mixed-integer plan above the same frequency PI, primary controls and protection. Every 60 s it solves four 15 min commitments on `core/milp`. Its forecast is the measured load minus measured solar power held constant for one hour. It has no access to future scenario events, weather, a chiller start or a reconnection time. The state is the BMS SOC, reinitialized from its measurement at each decision, and the commitment history is the supervisor's own run and rest timer. Section 1.6.3 describes the framework's predictive options in general. This one plans integer commitments on an energy model rather than moves of a linear dynamic model.

Table 7.8 The Added Island: Fuel and Ending Stored Energy

Dispatch	Fuel (L)	SOC (%)	Adj. (L)	Crit. (kWh)
Rules	476.30	77.98	476.30	0
Retuned rules	277.83	43.45	425.43	0.125
One-hour MPC	397.47	62.18	465.02	0
One-slot plan	326.58	45.33	466.15	0

For generation G_k, discharge D_k, charge C_k, end-slot SOC s_k and binary run command u_k, the nominal energy model is

$$G_k + D_k - C_k = P_{\text{net}}, \tag{7.14}$$

$$s_{k+1} = s_k - \frac{hD_k}{\eta E} + \frac{h\eta C_k}{E}, \tag{7.15}$$

where $h = 0.25$ h, $E = 2$ MWh and $\eta = \sqrt{0.92}$. A direction binary prevents simultaneous charge and discharge. The objective prices the nominal fuel curve, 110 L/h while running plus 0.205 L/kWh generated, a cranking allowance and a small positive battery-travel price. The generator's minimum and maximum output and the present BMS limits constrain the moves, with a ten-percent BMS reserve. End-slot SOC lies between 20 and 90 percent and terminal SOC is at least the low edge of the learner's band. These are planning limits, not replacements for the device protection. The model coefficients are the declared defaults, so an edited battery capacity or efficiency is model mismatch.

The commitment grid conservatively exceeds the physical 600 s minimum run and 300 s minimum rest, and the exact sequencer remains underneath. Startup reduces first-slot generation capacity. The energy model omits ramp and frequency dynamics. The applied generator share is corrected by the frequency PI's difference from the forecast net load, while the battery carries the remainder. A stop request first unloads the genset to near zero measured output and near-setpoint frequency before its breaker may open. A rejected, non-optimal, nonfinite, fractional or out-of-bound plan clears the future and returns to the selected rules through their existing ramp. Manual stations, a following-mode battery, a dead bus or grid connection suspend dispatch planning. A browser-only 25 ms frame budget preserves ordered physics steps while reducing playback throughput when necessary.

7.7.1 A Comparison That Prices the Battery Left Behind

Case 11, Rules versus a one-hour dispatch plan, runs a 90 min island from 50-percent SOC with the genset online. Table 7.8 compares by-SOC Rules with the plan, a fixed rule retune and a one-slot counterfactual, all on the same measurements and seed. The counterfactual uses the same constraints and objective but retains only one 15 min slot. No case-specific tuning is applied.

The adjustment values the difference in ending battery energy at the nominal marginal diesel cost of replenishing it, as shown in Equation (7.16).

$$F_{\text{adj}} = F + \frac{1000(0.205)E}{\eta}(s_{\text{Rules,end}} - s_{\text{end}}) \tag{7.16}$$

It excludes additional no-load fuel and is an accounting comparison, not a simulated replenishment. Most of the apparent 78.83 L saving is energy borrowed from the battery. The one-slot counterfactual recovers almost all of the adjusted benefit, so it does not support attributing that headline to an hour of prediction. Rules, MPC and the one-slot plan serve the critical load without loss and remain in the frequency band. The plan starts the genset once more, and the switching raises the peak ROCOF from about 0.0012 Hz/s under Rules to 0.741 Hz/s under MPC.

The fixed rule retune uses a 40–60 percent SOC band, a 0.9 MW charge and discharge cap and an SOC gain of 0.2 MW/%, and the frequency and primary gains remain shipped. These settings were

declared once before measuring this row and retained on every case, without a search for a winning result. Its smaller 425.43 L adjusted number is not an admissible improvement. The original by-SOC rule stops a loaded engine and loses 0.125 kWh of critical energy in the added case, with a peak ROCOF of 4.461 Hz/s. It also loses 0.130 kWh on resynchronization. The MPC's unload-before-stop path avoids that loss. This is one fixed retune, not a globally optimal classical controller.

On the unchanged two-hour challenge with its suggested settings, matched Rules consume 585.31 L and end at 70.32-percent SOC. MPC consumes 550.64 L and ends at 51.38 percent, and its adjusted 631.64 L loses to Rules. Both lose no critical energy and spend no time outside the frequency band. The one-slot plan uses 471.73 L, ends at 36.73 percent and has an adjusted 615.31 L. Lower raw fuel is therefore insufficient evidence of a better dispatch. The fixed rule retune gives an adjusted 585.14 L on this challenge, close to Rules, while its unsafe added-case stop prevents calling it a generally better configuration.

Across all twelve scenarios, every command and output remains finite, the generator command stays at or below 2 MW and the battery command at or below 1 MW in magnitude. The matched comparison loses no additional critical energy. Switching increases ROCOF on the cloud, SOC and added cases, and each stays below 1 Hz/s. The grid-following fault remains infeasible with the planner in standby. The cold, low-SOC case invokes fallback rather than relabeling an infeasible plan. The reduced two-slot MILP agrees with exhaustive binary enumeration over sixteen operating points, and injected non-optimal, nonfinite, fractional and infeasible solutions are rejected. The ten earlier cases and the challenge, shipped and with the suggested settings, remain identical on all their original logged columns. The persistence forecast, the current-limit hold and the coarse energy slots limit the interpretation. The section compares dispatch policies on this campus and makes no claim of economic superiority.

7.8 Technology in Practice

The layers of Section 7.5 run on instrument transformers, relays, breakers and inverters. Figure 7.5 draws the campus's loops with that hardware in place of the schematic's symbols.

What Measures It

Every reading of the microgrid controller starts at an instrument transformer. Potential transformers step the line voltage down to a safe level, typically 120 V, and current transformers step the line current down, typically to 5 A at full load, for the meters and the protective relays to read [20]. The chapter's controller reads the bus frequency meter FT-100 with its ROCOF, a power meter on every source and feeder and on the PCC, the voltage transformers on either side of the PCC, and the charge and limits that the battery's management system reports.

> **TECHNOLOGY** **Sensor: A Frequency Relay on the Bus Voltage**
>
> No instrument reads frequency directly. A relay or meter takes the bus voltage through a potential transformer and estimates the frequency from it, by a phase-locked loop, a Fourier transform, the zero crossings or the change of the voltage phasor's angle [19, 20]. The estimate needs a few cycles, typically 90 to 120 ms, and about 30 mHz of accuracy is typical for underfrequency load shedding. A ROCOF for protection is computed over 180 to 240 ms to within 50 mHz/s. The common failure is a false reading: a short circuit or a breaker operation makes the voltage angle jump, and the local frequency spikes, which can shed load or disconnect solar installations [19]. The chapter's meter reads the angle over each 20 ms step and ROCOF over 0.1 s.

What Moves It

The island is moved by setpoints and breakers. The genset's governor and the battery inverter take power setpoints, and the solar inverter takes a cap. Alternating-current power circuit breakers, device 52 in the IEEE C37.2 numbering, open and close the PCC, the sources and the feeders [20]. The

microgrid controller: 81, JY-100 and 25 every 20 ms, FIC-100 and the SOC manager every 1 s. Relays inside the genset and the inverters trip them apart from it.

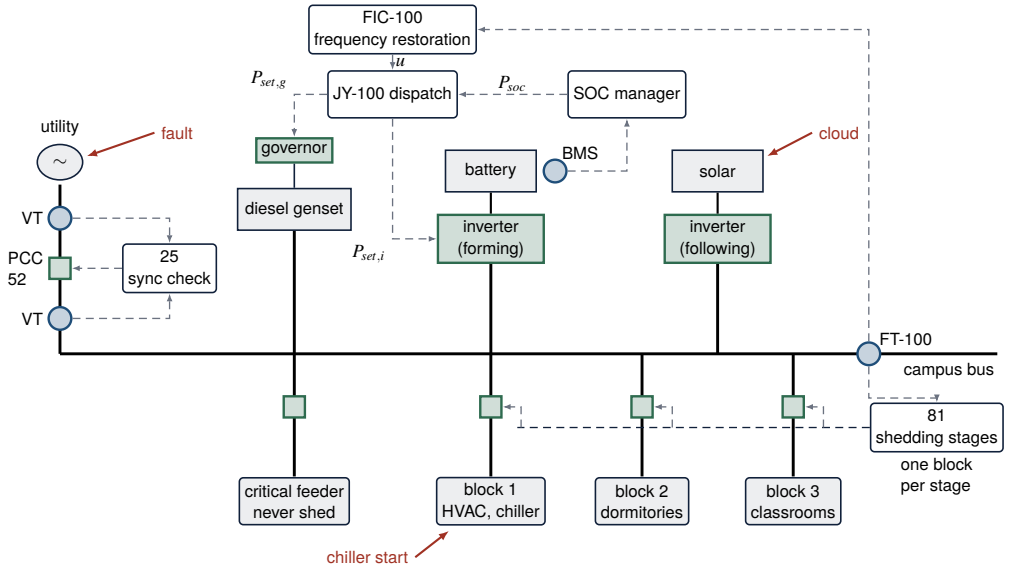

Figure 7.5 The campus's loops drawn with their hardware. The frequency meter FT-100 reads the bus voltage through a potential transformer. The frequency restoration FIC-100 sets the total u, the dispatch JY-100 splits it into the genset's setpoint, carried out by its governor, and the battery inverter's, with the share that the SOC manager asks from the battery management system's reading. The shedding stages 81 open one block breaker each on the same meter, and the sync check 25 compares the voltage transformers either side of the PCC breaker before it may close. Notice that one bus meter feeds both the slow restoration loop and the fast shedding stages.

fastest final element is the battery inverter itself, whose control law, not a moving part, decides the first swing.

> **TECHNOLOGY** **Final Element: The Grid-Forming Battery Inverter**
>
> The battery reaches the bus through an inverter, a power stage of semiconductor switches modulated by a digital controller. Run grid-forming, its terminals behave like a voltage source whose amplitude and frequency vary with its reactive power and with the load, so it adjusts its output nearly instantaneously [13]. Unlike a machine, it cannot give a large current for long. Its short-circuit current is normally held to 4 to 6 times rated for less than ten cycles, a limit set by the semiconductors, and then to less than twice rated [13]. Its behavior is set by firmware, which can change over its service life [13]. The chapter's inverter holds its current under 1.4 pu by raising a virtual impedance.

Where the Controller Runs

IEEE 2030.7 specifies a microgrid controller by the functions above the level of component control that let a microgrid manage itself, run on its own or connected to the grid, and connect to and disconnect from the distribution grid [21]. Supervisory and local control loops of this kind run at cycle times from milliseconds to minutes [22]. The chapter's controller scans every 20 ms for the shedding stages, the restoration sequence, the dispatch and the synchronizer, and every second for the frequency restoration, the SOC manager and the genset's start policy. The primary droops live in the devices, and the relays inside each device trip it apart from the controller, as Table 7.2 lists.

> **IN PRACTICE** **Closing Back Onto the Grid**
>
> Before a breaker joins two live systems, the voltages on its two sides are compared, by operators or by synchronizing relays, and the breaker is not closed until their difference is nearly zero. When the frequencies differ, the two voltages drift in and out of phase at the difference frequency, so a lamp across the open breaker brightens and dims at that rate [20]. Relays are known by the function numbers of IEEE C37.2: 25 is the synchronism check, 81 the under- and overfrequency and 81R the rate-of-change-of-frequency functions, 32 directional power, 27 undervoltage and 59 overvoltage, and technicians and engineers alike call relays by these numbers [20].

What Is Modernizing It

Three developments touch the loops of this chapter. Grid-forming control is expected to move from microgrids to bulk power systems in stages over roughly 10 to 30 years, once protection, controls and interoperability are researched and standards define what a grid-forming inverter must do [13]. Phasor measurement units time-stamp voltage phasors against a satellite clock with sub-microsecond uncertainty, so that angles across a grid can be compared at one instant, which helps to detect an island when breakers split a grid and its generators drift out of step [20]. Ethernet, Internet Protocol and wireless devices replacing older proprietary networks raise the likelihood of cybersecurity vulnerabilities and incidents in control systems [22], a concern for a controller that opens and closes a campus's breakers. The energy dispatch of Section 7.7 is a fourth development, and the chapter measures what it buys and what it costs.

7.9 Limitations

The model is one bus solved as phasors, so it cannot show the inverters' current-control transients in the first cycle after a breaker operates, a fault inside the campus, unbalance or harmonics, and a collapse that the model reports within a few steps may take several cycles on a real feeder. The genset is a classical machine with a timed start and a perfect synchronizer, the grid-forming inverter is one virtual synchronous machine with one current-limiting scheme, and the loads are static, with no motor dynamics or inrush. The relays act on an idealized meter with no filtering of their own, and the first swing's ROCOF depends on the meter's step, as Section 7.6 measured. The scripted days, a fault at 10:00 sharp and a cloud that halves the solar in exactly 10 s, would not arrive so cleanly on a real campus, although they let one layer at a time be seen.

> **OPEN CHALLENGE** **Protection When the Sources Are Inverters**
>
> Traditional protection relies on the high fault currents of synchronous generators, and inverters give much less. Whether grid-forming inverters in place of machines will make protection misoperate more often or less is an open question, and their effect on protection has not been studied extensively. Inverter makers provide no well-defined set of fault models and tests, and a firmware change can alter the fault current and invalidate a protection design [13]. The chapter models no fault inside the campus. Progress would be a consistent framework of models and tests for inverter fault behavior, and adaptive protection that knows which sources are online [13].

Every threshold in the simulator is a teaching value. The following inverters' trips and the ROCOF ride-through are the defaults that the application takes from the Category III settings of IEEE 1547-2018, the shedding stages at 59.3, 59.0 and 58.7 Hz after 0.2 s are patterned after bulk-system programs, and the sync-check window of 0.1 Hz, 10° and 5 % is close to what the application quotes as the standard's limits for a resource above 1.5 MVA, 0.1 Hz, 10° and 3 %. Real ride-through and trip settings come from IEEE 1547-2018 and the interconnection requirements of the utility [9], the design of an island system and its separation and reconnection from IEEE 1547.4 [17], the stages and the restoration of a shedding program from the guidance of IEEE C37.117 and the setting methods of the transmission literature [11, 12], synchronizing practice from the relaying literature [18], and

every number for a real campus from its own protection study. The battery's taper and allowance, the genset's trips and the challenge's targets are illustrative, and none of the numbers measures the performance of an installed microgrid.

OPEN CHALLENGE How Much Headroom Is Enough

With less inertia, the frequency falls faster after a disturbance and control must act sooner. The depth of the swing is tied to the imbalance, so untapped capacity must be held as headroom, in storage, spinning generators or renewable inverters run below their available power, and that unused capacity has an opportunity cost [13]. The same roadmap leaves open what reserves keep the frequency within bounds when grid-forming sources carry it, and what mix of machines and inverters guarantees stability [13]. This chapter prices headroom for one campus. Progress would be reserve rules for grids where inverters dominate.

7.10 Conclusion

An islanding campus needs a source that makes the voltage, and this chapter has measured what that source and each layer above it buy. A battery inverter that follows the grid leaves the bus dead 60 ms after the fault and costs the critical feeder 11.75 kWh, while the same battery forming the grid with 2 s of virtual inertia holds the island through a first swing of 2.48 Hz/s, under the solar's 3 Hz/s, and a nadir of 59.516 Hz. Virtual inertia sets the rate of a fall, while the droop and the headroom decide where it stops. The genset and the battery share a load in the ratio of their stiffnesses, 1.26 measured against 1.26 declared, the secondary returns the frequency to 60.000 Hz on the genset, and the SOC manager decides whether the battery or the diesel pays for the hours. Shedding saves the critical feeder where nothing else can, provided the battery's power-limit droop lets the relays hear the overload, and the sync check turns an 8.62 MVA close into a 0.327 MVA one. The two-hour challenge improves from 0.055 kWh of critical load lost and 13.4 s out of band to neither. The optional energy dispatch lowers the raw diesel only by leaving the battery emptier, and on the challenge its adjusted fuel loses to the rules.

The model is a single phasor bus with an idealized meter and teaching thresholds. A fuller model would add feeder impedances and faults inside the campus, a second grid-forming inverter, electromagnetic detail for the first cycle, and a dispatch that plans against a forecast of load and sun rather than on persistence. The simulator shows which source makes the grid when the utility is gone, and what every setting above that source costs.

KEY POINTS

- When the utility breaker opens, one of the island's own sources has to make the voltage that the others follow.
- A battery inverter that follows the grid leaves the bus dead 60 ms after the fault, while the same battery forming the grid with 2 s of virtual inertia holds the island through a first swing of 2.48 Hz/s.
- Virtual inertia sets the rate of a fall, while the droop and the headroom decide where it stops.
- Shedding saves the critical feeder where nothing else can, provided the battery's power-limit droop lets the relays hear the overload.

Where else this applies. Which loop owns a shared variable and how the other sources divide the load is also the question of the steam header of Volume 18, Chapter 9, and a governor on an islanded load runs Volume 1, Chapter 7.

References

[1] R. H. Lasseter, "MicroGrids," in *2002 IEEE Power Engineering Society Winter Meeting. Conference Proceedings (Cat. No.02CH37309)*, vol. 1, New York, NY, USA, 2002, pp. 305–308.

[2] J. M. Guerrero, J. C. Vasquez, J. Matas, L. G. de Vicuna, and M. Castilla, "Hierarchical control of droop-controlled AC and DC microgrids—a general approach toward standardization," *IEEE Transactions on Industrial Electronics*, vol. 58, no. 1, pp. 158–172, 2011.

[3] D. E. Olivares, A. Mehrizi-Sani, A. H. Etemadi, C. A. Canizares, R. Iravani, M. Kazerani, A. H. Hajimiragha, O. Gomis-Bellmunt, M. Saeedifard, R. Palma-Behnke, G. A. Jimenez-Estevez, and N. D. Hatziargyriou, "Trends in microgrid control," *IEEE Transactions on Smart Grid*, vol. 5, no. 4, pp. 1905–1919, 2014.

[4] J. Rocabert, A. Luna, F. Blaabjerg, and P. Rodríguez, "Control of power converters in AC microgrids," *IEEE Transactions on Power Electronics*, vol. 27, no. 11, pp. 4734–4749, 2012.

[5] M. C. Chandorkar, D. M. Divan, and R. Adapa, "Control of parallel connected inverters in standalone AC supply systems," *IEEE Transactions on Industry Applications*, vol. 29, no. 1, pp. 136–143, 1993.

[6] F. Katiraei, M. R. Iravani, and P. W. Lehn, "Micro-grid autonomous operation during and subsequent to islanding process," *IEEE Transactions on Power Delivery*, vol. 20, no. 1, pp. 248–257, 2005.

[7] R. H. Lasseter, Z. Chen, and D. Pattabiraman, "Grid-forming inverters: A critical asset for the power grid," *IEEE Journal of Emerging and Selected Topics in Power Electronics*, vol. 8, no. 2, pp. 925–935, 2020.

[8] Q.-C. Zhong and G. Weiss, "Synchronverters: Inverters that mimic synchronous generators," *IEEE Transactions on Industrial Electronics*, vol. 58, no. 4, pp. 1259–1267, 2011.

[9] *IEEE Standard for Interconnection and Interoperability of Distributed Energy Resources with Associated Electric Power Systems Interfaces*, IEEE Std 1547-2018, 2018.

[10] P. Kundur, *Power System Stability and Control*, ser. The EPRI Power System Engineering Series. New York: McGraw-Hill, 1994.

[11] H. Lokay and V. Burtnyk, "Application of underfrequency relays for automatic load shedding," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-87, no. 3, pp. 776–783, 1968.

[12] *IEEE Guide for the Application of Protective Relays Used for Abnormal Frequency Load Shedding and Restoration*, IEEE Std C37.117-2007, 2007.

[13] Y. Lin, J. H. Eto, B. B. Johnson, J. D. Flicker, R. H. Lasseter, H. N. Villegas Pico, G.-S. Seo, B. J. Pierre, and A. Ellis, "Research roadmap on grid-forming inverters," National Renewable Energy Laboratory, Golden, CO, Technical Report NREL/TP-5D00-73476, 2020.

[14] H.-P. Beck and R. Hesse, "Virtual synchronous machine," in *2007 9th International Conference on Electrical Power Quality and Utilisation*, 2007, pp. 1–6.

[15] S. D'Arco and J. A. Suul, "Equivalence of virtual synchronous machines and frequency-droops for converter-based MicroGrids," *IEEE Transactions on Smart Grid*, vol. 5, no. 1, pp. 394–395, 2014.

[16] *IEEE Standard for Interconnecting Distributed Resources with Electric Power Systems*, IEEE Std 1547-2003, 2003.

[17] *IEEE Guide for Design, Operation, and Integration of Distributed Resource Island Systems with Electric Power Systems*, IEEE Std 1547.4-2011, 2011.

[18] M. J. Thompson, "Fundamentals and advancements in generator synchronizing systems," in *2012 65th Annual Conference for Protective Relay Engineers*, 2012, pp. 203–214.

[19] ENTSO-E RG-CE System Protection & Dynamics Sub Group, "Frequency measurement requirements and usage," European Network of Transmission System Operators for Electricity, Brussels, Report Final Version 7, 2018. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/SOC%20documents/Regional_Groups_Continental_Europe/2018/TF_Freq_Meas_v7.pdf

[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] *IEEE Standard for the Specification of Microgrid Controllers*, IEEE Std 2030.7-2017, 2018.

[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.

7.11 Exercises

7.1 ⌕ Concept check · 2 min

In case 2 the battery inverter follows the grid. The PCC opens at 10:00 and the bus is dead 60 ms later. Which statement explains the collapse?

(a) No source on the island makes the voltage, so the bus voltage is the inverters' currents through the loads and each PLL chases an angle it turns itself

(b) The battery cannot add the 0.288 MW the campus was importing to its own output, so its current limit pulls the bus voltage down

(c) A PLL bandwidth of 2 Hz is too fast for an island, and at 0.5 Hz the PLLs would hold the bus until the genset closes

(d) The solar trips on its ROCOF relay at 20 ms, and the battery's island detection takes it off the bus with the solar

7.2 Concept check · 1 min

In case 4 the lone grid-forming battery takes three 0.2 MW load steps at virtual inertias of 0, 1.5 and 3 s. What does the virtual inertia change?

(a) The nadir, which rises by about 0.3 Hz from no inertia to 3 s

(b) The initial rate of change of frequency, while the fall stops at about the same depth

(c) The battery's headroom, which grows with the kinetic energy the inertia stands for

(d) The settled frequency, because the inertia adds to the stiffness of the droop

7.3 Concept check · 2 min

Why does the grid-forming inverter carry a power-limit droop m_{ov} in addition to its frequency droop m_p?

(a) To share the island's load with the genset in proportion to the battery's charge

(b) To hold the inverter's current under its hardware limit of 1.4 pu through a fault

(c) To turn an overdraw of the battery into a falling frequency that the shedding relays can see

(d) To slow the first swing after islanding so that the solar rides through its ROCOF limit

7.4 Theory and calculation · 6 min

The battery inverter forms the grid with $S_i = 1\,\text{MW}$, $f_0 = 60\,\text{Hz}$ and a droop $m_p = 3\,\%$. In case 3 it takes over an import of $\Delta P = 0.3293\,\text{MW}$ when the PCC opens.

(a) Use Equation (7.9) to compute the initial ROCOF at $H_v = 2$ and $H_v = 1.5\,\text{s}$. *Checkpoint:* 4.94 and 6.59 Hz/s.

(b) What virtual inertia would hold the bound at the 3.0 Hz/s the solar rides through? Compare it with the slider's range, and give two reasons from the chapter why the relay's reading at 2 s, 2.48 Hz/s, is smaller than the bound. *Checkpoint:* 3.29 s, beyond the 3 s top of the slider.

(c) With the battery alone on the bus, compute the time constant of Equation (7.10) at $H_v = 2\,\text{s}$ and the settled deviation after a 0.2 MW load step on the droop alone. Which of the two depends on H_v? *Checkpoint:* 120 ms and 0.36 Hz. Only the time constant.

7.5 Theory and calculation · 6 min

The genset has $S_g = 2\,\text{MW}$ and the rotor damping $D = 1.0\,\text{pu/pu}$, and the battery inverter has $S_i = 1\,\text{MW}$ on a droop $m_p = 3\,\%$. Both act on one frequency, with the secondary loop in Manual.

(a) Compute the stiffnesses $K_g = S_g/R + DS_g$ and $K_i = S_i/m_p$, in megawatts per unit of speed, and their ratio at a governor droop $R = 5\,\%$. *Checkpoint:* 42 MW, 33.3 MW and 1.26.

(b) Repeat for $R = 3\,\%$. *Checkpoint:* 68.7 MW and a ratio of 2.06.

(c) In case 5 the chiller makes the genset take 0.205 MW more and the battery 0.163 MW more at $R = 5\,\%$, and 0.251 MW and 0.122 MW at 3 %. Use Equation (7.11) to compute the frequency fall in each case, to the nearest millihertz. *Checkpoint:* 0.293 Hz and 0.219 Hz.

7.6 Theory and calculation · 6 min

Case 8 runs a 90 min island with the genset online. With the SOC manager off, the battery falls from 50.0 to 9.1 % and the run burns 399 L. With the manager on, the battery stays at 50.0 % and the run burns 560 L. The battery holds $E_{cap} = 2\,\text{MWh}$ with a round trip $\eta^2 = 92\,\%$, and the engine burns $q_0 = 110\,\text{L/h}$ running plus $q_1 = 0.205\,\text{L/kWh}$.

(a) Compute the DC energy drawn from the battery in the first run and the AC energy it delivered, using Equation (7.7). *Checkpoint:* 0.818 MWh DC and 0.785 MWh AC.

(b) How much diesel would the genset burn to make that AC energy at its marginal rate? Compare it with the difference between the two runs. *Checkpoint:* 161 L, equal to 560 − 399.

(c) With the manager on, the genset carries 1.28 to 1.30 MW for the 1.5 h. Estimate the run's diesel from the fuel line and compare it with the 560 L measured. *Checkpoint:* 559 to 565 L.

7.7 App challenge · 9 min

Load scenario 3, The fault with a grid-forming inverter. The model has no random input, so every run repeats exactly. Run each part at 5× past 10:00:05 ($t = 65\,\text{s}$, about 13 s of wall time) and read the Nadir · worst ROCOF tile and the Load shed tile.

(a) Run as shipped, with $H_v = 2.0\,\text{s}$. *Checkpoint:* 59.516 Hz, 2.48 Hz/s, 0 stage trips.

(b) Press Reset, set the Virtual inertia to 1.0 s on the VSM card before pressing Run, and run again. What happened to the solar, and why did all three stages shed? *Checkpoint:* 58.226 Hz, 5.83 Hz/s, 3 stage trips.

(c) Find the smallest virtual inertia on the slider's 0.1 s grid at which the solar still rides through, pressing Reset and setting the value before each run. State the tile readings at that setting and one step below. *Checkpoint:* 1.3 s rides through at 2.97 Hz/s, and 1.2 s reads 5.52 Hz/s and sheds three stages.

7.8 App challenge · 8 min

Load scenario 5, The genset starts: battery alone, then droop sharing. The frequency loop and the SOC loop are in Manual. Run to the end, 240 s, at 20× (12 s of wall time) and export the CSV file.

(a) From the columns `P_gen_MW`, `P_bat_MW` and `f_Hz`, read the row at `time_s` 180 and the last row. Compute each source's change, the ratio of the changes and the frequency fall, to within 0.002 MW, 0.01 and 0.002 Hz. *Checkpoint:* 0.205 MW and 0.163 MW, a ratio of 1.26, a fall of 0.294 Hz.

(b) Press Reset, set the Governor droop to 3 % on the GEN card before pressing Run, and repeat. *Checkpoint:* 0.251 MW and 0.122 MW, a ratio of 2.06, a fall of 0.219 Hz.

(c) Compare both with your answers to Exercise 7.5, and say why the frequency stays low after the chiller. *Checkpoint:* The droops share the change and nothing moves a setpoint.

7.9 App challenge · 9 min

Load scenario 7, Load shedding: a step the sources cannot meet. Run each part to the end, 240 s, at 20× (12 s of wall time).

(a) Run as shipped and read the Critical load unserved, Load shed and Nadir · worst ROCOF tiles. *Checkpoint:* 0.0 kWh, 22 kWh with 1 stage trip, 59.182 Hz.

(b) Press Reset, turn the Under-frequency load shedding off on the UFLS card before pressing Run, and run again. Read the Critical load unserved tile, and from the CSV file the `time_s` of the first row whose `status` reads BLACKOUT, within 0.5 s. *Checkpoint:* 17.2 kWh and 38.5 s.

(c) Reload the scenario with the reload button beside the scenario list, so the shedding is armed again, open Show Model, set the Power-limit droop to 0.02 and run once more. Read the same tile, the Load shed tile, and `f_Hz` in the rows from 33 to 36 s. Explain why no stage shed. *Checkpoint:* 17.6 kWh, 0 stage trips, `f_Hz` 59.509 to 59.511.

7.10 Application area · 3 min

Two voltages of 1 pu that differ in phase by $\Delta\varphi$ put $2\sin(\Delta\varphi/2)$ per unit across an open breaker.

(a) Compute it at the edge of the sync-check window, 10°, and at the 153.9° of case 10 when the utility returns. *Checkpoint:* 0.174 pu and 1.95 pu.

(b) The synchronizer aims the island 0.05 Hz above the utility rather than at it. Why must the slip be nonzero, and how long does the needle take to turn once? *Checkpoint:* With no slip the phase difference never moves into the window. 20 s.

7.11 Application area · 4 min

A commissioning engineer runs case 3 on the app, finds that the solar rides through at a virtual inertia of 1.3 s, and proposes 1.3 s for the real campus to leave more of the battery's headroom for load.

(a) Using the chapter's results on the step and the meter, say why 1.3 s is not a safe setting and what margin the 2 s default keeps. *Checkpoint:* The resolved cliff lies between 1.6 and 1.7 s. 2 s stays under the ride-through at every step.

(b) Name two effects the model leaves out, from its Assumptions and Limitations, that would make the first swing harder on the real campus. *Checkpoint:* For example motor inrush and the first cycle of current control.

7.12 Review the agent · 5 min

An AI assistant compares the two energy dispatches on the challenge with the suggested settings and replies as follows. "With Energy dispatch on MPC the campus burns 550.64 L against 585.31 L under Rules, a saving of 34.67 L, with no critical load lost and no time out of band. The one-hour forecast earns the saving, so the campus should run MPC." The runs end at 51.38 and 70.32 % of charge. Find the flaw and say how you would check it.

CHAPTER 8

Hunting, Reverse Flow and an Unmeasured Customer on a Solar Distribution Feeder

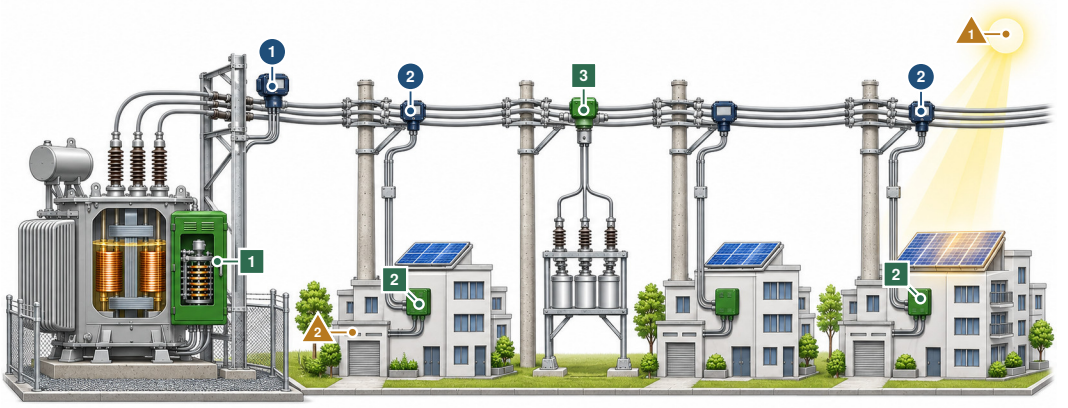

The controller moves 1 tap changer on the substation transformer 2 solar inverters on every block 3 capacitor bank switch

The controller reads 1 voltage and power at the substation 2 three meters along the line

What disturbs it 1 changing sunlight 2 household demand

Notice that the tap changer sits at the substation end while the sunlit roofs at the far end push the voltage up, and that only three meters see the line between them. Illustration, not to scale.

IN PLAIN TERMS **Holding Every Home's Voltage on a Sunny Street**

The main idea. A power line loses voltage along its length, so a transformer at its start steps the voltage up or down for every home at once. Rooftop solar pushes the far end up instead, and at noon power can flow backward.

The challenge. The transformer waits before each step and watches only its own end, so it misses a cloud over the far end. The panels' electronics answer in seconds, but if they react too hard, they and the transformer undo each other's work.

Why it matters. Every home needs a voltage inside a narrow range, all day and in any weather.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the history box of Section 8.2, the try-it and fact boxes of Section 8.6, and the key points
Study	50 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 8.6, Section 8.7, Section 8.8 and its open challenge

The control question A substation tap changer moves the whole feeder in steps after a timer, while rooftop solar reverses the flow at noon, clouds move the far end in seconds and inverters answer at every node. Which device should set the level and which should trim the spread, and what does a central controller that reads only the utility's measurements add?

Learning objectives

- Compute the line-drop compensator's correction under reverse flow and state when it raises the bus instead of lowering it.
- Relate the bandwidth and the time delay of a tap changer to its operations count and to the lowest customer voltage.
- Explain how a steep Volt-VAR slope and a reactance-weighted compensator make the tap changer hunt, and name the settings that remove it.
- Measure a conservation voltage reduction factor and say why it depends on the composition of the load.
- Compare a linear program and a mixed-integer predictive controller on the same estimate in minutes out of band and tap operations.

Prerequisites Sections 1.3 (loop action), 1.4.3 (Volt-VAR droop and the tap changer as the slow layer), 1.6.1 and 1.6.4 (engine, scenarios and scorecard). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/feeder-volt-var/

See also Volume 9, Chapter 11 for integer decisions by branch and bound and Volume 9, Chapter 5 for releases planned ahead of demand. Volume 5, Chapter 3 for a loop behind a late measurement and Volume 5, Chapter 11 for a limiting state that no instrument reads.

8.1 Introduction

At noon on a clear summer day in the simulator of this chapter, the twelve blocks of houses on a 9.6 km feeder make more power than they use, and 1.08 MW flows back through the substation transformer. The houses' air conditioners still draw reactive power, so 0.64 Mvar keeps flowing forward. The tap changer's line-drop compensator, set on the evening's load with $R_{set} = 0.01$ pu and $X_{set} = 0.15$ pu, reads that reactive power as load and holds the bus 0.0083 pu above its setpoint, more than one tap step of 0.006 25 pu. The far end of the feeder passes the 1.05 pu limit at 10:44 and spends 203 min above it before 15:00, while every instrument at the substation reads a voltage inside its band.

The scene contradicts three readings of the problem that a block diagram invites. The first is that the substation sees the feeder. When a bank of cloud reaches the far end, the far end falls 0.031 pu and the bus 0.001 pu, so the regulated voltage never stays outside its band for the 45 s a step needs. The second is that devices that each regulate correctly alone will regulate correctly together. With the inverters' Volt-VAR slope raised from the default 7.33 to 22 on a cloudy day, the tap changer makes 18 operations between 09:00 and 17:00 instead of 2, because its compensator reads the inverters' reactive swing as load coming and going. The third is that a compensator designed for forward flow fails in a known direction when the flow reverses. Its correction is the in-phase part of a phasor drop, and whether it raises or lowers the bus depends on the ratio of its own reactance to resistance against the ratio of exported to imported power.

The class of problem is general. A slow, discrete, integrating actuator with a deadband and a timer shares a plant with fast proportional devices that it cannot see, and each reads the other's action as a disturbance. It appears between substation tap changers and switched capacitors [1], and distributed generation added voltage rise, reverse flow and interaction with regulators and capacitors [2, 3]. Photovoltaic generation forces tap changers and line regulators to operate almost continuously, and uncoordinated reactive power control can drive a regulator to its limit [4]. In process terms the tap changer is a large valve with a slow positioner, the inverters are small fast trim valves, and the compensator is a feedforward computed from a flow meter that the trim valves also move.

These interactions are usually studied on detailed feeder models [4, 5], and the interconnection standard states the curves an inverter must offer rather than what they do to a regulator [6]. This chapter works on a smaller plant that turns the bandwidth, the timer, the compensator dials, the Volt-VAR slope and the meter delay into settings whose effects read in tap operations, minutes out of band and kilowatt-hours. It describes the feeder simulator in enough detail to rebuild it and measures what each local device and the central program contribute when solar power reverses the flow and clouds move it. The power flow, the curves, the compensator and conservation voltage reduction follow the cited literature and standards. The hunting of a reactance-weighted compensator against steep Volt-VAR curves belongs to a class of regulator runaway already described [4], and it is reported here as something measured on this teaching plant, within a stated window. The model has 29 states, and its energy ledger closes to 1.6×10^{-14} over a day. The tap changer and the capacitor are explicit discrete controllers, the estimator and the linear program read only what a utility measures, and a day runs in 24 s in a browser. Ten cases, graded in the application, range from a tap moved by hand to a cloudy day at 125 % solar. An optional mixed-integer model predictive controller on the same estimate and weights is measured against the program and against stronger forms of it.

Section 8.2 traces how feeder voltage regulation reached its present form. Section 8.3 gives the model, Section 8.4 four mechanisms that explain the cases, and Section 8.5 the controllers, with the predictive one in Section 8.5.5. Section 8.6 measures the nine cases and the challenge, Section 8.6.10 sets the predictive controller beside the program and its stronger forms, and Section 8.6.11 states what was verified and what is not claimed, Section 8.7 describes the hardware that does this work on a real feeder, and Section 8.8 lists the limitations.

8.2 Voltage Regulation on Radial Feeders

A radial feeder loses voltage along its length in proportion to the current it carries, so the voltage a distant customer receives falls as the load rises. The first remedy was to regulate the sending end on a measurement of the far end. A voltage-regulating relay combined with a line-drop compensator, described in 1942, combined the voltage at the instrument with a measured line drop so that the relay acted on the voltage at a remote point of the feeder, and it named tap-changing transformers and capacitor switching as the equipment it would control [7]. The step-voltage regulator with that compensator is now textbook material for distribution engineers [8]. Supervisory schemes later used the measurements available at substations to improve both the regulator and the capacitor controllers there [1].

Running the voltage low on purpose has its own history. A 1978 experimental investigation related the energy and demand a utility saved by voltage reduction to the composition of its load [9], and a national evaluation in 2010 simulated conservation voltage reduction on 24 prototypical feeders with appliance-level load models and found energy reductions of about 0.5 to 4 %, nearly all of it in the end-use loads [10]. Analysis of the feeder itself became systematic when the DistFlow branch equations were introduced for capacitor sizing on radial systems [11], and power-flow approximations built on them guided feeder reconfiguration for loss reduction [12]. A weighted least-squares state estimator for distribution systems was then proposed to improve forecast load data with real-time measurements [13].

Distributed generation changed the assumptions every one of these devices was set on. Interconnected resources raise the voltage, interact with capacitor, regulator and tap changer operations, and break the premise that the substation is the sole source of power [2]. Voltage control by on-load tap changers in medium-voltage feeders was then examined in the presence of distributed generation [3]. How inverters came to be allowed to hold the voltage is told in the history box below. In 2011 an industry collaboration defined Volt-VAR as one of the smart inverter functions for high penetration [14]. Local Volt-VAR control by photovoltaic inverters was shown to mitigate voltage rise on detailed feeders [5], an optimal coordination of reactive power was proposed to limit the tap operations and the regulator runaway that photovoltaics cause [4], and the optimality and stability of local reactive

Table 8.1 Variables by Role, With Values at 12:00 in Case 3

Variable	Role	Value	Range
V_j	CV of the challenge, node voltages (truth)	1.040–1.057	0.95–1.05
V_{reg}	CV of the tap changer, computed	1.0295	$V_{set} \pm BW/2$
V_{bus}	measured, bus voltage	1.0394	—
P, Q	measured, into the transformer	−1.093, 0.649	MW, Mvar
n	MV, tap position (integer)	+7	−16 to +16
V_{set}	setpoint, or MV of the program	1.025	0.95–1.05
u_j	MVs of the program, reactive offsets	0	± capability
b_{cap}	MV, capacitor in or out	out	0 or 0.06 pu
$p_{L,j}$, κ_j	DVs, load and sunlight share	—	shape, seeded
V_{grid}	DV, transmission voltage	1.00	0.97–1.03
V_3, V_7, V_{10}	measured, meters, 15 min old	1.045–1.053	—
q_j, w_j	states, inverter Q and Volt-Watt limit	0, 1	12 each
E	states, five energy ledgers	—	pu s
s_L, s_{PV}	estimator, load and solar scales	0.989, 0.994	see text

control under limited capability were analyzed [15]. The design of the curves within the ranges that IEEE 1547-2018 allows has itself become an optimization problem [16].

> ⧖ HISTORY **Inverters Join Voltage Regulation (2010 to 2018)**
>
> For years, United States interconnection standards largely kept inverters from supporting voltage regulation and required a distributed system to disconnect at set levels of grid disturbance instead [17]. A study first posted in 2010 noted that cloud-driven swings of photovoltaic output could outpace slowly responding utility equipment such as tap changers and switched capacitors, and that fast inverters able to inject or absorb reactive power, although not permitted under the standards of the day, could hold the voltage [18]. In May 2014 an amendment, IEEE 1547a, allowed voltage regulation support by distributed resources [17]. IEEE 1547-2018 then set reactive power categories for every resource it covers [6]: category A requires constant power factor, constant reactive power and Volt-VAR capability, and category B adds watt-VAR and Volt-Watt [19].

8.3 The Feeder as Modeled

The simulated plant is one generic 12.47 kV suburban feeder with synthetic teaching values, drawn in Figure 8.1(a) in the arrangement of the application's Feeder view. A substation transformer with an on-load tap changer feeds a bus, node 0, from which twelve overhead segments of 0.8 km reach twelve nodes. Each node stands for a block of 80 houses, and the share of houses with rooftop solar rises from 10 % at node 1, the older blocks, to 90 % at node 12, so that half the houses carry solar. A 600 kvar capacitor bank sits at node 6 and smart meters at nodes 3, 7 and 10. Every power is in per unit of $S_b = 10\,\text{MVA}$ and every voltage in per unit of 12.47 kV, so the base impedance is 15.55 Ω and the whole feeder is $0.2642 + j0.2485$ pu, an R/X of 1.06. Time is in seconds from the case's start, which a hidden parameter places at its time of day. Table 8.1 classifies the variables.

8.3.1 Branch Flows and Voltages

The flow entering segment k is the sum of the injections it feeds, and the voltage falls along each segment by its resistance times the active flow plus its reactance times the reactive flow, as shown in Equation (8.1).

$$\begin{aligned} P_k &= \sum_{j \geq k} p_j, \qquad Q_k = \sum_{j \geq k} q_j \\ V_0 &= V_{grid}(1 + 0.00625n) - r_t P_h - x_t Q_h \\ V_k &= V_{k-1} - (r_k P_k + x_k Q_k), \quad k = 1 \ldots 12 \end{aligned} \tag{8.1}$$

where p_j and q_j are the net active and reactive loads at node j, P_h and Q_h are the flows at the feeder's head including its losses, and the nominal voltage that divides the drop is 1 pu. This is the DistFlow

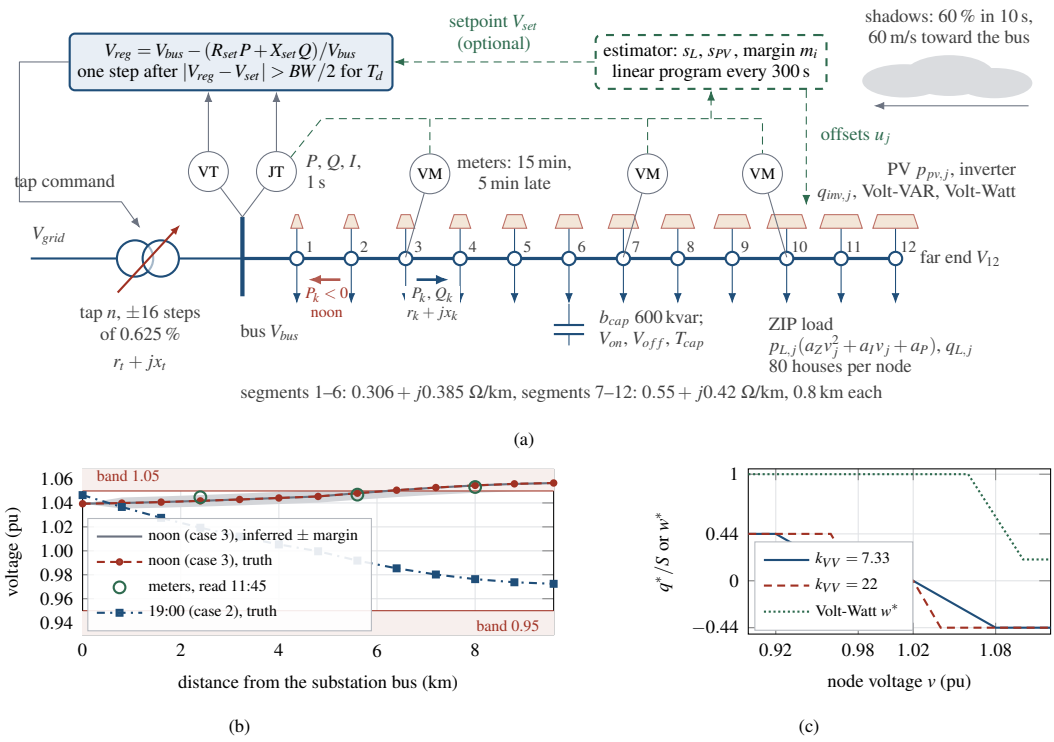

Figure 8.1 The feeder as modeled. (a) One-line diagram, not to scale: the tap n sets the source behind the transformer impedance $r_t + jx_t$, the voltage transformer (VT) and power transducer (JT) give the tap changer's control V_{bus}, P, Q and I, and the control steps the tap once V_{reg} has stayed outside $V_{set} \pm BW/2$ for T_d. Segment k carries P_k and Q_k through $r_k + jx_k$, each node j has a ZIP load and rooftop solar whose panels grow toward the far end, the capacitor b_{cap} sits at node 6 and the smart meters (VM) at nodes 3, 7 and 10. Dashed green: the optional estimator and linear program. (b) Voltage along the feeder at 12:00 in case 3 (reverse flow: the inferred profile, solid, inside its gray margin, the truth dashed with dots and the meters' last readings as circles) and at 19:00 in case 2 (the evening peak, no solar, dash-dotted with squares), against the 0.95 pu to 1.05 pu teaching band. (c) The Volt-VAR curve at the category B default slope (solid) and at the steepest the standard allows (dashed), and the Volt-Watt limit (dotted).

branch model [11] with the loss terms dropped from the voltage equation, the form usually called LinDistFlow. The losses still enter the power balance, so the substation supplies them, as shown in Equation (8.2).

$$P_h = P_1 + \sum_k r_k \frac{P_k^2 + Q_k^2}{V_{k-1}^2}, \quad Q_h = Q_1 + \sum_k x_k \frac{P_k^2 + Q_k^2}{V_{k-1}^2}$$
$$P = P_h + r_t \frac{P_h^2 + Q_h^2}{V_0^2}, \quad Q = Q_h + x_t \frac{P_h^2 + Q_h^2}{V_0^2} \tag{8.2}$$

where P and Q enter the transformer from the grid. Segments 1 to 6 are $0.306 + j0.385\,\Omega$/km and segments 7 to 12 are $0.55 + j0.42\,\Omega$/km, a heavier trunk and a lighter outer conductor, and the transformer is $r_t = 0.006$ pu, $x_t = 0.08$ pu. Against the full nonlinear DistFlow, solved by a backward and forward sweep that an independent complex phasor flow reproduces to 10^{-14}, the linearized far end reads high by 0.11 % of nominal at clear noon with 75 % solar, by 0.46 % with 125 % and by 0.16 % at the evening peak, because the linear form leaves out the drop that the losses cause.

8.3.2 Loads, Sun and Clouds

Each node's injection depends on its own voltage through a ZIP load and on the capacitor, as shown in Equation (8.3).

$$\begin{aligned} p_j &= p_{L,j}\, z(v_j) - p_{pv,j} \\ q_j &= q_{L,j}\, z(v_j) - q_j^{inv} - b_{cap,j} v_j^2 \\ z(v) &= a_Z v^2 + a_I v + a_P \\ p_{L,j} &= \tfrac{1}{12} P_{pk}\, \lambda(t)\, \big(1 + \omega_j(t)\big), \quad q_{L,j} = p_{L,j} \tan\varphi \end{aligned} \tag{8.3}$$

where $(a_Z, a_I, a_P) = (0.3, 0.2, 0.5)$ are the constant-impedance, constant-current and constant-power shares, $P_{pk} = 4.416\,\text{MW}$ is the evening peak of 960 houses at 4.6 kW each, $\lambda(t)$ is a declared hot-day shape of 24 hourly values with its peak of 1.00 at 19:00 and 0.40 at midday, cosine-smoothed between hours, ω_j is a slow random wander with a 3 % standard deviation and a 15 min correlation time, and $\tan\varphi = 0.329$ for a power factor of 0.95. Because the injections depend on the voltages, the flow is iterated three fixed times from a flat start. Three passes agree with thirty to 3.9×10^{-6} pu at the three check points, and the fixed count keeps the work per call constant, so the derivative function stays pure and runs replay bit for bit.

The solar a node's inverters can deliver is its rating times a clear-sky curve times the share of sunlight that reaches it, as shown in Equation (8.4).

$$p_{av,j} = S_j\, c_s(t)\, \kappa_j(t), \qquad c_s = 0.92 \sin^{1.3}\!\Big(\pi \frac{t_d - t_r}{t_s - t_r}\Big) \tag{8.4}$$

where the ratings S_j add to $\pi_{pv} P_{pk}$ with a penetration $\pi_{pv} = 75\,\%$ (3.31 MW, about 6.9 kW on each solar house) and are shared in proportion to $0.1 + 0.8(j-1)/11$, t_d is the time of day, sunrise t_r is 06:00, sunset t_s is 20:00 and solar noon 13:00. The sunlight share κ_j is 1 under a clear sky. A cloud field of strength c draws, from the case's seed, shadows between 08:00 and 18:00 that last 1 to 10 min and remove 55 to 80 % of the sunlight at $c = 1$ (up to 90 % above it). Each edge takes 10 s to cross a roof, so a 60 % shadow cuts a roof's output by 60 % in 10 s, and the shadows sweep from the far end toward the bus at 60 m/s, reaching node j $(12 - j)\,0.8\,\text{km}/(60\,\text{m/s})$ after node 12.

8.3.3 Inverters

Each node's inverters are one aggregate that follows its reactive power target through a first-order lag, and a second lag carries its Volt-Watt limit, as shown in Equation (8.5).

$$\begin{aligned} \tau \frac{\mathrm{d}q_j^{inv}}{\mathrm{d}t} &= \operatorname{clip}\big(S_j \phi(v_j) + u_j,\ \pm \bar{q}_j\big) - q_j^{inv} \\ \phi(v) &= -\operatorname{sgn}(v-1) \min\big(0.44,\ k_{VV}(|v-1| - \delta)^+\big) \\ \tau_w \frac{\mathrm{d}w_j}{\mathrm{d}t} &= w^*(v_j) - w_j, \qquad p_{pv,j} = \min(p_{av,j},\ w_j S_j) \end{aligned} \tag{8.5}$$

where $\tau = T_{ol}/\ln 10$ makes the open-loop response time $T_{ol} = 5\,\text{s}$ the time to 90 %, $\delta = 0.02\,\text{pu}$ is the deadband, $k_{VV} = 0.44/0.06 = 7.33$ puts the 44 % limit at 0.92 pu and 1.08 pu, u_j is the central program's offset (zero without it) and $\bar{q}_j = 0.44 S_j \min(1, p_{av,j}/0.05 S_j)$ takes the capability away at night. These are the category B defaults of IEEE 1547-2018, read through a paper that states the standard's table and its ranges, which allow the saturation as close as 0.02 pu beyond the deadband, a slope of 22 [16], as plotted in Figure 8.1(c). The Volt-Watt limit w^* falls linearly from 1 at 1.06 pu to 0.2 at 1.10 pu with $\tau_w = 10\,\text{s}/\ln 10$, a teaching value after the category B form that no source available here confirmed. In the convention of Section 1.3 the curve acts on the reactive power it injects as a reverse-acting proportional controller with a deadband: a node voltage above the deadband makes the inverter absorb.

Table 8.2 Model and Controller Parameters. A: Adjustable, F: Fixed, H: Hidden

Symbol	Meaning	Value	Unit	
S_b, V_b	base power, voltage	10, 12.47	MVA, kV	F
$r_k + jx_k$	segments 1–6, 7–12 (0.8 km)	see text	Ω/km	F
x_t, r_t	transformer	0.08, 0.006	pu	A, F
p_{pk}, pf	peak per house, power factor	4.6, 0.95	kW, —	A
a_Z, a_I	ZIP shares (a_P the rest)	0.3, 0.2	—	A
π_{pv}	solar rating / evening peak	75	%	A
c, v_c	cloud field, shadow speed	0, 60	—, m/s	A
n	tap, step	±16, 0.625	–, %	F
V_{set}, BW, T_d	setpoint, band, delay	1.025, 1.67, 45	pu, %, s	A
R_{set}, X_{set}	compensator (Forward)	0.01, 0.15	pu	A
δ, k_{VV}	deadband, slope (2.75–22)	0.02, 7.33	pu, pu/pu	A
T_{ol}	inverter response time (3–90)	5	s	A
w^*	Volt-Watt 1.06 to 1.10 pu	1 to 0.2	—	A
V_{on}, V_{off}, T_{cap}	capacitor rule (when on)	0.98, 1.035, 60	pu, s	A
—	meter period, delay	900, 300	s	F
m	program's added margin	0.005	pu	A
w_{tap}, w_q, w_s	program weights	3, 1, 300	—	F
T_s, N	MPC period, horizon	300, 24	s, steps	F
L_b	MPC blocks	1,1,1,1,2,2,4,6,6	steps	F
—	MPC node, band-row caps	100, 72	—	F
—	MPC preview	On	—	A
t_0, n_0	start time, starting tap	case	s, —	H

8.3.4 Tap Changer, Capacitor and Ledgers

The tap changer and the capacitor are discrete, and their steps are applied once per step after integration, with a historian mark. One tap step is 0.625 % of the source voltage, ±16 steps give ±10 %, and the capacitor adds $b_{cap} = 0.06$ pu of susceptance at node 6 with a 5 min discharge lockout after it opens. Five ledgers integrate the substation's energy, the losses, the loads' energy, the available solar and the delivered solar, so the balance $E_{sub} = E_{load} - E_{pv} + E_{loss}$ closes by construction. The 29 states (twelve reactive powers, twelve Volt-Watt limits and five ledgers) advance by fourth-order Runge-Kutta steps of 1 s with the power flow inside each stage and the weather and loads held at the step's midpoint, the local controllers run every second, the central program every 300 s, and a row of 93 columns is logged every 5 s, on the engine of Section 1.6.1.

8.3.5 Measurements and Initialization

Controllers see only a measurement packet. The substation publishes V_{bus} with 0.05 % noise and P and Q with 0.5 %, each a function of the seed and the whole second, so that halving the step hands the controllers the same readings. The capacitor has its own voltage sensor. The three meters record their node voltages every 15 min with 0.1 % noise, report 5 min later through the framework's analyzer, and hold the last reading through an outage, marked stale. A late sample of this kind limits the composition loop of Volume 5, Chapter 3 as well. Every other node voltage, the clouds, the loads and the inverters' outputs are teaching truth, logged under a `truth_` prefix, and a source test rejects any read of them by the controller, estimator or program. A run starts at the tap that puts the regulated voltage nearest its setpoint with the inverters at unity power factor, unless the case pins one, and with the inverters relaxed onto their curves. Table 8.2 lists the parameters.

8.3.6 Assumptions

The feeder is a balanced single-phase equivalent, so unbalance, single-phase laterals, three-phase line regulators and the service transformers are out of scope, and a node voltage stands for its houses' service voltage. LinDistFlow leaves out the losses' effect on the voltages, an error of 0.11 to 0.46 % of nominal at the check points, which makes it a teaching approximation rather than a planning tool. The loads follow one declared day, the ZIP shares are fixed, and thermostatic recovery after a voltage

change is absent. The inverters are one aggregate per node with no ride-through or trip settings, and the tap changer has no mechanism time or intertap delay. The regulator, estimator and program read P and Q entering the transformer, which include its own reactive losses, about 0.003 pu of correction at the evening peak with $X_{set} = 0.15$ pu. The band 0.95 pu to 1.05 pu is the ANSI C84.1 Range A service band used as a teaching value, the Volt-VAR defaults are IEEE 1547-2018 category B as teaching values, and nothing is a utility's or a manufacturer's setting or data.

8.4 Four Mechanisms Behind the Cases

8.4.1 A Tap Moves Everyone

A tap step scales the source, and to first order every node moves with the bus, because the drop along each segment depends on the flows, which one step barely changes. In case 1 at 11:30 one step down takes 0.0062 pu off the bus and 0.0060 pu off the far end. The profile's shape, a rise of 0.017 pu from the bus to node 12 at noon in case 3 and a fall of 0.074 pu at the evening peak in case 2, belongs to the flows, as shown in Figure 8.1(b). One setting of the tap therefore has to keep the far end under 1.05 pu at noon and over 0.95 pu at 19:00, and only a device at the nodes can change the spread.

8.4.2 The Compensator's Sign

The regulator controls the compensated voltage of Equation (8.6) rather than the bus.

$$V_{reg} = V_{bus} - c, \qquad c = \frac{R_{set}P + X_{set}Q}{V_{bus}} \tag{8.6}$$

The correction c is the in-phase part of the phasor drop $(R_{set} + jX_{set})I$ that a control computes from its current transformer. With $P < 0$ and $Q > 0$ it is positive, and holds the bus up, exactly when $X_{set}/R_{set} > -P/Q$. At 12:00 in case 3, $P = -1.081$ MW and $Q = 0.645$ Mvar put the threshold at 1.68, so the shipped X/R of 15 raises the bus, by 0.008 27 pu in the application's form against 0.008 14 pu from the phasor magnitude $|V - (R_{set} + jX_{set})I|$, while a compensator matched to the line, $R = X = 0.06$ pu, lowers it by 0.0025 pu. A design note for the simulator expected the opposite failure, a compensator that lowers the bus when it should raise it. On one feeder with solar spread along its length the far end is the high point, so that error needs generation concentrated beyond the loads, and the application teaches the sign rule instead. On forward flow at a steady power factor the correction reads only $R_{set} + X_{set}\tan\varphi$, so evening data cannot tell the shipped setting from a line-matched one. Case 2 makes 7 operations with $R = 0.01$, $X = 0.15$ and 5 with $R = X = 0.045$, whose largest corrections are 0.031 pu and 0.028 pu. The shipped dials equal 11.7, 15.6 or 23.3 V of reactance on a 120 V base for a 300, 400 or 600 A current transformer, inside a regulator's range, but the textbook procedure sets both dials in proportion to the impedance to the regulated point, which gives $R \approx X$ on this line. The shipped setting is a teaching choice.

🕮 **Pause.** Try Exercise 8.1 before reading on.

8.4.3 Steep Curves Against the Compensator

The inverters alone are stable at any slope. Linearizing Equation (8.5) with the tap fixed, the voltages respond to the reactive powers through the shared path reactance X_{ij} of the LinDistFlow model, $\partial v/\partial q = X$, and the Jacobian of the lags is $-(I + KX)/\tau$ with K the diagonal of the slopes times the ratings. KX is similar to $K^{1/2}XK^{1/2}$, which is positive semidefinite, so every eigenvalue lies at or left of $-1/\tau$. The stability condition on the norm of the sampled update that the Volt-VAR literature derives [16] applies to an inverter that jumps to its curve once per sample, which the application does not model.

The hunting comes from the coupling to the tap changer. The inverters' reactive power returns through the substation, where the compensator reads X_{set} times its change as a change of load. A common change Δv of the node voltages inside the curves' slope moves the correction by about $X_{set}k_{VV}\sum_j S_j\,\Delta v$, a factor of 0.36 at the default slope and 1.09 at 22 for $X_{set} = 0.15$ pu at 75 % solar.

On case 5's cloudy day the range of the correction over each 10 min has a median of 0.0031 pu and a 90th percentile of 0.0050 pu at slope 7.33, against 0.0082 pu and 0.0142 pu at 22, close to the 0.0167 pu band. At the steep slope the inverters swing their whole range with the cloud edges, with a median 10 min range of 0.664 Mvar against 0.302 Mvar and a lowest total of −1.41 Mvar against −0.50 Mvar, and all 8 steps down come with a shadow over the far end and all 10 steps up in sun, the wrong way each time.

8.4.4 Conservation Voltage Reduction

A ZIP load's power changes with its voltage at the rate $\partial \ln p / \partial \ln v = 2a_Z + a_I$, 0.8 for the declared shares, so lowering the voltage by 1 % lowers a constant-impedance and constant-current load by about 0.8 %. A constant-power load draws more current instead, and the losses rise. The feeder's measured factor is the energy saved over the load-weighted mean voltage reduction.

8.5 The Controllers as Implemented

No PID loop runs in this application, so the velocity form of Section 1.6.2 does not appear. Every controller is written explicitly in its own module and executes on the measurement packet every second, except the estimator, the program and the optional predictive controller, which execute every 300 s.

8.5.1 Tap Changer

At each second k the tap changer computes $V_{reg,k}$ from Equation (8.6) and the error $e_k = V_{reg,k} - V_{set}$, and runs the timer law of Equation (8.7).

$$
\begin{aligned}
d_k &= \begin{cases} -\operatorname{sgn}(e_k) & |e_k| > BW/2 \\ 0 & \text{otherwise} \end{cases} \\
\theta_k &= \begin{cases} \theta_{k-1} + \Delta t & d_k = d_{k-1} \neq 0 \\ \Delta t & d_k \neq d_{k-1},\ d_k \neq 0 \\ 0 & d_k = 0 \end{cases} \\
n_k &= \operatorname{clip}(n_{k-1} + d_k,\ \pm 16),\ \theta_k \leftarrow 0 \ \text{ if } \theta_k \geq T_d
\end{aligned}
\tag{8.7}
$$

where θ is the timer, which counts while V_{reg} stays outside the band on one side, restarts on a side change and clears inside the band, and the tap otherwise holds. In the convention of Section 1.3 the tap changer is reverse acting: a regulated voltage below the band steps the tap up. The compensator is Off, Forward (always on) or Reverse-aware, which drops the correction while the measured P is negative, the cogeneration-style mode of regulator controls. In Manual the tap is the operator's station, rounded, and the timer is held at zero. The application reports the tap the regulator asked for as its manual-follow value.

8.5.2 Capacitor

The voltage rule switches the bank in when its own sensor reads below V_{on} for T_{cap} and out above V_{off} for T_{cap}. A schedule keeps it in from 13:00 to 21:00. A closing command waits out the 5 min lockout, and an out-of-service event keeps it open.

8.5.3 Estimator

Every 300 s the estimator fits a load scale s_L and a solar scale s_{PV} on the declared shapes to the substation's 1 min averages now and to the meters' last readings, each compared with the model at the time the meter recorded it, as shown in Equation (8.8).

$$\min_{s_L,s_{PV}} \left(\frac{\hat{P}-\bar{P}}{\sigma_P}\right)^2 + \left(\frac{\hat{Q}-\bar{Q}}{\sigma_Q}\right)^2 + \sum_{m\in\{3,7,10\}} \left(\frac{\hat{V}_m(t_m)-V_m}{\sigma_m}\right)^2 \tag{8.8}$$

where $\sigma_P = \sigma_Q = 0.005\,\text{pu}$, $\sigma_m = 0.002 + 0.1|\bar{P}_{now} - \bar{P}(t_m)|$ discounts a meter read before a change at the substation, and four Gauss-Newton iterations with a finite-difference Jacobian solve it with s_L held in 0.3 to 2.5 and s_{PV} in 0 to 1.3. The model is the utility's: the declared impedances, shapes, ratings, curves and ZIP shares, with neither the clouds nor the loads' wander nor the inverters' outputs. The profile $\hat{V}$ is then drawn from the fit, anchored at the measured bus, and each node carries the margin of Equation (8.9).

$$m_i = 2\max(0.002,\ \rho) + 0.1|\hat{V}_i - V_{bus}| + \left|\frac{\partial \hat{V}_w}{\partial s_{PV}}\right| \frac{\Delta \bar{P}_{10}}{\sum_j S_j c_s} + 0.01\,\chi_{st} \tag{8.9}$$

where ρ is the root mean square of the meter residuals, $\Delta\bar{P}_{10}$ is the range of the substation's 1 min power over the last 10 min, the clouds the fit cannot place, and $\hat{V}_w$ is the inferred worst node, the one nearest or furthest past its band edge, and χ_{st} is 1 while the meters are stale and 0 otherwise.

8.5.4 Central Program

With the central controller on, every 300 s a linear program on the framework's bounded two-phase simplex chooses a bus move $d^+ - d^-$ and an offset $u_j^+ - u_j^-$ per node, as shown in Equation (8.10).

$$\begin{aligned} \min\ & \frac{w_{tap}}{0.00625}(d^+ + d^-) + \sum_j (w_q \mp \lambda_j) u_j^{\pm} \\ & + w_s \sum_i (s_i^+ + s_i^-) \\ \text{s.t. } & V_i = \hat{V}_i^0 + d^+ - d^- + \textstyle\sum_j X_{ij}(u_j^+ - u_j^-) \\ & 0.95 + m_i' - s_i^- \le V_i \le 1.05 - m_i' + s_i^+ \\ & 0 \le d^{\pm} \le 0.06, \quad 0 \le u_j^{\pm} \le \bar{q}_j \end{aligned} \tag{8.10}$$

where $\hat{V}_i^0 = \hat{V}_i - \sum_j X_{ij}\bar{u}_j$ is the inferred profile without the offsets $\bar{u}_j$ in force, which the new ones replace, $m_i' = m_i + m + \Delta n/2$, $\lambda_j = 2\sum_{k\le j} r_k \hat{Q}_k$ is the loss gradient, so offsets that inject where reactive power flows forward pay less, Δn is one tap step, the capabilities $\bar{q}_j$ are taken at the lower of the clear-sky output now and in 5 min and are zero while the curves are Off, when the inverters run at unity power factor and apply no offset, and conservation voltage reduction adds a cost of 600 per pu on the bus voltage. The bus move is rounded to whole steps and sent as a target tap, which the tap changer reaches one step per T_d, together with the setpoint that target implies. When the band slack the program needs exceeds 0.02 pu summed over the nodes, a sign that the margins fill the band, it holds its last target and sends no offsets. Without the central controller, conservation voltage reduction sets the setpoint every 300 s to put the inferred lowest node at 0.95 pu plus its margin, the added margin and half the band.

8.5.5 Mixed-Integer Predictive Controller

As a third setting of the central controller, a model predictive controller (MPC) takes the program's call. It decides on the same 300 s clock from the same estimate and drives the same tap target, setpoint and offsets, but it plans the tap as an integer over the next 2 h instead of rounding one bus move afterwards, and it sees the declared load shape and clear-sky curve coming. It is written for this application rather than taken from the framework's option of Section 1.6.3, because its main

decision is an integer tap. Receding-horizon control under constraints has a settled theory [20], a plant that mixes logical decisions with continuous dynamics turns the predictive problem into a mixed-integer program [21], and centralized predictive schemes that correct a multi-step plan with real-time measurements have been proposed for the voltages of distribution networks that host many generators [22].

The controller's model is the estimator's fitted feeder, so nothing is identified. At a decision at time t with tap n_0, the source behind the transformer is estimated as $\hat{v}_{src} = V_{bus} + r_t P + x_t Q$ and corrected twice so that the model's bus equals the measured bus with the offsets in force, and $\hat{v}_{grid} = \hat{v}_{src}/(1+0.00625 n_0)$. The base profile is evaluated at the 24 five-minute steps of the horizon with no offsets and linearized in the tap and the offsets, as shown in Equation (8.11).

$$\begin{aligned} B_{i,k} &= \mathscr{V}_i\big(s_L, s_{PV}, t+300k, \hat{v}_{src}, 0\big), \quad k = 0 \ldots 23 \\ V_{i,k} &\approx B_{i,k} + S_i(n_k - n_0) + \textstyle\sum_j G_{ij} u_{j,k} \end{aligned} \tag{8.11}$$

where $\mathscr{V}_i(s_L, s_{PV}, \tau, v, u)$ is the estimator's LinDistFlow profile at node i for the declared shapes at time of day τ, source voltage v and offsets u, with the inverters on their declared curves through four fixed-point passes, S_i is the change at node i for one tap step and $G_{ij} = \partial V_i/\partial u_j$, both taken once per decision at $\tau = t$ by finite differences (an offset of 0.002 pu) through the same profile. With the preview Off the shapes are held at their values now, $B_{i,k} = B_{i,0}$. The preview is only the declared day, so its error is every cloud and the loads' wander, which only the next fit sees. With no load and no sun the computed S_i equals $0.00625 \hat{v}_{grid}$ to a relative 2.3×10^{-14}. At noon in case 3 with the curves Off, $G_{ij} - X_{ij}$ lies between 0.074 and 0.080 against $x_t = 0.08$ pu, because this source sits behind the transformer while the program's X holds the bus. A tap step moves the far end 0.0039 pu with the curves On against 0.0061 pu with them Off, since the curves answer inside the profile. The model leaves out the clouds, the wander, the curves' saturation in G, the band acting between decisions, the capacitor's rule inside the horizon and the change of S and G over 2 h. It is offset-free in the output-bias sense [23], re-anchored on the measured bus at every decision while the estimator refits its two scales at the same call, and no Kalman filter is used.

The horizon is held in nine blocks of L_b = 1, 1, 1, 1, 2, 2, 4, 6 and 6 steps. In block b the tap moves by m_b at the block's start, an integer in the first four blocks (the first 20 min) with $|m_0| \le 1$, and a continuous ramp of at most one step per 5 min, $|m_b| \le L_b$, afterwards. The offsets $u_{j,g} = u^+_{j,g} - u^-_{j,g}$ are held in three groups $g(b)$: the first block, the next three and the rest. The program is Equation (8.12).

$$\begin{aligned} \min \sum_b \Big[& w_{tap} a_b + L_b \sum_j (w_q \mp \lambda_j) u^{\pm}_{j,g(b)} \\ & + L_b w_s \sum_i (s^+_{i,b} + s^-_{i,b}) \Big] \\ \text{s.t. } & a_b \ge |m_b|, \quad N_b = \textstyle\sum_{c \le b} m_c, \quad |n_0 + N_b| \le 16 \\ & \max_{k \in b} B_{i,k} + S_i N_b + \textstyle\sum_j G_{ij} u_{j,g(b)} - s^+_{i,b} \le 1.05 - m''_i \\ & \min_{k \in b} B_{i,k} + S_i N_b + \textstyle\sum_j G_{ij} u_{j,g(b)} + s^-_{i,b} \ge 0.95 + m''_i \\ & 0 \le u^{\pm}_{j,g} \le \bar{q}_{j,g}, \quad s^{\pm}_{i,b} \ge 0 \end{aligned} \tag{8.12}$$

where a_b counts the block's tap operations, $m''_i = m_i + m$ is the estimator's margin plus the added margin without the program's half step, since the tap is already whole, $\bar{q}_{j,g}$ is the capability of Equation (8.10) at the lowest clear-sky output of the group's steps, zero while the curves are Off, and the weights are the program's, $w_{tap} = 3$ per operation, $w_q = 1$ and $w_s = 300$ per pu and step, with conservation voltage reduction adding $600 \sum_b L_b S_0 N_b$. A block's decision is constant, so a node stays inside the band at every step of the block exactly when the block's highest and lowest base voltages do, and the band rows judge all 24 steps. One slack per node, side and block prices a violation per 5 min, as the program prices its own.

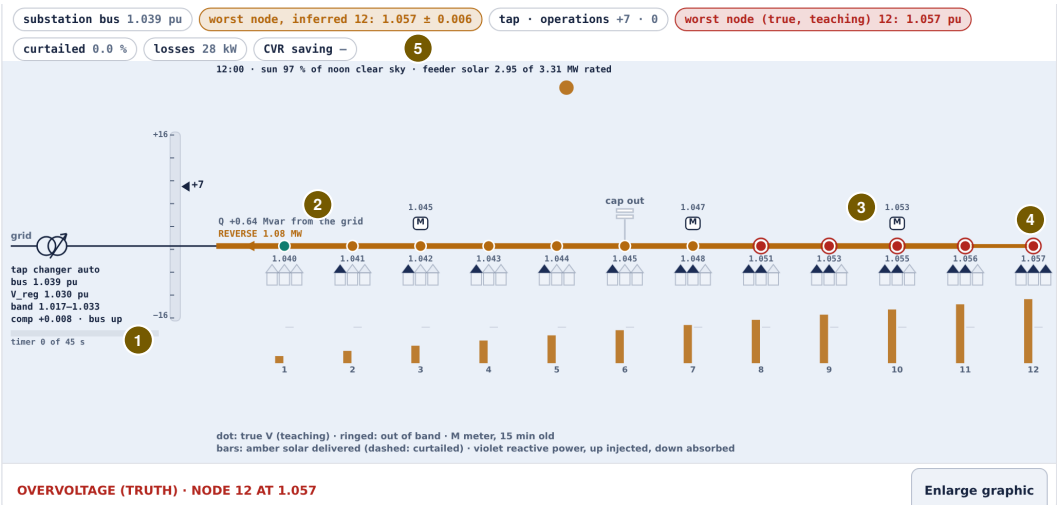

Figure 8.2 Reverse flow at solar noon on the Feeder view, case 3 at 12:00 as shipped, cropped to the plant view. The substation exports 1.08 MW while 0.64 Mvar still flows in (2), and the compensator ($R_{set} = 0.01$, $X_{set} = 0.15$ pu) holds the bus up by 0.008 pu, so the tap stays at +7 with $V_{reg} = 1.030$ pu inside its band of 1.017 to 1.033 pu and the bus at 1.039 pu (1). Nodes 8 to 12 are ringed over 1.05 pu, the far end at 1.057 pu (4), which the estimator infers as 1.057 ± 0.006 pu (5) from meter readings 15 min old (3, one of the three meters). Solar delivers 2.95 MW of the 3.31 MW rated, the losses are 28 kW, and the tap has made no operation since 09:00.

Depth-first branch and bound on the framework's bounded simplex solves the program from the incumbent with no move and no offset. The pumps of Volume 9, Chapter 11 are scheduled as integers by the same kind of search. Row generation starts from the three nodes nearest each limit per block and side and adds every row the plan crosses, for six rounds at most, and a final check verifies all 216 band rows. The first program has 144 variables, 81 rows and 4 integers. With all nine moves integer it cost 1.1 s a decision on average, which set the integer window at 20 min. The controller sends the tap target $n_0 + m_0$, the bus voltage the plan implies as the setpoint and the first offsets, and the band and timer still act between decisions.

The MPC first runs the program's hold test on the same estimate and, when the margins fill the band, keeps the last target and sends no offsets, as the program does. A plan is accepted only if it is optimal with a zero gap, finite, inside every box, integral, consistent and within every band row, including those left out. A plan that needs more than 100 branch-and-bound nodes or 72 band rows, or fails a check, hands the period to the program, marked in the historian, until the next valid plan. Replayed on the measurement packets of a run over a different hidden plant, the program and the MPC made 0 different calls in 46,800. In Chrome, under another simulation's load, a solve took 1.7 to 4.8 ms on average over six cases and at most 208 ms, at most 0.0015 % of a core at 1× and 5.2 % at 3600×, where a day plays in 24.1 s with 2 of 1,369 frames over 50 ms. A decision falls once per 300 simulated seconds, so no frame guard is used.

8.5.6 Three Models

The plant is Section 8.3, with the true ZIP shares, clouds and wander. The estimator's model is the utility's declared feeder without them. The program's model is linear, the bus move and the shared reactance X_{ij} with the bus held, and it omits the inverters' own curves, whose response the next call corrects. The MPC's model is the estimator's declared feeder again, evaluated over 2 h and linearized once per decision with the curves answering inside it. None of the local devices carries a model.

Table 8.3 The Cases as Shipped

#	Setting and scripted events	Teaching point
1	Manual from 10:00, tap +6, clear	a tap moves every node
2	12:00–24:00, no solar	band, timer, operations
3	09:00–15:00, clear, Forward 0.01 + *j*0.15	reverse flow, the compensator
4	12:00, cloud bank at 12:15 for 25 min	the cloud the tap misses
5	09:00–17:00, curves on, cloud field 1.5	the slope that hunts
6	16:00–22:00, band 1.25 %, capacitor rule	two timers
7	16:00–22:00, no solar	conservation voltage reduction
8	10:00, curves on, meter outage at 12:00	the unmeasured customer
9	24 h, 150 % solar, field 1.5, suggested local settings, MPC	the MPC against the program
C	Challenge: 24 h, 125 % solar, field 1.5	coordination

Table 8.4 Measured Results. Voltages in Per Unit

#	Setting	Quantity	Value
1	tap +6 held, +5	far-end peak	1.0516, 1.0455
2	band 1.67, 2.5, 1.25 %	operations	7, 5, 8
	band 1.67, 2.5 %	lowest node	0.9724, 0.9698
3	shipped, at 12:00	correction, far end	0.0083, 1.0567
	shipped, Off at 11:00, +center	highest	1.0575, 1.0514, 1.0453
4	curves off, on	far-end fall in 30 s	0.0144, 0.0105
5	slope 7.33, 22	operations to 17:00	2, 18
	slope 22, seeds 1–3	operations	10, 13, 18
6	shipped, tap delay 20 s	tap, capacitor ops	14 + 12, 5 + 2
7	baseline, CVR, CVR + central	energy (kWh)	24 053, 23 863, 23 720
	constant power, CVR	energy change	+16.9 kWh
8	before, after outage	fits bracketing	23/24, 24/24
9	program, MPC	min out (operations)	21.2 (17), 0 (15)

8.6 Nine Cases and a Challenge

The cases appear in Table 8.3 and the headline measurements in Table 8.4. The numbers in this section were taken headlessly from the same engine code the page loads, seeded with 20261026 and started from the case's own settings (Section 1.6.4), and a change a sentence names is made at the time it states. Clock times are rounded to the nearest minute.

8.6.1 The Tap by Hand

Case 1 starts at 10:00 in Manual with the tap at +6. The far end passes 1.05 pu at 11:55 and, held at +6, peaks at 1.0516 pu at 12:20 after 77 min over the band, while at +5 it peaks at 1.0455 pu. Lowering the tap one step at 11:30 moves the bus from 1.0331 pu to 1.0269 pu and the far end from 1.0486 pu to 1.0426 pu.

8.6.2 Band and Timer

Case 2 runs the afternoon and evening without solar. As shipped, with a 1.67 % band (2 V on 120 V) and a 45 s delay, the tap changer makes 7 operations from 12:00 to 24:00 with the lowest node at 0.9724 pu. Widening the band to 2.5 % cuts the count to 5 and lets the lowest node fall to 0.9698 pu, a 1.25 % band makes 8, and a 120 s delay makes 5. A band must exceed one step, or a step taken at one edge would cross the other. At 1.0 %, 1.6 steps, the regulator makes 8 operations without hunting. At 19:00 the tap stands at +10, the bus at 1.0465 pu and the far end at 0.9724 pu, as shown in Figure 8.1(b).

8.6.3 Reverse Flow and the Compensator

Case 3 runs a clear day from 09:00. Power flows back from 09:34, the far end passes 1.05 pu at 10:44, and it peaks at 1.0575 pu at 12:20 with 203 min over the band by 15:00. The state at 12:00 appears in Figure 8.2. Switching the compensator Off at 11:00 brings one step down to +6 by 11:15, after

which the far end peaks at 1.0514 pu because the bus is then inside its band. Pressing the application's Center the taps helper after switching it Off takes the tap to +5, and the far end peaks at 1.0453 pu. A Reverse-aware compensator would also stop the correction, but it switches with every change of sign and is no cure on a cloudy day, as case 5 measures.

> **TRY IT IN THE APP** **3 · Solar Noon: Reverse Flow and the Compensator**
>
> Load scenario 3 and press Run. Watch the far end of the feeder on the Feeder view. Power flows back to the substation from 09:34, and the far end passes 1.05 pu at 10:44. The tap changer's regulated voltage stays inside its band, and by 12:00 no tap has moved while nodes 8 to 12 are ringed over the band.

8.6.4 The Cloud

Case 4 adds a cloud bank at 12:15 that takes 60 % of the sunlight in 10 s, sweeps the feeder in 160 s and stays 25 min. Node 12's rooftops fall from 449 to 179 kW in 10 s. With the inverters at unity power factor, the far end falls from 1.0447 pu to 1.0303 pu in 30 s and to 1.0137 pu by 12:40, while the bus moves only from 1.0268 pu to 1.0258 pu. The regulated voltage stays between 1.0145 pu and 1.0284 pu against a band of 1.0166 to 1.0333 pu, never outside for 45 s, and the tap does not move. With the curves on, the far-end inverters had been absorbing 0.294 Mvar at the high noon voltage and fall to 0.154 Mvar within 30 s, so the far end starts lower, at 1.0369 pu, and falls 0.0105 pu instead of 0.0144 pu in the first 30 s.

Pause. Try Exercise 8.2 before reading on.

8.6.5 Hunting

Case 5 runs a heavily clouded day from 09:00 with the curves on and the evening compensator. At slope 7.33 the tap changer makes 2 operations by 17:00. At 22, set before Run, it makes 18, the score marks HUNTING (two reversals within 30 min) first at 11:54, and the tap steps between +9 and +10 with the clouds, as shown in Figure 8.3. On seeds 1 to 3 the counts are 10, 13 and 18 against 2 each, and halving the time step leaves 18 and 2 unchanged.

The hunting needs three settings at once, and the application states the window. Each of the following changes, made alone at slope 22, removes it: a cloud field of 1.0 instead of 1.5 (3 operations, while 1.2 gives 8), a 2.5 % band (1), a setpoint of 1.0125 pu (1), $R_{set} = 0.02$ pu (3), $X_{set} = 0.12$ pu (2), the compensator Off (0) and slopes of 18 or 16 (2 each), while slope 20 still hunts with 10. Settings in proportion to the line remove it too, $R = X = 0.06$ pu with 4 operations and the far end's uniform-load equivalent $0.124 + j0.132$ pu with 9 and no HUNTING mark, while a reactance-only $X_{set} = 0.15$ pu hunts harder, with 29. The steep slope, a compensator weighted on reactance (X/R 15 on a line whose X/R is 1.06) and a heavy cloud field are teaching choices that put the mechanism on the screen, not a claim that every steep curve hunts. The Reverse-aware mode makes 8 operations at 7.33, with HUNTING at 14:59, and 4 at 22, because it turns its compensation on and off with each shadow that reverses the flow.

> **TRY IT IN THE APP** **5 · The Slope: A Steep Curve That Hunts**
>
> Load scenario 5, set the Volt-VAR slope to 22 before pressing Run, and run to 17:00. The tap changer makes 18 operations instead of the 2 it makes at the default slope of 7.33. The status marks HUNTING first at 11:54, and the tap steps between +9 and +10 with the clouds.

8.6.6 Two Timers

Case 6 runs the evening with a 1.25 % band and the capacitor on a voltage rule of 1.008 pu and 1.016 pu, both with 45 s delays. The capacitor's own step, 1.16 % at node 6 and 0.48 % at the bus at 19:00, is wider than its 0.8 % rule, and with equal delays both timers start on the same change and run out together. The evening takes 14 tap operations and 12 capacitor switchings, against 4 tap

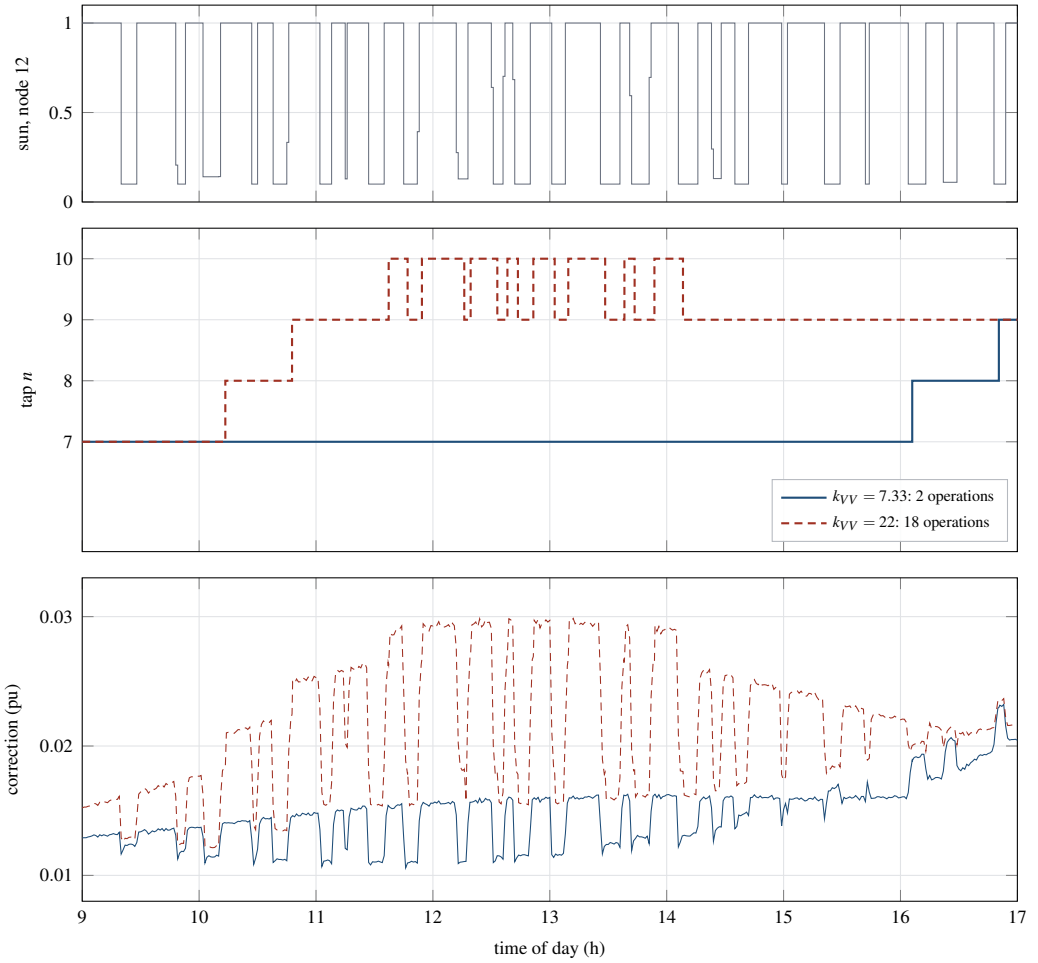

Figure 8.3 Case 5 from Reset on one seed, 09:00 to 17:00: the sunlight share at node 12 (teaching truth, the same at both slopes), the tap position and the compensator's correction c at the default Volt-VAR slope (solid) and at 22 (dashed). At 22 the correction swings close to the band's width with each shadow and the tap hunts.

operations with the capacitor out. A 20 s tap delay gives 5 and 2, a 120 s capacitor delay 5 and 3, a wider rule of 0.99 and 1.03 pu 4 and 0, and a schedule 6 and 2.

8.6.7 Conservation Voltage Reduction

Case 7 runs 16:00 to 22:00 without solar. The baseline delivers 24 053 kWh with the lowest node at 0.9731 pu. With conservation voltage reduction, which holds the inferred lowest node at the floor plus its margins, the feeder delivers 23 863 kWh, 190 kWh or 0.79 % less, with the lowest node at 0.9633 pu and 8 operations, and with the central program as well 23 720 kWh, 1.39 % less, at 0.9560 pu and 12 operations. The loads' energy falls by 192.9 kWh while the losses rise from 812.6 to 816.0 kWh. The load-weighted mean voltage falls 1.04 %, so the feeder's measured factor is 0.76, a little under the loads' 0.8, and the saving lies inside the national evaluation's 0.5 to 4 % [10]. With constant-power loads the same reduction delivers 16.9 kWh more than its own baseline as the losses rise from 818 to 835 kWh.

Table 8.5 The Challenge on Six Seeds (Default and 1–5)

Policy	pass	ops	out (min)	nodes, default seed
shipped: tap alone, Forward 0.01 + *j*0.15, 1.025, 1.67 %, 45 s	0/6	5	264.9–319.1	0.9754–1.0902
suggested: curves 7.33, Rev.-aware 0.06 + *j*0.06, 1.00, 2.5 %, 60 s	6/6	6–8	0	0.9522–1.0442
suggested, Forward compensator	6/6	8–10	0	0.9522–1.0419
suggested without the curves	0/6	6–8	102.0–160.4	0.9521–1.0597
suggested with 1.67 %, 45 s	0/6	23–29	0	0.9581–1.0442
suggested + central program	6/6	7–13	0	0.9557–1.0442
tap alone, Forward 0.09 + *j*0.09, 0.995, 4 %, 60 s	6/6	11–15	0	0.9579–1.0485

+ DID YOU KNOW? Voltage Reduction at National Scale

A 2010 national evaluation simulated conservation voltage reduction on 24 prototypical feeders and extrapolated the result. Deployed on every distribution feeder in the United States, it would cut annual energy consumption by 3.04 %, and deployed on the 40 % of feeders with the highest value, by 2.4 % [10].

8.6.8 The Unmeasured Customer

Case 8 runs a cloudy morning with the curves on and the meters failing at 12:00. Before the outage the true worst voltage lies inside the inferred worst node's margin at 23 of 24 fits, with a median margin of 0.0176 pu, although at 3 of those fits the true worst is a different node. The miss, at 11:50, is an error of 0.0057 pu against a margin of 0.0043 pu under a shadow's edge. After the outage all 24 fits hold the truth, with a median margin of 0.0191 pu, because the substation's P and Q alone fix the two scales and a stale meter adds 0.01 pu. The margin is conservative. Fit by fit it is a median of 44 times the error, whose own median is 0.0003 pu, and on seeds 1 to 3 every one of 48 fits holds the truth. The worst customer is inferred from measurements elsewhere on the plant, as the brick core temperature of Volume 5, Chapter 11 is.

Pause. Try Exercise 8.3 before reading on.

8.6.9 The Challenge

The challenge runs a day from midnight at 125 % solar with a cloud field of 1.5, graded on the truth: every node in 0.95 pu to 1.05 pu at every second, at most 20 tap operations and at most 2 % of the available solar curtailed, with the feeder, weather and meters as shipped. Table 8.5 lists the policies on six seeds from the application's benchmark file. The shipped tap changer alone spends 265 to 319 min out of band. Turning the curves on with the shipped compensator makes 67 operations and still leaves 222.9 min out, case 5's mechanism at the default slope on a sunnier feeder. The suggested settings coordinate locally, with the curves at the default slope, the compensator matched to the line, the setpoint at 1.00 pu and a slower, wider band, and pass on all six seeds with 6 to 8 operations.

A tap changer alone can just pass this day. Of about 1,300 tap-alone settings swept, 17 pass, every one with the compensator weighted like the line (X/R 1 to 2) and a slow or wide band, and the steadiest passes on all six seeds with 11 to 15 operations by letting its bus wander across a 4 % band. The infeasible variant is the same day at 150 % solar, the slider's top, where none of 1,536 tap-alone settings of setpoint, band, delay and compensator holds the band within 20 operations (the closest is 4.7 min out with 39 operations), while the curves at slope 11 with the suggested tap settings pass on all six seeds with 4 to 6. The clear-to-shadow spread of the profile, 0.075 pu at 125 % and 0.089 pu at 150 %, stays under the 0.10 pu band until about 170 %, so the tap changer fails there on operations and timing rather than on a static bound.

The central program also passes, with 7 to 13 operations, its five-minute estimate a step behind the clouds. From 09:00 to 17:00 its margins have a median of 0.054 pu and exceed half the band at 52 of 96 fits, up to 0.124 pu, so it holds at 56 of its 96 calls. As shipped, at 13:00, the inferred far end carries ±0.159 pu. The margins are conservative, a median of 6.4 times the error fit by fit, but too

Table 8.6 Central Structures on the Same Seed, Instruments and Noise: Minutes Any Node Is Out of Band · Tap Operations

Case	Local only	Program	Program, 5 min ahead	Program, 2 h ahead	Program, no half step	MPC	MPC, preview Off
3, reverse flow, 09:00 + 24 h	202.7 · 6	0 · 8	0 · 8	0 · 14	0 · 7	0 · 9	0 · 8
4, cloud bank, 12:00 + 24 h	7.5 · 9	0 · 19	0 · 19	0 · 18	0 · 15	0 · 15	0 · 13
7, CVR on, 6 h	0 · 8	0 · 12	0 · 12	0 · 6	0 · 11	0 · 13	0 · 13
energy (MWh)	23.86	23.72	23.73	23.80	23.66	23.71	23.71
8, meters out at 12:00, 24 h	0 · 13	0 · 31	0 · 25	0 · 28	0 · 37	0 · 21	0 · 19
challenge as shipped (curves Off)	264.9 · 5	30.1 · 21	30.1 · 21	10.2 · 21	65.2 · 19	11.1 · 21	30.1 · 21
challenge, suggested	0 · 6	0 · 11	0 · 13	0 · 9	0 · 11	0 · 11	0 · 11
challenge, suggested, clear sky	0 · 6	0 · 5	0 · 5	0 · 7	0 · 5	0 · 7	0 · 7
9, 150 % solar, default seed	46.2 · 6	21.2 · 17	50.5 · 19	30.0 · 11	17.2 · 13	0 · 15	15.9 · 11
9, six seeds summed	443.7 · 36	117.0 · 78	92.0 · 78	353.1 · 66	186.9 · 63	42.9 · 84	65.6 · 68

Local only: the case's tap changer, capacitor and curves with the central controller Off. The three program variants are measurement scripts, not settings: the program judged against the worse of now and the declared shapes 5 min or 2 h ahead, and the program without its half-step margin, the MPC's own rule. Curtailment is 0.00 % in every run. Periods the program decided after a rejected plan, MPC: 0, 0, 0, 6, 1, 3, 0, 8 and 39 over the seeds.

wide on this day to steer by. No passing run curtails any solar, because Volt-Watt starts at 1.06 pu, above the band, so the 2 % limit binds only on runs that already fail (0.43 % with the curves on and the shipped compensator).

8.6.10 Predictive Control Against the Program

Two defects of the program were found and fixed while the MPC was added, and both change the comparison. First, the estimate $\hat{V}$ is fitted with the offsets in force, which the new offsets replace, yet the program had added the new ones to it and so counted the old ones twice. On the challenge with the suggested settings its summed offsets alternated 6 times in 23 acting calls from 10:00 to 14:00, for example -1.18, 0, -1.30 and 0 Mvar at 11:55 to 12:10. Starting from $\hat{V}^0$ in Equation (8.10) removed every alternation and moved the program on case 9 from 27.3 min out with 13 operations to 21.2 min with 17, and over case 9's six seeds from 155.5 to 117.0 min. The same fix took the peak per house and the power factor for the utility's model from their declared 4.6 kW and 0.95 instead of the run's values, so a heavier load is a model error that the load scale absorbs ($s_L = 1.43$ after 4 h with 6.0 kW per house and a power factor of 0.88, against 1.00 before). Second, with the curves Off the inverters run at unity power factor, yet the program, the MPC and the estimator modeled the offsets they sent. On the challenge as shipped the program asked for up to 2.43 Mvar over 10 500 s and the MPC for up to 1.58 Mvar over 6900 s, and the inverters delivered none. With no capability planned while the curves are Off, the challenge as shipped went from 143.7 to 30.1 min out for the program and from 159.7 to 11.1 min for the MPC, and cases 3 and 4 from 11.7 and 25.2 min to 0 for the program. No shipped case runs a central controller with the curves Off, and the program's numbers in Tables 8.4 and 8.5 are those after both fixes, the second moving none.

Table 8.6 sets the MPC beside the program and three stronger forms of it on the same seed, instruments and noise, the MPC at its shipped settings in every case. On case 9, the 150 % cloudy day with the suggested local settings, the MPC keeps every node between 0.9520 and 1.0498 pu all day with 15 operations, while the program's far end passes 1.05 pu at 10:56 and reaches 1.0529 pu at 11:26, 21.2 min out with 17 operations, as shown in Figure 8.5. The MPC has stepped the tap from 4 to 2 by 09:41, ahead of the noon rise, and at 11:26 its highest node stands at 1.0475 pu, while the program, holding its last target, keeps the tap at 4 until 11:28. Over the six seeds the MPC sums 42.9 min against 117.0 min for the program, with 84 operations against 78. Figure 8.4 draws case 9's two rows of the table as bars.

No classical row dominates, and neither of the MPC's two ingredients alone reproduces its case 9 result. The program without the half step, retuned by the MPC's own rule, spends 17.2 min on case 9,

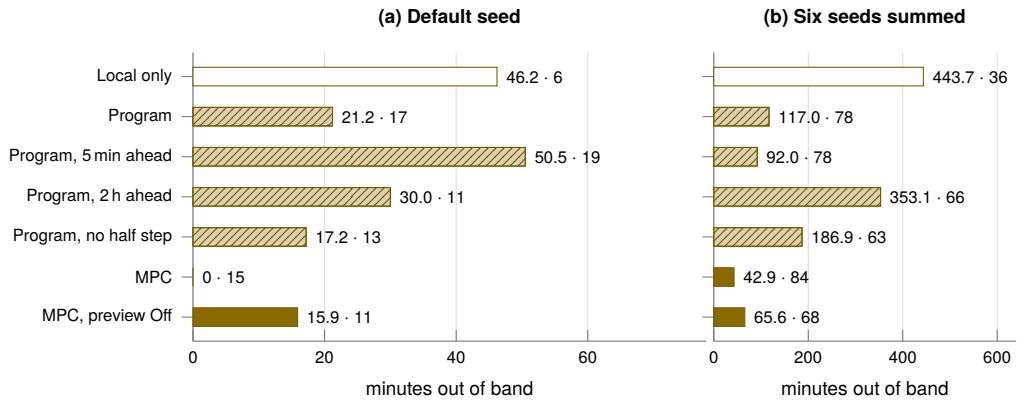

Figure 8.4 Case 9, the 150 % cloudy day with the suggested local settings, from Table 8.6: minutes any node is out of band (bars) and tap operations (after the dot), (a) on the default seed and (b) summed over six seeds. White: the local devices alone. Hatched: the program and its three stronger forms. Solid: the MPC. Notice that the MPC row is lowest in both panels, while the program's best variant changes from one panel to the other.

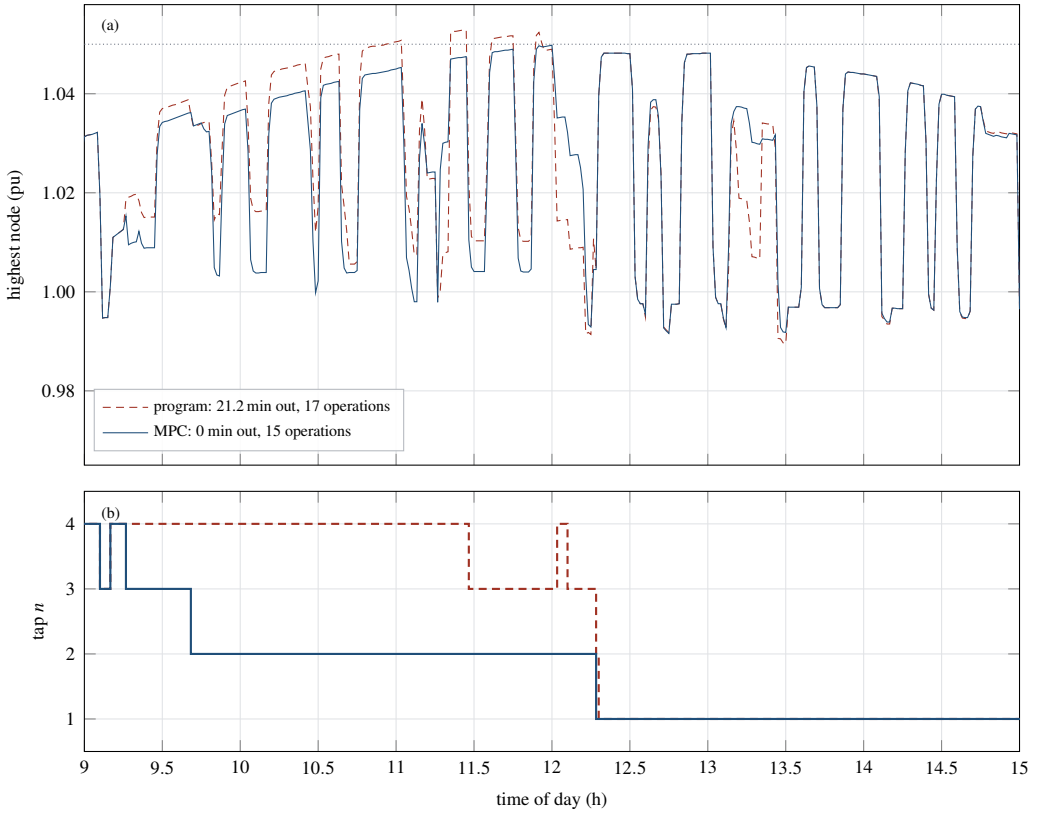

Figure 8.5 Case 9 on the default seed, 09:00 to 15:00, under the program (dashed) and the MPC (solid). (a) The highest true node voltage in each minute against the 1.05 pu limit. (b) The tap. The MPC steps the tap to 2 by 09:41, while the program keeps tap 4 until 11:28 and is out of band from 10:56.

Table 8.7 Worst Cases: Minutes Out · Tap Operations

Case	Local	Program	Prog. 5 min	MPC
3, reverse flow, 09:00 + 24 h	202.7 · 6	0 · 8	0 · 8	0 · 9
9, 150 %, cloudy	46.2 · 6	21.2 · 17	50.5 · 19	0 · 15
150 %, clear, suggested	80.6 · 6	0 · 5	0 · 5	0 · 13
challenge, suggested, shadows at 120 m/s	0 · 6	0 · 9	0 · 13	0 · 11
2, cold-load pickup at 19:00	40.2 · 9	0 · 5	0 · 5	0 · 5
6, capacitor out at 19:00	0 · 12	0 · 2	0 · 3	0 · 2
2, transmission 0.97 pu at 19:00	1.3 · 12	2.3 · 8	2.3 · 8	2.3 · 8
challenge, suggested, 1.03 pu at 12:00	3.0 · 11	1.0 · 18	1.0 · 18	1.0 · 16
challenge, suggested, meters out at 10:00	0 · 6	27.3 · 11	0 · 11	0 · 13

186.9 min over the seeds and 65.2 min on the challenge as shipped. Given the MPC's 2 h of declared shapes, the program spends fewer minutes outside the band than the MPC on the challenge as shipped, 10.2 min against 11.1 min, although neither meets its band requirement. It spends 30.0 min on case 9 and 353.1 min over the seeds, because one bus move must then hold the band against the worst of the next 2 h. With 5 min ahead it is the strongest classical row over the seeds, 92.0 min, and it holds the band on seeds 3 and 5, where the MPC spends 22.9 and 4.6 min. Without its preview the MPC still sums 65.6 min over the seeds and 15.9 min on case 9, below every program row, and the preview of the clear-sky rise takes it to 42.9 and 0. It spends fewer operations than the program on cases 4, 8 and 9 (15, 21 and 15 against 19, 31 and 17) and more on the clear days (7 against 5 at 125 %, 13 against 5 at 150 %).

The largest remaining losses come from the hold rule that the MPC shares with the program. Of the MPC's 42.9 min over the six seeds, 36.7 min fall while it holds (13.7 on seed 2 and 22.9 on seed 3), 6.2 min while a plan acts and none in a period the program decided. On seed 3 its last plan before a cloud set the bus target at 1.0070 pu with the tap at 2, and the hold kept them from 11:55 to 12:35, where the program's hold kept 0.9949 pu at tap 1 and the program spends 0.9 min out. The program loses 21.17 of its 21.2 min on case 9 while it holds. A program that keeps its offsets through a hold, a change to the shipped rule rather than a tuning, sums 14.5 min over the six seeds with 54 operations, 0 on the default seed with 7, below the MPC's 42.9 with 84, while the same change moves the MPC to 44.8 min, so the rule is left as shipped. On case 9 the MPC plans at 200 of its 288 calls, holds at 80 and hands 8 periods to the program, every one on the band-row cap, and no plan in any run needed more than 51 nodes. Without the caps the seeds sum 43.4 min and the largest solve grows from about 0.16 to about 1.5 s in Node, and holding the MPC's own last plan instead of handing over sums 27.3 min over the seeds but 29.9 min on the challenge as shipped, so the program remains the fallback.

Table 8.7 runs every structure through the worst cases the application can pose. In none of these cases is the MPC out of band longer than the program, no run curtails solar, no two tap steps come closer than T_d, the tap stays within ± 16 and no capacitor recloses inside its lockout, because the band, the timer, the tap clamp, the lockout and Volt-Watt act below either central controller. The one case in which both central structures lose to the local devices is the transmission sag to 0.97 pu at 19:00 in case 2, 2.25 min out against 1.33 min. Neither holds in that run. Before the sag both run the feeder lower, the lowest node at 0.9666 pu against 0.9725 pu, inside their margined band, so the sag takes it to 0.9370 pu instead of 0.9426 pu and the tap changer needs three 45 s steps to bring it back over 0.95 pu instead of two. The size of that loss is set by the band's timer, which neither central controller overrides.

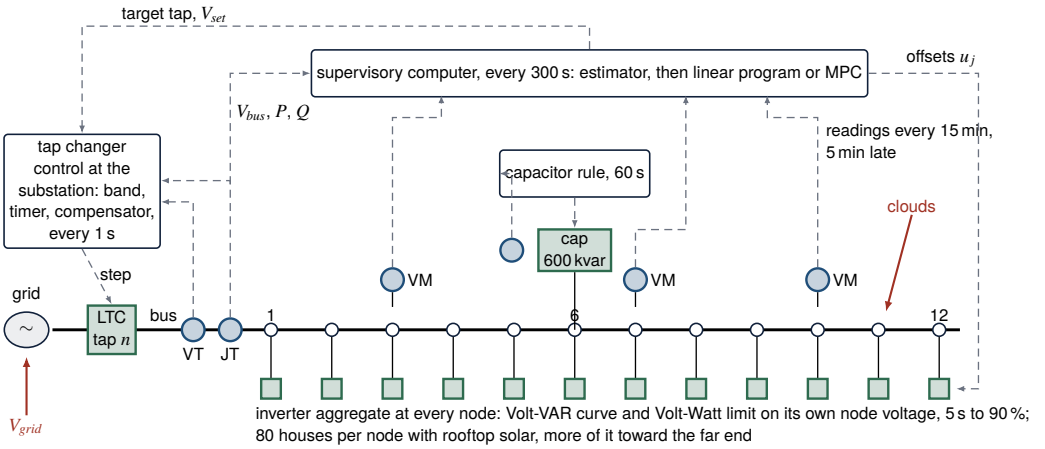

Figure 8.6 The feeder's loops drawn with their hardware. The tap changer's control at the substation reads only the voltage transformer VT and the power transducer JT on the bus and steps the load tap changer (LTC) after its timer. The capacitor at node 6 has its own voltage sensor and rule, and each node's inverter aggregate follows its curve on its own node voltage. The three smart meters VM reach only the supervisory computer, which sends a target tap and setpoint to the tap changer's control and offsets to the inverters. Notice that no device that moves the far end reads the far end except the inverters there.

> **DID YOU KNOW? When the Meters Fail, the Local Devices Win**
>
> From the simulator: on the challenge day with the suggested settings and the three meters out from 10:00, the central program spends 27.3 min out of band with 11 tap operations, while the local devices alone and the MPC spend none (Table 8.7).

8.6.11 Checks of the Code, Agreement and Limits of the Claims

Code verification tests the simulator against its own declared equations. Over the challenge day as shipped the ledger residual $E_{sub} - (E_{load} - E_{pv} + E_{loss})$ is 1.6×10^{-14} of the load energy, and the branch flows close node by node to 1.0×10^{-17} pu. One tap step is 0.6250 %, and the Volt-VAR function returns +0.44, 0, 0 and −0.44 at 0.92, 0.98, 1.02 and 1.08 pu, and −0.44 at 1.04 pu for slope 22. Every tap step in case 2 and on the hunting day follows T_d outside the band on its side, and none is decided inside it. Halving the step to 0.5 s leaves the operations of case 5 at both slopes and of case 2 unchanged (18, 2 and 7) and their energies equal to 10^{-4} MWh. The linear program's objective matches a second simplex at all 360 calls checked, within 1.9×10^{-11} relative. A day costs 0.70 s headless, and the challenge plays its 24 hours in about 24 s at 3600× in a browser.

Agreement with published relations is a separate claim. The Volt-VAR breakpoints are the category B defaults as stated in [16], the measured CVR saving of 0.79 % lies inside the simulated national range of 0.5 to 4 % [10], and the linearized flow stays within 0.46 % of the exact branch flow. The model was fitted to no feeder's records and is therefore not validated against one, and without a study of learners no learning gain is claimed.

8.7 Technology in Practice

The devices of Section 8.3 are, in the field, instrument transformers, a regulating transformer with its control, a switched capacitor bank, smart meters and inverters. Figure 8.6 draws the chapter's loops with that hardware in place of the schematic's symbols.

What Measures It

The tap changer's control reads the bus voltage and the current through the transformer, and the capacitor's control reads its own voltage. A customer's voltage is known only where a meter reports it. Smart meters give hourly demand for every customer, and engineers treat such measurements as the best basis for spreading load along a feeder, with billing data, load research statistics and transformer sizes filling the gaps [24]. Many voltage reduction schemes rely on remote measurements from a handful of locations and assume that they represent the whole system [10]. The chapter's three meters, read every 15 min and reported 5 min late, are such a handful.

> **TECHNOLOGY Sensor: Instrument Transformers at the Regulator**
>
> A tap changer's control never touches the primary voltage. A potential transformer steps the line voltage down to a safe level, typically 120 V, and a current transformer steps the line current down, typically to 5 A at full load, for meters and protective relays [25]. The control combines that local voltage and current with an impedance setting, entered in volts, to estimate the drop along the circuit [24]. Feeder-head data may be logged every minute or every second, and bad data, from communication failures, outages and unusual configurations, are to be expected. The rate of change of a measurement helps to find them [24]. The chapter's VT and JT carry 0.05 % and 0.5 % noise.

What Moves It

Three kinds of final element act on the feeder. The regulating transformer steps the source for every customer. A switched capacitor bank adds reactive power at one node under voltage, time, temperature, time-biased, reactive power or current control [24]. The inverters inject or absorb reactive power along a Volt-VAR curve, which may carry hysteresis and a deadband, with a response time that can be programmed between 300 ms and several seconds [24]. The chapter's inverters reach 90 % of a change in 5 s, while its tap changer waits 45 s before each step and its capacitor 60 s.

> **TECHNOLOGY Final Element: The On-Load Tap Changer**
>
> A load tap changer is a power transformer that changes its turns ratio under load. Common ranges are plus or minus 5 or 10 % in 16 or 32 discrete steps, and a substation unit normally moves all three phases together on the feedback of one [24]. Its control holds a setpoint inside a bandwidth and waits out a time delay before each step. A unidirectional unit assumes that power flows away from the substation, while a bidirectional one senses the direction and switches to a second group of settings [24]. Photovoltaic swings that multiply the steps can shorten the device's life and add maintenance [24]. The chapter's tap changer has ± 16 steps of 0.625 % and no mechanism time.

Where the Controller Runs

A tap changer's control sits beside the transformer and runs the band, the timer and the compensator on its own measurements, and each capacitor and each inverter carries its own rule. Above them, the inverters' Volt-VAR control can be coordinated at the level of supervisory control and data acquisition with capacitor banks and voltage regulators [24]. Supervisory and local loops of this kind run with cycle times from milliseconds to minutes [26]. In the chapter the local devices execute every second, and the estimator, the linear program and the predictive controller every 300 s, as the supervisory layer that sends a tap target, a setpoint and the inverters' offsets.

What Is Modernizing It

Three developments touch these loops. The first is the smart inverter. Many inverters already in service can provide advanced functions with a software upgrade or new parameters, and an inverter works below its maximum current rating more than 95 % of the time, which leaves capacity for reactive power [17]. In 2014 these functions had to be set inside each inverter by hand, and setting them remotely was foreseen [17]. The second is communication. Ethernet, Internet Protocol and wireless

devices are replacing older proprietary technologies, which raises the likelihood of cybersecurity vulnerabilities and incidents in control systems [26], a concern for a program that sends offsets to every node. The third is central optimization over a forecast. The predictive controller of Section 8.5.5 is one form of it, and Section 8.6.10 measures what it buys on this feeder.

> **IN PRACTICE** **Regulator Settings on a Solar Feeder**
>
> Planners coordinate regulators in series by their time delays [24]. A compensator set on the impedance to the first customer raises the regulator's output with load current while that customer stays below 126 V on a 120 V base, the first house setting, and a line that uses a compensator should be monitored [24]. With variable photovoltaic output the bandwidth may need widening, and some utilities disable the compensator and move the setpoint instead. When devices hunt, a nonunity inverter power factor is tried first, then the setpoint, bandwidth and delays [24]. The voltage limits are those of ANSI C84.1 [24].

8.8 Limitations

Every value is synthetic. The feeder is balanced, single-phase and linearized, so it carries none of the unbalance, the single-phase laterals or the secondary drops that decide many real voltage complaints, and its far end reads up to 0.46 % high. The impedances, the solar layout and the load shape were chosen so that noon overvoltage exists, and the shipped compensator and the hunting window are teaching choices that the text labels. The ZIP shares are fixed and the inverters are ideal aggregates with continuous response. The Volt-Watt breakpoints are not confirmed against the standard, and the band and the curves are teaching values from ANSI C84.1 and IEEE 1547-2018 rather than any utility's practice. The estimator's margins are conservative by design, so the central program is weaker here than local coordination on a cloudy day, and a real distribution management system has more meters and a different model. A simulated pass is not a certification or a planning study.

The MPC inherits the estimator's errors because its model is the estimator's declared feeder. It forecasts no cloud, holds the sensitivities S and G fixed over 2 h, plans neither the band's own steps between decisions nor the capacitor, which stays on its shipped rule under both central controllers, and treats the tap as an integer over only the first 20 min, with a continuous ramp after. It shares the program's hold rule, which carries 36.7 min of its 42.9 min over case 9's six seeds, and the comparison rests on one feeder and six weather seeds.

> **OPEN CHALLENGE** **Thousands of Volt-VAR Curves on One Network**
>
> As inverters replace machines, reactive power and voltage control move into thousands of devices with local Volt-VAR curves. Time delays and filters that slow those curves are known to cause voltage oscillations and interactions between inverters and the voltage-regulation equipment on the same distribution network, and piecewise linear curves have inherent instability points [27]. A 2020 national laboratory roadmap asks how the curves should be coordinated with capacitor banks and other devices, how the instability points are avoided, and whether reactive power should be controlled at each inverter, by local groups or over communication networks [27]. Progress would be coordination rules that answer these questions for any feeder.

8.9 Conclusion

A distribution feeder with rooftop solar is a plant whose slow discrete actuator sits at one end while its disturbances and its worst customer sit at the other, and the simulator of this chapter puts that on one screen with a linearized branch flow, ZIP loads, seeded clouds, category B inverter curves, a capacitor and a tap changer with a band, a timer and a compensator. At solar noon the compensator holds the bus 0.0083 pu high while 1.08 MW flows back, because the reactive power still flows forward and X/R exceeds 1.68. A cloud moves the far end 0.031 pu and the bus 0.001 pu. A Volt-VAR slope of 22 against a reactance-weighted compensator turns 2 tap operations into 18 on a cloudy day, inside a window the text states. Conservation voltage reduction saves 0.79 % on the declared loads and costs

energy on constant-power ones, and three meters with the substation's telemetry bracket the worst customer at 47 of 48 fits in case 8. A tap changer alone just passes the cloudy challenge at 125 % solar and passes nowhere at 150 %, where the inverter curves still do. A mixed-integer predictive controller on the same estimate and weights holds the 150 % cloudy day that the program misses by 21.2 min, although on the shipped challenge, which neither holds, a program given a 2 h look-ahead spends fewer minutes out of band, 10.2 min against 11.1 min. The model is balanced and linear and its settings are declared. It leaves out unbalanced phases, line regulators with reverse-power modes, a second feeder on the bus, inverter ride-through and measured load shapes. The simulator shows that on a solar feeder the slow device should set the level and the fast ones should trim the spread.

KEY POINTS

- A tap step moves every customer with the bus, so only devices at the nodes can change the spread of voltage along the feeder.
- At solar noon the compensator holds the bus 0.0083 pu high while 1.08 MW flows back, because the reactive power still flows forward.
- A cloud moves the far end 0.031 pu and the bus 0.001 pu.
- A Volt-VAR slope of 22 against a reactance-weighted compensator turns 2 tap operations into 18 on a cloudy day.
- A mixed-integer predictive controller holds the 150 % cloudy day that the program misses by 21.2 min.

Where else this applies. Integer decisions found by branch and bound also schedule the pumps of Volume 9, Chapter 11, and a late measurement limits the composition loop of Volume 5, Chapter 3.

References

[1] M. E. Baran and M.-Y. Hsu, "Volt/VAr control at distribution substations," *IEEE Transactions on Power Systems*, vol. 14, no. 1, pp. 312–318, 1999.

[2] R. A. Walling, R. Saint, R. C. Dugan, J. Burke, and L. A. Kojovic, "Summary of distributed resources impact on power delivery systems," *IEEE Transactions on Power Delivery*, vol. 23, no. 3, pp. 1636–1644, 2008.

[3] F. A. Viawan, A. Sannino, and J. Daalder, "Voltage control with on-load tap changers in medium voltage feeders in presence of distributed generation," *Electric Power Systems Research*, vol. 77, no. 10, pp. 1314–1322, 2007.

[4] Y. P. Agalgaonkar, B. C. Pal, and R. A. Jabr, "Distribution voltage control considering the impact of PV generation on tap changers and autonomous regulators," *IEEE Transactions on Power Systems*, vol. 29, no. 1, pp. 182–192, 2014.

[5] P. Jahangiri and D. C. Aliprantis, "Distributed volt/VAr control by PV inverters," *IEEE Transactions on Power Systems*, vol. 28, no. 3, pp. 3429–3439, 2013.

[6] IEEE, "IEEE standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces," IEEE Std 1547-2018, 2018.

[7] H. J. Carlin, "A new voltage-regulating relay plus line-drop compensator," *Electrical Engineering*, vol. 61, no. 2, pp. 63–66, 1942.

[8] W. H. Kersting, *Distribution System Modeling and Analysis*. Boca Raton, FL, USA: CRC Press, 2017.

[9] R. F. Preiss and V. J. Warnock, "Impact of voltage reduction on energy and demand," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-97, no. 5, pp. 1665–1671, 1978.

[10] K. P. Schneider, J. C. Fuller, F. K. Tuffner, and R. Singh, "Evaluation of conservation voltage reduction (CVR) on a national level," Pacific Northwest National Laboratory, Richland, WA, USA, Tech. Rep. PNNL-19596, 2010.

[11] M. Baran and F. F. Wu, "Optimal sizing of capacitors placed on a radial distribution system," *IEEE Transactions on Power Delivery*, vol. 4, no. 1, pp. 735–743, 1989.

[12] M. E. Baran and F. F. Wu, "Network reconfiguration in distribution systems for loss reduction and load balancing," *IEEE Transactions on Power Delivery*, vol. 4, no. 2, pp. 1401–1407, 1989.

[13] M. E. Baran and A. W. Kelley, "State estimation for real-time monitoring of distribution systems," *IEEE Transactions on Power Systems*, vol. 9, no. 3, pp. 1601–1609, 1994.

[14] J. W. Smith, W. Sunderman, R. Dugan, and B. Seal, "Smart inverter volt/var control functions for high penetration of PV on distribution systems," in *2011 IEEE/PES Power Systems Conference and Exposition*, Phoenix, AZ, USA, 2011, pp. 1–6.

[15] H. Zhu and H. J. Liu, "Fast local voltage control under limited reactive power: Optimality and stability analysis," *IEEE Transactions on Power Systems*, vol. 31, no. 5, pp. 3794–3803, 2016.

[16] I. Murzakhanov, S. Gupta, S. Chatzivasileiadis, and V. Kekatos, "Optimal design of volt/VAR control rules for inverter-interfaced distributed energy resources," *IEEE Transactions on Smart Grid*, vol. 15, no. 1, pp. 312–323, 2024.

[17] National Renewable Energy Laboratory, "Advanced inverter functions to support high levels of distributed solar: Policy and regulatory considerations," Brochure NREL/BR-6A20-62612, Golden, CO, OSTI 1164879, 2014.

[18] K. Turitsyn, P. Sulc, S. Backhaus, and M. Chertkov, "Options for control of reactive power by distributed photovoltaic generators," *Proceedings of the IEEE*, vol. 99, no. 6, pp. 1063–1073, 2011.

[19] M. Louie, C. Kelley, M. Williams, M. Ropp, and R. Darbali-Zamora, "A review of distributed wind interconnection technology and standards," Sandia National Laboratories, Albuquerque, NM, Sandia Report SAND2026-16105, 2026.

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

[21] A. Bemporad and M. Morari, "Control of systems integrating logic, dynamics, and constraints," *Automatica*, vol. 35, no. 3, pp. 407–427, 1999.

[22] G. Valverde and T. Van Cutsem, "Model predictive control of voltages in active distribution networks," *IEEE Transactions on Smart Grid*, vol. 4, no. 4, pp. 2152–2161, 2013.

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

[24] R. Seguin, J. Woyak, D. Costyk, J. Hambrick, and B. Mather, "High-penetration PV integration handbook for distribution engineers," National Renewable Energy Laboratory, Golden, CO, Technical Report NREL/TP-5D00-63114, 2016.

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

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

[27] Y. Lin, J. H. Eto, B. B. Johnson, J. D. Flicker, R. H. Lasseter, H. N. Villegas Pico, G.-S. Seo, B. J. Pierre, and A. Ellis, "Research roadmap on grid-forming inverters," National Renewable Energy Laboratory, Golden, CO, Technical Report NREL/TP-5D00-73476, 2020.

8.10 Exercises

8.1 ⌕ Concept check · 2 min

At 12:00 in case 3 the substation exports 1.08 MW while 0.64 Mvar still flows in from the grid. The compensator is Forward with $R_{set} = 0.01$ pu and $X_{set} = 0.15$ pu. What does it do to the bus?

(a) It holds the bus up, because its reactive term outweighs the active one while X_{set}/R_{set} exceeds $-P/Q$

(b) It pulls the bus down, because the current through the transformer has reversed together with the active power

(c) It does nothing, because a Forward compensator drops its correction while the measured active power is negative

(d) It pulls the bus down by the whole line drop, because the far end is the high point when the solar is spread out

8.2 ⌕ Concept check · 2 min

In case 4 a cloud bank removes 60 % of the sunlight in 10 s and stays 25 min. With the inverter curves off, the far end falls from 1.0447 pu to 1.0137 pu by 12:40, the bus moves only from 1.0268 pu to 1.0258 pu, and the tap never moves. Why?

(a) The time delay of 45 s is longer than the fall, so any step would have come too late to help

(b) The drop happens along the line beyond the bus, so the regulated voltage never stays outside its band for 45 s

(c) The compensator cancels the cloud in the regulated voltage, so the regulator sees no change at all

(d) The tap changer is locked out after the cloud's first edge, as the capacitor is after it opens

8.3 Concept check · 2 min

The three meters of case 8 fail at 12:00. Afterward the inferred worst node's margin still holds the true worst voltage at all 24 fits, with a median margin of 0.0191 pu against 0.0176 pu before. What keeps the estimate useful without the meters?

(a) The meters' last readings are replayed at every fit, so the profile stays fixed at the one of 12:00

(b) The estimator reads the logged truth while the meters are stale, which the wider margin then reports

(c) The substation's P and Q alone fix the load and solar scales, and the stale flag widens the margin

(d) The central program holds its last target during the outage, so the voltage profile stops moving

8.4 Theory and calculation · 6 min

The feeder is 12.47 kV on $S_b = 10\,\text{MVA}$. Segments 1 to 6 are $0.306 + j0.385\,\Omega$/km, segments 7 to 12 are $0.55 + j0.42\,\Omega$/km, and each segment is 0.8 km long. One tap step is 0.625 %.

(a) Find the base impedance, the whole feeder's impedance in per unit and its R/X. *Checkpoint:* $15.55\,\Omega$, $0.2642 + j0.2485$ pu, $R/X = 1.06$.

(b) Express the shipped band of 1.67 % and the bands of 1.25 % and 1.0 % in tap steps, and say why a band must be wider than one step. *Checkpoint:* 2.67, 2.0 and 1.6 steps.

(c) A shadow edge sweeps from node 12 toward the bus at 60 m/s. How long after node 12 does it reach node 6 and the bus? *Checkpoint:* 80 s and 160 s.

8.5 Theory and calculation · 6 min

At 12:00 in case 3, $P = -1.081\,\text{MW}$, $Q = 0.645\,\text{Mvar}$ and $V_{bus} = 1.039\,\text{pu}$ on $S_b = 10\,\text{MVA}$. Use Equation (8.6).

(a) Find the correction c with the shipped dials $R_{set} = 0.01\,\text{pu}$, $X_{set} = 0.15\,\text{pu}$, and the dial ratio above which c is positive. *Checkpoint:* $c = 0.008\,27\,\text{pu}$, ratio 1.68.

(b) Find c for a compensator matched to the line, $R_{set} = X_{set} = 0.06\,\text{pu}$. *Checkpoint:* $c = -0.0025\,\text{pu}$.

(c) On forward flow at a steady power factor of 0.95, the correction is proportional to $R_{set} + X_{set}\tan\varphi$. Evaluate it for the shipped dials and for $R_{set} = X_{set} = 0.045\,\text{pu}$. *Checkpoint:* 0.0593 and 0.0598.

8.6 Theory and calculation · 6 min

The category B Volt-VAR curve has a deadband $\delta = 0.02\,\text{pu}$ and saturates at 44 % of the rating. The installed solar is 75 % of the 4.416 MW evening peak, on $S_b = 10\,\text{MVA}$, and the compensator has $X_{set} = 0.15\,\text{pu}$.

(a) Find the default slope that reaches saturation at 1.08 pu, and the voltage at which the steepest allowed slope, 22, saturates above nominal. *Checkpoint:* 7.33 and 1.04 pu.

(b) The open-loop response time is $T_{ol} = 5\,\text{s}$ to 90 %. Find the time constant τ of the first-order lag. *Checkpoint:* 2.17 s.

(c) Find the coupling factor $X_{set} k_{VV} \sum_j S_j$ by which a common change of the node voltages inside the slope moves the compensator's correction, at both slopes. *Checkpoint:* 0.36 and 1.09.

8.7 App challenge · 6 min

Load case 2, *An evening without solar: the band and the timer*, with the default seed. It runs from 12:00 and pauses at 24:00. At 3600× that takes about 12 s. Read the scorecard tile *Customers in band (truth) · tap operations* at the pause.

(a) Run as loaded. Give the tap operations and the lowest node voltage. *Checkpoint:* 7 operations, lowest node 0.972 pu.

(b) Set Bandwidth to 2.50 %, press Reset, and run to the pause again. *Checkpoint:* 5 operations, lowest node 0.970 pu.

(c) Press the reload button to restore the case, set Time delay to 120 s before Run, and run to the pause. *Checkpoint:* 5 operations, lowest node 0.972 pu.

8.8 App challenge · 7 min

Load case 5, *The slope: a steep curve that hunts*, with the default seed. It runs a cloudy day from 09:00 with the curves on and pauses at 17:00, about 8 s at 3600×. Read the tap operations on the scorecard tile *Customers in band (truth) · tap operations* at the pause, and watch the status line of the scene.

(a) Run as loaded, at the default Volt-VAR slope of 7.33. *Checkpoint:* 2 operations.

(b) Set Volt-VAR slope to 22.00, press Reset, and run to the pause. Does the status line read HUNTING? *Checkpoint:* 18 operations, and yes, first shortly before noon.

(c) Keep the slope at 22.00, set Compensator X (Advanced) to 0.120 pu, press Reset and run to the pause. *Checkpoint:* 2 operations.

8.9 App challenge · 8 min

Load case 9, *150 % solar: the MPC against the linear program*, with the default seed. It runs a cloudy day from midnight with the suggested local settings and Central controller on MPC, and pauses at 24:00, about 24 s at 3600×. Read the tile *Customers in band (truth) · tap operations* at the pause.

(a) Run as loaded. *Checkpoint:* In band, nodes 0.952 to 1.050 pu, 15 operations.

(b) Press the reload button, set Central controller to On (the linear program) before Run, and run to the pause. *Checkpoint:* 21.2 min out, highest node 1.053 pu, 17 operations.

(c) Press the reload button, set MPC preview (Advanced) to Off before Run, and run to the pause. What share of the program's minutes out does the MPC remove without its preview, and what share with it? *Checkpoint:* 15.9 min out, 11 operations, about 25 % without the preview and 100 % with it, within 1 point.

8.10 Application area · 4 min

A utility considers conservation voltage reduction on a residential feeder and has the case 7 results of this chapter: over 16:00 to 22:00 the baseline delivers 24 053 kWh and the run with conservation voltage reduction 23 863 kWh, while the load-weighted mean voltage falls 1.04 %. The loads are 30 % constant impedance, 20 % constant current and 50 % constant power.

(a) Find the saving in kWh and in percent, the feeder's measured factor and the loads' own factor $2a_Z + a_I$. *Checkpoint:* 190 kWh, 0.79 %, 0.76 against 0.8.

(b) Explain why the measured factor is lower than the loads' factor, and what the result would be on a feeder of constant-power loads.

8.11 Application area · 4 min

A distribution engineer wants to stop the noon overvoltage of case 3 by switching the regulator's compensator from Forward to Reverse-aware, keeping the shipped dials $R_{set} = 0.01$ pu, $X_{set} = 0.15$ pu.

(a) State what the change fixes at noon and what it costs on a cloudy day, with the measured counts of case 5 at the default slope.

(b) Name the local settings that the chapter shows passing the cloudy 125 % challenge on all six seeds, and the dial ratio they use.

8.12 Review the agent · 5 min

An AI assistant reviews case 5 and writes: “At slope 22 the tap changer makes 18 operations against 2 at 7.33, so the inverters’ Volt-VAR loop is unstable above a slope of about 20. The stability condition on the sampled update in the Volt-VAR literature confirms it. Cap the slope at 7.33 on every feeder.” Find the flaw and say how you would check it in the application.

CHAPTER 9

Electric Vehicle Fast Charge Simulator for Selector Control of an Unmeasured Constraint

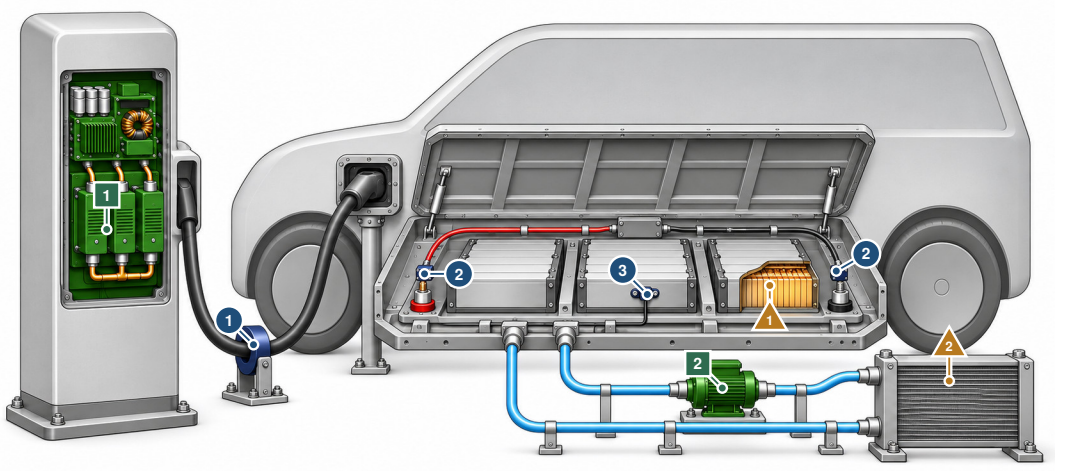

The controller moves 1 charger power stages that set the current 2 coolant pump

The controller reads 1 current sensor on the cable 2 voltage taps at the pack terminals 3 thermistor on a module case

What disturbs it 1 cell aging the stored table does not know 2 ambient air at the chiller

Notice that every sensor sits outside the cells, while the limit that ends the charge lies at the graphite surface inside them. Illustration, not to scale.

IN PLAIN TERMS **Charging Fast Without Hurting the Battery**

The main idea. A fast charger pushes current into a car's battery as fast as the battery can safely take it. The real limit is a tiny voltage at the carbon surface inside each cell, and no sensor in the car can read it.

The challenge. So the car works it out from a stored table and a few gauges, and keeps a safety gap. If the table is wrong, as it becomes when the battery ages, lithium metal coats the carbon while every gauge looks normal.

Why it matters. Each such charge quietly takes away some of the battery's driving range for good.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the history box of Section 9.2, the fact box of Section 9.3, the try-it and fact boxes of Section 9.6, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	25 min	The first page, Section 9.6, Section 9.7, Section 9.8 and the open challenges

The control question Four limits share one charging current, and the one that ends a fast charge, the anode surface potential, has no instrument. How should a selector defend a constraint it can only compute, and what margin does an error in that computation cost?

Learning objectives

- Compute the anode-limited current from the plateau potential and the kinetic resistance at a state of charge and temperature.
- Explain why the anode limit is an inverse map that cannot be given integral action.
- Size the anode margin as a budget for kinetic table error.
- Explain how external reset and a capped preload keep a low selector from ratcheting its output down.
- Compare the selector with a predictive controller on time, peak core temperature and plated lithium.

Prerequisites Sections 1.4.4 (storage state and tapered limits), 1.4.7 (a limit or selector joining the layers), 1.6.2 (velocity form, external reset), 1.6.3 (predictive control, Kalman filter) and 1.6.4 (scenarios and data).

The app https://apmonitor.com/apps/ev-fast-charge/

See also Volume 4, Chapter 3 for a min/max selector over limiters guarding a surge margin no sensor reads, and Volume 4, Chapter 2 for a protective override that rests on its clamp. Volume 5, Chapter 11 and Volume 5, Chapter 10 for control of a temperature and a constraint that are inferred rather than measured.

9.1 Introduction

A constraint in an introductory control problem arrives with an instrument behind it. A pressure has a transmitter, a level has a float, a metal temperature has a thermocouple welded to it, and the controller that defends the constraint acts on the difference between what the instrument reports and where the limit was drawn. Charging a lithium-ion vehicle pack breaks that arrangement at the only point where it matters. Three of the four limits on the charging current are instrumented, being the rating of the charging post, the terminal voltage of the pack and the cell temperature, and not one of the three is what ends a fast charge. What ends it is the electric potential of the graphite surface against lithium metal, and no vehicle in service carries a reference electrode that reports it [1].

That potential decides the chemistry. Lithium arriving at the negative electrode intercalates into the graphite while the surface stays above the potential of lithium metal, and below that value it deposits as metal on the surface instead [2]. Part of the deposit never rejoins the inventory, and what remains is a rough surface that deposits more readily on the next charge [3], with the simulator treating 62 % of it as permanently lost. The equilibrium potential of graphite falls onto a low plateau over most of its range of lithiation [4], taken here as 85 mV, so the whole margin available to a fast charge is a few tens of millivolts of overpotential, and the overpotential is the charging current multiplied by a kinetic resistance that rises steeply as the electrode fills and as the cell cools [5].

A block diagram of this plant has a comparator drawn in front of the anode controller, and drawing it is the mistake. What the application corrects is the belief that a limit is defended as soon as a setpoint has been written for it. A feedback loop is wrong by the amount its plant departs from its setpoint, and the error signal that measures the departure is the same signal that removes it. An inverse model carries no such signal. It is wrong by the amount the model departs from the plant, that discrepancy appears nowhere in the instrument set, and integral action cannot be added because there is no error to integrate. The application shows the two side by side on one screen, with the computed anode potential that a vehicle would display drawn against the true anode potential that a vehicle cannot obtain, and the distance between the two lines is the subject of this chapter.

The class of problem is constraint control on an inferred variable that shares an actuator with measured constraints, resolved by a selector, and it is general. Selective and override schemes are the standard answer wherever several limits compete for one handle, and the losing controllers have to be told what was actually applied or their integrals will wind up while they are out of service [6, 7]. Battery management is a well-documented instance of the same structure [8], and the shape recurs

whenever the binding quantity sits somewhere an instrument cannot go, such as a metal temperature inside a casting, a film boiling point on the inside of a tube, or a stress at a weld root. The core of a brick in the cathode kiln of Volume 5, Chapter 11 is another such place. In each case the constraint enters the selector as a computed number, stands beside constraints that do have transmitters, and is protected by a margin that somebody has to justify to a design review.

This chapter makes three kinds of statement. The first is a measurement rather than a method, a controlled comparison on one pack between a constant-current charge and an anode-limited charge that reports the four and a half minutes the second one costs alongside the 3.381 Ah of plated lithium the first one produces, and a sweep of the constraint margin against kinetic model error that turns the margin into a budget with numbers in it. The second is a nine-state electrochemical, thermal and aging model with a four-way low selector, written out below in enough detail to be rebuilt from the text. It runs in a browser at interactive speed beside an optional model predictive controller that is measured against the selector on seven cases. The third is the scenario set and the manual mode underneath it, where the plating threshold is found by pushing current into the pack and watching a line that a real vehicle does not have. Neither the cell model nor the selector arithmetic is new. The aim is to write down a fast charge whose binding limit is computed instead of measured, and to price what that substitution costs when the computation is wrong.

Section 9.2 follows the domain from constant-current charging to the charge current limit map. Section 9.3 states the plant, its parameters and its assumptions. Section 9.4 reads the four limits as functions of the operating point, and Section 9.5 gives the implemented scheme: the shared PID routine, the anode map, the selector with its preload and the predictive alternative. Section 9.6 carries the measured scenarios, Section 9.7 describes the hardware of a real vehicle and charger, Section 9.8 states what the model cannot support, Section 9.9 draws the lessons that carry beyond batteries and Section 9.10 collects the main numbers.

9.2 From Constant Current to a Charge Current Limit Map

The protocol that lithium-ion cells inherited is constant current followed by constant voltage. Current is held at a rate the cell vendor guarantees until the terminal voltage reaches its ceiling, and the voltage is then held while the current decays. Both numbers come from a data sheet rather than from the state of the cell, which is why the protocol is safe at a third of a C and unsafe at two, and why the first serious studies of fast charging were comparisons of protocols on cycle life rather than analyses of any mechanism. Cells charged to the same terminal voltage by different current histories were shown to lose capacity at markedly different rates, with the loss traced to the negative electrode rather than to the positive [9]. Pulse and boost protocols were proposed on the reasoning that a resting interval lets concentration gradients relax and raises the surface potential before the next pulse [10], which is a correct mechanism attacked without a measurement of the quantity it acts on.

The quantity itself was available in the modeling literature well before it reached a vehicle. The porous-electrode formulation of the lithium insertion cell resolves solid-phase and electrolyte concentrations and returns the local electrode potential at every point through the thickness [11], and a deposition side reaction was added to that framework with a criterion for the onset of metallic lithium expressed directly as the potential of the solid surface against a lithium reference [2]. The equilibrium half of that potential had been measured for graphite over its whole range of lithiation, with the staging plateaus that produce the flat region a fast charge has to work inside [4]. Low-temperature work then established the part that surprises drivers, that a cold cell is slow because the intercalation kinetics are slow rather than because a controller is being cautious [5], and a later review collected the evidence that plating is partly irreversible and that its consequences are morphological as well as coulombic [3].

Control formulations followed once the criterion existed. The estimation and control problems specific to a pack were set out as a class, with the observation that the quantities a management system would most like to bound are internal states rather than terminal measurements [8], and a charging

problem was then posed explicitly as maximizing the charge delivered subject to a lower bound on the computed side-reaction overpotential rather than to an upper bound on the terminal voltage [12]. Later work carried the formulation into a coupled electro-thermal-aging setting and traded the charge time against the predicted capacity loss in one objective [13]. Production practice took a simpler route than an on-line optimizer. The solution shipped in vehicles is a charge current limit map, a table of allowable current against state of charge and temperature computed offline from the same criterion, looked up every scan and passed to a selector alongside the voltage and temperature limits, and it is the reason a modern charge curve tapers continuously from the moment the charger ceiling is released. Reviews of fast charging from the end of the last decade describe the limit map as the standard answer and identify its weakness, which is that the table is fitted to particular cells at a particular age [14], and the same weakness appears in technology gap assessments alongside the thermal and infrastructure constraints on a ten to eighty percent charge in fifteen minutes [15]. Closing that gap by measurement rather than by margin is an active subject, with adaptive schemes that identify a virtual reference electrode on line and control the anode potential against it directly [1]. The application in this chapter implements the standard answer and the failures that motivate the adaptive one.

> ⧖ HISTORY **A Ten-Minute Target for Fast Charging (2017)**
>
> In 2017 a group of United States national laboratories assessed what it would take to recharge an electric vehicle about as quickly as a car is refueled. Direct-current fast chargers then on the market delivered up to 120 kW, and the assessment defined extreme fast charging as recharging up to 80 % of the battery capacity in 10 min or less [15]. Most packs of the day were rated at or below 400 V, and the companion vehicle study set 800 to 1000 V and more than 400 kW as the design point for future vehicles [16]. The assessment also warned that the conventional constant-current, constant-voltage protocol may not suit such charging, because of the high probability of lithium plating toward the end of the charge [15].

9.3 Pack and Cell Model

The simulated vehicle carries 96 cells in series of 200 Ah each, which is 316.8 V at the empty end of the open-circuit curve and 398.4 V at the full end, and 49.7 kWh of stored energy between 10 and 80 % state of charge. The post is rated 250 kW and the vehicle contactors are rated 500 A, which at these terminal voltages is 165 to 199 kW, so the current ceiling binds on this pack and the power rating never does. A liquid loop carries heat from the pack to a chiller whose capacity falls with ambient temperature. All of that appears in Figure 9.1, together with the instruments, the block that computes the anode potential and every symbol carried into the equations of this section.

The clock runs in minutes because a charge takes 17 to 45 of them. Currents are in amperes, voltages in volts, temperatures in degrees Celsius with the Arrhenius terms converted to kelvin internally, heat rates in kilowatts and charge in ampere hours. State of charge is a fraction in the equations and is displayed as a percentage.

9.3.1 Variables

Everything the application exposes is classified in Table 9.1. The single manipulated variable that matters is the charging current, and four candidate controllers write to it. The coolant flow is a second manipulated variable belonging to a loop of its own.

9.3.2 The Cell as a Circuit

Terminal behavior is an open-circuit curve, two polarization branches and an ohmic drop, as shown in Equation (9.1).

$$\begin{aligned} V &= n_s U_{ocv}(z) + v_1 + v_2 + I R_0(T) \\ \frac{\mathrm{d}v_k}{\mathrm{d}t} &= \frac{60\,(I R_k - v_k)}{R_k C_k}, \qquad k = 1,2 \end{aligned} \tag{9.1}$$

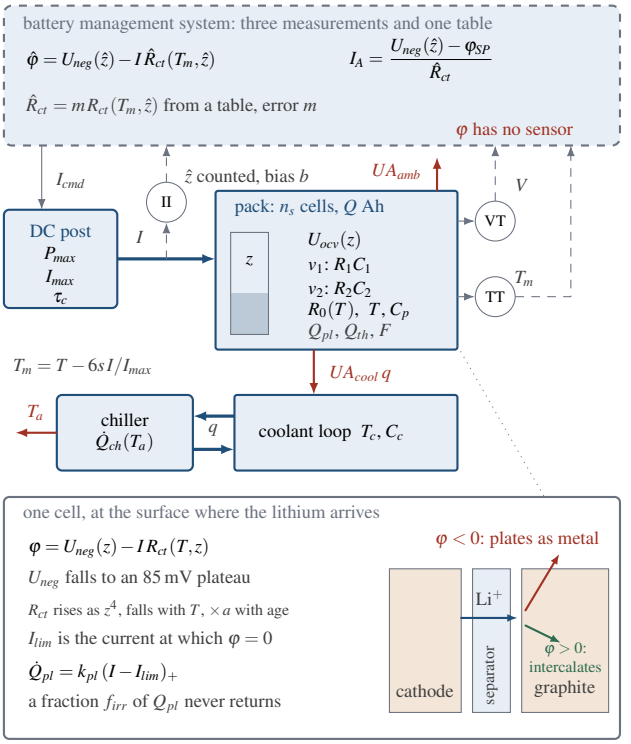

Figure 9.1 Vehicle pack on a direct-current post, with the coolant loop beneath it and the cell surface enlarged below. The management system reads the pack current, the terminal voltage V and one thermistor T_m on the module case, counts the state of charge as $\hat{z}$, and computes the anode surface potential $\hat{\varphi}$ from a kinetic table carrying error m. Solid heavy lines are power, dashed lines are signals, and the short thick arrows labeled with a conductance or the ambient temperature are heat. Nothing in the vehicle measures φ.

where n_s is the series count, v_k is the voltage across branch k of the pack in volts, R_k and C_k are that branch resistance and capacitance, and the factor 60 converts the branch time constant from seconds to minutes. The open-circuit curve carries a steep tail and a low-state-of-charge knee, as shown in Equation (9.2).

$$U_{ocv}(z) = 3.30 + 0.55z + 0.15z^8 + 0.15\left(1 - \mathrm{e}^{-20z}\right) \tag{9.2}$$

in volts per cell, giving 3.30 V empty, 3.73 V at half charge and 4.15 V full. Every resistance in the pack is Arrhenius in temperature and grows together with cell age, as shown in Equation (9.3).

$$R_j(T) = R_j^0\, a \exp\left[E_R\left(\frac{1}{T_K} - \frac{1}{T_{ref}}\right)\right], \quad j = 0,1,2 \tag{9.3}$$

where $T_K = T + 273.15\,\mathrm{K}$ is the absolute pack temperature, $T_{ref} = 298.15\,\mathrm{K}$, E_R is an activation temperature in kelvin and a is the aging multiplier. Equation (9.3) is why a cold pack is a resistive pack and a resistive pack is a hot pack, and 500 A through the nominal 23.5 mΩ deposits 5.88 kW in the cells, which is the same order as the chiller.

9.3.3 The Potential That Decides the Chemistry

The quantity the whole application is built around is the potential of the graphite surface against lithium metal, and it is the equilibrium potential of the electrode less the overpotential the current has

Table 9.1 Variable classification. Ranges are the limits the application enforces on a slider or a state.

Symbol	Meaning	Unit	Range
Constrained and controlled variables			
φ	anode surface potential, true	mV	not measured
$\hat{\varphi}$	anode potential, as computed	mV	SP 0 to 80
V	pack terminal voltage	V	limit $n_s V_{max}$
T	pack core temperature	°C	derate SP 30 to 55
z	state of charge	%	target 50 to 100
Manipulated variables			
I_{cmd}	charging current command	A	0 to 500
q	coolant flow	%	0 to 100
Disturbances and the post rating			
T_a	ambient temperature	°C	−15 to 45
m	kinetic table error	×	0.50 to 1.60
a	cell aging multiplier	×	1.00 to 2.20
s	thermistor offset scale	×	0 to 2
b	current sensor error	×	0.950 to 1.050
P_{sp}	power the post offers	kW	20 to 400
States			
z	state of charge	—	0 to 1.02
v_1, v_2	polarization voltages	V	≥ 0
T	pack core temperature	°C	−30 to 90
T_c	coolant temperature	°C	unbounded
I	current flowing	A	0 to I_{max}
Q_{pl}	plated lithium	Ah	≥ 0
Q_{th}	charge throughput	Ah	≥ 0
F	capacity fade	%	0 to 40
Measurements			
$\hat{z}$	counted state of charge	%	0 to 105
$\hat{V}$	filtered terminal voltage	V	—
T_m	thermistor on the case	°C	—

to drive through the intercalation kinetics, as shown in Equation (9.4).

$$\varphi = U_{neg}(z) - I R_{ct}(T,z), \qquad U_{neg}(z) = U_0 + U_A \mathrm{e}^{-k_U z} \tag{9.4}$$

where φ is in volts, U_0 is the plateau potential of lithiated graphite, U_A is the rise on an empty electrode and k_U sets how quickly that rise decays. Equation (9.4) returns 685 mV on an empty anode, 266 mV at 10 % and 85.4 mV at 60 %, so the margin available for the overpotential collapses within the first tenth of a charge. The kinetic resistance is Arrhenius in temperature and rises as the fourth power of lithiation, as shown in Equation (9.5).

$$R_{ct}(T,z) = R_{ct}^{0}\, a \left(1 + k_R z^4\right) \exp\left[E_c \left(\frac{1}{T_K} - \frac{1}{T_{ref}}\right)\right] \tag{9.5}$$

where E_c is the activation temperature of the anode kinetics and k_R sets how steeply the resistance grows as the electrode fills. Setting $\varphi = 0$ in (9.4) gives the plating current, and setting φ to any target margin gives the current a controller would allow, as shown in Equation (9.6).

$$I_{lim}(T,z) = \frac{U_{neg}(z)}{R_{ct}(T,z)}, \qquad I(\varphi^*) = \frac{U_{neg}(z) - \varphi^*}{R_{ct}(T,z)} \tag{9.6}$$

Equation (9.6) is the whole control problem in one line, because it is an algebraic inverse of the model rather than a measurement of anything, and its right-hand side is built from a state of charge

that was counted, a temperature read some distance from the reaction and a resistance taken from a table.

Current beyond I_{lim} deposits metal, and a fixed fraction of the deposit is treated as permanently lost, as shown in Equation (9.7).

$$\frac{\mathrm{d}Q_{pl}}{\mathrm{d}t} = \frac{k_{pl}}{60}\,(I - I_{lim})_{+} \tag{9.7}$$

where $(\cdot)_{+}$ is the positive part, k_{pl} is the fraction of the excess current that plates and the factor $1/60$ converts amperes to ampere hours per minute. Capacity fade accumulates from throughput with an Arrhenius temperature factor and from the irreversible part of the deposit, as shown in Equation (9.8).

$$\frac{\mathrm{d}F}{\mathrm{d}t} = \frac{k_{cyc} I}{60} \exp\left[E_f\left(\frac{1}{T_{ref}} - \frac{1}{T_K}\right)\right] + k_{pf} f_{irr} \frac{\mathrm{d}Q_{pl}}{\mathrm{d}t} \tag{9.8}$$

where F is in percent of nameplate capacity, k_{cyc} is the fade per ampere hour at 25 °C, E_f is the activation temperature of throughput aging, f_{irr} is the irreversible fraction of the plated lithium and k_{pf} is the fade per ampere hour of it. The two terms in (9.8) have very different characters. Throughput aging is slow and unavoidable, while the plating term can produce more damage in the last three minutes of one charge than the first term produces in the whole of it.

Coulomb counting closes the state of charge on the effective capacity, as shown in Equation (9.9).

$$\frac{\mathrm{d}z}{\mathrm{d}t} = \frac{I}{60\,Q\,(1 - F/100)}, \qquad \frac{\mathrm{d}Q_{th}}{\mathrm{d}t} = \frac{I}{60} \tag{9.9}$$

with Q in ampere hours and z dimensionless.

9.3.4 Thermal Side and the Chiller

Two thermal nodes are integrated, the pack and the coolant leaving the chiller, as shown in Equations (9.10) and (9.11).

$$\frac{\mathrm{d}T}{\mathrm{d}t} = \frac{60}{C_p}\left[\dot{Q}_{gen} - UA_{cool}\,\tfrac{q}{100}(T - T_c) - UA_{amb}(T - T_a)\right] \tag{9.10}$$

$$\frac{\mathrm{d}T_c}{\mathrm{d}t} = \frac{60}{C_c}\left[UA_{cool}\,\tfrac{q}{100}(T - T_c) - \dot{Q}_{ch} + UA_{cl}(T_a - T_c)\right] \tag{9.11}$$

where C_p and C_c are heat capacities in kJ/°C, the conductances are in kW/°C, q is the coolant flow in percent and the factor 60 converts kW/(kJ °C) to degrees per minute. Loss heat is ohmic plus polarization, as shown in Equation (9.12).

$$\dot{Q}_{gen} = \frac{1}{1000}\left(I^2 R_0 + \frac{v_1^2}{R_1} + \frac{v_2^2}{R_2}\right) \tag{9.12}$$

in kilowatts with the resistances in ohms. The chiller runs only while current is flowing and its capacity falls linearly with ambient temperature, as shown in Equation (9.13).

$$\dot{Q}_{ch} = \mathbb{1}[I_{cmd} > 2\,\mathrm{A}]\,\mathrm{clip}\big(k_{ch}(T_c - T_c^{set}),\ 0,\ \dot{Q}_{ch}^{max}\big) \tag{9.13}$$

with $\dot{Q}_{ch}^{max} = \max\big[0,\ \dot{Q}_0\,(1 - f_{ch}(T_a - 25))\big]$, where $\dot{Q}_0$ is the capacity at 25 °C and f_{ch} is the fractional loss per degree. The indicator function in (9.13) is why a parked pack sits at ambient rather than being quietly refrigerated, and it is what makes the resting scenario a genuine equilibrium.

> **DID YOU KNOW? The Heat of a Ten-Minute Charge**
>
> A 2017 thermal study of a 66 kWh pack on a 350 kW charger put the pack's heat generation at 39 to 116 kW, for charge efficiencies of 90 down to 70 %. It called a 15 kW cooling system substantially oversized compared with most vehicles then on the market [17].

9.3.5 Charger, Instruments and the Counted State of Charge

The post is a current source with a first-order response, as shown in Equation (9.14).

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \frac{I_{cmd} - I}{\tau_c} \tag{9.14}$$

with τ_c in minutes. Three signals reach the controller. The thermistor sits on the module case rather than in the jelly roll, so it reads low while the current is high, and it is filtered, as shown in Equation (9.15).

$$\begin{aligned} T_m^k &= \alpha_T T_m^{k-1} + (1-\alpha_T)\left(T - 6\,s\,\frac{I}{I_{max}}\right) \\ \alpha_T &= \mathrm{e}^{-\Delta t_c/\tau_{Tm}} \end{aligned} \tag{9.15}$$

where s is the offset scale, the constant 6 °C is the full-current gap at $s = 1$ and Δt_c is the controller period. The terminal voltage passes through the same exact zero-order-hold update with time constant τ_{Vm}. State of charge is counted from the current with whatever error the sensor carries, as shown in Equation (9.16).

$$\hat{z}_k = \hat{z}_{k-1} + \frac{b\,I\,\Delta t_c}{60\,Q} \tag{9.16}$$

where b is the current sensor error. Equation (9.16) uses the nameplate capacity rather than the faded capacity of (9.9), which is deliberate, because a vehicle counts against what it was told the pack holds.

Three different models are in play here, and confusing them is the usual error. The plant is Equations (9.1) to (9.14) with the true temperature T, the true state of charge z and the true aging multiplier a. The controller internal model is Equations (9.4), (9.5) and (9.6) evaluated at T_m, $\hat{z}$ and a table multiplier m in place of a, which is the map the management system actually carries. There is no identified control model, because nothing in the application fits a transfer function to data, and the gains in Section 9.5 are supplied rather than derived. The predictive strategy of Section 9.5.5 carries a second internal model, also written rather than identified, from the same equations at their declared values.

9.3.6 Parameters

Table 9.2 lists every constant the equations use. Three of them are hidden from the model panel and the rest are displayed and adjustable, which is a lighter concealment than the identification applications in this series use, because the problem here is structural rather than one of estimation. The post rating P_{max} caps the offered power P_{sp} of Table 9.1 in the charger candidate of Section 9.5, and at its 250 kW default it never binds on this pack.

9.3.7 Initialization and Numerics

A run opens with the state of charge and the pack temperature taken from two segmented controls, at 10, 30 or 55 % and at 0, 5, 15, 25, 35 or 38 °C. Both are read only at initialization, so a new arrival condition acts at once before a run and at the next Reset during one. Both polarization voltages, the current and the three counters open at zero, and the coolant starts at the pack temperature. The default arrival is 10 % and 25 °C. Four bounds are enforced after every step, with z clipped to 1.02, the current to I_{max}, the pack temperature to the range −30 to 90 °C and the fade to 40 %.

Table 9.2 Model parameters. Hidden constants are never displayed in the application.

Symbol	Meaning	Value	Unit	Visibility
n_s	cells in series	96	—	adjustable
Q	pack capacity	200	Ah	adjustable
R_0^0	ohmic resistance	0.0235	Ω	adjustable
R_1^0	charge-transfer branch	0.0090	Ω	adjustable
C_1	charge-transfer branch	2200	F	adjustable
R_2^0	diffusion branch	0.0055	Ω	adjustable
C_2	diffusion branch	42000	F	adjustable
E_R	activation, resistances	3600	K	adjustable
V_{max}	cell voltage limit	4.15	V	adjustable
U_0	graphite plateau	0.085	V	adjustable
U_A	rise on empty anode	0.60	V	adjustable
k_U	decay of that rise	12	—	adjustable
R_{ct}^0	kinetic resistance	1.2×10^{-4}	Ω	adjustable
k_R	lithiation coefficient	8.0	—	adjustable
E_c	activation, kinetics	4000	K	adjustable
k_{pl}	plated fraction of excess	0.55	—	adjustable
f_{irr}	irreversible fraction	0.62	—	adjustable
C_p	pack heat capacity	340	kJ/°C	adjustable
UA_{cool}	pack to coolant, full flow	1.55	kW/°C	adjustable
UA_{amb}	pack to cabin	0.09	kW/°C	adjustable
$\dot{Q}_0$	chiller at 25 °C	6.5	kW	adjustable
f_{ch}	capacity lost per degree	0.018	1/°C	adjustable
T_c^{set}	coldest coolant	15	°C	adjustable
T_{max}	cell temperature limit	45	°C	adjustable
k_{ch}	chiller control gain	1.4	kW/°C	hidden
C_c	coolant heat capacity	25	kJ/°C	hidden
UA_{cl}	coolant to ambient	0.05	kW/°C	hidden
I_{max}	vehicle current ceiling	500	A	adjustable
P_{max}	post power rating	250	kW	adjustable
τ_c	charger current response	0.02	min	adjustable
E_f	activation, throughput fade	5500	K	adjustable
k_{cyc}	fade per ampere hour	3.4×10^{-4}	%/Ah	adjustable
k_{pf}	fade per plated ampere hour	0.010	%/Ah	adjustable
τ_{Tm}	thermistor filter	0.25	min	adjustable
τ_{Vm}	voltage filter	0.02	min	adjustable

Each 0.005 min of simulated time, which is 0.3 s, is advanced by one classical Runge-Kutta step of fourth order, by the engine of Section 1.6.1. The management system executes on a 0.02 min period, which is 1.2 s and is a realistic scan for a pack controller, and the two sensor filters in (9.15) advance on that same period rather than on the physics step. The historian keeps one sample every 0.035 min, which is seven physics steps or 2.1 s. There is no stochastic term anywhere in this application, so every scenario is deterministic and the trends, the scorecard and the exported file are identical on any machine. Speed multipliers between 0.5 and 10 simulated minutes per wall second leave the trajectory untouched.

9.3.8 Assumptions

Six simplifications hold this model up, and the paragraph for each one names the condition that breaks it.

One equivalent cell, scaled to the pack. Series and parallel spread, balancing and the weakest cell are all absent. On a real pack the management system defends the weakest cell rather than the average one, and that cell is usually several percent away, so the margins reported here are optimistic by whatever the spread happens to be.

A lumped anode criterion rather than a porous electrode. The plateau potential and the lumped kinetic resistance of (9.4) and (9.5) reproduce the shapes that a distributed model gives [2, 11], which are a flat equilibrium region, an Arrhenius temperature dependence and a steep rise as the electrode fills, but there is no solid-phase diffusion, no exchange current density and no variation through the electrode thickness. The criterion therefore locates the onset of plating at one surface rather than at the separator face where a distributed model puts it first.

Plating is proportional to the overshoot with no threshold. Equation (9.7) starts depositing the instant φ crosses zero. Real deposition needs a nucleation overpotential, part of it strips back on rest, and the damage depends on morphology far more than on quantity [3]. Dendrite growth and separator penetration are outside the model entirely, so nothing here supports a statement about safety.

Two thermal nodes. There is no gradient inside a module, no cell-to-coolant film and no cold plate, and pack thermal design is largely about those three. The 6 °C core-to-case gap in (9.15) is a fitted offset proportional to current rather than a conduction path, so it responds instantly to a current change where a real gradient would lag.

Two aging terms. Calendar aging, depth of discharge, resting state of charge and the interactions among them are absent from (9.8), so a projection over five hundred charges illustrates a mechanism and does not predict a life. The ratios between strategies are the usable result and the absolute percentages are the arithmetic of this model.

The post is a current source with a ceiling. Real direct-current chargers run control loops of their own, negotiate continuously with the vehicle, derate on cable temperature and time out on communication faults. Here the post follows (9.14) with one time constant and a power limit.

9.4 The Constraint Surface

Before any controller is discussed it is worth reading the four limits as functions of the operating point, because the answer decides which of them is worth tuning. Figure 9.2 plots the anode-limited and voltage-limited currents against state of charge, with the two measured charge trajectories drawn over them.

Three readings follow from Figure 9.2. The anode limit is below the voltage limit everywhere on the working range, by a factor of 3.6 at half charge (480 A against 1734 A) and 5.8 at 80 % (166 A against 959 A), so a charge that respects the terminal voltage and nothing else is not protected at all. Among the four, the anode limit alone is a strong function of what the operator is doing, because it moves with both the state of charge the current is creating and the temperature the current is generating, while the charger rating and the cell voltage limit are properties of the equipment. Cold makes it worse by the amount (9.5) predicts, since the kinetic resistance at 0 °C is 3.41 times its value at 25 °C and the allowable current at half charge falls from 480 to 141 A.

The consequence for a charge curve is the shape drawn solid with circles. The charger ceiling holds the current for the first part of the charge, the anode limit descends through it, and from that crossing the current follows a continuous taper set by an electrode potential rather than a step change set by a voltage limit. The classical constant-voltage phase all but disappears on this pack once the anode is represented (on the charge to a hundred percent of scenario 10 the terminal voltage reaches its limit only at 98.4 % true charge and holds the current for the last 4.6 min), which is worth stating plainly, because the constant-voltage phase on a real charge curve is largely what a voltage loop does when it is standing in for a constraint that was never modeled.

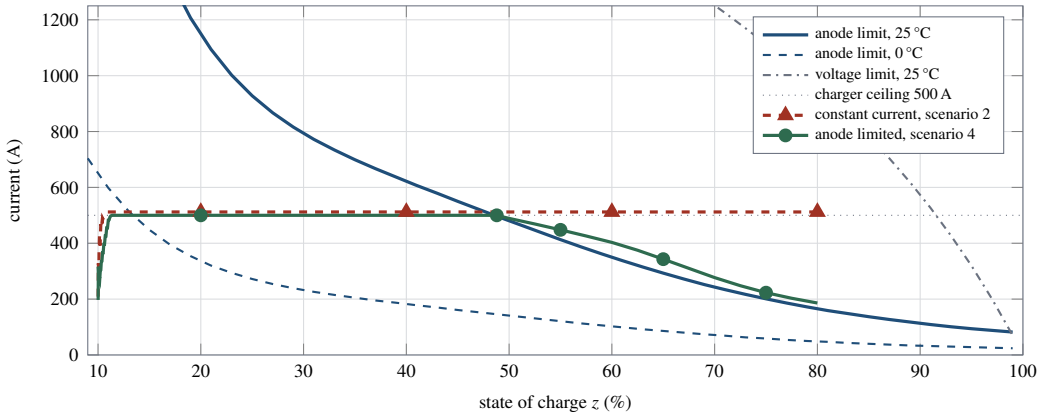

Figure 9.2 The four limits as functions of state of charge, with two measured charges drawn over them. At 25 °C the anode limit crosses the charger ceiling near 49 % and reaches 166 A at 80 %, while the voltage limit at the same temperature is still 959 A. The two isothermal curves both shift upward with the temperature a charge creates for itself. The measured traces are the current flowing at each state of charge. Both dip to 267 A in the first 0.1 min, which is the opening transient of Section 9.6.3, and are back at 500 A by 11.4 %. The constant-current trace is drawn 12 A high so that it separates from the anode-limited one, and it crosses the anode limit near 49 % and stays above it to the end. The anode limits are solid (25 °C) and dashed (0 °C), the voltage limit is dash-dotted and the charger ceiling dotted. The constant-current charge is dashed with triangles and the anode-limited charge solid with circles.

9.5 The Charge Controller

9.5.1 The Discrete Algorithm

One proportional-integral-derivative routine serves the voltage controller, the temperature controller and the coolant controller. It is the velocity-form routine of Section 1.6.2, with the derivative on the measurement and back-calculation on the integral alone. Equations (9.17) to (9.19) restate it in this chapter's symbols: γ is the action a of Chapter 1, Σ_k is the running sum S_k, u_k^u is the output before limiting v_k and u_{bias} is the bias u_b.

$$e_k = \gamma(\mathrm{SP}_k - \mathrm{PV}_k), \quad P_k = K_c e_k, \quad D_k = -K_c \tau_D \gamma \dot{\mathrm{PV}}_k \tag{9.17}$$

The direction flag γ is $+1$ when a rising measurement calls for a smaller output, which is reverse action in the ISA convention, and -1 when it calls for a larger one.

$$\begin{aligned} \Sigma_k &= \Sigma_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ u_k^u &= u_{bias} + \Sigma_k + I_{k-1} \\ u_k &= \mathrm{clip}\,(u_k^u,\ u_{min},\ u_{max}) \end{aligned} \tag{9.18}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - u_k^u}{T_t} \right] \Delta t_c \tag{9.19}$$

The tracking constant T_t equals τ_I, because no loop sets one of its own. All three loops are configured with $\tau_D = 0$, so the derivative path is inactive and (9.17) reduces to a proportional term. No loop carries a rate limit on its output.

Solving (9.18) for the integral at a specified output is the tracking operation, and it serves two purposes here. It supplies the operator actuator position in manual mode, and it is the external reset applied to whichever controllers lose the selector.

Table 9.3 Constraint controllers as implemented. All five are evaluated every 0.02 min, and derivative action is off in the three that have it.

Controller	Reads	Writes	K_c	τ_I (min)	Action
charger	$\hat{V}$, P_{sp}	I	—	—	algebraic
voltage	$\hat{V}$	I	60	0.15	reverse
temperature	T_m	I	120	2.0	reverse
anode	$\hat{z}$, T_m	I	—	—	inverse map
coolant	T_m	q	14	1.5	direct

9.5.2 Four Limits and One Handle

Table 9.3 lists the four candidate controllers, the coolant loop and their default settings. The signs deserve a sentence each. More current raises the terminal voltage and raises the temperature, so the voltage and temperature controllers face a positive process gain and are reverse acting. More coolant flow lowers the temperature, so the coolant loop faces a negative gain and is direct acting.

The charger controller is algebraic rather than a loop, because the post delivers whatever the vehicle asks for up to a current ceiling and a power that is the lower of the offered power and the post rating, as shown in Equation (9.20).

$$I_P = \min\left(I_{max},\ \frac{1000\ \min(P_{sp}, P_{max})}{\max(200, \hat{V})} \right) \tag{9.20}$$

with P_{sp} and P_{max} in kilowatts and $\hat{V}$ in volts. The voltage controller acts on the pack limit $n_s V_{max}$ rather than clamping the current directly, because the terminal voltage of (9.1) includes a resistive rise that moves with temperature and age, so the current that corresponds to a given cell voltage is not a constant. The temperature controller acts on the derate setpoint T_{derate} with a gain of 120 A/°C, which is a hard pullback rather than a trim, because its job is to protect a limit and not to hold a setpoint. Both are bounded to 0 to 900 A at the controller and the selected current is then clipped to the vehicle ceiling. The charger candidate never exceeds that ceiling, so a proportional-integral winner is always below it and the second clip never acts on an integral it does not know about.

9.5.3 The Anode Limit Is a Map, Not a Loop

The fourth candidate has no measurement, so it has no error, so it cannot be a controller in the sense the other three are. What the management system carries instead is the algebraic inverse of its own model, evaluated at the signals it does have and smoothed so that its output does not step, as shown in Equation (9.21).

$$\begin{aligned} I_A^{raw} &= \mathrm{clip}\left(\frac{U_{neg}(\hat{z}) - \varphi_{SP}}{\hat{R}_{ct}(T_m, \hat{z})},\ 0,\ 900\,\mathrm{A} \right) \\ I_A^k &= \alpha_A I_A^{k-1} + (1-\alpha_A) I_A^{raw}, \qquad \alpha_A = \mathrm{e}^{-\Delta t_c/\tau_A} \end{aligned} \tag{9.21}$$

where $\hat{R}_{ct}$ is Equation (9.5) evaluated with the table error m in place of the true aging multiplier a, and τ_A defaults to 0.12 min. A production system tabulates the first line of (9.21) against state of charge and temperature and interpolates it, which is the charge current limit map, and solving the expression each scan is the same thing without the interpolation error.

Two properties of (9.21) shape everything in Section 9.6. The first is that it has no integral action and cannot be given any. Over the part of the default charge where the map owns the selector the computed potential $\hat{\varphi}$ sits between 13.9 and 14.5 mV against a 15 mV setpoint, and the residual is the filter lagging a limit that is falling rather than an offset that tuning would remove. The second is that the accuracy of the output is the accuracy of $\hat{R}_{ct}$ and of the two signals feeding it, and nothing downstream reveals an error in either. The margin φ_{SP} is therefore not a safety factor in the usual

sense. It is an explicit budget for how wrong the table is allowed to be, and Section 9.6.4 measures the exchange rate between the two.

🕮 **Pause.** Try Exercise 9.1 before reading on.

9.5.4 The Selector and Its Preload

The applied command is the smallest of the candidates that the selected architecture includes, as shown in Equation (9.22).

$$I_{cmd} = \mathrm{clip}(\min\{I_P,\ I_V,\ I_T,\ I_A\},\ 0,\ I_{max}) \tag{9.22}$$

The three selector architectures the application offers include successively more of the set. Constant current and constant voltage uses I_P and I_V, the temperature variant adds I_T, and the anode-limited architecture adds I_A. The current is driven to zero at the scan on which the counted $\hat{z}$ of (9.16) reaches the target, so $b = 1.05$ ends scenario 4 at 76.7 % true charge, and times to target below are quoted at that scan.

Each proportional-integral candidate that does not own the selector is given an external reset, so both are reset while the charger or the map holds the current. The integral is solved from (9.18) so that its output equals the applied current plus a preload Δ, capped at the loser's own proportional term $P_{j,k}$, as shown in Equation (9.23) for each loser $j \in \{V, T\}$.

$$I_j \leftarrow I_{cmd} + \min(\Delta, P_{j,k}) - u_{bias} - \Sigma_{j,k} \tag{9.23}$$

with Δ defaulting to 25 A and adjustable from 0 to 120 A. Tracking a loser to the selected current alone, with $\Delta = 0$, parks it exactly on the boundary, where the change in its own proportional term across one scan is enough to make it the winner. What follows is not an oscillation, and Section 9.6.3 reports what it is instead.

The coolant loop is independent of the selector. It acts on the thermistor reading against a coolant setpoint that defaults to 32 °C, with an output of 0 to 100 % flow, and in manual it tracks the operator slider through the same mechanism as (9.23).

9.5.5 Model Predictive Control

A fourth strategy replaces the selector by one optimization over the charging current. Every $T_s = 0.1$ min, which is five controller scans, a quadratic program plans the current over $N = 100$ steps, or 10 min, and only the first move of the plan is applied, held until the following decision, as in the predictive option of Section 1.6.3 [18, 19]. The coolant loop stays outside the plan. The controller's model has five states, the counted state of charge, the two branch voltages of (9.1), the core temperature and the thermistor reading, written from the equations of Section 9.3 with the declared values of Table 9.2 and $a = 1$, as shown in Equation (9.24).

$$\begin{aligned} \frac{\mathrm{d}z}{\mathrm{d}t} &= \frac{I}{60\,Q}, \qquad G(q) = UA_{cool}\,\tfrac{q}{100} + UA_{amb} \\ \frac{\mathrm{d}T}{\mathrm{d}t} &= \frac{60}{C_p}\left[\dot{Q}_{gen} - G(q)\,(T - T_s)\right] \\ \frac{\mathrm{d}T_m}{\mathrm{d}t} &= \frac{1}{\tau_{Tm}}\left(T - \frac{6\,I}{I_{max}} - T_m\right) \end{aligned} \tag{9.24}$$

where T_s is one heat-sink temperature that stands in for the coolant, the chiller and the cabin, q is the flow the coolant loop commands at the decision, held over the horizon, and the thermistor offset is at its design scale $s = 1$. The model is written rather than identified by step tests, which the original design called for with three states, because every gain of this plant moves with its state through the Arrhenius resistances, the z^4 kinetic term and the steep top of (9.2), while the state of charge can

move 42 points over one horizon at 500 A. It is therefore relinearized at every decision along the previous plan shifted by one step, sampled exactly with a zero-order hold, with an affine term per step that makes the linear model reproduce the nonlinear one along that plan. Fed the current actually applied in anode-limited runs of five scenarios from six starting points, the model is within 0.14 V, 0.15 °C and 0.33 mV of the logged voltage, core temperature and anode estimate one minute ahead.

The predicted outputs are the voltage of (9.1) plus a bias b_V, the core temperature T and the anode headroom $I - \hat{I}_A$, where $\hat{I}_A$ is the current at which the vehicle's own estimate would sit on its setpoint, as shown in Equation (9.25).

$$\hat{I}_A(z, T_m) = \frac{U_{neg}(z) + b_U - \varphi_{SP}}{k_r R_{ct}(T_m, z)}, \qquad k_r = \frac{\hat{R}_{ct}}{R_{ct}(T_m, \hat{z})} \tag{9.25}$$

Here $\hat{R}_{ct}$ is the believed resistance of (9.21) at the decision, R_{ct} is (9.5) at the declared values, and b_U is the difference between the equilibrium potential inside the estimate and the declared one. The anode reaches the controller only through the vehicle's estimate, so a table error m or an aged pack misleads it exactly as it misleads the map, and an edited anode parameter reaches it through that estimate in the same way. An edit of any other plant parameter never reaches its model and becomes a mismatch for the estimator. A steady-state Kalman filter on the thermistor estimates T, T_m and T_s at every scan under every strategy, with T_s a random walk that makes the temperature row offset free [20]. Its inputs are the heat of (9.12) from the measured current and the offset of the thermistor, its variances per scan are 1×10^{-4} and 9×10^{-6} °C^2 on the core and the sink against 2.5×10^{-3} °C^2 on the reading, and its gain, which depends on the flow, is computed at every 10 % of flow and interpolated. The voltage bias b_V is the filtered voltage less the model's at each decision.

The decisions are moves of the current from the value applied at the decision, scaled by $I_{ref} = 500$ A, over twelve blocks whose lengths grow from a single step for the first four to 55 steps for the last, and the plan minimizes Equation (9.26).

$$\begin{aligned} J = & -c_I \sum_{k=0}^{N-1} \frac{I_k}{I_{ref}} + S \sum_{j=1}^{12} \left(\frac{\Delta I_j}{I_{ref}} \right)^2 \\ & + \sum_{r \in \{V,T,A\}} \left(\ell_r \sigma_r + \rho_r \sigma_r^2 \right) \end{aligned} \tag{9.26}$$

where $c_I = 1$ prices the charge delivered, ΔI_j is the move of block j, $S = 1$ is the move suppression, and σ_V in volts, σ_T in degrees and σ_A in amperes are one slack each on three soft rows, $V_k \le n_s V_{max} + \sigma_V$, $T_k \le T_{derate} + \sigma_T$ on the predicted core and $I_k \le \hat{I}_{A,k} + \sigma_A$. The prices are $\ell_V = 200$, $\rho_V = \ell_T = \rho_T = 1000$, $\ell_A = 100$ and $\rho_A = 10$, and while charging every logged slack stays below 10^{-15} in all seven cases of Section 9.6.7. The charger row is hard, $0 \le I_k \le \min(I_P, I_{max}, 900\,\text{A})$ with I_P of (9.20) read as the selector reads it, and the held move is capped by I_P again at every scan. A dual active-set solver [21] handles the twelve decisions. The move suppression and the sink variance were chosen once, on scenario 12, as the largest move suppression within 1 % of the fastest charge (41.31 min) among nine settings that held every limit and moved the current no more than the selector.

A plan is rejected when the solver does not return it as optimal, when it is not finite or when it leaves the charger bounds, and the anode-limited selector then takes the current until a later plan succeeds. While the plan acts, the voltage and temperature controllers track the applied current with $\Delta = 0$ in (9.23), so a hand-over moves the current only by the winning controller's own one-scan move, 1.9 A at 15 min in scenario 12 and at most 2.9 A at twelve forced instants on six cases, after which the preload applies again. The plant runs with the parameters of the run, the instruments are those of (9.15) and (9.16), and the controller reads only the counted state of charge, the applied and measured current, the filtered voltage, the thermistor, the anode estimate and its believed resistance, the commanded flow, the charger limit and the three setpoints. In Chrome a plan took a mean of 503 to 516 µs in three runs of scenario 12, at most 10.7 ms, which is about 0.5 % of one core in real time on a desktop with 20 cores whose load average was 1.5 to 1.8.

Table 9.4 The thirteen scenarios and what each isolates.

#	Scenario	Kind	What it isolates
1	The pack at rest	Experiment	four limits at one operating point
2	Constant current, constant voltage	Demonstration	a fast charge with every gauge inside its limit
3	What that charge cost	Experiment	where in a charge the damage occurs
4	Riding the limit	Experiment	the time cost of the anode-limited taper
5	Which limit is holding the current	Experiment	selector dwell and the preload
6	A cold pack	Experiment	cold kinetics and self-heating
7	A hot afternoon	Experiment	cooling that costs time and buys life
8	The model is optimistic	Demonstration	holding the estimate and plating anyway
9	An old pack and a new pack's model	Demonstration	aging the model does not know about
10	All the way to a hundred	Experiment	the tail the anode limit creates
11	A shared charger	Experiment	relieving a constraint that is not binding
12	Four limits, one QP	Experiment	one plan against four limits
13	Challenge: ten minutes at the services	Challenge	cold, aged, optimistic and short of power

9.6 Scenarios and Measured Results

The scenario sequence of Table 9.4 meets each constraint before it meets the controller that defends it. Every number reported below comes from the application running its own deterministic model at the settings the scenario loads, with any setting a sentence names changed before the run starts, which gives the same run as carrying it through Reset. Each scenario is credited separately on its own completion rule (Section 1.6.4), when the learner performs the actions its briefing lists and the simulation reaches the stated state. The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt.

9.6.1 What the Old Protocol Costs

Scenario 2 is a plain constant-current charge from 10 to 80 % at 25 °C, and the first thing to record is how well it appears to go. The charger holds 500 A for all but the first 0.21 min, when the voltage controller takes the selector on five alternate scans and the command dips to 238 A (the opening transient of Section 9.6.3). The 250 kW rating never binds and the power peaks at 194 kW. The highest cell voltage reached is 4.034 V against a 4.15 V limit, so the constant-voltage phase the protocol is named for never begins. The pack peaks at 38.6 °C against a 45 °C limit. Eighty percent arrives in 16.84 min with 51.6 kWh delivered, and every instrument in the vehicle is inside its limit for the whole stop.

> **TRY IT IN THE APP** **2 · Constant Current, Constant Voltage**
>
> Load scenario 2 and press Run. Watch the dotted true line on the Anode potential strip. It crosses zero at 13.57 min and 66.4 % state of charge and reaches −61.2 mV, while the pack voltage and temperature stay inside their limits. Eighty percent arrives at 16.84 min with 3.381 Ah plated.

The anode surface crosses zero at 13.57 min and 66.4 % state of charge, reaches −61.2 mV, and stays negative to the end. 3.381 Ah of lithium deposits as metal and the pack loses 0.1035 % of its capacity on one stop. Table 9.5 puts that charge against the anode-limited one in three arrival conditions.

The room-temperature row is the commercial argument about fast charging expressed in four numbers. Riding the anode limit costs 4.50 min, which is what a driver notices, and it saves 3.381 Ah of plated lithium and a quarter of the capacity fade, which is what nobody notices. The cold row is

Table 9.5 Constant-current against anode-limited charging, 10 to 80 %, measured at the moment the target is reached.

Arrival	constant current		anode limited	
	time (min)	Q_{pl} (Ah)	time (min)	Q_{pl} (Ah)
25 °C pack, 22 °C air	16.84	3.381	21.34	0.000
38 °C pack, 42 °C air	16.84	2.204	19.66	0.000
0 °C pack, 0 °C air	16.92	13.867	44.80	0.000
capacity fade on the same charges (%)				
25 °C		0.1035		0.0779
38 °C		0.1282		0.1123
0 °C		0.1232		0.0206

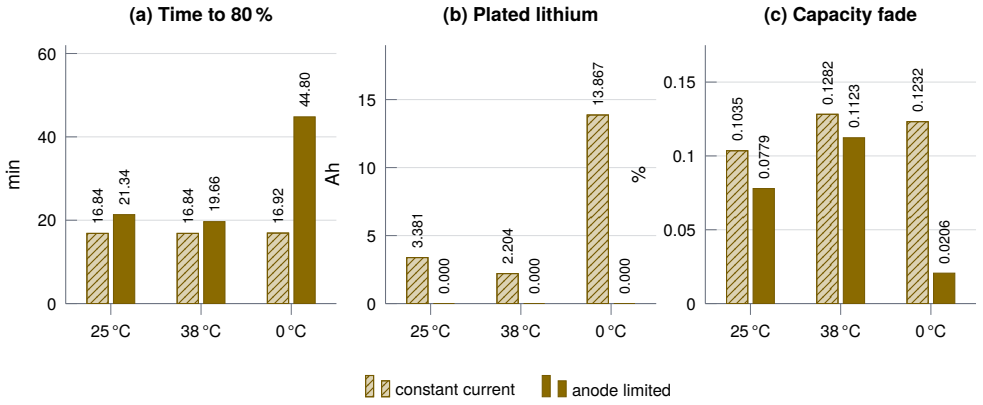

Figure 9.3 Constant-current against anode-limited charging from 10 to 80 % for the three arrivals of Table 9.5, labeled by pack temperature: (a) time to 80 %, (b) plated lithium and (c) capacity fade on the same charges. Hatched: constant current. Solid: anode limited. Notice that the constant-current charge takes nearly the same time from every arrival, while its plated lithium is largest from the cold pack, which is where the anode-limited charge pays most in time.

the one to carry away. A constant-current charge at 0 °C finishes in the usual 16.92 min and destroys 13.867 Ah of lithium inventory doing it, four times the room-temperature figure, and it looks identical on every instrument the vehicle has. The hot row runs the other way and is discussed in Section 9.6.5. Figure 9.3 draws the same table as bars.

Scenario 3 locates the damage in time by switching the map into the selector partway through the same charge, at a scripted 9 min or by hand after the scripted switch has been undone. Switching at 5, 9 or 12 min plates nothing, because scenario 2 does not cross zero until 13.57 min. Switching at 14 min costs 0.066 Ah and switching at 16 min, one minute before the end, costs 1.940 Ah, which is 57 % of the whole 3.381 Ah. The overshoot compounds because the current stays at 500 A while the true limit collapses from 524 A at 65 % to 436 at 70 and 291 A at 80 %, so nearly all the harm on a fast charge is done in its last three minutes (3.354 Ah of the 3.381 Ah, and 1.687 Ah in the final minute alone).

9.6.2 Which Constraint Is Load Bearing

Scenario 5 runs the anode-limited architecture to 95 % and reports which candidate owns the selector at every scan. The result is worth more than any tuning exercise on this plant. Of the 34.38 min

Table 9.6 Selector preload Δ swept on the scenario 4 charge to 80 %, set before the run, with a 60 min cutoff. Changes count transitions between scans.

Δ (A)	0	5	15	25	60	120
selection changes	2999	139	33	19	9	5
changes after 1 min	2950	90	1	1	1	1
held by I_V or I_T (min)	59.98	2.72	0.60	0.32	0.12	0.06
time to 80 % (min)	—	21.92	21.40	21.34	21.28	21.26
largest step after 1 min (A)	1.5	6.9	0.8	0.8	0.8	0.8
plated Q_{pl} (Ah)	0.000	0.000	0.000	0.000	0.000	0.000

charge, the charger ceiling holds the current for 9.08 min and hands over at 9.41 min and 48.8 % state of charge, the anode map holds it for 24.98 min, and the voltage and temperature controllers between them hold it for 0.32 min, all of that inside the first 0.37 min while the terminal voltage rises with the onset of current. The 250 kW power rating is never selected, and the current that would put the pack on its cell voltage limit is 1245 A at 80 % and 367 A at 95 %, both far above what the anode allows. Two of the four constraints are decorative on this pack, and the honest lesson of a four-way selector is that identifying the load-bearing constraint is worth more than tuning any of them.

Scenario 11 prices that observation. Halving the post to 120 kW between 3 and 14 min holds the current between 338 and 326 A instead of 500 A and pushes eighty percent from 21.34 to 25.32 min, a loss of 3.98 min. Firing the identical outage in scenario 4 at 14 min, in the second half of the charge, costs 0.02 min, since 21.36 against 21.34 min is the whole effect. In scenario 11 itself a second cut fired at 15 min, after the power has returned, binds only until 18.24 min and costs 0.34 min, and one fired at 18 min or later costs nothing. Relieving a constraint that is not binding is worth nothing, and a great deal of capital is spent on the wrong one.

Scenario 10 shows where the taper ends. Eighty percent arrives at 21.325 min, ninety at 29.1 min, ninety-five at 34.3 min and ninety-nine at 39.5 min on the true state of charge. The true anode limit falls from 246 to 115 A by 38.5 min and the current follows it from 185 down to 91 A. There, at 98.4 %, the voltage controller takes the current until it has tapered to $Q/20$ (10 A), where the charger ends the charge at 43.15 min and 99.75 %. The last twenty points take 21.8 min against 21.3 min for the first seventy, which is why charging networks quote ten to eighty and route planners stop there.

9.6.3 The Preload That Keeps the Selector Still

Table 9.6 sweeps the preload of (9.23) on the scenario 4 charge with everything else held.

With no preload the selection changes 2999 times in the 60 min the run is allowed, alternating between the voltage and temperature controllers at every scan, and the charge never finishes, reaching only 28.0 % by 40 min and 30.0 % by 60 min. The failure mode is not oscillation, and after the first minute the largest single-scan change in the command is only 1.5 A. Because every new winner starts from the value the previous one was tracking, the current ratchets steadily downward, from 212 A at one minute to 65 A at ten and 15 A at forty. A selector with no preload looks like a charger that is quietly giving up, which is a much harder fault to recognize on a trend than a limit cycle would be.

The remedy is cheap because the cap in (9.23) keeps the preload from standing between a controller and its own limit. Uncapped, a loser takes the current only if its proportional term moves by more than Δ in one scan (0.42 V in 1.2 s for the voltage controller), and constant current and constant voltage to a hundred percent ran 11.48 V over the pack limit, while capped it stays within 0.53 V. Twenty-five amperes costs 0.08 min on a 21.3 min charge against the largest value swept, and it reduces the selection count from 2999 to 19, with 18 of those 19 confined to the first 0.37 min. The residual activity in that window is an opening transient that no preload on the slider removes. The controller execution made when the run is initialized tracks both losers to the selected 500 A plus the preload, and in the first scan after the run starts the measured terminal voltage rises 4.8 V as the current comes on through R_0, which lowers the voltage controller’s proportional term by 288 A

Table 9.7 Kinetic table error swept at a 15 mV margin, scenario 4 conditions, 25 °C arrival.

table error m	1.00	0.90	0.80	0.70	0.60
lowest true φ (mV)	20.9	14.4	6.3	−4.1	−16.4
plated Q_{pl} (Ah)	0.000	0.000	0.000	0.187	1.628
time to 80 % (min)	21.34	20.22	19.18	18.26	17.58
fade F (%)	0.0779	0.0787	0.0797	0.0820	0.0919

at $K_c = 60\,\mathrm{A/V}$. The command dips to 238 A, which is 213 A plus the preload. The new winner's integral then climbs by about 470 A a scan against a 59 V error while the loser's proportional term falls again, and the two alternate until the voltage has settled. Raising the preload to 60 A shortens the alternation to 9 changes and 0.12 min, and 120 A to 5 changes and 0.06 min, but the first scan still dips to 273 and 333 A. The preload has to exceed the change in a losing controller's proportional term over one scan. On this plant 15 A is enough once the charge is under way, while the first scans of a charge would need more than the 120 A the slider allows. The app's design retains every constraint controller from the first scan and the shared PID tracking law. The opening dip remains a disclosed controller-startup transient rather than being hidden by temporarily gating the voltage or temperature constraints.

Pause. Try Exercise 9.2 before reading on.

9.6.4 Holding the Estimate Perfectly and Plating Anyway

Scenario 8 makes the kinetic table 30 % optimistic, which is an ordinary amount for a table fitted to a different batch three years earlier. The computed potential stays between 13.4 and 14.2 mV while the map owns the selector, so it does exactly what it was asked to do. The true anode potential reaches −12.2 mV and the pack plates 2.237 Ah. Nothing in the vehicle distinguishes this charge from a good one.

> **TRY IT IN THE APP** **8 · The Model Is Optimistic**
>
> Load scenario 8 and press Run. The kinetic table is 30 % optimistic. While the map holds the current, the computed potential stays between 13.4 and 14.2 mV. The true potential reaches −12.2 mV, and the pack plates 2.237 Ah with no warning on any instrument.

Table 9.7 converts the margin into a budget by sweeping the table error at the default 15 mV setpoint on the scenario 4 charge.

A 15 mV margin buys tolerance to a table roughly 20 % optimistic and no more. That is the sentence a design review needs, and it is a different sentence from the one a safety factor produces, because it names the model error the number is paying for. The exchange rate runs the other way as well. Under the challenge conditions, which stack a cold aged pack on a table 15 % optimistic, the suggested anode-limited configuration with a 5 mV margin plates 5.598 Ah and reaches eighty percent in 25.70 min, 15 mV plates 2.484 Ah in 27.66 min, 25 mV plates 0.024 Ah in 30.66 min and 35 mV plates nothing in 36.52 min, with the plated totals read at the 45 min judgment. Buying the last 0.024 Ah therefore costs 5.86 min, and choosing where on that curve to sit is the design decision the challenge is built around.

Scenario 9 reaches the same failure by a different road. Every resistance is 60 % higher after five years while the table is still the one calibrated when the pack was new, so the charge takes a nearly unchanged 21.38 min, runs 1.7 °C hotter and plates 3.412 Ah, which is as much as the reckless charge in scenario 2 from a strategy that plated nothing when the pack was new. Covering it with margin alone is expensive, because 30 mV brings the deposit to zero at 24.84 min, 45 mV takes 31.28 min and 60 mV takes 46.94 min. A margin that has to cover an unknown amount of aging has to cover

the worst of it on every pack all the time, which is the case for identifying the resistance on line [1] rather than widening the budget.

One result in this group is genuinely uncomfortable. The thermistor reads low while the current is high, which is an error, and it is a protective one, because the map then believes the pack is colder than it is, computes a higher kinetic resistance and allows less current. Holding the table error at 30 % and sweeping the offset scale s from 0 to 2 takes the plated lithium from 3.422 to 0.000 Ah, passing through 2.237 Ah at the scenario default of $s = 0.3$ and 0.187 Ah at the nominal $s = 1$. Fitting a better thermistor to this vehicle would make the pack worse. Two model errors of opposite sign were canceling, which is the argument for validating an inferred constraint end to end rather than instrument by instrument.

9.6.5 Cooling That Costs Time

Scenario 7 arrives at 38 °C into 42 °C air, where the chiller has lost 31 % of its capacity by (9.13). Running the coolant loop in service against holding the flow at 10 % gives 19.66 against 18.96 min to eighty percent, a peak of 40.5 against 42.7 °C, and a fade of 0.1123 against 0.1237 %. The peaks are the scorecard's, which records the hottest core temperature at any physics step since Reset. The cooled charge is the slower one, by 42 s, and that is the correct direction. A hot anode has faster kinetics by (9.5), so its plating limit is higher and the selector passes more current.

What the cooling buys is 2.2 °C of peak temperature and 9.2 % of the fade on an aging law that is Arrhenius. Judging a cooling policy on minutes gets the sign of the answer wrong, and the scorecard reports peak temperature and fade beside the time for that reason. The same asymmetry runs through the whole application. Speed is visible today and the price is paid in three years, and the controller is where somebody has to decide between them on the driver's behalf.

🕮 **Pause.** Try Exercise 9.3 before reading on.

9.6.6 The Cold Arrival and the Graded Challenge

Scenario 6 arrives at 0 °C. The anode-limited charge takes 44.80 min against 21.34 min from a warm pack, which makes the arrival temperature the single largest influence on the length of the stop. The current climbs from 216 A at half a minute to a peak of 414 A at 1.5 min and then falls, to 372 A at two minutes, 303 A at five and 261 A at ten. The climb is not the pack warming itself. Until 1.56 min the selector changes hands 76 times, almost all of them between the voltage and temperature controllers as in the opening transient of Section 9.6.3, and the climb is their recovery from the opening dip. The true plating limit and the map both fall from the first scan, because the equilibrium potential of (9.4) drops from 266 to 139 mV between 10 and 20 % state of charge faster than self-heating lowers R_{ct}, and the pack has reached only 11.0 °C by ten minutes. With the arrival temperature set to 15 °C and the run restarted with Reset, the same charge in the same 0 °C air takes 27.40 min. This is the quantitative case for preconditioning, because heat put in from the drivetrain on the way to the charger is cheaper than heat bought in charging time.

> ✦ **DID YOU KNOW? A Cold Pack Looks Exactly Like a Warm One**
>
> From the simulator: at a 0 °C arrival the constant-current charge reaches eighty percent in 16.92 min and plates 13.867 Ah, four times the room-temperature figure, while the anode-limited charge plates nothing and takes 44.80 min (Table 9.5).

The graded challenge stacks the failures. The pack arrives at 5 °C into 6 °C air, aged 30 %, on a table 15 % optimistic, with the post halving between 6 and 16 min, and it opens in the constant-current architecture with a 5 mV margin, the detuned settings, zero preload and a fixed 20 % coolant flow. As loaded it reaches eighty percent in 21.64 min and has plated 14.196 Ah by the 45 min judgment, 284 times the limit, and it projects a 62.6 % loss, although it peaks at only 28.0 °C. The target is eighty percent with less than 0.05 Ah plated, a peak under 45 °C and a projected loss under 20 % over five

Table 9.8 Anode-Limited Selector Against the Predictive Controller

Case	Structure	Time to target (min)	Hottest core °C (min)	Core over T_{derate} (min)	Largest $V - n_s V_{max}$ V (min)	Over by 0.1 V (min)	Lowest $\hat{\varphi} - \varphi_{SP}$ (mV)	Plated (Ah)	Current travel (A)
4	selector	21.345	37.40 (12.53)	0.00	−17.45 (21.32)	0.00	−1.14	0.000	796
	predictive	21.525	37.35 (12.53)	0.00	−17.49 (21.49)	0.00	−0.11	0.000	802
6	selector	44.805	15.70 (26.81)	0.00	−17.64 (44.80)	0.00	−8.04	0.000	1728
	predictive	44.585	15.94 (26.42)	0.00	−17.68 (44.56)	0.00	−0.06	0.000	872
7	selector	19.665	40.53 (12.39)	6.58	−17.40 (19.64)	0.00	−1.28	0.000	650
	predictive	20.405	40.03 (10.36)	3.82	−17.47 (20.37)	0.00	−0.11	0.000	835
9	selector	21.385	39.11 (9.84)	0.00	−14.21 (21.35)	0.00	−1.40	3.412	1015
	predictive	21.445	39.11 (9.56)	0.00	−14.30 (21.42)	0.00	−0.13	2.997	802
10	selector	43.145	37.40 (12.53)	0.00	0.10 (38.92)	0.00	−1.14	0.000	796
	predictive	43.605	37.35 (12.53)	0.00	0.01 (39.20)	0.00	−0.11	0.000	802
12	selector	38.405	38.71 (12.39)	12.81	0.09 (34.44)	0.00	−1.22	0.000	1072
	predictive	41.625	35.95 (6.37)	0.00	0.01 (37.10)	0.00	−0.10	0.000	957
13	selector	30.665	26.52 (23.97)	0.00	−15.84 (30.66)	0.00	−5.88	0.024	2103
	predictive	30.065	26.90 (23.24)	0.00	−15.91 (30.03)	0.00	0.01	0.002	955

Cases are the scenarios of Table 9.4 (4 riding the limit, 6 a cold pack, 7 a hot afternoon, 9 an old pack, 10 to a hundred, 12 four limits, one QP, and 13 the challenge at its suggested settings). The selector is the anode-limited selector. The derate setpoint T_{derate} is 36 °C in scenario 12 and 40 °C elsewhere. Time to target is the scorecard's, latched when the vehicle's count reaches the target or, in scenario 10, when the current has tapered to $Q/20$ on the voltage limit. The hottest core with its time, the minutes over T_{derate}, the largest voltage excess with its time, the minutes more than 0.1 V over and the lowest estimate margin are taken while charging, before the target is reached. Current travel is the scorecard's total movement of the charging current over the whole run.

hundred charges, and the answer is mostly structural. The map has to be in the selector, the margin has to be widened to 25 mV or beyond, and the coolant loop has to be put back on temperature, which on a pack this cold holds the flow at zero and keeps the heat in the cells. With the suggested settings (anode limited, 25 mV, the tuned gains and preload, the coolant loop on temperature) the charge takes 30.66 min, plates 0.024 Ah, peaks at 26.5 °C and projects an 18.9 % loss. It passes with 0.026 Ah and 1.1 percentage points to spare, and revealing the suggestion before the run and carrying it through Reset give the same run. Two tuning values matter as well. With the structure and the margin changed but the 1.4 min map filter and the zero preload kept, the charge stands at 79.05 % at 45 min with 0.186 Ah plated, while restoring those two passes even with the detuned proportional-integral gains, in 29.78 min with 0.021 Ah. The suggested settings with the shipped 20 % coolant flow plate 0.169 Ah and fail. What remains after that cannot be tuned away, because a cold pack is slow and the first five minutes are spent whatever anybody does.

9.6.7 Comparison With the Predictive Controller

Table 9.8 sets the anode-limited selector against the predictive controller of Section 9.5.5 on seven cases, each run for 60 min (the challenge for 45 min at its suggested settings) with the same events, instruments and coolant loop, and with the predictive settings of scenario 12 on every case. Neither receives a preview, which the application does not have, and the predictive controller solved 600 plans per hour with none rejected.

Scenario 12 was built so that every limit holds the predictive controller's current for a stretch. The pack arrives at 30 °C in 38 °C air with a 95 % target, a cell limit of 4.08 V (391.7 V for the pack), a derate setpoint of 36 °C and the post cut to 120 kW from 4 to 9 min. Under the predictive controller the charger holds the current from 0.32 to 5.81 min, the core temperature to 19.01, the anode estimate to 37.03 and the voltage to the end, and the target arrives at 41.625 min. The selector arrives at 38.405 min, but its core reaches 38.71 °C at 12.39 min and spends 12.81 min above 36 °C while its thermistor stays at least 2.58 °C below it, as shown in Figure 9.4. Its voltage controller holds the pack from 33.98 min, at most 0.09 V over the limit. The predictive controller holds its core at 35.85 °C with 320 A flowing at 12.4 min and rides the voltage limit only from 37.03 min, exceeding it by 0.01 V at most.

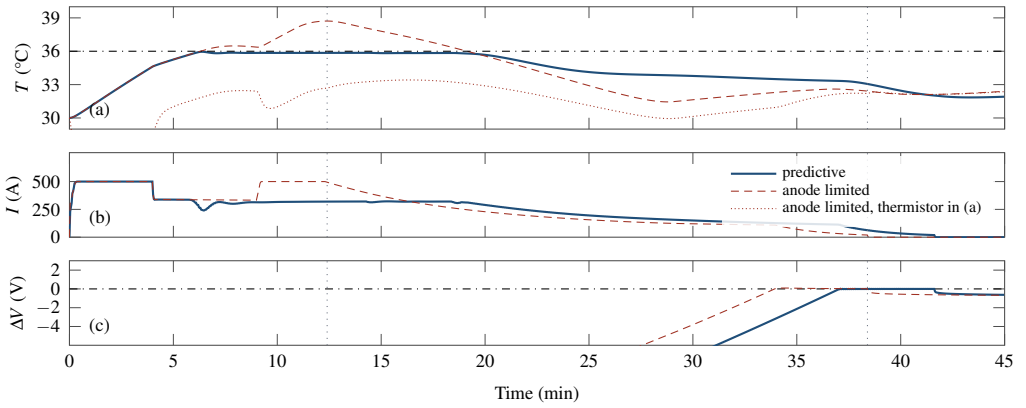

Figure 9.4 Scenario 12 under the predictive controller and under the anode-limited selector, from the same arrival and the same post cut from 4 to 9 min. The predictive controller is solid, the anode-limited selector densely dashed, its thermistor densely dotted, and the setpoint and zero lines dash-dotted. The vertical dotted lines mark 12.4 min, where the selector's core is at 38.7 °C and the predictive controller holds 35.9 °C, and 38.4 min, where the selector reaches 95 %. (a) Core temperature, the selector's thermistor and the 36 °C derate setpoint. (b) Current flowing. (c) Pack voltage less its limit, $\Delta V = V - n_s V_{max}$.

The selector passes the temperature limit because its temperature controller reads the thermistor, which sits below the core by the offset of (9.15). Relaxing the predictive controller's core row with a derate setpoint of 55 °C brings it to 38.79 min, so 0.38 min of its 3.22 min remains once it is allowed what the selector took.

The hot afternoon of scenario 7 is a different kind of loss. There the selector holds every limit on the instrument it reads, with its thermistor at least 2.00 °C under the 40 °C setpoint and its voltage 17.40 V under its limit, while its core, which no controller reads, reaches 40.53 °C at 12.39 min and spends 6.58 min of the charge above 40 °C. The predictive controller applies the setpoint to its estimate of the core instead, its core exceeds the setpoint by at most 0.03 °C, and it is 0.74 min slower. With that row relaxed it takes 19.81 min, so 0.60 min of the difference is the stricter reading of the setpoint and 0.14 min the remainder. Its peak temperature on the scorecard, 40.24 °C, is set at 60 min, after the charge, as the idle pack warms toward the 42 °C air with the chiller off. On the plain charge of scenario 4 neither structure reaches its voltage or temperature limit and the predictive controller is 0.18 min slower, of which 0.10 min is the move price, since a move suppression of 0.1 gives 21.425 min. From the cold arrival it is 0.22 min faster with half the current travel, and on the aged pack both plate (2.997 against 3.412 Ah) because both read the same new-pack estimate.

Charging to a hundred percent brings the two together. Both hold the voltage limit from 38.5 min (0.10 V and 0.01 V over at most), and since the open-circuit voltage of a full cell is the 4.15 V limit itself, both end at $Q/20$, at 43.145 min and 43.605 min. With the challenge's suggested settings the predictive controller reaches eighty percent at 30.065 min against 30.665 min and plates 0.002 Ah against 0.024 Ah. The predictive controller is therefore not a faster charger on this pack. In scenario 12 most of its extra time is a limit the selector does not keep on the quantity it protects, while on the hot afternoon the minutes are its stricter reading of the derate setpoint and on the plain charge they are its price on moves. Figure 9.5 shows the application during the anode-limited charge of scenario 4.

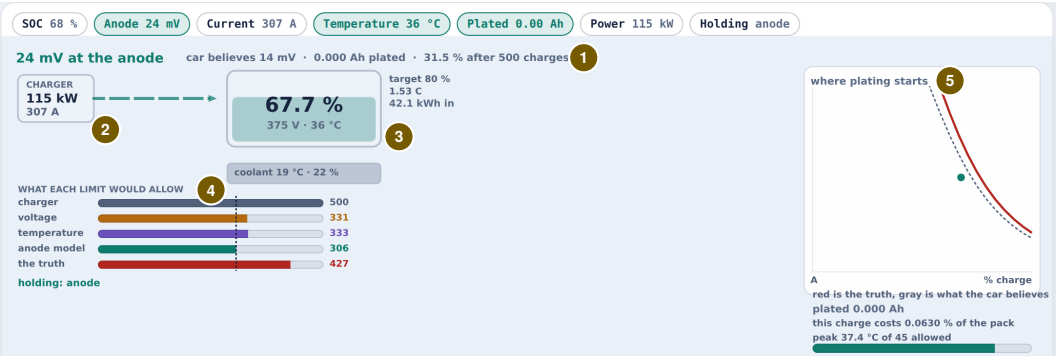

Figure 9.5 Anode-limited charging in scenario 4 at 15 min, some way past the handover from the charger ceiling, cropped to the plant scene. (1) The anode line: the true potential is 24.0 mV while the computed one is 13.9 mV, and nothing has plated. (2) The charger, with 307 A flowing. (3) The pack at 67.7 % state of charge, with the coolant loop beneath it. (4) The four candidate currents on one axis, 500 A from the charger, 331 on voltage, 333 on temperature and 306 from the anode model, with the true anode limit of 427 A drawn beside them and the anode named as the one holding. (5) The plating map with its believed and true curves, and beneath it the plated total of 0.000 A h.

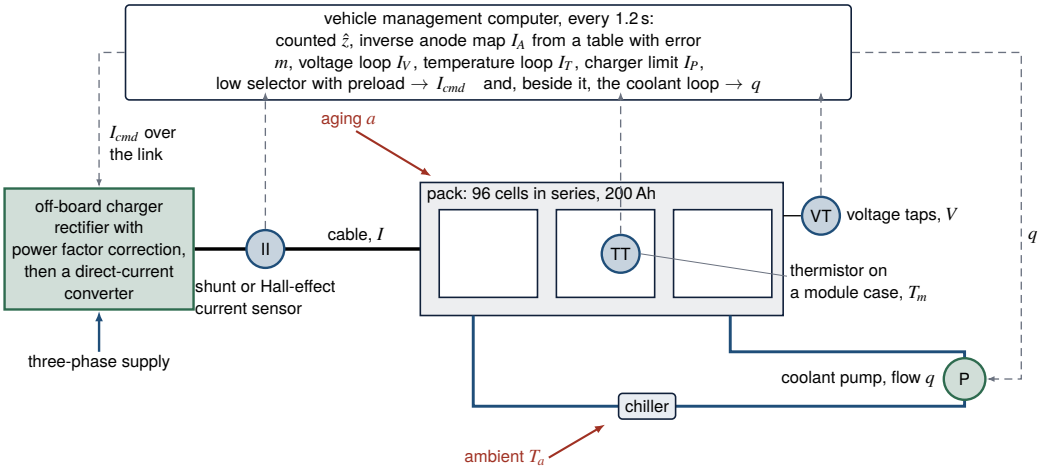

Figure 9.6 The chapter's loops drawn with their hardware. The current sensor II on the cable, the voltage taps VT and one thermistor TT on a module case are all the management computer reads. Every 1.2 s it computes four candidate currents, sends the lowest to the off-board charger over the communication link, and sets the coolant pump. Ambient temperature T_a acts on the chiller and cell aging a on the pack, and the table error m sits inside the computer. Notice that no device reads the anode surface potential, the quantity that ends the charge.

9.7 Technology in Practice

The instruments and the post of Section 9.3 are, in a vehicle, a current sensor, voltage taps, thermistors, a management computer, a coolant pump and an off-board charger. Figure 9.6 draws the chapter's loops with that hardware in place of the schematic's symbols.

What Measures It

The quantity that ends a fast charge has no field instrument. Plating can occur when the local potential of the negative electrode is below zero against lithium [15], and no vehicle in service carries a reference electrode that reports it (Section 9.1). What the management system does read is the pack current, the terminal voltage and a few temperatures. A current can be measured as the voltage across a precision shunt resistor of small value, or without breaking the circuit by a Hall-effect sensor that clamps around the conductor [22]. Whichever is fitted, the counted state of charge integrates its error, which is the bias b of Equation (9.16).

> **TECHNOLOGY** **Sensor: Thermistors on the Modules**
>
> A thermistor is a piece of metal oxide whose resistance changes with temperature. It is highly sensitive and nonlinear, and it is typically used where high accuracy is unimportant [22]. The current that reads its resistance also heats it, which adds a positive error [22]. A national laboratory study of fast charging noted that if temperature and voltage sensors could be placed on every cell, the weaker and hotter cells could be identified [15]. The chapter's thermistor sits on the module case, reads low in proportion to the current and is filtered with a 0.25 min time constant.

What Moves It

Two final elements act on the pack. The charger sets the current, and a pump sets the coolant flow. Liquid cooling, typically a mixture of ethylene glycol and water, is the preferred thermal management for most battery vehicles. Its many connections raise the potential for failure, and a failed system could let the liquid short out adjacent cells [17]. The chapter's coolant controller writes the flow q of Table 9.3, and its chiller of Equation (9.13) stands in for the refrigeration that removes the heat.

> **TECHNOLOGY** **Final Element: The Off-Board Direct-Current Charger**
>
> A direct-current fast charger converts three-phase alternating voltage in two power stages. A rectifier with power factor correction makes an intermediate direct voltage, and a second converter regulates the direct voltage the battery needs. Galvanic isolation comes from a line-frequency transformer ahead of the rectifier or from a high-frequency transformer inside the second stage, and identical modules run in parallel when one is not enough [23]. The vehicle's charge controller tells the off-board charger the battery's constraints, and the two coordinate the charging voltage and current [16]. The chapter's post follows the vehicle's command with a 0.02 min time constant up to a 500 A ceiling.

Where the Controller Runs

The selector of Section 9.5 runs in the vehicle's management computer, an embedded controller of the kind that replaced analog mechanical controls in machines and engines [24]. The vehicle and the off-board charger coordinate the charging voltage and current over a communication link [16]. Control loops of this kind run with cycle times from milliseconds to minutes [24]. In the chapter the management system executes every 0.02 min, which is 1.2 s, on a physics step of 0.3 s, and the predictive controller decides every 0.1 min. The inverse map, the voltage and temperature loops and the low selector with its preload are one block of that software, and the coolant loop runs beside it.

> **IN PRACTICE** **The Standards Between Vehicle and Post**
>
> Conductive charging in North America is defined in SAE J1772, and IEC 62196-3 defines four vehicle coupler configurations for direct-current fast charging [23]. IEC 61851-23 covers the direct-current charging station and requires that each output circuit be isolated from the others, and IEC 61851-24 covers the digital communication between station and vehicle for control of the charge [23]. Cable weight is a practical limit. A 50 kW cable weighs about 9 kg, and at 400 V a cable for 200 kW can exceed the 22.7 kg a single person may safely lift [23].

What Is Modernizing It

Four developments touch this loop. The first is voltage. A study in 2017 took 800 to 1000 V packs and more than 400 kW of charging power as the design point for future vehicles, and a higher charging voltage reduces the size of the cable between charger and vehicle [16]. The second is the cable. Liquid cooling of the cable reduces its thermal stress and makes a smaller and lighter cable feasible [23]. The third is the station. A station can share one central rectifier on a direct-current bus, or use a solid-state transformer that rectifies, steps down and isolates in one unit [23]. The fourth is communication. The vehicle must send its battery constraints to the off-board charger, and a 1000 V, 400 kW link could give an attacker an enticing target [16]. None of the four measures the anode, and the open challenges below concern that gap.

9.8 Limitations

Nothing in this application supports a claim about a real cell. The anode criterion is lumped, the plating law has no nucleation threshold and no morphology, the pack has no cell-to-cell spread, and the aging law has two terms where a life prediction needs many. The projections over five hundred charges are the arithmetic of Equation (9.8) repeated, and the comparisons between strategies are the result while the absolute ampere hours and percentages are not. Section 9.3.8 states where each simplification stops holding.

The scope is narrower than the domain in three further respects. The state estimator behind the selector is plain coulomb counting and a table lookup, with no Kalman filter, no on-line resistance identification and no adaptation, which is deliberate because the subject of the application is what happens when a model does not adapt [8]. Only charging is modeled, so regenerative braking, driving loads and the discharge side of the aging problem are absent. The vehicle and post negotiate nothing, since the post is Equation (9.14) and a slider.

Two properties of the numerics are convenient and unrealistic. A frozen integration step and a frozen scan period reproduce the same trajectory on any computer, while a real pack controller has jitter, quantized measurements and a scan order that moves. The absence of any noise term means the selector never has to arbitrate between two candidates whose measurements are rattling, which is precisely the condition under which the preload of (9.23) earns most of its keep in a plant, and the deterministic chatter reported in Section 9.6.3 understates it. Measurement noise on a loser near its limit would also make the cap of (9.23) rattle, which the deterministic model does not exercise.

The predictive controller’s model leaves out the coolant node, the chiller and the cabin, which it lumps into one sink temperature, and it holds the coolant flow over the horizon, so its prediction of the core ten minutes ahead runs up to 5.9 °C hot where the coolant loop opens later in the window. That prediction is not displayed. Its visible consequence is a predicted state of charge that reads low, by 3.8 points on average against the charge reached ten minutes later in scenario 12. The model also omits capacity fade, the true aging and the true kinetic resistance, so it inherits the error of the vehicle’s estimate in full.

OPEN CHALLENGE Detecting Plating on Board

Plating can occur when the local potential of the negative electrode falls below zero against lithium, and it grows with current density and as the temperature falls [15]. Methods that detect it without taking the cell apart exist in the literature: high-precision coulometry, small changes in cell volume, calorimetry, and a voltage plateau while the plated metal strips back [15]. None of them is an instrument a vehicle carries on every charge, and the 2017 assessment listed methods to detect and prevent plating as a research need, since plating is a safety and performance issue [15]. Progress would be a signal in service that tells a computed margin from a real one.

9.9 Discussion

The result that generalizes past batteries is the one about the preload, and it generalizes because it does not depend on the physics. Any low selector whose losing controllers are tracked exactly to the selected output has candidates sitting on the boundary, and any change in a proportional term across one scan is then enough to hand the selection over. Whether the outcome is a downward ratchet or a limit cycle depends on whether the new winner inherits the old winner's value, and with external reset it does, so the failure is a slow drift of the manipulated variable toward one end of its range. A drift is far harder to see on a trend than a cycle, and the plant looks starved rather than unstable. The limiter selector of a jet engine in Volume 4, Chapter 3, whose every integrator tracks the fuel commanded, has the same structure. The remedy is a preload larger than the per-scan movement of the fastest losing controller, which Table 9.6 puts between 5 and 15 A on this plant once a charge is under way. The opening scans are the exception, because the ohmic step moves the voltage controller by 288 A in one scan, and a selector that must be quiet from the first scan needs the step itself removed, by ramping the command or by holding the losers tracked for the first few scans.

The result that matters most inside the domain is the one about margin. An inferred constraint is defended by a number chosen in advance, and that number is paying for model error rather than for uncertainty in the plant. The distinction has consequences for how the number is justified, because a tolerance to a 20 % optimistic table is an auditable statement about a validation program while a safety factor is not. The sticky-margin override of the spray dryer in Volume 5, Chapter 10 defends another constraint that is inferred rather than measured. It also sets the economics of adaptation. On the aged pack of scenario 9 a 30 mV margin plates nothing and adds 3.50 min to the 21.34 min a new pack takes, while a margin chosen to cover an unknown amount of aging has to be wider still and 45 mV adds 9.94 min. Identifying the resistance from the voltage response would recover most of that [1], which is the trade a fleet operator would price.

OPEN CHALLENGE A Safe Envelope for Uneven Aging

Cells in one pack age at slightly different rates, so after a fast charge some sit at a higher state of charge than others, and the management system has to limit the current through the weaker cells to prevent overheating [15]. The 2017 assessment asked for a management system sensitive enough to compensate for non-uniform aging, and for research to identify a safe envelope beyond which major performance and safety concerns arise [15]. Progress would be limits that follow each pack's aging rather than a table fitted when it was new.

The plant is also a reasonable target for the dynamic optimization and estimation methods of other volumes in this series. Equations (9.1) to (9.8) with the plating rate as a soft constraint and the charge time as the objective is a small nonlinear program of the kind a dynamic optimization package handles [25], and posing it that way replaces the selector and the margin with one explicit constraint and one explicit horizon [13]. The estimation problem sits alongside it, since the same equations with the terminal voltage as the measurement and R_{ct} as an estimated parameter are a moving horizon estimator, and running the two together is the adaptive scheme the aged pack argues for.

9.10 Conclusion

A vehicle pack on a fast charger serves in this chapter as an interactive laboratory for constraint control in which the binding constraint has no instrument. Nine states carry the state of charge, two polarization voltages, the pack and coolant temperatures, the charger current and three counters, and four candidate controllers write to one current through a low selector. Three of the four are proportional-integral loops closed on a measured error. The fourth is the algebraic inverse of a kinetic model, filtered and looked up, and it holds its computed potential between 13.9 and 14.5 mV against a 15 mV setpoint whether or not the pack is anywhere near that value. On this pack the constant-current protocol reaches 80 % in 16.84 min and deposits 3.381 Ah of metallic lithium for 0.1035 % of capacity, while the anode-limited strategy takes 21.34 min, deposits nothing and fades 0.0779 %. A table 30 % optimistic still plates 2.237 Ah with the estimate held perfectly, and a 15 mV margin covers a table roughly 20 % optimistic and no more. Removing the selector preload changes the selection 2999 times in 60 min and ratchets the current from 212 to 15 A without finishing the charge. The optional predictive controller holds the four limits in one plan, and in scenario 12 it keeps the core at 35.95 °C and the voltage within 0.01 V of its limit for a charge 3.22 min longer than the selector’s. Two of the four constraints are never selected on a charge to eighty percent for more than 0.32 min, which is the general observation the application is built to deliver, since knowing which constraint carries the plant is worth more than tuning the ones that do not.

KEY POINTS

- No instrument in the vehicle reads the potential at the graphite surface, the limit that ends a fast charge.
- From 25 °C, constant current reaches 80 % in 16.84 min and plates 3.381 Ah, and the anode-limited charge takes 21.34 min and plates nothing.
- A 15 mV margin covers a table roughly 20 % optimistic and no more.
- Without a preload the selection changes 2999 times in 60 min and the charge never finishes.
- Knowing which constraint carries the plant is worth more than tuning the ones that do not.

Where else this applies. A low selector over limiters that guard a margin no sensor reads also sets the fuel of the turbofan in Volume 4, Chapter 3, and an unmeasured core temperature is controlled in Volume 5, Chapter 11.

References

[1] F. F. Oehler, X. Deuschl, K. Nürnberger, A. Graule, S. Kücher, T. Roth, A. Adam, J. Li, R. Mörtel, and A. Jossen, “Online adaptive anode potential-controlled fast charging of lithium-ion cells using a validated electrochemical model-based virtual reference electrode,” *Journal of Power Sources*, vol. 608, p. 234620, 2024.

[2] P. Arora, M. Doyle, and R. E. White, “Mathematical modeling of the lithium deposition overcharge reaction in lithium-ion batteries using carbon-based negative electrodes,” *Journal of The Electrochemical Society*, vol. 146, no. 10, pp. 3543–3553, 1999.

[3] T. Waldmann, B.-I. Hogg, and M. Wohlfahrt-Mehrens, “Li plating as unwanted side reaction in commercial Li-ion cells – a review,” *Journal of Power Sources*, vol. 384, pp. 107–124, 2018.

[4] T. Ohzuku, Y. Iwakoshi, and K. Sawai, “Formation of lithium-graphite intercalation compounds in nonaqueous electrolytes and their application as a negative electrode for a lithium ion (shuttlecock) cell,” *Journal of The Electrochemical Society*, vol. 140, no. 9, pp. 2490–2498, 1993.

[5] Y. Ji, Y. Zhang, and C.-Y. Wang, “Li-ion cell operation at low temperatures,” *Journal of The Electrochemical Society*, vol. 160, no. 4, pp. A636–A649, 2013.

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

[7] R. Hanus, M. Kinnaert, and J.-L. Henrotte, “Conditioning technique, a general anti-windup and bumpless transfer method,” *Automatica*, vol. 23, no. 6, pp. 729–739, 1987.

[8] N. A. Chaturvedi, R. Klein, J. Christensen, J. Ahmed, and A. Kojic, “Modeling, estimation, and control challenges for lithium-ion batteries,” in *Proceedings of the 2010 American Control Conference*, 2010, pp. 1997–2002.

[9] S. S. Zhang, "The effect of the charging protocol on the cycle life of a Li-ion battery," *Journal of Power Sources*, vol. 161, no. 2, pp. 1385–1391, 2006.

[10] P. H. L. Notten, J. H. G. Op het Veld, and J. R. G. van Beek, "Boostcharging Li-ion batteries: A challenging new charging concept," *Journal of Power Sources*, vol. 145, no. 1, pp. 89–94, 2005.

[11] M. Doyle, T. F. Fuller, and J. Newman, "Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell," *Journal of The Electrochemical Society*, vol. 140, no. 6, pp. 1526–1533, 1993.

[12] R. Klein, N. A. Chaturvedi, J. Christensen, J. Ahmed, R. Findeisen, and A. Kojic, "Optimal charging strategies in lithium-ion battery," in *Proceedings of the 2011 American Control Conference*, 2011, pp. 382–387.

[13] H. E. Perez, X. Hu, S. Dey, and S. J. Moura, "Optimal charging of Li-ion batteries with coupled electro-thermal-aging dynamics," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 9, pp. 7761–7770, 2017.

[14] A. Tomaszewska, Z. Chu, X. Feng, S. O'Kane, X. Liu, J. Chen, C. Ji, E. Endler, R. Li, L. Liu, Y. Li, S. Zheng, S. Vetterlein, M. Gao, J. Du, M. Parkes, M. Ouyang, M. Marinescu, G. Offer, and B. Wu, "Lithium-ion battery fast charging: A review," *eTransportation*, vol. 1, p. 100011, 2019.

[15] S. Ahmed, I. Bloom, A. N. Jansen, T. Tanim, E. J. Dufek, A. Pesaran, A. Burnham, R. B. Carlson, F. Dias, K. Hardy, M. Keyser, C. Kreuzer, A. Markel, A. Meintz, C. Michelbacher, M. Mohanpurkar, P. A. Nelson, D. C. Robertson, D. Scoffield, M. Shirk, T. Stephens, R. Vijayagopal, and J. Zhang, "Enabling fast charging – a battery technology gap assessment," *Journal of Power Sources*, vol. 367, pp. 250–262, 2017.

[16] A. Meintz, J. Zhang, R. Vijayagopal, C. Kreutzer, S. Ahmed, I. Bloom, A. Burnham, R. B. Carlson, F. Dias, E. J. Dufek, J. Francfort, K. Hardy, A. N. Jansen, M. Keyser, A. Markel, C. Michelbacher, M. Mohanpurkar, A. Pesaran, D. Scoffield, M. Shirk, T. Stephens, and T. Tanim, "Enabling fast charging – vehicle considerations," *Journal of Power Sources*, vol. 367, pp. 216–227, 2017.

[17] M. Keyser, A. Pesaran, Q. Li, S. Santhanagopalan, K. Smith, E. Wood, S. Ahmed, I. Bloom, E. Dufek, M. Shirk, A. Meintz, C. Kreuzer, C. Michelbacher, A. Burnham, T. Stephens, J. Francfort, B. Carlson, J. Zhang, R. Vijayagopal, K. Hardy, F. Dias, M. Mohanpurkar, D. Scoffield, A. N. Jansen, T. Tanim, and A. Markel, "Enabling fast charging – battery thermal considerations," *Journal of Power Sources*, vol. 367, pp. 228–236, 2017.

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

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

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

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

[22] 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] H. Tu, H. Feng, S. Srdic, and S. Lukic, "Extreme fast charging of electric vehicles: A technology overview," *IEEE Transactions on Transportation Electrification*, vol. 5, no. 4, pp. 861–878, 2019.

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

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

9.11 Exercises

9.1 ○ Concept check · 2 min

While the anode map holds the current in scenario 4, the computed anode potential sits between 13.9 and 14.5 mV against its 15 mV setpoint. A colleague proposes adding integral action to the anode candidate to remove the offset. Which assessment is right?

(a) It would remove the offset and protect the pack better, since the true potential would then sit on the setpoint

(b) It would remove the offset but wind up while the map loses, so it needs the same preload as the other two loops

(c) It would not help, since the offset comes from the thermistor reading low and a better thermistor removes it

(d) It cannot be done in a useful sense, since nothing measures the potential and the residual is the filter lagging a falling limit

9.2 Concept check · 1 min

Scenario 5 is run with the Selector preload set to 0 A before Run. What does the charge do?

(a) The selection changes at almost every scan and the current ratchets steadily down, so the charge never finishes

(b) The voltage and temperature controllers enter a growing limit cycle, and the charger trips on the current swings

(c) The charger keeps the selector for the whole charge, and the voltage limit ends the charge early

(d) Nothing changes after the first minute, because the preload matters only in the opening transient

9.3 Concept check · 2 min

Scenario 7 arrives at 38 °C into 42 °C air. The anode-limited charge is run once with the coolant loop on temperature and once with the flow fixed at 10 %. Which statement is right?

(a) The cooled charge is faster, because a cooler pack has a lower ohmic resistance and a smaller voltage rise

(b) The cooled charge is slower, because a cooler anode has slower kinetics and a lower plating limit

(c) The two take the same time, because the anode limit depends on the state of charge and not on temperature

(d) The uncooled charge is slower, because the temperature controller cuts the current back as the pack heats

9.4 Theory and calculation · 6 min

The default cell has $U_0 = 0.085\,\mathrm{V}$, $U_A = 0.60\,\mathrm{V}$, $k_U = 12$, $R_{ct}^{0} = 1.2 \times 10^{-4}\,\Omega$ and $k_R = 8.0$, with no aging ($a = 1$). Work at 25 °C, where the Arrhenius factor of Equation (9.5) is one.

(a) Find U_{neg} and R_{ct} at 50 % and at 80 % state of charge. *Checkpoint:* 86.5 mV and 0.180 mΩ, then 85.0 mV and 0.513 mΩ.

(b) Find the plating current I_{lim} at both states of charge, and say which limit binds at 50 % on a 500 A charger. *Checkpoint:* 480 A and 166 A, so the charger binds.

(c) Find the current the map allows with $\varphi_{SP} = 15\,\mathrm{mV}$ when its signals and its table are exact. *Checkpoint:* 397 A and 137 A.

9.5 Theory and calculation · 7 min

Use Table 9.2. Absolute temperature is $T + 273.15\,\mathrm{K}$ and $T_{ref} = 298.15\,\mathrm{K}$.

(a) Find the factor by which the kinetic resistance at 0 °C exceeds its value at 25 °C, and the plating current at half charge from a 0 °C pack. *Checkpoint:* 3.41 and 141 A.

(b) Find the same factor for the ohmic resistance, which has $E_R = 3600\,\mathrm{K}$. *Checkpoint:* 3.02.

(c) Find the heat 500 A deposits in the nominal R_0^0 at 25 °C, and the chiller capacity at 42 °C ambient from Equation (9.13). *Checkpoint:* 5.88 kW against 4.51 kW, a loss of 31 %.

9.6 Theory and calculation · 8 min

The map holds its estimate on φ_{SP}, using $\hat{R}_{ct} = m R_{ct}$ in place of the true R_{ct}. Take the counted state of charge and the temperature it reads as exact.

(a) Show that the true potential is then $\varphi = U_{neg} - (U_{neg} - \varphi_{SP})/m$. *Checkpoint:* The map's current is $(U_{neg} - \varphi_{SP})/(mR_{ct})$.

(b) On the plateau, $U_{neg} = 85.4\,\text{mV}$. Find the smallest m a 15 mV margin tolerates, and the true potential at $m = 0.70$. *Checkpoint:* $m = 0.824$, then $-15.2\,\text{mV}$.

(c) Find the margin that keeps $\varphi \geq 0$ at $m = 0.70$, and at the challenge's believed-to-true ratio $0.85/1.3$. *Checkpoint:* 25.6 mV, then 0.654 and 29.6 mV.

9.7 App challenge · 8 min

Use the default seed and the speed 5 (25 min takes 5 s of wall time). Run each case to 25 min and read the scorecard tiles Time to target, Plated, Fade this charge and After 500 charges. The model has no noise, so the tiles should match the answers in their last digit.

(a) Load scenario 2, Constant current, constant voltage, and run it as shipped. *Checkpoint:* 16.8 min, 3.383 Ah, 0.1037 %, 51.8 %.

(b) Load scenario 4, Riding the limit, and run it as shipped. *Checkpoint:* 21.3 min, 0.000 Ah, 0.0780 %, 39.0 %.

(c) From the tiles, find the minutes the anode limit costs, the share of the fade it saves, and the capacity of a 200 Ah pack it saves over five hundred charges. Give the answers within 2 %. *Checkpoint:* 4.5 min, 25 % and 25.6 Ah.

9.8 App challenge · 8 min

Load scenario 8, The model is optimistic, which sets the Kinetic model error to 0.70 and the Thermistor offset to 0.3. Use the default seed, the speed 5 and runs of 25 min. Read the Plated, Time to target and Peak temperature tiles, which should match the answers in their last digit.

(a) Run it as shipped. *Checkpoint:* 2.237 Ah, 18.5 min, 34.7 °C.

(b) Set the Thermistor offset to 1.0, press Reset and run again. Then set it to 0.0, press Reset and run again. *Checkpoint:* 0.187 Ah at 1.0 and 3.422 Ah at 0.0.

(c) Explain the sign of the result, and say what it means for a project that fits a better thermistor to this vehicle. *Checkpoint:* A reading that is low makes the map believe a colder, more resistive anode.

9.9 App challenge · 9 min

Load the challenge, scenario 13, Ten minutes at the services. Use the default seed, the speed 10, and let each run go to the 45 min judgment without pausing (about 5 s of wall time). Read Time to target, Plated, Peak temperature and After 500 charges at 45 min. The tiles should match the answers in their last digit.

(a) Run it as loaded. *Checkpoint:* 21.6 min, 14.196 Ah, 28.0 °C, 62.6 %.

(b) Reload, reveal the suggested settings before Run, and run again. *Checkpoint:* 30.7 min, 0.024 Ah, 26.5 °C, 18.9 %.

(c) Reload, reveal the suggestion, set Cooling to Fixed flow with the Coolant flow at 20 %, and run again. Say whether the attempt passes and why the flow matters. *Checkpoint:* 34.7 min, 0.169 Ah, a fail on plating.

9.10 Application area · 3 min

The model is one equivalent cell scaled to the pack. A supplier proposes to ship this chapter's 15 mV margin, justified by Table 9.7, on a pack whose cells differ from one another by several percent in kinetic resistance.

(a) Which cell must the management system defend, and how does the spread change the margin? *Checkpoint:* The weakest cell, and the margin must cover the spread as well as the table error.

(b) State the specification sentence the table supports, in the form a design review would accept. *Checkpoint:* A tolerance to a named model error, not a safety factor.

9.11 Application area · 3 min

A charging network quotes its charge times from ten to eighty percent, and an automaker offers to warm the pack from the drivetrain on the way to the post.

(a) Use scenario 10 to explain why the network stops the quote at eighty percent. *Checkpoint:* The last twenty points take 21.8 min against 21.3 min for the first seventy.

(b) Use scenario 6 to state what preconditioning is worth on a 0 °C day. *Checkpoint:* 44.80 min from a 0 °C arrival against 27.40 min from 15 °C.

9.12 Review the agent · 5 min

An AI assistant reviewing scenario 12, Four limits, one QP, writes: "The anode-limited selector reaches 95 % in 38.4 min and the predictive controller in 41.6 min. The selector held every limit on its instruments, so the predictive controller wastes 3.2 min and should not be used on this pack." Find the flaw, and say how to check it in the app with the default seed, at the speed 10 with runs of 50 min.

CHAPTER 10

A Vanadium Flow Battery Simulator for Inventory-Constrained Dispatch on a Duck Curve

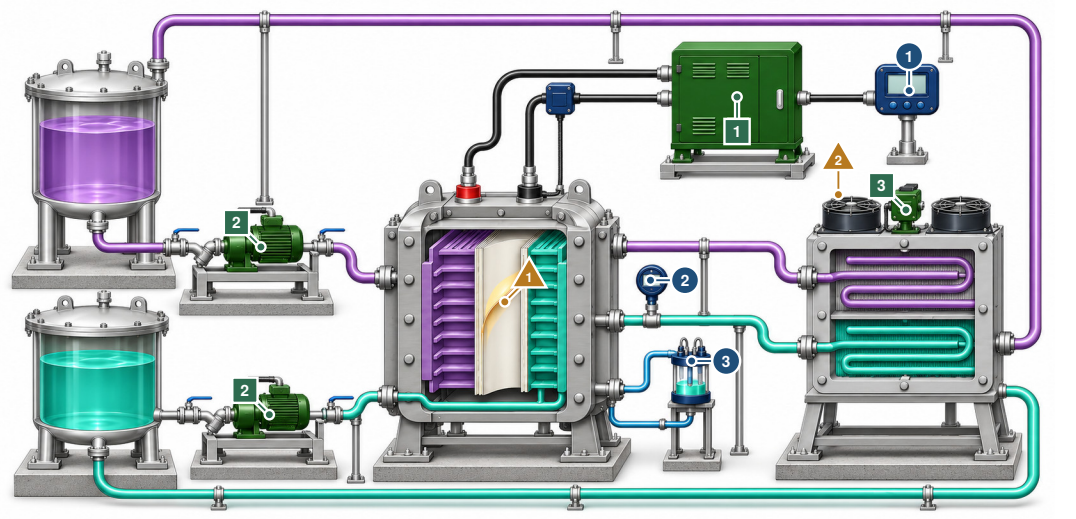

The controller moves 1 the power converter, which sets the current 2 the two electrolyte pumps, which set the flow 3 the cooler on the return lines

The controller reads 1 battery power on the line to the grid 2 electrolyte temperature at the stack outlet 3 the charge-state estimate

What disturbs it 1 vanadium crossing the membrane 2 hot weather at the cooler

Notice that the energy sits in the two tanks of liquid while the stack and the converter only set how fast it can flow in or out. Illustration, not to scale.

IN PLAIN TERMS **A Battery That Keeps Its Charge in Tanks**

The main idea. A flow battery stores its energy in two big tanks of liquid. Pumps carry the liquid through stacks of cells that turn it into electricity, so more tank means more hours and more cells mean more power.

The challenge. The tanks hold only so much, and the same power is worth far more at some hours than at others. Deciding when to empty them is the hard part. A forecast helps only on days when the price jumps at an hour the grid's demand does not show.

Why it matters. Spending the stored charge at the wrong hour can cost thousands of dollars in a single day.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the two history boxes of Section 10.2, the fact box of Section 10.3, the try-it box of Section 10.4, the try-it and fact boxes of Section 10.7, and the key points
Study	50 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 10.7, Section 10.8, Section 10.9 and the open challenges

The control question A flow battery buys its megawatts as stacks and its megawatt-hours as electrolyte, and a converter answers a power command in seconds. When should a finite inventory be spent on a duck-curve feeder, and when is a forecast worth more than feedback on the net load?

Learning objectives

- Compute the energy, the working band and the hours at rated power of a tank of vanadium electrolyte, and what an imbalance takes off them.
- Explain why the transport margin, set by the pump, limits rated power before the tanks are empty.
- Compare fixed, feedforward and margin-trimmed electrolyte flow by pumping energy and worst margin.
- Explain when forecast-based dispatch beats a net-load feedback controller, and why a finite-horizon plan needs a terminal condition.

Prerequisites Sections 1.4.4 (storage as an inventory), 1.4.6 (horizon and terminal value), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios and data).

The app https://apmonitor.com/apps/vanadium-flow/

See also Volume 9, Chapter 6 and Volume 9, Chapter 11 for a stored inventory scheduled against a price over a day, with a predictive or integer plan. Chapter 13 and Chapter 11 in this volume for a quadratic dispatch program and a water-value plan.

10.1 Introduction

A battery in a control course is usually a bucket. Charge goes in at one efficiency, comes out at another, and the state of charge is the integral of the difference. That model is adequate for a lithium-ion cell on a bench and it is misleading for a flow battery on a feeder, because a flow battery does not have one size. Its energy is a volume of dissolved vanadium in two tanks and its power is an area of membrane in the stacks, and the two are bought on separate purchase orders [1]. The consequence is a plant that can be full and unable to deliver its rating, or empty with every stack healthy, and neither condition is visible in a lumped-efficiency model.

A common misconception is that storage control is a power problem. It is not. A grid-scale converter answers a megawatt command in seconds and there is very little to tune, so a student who spends an afternoon on the power loop has learned almost nothing about storage. Everything difficult sits above that loop. The tanks hold a finite quantity of charge, the price of spending it moves by more than a factor of ten across a day, and the evening ramp arrives whether or not the midday surplus was used well. The decision of when to spend the inventory is the control problem, and it is the one that a transfer function and a step test cannot pose.

The class of problem is inventory-constrained economic dispatch under a moving price, and it is general. A pumped hydro reservoir, a compressed air cavern, a hydrogen buffer, a chilled water tank, a refinery intermediate tank and a warehouse of finished product all share the shape. Each has a state that integrates the manipulated variable, a capacity constraint at both ends, a time-varying value on the contents and a horizon shorter than the physics that fills it. The formal statement of that shape is economic model predictive control, in which the optimizer minimizes a cost rather than a tracking error [2]. Once a student has seen why a finite-horizon optimizer empties the tank at the edge of its horizon, the same reasoning transfers to every plant in that list.

The chapter rests on one controlled comparison on one plant, between feedback dispatch and forecast-based dispatch, and it reports the case where the forecast is worth nothing as carefully as the case where it is worth almost seventeen thousand dollars in a day. The plant is a twelve-state electrochemical and thermal model with a supervisory layer that runs at interactive speed in a browser. The supervisory layer is a quadratic program of twenty-four variables solved by a fixed count of projected-gradient steps rather than a nonlinear program. Eleven scenarios and a manual mode beneath them let the reader reproduce each result.

This chapter gives the model, the control structure and the measured behavior of the vanadium flow battery simulator in enough detail to rebuild it, and uses it to show why the inventory rather

than the converter sets the value of grid storage. Section 10.2 traces the domain from the NASA redox program to grid-scale vanadium plants. Section 10.3 gives the electrochemical, hydraulic and thermal model and the synthetic feeder. Section 10.4 measures the open-loop character in manual mode. Section 10.5 describes the four regulatory loops and the battery management limits, and Section 10.6 the dispatch layer above them. Section 10.7 reports the scenarios: the idle day, the flow policy, feedback against forecast dispatch, the terminal charge state, degradation and the graded challenge. Section 10.8 describes the hardware of a real plant, Section 10.9 states the limitations and Section 10.10 the conclusions. In Chapter 1 the plant is an example of storage as an inventory (Section 1.4.4) and of dispatch against a forecast with a terminal value (Section 1.4.6).

10.2 From the NASA Redox Program to a Feeder

The redox flow battery began as an iron-chromium storage program at NASA, in which shunt currents through the shared electrolyte, pumping power and cell performance had to be traded against one another [3].

> ⧖ HISTORY **Pumping Power in the NASA Redox Program (1982)**
>
> A NASA laboratory in Cleveland developed a redox storage system from the mid-1970s, with iron and chromium solutions pumped from tanks through stacks of up to forty cells [3]. Its 1982 report called the sizing of each cell's flow ports the single most important design feature. Narrow ports suppress the shunt currents that leak through the shared electrolyte, wide ports cut the pumping power, and the cell's performance sets a minimum flow, so no port size wins on all three [3]. A throttling valve had been used to vary the flow, which saved no pumping energy. The report suggested a variable-frequency motor and pump instead, and estimated that a variable flow schedule could cut the parasitic losses by about 65 % against a fixed flow sized for the worst case [3].

The iron-chromium chemistry was abandoned because the two half-cells contained different metals, so any transport across the membrane contaminated both sides permanently. The all-vanadium cell solved that by using one element in four oxidation states, so crossover causes an imbalance that can be corrected rather than a contamination that cannot [4]. The first full characterization of a working all-vanadium cell followed two years later and reported the coulombic and voltaic efficiencies that the technology still quotes [5]. Commercial development through the 1990s and 2000s was slow and largely Japanese and Australian, and the review literature of the period reads as a catalog of stack engineering rather than of system operation [6].

Control of these plants developed late and in a particular order. Stack-level electrochemical models capable of predicting the transient voltage appeared in 2008 [7], thermal models of a complete battery including self-discharge followed in 2012 [8], and quantitative membrane transport models were published in the same year [9]. Flow rate was treated as a design constant until pressure-drop and efficiency optimization studies showed that a fixed flow is wrong at almost every operating point [10], and gain-scheduled flow control was proposed once it was clear that the gain from flow to mass-transfer margin falls as the flow rises [11]. The shunt current that the NASA report identified was quantified for a modern multi-stack plant in 2013 [12].

The application context arrived from outside the battery literature, with the duck chart of a solar-heavy grid [13].

> ⧖ HISTORY **The Duck Chart (2013)**
>
> In 2013 California's grid operator published a chart of the net load, the demand minus wind and solar output, on a spring day as solar capacity grew [13]. The curve sagged at midday, where solar output peaks, and its name came from its resemblance to the profile of a duck. The chart raised the concern that the conventional power system could not meet the ramp rate and range needed to use all the solar energy, so output would be curtailed. Negative prices had already appeared in that market, driven in part by wind. In October 2013

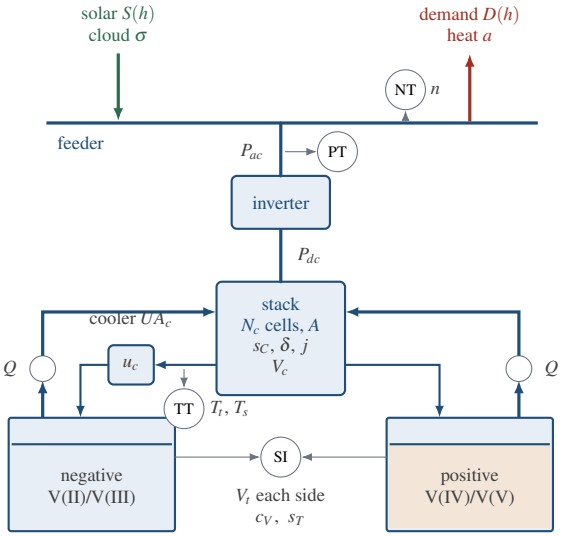

Figure 10.1 Plant arrangement and the symbols used in the equations. Two tanks of vanadium electrolyte hold the energy, the stacks hold the power, and one pump per circuit sets the rate at which reactant reaches the membrane. The cooler acts on the combined return header. PT is the power transducer, TT the electrolyte thermometer, SI the charge-state estimator and NT the feeder net-load measurement.

> the state's utility regulator set targets for 1325 MW of energy storage by 2020 [13].

Storage is the answer being built at scale, and reviews of the storage portfolio place flow batteries in the multi-hour bulk role rather than in the fast frequency role [14]. Predictive dispatch of a flow battery for peak shaving is now an active subject in its own right [15]. The plant of this chapter sits at that junction, where a well-understood electrochemical device is asked to answer a scheduling question.

10.3 Process Model

The simulated plant is a 100 MW converter connected to 176 000 cells of 4000 cm^2 active area, arranged as 1100 stacks of 160 cells. Each side of the plant holds 10 000 m^3 of 1600 mol/m^3 vanadium electrolyte in a tank, plus 150 m^3 of hold-up in the stacks and manifolds. At the formal cell potential of 1.40 V that inventory is 600.4 MWh between empty and full and 480.3 MWh inside the 10 % to 90 % working band. Two pumps circulate electrolyte, one cooler removes the loss heat, and one inverter connects the direct-current stacks to the feeder. Figure 10.1 shows the arrangement and every symbol used in the equations that follow.

The sign convention is positive for discharge throughout. A positive current density, a positive direct-current power and a positive alternating-current power all mean that the battery is supplying the feeder. The clock is in hours, power is in megawatts and energy in megawatt-hours, current density is in amperes per square centimeter, and charge state is a fraction rather than a percentage.

10.3.1 Variables

Table 10.1 classifies every quantity the app exposes. Four controlled variables are paired with four manipulated variables, and the pairing is one-to-one because each loop acts on a different time scale. The disturbances split into three that the dispatch layer is told about in its forecast, namely the feeder demand scale, the afternoon air conditioning load and the scarcity price adder, and eight that it is not.

Table 10.1 Variable Classification

Symbol	Meaning	Unit	Range
Controlled variables			
n	feeder net load	MW	−100 to 700
P_{ac}	battery power at the bus	MW	−100 to 100
m	transport margin $j_\ell/\lvert j\rvert$	—	1.2 to 8
T_s	stack outlet temperature	°C	18 to 40
Manipulated variables			
P_b	battery power request	MW	−100 to 100
j	stack current density	% rated	−100 to 115
Q	electrolyte flow per circuit	m³/s	2 to 30
u_c	cooler duty	%	0 to 100
r	rebalance cell flow	%	0 to 100
Disturbances			
σ_∞	sky clarity	× clear	0.10 to 1.05
ℓ	feeder demand scale	×	0.80 to 1.30
a_∞	afternoon air conditioning	MW	0 to 200
T_a	ambient temperature	°C	0 to 45
w_p	hydraulic resistance	×	0.8 to 3.0
w_c	membrane and contact resistance	×	0.8 to 2.5
χ	crossover rate	×	0 to 5
$\Delta\pi$	scarcity price adder	\$/MWh	0 to 700
ε	forecast error scale	×	0 to 1.5
A_{reg}	regulation signal, when in service	MW	0 to 30
σ_P	power transducer noise	MW	0 to 6
States			
s_T	tank charge state	—	0 to 1
s_C	charge state inside the stacks	—	0 to 1
δ	half-cell imbalance	—	−0.45 to 0.45
T_t	tank electrolyte temperature	°C	−5 to 80
j_a	applied current density	A/cm²	−0.120 to 0.138
Q_a	applied electrolyte flow	m³/s	0.5 to 32
u_c^a	applied cooler duty	—	0 to 1
R_m	membrane resistance growth	—	0 to 3
ϕ	capacity fade	—	0 to 0.5
$\bar{n}$	filtered daily mean net load	MW	—
σ	filtered sky clarity	—	0 to 1.05
a	filtered air conditioning load	MW	0 to 400

10.3.2 Cell Voltage

The open-circuit potential of one cell follows the Nernst expression in the two half-cell charge states, which differ from each other by the imbalance δ,

$$E = E^0 + \frac{RT}{F}\left[\ln\frac{x}{1-x} + \ln\frac{y}{1-y}\right] \tag{10.1}$$

with $x = s_C + \delta/2$ and $y = s_C - \delta/2$, with T in kelvin and both arguments clipped to $[10^{-4}, 1-10^{-4}]$. With no imbalance this reduces to the usual $2(RT/F)\ln[s_C/(1-s_C)]$ and returns E^0 exactly at half charge. The two half-cell charge states are carried separately because a plant with crossover has capacity set by whichever half runs out first, and that is invisible in a single charge-state variable.

Three loss terms are subtracted on discharge and added on charge. Activation loss uses the

symmetric Butler-Volmer inverse,

$$\eta_{act} = \frac{2RT}{F}\,\operatorname{asinh}\!\left(\frac{|j|}{2j_0}\right), \quad j_0 = j_{00}\,\mathrm{e}^{E_{a0}\left(\frac{1}{303.15}-\frac{1}{T}\right)} \tag{10.2}$$

ohmic loss uses an area-specific resistance that improves with temperature and degrades with membrane age and the fouling disturbance,

$$\eta_{ohm} = R_\Omega |j|, \quad R_\Omega = \mathrm{ASR}_0\, w_c\,(1+R_m)\,\mathrm{e}^{E_{aR}\left(\frac{1}{T}-\frac{1}{303.15}\right)} \tag{10.3}$$

and concentration loss follows the limiting-current logarithm,

$$\eta_{conc} = -\frac{RT}{F}\ln(1-u), \qquad u = \min\!\left(\frac{|j|}{j_\ell}, 0.985\right). \tag{10.4}$$

The clip at 0.985 keeps the logarithm defined and caps the term at about 0.11 V at 30 °C, so the concentration loss alone never stops the current. That is the job of the supply limit in Equation (10.10) below. The limiting current density is where the pump enters the polarization,

$$j_\ell = j_{\ell 0}\,\alpha_C \left(\frac{Q}{Q_0}\right)^{n_Q}, \quad \alpha_C = \begin{cases} \min(x,\,y) & j \geq 0 \\ \min(1-x,\,1-y) & j < 0 \end{cases} \tag{10.5}$$

where α_C is the availability of whichever species is being consumed, so the same current is harder to sustain near an empty tank on discharge and near a full tank on charge. The cell terminal voltage and the stack powers follow,

$$\eta_{tot} = \eta_{act} + \eta_{ohm} + \eta_{conc} \tag{10.6}$$

$$V = \max(0.15,\; E - \operatorname{sgn}(j)\,\eta_{tot}) \tag{10.7}$$

$$P_{dc} = 10^{-6} N_c A\, V\, j \tag{10.8}$$

$$P_{loss} = 10^{-6} N_c A\, \eta_{tot}\, |j|. \tag{10.9}$$

The transport margin that the flow loop controls is $m = j_\ell / \max(|j|, 0.08\, j_R)$, capped at 12 so that an idle plant reports a finite number.

The current density j in Equations (10.2) to (10.9) and in the inventory balance below is the current the stacks actually pass. It is the applied current j_a limited by Faraday's law to the reactant that the flow brings from the tanks, at one electron per vanadium ion,

$$j = \mathrm{clip}(j_a, -j_F, j_F), \quad j_F = \frac{Q\, c_V F}{N_c A}\left[\alpha_T - \frac{|\delta|}{2}\right]^{+} \tag{10.10}$$

with $\alpha_T = s_T$ on discharge and $1 - s_T$ on charge. At design flow and half charge j_F is 1.64 A/cm^2, almost fourteen times the rated current density, so the stoichiometric flow that just supplies rated current at half charge is 1.09 m^3/s per circuit and the limit is idle in normal operation. It binds at low flow, where 3 m^3/s cannot supply rated current below 18 % charge, and near an exhausted tank, where it tapers charge and discharge to zero in manual as well as automatic mode while the pumps and auxiliaries go on drawing power. The same limited current sets the power, the losses and the charge consumed, so no current delivers energy that the tanks could not supply.

10.3.3 Inventory and Transport

The charge state inside the stacks and the charge state in the tanks are separate states because they move on separate time scales. Electrolyte passes through the 150 m^3 of cell and manifold hold-up in

10 s at design flow and through the 10 000 m^3 tank in 11.1 min, so

$$\frac{\mathrm{d}s_C}{\mathrm{d}t} = 3600\left[\frac{Q}{V_c}(s_T - s_C) - \frac{N_c A\, j}{F c_V V_c}\right] \tag{10.11}$$

$$\frac{\mathrm{d}s_T}{\mathrm{d}t} = 3600\,\frac{Q}{V_t}(s_C - s_T). \tag{10.12}$$

The factor 3600 converts the second-based flow and Faraday terms to the hour clock. Vanadium is conserved by construction, because $V_t\,\dot{s}_T + V_c\,\dot{s}_C$ equals the reaction term alone and vanishes at zero current. The gap between s_T and s_C is the entire transport story, and it is what makes the charge state measured at the stack outlet a biased estimate of the charge state in the tanks whenever current is flowing.

Crossover moves vanadium across the membrane at a rate that grows with temperature and with current, and the rebalance cell is the only thing that reverses it,

$$\frac{\mathrm{d}\delta}{\mathrm{d}t} = \chi k_x\left(0.35 + 0.65\frac{|j|}{j_R}\right)(s_T - 0.5)\,\mathrm{e}^{E_{aX}\left(\frac{1}{303.15} - \frac{1}{T}\right)} - k_{rb}\frac{r}{100}\,\delta. \tag{10.13}$$

The driving term changes sign at half charge, which reproduces the observed behavior that a cell parked full or parked empty drifts in opposite directions [9].

10.3.4 Balance of Plant

Pumping power for both circuits is cubic in flow on a fixed hydraulic system, which is why the last few cubic meters per second are the expensive ones,

$$P_{pump} = \frac{2k_Q\, w_p\, Q^3}{\eta_p\, 10^6}. \tag{10.14}$$

Inverter loss carries a no-load term, a conduction term linear in the direct-current power and a resistive term quadratic in it,

$$P_{inv} = P_{nl} + k_1|P_{dc}| + \frac{k_2 P_{dc}^2}{P_R}. \tag{10.15}$$

The power delivered at the point of connection subtracts everything the plant consumes, so on charge the feeder pays for the losses as well,

$$P_{ac} = P_{dc} - P_{inv} - P_{pump} - P_{aux,0} - P_{reb}\,\frac{r}{100}. \tag{10.16}$$

At rated current density, half charge and 15 m^3/s of flow the plant produces 98.71 MW of direct-current power at 1.1685 V per cell against an open-circuit 1.3924 V, loses 18.91 MW in the stacks and 1.45 MW in the pumps, and delivers 94.90 MW. The same 100 MW requires 106.4 % of rated current on discharge and only 72.2 % on charge, because charging runs at a higher cell voltage. The converter rating and the current rating are therefore different constraints, and which one binds depends on the direction.

Heat is generated by the stack loss and the pumps and removed through the tank walls and the cooler,

$$C_T\frac{\mathrm{d}T_t}{\mathrm{d}t} = 3600\Big[10^6(P_{loss} + P_{pump}) - UA_{amb}(T_t - T_a) - UA_c\,u_c^a\,(T_t - T_a)^+\Big] \tag{10.17}$$

where $(\cdot)^{+}$ denotes the positive part, so the cooler cannot heat the inventory. With 2.9×10^{10} J/°C of thermally active electrolyte the open-loop time constant is 25.1 h and the closed-loop constant at full cooler duty is 3.63 h. The sensor reads the stack outlet rather than the tank, and the stacks run a fraction of a degree above the tanks because the flow through them is large,

$$T_s = T_t + \frac{10^{6} P_{loss}}{2 Q \rho_e c_{p,e}}. \tag{10.18}$$

Two health proxies integrate use. Membrane resistance grows and capacity fades in proportion to current, and the fade accelerates above 35 °C,

$$\frac{\mathrm{d}R_m}{\mathrm{d}t} = k_{Rm}\frac{|j|}{j_R}, \qquad \frac{\mathrm{d}\phi}{\mathrm{d}t} = k_{fade}\frac{|j|}{j_R}\left[1 + 0.08\,(T_t - 35)^{+}\right]. \tag{10.19}$$

10.3.5 The Feeder, the Price and the Target

The feeder is a synthetic day built from a base load and three gaussian shoulders, each evaluated at the hour and at the hour plus and minus a day so that the evening peak is continuous across midnight,

$$w(h;c,s) = \sum_{k\in\{-1,0,1\}} \mathrm{e}^{-\left(\frac{h+24k-c}{s}\right)^{2}} \tag{10.20}$$

$$D(h) = \ell\Big[L_b + L_m w(h;8,2.6) + L_e w(h;19.5,2.4) + L_c w(h;14,3.5) + a\,w(h;17,3)\Big] \tag{10.21}$$

$$S(h) = \sigma S_{max}\sin^{1.3}\!\left(\pi\frac{h-h_{up}}{h_{dn}-h_{up}}\right), \qquad h_{up} < h < h_{dn} \tag{10.22}$$

and the raw net load is $n_{raw} = D - S$. With the default parameters the raw net load reaches 585.2 MW at 19:27, falls to −35.2 MW at 12:28 when the feeder exports, and averages 282.3 MW over the day. The steepest evening climb is 2.75 MW/min measured over a ten-minute window and 2.73 MW/min measured over an hour, so the ramp is not an artifact of the measurement window.

> ✦ DID YOU KNOW? **Curtailment at a Modest Solar Share**
>
> A 2015 national laboratory study of California found that, without new flexibility, solar supplying as little as 20 % of annual energy could lead to marginal curtailment rates above 30 %. Letting distributed solar and storage provide grid services could greatly reduce it [13].

Energy price is built the same way and deliberately does not peak at the same hour as the scarcity adder,

$$\pi(h) = p_b + p_{pk} w(h;19,1.7) - p_{bl} w(h;13,2.6) + \Delta\pi\, w(h;h_{sp},w_{sp}) + 6\cos\frac{2\pi(h-6)}{24}. \tag{10.23}$$

The default day runs from \$−13.6 per megawatt-hour at 13:07 to \$158 per megawatt-hour at 19:01. The scarcity event is centered at 15:00 by default, away from the evening peak, because on an ordinary duck-curve day the highest net load and the highest price are the same hour and a controller that watches only the net load is already pointing the right way.

The trajectory the feeder asks the plant to follow is peak shaving and valley filling and nothing more ambitious,

$$\begin{aligned} \hat{n} &= \mathrm{clip}(n_{raw},\ \bar{n} - W,\ \bar{n} + W) \\ n^{*} &= \mathrm{clip}(\hat{n},\ n_{raw} - 0.85P_R,\ n_{raw} + 0.85P_R) \end{aligned} \tag{10.24}$$

where $\bar{n}$ is a one-day first-order average of the raw net load, carried as a state, and W is a deadband inside which the feeder is left alone. The outer clip to within 85 % of the converter rating matters. A target the storage could not meet with full tanks is a statement about the size of the purchase order rather than a setpoint, and a controller scored against one is being measured on somebody else's decision. Flattening to the daily mean is the tempting alternative and it is wrong on this shape, because the overnight net load sits above the daily mean and a target of the mean asks the battery to discharge all night and arrive at the evening peak empty.

10.3.6 Actuators, Instruments and Filtered Disturbances

Three actuators are first-order lags on their commands, with time constants of 4 s for the inverter current, 20 s for pump speed and 60 s for the cooler. Three measurements are first-order filters, with time constants of 8 s for the power transducer, 3 min for the thermometer and 72 s for the charge-state estimator. Two disturbances are lagged rather than stepped, because a cloud front takes about a quarter hour to cross a feeder and an afternoon air conditioning load takes most of an hour to build,

$$\frac{\mathrm{d}\sigma}{\mathrm{d}t} = \frac{\sigma_{\infty} - \sigma}{\tau_{sky}}, \qquad \frac{\mathrm{d}a}{\mathrm{d}t} = \frac{a_{\infty} - a}{\tau_{ac}}. \tag{10.25}$$

A step in either would give the net load an unbounded derivative, which makes a ramp metric meaningless and does not happen on a real feeder. The net-load ramp reported to the operator is measured against a ten-minute lag rather than differentiated,

$$\dot{n} = \frac{n - \mathrm{lag}(n, \tau_r)}{60\,\tau_r}, \qquad \tau_r = 0.1667\,\mathrm{h}. \tag{10.26}$$

10.3.7 Parameters

Table 10.2 gives every parameter with its default value. The column marked V indicates whether the value is visible in the app's parameter panel, where A means adjustable and shown, H means hidden from the beginner view so that the reader must infer it from behavior, and F means fixed physical constant.

10.3.8 Initialization and Numerics

The initial state places the tanks at the scenario charge state, sets the imbalance and both health proxies to zero, and computes the temperature at which the loss heat produced by the initial commands is exactly removed,

$$T_t(0) = T_a + \frac{10^6\,[P_{loss}(0) + P_{pump}(0)]}{UA_{amb} + UA_c u_c(0)}. \tag{10.27}$$

The cell charge state is initialized to its algebraic steady value rather than to the tank value, because the hold-up empties in 10 s and starting it equal produces a transient that is a numerical artifact.

Integration is classical fourth-order Runge-Kutta on a fixed step of 0.001 h, which is 3.6 s and about one time constant of the fastest actuator. The regulatory controllers execute every 0.002 h, the dispatch layer every 0.25 h, and the historian logs one sample every 0.017 h. The declared and realized rate is 58.8 samples per simulated hour, with a logging interval of 61.2 s. The engine keeps these clocks apart from the screen, as Section 1.6.1 describes. A post-step clip holds every state inside its physical range, including the charge states (0 to 1), the imbalance and the cooler fraction.

10.3.9 Assumptions

1. **The electrolyte is perfectly mixed in each tank and in the stack hold-up.** A real tank stratifies, and the charge-state estimate at the return header then lags the bulk by more than the 11.1 min residence time used here. The assumption breaks down at very low flow, which is where the model is least trustworthy.

Table 10.2 Model Parameters and Default Values

Symbol	Meaning	Value	Unit	V
N_c	cells in the plant	176 000	—	A
A	active area per cell	4000	cm^2	A
j_R	rated current density	0.12	$\mathrm{A/cm}^2$	A
E^0	formal cell potential	1.40	V	A
ASR_0	area resistance at 30 °C	1.00	$\Omega\,\mathrm{cm}^2$	A
E_{aR}	resistance temperature coefficient	1800	K	A
j_{00}	exchange current at 30 °C	0.020	$\mathrm{A/cm}^2$	A
E_{a0}	kinetic temperature coefficient	2500	K	A
$j_{\ell 0}$	limiting current, design flow	0.90	$\mathrm{A/cm}^2$	A
n_Q	mass-transfer exponent on flow	0.5	—	H
V_t	tank volume, each side	10 000	m^3	A
V_c	stack and manifold hold-up	150	m^3	A
c_V	vanadium concentration	1600	$\mathrm{mol/m}^3$	A
ρ_e	electrolyte density	1350	$\mathrm{kg/m}^3$	H
$c_{p,e}$	electrolyte heat capacity	3000	J/(kg K)	H
Q_0	design flow, each side	15	m^3/s	A
k_Q	hydraulic resistance, one circuit	155	$\mathrm{Pa\,s}^2/\mathrm{m}^6$	A
η_p	pump and drive efficiency	0.72	—	A
C_T	thermal inventory	2.9×10^{10}	J/°C	A
UA_{amb}	tank and pipework heat loss	3.2×10^5	W/°C	A
UA_c	cooler conductance at full duty	1.9×10^6	W/°C	A
P_R	converter rating	100	MW	A
η_{plan}	one-way efficiency in the plan	0.84	—	A
P_{nl}	inverter no-load loss	0.25	MW	A
k_1	inverter conduction coefficient	0.0075	—	A
k_2	inverter resistive coefficient	0.011	—	A
$P_{aux,0}$	balance-of-plant auxiliary load	0.30	MW	A
P_{reb}	rebalance cell load at full flow	0.15	MW	A
k_x	crossover imbalance rate	0.020	1/h	A
E_{aX}	crossover temperature coefficient	3500	K	H
k_{rb}	rebalance cell effectiveness	0.55	1/h	A
k_{fade}	capacity fade at rated current	2.0×10^{-4}	1/h	A
k_{Rm}	membrane resistance growth	3.0×10^{-4}	1/h	A
L_b	overnight feeder demand	340	MW	A
L_m	morning shoulder	90	MW	A
L_e	evening peak above base	240	MW	A
L_c	afternoon cooling load	60	MW	A
S_{max}	photovoltaic capacity	430	MW	A
h_{up}, h_{dn}	sunrise, sunset	6.3, 19.0	h	A
p_b	overnight energy price	34	\$/MWh	A
p_{pk}	evening price adder	130	\$/MWh	A
p_{bl}	midday price depression	46	\$/MWh	A
π_{ref}	price scale in the objective	45	\$/MWh	H
c_{curt}	cost charged for curtailed solar	12	\$/MWh	A
c_{deg}	degradation cost on throughput	4.0	\$/MWh	A
s_{min}, s_{max}	working charge-state band	0.10, 0.90	—	A
T_{min}, T_{max}	precipitation limits	15, 40	°C	A
T_{trip}	stack high-temperature trip	45	°C	A
W	shaving deadband	110	MW	A
τ_{tgt}	averaging time behind $\bar{n}$	24	h	A
τ_j, τ_Q, τ_c	actuator lags	4, 20, 60	s	A
τ_{sky}, τ_{ac}	disturbance lags	0.25, 0.6	h	A
F	Faraday constant	96 485	C/mol	F
R	gas constant	8.314	J/(mol K)	F

2. **One lumped thermal capacitance.** The tanks, the pipework and the stack inventory are one temperature with a fraction of a degree of stack rise from Equation (10.18). This is reasonable at the flows modeled and it is wrong during a pump trip, when the stacks would heat independently in seconds.

3. **Shunt currents are folded into the area-specific resistance rather than modeled.** A multi-stack plant loses current through shared manifolds in a way that depends on stack count and manifold geometry [12], and no such term appears here. The consequence is that the modeled efficiency does not degrade with plant size.

4. **All cells are identical.** The model carries one voltage for 176 000 cells, so cell-to-cell distribution, the reason a real battery management system exists, is absent.

5. **Crossover is a single scalar.** The four vanadium species diffuse at different rates through a real membrane [9], and the single imbalance state cannot reproduce the resulting composition drift. It reproduces the effect that matters for dispatch, which is a loss of usable capacity that no controller reports.

6. **The feeder, the solar output and the price are smooth synthetic profiles.** They are teaching data shaped like a spring day in California and they are never labeled as measured. Real net-load data has structure at every time scale that no smooth profile reproduces.

7. **The dispatch layer plans with one constant round-trip efficiency and constant megawatt-hours per charge point.** That is a linearization of a logarithm, and it is corrected every fifteen minutes by re-planning from the measured state rather than by making the plan more accurate.

8. **Capacity fade and membrane growth are first-order in current and do not feed back into the charge-state capacity.** They are indicators rather than a life model, and a plant run for a simulated month would report a fade number that no accelerated-aging study supports.

10.4 Open-Loop Character

Manual mode hands the operator the stack current, the electrolyte flow, the cooler duty and the battery power request, and the first useful experiment is to decide which of those four moves megawatts, which moves megawatt-hours and which does neither. Current moves power at once and charge state slowly, because the tanks hold 600 MWh and the stacks hold about 9 MWh. Flow moves almost nothing until it gets low, because the concentration term in Equation (10.4) is flat until $|j|$ approaches j_ℓ. In the first scenario, cutting the flow from 10 m³/s to 3 m³/s at 60 % of rated current and 44 % charge lowers the transport margin from 4.2 to 1.7 within a quarter hour and the power only from 59.6 MW to 58.2 MW. The margin is still above its starvation line of 1.15, which takes a lower flow or an emptier tank to reach.

Three time scales separate cleanly. The converter answers a command in 4 s and the pump follows in 20 s, so under the power loop of the fourth scenario a 140 MW swing from 60 MW of discharge to 80 MW of charge is 63 % complete in 14 s and 90 % complete in 65 s. The tank charge state moves at 0.34 % per minute when discharging at 100 MW and 0.23 % per minute when charging at that power, so the working band lasts 3.9 h at rated discharge and 5.7 h at rated charge. The electrolyte temperature has a 25.1 h open-loop time constant, which is longer than the day the plant is being dispatched over. A student who sets the temperature setpoint at two in the afternoon has already lost, because the cooler cannot change anything before six.

The inventory constraint is the lesson manual mode exists to deliver, and it is delivered by omission. Nothing in manual mode holds the plant inside its working band. In the third scenario, holding −70 % of rated current from thirty percent charge at 10:00 brings the tank to the working ceiling of 90 % at 14:24, with the 100 MW converter rating having trimmed the command to −64.5 % on the way.

Charging continues at full current to 96.7 % and then tapers as the supply limit of Equation (10.10) runs out of uncharged reactant. It stops by 16:00 with the tank at 99.0 % and the positive half-cell full, which is where the positive electrolyte is nearly all V(V) and vanadium pentoxide precipitates on a warm day. From then on the plant draws 1.29 MW for its pumps and auxiliaries and stores nothing. Current is a command and charge state is a consequence, and only the operator is watching the consequence.

> **TRY IT IN THE APP 3 · Charge, Discharge, and the Inventory Nobody Mentioned**
>
> Load scenario 3 and press Run. Watch the tank charge state. The plant charges at −70 % of rated current from 30 % at 10:00 and reaches the 90 % working ceiling at 14:24, and nothing in manual mode stops it there. It stops by 16:00 at 99.0 % and then draws 1.29 MW for its pumps and auxiliaries.

Coming back down, the interesting number is not where the plant stops but where it stops being a 100 MW battery. Reversed to 100 % of rated current at the ceiling alarm, with the flow held at 12 m^3/s, the plant makes 106 MW at first and falls below 100 MW at 63 % charge as the open-circuit voltage falls. The limiting current in Equation (10.5) is proportional to the availability of the species being consumed in the cell, so the transport margin reaches its starvation line of 1.15 at 22 % tank charge (17:53) with 88 MW still flowing, and the concentration overpotential takes a growing share from there. The floor alarm sounds at 10.4 % with 78 MW still flowing, the supply limit cuts the current only below 4.7 %, and the plant stops exporting at 19:39 with the tank at 0.2 %. The tanks still held about 130 MWh when the cell began to starve. Usable energy and usable energy at rated power are different numbers, and the gap between them is set by the pump.

Pause. Try Exercise 10.2 before reading on.

The pump trade is worth measuring directly. At rated current, half charge and 30 °C, stack loss falls from 20.90 MW at 4 m^3/s to 18.70 MW at 25 m^3/s while pumping rises from 0.03 MW to 6.73 MW. The sum has a minimum at 8.92 m^3/s, which the app reports from a golden-section search of twenty-four iterations and twenty-six evaluations. That optimum is well below the design flow of 15 m^3/s, although still about eight times the stoichiometric flow, and it moves with charge state, rising to 16.09 m^3/s at twenty percent charge and falling to 7.20 m^3/s at eighty percent. At twenty percent charge the sum is lower again below 2.74 m^3/s, but only because the supply limit then cuts the current, and the golden-section search settles on the local minimum at 16.09 m^3/s, where the full current flows. A little concentration overpotential is cheaper than a lot of pumping, and how much cheaper depends on the inventory.

10.5 Regulatory Control

10.5.1 The Discrete Algorithm

Four loops share the velocity-form PID of Section 1.6.2, Equations (1.10) to (1.13), with the derivative on the measurement and back-calculation on the integral alone. The chapter writes the action flag as γ, the framework's a, with $+1$ for a reverse-acting loop and -1 for a direct-acting one in the ISA sense. The settings specific to this app are few. The tracking constant T_t is τ_I on every loop, no loop sets a slew rate, and all four loops run with $\tau_D = 0$, so derivative action is off everywhere in this app. The regulatory sample period is $\Delta t_c = 0.002\,\mathrm{h}$, which is 7.2 s. In manual mode the integral is set so that the controller output equals the operator's actuator value, so transfer is bumpless in both directions with no separate logic. The supply limit of Equation (10.10) acts outside the power loop's output bounds, so the power integral would wind against it, but no automatic case reaches it.

10.5.2 Four Loops on Three Time Scales

Table 10.3 gives the pairings and the default tunings. Two of the four are direct acting. More cooling makes the electrolyte colder, and more discharge lowers the feeder net load, so both process gains are negative.

Table 10.3 Regulatory Loops as Implemented

Loop	CV	MV	K_c	τ_I (h)	Action
net load	n	P_b	0.9	0.08	direct
power	P_{ac}	j	0.6	0.006	reverse
flow	m	Q	3.0	0.03	reverse
temperature	T_s	u_c	18	1.5	direct

The power loop is the trivial one and it is included because a student needs to see how trivial it is. One megawatt of alternating-current power costs about one percent of rated current on discharge (0.98 MW per percent at 60 MW) and less on charge (1.37 MW per percent at −80 MW), through a four-second converter. In the fourth scenario the loop reaches 63 % of a 60 MW step in 18 s and 90 % in 83 s, and there is no interaction to manage. Few plants answer a load change that fast, the grid-forming inverter of Chapter 7 being one.

The flow loop is a feedforward path with an optional feedback trim. The feedforward takes the flow the cell needs from the current command, which is known exactly,

$$Q_{ff} = \mathrm{clip}\left(k_{ff} \frac{|j_{cmd}|}{100} Q_0 + Q_{min},\ 2,\ 30 \right) \tag{10.28}$$

and the trim, when it is selected, adds a margin-controlled correction bounded so that the sum stays inside the pump range. That structure is the standard division of labor. Feedforward handles the current, which is measurable and instantaneous, and feedback handles the charge state, which is not. The gain from flow to transport margin falls as the flow rises and as the tanks empty, so a trim tuned at high flow is too aggressive at low flow and one tuned at low flow is sluggish at high flow. Gain scheduling on that loop is the answer the flow battery literature gives [11], and the app reproduces the problem it solves.

The temperature loop uses a setpoint rather than a limit because both precipitation boundaries matter. V(II) precipitates below about 15 °C and V(V) above about 40 °C, so the electrolyte has to be held in a band rather than kept cold. With a 1.5 h integral time against a 3.63 h closed-loop plant the loop is deliberately slow, and the practical consequence is that its setpoint is a planning decision made in the morning.

10.5.3 Battery Management as Cascade Anti-Windup

Charge and discharge limits are computed before the supervisory layer runs, not after it. Charging tapers to nothing as the fuller half-cell approaches the working ceiling and discharging as the emptier half approaches the floor, over a three-point band so that the limit arrives as a derate rather than as a trip. The halves are the tank values $s_T \pm \delta/2$,

$$\begin{aligned} \beta_{up} &= \mathrm{clip}\left(\frac{s_{max} - s_T - |\delta|/2}{0.03},\ 0,\ 1 \right) \\ \beta_{dn} &= \mathrm{clip}\left(\frac{s_T - |\delta|/2 - s_{min}}{0.03},\ 0,\ 1 \right) \end{aligned} \tag{10.29}$$

and the same range is handed to the net-load controller as its output bounds. That is ordinary cascade anti-windup and it is the difference between a battery that recovers the moment the tanks refill and one that spends an hour unwinding an integrator. Section 1.4.4 treats tapered limits in the same terms, and Chapter 7 tapers its battery's limits over a wider band. The taper applies in automatic mode only, because in manual the operator is allowed to drive the plant past its own band, which is the point of the third scenario. There only the supply limit of Equation (10.10) stops the plant, at exhaustion rather than at the band.

The current limits are recomputed every scan rather than fixed, because the current that produces the rated charging power moves with charge state, temperature and membrane age. The lower limit is found by bisecting Equation (10.16) for the current density that produces $-P_R$, using sixty halvings on a bracket of $\pm 1.35 j_R$. The same limit clips the operator's current in manual mode, which is why the third scenario's $-70\,\%$ command is trimmed as the tank fills and the cell voltage rises.

Pause. Try Exercise 10.3 before reading on.

10.6 The Dispatch Layer

The supervisory layer answers the only question the regulatory loops cannot, which is when to spend the inventory. It uses twenty-four one-hour blocks, one decision variable per block, and a quadratic objective in the battery power u_k,

$$J = \sum_{k=1}^{N} \left[\frac{w_t (n_k - n_k^{*})^2}{P_R^2} + \frac{w_r (n_k - n_{k-1})^2}{P_R^2} - \frac{w_p \pi_k u_k}{\pi_{ref} P_R} \right] + w_f (s_N - s_f)^2 \tag{10.30}$$

with $n_k = f_k - u_k$, where f_k is the forecast net load in block k and s_N is the charge state at the end of the horizon predicted by the plan. Every term is quadratic or linear in u, so the gradient is affine and no line search is needed. The step size comes from a Gershgorin bound on the Hessian computed once per execution,

$$L = \frac{2 w_t}{P_R^2} + \frac{8 w_r}{P_R^2} + 2 w_f N c_T^2, \quad c_T = \frac{\Delta}{\eta_{plan} E_{cap}} \tag{10.31}$$

with a step size $\alpha = 1/L$, where the tracking term is diagonal, the ramp term has a spectral radius of at most four and the terminal term is a dense rank-one block whose largest eigenvalue is N times its diagonal entry. The layer then takes twenty projected-gradient steps with no convergence test. The cost per execution is therefore fixed and known in advance, which is what makes the layer run at interactive speed in a browser. The plant beneath it is already closed loop, and the only thing this layer can see is an integrator with a price on it, so a nonlinear model integrated over a horizon would buy nothing. The layer is the app's own and not the predictive control option of Section 1.6.3, which this app does not offer. Chapter 13 solves a box-constrained quadratic program of twice as many variables on the same fifteen-minute cadence.

The projection is where the physics lives. Each iteration restores feasibility in three passes. The box constraint clips every block to the converter rating, the ramp constraint clips forward from the power currently in force, and a causal forward pass on the charge state reduces any block that would spend energy an earlier block has not yet stored,

$$s_{k+1} = s_k + \frac{1}{E_{cap}} \begin{cases} -u_k \Delta / \eta_{plan} & u_k > 0 \\ -u_k \Delta \, \eta_{plan} & u_k \le 0 \end{cases} \tag{10.32}$$

with u_k reduced to whatever hits s_{min} or s_{max} exactly. One efficiency is used in each direction because the layer below already carries the real polarization and the plan is not permitted to pretend that it knows it. The forward pass is greedy feasibility restoration rather than an exact projection onto the reachable set, it is order N, and it is what a production dispatch code does. The planning capacity is Faraday capacity at the formal potential,

$$E_{cap} = \frac{c_V V_t F E^0}{3.6 \times 10^9} = 600.4\,\mathrm{MWh}. \tag{10.33}$$

The plan sets a reference and the net-load controller trims it. A one-hour block cannot follow a peak that is two hours wide, and the plan also steers toward its own terminal charge state, so the

Table 10.4 Effect of the Trim Authority on the Eighteen-Hour Dispatch Day of the Seventh Scenario

P_{trim} (MW)	Peak net (MW)	Off target (h)	Margin ($)	Final SOC
0	571.6	6.926	27 920	0.182
10	551.0	4.706	31 545	0.167
50	502.8	0.607	35 067	0.101
100	502.5	0.002	34 675	0.101

Table 10.5 Scenarios

#	Start, SOC	Architecture	Subject
1	00:00, 0.50	manual	which handle does what
2	00:00, 0.50	manual	the duck curve, battery idle
3	10:00, 0.30	manual	charging past the ceiling
4	00:00, 0.50	power	a fast power loop
5	16:00, 0.65	power	fixed, feedforward and trimmed flow
6	13:00, 0.40	net load	feedback and an unseen price event
7	06:00, 0.20	dispatch	the plan meets the midday belly
8	06:00, 0.20	dispatch	the terminal charge state
9	09:00, 0.30	dispatch	a cloud front the plan did not know
10	06:00, 0.30	dispatch	heat, crossover and lost capacity
11	04:00, 0.45	power	graded challenge, twenty hours

reference and the target part company. On the dispatch day of the seventh scenario the interpolated plan sits 13 MW above what the target asks for at the 19:27 peak and falls 56 MW short by 20:48 and 80 MW short by 21:00, as it holds charge for the end of its horizon. The base is interpolated across the block rather than held, because consecutive blocks on that day differ by a median of 38 MW and by up to 100 MW, and a block held flat and then stepped would put that discontinuity into the feeder every hour, which is a ramp the plan created rather than one it removed.

The trim authority is bounded to $\pm P_{trim}$, defaulting to 50 MW. Given the whole range the trim simply tracks the target and the plan becomes decoration, and the energy policy is then written by a loop that cannot see past the present instant. Table 10.4 shows what the bound is worth in both directions on the seventh scenario's day. With no trim the plan alone leaves the feeder 6.9 h off target and books $27 920. At the default bound it books $35 067 with 0.61 h off target. A 100 MW bound removes almost all of the off-target time and gives back $392, because the inventory is then spent on regulation rather than on price. Figure 10.4 catches the default trim at its bound late in the evening, when the plan has begun to hold charge back.

The forecast is deliberately imperfect and deliberately biased. A multiplicative error is drawn once per reset from the seeded generator and attached to the hour of the day rather than to the execution, so a forecast that is wrong at three in the afternoon stays wrong every time the day is re-planned. The sky disturbance is not in the forecast at all, which is the subject of the ninth scenario.

10.7 Scenarios and Measured Results

Table 10.5 lists the eleven scenarios. Each sets a clock start, an initial charge state, a control architecture and a tuning, and the first three run in manual mode. Table 10.6 gives the scorecard at the end of each scenario window with the shipped settings and no learner action. Pumping is the logged pumping energy and the other columns are the app's scorecard tiles. The ninth scenario is the only row that depends on the seeded generator, through its forecast error. Across five other seeds its peak net load ranges from 504 MW to 516 MW and its margin from $30 042 to $32 704.

Table 10.6 Scorecard at the End of Each Scenario Window

#	Window (h)	Peak net (MW)	Worst ramp (MW/min)	Off target (h)	Pumping (MWh)	Margin ($)
1	4	349.4	0.14	0.00	1.7	−160
2	24	585.9	2.74	10.74	2.2	−1468
3	12	586.5	2.74	4.80	8.9	−2290
4	3	420.8	11.66	1.50	1.1	−1269
5	6	585.7	6.81	4.71	7.2	21470
6	12	500.2	4.80	0.94	13.7	14035
7	18	502.8	5.55	0.61	17.0	35067
8	18	503.2	5.64	0.84	16.0	34740
9	14	503.3	7.36	2.91	14.5	32806
10	18	574.3	5.82	1.98	21.8	33852
11	20	618.3	3.50	10.73	81.3	−320

Table 10.7 Six Hours From 65 % Charge at an 85 MW Request, Each Flow Policy Held Throughout

Policy	Pumping (MWh)	Peak flow (m^3/s)	Worst margin	Margin ($)
fixed at 18 m^3/s	15.04	18.00	0.69	20720
feedforward	6.34	20.18	0.75	21598
feedforward and trim	13.47	30.00	1.01	20546

Peak flow and worst margin are taken at the 3.6 s physics step. The logged trend reads worst margins of 0.71, 0.76 and 1.03. Settings made before Run and carried through Reset give identical results.

10.7.1 The Day with Nothing Smoothing It

The second scenario runs twenty-four hours with the battery idle, and it establishes the numbers every later scenario is measured against. The net load reaches 585.9 MW in the evening, the worst ten-minute ramp is 2.74 MW/min, 63.9 MWh of solar is curtailed because the feeder exports at midday, and the feeder spends 10.74 h more than 25 MW away from the shaving target. The margin is −$1 468, which includes $766 of curtailment and the 15.42 MWh the idle plant buys for its pumps, its auxiliaries and its inverter. A plant that does nothing is not neutral on a feeder like this one.

> **TRY IT IN THE APP** **2 · The Duck, with Nothing to Smooth It**
>
> Load scenario 2 and press Run for the full twenty-four hours with the battery idle. Watch the feeder net load. It reaches 585.9 MW in the evening, its worst ten-minute ramp is 2.74 MW/min, and 63.9 MWh of solar is curtailed while the feeder exports at midday.

10.7.2 Flow Policy Is an Economic Decision

The fifth scenario requests 85 MW of discharge from 65 % charge at 16:00 and switches the flow policy from fixed to feedforward at 1.5 h and to feedforward with trim at 3.0 h. The power cannot be held past about 2.86 h, and at 3.0 h the tank is at 11.3 % and the power is down to 36 MW, so the margin recovery after the second switch is the current collapsing at the floor rather than the trim at work. Table 10.7 therefore holds each policy for the whole six hours, which a learner does by putting the policy back, with the 18 m^3/s flow for the fixed policy, after each scheduled switch. The differences are large and they do not all point the same way.

Feedforward alone cuts pumping energy by 58 % against fixed flow and raises the margin by $878. Only 1.2 MWh of that saving comes while the plant is discharging, because the flow the current calls for at 85 MW is 16.5 m^3/s to 20 m^3/s, close to the fixed 18 m^3/s. The other 7.5 MWh comes after the tanks reach their floor, when the current falls to zero and the feedforward drops to its 3 m^3/s minimum



Figure 10.2 Six hours from 65 % charge at an 85 MW request with each flow policy held throughout, from Table 10.7: (a) pumping energy, (b) worst transport margin and (c) margin. Hatched: fixed at 18 m^3/s. Solid: feedforward. Light: feedforward with the margin trim. Notice that the trim buys the only worst margin above one at more than twice the feedforward's pumping energy.

Table 10.8 Eighteen Hours From 06:00, Feedback Against Forecast Dispatch

Day	Architecture	Peak net (MW)	Worst ramp (MW/min)	Off target (h)	Margin ($)
clean	net load	500.2	2.61	0.000	35 083
clean	dispatch	502.8	5.55	0.607	35 067
scarcity	net load	500.2	2.61	0.000	−2 242
scarcity	dispatch	506.6	5.67	2.968	14 599

while the fixed policy goes on pumping 2.5 MW. Adding the transport-margin trim at its default setpoint of 3.0 raises pumping back to 13.47 MWh and gives back most of the gain. The trim is not failing. It is buying what it was asked to buy, which is transport margin, and it lifts the worst margin of the run from 0.75 to 1.01 by driving the pump to its 30 m^3/s limit as the tanks empty. Pumping power is cubic, so a flow excursion costs far more than its duration suggests. The lesson is that the margin setpoint is a purchased quantity rather than a safety setting, and at 3.0 on this plant at this charge state it is expensive. Figure 10.2 draws the three policies side by side.

The pump-wear disturbance sharpens the same point. Raising the hydraulic resistance by 70 % costs the fixed-flow plant 10.5 MWh of extra pumping, the trimmed plant 9.6 MWh and the feedforward plant 4.8 MWh. Pumping power at a given flow is proportional to the resistance, so each policy pays roughly seventy percent of its own pumping again. The trim does not avoid the loss, because it holds the margin rather than the pumping, and the feedforward pays least because it spends the empty-tank hours at its flow floor.

10.7.3 Feedback Dispatch Against Forecast Dispatch

The native sixth scenario runs twelve hours through the scheduled scarcity event. Feedback books $14 035 and dispatch $18 863, a $4 828 difference. The comparison the app is built around runs the same eighteen-hour day from six in the morning at twenty percent charge, once with a net-load controller that has no forecast and once with the dispatch layer. This is the seventh scenario, run as shipped and again with the architecture switched to the net-load controller before Run. Table 10.8 gives both, with and without a $480 per megawatt-hour scarcity event at three in the afternoon, set before Run with the scarcity price control. Each of these settings gives the same result to the dollar whether it is made before Run or carried through Reset. Figure 10.3 draws the same table as bars.

On the clean duck-curve day the forecast is worth nothing. The two architectures book $35 083 and $35 067, a difference of $16 or 0.05 %, and the feedback controller is the better of the two on every regulation metric, holding a worst ramp of 2.61 MW/min against 5.55 MW/min and spending

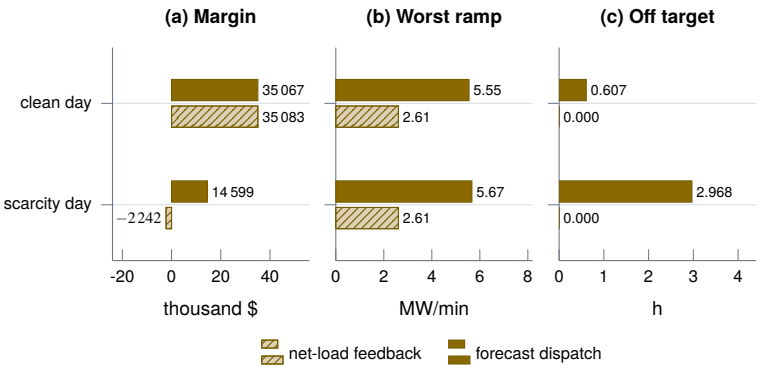

Figure 10.3 Eighteen hours from 06:00 under net-load feedback and forecast dispatch, from Table 10.8: (a) margin, (b) worst ten-minute ramp and (c) time off target, on the clean day and on the day with a scarcity hour at 15:00. Hatched: net-load feedback. Solid: forecast dispatch. Notice that on the clean day the two margins are equal while feedback regulates better, and that only the scarcity day separates them.

no time off target at all against 0.61 h. That result is half of the comparison and it carries as much weight as the other half. On an ordinary duck-curve day the highest net load and the highest price fall in the same hour, so a controller that watches only the net load is already pointing the right way.

🕮 **Pause.** Try Exercise 10.1 before reading on.

The scarcity day separates them completely. The feedback controller does nothing about a $480 per megawatt-hour event at three in the afternoon, because nothing it measures has changed. It does worse than nothing, in fact, since the midday net load is below its target and the controller is charging through the most expensive hour of the day. The same eighteen hours book −$2 242 under feedback and $14 599 under the dispatch layer, a swing of $16 841 produced by one hour of one day. A controller can act only on what it measures, so a decision that depends on something it does not measure has to be made somewhere else.

The dispatch layer pays for that capability in regulation quality. Its worst ramp is more than twice the feedback controller's on both days and it spends 3.0 h off target on the scarcity day against none for feedback, because the block schedule creates ramps at its own boundaries and the bounded trim cannot remove all of them. A plant that must meet a ramp contract and a price signal at once needs both layers, which is how the app is wired.

> ✦ **DID YOU KNOW? A Forecast Worth Sixteen Dollars**
>
> From the simulator: on the clean eighteen-hour day the net-load controller books $35 083 and the dispatch layer $35 067, and the controller with no forecast holds the worst ramp to 2.61 MW/min against 5.55 MW/min (Table 10.8). One scarcity hour turns the same pair into −$2 242 and $14 599.

10.7.4 The Terminal Charge State

An optimizer with a finite horizon and no terminal condition sells everything at the edge of its horizon, because energy left in the tank beyond the horizon is worth zero to the objective. The eighth scenario runs the seventh scenario's day with a terminal target of 55 % at a weight of 12, and the learner repeats it at weights of zero and 60. At midnight the plan's end-of-horizon charge state, the tank and the margin read 11.6 %, 10.1 % and $34 998 at weight zero, 18.4 %, 11.7 % and $34 740 at the shipped weight, where the net-load trim spends most of what the plan holds back, and 38.8 %, 26.2 % and $32 088 at the maximum. Holding a quarter of the tank for tomorrow morning costs $2 910

against the plan that sells everything, and it should. The terminal term is in charge-fraction units, so even the maximum weight does not bring the plan to its 55 % target against tracking and ramp terms summed over twenty-four blocks. The terminal condition is where the horizon's ignorance is priced, and every storage plant has one whether it is stated or not. Section 1.4.6 prices the same edge with a water value taken from the merit order of the forecast prices, and Chapter 11 plans its cavern with a rolling thirty-day water-value plan for the same reason.

10.7.5 Degradation That No Controller Reports

The tenth scenario runs a hot day with 34 °C ambient, 120 MW of extra air conditioning, crossover at three and a half times nominal and an imbalance of 9 % already present. The imbalance costs the most and says the least. The two half-cells no longer hold the same charge, so whichever reaches its limit first sets the capacity of the whole plant, and a nine percent imbalance takes 4.5 points of charge state off each end of the working band, cutting it from 480 MWh to 426 MWh. Crossover drives the imbalance up to 15.6 % by 19:00, while the tanks are above half charge, and back to 9.2 % by midnight as they empty. The cooler runs at full duty from 07:49, yet the electrolyte reaches 37.0 °C by 13:56 and 39.7 °C at 20:58, against a 32 °C setpoint and 0.3 °C short of the V(V) precipitation limit. A 26 °C setpoint from the start gives an identical run. The peak net load is 574.3 MW against 502.8 MW on the clean day. The power and flow loops hold their setpoints through all of it and nothing on the screen reports the lost capacity. It is a fault that degrades the plant while the control system reports success.

10.7.6 The Graded Challenge

The eleventh scenario opens twenty hours from four in the morning with weak tuning, fixed flow at 20 m^3/s, the battery on a hand-set power request, and four timed events including a cloudy midday, a hot afternoon and a pump losing its impeller. A run passes when, over the twenty hours, the net load stays within 50 MW of target over the last third, the peak net load stays below 560 MW, pumping stays below 45 MWh, the electrolyte stays within 15 °C to 40 °C and the margin ends above $25 000. The model parameters and disturbances must stay as given: changing one voids the attempt, as does an extra learner disturbance event when it fires. The default configuration books −$320, reaches 618.3 MW of peak net load, spends 81.3 MWh on pumping and misses the target by up to 86.9 MW in the evening, so it fails every criterion except temperature. The suggested answer is the dispatch layer, feedforward flow with trim at a margin setpoint of 2.6, a 28 °C temperature setpoint chosen in the morning and a 32 % terminal target. It books $34 713, holds the peak to 535.8 MW, spends 25.1 MWh on pumping and keeps the evening error to 47.0 MW, 3.0 MW inside the band and 4.6 MW to 6.9 MW inside with five other seeds of the forecast error. Settings made before Run and carried through Reset give identical results. The structural decisions are worth $35 033 and 56.2 MWh, and no gain change in the default structure comes close.

10.8 Technology in Practice

The instruments and actuators of Section 10.3 are, in a plant, instrument transformers with a power transducer, a temperature sensor, a charge-state computation, pumps on variable-speed drives, a cooler and a power converter. Figure 10.5 draws the chapter's loops with that hardware in place of the schematic's symbols.

What Measures It

The net load and the battery's own power are electrical quantities. They are read through instrument transformers and turned into megawatts by a transducer. Utilities collect such readings from dispersed assets with supervisory systems that display them to an operator in near real time [16]. The charge state has no direct sensor in the chapter. Its estimator SI is a filter with a 72 s time constant, and a charge state read at the stack outlet is a biased estimate of the tanks whenever current flows.

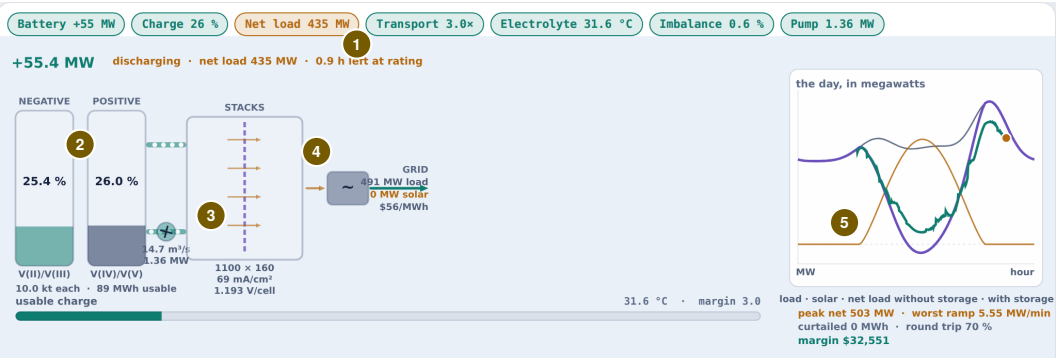

Figure 10.4 The app at 21:09 on the seventh scenario, 15.15 h into an eighteen-hour dispatch day that began at six in the morning at twenty percent charge. The plan filled the tanks to 86.8 % through the midday belly and has spent them into the evening neck, so the tank holds 25.7 %, 89 MWh above the floor (2). The plan's own request has fallen to 5.4 MW as it holds charge for the end of its horizon, and the net-load trim sits at its 50 MW bound, so the battery delivers 55.4 MW through the converter (4) and the feeder net load reads 435.4 MW (1). The trim holds the transport margin at its setpoint of 3.0 with 14.7 m³/s of flow and 1.36 MW of pumping (3). The scorecard line under the day inset (5) shows a peak net load of 503 MW and a margin of $32 551.

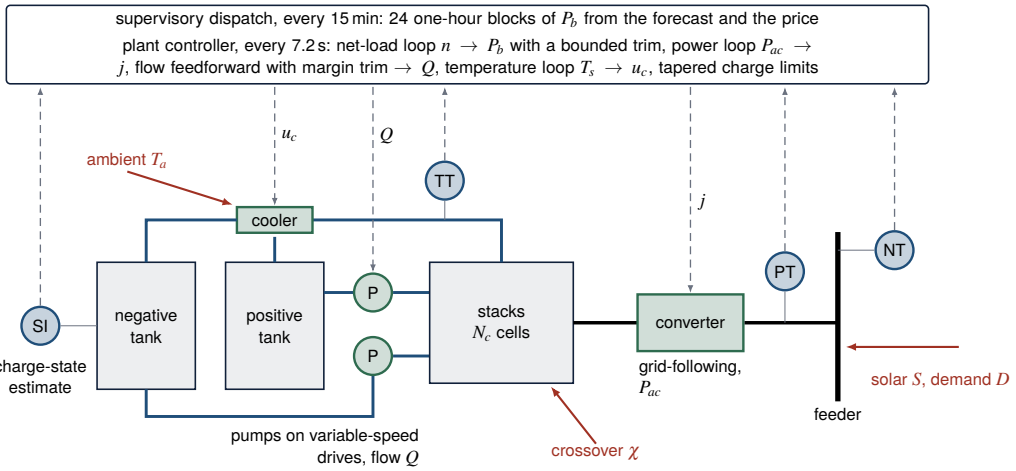

Figure 10.5 The chapter's loops drawn with their hardware. The power transducer PT and the feeder measurement NT read through instrument transformers, TT reads the stack outlet and SI estimates the charge state. Every 7.2 s the plant controller sets the converter current j, the pump flow Q and the cooler duty u_c, and every 15 min the dispatch layer above it re-plans the battery power. Crossover acts on the stacks, ambient temperature on the cooler, and solar and demand on the feeder. Notice that no device measures the energy left in the tanks directly.

TECHNOLOGY **Sensor: Instrument Transformers and a Power Transducer**

A potential transformer steps the line voltage down to about 120 V, and a current transformer steps the line current down to about 5 A at full load [17]. A metering-class current transformer rated 0.3B1.8 is accurate to 0.3 % of its ratio while it drives a 1.8 Ω burden, and a larger burden makes the signal sag, so the power reads low [17]. An open secondary on an energized current transformer can develop an extremely high

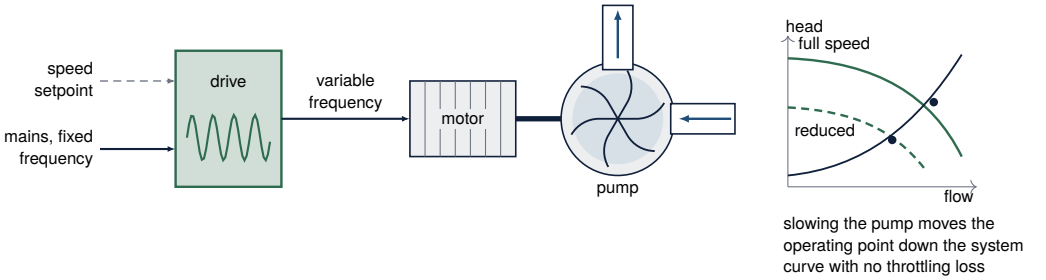

Figure 10.6 Principle of a variable-speed drive on a centrifugal pump. The drive changes the motor's speed, and the pump curve moves along one system curve. Notice that a lower speed reaches a lower flow without a throttling valve, which is how the feedforward policy of Section 10.5 saves pumping energy.

voltage, which is why its circuits carry shorting switches [17]. The chapter's transducer PT is a first-order filter of 8 s, with a noise disturbance of up to 6 MW.

What Moves It

Three final elements act on the plant. The converter sets the current. In the common grid-following design it estimates the angle of the grid voltage with a phase-locked loop and regulates the current it injects, so for fixed power commands it behaves as a constant real and reactive power source [18]. That is why the chapter's power loop is easy: the converter follows its command with a 4 s lag. The cooler is the second element, with a 60 s lag, and the pumps are the third and the one with an economic trade. Figure 10.6 shows the principle of their drive.

TECHNOLOGY Final Element: Electrolyte Pumps on Variable-Speed Drives

A variable frequency drive sets the speed of an alternating-current motor by varying the frequency of its supply, and its torque by varying the voltage and current [17]. On a flow battery it replaces the throttling valve, which in the NASA tests varied the flow without saving any pumping energy [3]. The same report found the small chemically resistant pumps of kilowatt systems only 5 to 20 % efficient [3]. Pumping power is cubic in flow in Equation (10.14), so a worn pump or fouled piping shows on the trend as more pumping power at the same flow. The chapter's pumps follow their command with a 20 s lag between 2 and 30 m^3/s.

Where the Controller Runs

The four loops of Table 10.3 run in the plant's own controller, and the dispatch layer runs above them in supervisory software that also gathers the feeder data. Supervisory and lower-level loops of this kind run with cycle times from milliseconds to minutes [16]. In the chapter the regulatory loops execute every 7.2 s on a 3.6 s physics step, and the dispatch layer re-plans every 15 min. The tapered limits of Section 10.5.3 belong to the battery management layer between the two.

IN PRACTICE **Spill Containment, Permits and the Queue**

A water-based electrolyte cannot burn and has no thermal runaway, so the main operating risk of a flow battery is a leak, handled with containment and neutralizing spill kits as NFPA 855 describes [19]. A battery plant also answers to the local authority having jurisdiction, often the fire department, whose setback and siting rules vary from place to place. Following current codes and standards and engaging that authority early minimizes delays, and an interconnection request typically takes over three years to complete [19].

What Is Modernizing It

Three developments touch this plant. The first is the pump drive. The NASA report already pointed to a variable-frequency motor and pump as the way to make variable flow efficient [3], and the chapter's feedforward flow schedule is that idea with a modern drive. The second is the converter. Grid-forming inverters act as voltage sources and can black-start a power system. Replacing synchronous machines with them on the bulk grid is a 10 to 30 year prospect that first needs protection, controls, interoperability and standards [18]. The third is the network. Operational systems increasingly resemble information systems as they adopt standard computers, operating systems and protocols for remote access, which leaves them less isolated from the outside world and in greater need of security [16]. A plant dispatched from a forecast depends on that network every fifteen minutes.

10.9 Limitations

The feeder, the solar profile and the price curve are synthetic. They are shaped like a spring day on a California distribution system and no claim is made that they reproduce one. Any conclusion about how much a forecast is worth on a real system requires real data, and the numbers here should be read as the shape of an answer rather than as its magnitude.

OPEN CHALLENGE **Proving Long Life at Scale**

A 2024 national laboratory assessment put a 100 MW, ten-hour vanadium flow battery at about $446 per kilowatt-hour installed, against about $379 for a lithium iron phosphate system of the same size. The electrolyte allows projected lives of 20 years or more, but few plants in the United States have run long enough to prove it, and some investors wait for operating data from the first deployments [19]. Progress would be successful demonstration plants, which the assessment names as the key to confidence and lower cost [19].

The dispatch layer runs a fixed twenty projected-gradient iterations with no convergence test, so the plan is a bounded-effort approximation to the solution of Equation (10.30) rather than the solution. Receding-horizon re-planning every fifteen minutes makes the approximation adequate, and it also hides it. A student who reads the objective and expects the plan to be optimal will find small inconsistencies between successive plans that the framework does not explain.

The electrochemical model is a lumped single-cell model scaled by a cell count, so everything that makes a real stack differ from the sum of its cells is missing. Shunt currents, cell-to-cell voltage distribution, electrode compression, flow-field maldistribution and gas evolution at the extremes of charge are all absent. The concentration overpotential in Equation (10.4) is a limiting-current form with a flow exponent fitted to nothing in particular, and it should be read as the right qualitative shape rather than as a prediction.

Crossover is one scalar where the physical process is four coupled diffusion fluxes at different rates. The model reproduces the operational consequence, which is capacity that disappears without an alarm, and not the composition drift that a plant operator would actually measure through electrolyte sampling.

Round-trip efficiency is reported only when the run is close to a closed cycle, because an uncorrected ratio of energy out to energy in exceeds unity whenever a run ends emptier than it started. The

application shows a dash instead of a number when the stored-energy change exceeds forty percent of the energy moved, which is honest and occasionally frustrating.

> **OPEN CHALLENGE** **Where the Vanadium Comes From**
>
> The same assessment rated vanadium supply a moderate risk, because the ore and the processed precursor are both geographically concentrated and the price is volatile. China holds 39 % of global vanadium reserves and leads production, largely from the slag of steelmaking, and supply deficits are predicted through 2030 if flow batteries are deployed widely [19]. Progress would be chemistries built from abundant elements such as iron or zinc, which the assessment rates low risk [19].

10.10 Conclusion

A vanadium flow battery separates the purchase of power from the purchase of energy, and the control structure has to respect that separation. The power loop on a 100 MW converter reaches 63 % of a 60 MW step in 18 s and 90 % in 83 s, and offers nothing to a control course beyond a demonstration of how easy it is. The flow loop carries a real trade-off, because pumping is cubic and the loss-minimizing flow of 8.92 m^3/s at rated current is well below the 15 m^3/s design flow and moves with charge state. The temperature loop is slower than the day it is dispatched over, so its setpoint is a planning decision. Above all three, the question of when to spend 480 MWh of usable inventory is the one that decides whether the plant is worth what it cost.

The measured comparison between feedback dispatch and forecast dispatch is the result worth carrying away, and it has two halves. On a clean duck-curve day the two are within 0.05 % of each other and the feedback controller regulates better, holding 2.61 MW/min of worst ramp against 5.55 MW/min. On a day with a price event away from the evening peak the same two architectures book −$2 242 and $14 599. A forecast is not worth having because it is a forecast. It is worth having exactly when the quantity that decides the right action is one the feedback path cannot measure, and naming that condition is more useful to a student than any comparison of architectures in the abstract.

> **KEY POINTS**
>
> - A flow battery's energy is the liquid in its tanks and its power is the membrane in its stacks.
> - The power loop reaches 63 % of a 60 MW step in 18 s, so the hard question is when to spend the inventory.
> - The loss-minimizing flow of 8.92 m^3/s at rated current is well below the 15 m^3/s design flow and moves with charge state.
> - Feedback and forecast dispatch are within 0.05 % on a clean day, and a price event splits them to −$2 242 and $14 599.
> - A forecast is worth having when the quantity that decides the right action is one the feedback path cannot measure.
>
> **Where else this applies.** A stored inventory scheduled against a price over a day also sets the plan in Volume 9, Chapter 6 and Volume 9, Chapter 11, and Chapter 11 prices its terminal state with a water value.

References

[1] A. Z. Weber, M. M. Mench, J. P. Meyers, P. N. Ross, J. T. Gostick, and Q. Liu, "Redox flow batteries: a review," *Journal of Applied Electrochemistry*, vol. 41, no. 10, pp. 1137–1164, 2011.

[2] J. B. Rawlings, D. Angeli, and C. N. Bates, "Fundamentals of economic model predictive control," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 2012, pp. 3851–3861.

[3] N. Hagedorn, M. Hoberecht, and L. Thaller, "NASA-redox cell-stack shunt current, pumping power, and cell-performance tradeoffs," NASA Lewis Research Center, Cleveland, OH, USA, Tech. Rep., 1982.

[4] M. Skyllas-Kazacos, M. Rychcik, R. G. Robins, A. G. Fane, and M. A. Green, "New all-vanadium redox flow cell," *Journal of The Electrochemical Society*, vol. 133, no. 5, pp. 1057–1058, 1986.

[5] M. Rychcik and M. Skyllas-Kazacos, “Characteristics of a new all-vanadium redox flow battery,” *Journal of Power Sources*, vol. 22, no. 1, pp. 59–67, 1988.

[6] M. Skyllas-Kazacos, M. H. Chakrabarti, S. A. Hajimolana, F. S. Mjalli, and M. Saleem, “Progress in flow battery research and development,” *Journal of The Electrochemical Society*, vol. 158, no. 8, p. R55, 2011.

[7] A. A. Shah, M. J. Watt-Smith, and F. C. Walsh, “A dynamic performance model for redox-flow batteries involving soluble species,” *Electrochimica Acta*, vol. 53, no. 27, pp. 8087–8100, 2008.

[8] A. Tang, J. Bao, and M. Skyllas-Kazacos, “Thermal modelling of battery configuration and self-discharge reactions in vanadium redox flow battery,” *Journal of Power Sources*, vol. 216, pp. 489–501, 2012.

[9] M. Skyllas-Kazacos and L. Goh, “Modeling of vanadium ion diffusion across the ion exchange membrane in the vanadium redox battery,” *Journal of Membrane Science*, vol. 399-400, pp. 43–48, 2012.

[10] A. Tang, J. Bao, and M. Skyllas-Kazacos, “Studies on pressure losses and flow rate optimization in vanadium redox flow battery,” *Journal of Power Sources*, vol. 248, pp. 154–162, 2014.

[11] Y. Li, X. Zhang, J. Bao, and M. Skyllas-Kazacos, “Control of electrolyte flow rate for the vanadium redox flow battery by gain scheduling,” *Journal of Energy Storage*, vol. 14, pp. 125–133, 2017.

[12] A. Tang, J. McCann, J. Bao, and M. Skyllas-Kazacos, “Investigation of the effect of shunt current on battery efficiency and stack temperature in vanadium redox flow battery,” *Journal of Power Sources*, vol. 242, pp. 349–356, 2013.

[13] P. Denholm, M. O’Connell, G. Brinkman, and J. Jorgenson, “Overgeneration from solar energy in california: A field guide to the duck chart,” National Renewable Energy Laboratory, Golden, CO, USA, Tech. Rep., 2015.

[14] X. Luo, J. Wang, M. Dooner, and J. Clarke, “Overview of current development in electrical energy storage technologies and the application potential in power system operation,” *Applied Energy*, vol. 137, pp. 511–536, 2015.

[15] T. Ouyang, M. Zhang, P. Qin, and X. Tan, “Flow battery energy storage system for microgrid peak shaving based on predictive control algorithm,” *Applied Energy*, vol. 356, p. 122448, 2024.

[16] 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.

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

[18] Y. Lin, J. H. Eto, B. B. Johnson, J. D. Flicker, R. H. Lasseter, H. N. Villegas Pico, G.-S. Seo, B. J. Pierre, and A. Ellis, “Research roadmap on grid-forming inverters,” National Renewable Energy Laboratory, Golden, CO, Technical Report NREL/TP-5D00-73476, 2020.

[19] A. Hollas, A. Tuan, V. Viswanathan, and I. Ragazzi, “Adoption readiness level assessment of redox flow batteries,” Pacific Northwest National Laboratory, Richland, WA, Report PNNL-36780, 2024.

10.11 Exercises

10.1 ○ Concept check · 2 min

On the clean eighteen-hour day of scenario 7 the net-load feedback controller books \$35 083 and the dispatch layer, which has a forecast, books \$35 067. Why is the forecast worth nothing on that day?

(a) The twenty fixed projected-gradient steps stop well short of the optimum of the dispatch objective

(b) The terminal charge weight makes the plan hold back energy that it would otherwise sell in the evening

(c) The highest net load and the highest price fall in the same hour, so feedback on the net load already points the right way

(d) The shaving target is built from the price curve, so the net-load controller sees the price through its setpoint

10.2 ○ Concept check · 2 min

In scenario 3 the plant is reversed to full discharge current at the ceiling alarm with the flow held at 12 m^3/s. The transport margin reaches its starvation line of 1.15 at about 22 % tank charge, with 88 MW still flowing. What sets the charge state at which this happens?

(a) The battery management taper, which derates discharge over the last three points above the floor

(b) The Faradaic supply limit, which cuts the current once the flow cannot carry enough reactant

(c) The converter rating, which trims the current as the open-circuit voltage falls with the charge

(d) The limiting current, which is proportional to the availability of the consumed species and rises with the flow

10.3 Concept check · 1 min

The charge and discharge limits of Equation (10.29) are handed to the net-load controller as its output bounds. What does that buy?

(a) The controller's integral does not wind up at the limit, so the battery responds as soon as the tanks have room again

(b) The plant can no longer charge past its working ceiling in manual mode, which protects the positive electrolyte

(c) The dispatch plan no longer needs its causal pass on the charge state, because the bounds keep it feasible

(d) The loop changes its action near the floor, so it charges the tanks back toward the middle of the band

10.4 Theory and calculation · 6 min

Use the parameters of Table 10.2: $c_V = 1600\,\mathrm{mol/m^3}$, $V_t = 10\,000\,\mathrm{m^3}$ a side, $V_c = 150\,\mathrm{m^3}$ of hold-up, $E^0 = 1.40\,\mathrm{V}$, $Q_0 = 15\,\mathrm{m^3/s}$ and a working band from 0.10 to 0.90.

(a) Find the energy of the tanks between empty and full from Equation (10.33), and the energy inside the working band. *Checkpoint:* 600.4 MWh and 480.3 MWh.

(b) Find the energy held in the stack hold-up and the time the electrolyte takes to pass through the hold-up and through one tank at design flow. *Checkpoint:* About 9.0 MWh, 10 s and 11.1 min.

(c) An imbalance $\delta = 0.09$ is present. Using the half-cell charge states $s_T \pm \delta/2$ of Equation (10.29), find the energy left inside the working band. *Checkpoint:* About 426 MWh.

10.5 Theory and calculation · 7 min

The plant has $N_c = 176\,000$ cells of $A = 4000\,\mathrm{cm^2}$ with a rated current density $j_R = 0.12\,\mathrm{A/cm^2}$. At half charge and design flow, 100 MW at the connection point takes 106.4 % of rated current on discharge and 72.2 % on charge. Take $F = 96\,485\,\mathrm{C/mol}$ and $c_V = 1600\,\mathrm{mol/m^3}$.

(a) Find the plant current at rated current density. *Checkpoint:* 84.48 MA.

(b) The tank and the hold-up move together once the transient of the stacks has passed, so the vanadium in $V_t + V_c = 10\,150\,\mathrm{m^3}$ changes state together. Find the rate of change of the tank charge state in percent per minute at 100 MW of discharge and of charge. *Checkpoint:* 0.344 %/min and 0.234 %/min, within 0.002.

(c) Find how long the working band of 80 points lasts at each rate, and say why charging takes less current than discharging for the same power. *Checkpoint:* 3.87 h and 5.71 h, within 0.05 h.

10.6 Theory and calculation · 6 min

Pumping power follows Equation (10.14) with $k_Q = 155\,\mathrm{Pa\,s^2/m^6}$, $\eta_p = 0.72$ and $w_p = 1$.

(a) Find the pumping power for both circuits at the design flow of 15 $\mathrm{m^3/s}$ and at the pump limit of 30 $\mathrm{m^3/s}$. *Checkpoint:* 1.45 MW and 11.6 MW.

(b) From Equation (10.10) with no imbalance, find the flow per circuit that just supplies rated current at half charge. Compare the design flow and the loss-minimizing flow of 8.92 $\mathrm{m^3/s}$ with it. *Checkpoint:* 1.09 $\mathrm{m^3/s}$, about 13.7 and 8.2 times.

(c) Over six hours of scenario 5 the fixed policy at 18 m^3/s pumps 15.04 MWh. Predict the extra pumping when the hydraulic resistance rises by 70 %, and say why the prediction is exact for this policy and only approximate for the other two. *Checkpoint:* 10.5 MWh.

10.7 App challenge · 7 min

Load scenario 2, *The duck, with nothing to smooth it*, and change no setting. Run the full 24 h at speed 5 with the default seed, then export the CSV file.

(a) Read the Peak net load, Worst ramp, Curtailed and Margin tiles. *Checkpoint:* 586 MW, 2.74 MW/min, 64 MWh and −$1 468.

(b) In the last row of the file, at `time_h` 23.987, read `metric_curtail_MWh` and `metric_in_MWh`. With the curtailment charge $c_{curt} = \$12$ per megawatt-hour, find the share of the margin's loss that the curtailment accounts for, as a percentage of $1 468, within one point. *Checkpoint:* 63.87 MWh and 15.42 MWh; $766, about 52 %.

(c) From the columns `net_raw_MW` and `clock_h`, find the clock times of the lowest and the highest raw net load, within two minutes. *Checkpoint:* 12:28 and 19:27.

10.8 App challenge · 9 min

Load scenario 7, *The dispatch layer meets the belly*, with the default seed. Each run lasts 18 h at speed 5. Read the Margin, Worst ramp and Off target tiles at the end of each run, and give dollar values to within $2.

(a) Run the scenario as loaded, on the Dispatch architecture. *Checkpoint:* $35 067, 5.55 MW/min, 0.61 h.

(b) Pause, Reset, set Power architecture to Ramp before Run and run again. Find what the forecast is worth on this day. *Checkpoint:* $35 083, 2.61 MW/min, 0.00 h. The forecast is worth −$16.

(c) Reload the scenario, set the Scarcity price event slider to $480 per megawatt-hour before Run, and run on Dispatch. Then Reset, set Power architecture to Ramp and run again. Find the swing between the two margins and explain it. *Checkpoint:* $14 599 and −$2 242, a swing of $16 841.

10.9 App challenge · 8 min

Load scenario 8, *The terminal charge state, and why it is not optional*, with the default seed. It runs the day of scenario 7 with a terminal target of 55 % at a weight of 12. Switch to the Advanced level. Run 18 h at speed 5, then Pause, Reset, set Dispatch: terminal charge weight to 0 and run again, then Reset, set it to 60 and run a third time. After each run read the Charge and Margin tiles and, from the exported CSV file, `soc_plan_pct` in the last row (`time_h` 17.986).

(a) Record the planned end-of-horizon charge, the tank charge and the margin for the three weights, the charges within 0.2 points and the margins within $2. *Checkpoint:* Weight 12: 18.4, 11.7, $34 740. Weight 0: 11.6, 10.1, $34 998. Weight 60: 38.8, 26.2, $32 088.

(b) Find what holding charge for the next morning costs at the maximum weight, against the plan with no terminal term. *Checkpoint:* $2 910.

(c) Explain why even the maximum weight leaves the plan well short of its 55 % target. *Checkpoint:* The terminal term is in charge-fraction units.

10.10 Application area · 4 min

Scenario 10 runs a hot day at an ambient of 34 °C. The cooler goes to full duty in the morning, the electrolyte still reaches 39.7 °C against a 32 °C setpoint, and a 26 °C setpoint from the start gives an identical run.

(a) From $C_T = 2.9 \times 10^{10}\,\mathrm{J/°C}$, $UA_{amb} = 3.2 \times 10^{5}\,\mathrm{W/°C}$ and $UA_c = 1.9 \times 10^{6}\,\mathrm{W/°C}$, find the time constant of the electrolyte temperature at full cooler duty. *Checkpoint:* 3.63 h.

(b) Using Equation (10.17), explain why no setpoint below the ambient can change this run, and why the setpoint is held in a band rather than kept as low as possible. *Checkpoint:* The cooler rejects heat to the ambient air.

10.11 Application area · 3 min

A planner proposes to use the chapter's $16 841 swing between feedback and forecast dispatch to justify buying a price-forecast service for a real 100 MW flow battery. Name two limitations stated in the chapter that stop the number from transferring, and the data that would be needed to replace it.

10.12 Review the agent · 5 min

An AI assistant reviews scenario 5, *Fixed flow, feedforward flow, and the pump bill*, and writes: "Run the flow on FF + trim at the default transport margin setpoint of 3.0. The trim adds flow only when the margin falls, so it pumps less than feedforward alone and pays for itself, and when a pump loses its impeller the trim finds that it needs less flow and recovers most of the loss." Find the flaw and say how you would check it in the app.

CHAPTER 11

A Salt-Cavern Air Storage Simulator for Inventory Policy Under a Capacity Commitment

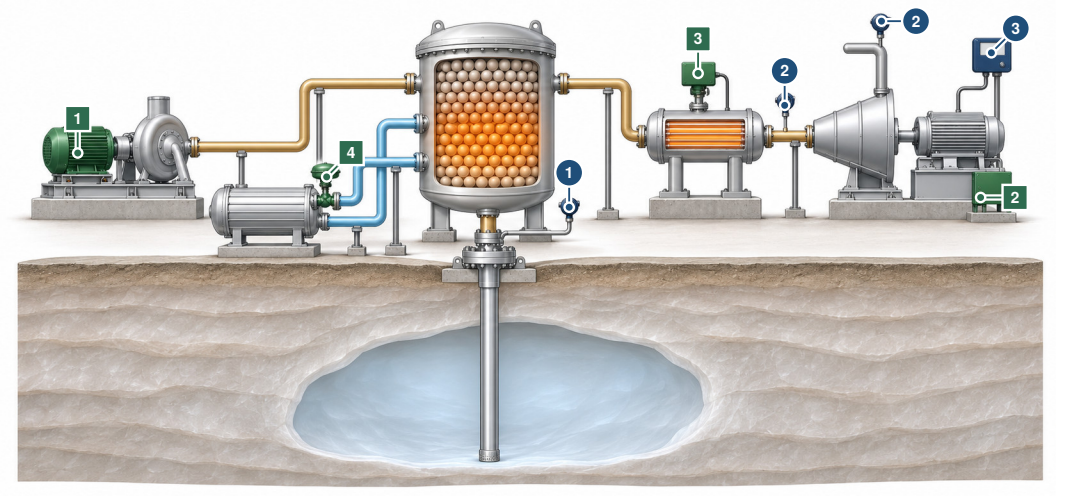

The controller moves 1 compressor and its motor 2 expander load on the generator 3 supplemental heater 4 heat-rejection valve

The controller reads 1 cavern pressure at the wellhead 2 thermometers at the expander inlet and exhaust 3 power meter

Notice that the air passes through the heat store on the way in and on the way out, so the store can run cold while the cavern still holds most of its air. Illustration, not to scale.

IN PLAIN TERMS **Storing Spare Wind as Squeezed Air**

The main idea. When wind power is spare, a compressor squeezes air into a huge hollow in underground salt. Later the air flows back out through a turbine, a machine that the rushing air spins to drive a generator.

The challenge. The air must be warmed before it reaches the turbine. The heat saved from squeezing it runs out in hours, while the air lasts for days. A plant can break its winter promise with most of its air still underground.

Why it matters. The plant is paid to deliver power on the worst winter days, so it must save heat as well as air.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the two history boxes of Section 11.2, the try-it and fact boxes of Section 11.4, the fact and try-it boxes of Section 11.7, and the key points
Study	45 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 11.7, Section 11.8, Section 11.9 and the open challenge

The control question A compressed-air plant stores mass in a salt cavern but sells a capacity commitment in megawatts. Which inventory, the gas or the heat, decides whether a winter commitment is kept, and what must a reserve count to protect it?

Learning objectives

- Compute the deliverable power of a fixed-geometry expander from the cavern pressure and the store temperature.
- Compare the gas and heat inventories by capacity, time constant and duration at rating.
- Explain the budget-and-trim split, the capacity floor and the back-calculation limit of an hourly scan.
- Compare three inventory policies and explain why the heater, not the reserve, removes the default.
- Judge a reserve policy against a commitment written in megawatts.

Prerequisites Sections 1.4.4 (storage as an inventory), 1.4.6 (water value), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios and data). Chapter 10 is a useful first store with two inventories.

The app https://apmonitor.com/apps/salt-cavern-caes/

See also Volume 9, Chapter 7 (a reservoir held against a forecast), Volume 9, Chapter 6 (a price of heat), Volume 4, Chapter 14 and Volume 4, Chapter 8 (an inventory defended at a floor) and Volume 4, Chapter 2 (the surge override).

11.1 Introduction

The phrase seasonal storage does a great deal of quiet work. A salt cavern of 8×10^6 m^3 cycled between 45 bar and 75 bar holds 253.8 kt of air, which is 24.6 GWh of electricity through the expander train modeled here. That number reads as seasonal and it is 81.9 h at the expander rating, which is three and a half days. The plant is therefore not carrying summer wind into winter. It is being asked to be in the right state on the worst fortnight of the year, while every profitable day between now and then argues for spending the inventory, and nothing in the physics decides that argument.

Underneath the scheduling question is a physical one that is easier to miss and harder to recover from. Stored energy is mass, and deliverable power is a fixed-geometry expander's swallowing capacity multiplied by the specific work available at the cavern pressure. Both fall as the cavern empties, and they fall together. At 75 bar the plant can deliver 282.6 MW. At 50 bar it still holds 3.99 GWh, nearly a day of output at the 175.2 MW it can then deliver. A capacity commitment is sold in megawatts and an inventory policy is written in megawatt-hours, and the constraint that binds first is the one nobody wrote down.

The misconception the application corrects is that a storage plant has a state of charge in the sense a battery does (Section 1.4.4). It has two, the gas in the ground and the heat in the packed bed, they empty on different time scales, and neither of them is the quantity a commitment was made against. The class of problem is inventory management with a state-dependent delivery constraint and a horizon longer than any forecast. Hydropower with a seasonal reservoir, a gas storage field with a contracted withdrawal profile, a refinery tank farm under a shipping schedule and a hospital oxygen system under a surge commitment all share the shape, and the reservoir of Volume 9, Chapter 7 holds empty volume against a storm on the same terms. The classical answer is a marginal water value on the inventory [1], which Section 1.4.6 derives by merit order, and the modern framing is economic model predictive control with the risky part expressed as a constraint [2].

This chapter works with an interactive salt-cavern storage simulator. It gives the model, the control structure and the measured behavior in enough detail that a reader can rebuild them, and it uses them to show that a capacity commitment is defended by the heat in the store as much as by the gas in the ground. The model has ten states for the cavern, the rock and the thermal store, and it computes a capacity floor by bisection on the delivery physics at every scan. The central measurement compares three inventory policies and a seasonal percentage reserve over a simulated hundred and fifty days, and it finds that the heat available to the expander, more than the gas in the ground, decides whether the commitment is kept. With the supplemental heater armed every policy keeps it, at a cost

of 13 GWh to 27 GWh of heat. With the heater off, the threshold rule and the rolling plan default on 0.56 GWh and 0.47 GWh, the operator curve that holds the most gas defaults on 2 MWh, and raising the reserve peak from 25 % to 85 % cuts the plan's default from 1.15 GWh to 0.13 GWh for $1.70 M without meeting the commitment.

Section 11.2 traces the technology from Huntorf. Section 11.3 gives the model, with the deliverability relation at its center. Sections 11.4 and 11.5 give the open-loop character and the regulatory scheme, including the capacity floor and the two overrides. Section 11.6 gives the three inventory policies and the seasonal reserve that bounds them. Section 11.7 reports the scenarios, the policy comparison, the reserve sweep, the drought and the graded challenge. Section 11.8 describes the devices and platforms of a real plant, Section 11.9 states the limitations and Section 11.10 draws the conclusions.

11.2 From Huntorf to a Capacity Market

Compressed-air energy storage has exactly two commercial ancestors and both are still running. The Huntorf plant in Lower Saxony started in 1978 with two solution-mined caverns of about 150 000 m^3 each, operated between 43 bar and 70 bar, and the engineering account of those caverns is the founding document of the subject [3]. Huntorf and the later McIntosh plant in Alabama are diabatic, which means the heat of compression is rejected to atmosphere and natural gas is burned to reheat the air before expansion. That choice is the reason the round-trip efficiency of the first generation is usually quoted at around 42 %, and it is also why the technology is described in the review literature as a hybrid rather than as storage [4].

> ⧖ HISTORY **Huntorf, Built to Restart a Grid (1978)**
>
> The first compressed-air storage plant was completed in 1978 near Bremen in Germany. It was built to provide black-start power, the power that nuclear units near the North Sea need to restart after a complete grid outage, and cheap peak power [5]. Its caverns were first sized for two hours of rated output, later stretched to three, and the plant has been used more and more to balance the growing wind output of northern Germany. Its high-pressure expansion stage came from steam turbine practice, because gas turbines of the time did not suit that pressure range. A recuperator to preheat the air with exhaust heat was left out to keep the start-up short [5].

Adiabatic designs keep the heat. Compressing air to 60 bar in two stages at a polytropic efficiency of 85 % leaves it at 557 K, and a packed bed of rock or ceramic can hold that heat for weeks and give it back before expansion. The concept removes the fuel and adds a second inventory, and the engineering literature on packed-bed stores for this duty is recent and active [6]. A low-temperature variant that accepts a smaller store and a modest fuel burn has been proposed as the practical middle ground [7], which is precisely the trade the simulator's supplemental heater is meant to expose.

> ⧖ HISTORY **McIntosh, a Day of Storage (1991)**
>
> A second plant, of 110 MW, began operating on a salt dome in southwestern Alabama in 1991. It was designed for 26 hours of generation at full power from a single cavern of 560 000 m^3 worked between 45 and 74 bar [5]. Unlike Huntorf it carries a recuperator, which recovers exhaust heat and cuts the fuel burned at full load by about 22 %. Early outages were cured by changes to its combustors, and over ten years of operation its average starting reliability was 91.2 % for generation and 92.1 % for compression [5].

The cavern itself is a thermodynamic object and not a tank. Air entering at aftercooled temperature warms it, withdrawal cools it, and the surrounding salt exchanges heat on a time scale of weeks. Models of that behavior were developed for Huntorf and then generalized, and the review of cavern thermodynamic and hydrodynamic response collects them [8–10]. The mechanical side is a separate literature. Salt creeps, the cavern closes slowly at low pressure, and both the minimum working

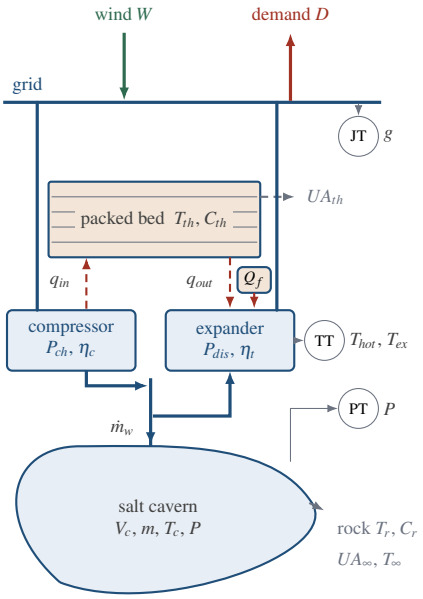

Figure 11.1 Plant arrangement and the symbols used in the equations. Air paths are thick solid arrows and heat paths dashed arrows. A single wellbore carries air in both directions between the compressor train and the cavern, and the rock around the cavern is a thermal capacitance of its own. The packed bed takes the heat of compression on the way in and returns it on the way out, and a supplemental heater Q_f reheats the expander inlet air directly when the bed is too cold. PT measures cavern pressure, TT the expander inlet and exhaust, and JT the grid exchange that the delivery loop controls.

pressure and the permitted rate of pressure change come from a creep analysis rather than from a process requirement [11].

What changed recently is the question being asked. When storage was arbitrage the operating problem was a price forecast. When storage is sold into a capacity market the operating problem is a state constraint, because the plant has undertaken to be able to deliver a stated power when called, and being unable to is a default rather than a trade. The system-level studies of what duration is worth having make the same point from the outside, which is that the value of long-duration storage is concentrated in a small number of hours that cannot be identified in advance [12, 13]. The plant modeled here sits at that junction.

11.3 Process Model

The simulated plant is 850 MW of wind, a 250 MW two-stage intercooled compressor, a 300 MW two-stage reheat expander, a packed-bed thermal store of 6.1×10^{10} J/K, and 8×10^6 m^3 of solution-mined cavern worked between 45 bar and 75 bar. Figure 11.1 shows the arrangement and every symbol used below. The clock is in days, because the horizon is a year. Power is in megawatts, energy in megawatt-hours, mass in kilograms, pressure in bar and temperature in kelvin.

11.3.1 Variables

Table 11.1 classifies what the application exposes. The unusual entry is the inventory reference, which is a state rather than a control, because the supervisory policy writes it through a first-order lag of time constant τ_{ref} rather than stepping it.

Table 11.1 Variable classification

Symbol	Meaning	Unit	Range
Controlled variables			
g	grid exchange, positive short	MW	−900 to 900
S	gas state of charge	—	0 to 1
T_{hot}	expander inlet temperature	K	300 to 590
T_{th}	thermal store temperature (override)	K	320 to 590
Manipulated variables			
P_{ch}	compressor load	MW	0 to 250
P_{dis}	expander load	MW	0 to 300
Q_f	supplemental heat	MW	0 to 120
r	heat rejection	%	0 to 100
Disturbances			
κ_w	wind resource scale	×	0.15 to 1.5
κ_d	demand scale	×	0.7 to 1.4
ΔD	cold snap adder	MW	0 to 200
ε_{fc}	forecast error scale	×	0 to 1.5
States			
m	cavern air mass	kg	to 643 kt
T_c	cavern gas temperature	K	200 to 420
T_r	rock temperature	K	250 to 400
T_{th}	thermal store temperature	K	250 to 620
P^a_{ch}	applied compressor load	MW	0 to 250
P^a_{dis}	applied expander load	MW	0 to 300
Q^a_f	applied heater duty	MW	0 to 120
w	wind variability	—	−1 to 1.6
S_{ref}	inventory reference	—	0 to 1
Φ	cyclic stress accumulator	bar	—

11.3.2 The Cavern

Pressure follows the real gas law with a single linear compressibility correction, solved by three fixed-point passes because the correction depends on the pressure it corrects,

$$Z = 1 + Z_1 P, \qquad P = \frac{m R_s T_c}{V_c}\,\frac{Z}{10^5}\ \text{[bar]}. \tag{11.1}$$

At 70 bar the correction is 2.1 % of the inventory, which over a year is about half a gigawatt-hour and is therefore not negligible. The gas state of charge is defined between the working pressures at the present temperature,

$$S = \frac{m - m(P_{min}, T_c)}{m(P_{max}, T_c) - m(P_{min}, T_c)}. \tag{11.2}$$

The mass balance carries the compressor inflow, the expander outflow and a small cavern loss,

$$\frac{\mathrm{d}m}{\mathrm{d}t} = 86400\,(\dot{m}_{in} - \dot{m}_{out} - \lambda m)\,. \tag{11.3}$$

The cavern energy balance is where pressure and temperature separate. Charging pushes aftercooled air in at 313 K and the cavern warms, withdrawal cools it, and the rock takes weeks to argue,

$$mc_v \frac{\mathrm{d}T_c}{\mathrm{d}t} = 86400\Big[\dot{m}_{in}(c_p T_{after} - c_v T_c) - \dot{m}_{out}(c_p - c_v)T_c + UA_r(T_r - T_c)\Big] \tag{11.4}$$

$$C_r \frac{\mathrm{d}T_r}{\mathrm{d}t} = 86400\left[UA_r(T_c - T_r) + UA_\infty(T_\infty - T_r)\right]. \tag{11.5}$$

The two rock time constants are 40.1 d for the gas to rock exchange and 521 d for the rock to the far field, which is why a cavern that has been cycled hard through a winter is still recovering in the spring.

11.3.3 The Machines

Compression is staged with intercooling. The specific work and the stage discharge temperature follow the polytropic relations,

$$r = \left(\frac{P}{P_{amb}}\right)^{1/n_c}, \quad w_c = \frac{n_c c_p T_{amb}\left(r^{\kappa} - 1\right)}{\eta_c}, \tag{11.6}$$

$$T_{dis} = T_{amb}\left[1 + \frac{r^{\kappa} - 1}{\eta_c}\right], \qquad \kappa = \frac{\gamma - 1}{\gamma} = 0.2857. \tag{11.7}$$

At 60 bar the work is 541.3 kJ/kg, the stage discharge is 557 K, and the compressor at its 250 MW rating passes 462 kg/s. The heat recovered into the packed bed is taken from the *actual* discharge temperature rather than the isentropic one,

$$q_{st} = \eta_{hx}\left[(n_c - 1)c_p(T_{dis} - T_{amb}) + c_p(T_{dis} - T_{after})\right], \tag{11.8}$$

which is 464.5 kJ/kg at 60 bar. That expression has a consequence worth noticing before anybody optimizes the compressor in isolation. Everything the compressor does to the air becomes enthalpy, and the inefficiency makes the air hotter rather than making it disappear, so a compressor with a poorer polytropic efficiency hands the thermal store more heat, not less.

Expansion is staged with reheat to whatever temperature the store can supply,

$$r_e = \left(\frac{P}{P_{amb}}\right)^{1/n_e}, \quad x = r_e^{-\kappa}, \quad w_t = n_e\,\eta_t\,c_p\,T_{hot}(1 - x) \tag{11.9}$$

$$\begin{aligned} T_{hot} &= T_b + \min\left(\frac{10^6 Q_f^a}{\dot{m}_{out} c_p},\; T_{hot,max} - T_b\right) \\ T_b &= \max(T_{amb},\; T_{th} - \Delta T_p) \\ T_{ex} &= T_{hot}\left[1 - \eta_t(1 - x)\right] \end{aligned} \tag{11.10}$$

The heater fires only into flowing air and only up to the inlet limit $T_{hot,max} = 650$ K, and the duty it actually delivers is what is charged as fuel, together with the additional interstage duty. Specifically, $Q_{reheat} = Q_f(n_e - 1)(1 - x)$ and $Q_{total} = Q_f + Q_{reheat}$, where Q_f is the heat actually delivered to the first-stage inlet. Ideal interstage reheaters hold the same fired temperature, and the fuel ledger and the Fuel tile use Q_{total}. The exhaust temperature is a constraint in its own right, because expanding air that was not reheated enough puts the last stage below freezing and what freezes is the moisture in it. Measured at 60 bar, the exhaust falls from 345 K at a store temperature of 590 K to 259 K at 450 K, crossing the 260 K icing limit while the plant can still deliver 207.7 MW. The icing constraint therefore binds before the power constraint does.

11.3.4 Deliverability

A fixed-geometry turbine swallows in proportion to its inlet pressure, so

$$\begin{aligned} \dot{m}_{max} &= \min\left(\dot{m}_w,\; \frac{\dot{m}_e P}{P_{max}}\right) \\ P_{del} &= \min\left(P_{exp},\; \frac{\dot{m}_{max} w_t}{10^6}\right) \end{aligned} \tag{11.11}$$

Table 11.2 Deliverable power and stored energy across the working band

P (bar)	S (—)	Swallow (kg/s)	w_t (kJ/kg)	P_{del} (MW)	Stored (GWh)
75.0	1.000	780	362.3	282.6	25.54
70.0	0.835	728	358.1	260.7	21.07
65.0	0.669	676	353.5	239.0	16.66
60.0	0.502	624	348.6	217.5	12.34
55.0	0.335	572	343.1	196.2	8.11
50.0	0.168	520	337.0	175.2	3.99
45.2	0.007	470	330.4	155.3	0.16

Equation (11.11) is the center of the application, because both factors fall together as the cavern empties. Table 11.2 evaluates it across the working band at a store temperature of 470 K.

Spending 84 % of the stored energy costs 38 % of the deliverable power. That is the sentence the whole application exists to make measurable, and it is the reason the plant carries a capacity floor rather than a percentage reserve.

🕮 **Pause.** Try Exercise 11.1 before reading on.

11.3.5 The Thermal Store

The packed bed is one lumped capacitance charged by the compressor and discharged by the expander, with a standing loss,

$$C_{th}\frac{\mathrm{d}T_{th}}{\mathrm{d}t} = 86400\,[q_{in} - q_{out} - UA_{th}(T_{th} - T_{amb})] \tag{11.12}$$

$$\begin{aligned} q_{in} &= \dot{m}_{in}\, q_{st}\, \Pi\!\left(\frac{T_{dis} - \Delta T_p - T_{th}}{30}\right)(1 - r) \\ q_{out} &= \dot{m}_{out}\, q_e (T_{th} - \Delta T_p)\; \mathbf{1}[T_{th} - \Delta T_p > T_{amb}] \end{aligned} \tag{11.13}$$

where Π clips to the unit interval, r is the heat rejection override and q_e is the expander train's heat draw per kilogram evaluated at the inlet temperature the store alone can reach. The heater's inlet and additional interstage heat are paid for as fuel and are not taken from the store, and a store too cold to warm the air above ambient gives nothing. Heat only moves into the store while the compressed air is hotter than it, and whatever the rejection override dumps never arrives, which is pure round-trip efficiency thrown at the sky and is what a store at its material temperature limit costs. Between its limits of 320 K and 590 K the bed holds 4.58 GWh of heat. Rated discharge from a full cavern with the bed at 590 K draws 332 MW of it, so the useful heat lasts about fourteen hours at rating against 81.9 h of gas. The standing loss time constant is 101 d. The second inventory therefore empties within a day of the first one starting to be spent, and in the third scenario the bed falls 195 K in 16.5 h while the cavern loses 22 points of charge.

11.3.6 The Year

Wind capacity factor is a sinusoid on the day of the year multiplied by a bounded stochastic shape carried as a state, so that a run reproduces exactly,

$$c_f = \bar{c} + \hat{c}\cos\frac{2\pi(d - d_{peak})}{365} \tag{11.14}$$

$$W = \min\!\left\{P_w,\; P_w\,\Pi_{0.02}^{0.95}(c_f\,\kappa_w)\,\Pi_0^{2.6}(1 + w)\right\} \tag{11.15}$$

$$\frac{\mathrm{d}w}{\mathrm{d}t} = -\frac{w}{\tau_w} + \sigma_w\sqrt{\frac{2}{\tau_w \Delta t_c}}\,\xi \tag{11.16}$$

with ξ a seeded standard normal drawn once per controller execution and held across it, so that the derivative function stays pure. Demand carries a winter season and a two-harmonic daily cycle, and the price is affine in the net load between a floor and a cap,

$$\pi = \mathrm{clip}(\pi_0 + \beta(D - W),\ \pi_{min},\ \pi_{cap})\,. \tag{11.17}$$

At day 20 the seasonal wind before the stochastic shape is 468 MW against a daily mean demand of 453 MW, and at day 200 it is 179 MW against 226 MW. The shape itself has a mean of about 1.04, because the clip at zero removes more than the clip at 2.6.

11.3.7 Parameters

Table 11.3 gives every parameter the model and controller use, with its default. The V column marks A for adjustable and shown, H for hidden from the beginner view, and F for a fixed physical constant. Four further adjustable entries set only display bands and are omitted.

11.3.8 Initialization and Numerics

Each scenario supplies a starting state of charge, cavern temperature and store temperature. The cavern mass follows from Equation (11.2) inverted at the starting temperature, the rock starts at the undisturbed temperature, and all three actuator states start at zero.

The engine of Section 1.6.1 integrates with classical fourth-order Runge-Kutta on a fixed step of $1/48$ d, which is thirty minutes. The regulatory loops execute every $1/24$ d, which is one hour, the plan re-solves every 0.5 d, and the historian logs the initial state and then one row per simulated hour. A post-step clip holds the cavern mass, the three temperatures, the three actuator states and the inventory reference inside their physical ranges.

The three actuator lags are shorter than the step, with step to time constant ratios of 2.60 for the machines and 5.21 for the heater, and they are not integrated by the Runge-Kutta stages. Each is advanced over the step by its exact first-order solution, and the derivative function evaluates that solution at each stage time, so a commanded load step reaches 92.6 % of its value after one step and a heater step 99.5 %. Carried as ordinary states, the heater lag would lie outside the real-axis stability limit of 2.785 for fourth-order Runge-Kutta and the machine lag would behave as a 1.8 h lag.

11.3.9 Assumptions

1. **The cavern gas is uniform.** A real cavern stratifies during withdrawal and the brine sump at its base exchanges heat and moisture with the air above it [9]. Neither appears here, so the modeled gas temperature swings more than a real one and the wall heat transfer coefficient is a lumped fit rather than a correlation.

2. **The rock is one capacitance and one conductance.** Real cavern wall heat transfer is a diffusion problem whose effective conductance depends on the frequency of the cycling [8], so a single UA_r is right for one cycling rate and wrong for the others.

3. **Salt mechanics appear only as limits.** The minimum working pressure, the maximum pressure and the rated pressure slew come from a creep analysis that is not modeled [11]. The cyclic stress accumulator counts slew beyond the rating and does not feed back into anything.

4. **The packed bed is one temperature.** A real bed has a thermocline whose position matters more than its mean temperature, and the exit temperature during discharge falls in a characteristic way that a lumped capacitance cannot reproduce [6].

5. **The machines are steady-state maps behind first-order lags.** Compressor surge appears only as a minimum load below which the command is zero, and no start cost, minimum run time or thermal stress limit is represented.

Table 11.3 Model parameters and default values

Symbol	Meaning	Value	Unit	V
V_c	total cavern volume	8.0×10^6	m^3	A
P_{min}, P_{max}	working pressure limits	45, 75	bar	A
Z_1	compressibility slope	3.0×10^{-4}	1/bar	H
UA_r	gas to rock conductance	2.6×10^6	W/K	A
C_r	rock heat capacity	9.0×10^{12}	J/K	A
UA_∞	rock to far field	2.0×10^5	W/K	A
T_∞	undisturbed rock temperature	313	K	A
λ	cavern loss rate	2.0×10^{-5}	1/d	A
$\dot{m}_w$	well flow limit	820	kg/s	A
$\dot{\sigma}_{max}$	rated pressure slew	6	bar/d	A
P_{comp}	compressor rating	250	MW	A
P_{exp}	expander rating	300	MW	A
n_c, n_e	compression, expansion stages	2, 2	—	A
η_c	compressor polytropic efficiency	0.85	—	A
η_t	expander isentropic efficiency	0.88	—	A
c_p	air heat capacity	1005	J/(kg K)	H
P_{amb}	ambient pressure	1.0	bar	H
T_{amb}	compressor inlet temperature	288	K	A
T_{after}	aftercooled air into the cavern	313	K	A
f_{surge}	compressor minimum load	0.30	—	A
$\dot{m}_e$	expander swallowing at P_{max}	780	kg/s	A
$x_{e,min}$	expander minimum load	0.15	—	A
P_{aux}	auxiliary load when running	6	MW	A
C_{th}	thermal store heat capacity	6.1×10^{10}	J/K	A
η_{hx}	compression heat recovered	0.90	—	A
ΔT_p	store to air approach	25	K	A
UA_{th}	store standing loss	7.0×10^3	W/K	A
$T_{th,min}$	store useless below	320	K	A
$T_{th,max}$	store material limit	590	K	A
$T_{ex,min}$	lowest allowed exhaust	260	K	A
$Q_{f,max}$	supplemental heater rating	120	MW	A
$T_{hot,max}$	highest fired inlet temperature	650	K	A
d_0	day of the year the run starts	0	d	A
P_w	wind farm rating	850	MW	A
$\bar{c}, \hat{c}$	mean, seasonal swing in c_f	0.38, 0.17	—	A
d_{peak}	day of peak seasonal wind	20	d	A
σ_w, τ_w	wind variability, correlation time	0.85, 0.55	—, d	A
$\bar{D}$	annual mean demand	340	MW	A
D_{win}, D_{diu}	winter adder, daily swing	120, 80	MW	A
π_0, β	price intercept, slope	20, 0.30	$/MWh	A
π_{min}, π_{cap}	price floor, cap	−32, 420	$/MWh	A
c_{unmet}	value of lost load	900	$/MWh	A
c_{fuel}	cost of supplemental heat	46	$/MWh	A
c_{wear}	wear charge on throughput	1.6	$/MWh	A
P_{firm}	firm winter commitment	150	MW	A
d_{from}, d_{to}	firm window	330, 60	d	A
S_{term}	plan's terminal inventory	0.55	—	A
τ_{mach}	machine load response	0.008	d	A
τ_{fuel}	heater response	0.004	d	A
τ_{ref}	inventory reference rate	2.0	d	A
R_s	specific gas constant for air	287	J/(kg K)	F

6. **Air is ideal with one compressibility correction.** At 75 bar and 320 K the single linear term in Equation (11.1) is adequate for inventory accounting and is not a substitute for a real equation of state.

7. **The wind, demand and price are synthetic.** The resource is a sinusoid on an Ornstein-Uhlenbeck process, the demand is two harmonics, and the price is affine in the net load. They are teaching data and are never labeled as measured.

11.4 Open-Loop Character

Manual mode gives the compressor load, the expander load, the heater and the rejection valve directly, and the first experiment shows that the two inventories are not one. The application's third scenario starts a discharge on day 320 with the cavern at 85 % and the store at 560 K and holds an expander command of 280 MW. The expander inlet falls below 500 K after 2.5 h, the exhaust crosses the 260 K icing limit after 7 h (the alarm is raised on the step that ends at 7.5 h), and the store reaches 365 K after 16.5 h, while the cavern pressure goes only from 70.3 bar to 63.8 bar. Delivered power holds the 280 MW command for the first hours and is down to 184 MW at 15 h, and the December reserve floor of 61.4 % ends the discharge at about 18.5 h, before the store reaches its 320 K useful limit. The scorecard at the end of three days reports 3.96 GWh delivered with the cavern still at 62.4 % and the store at 353 K. Deliverable power at that point is 172.0 MW against a nameplate of 300 MW, and nothing about the cavern explains it, because the same 63.7 bar with a 470 K store would deliver 233.2 MW. The scenario's second run gives the expander 60 MW of supplemental heat from the manual station, which keeps the exhaust above its limit for the day and raises the day's delivery from 3.82 GWh to 4.78 GWh for 1.20 GWh of first-stage heat. The Fuel tile, which adds the interstage duty Q_{reheat}, reads more than that.

> **TRY IT IN THE APP** **3 · The Second Inventory, and It Empties First**
>
> Load scenario 3 and press Run. Watch the thermal store temperature beside the cavern pressure. In about 16.5 h the store falls from 560 to 365 K while the cavern goes only from 70.3 to 63.8 bar, and the exhaust alarm comes on after about seven hours.

The second scenario makes the other half of the point. A cavern at 30 % is a cavern at 53.9 bar, and with the scenario's 500 K store that is 204.7 MW of deliverable power (191.8 MW at the 470 K of Table 11.2). It still holds 7.76 GWh. Asked for 300 MW, the expander peaks at 199 MW in the first hour and falls below 150 MW after 11 h, by which time the pressure is 50.5 bar and the store 399 K, and the summer reserve floor of 13.9 % ends the discharge at about 16.5 h, at 49.5 bar and 15.0 %. A fixed-geometry expander swallows in proportion to its inlet pressure, and the specific work falls with the pressure ratio and with the store temperature at the same time.

The time scales are widely separated and that separation is the reason the control problem decomposes. The machine lag is 11.5 min, solved exactly across the half-hour step. The grid exchange the delivery loop controls moves on the hour. The cavern holds 81.9 h at rating, so the inventory moves over days. The rock responds over 40 d and the far field over 521 d. The reserve policy is written for a year. Four decades of time scale between the hourly loop and the yearly decision is more than any single controller should be asked to span, and the architecture in Section 11.5 follows that split exactly.

> **DID YOU KNOW?** **A Storage Plant That Ramps Fast**
>
> The Alabama plant can change its output by about 18 MW per minute, about 60 % faster than a typical gas turbine, and its heat rate rises only from 4330 to 4750 kJ h/kW between full load and 20 % load [5]. Hours of storage, not minutes of response, are the scarce quantity.

Table 11.4 Regulatory loops as implemented

CV	MV	K_c	τ_I (d)	Character
g	P_{dis}	1.1	0.20	direct, trim only
S	P_{ch}	10	2.50	reverse, Any only
T_{hot}	Q_f	0.2	0.15	reverse, costs fuel
T_{th}	r	6	0.40	direct, override

11.5 Regulatory Control

11.5.1 The Discrete Algorithm

Four loops share one proportional-integral-derivative implementation in velocity form, with the derivative taken on the measurement,

$$e_k = \gamma(\mathrm{SP}_k - \mathrm{PV}_k), \quad P_k = K_c e_k, \quad D_k = -K_c \tau_D \gamma \dot{\mathrm{PV}}_k \tag{11.18}$$

$$\begin{aligned} \Sigma_k &= \Sigma_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ c_k &= u_{bias} + \Sigma_k + I_{k-1}, \qquad u_k = \mathrm{clip}(c_k, u_{min}, u_{max}) \end{aligned} \tag{11.19}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t} \right] \Delta t_c \tag{11.20}$$

These are Equations (1.10) to (1.13) of Section 1.6.2 in this chapter's symbols, with $\Sigma_k = S_k$, $c_k = v_k$ and the action $\gamma = a$, and that section explains the back-calculation and the bumpless transfer from Manual. Here u_k is the demand c_k clipped to the loop limits with no slew limit, the tracking time T_t is τ_I on all four loops, $\Delta t_c = 1/24$ d, and all four loops run with $\tau_D = 0$.

One tuning consequence of the one-hour scan is worth recording. While an output is clipped, Equation (11.20) is a first-order recursion in I with pole $1 - \Delta t_c / T_t$, so the back-calculation is stable only for τ_I above $\Delta t_c / 2$, which is half an hour, and it oscillates for τ_I below an hour. The integral-time sliders stop at 0.08 d for the grid and inlet loops, 0.3 d for the inventory loop and 0.04 d for the rejection override, all above the limit, the last with a pole of -0.04. A value of 0.02 d, below the slider's range, shows what the limit means. In a counterfactual run of the fifth scenario at that value the override's integral, which sits exactly at zero while the store is below its margin, picks up rounding noise and grows by 8 % per scan, the override starts dumping compression heat on alternate scans on day 19.6 of the run, and $|I|$ passes 10^6 by day 24.2. The tuned values are 0.15 d or more.

11.5.2 Four Loops on Four Time Scales

Table 11.4 gives the pairings and default tunings. Two loops are direct acting ($\gamma = -1$ in the convention of Section 1.6.2), because more expander power reduces the grid shortfall and more rejection cools the store. The other two are reverse acting ($\gamma = +1$). The grid loop measures g through a first-order filter of 0.01 d, and the inventory loop measures S through one of 0.02 d against the lagged reference S_{ref}. The inventory loop's output reaches the compressor only when the charging source is set to Any. Every Auto scenario ships with Surplus only, which gives the compressor the wind surplus $W - D - P_{aux}$ up to its rating before the overrides and discards the loop's output, so $K_{c,S}$ and $\tau_{I,S}$ change nothing in those runs and a run with $K_{c,S} = 0$ is bit-identical. The inlet-temperature loop runs only in Auto with the heater armed and the expander running.

The hourly sampled inlet loop uses $K_{c,H} = 0.2\,\mathrm{MW/K}$. A gain of 1.6 MW/K exceeds the sampled-loop stability limit, and its command alternates between zero and the heater rating. The exact actuator lag does not cure a controller scanned too slowly. The gain of the table settles at both the minimum reference flow and the design flow, and all heater-armed results below use it.

Table 11.5 Capacity floor for a 150 MW commitment

Store temperature (K)	S_{cap}, no heat	S_{cap}, heater armed
590	0.000	0.000
500	0.000	0.000
450	0.021	0.000
400	0.190	0.000
350	0.407	0.000

11.5.3 Budget and Trim

The expander is commanded by two numbers that are deliberately kept apart. The budget says how much energy today is worth and comes from the inventory policy. The trim says how much of that budget this hour actually needs, and it is the grid loop. The trim is allowed to reduce the budget and never to exceed it,

$$P_{dis} = \mathrm{clip}(B + u_g,\ 0,\ P_{lim}), \qquad u_g \in [-B, 0]. \tag{11.21}$$

A command below the 15 % expander minimum load is then set to zero, and a scan that asks for charge and discharge together keeps only the larger. The grid loop's own limits see neither cut. Giving the grid loop authority to exceed the budget is the instructive failure, and the application does not offer it as a control. In a counterfactual run with the trim's upper limit raised to the expander rating, the taper of Equation (11.22) still holds the challenge year's minimum inventory at 24.9 %, and the failure arrives as money instead. The loop discharges into deficit hours whose price the plan had refused, and over the hundred and fifty days the margin falls by $2.87 M and curtailment rises by 24.9 GWh.

The discharge limit P_{lim} is not the expander rating. It is the deliverable power of Equation (11.11) multiplied by a taper that closes as the state of charge approaches the binding floor,

$$P_{lim} = \min(P_{del}, P_{exp})\,\Pi\!\left(\frac{S - S_{floor}}{0.04}\right). \tag{11.22}$$

A tapered limit of this kind is the cascade anti-windup of Section 1.4.4, and Chapter 10 tapers the charge and discharge limits of a flow battery in the same way.

11.5.4 The Capacity Floor

Inside the firm window the binding floor is the larger of two quantities. One is the seasonal reserve the year has to end on. The other is the state of charge at which the deliverable power would fall to the commitment, which is computed by twelve bisections on Equation (11.11) at every scan inside the window and is zero outside it,

$$S_{cap} = \min\left\{S : P_{del}\left(P(S, T_c),\, T_{th},\, Q_f\right) \geq P_{firm}\right\}. \tag{11.23}$$

The distinction between S_{cap} and a percentage reserve is the point of the application, because the commitment was sold in megawatts and a percentage is only a proxy for them. Table 11.5 shows how far the proxy drifts.

The same commitment is defended at a state of charge of zero with a hot store and at 40.7 % with a cold one, and a fixed percentage reserve cannot express that. The application draws S_{cap} on the inventory trend alongside the seasonal curve so that the reader can watch the two cross. With the heater armed, which means plant and heater both in Auto, the calculation and the Deliverable tile count the heater up to the inlet setpoint its loop would hold, 500 K by default. At that inlet even the bottom of the band makes 174.5 MW, so the floor is zero at every store temperature in the table, and the heater does deliver that heat when it is called. The floor also uses the store temperature of the present scan, so it rises behind a discharge that cools the store rather than ahead of it.

11.5.5 Two Overrides

The compressor carries a pressure override. Its command is multiplied by a headroom factor that closes over the last 1.5 bar below the maximum working pressure, and below the anti-surge minimum load the command is zero rather than small, because a machine on recycle is pure loss. The surge line that sets such a minimum load is the subject of Volume 4, Chapter 2. Salt does not negotiate about the upper pressure and the override is not a controller with a setpoint, it is a multiplier.

The thermal store carries a rejection override that sits against zero until the bed approaches its material temperature limit. What it rejects never reaches the store, which is round-trip efficiency thrown at the sky, and it is the correct action because a bed above its material limit is a capital loss rather than an efficiency loss.

11.6 Three Inventory Policies

Every architecture in this application is a different way of writing the inventory reference trajectory, and all three are bounded below by the same seasonal reserve, which is a squared raised cosine on the day of the year,

$$S_{res}(d) = \frac{R_{min}}{100} + \frac{R_{peak} - R_{min}}{100} \left[\frac{1 + \cos \frac{2\pi (d - d_{res})}{365}}{2} \right]^2 . \tag{11.24}$$

With the defaults it runs from 12.0 % at day 147.5 in late May to 62.0 % at day 330, and it is 12.2 % in late April. The squared raised cosine is chosen so that the floor is expensive only where it is needed, relaxing quickly through the spring and tightening through the autumn. It is enforced by projection rather than as a penalty term, because a winter reserve that can be bought off with a weight is not a reserve.

The first policy is a pair of price thresholds and nothing else. Empty at rating whenever the price is above the discharge threshold, fill whenever it is below the charge threshold, and never go below the seasonal floor. There is no forecast anywhere in it. With the shipped Surplus only charging source, which every policy uses, the compressor takes the wind surplus whatever the price, and the charge threshold acts only when grid charging is allowed.

The second policy is the operator's own reference curve. The expander budget rises from zero to rating as the measured state of charge climbs from the curve to five points above it, and the compressor follows the error to the curve only when grid charging is allowed.

The third policy is a water value. The marginal value of a megawatt-hour in the ground is the classical answer to this problem [1] and it is far cheaper than an optimizer. Section 1.4.6 reads it from a merit order of forecast prices. The plan here finds it by bisection on a forward simulation instead, as Chapter 6 does for its own store. The plan simulates the next thirty daily blocks forward under a trial threshold price, charging from whatever surplus the wind forecast offers and discharging the forecast deficit plus a quarter of the expander rating, up to the rating, in every block whose price clears the threshold divided by the one-way efficiency $\eta_{plan} = 0.72$, while the seasonal floor stops it selling,

$$\begin{aligned} S_{k+1} &= \mathrm{clip}\left(S_k + \frac{\Delta E_k}{E_{cap}}, 0, 1 \right) \\ \Delta E_k &\geq -\left(S_k - S_{res,k} \right) E_{cap} \end{aligned} \tag{11.25}$$

Raising the threshold sells into fewer days and ends with more in the ground, so the end-of-horizon inventory is monotone in the threshold, and sixteen bisections find the lowest threshold that still meets the terminal target, which moves from the floor at the horizon's end toward S_{term} as the terminal weight rises to 10. That threshold is the marginal value of the inventory. It is computed rather than chosen, it moves by itself when the month turns windy, and dividing it by η_{plan} gives the sell price the expander budget is dispatched against. Nobody has to pick a threshold.

One detail of Equation (11.25) was found by the model failing and is worth stating. The floor stops the plant selling, and it does not manufacture gas. Written instead as a clamp that raises the state up onto the floor, the forward simulation invents energy out of the constraint, the end-of-horizon inventory stops depending on the threshold at all, and the bisection then reports the price floor as the value of everything in the ground.

🕮 **Pause.** Try Exercise 11.3 before reading on.

11.7 Scenarios and Measured Results

Table 11.6 lists the eleven scenarios in selector order with the day of the year each one starts on, the dispatch policy it loads and what a learner must do to complete it. Credit, the challenge and its voiding rule work as Section 1.6.4 describes. Table 11.7 gives the scorecard at the end of each scenario's trend window with the settings as shipped. Every result in this section is for the shipped random seed. Reset keeps the learner's settings and reloading restores the case's own (Section 1.6.1), and a case loaded from the menu or rerun after Reset reproduces the case loaded first exactly. Figure 11.3 shows the eighth scenario sixty-three days into its run.

11.7.1 Three Policies on the Same Hundred and Fifty Days

The comparison the application is built around runs the challenge scenario's year, a hundred and fifty days from day 245 including a sixteen-day December wind drought from day 337 and a 70 MW cold snap that overlaps its end, under each of the three inventory policies with everything else at the challenge's suggested configuration. That configuration arms the supplemental heater, so each policy is also run with the heater left in Manual at zero, which is the purely adiabatic plant. Table 11.8 gives the result.

With the heater armed, all three policies keep the commitment with zero default and no day below it, burning 21.0 GWh, 13.0 GWh and 23.1 GWh of heat for the rule, the curve and the plan, and a pair of price thresholds books \$0.77 M more than the rolling plan, which delivers the most energy. The explanation is in the price model. Price here is affine in the net load, so a price threshold is a net-load threshold, and the net load is exactly the quantity the water value is trying to forecast. A policy that reacts to the present value of a variable it can measure perfectly is hard to beat with a forecast of the same variable. The plan earns its place on the days when price and net load come apart, and this price model never lets them.

The reliability in the upper half of the table is the heater's. Its applied duty stays below the 120 MW limit and is small in the drought itself, but the armed heater counts in the capability whether it fires or not, and measured against the unfired capability the plan would be below the commitment on 28.4 d of the firm window, the rule on 25.5 d and the operator curve on 0.3 d. With the heater off, the rule and the plan default on 0.56 GWh and 0.47 GWh and spend 17.0 and 15.3 days short, while the operator curve, which never lets the inventory below 33.9 %, defaults on 2 MWh and is short for 7 h. Once the stored heat is gone the commitment is protected either by gas held back or by heat bought, and the operator curve, which holds the gas, books the least. Figure 11.2 draws the table as bars.

The year is one seeded realization, so the comparison was repeated on seven other seeds. The margins move by about \$6 M between seeds, the rule books more than the plan with the heater armed in six of eight, the operator curve books the least in all eight, and with the heater off the operator curve has the fewest days short in all eight. The ordering is robust and the size of each gap is not.

> ✦ DID YOU KNOW? **No Forecast, More Money**
>
> From the simulator: with the heater armed, a pair of price thresholds that uses no forecast at all books \$16.04 M over the 150 days, and the rolling plan that looks thirty days ahead books \$15.28 M while delivering the most energy, 128.46 GWh (Table 11.8).

Table 11.6 Scenarios and Completion Requirements

#	Scenario	Day	Policy	Kind	Required to complete
1	Two machines	100	manual	Experiment	In Manual, charge at 250 MW until the cavern pressure exceeds 74 bar, then discharge at 300 MW for one day.
2	Full cavern short of rating	200	manual	Demonstration	At a 300 MW expander command, deliver more than 50 MW for 6 h, then continue until the reserve floor stops the expander (delivery below 5 MW within 2 points of the floor).
3	Second inventory	320	manual	Demonstration	Discharge at 280 MW for one day with no supplemental heat, then Reset, set 60 MW of supplemental heat and discharge for another day.
4	Spring surplus	105	schedule	Experiment	Two 30-day runs, with the inventory reference at 80% and, after Reset, held at 30% from the start (logged reference never above 40%).
5	Price thresholds	200	rule	Experiment	Three 60-day runs with the thresholds at $18 and $62 as shipped, at least $60 apart, and at most $28 apart, each set before the run.
6	Reserve floor in April	95	rule	Experiment	Three 60-day runs with the reserve peak held from the start at 62%, at 90% or more, and at 20% or less.
7	Winter drought	340	rule	Experiment	With the heater armed, run until the inventory has stayed within 3 points of the binding floor for one day of the drought, then on to run day 30.
8	Rolling plan	250	plan	Experiment	Two 90-day runs with the terminal weight held at 8 and, after Reset, at 0.
9	Compressor trip	110	plan	Experiment	Forty days through the scripted loss of half the compressor on run day 4 and its return on run day 18.
10	Forecast bias	300	plan	Experiment	Two 60-day runs with forecast bias scale 1.1 and wind 0.72, with the reserve peak held at 62% and at 77% or more.
11	Challenge	245	rule	Challenge	Over 150 days, default below 0.5 GWh, deliverable power below 150 MW on less than 4 days of the firm window, inventory above 20% and margin above $14.8 M, with the model parameters and disturbances as given.[a]

[a] As shipped (threshold rule at $40 and $45, reserve peak 25%, heater off, weak tuning) the case fails all four clauses, with 1.11 GWh of default, 25.08 days short and a lowest inventory of 19.5%, and its $14.36 M misses the margin. With the suggested settings (rolling plan, reserve peak 50%, heater armed, tuned loops) it passes with 0.00 GWh, 0.0 days, 24.9% and $15.28 M, the same whether entered before Run or carried through Reset, $0.48 M and 4.9 points inside the margin and inventory limits, burning 23.1 GWh of heat. The same settings with the heater off fail with 0.47 GWh and 15.33 days short.

11.7.2 What the Seasonal Reserve Costs

The seasonal reserve peak was swept across its useful range with everything else at the suggested rolling-plan configuration, once with the heater armed and once with it off. Table 11.9 gives the result.

With the heater armed, raising the seasonal reserve from 25 % to 85 % costs $3.11 M of margin, buys 15.5 points of minimum inventory, and leaves the unserved counter at zero at every setting. That zero is not the capacity floor at work, because with the heater armed the capacity floor is zero at every store temperature. It is the heater, which keeps the capability above the commitment through the drought. With the heater off, the same sweep cuts the default from 1.15 GWh to 0.13 GWh and the shortfall from 575 h to 227 h for $1.70 M of margin, and no setting meets the commitment. Across the eight seeds the armed sweep always costs between $3.0 M and $5.0 M, while the reduction in default with the heater off ranges from −0.01 to 1.03 GWh, the shipped seed giving the largest.

🕮 **Pause.** Try Exercise 11.2 before reading on.

The practical lesson survives in a sharper form. A floor in megawatts is in the right units, and

Table 11.7 Scorecard at the end of each scenario window, as shipped

#	Days	Delivered (GWh)	Charged (GWh)	Curtailed (GWh)	Min S (—)	Margin ($M)
1	4	0.00	0.00	32.70	0.472	0.00
2	6	2.65	0.00	21.47	0.150	-0.03
3	3	3.96	0.00	23.56	0.621	-0.11
4	30	20.81	55.92	81.44	0.349	2.59
5	60	34.89	61.70	19.64	0.236	3.93
6	60	57.17	90.61	109.43	0.130	6.60
7	40	40.77	47.03	18.15	0.237	5.22
8	90	81.87	133.14	78.02	0.239	8.98
9	40	52.77	79.71	75.25	0.132	4.70
10	60	39.24	60.26	30.50	0.239	3.86
11	150	115.06	229.62	111.99	0.195	14.36

Table 11.8 Three inventory policies over the same 150 days

Policy	Heater	Delivered (GWh)	Unserved (GWh)	Min S (—)	Hours short	Margin ($M)
threshold rule	armed	118.16	0.00	0.249	0	16.04
operator curve	armed	118.21	0.00	0.339	0	12.34
rolling plan	armed	128.46	0.00	0.249	0	15.28
threshold rule	off	105.87	0.56	0.296	408	15.05
operator curve	off	111.23	0.00	0.339	7	12.19
rolling plan	off	115.52	0.47	0.250	368	14.55

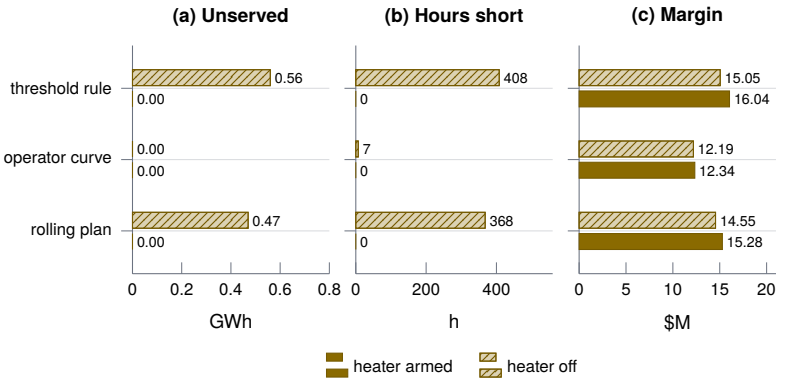

Figure 11.2 The three inventory policies over the same 150 days, from Table 11.8: (a) unserved energy, (b) hours below the commitment and (c) margin. Solid: supplemental heater armed. Hatched: heater off. Notice that with the heater armed every policy keeps the commitment, and that with it off the operator curve, which books the least margin, comes closest to keeping it.

Table 11.5 shows how far a percentage drifts from it. It is computed from the store temperature of the present scan, though, and a discharge into a drought cools the store faster than it empties the cavern, so the floor rises behind the discharge rather than ahead of it. In the challenge year with the heater off, the plan spends from 51 % on day 330 to 31.5 % by day 336, leaves the store near 365 K, and then holds between 147 and 151 MW of capability through the drought while the capacity floor climbs past an inventory the plant can no longer refill. A reserve that protects this commitment has to count

Table 11.9 Seasonal reserve peak swept over 150 days, rolling plan

	Heater armed		Heater off		
R_{peak} (%)	Min S (—)	Margin ($M)	Unserved (GWh)	Hours short	Margin ($M)
25	0.184	16.25	1.15	575	14.65
50	0.249	15.28	0.47	368	14.55
68	0.296	14.02	0.38	333	13.59
85	0.339	13.14	0.13	227	12.95

the heat as well as the gas, and the configurations here that do so are the one that holds more gas and the one that buys heat.

11.7.3 The Drought

The seventh scenario opens on day 340 with the caverns at 80 %, 22.40 GWh on the stored-energy trace, and drops the wind to 30 % of normal for fifteen days with 60 MW of cold-snap demand on top, under the threshold rule with the heater armed. The fortnight needs several times what is stored, so it is not covered and was never going to be. The thermal store goes first. It falls from 520 K to 380 K in twelve hours and to 313 K in a day and a half, near the temperature below which it can no longer warm the air above ambient. The heater rises to about 113 MW by a day and a half, then stops when discharge is no longer scheduled. By day 3 the inventory is near 24.6 % at 52.3 bar. The armed heater is counted in the capability even when it is idle. Over forty days the plant delivers 40.8 GWh, imports 178.2 GWh, burns 12.2 GWh of heat and ends near 24.1 %.

> **TRY IT IN THE APP** **7 · A Two-Week Winter Wind Drought**
>
> Load scenario 7 and press Run. Watch the thermal store temperature for the first two days. It falls from 520 to 380 K in twelve hours and to 313 K in a day and a half, while the supplemental heat rises to about 113 MW to hold the expander inlet.

The instructive trace is the deliverable capacity. The Deliverable tile reads 207.3 MW on the floor, counting the armed heater up to the 500 K inlet setpoint, while the unfired deliverable power of the same cavern and store is 119.5 MW, below the 150 MW commitment. The unserved counter stays at zero because the heater is there to be fired. With the heater left in Manual the capacity floor climbs to 54 % within a day as the store cools to 325 K, and the plant holds 56.5 % of its gas at 61.9 bar while it can deliver only about 149 MW, booking 0.11 GWh of default by day 15 and 0.12 GWh by day 40. The gas is there and the heat is not.

11.7.4 The Graded Challenge

The eleventh scenario opens on day 245 for a hundred and fifty days with weak tuning, the threshold rule at $40 and $45, a reserve peak of 25 %, the heater in Manual and a forecast error scale of 0.6. As shipped it books $14.36 M, defaults on 1.11 GWh of its firm commitment, falls below the commitment on 25.08 d of the firm window and takes the cavern to 19.5 %. The suggested configuration, which is the rolling plan with a reserve peak of 50 % and a spring minimum of 15 % peaking at day 335, the heater armed with a 505 K inlet setpoint, a terminal weight of 10 and the tuned gains, books $15.28 M with no default and a minimum inventory of 24.9 %, burning 23.1 GWh of heat, and the app gives the same run whether these settings are entered before Run or carried through Reset. The difference is $0.92 M and the whole of the reliability gap, and the reliability part of it is the heater. Arming the heater alone on the shipped configuration removes the default and the shortfall for 10.0 GWh of heat and raises the margin to $15.71 M, failing only on its 15.7 % minimum inventory, while the suggested configuration with the heater left in Manual fails with 0.47 GWh and 15.33 d short. Across

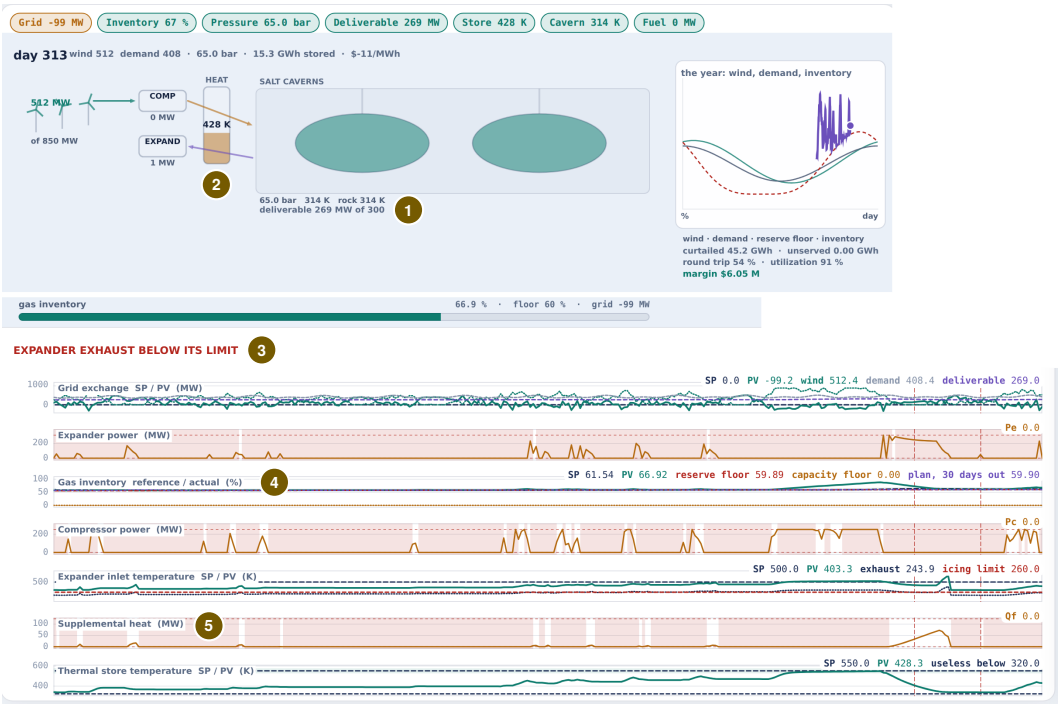

Figure 11.3 The application on day 313 of the eighth scenario, sixty-three days into a rolling-plan run that began on day 250, cut to the plant view and the trends. The callouts mark (1) the deliverable power under the caverns, (2) the thermal store, (3) the exhaust alarm, (4) the inventory trend and (5) the supplemental heat trend. The inventory trend carries the three curves the reader is asked to distinguish, which are the measured state of charge at 66.9 %, the seasonal reserve floor at 59.9 % on its way to its 62 % peak at day 330, and the capacity floor at 0.0 %, because it is computed only inside the firm window that opens on day 330 (and with the heater armed it would read zero there too). The plant has curtailed 45.2 GWh, left 0.00 GWh unserved and booked $6.05 M. The store is at 428 K and the expander, idling at 1 MW, sees 403 K at its inlet against the 500 K setpoint, so the exhaust alarm is active at 243.9 K against the 260 K icing limit, an alarm the dispatch does not act on. The supplemental heat trend shows the heater firing in bursts whenever the expander runs and idle now, and the 269 MW Deliverable badge counts the armed heater up to the inlet setpoint against an unfired 217 MW.

the shipped seed and seven others, the suggested configuration passes six of eight times. Seeds 2 and 3 meet every physical clause but miss the margin target at $14.14 M and $13.46 M. The shipped configuration fails at the default seed. The absolute margin target is a benchmark for that scenario, rather than a guarantee across different synthetic winters.

11.8 Technology in Practice

The instruments and machines of Figure 11.1 are, in a plant, a pressure transmitter on the wellhead, thermometers on the store and the expander, a power measurement at the grid tie, a motor-driven compressor, a heat-rejection valve, a fired heater and an expander train. Figure 11.4 draws the chapter's loops with that hardware in place of the schematic's symbols.

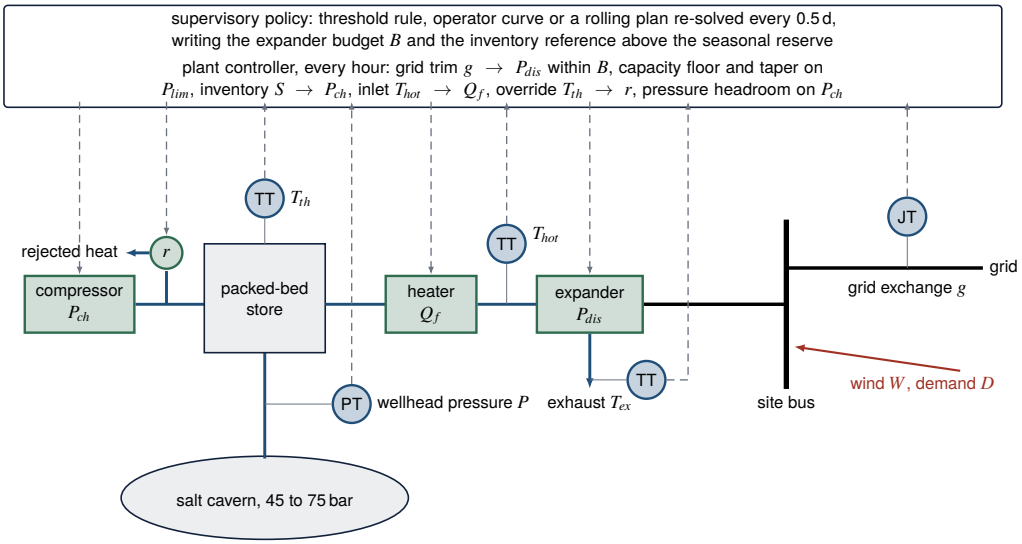

Figure 11.4 The chapter's loops drawn with their hardware. PT reads the cavern pressure at the wellhead, the thermometers TT read the store, the expander inlet and the exhaust, and JT reads the grid exchange at the site's tie. Every hour the plant controller sets the expander load, the compressor load, the supplemental heat and the heat rejection, under a budget that the supervisory policy writes. Wind and demand act on the site bus. Notice that the pressure measures only the gas inventory: the heat inventory is read separately, in the store.

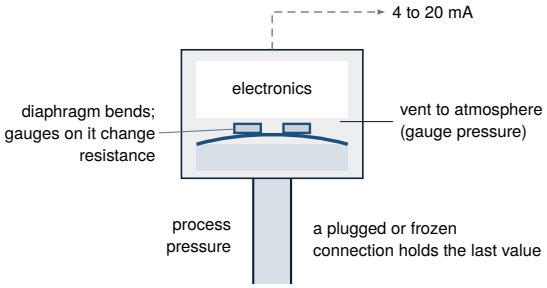

Figure 11.5 Principle of a gauge pressure transmitter. Pressure deflects a diaphragm, strain gauges on it change resistance and the electronics send the reading. Notice that a plugged connection leaves the signal steady while the process moves, which on a cavern looks like an inventory that stopped changing.

What Measures It

The gas inventory is never measured directly. The plant reads the cavern pressure at the wellhead and computes the mass from it and the cavern temperature through Equation (11.1), which is what the chapter's state of charge does. Figure 11.5 shows the principle of the transmitter. The heat inventory needs its own thermometer in the store, and the deliverable power is a calculation from both.

> **TECHNOLOGY** **Sensor: A Pressure Transmitter at the Wellhead**
>
> An electronic pressure transmitter senses the deflection of a diaphragm, with bonded strain gauges or a differential capacitance cell, and sends the pressure as a 4 mA to 20 mA or digital signal [14]. A silicon strain gauge that is overstressed fails completely, while a metal one can drift and still look healthy. A snubber, a small restriction that damps pulsations, can plug with debris, and the reading then responds too slowly or

not at all [14]. On a trend that looks like a cavern pressure frozen while the compressor runs. The chapter's PT is exact, and the inventory is computed from it and the cavern temperature.

What Moves It

Four final elements act on the plant. The compressor's motor sets the charging load, and the chapter's override closes it over the last 1.5 bar below the maximum pressure. A valve rejects compression heat when the store reaches its material limit. A fired heater lifts the expander inlet when the store is too cold, which is the design the two commercial plants use all the time: they burn fuel in the air withdrawn from storage, partly because air expanded without reheating would leave the turbine cold enough to risk icing on its blades [5]. The fourth element, the expander, carries the chapter's lesson about deliverable power.

TECHNOLOGY **Final Element: The Expander and Its Inlet Throttle**

A compressed-air plant sets its output by the air flow through the expander while it holds the inlet temperatures of both stages constant, which keeps the efficiency high at part load [5]. Both commercial plants throttle the cavern air to a fixed turbine inlet pressure, 46 bar at Huntorf and 45 bar at McIntosh, because the turbine is more efficient at a constant inlet pressure, at the cost of a larger cavern. The other option lets the inlet pressure follow the cavern and accepts a falling output [5]. The chapter's expander takes that option, which is why its deliverable power falls with the pressure.

Where the Controller Runs

The four loops of Table 11.4 and the two overrides run in the plant's own controller, and the inventory policy runs above it in supervisory software. Supervisory and lower-level loops of this kind run with cycle times from milliseconds to minutes [15]. The chapter's loops scan once an hour on a thirty-minute physics step and the plan re-solves every half day, because the inventory moves over days. At Huntorf the switch between compressing and generating is fully automated and takes at least twenty minutes, with the plant neither generating nor compressing [5], which the chapter's rule that keeps only the larger of a charge and a discharge command does not represent.

IN PRACTICE **Written Procedures and Integrity Management**

In the United States a solution-mined salt cavern that stores natural gas must meet the recommended practice API RP 1170 under 49 CFR 192.12 [16]. Its operator keeps written manuals for operations, maintenance and emergency response, reviews them at intervals of no more than 15 months and at least once each calendar year, and runs an integrity management program with risk assessments of every cavern and every well, wellhead included [16]. Huntorf itself was built beside a natural gas storage facility of four solution-mined caverns [5].

What Is Modernizing It

Three developments touch this plant. The first is the heat store. Advanced adiabatic designs keep the heat of compression in a thermal store and use it in place of fuel. One concept reached a round-trip efficiency of about 70 % with no fuel, and the stores studied for it range from beds of natural stone and ceramics to molten nitrate salt and oil [5]. The second is the machine train. Separating the compressor and expander, which share one shaft through a clutch at both existing plants, would allow direct switching between charging and generating and standard production compressors [5]. The third is the network. Operational systems increasingly resemble information systems as they adopt standard computers, operating systems and protocols for remote access, which leaves them less isolated and in greater need of security [15]. A plant that sells a capacity commitment depends on that network every time it is called.

11.9 Limitations

The first-stage supplemental heater has a 650 K firing limit rather than a burner model. Ideal interstage reheaters reach the same fired temperature, and their additional duty is charged at the fuel price. The rating and the actuator apply to the first-stage heater, while the Fuel tile and the margin include both it and the additional interstage duty. The model has no separate interstage burner dynamics or heat-transfer limits, and its thermal-store and reheat laws are reduced-order approximations.

The price model makes the forecast-based policy look worse than it should. Price is affine in the net load, so the quantity a price threshold responds to and the quantity the water value forecasts are the same variable, and the comparison in Table 11.8 is therefore a weak test of forecasting. A price model with scarcity events decoupled from the net load would separate them, and the result reported here should be read as a statement about this price model rather than about forecasting in general.

The thermal store is one temperature where a packed bed has a thermocline. The modeled store cools from the first hour of a discharge, by 30 K in the first 2.5 h of the third scenario, where a real bed's exit temperature follows the position of its thermocline [6], so where the icing constraint of Equation (11.10) binds in a real plant depends on the bed's design and not on its mean temperature. A store too cold to warm the air above ambient gives no heat, and the inlet is then held at ambient, an idealization of the air-side exchangers that would be needed.

Salt mechanics are limits rather than a model. The cyclic stress accumulator counts pressure slew beyond the rated 6 bar per day and feeds back into nothing, so a reader can cycle the cavern as hard as the machines allow and see a number rise without consequence. Cavern closure by creep, which sets the minimum pressure in the first place, is absent [11].

The year is one seeded realization. The wind variability reproduces exactly, which is what makes the paired comparisons in this chapter meaningful, and it is a single sample from a distribution. Over seven other seeds the challenge margins move by about \$6 M and the heater-off default of the rolling plan ranges from 0.02 GWh to 0.59 GWh, with 0.47 GWh on the shipped seed. Halving the physics step from 30 to 15 minutes changes the suggested challenge's fuel heat from 23.08 to 22.25 GWh and its margin from \$15.28 M to \$15.36 M, while preserving all four verdict clauses. The hourly controller scan is unchanged; these default-step measurements should not be interpreted as numerically converged energy predictions. Nothing here supports a statement about expected value across weather years.

> ⚠ **OPEN CHALLENGE** **Keeping the Heat of Compression**
>
> A 2022 feasibility study named the energy lost as heat a significant gap. A plant that dissipates the heat of compression is limited to about 50 % efficiency, while recovering it could raise that to about 80 %. A molten-salt store helps, but its finite heat capacity limits the storage and adds cost [17]. Early studies found adiabatic plants too costly unless fuel was very expensive [5]. Progress would be a heat store large and cheap enough to outlast the gas, which is the gap this chapter's drought exposes.

11.10 Conclusion

A salt cavern holding 24.6 GWh is three and a half days at the expander rating, and calling it seasonal storage describes the calendar rather than the plant. What the plant is actually asked to do is be in the right state on the worst fortnight of the year, and the state that matters is not the one on the inventory gauge. Spending 84 % of the stored energy costs 38 % of the deliverable power, and a store cooled to its 320 K useful limit costs another third with the cavern full.

The measured comparison has two results and both are about heat. With the heater armed, every inventory policy and every reserve peak from 25 % to 85 % reports zero default over the hundred and fifty days, at a cost of 13 GWh to 27 GWh of heat, and a pair of price thresholds books \$0.77 M more than the rolling plan because in this price model the threshold responds to the same variable the plan forecasts. With the heater off, the threshold rule and the plan default on 0.56 GWh and 0.47 GWh while the operator curve, which holds the most gas, misses by 2 MWh, and the percentage reserve

reduces the plan's default without removing it. A floor written in megawatts is in the right units and computed at the wrong time, because a drought spends the heat before it spends the gas, and the reserve that protects the commitment is the one held before the store goes cold.

KEY POINTS

- A salt-cavern plant holds two inventories, air in the ground and heat in a packed bed, and the heat runs out first.
- Spending 84 % of the stored energy costs only 38 % of the deliverable power, so a commitment in megawatts is not a percentage of the cavern.
- With the supplemental heater armed every inventory policy keeps the commitment, and with it off the threshold rule and the plan default.
- A floor written in megawatts is in the right units but computed at the wrong time, because a drought spends the heat before the gas.

Where else this applies. An inventory held against a commitment also sets the policy of Volume 9, Chapter 7, and Chapter 10 tapers the limits of a store with two inventories in the same way.

References

[1] M. Abbaspour, M. Satkin, B. Mohammadi-Ivatloo, F. Hoseinzadeh Lotfi, and Y. Noorollahi, "Optimal operation scheduling of wind power integrated with compressed air energy storage (CAES)," *Renewable Energy*, vol. 51, pp. 53–59, 2013.

[2] J. B. Rawlings, D. Angeli, and C. N. Bates, "Fundamentals of economic model predictive control," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 2012, pp. 3851–3861.

[3] F. Crotogino and P. Quast, "Compressed-air storage caverns at Huntorf," in *Subsurface Space*. Elsevier, 1981, pp. 593–600.

[4] M. Budt, D. Wolf, R. Span, and J. Yan, "A review on compressed air energy storage: Basic principles, past milestones and recent developments," *Applied Energy*, vol. 170, pp. 250–268, 2016.

[5] S. Succar and R. H. Williams, "Compressed air energy storage: Theory, resources, and applications for wind power," Princeton Environmental Institute, Princeton University, Princeton, NJ, Report, Apr. 2008.

[6] E. Barbour, D. Mignard, Y. Ding, and Y. Li, "Adiabatic compressed air energy storage with packed bed thermal energy storage," *Applied Energy*, vol. 155, pp. 804–815, 2015.

[7] D. Wolf and M. Budt, "LTA-CAES: A low-temperature approach to adiabatic compressed air energy storage," *Applied Energy*, vol. 125, pp. 158–164, 2014.

[8] R. Kushnir, A. Dayan, and A. Ullmann, "Temperature and pressure variations within compressed air energy storage caverns," *International Journal of Heat and Mass Transfer*, vol. 55, no. 21-22, pp. 5616–5630, 2012.

[9] R. Kushnir, A. Ullmann, and A. Dayan, "Thermodynamic and hydrodynamic response of compressed air energy storage reservoirs: a review," *Reviews in Chemical Engineering*, vol. 28, no. 2-3, 2012.

[10] M. Raju and S. K. Khaitan, "Modeling and simulation of compressed air storage in caverns: A case study of the Huntorf plant," *Applied Energy*, vol. 89, no. 1, pp. 474–481, 2012.

[11] W. Li, X. Miao, and C. Yang, "Failure analysis for gas storage salt cavern by thermo-mechanical modelling considering rock salt creep," *Journal of Energy Storage*, vol. 32, p. 102004, 2020.

[12] N. A. Sepulveda, J. D. Jenkins, A. Edington, D. S. Mallapragada, and R. K. Lester, "The design space for long-duration energy storage in decarbonized power systems," *Nature Energy*, vol. 6, no. 5, pp. 506–516, 2021.

[13] J. D. Jenkins and N. A. Sepulveda, "Long-duration energy storage: A blueprint for research and innovation," *Joule*, vol. 5, no. 9, pp. 2241–2246, 2021.

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

[15] K. 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.

[16] Pipeline and Hazardous Materials Safety Administration, "49 CFR 192.12, underground natural gas storage facilities," Code of Federal Regulations, Title 49, Part 192, edition of October 1, 2024, govinfo.gov, 2024.

[17] H. Leetaru, "Illinois compressed air energy storage: Final report," Illinois State Geological Survey, University of Illinois, Report DE-FE0032019-FINAL, 2022.

11.11 Exercises

11.1 Concept check · 2 min

The thermal store is held at 470 K while the cavern pressure falls from 75 to 50 bar, and the deliverable power falls from 282.6 to 175.2 MW. Which explanation is right?

(a) The well flow limit of 820 kg/s binds at low pressure and caps the flow

(b) The expander swallows less air and each kilogram yields less work, and both fall with the pressure

(c) The thermal store cools as the cavern empties, so the expander inlet falls with the pressure

(d) The compressibility correction removes more air at low pressure than at high pressure

11.2 Concept check · 1 min

With the supplemental heater armed, every inventory policy keeps the challenge year's commitment with zero unserved energy, and the capacity floor reads zero at every store temperature. What keeps the commitment?

(a) The capacity floor, which holds back enough gas to make the commitment at any store temperature

(b) The seasonal reserve, which holds enough gas in the cavern to carry the plant through the drought

(c) The supplemental heater, which lifts the expander inlet so the capability stays above the commitment

(d) The threshold rule, which stops selling before the drought because the price follows the net load

11.3 Concept check · 2 min

The plan's forward simulation in Equation (11.25) lets the seasonal floor stop the plant selling. Suppose it were written instead as a clamp that lifts the state back up onto the floor whenever a block ends below it. What would the bisection for the water value report?

(a) A higher threshold, because the plan would then keep more gas in the ground

(b) The same threshold, because the floor binds only in the last days of the horizon

(c) A threshold that swings between re-solves, because the lifted state moves every twelve hours

(d) The price floor, because the end-of-horizon inventory would no longer depend on the threshold

11.4 Theory and calculation · 8 min

Use Equations (11.9) to (11.11) with $n_e = 2$, $\eta_t = 0.88$, $c_p = 1005\,\mathrm{J/(kg\,K)}$, $\kappa = 0.2857$, $P_{amb} = 1\,\mathrm{bar}$, a swallowing capacity of 780 kg/s at $P_{max} = 75\,\mathrm{bar}$, a store-to-air approach of 25 K, an expander rating of 300 MW and no supplemental heat.

(a) The cavern is at 60 bar and the store at 470 K. Find the swallowing capacity, the specific work and the deliverable power. *Checkpoint:* 624 kg/s, 348.6 kJ/kg and 217.5 MW.

(b) Scenario 2 starts at 53.9 bar with the store at 500 K. Find the deliverable power. *Checkpoint:* 204.5 MW, within 0.5 MW.

(c) At 60 bar, find the exhaust temperature with the store at 590 and at 450 K, and say whether the 260 K icing limit or the 150 MW commitment binds first as the store cools. *Checkpoint:* 345 K and 259 K, the icing limit, with 207.7 MW still deliverable.

11.5 Theory and calculation · 6 min

The plant holds two inventories with their own clocks. Use $C_r = 9.0 \times 10^{12}\,\mathrm{J/K}$, $UA_r = 2.6 \times 10^{6}\,\mathrm{W/K}$, $UA_\infty = 2.0 \times 10^{5}\,\mathrm{W/K}$, $C_{th} = 6.1 \times 10^{10}\,\mathrm{J/K}$ and $UA_{th} = 7.0 \times 10^{3}\,\mathrm{W/K}$.

(a) Find the time constants of the gas-to-rock and rock-to-far-field exchanges of Equation (11.5), in days. *Checkpoint:* 40.1 d and 521 d.

(b) Find the heat the packed bed holds between its limits of 320 and 590 K, and how long it lasts against the 332 MW that rated discharge draws. Compare with the 81.9 h of gas at rating. *Checkpoint:* 4.58 GWh and 13.8 h, about a sixth of the gas duration.

(c) Find the standing-loss time constant of the store and say which of the two drains empties it in practice. *Checkpoint:* 101 d. The expander's draw empties it.

11.6 Theory and calculation · 6 min

The loops scan every $\Delta t_c = 1/24$ d with a tracking time $T_t = \tau_I$, and the plant step is $h = 1/48$ d.

(a) While an output is clipped and the error and the running sum are held, Equation (11.20) is a first-order recursion in I with pole $1 - \Delta t_c/T_t$. Find the smallest τ_I for which it is stable and the τ_I below which it oscillates. *Checkpoint:* 0.0208 d (half an hour) and 0.0417 d (one hour).

(b) Find the pole for $\tau_I = 0.02$ d and for the override's slider minimum of 0.04 d. *Checkpoint:* -1.083 and -0.042.

(c) The heater lag is $\tau_{fuel} = 0.004$ d. Find the ratio h/τ_{fuel}, say whether classical Runge-Kutta, whose real-axis stability limit is 2.785, could carry this lag as a state, and find the share of a commanded step the exact solution reaches after one step. *Checkpoint:* 5.21, no, and 99.5 %.

11.7 App challenge · 7 min

Load scenario 2, A full cavern that cannot make its rating, and change no setting: the plant is in Manual with Expander power at 300 MW. Use the default seed.

(a) Before pressing Run, read the Deliverable tile and compare it with part (b) of Exercise 11.4. *Checkpoint:* 205 MW.

(b) Run one simulated day at speed 1 and export the CSV file. Read `dis_MW` in the row at `time_d` 0.0417 (one hour), then find the first row with `dis_MW` below 150 and read its time, `P_bar` and `Tth_K`. *Checkpoint:* 199.2 MW; 11 h (0.458 d), 50.5 bar and 399 K, within 1 MW, 0.1 bar and 1 K.

(c) Find the first row with `dis_MW` below 5 and read `soc_pct` and `floor_pct`. Explain why the expander stops with the inventory still above the floor. *Checkpoint:* 17 h (0.708 d), 15.0 % against 13.9 %.

11.8 App challenge · 8 min

Load scenario 3, The second inventory, and it empties first, and change no setting: Manual, Expander power 280 MW, Supplemental heater on Manual and Supplemental heat at 0 MW. Use the default seed and speed 1.

(a) Run one simulated day and export the CSV file. Find the first row with `Texh_K` below 260, and read `dis_MW` in the row at 15 h (`time_d` 0.625) and `metric_discharge_MWh` in the last row. *Checkpoint:* 7 h (257.7 K), 184.3 MW and 3820 MW h, within 1 MW and 10 MW h.

(b) Pause, press Reset, set Supplemental heat to 60 MW and run one day again. Read `metric_discharge_MWh` in the last row and the Fuel tile, and find the extra energy delivered per unit of fuel heat. *Checkpoint:* 4780 MW h, 1.74 GWh and 0.55.

(c) The chapter reports 1.20 GWh of first-stage heat for this run. Use $Q_{total} = Q_f[1+(n_e-1)(1-x)]$ with the cavern between 70 and 64 bar to explain the Fuel tile. *Checkpoint:* The factor is about 1.45, and $1.20 \times 1.45 = 1.74$.

11.9 App challenge · 9 min

Load the Challenge, From September to the far side of winter. Use the default seed and speed 30, at which a 150-day run takes about five seconds.

(a) Run all 150 days as shipped. Read the Unserved and Margin tiles, read `metric_minsoc` in the last row of the CSV file, and count the rows with `in_firm` equal to 1 and `available_MW` below 150. Say which of the four clauses fail. *Checkpoint:* 1.11 GWh, $14.36 M, 0.195 and 602 rows (25.08 d). All four fail.

(b) Reload, press Reveal suggested tuning and run 150 days. Read the same four values and the Fuel tile. *Checkpoint:* 0.00 GWh, $15.28 M, 0.249, 0 rows and 23.08 GWh.

(c) Reload, apply the suggested tuning, set Supplemental heater to Manual before pressing Run, and run 150 days. Read the Unserved and Margin tiles and count the short rows again. *Checkpoint:* 0.47 GWh, $14.55 M and 368 rows (15.33 d).

11.10 Application area · 4 min

The Huntorf and McIntosh plants are diabatic and burn gas before expansion, while the plant of this chapter is adiabatic with a supplemental heater.

(a) Explain why the supplemental heater turns the adiabatic plant into a hybrid when the store is cold, and what it costs at the challenge's suggested settings at the fuel price of $46 per megawatt-hour of heat. *Checkpoint:* About $1.06 M for 23.08 GWh.

(b) Name two consequences of an exhaust below its icing limit for a real expander, and explain why a lumped store temperature misplaces the moment the limit binds. *Checkpoint:* Ice from the moisture in the air on the last stage, and a thermocline that a lumped model cannot follow.

11.11 Application area · 3 min

A plant owner proposes to protect the winter commitment with a seasonal percentage reserve alone and to retire the supplemental heater.

(a) Use Table 11.9 to say what the reserve achieves with the heater off. *Checkpoint:* The default falls from 1.15 to 0.13 GWh for $1.70 M, and no setting meets the commitment.

(b) Say what a reserve would have to measure to protect a commitment written in megawatts, and why salt mechanics also limit how the cavern is worked. *Checkpoint:* The deliverable power, through gas and heat together, and creep-based pressure and slew limits.

11.12 Review the agent · 5 min

On day 5 of scenario 7, A two-week winter wind drought, an AI assistant reviews the screen and replies as follows. "The cavern sits on its floor at 24 % and the Deliverable tile reads 207 MW, which is 57 MW above the 150 MW commitment. The plant has headroom, so set the Supplemental heater to Manual at zero for the rest of the drought and save the fuel." Find the flaw in the advice and say how it could be checked in the app.

CHAPTER 12

A Solar Hydrogen Simulator for Safety Interlocks and the Question of Keeping the Oxygen

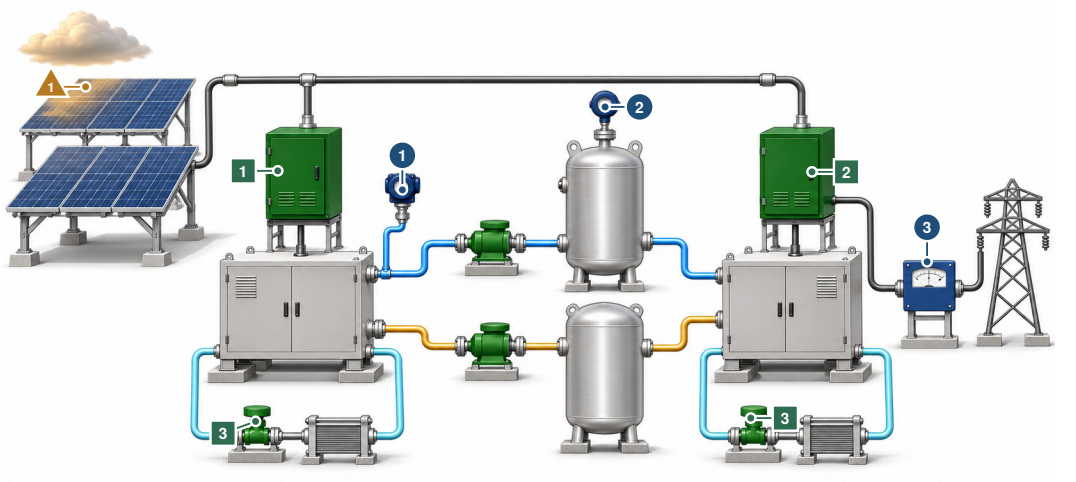

The controller moves 1 electrolyzer current (rectifier) 2 fuel cell current (converter) 3 cooling of each stack

The controller reads 1 analyzer for oxygen in the hydrogen 2 hydrogen store pressure 3 power delivered to the grid

What disturbs it 1 clouds over the array

Spare solar power splits water, and the hydrogen and oxygen travel on separate lines to separate tanks and meet again only inside the fuel cell. Illustration, not to scale.

IN PLAIN TERMS Saving Spare Sunshine as Hydrogen

The main idea. At midday a big solar farm makes more power than its line to the grid can carry. The spare power splits water into hydrogen and oxygen, which are stored. In the evening a fuel cell joins the two gases again to make power.

The challenge. The two gases must never mix, so safety checks, not prices, decide when the plant runs. A small leak between them can stay hidden for hours, until the fading sun makes less gas.

Why it matters. Only about a third of the power comes back, so control must waste no spare sunshine and must stop the plant before the gases meet.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the fact box of the Introduction, the two history boxes of Section 12.2, the try-it box of Section 12.5, the fact and try-it boxes of Section 12.7, and the key points
Study	40 min	The whole chapter with its pause points and the twelve exercises
Practice	10 min	The first page, Section 12.7, Section 12.8, Section 12.9 and the open challenge

The control question A solar hydrogen plant returns about a third of the electricity it takes in, and no controller changes that. What does control decide on such a plant, why must interlocks choose the mode while economics only choose setpoints inside it, and is the oxygen worth keeping?

Learning objectives

- Explain why the oxygen in the hydrogen product rises as the electrolyzer is turned down, and why a stack has a minimum load.
- Trace the supervisory machine through a trip, a start and a cloud, and name the role of each interlock, latch, dwell and hysteresis band.
- Decompose the advantage of stored oxygen into cell voltage and blower load, and weigh it against the compression it costs.
- Compare solar following, firm and economic dispatch on a net margin that prices hydrogen inventory and starts.
- Read a membrane fault from an analyzer whose reading depends on the load.

Prerequisites Sections 1.4.4 (inventories), 1.4.6 (water value), 1.4.7 (fast and slow layers), 1.6.2 (velocity form) and 1.6.4 (scenarios and CSV).

The app https://apmonitor.com/apps/solar-hydrogen/

See also Volume 8, Chapter 2 for interlocks as the safety layer beneath the feed, Volume 8, Chapter 13 for a supervisor that diverts and latches, and Volume 8, Chapter 14 for start-stop flapping. Volume 4, Chapter 6 for a limit set by an explosive mixture and Volume 4, Chapter 2 for a protective override.

12.1 Introduction

A water electrolyzer produces two gases in a fixed stoichiometric ratio and almost every commercial installation vents the oxygen. That is not an oversight and it is not obviously right either. Keeping the oxygen costs a second compressor, a second pressure vessel and a second set of safety rules that are not the ones hydrogen needs. It buys about 46 mV per cell at the fuel cell cathode from the Nernst term, the kinetics and the relieved concentration overpotential, and it removes the cathode air blower entirely. Measured on this plant at the rated point the two together are worth 14.3 % of the electricity returned per kilogram of hydrogen, of which 6.0 % is cell voltage and the rest is the blower that disappears. Whether that pays depends on how many hours the fuel cell runs, which makes it a dispatch question rather than a design one.

> **DID YOU KNOW? Ten Million Tonnes a Year**
>
> About 10 million metric tons of hydrogen are produced in the United States each year, mostly from natural gas, for petroleum refining, ammonia and other industrial uses [1]. Electrolyzers that split water with electricity were available in 2020 at the multi-megawatt scale [1].

The misconception this chapter corrects is that a hydrogen plant is an efficiency problem. A round trip of about a third is not going to be improved by control, and a student who spends an afternoon tuning the rectifier loops has learned nothing that matters. What control decides on this plant is narrower and more consequential. It decides whether the conversions that do happen are wasted, which means not curtailing array output the stack could have absorbed, not running the electrolyzer below the current at which crossover takes the product out of specification, not restarting a stack through every passing cloud, and not tripping.

The class of problem is supervisory sequencing under hard safety constraints with an economic layer subordinated to it. Everything dangerous about a hydrogen plant is a mixture, and every interlock in it exists to keep two streams that must never meet on opposite sides of a wall. A batch reactor with a runaway exotherm (Volume 8, Chapter 2), a furnace with a purge sequence, a compressor with an anti-surge lineup (Volume 4, Chapter 2) and an air separation unit all have the same shape.

The formal statement is that the mode is a finite state whose transitions are determined by interlocks alone, and the economics may only choose setpoints inside the mode that is currently legal. Economic model predictive control has a well-developed theory for the setpoint layer [2], and none of it licenses an optimizer to choose a valve lineup.

The chapter answers the oxygen question for one plant and splits the answer into its cell-voltage and balance-of-plant parts. A measurement then shows that the advantage in electricity per kilogram does not by itself decide the case. The simulator behind the answer carries a fifteen-state model of the two gases, a seven-state supervisory machine and a schedule solved by bisection, and it runs at interactive speed in a browser. One of its scenarios is a purity excursion whose fault and whose alarm are four and a half hours apart.

The chapter gives the model, the control structure and the measured behavior in enough detail that a reader can rebuild the simulator, and it uses them to show why the mode must be chosen by interlocks and the setpoint by economics and never the other way round. Section 12.2 traces the domain from the first wind hydrogen demonstrations. Section 12.3 gives the model of the array, the two stacks, the crossover, the stores and the water loop. Section 12.4 describes the open-loop character, and Section 12.5 the four regulatory loops and the minimum current. Section 12.6 gives the supervisory machine and the economic schedule, and Section 12.7 reports the scenarios: the oxygen decision in operation, the membrane fault, three architectures on three cloudy days and the graded challenge. Section 12.8 sets the real instruments, final elements and controllers beside the model, Section 12.9 states the limitations and Section 12.10 the conclusions, and the exercises follow.

12.2 From Utsira to a Firm Evening

Renewable hydrogen systems were demonstrated long before they were economic. The system simulation approach that most later work builds on was developed for advanced alkaline electrolyzers and published as a set of empirical current voltage and Faraday efficiency relations fitted to operating data [3]. That model was then used to design and evaluate the wind hydrogen demonstration on the island of Utsira in Norway, which supplied ten households from a wind turbine, an electrolyzer, a 2400 bar m^3 hydrogen store and a fuel cell, and whose published evaluation is unusually candid about what did not work [4]. The recurring finding from those projects is that availability and part-load behavior, rather than stack efficiency, set the usefulness of the system.

Proton exchange membrane electrolysis changed the operating envelope rather than the efficiency. A PEM stack follows a varying input in seconds where an alkaline stack takes minutes, tolerates differential pressure across the membrane, and delivers hydrogen at tens of bar without a first compression stage. The comprehensive review of the technology collects the cell voltage breakdown, the materials and the cost structure [5], and detailed electrical models of commercial units are available for the alkaline comparison [6].

What PEM does not remove is crossover. Hydrogen and oxygen both permeate the membrane, the flux is set by the partial pressure difference and the membrane rather than by the current, and the consequence is the single most important operating constraint on the technology. As a fraction of the gas being made the crossover climbs as the stack is turned down, so a stack idling at a few percent of rating is the dangerous place on the operating curve rather than the safe one. The pressurized case was measured and modeled directly, and the study that did so is the reason a modern electrolyzer specification quotes a minimum load rather than a turndown ratio [7].

The fuel cell side has its own control literature, developed largely for automotive stacks, whose central problem is oxygen starvation during a load transient and whose standard treatment is a modeling and feedback design text that remains the reference [8, 9]. A plant that stores its own oxygen removes that problem and replaces it with an inventory problem. The unitised regenerative concept, in which one stack does both jobs on a bifunctional electrode, has been pursued for decades and remains a materials problem rather than a control one [10], which is why the plant modeled here has two separate stacks.

HISTORY **Fuel Cells on Pure Oxygen in Space (1960s)**

Fuel cells first went to work on pure hydrogen and pure oxygen. A 1 kW polymer electrolyte fuel cell flew in the Gemini program and a 1.5 kW alkaline unit in the Apollo program, both in the 1960s, and these space units used pure reactant gases [11]. The Gemini stack had 32 cells, typically ran at 50 °C and 2 atmospheres, and gave 0.78 V at 37 mA/cm^2. Decades later a laboratory cell on pressurized hydrogen and air gave the same 0.78 V at about 200 mA/cm^2 and 80 °C [11]. A plant that stores its own oxygen feeds its fuel cell the way the spacecraft did.

The dispatch question arrived with the grid. Electrolyzer operation scheduling is now an active subject with a review of its own [12], and the system studies that size long-duration storage put hydrogen in the hundred-hour role rather than the four-hour one [13]. The plant modeled here sits between those two literatures, with a connection limit that makes curtailment the default and a firm evening commitment that makes the inventory matter.

HISTORY **Electrolyzers Tested as Grid Loads (2009 to 2014)**

A wind-to-hydrogen pilot project in North Dakota, reported in 2009, explored how an electrolyzer could soften the swings of wind output and the limits of the transmission lines, while it made hydrogen for two trucks and a tractor [14]. Its electrolyzer answered load changes in 3 to 7 s, and its maker recommended a minimum load of 43 % of rated power. Laboratory tests reported in 2014 found that electrolyzers of about 40 kW begin to change their demand within milliseconds of a new setpoint and settle within seconds, while a cold start of a PEM unit to full power took 6 minutes and 27 seconds [14]. The plant of this chapter has all three features: a connection limit, a minimum load and a start that costs time.

12.3 Process Model

The simulated plant is a 100 MW photovoltaic array on a 40 MW grid connection, an electrolyzer of 12 700 cells of 1000 cm^2 rated at 2.0 A/cm^2, a fuel cell of 29 000 cells rated at 1.2 A/cm^2, 1540 m^3 of hydrogen storage and 770 m^3 of oxygen storage at 200 bar. At the rated point the electrolyzer draws 49.8 MW and makes 941 kg/h, which is 52.9 kWh/kg. The fuel cell returns 21.5 kWh/kg on stored oxygen. Figure 12.1 shows the arrangement and every symbol used below. The clock is in hours, power in megawatts, mass in kilograms, pressure in bar and temperature in degrees Celsius except inside the Nernst and Arrhenius expressions, where it is in kelvin.

12.3.1 Variables

Table 12.1 classifies what the application exposes. The oxygen mode is an option rather than a continuous manipulated variable, and it is the decision the plant is built around.

12.3.2 The Array

Plane-of-array irradiance is a sine on the solar day raised to a power, and the array output carries the cell temperature derate that makes a clear noon in July worth less than a clear noon in April,

$$G = \sigma G_{max} \sin^{1.25}\left(\pi \frac{h - h_{up}}{h_{dn} - h_{up}}\right) \tag{12.1}$$

$$\begin{aligned} T_{cell} &= T_{amb} + \frac{G}{800}(T_{noct} - 20) \\ P_{pv} &= \frac{P_{pv,0} G}{1000}\left[1 - \beta (T_{cell} - 25)\right] \end{aligned} \tag{12.2}$$

With the defaults the array peaks at 89.3 MW at 12:48 on a 40 MW connection, so almost 50 MW of midday output has nowhere to go except the electrolyzer. Price is a gaussian evening peak on a base with a midday depression, running from −14.0 dollars per megawatt-hour at 12:30 to 227.9 at 19:30. The firm commitment is 18 MW from 17:00 to 23:00, which is 108 MWh.

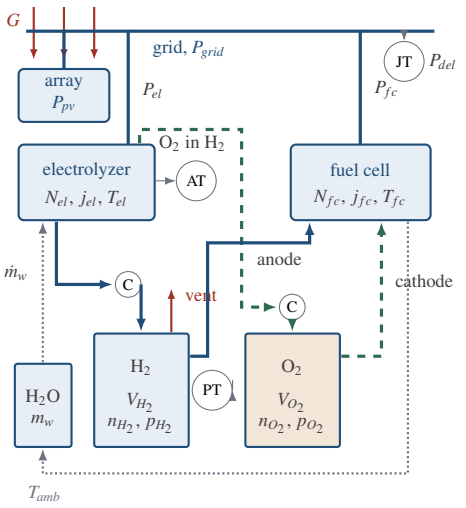

Figure 12.1 Plant arrangement and the symbols used in the equations. The array feeds the grid connection and the electrolyzer, and the part of the array output that exceeds the connection limit is what the electrolyzer is there to absorb. Both product gases pass through their own compressor into their own vessel, and the fuel cell draws hydrogen from one and either stored oxygen or air at the cathode. AT is the analyzer that measures oxygen in the hydrogen product, PT the hydrogen storage pressure, and JT the delivered power that the fuel cell loop controls. The water loop is nearly closed and the difference is purchased and deionized. Hydrogen streams are solid arrows, oxygen streams dashed arrows and water dotted arrows. The thick lines without arrowheads are the electrical connections.

12.3.3 The Electrolyzer

The cell voltage is the reversible potential plus an activation and an ohmic term. No concentration term appears because a PEM electrolyzer fed liquid water does not run out of reactant,

$$E_{rev} = E^0 - 9\times10^{-4}(T_k - 298.15) + \frac{RT_k}{2F}\ln\left(p_{H_2}\sqrt{p_{O_2}}\right) \tag{12.3}$$

$$V_{el} = E_{rev} + \frac{RT_k}{\alpha 2F}\operatorname{asinh}\frac{j}{2j_{0,el}} + R_{\Omega,el}\,j \tag{12.4}$$

$$R_{\Omega,el} = \frac{\mathrm{ASR}_{el}\,\omega_{el}}{\lambda}\exp\left[1600\left(\frac{1}{T_k} - \frac{1}{333.15}\right)\right]. \tag{12.5}$$

The hydration factor λ divides the resistance, so a dry membrane is a resistive one. At the rated 2.0 A/cm^2 and 62 °C the cell sits at 1.9612 V, made up of 1.2694 V reversible, 0.3809 V activation and 0.3110 V ohmic, and the stack draws 49.82 MW for 941 kg/h.

Production and heat follow from Faraday's law and the thermoneutral voltage,

$$\begin{aligned}\dot{n}_{H_2} &= \frac{\eta_F N_{el}\,jA_{el}}{2F}\\ \dot{Q}_{el} &= \frac{N_{el}(V_{el} - V_{tn})\,jA_{el}}{10^6}, \quad V_{tn} = 1.481\,\mathrm{V}\end{aligned} \tag{12.6}$$

Everything above the thermoneutral voltage is heat and everything below it is heat the stack has to be given, which is why a cold stack needs a heater and not a cooler. Temperature matters economically as well as thermally. At 30 °C the same rated current costs 58.0 kWh/kg against 52.9 kWh/kg at 62 °C, a difference of 9.6 % that the control system can do nothing about until the stack has warmed.

Table 12.1 Variable classification

Symbol	Meaning	Unit	Range
Controlled variables			
P_{del}	delivered power	MW	−40 to 40
P_{el}	electrolyzer load	MW	0 to 50
T_{el}	electrolyzer temperature	°C	40 to 80
T_{fc}	fuel cell temperature	°C	40 to 80
Manipulated variables			
j_{el}	electrolyzer current density	A/cm^2	0 to 2.7
j_{fc}	fuel cell current density	A/cm^2	0 to 1.8
u_{el}	electrolyzer cooling duty	%	0 to 100
u_{fc}	fuel cell cooling duty	%	0 to 100
χ_{O_2}	oxygen mode (store, reserve, vent)	—	—
Disturbances			
σ_∞	sky clarity	× clear	0.05 to 1.05
T_{amb}	ambient temperature	°C	−5 to 48
γ	membrane crossover multiplier	×	0.5 to 15
ω_{el}, ω_{fc}	stack wear multipliers	×	0.9 to 2.5
ε_{fc}	forecast error scale	×	0 to 1.5
States			
n_{H_2}	hydrogen in store	mol	to 200 bar
n_{O_2}	oxygen in store	mol	to 200 bar
T_{H_2}, T_{O_2}	vessel temperatures	°C	−40 to 120
T_{el}, T_{fc}	stack temperatures	°C	−10 to 110
λ	membrane hydration	—	0.2 to 1.15
m_w	deionized water inventory	kg	0 to 90 000
j^a_{el}, j^a_{fc}	applied current densities	A/cm^2	—
u^a_{el}, u^a_{fc}	applied cooling duties	—	0 to 1
c	oxygen in hydrogen	ppm	0 to 50 000
σ	filtered sky clarity	—	0 to 1.05
ϕ^a	applied compressor duty	—	0 to 1

12.3.4 Crossover and the Minimum Load

Oxygen permeates the membrane at a rate set by the pressure difference and the membrane rather than by the current, so as a fraction of the hydrogen being made it climbs as the stack is turned down [7],

$$c^* = \begin{cases} c_0\,\gamma\,\dfrac{j_{el,R}}{\max(j^a_{el},\,0.06\,j_{el,R})} & j^a_{el} > 0.01 \\ 0.35\,c_0 & \text{otherwise,} \end{cases} \tag{12.7}$$

with the analyzer and mixing volume as a first-order lag, $\tau_c\,\dot{c} = c^* - c$. The second branch matters as much as the first. With the stack stopped there is no gas being made and nothing to contaminate, so the analyzer reads down toward the header background rather than off the top of its scale, which is the difference between a plant that is idle and a plant that is about to be tripped. The reciprocal in the first branch is the whole reason a modern electrolyzer specification quotes a minimum load rather than a turndown ratio.

12.3.5 The Fuel Cell and the Oxygen

The fuel cell polarization carries the same reversible, activation and ohmic terms plus a concentration term, the four terms of the solid oxide stack in Chapter 3, and the oxygen mode changes three things

Table 12.2 Fuel cell at rated current and 66 °C on stored oxygen and on air

Quantity	Stored O_2	Air	Difference
Cell voltage (V)	0.8100	0.7639	+0.0461
reversible (V)	1.2182	1.2018	+0.0164
concentration (V)	0.0049	0.0146	+0.0097
flat kinetic term (V)	0.0200	0	+0.0200
Gross power (MW)	28.19	26.58	
Air blower (MW)	0	1.914	
Net power (MW)	28.19	24.67	
Net electricity (kWh/kg)	21.54	18.85	+14.3 %

at once,

$$p_a = \min(p_{H_2}, p_e), \quad p_c = \begin{cases} \min(p_{O_2}, p_e) & \text{stored } O_2 \\ 0.21 \min(1.5, p_e) & \text{air} \end{cases} \tag{12.8}$$

$$V_{fc} = E_{rev} - \eta_{act} - \eta_{ohm} - \eta_{conc} + \Delta V_{O_2} \mathbb{1}[\text{on } O_2] \tag{12.9}$$

$$\eta_{conc} = -\frac{RT_k}{2F} \ln\left(1 - \frac{j}{j_\ell}\right), \quad j_\ell = \begin{cases} 4.2\,\text{A/cm}^2 & \text{on } O_2 \\ j_{\ell,air} & \text{on air} \end{cases} \tag{12.10}$$

The electrode pressure clamp in Equation (12.8) is the detail that keeps the answer honest. A regulator drops both gases to a few bar before they reach the stack, so the Nernst benefit of stored oxygen is a couple of tens of millivolts and not the hundred a reader obtains by putting 200 bar into the logarithm. The cathode air blower is the third effect,

$$P_{blow} = \begin{cases} 0 & \text{on stored oxygen} \\ f_{blow} P_{gross} & \text{on air} \end{cases} \tag{12.11}$$

and it is the largest of the three. Table 12.2 decomposes the comparison at the rated point.

The decomposition is the useful part. Of the 14.3 %, 6.0 % is cell voltage and the remainder is a balance-of-plant load that stops existing. Against that, storing the oxygen costs 0.458 MW of compression at the rated electrolysis flow and a second pressure vessel of 770 m^3, and the compression is paid whenever the electrolyzer runs while the benefit is collected only when the fuel cell runs.

Pause. Try Exercise 12.3 before reading on.

12.3.6 Storage, Water and the Balance of Plant

Both vessels use the same real gas relation with one linear compressibility correction solved by three fixed-point passes, which at 200 bar is worth 12 % of the inventory and is not optional,

$$Z = 1 + Z_1 p, \qquad p = \frac{nR(T + 273.15)}{V} \frac{Z}{10^5} \ \text{[bar]}. \tag{12.12}$$

The hydrogen vessel holds 21 997 kg at 200 bar and 1468 kg at the 12 bar regulator limit, so the usable inventory is 20 529 kg, which is 15.7 h at the fuel cell rating, the duration of Section 1.4.4. Filling a pressure vessel warms it and the vessels are carried as separate temperature states with a 12.3 h loss time constant.

Compression is staged with intercooling from the electrolyzer delivery pressure to the vessel,

$$P_{comp} = \frac{n_s \dot{m} c_p (T_{amb} + 288.15)\,(r^\kappa - 1)}{\eta_{comp}\, 10^6}, \qquad r = \left(\frac{p_{ves}}{p_{stack}}\right)^{1/n_s} \tag{12.13}$$

which is 0.984 MW for hydrogen and 0.458 MW for oxygen at the rated flow into a full vessel.

The water loop is nearly closed and it is the constraint that stops the plant. Electrolysis consumes nine kilograms of water for every kilogram of hydrogen, which at the rated point is 8401 kg/h, and the fuel cell returns it only when it runs,

$$\frac{\mathrm{d}m_w}{\mathrm{d}t} = 3600\big(0.018\,\dot{n}_{H_2,fc} - 0.018\,\dot{n}_{H_2,el} + \dot{m}_{mk} - 0.4\,f_{purge}\big) \quad (12.14)$$

The default makeup rate is 0.12 kg/s, which is 432 kg/h against 8401 kg/h of consumption, so a full tank supports 10.2 h of rated electrolysis and no more.

12.3.7 Thermal and Hydration Dynamics

Both stacks are lumped capacitances with an ambient loss and a cooling term that fades as the stack approaches ambient,

$$C\frac{\mathrm{d}T}{\mathrm{d}t} = 3600\left[10^6\dot{Q} - UA(T - T_{amb}) - u^a Q_{cool}\,10^6\,\Pi\left(\frac{T - T_{amb}}{35}\right)\right] \quad (12.15)$$

The open-loop time constants are 2.78 h for the electrolyzer and 2.43 h for the fuel cell, both far longer than the current loops beneath them.

Membrane hydration is a first-order state driven toward a target that falls with current and with temperature and rises with purge,

$$\lambda^* = \zeta\,\Pi_{0.35}^{1.1}\left[1.05 - 0.22\frac{j_{el}^a}{j_{el,R}} - 0.012(T_{el} - 70)^+ + 0.10 f_{purge}\right] \quad (12.16)$$

with ζ the hydration disturbance and a time constant of 0.25 h. A dry membrane raises both stack resistances through Equation (12.5), so running hard and hot costs voltage twice.

12.3.8 Parameters

Table 12.3 gives the parameters the equations use, with their defaults. The V column marks A for adjustable and shown, H for hidden from the beginner view, and F for a fixed physical constant.

12.3.9 Initialization and Numerics

Each scenario supplies starting hydrogen and oxygen pressures and a starting stack temperature. The vessel temperatures start at ambient plus 8 °C, the hydration at unity, the analyzer at the rated crossover value and the water tank at 55 %.

The engine of Section 1.6.1 integrates by classical fourth-order Runge-Kutta on a fixed step of 0.001 h, which is 3.6 s and about two time constants of the rectifier lag. The regulatory loops execute every 0.002 h, which is 7.2 s, the schedule re-solves every 0.1 h, and the historian logs every 0.017 h, which is seventeen steps or 61.2 s. A post-step clip holds both inventories, the four temperatures, the hydration, the water inventory and the analyzer inside their physical ranges.

12.3.10 Assumptions

1. **Both stacks are one cell scaled by a cell count.** Cell-to-cell voltage distribution, which is what a real stack's monitoring system exists to watch, is absent. The pinhole in scenario ten is therefore a change in a plant-wide crossover multiplier rather than a defect in one cell.

Table 12.3 Model parameters and default values

Symbol	Meaning	Value	Unit	V
$P_{pv,0}$	array rating	100	MW	A
T_{noct}	nominal cell operating temperature	45	°C	H
β	cell power temperature coefficient	0.0038	1/°C	A
G_{max}	peak plane-of-array irradiance	1000	W/m^2	A
h_{up}, h_{dn}	sunrise, sunset	6.4, 19.2	h	A
N_{el}	electrolyzer cells	12 700	—	A
A_{el}	electrolyzer active area	1000	cm^2	A
$j_{el,R}$	rated electrolyzer current density	2.0	A/cm^2	A
ASR_{el}	electrolyzer resistance at 60 °C	0.16	Ω cm^2	A
$j_{0,el}$	anode exchange current density	1.0×10^{-6}	A/cm^2	H
α	anode charge transfer coefficient	0.55	—	H
η_F	Faradaic efficiency	0.985	—	A
C_{el}	electrolyzer thermal mass	2.6×10^8	J/K	A
UA_{el}	electrolyzer ambient loss	2.6×10^4	W/K	A
$Q_{cool,el}$	electrolyzer cooling capacity	14	MW	A
$T_{el,max}$	membrane temperature limit	80	°C	A
$j_{el,min}$	minimum electrolyzer current density	0.12	A/cm^2	A
N_{fc}	fuel cell cells	29 000	—	A
A_{fc}	fuel cell active area	1000	cm^2	A
$j_{fc,R}$	rated fuel cell current density	1.2	A/cm^2	A
ASR_{fc}	fuel cell resistance at 65 °C	0.13	Ω cm^2	A
$j_{0,fc}$	cathode exchange current on air	1.2×10^{-4}	A/cm^2	H
ΔV_{O_2}	flat gain from pure oxygen	0.020	V	A
$j_{\ell,air}$	limiting current density on air	1.9	A/cm^2	A
C_{fc}	fuel cell thermal mass	1.4×10^8	J/K	A
UA_{fc}	fuel cell ambient loss	1.6×10^4	W/K	A
$Q_{cool,fc}$	fuel cell cooling capacity	30	MW	A
$T_{fc,hot}$	nominal fuel cell ready temperature	62	°C	A
$x_{fc,min}$	minimum fuel cell load	0.18	—	A
f_{blow}	cathode blower, fraction of gross	0.072	—	A
V_{H_2}, V_{O_2}	storage volumes	1540, 770	m^3	A
p_{max}	storage design pressure	200	bar	A
p_{min}	regulator limit	12	bar	A
p_e	electrode operating pressure	3.0	bar	H
p_{stack}	electrolyzer delivery pressure	30	bar	A
Z_1	compressibility slope	6.0×10^{-4}	1/bar	H
UA_{ves}	vessel heat loss, each	900	W/K	A
C_{ves}	vessel heat capacity, each	4.0×10^7	J/K	A
n_s	compressor stages	3	—	A
η_{comp}	compressor isentropic efficiency	0.70	—	A
$m_{w,max}$	water tank capacity	90 000	kg	A
$m_{w,min}$	minimum working water	9000	kg	A
c_0	oxygen in hydrogen at rated current	40	ppm	A
c_{max}	purity limit before venting	1800	ppm	A
τ_c	analyzer and mixing response	0.08	h	A
P_{grid}	grid connection limit	40	MW	A
P_{firm}	firm power commitment	18	MW	A
h_{f1}, h_{f2}	firm window	17, 23	h	A
c_{H_2}	value of hydrogen	4.2	$/kg	A
c_{unmet}	penalty for missing the commitment	600	$/MWh	A
τ_j	rectifier current response	4.2×10^{-4}	h	A
τ_{cool}	cooling response	0.014	h	A
τ_λ	hydration response	0.25	h	A
F	Faraday constant	96 485	C/mol	F
R	gas constant	8.314	J/(mol K)	F

2. **Crossover is a scalar in the product stream.** The physical process is a permeation flux driven by partial pressure across the membrane [7], and the reciprocal-current form in Equation (12.7) reproduces its operational consequence without its pressure dependence. A pressurized stack would show a stronger effect.

3. **The gases are ideal with one compressibility correction.** At 200 bar and ambient temperature the single linear term in Equation (12.12) is adequate for inventory accounting and is not a real equation of state. Hydrogen's negative Joule-Thomson coefficient at these conditions is not modeled.

4. **The fuel cell has no reactant transport dynamics.** Oxygen starvation during a load transient is the central control problem in the automotive fuel cell literature [8] and it does not appear here, because the plant is fed from pressure vessels rather than from a compressor with its own dynamics.

5. **Membrane hydration is one state for two stacks.** A real plant has separate water management on each, and the shared state overstates the coupling between them.

6. **Degradation is a disturbance slider rather than a mechanism.** The stack wear multipliers are set by hand and do not accumulate with operation, so nothing in the model connects the start count to the stack life it is meant to stand for.

7. **The irradiance and the price are synthetic.** They are teaching data shaped like a clear spring day with an evening peak and are never labeled as measured.

12.4 Open-Loop Character

The round trip is the first thing manual mode makes concrete and the least interesting. Fifty kilowatt-hours of electricity make a kilogram of hydrogen and the kilogram returns about twenty-one, which is a round trip near a third and is a property of the electrochemistry rather than of the control system. No scenario in this application improves it and none of them tries.

The minimum load is the second thing and it is the one worth an hour of a student's time. Turning the electrolyzer down is not free and it is not linear. At the rated 2.0 A/cm^2 the analyzer reads 40 ppm, at 0.4 A/cm^2 it reads 200 ppm, and at 0.12 A/cm^2 it reads 667 ppm, all of it from a crossover flux that did not change. The specification limit is 1800 ppm and the model floors the current ratio at 6 % of rating, so on a stack with nominal membranes the worst reading is 667 ppm and the limit is never reached. A membrane defect changes that completely, and scenario ten is built on it.

🕮 **Pause.** Try Exercise 12.1 before reading on.

The time scales separate cleanly and the separation is what makes the architecture obvious. The rectifiers answer in 1.5 s. The analyzer and the mixing volume answer in 0.08 h, which is about five minutes. The stacks have thermal time constants near 2.5 h. The hydrogen inventory is 15.7 h at the fuel cell rating and the water inventory is 10.2 h at the electrolyzer rating. The day is the outer loop.

Curtailment is the third thing and it is a definition rather than a fault. The array peaks at 89.3 MW on a 40 MW connection, so at solar noon almost 50 MW has nowhere to go except the electrolyzer, and everything the electrolyzer cannot take is thrown away. A plant of this shape is sized so that curtailment is normal operation, which is why the electrolyzer's setpoint in every automatic architecture is the surplus above the connection limit rather than a demand.

12.5 Regulatory Control

12.5.1 The Discrete Algorithm

Four loops share the velocity-form proportional-integral-derivative implementation of Section 1.6.2, Equations (1.10) to (1.13), with the derivative taken on the measurement and back-calculation against

Table 12.4 Regulatory loops as implemented

CV	MV	K_c	τ_I (h)	Action
P_{del}	j_{fc}	0.09	0.004	reverse, +1
P_{el}	j_{el}	0.05	0.004	reverse, +1
T_{el}	u_{el}	9	0.06	direct, −1
T_{fc}	u_{fc}	9	0.05	direct, −1

the applied output. The settings particular to this plant are a tracking time $T_t = \tau_I$, an output bias of zero, no rate limit and $\Delta t_c = 7.2\,\text{s}$. The running sum is never clipped, so anti-windup acts on the integral alone. A loop that is not in service (a cooling loop in Manual, or a current loop whose stack is down or held at its start-up load) solves Equation (1.11) for its integral so that its output equals the value actually applied, which makes the return to service bumpless. The delivered power reaches its loop through a 0.002 h first-order filter and the stack temperatures through 0.01 h filters. All four loops run with $\tau_D = 0$, so there is no derivative action anywhere in this application [15].

12.5.2 Four Loops and Two Parked Setpoints

Table 12.4 gives the pairings and the default tunings, with the action in the convention of Section 1.3 (+1 reverse acting, −1 direct acting), which is also the sign the app's controller carries. The two current loops are reverse acting because more current raises the delivered or the absorbed power. The two cooling loops are direct acting because more cooling makes a colder stack.

Both current loops have output limits that are recomputed every scan rather than fixed, because the current that produces a requested power moves with temperature, hydration and storage pressure. Each limit is found by forty-eight bisections on the corresponding polarization relation.

The two cooling setpoints are parked on their measurements whenever the stack they belong to is down. A temperature setpoint on a stack with no current through it is not a control target but a number, because the stack cools to ambient and the loop can do nothing about it. Parking the setpoint is what a real plant's temperature controller does when its unit is off, and it is the difference between an idle loop and a thirty-degree error that nobody can act on and that contaminates every integrated performance metric in the scorecard. The application logs and scores the parked value for both stacks, so on the fifth scenario the fuel cell loop reports an integrated error of 26.7 °C h by 19:04, of which only 0.3 °C h accumulated in the five hours before the stack was started and the rest while it warmed from 26 °C toward its 66 °C setpoint. Figure 12.4 shows the application at that moment.

12.5.3 The Minimum Current as a Discontinuity

Below the minimum current density the electrolyzer is not merely inefficient, it is on its way out of specification. The loop output is therefore snapped rather than clipped,

$$j_{el} = \begin{cases} j_{el}^{PI} & j_{el}^{PI} \geq j_{el,min} \\ j_{el,min} & j_{el}^{PI} < j_{el,min},\ P_{el}^{SP} > 0.5\,\text{MW} \\ 0 & \text{otherwise} \end{cases} \tag{12.17}$$

which turns a continuous manipulated variable into one with a forbidden band. The snap is applied after the controller, whose own lower limit is zero, so while the floor holds the current above what the loop asks for, the integral winds down to that zero limit. In the second scenario the 2 MW request leaves the controller output at zero and the applied current at the 0.12 A/cm^2 floor, and a later rise in the request has to carry the controller output back up through the floor before the current moves. The same treatment is applied to the fuel cell at 18 % of its rated current. A student who has met this on a real plant recognizes it immediately and a student who has not is usually surprised that it is done deliberately.

Table 12.5 Supervisory state transitions

From	Condition	To
TRIP	interlocks clear	SHUTDOWN
SHUTDOWN	always	OFF
OFF	any demand and clear	STARTUP
STARTUP	not clear	TRIP
STARTUP	fuel cell wanted, warm, H_2 present	FUELCELL
STARTUP	electrolyzer wanted, warm, water	ELECTROLYZE
STARTUP	no demand	SHUTDOWN
ELECTROLYZE	not clear	TRIP
ELECTROLYZE	fuel cell wanted	STANDBY
ELECTROLYZE	no room, no water or no demand	STANDBY
STANDBY	not clear	TRIP
STANDBY	fuel cell wanted and warm	FUELCELL
STANDBY	fuel cell wanted and cold	STARTUP
STANDBY	electrolyzer wanted, room, water	ELECTROLYZE
STANDBY	no demand	SHUTDOWN
FUELCELL	not clear	TRIP
FUELCELL	no H_2 or not wanted	STANDBY

TRY IT IN THE APP **2 · The Minimum Load, and Why It Is Not a Turndown Ratio**

Load scenario 2 and press Run. Watch the electrolyzer current strip. When the scenario asks for 2 MW, the current stops at its floor of 0.12 A/cm^2 and goes no lower, while the controller output beneath it has wound down to zero.

12.6 The Supervisory Machine

The mode is a finite state with seven values and explicit transitions. The transition function depends only on an interlock vector and a demand vector, and the economics contribute to the demand vector alone,

$$m_{k+1} = \Phi(m_k,\ \mathbf{ok},\ \mathbf{want}). \tag{12.18}$$

The interlock vector is evaluated from measurements alone. It carries whether the plant is clear, which requires the analyzer inside specification whenever the stack is running, both stack temperatures below 82 °C (two degrees above their limits), the hydrogen vessel below 212 bar (6 % over its design pressure) and the water above a hard floor of 2700 kg. It also carries whether there is water, which is latched off at the 9000 kg working minimum until the makeup has refilled the tank to 10 800 kg, whether there is room for more hydrogen, latched off at 199 bar until the store falls below 194 bar, whether there is hydrogen above the regulator limit, and whether each stack is warm enough to be loaded, which for the electrolyzer means above 8 °C and for the fuel cell above $T_{fc,hot} - 14\,°\mathrm{C}$, or 48 °C. On the fifth scenario the fuel cell is started at 17:04 and loaded at the scan at 17:46, at 48.0 °C, and it reaches the 62 °C of $T_{fc,hot}$ only at 18:17. Table 12.5 gives the transitions.

Three properties of that table are worth stating because they are design decisions rather than consequences. Every path into TRIP depends on interlocks alone, so no economic setting can prevent a trip and no economic setting can cause one. Every path out of TRIP goes through SHUTDOWN and OFF, so recovery is a sequence and not a toggle. And STARTUP is not an idle state. A cold fuel cell cannot be loaded and the only heat available to it is its own inefficiency, so STARTUP holds the stack at its minimum load until it is warm enough to be useful. That is what a start costs and it is why the scorecard counts them.

Pause. Try Exercise 12.2 before reading on.

Table 12.6 Scenarios

#	Start	Architecture	Subject
1	10:00	manual	two stacks and the round trip
2	09:00	manual	the minimum load
3	08:00	follow	the connection limit and curtailment
4	06:00	follow	solar following to a full store
5	12:00	firm	the firm evening commitment
6	12:00	firm	should the oxygen be kept
7	09:00	follow	a cloud front and the state machine
8	05:00	economic	what a stored kilogram is worth
9	05:00	economic	three days of low solar
10	10:00	follow	a purity excursion and a trip
11	05:00	open	the graded challenge

A minimum dwell of 0.08 h applies to every transition except one into TRIP. Real supervisory machines all have one, for the same reason, which is that the physical act of starting a stack is not instantaneous and a machine that can leave a state faster than the plant can enter it is a machine that will.

The demand vector carries hysteresis for the same reason. The electrolyzer start threshold is a fraction of the array rating, and once the stack is running the threshold drops to 55 % of that value. Without it a surplus wandering across the threshold starts and stops the stack every few seconds, and the start counter fills with chatter instead of with the effect it was built to show. A server fleet scaled without such a band flaps in the same way in Volume 8, Chapter 14.

Three architectures write the demand vector. Solar following runs the electrolyzer on surplus and the fuel cell only when the array cannot meet the commitment. Firm runs the fuel cell through the whole commitment window. Economic adds a schedule that computes what a kilogram in the bottle is worth, by simulating thirty-two half-hour blocks forward at a trial threshold price and bisecting sixteen times for the lowest threshold that still ends on the terminal inventory,

$$\text{run the fuel cell in block } k \iff \pi_k \, \varepsilon_{kg} > \theta \tag{12.19}$$

where ε_{kg} is the megawatt-hours a kilogram returns, measured at the rated point rather than assumed and therefore different on oxygen and on air. The threshold θ is the value of a stored kilogram, the water value of Section 1.4.6, found here by forward simulation and bisection rather than by a merit order.

12.7 Scenarios and Measured Results

Table 12.6 lists the eleven scenarios and Table 12.7 gives the scorecard at the end of each window with every setting as loaded. In the two manual cases the cooling is left at zero, so the electrolyzer reaches its 110 °C model limit in the first. In the third and eighth the water tank reaches its working minimum at 15:37 and 15:38, and in the fourth the store fills at 15:32, so the electrolyzer stops for the rest of the afternoon and most of their curtailment comes after it.

> **DID YOU KNOW? Cloudy Days Throw Away Almost Nothing**
>
> From the simulator: seventy-two hours at 30 % of clear-sky sunshine in scenario 9 curtail only 1.3 MWh, while the ten-hour run of scenario 3 on a clear day curtails 20.9 MWh, most of it after the water tank has reached its working minimum (Table 12.7).

12.7.1 The Oxygen Decision in Operation

The sixth scenario switches the fuel cell from stored oxygen to an air cathode at 18:00, an hour into the evening commitment, and its briefing asks for the twelve-hour run of the fifth scenario to

Table 12.7 Scorecard at the end of each scenario window

#	Hours	Unserved (MWh)	Curtailed (MWh)	H_2 made (kg)	Starts	Margin ($)
1	3	0.00	12.4	2196	1	6 399
2	3	0.00	78.6	483	1	483
3	10	0.00	20.9	4867	1	25 504
4	14	2.44	23.1	4820	2	25 241
5	12	0.00	9.0	3240	2	13 287
6	12	0.00	9.0	3240	2	11 066
7	8	0.00	4.2	3864	2	15 187
8	20	0.01	22.2	4867	2	22 533
9	72	2.83	1.3	4820	4	15 656
10	6	0.00	4.6	4536	1	17 202
11	20	57.60	56.4	571	9	−44 664

be repeated once on each policy. The rated-point comparison in Table 12.2 says stored oxygen is worth 14.3 % more electricity per kilogram, and the twelve-hour pair says something more useful. On stored oxygen the run consumes 3684 kg of hydrogen and 3.41 MWh of compression to deliver its commitment. On air it consumes 4200 kg and 2.46 MWh of compression, and the cathode blower draws 6.83 MWh and raises the auxiliary load from 7.87 MWh to 13.75 MWh. Storing the oxygen therefore costs 0.95 MWh of extra compression, saves 5.88 MWh of auxiliary load net of it over the same window, and saves 516 kg of hydrogen.

The net margin includes delivered electricity and oxygen sales, missed-commitment penalties, the value of hydrogen inventory changes and stack-start costs. At $4.20 per kilogram, the extra hydrogen consumed by the air cathode reverses its cash advantage. Stored oxygen nets $13 287 against $12 830 on air, an advantage of $457. The inventory ledger includes vents and storage losses, and each movement is valued at the configured price when it occurs. Start costs are incurred at each counted start, so changing a price mid-run does not reprice earlier activity.

12.7.2 The Fault and the Alarm

The tenth scenario opens a pinhole at 11:30 that raises the crossover rate twenty-fold, and the plant does not trip for four and a half hours. At noon the stack is at 1.87 A/cm^2, 94 % of rated current, and the analyzer reads 858 ppm, close to the 800 ppm that Equation (12.7) gives for a twenty-fold crossover at rated current and comfortably inside the 1800 ppm limit. The reading climbs only when the surplus falls in the afternoon, to 1664 ppm at 16:00 with the current at 0.91 A/cm^2, and the plant trips at the scan at 16:06, when the same leak has become a much larger fraction of a much smaller hydrogen flow. The machine restarts the stack and trips twice more, at 16:30 and 16:51, before the surplus is gone.

> **TRY IT IN THE APP** **10 · A Purity Excursion, and a Trip**
>
> Load scenario 10 and press Run. Watch the purity trace beside the electrolyzer current. The pinhole opens at 11:30, yet the reading is 858 ppm at noon. It climbs only as the current falls in the afternoon, and the plant trips at 16:06.

Nothing about the plant changed between the fault and the alarm except its load. A monitoring calculation that normalized the analyzer reading by the current would have seen the step immediately, and the reason plants do not usually have one is that nobody asked the analyzer what it was for. That observation transfers directly to any composition measurement on a stream whose flow rate is a manipulated variable. Figure 12.2 draws Equation (12.7) for the healthy and the leaking membrane with the readings of this section.

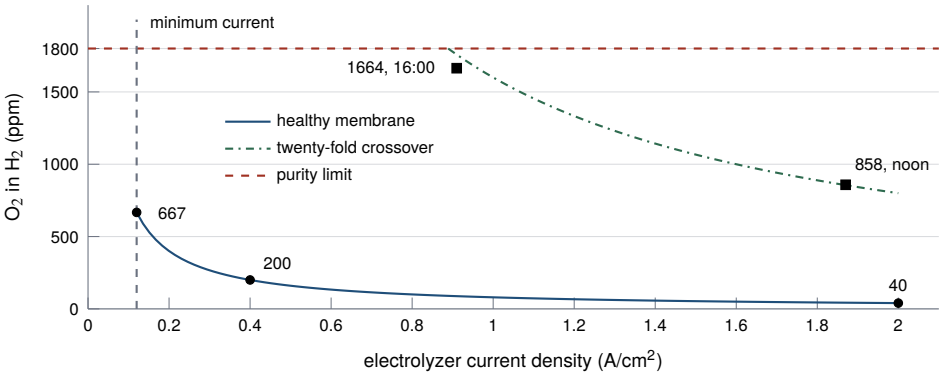

Figure 12.2 Oxygen in the hydrogen product against current density at steady state, from Equation (12.7), for a healthy membrane (solid) and for the twenty-fold crossover of scenario 10 (dash-dot), with the purity limit (dashed). Dots are the readings of Section 12.4 and squares the readings of scenario 10 at noon and at 16:00, which trail the curve by the analyzer's lag. Notice that the same leak sits far inside the limit at full current and reaches it only when the current falls.

Table 12.8 Three architectures over seventy-two hours of low solar

Architecture	Unserved (MWh)	Fuel cell (MWh)	Final H_2 (bar)	Margin ($)
solar following	5.08	273.05	92.2	15 468
firm	2.83	281.62	89.1	16 508
economic	2.83	321.85	70.3	15 656

12.7.3 Three Architectures on Three Cloudy Days

The ninth scenario runs seventy-two hours at 30 % of clear-sky irradiance with a firm evening commitment on each day and the bottle starting at 150 bar. Table 12.8 compares the three architectures on it.

The economic schedule delivers 321.85 MWh from the fuel cell against 273.05 MWh on solar following and leaves 2.25 MWh less unserved. Its cash advantage is \$10 101, but pricing the depleted store reduces its net-margin advantage to \$188. Firm dispatch earns \$16 508, compared with \$15 656 for the economic schedule and \$15 468 for solar following. The schedule forecasts energy and inventory, but does not optimize the complete plant ledger including starts. More electricity sold is therefore not necessarily more profit. Figure 12.3 draws the table as bars.

12.7.4 The Graded Challenge

The eleventh scenario opens twenty hours from five in the morning with weak tuning, vented oxygen, an 18 % start threshold and stack setpoints of 74 °C and 76 °C against 33 °C ambient. It trips eight times on stack temperature, with a logged electrolyzer peak of 81.8 °C, leaves 57.6 MWh unserved, curtails 56.4 MWh and nets −\$44 664. Retuning the four loops removes the trips and gives 0.04 MWh unserved, a 78.6 °C peak and \$5 467 net margin, below the \$6 000 target.

The suggested configuration uses solar following, stored oxygen, a 15 % start threshold, temperature setpoints of 62 °C and 66 °C and the default tunings. It leaves 0.068 MWh unserved, trips not at all, peaks at 67.6 °C, curtails 10.42 MWh, starts twice and nets \$6 959. Its electricity and oxygen cash margin is \$20 012, reduced by \$8 653 of hydrogen inventory depletion and \$4 400 of starts. Five additional forecast seeds give the same result because solar following does not use that forecast.

With the remaining settings as suggested, firm dispatch nets \$6 601 and the economic schedule \$6 277. The latter sells more electricity but also burns more stored hydrogen. Lowering the start

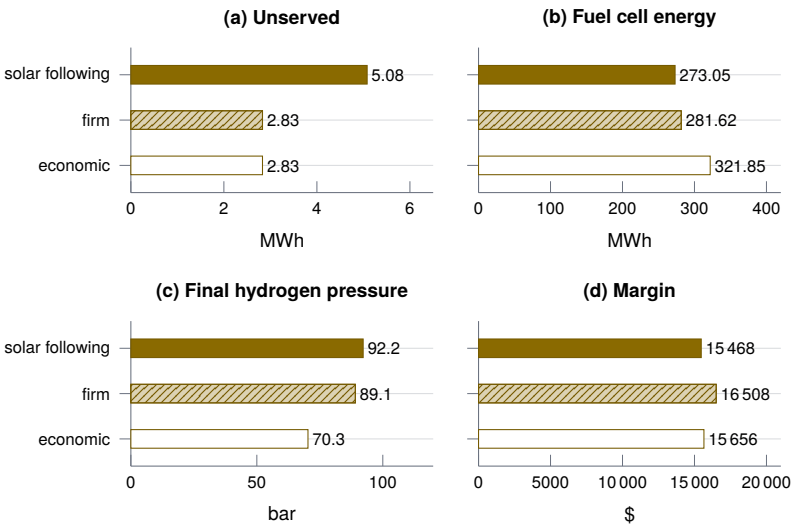

Figure 12.3 The three architectures over seventy-two hours of low solar, from Table 12.8: (a) unserved energy, (b) energy from the fuel cell, (c) hydrogen pressure at the end and (d) net margin. Solid: solar following. Hatched: firm. White: economic. Notice that the economic schedule sells the most energy and leaves the emptiest store, and that firm dispatch, not the economic schedule, earns the most margin.

threshold to 4 % on solar following reduces curtailment to 2.81 MWh, adds a third start and lowers net margin to $5 436. Venting oxygen lowers net margin to $5 546. Applying the suggested structure to weak loops still fails: 5.64 MWh unserved, an 81.2 °C peak and $1 480 margin. Neither tuning alone nor structure alone passes the revised challenge.

12.8 Technology in Practice

The symbols of Figure 12.1 are, in a plant, an analyzer on the hydrogen product, a pressure transmitter on the store, thermometers in both stacks, power measurements on the electrolyzer feed and at the grid connection, two power converters, two cooling systems and a valve that stores or vents the oxygen. Figure 12.5 draws the chapter's loops with that hardware in place of the schematic's symbols.

What Measures It

The measurement that decides whether the plant runs is the oxygen in the hydrogen product, a composition read by an analyzer with a sample system. The chapter lumps the analyzer and the mixing volume into one 0.08 h lag and gives no sample system that can fail. The stack temperatures and the two power measurements serve the regulatory loops, and the store pressure is the inventory of Equation (12.12).

> **TECHNOLOGY** **Sensor: The Oxygen-in-Hydrogen Analyzer**
>
> Most industrial analyzers sit away from the process and receive a sample through tubes, filters, regulators and heaters, and many analyzer problems start in that sample system [16]. A gas sample is held at steady pressure and temperature, kept clean and dry, and carried in metal tubing when the species of interest is at parts per million, because some plastic tubes let gases through. One oxygen sensing principle is the micro fuel cell, which makes a current from oxygen by oxidizing a lead fuel inside it and must be replaced as that fuel is used up [16]. A plugged sample line leaves an analyzer reading old gas with no way of knowing, which on a trend is a purity reading that stops following the current. The chapter's AT is exact behind a 0.08 h lag.

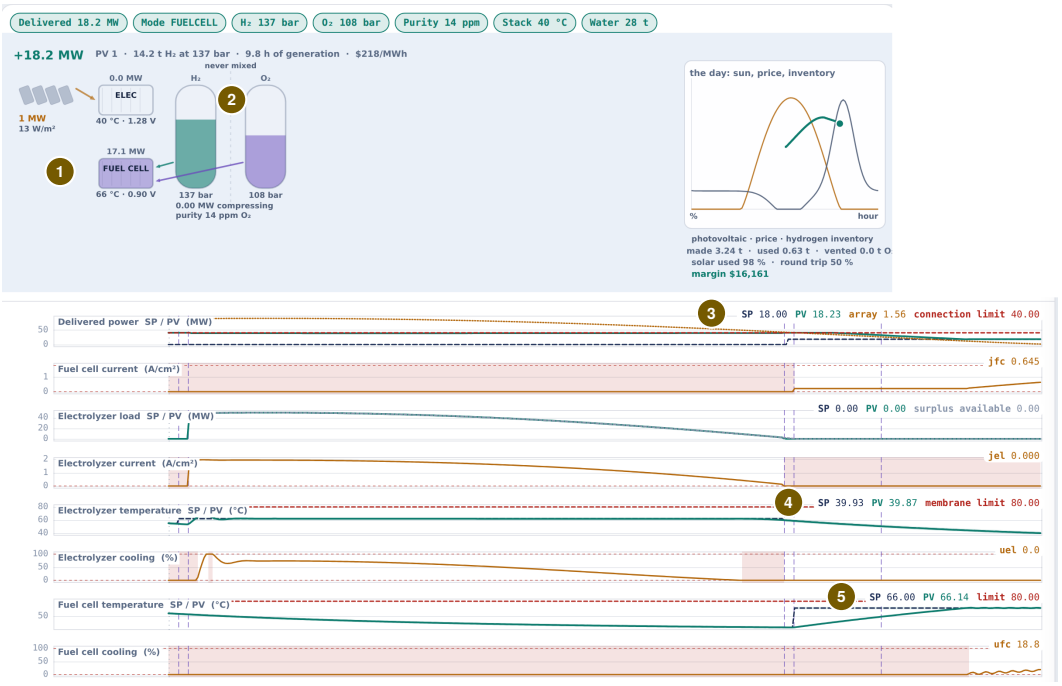

Figure 12.4 The application at 19:04 on the fifth scenario, seven hours into the run and two hours into the firm evening commitment, with the fuel cell running on stored oxygen. The delivered power strip reads 18.23 MW against a setpoint of 18 MW while the array contributes 1.56 MW, and the two current strips show the electrolyzer at zero and the fuel cell at 0.645 A/cm^2. The electrolyzer temperature strip shows its setpoint parked on the measurement (39.93 °C against 39.87 °C) since the stack stopped at 16:59. The fuel cell setpoint follows its measurement down until STARTUP at 17:04, then steps to 66 °C, and the loaded stack now reads 66.14 °C. The mode badge reads FUELCELL, which is the only state in which the economics are permitted to choose anything, and the scene labels the wall between the two vessels *never mixed*. The plant view and the trend panel are cut from the full window. Callouts: (1) the running fuel cell, (2) the wall marked *never mixed*, (3) the delivered power strip, (4) the parked electrolyzer setpoint, (5) the fuel cell setpoint after its step.

What Moves It

Five final elements act on the plant. The rectifier and the fuel cell converter set the two stack currents within limits recomputed every scan. The two cooling duties hold the stack temperatures. The oxygen valve chooses between the store and the vent. The power electronics between the grid and an electrolyzer can also offer the grid services of power quality and power factor correction, which the laboratory tests of 2014 could not assess [14].

> **TECHNOLOGY** **Final Element: The Rectifier Under the Current Loop**
>
> An electrolyzer takes its direct current from rectifiers and other power electronics that convert the alternating current of the grid [14]. In a controlled rectifier, thyristors take the place of plain diodes, and a phase-control circuit sets the output by timing the trigger pulse of each thyristor against the alternating waveform. A drive with a feedback loop of its own acts as the slave of a cascade and takes its setpoint from the process controller [16]. Tested units began to change their demand within milliseconds and settled within seconds [14]. The chapter's rectifier lag of 1.5 s stands for that response, and its current loops are the masters that write the setpoint.

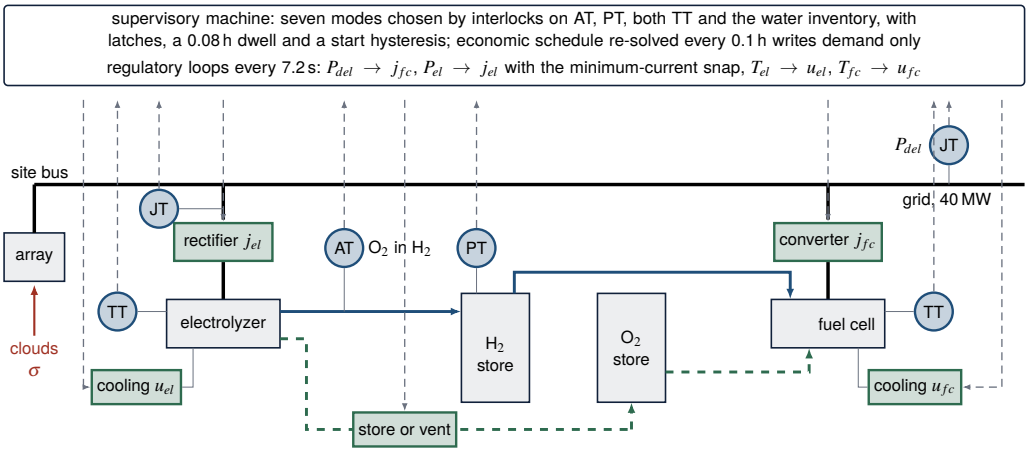

Figure 12.5 The chapter's loops with their hardware. Sensors are the round tags: the analyzer AT on the hydrogen product, the pressure transmitter PT on the hydrogen store, thermometers TT on both stacks, and power measurements JT on the electrolyzer feed and at the grid connection. Final elements are the square tags: the electrolyzer rectifier, the fuel cell converter, two cooling duties and the valve that stores or vents the oxygen. Solid arrows carry hydrogen and dashed arrows oxygen. Notice that the analyzer and the pressure transmitter feed only the supervisory machine, which chooses the mode, while the four regulatory loops act inside it.

Where the Controller Runs

A programmable logic controller holds instructions for logic, timing, counting and proportional-integral-derivative control, and in a distributed system it acts as a local controller under a supervisory scheme [17]. Such loops run with cycle times from milliseconds to minutes. The chapter's regulatory loops run every 7.2 s on a 3.6 s physics step, and the schedule re-solves every 0.1 h. The interlocks belong in a safety instrumented system, made of sensors, logic solvers and final elements that bring the plant to a safe state when a threshold is crossed. It is designed to IEC 61511 and is often independent of the basic control system, so that a failure of the control system does not weaken it [17]. The chapter's machine, whose transitions to TRIP depend on interlocks alone, is that separation written as a table.

> ⚠ **IN PRACTICE** **Two Gases, Two Sets of Rules**
>
> United States workplace rules treat the two products of this plant in separate sections, 29 CFR 1910.103 for hydrogen systems and 1910.104 for bulk oxygen systems on a user's premises [18, 19]. Hydrogen containers carry relief devices arranged to discharge upward and unobstructed to the open air [18]. Bulk oxygen equipment is cleaned of oil, grease and other readily oxidizable materials before it enters service, its gaskets, valves, gauges and regulators must be suitable for oxygen service, and its storage must not be exposed by flammable gas lines [19]. Keeping the oxygen brings a second vessel under a second rule.

What Is Modernizing It

Three developments touch this plant. The first is the grid role of the electrolyzer. Its subsecond dynamic response, validated in national laboratory projects that simulate grid fluctuations, enables services such as voltage and frequency stabilization [1]. In one test on an isolated grid of a diesel generator, a load bank and an electrolyzer, the frequency after a load step was restored faster with the electrolyzer's help than without it [14]. The second is the stack itself. Reversible fuel cells, still under development, would split water and make electricity in one unit [1], the role that the chapter's two separate stacks share. The third is the network. Operational systems increasingly resemble

information systems as they adopt standard computers and protocols, which leaves them less isolated and in greater need of security, and some modern safety systems are designed to communicate with the control system [17].

12.9 Limitations

The scorecard is an operating-margin ledger rather than a project valuation. It counts electricity and oxygen cash flows, missed-commitment penalties, actual hydrogen inventory changes and starts. It omits capital, financing and maintenance beyond the synthetic per-start charge. Hydrogen is valued at the parameter's configured price when its inventory moves, including leaks and vents. The economic dispatcher remains a forecast-based energy schedule, rather than a complete optimization of this ledger. The round-trip tile compares electricity per kilogram in each direction and is independent of the initial inventory.

The crossover model has no pressure dependence, where the measured process is driven by the partial pressure difference across the membrane [7]. The reciprocal-current form reproduces the operational consequence at one delivery pressure, and a student who changes the delivery pressure slider will see nothing happen to the analyzer, which is wrong.

Both stacks are single cells scaled by a count, so the pinhole in scenario ten is a change in a plant-wide multiplier rather than a defect in one cell of twelve thousand. Real plants detect that fault through cell voltage monitoring long before the product analyzer moves, and nothing in this model can pose that detection problem.

The irradiance and price are synthetic and the results should be read as the shape of an answer rather than as its magnitude. In particular, the economic schedule delivers the most fuel cell energy in Section 12.7 partly because the price here is a fixed daily shape independent of the array output, which is a favorable case for a forecast.

> **OPEN CHALLENGE** **Cheap Hydrogen From Cheap Electricity**
>
> In 2020 hydrogen from low-temperature electrolysis cost $5 to $6 per kilogram at electricity prices of 5 to 7 cents per kilowatt-hour. Electricity at 2 to 3 cents from new wind and solar plants, with better electrolyzers, offers a path below $2 per kilogram, but consistent, widely available low-cost electricity still needs work, and research is still needed to cut the cost and raise the efficiency and durability of electrolyzers [1]. The cheap midday power of this chapter's plant comes in bursts, and its stack wear does not grow with its starts. Progress would be a stack life known as a function of starts and load, so that a start could be priced from wear.

12.10 Conclusion

A solar hydrogen plant returns about a third of what it takes in, and no control system is going to change that. What the control system decides is whether the conversions that do happen are wasted, and the decisions that matter are discrete rather than continuous. The mode is chosen by interlocks, the setpoint is chosen by economics inside that mode, and the ordering is the architecture rather than a detail of it. An optimizer that can select a valve lineup can select an unsafe one.

Two measured results carry beyond this plant. Keeping the oxygen is worth 14.3 % of the electricity returned per kilogram of hydrogen, of which only 6.0 % is cell voltage and the rest is a cathode blower that stops existing, so the case for it is an operating case decided by a duty cycle and not a thermodynamic one. And a crossover flux that is constant in absolute terms is a rising fraction of a falling product flow, so a membrane defect that appears at half past eleven in the morning raises its first alarm just after four in the afternoon, with nothing having changed in between except the load. A specification written as a concentration and a plant operated on a flow are two different statements about the same stream, and the gap between them is where the alarm was.

KEY POINTS

- A solar hydrogen plant returns about a third of the electricity it takes in, and control decides only whether the conversions that do happen are wasted.
- The mode is chosen by interlocks and the setpoint by economics inside that mode. An optimizer that can select a valve lineup can select an unsafe one.
- Keeping the oxygen is worth 14.3 % of the electricity returned per kilogram, and only 6.0 % of it is cell voltage.
- A constant crossover is a rising fraction of a falling hydrogen flow, so the membrane fault raised its alarm hours after it appeared.

Where else this applies. A limit set by an explosive mixture also governs Volume 4, Chapter 6, and a band against start-stop flapping serves Volume 8, Chapter 14.

References

[1] U.S. Department of Energy, "Department of energy hydrogen program plan," U.S. Department of Energy, Washington, DC, Report, 2020.

[2] J. B. Rawlings, D. Angeli, and C. N. Bates, "Fundamentals of economic model predictive control," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 2012, pp. 3851–3861.

[3] Ø. Ulleberg, "Modeling of advanced alkaline electrolyzers: a system simulation approach," *International Journal of Hydrogen Energy*, vol. 28, no. 1, pp. 21–33, 2003.

[4] Ø. Ulleberg, T. Nakken, and A. Eté, "The wind/hydrogen demonstration system at Utsira in Norway: Evaluation of system performance using operational data and updated hydrogen energy system modeling tools," *International Journal of Hydrogen Energy*, vol. 35, no. 5, pp. 1841–1852, 2010.

[5] M. Carmo, D. L. Fritz, J. Mergel, and D. Stolten, "A comprehensive review on PEM water electrolysis," *International Journal of Hydrogen Energy*, vol. 38, no. 12, pp. 4901–4934, 2013.

[6] A. Ursúa and P. Sanchis, "Static-dynamic modelling of the electrical behaviour of a commercial advanced alkaline water electrolyser," *International Journal of Hydrogen Energy*, vol. 37, no. 24, pp. 18 598–18 614, 2012.

[7] M. Schalenbach, M. Carmo, D. L. Fritz, J. Mergel, and D. Stolten, "Pressurized PEM water electrolysis: Efficiency and gas crossover," *International Journal of Hydrogen Energy*, vol. 38, no. 35, pp. 14 921–14 933, 2013.

[8] J. T. Pukrushpan, A. G. Stefanopoulou, and H. Peng, *Control of Fuel Cell Power Systems*, ser. Advances in Industrial Control. London: Springer London, 2004.

[9] ——, "Modeling and control for PEM fuel cell stack system," in *Proceedings of the 2002 American Control Conference*, 2002, pp. 3117–3122.

[10] S. Altmann, T. Kaz, and K. A. Friedrich, "Bifunctional electrodes for unitised regenerative fuel cells," *Electrochimica Acta*, vol. 56, no. 11, pp. 4287–4293, 2011.

[11] U.S. Department of Energy, National Energy Technology Laboratory, "Fuel cell handbook," U.S. Department of Energy, Office of Fossil Energy, National Energy Technology Laboratory, Morgantown, WV, Report, Seventh Edition, Contract DE-AM26-99FT40575, Nov. 2004. [Online]. Available: https://netl.doe.gov/sites/default/files/netl-file/FCHandbook7.pdf

[12] J. Wang, J. Wen, J. Wang, B. Yang, and L. Jiang, "Water electrolyzer operation scheduling for green hydrogen production: A review," *Renewable and Sustainable Energy Reviews*, vol. 203, p. 114779, 2024.

[13] N. A. Sepulveda, J. D. Jenkins, A. Edington, D. S. Mallapragada, and R. K. Lester, "The design space for long-duration energy storage in decarbonized power systems," *Nature Energy*, vol. 6, no. 5, pp. 506–516, 2021.

[14] J. Eichman, K. Harrison, and M. Peters, "Novel electrolyzer applications: Providing more than just hydrogen," National Renewable Energy Laboratory, Golden, CO, Technical Report NREL/TP-5400-61758, 2014.

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

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

[18] Occupational Safety and Health Administration, "29 CFR 1910.103, hydrogen," Code of Federal Regulations, Title 29, Part 1910, edition of July 1, 2024, govinfo.gov, 2024.

[19] ——, "29 CFR 1910.104, oxygen," Code of Federal Regulations, Title 29, Part 1910, edition of July 1, 2024, govinfo.gov, 2024.

12.11 Exercises

12.1 Concept check · 2 min

In scenario 2 the electrolyzer is turned down from 30 MW to 6 MW on a healthy membrane, and the oxygen-in-hydrogen analyzer settles at a reading more than four times higher. Which explanation is right?

(a) The membrane grows more permeable at low current, so the crossover flux rises
(b) The crossover flux stays the same while the hydrogen flow falls with the current
(c) The analyzer lag overshoots after a load reduction and has not yet settled back
(d) A cooler stack at low load dissolves more oxygen into the hydrogen product

12.2 Concept check · 1 min

Which statement about the supervisory machine of Table 12.5 is right?

(a) The economic schedule may send the plant to TRIP when the price is negative
(b) Recovery from TRIP returns the plant directly to the state it was in before
(c) Every path into TRIP depends on interlocks, and every path out passes SHUTDOWN and OFF
(d) STARTUP holds a cold stack at zero current until it has warmed to its ready point

12.3 Concept check · 2 min

A reader puts the 200 bar storage pressure into the Nernst term and predicts about a hundred millivolts per cell from running the fuel cell on stored oxygen. Table 12.2 shows 16.4 mV of reversible gain. Which statement explains the difference and names the largest part of the 14.3 % advantage?

(a) The vessels are cooler than the stack, and most of it is the concentration term
(b) The compressibility removes most of the gain, and most of it is the kinetic term
(c) The analyzer limits the oxygen purity, and most of it is the reversible potential
(d) A regulator drops both gases to a few bar, and most of it is the blower that is removed

12.4 Theory and calculation · 6 min

Use Equation (12.7) in steady state ($c = c^*$) with $c_0 = 40\,\text{ppm}$, $j_{el,R} = 2.0\,\text{A/cm}^2$ and the purity limit $c_{max} = 1800\,\text{ppm}$.

(a) With nominal membranes ($\gamma = 1$), find the reading at 0.4 A/cm^2 and at the 0.12 A/cm^2 minimum current. *Checkpoint:* 200 ppm and 667 ppm.

(b) Without the 6 % floor in the equation, at what current density would a nominal stack reach the limit? Say why the model can never show it. *Checkpoint:* 0.044 A/cm^2, below the 0.12 A/cm^2 floor.

(c) After the pinhole of scenario 10 ($\gamma = 20$), find the reading at the rated current and the current density below which the steady reading exceeds the limit. *Checkpoint:* 800 ppm and 0.889 A/cm^2.

12.5 Theory and calculation · 6 min

The electrolyzer makes 941 kg/h of hydrogen at its rated point, and Equation (12.14) charges 0.018 kg of water for each mole of hydrogen made (molar mass 2.016 g/mol). The makeup is 0.12 kg/s, the tank holds 90 000 kg, the working minimum is 9000 kg and the water interlock releases at 10 800 kg. The fuel cell is off and the purge is zero.

(a) Find the water consumption at the rated point and the net drain rate. *Checkpoint:* 8402 kg/h and 7970 kg/h.

(b) Find how long a full tank supports rated electrolysis, and how long the 55 % starting inventory does. *Checkpoint:* 10.2 h and 5.08 h.

(c) Once the interlock has stopped the stack, how long does the makeup alone take to release it? *Checkpoint:* 4.17 h.

12.6 Theory and calculation · 7 min

Use the rated-point values of Table 12.2 and the oxygen compression of 0.458 MW at the rated electrolysis flow of 941 kg/h.

(a) From the table, find the gain in cell voltage, the gain in net electricity per kilogram, both as percentages of the air values, and the blower load as a fraction of gross power on air. *Checkpoint:* 6.0 %, 14.3 % and 0.072.

(b) Find the oxygen compression energy per kilogram of hydrogen made and the extra electricity per kilogram burned on stored oxygen. *Checkpoint:* 0.487 kW h and 2.69 kW h.

(c) Storing the oxygen costs compression on every kilogram made and pays only on the kilograms burned. Find the fraction of the hydrogen made that must later be burned on stored oxygen for the energy to break even. *Checkpoint:* 0.181.

12.7 App challenge · 8 min

Load scenario 2, The minimum load, and change no setting. The schedule asks for 6 MW at 0.5 h and for 2 MW at 1.6 h. Run at 1x with the default seed past 2.6 h, pause and export the CSV file. Use the last rows at or before 0.45 h, 1.5 h and 2.5 h.

(a) Read `jel_A_cm2` and `purity_ppm` in the three rows. *Checkpoint:* 1.296, 0.288 and 0.120 A/cm^2; 61.4, 278.2 and 666.7 ppm, within 1 ppm.

(b) Multiply each reading by its current density and compare the products with $c_0 j_{el,R}$. *Checkpoint:* About 80 in every row.

(c) Read `Pel_MW` in the last row against the 2 MW request, and explain the difference. *Checkpoint:* 2.43 MW.

12.8 App challenge · 8 min

Load scenario 5, The firm evening commitment, and run at 5x with the default seed to about 12 h, then pause and export the CSV file. Pause, Reset, set Oxygen policy to Vent before Run, and repeat. In each file use the last row at or before 12 h.

(a) Read the Margin tile at the pause in each run. *Checkpoint:* \$13,287 on Store and \$12,830 on Vent, within \$5.

(b) Read `metric_h2_used_kg` in each file, and compare the ratio of the two with the 14.3 % rated-point advantage. *Checkpoint:* 3684 kg and 4200 kg, within 2 kg, a ratio of 1.140.

(c) Read `metric_cash_margin` and `metric_hydrogen_value` in each file, and explain which policy the cash alone would favor and why the Margin tile ranks them the other way. *Checkpoint:* \$19,610 and −\$1,924 on Store; \$21,320 and −\$4,090 on Vent, within \$5.

12.9 App challenge · 10 min

Load scenario 11, the graded challenge, and run all 20 h at 5x with the default seed. Export the CSV file at the end of each run.

(a) Run the case as given. Read the Unserved and Margin tiles, the last value of `metric_trips`, and the clock time of the first trip from that column. *Checkpoint:* 57.60 MWh, −$44,664, 8 trips, the first near 13:50.

(b) Pause, Reset, and set the eight tunings to the defaults of Table 12.4 before Run. Read Unserved, Curtailed and Margin, and the largest `Tel_C` in the file. *Checkpoint:* 0.04 MWh, 10.7 MWh, $5,467 and 78.6 °C.

(c) Reset again and keep the tunings. Set Oxygen policy to Store, the start threshold to 15 % and the two stack temperature setpoints to 62 °C and 66 °C. Read the same values, then check that the last-row `metric_cash_margin` less the hydrogen value loss and `metric_start_cost` gives the margin. *Checkpoint:* 0.07 MWh, 10.4 MWh, $6,959 and 67.6 °C.

Give money within $5 and temperatures within 0.1 °C.

12.10 Application area · 4 min

A plant engineer is asked whether a new electrolyzer, whose fuel cell will run only as backup for a few hours a month, should keep its oxygen. Using the chapter, list what keeping the oxygen costs and what it buys, and give a recommendation with its reason.

12.11 Application area · 3 min

A vendor quotes a PEM electrolyzer with a turndown of ten to one. Explain why a specification should quote a minimum load instead, and what a supervisor should do with a surplus below that load.

12.12 Review the agent · 5 min

An AI assistant reviews the trend of scenario 10 and writes: “After the pinhole at 11:30 the analyzer read 858 ppm at noon, under half of the 1800 ppm limit. The membrane is therefore within specification, and the plant can run until the analyzer says otherwise. Dividing the reading by the current would add nothing, since the analyzer already measures the concentration directly.” Find the flaw and say how you would check it with the app.

CHAPTER 13

Supervisory Economic Dispatch of Nuclear Heat, Storage and Hydrogen Production

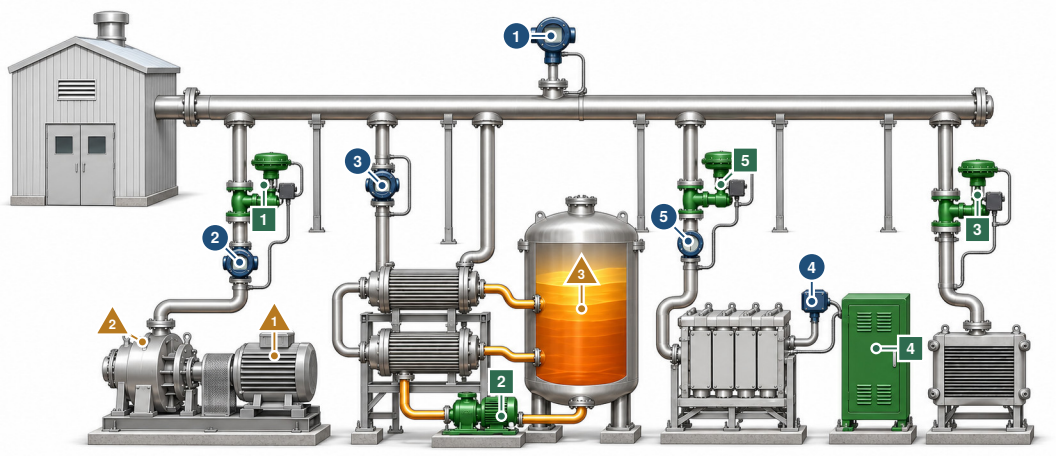

The controller moves 1 turbine throttle valve 2 salt pump, charging or discharging the store 3 dump valve to the condenser 4 electrolyzer power supply 5 superheater on the stack steam

The controller reads 1 steam header pressure 2 steam flow to the turbine 3 heat into or out of the store 4 electrolyzer power 5 stack temperature

What disturbs it 1 electricity price and its forecast 2 turbine outage 3 smaller usable tank

Notice that every branch leaves one header fed by steady reactor heat, so whatever the turbine, the salt store and the electrolyzer do not take goes to the dump condenser at the right. Illustration, not to scale.

IN PLAIN TERMS Steady Reactor Heat, Sold at the Best Hour

The main idea. A nuclear reactor makes the same heat all day, but the price of electricity changes every hour. The plant can sell the heat as electricity, keep it in a tank of hot melted salt, or use it to make hydrogen fuel.

The challenge. All of the heat must go somewhere at every moment, and heat that nobody buys is thrown away. On a normal day, a plan that looks a whole day ahead earns only a little more than three simple price rules.

Why it matters. Knowing on which days a smarter planner pays off keeps a plant from buying one for the wrong reason.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the fact box of the Introduction, the two history boxes of Section 13.2, the fact and try-it boxes of Section 13.6, and the key points
Study	55 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 13.6, Section 13.7, Section 13.8 and the open challenges

The control question A reactor that cannot be turned down sells into a market whose price moves sixteen to one across a day. How much does a supervisory economic optimizer earn over three price thresholds an engineer can defend, and on which days does its horizon pay for itself?

Learning objectives

- Write the heat balance of the steam header and compute how fast a residual moves the header pressure.
- Derive the electrolyzer switching price and the arbitrage value of a thermal megawatt-hour from the plant's parameters.
- Explain how the dispatch program eliminates the turbine, keeps its Hessian constant and carries the contract as a one-sided penalty.
- Compare fixed, threshold and optimizing supervisors on one day, with the heat left in the tank at midnight counted.
- Identify the constraint or forecast error that makes a horizon pay on a hard day.

Prerequisites Sections 1.4.2 (the header as stored inertia), 1.4.4 (the store balance), 1.4.5 (constant heat downstream), 1.4.6 (water value), 1.6.2 and 1.6.4. Chapter 10 for storage dispatch on a forecast.

The app https://apmonitor.com/apps/nuclear-hydrogen/

See also Volume 9, Chapter 6 for a price of heat set by predictive control and Volume 9, Chapter 7 for a release planned on an ensemble forecast. Volume 18, Chapter 9 for one header shared by three reactor modules and Volume 18, Chapter 5 for the steam dump after a load rejection.

13.1 Introduction

The usual justification for a supervisory optimizer is that it beats the heuristic underneath it. That claim is rarely presented with both numbers attached, and on the plant of this chapter it is only half true. Over an ordinary day, three price thresholds an engineer can defend at a meeting book $516,300 against a fixed allocation's $470,311, and a 48-variable economic optimization over a twenty-four hour horizon books $529,303. The heuristic captures 78 % of the available improvement and the optimization collects the remaining $13,003. Raise the evening price to five and a half times the day-ahead curve, leave the forecast blind to it, and the same two controllers separate by $85,447. A horizon is insurance, and insurance looks expensive on the days it is not needed.

The plant that produces that result is a reactor whose output cannot move. Nuclear heat is capital rather than fuel, so no price justifies turning a reactor down, and the flexibility has to be built downstream of it [1, 2]. Four sinks take 860 MWth of reactor heat in this application. A steam turbine rated at 1180 MWth of swallowing capacity can absorb more than the reactor makes, which is what gives a store somewhere to discharge into. A two-tank molten salt store holds 2400 MWh at a 95.0 % thermal round trip [3]. A solid-oxide electrolyzer draws 180 MWe of stacks at 800 °C and takes header steam as well as electricity [4]. A dump condenser takes whatever is left and is worth nothing. Grid export swings from −102.7 MWe to 431.7 MWe off a heat source that never moves.

> **DID YOU KNOW? A Quarter of a Million Tonnes From One Reactor**
>
> A feasibility study for an existing United States plant found that a standard four-loop pressurized water reactor of about 1095 MWe could make at most about 240 thousand tonnes of hydrogen a year by high-temperature steam electrolysis. Piped to some six users, that output needed about 150 miles of pipeline [5].

What a simulator of that plant offers over a block diagram is the constraint itself. The header stores 3.0 MWh of steam and drum inventory per bar against 860 MW of supply, which is thirteen seconds of reactor output per bar, so the allocation equation is not a modeling convenience but a physical identity that closes every instant. An operator moving three sliders discovers that they are not three degrees of freedom, that turning the turbine down saves nothing because the heat still has to go somewhere, and that whatever the three do not account for leaves through a valve nobody is paid for.

Neither the identity nor its consequences appear in a transfer function from one input to one output, and the misconception the plant corrects is the one stated above, that a better optimizer is always worth having. A three-threshold rule is very hard to beat on a smooth day with no active constraints, and published assessments of control practice make the same point about advanced control in general, which is that the value comes from the constrained and the coupled cases rather than from the average one [6, 7].

The class of problem generalizes. A supervisory layer allocates one inflexible resource among competing consumers, against a forecast that is wrong, over a horizon long enough that terminal constraints bind, with a fast regulatory layer underneath. A refinery hydrogen network distributes make-up hydrogen from a reformer that does not turn down among hydrotreaters that compete for it [8]. A combined heat and power plant allocates one boiler between a steam host and an electricity market whose price moves [9]. A data center shifts a flexible fraction of its compute load in time against a carbon or price signal it must forecast first [10]. A reservoir operator spends a finite inflow across a season and prices the water left in storage against what the rest of the horizon will pay for it [11], which is exactly the marginal value that sets the charge and discharge thresholds here. In every one of them the supervisory problem is a linear or quadratic program over hours to months and the regulatory problem is PID over seconds [6, 12].

However, the layer that decides the allocation is normally met as a formulation on paper, with no means of watching a plan being made, watching it be wrong, and watching feedback recover part of it. This chapter treats a browser simulation of an integrated nuclear energy system in enough detail to be rebuilt from the text, and measures what a supervisory economic optimizer is and is not worth against the heuristic it replaces. Section 13.2 traces nuclear cogeneration and flexible operation. Section 13.3 collects the equations and constants that reproduce the simulator, and Section 13.4 derives the three prices that decide what a good allocation looks like. Section 13.5 describes the six regulatory loops and the three supervisors, among them an economic quadratic program whose Hessian is constant, whose turbine allocation is eliminated rather than optimized, and whose fixed eighteen projected-gradient iterations cost 23.4 µs per execution. Section 13.6 measures the three supervisory architectures on identical deterministic days, decomposes the improvement into what the horizon buys and what is left sitting in the tank at midnight, and puts a wrong forecast in front of a controller that acts on it with complete confidence. Section 13.7 names the instruments, final elements and control platform behind the loops, and Sections 13.8 and 13.9 state the limitations and what the results mean for practice.

13.2 Nuclear Cogeneration and Flexible Operation

Using reactor heat for something other than electricity is as old as civil nuclear power, and desalination was the first serious proposal [13], as the history segment after this paragraph tells. Chemical process heat followed the same reasoning at higher temperature, and the high-temperature gas-cooled reactor was developed in part to reach the 700 °C to 950 °C that industrial reforming, gasification and hydrogen production need. The economic argument in every one of those studies was the same as the one this application makes, which is that a capital-intensive heat source is cheapest when its output is constant and therefore wants a second customer whenever the first one does not want the heat.

⧖ HISTORY **Reactor Heat for Water and Industry (1973 to 2008)**

On the eastern shore of the Caspian Sea a 750 MWth sodium-cooled fast reactor started up in 1973 and, until 1999, supplied electricity, drinking water and heat to the city of Aktau. It sent up to 125 MWe to the grid and made 80 000 m^3 of potable water a day [14]. In Switzerland a 970 MWe pressurized water reactor has sent process steam to a nearby cardboard factory since 1979, about 2 % of its thermal power. At ten coastal sites in Japan, desalination units beside the reactors have made make-up and drinking water for more than 30 years, and in 2008 a heavy water reactor in India was retrofitted with a desalination plant [14].

Light-water plants were nevertheless built and licensed to run flat out, for reasons that are partly economic and partly physical. With fuel a small fraction of the busbar cost, every hour spent below full power is revenue given away against a cost that does not fall, so the merit order places nuclear at the bottom and leaves it there. The physical reasons are the ones that make load following an engineering exercise rather than a dispatch instruction. Reactivity feedback, control rod worth, xenon transients after a power reduction and fatigue on thick-walled components all restrict how far and how often a plant may be moved, and a profitability evaluation that carries those constraints explicitly finds the flexible operating envelope narrower and more expensive than a purely economic model suggests [1]. Fleets in France and Germany did operate flexibly at scale, and did so within exactly those limits.

The economic pressure changed when zero-marginal-cost generation reached a large share of the market. The value of a megawatt-hour is no longer a single number, because the average price a technology captures falls as its own penetration rises, and the effect is strong enough to be measured across European markets [15]. A firm generator therefore faces a price that varies by more than an order of magnitude inside one day and goes negative when solar output exceeds demand. Two answers appeared, and both put storage or a second product between the reactor and the market rather than asking the reactor to move. Thermal storage on the steam side lets a constant-output reactor sell a variable amount of electricity [2], and sizing that store for assured peak production rather than for arbitrage alone changes both the store and the turbine [16], which is why the turbine in this application swallows 1180 MWth against a reactor that makes 860 MWth.

The current programs make this a coordinated system problem. Nuclear-renewable hybrid energy systems are defined by the coupling between the reactor, the grid and an industrial process, and the research needs identified for them are led by the supervisory control question of how to dispatch the flexible elements rather than by the reactor itself [17]. Hydrogen is the industrial product that has attracted most of the effort, because high-temperature steam electrolysis consumes both electricity and heat and its electrical demand falls as its operating temperature rises [4]. Feasibility evaluations for hydrogen production at existing light-water plants set the scales, the electrolyzer specific energies and the hydrogen prices that this application uses [5]. What those studies leave to the control engineer is the layer above the loops, which is where a day-ahead price curve, a state estimate and a contract quantity have to become setpoints.

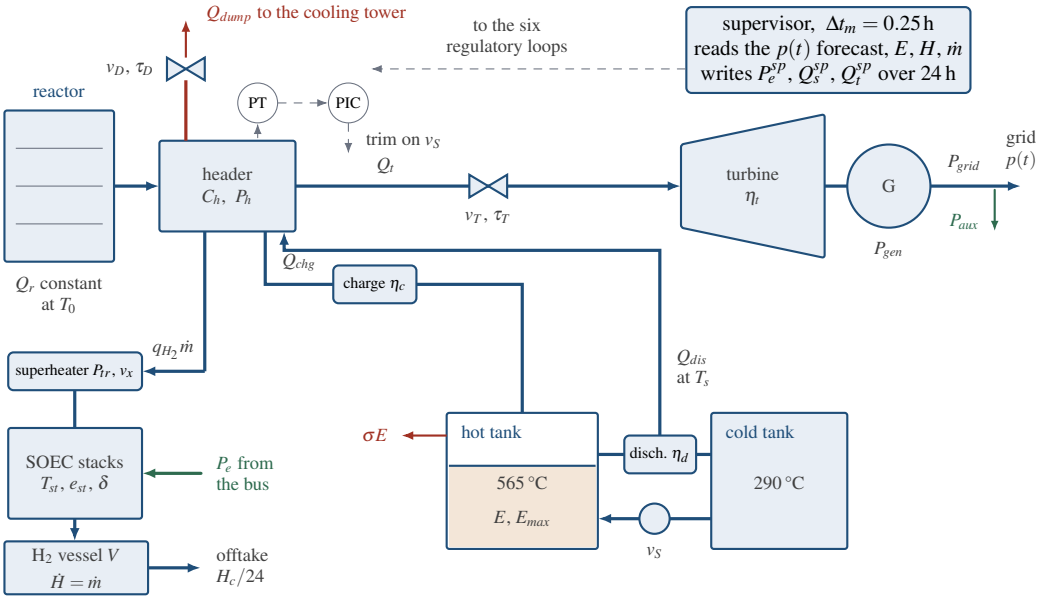

Figure 13.1 The integrated energy system, its four heat sinks and the two control layers. Constant reactor heat Q_r enters a steam header of inventory C_h holding pressure P_h, and leaves through the turbine throttle v_T, the charging exchanger, the electrolyzer steam draw $q_{H_2}\dot{m}$ and the dump valve v_D. Hot salt at E raises steam again through the discharge exchanger at T_s and returns it to the header, where the same throttle passes it to the turbine, which is why the store can lift generation above the reactor's own output. PIC trims the salt pump v_S to close the heat balance, and the supervisor writes the three allocation setpoints every 0.25 h from a price forecast and the measured state.

> ⧖ HISTORY **A Test Reactor Built for High Heat (1998 to 2010)**
>
> Light water reactors give steam at 200 to 300 °C, while many industrial processes want 500 to 900 °C [17]. Japan built a 30 MWth engineering test reactor of the helium-cooled, graphite-moderated type to reach those temperatures. It reached first criticality in 1998, ran at full power with an 850 °C outlet from 2001 and at 950 °C from 2004, and in a 50-day test in 2010 it held full power above 930 °C [14]. Its tests showed that it rode through a loss of coolant circulation without active control or emergency measures, and a thermochemical hydrogen process was to be connected to it. Today's designs of the type are passively safe modules of under 600 MWth each, added several to a site as capacity is needed [14], as the four modules of this chapter are.

13.3 Plant Model

The plant is four high-temperature reactor modules delivering a combined 860 MWth as 565 °C steam into one header. Steam leaves the header through four paths. The turbine throttle takes 840.8 MWth at the design point for 333.0 MWe gross at 39.60 % cycle efficiency. A charging heat exchanger moves heat into a hot salt tank at 565 °C over a cold tank at 290 °C, and a discharge exchanger raises 550 °C steam back out of it. The electrolyzer island draws 106.5 MWe and 19.2 MWth of header steam for 3.003 t/h of hydrogen, which is 72.07 t a day against a 72 t contract. The dump condenser takes the remainder. Station service is 12.0 MW, so 214.5 MWe reaches the grid. Figure 13.1 carries the layout and every symbol used in the equations, and Section 1.4.1 places the plant in the volume's flow scheme. Time is in hours, heat and power in megawatts, stored heat in megawatt-hours, hydrogen in tonnes, header pressure in bar and price in dollars per megawatt-hour.

Table 13.1 classifies what the two control layers interact with, and twelve of those entries are

Table 13.1 Classification of Variables. Nominal Values Are Those of the Design Allocation at Which Every Scenario Starts.

Symbol	Role	Nominal	Range
Q_t (FIC)	CV, measured	840.8 MWth	250–1180 MWth
P_h (PIC)	CV, measured	120.0 bar	120 bar SP
Q_s (FIC)	CV, measured	0 MWth	−340–320 MWth
P_e (PIC)	CV, measured	106.5 MWe	18–180 MWe
T_{st} (TIC)	CV, measured	800 °C	800 °C SP
v_T throttle	MV	71.3 %	0–100 %
v_S salt pump	MV	0 %	−100–100 %
v_D dump valve	MV	0 %	0–100 %
P_e^{cmd}	MV	106.5 MWe	0–180 MWe
v_x superheater	MV	34.9 %	0–100 %
p price	DV, measured	table	−25–$528/MWh
$\hat{p}$ forecast	DV, measured	table	spike fraction f
H_c contract	DV, measured	72 t/day	30–115 t/day
a_T turbine avail.	DV, measured	1.00	0–1.00
r reactor trim	DV, measured	1.00	0.85–1.03
c_E tank derate	DV, measured	1.00	0.20–1.00
δ_0 stack wear	DV, unmeasured	0.00	0–1.00
E	state	600 MWh	0–2400 MWh
P_h	state	120.0 bar	20–220 bar
v_T, v_S, v_x, v_D	states	see above	clipped
P_e	state	106.5 MWe	0–180 MWe
T_{st}	state	800 °C	300–950 °C
V	state	14.0 t	0–40 t
H, H_d	states	0 t	≥ 0
δ	state	0.00	0–1.00

integrated states. Operation is bounded by five latching annunciations: a dump condenser passing more than 25 MWth, a hot tank at its capacity, a hot tank empty, a header below the 104 bar low alarm, and a declaration that the hydrogen contract can no longer be reached even with the electrolyzer at nameplate for every hour that remains. None of the five alters the model, which leaves the transient that raised it readable to the end of the run.

13.3.1 The Heat Balance and the Header

The reactor cannot store heat and the header holds very little, so the allocation is an identity rather than a design target, as shown in Equation (13.1).

$$C_h \frac{\mathrm{d}P_h}{\mathrm{d}t} = r\,Q_r + Q_{dis} - Q_t - Q_{chg} - q_{H_2}\dot{m} - Q_{dump} \tag{13.1}$$

where C_h is the header and drum steam inventory in MWh/bar, P_h is header pressure in bar, Q_r is reactor thermal power, r is a reactor trim multiplier, Q_{dis} and Q_{chg} are the salt discharge and charge duties, Q_t is heat to the turbine, $q_{H_2}\dot{m}$ is the electrolyzer steam draw and Q_{dump} is the dump condenser duty, all in MWth. Both sides carry MWth/bar times bar/h. With $C_h = 3.0\,\mathrm{MWh/bar}$ a residual of 100 MWth moves the header at 33 bar/h, so nothing on the right of (13.1) can be left unaccounted for. The state is clipped to 20–220 bar after each step. Section 1.4.2 sets this header beside the other stored inertias of the volume, and Volume 18, Chapter 9 shares one header of the same kind among three small reactor modules.

The three duties in (13.1) come from three actuators, and each carries the pressure or tank end effects that a linear dispatch model cannot see, as shown in Equations (13.2) and (13.3).

$$Q_t = \frac{v_T}{100}\,Q_{t,\max}\,\phi\,a_T, \quad \phi = \mathrm{clamp}\big(P_h/P_{h,0}, 0.25, 1.20\big) \tag{13.2}$$

$$\begin{aligned} Q_{chg} &= \tfrac{1}{100}\max(0, v_S)\, P_{c,\max}\, f_{hi} \\ Q_{dis} &= \tfrac{1}{100}\max(0, -v_S)\, P_{d,\max}\, f_{lo} \end{aligned} \tag{13.3}$$

where v_T is throttle position in percent, $Q_{t,\max}$ is the swallowing capacity, a_T is turbine availability, v_S is the signed salt pump command, and $f_{hi} = \text{clamp}((c_E E_{max} - E)/40, 0, 1)$ and $f_{lo} = \text{clamp}(E/40, 0, 1)$ throttle the exchangers over the last 40 MWh of tank travel so that a full tank cannot be charged and an empty one cannot be discharged. The throttle gain in (13.2) falls with header pressure, which is why a starved header settles rather than running away.

13.3.2 The Two-Tank Store

The store is one energy state with fixed exchanger efficiencies and a standing loss, as shown in Equation (13.4), which is the inventory balance of Equation (1.4) written in stored heat rather than in state of charge.

$$\frac{\mathrm{d}E}{\mathrm{d}t} = \eta_c Q_{chg} - \frac{Q_{dis}}{\eta_d} - \sigma E, \qquad 0 \le E \le c_E E_{max} \tag{13.4}$$

where E is stored heat in MWh, η_c and η_d are the charging and discharge exchanger efficiencies, σ is the standing heat loss per hour and c_E derates the usable capacity. Charge and discharge are kept as separate non-negative flows so that the round-trip loss stays linear in the allocation, and no optimum ever asks for both at once because both cost money. The thermal round trip is $\eta_c \eta_d = 95.04\,\%$.

13.3.3 Turbine Cycle and the Electric Balance

The gross cycle carries two penalties, one for part load and one for salt steam that arrives at 550 °C rather than 565 °C with no reheat behind it, as shown in Equation (13.5).

$$\eta_t = \eta_0 \left[1 - c_p (1 - x)^2\right] \left[1 - c_d f_d\right], \qquad x = \frac{Q_t}{Q_{t,\max}} \tag{13.5}$$

where x is clipped to 0–1.05 and $f_d = Q_{dis}/(Q_r + Q_{dis})$ is the salt-raised fraction of the steam, taken against the untrimmed reactor rating. At the design allocation (13.5) returns 39.60 % against a rated 40.00 %. The electric balance is Equation (13.6).

$$\begin{aligned} P_{gen} &= \eta_t Q_t, \qquad P_{aux} = P_{a,0} + k_p \left(Q_{chg} + Q_{dis}\right) \\ P_{grid} &= P_{gen} - P_e - P_{aux} \end{aligned} \tag{13.6}$$

where $P_{a,0}$ is station service, k_p is salt pumping power per unit of salt duty and P_e is the whole electrolyzer island load. A negative P_{grid} means the plant is importing, which at a negative price is a revenue stream rather than a fault.

13.3.4 The Electrolyzer Island

The electrolyzer is represented by a specific-energy correlation rather than an electrochemical model, and the island load covers three consumers, as shown in Equations (13.7) to (13.9).

$$P_{bop} = b P_e, \qquad P_{tr} = \frac{v_x}{100} P_{tr,\max} \tag{13.7}$$

$$P_{st} = \max\left(0, P_e - P_{tr} - P_{bop}\right)$$

$$e_{st} = \max(24,\ e_0 + e_T\,(T_{sp} - T_{st}) + k_\delta \delta) \tag{13.8}$$

$$\dot{m} = P_{st}/e_{st} \tag{13.9}$$

where b is the balance-of-plant share, P_{tr} is the electric superheater duty, P_{st} is stack DC power, $\dot{m}$ is hydrogen production in t/h, e_{st} is stack DC energy per tonne, e_T is its temperature coefficient, T_{sp} is the 800 °C stack setpoint and δ is a scalar degradation state. At the design allocation (13.7) splits 106.5 MWe into 97.90 MWe of stack, 4.88 MWe of superheater and 3.73 MWe of balance of plant,

giving 3.003 t/h at 32.60 MWh/t of DC and 35.47 MWh/t of plant electricity. Holding the same load with the stacks at 750 °C raises that figure to 36.50 MWh/t, which is why the temperature loop is not decorative.

Stack temperature follows the enthalpy the feed steam and the superheater supply plus the stack's own irreversible heat, as shown in Equation (13.10).

$$\begin{aligned} T_f &= \frac{rQ_rT_0 + Q_{dis}T_s}{rQ_r + Q_{dis}}, \qquad \dot{m}_s = \max(1, s_{H_2}\dot{m}) \\ T^{tgt} &= \min\left(950,\ T_f + \frac{1636\,P_{tr}}{\dot{m}_s} + T_J\frac{P_{st}}{P_{e,\max}}\right) \\ \tau_{st}\frac{\mathrm{d}T_{st}}{\mathrm{d}t} &= T^{tgt} - T_{st} \end{aligned} \tag{13.10}$$

where T_f is the mixed feed steam temperature, $\dot{m}_s$ is steam flow in t/h at s_{H_2} tonnes of steam per tonne of hydrogen, T_J is stack self-heating at full load, and the constant 1636 °Ct/MWh is the inverse specific heat of steam at 2.2 kJ/(kg K). Turning the electrolyzer down reduces $\dot{m}_s$ faster than the superheater duty falls, so the stack runs hot unless the trim backs off.

13.3.5 Hydrogen Inventory, Degradation and Actuators

Production accumulates against a daily contract, and the offtake draws from a vessel that can run dry, as shown in Equation (13.11).

$$\frac{\mathrm{d}V}{\mathrm{d}t} = \dot{m} - \dot{m}_o, \quad \dot{m}_o = \frac{H_c}{24}\mathrm{clamp}\left(\frac{V}{0.5}, 0, 1\right), \quad \frac{\mathrm{d}H}{\mathrm{d}t} = \dot{m} \tag{13.11}$$

where V is vessel inventory in tonnes, H_c is the contract quantity in tonnes per day, $\dot{m}_o$ is the offtake rate and H is cumulative production. A third accumulator H_d integrates $\dot{m}_o$ and is reported but not used by any controller. Stack degradation accrues from three mechanisms, as shown in Equation (13.12).

$$\frac{\mathrm{d}\delta}{\mathrm{d}t} = g_0 + g_c\frac{\max(0, T_c - T_{st})}{100} + g_r\frac{\left|P_e^{cmd} - P_e\right|}{P_{e,\max}} \tag{13.12}$$

where T_c is the cold limit below which degradation accelerates and g_0, g_c and g_r are the baseline, cold and ramp coefficients. Thermal cycling and hard ramps dominate the total, which is the reason the dispatch layer carries an electrolyzer ramp limit it does not strictly need. The five actuators are first-order, as shown in Equation (13.13).

$$\tau_i\frac{\mathrm{d}u_i}{\mathrm{d}t} = u_i^{cmd} - u_i, \qquad u_i \in \{v_T,\ v_S,\ P_e,\ v_x,\ v_D\} \tag{13.13}$$

with time constants τ_T, τ_{pump}, τ_e, τ_{tr} and τ_D. Every state is clipped to its bound after each step, and the electrolyzer command additionally carries the 72 MW/h ramp limit on the controller output rather than on the actuator.

13.3.6 Price and the Money Accumulators

Price is an exogenous hourly table with the shape a market carrying a large solar share now has, cheap overnight, a morning ramp, a midday trough and an evening peak, as shown in Equation (13.14).

$$p(h) = [(1-g)\,p_b(h) + g\,p_{neg}] \times m(h) + p_{bias} \tag{13.14}$$

where p_b is the base curve, the solar glut g applies to hours 11 through 14 only and pulls them toward p_{neg} of −\$25/MWh, and $m(h) = s$ over the three hours starting at hour 17 and unity elsewhere. The base curve runs from \$6/MWh at hour 13 to \$96/MWh at hour 18, a ratio of sixteen to one before any

spike. The forecast the controller is given replaces s by $1+f(s-1)$, so $f=1$ is a perfect forecast and $f=0$ leaves the controller planning an ordinary evening. Money accrues as Equation (13.15).

$$\begin{aligned} \Pi_e &= \int p\,P_{grid}\,dt, \qquad \Pi_v = \int v\,P_{gen}\,dt \\ \Pi_H &= \int \left[c_c\,\dot{m}_c + c_m\,(\dot{m}-\dot{m}_c)\right] dt, \quad \dot{m}_c = \dot{m}\big|_{H\le H_c} \\ \Pi &= \Pi_e + \Pi_H - \Pi_v - c_s S\,\mathrm{clamp}(1-h_\ell/h_r, 0, 1) \end{aligned} \tag{13.15}$$

where v is variable operating cost on generation, c_c and c_m are the contract and merchant hydrogen prices, c_s is the take-or-pay penalty, $S=\max(0,H_c-H)$ is the shortfall, h_ℓ is the hours left before midnight and h_r is the window over which the accrual phases in. The accrual is what makes a missed contract visible as a step down in the margin trend during the last six hours rather than as a surprise at midnight.

13.3.7 Parameters, Initialization and Numerics

The fifty-six constants are listed in Table 13.2. Fifty-two are adjustable from the model dialog and four are hidden, and the four hidden ones are the standing tank loss and the three degradation coefficients, none of which a learner can read off a trend inside one day.

Every scenario opens at an allocation that closes (13.1) exactly. The three allocation controls set the throttle, the electrolyzer load and the salt duty, the electrolyzer island is evaluated through (13.7) and (13.8) at the stack setpoint, and whatever residual remains is placed on the dump valve so that the header does not have to swing on the first step, as shown in Equation (13.16).

$$\begin{aligned} \varepsilon_0 &= rQ_r + Q_{dis} - Q_t - Q_{chg} - q_{H_2}\dot{m} \\ v_D(0) &= \mathrm{clamp}\big(100\,\varepsilon_0/Q_{d,\max},\ 0,\ 100\big) \end{aligned} \tag{13.16}$$

The tank opens at $E(0)=\mathrm{soc}_0 c_E E_{max}$, the header at $P_{h,0}$, the stacks at T_{sp} and the vessel at V_0. A run that opens at a clock hour $h_0>0$ opens on a day already partly delivered, so cumulative production is initialized to $H(0)=(H_c/24)h_0 k_s$ where k_s is the fraction of schedule already met, and $k_s<1$ starts the day behind. Verification of the design point is one of fifteen self-checks the application runs, and it confirms 840.8 MWth, 214.5 MWe of export, 3.003 t/h and zero dump duty from (13.5) to (13.9).

The twelve states advance under fourth-order Runge-Kutta at a fixed 0.002 h step, which is 7.2 s of plant time. Six regulatory loops run on a 0.01 h (36 s) period, the supervisory quadratic program runs on a 0.25 h period, and the historian retains twenty logged samples per simulated hour. Not one disturbance in this application is stochastic, so the seed 20261118 never comes into play and each dollar figure quoted below is exactly repeatable. The speed selector sets simulated hours per second of wall clock and touches nothing else.

13.3.8 Assumptions

Detail is spent on whatever a controller has to be designed around and withheld elsewhere. Each of the seven simplifications below has a place where it stops holding.

The turbine is one efficiency curve. Equation (13.5) has no shaft, no inertia, no frequency, no governor droop and no minimum up or down time, so the turbine outage in scenario 10 is a heat-balance event rather than the mechanical and grid event it would really be. The 250 MWth minimum stable load is honored by the supervisor and by the operating map but is not enforced on the plant, so a manual allocation may carry the turbine below it.

The store is a single energy state. Equation (13.4) has no salt temperature, no tank level, no thermocline and no freeze protection, and the discharge steam temperature T_s is a constant rather than a function of the hot tank's condition. Two-tank molten salt systems in service depart from that on both counts [3]. The assumption understates how the last of a tank behaves.

Table 13.2 Model Parameters. Hidden Constants, Marked †, Are Never Displayed. Every One of the Fifty-Six Is Read by the Model, the Controller or Both.

Sym.	Meaning	Value	Unit	Sym.	Meaning	Value	Unit
Q_r	reactor thermal power	860	MWth	q_{H_2}	header steam per tonne H_2	6.4	MWh/t
h_0	clock hour at start	0	h	b	balance-of-plant share	0.035	—
T_0	header steam temperature	565	°C	$P_{tr,\max}$	superheater rating	14	MWe
$Q_{t,\max}$	turbine swallowing capacity	1180	MWth	s_{H_2}	steam fed per tonne H_2	12.0	t/t
$Q_{t,\min}$	turbine minimum stable load	250	MWth	T_J	stack self-heating at full load	25	°C
η_0	cycle efficiency at rated flow	0.400	—	T_{sp}	stack temperature setpoint	800	°C
c_p	part-load penalty	0.12	—	T_c	cold degradation limit	720	°C
c_d	salt-steam penalty	0.055	—	τ_e	rectifier response	0.008	h
τ_T	throttle travel time	0.020	h	τ_{tr}	superheater response	0.005	h
E_{max}	hot tank capacity	2400	MWh	τ_{st}	stack thermal time constant	0.15	h
$P_{c,\max}$	charging rate limit	320	MWth	$g_0{}^\dagger$	baseline degradation rate	8×10^{-6}	1/h
$P_{d,\max}$	discharge rate limit	340	MWth	$g_c{}^\dagger$	degradation, 100 °C cold	1.2×10^{-4}	1/h
η_c	charging exchanger efficiency	0.99	—	$g_r{}^\dagger$	degradation at a full-range ramp	9×10^{-5}	1/h
η_d	discharge exchanger efficiency	0.96	—	$P_{h,0}$	header pressure setpoint	120	bar
$\sigma^\dagger$	tank standing loss	1.458×10^{-4}	1/h	C_h	header steam inventory	3.0	MWh/bar
T_s	steam temperature from salt	550	°C	$P_{h,D}$	dump condenser opens at	126	bar
τ_{pump}	salt pump response	0.015	h	$P_{h,L}$	low header pressure alarm	104	bar
k_p	pumping power per salt duty	0.010	MWe/MWth	$Q_{d,\max}$	dump condenser capacity	1200	MWth
$P_{e,\max}$	electrolyzer nameplate	180	MWe	τ_D	dump valve travel time	0.010	h
$P_{e,\min}$	electrolyzer turndown	18	MWe	$P_{a,0}$	station service load	12	MWe
R_e	electrolyzer ramp limit	72	MW/h	$V_{\max}$	H_2 vessel capacity	40	t
e_0	stack DC energy at 800 °C	32.6	MWh/t	V_0	vessel inventory at midnight	14	t
e_T	stack energy per °C below T_{sp}	0.020	MWh/t/°C	c_c	contract hydrogen price	4200	$/t
k_δ	extra stack energy, full wear	3.0	MWh/t	c_m	merchant hydrogen price	1550	$/t
c_s	take-or-pay penalty	6000	$/t	h_r	shortfall phase-in window	6	h
v	variable O&M on generation	1.2	$/MWh	h_s	first hour of the evening peak	17	h
n_s	length of the evening peak	3	h	p_{neg}	midday floor at full solar glut	−25	$/MWh
τ_{Pm}	header pressure transmitter	0.004	h	τ_{Qm}	steam flow transmitter	0.006	h

High-temperature electrolysis is a correlation. Equation (13.8) has no polarization curve, no sweep gas, no pressurization and no stack-to-stack variation, and degradation is one scalar. The 35.47 MWh/t of plant electricity at 800 °C is a defensible figure for a nuclear-coupled solid-oxide unit [4, 5] and is not a measurement.

The dispatch model is deliberately a linearization of the plant beneath it, with a constant cycle efficiency, a constant electrolyzer specific energy, a linear steam draw per unit of load and no header dynamics at all. The mismatch is measured every execution and enters the linear cost vector as a bias, which is what keeps the layer offset free and the Hessian constant.

Price is an exogenous table and the plant is a price taker. There is no bid curve, no ancillary-service revenue, no imbalance settlement and no uncertainty band, and the forecast error is a fixed fraction of the spike chosen with a slider rather than a distribution.

The hydrogen contract is a daily quantity with a phased penalty. Real offtake agreements carry hourly delivery pressure obligations, monthly true-ups and force-majeure terms, none of which are here.

The numerics are more repeatable than a control system. A frozen step, a frozen scan and a frozen dispatch period give the identical trajectory everywhere, which is convenient and false, because scan order and execution time move on a real distributed control system and optimization server.

13.4 What an Hour Is Worth

Before any controller is chosen, three prices decide what a good allocation looks like, and all three are arithmetic on the parameters of Section 13.3.

13.4.1 The Electrolyzer Switching Price

Merchant hydrogen beyond the contract is worth $c_m = \$1{,}550/t$ against 35.46 MWh/t of plant electricity, which is \$43.71/MWh of electrolyzer load. That figure alone is not the switching price, because the electrolyzer also diverts header steam that the turbine would have converted, as shown in Equation (13.17).

$$p^* = \frac{c_m / e_{tot}}{1 + \kappa \eta_t}, \qquad \kappa = \frac{q_{H_2}}{e_{tot}} = 0.1805 \tag{13.17}$$

where $e_{tot} = 35.46\,\mathrm{MWh/t}$ is the linear specific energy the dispatch layer carries and κ is the thermal megawatts diverted per electrical megawatt of load. At the design cycle efficiency of 39.60 %, Equation (13.17) returns \$40.80/MWh, and at the rated 40.00 % it returns \$40.77/MWh. The price-threshold architecture is offered a default of \$44/MWh, which is deliberately the wrong number by \$3.20/MWh and is one of the three thresholds scenario 3 asks the learner to tune.

13.4.2 Storage Arbitrage in Thermal Terms

A megawatt-hour charged into the tank is a megawatt-hour the turbine did not convert, so the arbitrage has to be evaluated at the cycle efficiency on each side of the trade rather than at a single number. Charging at 320 MWth drives the turbine to 520.8 MWth where Equation (13.5) gives 38.50 %, and discharging at 340 MWth drives it to 1180 MWth where the same equation gives 39.38 % once the salt-steam penalty is applied. Charging one thermal megawatt-hour during the \$6/MWh hour therefore forgoes 0.385 MWh of generation and costs \$2.31, and returning $\eta_c \eta_d = 0.9504$ of it into the \$96/MWh hour yields 0.374 MWh worth \$35.93, a net \$33.62 per thermal megawatt-hour. Between \$24/MWh and \$35/MWh the same trade is worth \$3.86 and does not pay for the cycle. Salt pumping is left out of both figures.

The two efficiencies in that calculation are not equal, and the inequality runs the useful way. Evaluated at those endpoints the electrical round trip is $0.9504 \times 0.3938/0.3850 = 97.2\,\%$, which is higher than the 95.04 % thermal round trip, because the part-load penalty makes the displaced heat worth less per megawatt-hour than the heat the discharge adds. Measuring instead with the marginal efficiency over the full 320 MWth charge and 340 MWth discharge steps from the design allocation

gives 0.414 MWh forgone per thermal megawatt-hour charged and 0.387 MWh returned per thermal megawatt-hour discharged, an electrical round trip of 88.9 %. Charging the 0.010 MWe/MWth of salt pumping to both legs moves those to 0.424 MWh and 0.377 MWh, a round trip of 84.6 % and an arbitrage of $31.87 per thermal megawatt-hour. The dispatch layer of Section 13.5 uses one filtered efficiency for both directions, so within its own model the electrical round trip is the thermal one less pumping.

13.4.3 The Water Value of Stored Heat

Charge and discharge thresholds are not chosen in the optimizing architecture. They follow from a merit order over the forecast horizon, as shown in Equation (13.18).

$$\begin{aligned} E^{*} &= \min\left(c_E E_{max},\; E + \eta_c P_{c,\max} \Delta_b N_b\right) \\ n &= \operatorname{clamp}\left(\left\lceil \frac{\eta_d E^{*}}{P_{d,\max}\Delta_b} \right\rceil, 1, N_b\right), \qquad p_{ref} = \hat{p}_{[n]} \end{aligned} \tag{13.18}$$

where $\hat{p}_{[1]} \geq \hat{p}_{[2]} \geq \cdots$ are the forecast block prices sorted downward, Δ_b is the block length, N_b is the number of blocks, E^{*} is the fullest the tank could be made over the horizon, and n is how many blocks of full discharge that inventory would supply. The price at which the tank runs out is the marginal value of stored heat, and dividing the exchanger efficiencies into it produces the thresholds directly, charge below $p_{ref}\eta_c$ and discharge above p_{ref}/η_d. This is the same reservoir water value that season-long hydro scheduling is built on [11], evaluated over twenty-four hours instead of a year. It is the merit order of Equation (1.9) with two differences. The tank is first made as full as the cheap blocks of the horizon allow, and p_{ref} is the block price itself rather than that price multiplied by η_d, so both thresholds sit a factor $1/\eta_d$ above those of Chapter 1.

13.5 Two Control Layers

13.5.1 The Regulatory Layer

Six loops run on a 0.01 h scan and never see a price. Each is the velocity-form PID of Section 1.6.2 with back-calculation anti-windup, written with this chapter's symbols in Equation (13.19).

$$\begin{aligned} e_k &= a\left(\mathrm{SP}_k - \mathrm{PV}_k\right), \qquad \Delta u_k^{P} = K_c\left(e_k - e_{k-1}\right) \\ \Delta u_k^{I} &= \frac{K_c}{\tau_I} e_k \Delta t_c, \qquad s_k = s_{k-1} + \Delta u_k^{P} + \Delta u_k^{D} \\ c_k &= u_b + s_k + I_{k-1} \\ u_k &= \operatorname{slew}\left(\operatorname{clamp}\left(c_k, u_{\min}, u_{\max}\right),\; u_{k-1},\; R\,\Delta t_c\right) \\ I_k &= I_{k-1} + \Delta u_k^{I} + \frac{\Delta t_c}{T_t}\left(u_k - c_k\right) \end{aligned} \tag{13.19}$$

where $a = \pm 1$ carries the loop action, s_k is the unlimited running sum of Equation (1.11), c_k is the output the algorithm demands, u_k is that demand clipped to the output span and then rate limited at R, and T_t is the tracking time of the back-calculation in Equation (1.13). Only the electrolyzer loop has a finite R, its 72 MW/h ramp limit. No loop in this application uses derivative action, so $\tau_D = 0$ and $\Delta u_k^{D} = 0$ in all six, because every measured variable here is either an actuator-following signal or a pressure on a thirteen-second inventory and neither rewards anticipation. The six loops never leave automatic. Manual operation changes only where the three allocation setpoints come from. The one use of the tracking form is the stack trip, during which the rectifier loop's integral is reset every scan so that its output reads zero, and the load then returns from zero at the ramp limit when the trip clears.

Table 13.3 lists the six as they are configured. Two of them are worth separating from the other four. The header pressure controller is *direct* acting because charging takes steam out of the header, so more charge means less pressure, and its output straddles zero, which is why it can be left in

Table 13.3 The Six Regulatory Loops as Implemented. Action +1 Is Reverse Acting and −1 Is Direct Acting in the ISA Sense.

CV	MV	K_c	τ_I (h)	Act.	Output span
Q_t	throttle v_T	0.30	0.040	+1	0–100 %
P_h	salt trim	45	0.10	−1	−340–320 MWth
Q_s	salt pump v_S	0.30	0.030	+1	−100–100 %
P_h	dump valve v_D	12	0.060	−1	0–100 %
P_e	rectifier P_e^{cmd}	1.6	0.030	+1	0–180 MWe
T_{st}	superheater v_x	1.2	0.15	+1	0–100 %

service all day without parking against a limit. It does not drive an actuator directly. Its output is a signed trim added to whatever the supervisor asked the salt loop for, as shown in Equation (13.20).

$$Q_s^{eff} = \mathrm{clamp}\big(Q_s^{sp} + t_h,\ -P_{d,\max},\ P_{c,\max}\big) \tag{13.20}$$

where t_h is the pressure controller output in MWth. That arrangement is feedforward plus feedback with the store as the buffer. When the supervisor's plan is right the trim sits near zero, and when it is wrong the store silently carries the difference, which is what a thermal buffer exists to do. The dump condenser controller is the second. Its setpoint is 126 bar, six bar *above* the header setpoint, so it is an override rather than a regulator, it sits clamped at zero with a large negative error for most of every day, and the back-calculation term in (13.19) with $T_t = 0.02$ h is what stops it winding down while it waits. It holds the demand c_k near $(K_c/\tau_I)\,e\,T_t$, which is −24 % at a −6 bar error, rather than letting the integral run away. The tracking constants are 0.02 h on the throttle, dump and salt loops, 0.03 h on the header loop, 0.01 h on the electrolyzer and 0.05 h on the stack temperature loop.

Pause. Try Exercise 13.1 before reading on.

13.5.2 Three Supervisors

The regulatory layer is identical under all four architectures, and only the source of the three allocation setpoints changes. Manual operation passes the operator's three allocations. The fixed split passes one allocation chosen once and held for twenty-four hours, 840.8 MWth and 106.5 MWe by default, which meets the contract exactly and never dumps. The price-threshold architecture is three comparisons and a catch-up term, as shown in Equation (13.21).

$$\begin{aligned}
P_e^{thr} &= \begin{cases} P_{e,\max}, & p < p_{ely} \\ P_{e,\min}, & \text{otherwise} \end{cases} \\
P_e^{sp} &= \max\left(P_e^{thr},\ \mathrm{clamp}\left(\frac{H_c - H}{h_\ell} e_{tot}, P_{e,\min}, P_{e,\max}\right)\right) \\
Q_s^{sp} &= \begin{cases} +P_{c,\max}, & p < p_{chg},\ E < 0.995\,c_E E_{max} \\ -P_{d,\max}, & p > p_{dis},\ E > 0.005\,c_E E_{max} \\ 0, & \text{otherwise} \end{cases} \\
Q_t^{sp} &= \mathrm{clamp}(Q_r - Q_s^{sp} - \kappa P_e^{sp},\ Q_{t,\min}, Q_{t,\max})
\end{aligned} \tag{13.21}$$

with defaults $p_{chg} = \$22/MWh$, $p_{dis} = \$62/MWh$ and $p_{ely} = \$44/MWh$. The salt setpoint is then recomputed from the clipped turbine allocation so that the store absorbs whatever the turbine could not take. Equation (13.21) reads only the price on the meter at this instant and the contract position, and it has no view of any future hour at all. Its Q_r and κ are the declared values of Table 13.2.

13.5.3 The Dispatch Quadratic Program

The optimizing architecture is this application's own program rather than the framework option of Section 1.6.3, and its objective is money rather than a tracking error. It divides twenty-four hours into

$N_b = 12$ blocks of $\Delta_b = 2\,\mathrm{h}$ and carries four manipulated variables per block, the electrolyzer load, the charge duty, the discharge duty and the dump duty, for $N_u = 48$ decision variables. Turbine heat is not among them. It is eliminated through the heat balance, as shown in Equation (13.22).

$$\begin{aligned} \tilde{Q}_{t,k} &= rQ_r + Q_{dis,k} - Q_{chg,k} - \kappa P_{e,k} - b_Q - Q_{dump,k} \\ Q_{t,k} &= \mathrm{clamp}\big(\tilde{Q}_{t,k},\ Q_{t,\min},\ Q_{t,\max}\big) \end{aligned} \tag{13.22}$$

where Q_r and κ are the declared values of Table 13.2, r is the trim the reactor was dispatched to, and $b_Q = q_{H_2}\dot{m} - \kappa P_e$ is the measured error in the linear steam draw, with $q_{H_2}\dot{m}$ the metered steam to the stacks. Eliminating the turbine rather than optimizing it removes the only equality constraint in the problem, which leaves a genuine box-constrained program, and it makes every economic gradient exact instead of weighted. Charging costs the turbine heat it displaces at $p\eta_t$, discharging earns the same, dumping costs $p\eta_t$ and therefore *earns* that at a negative price, and the electrolyzer costs $p(1+\kappa\eta_t)$ while earning c_m/e_{tot}. No weight has to be chosen for any of the four.

Each decision is scaled to the unit interval, $u_{k,m} = o_m + s_m z_{k,m}$ with $z \in [0,1]^{48}$, where the offsets and spans are $(P_{e,\min}, P_{e,\max} - P_{e,\min})$, $(0, P_{c,\max})$, $(0, P_{d,\max})$ and $(0, 950\,\mathrm{MWth})$. The objective is $J(z) = \frac{1}{2} z^{\mathsf{T}} H z + g^{\mathsf{T}} z$ with the Hessian of Equation (13.23).

$$H = 2\left(w_H\, hh^{\mathsf{T}} + w_E\, ss^{\mathsf{T}} + \Lambda D^{\mathsf{T}} D + \rho I\right) \tag{13.23}$$

where h has the constant $\Delta_b s_{P_e}/(e_{tot} H_n)$ in each of the twelve electrolyzer positions and zero elsewhere, because cumulative hydrogen is a pure integrator of the electrolyzer decisions and H_n normalizes it, s carries $\Delta_b \eta_c s_{chg}/E_{max}$ in the charge positions and $-\Delta_b s_{dis}/(\eta_d E_{max})$ in the discharge positions, D is a first difference between consecutive blocks, Λ weights move suppression at λ on the electrolyzer and 0.4λ on the other three, and $\rho = 0.02$ is a ridge. Nothing in (13.23) depends on the price, the state or the contract, so H is assembled once at reset and rebuilt only when a weight or a plant limit moves. Default weights are $w_H = 8$, $w_E = 1.2$ and $\lambda = 0.06$.

Everything that changes every execution enters the linear term, as shown in Equation (13.24).

$$\begin{aligned} g_{k,P_e} &= 2w_H h_k c_H - 2\lambda\,[k{=}0]\, z_{P_e}^{now} \\ &\quad + \mu s_{P_e}\left[\hat{p}_k\,(1+\kappa\eta) - c_m/e_{tot} - v\eta\,\kappa\right] \\ g_{k,chg} &= 2w_E s_k c'_E + \mu s_{chg}\left[\hat{p}_k\,(\eta + k_p) - v\eta - v_E\eta_c\right] \\ g_{k,dis} &= 2w_E s_k c'_E + \mu s_{dis}\left[-\hat{p}_k\,(\eta - k_p) + v\eta + v_E/\eta_d\right] \\ g_{k,dmp} &= \mu s_{dmp}\left[\hat{p}_k\eta - v\eta\right] \end{aligned} \tag{13.24}$$

where $\mu = \Delta_b/M$ with $M = \$40{,}000$ scales the economic term against the dimensionless penalties, η is the filtered measured cycle efficiency, $v_E = w_V\,\eta\,p_{ref}$ is the water value of Equation (13.18) in dollars per thermal megawatt-hour, c_H is the offset of the hydrogen terminal state from its target and c'_E the offset of the terminal tank from its reference. The move-suppression terms carry the allocation actually in force, and the other three decisions carry $-0.8\lambda\,[k{=}0]\,z^{now}$ in the same way. The hydrogen target itself absorbs the urgency of the approaching deadline, as shown in Equation (13.25).

$$\begin{aligned} \gamma &= 1 + 2\,\mathrm{clamp}(1 - h_\ell/24,\ 0,\ 1) \\ H^{tgt} &= \frac{H_c \Delta_b N_b}{24} + \gamma\left(\frac{H_c h}{24} - H\right) \end{aligned} \tag{13.25}$$

Putting the urgency in the target rather than in w_H is what allows the Hessian to stay constant as midnight approaches. The measured model error in the hydrogen rate, $b_H = \dot{m} - P_e/e_{tot}$, enters the same target, which is the bias correction that makes the layer offset free.

Each execution runs a fixed $N_{it} = 18$ projected-gradient iterations, as shown in Equation (13.26).

$$z \leftarrow \Pi\big(z - \alpha\left[Hz + g - 2w_H h\,\max\big(0,\ h^{\mathsf{T}} z + c_H\big)\right]\big) \tag{13.26}$$

where Π is the feasibility restoration described below and the third term inside the bracket makes the contract penalty one sided. Falling short costs \$6,000/t while going over is merchant hydrogen already priced in the economic term, so when the plan is above target the terminal gradient is dropped by twelve subtractions while the Hessian keeps it, because the step size still has to be safe for the case where it binds.

The step size is not the textbook $1/\max_i H_{ii}$. Both terminal states in Equation (13.23) are sums over every block and enter the Hessian as dense rank-one blocks whose largest eigenvalue is N_b times their diagonal entry, so the row-sum bound of Equation (13.27) is used instead.

$$\alpha = 1 \Big/ \max_i \sum_j \left| H_{ij} \right| \tag{13.27}$$

Equation (13.27) is a Gershgorin bound on the spectrum, it is one extra loop, and it is computed once with the Hessian. At the default weights it returns $\alpha = 0.2066$ against $1/\max_i H_{ii} = 1.860$, so the diagonal step is 9.00 times longer than the safe one, and Section 13.6 measures what taking it costs.

Feasibility restoration Π runs in a fixed order after every gradient step. The four boxes are applied first. The electrolyzer ramp limit is then applied forward from the load in force, at $R_e \Delta_b$ per block. Simultaneous charge and discharge is netted, because it is always strictly worse than doing neither yet a gradient step can ask for it. The state of charge is restored in two passes, a merit-order trim that removes the excess from the blocks that valued it least, the dearest charging and the cheapest discharging, followed by a causal forward pass because energy cannot be borrowed from a later block. The turbine limits of (13.22) are restored last, by moving the dump valve and then the charge duty. That is greedy feasibility restoration rather than the exact Euclidean projection onto the chain, which is what a production dispatch code does, it is $O(N_b)$, and doing the merit trim before the causal pass is what stops the tank being spent on the first profitable block instead of the best one.

Warm starting slides the trajectory forward by the fraction of a block that has elapsed, $z_k \leftarrow (1-\beta) z_k + \beta z_{k+1}$ with $\beta = \Delta t_m / \Delta_b = 0.125$, so eight executions make one whole block shift and a block is refined ninety-six times before it becomes the present. Whenever another architecture is driving the plant, the first block of the trajectory is overwritten with the allocation actually in force before the solve, which is the same substitution dynamic matrix control makes when it charges its prediction with the move that was delivered rather than the move that was requested, and it is what makes selecting the optimizer in mid-run bumpless. One execution is a shift, a linear-term build and eighteen matrix-vector products of size 48×48. Timed over 400 executions after a 200-execution warm-up, which the application does at reset when its solver benchmark is switched on, the cost is 23.4 µs with a spread of 23.2 to 26.0 µs over six independent measurements on an Apple M-series laptop. The timing runs in the client, so the microsecond figure the application reports belongs to whatever machine it was opened on. A fixed iteration count with a precomputed step is the same design choice that makes first-order predictive control certifiable in real time [18], and the economic objective and terminal penalty follow the standard economic model predictive control formulation [19]. Every coefficient of the dispatch model, Q_r, κ, η_c, η_d, k_p and the rated η_0 that replaces the measured efficiency while the turbine is off, is the declared value of Table 13.2, frozen when the application loads. A change to one of them in the model dialog is therefore plant and model mismatch, absorbed by the biases b_Q and b_H, the measured tank level and the header pressure trim, while the plant limits, the contract, the prices, the forecast and the three measured disturbances a_T, r and c_E of Table 13.1 are read as the plant has them.

13.6 Results

Table 13.4 lists the eleven scenarios in the order a laboratory would work through them, the mechanism each exercises and what completes it. Each one pins the starting clock hour, the supervisory architecture, the three price thresholds, the four optimizer weights, the disturbance settings and any timed events. A case is credited when the listed actions are taken and the simulation reaches the listed

Table 13.4 Scenarios and Completion Requirements

#	Scenario (kind)	What it exercises	Required to complete
1	Allocate by hand at hour 11 (exploration)	the heat balance as an identity	In Manual, open the dump condenser past 25 MWth (a turbine near its 250 MWth minimum does it), raise the electrolyzer load to at least 150 MWe, and find an allocation that takes the header below its 104 bar low alarm.
2	Fixed split, one whole day (experiment)	a good plant and a poor merchant	Run the 24 h day on the fixed split at 106.5 MWe, then a second day with the fixed electrolyzer load set at least 5 MWe away before the run, and compare the hydrogen delivered with the 72 t contract.
3	Price thresholds, same day (experiment)	a heuristic with no horizon	Run the day on the shipped thresholds ($22, $62 and $44/MWh), then a second day with at least one threshold changed before the run and held for all 24 h.
4	Economic MPC, same day (experiment)	what a horizon is worth	Run the optimizer to 06:00 and read the plan ahead of the present, then run on to midnight and compare what was done with what was planned.
5	Afternoon spike, $f = 0.35$ (experiment)	three constraints binding at once	Run the optimizer from 08:00 to midnight with the 4.5× peak and an 82 t contract, then the same window on the price rules, the architecture changed before the run.
6	Negative midday prices (experiment)	every sign in the objective flips	On the optimizer, import more than 20 MW with the dump valve more than 5 % open for 0.5 h and run from 09:00 to midnight, then repeat on the price rules chosen before the run.
7	Spike at 5.5×, $f = 0$ (demonstration)	acting on a wrong forecast	Run the whole day on the optimizer with the 5.5× evening peak absent from the forecast, then the same day on the price rules or the fixed split chosen before the run.
8	Stack trip, contract behind (experiment)	a terminal penalty, not a constraint	On the optimizer, see the electrolyzer drop below 1 MWe after the 13:00 trip and pass 50 MWe again after the stacks return at 14:30, then run to midnight (12 h).
9	Derated tank, thresholds (demonstration)	a capacity limit felt hours later	On the price rules with an 840 MWh tank, run until the hot tank reports full and on to midnight (23 h), then repeat the day on the optimizer chosen before the run.
10	Turbine outage for 6 h (demonstration)	the dump condenser earns its name	On the optimizer, see the dump valve open past 10 % after the 07:00 trip, an hour of import above 100 MW with more than 300 MWth dumped, and after 13:00 the turbine above 500 MWth with the store discharging more than 50 MWth.
11	Graded challenge (challenge)	the whole day under a target	Over 24 h deliver at least 90 t, hold the header within ±6 bar of 120 bar whenever the dump valve is shut, and finish above $975,000, with the model parameters and disturbances as given.[a]

[a] As loaded (a fixed split at 780 MWth and 140 MWe, optimizer weights deliberately wrong) the day books $641,543 with 94.6 t delivered and the header within 1.07 bar while the dump valve is shut, so it fails on margin alone, by $333,457. The suggested configuration (economic MPC with $\lambda = 0.06$, $w_V = 1$, $w_H = 8$, $w_E = 1.2$ and a terminal state of charge of 0.25) books $1,003,720 with 94.3 t and 5.66 bar on the logged record, a pass by $28,720 and 0.34 bar, and the same whether it is revealed before Run or carried through a Reset.

state. The challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt. Section 1.6.4 describes the briefings, the credit and the certificates common to every application. Reset restarts a run from its starting clock hour and keeps the learner's own settings, while reloading a scenario restores its shipped ones, so a comparison of two architectures on one day needs only the architecture changed between runs. A run exports as a CSV record with 108 columns at twenty samples per simulated hour, holding the realized allocation beside the plan, the settled price beside the forecast, and the money accumulators beside the heat balance residual.

13.6.1 The Heat Balance Before Any Controller

Scenario 1 makes the point with arithmetic. At the design allocation the turbine takes 840.8 MWth, the electrolyzer steam 19.2 MWth and the store nothing, which sums to 860.0 MWth. Asking the store to charge at its 320 MWth limit forces the turbine to 520.8 MWth, and turning the turbine down

Table 13.5 Measured Outcomes at Midnight. The Ordinary Day Is Scenarios 2, 3 and 4 and the Spike Day Is Scenario 7 With the Architecture Changed. Every Figure Is Deterministic.

	Fixed	Thresholds	Economic MPC
Ordinary day			
Net margin ($)	470 311	516 300	529 303
Electricity ($)	177 392	178 678	203 243
Hydrogen ($)	302 508	346 641	335 312
H_2 delivered (t)	72.07	100.55	93.24
Heat dumped (MWh)	0.0	14.7	0.0
Tank charged (MWh)	0	1896	2976
Tank at midnight (%)	24.9	59.3	35.1
Spike day, peak 5.5×*, forecast blind*			
Net margin ($)	709 818	924 749	1 010 196

to its 250 MWth minimum without moving anything else places 590.8 MWth on the dump condenser. No combination of the three allocations avoids the identity, and the pressure controller closes it whether or not the operator intended the closure.

Over an economically optimized day the residual in (13.1) never exceeds 100.8 MWth, is larger than 40 MWth for 0.310 h out of twenty-four, holds the header within 2.02 bar of setpoint, and integrates to 25.3 ppm of the heat supplied over the day. A residual during a setpoint move is real and is what the buffer exists for, because 100 MWth held for one 36 s regulatory scan is 0.33 bar. What has to be true is that it is bounded, brief and integrates to nothing.

13.6.2 Three Architectures on One Day

Table 13.5 reports the same twenty-four hour price curve dispatched three ways, and then the same comparison with the evening peak at five and a half times the day-ahead curve and a forecast that saw none of it.

Three thresholds book $45,989 more than the fixed split and the optimization books a further $13,003, so the heuristic captures 78 % of the available improvement and the optimization is worth 2.5 % on top of it. That is the honest result on an ordinary day, and it is worth stating plainly, because an ordinary day's price curve is smooth and none of its constraints bind, and smooth unconstrained problems are what heuristics are good at.

🕮 **Pause.** Try Exercise 13.2 before reading on.

The comparison also has to account for what is left in the tank, because a run that ends with the store fuller than it started has put some of its value into inventory rather than into cash. The threshold architecture finishes the ordinary day at 59.3 % state of charge against the optimizer's 35.1 %, which is 579 MWh of stored heat. Discharged at η_d and the 39.4 % high-load efficiency into the overnight $26/MWh hour, that inventory is worth $5,693, which would cut the optimizer's advantage on this day from $13,003 to about $7,300 and raise the heuristic's share of the improvement from 78 % to 88 %. Carried into the following evening's $96/MWh peak instead, the same inventory is worth $21,000 and the ranking reverses. The ordering of the three architectures is robust and the size of the gap is not, which is why the application reports the state of charge at midnight beside the margin.

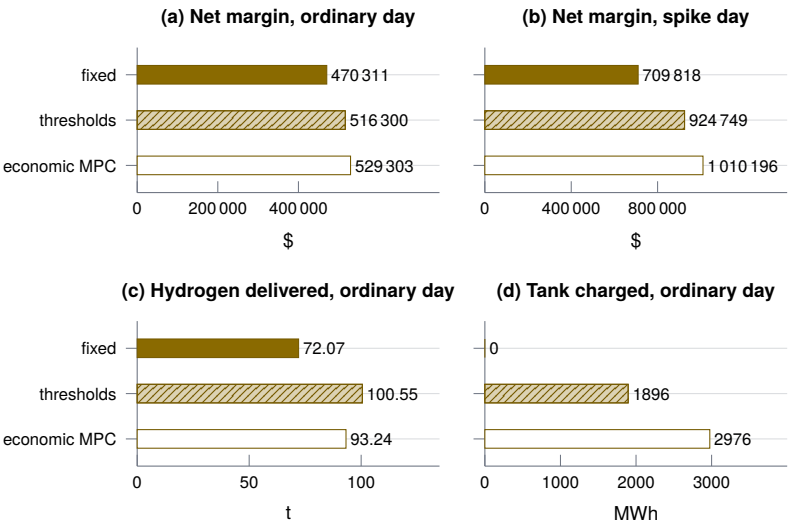

Figure 13.2 The three supervisory architectures of Table 13.5: (a) net margin on the ordinary day, (b) net margin on the spike day with a blind forecast, (c) hydrogen delivered and (d) heat charged into the tank on the ordinary day. Solid: fixed split. Hatched: price thresholds. White: economic MPC. Notice how close the thresholds come to the optimizer in (a) and how far apart they stand in (b), and that the thresholds make the most hydrogen while the optimizer moves the most heat through the store.

> ✦ **DID YOU KNOW?** **More Hydrogen, Less Money**
>
> From the simulator: on the ordinary day the threshold rule delivers 100.55 t of hydrogen and books $346,641 from it, more than the optimizer's 93.24 t and $335,312. The optimizer still books more in all, because its electricity earns $203,243 against $178,678 (Table 13.5).

On the spike day the arithmetic changes character. The optimization books $85,447 more than the thresholds, six and a half times its ordinary-day advantage, against a fixed split that books $300,378 less than the optimizer. Neither controller saw the spike in its forecast, and neither changes its dispatch when the meter shows it. Both allocations are those of the ordinary day to the megawatt, because the rules' discharge and electrolyzer thresholds are already crossed at the forecast peak and the optimizer prices every block, the present one included, from the forecast alone. The whole difference is therefore what each was already doing through the three peak hours, and the optimizer arrives at them with the tank it built for the forecast peak and spends it with the electrolyzer at turndown. The larger margin gap does not establish forecast-error recovery by MPC. Figure 13.2 draws Table 13.5 as bars.

The mechanism behind the gap is visible in the state of charge rather than in the allocation, as shown in Figure 13.3. The threshold rule charges whenever the price is below $22/MWh, which saturates the tank at hour 14 while the price is still $11/MWh, and it discharges only above $62/MWh, which begins at hour 17 and immediately meets the 340 MWth rate limit. Three hours of full discharge leave 58 % in the tank at hour 20 with the peak over, and the rule moves 1008 MWh out of the store over the whole day. The optimizer charges overnight to 71.6 %, discharges into the morning ramp, recharges through the midday trough to 93.4 %, begins the evening discharge at 14:31, two and a half hours before the rule, and empties the tank to 10.0 % at 21:33, moving 2591 MWh against the rule's 1008 MWh. It also holds the electrolyzer at its 18 MWe turndown from hour 17 to hour 20 while the threshold rule takes until 19:17 to get there, and it puts the turbine on its 1180 MWth swallowing

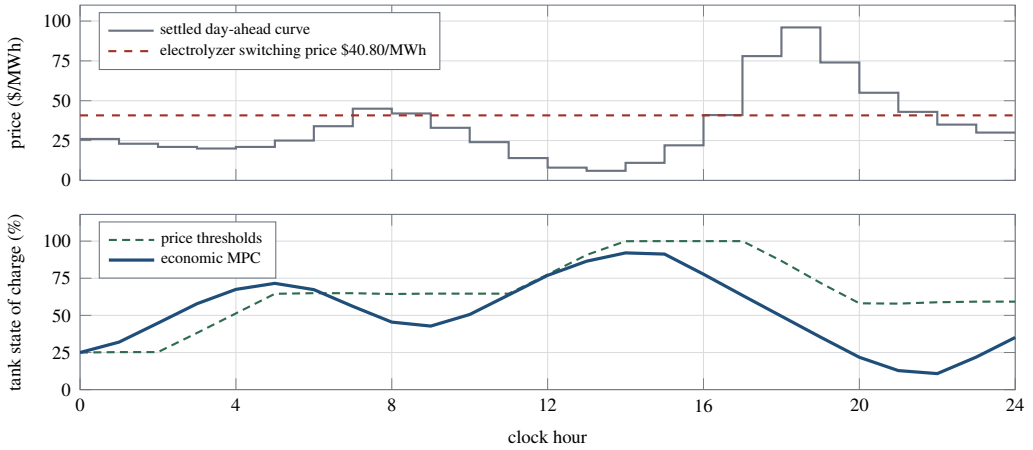

Figure 13.3 Where the twenty-four hour horizon shows up. The upper panel carries the settled price (solid, stepped) and the switching price (dashed), the lower panel the tank under the threshold rule (dashed) and the optimizer (solid). The threshold rule saturates the tank at hour 14 while the price is still $11/MWh, waits for a price above $62/MWh at hour 17 and then meets the 340 MWth discharge limit, so it moves 1008 MWh out of the store over the day and still holds 58 % at hour 20. The optimizer begins its evening discharge at 14:31, empties the tank to 10.0 % at 21:33 and moves 2591 MWh. Hourly samples of the physics-step state.

limit for 431.8 MWe of export at hour 19.5. Nothing told it to do any of that. There is no hydrogen setpoint anywhere either, and 93.24 t against a 72 t contract falls out of the terminal penalty of (13.25) and the merchant price of (13.17), with the 21.24 t of surplus made in the hours below $40.80/MWh.

> **TRY IT IN THE APP** **4 · Linear Economic MPC, the Same Day**
>
> Load scenario 4 and press Run. Watch the hot tank state of charge. The optimizer charges overnight to 71.6 %, spends some of it into the morning ramp, refills through the midday trough to 93.4 %, starts its evening discharge at 14:31 and empties the tank to 10.0 % at 21:33.

13.6.3 Two Formulation Details That Cost Money

Two lines of the dispatch layer are easy to write the obvious way and expensive to get wrong, and both are measurable on the ordinary day with every other setting held. Neither is a setting a learner can change, so the figures in this subsection come from counterfactual model runs on a modified copy of the controller.

Replacing the row-sum step of (13.27) with the textbook $1/\max_i H_{ii}$, a step 9.00 times longer, drops the day from $529,303 to $515,156, a loss of $14,147 that is larger than the entire advantage the optimizer had over the thresholds. The mean value of J after each of the 96 executions of the day rises from 2.36 to 3.06, so the iteration is not converging within its budget. What it does not do is oscillate between turndown and nameplate on alternate iterations. The projection absorbs most of the overshoot, and the largest back-and-forth move in a normalized electrolyzer decision (the smaller of two consecutive iteration moves of opposite sign) grows only from 0.0112 to 0.0274, so the symptom is a worse plan rather than a visible instability, which is the harder failure to notice.

Making the contract penalty symmetric, by keeping the terminal gradient in (13.26) when the plan is already above target, costs $20,463 on the same day and drops production from 93.24 t to 70.45 t. The symmetric penalty forbids profitable surplus, so the optimizer stops making hydrogen in the hours where electricity is below the switching price and then finishes 1.55 t short of the contract it was trying to hit exactly.

Table 13.6 Effect of the Iteration Count on the Ordinary Day. The Shipped Value Is 18.

N_{it}	2	4	10	18	120	600
Net margin (k$)	531.5	533.4	531.1	529.3	526.0	525.9
H_2 (t)	90.90	92.66	93.16	93.24	93.57	93.59

The third detail runs the other way and is worth reporting because it is counterintuitive. Table 13.6 sweeps the iteration count with everything else fixed.

Solving the quadratic program thirty-three times harder makes the day $3,396 *worse*, and the whole range from two iterations to six hundred spans $7,476, or 1.4 %. The reason is that the program being solved is a linearization whose bias is corrected every fifteen minutes, so converging it exactly buys accuracy against a model that is already approximate while the move-suppression term keeps pulling a fully converged plan closer to the previous one. A capped iteration count is therefore not a compromise on this plant, and the honest reading of Table 13.6 is that eighteen iterations are sufficient and that no iteration count recovers a bad forecast.

🕮 **Pause.** Try Exercise 13.3 before reading on.

13.6.4 The Hard Days

Scenario 6 inverts every sign in the objective. With the midday hours at −$25/MWh the optimizer holds the turbine at its 250 MWth minimum, the electrolyzer at its 180 MWe nameplate and the plant at −102.7 MW at the meter, so it is paid to import, and it sends 1529 MWh to the cooling tower over the run on purpose because heat that cannot be used is better dumped than converted at a negative price. It finishes $13,719 ahead of the threshold rule, which reaches only 502.5 MWth on the turbine, never imports more than 1.1 MW and dumps 15 MWh.

Scenario 5 binds three constraints inside one hour, the 72 MW/h electrolyzer ramp, the 340 MWth discharge limit and the turbine's swallowing capacity. The optimizer gets the stacks to their 18.0 MWe turndown where the thresholds reach only 42.6 MWe, and it holds 431.7 MWe of export against the rule's 407.4 MWe. It books $734,488 against $706,281, so the horizon is worth $28,207 here, and it is not free, because the optimizer finishes on 80.15 t against an 82 t contract and pays take-or-pay on the last 1.85 t. Two tonnes of hydrogen were worth less than the peak hours they would have cost, which is the trade that writing the contract as a penalty rather than a constraint was for.

Scenario 8 trips the electrolyzer for ninety minutes with the day already behind. Against a companion model run of the identical day with the scripted trip removed, which a learner cannot do, the trip costs 2.90 t of hydrogen and $46,437 of margin. The deficit enters the target of (13.25) rather than the weight, so the Hessian never changes and the solve stays exactly as cheap as it was, and the optimizer buys the deficit back in whichever remaining hours are cheapest instead of declaring the problem infeasible.

Scenario 9 derates the tank to 840 MWh and runs the threshold rule. Charging at 320 MWth through every hour under $22/MWh brings the exchanger derate factor f_{hi} of (13.3) into action, the dump valve opens at 03:10, the dump annunciation follows at 03:11, the tank reaches its capacity at 03:32, the header climbs to 126.54 bar and the dump condenser peaks at 293 MWth, sending 51.4 MWh of reactor heat to the cooling tower in the cheapest hours of the day. The same day on the optimizer never opens the dump valve at all and never takes the header past 120.80 bar, because the capacity constraint is enforced over the whole horizon rather than over the present instant. That is also the clearest statement in the application of what a horizon does not buy, because the optimizer ends the run holding 91.1 % of a small tank it never got to spend and books $482,537 against the rule's $501,272. A horizon bought feasibility here and not money.

> **TRY IT IN THE APP** **9 · The Hot Tank Hits Its Limit**
>
> Load scenario 9 and press Run on the price rules it opens with. Watch the dump valve and the header pressure. The dump valve opens at 03:10, the tank is full at 03:32 and the header climbs to 126.54 bar. Run the day again with the optimizer chosen before Run: the dump valve never opens and the header never passes 120.80 bar.

Scenario 10 opens the generator breaker at 07:00 with the reactor at full power. The header carries the whole imbalance while the salt loop and the dump valve travel, and because the store happened to be discharging 148.5 MWth at that instant the residual reaches 985.4 MWth rather than 860 MWth, which against 3.0 MWh/bar is 5.47 bar/min, and the header peaks at 129.22 bar at 07:03. The reactor alone would give 4.78 bar/min. The plant then settles into the only configuration available, the store charging at its limit, the electrolyzer at nameplate on imported power at −195.2 MW at the meter, and 507 MWth straight to the cooling tower, and 3242 MWh of reactor heat is thrown away over the six-hour outage. That figure is the argument for building the store larger than arbitrage alone would justify. Volume 18, Chapter 5 treats the same event on a pressurized water reactor with the reactor kinetics and the secondary side in the model, where a secondary shut in at full power reaches its first safety valve within seconds.

The graded challenge asks for the full 90 t contract, the header held inside ±6 bar whenever the dump condenser is shut, and a margin above $975,000 through a midday collapse to −$18.8/MWh and an evening spike five times the day-ahead curve that the forecast saw a third of. As loaded, on a fixed split with the optimizer weights deliberately wrong, the day books $641,543. The suggested configuration books $1,003,720 and lands 94.33 t, and the threshold architecture with the case's own thresholds books $935,748 on the same day. The suggested configuration reaches 126.35 bar at its worst, 6.35 bar above setpoint, at 10:32 with the dump condenser 21.1 % open, and stays inside 5.66 bar of setpoint whenever that valve is shut. The carve-out is what makes the target consistent with its own answer. The dump controller holds its setpoint at 126 bar, so deliberate dumping puts the header 6 bar high by construction, and a flat band would ask the student to undo the very action this scenario teaches. The application judges the band on its logged record of header pressure and dump valve position, one sample every three simulated minutes, on which the suggested configuration also reads 5.66 bar.

Figure 13.4 shows the Dispatch view during scenario 4. The price band runs across the top with the forecast the controller was given dashed over it, the realized allocation is drawn to the left of the vertical marker and the optimizer's plan for the rest of the day is ghosted to the right of it, and the state of charge appears realized and planned below. That ghosted trajectory is the controller's intention, it is visible before it happens, and it is the single most useful object on the screen for the purpose this application exists to serve.

13.7 Technology in Practice

The symbols of Figure 13.1 are, in a plant, a pressure transmitter on the steam header, flow measurements on the turbine steam and on the salt circuit, a thermometer in the stacks, a power measurement on the electrolyzer feed, and five final elements. Figure 13.5 draws the chapter's loops with that hardware in place of the schematic's symbols.

What Measures It

The measurement the whole allocation answers to is the header pressure, because Equation (13.1) closes through it. Both the pressure trim and the dump override read it. The chapter gives its transmitter the 0.004 h lag τ_{Pm} and the steam flow transmitters the 0.006 h lag τ_{Qm}. In a pressurized water plant the turbine control valve is set from the turbine impulse pressure, a surrogate for turbine power [5], where the chapter's throttle loop reads the heat to the turbine directly.



Figure 13.4 The Dispatch view of scenario 4 at $t = 6.00$ h, with the optimizer's plan for the rest of the day ghosted ahead of the vertical marker, cut to the plant view and the heat readout of the control panel. Callouts: (1) the price ribbon, (2) the present, (3) the planned allocation, (4) the tank state of charge, realized and planned, (5) the heat balance readout. The price ribbon at the top carries the settled hourly curve with the forecast the controller was given dashed over it, with the dearest hours shaded amber and the cheapest violet on screen. Beneath it the allocation shows what has happened as solid traces to the left of the marker and what the controller intends as dashed traces to the right: turbine heat is planned down through the midday trough and up to its 1180 MWth swallowing limit from 15:45 to 21:45, the electrolyzer is planned in the opposite sense, down to its 18 MWe turndown over the peak, and the dump trace is flat on zero because at a positive price dumping costs the electricity the heat would have made. The hot tank state of charge below that stands at 67 % and is planned to fall through the morning ramp, refill to 96 % by mid-afternoon and empty into the evening peak. At this instant the heat balance reads 1000 MWth to the turbine, 170 MWth out of the store, 31 MWth of electrolyzer steam and nothing dumped, against the constant 860 MWth the reactor supplies, and the application reads $131,365 booked with 29.4 t of the 72 t daily contract delivered, 212 MW of export at the $25/MWh price of the hour then ending and the header on 120.1 bar.

TECHNOLOGY **Sensor: The Header Pressure Transmitter**

A common electronic pressure transmitter senses with a taut metal diaphragm between two fixed plates, so that a pressure difference flexes the diaphragm and changes two capacitances [20]. Hot steam would damage the sensing element, so the impulse line gets a low point or a pigtail siphon where condensate collects as a liquid barrier. A transmitter that follows the NAMUR signal levels drives its current to 3.6 mA or below, or to 21 mA or above, when its sensing element fails, and the control system can then force the loop into manual [20]. On a trend that is a jump off scale, not a drift.

What Moves It

Five final elements act on the plant: the turbine throttle, the salt pumps, the dump valve, the rectifier of the electrolyzer and its electric superheater. The salt pumps carry the chapter's lesson, because the pressure controller trims their duty around whatever the supervisor asked for. A review of hybrid systems lists remote flow monitoring and control valves and pumps for molten salts among the hardware such plants will need [17]. The electrolyzer moves faster than any of them, because 75 to

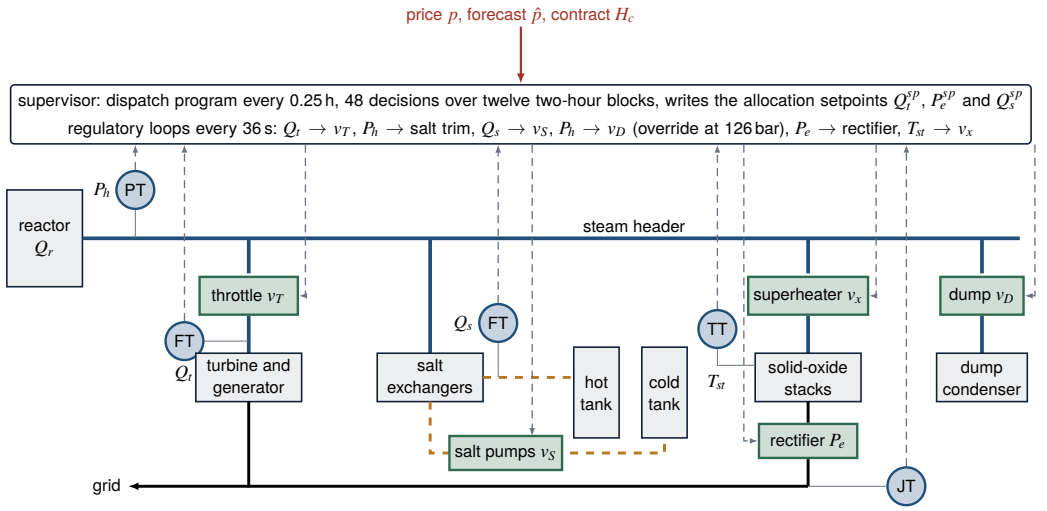

Figure 13.5 The chapter's loops with their hardware. Sensors are the round tags: the pressure transmitter PT on the steam header, flow transmitters FT on the turbine steam and on the salt duty, the thermometer TT in the stacks and the power measurement JT on the electrolyzer feed. Final elements are the square tags: the turbine throttle, the salt pumps, the dump valve, the rectifier and the electric superheater. Thick solid lines carry steam, dashed lines salt and plain black lines electricity. Notice that the supervisor reads only the price, the forecast, the contract and the plant state, and acts through three setpoints, while the header pressure reaches the salt pumps through the regulatory layer.

90 % of its power is electricity for electrolysis, which can be changed almost immediately [5].

> **TECHNOLOGY** **Final Element: Long-Shafted Salt Pumps**
>
> No mechanical seal suitable for molten salt has yet been identified, so salt pumps are vertical machines sealed by a throttle bushing, with the leakage running back to the suction by gravity [21]. In storage plants they hang from a structure above each tank and reach through its roof on shafts of 12 to 15 m, with bearings lubricated by the salt, and the tank keeps a minimum heel of salt over the impeller. Variable-frequency drives on the salt pumps allow the salt flow to be trimmed finely, and plants report few major problems apart from installation and alignment [21]. The chapter's one signed command v_S stands for the pumps of both tanks, charging one way and discharging the other, behind the 0.015 h lag τ_{pump}.

Where the Controller Runs

In a distributed control system a supervisory controller sends setpoints to the field controllers and requests data from them, and the field controllers move their actuators from those commands and their own sensors [22]. Supervisory and lower-level loops run with cycle times from milliseconds to minutes. The chapter's six loops run on a 36 s scan above a 7.2 s physics step, and the dispatch program writes its three allocation setpoints every 0.25 h. The dispatch program is the supervisory controller of that description, the six loops are its field controllers, and the pressure trim of Equation (13.20) is a summing block between the two layers.

IN PRACTICE **Keep the Stacks Hot**

Solid-oxide cells crack under steep thermal gradients. The rule of thumb that makers recommend for start-up and shutdown is no more than 2 °C per minute, so reaching 800 °C takes about 6.5 hours, too slow for a plant that moves between markets [5]. Makers therefore offer a hot standby mode. The electrolysis power is cut while the steam heat and the electric topping heat stay on, so no gradients form and production can resume at any moment [5]. In the chapter the dispatch never takes the stacks below their 18 MWe turndown, and the degradation term grows whenever they run colder than 720 °C.

What Is Modernizing It

Three developments touch this plant. The first is grid service from the electrolyzer. Nuclear plants have historically been unable to sell ancillary services because of their ramp limits, but a plant coupled to steam electrolysis could, because the electrical part of the electrolyzer's power can be moved back and forth almost at once [5]. The second is online optimization itself. More data generated, collected and used now makes continuous online optimization of a coupled plant possible [17], and the dispatch layer of this chapter is such an optimization. The third is the network. Operational systems increasingly resemble information systems and are less isolated from the outside world, so they need more security [22], and a supervisor that reads market prices sits at that boundary.

13.8 Limitations

The application is not a techno-economic study and cannot be used as one. With a surrogate steam cycle, a specific-energy electrolyzer and one deterministic price curve it cannot support a claim about whether a real integrated energy system pays for itself, what tank size is optimal, or what a hydrogen contract should be priced at, and the published feasibility work that can support such claims carries a full plant model and a market simulation [5, 17]. The reactor is a constant with a trim multiplier, so none of the neutronics, xenon or thermal-fatigue constraints that actually decide whether a nuclear unit may be moved appear anywhere [1]. The dispatch layer is a linearization with a measured bias, the feasibility restoration is greedy rather than an exact projection, and eighteen iterations return a good feasible plan rather than a solved program, so the trajectory drawn on the screen should be read as an intention and not as an optimum. Two-hour blocks smear one-hour prices, and a block averaging $39.50/MWh receives a partial electrolyzer allocation where hourly blocks would have given a full one in the cheap hour and none in the dear one. Hourly blocks over the same horizon would be ninety-six decision variables and roughly four times the solve cost, which is the trade the application makes on purpose. Section 13.3.8 states where each individual assumption stops holding.

OPEN CHALLENGE **Several Plants, One Supervisor**

A reactor, a grid connection and a hydrogen plant are separate systems with their own rules and time scales. A review of nuclear-renewable hybrid systems calls for techniques that coordinate linked but independent systems with differing dynamic time scales, and for new schemes of supervisory monitoring and control [17]. In 2019 high-temperature steam electrolysis had yet to be validated with a large generator at commercial scale, and how its stacks should be ramped thermally was still undecided [5]. Progress would be a coupled plant run for a season under a supervisor like the one of this chapter, with its constraints and its value measured.

13.9 Discussion

The result that matters for practice is the one in Table 13.5 read across both halves. An optimizer justified on average-day performance is justified on the wrong evidence, because the average day is where a heuristic is closest to it. The case for the horizon is built on the days where a constraint binds hours after the decision that caused it, which is scenario 9, where three constraints bind inside one hour, which is scenario 5, and where the price does something the rule was not written for, which is scenario 6. A plant that cannot articulate which of those it expects to meet, and how often, cannot

size the benefit of the layer it is buying, and this is the same conclusion that assessments of real-time optimization practice reach from industrial experience [6].

The second result is about forecasts. Scenario 7 gives the controller a day-ahead curve saying the evening peak is \$96/MWh when it will be \$528/MWh, and the controller acts on that curve with complete confidence. Re-planning every fifteen minutes rebuilds the plan from the state the plant is actually in, but the plan prices every block from the forecast, so the settled price never reaches it and the day is dispatched exactly as the ordinary day was. Feedback here corrects the state and not the price, and no re-plan can recover a storage decision taken six hours earlier. Chapter 10 measures the same pattern on a flow battery, with a forecast worth little on an ordinary day and a great deal on a day with one scarce hour. The engineering answer is not a better optimizer, and Table 13.6 shows that solving the same program harder makes the day slightly worse. The answer is a forecast with an uncertainty band and a policy that keeps some tank in reserve against the tail, which costs money on every ordinary day and pays for itself on a handful. That is a risk decision rather than a control decision, and it is the point at which a control engineer should say so.

Methods other than the implemented layer fit the same documented model. The twelve states admit a nonlinear economic formulation carrying the part-load curve and the temperature-dependent stack energy inside the optimization rather than as a measured bias [19], and a scenario-based formulation over the same horizon would price the uncertainty that scenario 7 exposes, as the scenario-tree program of Volume 9, Chapter 7 does for a reservoir on an ensemble forecast. Moving horizon estimation on (13.4) and (13.8) would recover the standing tank loss and the degradation state that the parameter table hides, turning two hidden constants into a soft sensor, and the block structure transfers directly to the cogeneration scheduling formulations the same problem class uses at larger scale [8, 9].

> ⚠ **OPEN CHALLENGE** **Hydrogen Beside a Licensed Reactor**
>
> Hydrogen has its own codes and standards, and a flammable gas near a reactor brings closer scrutiny under the United States nuclear power rules 10 CFR 50 and 52. The hydrogen plant may have to stand outside the exclusion area boundary, which adds security problems, heat lost on the way and concern over losing a heat sink [17]. One path keeps it a non-nuclear-grade chemical plant: the reactor stays safe whatever the hydrogen plant does, events there cannot affect the reactor, and the heat transfer loop between them carries no safety function [14]. The chapter's model has neither the distance nor the safety case.

13.10 Conclusion

This chapter has treated an integrated nuclear energy system as an interactive laboratory for supervisory economic dispatch. Twelve states are integrated, with a steam header of 3.0 MWh/bar that closes the reactor's heat balance continuously, a two-tank store at a 95.04 % thermal round trip, a solid-oxide electrolyzer taking 35.47 MWh/t of plant electricity and 6.4 MWh/t of header steam, and a dump condenser sized for the whole reactor. Six PI loops close the plant on a 36 s scan, and the header pressure controller is direct acting and trims the salt loop around whatever the supervisor asked for. Above them, a 48-variable box-constrained quadratic program over twelve two-hour blocks eliminates the turbine allocation through the heat balance, carries the hydrogen contract as a one-sided terminal penalty, and costs 23.4 µs per execution because its Hessian is constant and its eighteen projected-gradient iterations use a precomputed row-sum step. Measured on one deterministic day, a fixed split books \$470,311, three price thresholds book \$516,300 and the optimization books \$529,303, so the heuristic captures 78 % of the improvement and the optimizer earns \$13,003, while on a day whose evening peak is five and a half times the forecast the same two controllers separate by \$85,447. Replacing the row-sum step with the textbook diagonal step costs \$14,147 and making the contract penalty symmetric costs \$20,463, both larger than the advantage the optimizer had. Extensions of the simulator would add a price forecast with an uncertainty band, hourly blocks with the solve cost measured against them, a reserve policy scored on the tail rather than the mean, and

a reactor model carrying the constraints that decide whether a nuclear unit may be moved at all. Measuring what a supervisory optimizer is worth against the heuristic it would replace is what the simulator makes accessible.

KEY POINTS

- A reactor's heat cannot be turned down, so at every instant all of it goes to the turbine, the salt store, the electrolyzer or the dump condenser.
- On an ordinary day three price thresholds capture 78 % of the improvement over a fixed split, and the optimizer adds $13,003.
- The horizon pays where a constraint binds hours after the decision, where several bind at once, or where the price does something the rule was not written for.
- Re-planning every fifteen minutes corrects the state and not the price, so a spike the forecast missed is dispatched exactly as an ordinary day.

Where else this applies. A plan that prices a store against a forecast also sets the heat price of Volume 9, Chapter 6 and the reservoir release of Volume 9, Chapter 7.

References

[1] R. Ponciroli, Y. Wang, Z. Zhou, A. Botterud, J. Jenkins, R. B. Vilim, and F. Ganda, "Profitability evaluation of load-following nuclear units with physics-induced operational constraints," *Nuclear Technology*, vol. 200, no. 3, pp. 189–207, 2017.

[2] P. Denholm, J. C. King, C. F. Kutcher, and P. P. H. Wilson, "Decarbonizing the electric sector: Combining renewable and nuclear energy using thermal storage," *Energy Policy*, vol. 44, pp. 301–311, 2012.

[3] U. Herrmann and D. W. Kearney, "Survey of thermal energy storage for parabolic trough power plants," *Journal of Solar Energy Engineering*, vol. 124, no. 2, pp. 145–152, 2002.

[4] M. A. Laguna-Bercero, "Recent advances in high temperature electrolysis using solid oxide fuel cells: A review," *Journal of Power Sources*, vol. 203, pp. 4–16, 2012.

[5] K. Frick, P. Talbot, D. Wendt, R. Boardman, C. Rabiti, S. Bragg-Sitton, M. Ruth, D. Levie, B. Frew, A. Elgowainy, and T. Hawkins, "Evaluation of hydrogen production feasibility for a light water reactor in the midwest," Idaho National Laboratory, Idaho Falls, ID, Tech. Rep., 2019.

[6] M. L. Darby, M. Nikolaou, J. Jones, and D. Nicholson, "RTO: An overview and assessment of current practice," *Journal of Process Control*, vol. 21, no. 6, pp. 874–884, 2011.

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

[8] J. J. Alves and G. P. Towler, "Analysis of refinery hydrogen distribution systems," *Industrial & Engineering Chemistry Research*, vol. 41, no. 23, pp. 5759–5769, 2002.

[9] S. Mitra, L. Sun, and I. E. Grossmann, "Optimal scheduling of industrial combined heat and power plants under time-sensitive electricity prices," *Energy*, vol. 54, pp. 194–211, 2013.

[10] A. Radovanović, R. Koningstein, I. Schneider, B. Chen, A. Duarte, B. Roy, D. Xiao, M. Haridasan, P. Hung, N. Care, S. Talukdar, E. Mullen, K. Smith, M. Cottman, and W. Cirne, "Carbon-aware computing for datacenters," *IEEE Transactions on Power Systems*, vol. 38, no. 2, pp. 1270–1280, 2023.

[11] J. W. Labadie, "Optimal operation of multireservoir systems: State-of-the-art review," *Journal of Water Resources Planning and Management*, vol. 130, no. 2, pp. 93–111, 2004.

[12] 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.

[13] B. M. Misra and J. Kupitz, "The role of nuclear desalination in meeting the potable water needs in water scarce areas in the next decades," *Desalination*, vol. 166, pp. 1–9, 2004.

[14] International Atomic Energy Agency, *Opportunities for Cogeneration with Nuclear Energy*, ser. IAEA Nuclear Energy Series. Vienna: International Atomic Energy Agency, 2017, no. NP-T-4.1.

[15] L. Hirth, "The market value of variable renewables," *Energy Economics*, vol. 38, pp. 218–236, 2013.

[16] C. Forsberg, "Variable and assured peak electricity production from base-load light-water reactors with heat storage and auxiliary combustible fuels," *Nuclear Technology*, vol. 205, no. 3, pp. 377–396, 2018.

[17] M. F. Ruth, O. R. Zinaman, M. Antkowiak, R. D. Boardman, R. S. Cherry, and M. D. Bazilian, "Nuclear-renewable hybrid energy systems: Opportunities, interconnections, and needs," *Energy Conversion and Management*, vol. 78, pp. 684–694, 2014.

[18] S. Richter, C. N. Jones, and M. Morari, "Computational complexity certification for real-time MPC with input constraints based on the fast gradient method," *IEEE Transactions on Automatic Control*, vol. 57, no. 6, pp. 1391–1403, 2012.

[19] J. B. Rawlings, D. Angeli, and C. N. Bates, "Fundamentals of economic model predictive control," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 2012, pp. 3851–3861.

[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] M. Mehos, H. Price, R. Cable, D. Kearney, B. Kelly, G. Kolb, and F. Morse, "Concentrating solar power best practices study," National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5500-75763, 2020.

[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.

13.11 Exercises

13.1 Concept check · 2 min

Table 13.3 gives the header pressure controller the action -1, direct acting. Why is that the right action for this loop?

(a) A rising pressure calls for more charging, and charging takes steam out of the header, so the output must rise with the measurement

(b) The salt pump fails open, so a larger signal closes it and the sign of the loop has to be turned around

(c) Direct action keeps the loop out of windup, because its output then never reaches either end of its span

(d) The dump controller is reverse acting, so the trim must act the opposite way to avoid fighting it

13.2 Concept check · 1 min

On the ordinary day the threshold rule captures 78 % of the improvement the optimizer makes over the fixed split. Which reason does the chapter give?

(a) The default thresholds were tuned on this day, so the rule starts close to the optimum

(b) The price curve is smooth and no constraint binds, which is the case a heuristic handles well

(c) The two-hour blocks smear the prices so badly that the optimizer loses most of its advantage

(d) Eighteen iterations stop the optimizer far short of the optimum of its own program

13.3 Concept check · 2 min

Raising the iteration count of the dispatch program from 18 to 600 makes the ordinary day \$3,396 worse. What explains it?

(a) The row-sum step becomes unstable when it is repeated that many times

(b) Six hundred iterations overrun the fifteen-minute period, so each plan arrives late

(c) The program is a linearization corrected every execution, and move suppression pulls a converged plan toward the last one

(d) The one-sided contract penalty turns symmetric once the iteration has converged

13.4 Theory and calculation · 5 min

Merchant hydrogen sells at $c_m = \$1{,}550/\mathrm{t}$. The dispatch layer carries $e_{tot} = 35.46\,\mathrm{MWh/t}$ of plant electricity and $q_{H_2} = 6.4\,\mathrm{MWh/t}$ of header steam per tonne.

(a) Compute the hydrogen value per megawatt-hour of electrolyzer load, c_m/e_{tot}, and the steam ratio $\kappa = q_{H_2}/e_{tot}$. *Checkpoint:* \$43.71/MWh and 0.1805.

(b) Use Equation (13.17) to compute the switching price at the design cycle efficiency of 39.60 % and at the rated 40.00 %. By how much is the default threshold of \$44/MWh too high? *Checkpoint:* \$40.80/MWh and \$40.77/MWh, so \$3.20/MWh too high.

(c) Explain in one sentence why the switching price is lower than c_m/e_{tot}. *Checkpoint:* Each megawatt of load also diverts κ megawatts of heat that the turbine would have converted.

13.5 Theory and calculation · 6 min

The header stores $C_h = 3.0\,\text{MWh/bar}$, Equation (13.1).

(a) A residual of 100 MWth persists. How fast does the header move, and how far does it move in one regulatory scan of 0.01 h? *Checkpoint:* 33.3 bar/h, 0.33 bar per scan.

(b) When the turbine trips at full power the residual is the reactor's 860 MWth, or 985.4 MWth with the store discharging at the instant of the trip. Convert both to bar per minute. *Checkpoint:* 4.78 bar/min and 5.47 bar/min.

(c) The dump controller acts from 126 bar and the dump valve has a travel time $\tau_D = 0.010\,\text{h}$. Estimate how far above 126 bar the header rises during one travel time at the larger rate, and compare with the measured peak of 129.22 bar. *Checkpoint:* About 3.3 bar, to about 129.3 bar.

13.6 Theory and calculation · 7 min

Use Equation (13.5) with $\eta_0 = 0.400$, $c_p = 0.12$, $c_d = 0.055$, $Q_{t,\max} = 1180\,\text{MWth}$ and $Q_r = 860\,\text{MWth}$.

(a) Charging at 320 MWth puts the turbine at 520.8 MWth, and discharging at 340 MWth puts it at 1180 MWth with $f_d = 340/(860+340)$. Compute the cycle efficiency at each point. *Checkpoint:* 38.50 % and 39.38 %.

(b) With these endpoint efficiencies and the thermal round trip $\eta_c\eta_d = 0.9504$, compute the net value of one thermal megawatt-hour charged in the \$6/MWh hour and returned in the \$96/MWh hour, ignoring salt pumping. *Checkpoint:* \$2.31 forgone, \$35.93 returned, \$33.62 net.

(c) Repeat for a charge at \$24/MWh and a return at \$35/MWh, and compute the electrical round trip on the same convention. *Checkpoint:* \$3.86 net, 97.2 %.

13.7 App challenge · 8 min

Load scenario 2, Fixed allocation, one whole day, change no setting, and run at speed 6 (six simulated hours per wall second) with the default seed past 24 h. Pause and export the CSV file. Do the same for scenario 3, Price-threshold rules, the same day, and scenario 4, Linear economic MPC, the same day. In each file use the row at `time_h` = 24.00.

(a) Read `metric_margin` and `h2_cum_t` in each file. *Checkpoint:* \$470,311, \$516,300 and \$529,303, within \$1, and 72.07, 100.55 and 93.24 t.

(b) Compute the share of the optimizer's improvement over the fixed split that the rules capture, as a percentage of that improvement. *Checkpoint:* 78.0 %.

(c) Read `soc_pct` in the rules and optimizer files. Convert the difference to megawatt-hours of a 2400 MWh tank, value it discharged at $\eta_d = 0.96$ and a cycle efficiency of 39.4 % into the \$26/MWh hour, and recompute the optimizer's advantage and the rules' share. *Checkpoint:* 59.25 and 35.13 %, 579 MWh, \$5,693, an advantage of about \$7,300 and a share of 88 %.

13.8 App challenge · 8 min

Load scenario 7, The forecast is wrong and the controller acts on it anyway, and run at speed 6 with the default seed past 24 h. Pause and export the CSV file. Press Reset, set Dispatch architecture to Price rules before pressing Run, and repeat. Press Reset again, set Fixed split, and repeat. In each file use the row at `time_h` = 24.00.

(a) Read `metric_margin` in the three files. *Checkpoint:* $1,010,196, $924,749 and $709,818, within $1.

(b) Compute the optimizer's advantage over the rules, and divide it by its ordinary-day advantage of $13,003. *Checkpoint:* $85,447, about 6.6 times.

(c) Read `h2_cum_t` and `metric_charged_MWh` in the optimizer and rules files and compare them with the ordinary-day files of scenarios 4 and 3. Explain what you find. *Checkpoint:* 93.24 t and 2976 MWh for the optimizer, 100.55 t and 1896 MWh for the rules, both identical to the ordinary day.

13.9 App challenge · 8 min

Load scenario 9, The hot tank hits its limit, change no setting, and run at speed 6 with the default seed past 23 h, which is midnight on this run. Pause and export the CSV file. Press Reset, set Dispatch architecture to Economic MPC before pressing Run, and repeat. In each file use the rows up to `time_h` = 23.00.

(a) In the rules file read `metric_dumped_MWh`, `metric_margin` and `soc_pct` at $t = 23.00$ h, and the largest `hdr_bar`. *Checkpoint:* 51.4 MWh, $501,272, 0.0 % and 126.42 bar, within 0.05 bar.

(b) Read the same four values in the optimizer file. *Checkpoint:* 0.0 MWh, $482,537, 91.09 % and 120.80 bar.

(c) The optimizer books less money. Say what its horizon bought instead, and how much stored heat it holds at midnight in a tank derated to 840 MWh. *Checkpoint:* No dumping and a header near setpoint, with about 765 MWh left in the tank.

13.10 Application area · 4 min

An engineer at a utility asks why the plant needs a salt store, an electrolyzer and a turbine larger than the reactor when the reactor itself could follow the price. Give two economic and two physical reasons from the chapter, and say what the 1180 MWth swallowing capacity has to do with the answer.

13.11 Application area · 3 min

During the six-hour turbine outage of scenario 10 the plant throws 3242 MWh of reactor heat away with 507 MWth settled on the dump condenser. Why is the dump condenser sized for the whole reactor, and what does the assumption that the turbine is one efficiency curve leave out of this event?

13.12 Review the agent · 5 min

An AI assistant reviews scenario 4 and writes: "The optimizer delivers 93.24 t against a 72 t contract. The 21.24 t of surplus is a tracking error that wastes electricity. Make the contract penalty symmetric, so that the plan lands on 72 t exactly and the margin rises." Find the flaw and say how you would check it with the app.

CHAPTER 14

Summary and Outlook

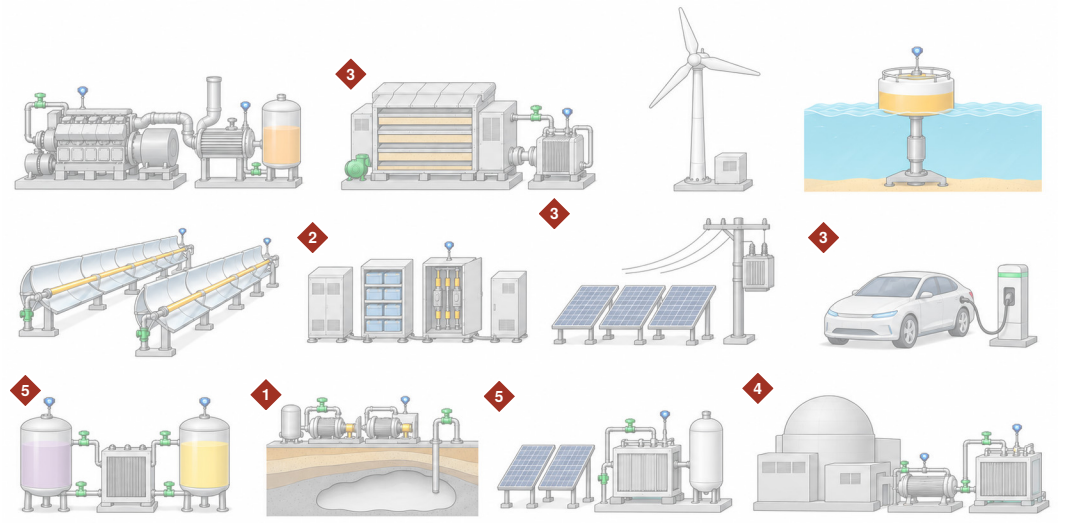

Open gaps 1 slow clock 2 inertia and meter 3 unmeasured limit 4 forecast error 5 silent degradation

Notice that the fuel cell, the feeder and the charger share one gap, and that no gap marks the gas engine, the wind turbine, the wave buoy or the trough field. Illustration, not to scale.

IN PLAIN TERMS **What Twelve Energy Plants Taught**

The main idea. Each chapter of this book ran one plant that makes or stores energy in a simulator, from a gas engine to a cave full of air. Every plant had fast controls that act within seconds and a slow plan that looks hours or days ahead.

The challenge. A fast control can hold its reading steady while the plant wastes money or wears out. The limit that ends a run is often one that no instrument can see.

Why it matters. The plants that did best kept a store between the two layers, judged every plan by what was left in the store, and let safety rules act first.

Twelve chapters controlled twelve plants that convert or store energy, from a gas engine to a salt cavern. Each joined a fast regulating layer to a slow deciding one. Table 1.2 of Chapter 1 listed the two layers of every chapter, and Section 1.4.7 named the four ways they are joined: a store that lets them run at different rates, a setpoint written by the slow layer and held by the fast one, a limit or interlock that lets protection take the actuator from either, and a forecast that lets the slow layer act before the fast one sees anything. This chapter adds what each arrangement achieved and where it failed, draws a short guide for choosing a method, looks at where practice is heading, reports how the volume itself was made with AI agents and what they got wrong, and lists the problems that remain open.

14.1 What the Chapters Showed

Tables 14.1 and 14.2 compare the chapters. The second column condenses Table 1.2. The other columns name an arrangement that worked, one that failed, and the measured result with the table or

Table 14.1 Fast Regulation and Slow Dispatch, Chapters 2 to 7

Chapter	Fast and slow layers	What worked	What failed	Measured result
2, gas engine CHP	Power in 2.3 s, recovered heat in 68 s. A hot-water buffer and a UPS between them.	Metered-load feedforward. A short tracking time. Fast jacket and supply loops in the challenge.	Feedback alone. Any tuning once demand exceeds two engines. A slow cooling loop that the buffer hides.	Import IAE 18.82 against 4.14 MW s (Table 2.7). Challenge UPS low 51.97 shipped against 79.19 % suggested (Table 2.8).
3, fuel cell CHP	Current on the inverter's 50 ms. Anode in 1.38 s, fuel arriving in 17.3 s.	A governor on the arriving fuel with the feedforward. The UPS paying for the wait. An estimate of the gas.	Load-to-current control, which leaves the inverter's foldback to stop the current. Feedforward alone.	Anode margin 0.157 naive against 0.623, power IAE 2.638 against 0.797 kWh (Table 3.6). Challenge passed on 31 of 31 seeds (Section 3.8).
4, wind turbine	Torque and pitch on one speed. Region logic and a gain schedule. A wind preview seconds ahead.	Torque below rated, pitch above, the pitch gain scheduled. A tight objective with a preview.	The schedule off, which holds speed tightly and rides the pitch rate limit. The challenge's weak gains.	IAE 3.058 with the schedule against 0.865 rad without it, pitch travel 108.4 against 255.2° (Section 4.6.2). MPC energy cost 2.6 to 5.2 % over ten seeds (Table 4.6).
5, wave buoy	Force law every 0.1 s. Sea state estimated over 120 s.	MPC on a forecast built from its own force estimate. A survival lock on the estimated sea.	A damper, which resists only motion already there. The MPC without survival in the storm.	Five minutes: damper 0.767, MPC 2.522 on its forecast and 2.713 kWh on the truth (Table 5.4). Challenge 10.889 kWh with no contact (Table 5.5).
6, trough field	Outlet loop on field flow. Setpoint optimizer and day schedule. A two-tank store.	A moving outlet setpoint on a broken day. A schedule into the evening peak. Tuned gains.	A fixed setpoint on a broken day. Gains alone or the schedule alone in the challenge.	Broken day 25.4 against 33.5 MWh, clear day 120.8 against 121.1 (Table 6.6). Challenge 3.2 shipped against 72.8 MWh suggested (Section 6.7).
7, campus microgrid	Virtual inertia and droops within milliseconds. Secondary, SOC manager and genset start over minutes.	A battery that forms the grid with 2 s of virtual inertia. Shedding. The genset started by charge.	A battery that follows the grid. Too little inertia. The one-hour dispatch on adjusted fuel.	Following: bus dead at 60 ms, 11.75 kWh critical lost. Forming: 2.48 Hz/s, nadir 59.516 Hz (Table 7.6). Challenge MPC adjusted 631.64 against 585.31 L (Section 7.7).

section of the chapter that reports it.

The columns make five points, and Figure 14.1 draws one pair of results from each row.

The fast layer holds what it measures, and the slow layer decides the value. In Chapter 10 the power loop reaches 63 % of a step in 18 s, and the chapter finds nothing in it to teach beyond how easy it is. On the day with one scarce hour the same plant booked −$2,242 under feedback and $14,599 under the dispatch layer (Table 10.8). In Chapter 6 the outlet loop held 393 °C through a broken day while the receiver went on losing heat, and a setpoint chosen by the optimizer raised net generation from 25.4 to 33.5 MWh (Table 6.6). On a clear day the same optimizer was worth 0.3 %. Through a single cloud the temperature trend said the controller succeeded and the energy trend said the plant did not. A fast loop that works is necessary. It says little about whether the plant is worth running.

Table 14.2 Fast Regulation and Slow Dispatch, Chapters 8 to 13

Chapter	Fast and slow layers	What worked	What failed	Measured result
8, solar feeder	Inverter curves in seconds. A tap changer on a 45 s timer and a central program.	Curves at the default slope, a compensator matched to the line, a slower band. MPC with a clear-sky preview.	A steep slope against a reactance-weighted compensator. The tap changer alone.	Tap operations 2 against 18 (Table 8.4). Challenge 0 of 6 seeds shipped, 6 of 6 suggested (Table 8.5). Case 9 out of band 21.2 against 0 min (Table 8.6).
9, fast charger	Four current limits in one selector. A charge to target over half an hour. The anode has no instrument.	The anode map in the selector with a margin and a preload. The coolant loop on temperature.	Constant current, which plates unseen. No preload. A margin smaller than the table's error.	To 80 %: 16.84 min and 3.381 Ah plated against 21.34 min and none (Table 9.5). Margin covers a table about 20 % optimistic (Table 9.7).
10, flow battery	Power loop in 18 s. A 24-hour dispatch program with a bounded net-load trim.	Feedback on a clean day. Dispatch on a day with one scarce hour. Feedforward flow.	Feedback through a price it cannot measure. Dispatch judged on regulation.	Clean day \$35,083 against \$35,067, worst ramp 2.61 against 5.55 MW/min. Scarcity day −\$2,242 against \$14,599 (Table 10.8).
11, air cavern	Grid trim and heat loops hourly. Inventory policy and seasonal reserve over months.	The heater armed. A rolling plan with a reserve. The operator curve, which holds the most gas.	The threshold rule and the plan with the heater off. A floor in megawatts computed at the present scan.	Heater off: default 0.56 for the rule, 2 MWh for the curve, 0.47 GWh for the plan (Table 11.8). Challenge 1.11 GWh default shipped, none suggested (Table 11.6).
12, solar hydrogen	Rectifier loops in 1.5 s. A state machine and an evening schedule. Gas stores behind interlocks.	Interlocks choosing the mode and economics the setpoint. Stored oxygen. Structure and tuning together.	Tuning alone or structure alone in the challenge. The economic schedule judged on net margin.	Three cloudy days: firm \$16,508, economic \$15,656, following \$15,468 (Table 12.8). Challenge −\$44,664 shipped, \$5,467 tuned only, \$6,959 suggested (Section 12.7).
13, nuclear heat and H_2	Six PI loops on the header, a 36 s scan. An economic plan every 15 min over a day.	Price thresholds on an ordinary day. The economic plan on a day with a spike.	A fixed split. Re-planning on a forecast blind to the spike, which prices the forecast only.	Ordinary day \$470,311 fixed, \$516,300 thresholds, \$529,303 MPC. Spike day \$709,818, \$924,749 and \$1,010,196 (Table 13.5).

A store lets the layers run apart, and its state enters every comparison. Every plant in the volume has a store between its layers, and the store changes what a comparison means. In Chapter 2 the hot-water buffer hid a slow cooling loop for eight minutes. In Chapter 3 the UPS paid for the fuel's 17.3 s wait. In Chapter 7 the one-hour dispatch burned 397.47 against 476.30 L of diesel, but it ended with the battery at 62.18 against 77.98 %, and priced at the replenishing fuel it used 465.02 against 476.30 L. A plan with only one slot recovered almost all of that (Table 7.8). In Chapter 13 the thresholds left the tank at 59.3 against 35.1 %, worth \$5,693 if sold overnight and \$21,000 at the next evening's peak, enough to reverse the ranking (Section 13.6). A dispatch comparison without the end state of its store is not a comparison.

A forecast pays where feedback cannot see the cause. On a clean day the flow battery's forecast was worth nothing, \$16 less than feedback, because the highest net load and the highest price fall in the same hour. One scarce hour at three in the afternoon separated them by \$16,841. The nuclear plant's thresholds captured 78 % of what its optimizer added on an ordinary day, and the optimizer

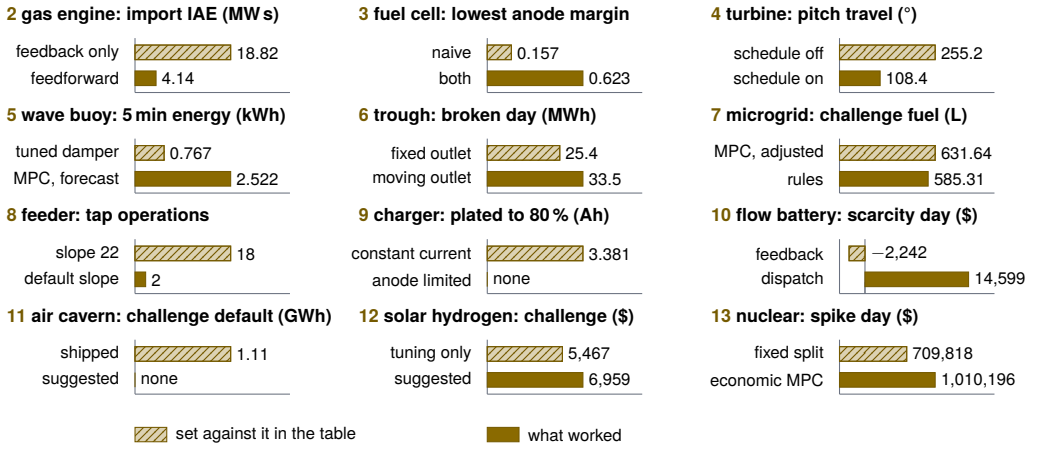

Figure 14.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 14.1 and 14.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 3, 5, 6, 10, 12 and 13, and a smaller one in the others. The solar hydrogen panel sets the tuning alone, which falls short of the challenge's $6,000 target, against the suggested answer. Notice that in most rows the better result came from a change of structure, measurement or information, not of tuning alone.

booked $85,447 more on a day with an evening spike (Table 13.5). In Chapter 11 a pair of price thresholds booked $0.77 M more than the rolling plan, because the price is affine in the net load that the plan forecasts. At a shorter scale the turbine's preview supplied 2.1 of the 5.3 points of its gust advantage, and the rest came from a tighter objective that a shipped option of the loops matched on that gust (Section 4.6.5).

Protection and the unmeasured limit act below every layer. The limits that ended a run were rarely the ones on a gauge. On the fast charger the charger ceiling and the anode map held the current for all but 0.32 min of a charge to 95 %, so two of the four limits were decorative, and the one that mattered has no instrument. A constant-current charge looked identical on every instrument while it plated 3.381 Ah, and 13.867 Ah from a cold pack (Table 9.5). In Chapter 8 the band, the timer, the tap clamp and Volt-Watt act below either central controller. Through the worst cases the application can pose no run curtailed solar or stepped the tap faster than its timer, and in a transmission sag both central structures lost to the local devices because the band's timer set the loss (Table 8.7). The buoy passed its storm only with a survival lock on the estimated sea, and in Chapter 12 the interlocks choose the mode before any economics may choose a setpoint.

Tuning repairs the fast layer, and structure repairs the slow one. The challenges separated the two kinds of change. In the trough field's challenge the tuned gains alone generated 68.7 MWh and booked $2,663, while the schedule with the weak gains booked $4,751 on 41.1 MWh. Together they generated 72.8 MWh and booked $6,780 (Section 6.7). The gains recovered the energy and the schedule earned the money. In Chapter 12 neither tuning alone nor structure alone passed. On the flow battery the structural decisions were worth $35,033 and 56.2 MWh of pumping, and no gain change in the default structure came close. In the gas engine's challenge each loop guarded a different part of the target, so either one tuned alone failed the other part (Table 2.8). A fast loop is tuned against its disturbance. A slow layer is chosen against the question the plant is paid to answer.

14.2 Choosing a Method

The chapters suggest an order of questions for an energy plant that is new to the reader. Each question names the chapter where its answer decided the result, and Figure 14.2 draws the list as a chart.

1. *Which clocks does the plant run on?* Place its loops and layers on one axis, as Figure 1.3 does, and check that each layer is several times slower than the one below it.
2. *What makes the voltage or the frequency when the grid is gone?* A source that forms the grid, with enough inertia for the meter that protection uses, comes before any dispatch (Chapter 7).
3. *Is the binding limit measured?* If not, infer it and size the margin by the model error it must cover, as the charger's 15 mV covers a table 20 % optimistic (Chapter 9). Find which limit carries the plant before tuning the others.
4. *What does the fast layer wait for?* A fuel that arrives 17.3 s late or heat that answers in a minute needs a store to pay for the wait and a governor that admits only what can be supplied (Chapter 3, Chapter 2).
5. *What is the store's state worth at the end of the horizon?* Give the plan a terminal value, such as the merit-order water value of Equation (1.9), and judge every comparison with the end state priced (Chapter 10, Chapter 7).
6. *Does the right action depend on something the fast layer cannot measure?* If it does, as a price away from the load peak does, plan on a forecast. If it does not, feedback is as good and regulates better (Chapter 10, Chapter 11).
7. *Which layer must lose?* Write the interlocks, overrides, shedding and survival modes first, and let the economics act only inside the mode they leave (Chapter 12, Chapter 5).

A predictive controller paid where it held information or a model the loops did not have. On the buoy it captured 2.522 kWh in five minutes on its own forecast against 0.767 kWh for the tuned damper (Table 5.4). On the feeder's cloudy day at 150 % solar it held every node in band where the program spent 21.2 min out, and its preview of the clear-sky rise took it from 15.9 to 0 min (Table 8.6). Its advantage often belonged to the information more than to the horizon. The fuel cell's governor given the same estimate of the gas reached a peak utilization of 0.805 against the MPC's 0.807 (Table 3.8), and a trough loop with a feedforward from the same pyrheliometer did as well through passing clouds, 45.7 against 46.8 °C min (Table 6.7). Prediction cost something elsewhere. The turbine paid 2.6 to 5.2 % of its energy in turbulence, the charger took 3.22 min longer in scenario 12 to keep its core cool (Table 9.8), and the microgrid's dispatch lost on adjusted fuel. Before reaching for prediction, give the loops the measurement or the estimate the predictive controller would use. Volume 9 treats that choice in full.

14.3 Where Practice Is Heading

Inverters that form the grid. Power systems in Hawaii, South Australia, Tasmania, Texas and Ireland commonly run with inverter-based sources supplying more than 50 to 60 % of demand at an instant [1]. A survey of grid-forming inverters describes the concept as drawing increased attention from academia and industry, takes recently commissioned projects in Australia, the United Kingdom and the United States as examples of the trend, and names the transition between grid-connected and standalone operation among the open challenges [2]. On the campus of Chapter 7 a battery that followed the grid left the bus dead 60 ms after the fault, and the same battery forming it held the island. In the editor's view the virtual inertia of an inverter will be set, documented and audited like a governor's droop, together with the meter that protection reads it through.

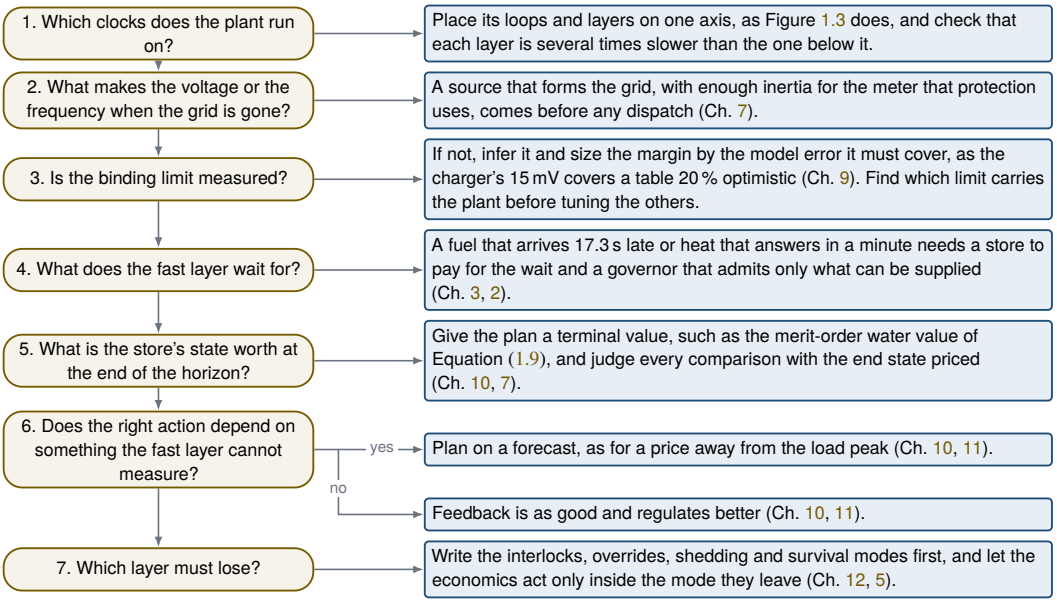

Figure 14.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the grid's own voltage and the unmeasured limit come before any question about the plan.

> ⧖ HISTORY **2014 to 2020: Inverters Asked to Hold the Grid**
>
> The second history box of Chapter 1 ended with an amendment of 2014 that let inverters support the voltage of the grid they feed. In 2018 a European guidance document asked generating units to stay connected while the frequency changed by up to 2 Hz/s over a 500 ms window, 1.5 Hz/s over one second and 1.25 Hz/s over two [3]. In 2020 a research roadmap saw early grid-forming inverters emerging in islanded microgrids with high shares of renewable generation, and judged that replacing synchronous machines with them at the transmission level would take about 10 to 30 years and a research base of protection, controls and interoperability [4]. In the editor's reading the droop of the early 1990s has returned as the inverter that forms the grid, and the campus of Chapter 7 shows why the window of its meter matters.

Preview in the fast layer. Lidar lets a turbine controller act in anticipation of the incoming wind rather than after it has struck the rotor. A review of lidar-enhanced control reports its use for load reduction through pitch and for energy through yaw, and ongoing work that combines pitch and torque in a feedforward nonlinear predictive controller to reduce loads and raise energy at once [5]. Wave energy has not reached commercial viability, and a review of its control holds that control technology can play a major part in getting there [6]. The buoy of Chapter 5 earned 88.5 % of what a 4 s perfect forecast adds within its first 2 s (Section 5.6), and the turbine's preview supplied about two fifths of its gust advantage. In the editor's view a claim for a previewing controller should report the preview's share apart from the share of its objective.

One store on two clocks. A battery used for peak shaving and frequency regulation together, in a joint optimization with degradation and uncertain load and regulation signals, cut a user's bill by up to 12 % on real data. The joint saving was often larger than the sum of the two services alone, and a simple threshold rule in real time achieved it [7]. The flow battery of Chapter 10 is wired the same way, a day-ahead program with a net-load trim beneath it. In the editor's view the stacking of fast and

slow services on one store will make its end state, not its profit on one day, the quantity a dispatch is judged by.

Charging on an estimated limit. Charging an electric vehicle in fifteen minutes is a goal that present lithium-ion cells with graphite anodes cannot reach without harming their performance and safety [8]. Optimal control on a coupled electro-thermal-aging model has been used to trade charge time against degradation under electrical and thermal constraints, with experiments comparing a balanced protocol with a traditional one [9]. On the charger of Chapter 9 a pack whose resistances had risen 60 % plated 3.412 Ah on a strategy that plated nothing when the pack was new. In the editor's view a charging protocol will be delivered with the model error its margin covers, and with an estimate of aging that keeps that error inside the budget.

Electrolysis as a flexible load. Water electrolysis is expected to couple electricity with mobility, heating and chemicals, and a review compares alkaline, proton exchange membrane and solid oxide electrolyzers on the properties that decide flexible operation: load range, load gradients, start-up time and stand-by losses [10]. The two hydrogen chapters meet those properties as control constraints. The solar plant's suggested challenge run paid $4,400 in starts, and the nuclear plant's optimizer held its electrolyzer at its 18 MWe turndown through the evening peak. In the editor's view the dispatch of electrolyzers will be limited as much by starts, turndown and gas purity as by the price of electricity.

Data centers on their own generation. United States data centers used an estimated 192 TWh of electricity in 2024, 4.7 % of the national total, and a reference case puts them at 649 TWh and 11.8 % in 2030, a third of all load growth over those six years. The same report asks that the variability of their demand be characterized, to learn how the grid or generation on site can serve it, and that demand flexibility be explored [11]. Fuel cells that power data centers are slow to follow a sudden rise in demand because of fuel delivery. Capping the servers limits their performance, and storage alone must be greatly oversized for the largest surge [12]. The engines of Chapter 2 and the fuel cell of Chapter 3 both carry such a load and pay for the wait with a UPS. In the editor's view the ramp a data center may ask of its own generation will become a stated term of its supply, set with the fuel's delay in view.

Hydrogen in the pipeline gas. Natural gas utilities have begun to blend hydrogen into their existing pipelines. Hydrogen carries less energy per volume, so at a blend of 20 % the volumetric fuel flow must rise by 15 % to deliver the same energy. On a stoichiometric engine whose controller had no warning of the change, sudden steps in hydrogen swung the air-fuel ratio and released bursts of nitrogen oxides, and better-tuned controllers responded faster and emitted less [13]. The fuel cell plant of Chapter 3 meets the same disturbance as a fall in methane content, and an analyzer read from the start of the run cut the energy drawn from its UPS from 48.37 to 36.50 kWh (Table 3.6). In the editor's view a reading of gas quality, fed forward, will become part of the fuel system of every engine and fuel cell that burns blended gas.

Estimates in place of meters. A study of best practices at concentrating solar plants found the heat transfer fluid's flow meters the one noted exception to otherwise reliable instruments. Plants estimate the flow from pump current, drive speed and pressure drop instead, two plants were pleased with newer schemes for automatic flow control, and the study asks for performance forecasts from irradiance data at 5 to 15 min rather than hourly, because hourly data miss short cloud transients [14]. The outlet loop of Chapter 6 acts through that same field flow, and its passing clouds are the transients the study names. In the editor's view such an estimate will be judged, like the charger's anode map, by the error it carries. Figure 14.3 sets developments of this section and of the volume's chapters on three lanes.

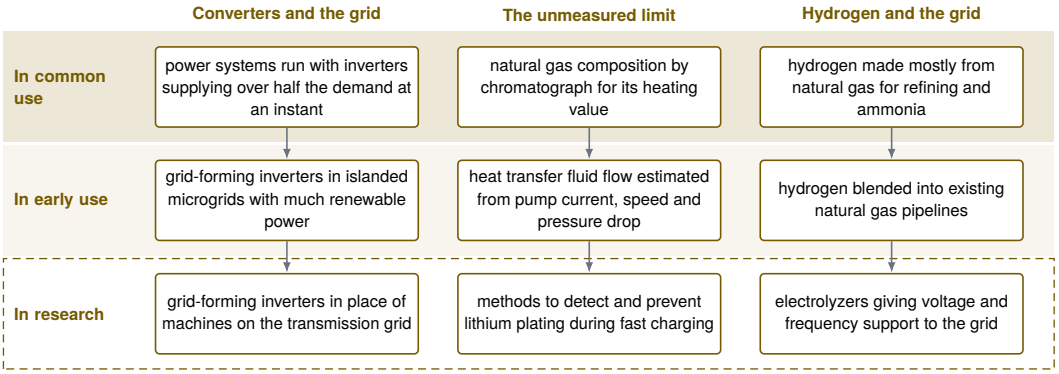

Figure 14.3 Developments in three threads. The top lane holds what the cited source calls standard or typical practice, or reports in common operation, the middle lane what it reports in service at some sites or in early use, and the bottom lane what it reports from simulation, from laboratory projects or as a need, in the editor's reading. Sources by thread: converters and the grid [1, 4], the unmeasured limit [14–16] and hydrogen and the grid [13, 17]. Notice that each thread moves from what a plant measures or follows toward what it must estimate or hold for others.

14.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [18]. In the editor's view, agents can now carry out much of the work that used to fill the hours of a control project. They derive a model from balances or fit one to plant data. They write the simulator and its tests, code the controller, tune it against stated criteria, script commissioning checks, write the documentation and watch running loops for degradation. The apps and chapters of this series were produced that way under the editor's direction, and the account that follows this section reports what the agents got wrong in this volume and how it was caught.

In the editor's view, this changes where an engineer's value lies. When a model, a controller and a report can each be produced in minutes, producing them is no longer the scarce skill. The work moves to five things.

Quality control. An agent's output arrives as a claim: this model fits, this tuning settles in four minutes, this code implements the equation. The engineer checks claims against the plant or the simulator. A reported number is measured again on a fresh run at the stated seed. The code is read for what it clamps, what it initializes and what it reads, because a display that shows the controller's model in place of the plant looks correct until the two differ. A tuning is tried on a disturbance the agent did not test. The most useful checks are independent of the author, made by a second agent or a person who has not seen the first answer, and they are checks that can fail.

Supervision. The engineer sets the scope of each task and decides what an agent may touch. Reading a historian and proposing a tuning is one level of access. Writing a setpoint, downloading a configuration or changing an interlock is another, and it passes through the plant's management of change with a named approver. Supervision also fixes the order of work: a baseline before a change, one change at a time and a comparison that shows what moved.

Creativity. An agent answers the question it is given. Choosing the question remains the engineer's task: which loop is worth improving, whether a disturbance calls for feedforward, a cascade or a

change to the process, which step test the operation can tolerate, and which comparison would show that a predictive controller earns its place over a retuned PID loop. The experiment that separates two explanations is designed by someone who can imagine both.

Process insight. A model contains what its author put in. The equipment does more. A valve sticks, the dead time of an analyzer grows when its sample flow drops, a thermowell lags, a heat exchanger fouls and an operator intervenes at the same hour each night. The engineer who knows the unit recognizes a result that is too good, asks what the model left out and knows which measurement to distrust first.

Accountability. Responsibility does not transfer to a tool. Functional safety standards assign the activities of the safety lifecycle to persons and organizations of stated competence, and they require an assessment that includes someone independent of the design team [19]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

These five are the professional form of the habit introduced in Section 1.7: check before trusting, and know enough to check.

14.4.1 How This Volume Was Made

Each of the twelve apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace each number to a run. Four of the twelve referee reports call themselves self-reviews. The other eight describe a reviewer's stance or say only that they were written against the draft and then applied, and none names a reviewer independent of the writing. Six carry a second pass on predictive-control material added later. Eight of the papers were measured again in a consistency round and an app-fix round, and the other four were written after those rounds. The defect logs hold 131 entries with a severity. Of these, 42 are wrong output or behavior of an app or its model, 70 documentation or briefing content, 17 cosmetic or display, and 2 something else, a completion rule and a check. Some round notes list fixes without a severity, so the count is a floor. For this book one agent per chapter reran the numbers its exercises use and checked the app's briefings, helpers and completion steps. The editor set the scope of each round and what each agent could change. Four cases show what went wrong and what caught it. In each a fast part met a slow clock or a slow layer misjudged a fast one.

A heater that never fired. The air cavern's supplemental heater had a lag of about six minutes, and the engine integrated it with a fourth-order Runge-Kutta step of thirty minutes. The step was 5.21 times the time constant, against a stability limit of 2.785, so the heater's state diverged on every step and a clip returned it to zero. Over the 150-day challenge the command reached 120 MW and integrated to 71 048 MWh while the heat delivered was exactly zero. The trend plotted the command, so the strip showed a working heater, and only a badge on the scene read 0 MW. The agent writing the paper found it, and a consistency round found the expander answering in about 1.8 h where its label said 12 min for the same reason. The app-fix round replaced the lags with their exact solution over the step, the challenge then burned 26.0 GWh of heat, and every run that moved a machine or fired the heater moved.

A program that counted its offsets twice. The feeder's central program fitted its voltage estimate with the inverter offsets in force, predicted the effect of new offsets on top of that estimate, and then replaced the old offsets with the new ones, so the old ones counted twice. On a cloudy day the summed offset alternated six times between 10:00 and 14:00. The stage that added the predictive controller beside the program found it, and a regression test now fails on the old code with six alternations and passes with none. The same round found that with the curves off the inverters run at

unity power factor while the program planned offsets for them. On the challenge as shipped it asked for up to 2.43 Mvar over 10 500 s, and the inverters delivered none. The first fix moved the program's case 9 from 27.3 to 21.2 min out of band, and the second moved the challenge as shipped from 143.7 to 30.1 min. The slow layer had modeled the fast devices doing what they could not.

A law that always matched the plant. The turbine's torque law below rated, $K\omega^2$, took its constant from the run's own air density and rotor radius. At 6 m/s the loop still settled at a tip-speed ratio of 7.53 to 7.54 for densities of 1.0 and 1.4 and radii of 60 and 70 m, so the Model dialog's warning that the law is right only if K is right could not be shown. A truth audit found it. The law now takes the declared design, a density of 1.0 gives a ratio of 7.02, a 70 m radius gives 8.87 and 3.9 % less power, and no shipped run moved. A test fails on the old code. Nine entries in six of the twelve logs record a controller or an operator aid that read the plant's truth or the run's hidden settings in place of a measurement or a declared value.

A re-plan that never saw the price. The nuclear plant's paper, and five texts of its app, said that re-planning every fifteen minutes puts the metered price into the objective, so that on the spike day the optimizer discharges at its limit once the meter shows the spike. The chapter agent compared the logged allocations of the spike day with those of the ordinary day and found them identical at every sample. The program builds its block prices from the forecast alone. The chapter now says that feedback corrects the state and not the price, and the five texts were logged with a recommendation to change the texts rather than the controller, so that the paper's numbers for the case stand. A repair round has since changed them.

The book round. The twelve chapter records list 156 rows of numeric verification, each a stated value or group of values set beside its reproduction. Seven records report no drift. The other five name a printed value that a run does not reproduce exactly, and all stand as printed. Every record tested its chapter's causal claims by run. Three explanations were contradicted and corrected without a new number. One is the nuclear case above. The flow battery's idle-day loss was said to be curtailment and nothing else, when curtailment is \$766 of the \$1,468 and the idle plant also buys 15.42 MWh for its pumps, auxiliaries and inverter. The cavern's heater was said to fire at its full 120 MW through the first days of the drought, when its command alternated on successive hourly scans, a limit cycle the repair round removed. Two summaries were aligned with their own tables, a feeder program said to win a challenge in which neither run holds the band and an economic schedule said to win where firm dispatch earns the most. The turbine's energy cost, one value in the paper's abstract, is the ten-seed range of 2.6 to 5.2 % in the chapter. Nine of the twelve chapter passes logged an app defect, most of them in briefings and helpers, such as a challenge hint that said tuning alone passes when it nets \$5,467 against a \$6,000 target. No chapter was found naming a loop's action against the convention of Section 1.6.2. Chapter 1 was written from the converted chapters, and seven of its statements, in its rows for five chapters, were corrected from the chapter reports, among them the turbine's fastest clock and the fluid in the trough field. One finding of the round is still open, and Section 14.5 describes it.

Every error above was caught by a check that could fail: a delivered heat compared with its command, a decision count on a cloudy day, a hidden value moved while the measurements were held, a logged plan compared with the plan of another day. None was caught by rereading the text that contained it. Four working rules follow. Integrate every fast lag at the step that will run it, and check each loop's gain at the scan it is given. Check that what a slow layer asks of a fast device the device can deliver. Give a controller only what the plant measures and the design declares. Test the event a plan claims to answer by comparing its moves with and without the event.

14.5 Gaps in the Control of Energy Conversion and Storage

Five gaps belong to this volume's topic, and the gaps common to agent-era practice follow them in Section 14.5.1. In the editor's view none is closed. A sentence on difficulty or progress that carries no citation is the editor's view. Figure 14.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the plants in the picture that opens this chapter.

1. *A fast loop on a slow clock.* In Chapter 11 the heater loop runs on the dispatch layer's hourly scan. The heater's six-minute lag is over within the scan, so the loop sees a static process. At the shipped gain of 1.6 MW/K its discrete loop gain was above the stable limit, and the heater command flipped between 0 and 120 MW every other scan. No tile and no number of the paper showed it. The chapter agent found it in the hourly rows of the exported record. A repair round set the gain to 0.2 MW/K, which settles, and every heater-armed number of the paper and the chapter was measured again. Supervisory and lower-level loops run continuously with cycle times from milliseconds to minutes [20], and a study of hybrid nuclear and renewable plants held that techniques are needed to coordinate linked but independent systems with differing dynamic timescales [21]. The gap is hard because a supervisory layer chooses its scan for its own decisions, and a loop placed on it inherits that period without a check of its gain. What stays open is the tooling. Tools that check every loop's gain at the scan a supervisory layer imposes on it are not routine.

2. *Inertia stated with its meter.* The microgrid's first swing read 2.48 at the shipped 20 ms step and 2.75 Hz/s at 2.5 ms, and the virtual inertia below which the solar trips moved from between 1.25 and 1.3 s to between 1.6 and 1.7 s once the meter was resolved (Section 7.6). European guidance states the rate of change of frequency that generating units must ride through separately for each averaging window, from 2 Hz/s over 500 ms to 1.25 Hz/s over two seconds [3], and a report on frequency measurement from the same body notes that the shorter the time for a frequency calculation, the noisier the measurement [22]. Reduced inertia gives a larger rate of change of frequency [4]. The gap is hard because an inverter's first swing is over within the window that a relay averages, so the same inertia passes or fails with the meter. A method that sizes virtual inertia against the protection's own meter is missing.

3. *Estimating the limit that has no instrument.* The margin of Chapter 9 covered a table 20 % optimistic and no more, and an aged pack plated as much as the reckless charge. In Chapter 8 the program's margins on the worst customer were a median of 6.4 times the error and too wide to steer by on a cloudy day. The fuel cell of Chapter 3 runs on an estimate of the gas arriving at its anode, and its governor matched the predictive controller only when given the same estimate, 0.805 against 0.807 (Table 3.8). Estimators that track an anode potential, an arriving fuel or a far-end voltage and state their own uncertainty online are not standard. The Open challenge box after this list takes the gap further.

4. *Forecasts with a stated uncertainty.* The plan of Chapter 13 priced every block from a forecast that missed a spike of 5.5 times, and no re-plan could recover a storage decision taken six hours earlier. A roadmap for grid-forming inverters lists the reserve requirements that keep frequency within bounds as an open question, and notes that capacity held back as headroom can be an opportunity cost [4]. A study of solar plants asks for forecasts from data at 5 to 15 min, because hourly data miss short cloud transients [14]. The gap is hard because a reserve against the tail costs money on every ordinary day and pays on a few. A dispatch layer that carries the forecast's error and prices a reserve against it is not standard.

5. *Degradation that no loop reports.* The flow battery of Chapter 10 lost 54 MWh of working band to an imbalance while its power and flow loops held their setpoints. In Chapter 12 a membrane defect raised its first alarm 4.6 hours after it appeared, because the crossover is a rising fraction of

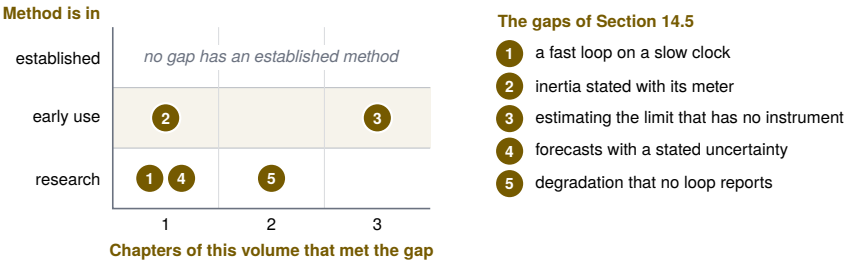

Figure 14.4 The five gaps of this section. Across is the number of chapters that each item names as evidence: one for gap 1 (Chapter 11), one for gap 2 (7), three for gap 3 (3, 8 and 9), one for gap 4 (13) and two for gap 5 (10 and 12). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that no gap has a method in established use.

a falling flow. In simulations of a wave energy converter, heavy biofouling cut the take-off power by 15 % in one sea and retuning the control recovered only about 6 % of the loss, and the report calls for adaptive control and models that account for environmental degradation [23]. Flow batteries promise lives of 20 years or more, yet few United States projects have stayed online for several years to prove it at scale [24]. The gap is hard because lost capacity drifts slowly, the loops absorb it, and its trend looks like a change of load. Monitoring that normalizes a reading by the load that the dispatch sets, and reports lost capacity as a measured quantity, is not routine.

⚠ **OPEN CHALLENGE** **Estimating the Limit That Has No Instrument**

Gap 3 is the gap most chapters met. A gap assessment of fast charging held that methods are needed to detect and prevent lithium plating, and that voltage and temperature sensors on every cell would identify the weaker and hotter cells [16]. Trough plants with unreliable flow meters estimate the flow from pump current, drive speed and pressure drop, with accuracy typically not nearly as good as a working meter's [14]. An online gas chromatograph can take minutes per analysis [15]. The gap is hard because the estimate is checked only when the limit is crossed, and its model ages with the plant. Progress would be an estimator that reports its own uncertainty online and a margin sized from that report.

14.5.1 Gaps Common to Agent-Era Practice

The final chapter of each volume lists gaps specific to its topic. Seven more are common to every topic in the series. In the editor's view none of them is closed.

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [25]. In the editor's view, code written by an agent has no comparable body of method yet. Practice needs a routine path from an agent's draft to code a plant can trust: independent derivation, tests of what the code reads and clamps, and trials with hardware in the loop.

2. *Trustworthy models from sparse or poor plant data.* Plant records are short, taken in closed loop, noisy and marked by unmeasured disturbances. A fitting routine returns a model from any record. What is missing is a dependable statement, delivered with the model, of where it can be trusted, and a method that says so when the data cannot identify a parameter.

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [19]. In

the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

4. *Security of agent access to control systems.* An agent with credentials to a control system is a new path into it. Industrial security standards state the requirements a control system must meet at each security level [26]. An agent that reads documents, messages or web pages can also be steered by text placed there. In the editor's view, least privilege, read-only access by default and human approval of every write are a starting point and not a solution.

5. *Benchmarks and reproducibility.* An agent asked twice may deliver two different controllers, and a result reported without its plant, seed and scoring cannot be checked. Control needs shared benchmark plants with fixed disturbances and published measures, on which agent-built and conventional designs are compared like for like. The apps of this series, with a default seed and an exported record of every run, are a start.

6. *Transfer from simulation to hardware.* A controller developed on a simulator inherits what the simulator omits. The robotics literature documents the gap between simulated and real performance and the methods used to narrow it, such as randomizing the simulated plant and adapting on the real one [27]. Process plants add slow dynamics, rare upsets and little freedom to experiment, so each transfer still needs a staged plan from the simulator to a bench rig to the unit.

7. *Keeping engineers' judgment sharp when agents do the routine work.* Automation removes the routine work through which skill is built and leaves the person the hard cases, an irony described for process operators in 1983 [28]. The same now applies to engineers. An engineer who has never tuned a loop by hand or closed a balance on paper is poorly placed to review an agent that has. Deliberate practice, of the kind the exercises in this book provide, has to be designed into education and into work.

KEY POINTS

- A fast loop that works is necessary. It says little about whether the plant is worth running, which the slow layer decides.
- A store lets the layers run apart, and a dispatch comparison without the end state of its store is not a comparison.
- A forecast pays where feedback cannot see the cause, and is worth little where the price follows the load it measures.
- The limits that ended a run were rarely the ones on a gauge.
- Tuning repairs the fast layer, and structure repairs the slow one.

References

[1] J. Matevosyan, B. Badrzadeh, T. Prevost, E. Quitmann, D. Ramasubramanian, H. Urdal, S. Achilles, J. MacDowell, S. H. Huang, V. Vital, J. O'Sullivan, and R. Quint, "Grid-forming inverters: Are they the key for high renewable penetration?" *IEEE Power and Energy Magazine*, vol. 17, no. 6, pp. 89–98, 2019.

[2] D. B. Rathnayake, M. Akrami, C. Phurailatpam, S. P. Me, S. Hadavi, G. Jayasinghe, S. Zabihi, and B. Bahrani, "Grid forming inverter modeling, control, and applications," *IEEE Access*, vol. 9, pp. 114 781–114 807, 2021.

[3] ENTSO-E, "Rate of change of frequency (RoCoF) withstand capability: ENTSO-E guidance document for national implementation for network codes on grid connection," European Network of Transmission System Operators for Electricity, Brussels, Guidance document, Jan. 2018. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/Network%20codes%20documents/NC%20RfG/IGD_RoCoF_withstand_capability_final.pdf

[4] Y. Lin, J. H. Eto, B. B. Johnson, J. D. Flicker, R. H. Lasseter, H. N. Villegas Pico, G.-S. Seo, B. J. Pierre, and A. Ellis, "Research roadmap on grid-forming inverters," National Renewable Energy Laboratory, Golden, CO, Technical Report NREL/TP-5D00-73476, 2020.

[5] A. Scholbrock, P. Fleming, D. Schlipf, A. Wright, K. Johnson, and N. Wang, "Lidar-enhanced wind turbine control: Past, present, and future," in *Proceedings of the 2016 American Control Conference (ACC)*, 2016, pp. 1399–1406.

[6] J. V. Ringwood, "Wave energy control: Status and perspectives 2020," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 12 271–12 282, 2020.

[7] Y. Shi, B. Xu, D. Wang, and B. Zhang, "Using battery storage for peak shaving and frequency regulation: Joint optimization for superlinear gains," *IEEE Transactions on Power Systems*, vol. 33, no. 3, pp. 2882–2894, 2018.

[8] Y. Liu, Y. Zhu, and Y. Cui, "Challenges and opportunities towards fast-charging battery materials," *Nature Energy*, vol. 4, no. 7, pp. 540–550, 2019.

[9] H. E. Perez, X. Hu, S. Dey, and S. J. Moura, "Optimal charging of li-ion batteries with coupled electro-thermal-aging dynamics," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 9, pp. 7761–7770, 2017.

[10] A. Buttler and H. Spliethoff, "Current status of water electrolysis for energy storage, grid balancing and sector coupling via power-to-gas and power-to-liquids: A review," *Renewable and Sustainable Energy Reviews*, vol. 82, pp. 2440–2454, 2018.

[11] S. J. Smith, A. Hubbard, A. Newkirk *et al.*, "United states data center energy usage report: 2025 update," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-2001758, 2026.

[12] Y. Li, D. Wang, S. Ghose, J. Liu, S. Govindan, S. James, E. Peterson, J. Siegler, R. Ausavarungnirun, and O. Mutlu, "Techniques for efficiently handling power surges in fuel cell powered data centers: Modeling, analysis, results," arXiv preprint arXiv:1806.04506, 2018. [Online]. Available: https://arxiv.org/abs/1806.04506

[13] N. Katsampes, D. Montgomery, G. Arney, and D. B. Olsen, "Hydrogen-natural gas fuel blending in a "rich burn" engine with 3-way catalyst," *Frontiers in Fuels*, vol. 2, p. 1416716, 2024.

[14] M. Mehos, H. Price, R. Cable, D. Kearney, B. Kelly, G. Kolb, and F. Morse, "Concentrating solar power best practices study," National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5500-75763, 2020.

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

[16] S. Ahmed, I. Bloom, A. N. Jansen, T. Tanim, E. J. Dufek, A. Pesaran, A. Burnham, R. B. Carlson, F. Dias, K. Hardy, M. Keyser, C. Kreuzer, A. Markel, A. Meintz, C. Michelbacher, M. Mohanpurkar, P. A. Nelson, D. C. Robertson, D. Scoffield, M. Shirk, T. Stephens, R. Vijayagopal, and J. Zhang, "Enabling fast charging – a battery technology gap assessment," *Journal of Power Sources*, vol. 367, pp. 250–262, 2017.

[17] U.S. Department of Energy, "Department of energy hydrogen program plan," U.S. Department of Energy, Washington, DC, Report, 2020.

[18] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, W. X. Zhao, Z. Wei, and J. Wen, "A survey on large language model based autonomous agents," *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024.

[19] International Electrotechnical Commission, "Functional safety – Safety instrumented systems for the process industry sector – Part 1: Framework, definitions, system, hardware and application programming requirements," IEC 61511-1:2016, Edition 2.0, 2016.

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

[21] M. F. Ruth, O. R. Zinaman, M. Antkowiak, R. D. Boardman, R. S. Cherry, and M. D. Bazilian, "Nuclear-renewable hybrid energy systems: Opportunities, interconnections, and needs," *Energy Conversion and Management*, vol. 78, pp. 684–694, 2014.

[22] ENTSO-E RG-CE System Protection & Dynamics Sub Group, "Frequency measurement requirements and usage," European Network of Transmission System Operators for Electricity, Brussels, Report Final Version 7, 2018. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/SOC%20documents/Regional_Groups_Continental_Europe/2018/TF_Freq_Meas_v7.pdf

[23] T. Brekken, "Applied research and development to support open-water testing at PacWave, task 7: WEC system control for marine growth and electrical faults," Oregon State University, Corvallis, OR, Final Technical Report DOE-OSU-EE0009969-7, 2025.

[24] A. Hollas, A. Tuan, V. Viswanathan, and I. Ragazzi, "Adoption readiness level assessment of redox flow batteries," Pacific Northwest National Laboratory, Richland, WA, Report PNNL-36780, 2024.

[25] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.

[26] International Electrotechnical Commission, “Industrial communication networks – Network and system security – Part 3-3: System security requirements and security levels,” IEC 62443-3-3:2013, Edition 1.0, 2013.

[27] W. Zhao, J. P. Queralta, and T. Westerlund, “Sim-to-real transfer in deep reinforcement learning for robotics: A survey,” in *Proc. 2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2020, pp. 737–744.

[28] L. Bainbridge, “Ironies of automation,” *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

14.6 Exercises

14.1 Discussion · 8 min

An agent reviewing Chapter 10 proposes to retire the dispatch layer and run the battery on the net-load controller alone, “because on the clean day it books the same margin and regulates twice as well”. Using Table 10.8 and Table 13.5, say what the proposal gets right, what it misses, and what you would keep.

14.2 Discussion · 8 min

An agent reports that the one-hour energy dispatch of Chapter 7 “saves 78.83 L of diesel against the rules on the added case”. Using Table 7.8 and the end-of-day tank in Table 13.5, explain what the claim leaves out and what a fair comparison of two dispatch policies on a plant with a store must report.

14.3 Discussion · 8 min

On the fast charger of Chapter 9 an agent proposes two simplifications: remove the voltage and power limits from the selector, “since they hold the current for at most 0.32 min”, and keep the 15 mV anode margin, “since it plated nothing in scenario 4”. Using Tables 9.5, 9.7 and 9.8, say what is wrong with each and what the charger’s design should state.

14.4 Discussion · 8 min

An agent proposes to close the air cavern’s heater defect of Section 14.5 by halving the shipped gain of 1.6 MW/K. Using Section 1.4.2, explain the cause, judge the proposal and say what else must happen first.

14.5 Discussion · 7 min

Chapter 8 closes with “the slow device should set the level and the fast ones should trim the spread”. Using Tables 8.4, 8.5 and 8.6, explain the sentence, give a result where the layers fought, and say what the central controllers add and cannot override.

14.6 Discussion · 9 min

An agent reports that its new dispatch layer for one plant of this volume “passes all scenarios and raises the margin 4 %”. Choose the plant and, using Section 14.4.1, name four checks to run before accepting it and the result of each that would reject it.

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) 200 MJ and 0.30 Hz/s. (b) 2.40 s.

1.3 (b). Inertia sets the rate of the first fall, and the droops set where it stops.

1.5 (a) 0.040 pu, 0.333 pu. (b) 0.0303 pu, 0.081 pu.

1.7 (a) $\psi = 1.211$, 7.89 MW of fuel, 3.63 MW of heat. (b) 5.3 h, then 1.13 MW is dumped.

1.9 (a). A plan that cannot see the dear hours sees no reason to store energy in the cheap ones.

Chapter 2. Gas Engine Cogeneration on Two Clocks

2.1 (a). Heat reaches the water in proportion to how much warmer the metal is, and 8 MJ/K of jacket and 2 MJ/K of exchanger metal must warm first.

2.3 (c). Rejection begins only when warm water reaches the buffer bottom, at 399 s, after the buffer has filled from 43 to 81 %.

2.5 (a) 0.227 MWh. (b) 252 s. (c) 43 %, the value of Table 2.1. Without the clip the two cool nodes would count below zero and give 39.2 %.

2.7 (a) 144 s and 31 s. (b) 0.073 MW at 68 s, 0.127 MW below the target. (c) The full 0.6 MW lands on the grid before any meter reports it, and a delay of a second or two is small beside a feedback recovery of more than two minutes.

2.9 (a) 52.0 % (8.0 points short), 570 s, 408 s, not met on the UPS reserve and the supply band. (b) 79.2 %, 570 s, 0 s, not met on the supply band. (c) 79.2 %, 0 s, 0 s, met. The jacket loop guards the engines' capacity and so the UPS, the supply loop guards the band.

2.11 (a) 121 °C, close to the model's 120 °C minimum. Below it the exhaust's water vapor condenses in the piping. (b) A real block has hot spots where flow is low, and the single jacket node, whose temperature is the outlet, cannot show them.

Chapter 3. Fuel Cell Current That Waits for Its Fuel

3.1 (d). The request is power divided by voltage, so a falling voltage raises the current, which draws the anode down further.

3.3 (b). The shortfall is missing capacity, and only less load removes it.

3.5 (a) 16 s. (b) $x \simeq 3.26$, 16.3 s for the stages and about 17.3 s with the valve. (c) 56 s and about 61 s, against the 17.3 s the estimator expects.

3.7 (a) 0.16, 0.993, 2.42 kWh. (b) 0.62, 0.800, 0.65 kWh. (c) The naive run drives the cell to the inverter's floor and then heats the stack, so the foldback and the thermal limiter hold its current for longer than the governor holds the governed one.

3.9 (a) 1.048, 53.1 % lowest and 60.3 % at the end, not met. (b) 0.844, 46.6 % lowest and at the end, not met on the UPS at the end. (c) 0.833, 59.4 % and 63.2 %, met.

3.11 (a) 47.6 % and 62.8 %, within rounding of the chapter's 47.55 and 62.78 %. (b) The thermal lessons come later or not at all; the fuel lessons and the challenge verdicts stay.

Chapter 4. Wind Turbine Torque-to-Pitch Handover

4.1 (d). With different setpoints the two integral loops cannot both be satisfied, and the pitch loop can begin regulating before the region index declares Region 3.

4.3 (b). With $K = \rho\pi R^5 C_{P,max}/(2\lambda^{*3})$ the rotor is in equilibrium exactly at λ^*, using only the measured speed.

4.5 (a) $K_{c,\beta} = 3.769\,\text{s}$, $\tau_{I,\beta} = 2.333\,\text{s}$. (b) 33.6 and 175.6, $GK = 5.41$, 32.4. (c) 1.37 rad/s (0.218 Hz) and $\zeta = 1.60$.

4.7 (a) 0.746 MW. (b) $\lambda = 5.50$, $C_P = 0.418$, 0.649 MW. (c) 13.0 %.

4.9 (a) 100.29 %. (b) 105.59 %. (c) 102.39 %. The preview supplies 2.1 of the 5.3 points, about 40 %.

4.11 (a) It sits in the wake of the rotor and measures one point rather than the rotor-effective average over the disc. (b) K should fall to 0.871 of its commissioned value, and the old K holds the rotor below its new optimal ratio.

Chapter 5. Heaving Buoy on a Fading Wave Forecast

5.1 (c). The optimum velocity is $\hat{F}_e(\omega)/[2B(\omega)]$, and a real, even operator has an impulse response that is even in time.

5.3 (a). The absorbed energy is rewritten as the excitation's work less the radiated energy and the rise in stored energy, and passivity makes every quadratic weight non-negative.

5.5 (a) 4.48 kN s/m, 24.7 t, –209.4 kN s/m. (b) 209.4 kN s/m, –146.2 kN/m. (c) 0.042.

5.7 (a) 0.47 kWh, 13 %, 44 kN, 1.31 m. (b) 215 kN s/m, 0.77 kWh, 21 %, 123 kN, 0.79 m. (c) The helper tunes at the probe's estimate of 9.1 s, not at the true 9 s.

5.9 (a) 8.91 kWh, 23 contacts, not met. (b) 10.89 kWh, 0 contacts, met, 1.19 kWh. (c) 8.02 kWh, 0 contacts, not met. The lock engages in the 2.5 m sea.

5.11 (a) 1.97 s and 1.84 s. (b) K_{pto} moves toward zero. The ranking rests on the model's own consistency, not on these values.

Chapter 6. Parabolic Trough and Its Moving Setpoint

6.1 (b). The loss of Equation (6.6) holds no irradiance, so a cooler field keeps more of a small harvest.

6.3 (d). Removing the preview changes almost nothing, and a feedforward loop with the same measurement and model does as well.

6.5 (a) 4.5 min and 27 min. (b) 31.4 MJ/K and 146 MJ/K. (c) –0.384 °C per kg/s, and four times that, about –1.54 °C per kg/s, at a quarter of the heat and flow.

6.7 (a) 13.8 MWh, with the level below 58 %. (b) 3.75 MW lost of 11.9 MW absorbed, 31.5 %. (c) The protection threw it away, nearly all of it after 12:36.

6.9 (a) 46.8, 883 and 47.1 °C min (tiles 0.781, 14.714 and 0.785). (b) 0.0, 17.9 and 0.0 MWh. (c) 95 % lower under the MPC, about 0.6 % higher without the preview.

6.11 One row stands for the field, so the hottest loop and the flow distribution are missing. Both numbers are biased low. The engineer needs loop-by-loop flows and outlet temperatures.

Chapter 7. Forming the Grid When a Microgrid Islands

7.1 (a). A current source needs a voltage to push into, and with the utility gone nothing on the bus makes one.

7.3 (c). A voltage source on its limit holds its frequency, so without the power-limit droop the relays never hear the overload.

7.5 (a) 42, 33.3, 1.26. (b) 68.7, 2.06. (c) 0.293 Hz and 0.219 Hz (measured 0.294 and 0.219).

7.7 (a) 59.516 Hz, 2.48 Hz/s, no shed. (b) 58.226 Hz, 5.83 Hz/s, three stage trips: the solar tripped, and the battery had to replace it too. (c) 1.3 s (2.97 Hz/s), while 1.2 s reads 5.52 Hz/s.

7.9 (a) No critical load lost, 22 kWh shed in one stage trip, nadir 59.182 Hz. (b) 17.2 kWh lost, BLACKOUT from the row at 38.5 s. (c) 17.6 kWh lost with no stage trip: the frequency held at about 59.51 Hz, above stage 1.

7.11 (a) The ride-through at 1.3 s holds only on the 20 ms meter. Resolved, the cliff moves to 1.6 to 1.7 s, while 2 s reads 2.48 to 2.75 Hz/s, under 3.0 Hz/s, at every step. (b) Two of: compressor inrush, the inverters' first-cycle current control, relay filtering and breaker time, feeder impedance and faults inside the campus.

Chapter 8. Tap Changer and Inverters on a Solar Feeder

8.1 (a). The correction $(R_{set}P + X_{set}Q)/V_{bus}$ is positive while $X_{set}/R_{set} = 15$ exceeds $-P/Q = 1.68$, so the bus sits 0.0083 pu high.

8.3 (c). Two scales are fitted to two substation measurements, and a stale meter adds 0.01 pu to the margin.

8.5 (a) 0.008 27 pu, positive above $X_{set}/R_{set} = 1.68$. (b) –0.0025 pu. (c) 0.0593 and 0.0598.

8.7 (a) 7, 0.972 pu. (b) 5, 0.970 pu. (c) 5, 0.972 pu.

8.9 (a) In band, 15 operations. (b) 21.2 min out, 17 operations. (c) 15.9 min, 11 operations, 25 % and 100 %.

8.11 (a) No correction while the flow is reversed, so the bus is no longer held 0.0083 pu high. On case 5 it makes 8 operations with HUNTING against 2. (b) Curves at 7.33, $R_{set} = X_{set} = 0.06$ pu ($X/R = 1$), setpoint 1.00 pu, 2.5 % band, 60 s.

Chapter 9. EV Fast Charge Against an Unmeasured Limit

9.1 (d). The anode candidate is an inverse map with no measured error, and its small residual is filter lag, not an offset.

9.3 (b). The kinetic resistance of Equation (9.5) falls with temperature, so a hotter anode allows more current.

9.5 (a) 3.41 and 141 A. (b) 3.02. (c) 5.88 kW against 4.51 kW of chiller capacity.

9.7 (a) 16.8 min, 3.383 Ah, 0.1037 %, 51.8 %. (b) 21.3 min, 0.000 Ah, 0.0780 %, 39.0 %. (c) 4.5 min, a quarter of the fade, about 25.6 Ah.

9.9 (a) 21.6 min, 14.196 Ah, 28.0 °C, 62.6 %: fails. (b) 30.7 min, 0.024 Ah, 26.5 °C, 18.9 %: passes. (c) 0.169 Ah plated, so it fails, because the flow carries away heat the cold anode needs.

9.11 (a) The anode limit falls with lithiation, so the last twenty points take longer than the first seventy. (b) About 17.4 min on this pack.

Chapter 10. Vanadium Flow Battery Inventory Dispatch

10.1 (c). When the net-load peak and the price peak coincide, the action feedback takes for the net load is also the right action for the price.

10.3 (a). It is cascade anti-windup: the outer loop's bounds follow what the inner plant can actually do.

10.5 (a) 84.48 MA. (b) 0.344 %/min discharging and 0.234 %/min charging. (c) 3.87 h and 5.71 h. Charging runs at a higher cell voltage, so each ampere carries more power.

10.7 (a) 586 MW, 2.74 MW/min, 64 MWh, −$1 468. (b) $766 of curtailment, about 52 %. The rest is the 15.42 MWh the idle plant buys. (c) 12:28 and 19:27.

10.9 (a) Plan, tank and margin: 18.4, 11.7 %, $34 740 at weight 12; 11.6, 10.1 %, $34 998 at 0; 38.8, 26.2 %, $32 088 at 60. (b) $2 910. (c) A miss of a few tenths squared is small against tracking and ramp terms summed over twenty-four blocks.

10.11 The feeder, solar and price are smooth synthetic profiles, and the scarcity event was placed at 15:00 by design, away from the evening peak. Real hourly net-load and price records with their scarcity events are needed.

Chapter 11. Salt-Cavern Air Storage Inventory Policy

11.1 (b). In Equation (11.11) the swallowing capacity is proportional to the pressure, and the specific work falls with the pressure ratio.

11.3 (d). The clamp manufactures gas, so every threshold meets the terminal target and the lowest one wins.

11.5 (a) 40.1 d and 521 d. (b) 4.58 GWh, which lasts 13.8 h at the rated draw, against 81.9 h of gas. (c) 101 d. The expander's draw, not the standing loss, empties the store.

11.7 (a) 205 MW, the calculated 204.7 MW rounded. (b) 199.2 MW; 11 h, 50.5 bar, 399 K. (c) 17 h, 15.0 % against a 13.9 % floor. The taper of Equation (11.22) and the expander's minimum load stop it.

11.9 (a) 1.11 GWh, $14.36 M, a minimum of 0.195 and 25.08 d short. All four clauses fail. (b) 0.00 GWh, $15.28 M, 0.249, no hour short, 23.08 GWh of fuel heat. It passes. (c) 0.47 GWh, $14.55 M, 15.33 d short. It fails.

11.11 (a) It reduces the default without removing it: 1.15 GWh at a 25 % peak to 0.13 GWh at 85 %, for $1.70 M. (b) The capability in megawatts, counting the heat that will be left when the drought arrives, held within the creep limits on pressure and slew.

Chapter 12. Solar Hydrogen Plant Safety Interlocks

12.1 (b). In Equation (12.7) the reading is a fixed crossover divided by a production that is proportional to the current.

12.3 (d). The electrode pressure clamp of Equation (12.8) holds the Nernst term to tens of millivolts, and 6.0 % of the 14.3 % is cell voltage.

12.5 (a) 8402 kg/h consumed, 7970 kg/h net. (b) 10.2 h from full, 5.08 h from 55 %. (c) 4.17 h.

12.7 (a) 1.296 A/cm^2 and 61.4 ppm, 0.288 A/cm^2 and 278.2 ppm, 0.120 A/cm^2 and 666.7 ppm. (b) 79.6, 80.1 and 80.0, all near 80 ppm A/cm^2. (c) 2.43 MW, the power at the minimum current.

12.9 (a) 57.60 MWh, −$44,664, 8 trips, the first at about 13:50 on electrolyzer temperature. (b) 0.04 MWh, 10.7 MWh, $5,467, 78.6 °C: no trip, but below the $6,000 target. (c) 0.07 MWh, 10.4 MWh, $6,959, 67.6 °C, and 20012 − 8653 − 4400 = 6959.

12.11 Crossover is a rising fraction of a falling hydrogen flow, so low load is the unsafe end. Below the minimum the stack should stop, not idle.

Chapter 13. Supervisory Economic Dispatch of Nuclear Heat

13.1 (a). The trim moves heat into the store, and more charge means less pressure, so a high pressure must raise the output.

13.3 (c). Exact convergence buys accuracy against an approximate model, while the move term holds the plan near the previous one.

13.5 (a) 33.3 bar/h and 0.33 bar. (b) 4.78 bar/min and 5.47 bar/min. (c) About 3.3 bar, which puts the peak near the measured 129.22 bar.

13.7 (a) $470,311, $516,300 and $529,303, with 72.07, 100.55 and 93.24 t. (b) 78.0 %. (c) 579 MWh worth $5,693, which cuts the advantage to about $7,300 and raises the share to 88 %.

13.9 (a) 51.4 MWh, $501,272, 0.0 % and 126.42 bar. (b) Nothing dumped, $482,537, 91.09 % and 120.80 bar. (c) Feasibility rather than money, with about 765 MWh unspent.

13.11 A reactor that cannot be turned down needs a sink for all of its heat when the turbine is lost. The model leaves out the shaft, inertia, frequency, governor and up and down times, so the outage is only a heat-balance event.

Chapter 14. Summary and Outlook

14.1 A good answer grants the clean-day facts: $35,083 against $35,067, a worst ramp of 2.61 against 5.55 MW/min, and no time off target against 0.61 h. It names the reason, that the highest net load and the highest price fall in the same hour on an ordinary day. It then points to the scarcity day, −$2,242 under feedback against $14,599, because feedback cannot act on a price it does not measure. It keeps both layers, the plan for value and the trim for regulation, as the app is wired. It adds that a plan is only as good as its forecast, since the nuclear plant's optimizer did not change its dispatch when a spike its forecast missed appeared on the meter.

14.3 A good answer notes that a limit that rarely binds on one pack and one charge still binds elsewhere: scenario 10's tail ends on the voltage limit and scenario 12 is built so that every limit binds, including the voltage from 33.98 min under the selector. Identifying the load-bearing limit is the lesson, not a licence to delete the others, which are protection. The 15 mV margin covers a table about 20 % optimistic and no more: at 30 and 40 % optimistic it plates 0.187 and 1.628 Ah, and a pack whose resistances have risen 60 % plates 3.412 Ah. The design should state the model error its margin covers and how aging is tracked.

14.5 A good answer says the tap changer moves every node together and slowly, so it suits the level, while the inverter curves act at each node within seconds against the spread a cloud makes. The layers fought in case 5: a Volt-VAR slope of 22 against a reactance-weighted compensator turned 2 tap operations into 18. The suggested local settings pass the challenge on 6 of 6 seeds where the tap changer alone fails on all six. The MPC holds case 9, where the program spends 21.2 min out, and neither holds the shipped challenge (10.2 against 11.1 min for the program with a 2 h look-ahead and the MPC). The band, the timer and the clamps act below both.

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 heat instead of an electric compressor, so recovered engine heat can cool a building. *Chapter 2*.

back-calculation A way to stop windup: when a controller output is limited, its integral is pulled toward the output actually applied. *Chapter 2*.

battery management system The computer in a battery that tracks its charge, temperature and limits and tells the controller how much power it may take. *Chapter 1*.

Betz limit The largest share of the wind's energy that any rotor in open air can capture, about 59 percent. *Chapter 4*.

capacity commitment A promise, sold in megawatts, to deliver a stated power whenever the grid calls for it. *Chapter 11*.

capacity fade The loss of the charge a battery can hold as it ages with use. *Chapter 9*.

capture width ratio The energy a wave device captures divided by the wave energy arriving across its own width. *Chapter 5*.

cogeneration (CHP) Making electricity and useful heat from one fuel at once, also called combined heat and power. *Chapter 1*.

collective pitch Turning all the blades of a wind turbine together by the same angle. *Chapter 4*.

compressed air energy storage (CAES) Storing energy as air pumped into a cavern and burning fuel in that air later to drive a turbine. *Chapter 11*.

concentrating solar power Using mirrors to focus sunlight onto a receiver whose hot fluid then drives a steam turbine. *Chapter 6*.

conservation voltage reduction Holding the voltage of a feeder near the low end of its band so that customers use slightly less energy. *Chapter 8*.

constant-current, constant-voltage charging The usual way to charge a lithium battery: a fixed current until a voltage limit, then that voltage held while the current falls. *Chapter 9*.

controller action Whether a controller raises or lowers its output when its measurement rises, called reverse or direct acting. *Chapter 1*.

coulomb counting Estimating a battery's state of charge by adding up the current that flows in and out. *Chapter 9*.

crossover Leakage of a chemical through the membrane that should keep the two sides of a cell apart. *Chapter 10*.

current transformer An instrument transformer that brings a large line current down to a small one that a meter or relay can read. *Chapter 1*.

curtailment Throwing away wind or solar energy that the grid or the plant cannot use or store. *Chapter 10*.

deadband A band around a setpoint inside which a device makes no correction. *Chapter 1*.

droop control A proportional rule by which each unit changes its power in step with the frequency or voltage error, so several units share a load. *Chapter 1*.

duck curve The daily shape of net load on a grid with much solar, low at midday and rising steeply in the evening. *Chapter 10*.

dump radiator A cooler that throws away heat that can be neither used nor stored. *Chapter 1*.

economic dispatch Deciding how much each unit or store should produce or absorb so that the system runs at the least cost or greatest margin. *Chapter 1*.

electrical following Running a cogeneration plant to meet the electric demand and taking whatever heat comes with it. *Chapter 1*.

electrolyzer A device that splits water into hydrogen and oxygen with an electric current. *Chapter 12*.

end-stop The mechanical limit of travel of a wave buoy, which it strikes if it moves too far. *Chapter 5*.

energy storage A plant that takes energy in and gives it back later, such as a battery, a tank of hot salt or a cavern of air. *Chapter 1*.

expander The turbine of an air storage plant through which stored air expands to make power. *Chapter 11*.

extreme fast charging Recharging up to 80 percent of a battery's capacity in ten minutes or less. *Chapter 1*.

forecast skill How much better a forecast is than simply using the average; it falls as the forecast looks further ahead. *Chapter 1*.

fuel cell A device that turns the chemical energy of a fuel such as hydrogen directly into electricity, without burning it. *Chapter 3*.

fuel processor The equipment that turns natural gas

into a hydrogen-rich gas that a fuel cell can use. *Chapter 3*.

fuel utilization The share of the fuel fed to a fuel cell that it actually uses; too high a share starves the cell. *Chapter 3*.

gain scheduling Changing a controller's tuning as the operating point moves, following a measured variable. *Chapter 4*.

gas chromatograph An analyzer that separates a gas sample into its parts in a column and reports each, once per cycle. *Chapter 3*.

grid-following inverter An inverter that locks onto a voltage made by others and injects the power it is told to. *Chapter 7*.

grid-forming inverter An inverter that makes its own voltage and frequency, so it can hold up a grid on its own. *Chapter 7*.

headroom Spare power a unit can still add quickly when the load rises or another unit trips. *Chapter 7*.

heat-to-power ratio The heat a cogeneration unit recovers for each unit of electricity it makes. *Chapter 1*.

hierarchical control Control arranged in layers, each acting on a longer clock and setting the targets of the layer below. *Chapter 1*.

hunting A cycling in which control devices keep correcting one another without settling. *Chapter 8*.

IEEE 1547 The United States standard for connecting distributed generators and stores, such as solar inverters, to the utility grid. *Chapter 7*.

inertia constant The kinetic energy stored in rotating machines divided by their rating, in seconds of full output. *Chapter 1*.

instrument transformer A transformer that steps a line voltage or current down to a safe level for meters and relays. *Chapter 1*.

interlock A protective rule that stops or blocks an action whenever a safety condition is not met. *Chapter 6*.

inventory control Holding a stored amount, such as a tank level or a battery's charge, inside its limits. *Chapter 1*.

islanding A part of the grid running on its own, cut off from the main grid. *Chapter 7*.

jacket water The cooling water around an engine's cylinders, from which much of its heat is recovered. *Chapter 2*.

Kalman filter An estimator that blends a model's prediction with noisy measurements, each weighted by how much it is trusted. *Chapter 1*.

latching Holding a wave buoy still at the end of each stroke and releasing it later, so it keeps time with longer waves. *Chapter 5*.

LIDAR preview Measuring the wind ahead of a turbine with a laser so the controller can act before the gust arrives. *Chapter 4*.

line-drop compensation A setting that lets a tap changer estimate the voltage down the feeder from the current it measures. *Chapter 8*.

lithium plating Metallic lithium deposited on a battery's negative electrode during fast or cold charging, which ages the cell and can be unsafe. *Chapter 9*.

microgrid A small local power system with its own generators and stores that can run with the grid or alone. *Chapter 1*.

mirror defocus Turning collectors partly off the sun to cut the heat they deliver. *Chapter 6*.

mixed-integer linear programming An optimization with straight-line relations in which some decisions must be whole numbers, such as on or off. *Chapter 8*.

model predictive control (MPC) Control that predicts the plant over a horizon with a model, picks the best moves, applies the first and repeats. *Chapter 1*.

molten salt storage Storing heat in tanks of molten salt to be used hours later. *Chapter 13*.

moving horizon estimation An estimator that fits a model to a window of recent measurements to find the present state. *Chapter 1*.

net load The demand on a grid less the wind and solar it receives, which the other plants must supply. *Chapter 10*.

on-load tap changer A switch on a transformer that changes its turns ratio in steps, under load, to adjust the voltage. *Chapter 8*.

override control A protective loop that takes over an actuator from the normal loop when a limit is near. *Chapter 2*.

parabolic trough A long curved mirror that focuses sunlight on a pipe carrying a hot fluid. *Chapter 6*.

PEM electrolysis Splitting water in a cell with a proton exchange membrane, which follows a changing power within seconds. *Chapter 12*.

phase-locked loop A small control loop that tracks the angle and frequency of a measured voltage. *Chapter 7*.

phasor measurement unit An instrument that measures voltage and current with a time stamp from a satellite clock, so angles can be compared across a grid. *Chapter 7*.

pitch system The drives that turn each blade of a wind turbine about its long axis. *Chapter 4*.

point absorber A wave energy device much smaller than a wavelength that captures energy by moving with the waves. *Chapter 5*.

point of common coupling The point where a microgrid or a customer connects to the utility grid. *Chapter 7*.

power coefficient The share of the wind's energy that a turbine rotor captures. *Chapter 4*.

power converter Power electronics that change electricity between direct and alternating current or between voltages, controlled within milliseconds. *Chapter 1*.

power take-off The machine that turns the motion of a wave device into electricity. *Chapter 5*.

primary control The fastest layer of frequency control, a droop in each unit that stops a frequency fall within seconds. *Chapter 1*.

rate of change of frequency (ROCOF) How fast the grid frequency moves after an upset, in hertz per second. *Chapter 1*.

receding horizon Planning over a window ahead, applying only the first step and planning again at the next step. *Chapter 1*.

reference governor A device that alters a command so that the plant's response stays inside its limits. *Chapter 3*.

round-trip efficiency The share of the energy put into a store that comes back out. *Chapter 1*.

salt cavern A large cavity dissolved out of underground rock salt and used to store gas or compressed air. *Chapter 11*.

secondary control The layer that slowly returns frequency or voltage to its nominal value after droop has stopped the fall. *Chapter 1*.

selector control Control in which a selector passes the highest or lowest of several controller outputs to one actuator. *Chapter 3*.

smart inverter A solar or battery inverter that can also hold the voltage by trading reactive power. *Chapter 8*.

solid oxide fuel cell (SOFC) A fuel cell with a ceramic electrolyte that runs at several hundred degrees Celsius and can use reformed natural gas. *Chapter 3*.

state of charge (SOC) How full a store is, as a fraction or a percentage of what it can hold. *Chapter 1*.

state-of-charge manager A slow loop that nudges a store's power so its charge returns toward the middle of its band. *Chapter 1*.

steam header A large pipe that collects steam from a source and passes it to several users; its pressure shows whether supply and use balance. *Chapter 13*.

supervisory state machine Logic that moves a plant between defined modes, such as start, run and trip, by fixed rules. *Chapter 12*.

swing equation The relation between a power imbalance and how fast the speed of the grid's rotating machines, and so its frequency, changes. *Chapter 1*.

thermal following Running a cogeneration plant to meet the heat demand and selling or buying the electricity that results. *Chapter 1*.

thermocouple A temperature sensor made of two different metals joined at a tip, which give a small voltage that grows with temperature. *Chapter 1*.

time-scale separation Making each control layer several times slower than the one below, so the faster layer can be treated as settled. *Chapter 1*.

tip-speed ratio The speed of a blade tip divided by the wind speed; one value gives a rotor its best efficiency. *Chapter 4*.

uninterruptible power supply (UPS) A battery and converter that carry a critical load at once when its normal supply fails or lags. *Chapter 2*.

vanadium redox flow battery A battery that stores energy in liquid electrolytes held in tanks and pumped through a stack of cells. *Chapter 10*.

variable-speed drive An electronic drive that sets a motor's speed by changing the frequency of its supply. *Chapter 10*.

virtual synchronous machine An inverter controlled to behave like a rotating generator, with an inertia and droop of its own. *Chapter 7*.

Volt-VAR control Holding voltage by absorbing or supplying reactive power according to how far the voltage is from its target. *Chapter 1*.

Volt-Watt curtailment Cutting an inverter's active power when the voltage rises too high. *Chapter 8*.

water value The price at which energy kept in a store would be sold later, used to value what is left at the end of a plan. *Chapter 1*.

wave energy converter A device that captures energy from ocean waves and turns it into electricity. *Chapter 5*.

Index of Control Topics

active constraint, 233
AI agents
  causal claims, 361
  engineering practice, 359
  open problems, 363
  reproduction, 361
  truth audit, 361
  verification record, 360
  while studying, 17
anticausal optimum, 119

back-calculation, 33, 59, 89, 175, 228, 334
back-off, 235
branch and bound, 202
budget and trim, 286
bumpless transfer, 16, 33, 89, 175, 310, 337

cascade control
  anti-windup, 260
code verification, 126
combined heat and power, 9, 23, 49
complex-conjugate control, 119
constraint control
  unmeasured constraint, 218
  unmeasured limit, 355
control allocation, 33
control layers, 13
control structure
  against tuning, 355
controller action
  direct acting, 90

data export, 17
deadband, 199
discrete actuator, 197
DistFlow branch model, 195
disturbance model, 34, 62
  output bias, 146
droop control, 7, 163
  governor, 165
  inverters, 166
  load sharing, 173
  power-limit droop, 169
dwell timer, 91

economic dispatch, 10, 248, 300, 323
  inventory-constrained, 249
economic model predictive control, 112, 249, 323
  wave energy, 120
electrical following, 26
energy storage, 8, 248
  multiple services, 357
  reserve loop, 61
  state of charge in comparisons, 354
  state-of-charge manager, 175
energy systems
  control design procedure, 356
  open problems, 362
energy-maximizing control, 111
external reset, 228

fault detection, 313
feedforward, 60
  measured load, 33
  switched load, 174
feedforward control
  electrolyte flow, 260
final control element, 12
flow scheme, 4
forecasting, 10
  autoregressive, 120
  error, 262
  value of a forecast, 123, 354
fuel quality
  hydrogen blending, 44

gain scheduling, 80, 260
  pitch loop, 88
Gershgorin bound, 261
golden-section search, 147

headroom
  reserve, 186
hierarchical control, 12, 163, 191
  fast and slow layers, 353
hunting, 198
hysteresis, 34, 91, 312

inertia, 5, 85
  virtual, 169
infeasibility, 124

inferential control, 220
  inverse model, 229
inferential measurement, 199
instrument transformer, 12
integrated absolute error, 35
integrator windup, 33
interlock
  gain sign change, 144
inventory control, 8, 165, 248
  long-duration storage, 275
  price thresholds, 287
  reference curve, 287
  water value, 287

Kalman filter, 16, 62
  disturbance state, 92, 120, 231

linear programming, 200
loop interaction, 198

maintenance
  engine cogeneration, 43
  fuel cell plant, 74
  pitch bearing, 104
measurement delay, 197
minimum dwell time, 312
mixed-integer linear programming, 201
mixed-integer programming, 181
model mismatch, 37, 64, 235
model predictive control, 34, 61, 230
  against simpler control, 356
  collector field, 145
  energy dispatch, 181
  field demonstration, 105
  in the apps, 16
  mixed-integer, 201
  multi-stage, 93
  wind turbine, 91
move suppression, 336
moving horizon estimation, 16

numerical integration
  stiff actuator lag, 360

offset-free tracking, 336
override control, 33, 144, 275, 335
  pressure headroom, 287

passivity, 115
PID control
  anti-windup, 15, 285
  velocity form, 15, 59, 89, 143, 174, 228, 285, 310, 334
pole placement, 91
power converter, 12
power flow
  LinDistFlow, 195
preview control, 93
projected gradient, 261
projected gradient method, 336
protection
  inverter-based resources, 185
protective trip, 31

quadratic programming, 336

rate of change of frequency, 170
  measurement, 183
reactive control, 119
real-time optimization, 134
reference governor, 49, 59
rule-based control, 335

safety interlocks, 300
  trip, 311
saturation
  capacity limit, 38, 63
scenario tree, 93
scenarios, 16
  completion rules, 17
scheduling
  threshold price, 148
scorecard, 17
secondary control, 7, 174
selector
  min/max, 218, 230
  preload, 230
selector control, 59
  high selector, 144
setpoint offset, 90
setpoint parking, 310
simulation engine, 14
simulator framework, 13
soft constraint, 35, 121, 231
soft sensor
  flow, 358
split-range control
  two actuators on one measurement, 80
staging, 34
state estimation, 60
  distribution feeder, 199
  unmeasured limit, 363

supervisory control
    setpoint selection, 135
supervisory optimization, 134, 261, 323
supervisory state machine, 300
synchronizing
    practice, 185

terminal constraint, 265
terminal penalty, 336
thermal following, 26
time-scale separation, 5
transport delay
    flow-dependent, 143

Volt-VAR control, 8, 196
    history, 194
    open problems, 212
voltage measurement
    instrument transformers, 211
voltage regulation, 191
    practice, 212

water value, 312, 334
Willans line, 25

Index of Applications

absorption chiller, 26
actuator handover, 90
added mass, 113
air blower, 306
anode fuel margin, 55

battery, 8
battery aging, 235
    uneven, 242
battery fast charging, 358
battery management system, 170, 221, 260
battery preconditioning, 236
Betz limit, 84
biofouling, 129
Butler-Volmer equation, 55, 253

capacitor bank, 197
capacity commitment, 276
capacity fade, 224, 255
capacity floor, 286
capacity market, 278
capture width ratio, 117
cell voltage monitoring, 318
charge current limit map, 221
chiller, 224
cogeneration, 9, 23
collective pitch, 81
compressed air energy storage, 275
    adiabatic, 277
    diabatic, 277
    heat recovery, 295
compressibility factor, 306
concentrating solar power, 134
conservation voltage reduction, 199
constant-current, constant-voltage charging, 220
coulomb counting, 224
crossover, membrane, 254
curtailment, 263, 309

data center, 23, 49
    demand flexibility, 44
    electricity demand, 358
    power capping, 75
data hall, 26, 57
deliverability, 280
diesel genset, 168
distribution feeder, 8, 191
duck curve, 250
    origin, 250
dump radiator, 26

electrolysis
    flexible operation, 358
electrolyzer
    cost and durability, 318
    grid response, 303
electrolyzer switching price, 333
end-stop, 114
energy conversion, 1
    milestones, 2
equivalent circuit model, 221
exhaust heat recovery, 26, 56
expander
    fixed geometry, 280
    inlet throttle, 294

Faraday's law, 55
fast charging, 218
    charger power stages, 240
    extreme fast charging, 221
    standards, 240
flow battery, 248
    adoption, 269
    siting and permitting, 269
freeze protection, 145
frequency control, 5
fuel cell
    anode inventory, 55
    history, 3, 52, 303
    PEM, 300
    stack, 52
    stack life, 74
fuel processor, 52
fuel utilization, 55

gas analyzer, 58
gas chromatograph, 72
gas crossover, 305
generator speed measurement, 103
generator torque control, 81
graphite anode, 219

grid connection, 26
  history, 26
grid-following inverter, 164
grid-forming inverter, 166, 356
  hardware, 184
  history, 357

heat transfer fluid, 137
  pump, 156
high-temperature reactor, 327
  history, 327
Huntorf plant, 277
  black start, 277
hydrogen blending, 358
hydrogen contract, 323
hydrogen storage, 300

icing limit, 280
IEEE 1547, 166
instrument, 12
instrument transformer, 267
integrated energy system, 323
  open problems, 346
intrinsic impedance, 119
inverter
  foldback, 56
  history, 4
  power conditioning, 73
islanded bus, 57
islanding, 163

jacket water, 26
JONSWAP spectrum, 114

latching, 114
LIDAR preview, 93
lidar-assisted control, 357
limiting current density, 253
line-drop compensation, 198
linear generator, 114
  direct drive, 128
lithium plating, 219
  detection, 242
lithium-ion battery, 218
load curtailment, 34, 61

mass transport limitation, 253
McIntosh plant, 277
membrane hydration, 307
microgrid, 163
minimum load, 310
mirror defocus, 144

molten salt storage, 323
  pumps, 345
  two-tank, 329

Nernst equation, 252
Nernst potential, 55
net load, 250
nickel anode
  redox cycling, 51
NREL 5 MW reference turbine, 83
nuclear cogeneration, 323
  history, 326

off-gas burner, 56
on-load tap changer, 191
  bandwidth and time delay, 199
  hardware, 211
open-circuit voltage, 221
overpotential, 253
overspeed, 100
oxygen analyzer
  sample system, 315
oxygen storage, 300
  regulations, 317

packed-bed thermal store, 277
parabolic trough, 134
  automation, 158
  history, 136
  minimum flow, 156
peak shaving, 255
PEM electrolysis, 300
phase-locked loop, 164
photovoltaic array, 303
pitch sensitivity, 88
pitch system, 103
point absorber, 113
  absorption limit, 116
  heaving buoy, 111
  history, 114
point of common coupling, 164
power block, 140
power coefficient, 83
power system, 4
power take-off, 114
pressure transmitter
  wellhead, 293
pumping power, 254

radiation force
  Cummins form, 113
rebalance cell, 254

receiver tube, 137
    hydrogen, 157
reciprocating engine, 23
rectifier
    electrolyzer, 316
recuperator, 56
redox flow battery
    all-vanadium, 250
    NASA program, 250
resistance temperature detector, 42
resynchronization, 173
reverse power flow, 198
ride-through, 166
rotor inertia, 85
round-trip efficiency, 309

salt cavern, 275
    integrity management, 294
salt creep, 277
scarcity price, 265
seasonal reserve, 287
shunt current, 250
smart inverter, 196
smart meter, 197
solar collector field, 134
solar hydrogen plant, 300
solid oxide fuel cell, 49
solid-oxide electrolyzer, 329
    hot standby, 346
steam electrolysis, 323
    siting, 347
steam generator, 140
steam header, 328
    pressure transmitter, 344
steam reforming, 52
supplemental heater, 277
survival mode, 121
sync-check relay, 166
synchroscope, 166

take-or-pay penalty, 331
thermal energy storage
    history, 136
    two-tank, 137
thermal storage
    hot-water buffer, 29
thermal store, 281
thermistor, 225
    in a battery pack, 240
thermocouple, 155
thermoneutral voltage, 304
three-way valve, 43
tip-speed ratio, 83
torque law $K\omega^2$, 90
transport margin, 253
turbine trip, 343

under-frequency load shedding, 165
uninterruptible power supply, 26, 57

vanadium
    supply, 270
vanadium pentoxide precipitation, 259
vanadium redox flow battery, 248
variable-speed drive
    electrolyte pump, 268
virtual synchronous machine, 166
Volt-Watt curtailment, 196

water makeup, 307
water value, 276
wave energy converter, 111
    data quality, 128
    open problems, 129
wave measurement buoy, 127
wave probe, 114
wind turbine, 80
    aeroelasticity, 104
    history, 83
    operating regions, 80
    stall-regulated, 82

ZIP load model, 196

About the Editor

John D. Hedengren is a professor in the Department of Chemical and Biological Engineering at Brigham Young University, where he leads the Process Research and Intelligent Systems Modeling (PRISM) group. He received the B.S. and M.S. degrees in chemical engineering from Brigham Young University and the Ph.D. degree from the University of Texas at Austin in 2005. He developed advanced process control applications for polymer production at ExxonMobil Chemical Company and commissioned nonlinear model predictive controllers with PAS, Inc. before joining the faculty at Brigham Young University in 2011.

His teaching covers process dynamics and control, engineering computing, data science, machine learning, and dynamic optimization, with course material published online for use inside and outside the classroom. He developed the Temperature Control Lab, a portable heat transfer device now used by many universities for process control and data science instruction, and he maintains the APMonitor and GEKKO optimization software used in those courses. He serves as Chair of Control Education for the IEEE Control Systems Society and as a trustee of the CACHE Corporation. He received the AIChE David Himmelblau Award in 2014, the AIChE Computing Practice Award in 2018, and the John R. Ragazzini Award in 2024, and he was inducted into the Process Automation Hall of Fame in 2025.

Volumes in This Series

Interactive Control Engineering: Principles and Applications has 20 volumes. Nine are organized by control technology and eleven by industry or application area. Every chapter in every volume has a simulator at https://apmonitor.com/apps/.

1 **Process Dynamics, Modeling, and Identification.** How plants respond, and how to find the model.
2 **Practical PID Control.** Tuning, anti-windup and implementation that works on real loops.
3 **Feedforward, Cascade, and Ratio Control.** Choosing a structure from where the disturbance enters.
4 **Override, Split-Range, and Constraint Control.** When limits decide what the controller may do.
5 **Delays, Slow Measurements, and State Estimation.** Controlling what arrives late or cannot be measured.
6 **Multivariable and Plantwide Control.** Pairing loops, managing inventories, closing the recycle.
7 **Nonlinear Dynamics, Oscillations, and Stabilization.** Plants that will not stay put without control.
8 **Batch, Run-to-Run, and Hybrid Control.** Recipes, endpoints and learning from the last run.
9 **Model Predictive Control and Economic Optimization.** Planning moves ahead, and choosing where to operate.
10 **Control in Upstream and Midstream Oil and Gas.** From the wellbore to the terminal.
11 **Control in Petroleum Refining.** Throughput, quality and limits from crude unit to blending.
12 **Control in Petrochemicals and Polymers.** Severity, selectivity and grade from cracker to pellet.
13 **Control in Chemical Plants.** Recycle, equilibrium and energy in continuous chemical processes.
14 **Control in Semiconductor Manufacturing.** Run-to-run control where little can be measured in the chamber.
15 **Control in Mining and Mineral Processing.** Variable ore, late assays and equipment at its limit.
16 **Control in Iron, Steel, and Metals Production.** Slow furnaces, fast strip, and measurements that arrive once.
17 **Control of Energy Conversion and Storage.** Fast regulation and slow dispatch on a changing grid.
18 **Control of Nuclear and Fusion Systems.** Reactor physics, protection limits and coordinated control.
19 **Control of Biological Systems and Human Performance.** Variable subjects, sparse samples and limits that cannot be crossed.
20 **Motion Control, Robotics, and Autonomous Systems.** Tracking, navigation, contact and allocation.