INTERACTIVE CONTROL ENGINEERING

Control in Iron, Steel, and Metals Production

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 16

Control in Iron, Steel, and Metals Production

Slow furnaces, fast strip, and measurements that arrive once

John D. Hedengren
Editor

Advanced Process Solutions, LLC

Interactive Control Engineering: Principles and Applications
Volume 16: Control in Iron, Steel, and Metals Production
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/v16-iron-steel-metals/

How to cite this volume. John D. Hedengren, Ed., *Control in Iron, Steel, and Metals Production*, vol. 16 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026.

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

Preface

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

How to read this book with the apps

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

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

How AI agents were used

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

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

About this volume

This volume follows metal from ironmaking to coated strip. Chapters 2 to 5 cover the furnaces: a blast furnace, a hydrogen shaft furnace, an oxygen converter with one late measurement and an electric arc furnace. Chapters 6 to 10 cover casting, reheating, rolling, annealing and galvanizing, where material is tracked along a line and the delay is fixed in length. Chapter 11 covers an aluminum smelting cell.

A furnace answers in hours and a mill in milliseconds, and both serve the same order. Each chapter names the technology volume that teaches its main technique.

John D. Hedengren
2026

Contents

Preface v

Notation ix

1 Introduction: From Slow Furnaces to Fast Strip 1
- 1.1 The Question This Volume Answers . . . 1
- 1.2 A Short History . . . 2
- 1.3 Notation and Conventions . . . 3
- 1.4 From Furnace to Strip . . . 4
- 1.5 Instruments and Final Elements in This Volume . . . 12
- 1.6 The Simulator Framework . . . 14
- 1.7 Working With AI Agents While Studying . . . 18
- 1.8 Chapter Roadmap . . . 18
- 1.9 Exercises . . . 21

2 Iron Blast Furnace Simulator for Mid-Ranging Coal and Coke on a Sampled Hearth 23
- 2.1 Introduction . . . 24
- 2.2 Thermal Control of the Iron Blast Furnace . . . 25
- 2.3 Process Model . . . 26
- 2.4 Implemented Control System . . . 29
- 2.5 Results . . . 32
- 2.6 Technology in Practice . . . 40
- 2.7 Limitations . . . 42
- 2.8 Conclusion . . . 42
- 2.9 Exercises . . . 44

3 A Hydrogen Shaft Furnace Whose Product Answers Six Hours After Its Gas 48
- 3.1 Introduction . . . 49
- 3.2 Direct Reduction and Its Control . . . 50
- 3.3 The Shaft and Its Gas Loop as Modeled . . . 51
- 3.4 Dynamic Character . . . 57
- 3.5 Implemented Control System . . . 59
- 3.6 Nine Cases, Measured . . . 61
- 3.7 Technology in Practice . . . 64
- 3.8 Limitations . . . 67
- 3.9 Conclusion . . . 67
- 3.10 Exercises . . . 69

4 Basic Oxygen Furnace Simulator for Endpoint Control With One Late Measurement 72
- 4.1 Introduction . . . 73
- 4.2 Endpoint Control in Oxygen Steelmaking . . . 74
- 4.3 Model of One Heat . . . 75
- 4.4 Endpoint Strategies as Implemented . . . 78
- 4.5 Results . . . 80
- 4.6 Technology in Practice . . . 84
- 4.7 Limitations . . . 86
- 4.8 Conclusion . . . 87
- 4.9 Exercises . . . 88

5 Three Arcs on One Floating Neutral: Electrode Regulation in an Arc Furnace 92
5.1 Introduction . . . 93
5.2 Arc Furnaces and Their Electrode Regulators . . . 94
5.3 The Furnace as Modeled . . . 95
5.4 How the Neutral Couples the Phases . . . 100
5.5 Electrode Regulation as Implemented . . . 101
5.6 Nine Cases on the Running Furnace . . . 103
5.7 Code Verification and Agreement With Published Data . . . 108
5.8 Technology in Practice . . . 108
5.9 Limitations . . . 110
5.10 Conclusion . . . 111
5.11 Exercises . . . 112

6 Mold Level Simulator for Telling a Sticking Gate From Bulging and Overtuning 116
6.1 Introduction . . . 117
6.2 Mold Level Control on Continuous Casters . . . 118
6.3 Process Model . . . 119
6.4 Implemented Controller . . . 121
6.5 Results . . . 124
6.6 Technology in Practice . . . 128
6.7 Limitations . . . 130
6.8 Conclusion . . . 131
6.9 Exercises . . . 133

7 Walking-Beam Reheat Furnace Simulator for Slab Tracking and Level 2 Setpoint Control 136
7.1 Introduction . . . 137
7.2 From Zone Loops to Slab Tracking . . . 138
7.3 Furnace and Slab Model . . . 139
7.4 Two Levels of Control . . . 142
7.5 Results . . . 145
7.6 Technology in Practice . . . 150
7.7 Limitations . . . 152
7.8 Conclusion . . . 153
7.9 Exercises . . . 155

8 Hot Strip Mill Simulator for Inferential Gauge Control With a Sign-Inverted Error 159
8.1 Introduction . . . 160
8.2 Gauge Control on a Finishing Train . . . 161
8.3 Process Model . . . 162
8.4 The Two Measurements . . . 169
8.5 Controllers as Implemented . . . 171
8.6 Scenarios . . . 173
8.7 Results . . . 173
8.8 Technology in Practice . . . 179
8.9 Limitations . . . 181
8.10 Discussion . . . 182
8.11 Conclusion . . . 182
8.12 Exercises . . . 184

9 Annealing Line Entry Looper Simulator for Averaging Level Control and Speed Feedforward 187

9.1 Introduction . . . 188
9.2 Continuous Annealing and Its Control . . . 189
9.3 Model of the Line . . . 190
9.4 Controllers as Implemented . . . 193
9.5 Results . . . 197
9.6 Technology in Practice . . . 202
9.7 Limitations . . . 205
9.8 Conclusion . . . 206
9.9 Exercises . . . 208

10 Galvanizing Line Simulator for Coating Weight Control With a Delay Fixed in Length 211
10.1 Introduction . . . 212
10.2 Coating Weight Control on Galvanizing Lines . . . 213
10.3 Model of the Wiping Line . . . 214
10.4 The Coating Loop as Implemented . . . 217
10.5 The Linearized Loop . . . 219
10.6 Scenarios and Measured Results . . . 219
10.7 Technology in Practice . . . 224
10.8 Limitations . . . 226
10.9 Conclusion . . . 227
10.10 Exercises . . . 228

11 Alumina Feed Control Through a Non-Monotone Resistance in a Smelting Cell Simulator 232
11.1 Introduction . . . 233
11.2 Pots, Point Feeders and Process Computers . . . 234
11.3 One Prebake Cell at 300 kA . . . 235
11.4 A Reading With Two Answers . . . 242
11.5 The Cell Controller, Loop by Loop . . . 243
11.6 Measured Runs of the Cell . . . 245
11.7 Technology in Practice . . . 252
11.8 Limitations . . . 253
11.9 Conclusion . . . 254
11.10 Exercises . . . 256

12 Summary and Outlook 259
12.1 What the Chapters Showed . . . 259
12.2 Choosing a Method . . . 262
12.3 Where Practice Is Heading . . . 263
12.4 Engineering With AI Agents . . . 265
12.5 Gaps in Metals Process Control . . . 268
12.6 Exercises . . . 273

Answers to Selected Exercises 274

Glossary 278

Index of Control Topics 281

Index of Applications 284

About the Editor 287

Volumes in This Series 288

Notation

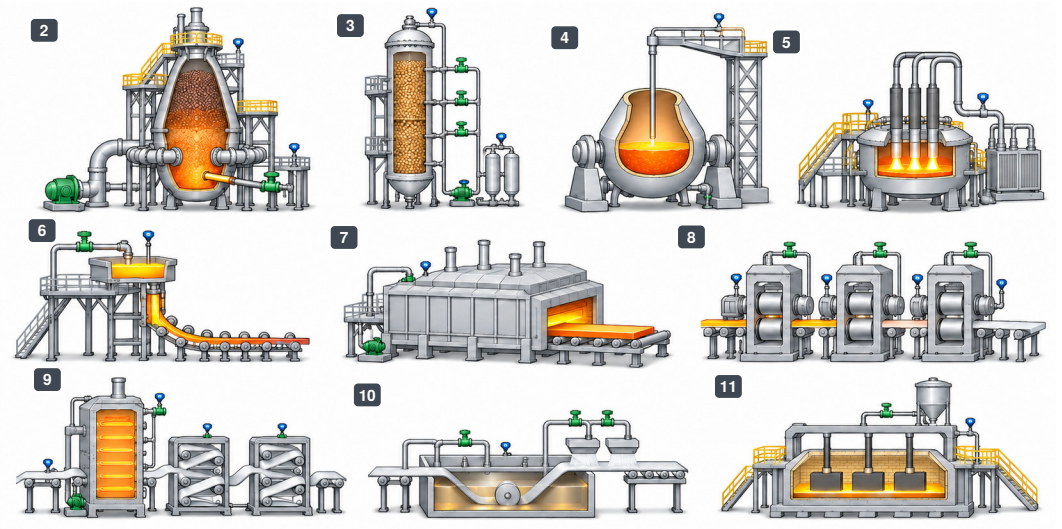
The chapters share the symbols below. A symbol used by one chapter only is defined where it first appears, and each chapter states its own units.

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 Slow Furnaces to Fast Strip



The chapters 2 blast furnace 3 hydrogen shaft furnace 4 oxygen converter 5 arc furnace 6 caster mold 7 reheat furnace 8 hot strip mill 9 annealing line 10 galvanizing line 11 aluminum smelting cell

Notice that the top row makes iron and steel, the middle row casts, reheats and rolls it, and the bottom row anneals and coats the strip and smelts aluminum. Illustration, not to scale.

IN PLAIN TERMS **Slow Furnaces Feeding Fast Strip**

The main idea. A steel works turns ore into iron, refines the iron into steel and rolls the steel into thin sheet. Some of its furnaces take hours to answer a change of fuel, while the rolling and coating lines move sheet past each machine in seconds.

The challenge. The slow furnaces must still deliver what the fast lines need. The things that matter most, such as the heat inside a slab or the thickness of the sheet, are seldom measured where the work is done.

Why it matters. Good control makes metal of the right heat, size and coating without wasting fuel, zinc or the metal itself.

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 1.4.1, 1.4.3 and 1.4.6, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, working each number of the example W-1, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.7 with Table 1.1, Section 1.5 with Figures 1.7 and 1.10, and Section 1.8.

1.1 The Question This Volume Answers

A steel works turns ore into coated strip through a chain of units that run on very different clocks. A blast furnace holds its burden for hours and answers a change of coke five hours later. A converter refines a heat in about twenty minutes and must stop it once, at the right carbon and temperature. A reheat furnace keeps a slab for two and a half hours, and the mill that receives it corrects its gauge every few milliseconds. Annealing and coating lines move strip at a few meters a second through furnaces that need minutes to change their temperature. This volume answers one question: how are

slow furnaces and fast strip controlled so that each unit delivers what the next one needs, and which techniques recur from the furnace to the coil?

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

- *Furnaces are slow.* A hearth, a slab or a smelting cell's frozen ledge stores so much heat that a change of firing shows in the product minutes to days later.
- *The product is seldom measured where it matters.* A hearth is dipped once a cast, a slab's mean temperature never, a strip's gauge at the stand only by inference and its coating a hundred meters downstream.
- *Steelmaking runs in batches judged at the end.* A heat is blown once, and carbon removed or iron burned cannot be put back.
- *Material moves, so delay is a distance.* A gauge reads strip that left the actuator meters ago, and the time that takes changes with every change of speed.
- *A fast, limited actuator works beside a slow, wide one.* Coal and coke, line speed and firing, and the electrode and the transformer tap each share one load.
- *Continuous lines are fed in pieces.* Heats, slabs and coils arrive one at a time, and the stops between them must be absorbed by an inventory or felt by the product.
- *Gains move and loops share nodes.* A heavier coating multiplies the gain of the coating loop, and three arc furnace electrodes share one floating neutral.

Ten chapters answer the question on ten plants, each with an app that the reader runs in a browser. Chapters 2 and 3 make iron, in a blast furnace and in a hydrogen shaft furnace. Chapters 4 and 5 make steel, by oxygen and by electric arcs. Chapter 6 casts it, Chapters 7 and 8 reheat and roll it, and Chapters 9 and 10 anneal and coat the strip. Chapter 11 adds a second metal, aluminum, made in an electrolytic cell whose heat balance is as slow as any furnace's.

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

- Estimate how fast a furnace and the product inside it answer, and say why a loop on the furnace does not hold the product.
- Divide a load between a fast, limited actuator and a slow, wide one, and check that the slow loop does not fight the fast one.
- Plan a batch from its measured charge, set the aim that a one-way late correction needs, and learn a drifting constant between batches.
- Write a transport delay as a length, track a piece of material along a line, and size an inventory against the stops it must carry.
- Say where an inferred gauge or coating is exact and where it is wrong, and which slower measurement must correct it.

This chapter develops each of these on one small example and names the chapter that uses it. PID forms and tuning are taken as known from Volume 2, feedforward, ratio and cascade from Volume 3, selectors, override and mid-ranging in general from Volume 4, dead time, Smith predictors and estimation from Volume 5, interaction and the relative gain array from Volume 6, batch and run-to-run control from Volume 8 and predictive control from Volume 9.

1.2 A Short History

The milestones of Figure 1.1 follow the plants of this volume, from the first fuel blown into a blast furnace to iron reduced by hydrogen made without fossil fuel.

Iron and steel came first. Fuel injection through the tuyeres of a blast furnace dates back to experiments with coal in France in 1840 [1], the practice whose rate Chapter 2 controls. Steel by blowing followed sixteen years later.

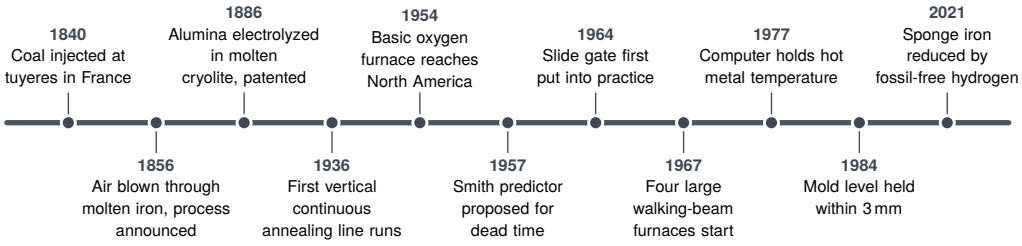


Figure 1.1 Milestones of iron, steel and aluminum making and their control. The dates are from [1] (1840), [2] (1856), [3] (1886), [4] (1936), [5] (1954), [6] (1957), [7] (1964), [8] (1967), [9] (1977), [10] (1984) and [11] (2021). The mold level of 1984 is dated by its report. Notice that each process came first and the loops that control its product decades later. Events are evenly spaced, not to a time scale.

> ⧖ **HISTORY** **1856 to 1969: From Air to Oxygen**
>
> In 1856 a process was announced at a meeting of the British Association that decarbonized cast iron by forcing a blast of air through the molten metal. Many metallurgists doubted it until they saw melted cast iron converted to malleable iron in a perfectly fluid state. Early trials by ironmasters failed, and only after two more years of experiments did the process turn out a product as good as that of the older methods [2]. A century later oxygen took the place of air. The basic oxygen furnace came to North America in 1954 and runs on its own heat when 70 percent or more of its charge is hot metal. By 1969 it made as much steel there as the open hearth, 60 million tons a year each [5].

Aluminum came by another route. A French patent of April 1886 covered a process based on alumina dissolved in molten cryolite, and in July of the same year a comparable patent was applied for in the United States. Apart from scale, hooding and fume recovery, the prebake cell of today keeps the basic design of the first commercial cells [3], the cell of Chapter 11.

The strip lines and their furnaces grew in the twentieth century. The first vertical continuous annealing line went into operation in 1936 at 25 m/min [4]. The sliding gate under a ladle, patented in the United States in 1885, was first put into practice in West Germany in 1964 [7], and four large walking-beam furnaces began to run at one Japanese works in June 1967 [8]. Control theory had supplied a tool for the delays of these lines in 1957, when the Smith predictor was proposed for processes with dead time [6].

> ⧖ **HISTORY** **Before 1960 to 1984: Measuring What the Loop Needs**
>
> Before 1960 automatic gauge control by load cell brought a gauge meter formula. The exit thickness at the head of each strip was computed from gap and force and stored, a step called lock-on, and the gap was then moved to hold the rest of the strip to that value [12]. In 1977 a blast furnace put its hot metal temperature under a computer that moved the injected oil, and 92 percent of the temperatures fell within 15 °C of the aim, against about 70 percent by hand [9]. In 1982 a level sensor for the caster mold was still called one of the most important and difficult tasks of the process [13], and in 1984 one slab caster held its level within 3 mm [10].

Iron is now being made in a new way. A pilot plant in Sweden, commissioned in August 2020, made the first sponge iron reduced with fossil-free hydrogen in June 2021 [11], the route of Chapter 3.

1.3 Notation and Conventions

Temperatures are in degrees Celsius, strip speeds in meters per second or per minute, thicknesses in millimeters, roll forces in meganewtons and coating weights in grams per square meter on each

side. A furnace or zone has a heat capacity C and a conductance G. A fast actuator u_f and a slow actuator u_s have gains K_f and K_s. In a converter C with a subscript is a carbon content in percent, λ the weight of a run-to-run update and δ the drift it follows. Strip runs at speed v, a gauge sits a length L downstream and θ is a delay. In a rolling stand d is the strip thickness, S the unloaded roll gap, F the roll force, M the mill modulus and Q the plastic modulus of the pass. The framework's h is its integration step and its a the controller action. Neither is used for anything else in this chapter, and the framework's Q in a predictive controller's objective is unrelated to the plastic modulus. 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 zone or strip temperature loop on the firing is reverse acting. A coating loop on the knife pressure is direct acting, because more pressure leaves less zinc.

The chapters keep their own symbols, and several collide. φ is the radiant exchange factor in Chapter 7, the heat-loss factor in Chapter 4 and the back-up roll angle in Chapter 8, while Chapter 6 writes the bulging phase ϕ. λ is the learning weight in Chapter 4, the laboratory bias weight in Chapter 3 and the roll pitch in Chapter 6, while Chapter 2 writes its bias weight α. Chapter 9 writes the controller action as α and Chapter 3 as σ, which is also a strip tension in Chapter 8 and the Stefan-Boltzmann constant in Chapter 7. h is the exit thickness in Chapter 8 and the slab thickness in Chapter 7, where this chapter writes d and 2ℓ. g is the pellet grade in Chapter 3, the exit factor of the strip in Chapter 9, the correction gain in Chapter 7, the feedforward multiplier in Chapter 10 and gravity in Chapter 6. M is the bath mass in Chapter 4 and the metallization in Chapter 3. K_I is the integrating gain of the mold in Chapter 6 and the current regulator's gain in Chapter 5. τ_c is the coke lag in Chapter 2 and a closed-loop time constant elsewhere, and T is the slab thickness in Chapter 6. In Chapter 11 φ is the share of a feeder shot that sinks as sludge, λ the factor on the metal pad's roll during an instability, g the slowdown of dissolution in a cold bath, h the bath-to-ledge heat transfer coefficient and M_b the bath mass.

1.4 From Furnace to Strip

1.4.1 The Flow Scheme

Figure 1.2 places each chapter's plant on the path from ore to coated strip. A blast furnace reduces sintered ore and pellets with coke and injected coal and taps liquid hot metal, which a basic oxygen converter refines to steel with scrap as coolant [14, 15]. This route and the one described next are the two main routes from ore to steel [16]. A shaft furnace can instead reduce pellets with hydrogen to solid direct reduced iron, which an electric arc furnace melts with scrap [17]. Either steel is cast continuously into slabs [18]. A reheat furnace brings the slabs back to rolling temperature, a hot strip mill rolls them to a few millimeters, and after cold rolling, which this volume does not cover, a continuous line anneals the strip and coats it with zinc. Aluminum has a chain of its own, in which electrolytic cells reduce alumina dissolved in a molten salt [19].

The techniques the chapters share are developed below on one running example, a works in miniature called W-1. Its hearth furnace makes hot metal on two fuels, its converter blows heats to a carbon and temperature aim, its reheat furnace heats 200 mm slabs for a hot mill, and one strip line anneals and galvanizes the rolled product. Every number of W-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.

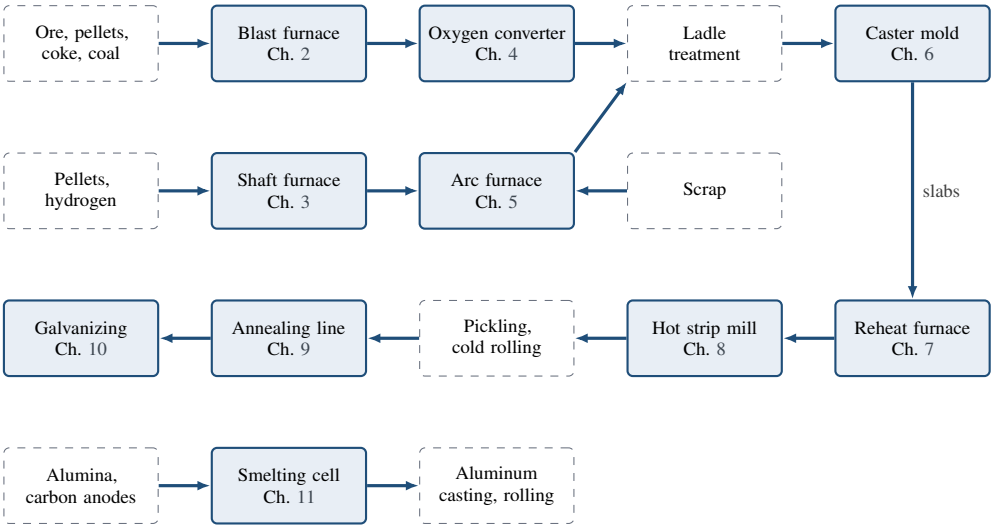


Figure 1.2 The Plants of the Volume on the Path From Ore to Coated Strip. Filled Blocks Are the Plants of Chapters 2 to 11, Dashed Blocks Lie Outside the Volume. Iron Is Made in the Top Two Rows, Steel Is Cast at the Right and Rolled, Annealed and Coated Along the Third Row, and Aluminum Has a Chain of Its Own. The Arrows Carry Hot Metal, Direct Reduced Iron (DRI), Steel, Slabs and Coils. Scrap Is Also Charged to the Converter as a Coolant (Not Drawn).

✦ **DID YOU KNOW? Two Billion Tonnes a Year**

The world made an estimated 1.9 billion metric tons of raw steel and 1.3 billion tons of pig iron in 2024 [20]. Its smelters made about 72 million tons of primary aluminum the same year [21], each ton of it in cells like the one of Chapter 11.

1.4.2 Thermal Inertia of Furnaces

A furnace zone stores heat in its walls, its gas and its load. If a zone of heat capacity C is fired at a rate Q_f and loses heat in proportion to its excess over the surroundings through a conductance G that lumps the losses and the linearized load, its temperature follows Equation (1.1).

$$C\frac{\mathrm{d}T_z}{\mathrm{d}t} = Q_f - G(T_z - T_0), \qquad \tau_z = \frac{C}{G} \tag{1.1}$$

A W-1 zone with $C = 90\,\mathrm{MJ/K}$ and $G = 0.15\,\mathrm{MW/K}$ answers a firing step with $\tau_z = 600\,\mathrm{s}$, ten minutes. A zone loop on that lag is an ordinary PI loop of Volume 2.

The slab inside the zone is slower. Heated from both faces, a slab of thickness 2ℓ conducts heat to its mid-plane on the time scale t_k of Equation (1.2), and its surface runs ahead of its center by an amount that the Biot number Bi measures.

$$t_k = \frac{\ell^2 \rho c_p}{k}, \qquad \mathrm{Bi} = \frac{h_r \ell}{k}, \qquad h_r = 4\varphi\sigma_{SB}T^3 \tag{1.2}$$

Here k is the conductivity, ρc_p the volumetric heat capacity, φ an exchange factor and h_r the radiant heat transfer coefficient at the absolute temperature T. W-1's 200 mm slab has $\ell = 0.1\,\mathrm{m}$, $k = 30\,\mathrm{W/(m\,K)}$ and $\rho c_p = 5.4\,\mathrm{MJ/(m^3\,K)}$, so $t_k = 30\,\mathrm{min}$. At 1250 °C with $\varphi = 0.6$, $h_r = 481\,\mathrm{W/(m^2\,K)}$ and $\mathrm{Bi} = 1.6$. The slab spends 120 min in three zones held at 1150, 1320 and 1235 °C. A five-cell conduction model of half the slab, radiated from its zones, discharges it with a mean of 1200 °C, a surface of 1212 °C and a core of 1194 °C, as Figure 1.3a shows. Three clocks are now in view: the zone answers in ten minutes, the slab's center in half an hour, and the product leaves after two hours.

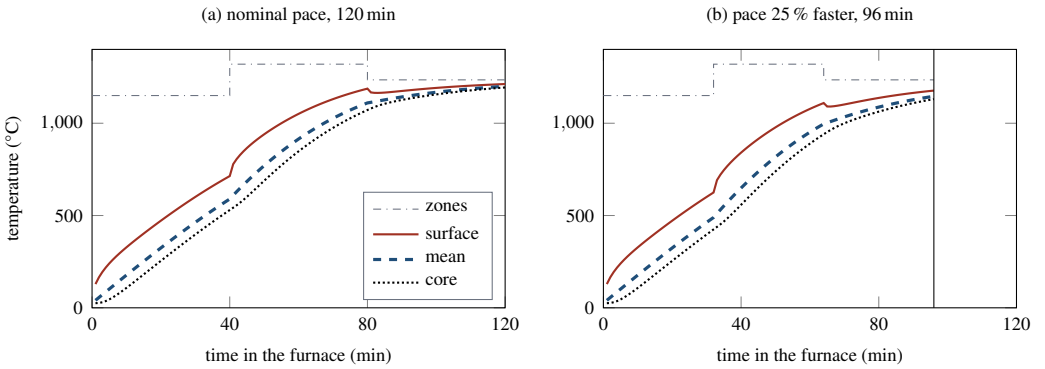


Figure 1.3 A W-1 Slab, 200 mm Thick and Charged Cold, Through Three Zones Held at Their Setpoints. The Surface (Solid), Mean (Dashed) and Core (Dotted) Against Time, With the Zone Temperatures (Dash-Dotted). (a) At the Nominal Pace the Slab Leaves With a Mean of 1200 °C. (b) At a Pace 25 % Faster, With the Zones Still on Setpoint, It Leaves at 96 min (Thin Vertical Line) With a Mean of 1146 °C.

The mill then asks for slabs 25 % faster. Each slab spends 96 min in the furnace, and with every zone loop holding its setpoint exactly the slab leaves with a mean of 1146 °C, 54 °C cold, as Figure 1.3b shows. Its surface is at 1177 °C, and the gap between surface and mean has grown from 12 to 31 °C. Nothing in the zone loops can see this, because they hold the furnace and not the product. A loop that raised every zone until the surface read its old 1212 °C would need 28 °C and would still discharge a mean of 1184 °C, because the wider gap breaks the fixed relation between surface and mean. A model that tracks the slab finds that a uniform offset of 40 °C restores the mean of 1200 °C, with a surface of 1228 °C.

Chapter 7 measures exactly this on a walking-beam furnace, whose 250 mm slabs have a conduction time of about 44 min and a Biot number of about 1.9 [22]. When the pace rises from one walk every 300 to every 240 s, zone loops that hold their setpoints deliver slabs 55 °C cold, a pyrometer trim on the surface settles the mean between 1174 and 1177 °C for every gain tried, and a slab-tracking model puts all 117 slabs on target, the approach that model-based furnace control takes [23]. The other furnaces of the volume add their own clocks. The coke of Chapter 2 reaches the hearth after five hours of burden descent and a three-hour lag. The product of Chapter 3 turns over a residence of six hours, while its top gas answers within three minutes. The furnace of Chapter 9 follows its firing with a lag of 240 s while strip crosses it in 80 s. The cell of Chapter 11 has a bath that answers in 1.30 h and a frozen ledge that integrates a heat imbalance for days, checked by one dip a day. Chapter 5 carries its bath temperature between dips forward by the metered energy.

1.4.3 Two Actuators on Two Clocks: Mid-Ranging

W-1's hearth is held at 1500 °C by two fuels. Injected fuel u_f acts with a gain $K_f = 1.0\,^{\circ}\mathrm{C}$ per kg/t of hot metal after a delay of 1 h and a lag of 1.5 h, but its injectors allow only 100 to 150 kg/t. Coke u_s acts with $K_s = 2.0\,^{\circ}\mathrm{C}$ per kg/t after a delay of 4 h and a lag of 2 h, over a wide range of 300 to 420 kg/t. A PI loop on the injected fuel, tuned by the SIMC rule of Volume 2 with a closed-loop time constant equal to the delay, has $K_c = 0.75\,\mathrm{kg/t}$ per °C and $\tau_I = 1.5\,\mathrm{h}$, and it is reverse acting. At 2 h a harder burden starts to demand 40 °C more heat through a lag of 2 h.

At steady state the two fuels must together supply the demand d, as Equation (1.3) states.

$$K_f\,\Delta u_f + K_s\,\Delta u_s = d, \qquad d \le K_f\,r_f \text{ (fast fuel alone)}, \qquad \Delta u_s \to \frac{d}{K_s} \text{ (mid-ranging)} \tag{1.3}$$

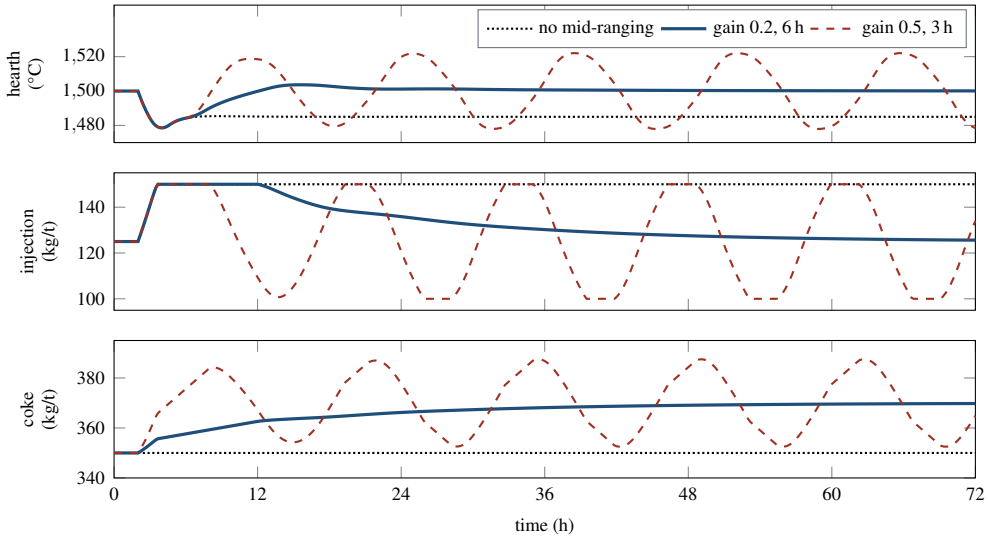


Figure 1.4 W-1's Hearth After a Burden Change That Demands 40 °C More Heat From 2 h. Hearth Temperature, Injected Fuel and Coke With the Injection Loop Alone (Dotted), With a Slow Mid-Ranging Loop on Coke, Gain 0.2 and Integral Time 6 h (Solid), and With a Faster One, Gain 0.5 and 3 h (Dashed). The Injection Saturates at 150 kg/t Without Mid-Ranging.

where r_f is the fast actuator's room above its operating point. W-1 runs its injection at 125 kg/t, so $r_f = 25\,\mathrm{kg/t}$ is worth only 25 °C. Alone, the injection loop reaches its limit at 3.7 h and stays there, and the hearth settles at 1485 °C, 15 °C cold, as the dotted curve of Figure 1.4 shows. Integral action cannot remove that offset, because the actuator cannot move.

Mid-ranging [24, 25] adds a second, slow PI loop whose measurement is the fast actuator's output and whose setpoint is the operating point the fast actuator should rest at, here 125 kg/t. Its output is the coke rate. More coke lets the injection fall, so the loop raises its output when its measurement rises, and it is direct acting. With a gain of 0.2 kg of coke per kg of injection and an integral time of 6 h, the coke rises to 369.8 kg/t, close to the 370 kg/t that Equation (1.3) requires, the injection leaves its limit at 12.0 h, the hearth is within 2 °C of setpoint from 19.5 h with a peak of 1503.6 °C, and the injection is back to 125.6 kg/t at 72 h. Neither structure prevents the first dip to 1478.5 °C at 4 h, because the coke arrives four hours after it is charged. A mid-ranging loop five times faster in integral action, with a gain of 0.5 and an integral time of 3 h, fights the injection loop through the coke's delay. The hearth then cycles between 1478 and 1522 °C with a period of about 13.6 h and never settles.

Three rules follow. The fast loop must be the faster of the two. The mid-ranging loop must be slow against the slow actuator's own delay and lag, here six hours, in gain as well as in integral time: in Chapter 2 a gain of 1.0 spends 7.0 h below the alarm even with a 12 h integral time. The resting point of the fast actuator is an economic choice. Resting it near its limit saves money when the fast fuel is the cheaper one, as coal is beside coke, but leaves less room for the next disturbance, which can then reach the limit before the slow fuel arrives. Anti-windup between the two loops is a refinement of its own [26], and Volume 4 treats mid-ranging in general.

Chapter 2 applies the structure with coal as the fast fuel and coke as the slow one, on a hearth dipped once every two hours. Two burden changes there hold the coal on its limit for about 26 h and leave the last half day at about 1477 °C. Mid-ranging raises the coke to 382 kg/t, keeps the coal off its limit and holds the last half day at about 1504 °C, for $1.73 against $1.68 million of fuel a day. A mid-ranging gain of 1.0 makes the two loops fight and spends 7.0 h below the cold-hearth alarm.

The chapter's predictive controller plans the coke five hours ahead instead. Chapter 9 uses the same shape without a mid-ranging loop. Line speed, borrowed from the entry looper, is its fast handle on strip temperature, the firing is its slow one, and a feedforward from speed to firing joins them. In Chapter 5 the electrode regulators correct each arc continuously while the transformer tap moves the operating point in coarse 50 V steps.

> **DID YOU KNOW? Two Handles on One Beam**
>
> Mid-ranging is not only a furnace technique. The storage ring of Volume 6, Chapter 6 steers its electron beam with a fast, small set of correctors and a slow, large one. It offloads the fast set to the slow one by the same mid-ranging [24] that lets coke take over from coal in Chapter 2.

1.4.4 Batch Endpoints in Steelmaking

A converter heat is a batch with one chance at its end. W-1 charges hot metal with 4.4 % carbon and blows oxygen until the bath reaches 0.05 % carbon within ±0.01 % and 1650 °C within ±10 °C. A charge model plans the oxygen and the coolant from the measured charge by mass and heat balances, which is feedforward from the initial state. Its plan is only as good as what it is told, and the scatter of what it is not told, the true carbon, the oxygen efficiency and the vessel's heat loss, leaves W-1's planned end temperature in error by a standard deviation of 18 °C.

Carbon has a measurement the temperature lacks. While carbon is plentiful, decarburization is limited by the oxygen supply. Below a critical carbon C^* the rate is limited by carbon transport and falls in proportion to the carbon left above an equilibrium value C_{eq}, as Equation (1.4) states.

$$r_C = r_{max}\,\frac{C - C_{eq}}{C^* - C_{eq}} \quad (C < C^*), \qquad \hat{C} = C_{eq} + (C^* - C_{eq})\,\frac{r_{C,m}}{r_{max}} \tag{1.4}$$

The off-gas flow and analysis give the measured rate $r_{C,m}$. Integrated from the charge, it estimates the carbon removed, and a 3 % meter error on W-1's 4.35 % of carbon removed is 0.13 %, thirteen times the half-width of the window. Read below C^* as an instantaneous rate, the same signal estimates the carbon by the second form of Equation (1.4), where the same relative error multiplies only $C - C_{eq}$. With $C^* = 0.40\,\%$ and $C_{eq} = 0.02\,\%$ that is 0.0009 % at the aim, 145 times smaller. Chapter 4 measures three hits and nine reblows in ten heats when it stops on the integral, against nine hits and no reblow on the rate with a 25 °C hot aim.

Temperature has one late measurement and a correction that works in one direction only. A sublance reading near the end of the blow lets the model re-plan, and coolant can then take heat out. Nothing puts heat back without blowing carbon below its aim. If e is the static plan's error, m a margin added to the aim and r the error left after the late re-plan, the heat ends at the temperature of Equation (1.5).

$$T_{end} - T_{aim} = \begin{cases} r, & m + e > 0 \text{ (too hot, cooled)} \\ m + e + r, & m + e \le 0 \text{ (too cold, left)} \end{cases} \tag{1.5}$$

With e and r normal, of standard deviations 18 and 4 °C, Figure 1.5 gives the share of W-1's heats inside the 10 °C window. Without the late reading, 42 % hit. With it and no margin, half the heats are cold and cannot be helped, so 70 % hit at 13.3 °C rms. A margin of one standard deviation raises the hits to 92.5 %, and two standard deviations, 36 °C, to 98.2 % at 4.2 °C rms, close to the 98.8 % that a two-way correction would reach. The margin is paid in coolant and in heats that must be cooled. In Chapter 4 all three misses of the sublance campaign without a margin are cold, and a hot aim of 25 °C cuts the rms error from 18.2 to 7.2 °C.

Between heats the model can learn. If one constant, such as the vessel's heat-loss factor, drifts by δ per heat as the lining wears, and each heat moves the model's estimate a fraction λ of the way toward the value that would have explained it, the estimate settles behind the truth by the lag of

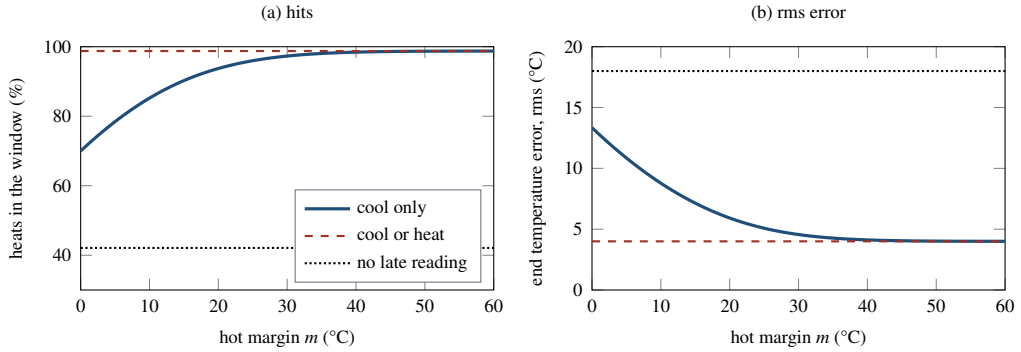


Figure 1.5 W-1's Converter Heats Against the Margin Added to the Temperature Aim. Static Error 18 °C, Error After the Late Re-Plan 4 °C, Window ±10 °C, All Normal. (a) The Share of Heats in the Window and (b) the rms End Temperature Error When the Late Correction Can Only Cool (Solid), When It Could Also Heat (Dashed) and Without the Late Reading (Dotted).

Equation (1.6) [27].

$$\hat{\varphi}_{k+1} = \hat{\varphi}_k + \lambda\,(\varphi_k - \hat{\varphi}_k) \quad \Rightarrow \quad \varphi_{k+1} - \hat{\varphi}_{k+1} \to \frac{\delta}{\lambda} \tag{1.6}$$

A W-1 lining whose loss factor grows by 0.02 a heat, learned at $\lambda = 0.5$, leaves a lag of 0.04, about 8 °C when each 0.01 of the factor is worth 2 °C. A larger λ shortens the lag and copies more of each heat's random error into the next plan. Chapter 4 derives a lag of 0.06, about 13 °C, and learning raises its hits on a wearing lining from five to seven in ten. Run-to-run control is the subject of Volume 8. The other batches of the volume have the same shape. Chapter 5 steps its power program on the metered energy per tonne and calls the tap when its estimate of the bath reaches 1630 °C. Chapter 11 ends each underfeed when a slope of the cell's resistance passes a threshold, a derivative read where a level says too little, as in the rate endpoint.

1.4.5 Tracking Material Along a Line

On a strip line the delay between an actuator and a gauge is a length. A gauge L downstream reads, at time t, the strip that passed the actuator at the earlier time $t - \theta$ defined by Equation (1.7).

$$\int_{t-\theta}^{t} v(t')\,\mathrm{d}t' = L, \qquad \theta = \frac{L}{v} \text{ at constant speed} \tag{1.7}$$

W-1's coating gauge sits 120 m after the air knives, a delay of 60 s at 2.0 m/s and 80 s at 1.5 m/s. A simulator or a Smith predictor [28] that stores the delay as 60 s is 20 s short after the slowdown, while one that records each piece of strip against its position along the line is right at every speed and through every ramp. Chapter 10 measures the difference. Its predictor with a delay stored in seconds tolerates a true delay 9 s longer than its own but not one 10 s longer, while the same predictor keeping its delay as 100 m of strip stays stable after a slowdown, at a spread of 1.24 g/m^2. Chapter 8 keys its X-ray gauge and the span between its stands to strip length in the same way. A disturbance can carry the speed as well. The strand of Chapter 6 bulges between support rolls of pitch λ, so the mold level sees a period λ / v, 12.5 s at 1.2 m/min on 0.25 m rolls, and a canceller locked to that computable frequency removes it.

Tracking is the same bookkeeping applied to identity. The model of Chapter 7 shifts its copy of every slab with each walk of the beams, so it knows each slab's position, charge temperature and elapsed time. Chapter 9 starts each coil's record at the strip coordinate it will reach after the strip

stored ahead of it in the looper, and the pellets of Chapter 3 charged at a step of the gas reach the bustle one residence later, at hour 7 of a step made at hour 1.

An inventory decouples a line fed in pieces from a furnace that wants a steady speed. W-1's entry stops for 60 s to weld each coil to the next while the process section runs at $v = 2.0\,\mathrm{m/s}$, and the entry decelerates and accelerates at $\dot{v}_E = 0.1\,\mathrm{m/s^2}$ toward a top speed of $v_E = 2.5\,\mathrm{m/s}$. One weld draws from the looper the length ΔL_w of Equation (1.8), and coils of length ℓ_c supply on average the speed $\bar{v}_E$.

$$\Delta L_w = v\left(t_w + \frac{v}{\dot{v}_E}\right), \qquad \bar{v}_E = \frac{\ell_c}{\ell_c/v_E + t_w + v_E/\dot{v}_E} \tag{1.8}$$

A weld draws 160 m, and a 500 m looper carries a stop of up to 230 s at full speed [29]. Coils of 2000 m supply 2.26 m/s on average, more than the furnace takes, but coils shorter than 850 m supply less, and coils of 800 m supply 1.975 m/s. No tuning of the looper loop can raise that average. Chapter 9 shows both halves. A looper loop tuned to hold its level passes every weld into the process speed and puts 71.3 % of the strip outside its temperature window, while an averaging loop that lets the level swing puts none outside. When short coils lower the supply to 277 m/min, the averaging loop lets the looper run empty seven times, and a cap on the speed at 97 % of that supply removes every stop.

1.4.6 Inferred Gauge and Coating

A rolling stand stretches under its own force, so the strip leaves thicker than the unloaded roll gap S by the stretch F/M, where M is the mill modulus [30]. The force in turn falls as the strip thickens, with a plastic modulus $Q = -\partial F/\partial d$, and an offset c such as thermal crown or roll wear moves the roll surface without moving the gap transducer. The thickness and its sensitivity to the gap are given by Equation (1.9).

$$d = S + c + \frac{F(d)}{M}, \qquad \frac{\partial d}{\partial S} = \frac{M}{M+Q} \tag{1.9}$$

W-1's last stand has $M = 5\,\mathrm{MN/mm}$ and $Q = 20\,\mathrm{MN/mm}$ and rolls 3.00 mm from a gap of 1.00 mm at 10 MN. Only a fifth of a gap move reaches the strip, so a gap move of five times a gauge error cancels it in one step.

The gaugemeter turns Equation (1.9) into a thickness measured at the stand, $\hat{d} = S + F/\hat{M}$, from the gap transducer, the load cell and a calibrated modulus $\hat{M}$. Its error is Equation (1.10).

$$\hat{d} - d = -c + F\left(\frac{1}{\hat{M}} - \frac{1}{M}\right) \tag{1.10}$$

Nothing on the right depends on the strip. A harder or colder piece moves the force and the thickness together along the mill line, and the gaugemeter reads it exactly. An offset is the opposite case, as Figure 1.6 shows. Crown of $c = -0.10\,\mathrm{mm}$ thins the strip by 0.02 mm and raises the force to 10.4 MN, and the gaugemeter reads 3.08 mm. Its error is the whole offset with the opposite sign, five times the true change. A loop that holds the gaugemeter at 3.00 mm closes the gap to 0.60 mm and delivers 2.90 mm, five times worse than no loop at all. A modulus calibrated 5 % high adds a steady bias of −95 µm.

> **+ DID YOU KNOW? Most of the Stretch Is in the Rolls**
>
> Of the elastic stretch of a rolling mill under load, the rolls take 40 to 70 percent, the housing 10 to 16 percent and the screw-down device 4 to 20 percent [12]. The mill modulus M of Equation (1.9) lumps all three into one number.

The X-ray gauge measures the product itself, 6 m after the stand and 0.5 s late at 12 m/s. A proportional loop on a pure delay in series with a static gain of 0.2 reaches its stability limit at a

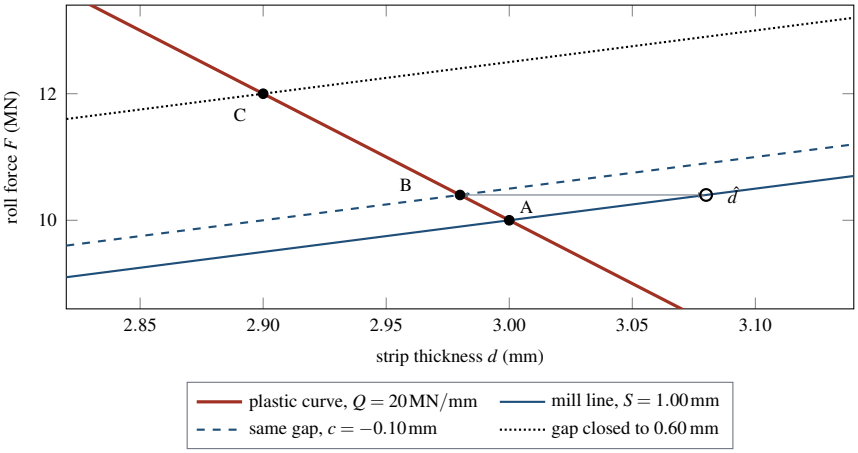


Figure 1.6 W-1's Last Stand. The Strip Leaves Where the Mill Line Crosses the Plastic Curve (Thick Solid), at A. Crown of −0.10 mm Moves the Mill Line (Dashed) and the Strip to B, 0.02 mm Thinner, While the Gaugemeter Reads the Higher Force Back Along the Undisplaced Line (Thin Solid) as $\hat{d} = 3.08$ mm (Open Circle). A Loop That Holds $\hat{d}$ at 3.00 mm Closes the Gap to 0.60 mm (Dotted) and Rolls 2.90 mm, at C.

gain of $1/0.2 = 5$ and a period of 1.0 s, so the delayed loop cannot use the gain the stand's own measurement allows. The two are combined by frequency: the gaugemeter closes the fast loop, and a slow trim from the X-ray removes its bias. Chapter 8 has a gap gain of 0.2104 and a deadbeat gain of 4.752. There, 200 µm of crown leaves the gaugemeter loop 198.6 µm thin against 41.9 µm with the loop open. On an entry wave the inferred loop holds 11.6 µm rms where the delayed loop at the same gains gives 173.8 µm, and a monitor trim at an eighth of the gauge loop's speed removes the crown.

A coating is inferred in the same spirit. Air knives wipe the zinc that a strip drags out of the bath, and the coating left on each side follows a power law in strip speed, knife distance D and knife pressure P [31], written in Equation (1.11) with the exponents Chapter 10 uses.

$$W = W_0 \left(\frac{v}{v_0}\right)^{0.6} \frac{D}{D_0} \left(\frac{P}{P_0}\right)^{-0.6}, \qquad \frac{\partial W}{\partial P} = -0.6\,\frac{W}{P} \tag{1.11}$$

W-1 coats 66 g/m^2 against a minimum of 60 g/m^2, with $W_0 = 60\,\mathrm{g/m^2}$ at $v_0 = 2.0\,\mathrm{m/s}$, $D = D_0$ and $P_0 = 25\,\mathrm{kPa}$. It needs 21.3 kPa, where the loop gain is −1.86 g/m^2 per kPa. Slowing to 1.5 m/s lowers the pressure needed to 16.0 kPa, which multiplies the gain by 1.33, and multiplies the delay by the same factor. Equation (1.11) holds the coating when P is proportional to v, so a speed feedforward removes the coating step and the gain change, but not the longer delay. A heavy grade of 140 g/m^2 at the same speed multiplies the gain by $(140/66)^{1+1/0.6} = 7.43$. Every gram per square meter on both sides of W-1's 1.25 m strip is 18 kg of zinc an hour, \$54 at \$3 per kilogram, so the spread of the coating sets the margin above the minimum and the zinc it costs. Chapter 10 finds a gain 5.6 times the nominal on its heavy grade, a spread cut from 5.08 to 1.25 g/m^2 by speed feedforward through a slowdown, and a coating tail heavy enough that 1 % below the minimum needs a margin of 3.3 standard deviations rather than 2.33.

The other chapters infer what they cannot measure. Chapter 2 reads a hearth soft sensor corrected at every cast [32], Chapter 3 closes an oxygen balance on the top gas and corrects it from a laboratory sample every two hours, Chapter 7 computes each slab's mean and corrects its exchange factor from a pyrometer, and Chapter 11 reads the alumina through a pseudo-resistance whose minimum hides which side of it the cell is on. The rule is the same in each [33]. An inferred measurement is exact for what its model contains, wrong and sometimes inverted for what the model leaves out, and it needs a

Table 1.1 For Each Chapter, the Slow Part, the Fast Part and What Joins Them

Ch.	Slow part	Fast part	What joins them
2	Coke, 5 h of descent	Coal, limited, hearth dipped every 2 h	Mid-ranging on coke, or a coke plan
3	Product over a 6 h residence	Top gas within minutes	Soft sensor with laboratory bias, buffer override
4	Lining wear from heat to heat	One blow of about 21 min	Charge model, late re-plan, run-to-run learning
5	Bath heat, taps in 50 V steps	Electrodes every 0.02 s	Power program on metered energy
6	Casting speed and strand	Mold level, sensor 0.3 s late	Canceller at the roll-passing frequency
7	Slab, 2.5 h in the furnace	Mill pace, zone loops	Slab-tracking setpoints every 120 s
8	X-ray 0.5 s late, crown drift	Gaugemeter, 4 ms scan	Monitor trim at an eighth the speed
9	Furnace, 240 s lag	Line speed, one 80 s residence	Looper, speed and firing feedforward
10	Gauge 100 m downstream	Knives, line speed ramps	Speed feedforward, predictor in meters
11	Heat and ledge, one dip a day	Resistance read every second	Target corrected from the daily dip

slower true measurement to correct its bias, with each sample compared to the inference at the instant the sample was taken, as Volume 5 develops.

1.4.7 Slow Furnaces and Fast Strip

A strip in a continuous furnace heats toward the furnace temperature T_f with a time constant proportional to its thickness, $\tau_1 d/d_1$, while it crosses a path of length L_f. Treated as plug flow, it leaves at the temperature of Equation (1.12).

$$T_{exit} = T_f - \left(T_f - T_{in}\right) \exp\left(-\frac{L_f}{v\,\tau_1 d/d_1}\right) \tag{1.12}$$

The exit temperature depends on speed and thickness only through the product vd. W-1's annealing section has $L_f = 120\,\text{m}$ and $\tau_1 = 30\,\text{s}$ for 1 mm strip, so at 2.0 m/s the strip stays 60 s and a furnace at 921.3 °C delivers 800 °C. A coil of 1.2 mm strip welded to it leaves at 752 °C, 48 °C cold. The furnace would need to rise by 59 °C, which takes minutes, while slowing the line to 1.67 m/s keeps vd and restores 800 °C within one residence. The looper absorbs the slower draw. Speed is the fast handle, firing the slow one, and a feedforward from speed to firing lets the furnace follow. Chapter 9 measures this. Strip of 1.3 mm with the speed held leaves 77.5 % of it outside its temperature window, against 2.5 % with the speed scaled by thickness.

Every chapter of the volume joins a slow part and a fast part in one of a few ways: a setpoint computed by a model of the slow part, a feedforward from the fast part's measured changes, an inventory that lets the two run at different rates, or a slow loop that returns a fast actuator to its range. Table 1.1 lists, for each chapter, what is slow, what is fast and what joins them. Chapter 12 adds the measured results and compares how slow furnaces and fast strip are coordinated.

1.5 Instruments and Final Elements in This Volume

Figure 1.7 lists what the controllers of each chapter read and what they move. A thermocouple, dipped into the metal or set in a furnace wall, is read in five chapters and a radiation pyrometer in two. Seven chapters control a value that only a model computes: a soft sensor, a slab mean, a gaugemeter, a carbon content from the off-gas rate, a bath temperature from the metered energy or a cell resistance. On the other side, solids are fed or added in five chapters, a hydraulic cylinder or a positioning drive moves the handle in four, and burners and valves in three each.

The hot metal of Chapter 2, the converter bath of Chapter 4 and the bath of Chapter 11 are read by a thermocouple dipped once a cast, once a heat or once a day. Thermocouples are typically less

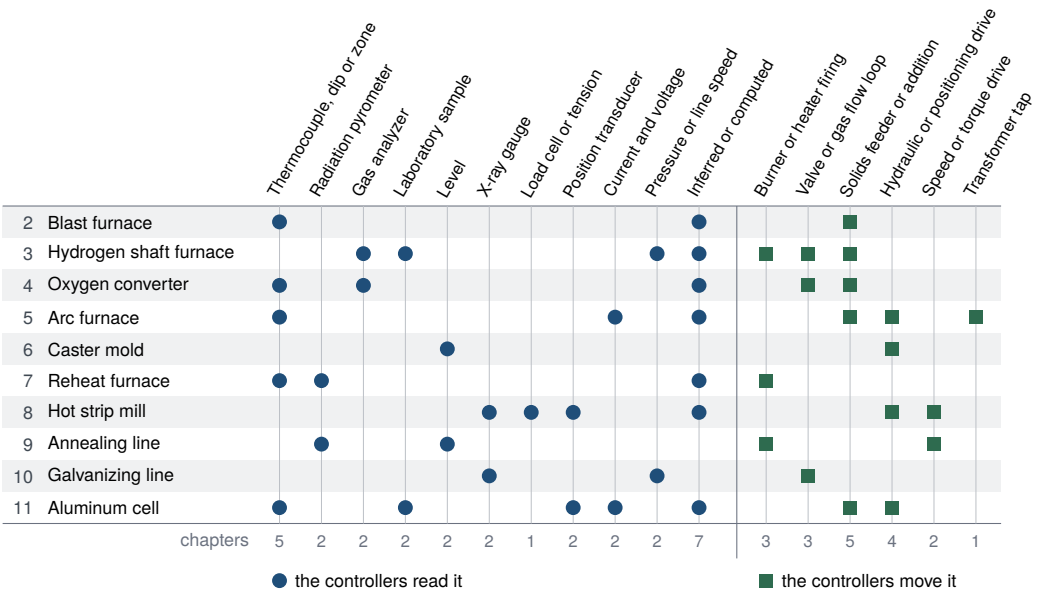


Figure 1.7 Sensors and final elements by chapter. A round mark means that the chapter's controllers read the device, a square mark that they move it. A mark is set only where the chapter's variables table, control structure or technology section names the device, and a flow station counts as a valve or flow loop. A blank says that the model leaves the device out. Notice that seven of the ten chapters control a computed value and that only the caster reads a single instrument.

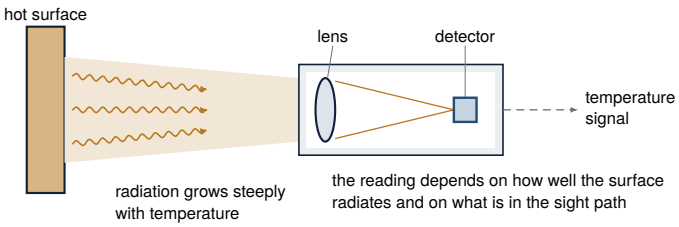


Figure 1.8 Principle of a radiation pyrometer. Notice that the instrument sees only the surface it points at and only through its emittance, so it reads the surface of a slab and never its mean, the quantity Chapter 7 must deliver [34].

accurate than resistance thermometers but more rugged, with greater spans [34], which suits a bath of molten metal. The reading is a sample, and each chapter holds it until the next one.

Hot product leaving a furnace is read without contact. A radiation pyrometer concentrates the radiation of a hot surface onto a detector, and the energy radiated grows with the fourth power of absolute temperature. Its main weakness is emittance, which depends on the material, its finish and its shape and is known for certain only by testing the object at a known temperature. Radiation reflected from other bodies adds to the reading [34]. Figure 1.8 draws the principle. Chapter 7 trims its slab model from a pyrometer, and Chapter 9 reads bright strip, the hard case.

The gauge of Chapter 8 is computed from two instruments at the stand, a transducer on the roll gap and load cells under the roll stack. A load cell is a steel element of precisely known modulus with strain gauges bonded to it, which compresses on a microscopic scale and changes their resistance [34]. Figure 1.9 draws one under the leg of a vessel. The gaugemeter of Equation (1.10) is only as good as this force and the modulus it is divided by.

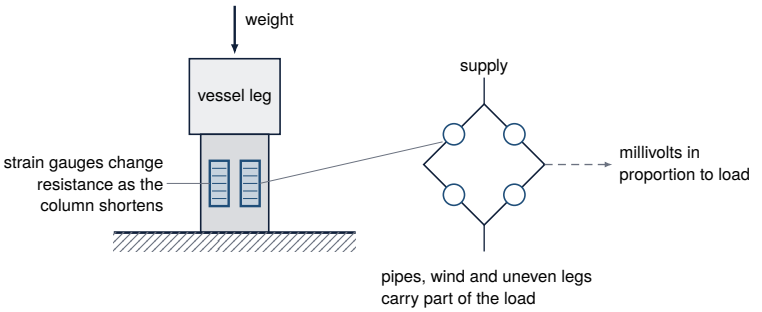


Figure 1.9 Principle of a strain-gauge load cell, drawn under the leg of a vessel. Notice that the cell reads only the load that passes through it. In a rolling stand the cells sit under the roll stack and carry the whole roll force [34].

Layer	What it does in this volume	Where it runs	How often
Setup, tracking and predictive control	plans a heat or a coil, tracks each slab, writes zone and fuel setpoints	**Level 2 process computer** that writes setpoints to the loops below	per heat or coil; 4 s to 30 min in the apps
Trims from late measurements	corrects a model or an inferred value from a dip, an assay or a gauge downstream		per sample: a cast, a heat, a laboratory report, a day
Regulatory control	PI loops on temperatures, flows, pressures, coating and the looper	**Distributed control system** or programmable controller	milliseconds to minutes; 0.2 s to 5 min in the apps
Fast servo loops	roll gap, electrode arms and the slide gate, inside the loops above	**Programmable controller** or a dedicated fast task	1 ms in one published mill; 4 ms to 0.05 s in the apps
Field devices	thermocouples, pyrometers, gauges, load cells, burners, cylinders and drives	on wires or a bus to the controller	continuously

Figure 1.10 The control layers of this volume, the platform that hosts each and how often each acts. The range of loop periods is that of [35], the 1 ms is the gap control of the published mill model that Chapter 8 cites, and the ranges in the apps are those of the chapters' technology sections. Notice that a trim acts once per cast, heat, report or day, while the servo loops below it act thousands of times a minute.

Figure 1.10 places the volume's loops on their platforms. Control loops in a plant run with cycle times from milliseconds to minutes, and a distributed control system uses a supervisory loop to mediate a group of localized controllers [35]. The apps run their fastest loops every 4 ms at the mill stand of Chapter 8 and every 0.02 s on the electrodes of Chapter 5, their regulatory loops from every 0.2 s in Chapter 10 to every 5 min in Chapter 2, and their setup and predictive layers from every 4 s in Chapter 9 to every 30 min in Chapter 3. The converter of Chapter 4 decides only at the events of each heat.

1.6 The Simulator Framework

Every chapter of this book has an app that runs in a web browser, collected with the others at https://apmonitor.com/apps/. About 240 apps exist across the series. All but a few are built on one shared framework, so the engine, the PID algorithm, the trends and the data export behave the same way in every chapter. The exceptions are standalone studios, such as the cascade studio and the two Temperature Control Lab apps, which follow the same ideas in their own code. This section describes the framework once and states where the standalone apps differ. Table 1.2 lists what a reader meets.

Table 1.2 What Every App Provides. Standalone Apps Are Described in Their Own Chapters.

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

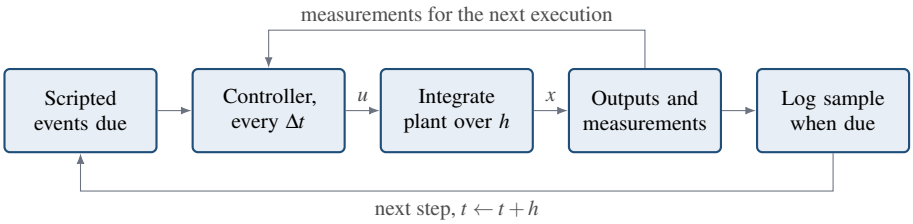


Figure 1.11 Work of the Engine in One Integration Step

1.6.1 The Simulation Engine

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

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

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

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

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

the scenario list restores the scenario's own settings instead.

1.6.2 The Velocity-Form PID Algorithm

One PID implementation serves the PID loops of every framework app. It is the velocity (incremental) form: each execution adds the change in the proportional and derivative terms to a running sum, and the integral is a separate term, as shown in Equations (1.13) to (1.16) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.13}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.14}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.15}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.16}$$

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 (1.16) drives I toward the value that reproduces the output applied, which is back-calculation [36]. With no integral action, Equation (1.16) 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 (1.17).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.17}$$

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 (1.18) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q (y_k - r_k) + \sum_{j} \Delta u_j^{\top} R \, \Delta u_j + \sum_{i} \rho_i s_i^2 \tag{1.18}$$

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 [37]. 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 [38]. 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 [39].

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

1.6.4 Scenarios, Progress and Data

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

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

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

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

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

An app exercise in this book names the scenario by its number and name as the app shows them, every setting that differs from the scenario's defaults, the run length in simulated time, the seed

(the default seed unless stated), what to read and where (a trend, a scorecard tile or a column of the exported file) and a tolerance on a numeric answer. Settings are made before Run unless the exercise gives a time.

1.7 Working With AI Agents While Studying

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

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

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

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

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

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

1.8 Chapter Roadmap

Table 1.3 lists the chapters with the lesson each adds to this one. The level says how much of Section 1.4 a chapter assumes. The time is an estimate for reading the chapter once and running each case of its app once, at about eight minutes a case.

Chapters 6, 9 and 2 form a first course: an integrating mold whose oscillation has three causes, an inventory that swings so the product does not, and two fuels on two clocks. They rest on Sections 1.4.3

Table 1.3 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Iron blast furnace	A slow, wide fuel returns a fast, limited one to its range.	intermediate	2.5 h
3	Hydrogen shaft furnace	The gas answers in minutes and the product over six hours.	advanced	2 h
4	Oxygen converter	One late reading and a one-way trim set the aim.	intermediate	2 h
5	Arc furnace	What a loop measures decides how its phases interact.	advanced	2 h
6	Caster mold	Diagnose an oscillation before treating it.	introductory	2 h
7	Reheat furnace	Holding the zones is not heating the slab.	intermediate	2.5 h
8	Hot strip mill	An inferred gauge is exact or inverted.	advanced	3 h
9	Annealing line	A looper is there to swing, not to hold its level.	introductory	2.5 h
10	Galvanizing line	A delay fixed in meters needs a model kept in meters.	intermediate	2 h
11	Aluminum smelting cell	A reading with a minimum cannot tell its side.	advanced	2.5 h

and 1.4.5. Chapters 4, 7 and 10 add a batch endpoint, a tracked product and a delay set by distance. Chapters 3, 5, 8 and 11 combine inferred measurements with interaction, inversion or an extremum and are best left until Section 1.4.6 is familiar. A reader who follows the metal can read the chapters in their order, which is the order of Figure 1.2.

KEY POINTS

- A steel works joins units on very different clocks, and each slow unit must deliver what the next one needs.
- Zone loops hold the furnace, not the product.
- A slow loop returns a fast, limited actuator to its range, and it must be slow against the slow actuator's delay.
- A one-way late correction needs a margin on the aim, paid in coolant.
- An inferred measurement is exact for what its model contains and wrong for what it leaves out, so a slower true measurement corrects it.

Where else this applies. Mid-ranging returns in Volume 4, delay and soft sensors in Volume 5 and run-to-run control in Volume 8.

References

[1] U.S. Department of Energy, "Comprehensive report to congress, clean coal technology program: Blast furnace granulated coal injection system demonstration project," U.S. Department of Energy, Washington, DC, Tech. Rep., 1990.

[2] "Bessemer, Sir Henry," in *Encyclopædia Britannica*, 11th ed. Cambridge: Cambridge University Press, 1910, vol. 3, p. 823, read in the Wikisource transcription (en.wikisource.org, 1911 Encyclopædia Britannica/Bessemer, Sir Henry), accessed November 19, 2026.

[3] T. R. Beck, "ECS classics: Hall and Héroult and the discovery of aluminum electrolysis," *The Electrochemical Society Interface*, vol. 23, no. 2, pp. 36–37, 2014.

[4] T. Fukushima, "Recent Technological Progress in High Speed Continuous Annealing," *Transactions of the Iron and Steel Institute of Japan*, vol. 25, no. 4, pp. 278–293, 1985.

[5] J. Stubbles, "Energy use in the U.S. steel industry: A historical perspective and future opportunities," U.S. Department of Energy, Office of Industrial Technologies, Tech. Rep., 2000.

[6] J. E. Normey-Rico, T. L. M. Santos, R. C. C. Flesch, and B. C. Torrico, "Control of dead-time process: From the Smith predictor to general multi-input multi-output dead-time compensators," *Frontiers in Control Engineering*, vol. 3, p. 953768, 2022.

[7] M. Ohnishi, T. Nishio, and T. Terashima, "Utilization of sliding gate nozzle in Japan," *Transactions of the Iron and Steel Institute of Japan*, vol. 16, no. 3, pp. 164–167, 1976.

[8] K. Kono, M. Fujiura, K. Goto, and J. Yamasaki, "Walking beam slab reheating furnace," *Tetsu-to-Hagane*, vol. 55, no. 14, pp. 1392–1398, 1969, in Japanese.

[9] M. Hatano, Y. Matoba, K. Otsuka, M. Yoshiki, and T. Miyagi, "Automatic control of hot metal temperature of blast furnace," *Transactions of the Iron and Steel Institute of Japan*, vol. 22, no. 7, pp. 534–542, 1982.

[10] "Technology of automatic mold bath level control," *Transactions of the Iron and Steel Institute of Japan*, vol. 24, no. 9, p. 777, 1984, new Technology note.

[11] F. Superchi, A. Mati, C. Carcasci, and A. Bianchini, "Techno-economic analysis of wind-powered green hydrogen production to facilitate the decarbonization of hard-to-abate sectors: A case study on steelmaking," *Applied Energy*, vol. 342, p. 121198, 2023.

[12] M. Ataka, "Rolling technology and theory for the last 100 years: The contribution of theory to innovation in strip rolling technology," *ISIJ International*, vol. 55, no. 1, pp. 89–102, 2015.

[13] H. Imada and Y. Nimura, "A new electrode-type mold level sensor for continuous casting machine," *Transactions of the Iron and Steel Institute of Japan*, vol. 22, no. 12, pp. 991–996, 1982.

[14] S. Ueda, S. Natsui, H. Nogami, J.-i. Yagi, and T. Ariyama, "Recent progress and future perspective on mathematical modeling of blast furnace," *ISIJ International*, vol. 50, no. 7, pp. 914–923, 2010.

[15] N. Dogan, G. A. Brooks, and M. A. Rhamdhani, "Comprehensive model of oxygen steelmaking part 1: Model development and validation," *ISIJ International*, vol. 51, no. 7, pp. 1086–1092, 2011.

[16] L. Holappa, "A general vision for reduction of energy consumption and CO_2 emissions from the steel industry," *Metals*, vol. 10, no. 9, p. 1117, 2020.

[17] F. Patisson and O. Mirgaux, "Hydrogen ironmaking: How it works," *Metals*, vol. 10, no. 7, p. 922, 2020.

[18] B. G. Thomas, "Modeling of the continuous casting of steel—past, present, and future," *Metallurgical and Materials Transactions B*, vol. 33, no. 6, pp. 795–812, 2002.

[19] P. Homsi, J.-M. Peyneau, and M. Reverdy, "Overview of process control in reduction cells and potlines," in *Essential Readings in Light Metals: Volume 2, Aluminum Reduction Technology*. Cham: Springer, 2016, pp. 739–746.

[20] U.S. Geological Survey, "Mineral commodity summaries 2025: Iron and steel," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[21] ——, "Mineral commodity summaries 2025: Aluminum," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[22] A. Steinboeck, D. Wild, T. Kiefer, and A. Kugi, "A mathematical model of a slab reheating furnace with radiative heat transfer and non-participating gaseous media," *International Journal of Heat and Mass Transfer*, vol. 53, no. 25–26, pp. 5933–5946, 2010.

[23] A. Steinboeck, D. Wild, and A. Kugi, "Nonlinear model predictive control of a continuous slab reheating furnace," *Control Engineering Practice*, vol. 21, no. 4, pp. 495–508, 2013.

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

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

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

[27] A. Ingolfsson and E. Sachs, "Stability and sensitivity of an EWMA controller," *Journal of Quality Technology*, vol. 25, no. 4, pp. 271–287, 1993.

[28] Z. Palmor, "Stability properties of Smith dead-time compensator controllers," *International Journal of Control*, vol. 32, no. 6, pp. 937–949, 1980.

[29] P. R. Pagilla, S. S. Garimella, L. H. Dreinhoefer, and E. O. King, "Dynamics and control of accumulators in continuous strip processing lines," *IEEE Transactions on Industry Applications*, vol. 37, no. 3, pp. 934–940, 2001.

[30] W. C. F. Hessenberg and R. B. Sims, "Principles of continuous gauge control in sheet and strip rolling," *Proceedings of the Institution of Mechanical Engineers*, vol. 166, no. 1, pp. 75–90, 1952.

[31] J. A. Thornton and H. F. Graff, "An analytical description of the jet finishing process for hot-dip metallic coatings on strip," *Metallurgical Transactions B*, vol. 7, no. 4, pp. 607–618, 1976.

[32] H. Saxén, C. Gao, and Z. Gao, "Data-driven time discrete models for dynamic prediction of the hot metal silicon content in the blast furnace—a review," *IEEE Transactions on Industrial Informatics*, vol. 9, no. 4, pp. 2213–2225, 2013.

[33] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

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

[35] 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.

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

[37] 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.

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

[39] 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

A furnace zone has a heat capacity of 144 MJ/K, and its losses and load together act as a conductance of 0.2 MW/K. It heats slabs 250 mm thick with $k = 30\,\mathrm{W/(m\,K)}$ and $\rho c_p = 5.4\,\mathrm{MJ/(m^3\,K)}$.

(a) Find the zone's time constant. *Checkpoint:* 720 s, 12 min.

(b) Find the slab's conduction time and its Biot number for a radiant coefficient of $450\,\mathrm{W/(m^2\,K)}$. *Checkpoint:* 46.9 min and 1.88.

1.2 Concept check · 2 min

A reheat furnace runs with three zone loops that hold their setpoints exactly. The mill asks for slabs 25 % faster, and nobody changes the setpoints. What happens to the slabs that reach the mill?

(a) nothing, because the zone loops hold every zone on its setpoint

(b) they leave colder, with a larger gap between surface and mean

(c) they leave colder, with a smaller gap between surface and mean

(d) they leave hotter, because the zones fire harder for the larger load

1.3 Theory and calculation · 4 min

A hearth is held by a fast fuel with a gain of 1.5 °C per kg/t that rests at 20 kg/t below its upper limit, and a slow fuel with a gain of 3.0 °C per kg/t. A burden change demands 45 °C more heat.

(a) With the fast loop alone, how much of the demand can it cover, and what offset is left? *Checkpoint:* 30 °C, 15 °C cold.

(b) With a mid-ranging loop on the slow fuel, how far must the slow fuel move at steady state, and where does the fast fuel end? *Checkpoint:* 15 kg/t, back at its resting point.

1.4 Concept check · 2 min

A mid-ranging loop measures the output of a fast fuel loop and moves a slow fuel that answers after a 4 h delay and a 2 h lag. More slow fuel lets the fast fuel fall. Which statement about the mid-ranging loop is right?

(a) it is reverse acting and should be faster than the fast loop

(b) it is reverse acting and should be slower than the slow fuel's dynamics

(c) it is direct acting and should be faster than the fast loop

(d) it is direct acting and should be slower than the slow fuel's dynamics

1.5 Theory and calculation · 3 min

A converter charges hot metal with 4.2 % carbon and aims at 0.04 % within ±0.01 %. The equilibrium carbon is 0.02 %, and the off-gas meters read the decarburization rate 2 % high.

(a) What carbon error does the meter error give if the rate is integrated from the charge? *Checkpoint:* 0.083 %.

(b) What error does it give if the rate is read below the critical carbon as in Equation (1.4)? *Checkpoint:* 0.0004 %.

1.6 Theory and calculation · 4 min

A converter model learns its heat-loss factor after every heat with weight $\lambda = 0.3$, while the true factor grows by 0.03 a heat as the lining wears. Each 0.01 of the factor is worth 2.5 °C of end temperature.

(a) How far behind the truth does the estimate settle? *Checkpoint:* 0.10.

(b) What end temperature error does that lag leave, and in which direction? *Checkpoint:* About 25 °C, cold.

1.7 Theory and calculation · 4 min

A process section runs at 2.5 m/s from a looper of 600 m. The entry stops for 40 s at each weld and decelerates and accelerates at 0.125 m/s^2.

(a) How much strip does one weld draw from the looper, and what is the longest stop a full looper carries? *Checkpoint:* 150 m and 220 s.

(b) Coils are 1500 m long and the entry can run at 3.0 m/s. Can the entry keep up on average? *Checkpoint:* Yes, 2.66 m/s against 2.5 m/s.

1.8 Theory and calculation · 3 min

A coating gauge sits 90 m after the air knives.

(a) Find the delay at 1.8 and at 1.2 m/s, and the error of a Smith predictor that keeps its delay at the first value after the line slows to the second. *Checkpoint:* 50 and 75 s, 25 s short.

(b) A 1.5 mm coil follows a 1.0 mm coil through an annealing furnace run at 2.4 m/s. At what speed does the thicker strip leave at the same temperature, with the furnace unchanged? *Checkpoint:* 1.60 m/s.

1.9 Theory and calculation · 5 min

A stand has a mill modulus $M = 4\,\mathrm{MN/mm}$ and a plastic modulus $Q = 16\,\mathrm{MN/mm}$. It rolls with an unloaded gap of 1.2 mm at a force of 8 MN.

(a) Find the thickness and the fraction of a gap move that reaches it. *Checkpoint:* 3.20 mm and 0.20.

(b) Roll wear opens the gap by $c = +0.05\,\mathrm{mm}$ without moving the gap transducer. Find the true change and the gaugemeter's error. *Checkpoint:* $+10\,\mu\mathrm{m}$ and $-50\,\mu\mathrm{m}$.

(c) A loop holds the gaugemeter at 3.20 mm. What does it deliver? *Checkpoint:* 3.25 mm.

1.10 Concept check · 2 min

A coating loop on knife pressure is tuned for 60 g/m^2. The line switches to 120 g/m^2 at the same speed, and the coating follows Equation (1.11). How does the loop's gain change?

(a) it rises about sixfold, so the loop may go unstable

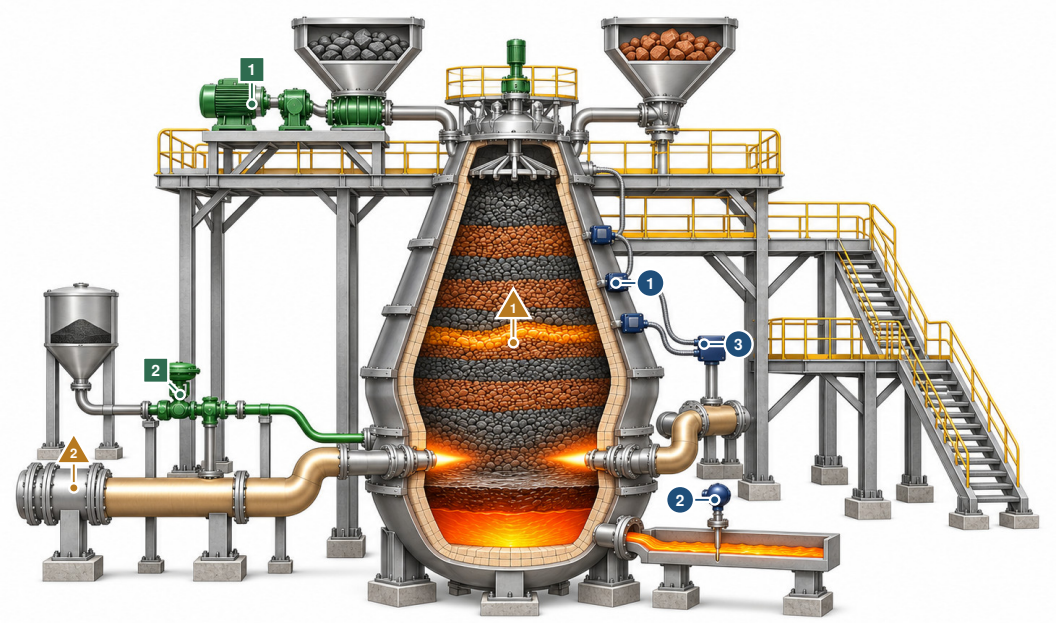
(b) it doubles, in proportion to the coating

(c) it halves, because the knives need less pressure

(d) it stays the same, because the exponents are fixed

CHAPTER 2

Iron Blast Furnace Simulator for Mid-Ranging Coal and Coke on a Sampled Hearth



The controller moves 1 coke rate, charged at the top 2 coal injected at the tuyeres

The controller reads 1 wall stave heat, feeding the thermal index 2 temperature dip and sample at each cast 3 hot blast temperature

What disturbs it 1 harder ore in the burden 2 a hot stove fault that cools the blast

Notice that the coal enters just above the hearth while the coke must ride the layers down from the top, and that the hearth is read only at the runner when the furnace is cast. Illustration, not to scale.

IN PLAIN TERMS **Two Fuels That Arrive at Different Times**

The main idea. A blast furnace melts iron ore with burning coke. The liquid iron collects in a pool at the bottom called the hearth. Its temperature is checked only when the furnace is drained, about every two hours.

The challenge. Coal blown in near the bottom acts within hours but first cools the hearth a little. Coke added at the top takes five hours to arrive. The controller makes quick fixes with coal, hands lasting changes to coke, and checks any estimate it uses in between against each reading.

Why it matters. A hearth that runs cold drains poorly and can freeze, and one that runs hot wastes costly coke.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	The first page, the fact box of Section 2.1, the two history boxes of Section 2.2, the try-it boxes and the fact box of Section 2.5, 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 the open challenge in it

The control question A blast furnace hearth is dipped once a cast and answers coal within the hour but coke only after five hours of burden descent. How should the coal loop be tuned for the cast, kept honest by a drifting soft sensor and relieved by the coke when the coal runs out of room?

Learning objectives

- Derive the net gain and inverse response of the coal path and the worth of coal against coke.
- Tune the coal loop by SIMC with the dead time a held cast adds.
- Compute the lag of a soft sensor's bias update against a drift.
- Decide from the capacity balance when the coal saturates and where mid-ranging on coke settles.
- Compare mid-ranging with a predictive controller that plans the coke on its delay.

Prerequisites Sections 1.4.3 (two actuators on two clocks), 1.6.2 (velocity form and loop action), 1.4.4 (the EWMA lag) and 1.4.6 (soft sensors). Section 1.6.3 only for the predictive controller.

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

See also Volume 4, Chapter 10 and Volume 4, Chapter 9 for mid-ranging and valve-position control. Volume 5, Chapter 3 for a loop tuned through a slow analyzer, Volume 5, Chapter 10 and Volume 5, Chapter 4 for inferred qualities corrected by late samples.

2.1 Introduction

The loops in a first control course have one actuator and a measurement that arrives continuously. Many industrial loops have neither. A fast actuator with a narrow range often works beside a slow actuator with a wide one, and the standard way to coordinate them, mid-ranging or valve position control, returns the fast actuator to the middle of its range while the slow one carries the sustained load [1, 2]. Quality variables are often known only from laboratory samples, and continuous inferential estimates stand in between the samples on the condition that they are corrected against them [3, 4].

The iron blast furnace combines both problems in one loop. The thermal state of its hearth, the hot metal temperature and the silicon content that follows it, is measured when the furnace is tapped, and the prediction of silicon from routine process data has a literature of its own because the sample arrives so seldom [5]. Coal injected at the tuyeres changes the heat supply within hours, but it first lowers the raceway flame temperature [6], while coke reaches the hearth only after the burden has descended from the top. Coal is the cheaper fuel, and raising the injection rate toward what the raceway can burn has long been an aim of furnace operation [7]. The steel industry's carbon dioxide emissions also make each kilogram of coke per tonne a matter of policy as well as cost [8].

DID YOU KNOW? 1.3 Billion Tonnes of Pig Iron

The world made an estimated 1.3 billion metric tons of pig iron in 2024, beside 1.9 billion tons of raw steel [9]. A furnace blown in in 1974, with a hearth 9.6 m across, made 3800 t of pig iron a day [10]. The simulated furnace of this chapter makes 10 000 t.

The cases of this chapter are built around three misconceptions. The first is that a loop can be tuned without regard to how often it sees the plant, when a correction held for two hours adds about an hour of dead time [11]. The second is that integral action removes any offset, when an actuator on its limit leaves the loop open. The third treats a soft sensor as though it were a measurement, when it is a model whose error grows until something checks it. The class of problem is general. Any plant with a cheap, fast and bounded input beside an expensive, slow and wide one, such as a trim valve beside a large valve, a letdown beside a boiler, or a fine reagent valve beside a coarse one, needs the same division of labor, and the saturation and windup that follow when it is missing have been studied in their own right [12, 13]. Volume 4, Chapter 9 returns a cooling load to an absorption chiller with a valve-position controller, and Volume 4, Chapter 10 adds a mid-ranging layer above a fast turbine bypass and a slow firing rate.

This chapter works with a blast furnace thermal simulator in which the coal loop can run on cast dips or on a soft sensor, the soft sensor can be corrected or left to drift, and a mid-ranging loop on coke can be switched in, with each choice priced in hearth temperature and fuel cost. Sampled PI control, bias updating and mid-ranging are established methods, and so is the model predictive controller (MPC) that the application also offers, which plans both fuels over a twelve-hour horizon in the receding-horizon manner of industrial practice [14]. What is particular to this plant is the thermal model of the hearth, with an inverse response on the coal path and a burden-descent delay on the coke path, written out completely enough to be rebuilt. The scenarios first isolate the measurement problem and the actuator problem and then pose them together. Section 2.2 traces the history of furnace thermal control and Section 2.3 gives the model. Section 2.4 describes the coal loop, the mid-ranging loop on coke and the MPC. Section 2.5 reports the measured results case by case, from the two response times through the sampled loop, the soft sensor, saturation and mid-ranging, the coal target and a stove fault to the graded challenge and the comparison with the MPC in Section 2.5.7. Section 2.6 sets the loops in their hardware, Section 2.7 states the limitations, Section 2.8 concludes, and Section 2.9 holds the exercises.

2.2 Thermal Control of the Iron Blast Furnace

A blast furnace is a countercurrent shaft reactor. Coke and iron ore are charged in alternating layers at the top, hot air is blown in at the tuyeres near the bottom, and the gas from the burning coke rises through the burden, reducing and melting the ore on its way. Liquid iron and slag collect in the hearth and are drained through a taphole at intervals, and every drainage is a cast. The phenomena inside the shaft, from burden descent and gas flow to heat exchange and reduction, have been the subject of mathematical models of increasing detail for several decades [15].

Operators judge the furnace by its thermal level. A hearth that runs cold drains poorly and in the extreme chills, and a hearth that runs hot wastes coke and raises the silicon in the iron. Because the thermal level is known directly only from the temperature and silicon of each cast, the historical practice was to read the recent casts, compare them with experience, and adjust the fuel rate, the blast temperature or the blast moisture. Expert systems that encoded that practice were applied to furnace operation by the late 1980s, with rules that diagnosed the thermal trend and proposed actions [16]. Statistical and neural predictors followed, trained on routine process data to forecast the silicon or the hot metal temperature of the next casts [17, 18], and the data-driven silicon models have been reviewed in detail, including the long and variable delays between the inputs and the cast [5].

> ⧖ HISTORY **The Hearth Under Computer Control (1975 to 1977)**
>
> In the 1970s one furnace that held its hot metal temperature by the rate of injected heavy oil put that rate under a computer. Its model was updated against the measured top gas composition and hot metal temperature, predicted the temperature ahead by simulating the furnace forward, and moved the oil by an integral law on the predicted error. Automatic control began in November 1977. Over a week of 78 casts, 92 % of the temperatures fell within 15 °C of the aim, against about 70 % by hand, and the steadier furnace was run 15 °C cooler, which cut the fuel rate by about 4 kg per tonne [10].

Tuyere injection changed the actuator set. Pulverized coal blown in with the hot blast replaces part of the coke as a fuel and reductant, and its combustion in the short residence time of the raceway was studied extensively as injection rates rose [7]. Injected coal lowers the raceway flame temperature, which is why a step in coal first cools the lower furnace before its extra reducing gas warms the hearth [6]. With coal at the tuyeres and coke at the top, the furnace acquired exactly the pair of actuators that mid-ranging is designed for. Mid-ranging structures were analyzed as a design problem [1], tuning rules for the valve position loop were developed and applied industrially [2], anti-windup was added for the case in which the fast actuator saturates before the slow one has taken over [12], and a feedforward formulation has since been proposed [13].

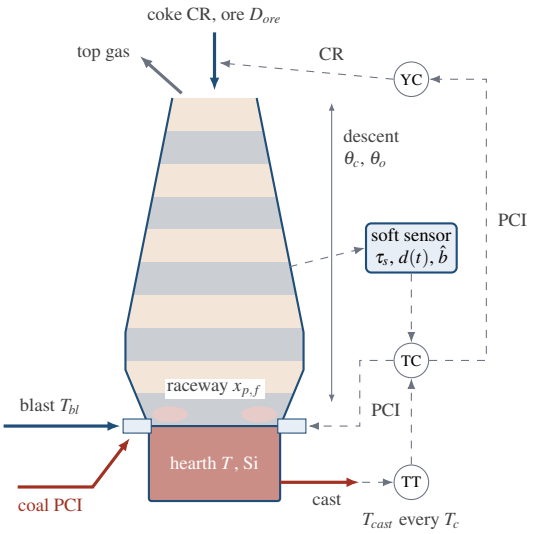


Figure 2.1 Iron blast furnace and its thermal control loops. Coke at rate CR and ore of thermal demand D_{ore} are charged at the top and reach the tuyeres after the burden descent times θ_c and θ_o. Hot blast at T_{bl} and pulverized coal at rate PCI enter at the tuyeres, where the coal first cools the raceway ($x_{p,f}$). The hearth temperature T and silicon are dipped at every cast, once every T_c, by TT, and a soft sensor built from stave heat loads lags the hearth by τ_s, drifts by $d(t)$ and is corrected by the bias estimate $\hat{b}$. The controller TC sets the coal rate and the mid-ranging controller YC moves the coke rate to return the coal to its target.

> ⧖ HISTORY **Fuel at the Tuyeres (1840 to the 1980s)**
>
> Fuel injection through blast furnace tuyeres dates back to coal experiments in France in 1840. Trials in the United Kingdom in 1962 blew coarsely ground coal in pneumatically at up to 360 pounds per ton of hot metal, and the coal replaced coke about pound for pound, but the trials stopped when oil became cheaper [19]. Oil had been the preferred injectant in the 1950s. When fuel oil grew scarce and costly around 1979, coal was taken up again, and injection spread quickly in Europe and Japan during the 1980s [20]. A demonstration proposed in the United States in 1990 fed each of 28 tuyeres per furnace from pressurized storage through metering equipment [19].

The soft sensor has its own lineage. Inferential control, in which a secondary measurement estimates an unmeasured output for use in feedback, was formalized in the late 1970s [3]. Data-driven soft sensors are now common in the process industries, and the review literature treats their drift as a maintenance problem, answered by adaptation mechanisms of which the simplest is a bias update against each laboratory result [4, 21]. Volume 5, Chapter 10 keeps a psychrometric soft sensor in service the same way, against a sampled and delayed laboratory moisture.

2.3 Process Model

The simulated furnace makes 10 000 tHM/d of hot metal on 350 kg/tHM of coke and 150 kg/tHM of injected coal with a 1180 °C blast. At these base rates the hearth sits at 1500 °C with 0.45 % silicon, and at $400 per tonne of coke and $150 per tonne of coal the fuel costs $162.5 per tonne of hot metal, or $1.625 million a day. The furnace and its two loops are drawn in Figure 2.1.

Table 2.1 lists the variables. The model integrates five thermal states, one for coke, two for coal, one for the blast and one for the burden, together with three accumulators for the fuel cost and the tonnes of coke and coal consumed. Every thermal state is a deviation from the base operation and

Table 2.1 Variables, Their Roles and Ranges Around the Base Operation

Symbol	Role	Base	Range
T	CV, hot metal temperature	1500 °C	SP 1460–1540
Si	hot metal silicon	0.45 %	≥ 0.05
T_{cast}	measured, cast dip	1500 °C	held 1 cast
$\hat{T}_s$	measured, soft sensor	1500 °C	—
PCI	MV, coal injection	150 kg/tHM	80–200
CR	MV, coke rate	350 kg/tHM	300–420
x_c, $x_{p,f}$, $x_{p,s}$, x_b, x_o	states, thermal	0	—
D_{ore}	DV, burden demand	0 °C	−90 to 30
T_{bl}	DV, blast temperature	1180 °C	1050–1250
σ_w	DV, furnace wander (1σ)	6 °C	0–15
r_d	DV, soft sensor drift	0 °C/h	±10
σ_s	DV, soft sensor noise	3 °C	0–10

starts at its steady value for the coke rate, coal rate, blast and burden in force, zero in every case as loaded. Each scenario therefore begins at a true steady state, and a rate chosen before the run starts is a starting condition rather than a step.

2.3.1 The Thermal State as a Sum of Delayed Responses

The hearth temperature is the base temperature plus five deviations and a slow wander, as shown in Equation (2.1).

$$T = 1500 + x_c + x_{p,f} + x_{p,s} + x_b + x_o + w(t) \tag{2.1}$$

Each deviation is a first-order lag behind a dead time, driven by the deviation of one input from its base value. Coke must descend with the burden, Equation (2.2).

$$\tau_c \frac{\mathrm{d}x_c}{\mathrm{d}t} = G_k \,\Delta\mathrm{CR}(t - \theta_c) - x_c \tag{2.2}$$

where $\Delta\mathrm{CR} = \mathrm{CR} - 350$ is the coke rate deviation in kg/tHM, $G_k = 2.0\,^\circ\mathrm{C}$ per kg/tHM is the thermal gain of coke, $\theta_c = 5\,\mathrm{h}$ is the burden descent time and $\tau_c = 3\,\mathrm{h}$ is the lag. A coke change therefore does nothing at the hearth for five hours.

2.3.2 Coal and the Inverse Response

Coal acts through two parallel paths with opposite signs, as shown in Equations (2.3) and (2.4).

$$\tau_r \frac{\mathrm{d}x_{p,f}}{\mathrm{d}t} = -G_r \,\Delta\mathrm{PCI}(t) - x_{p,f} \tag{2.3}$$

$$\tau_p \frac{\mathrm{d}x_{p,s}}{\mathrm{d}t} = G_p \,\Delta\mathrm{PCI}(t - \theta_p) - x_{p,s} \tag{2.4}$$

The fast path, with $G_r = 0.4$ and $\tau_r = 0.4\,\mathrm{h}$, has no dead time and represents the cold coal cooling the raceway flame. The slow path, with $G_p = 1.6$, $\theta_p = 1\,\mathrm{h}$ and $\tau_p = 1.5\,\mathrm{h}$, represents the extra reducing gas reaching the hearth. The steady-state gain is $G_p - G_r = 1.2\,^\circ\mathrm{C}$ per kg/tHM, positive, but the first response to a coal increase is a fall, an inverse response. In the resulting transfer function, Equation (2.5),

$$\frac{T(s)}{\Delta\mathrm{PCI}(s)} = -\frac{G_r}{\tau_r s + 1} + \frac{G_p\,\mathrm{e}^{-\theta_p s}}{\tau_p s + 1} \tag{2.5}$$

the fast negative term dominates until the delayed positive term arrives. In the notation of Section 1.4.3, coal is the fast actuator with $K_f = G_p - G_r$ and coke the slow one with $K_s = G_k$. Thermally, one kilogram of coal is worth $1.2/2.0 = 0.6$ kg of coke, so a degree of hearth temperature costs \$0.125 per tonne of hot metal from coal and \$0.20 from coke at the default prices.

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

2.3.3 Blast, Burden and Wander

The blast temperature and the burden act as Equations (2.6) and (2.7).

$$\tau_b \frac{\mathrm{d}x_b}{\mathrm{d}t} = G_b \,\Delta T_{bl}(t-\theta_b) - x_b \tag{2.6}$$

$$\tau_o \frac{\mathrm{d}x_o}{\mathrm{d}t} = D_{ore}(t-\theta_o) - x_o \tag{2.7}$$

with $G_b = 0.25$, $\theta_b = 0.3\,\mathrm{h}$, $\tau_b = 0.5\,\mathrm{h}$, $\theta_o = 4\,\mathrm{h}$ and $\tau_o = 2\,\mathrm{h}$. The burden demand D_{ore} is expressed directly in degrees of hearth temperature, so a harder ore that needs 30 °C more heat is $D_{ore} = -30$. The wander $w(t) = \sigma_w\, n(t)$ is first-order colored noise n of unit standard deviation and a 6 h time constant, standing in for coke quality, permeability and moisture. Silicon follows the temperature linearly, Equation (2.8).

$$\mathrm{Si} = \max\left[0.05,\; 0.45 + 0.008\,(T - 1500)\right] \;\%\tag{2.8}$$

The engine is that of Section 1.6.1. The dead times are realized by a ring buffer of the four input deviations, written at every 5 min controller scan and read with linear interpolation. The inputs are held between scans, and the states are integrated by fourth-order Runge-Kutta with a 60 s step. The fuel cost accumulates at $\dot{C} = (P/86400)(\mathrm{CR}\,p_k + \mathrm{PCI}\,p_p)/1000$ dollars per second, with the production P in tHM/d and the prices p_k and p_p in dollars per tonne.

2.3.4 The Cast and the Soft Sensor

The furnace is cast every $T_c = 2\,\mathrm{h}$, first at $t = T_c$. At cast k the dip reads $T_{cast,k} = T(kT_c) + \sigma_d\,\varepsilon_k$ with $\sigma_d = 6\,^\circ\mathrm{C}$ and ε_k a standard normal deviate, and the laboratory silicon reads $\mathrm{Si} + 0.03\,\varepsilon'_k$. Both are held until the next cast. The soft sensor is the hearth temperature through a first-order filter, plus a drift and colored noise, Equation (2.9).

$$T_s = \mathscr{F}_{\tau_s}[T] + d(t) + \sigma_s\,\nu(t), \qquad \frac{\mathrm{d}d}{\mathrm{d}t} = r_d \tag{2.9}$$

with $\tau_s = 0.5\,\mathrm{h}$ and ν of unit standard deviation and a 10 min time constant. When correction is on, the corrected reading is $\hat{T}_s = T_s - \hat{b}$, and at every cast the bias estimate moves a fraction α of the way toward the observed discrepancy, Equation (2.10).

$$\hat{b} \leftarrow \hat{b} + \alpha\left(T_s - \hat{b} - T_{cast,k}\right) \tag{2.10}$$

The default weight is $\alpha = 0.5$. A hearth below 1450 °C raises the cold-hearth alarm and one below 1420 °C the chilled-hearth alarm, each cleared with a hysteresis of 5 and 10 °C.

The parameters are listed in Table 2.2. The base rates and temperatures that define the deviations, 350 and 150 kg/tHM, 1180 °C and 1500 °C, are written in the code.

2.3.5 Assumptions

The thermal state is linear in the fuel rates with fixed dead times and lags. On a real furnace the descent time changes with the production rate, and the replacement of coke by coal becomes less effective at high injection rates [7], so the gains and delays here hold only near the base operation.

The hearth is one lumped temperature. Radial and circumferential differences, the deadman and the drainage behavior of the hearth are outside the model, as are burden distribution and gas flow [15].

Silicon is locked to the temperature by Equation (2.8). In practice the two are correlated but not locked, and many furnaces control from the laboratory silicon rather than from the temperature [5].

Casts are instantaneous samples on a fixed schedule. A real cast lasts an hour or more, the temperature is dipped partway through it, and its timing varies with taphole practice.

The soft sensor is a lagged copy of the truth with a drift and noise. It stands in for estimators built from stave heat loads, top gas analysis and burden descent rate, whose errors are structured rather than additive.

Table 2.2 Model Parameters. Visible Ones Can Be Changed in the Application, Hidden Ones Are Not Shown, and Fixed Ones Are Constants in the Code

Symbol	Meaning	Value	Access
G_k	thermal gain of coke	2.0 °C per kg/tHM	visible
θ_c, τ_c	burden descent, coke lag	5, 3 h	visible
G_p	coal gain, reducing gas	1.6 °C per kg/tHM	visible
θ_p, τ_p	coal dead time, lag	1, 1.5 h	visible
G_r	raceway cooling by coal	0.4 °C per kg/tHM	visible
τ_r	raceway cooling lag	0.4 h	hidden
G_b	thermal gain of blast	0.25	visible
θ_b, τ_b	blast dead time, lag	0.3, 0.5 h	fixed
θ_o, τ_o	burden demand delay, lag	4, 2 h	fixed
—	wander time constant	6 h	fixed
T_c	time between casts	2 h	visible
σ_d	cast dip noise (1σ)	6 °C	visible
—	laboratory silicon noise	0.03 %	fixed
τ_s	soft sensor lag	0.5 h	visible
—	soft sensor noise time constant	10 min	fixed
p_k, p_p	coke, coal price	400, $150/t	visible
P	production	10 000 tHM/d	visible
—	cold, chilled hearth	1450, 1420 °C	visible

2.4 Implemented Control System

2.4.1 Coal Loop on the Hot Metal Temperature

The coal controller TC is the framework's velocity-form PID with no derivative term, executed as Equations (2.11) to (2.14).

$$e_k = T_{SP} - T_{m,k}, \qquad P_k = K_c\, e_k \tag{2.11}$$

$$s_k = s_{k-1} + (P_k - P_{k-1}), \qquad c_k = s_k + I_{k-1} \tag{2.12}$$

$$u_k = \mathrm{clamp}\,(c_k, 80, 200) \tag{2.13}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\, e_k + \frac{u_k - c_k}{\tau_t}\right] \Delta t_k \tag{2.14}$$

where T_m is either the cast dip T_{cast} or the corrected soft sensor $\hat{T}_s$, c_k is the coal rate the algorithm demands, u_k the coal rate applied in kg/tHM and $\tau_t = \tau_I$ the tracking time. These are Equations (1.13) to (1.16) of Section 1.6.2 with no derivative and no output bias, with s_k for S_k and c_k for v_k. The sum is never clamped, the integral starts from the initial coal rate and so carries the bias, and only the integral is corrected for clipping, by back-calculation against the coal rate actually applied. The loop is reverse acting, with the action $a = +1$ of Section 1.6.2: a cold hearth raises the coal, because the steady-state gain of Equation (2.5) is positive even though the first response is not.

The execution depends on the measurement. On the soft sensor the loop runs at every 5 min scan with $\Delta t_k = 300$ s. On cast dips it runs only when a new cast arrives, with Δt_k equal to the time since its last execution, normally $T_c = 2$ h, and its output is held in between. Each cast therefore adds $(K_c/\tau_I) e_k T_c$ to the integral in one step. The defaults are $K_c = 0.30$ kg/tHM per °C and $\tau_I = 2$ h, with sliders from 0 to 3 and from 0.25 to 48 h, where 48 h removes integral action and with it the back-calculation. In manual mode the integral tracks the operator's coal rate at every scan and each cast still counts as an execution, so transfer to automatic is bumpless. After six hours of manual operation in case 2, the first cast in automatic adds 2.4 kg/tHM to the integral, one cast interval's worth at its 8 °C error.

The default tuning is close to the SIMC rule once the hold is counted as dead time [11]. Treating the raceway lag as effective delay because it produces the inverse response, and adding half the sample interval for the zero-order hold, the coal path has $\theta_e \approx \theta_p + \tau_r + T_c/2 = 2.4$ h, a lag $\tau_p = 1.5$ h

and a gain of 1.2. With the closed-loop time constant $\tau_{cl} = \theta_e$ the rule gives Equation (2.15).

$$\begin{aligned} K_c &= \frac{\tau_p}{1.2\,(\tau_{cl} + \theta_e)} \approx 0.26 \\ \tau_I &= \min\left[\tau_p,\, 4(\tau_{cl} + \theta_e)\right] = 1.5\,\mathrm{h} \end{aligned} \tag{2.15}$$

Without the hold term the same rule gives $K_c \approx 0.45$, so the two-hour sample alone should cut the gain by about a factor of 1.7.

2.4.2 Mid-Ranging on Coke

The mid-ranging controller YC is a second velocity-form PI, direct acting with $a = -1$ in the convention of Section 1.6.2, as the app's Show Model note also calls it. Its measurement is the coal command and its setpoint is the coal target PCI_{SP}, Equation (2.16).

$$\mathrm{CR}_k^{*} = \mathrm{clamp}[\mathrm{PI}_{K_v,\tau_V}\left(\mathrm{PCI}_k - \mathrm{PCI}_{SP}\right),\, 300,\, 420] \tag{2.16}$$

A coal rate above its target raises the coke rate, which warms the hearth five hours later, which in turn lets the coal loop take coal back out. The loop runs at every 5 min scan, so on cast dips it sees a coal command that changes once per cast. The scenarios use $K_v = 0.2$ kg of coke per kg of coal and $\tau_V = 12\,\mathrm{h}$. The applied coke rate follows the command at no more than 5 kg/tHM per hour, a limit on how fast the burden recipe may change. That rate limit acts after the controller and is not fed back to its integral. When mid-ranging is off the coke rate is set by hand and the controller tracks it.

The structure is the standard one of Section 1.4.3, and its two tuning choices are the ones the literature emphasizes [1, 2]. The coal loop must be the faster of the two, and the mid-ranging loop must be slow enough that the coke's five-hour delay does not make the two fight. There is no feedforward from the blast temperature or the burden, and no anti-windup coupling between the loops of the kind proposed for a saturating fast actuator [12].

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

2.4.3 Model Predictive Control of Coal and Coke

A second thermal-control structure, never the default and selected as loaded only by scenario 10, gives both fuels to one model predictive controller. Every $T_{\mathrm{MPC}} = 15\,\mathrm{min}$, which is three controller scans, and again after each cast, a quadratic program plans the coal rate and the coke rate over the next $N = 48$ periods (12 h), in the receding-horizon manner of Section 1.6.3. Predictive control of the silicon content and the slag basicity, the quality indices that follow the thermal state, has been studied on a large industrial furnace whose slow responses take more than eight hours [22]. Here the coal is decided as a rate held over each period and the coke as a ramp v in kg/tHM per hour, bounded by the recipe change limit $r_c = 5\,\mathrm{kg/tHM}$ per hour that already acts after the mid-ranging loop, so the coke the plan counts on is the coke the burden receives. While a plan acts, the coal loop and the mid-ranging loop track the applied rates at every scan with no preload.

The controller's model is Equations (2.2) to (2.4) with the parameter values of Table 2.2, frozen when the application loads, so a parameter changed by the learner becomes model error, as it would on a furnace. Its estimate of the hearth is Equation (2.17),

$$\hat{T} = 1500 + x_c + x_{p,f} + x_{p,s} + \hat{d} \tag{2.17}$$

where the three states are driven by the rates the controller itself applied, kept in its own five-hour record, and $\hat{d}$ is a single output disturbance that carries everything the model leaves out: the blast and burden paths of Equations (2.6) and (2.7), the wander, the soft sensor drift left after correction, and any changed parameter. The balance is linear. Stepped by 10 kg/tHM in manual with the noise off, the simulated hearth follows the model to 4×10^{-8} °C, while a first-order-plus-dead-time fit misses the coal step by 5.2 % RMS of its response and cannot reproduce its inverse response. Sampled at

T_{MPC}, the coke and coal dead times become chains of 20 and 4 periods, and the program's model has 28 states. A plan made at 60 h of scenario 10 with the setpoint raised 20 °C, then applied open loop with the wander and noise off, predicts the true hearth at the end of every period to within 0.005 °C.

On cast dips the disturbance is corrected once per cast, as shown in Equation (2.18),

$$\hat{d} \leftarrow \hat{d} + \alpha \left[T_{cast,k} - T_{mod}(t_k) - \hat{d} \right] \tag{2.18}$$

with the weight α of Equation (2.10) and the model's own hearth $T_{mod} = \hat{T} - \hat{d}$ read from its record at the tap time t_k, when the dip was taken. Between casts $\hat{d}$ is held. On the soft sensor a steady-state Kalman filter acts at every scan on two states, the index's declared 0.5 h lag and $\hat{d}$ as a random walk, and reads the index exactly as the coal loop reads it, bias-corrected at each cast when the correction is on. Its variances come from the declared noise settings rather than from tuning, 1 °C^2 per scan for $\hat{d}$ from the 6 °C, 6 h wander and 36 °C^2 for the index, whose 3 °C colored noise is replaced by white noise of equal low-frequency power, and the resulting gains are 0.115 on the lag state and 0.157 on $\hat{d}$ per scan. One integrating disturbance on the one measured output makes the plan offset-free [23], but on the measurement rather than on the hearth, so an index drift that the cast correction has not yet removed becomes hearth error.

Each plan minimizes Equation (2.19),

$$\begin{aligned} J = & \sum_{k=1}^{N} \left[\left(\hat{T}_k - T_{SP} \right)^2 + w_m \left(\mathrm{PCI}_k - \mathrm{PCI}_{SP} \right)^2 + r_v v_k^2 \right] \\ & + \sum_{j=1}^{12} \left(\lambda \, \Delta \mathrm{PCI}_j^2 + s_v \, \Delta v_j^2 \right) \end{aligned} \tag{2.19}$$

where $\hat{T}_k$ is the predicted hearth at the end of period k, PCI_{SP} is the coal target of Equation (2.16), Δ is the change from one move block to the next with the first change measured from the rates applied now, $w_m = 0.03$ per (kg/tHM)^2 against 1 per °C^2, $\lambda = 1$ per (kg/tHM)^2, and $s_v = 2$ and $r_v = 0.02$ per (kg/tHM per hour)^2. There is no terminal weight. The moves are held in twelve blocks, four of one period and two of two, then 3, 4, 6, 9 and 15 periods and a final block of 3, which gives 24 decisions. Because the coal target enters as a resting rate for the coal, a sustained load can be carried only by the coke, which is the mid-ranging idea written as a cost instead of a second loop. Every constraint is hard: the coal stays within the 80 to 200 kg/tHM clamps of Equation (2.13), $|v_k| \le r_c$, and the coke stays within 300 to 420 kg/tHM at the end of every block. The coke rate advances by $v_k T_{\mathrm{MPC}}$ over a period, and the staircase the ramp makes at the three scans enters the dead-time chain as its mean. The temperature has no soft band. A cold-hearth band at 1465 °C, carried by the solver's single slack over the whole horizon, raised the IAE of case 6 from 583 to 605 °C h and its coal travel from 399 to 444 kg/tHM, with the coldest hearth at 1457.4 against 1457.2 °C.

The original design asked for a 5 min period and 144 steps. That horizon exceeds the 120 steps the shared solver supports, and with a 5 min period and 120 steps the model needs 76 states and about 66 ms per plan in Node, about three processor cores at the top playback speed of 14400 times real time, for an IAE of 621 instead of 644 °C h on scenario 10 with half again as much coal and coke travel. The weights were chosen once, on scenario 10, from 216 combinations of horizon, λ, w_m and s_v. Of the 43 that moved neither fuel further than the coal loop does, rejected no plan and never held the coal at its limit, the best reached 627 °C h, and among those within 5 % of it the one that kept the coal closest to its target was taken, the shorter horizon breaking ties.

The controller keeps a plan when the solver's own status is optimal, the returned moves, predictions and cost are all finite, each planned coal rate, ramp and coke rate over the horizon respects its bounds, and the first planned move equals the one applied. Otherwise the coal loop and mid-ranging take the fuels from the rates they have been tracking until a later plan passes the same checks. With the wander and noise switched off, forcing a rejection at 20.25 h of scenario 10 moves the coal 0.09 kg/tHM

on the soft sensor, which is the coal loop's own first move, and not at all on cast dips, where the loop holds the coal until the next cast. None of the runs reported here rejected a plan, and each makes 288 plans in three days. A return from manual is a decision rather than a transfer. After three hours of manual operation across the arrival of the first burden change, the first plan raises the coal 16.0 kg/tHM at once.

The plant runs on the run's parameters, the instruments add their own lag, noise and drift, and the controller's model keeps the declared values, so the three models stay separate. Replaying the instrument record of six three-day runs, each over a different hidden plant, reproduces every applied coal and coke rate bit for bit. With no bound active the plans agree with an independently written condensed solution to 6.4×10^{-14}, and Karush-Kuhn-Tucker certificates hold at eight instants that exercise every kind of constraint. A plan takes 4.2 ms on average in Chrome, which is 1.7 % of one core at the default playback of 3600 times real time and 6.8 % at the top speed.

2.5 Results

The ten teaching scenarios and the challenge are listed in Table 2.3 with what each teaches and what a learner must do to complete it, and their measured outcomes are given in Table 2.4. Each case is credited when the learner performs the listed actions and the simulation reaches the listed state. The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt, and the credit from the eleven cases (ten scenarios and the challenge) accumulates toward certificates of progress and of completion, as Section 1.6.4 describes. Every value below comes from the application's engine at its default seed, with each setting changed after the case is loaded and before the run starts, which gives the same run as the setting carried through Reset.

2.5.1 Two Actuators on Two Clocks

Case 1 runs in manual with the wander off. The coke step of 10 kg/tHM reaches about 19 °C fifteen hours later, close to the steady 20 °C of Equation (2.2), and the coal step of the same size dips before it rises, as listed in Table 2.4. The coke step can be fitted by a first-order model with dead time to 0.97 % RMS error, a little later and slower than Equation (2.2) alone because the recipe change limit ramps the step in. The coal step cannot be described by that model, because its first response has the wrong sign.

> **TRY IT IN THE APP** **1 · Step the Coke and the Coal by Hand**
>
> Load the scenario and press Run, then watch the Hot metal temperature trend. The coke step at one hour does nothing for five hours, and fifteen hours later the hearth is up about 19 °C. After the coal step at twenty hours the hearth first falls about 3.6 °C within the hour, then rises about 12 °C.

2.5.2 A Loop That Sees the Hearth Once a Cast

In case 2 a harder ore charged at twelve hours reaches the hearth four hours later. The loop cannot see it earlier and the coal needs an hour to act, so the hearth falls to about 1475 °C before the loop recovers it, and the coal ends the two days about 20 kg/tHM higher. The true temperature at the casts is held to about 10.3 °C RMS. Case 3 raises the gain from 0.3 to 0.8 and halves the integral time, as a loop on a continuous measurement might be tuned, and the hearth swings from about 1439 to 1543 °C, spends 2.9 h below the cold alarm and holds the coal on a limit for 8 h.

The sampling and the coal dynamics each contribute to that ceiling, and Table 2.5 separates them. The continuous measurement is the soft sensor with no lag, no drift, no bias correction and a noise of 10^{-9} °C, which keeps the random sequence of the cast run, so it reads the true hearth every five minutes. At the default tuning the two-hour cast costs 10.3 against 7.8 °C RMS, and the SIMC estimate of Equation (2.15) performs about the same as the default, with 9.9 °C RMS on the casts. At the case 3 tuning the hold decides how cold the hearth goes. The continuous loop stays above the alarm, with a minimum of 1458 °C, while the cast loop falls to 1439 °C and spends 2.9 h below it.

Table 2.3 Scenarios and Completion Requirements

#	Scenario, teaching point and kind	Required to complete
1	Coke and coal stepped by hand *two response times* (experiment)	Run in manual past 16 h, by which time the coke step at 1 h has raised the hearth above 1510 °C, and on past 21.5 h to cover the first 1.5 h after the coal step at 20 h.
2	Coal loop on casts, new ore blend *sampled feedback* (experiment)	Run in automatic through the ore change at 12 h and past 36 h, by which time the coal rate exceeds 165 kg/tHM.
3	Tuned as if continuous *the gain ceiling* (demonstration)	Run past 24 h at $K_c \geq 0.75$ with the hearth falling below 1460 °C, then Reset, return to $K_c = 0.3$ and $\tau_I = 2$ h before 1 h and run past 24 h again.
4	Soft sensor drifting, uncorrected *an unchecked model* (demonstration)	Run on the uncorrected soft sensor through the drift start at 6 h to 24 h, by which time the hearth is more than 15 °C below setpoint.
5	Soft sensor corrected each cast *bias updating* (experiment)	Run to 24 h with the correction weight $\alpha = 0.5$, then Reset, set $\alpha \leq 0.25$ before the drift starts at 6 h and run to 24 h again.
6	Burden change beyond the limit *saturation* (demonstration)	Hold the coal on its 200 kg/tHM limit for 2 h after the second burden change at 36 h, and run the full 72 h.
7	Mid-ranging on coke *the slow actuator* (experiment)	Run 72 h with mid-ranging at $K_v \leq 0.21$ until the coke rate exceeds 365 kg/tHM, then Reset, set $K_v = 0.5$ before the first burden change at 2 h and run 72 h again.
8	A high coal target *economics and margin* (experiment)	Run 72 h with mid-ranging and a coal target of at least 189 kg/tHM, then Reset, set the target to 150 kg/tHM before 1 h and run 72 h again.
9	Hot stove fault *a fast disturbance* (demonstration)	Run past 14 h at the default gain, through the fault at 12 h, then Reset, raise K_c to 1.0 before 1 h and run past 14 h again.
10	Coke planned five hours ahead, MPC *planning on a delay* (experiment)	Run past 36 h in automatic with the MPC, then Reset, select the coal loop before the first burden change at 2 h and run past 36 h again.
11	Challenge *graded run* (challenge)	At 72 h the coldest hearth of the run is at or above 1450 °C and each of the last 12 cast dips, from 48 to 70 h, is within ±15 °C of 1500 °C, with the model parameters and disturbances as given.[a]

[a] With the shipped settings (uncorrected soft sensor, mid-ranging off) the hearth falls to 1337 °C, 113 °C below the alarm, and the worst last-day dip is 162.0 °C off setpoint. With the suggested settings (bias correction at $\alpha = 0.8$, mid-ranging on, $K_c = 0.25$ and $\tau_I = 3$ h) the coldest hearth is 1459.1 °C, 9.1 °C above the alarm, and the worst dip is 12.8 °C off, 2.2 °C inside the band.

Setting the soft sensor noise to exactly zero, as a learner would, removes its draws from the random sequence and so changes the furnace wander, and the continuous loop then just touches the alarm, at 1449 °C for 6 min, with 18.6 °C RMS. Neither loop is stable at that tuning. With the wander and the noise off and a 5 °C burden step, the loop with $\tau_I = 1$ h settles at $K_c = 0.65$ on the continuous measurement but not at 0.70, and on the casts it settles at 0.45 but not at 0.50, so the hold lowers the stable gain by a factor of about 1.4, against the 1.7 of the SIMC estimate. At $K_c = 0.8$, with the case's own ore change and the wander and noise still off, both loops lock into a cycle that the 200 kg/tHM coal limit bounds, between 1464 and 1528 °C on the continuous measurement and between 1446 and 1537 °C on the casts from 24 to 96 h, and only the larger cycle reaches the alarm. A gain of 1.2 swings the hearth to 1400 °C even on the ideal measurement. With the default tuning the cast error is

Table 2.4 Measured Scenario Outcomes. RMS Is of the True Temperature Minus Setpoint at Every Cast

Case and condition	Result
1, coke +10 kg/tHM	no change for 5 h, +19 °C after 15 h
1, coal +10 kg/tHM	−3.6 °C within 1 h, then +12 °C
2, default, ore −30	min 1475 °C, coal +20, 10.3 °C RMS
3, $K_c = 0.8$, $\tau_I = 1$ h	1439–1543 °C, 2.9 h cold, 21.6 °C RMS
4, drift, uncorrected, 24 h	index 1514, hearth 1464 °C, coal 114
4 and 5, cast error	15.1 against 6.5 °C RMS
5, corrected, 24 h	hearth 1492 °C, index 1503 °C
6, two changes, no mid-ranging	coal on limit 26 h, last 12 h 1477 °C
7, mid-ranging	never on limit, last 12 h 1504 °C
7, end of run	coke 382, coal 171 kg/tHM
6 and 7, fuel	$1.68 against $1.73 million per day
6 and 7, first change	min 1448 °C in both, 0.3 h cold
8, targets 190 and 150, no ore change	$1.59 against $1.63 million per day
8, ore −25	coal on limit 2 h, min 1462 against 1464 °C
9, stove −80 °C	min 1473 (loop), 1480 ($K_c = 1$), 1470 (none)
9, stove back	overshoot to 1522 °C

Table 2.5 Case 2 Over 48 h With Cast Dips and With an Ideal Continuous Measurement. Same Random Sequence in Each Pair

		Cast dips		Continuous	
K_c	τ_I (h)	RMS	min	RMS	min
0.30	2.0	10.3	1475	7.8	1479
0.26	1.5	9.9	1475	7.7	1479
0.60	1.5	9.5	1473	7.6	1474
0.80	1.0	21.6	1439	18.1	1458
1.20	1.0	31.1	1387	36.9	1400

Case 3 is the $K_c = 0.8$ row.

7.7, 10.3 and 10.2 °C RMS for casts every 1, 2 and 4 h, and the hearth error over the whole run ranks the three the same way, at 7.6, 11.1 and 10.3 °C RMS. Casts every hour help, while on this seed a four-hour interval costs no more than a two-hour one. Volume 5, Chapter 3 meets the same limit with a composition loop closed through a slow analyzer.

2.5.3 A Soft Sensor Is a Model

Case 4 puts the loop on the soft sensor and starts a drift of 3 °C per hour at six hours. The loop pulls the index toward setpoint by removing coal, and after a day the index reads about 1514 °C while the hearth is at about 1464 °C and the coal has been cut to about 114 kg/tHM. The cast dips show the truth throughout and nothing uses them. In case 5 the same drift is corrected at every cast with $\alpha = 0.5$, the hearth ends the day at about 1492 °C, and the cast error is 6.5 °C RMS against 15.1 °C uncorrected. Over the same day a weight of 0.2 gives 9.6 °C RMS and a weight of 1.0 gives 5.1 °C, so with this drift rate the drift outweighs the 6 °C dip noise that a large weight passes into the estimate.

> **TRY IT IN THE APP** **4 · A Soft Sensor That Drifts, Uncorrected**
>
> Run the scenario for a day. From six hours the soft sensor reads 3 °C per hour high. On the Hot metal temperature trend, compare the soft sensor with the cast dips: at 24 h the index reads about 1514 °C and the hearth about 1464 °C, and the Coal injection trend shows the coal cut to about 114 kg/tHM. Then load scenario 5, where each cast corrects the index.

Pause. Try Exercise 2.2 before reading on.



Figure 2.2 Mid-ranging on coke, case 7, at the end of the three-day run ($t = 72$ h), cropped to the plant view and the trends. The furnace section on the left shows the coke and ore layers with a 5.0 h descent (1), the coal at the tuyeres and the hearth at 1506 °C with 0.49 % silicon (2). Beside it the last twelve casts are plotted as dips and soft sensor readings against the ±15 °C band and the cold-hearth line (3), above gauges for the coal at 171 kg/tHM against its target and limit and the coke at 382 kg/tHM (4). The trends show the hearth dipping after each burden change, the deeper dip near 11 h marked as a cold-hearth event (5), the coal stepping at every cast to a peak of 192 kg/tHM, below its 200 kg/tHM limit, and then easing back toward its target, and the coke climbing from 350 to 382 kg/tHM.

2.5.4 Saturation and Mid-Ranging

Case 6 charges two burden changes a day and a half apart, 45 °C and then 80 °C of extra demand in all. At the net coal gain the two together need about 67 kg/tHM of extra coal against 50 kg/tHM of room, so the capacity test of Equation (1.3) fails and the coal sits on its 200 kg/tHM limit for about 26 h. Over the last 12 h the hearth averages about 1477 °C, and integral action cannot remove that offset because the actuator cannot reach it. Case 7 switches on mid-ranging. The coal peaks at 192 kg/tHM, below its limit, the hearth averages about 1504 °C over the last half day, and at the end the coke has risen to 382 kg/tHM with the coal back to 171 kg/tHM and still returning, as shown in Figure 2.2. The fuel costs about $1.73 million a day against $1.68 million without mid-ranging, which is the price of a hearth that is not cold. Neither case prevents the first dip. The first burden change reaches the hearth at six hours, and the coal, an hour late itself, cannot answer it in time, so in both cases the hearth falls to about 1448 °C and spends 20 min below the cold alarm.

Table 2.6 Mid-Ranging Tuning in Case 7 Over 72 h

K_v	τ_V (h)	RMS	max T	peak CR	time cold (h)
0.2	12	15.9	1519	382	0.33
0.2	4	18.2	1530	393	0.35
0.5	12	19.0	1541	390	0.52
1.0	12	37.6	1563	420	6.97

A faster mid-ranging loop overshoots, as reported in Table 2.6. At $K_v = 0.5$ the hearth overshoots to 1541 °C and the cast error rises from 15.9 to 19.0 °C RMS. At $K_v = 1.0$ the two loops fight through the coke delay, the coke reaches its 420 kg/tHM limit, the coal spends 8 h on a limit and the hearth spends 7.0 h cold. The coke rate limit plays almost no part at the default tuning. It is active for 3.7 h of the run, and raising it to 20 kg/tHM per hour moves the last half-day mean by 0.02 °C, from 1503.78 to 1503.76 °C.

> **+ DID YOU KNOW? A Faster Helper Makes the Hearth Colder**
>
> From the simulator: in Table 2.6, raising the mid-ranging gain from 0.2 to 1.0 at the same 12 h integral time raises the cast error from 15.9 to 37.6 °C RMS and the time below the cold-hearth alarm from 0.33 to 6.97 h. A coke loop that hurries fights the coal loop through the five-hour delay.

2.5.5 The Coal Target and a Fast Disturbance

The coal target is an economic setting. With the burden change removed and the wander left on, over three days a target of 190 kg/tHM runs the furnace on about $1.59 million of fuel a day against $1.63 million at 150 kg/tHM, a saving of about $33,000 a day. It also leaves the coal only 10 kg/tHM of room, while a 25 °C burden change asks for about 21. With that change the coal sits on its 200 kg/tHM limit for one cast, 2 h, and ends the run at about 193 kg/tHM while the coke catches up, and the coldest hearth is about 1462 °C against 1464 °C at the lower target, for $1.62 against $1.65 million of fuel a day. The target is set where the price of fuel meets the size of the disturbances the furnace actually sees.

A stove fault that drops the blast by 80 °C for eight hours shows the limit of feedback on this plant. Before the stove returns the hearth falls to about 1470 °C with no loop, 1473 °C with the default loop and 1480 °C with a gain of 1.0, because the blast reaches the hearth within the hour and the loop sees it only at the next cast. When the stove returns, the coal the loop has added is still arriving and the hearth overshoots to about 1522 °C. The hotter loop overshoots to 1524 °C and then swings down to 1461 °C at 28.75 h, colder than anything the fault itself did. The blast temperature is measured and its effect is known from Equation (2.6), which makes it a candidate for feedforward rather than for tighter feedback.

2.5.6 The Graded Challenge

The challenge runs three days with a 40 °C burden change at six hours, a 3 °C per hour soft sensor drift from twelve hours, and a second change to 75 °C in all at 24 hours. It starts on the uncorrected soft sensor with mid-ranging off. It is met when the hearth never falls below 1450 °C and every cast dip in the last day, the twelve casts from 48 to 70 h, is within 15 °C of setpoint. Table 2.7 reports the outcomes.

As loaded, the drifting index walks the hearth to about 1337 °C, a chilled hearth. Correcting the index alone keeps the hearth above 1453 °C but leaves the coal on its limit for 28.5 h and the last half day 18 °C cold. Mid-ranging alone makes matters slightly worse, because the coke then follows a coal rate that is itself answering a false measurement. The two settings together keep the coldest hearth at 1456 °C, but at a weight of 0.5 the bias estimate trails the drift, the last day runs 7.8 °C cold on average, and the worst last-day dip is 20.6 °C off. The suggested settings raise the weight to 0.8 and

Table 2.7 Challenge Outcomes Over 72 h. Worst Dip Is the Largest Deviation of a Cast Dip From Setpoint in the Last Day, as the Challenge Scores It

Configuration	min T (°C)	on limit (h)	last 12 h (°C)	worst dip (°C)	fuel ($M/d)	met
as loaded	1337	3.7	1353	162.0	1.60	no
bias correction	1453	28.5	1482	25.2	1.68	no
mid-ranging	1328	0	1344	170.7	1.59	no
both	1456	0	1493	20.6	1.71	no
suggested	1459	0	1500	12.8	1.72	yes
cast dips, mid-ranging	1454	0	1502	12.1	1.72	yes

Suggested: bias correction at $\alpha = 0.8$, mid-ranging on, $K_c = 0.25$, $\tau_I = 3$ h. Both: bias correction at $\alpha = 0.5$ and mid-ranging, with the coal loop at its default tuning.

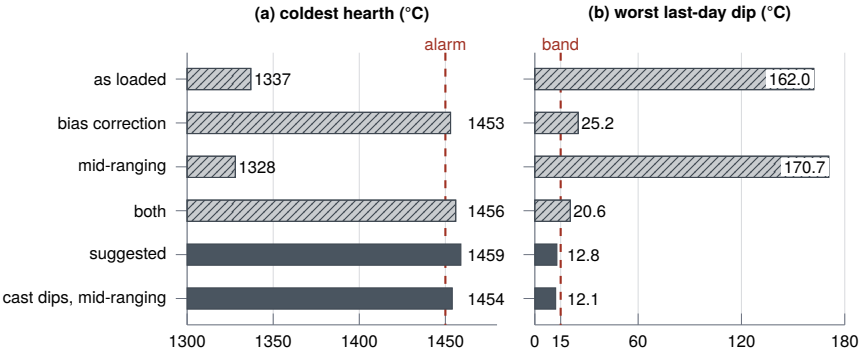


Figure 2.3 The graded challenge of Table 2.7 as a chart, runs that meet the target solid and runs that miss it hatched. The dashed lines are the 1450 °C cold-hearth alarm in (a) and the 15 °C band in (b). Notice that bias correction lifts the coldest hearth above the alarm but leaves the last day outside the band, that mid-ranging alone makes both slightly worse, and that only the suggested settings and the cast dips with mid-ranging clear both lines.

soften the coal loop to $K_c = 0.25$ and $\tau_I = 3$ h, which at the same weight cuts the total coal travel over the run from 611 to 508 kg/tHM, and they meet the target with a minimum of 1459 °C and a worst last-day dip of 12.8 °C. Either refinement alone leaves a worst dip of 15.8 to 16.6 °C. Abandoning the soft sensor for the cast dips, with mid-ranging on, also meets the target at 1454 °C and 12.1 °C, which shows that the corrected index buys margin above the alarm rather than feasibility. The dips carry 6 °C of noise, so the last-day criterion is decided by a few degrees. The true temperature at those casts is never more than 7.3 °C from setpoint with the suggested settings and 8.1 °C with the cast dips. Figure 2.3 draws the coldest hearth and the worst dip of each configuration as bars.

2.5.7 Comparison of the Predictive Controller With Mid-Ranging

The MPC is compared with the coal loop in Table 2.8. Every run lasts 72 h at the default seed with the instruments and noise as the case loads them. The comparator is the challenge's suggested coal loop, $K_c = 0.25$ and $\tau_I = 3$ h, on the soft sensor corrected at $\alpha = 0.8$, with mid-ranging at its default tuning, which is the line-up scenario 10 runs when the coal loop is selected. The MPC runs on the same corrected index with scenario 10's settings and is not retuned for any case, and neither structure is told about the burden. The hindsight row is the best of 768 coal-loop tunings (K_c, τ_I, K_v, τ_V and α) at the case's own coal target, chosen after seeing that case. It bounds what a fixed loop could do instead of offering a design, and a wider search inside the same ranges reaches 494 °C h on scenario 10. IAE

Table 2.8 The MPC Against the Suggested Coal Loop With Mid-Ranging and a Coal Loop Retuned in Hindsight. Each Run 72 h at the Default Seed

Case	Structure	IAE (°C h)	Coldest (°C)	At 200 (h)	Off target (kg h/tHM)	Travel (kg/tHM)	Coke (t)
10	coal loop	1016	1453.5	0.50	2138	506/106	11199
	hindsight	512	1456.6	0	1622	1168/296	11363
	MPC	644	1457.2	0	670	385/104	11594
	MPC, coal only[a]	667	1457.2	5.50	2018	360/94	11184
	MPC, coke delay 15 min[b]	1367	1450.9	0	476	275/294	11612
6, 7	coal loop	699	1458.0	0	1629	504/104	10984
	hindsight	504	1456.6	0	2006	1170/118	10915
	MPC	583	1457.2	0	605	399/103	11292
8	coal loop	648	1463.9	0	948	500/105	10094
	hindsight	385	1470.4	0	1384	1150/116	10331
	MPC	403	1473.0	7.25	534	322/105	10034
5	coal loop	459	1479.5	0	154	504/99	10471
	hindsight	351	1485.3	0	99	312/31	10496
	MPC	467	1472.9	0	267	335/69	10451
2	coal loop	470	1477.4	0	688	496/99	10709
	hindsight	349	1477.7	0	375	1147/295	10875
	MPC	415	1476.5	0	344	327/77	10840
11[c]	coal loop, met	694	1459.1	0	1520	508/103	10941
	hindsight, not met	558	1460.8	0	1338	838/88	11053
	MPC, not met	709	1458.6	0	534	374/94	11229
	MPC on cast dips, met	644	1452.3	0	536	302/87	11267

At 200 is the time with the coal at its injection limit, travel is the summed coal and coke movement, and coke is the tonnage charged. Every MPC run made 288 plans and rejected none. [a]The MPC planning the coal with mid-ranging moving the coke. [b]The MPC's model, free response and record with 15 min of coke dead time instead of 5 h. [c]The challenge with its suggested settings, judged by its own rule.

integrates the absolute error of the true hearth temperature, and the coal's distance from its target is integrated in the same way.

Scenario 10 charges a burden 45 °C harder at 2 h and 90 °C harder in all at 10 h. At the net coal gain the second change asks for about 75 kg/tHM of extra coal against 50 kg/tHM of room below the injection limit, so the coke must take part of the load. Under the MPC the hearth first falls below 1495 °C at 6.18 h, the coke ramp reaches the recipe limit at 8.02 h and stays there for 3.75 h in all, and when the second change reaches the hearth at 14 h the coke is at 375.0 against 357.2 kg/tHM under mid-ranging, as shown in Figure 2.4. The coal peaks at 192.0 kg/tHM at 10.0 h, never reaches its limit, and is back to 157.1 kg/tHM at 24 h and 147.9 kg/tHM at 36 h. Under the coal loop the coal reaches the limit at 21.7 h, still stands at 191.8 kg/tHM at 36 h, and ends the three days at 167.9 against 146.3 kg/tHM under the MPC. The IAE is 644 against 1016 °C h and the coal's distance from its target 670 against 2138 kg h/tHM. The earlier coke costs 395 t more over the three days, $18,532 a day at the default prices.

Removing one mechanism at a time locates these gains. An MPC that plans the coal alone while mid-ranging moves the coke reaches 667 °C h, so most of the temperature gain over the coal loop belongs to the coal plan. What the coke plan buys is room for the coal, because without it the coke reaches 370 kg/tHM at 22.4 h instead of 11.6 h and the coal spends 5.5 h at its limit, and with nobody moving the coke the IAE is 2020 °C h with 56.5 h at the limit. The price on the coal's distance from its target is what hands the load to the coke. With $w_m = 0$ the plan raises the coke less, holds the coal at its limit for 7.25 h and reaches an IAE of 742 °C h. Knowing the delay matters more than the horizon. A model that gives the coke 15 min of dead time instead of five hours roughly doubles the IAE, to 1367 on scenario 10 and 1329 on cases 6 and 7, and nearly triples the coke travel, while horizons of 8, 20 and 30 h all give 644 °C h.

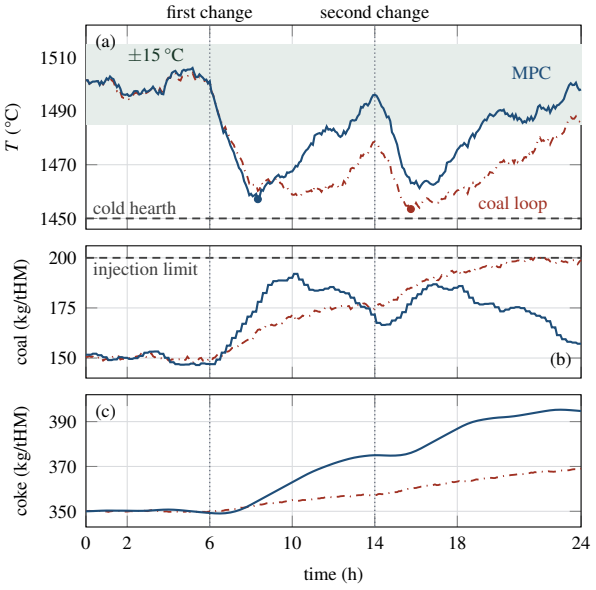


Figure 2.4 Scenario 10 over its first day under the MPC and under the coal loop with mid-ranging, at the default seed. The MPC is the thick solid line and the coal loop the dash-dotted line. The cold-hearth alarm and the injection limit are dashed, and the dotted verticals mark the arrival of each burden change. (a) True hot metal temperature against the ±15 °C band and the cold-hearth alarm, with each coldest point marked. The burden changes charged at 2 and 10 h reach the hearth at 6 and 14 h. (b) Coal rate against the injection limit. (c) Coke rate. At 24 h the MPC holds the coal at 157.1 kg/tHM with the coke at 394.7 kg/tHM, while the coal loop holds it at 198.9 with the coke at 369.1 kg/tHM.

The events of scenario 10 were chosen from nineteen candidate sets built around the same two changes, with second changes of 80 or 90 °C at 10 to 16 h, coal targets of 150 to 180 kg/tHM, and days of three and four changes. On all nineteen the MPC has the lower IAE, by 16 to 49 % with a median of 34 %, and scenario 10's 37 % ranks sixth, so the case is typical of the family rather than its best member. On three further seeds of the wander and the noise the MPC stays ahead on scenario 10, at 665 to 757 against 849 to 1075 °C h, with its coal never at the limit.

The coal loop retuned in hindsight holds the hearth closer than the MPC on every case of Table 2.8, with 2.2 to 3.6 times the MPC's coal travel except in case 5, where it moves the coal less, 312 against 335 kg/tHM, and still wins. Held to the MPC's own coal and coke travel, the best retuned loop still wins cases 2, 5 and 10 and the challenge (399, 351, 620 and 612 °C h), ties case 9 (423 against 420) and loses cases 6 and 7 (666) and case 8 (438), and on scenario 10 it then spends 5.4 h with the coal at its limit. In case 8 the MPC rests the coal at the 190 kg/tHM target it is given, cutting the coke to between 324.5 and 330.6 kg/tHM from 6 to 24 h to do so, which leaves the coal 10 kg/tHM of room, and after the burden change at 24 h the coal spends 7.25 h at its limit where the coal loop spends none, although the MPC's IAE is lower, 403 against 648 °C h. Case 5 is lost narrowly, 467 against 459 °C h, on all four seeds. Both structures act there on an index that reads 6.5 to 6.6 °C above the hearth over 24 to 72 h and so hold the hearth 6.6 to 6.8 °C cold, and the MPC trails by 9 and 14 °C h over 6 to 24 and 24 to 48 h before leading by 15 over the last day. On cases 2 and 9 the order changes with the seed, so neither is counted as a win. After the second change of cases 6 and 7 the MPC's hearth dips to 1464.7 against 1476.9 °C, with window IAEs of 237 and 236, because the coal loop is still running hot from the first change, 14.6 °C above setpoint on average over the four hours before the second one arrives against 9.6 under the MPC.

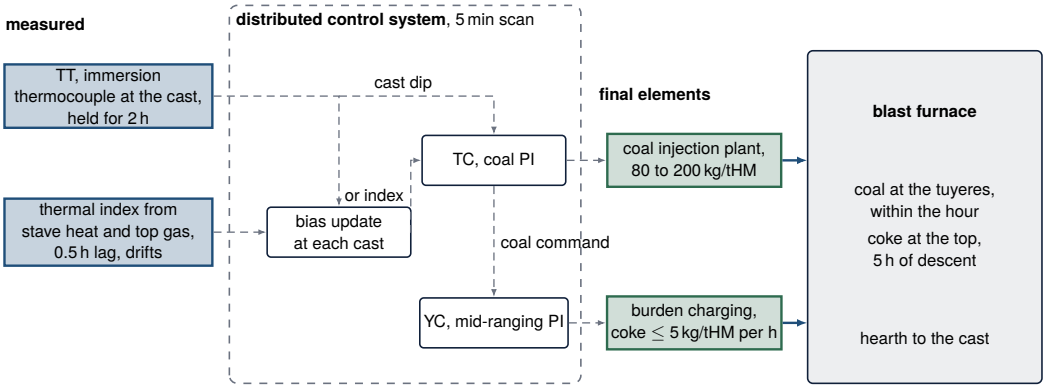


Figure 2.5 The loops of the simulated furnace drawn with their hardware. The coal controller TC reads either the cast dip or the corrected thermal index and sets the coal injection plant. The mid-ranging controller YC takes the coal command as its measurement and moves the coke through the burden recipe, so the coke takes over a lasting load while the coal returns to its target.

Run under the MPC with the suggested settings of the challenge, the furnace misses its target. The coldest hearth is 1458.6 °C and no unscored limit is broken, but the worst last-day dip is 23.7 °C off setpoint. Over the last day the corrected index reads 6.7 °C above the hearth on average, because the drift adds 6 °C between casts and $\alpha = 0.8$ removes 80 % of the error at each, and the MPC holds that index 0.4 °C under setpoint, so the hearth runs 7.1 °C cold. The suggested coal loop meets the target because it is slower. Its index is still 6.0 °C above setpoint over the last day, which leaves the hearth 0.6 °C cold, and continued to 144 h it settles with the index 0.1 °C under setpoint and the hearth 5.4 °C cold, where the MPC ends 0.9 °C under setpoint on the index and 6.2 °C cold. On cast dips, which do not drift, the MPC meets the challenge at the default seed, with a coldest hearth of 1452.3 °C and a worst dip of 12.7 °C, where the suggested coal loop on cast dips misses it at 16.3 °C. The verdicts turn on a few degrees of dip noise, since over four seeds the suggested coal loop meets the target three times, the MPC on the corrected index once and the MPC on cast dips twice. As loaded, on the uncorrected index, the MPC walks the hearth down as the coal loop does, to 1332.6 against 1337.5 °C.

2.6 Technology in Practice

Figure 2.5 draws the two loops of Section 2.4 with their hardware.

What Measures It

A thermocouple generates its own voltage at the junction of two metals. Thermocouples are typically less accurate than resistance thermometers but more rugged, with wider spans, and the types with platinum and rhodium wires are the ones listed for the highest temperatures [24]. The model's dip noise of 6 °C and the reading held until the next cast stand for the scatter of the dips and the stepwise record a plant computer keeps.

> **TECHNOLOGY** **Sensor: The Immersion Thermocouple at the Cast**
>
> The hot metal temperature is read by an immersion thermocouple dipped into the iron runner at the exit of the taphole. Two or three readings are taken in a cast, and the first tends to be low, because the first metal comes from the cold lower hearth near the walls [18]. The trough cools the metal as well. A reading past the skimmer ran about 30 °C low early in a cast, so one furnace took its dip only after a set quantity of metal had passed [10]. A thermocouple fails most often open, which reads as wild noise unless the instrument has a burnout mode [24].

The thermal index of the soft sensor is computed from stave heat loads and top gas analysis. The model represents it as a copy of the hearth temperature with a 0.5 h lag, a drift and noise, and the laboratory silicon arrives with each cast.

What Moves It

The coal is the fast final element, and the controller's command reaches the tuyeres through the injection plant.

> **TECHNOLOGY** **Final Element: The Coal Injection Plant**
>
> Dried and ground coal is conveyed by air to pressurized storage beside the furnace, and metering equipment there sends it pneumatically to each tuyere, 28 per furnace in one demonstration [19]. At the tuyere the coal enters through a lance inserted into the blowpipe that carries the hot blast. Coal that does not burn in its short stay in the raceway leaves char that rises and can block the void spaces of the burden, and a black layer was found inside lances after a period of service [20]. The coal rate a controller sets is delivered only while every line and lance stays clear.

The model applies the commanded coal rate at once, clamped between 80 and 200 kg/tHM. The coke rate is moved through the burden recipe at the top, and the model lets it follow its command at no more than 5 kg/tHM per hour. The hot blast enters as a disturbance.

Where the Controller Runs

Control loops in a plant are often nested or cascaded, with cycle times from milliseconds to minutes, and a distributed control system uses a supervisory loop to mediate a group of localized controllers [25]. In the app the coal loop and the mid-ranging loop are scanned every 5 min, the coal loop on cast dips executes once a cast, 2 h apart, and the predictive controller plans every 15 min. On such a platform the two loops are PID function blocks of the regulatory layer, the mid-ranging block taking the coal block's output as its measurement, and the bias update is a calculation block triggered by each cast. The predictive controller sits in the supervisory layer and writes both fuel rates.

> **IN PRACTICE** **Bringing a Hearth Loop Into Service**
>
> The computer control of the 1970s furnace in Section 2.2 went into service in three steps: an operation guide that showed the predicted temperature from September 1975 to January 1977, a test of automatic control from February to October 1977, and practical use from November 1977. The reasons given were that a blast furnace needs a long time to recover once it is in trouble, and that the experienced operators had to gain confidence in the system. The controlled temperature was the average of three consecutive casts, which smoothed the scatter between tapholes and the measurement errors [10].

What Is Modernizing It

The first development is data-driven prediction of the thermal state. Neural network models of the hot metal temperature take the elapsed times of the casting cycle as inputs, so that the irregular dips can be used [18], and a gradient-boosted model predicted the average silicon of the coming cast from the previous shift's operating data with an accuracy of 91 % over 200 predictions [26]. The second is simulation of the injection itself. Three-dimensional flow models of the blowpipe, lance, tuyere and raceway have been built to find how far the coal rate can rise, with long-term targets near 200 kg per tonne, 40 % of the fuel [20], the same figure as the injection limit of this chapter's coal. The third is a look inside the shaft. Muon radiography uses cosmic-ray muons to image the average density of the furnace and is being tested to follow the cohesive zone, where the descending burden softens and melts [27].

2.7 Limitations

The thermal model is linear with fixed delays. The descent time varies with production rate on a real furnace, and the coke replacement ratio of coal is not constant at high injection rates [7].

The coal path is written as two lags, one of them delayed. The raceway physics, including oxygen enrichment and the flame temperature limit that bounds injection in practice, is represented only by the fast negative path and the fixed 200 kg/tHM limit [6].

Silicon is a function of temperature, so the simulator cannot teach the separate prediction and control of silicon that much of the literature addresses [5, 17].

The soft sensor error is a ramp plus noise. A real thermal index errs when the furnace changes in ways its model did not see, and a bias update then corrects some of that error and not the rest [21].

There is no feedforward from the stoves or the burden, and no anti-windup coupling between the coal and coke loops [12].

The MPC's model leaves out what its single output disturbance must absorb. It does not read the burden, which on a real furnace is decided when it is charged, four hours before it reaches the hearth, so like the loops it cannot prevent the first cooling a burden change brings. It plans with the declared gains and delays, so a changed parameter is absorbed only as an offset. On the soft sensor it is offset-free on the index, so the drift that the cast correction leaves between casts becomes hearth error, 7.1 °C on average over the challenge's last day. A coal target near the injection limit leaves the plan no room, as case 8 shows, and the shared solver's single slack over the whole horizon is why the plan carries no soft temperature band.

> ⚠ **OPEN CHALLENGE** **Reading the Hearth From Everything Measured**
>
> Silicon in the hot metal serves as a proxy for the heat in the hearth, because the hostile interior rules out measuring its temperature directly. The furnace answers an operator's move late, and the laboratory analysis adds its own delay, which hampers real-time control. Physical models can be slow or not accurate enough, data-driven predictors mostly use a few hand-picked variables, and each furnace offers a different set [26]. Progress would be a predictor that a new furnace can adopt with little hand work, trusted enough to close the loop on, which the soft sensor of this chapter only stands in for.

2.8 Conclusion

The blast furnace hearth is held by a fast, cheap and bounded fuel and a slow, expensive and wide one, and it is seen once a cast. In the simulator of this chapter a coal step cools the hearth 3.6 °C before warming it 12 °C, and a coke step waits five hours. A coal loop tuned for the two-hour cast holds 10.3 °C RMS, while a gain of 0.8 with a one-hour integral time swings the hearth from 1439 to 1543 °C on the casts, against a minimum of 1458 °C on an ideal continuous measurement, and the hold lowers the stable gain by a factor of about 1.4. An uncorrected soft sensor drifting 3 °C per hour takes the hearth to 1464 °C while reading 1514 °C, and a bias update at each cast holds it at 1492 °C. When burden changes exceed the injection limit the hearth runs 23 °C cold, and a slow mid-ranging loop on coke restores it to about 1504 °C for about $44,000 a day of extra fuel. In the challenge a bias update weighted at 0.8, mid-ranging and a gentler coal loop together lift the coldest hearth from 1337 °C to 1459 °C. The predictive controller of Section 2.4.3, which schedules the coke on its five-hour delay, keeps the coal off its injection limit in scenario 10, with an IAE of 644 against 1016 °C h and no time at the limit against 0.5 h, for 395 t more coke over three days, while a coal loop retuned in hindsight reaches 512 °C h by moving the coal three times as far. The model keeps the delays fixed and locks silicon to temperature. A fuller model would add production-rate dependent delays, a silicon model, stove feedforward and anti-windup between the two loops. The furnace remains a compact case of dividing a load between a fast actuator and a slow one.

KEY POINTS

- The hearth is warmed by a fast, cheap fuel with little room, coal, and a slow, costly one with plenty, coke, and is checked once a cast.
- The two-hour hold of the cast lowers the stable gain of the coal loop by a factor of about 1.4.
- An uncorrected soft sensor drifting 3 °C per hour takes the hearth to 1464 °C while reading 1514 °C, and a bias update at each cast holds it at 1492 °C.
- When burden changes exceed the injection limit, a slow mid-ranging loop on coke restores the hearth to about 1504 °C.

Where else this applies. Volume 4, Chapter 10 and Volume 4, Chapter 9 divide a load the same way, and Volume 5, Chapter 10 corrects a soft sensor against late samples.

References

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

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

[3] B. Joseph and C. B. Brosilow, “Inferential control of processes: Part I. Steady state analysis and design,” *AIChE Journal*, vol. 24, no. 3, pp. 485–492, 1978.

[4] P. Kadlec, B. Gabrys, and S. Strandt, “Data-driven Soft Sensors in the process industry,” *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

[5] H. Saxen, C. Gao, and Z. Gao, “Data-Driven Time Discrete Models for Dynamic Prediction of the Hot Metal Silicon Content in the Blast Furnace—A Review,” *IEEE Transactions on Industrial Informatics*, vol. 9, no. 4, pp. 2213–2225, 2013.

[6] I. Cameron, M. Sukhram, K. Lefebvre, and W. Davenport, “Raceway Flame Temperature With Pulverized Carbon Injection,” in *Blast Furnace Ironmaking*. Elsevier, 2020, pp. 151–158.

[7] T. Kamijou and M. Shimizu, “PC Combustion in blast furnace,” in *Advanced Pulverized Coal Injection Technology and Blast Furnace Operation*. Elsevier, 2000, pp. 63–82.

[8] L. Ren, S. Zhou, T. Peng, and X. Ou, “A review of CO2 emissions reduction technologies and low-carbon development in the iron and steel industry focusing on China,” *Renewable and Sustainable Energy Reviews*, vol. 143, p. 110846, 2021.

[9] U.S. Geological Survey, “Mineral commodity summaries 2025: Iron and steel,” U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[10] M. Hatano, Y. Matoba, K. Otsuka, M. Yoshiki, and T. Miyagi, “Automatic control of hot metal temperature of blast furnace,” *Transactions of the Iron and Steel Institute of Japan*, vol. 22, no. 7, pp. 534–542, 1982.

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

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

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

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

[15] S. Ueda, S. Natsui, H. Nogami, J.-i. Yagi, and T. Ariyama, “Recent Progress and Future Perspective on Mathematical Modeling of Blast Furnace,” *ISIJ International*, vol. 50, no. 7, pp. 914–923, 2010.

[16] M. Shibata, K. Hashimoto, and K. Wakimoto, “Application of Expert System for Blast Furnace Operation Control,” *IFAC Proceedings Volumes*, vol. 20, no. 5, pp. 287–292, 1987.

[17] V. Radhakrishnan and A. Mohamed, “Neural networks for the identification and control of blast furnace hot metal quality,” *Journal of Process Control*, vol. 10, no. 6, pp. 509–524, 2000.

[18] J. Jiménez, J. Mochón, J. S. d. Ayala, and F. Obeso, “Blast Furnace Hot Metal Temperature Prediction through Neural Networks-Based Models,” *ISIJ International*, vol. 44, no. 3, pp. 573–580, 2004.

[19] U.S. Department of Energy, "Comprehensive report to congress, clean coal technology program: Blast furnace granulated coal injection system demonstration project," U.S. Department of Energy, Washington, DC, Tech. Rep., 1990.

[20] C. Zhou, "Computational fluid dynamics (CFD) modeling for high rate pulverized coal injection (PCI) into the blast furnace," Purdue University Calumet, final report to the U.S. Department of Energy, Hammond, IN, Tech. Rep., 2008.

[21] P. Kadlec, R. Grbić, and B. Gabrys, "Review of adaptation mechanisms for data-driven soft sensors," *Computers & Chemical Engineering*, vol. 35, no. 1, pp. 1–24, 2011.

[22] P. Azadi, R. Klock, and S. Engell, "Model Predictive Control of Molten Iron and Slag Quality Indices in a Large-scale Ironmaking Blast Furnace using a Hybrid Dynamic Model," *IFAC-PapersOnLine*, vol. 55, no. 21, pp. 138–143, 2022.

[23] K. R. Muske and T. A. Badgwell, "Disturbance modeling for offset-free linear model predictive control," *Journal of Process Control*, vol. 12, no. 5, pp. 617–632, 2002.

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

[25] 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] O. Raza, N. Walla, T. Okosun, K. Leontaras, J. Entwistle, and C. Zhou, "Prediction of silicon content in a blast furnace via machine learning: A comprehensive processing and modeling pipeline," *Materials*, vol. 18, no. 3, p. 632, 2025.

[27] L. Bonechi *et al.*, "BLEMAB european project: Muon imaging technique applied to blast furnaces," arXiv preprint arXiv:2110.10327, 2022.

2.9 Exercises

2.1 ○ Concept check · 1 min

In case 1 the coal injection rate is raised by 10 kg/tHM in Manual with the furnace wander off. What does the hearth do in the first hour after the step, and why?

(a) It cools by a few degrees, because the cold coal lowers the raceway flame temperature first

(b) It holds steady for five hours, because the coal must descend with the burden before it acts

(c) It warms at once by about 12 °C, because the coal burns just above the hearth

(d) It cools by a few degrees, because the cast dip is taken before the coal has burned

2.2 ○ Concept check · 2 min

In case 4 the coal loop runs on the soft sensor with no bias correction, and from 6 h the index reads 3 °C per hour high. What does the chapter find a day into the run?

(a) The index and the hearth both return to setpoint, because integral action removes any steady offset

(b) The index stays within about 15 °C of setpoint while the hearth cools far below it

(c) The hearth runs hot, because the loop adds coal to answer an index that keeps rising

(d) The loop sets the index aside once it disagrees with the cast dips and follows the dips

2.3 ○ Concept check · 2 min

Which description of the mid-ranging controller YC is right in the loop-action convention of Chapter 1, where $a = +1$ is reverse acting and $a = -1$ direct acting?

(a) It measures the hearth temperature and moves the coke rate, and it is reverse acting like the coal loop

(b) It measures the coke rate and moves the coal target, and it is direct acting

(c) It measures the coal command and moves the coke rate, and it is direct acting

(d) It measures the coal command and moves the coke rate, and it is reverse acting

2.4 Theory and calculation · 5 min

The soft sensor of case 5 drifts at $r_d = 3\,^{\circ}\mathrm{C}$ per hour and is corrected at every cast, $T_c = 2\,\mathrm{h}$, by Equation (2.10). Ignore the dip noise, the index noise and the index lag, and take the hearth as steady, so the corrected index reads the hearth plus the bias error $\varepsilon = d - \hat{b}$.

(a) Show that ε just before a cast settles at $r_d T_c/\alpha$, and evaluate it at $\alpha = 0.5$ and $\alpha = 0.8$. *Checkpoint:* 12 °C and 7.5 °C.

(b) Find ε just after a cast and its average over a cast interval at both weights. *Checkpoint:* After: 6 and 1.5 °C. Average: 9 and 4.5 °C.

(c) A loop that holds the corrected index at setpoint runs the hearth cold by that average. Over the challenge's last day at $\alpha = 0.8$, under the MPC, the chapter measures the index 6.7 °C above the hearth. Say why the run differs from (b), and what a larger weight trades. *Checkpoint:* 6.7 °C measured against 4.5 °C from (b).

2.5 Theory and calculation · 6 min

The coal path of Equation (2.5) has $G_r = 0.4$, $\tau_r = 0.4\,\mathrm{h}$, $G_p = 1.6$, $\theta_p = 1\,\mathrm{h}$ and $\tau_p = 1.5\,\mathrm{h}$, and the hearth is dipped every $T_c = 2\,\mathrm{h}$.

(a) Approximate the coal path as a first-order lag with dead time. Count the raceway lag as delay because it causes the inverse response, and add half the cast interval for the hold. Give the gain and the effective delay. *Checkpoint:* Gain 1.2, $\theta_e = 2.4\,\mathrm{h}$.

(b) Apply the SIMC rule with $\tau_{cl} = \theta_e$ for K_c and τ_I. *Checkpoint:* $K_c \approx 0.26\,\mathrm{kg/tHM}$ per °C, $\tau_I = 1.5\,\mathrm{h}$.

(c) Repeat (b) without the hold term, give the ratio of the two gains, and compare it with the factor the chapter measures from the stability boundaries. *Checkpoint:* $K_c \approx 0.45$, ratio about 1.7, against about 1.4 measured.

2.6 Theory and calculation · 5 min

Use the steady gains of the model, 1.2 °C per kg/tHM of coal and 2.0 °C per kg/tHM of coke, the coal limit of 200 kg/tHM and the coke range of 300 to 420 kg/tHM. Both fuels start at their base rates of 150 and 350 kg/tHM.

(a) Case 6 adds 80 °C of burden demand in all, with the coke held at 350 kg/tHM. Find the extra coal needed, the room the coal has, and the steady hearth temperature with the coal on its limit. *Checkpoint:* 66.7 kg/tHM against 50 kg/tHM; about 1480 °C.

(b) In case 7 mid-ranging returns the coal to its 150 kg/tHM target. Where must the coke settle, and is that inside its range? *Checkpoint:* 390 kg/tHM, inside the 420 kg/tHM limit.

(c) Case 8 sets the coal target to 190 kg/tHM. How large a burden change can the coal absorb alone, and how much coal does the case's 25 °C change ask for? *Checkpoint:* 12 °C; about 20.8 kg/tHM.

2.7 App challenge · 8 min

Load scenario 1, "Step the coke and the coal by hand (MANUAL)," at the default seed with no changes. The furnace wander is off in this case. Run to 36 h at the default speed of 3600× (one furnace hour per second), and export the CSV file. Read the column `hmt_C` against `Time (h)`. Give temperatures to ±0.1 °C.

(a) Read the hearth at 6.0 h and at 16.0 h. When does the coke step made at 1 h first move the hearth? *Checkpoint:* 1500.0 and 1519.0 °C; not before 6 h.

(b) Find the lowest hearth between 20 and 24 h, its time, and its difference from the value at 20.0 h. *Checkpoint:* 1516.1 °C at 21.0 h, 3.6 °C below the 1519.7 °C at 20 h.

(c) Read the hearth at 36.0 h and estimate the steady gains of coke and coal from your readings. *Checkpoint:* 1532.0 °C; about 2.0 and 1.2 °C per kg/tHM.

2.8 App challenge · 9 min

Load scenario 6, "A burden change beyond the injection limit," at the default seed with no changes, and run the full 72 h, at 14400× if you prefer (about 18 s). Export the CSV file. Then load scenario 7, "Mid-ranging: coke takes the load so the coal can return," and do the same.

(a) For each run, count the rows with `pci_saturated` equal to 1 and convert to hours (one row per 5 min), and average `hmt_C` over the rows from 60 to 72 h. Give hours to ±0.1 h and temperatures to ±0.1 °C. *Checkpoint:* Case 6: 26.0 h and 1477.3 °C. Case 7: 0 h and 1503.8 °C.

(b) Read the scorecard tiles Fuel cost, run average and Coke, coal at 72 h in each run. *Checkpoint:* $1681.6k and $1725.9k a day; coke 350 and 382 kg/tHM, coal 200 and 171 kg/tHM.

(c) Price the steady states the two structures head for, the coal on its limit with the coke at 350 kg/tHM and the coal back at 150 kg/tHM with the coke at 390 kg/tHM, at $400 per tonne of coke, $150 per tonne of coal and 10 000 tHM/d. Compare the difference with the three-day averages of (b). *Checkpoint:* $1.700 and $1.785 million a day, $85,000 apart, against $44,300 measured.

2.9 App challenge · 8 min

Load scenario 10, "Coke takes the load five hours ahead: MPC against mid-ranging," at the default seed. Run to 36 h with Thermal control on MPC as loaded, and export the CSV file. Press Reset, set Thermal control to Coal loop before pressing Run, run to 36 h again and export a second file. Skip any comment lines that start with `#`. Give rates to ±0.2 kg/tHM.

(a) In each file read `coke_kg` at 14.0 h, when the second burden change reaches the hearth, and `pci_kg` at 24.0 h. *Checkpoint:* MPC 375.0 and 157.1 kg/tHM. Coal loop 357.2 and 198.9 kg/tHM.

(b) In each file find the largest `pci_kg` and the first row at the 200 kg/tHM limit, if any. *Checkpoint:* MPC 192.0 kg/tHM, never at the limit. Coal loop at the limit from the 21.75 h row.

(c) Explain why the plan moves the coke earlier than mid-ranging, and what the earlier coke costs over the chapter's three-day run. *Checkpoint:* About 395 t more coke, $18,532 a day.

2.10 Application area · 4 min

The coal target is the rate the mid-ranging loop returns the coal to. At the base operation the furnace burns 350 kg/tHM of coke at $400 per tonne and 150 kg/tHM of coal at $150 per tonne, making 10 000 tHM/d, and one kilogram of coal is worth 0.6 kg of coke.

(a) With no disturbance, find the steady coke rate and fuel cost per tonne of hot metal at targets of 150 and 190 kg/tHM, and the daily saving. *Checkpoint:* Coke 350 and 326 kg/tHM; $162.5 and $158.9 per tonne; $36,000 a day.

(b) The chapter measures a saving of about $33,000 a day over three days. Give the reason for the difference, and say what the higher target costs in disturbance margin. *Checkpoint:* A margin of 12 °C for the coal alone.

2.11 Application area · 5 min

A hot stove fault drops the blast temperature by 80 °C for eight hours, as in case 9. The blast temperature is measured, so the coal could be moved as soon as the stove trips.

(a) From Equation (2.6) find the steady hearth loss and the coal step that offsets it. *Checkpoint:* 20 °C; about 16.7 kg/tHM.

(b) Compare the timing of the blast and coal paths, and say what a static feedforward of that size does in the first hour after the trip and when the stove returns. *Checkpoint:* Blast 0.3 h delay and 0.5 h lag. Coal cooling at once, heating after 1 h and a 1.5 h lag.

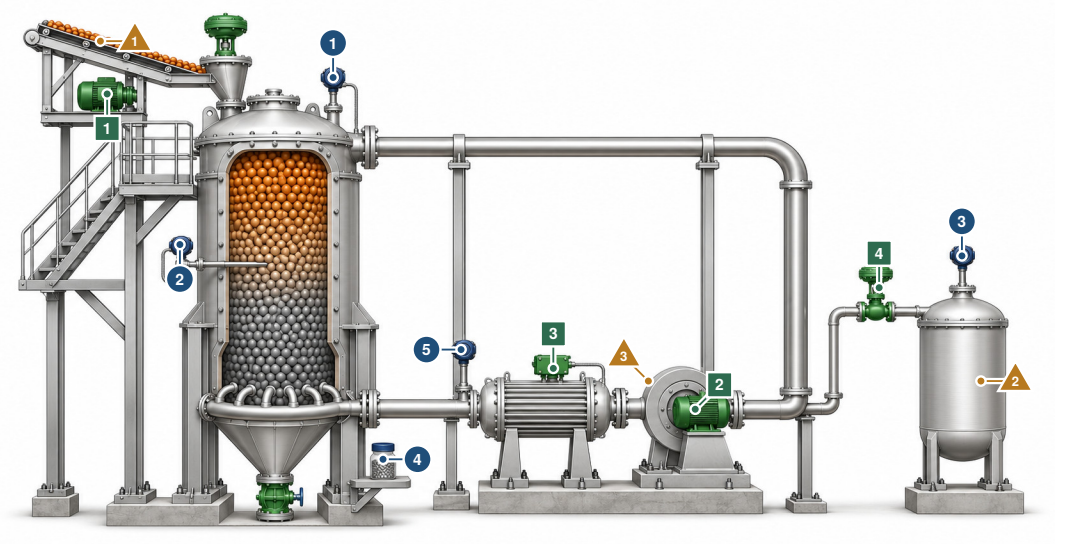
(c) The return of a stove is planned. How should the feedforward use that? *Checkpoint:* Compare the coal path's delay and lag with the blast's.

2.12 Review the agent · 5 min

In case 7 the coal takes days to return to its target after the second burden change. An AI assistant proposes: "Raise the mid-ranging gain from 0.2 to 1.0. The coke will then take the load five times faster, the coal will be back at its target within a day, and it will have more room for the next change." Find the flaw, and say how you would check it in the app.

CHAPTER 3

A Hydrogen Shaft Furnace Whose Product Answers Six Hours After Its Gas



The controller moves 1 pellet feed rate 2 recycle compressor, which sets the gas ratio 3 gas heater, which sets the bustle temperature 4 makeup hydrogen from the buffer

The controller reads 1 top pressure and top-gas analysis 2 burden thermocouple for the sticking margin 3 hydrogen left in the buffer 4 laboratory sample of the product every two hours 5 bustle gas flow and temperature

What disturbs it 1 a richer pellet lot 2 a wind lull that starves the buffer 3 a recycle compressor trip

Notice that the pellets take hours to fall from the feed at the top to the sample at the bottom while the gas loops through the bed in seconds, and that the buffer at the right is all the plant has when the wind drops. Illustration, not to scale.

IN PLAIN TERMS **Iron Made With Hydrogen, Checked Hours Later**

The main idea. Iron ore pellets slide slowly down a tall furnace. Hot hydrogen rises through them and carries off their oxygen as steam. A pellet takes six hours to reach the bottom, but the gas crosses the furnace in seconds.

The challenge. The controller must judge the finished iron from the gas that leaves the top. A laboratory test arrives only every two hours. The hydrogen is made with wind power, so when the wind drops the plant must choose. It can make less iron, or it can make poorer iron.

Why it matters. Making iron this way could remove most of the carbon dioxide that steelmaking puts into the air.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the fact box of Section 3.1, the two history boxes of Section 3.2, the try-it boxes of Sections 3.4 and 3.6, the fact box of Section 3.6, and the key points
Study	55 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 3.6, Section 3.7, Section 3.8 and the open challenge after it

The control question How can a hydrogen shaft furnace hold the metallization of a product that leaves six hours after its gas, is inferred from a top-gas balance and is checked by a late laboratory sample, while its hydrogen follows the wind?

Learning objectives

- Explain why the top gas answers a gas step within minutes while the product answers over the residence, and why the soft sensor leads the product.
- Compute the residence, the wüstite ceiling and the hydrogen a lull demands from the model's equations and constants.
- Judge when a laboratory report may correct the soft sensor's bias, and run the correction against a wrong pellet analysis.
- Compare a turndown override and a forecast-based MPC on the buffer, and say what each costs in tons or in quality.
- Use a temperature push to buy rate up to a measured sticking margin.

Prerequisites Sections 1.4.2, 1.4.5 and 1.4.6 of Chapter 1, the framework and the velocity-form PI of Sections 1.6 and 1.6.2, and for the buffer MPC Section 1.6.3. Chapter 2 shows the bias update on a hearth soft sensor.

The app https://apmonitor.com/apps/hydrogen-dri/

See also Volume 5, *Delays, Slow Measurements, and State Estimation*: Volume 5, Chapter 4 (a charge that descends for hours ahead of its analyzer) and Volume 5, Chapter 10 (a laboratory-corrected soft sensor and a push against a sticky margin). Volume 9, *Model Predictive Control and Economic Optimization*: Volume 9, Chapter 7 (decisions on a forecast) and Volume 9, Chapter 6 (a delay that follows the flow).



3.1 Introduction

In the simulated shaft, 140 t/h of iron ore pellets enter cold at the top and leave as direct reduced iron (DRI) at the bottom of the reduction zone six hours later, while 10 556 kmol/h of hot hydrogen enters at the bustle and crosses the bed in seconds. Raising the gas from 1690 to 1800 Nm^3 per tonne of pellets warms the top gas from 440.4 to 450.8 °C within three minutes. The metallization of the iron that leaves needs 2.4 h to make a tenth of its move and 6.7 h to make half of it, from 94.5 to 98.5 %. Between those two clocks sits everything the plant can measure about its product, an oxygen balance on the top gas and a laboratory sample drawn every two hours and reported half an hour late.

Steelmaking emits between 4 and 7 % of the world's anthropogenic carbon dioxide, and reducing the ore with hydrogen instead of carbon cuts that by roughly 90 % on the route's own accounting [1]. A route of this kind reduces pellets in a shaft furnace and melts the product in an electric arc furnace for 3.48 MWh per tonne of liquid steel, mostly to run the electrolyzer [2]. Shaft furnaces already make most of the world's DRI with natural gas [3], and their reduction zone is a counter-current moving bed whose models range from one dimension [4] to the detailed and multiscale models of the last decade [3, 5]. Shafts run on pure hydrogen have been simulated as well, with a gas flow several times the stoichiometric need [6, 7]. The control problem belongs to a wider class, a quality known only from an inference corrected by sparse laboratory results [8], on a plant with a long transport delay, under constraints that a selector or an override must arbitrate [9]. Kilns, grain dryers, polymer reactors with grade transitions and any plant whose laboratory reports after the next decision share that shape. The kraft digester of Volume 5, Chapter 4 is another vessel whose charge descends for hours before an analyzer and a later laboratory grade it.

DID YOU KNOW? 21.4 Kilograms an Hour Against 140 Tonnes

The world made 86 million tonnes of direct reduced iron in 2017, against 1180 million tonnes of blast furnace pig iron [3]. The first pilot shaft in the United States run on hydrogen was fed 21.4 kg of pellets an hour [10]. The simulated shaft of this chapter is fed 140 t an hour.

The cases of this chapter are built against two beliefs. The first is that a fast measurement of the gas reports the product. In the model the soft sensor, which counts the oxygen the whole bed is losing now, makes half the product's move in 2.2 h, and at hour 8 it reads 99.0 % against a true 96.6 %. The second is that the equilibrium sets the gas flow. The wüstite equilibrium caps the steam a hydrogen stream can carry at 36.9 % at 900 °C, but the designed top gas leaves at 27.0 %, because the hydrogen must also carry the heat that warms the pellets and drives an endothermic reduction.

However, the published shaft models are steady designs or plant validations, the hydrogen-route studies are flowsheets and energy balances, and none puts the inferred quality, the laboratory, the moving bed and the variable hydrogen supply in one loop that a learner can operate. This chapter works with a browser model of a shaft furnace fed by wind hydrogen and prices on its own scorecard what a soft sensor, a laboratory correction, a temperature push, a turndown override and a forecast-based model predictive controller (MPC) each contribute. The equilibria come from published tables, the kinetics are declared teaching values, and the controllers are established forms. What is particular to this plant is a moving-bed cell stack whose oxygen, energy and buffer ledgers close to the solver's precision. Every equation and constant of that stack appears below, and its controllers see only instruments, never the bed's states. Nine cases run from a hand step of the gas to a day on a wind forecast with an honestly infeasible lull. Section 3.2 follows the domain. Section 3.3 gives the bed, the gas loop, the buffer and the instruments, and Section 3.4 measures the design point and the two clocks of the open loop. Section 3.5 describes the soft sensor, the three PI loops, the turndown override and the optional MPC. Section 3.6 measures the nine cases, the challenge and the ledgers. Section 3.7 sets the loops in their hardware, Section 3.8 states the limitations, Section 3.9 concludes, and Section 3.10 holds the exercises.

3.2 Direct Reduction and Its Control

Direct reduction removes the oxygen from iron ore below its melting point and leaves a porous metallic product that an arc furnace can melt. Its gas-based form began with batch retorts and became a continuous shaft, as the history box below tells. Shaft furnaces make over 82 % of the world's DRI, two processes known by their trade names, MIDREX and HYL-ENERGIRON, holding 65 and 17 % [3]. Both reduce with a mixture of hydrogen and carbon monoxide reformed from natural gas, at up to about 950 °C [3, 5].

HISTORY From Batch Retorts to the Continuous Shaft (1957 to the 1960s)

The first generation of one gas-based direct reduction process, developed in 1957 in Monterrey, Mexico, reduced its charge in batches in four fixed-bed retorts. Its next generation was built around a shaft furnace, and the one after that kept the shaft. A second process was developed in the United States in the 1960s, and its first demonstration plant was built in Portland, Oregon. Its pellets or lumps are charged and preheated at the top, reduced as they descend by gas supplied in the middle of the furnace, cooled below the gas inlets and discharged, and ports sealed with gas let the furnace be charged and discharged continuously [11].

Graphical balances of the furnace predate the numerical models of the shaft. A dual graphic representation of the blast furnace's mass and heat balances was published by Rist and Meysson in 1967 [12], and the simulator's Map view, which plots the gas's steam share against the solids' reduction, is drawn in the spirit of such operating diagrams. One-dimensional models of the counter-current moving bed followed, with shrinking-core kinetics for one overall reaction [4], and later models resolved the three steps from hematite through magnetite and wüstite to iron and validated

them against an industrial shaft at 119.2 t/h, 93.8 % metallization and a 285 °C top gas [3]. Multiscale models link the pellet, the bed and the plant, with pellets of 7 to 15 mm reduced between 600 and 950 °C [5]. The kinetics of hydrogen reduction were measured in the laboratory step by step [13] and have been reviewed as a whole [14].

Hydrogen changes the balance of the shaft. The overall reduction is endothermic with hydrogen and exothermic with carbon monoxide, the stoichiometric need for hematite is 54 kg of hydrogen per tonne of iron, and the top gas is cooled to condense its water before the hydrogen is recycled with fresh makeup and reheated [1]. A shaft run on pure hydrogen was simulated at 3.8 times the stoichiometric gas flow with gas entering at 800 °C, and a reduction zone about 4 m high was enough to reach full metallization with 12 mm pellets [6, 7]. The doctoral work behind those studies read the wüstite equilibrium at 36 % steam for 800 °C and found that pellets reduced in pure hydrogen stuck together from 900 °C under load, from 1000 °C without it, and only once metallic iron had formed [15].

> ⧖ HISTORY **Hydrogen Alone (1999 to 2021)**
>
> A two-stage fluidized-bed process that reduced fine ores with pure hydrogen, made by reforming natural gas, is the only direct reduction process on pure hydrogen known to have operated commercially. It was decommissioned for economic rather than technical reasons [1], after running in Trinidad and Tobago from 1999 to 2001 and from 2004 to 2006 [10]. The first comprehensive study of steelmaking with hydrogen was a subproject of a European program on ultra-low carbon dioxide steelmaking, 2004 to 2010, which compared fluidized beds of fine ore with a shaft furnace of pellets and selected the shaft [1]. A pilot plant in Sweden was commissioned in August 2020 and made the first sponge iron reduced with fossil-free hydrogen in June 2021 [16].

The newest problem is the supply. An electrolyzer on renewable electricity makes hydrogen when the wind blows, while the shaft is the one part of the route designed to run continuously, so hydrogen storage acts as a buffer that detaches the electrolyzer from the rest of the process [2]. Techno-economic studies of wind-powered hydrogen for steelmaking size that storage to turn an intermittent production into the constant flow a steelworks needs [16]. What the operator does when the buffer runs low is a control decision. Selector logic is a simple and effective way to switch a plant between the constraints that bind [9], model predictive control is established industrial practice [17], and soft sensors corrected against laboratory results carry the quality between samples [8].

3.3 The Shaft and Its Gas Loop as Modeled

The simulated plant is the reduction zone of a generic hydrogen shaft, its gas loop and its hydrogen supply, as shown in Figures 3.1 and 3.2. The model is built for teaching from published sources. It ports no supplier's or licensed simulation and uses no plant data. Pellets of 67.5 % iron enter at 25 °C at 140 t/h, 1692.2 kmol/h of iron, and the zone holds $n_{Fe} = 10\,153$ kmol of it, a residence of 6.000 h. Gas leaving the top is dried in a condenser, purged of 2 %, compressed, mixed with makeup hydrogen from a 30 t buffer and reheated to the bustle temperature. A 350 MW electrolyzer follows a 600 MW wind farm into the buffer. The model runs in hours, kilomoles and degrees Celsius. The variables are classified in Table 3.1.

3.3.1 Solids Moving Down

The solids carry four states in each of $N = 80$ cells of equal iron holdup $n_c = n_{Fe}/N$. The fraction a of the cell's iron is still in the higher oxides, hematite and magnetite lumped at an oxygen-to-iron ratio of 1.5, m is metallic, and $w = 1 - a - m$ is wüstite, taken as FeO in the balances. T is the solids temperature and g the pellet grade the cell carries, a tracer that nothing else touches. The solids move

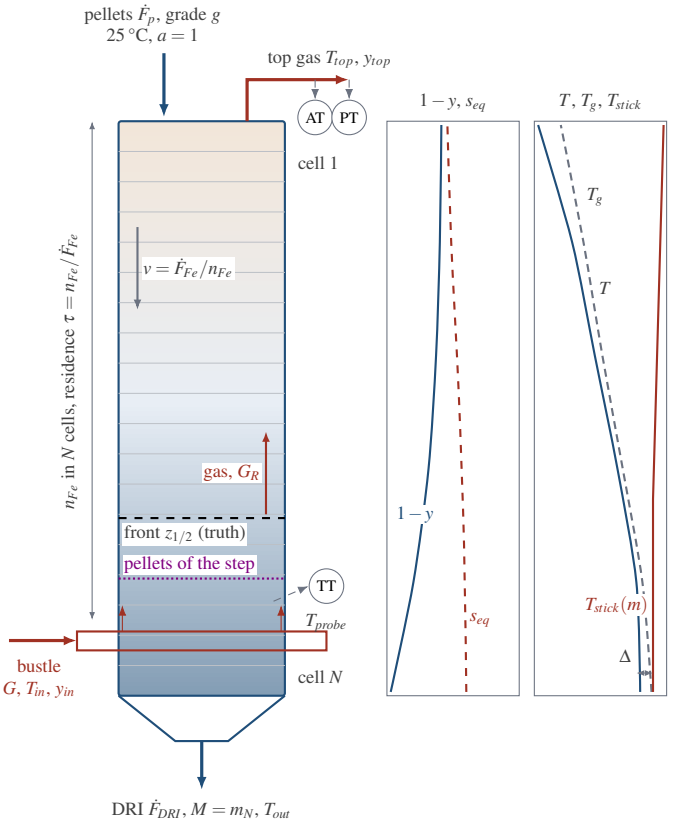


Figure 3.1 The reduction zone of the hydrogen shaft furnace as modeled, after the application's Shaft view. The zone of N cells holds n_{Fe} of iron, with pellets $\dot{F}_p$ of grade g entering at the top, solids descending at v, the bustle gas G at T_{in} and steam share $1 - y_{in}$ rising to the top gas at T_{top}, the reduction front $z_{1/2}$, the pellets charged at a step of the gas, the burden thermocouple T_{probe} and the product $M = m_N$ at T_{out}. Beside the bed, the profiles of the gas's steam share $1 - y$ against the wüstite wall s_{eq} and of the solids and gas temperatures against the sticking curve $T_{stick}(m)$ with the margin Δ. Process model only. Not to scale.

as a plug written in flux form, as shown in Equation (3.1).

$$\begin{aligned} \frac{\mathrm{d}\phi_i}{\mathrm{d}t} &= \frac{\dot{F}_{Fe}}{n_c}\left(\phi_{i-\frac{1}{2}} - \phi_{i+\frac{1}{2}}\right) + S_{\phi,i} \\ \phi_{i+\frac{1}{2}} &= \phi_i + \tfrac{1}{2}\,\psi(r_i)\,(\phi_i - \phi_{i-1}), \qquad \psi(r) = \frac{2r}{1+r} \\ S_a &= -r_1, \quad S_m = r_2, \quad S_g = 0 \\ S_T &= \frac{Q_i - n_c\left(\frac{1}{2}r_1\Delta H_1 + r_2\Delta H_2\right)}{n_c c_s} \end{aligned} \tag{3.1}$$

where $\dot{F}_{Fe} = 1000\,\dot{F}_p g/M_{Fe}$ is the iron fed in kmol/h, $r_i = (\phi_{i+1} - \phi_i)/(\phi_i - \phi_{i-1})$, $\psi = 0$ for $r \le 0$, the feed is the ghost value above cell 1 ($a = 1$, $m = 0$, 25 °C, the grade g), and the discharge face is first order, so the product is the bottom cell. The van Leer limiter keeps a front sharp without oscillation, and because each cell receives exactly what the one above gives, the oxygen and heat inventories close. A tracer charged at the top leaves with a mean delay of 6.000 h, equal to $n_{Fe}/\dot{F}_{Fe}$.

Table 3.1 Signals and States of the Shaft, Classified, at the Start of Case 1

Symbol	Role	Design	Range
M	CV, metallization of the DRI (truth)	94.54 %	SP 88–99
$\hat{M}_c$	measured, corrected soft sensor	94.5 %	0–100
M_{lab}	measured, laboratory, 2 h, 0.5 h late	—	held
R	MV, gas ratio	1690 Nm^3/t	900–2600
T_{in}	MV, bustle temperature	900 °C	800–1000
$\dot{F}_p$	MV, pellet feed	140 t/h	40–170
$\dot{H}_{mu}$	MV of the pressure loop, makeup	2670 kmol/h	0–6000
P	CV, top pressure	3.00 bar	SP 2.0–3.5
Δ	CV, measured sticking margin	51 K	SP 0–60
B	CV of the override, buffer	100 %	floor 20
g, g_m	DV, pellet grade and its analysis	67.5 % Fe	64–69.5
$f(t)$	DV, wind capacity factor	60 %	0–100
—	DVs, lull, trip, lot, outlier	events	—
a_i, m_i, T_i, g_i	states, 80 cells	—	—
n_L, B, P_{el}, T_{in}, $\dot{F}_p$, χ	states, plant	—	—
y_i, $T_{g,i}$	algebraic, gas in each cell	—	—

With first-order upwind cells, going from 20 to 40 cells moved the design metallization 3.1 points, which is why the model carries 80 cells of the second-order scheme.

3.3.2 Reduction Rates and the Equilibrium Walls

Two lumped steps remove the oxygen. The higher oxides go to wüstite with half an oxygen per iron and wüstite goes to iron with one, at shrinking-core rates under chemical control with an equilibrium driving force, as shown in Equation (3.2).

$$\begin{aligned} r_1 &= k_1(T)\, a^{2/3}\, P' x_R \Big[y - \frac{1-y}{K_1(T)} \Big]^+ \\ r_2 &= k_2(T)\, \frac{w}{(1-m)^{1/3}}\, P' x_R \Big[y - \frac{1-y}{K_2(T)} \Big]^+ \\ k_j(T) &= k_j \exp\Big[-\frac{E_j}{\mathscr{R}} \Big(\frac{1}{T} - \frac{1}{T_{ref}} \Big) \Big], \qquad K_j = \mathrm{e}^{A_j - B_j/T} \end{aligned} \tag{3.2}$$

where $y = H_2/(H_2 + H_2O)$ in the cell's gas, $x_R = 0.95$ is the reducing share, P' the top pressure over its 3 bar reference, T in kelvin, $T_{ref} = 1173.15\,\mathrm{K}$ and $[\cdot]^+ = \max(0, \cdot)$, so nothing reoxidizes. The factor $a^{2/3}$ is the unreacted core's surface and $w(1-m)^{-1/3}$ the iron front's surface times the share of the unconverted iron that is wüstite. The constants $A_2 = 1.0836$, $B_2 = 1899\,\mathrm{K}$ fit the NIST-JANAF Gibbs energies [18] of $Fe_{0.947}O + H_2 = 0.947\,Fe + H_2O$ between 900 and 1300 K within 0.31 %. The tables give $K_2 = 0.5026$, 0.5859 and 0.6657 at 800, 900 and 1000 °C, so the steam share over wüstite, $s_{eq} = K_2/(1+K_2)$, cannot exceed 33.45, 36.94 and 39.96 %, and the fit gives 33.5, 36.9 and 39.9 %. The iron-oxide stability (Chaudron) diagram read in a doctoral study of hydrogen reduction gives 36 % at 800 °C [15], about 2.5 points above the tables, and a reading offset by as much would put the ceiling near 39.5 % at 900 °C. The "40 % at 900 °C" of the model's first specification matches the tables only at 1000 °C, so the model keeps the tables' 36.9 %. The magnetite–wüstite step ($A_1 = 7.2464$, $B_1 = 7354.3\,\mathrm{K}$) allows 72.7 % steam at 900 °C. The rate constants $k_2 = 1.0\,\mathrm{h}^{-1}$ with $E_2 = 70\,\mathrm{kJ/mol}$ and $k_1 = 6\,\mathrm{h}^{-1}$ with $E_1 = 60\,\mathrm{kJ/mol}$ are teaching values chosen so that a pellet held in the bustle gas at 900 °C reaches 95 % metallization in 2.29 h. The activation energies lie in the range measured for hydrogen reduction of hematite and magnetite, 39 to 88 kJ/mol between 220 and 680 °C [13].

3.3.3 Gas Rising Through the Cells

The gas crosses the bed in seconds, so each cell's gas is algebraic. Hydrogen swaps for steam mole for mole, so the reducing flow $G_R = x_R G$ is the same in every cell and only y changes. Each cell's

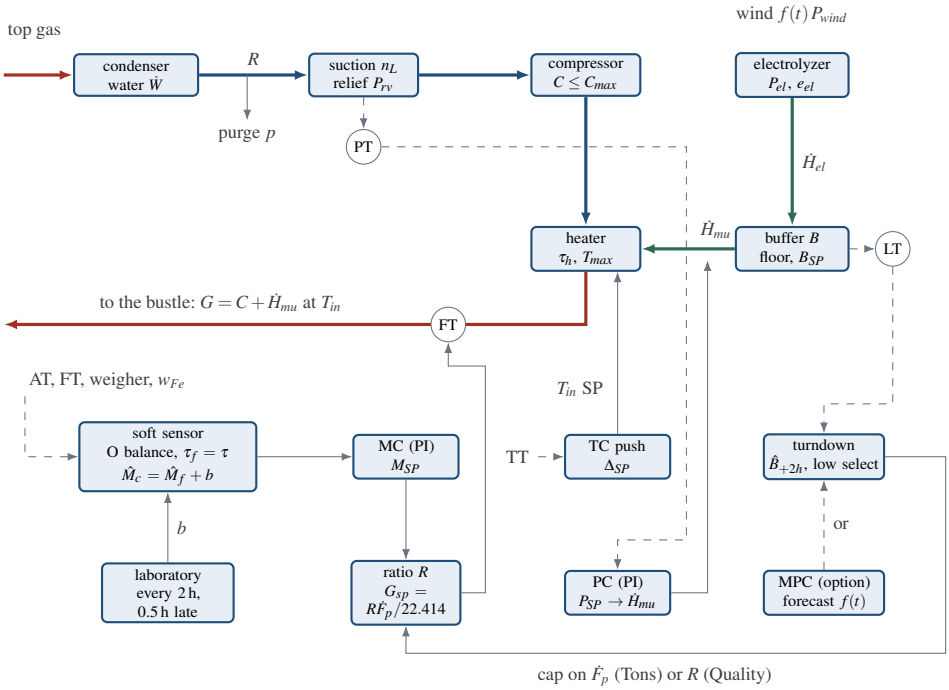


Figure 3.2 The gas loop and its control layers as modeled, after the application's Gas loop view: condenser, purge p, suction header n_L, compressor C, heater, makeup $\dot{H}_{mu}$ from the buffer B, electrolyzer P_{el} on the wind $f(t)P_{wind}$. Below, the control layers: the soft sensor corrected by the laboratory, the metallization loop MC on the gas ratio R, the temperature push TC, the pressure loop PC on the makeup and the turndown override with its optional MPC. Process model only. Not to scale.

gas is well mixed at its outlet value, which makes the oxygen it takes linear in y, and the gas gives heat to the solids through one conductance, as shown in Equation (3.3).

$$\begin{aligned} G_R\,(y_{i+1} - y_i) &= n_c\left(\tfrac{1}{2}r_1 + r_2\right)\Big|_{y_i}, \qquad y_{N+1} = \frac{x_R - x_W}{x_R} \\ T_{g,i} &= T_i + (T_{g,i+1} - T_i)\,\mathrm{e}^{-UA/(NGc_g)} \\ Q_i &= Gc_g\,(T_{g,i+1} - T_{g,i}) \end{aligned} \tag{3.3}$$

where $x_W = 0.025$ is the bustle's steam share, $T_{g,N+1} = T_{in}$ and the sweep runs from the bottom cell to the top. The first equation is solved in closed form with the active set of the two reactions found in at most three passes, so the sweep needs no iteration and is a pure function of the state. The top gas leaves at y_1 and $T_{g,1}$. The reaction heats $\Delta H_1 = 34$ and $\Delta H_2 = 16$ MJ per kilomole of oxygen are the JANAF enthalpies at 1100 K with gas-phase water, 33.93 and 16.18 MJ, which sum to +64.3 kJ per mole of hematite against +100.0 kJ at 298 K.

3.3.4 Gas Loop, Buffer and Electrolyzer

The bustle receives the compressor's recycle and the makeup $\dot{H}_{mu}$, and the top gas returns to a suction header whose inventory sets the pressure, as shown in Figure 3.2 and Equation (3.4).

$$\begin{aligned} P &= P_{ref}\, n_L / n_{L,0}, \qquad G = C + \dot{H}_{mu} \\ C &= \min\!\big(\alpha C_{max} \min(1.2, P/P_{ref}),\; G_{sp} - \dot{H}_{mu}\big) \\ S_{top} &= x_R G (1 - y_1), \qquad G_c = \frac{G - S_{top}}{1 - x_W} \\ \dot{W} &= S_{top} - (G_c - G + S_{top}), \qquad R_c = (1 - p)\, G_c \\ \frac{\mathrm{d}n_L}{\mathrm{d}t} &= R_c - C - V, \quad V = \frac{\max(0,\; n_L - n_{rv})}{0.02\,\mathrm{h}} \end{aligned} \tag{3.4}$$

where α is 1, or 0 while the compressor is tripped, G_c is the gas after the condenser at the bustle's steam share, $\dot{W}$ the condensed water, $p = 0.02$ the purge and n_{rv} the inventory at the 3.6 bar relief. The buffer and the electrolyzer follow Equation (3.5).

$$\begin{aligned} \frac{\mathrm{d}B}{\mathrm{d}t} &= \dot{H}_{el} - \dot{H}_{mu} - \dot{H}_{cu}, \qquad \dot{H}_{el} = \frac{1000\, P_{el}}{e_{el} M_{H_2}} \\ \frac{\mathrm{d}P_{el}}{\mathrm{d}t} &= \mathrm{clip}\!\left(\frac{\min(P_{el,max},\, f P_{wind}) - P_{el}}{0.02\,\mathrm{h}},\; \pm\rho_{el}\right) \\ \dot{H}_{mu} &\le \min\!\big(\dot{H}_{mu,max},\; \dot{H}_{el} + B/\Delta t\big) \end{aligned} \tag{3.5}$$

where $\dot{H}_{cu}$ curtails what would overfill a full buffer in one step and $e_{el} = 52\,\mathrm{kWh/kg}$. The makeup is written $\dot{H}_{mu}$ here because M is the metallization. The wind factor is a declared forecast plus a colored deviation of standard deviation 0.015 and correlation time 1 h. The forecast day falls from 0.78 at midnight to 0.06 for a lull from hour 10 to hour 16 and recovers to 0.66 by hour 24. The heater and the feeder are lags, $\tau_h \dot{T}_{in} = T_{in,cmd} - T_{in}$ with $\tau_h = 0.1\,\mathrm{h}$ and $T_{in} \le 1000\,^{\circ}\mathrm{C}$, and the same 0.1 h lag on the pellet feed. The heater duty is reported as $G c_g (T_{in} - 40\,^{\circ}\mathrm{C})$, 78.2 MW at design.

3.3.5 Sticking, Pressure Drop and the Product

Freshly metallized pellets sinter together above a temperature that depends on their composition. The teaching curve and its consequences are Equation (3.6).

$$\begin{aligned} T_{stick}(m) &= T_{st} + 150\,\mathrm{clip}\!\left(\frac{0.9 - m}{0.4}, 0, 1\right) \\ \frac{\mathrm{d}\chi}{\mathrm{d}t} &= \frac{0.05}{N} \sum_i \max\big(0,\; T_i - T_{stick}(m_i)\big) - \frac{\chi}{2\,\mathrm{h}} \\ \Delta P &= 0.9 \left(\frac{G}{10^4}\right)^{1.8} \frac{P_{ref}}{P}\,(1 + \chi)\ \mathrm{bar} \end{aligned} \tag{3.6}$$

with $T_{st} = 950\,^{\circ}\mathrm{C}$, inside the range in which hydrogen-reduced pellets were found to stick in the laboratory [15]. The rise toward oxide is a teaching shape and the cluster index χ a teaching state. The product is the bottom cell. Its metallization is $M = m_N$ and its mass per kilomole of iron is $M_{Fe} + 16(1 - m + a/2) + \gamma(g)$ with gangue $\gamma = M_{Fe}/g - M_{Fe_2O_3}/2$, which gives the DRI rate, its total iron and its gangue. The model has no transition or cooling zone, so the product leaves at the temperature of the bottom cell.

3.3.6 Measurements

The controllers read only a measurement packet: the pellet weigher and the bustle flow meter (0.3 % noise each), the pellet analysis $g_m = g + \delta g$ with a declared error, the bustle and top-gas thermocouples (1 and 2 K), a top-gas analyzer of hydrogen and steam with a 2 min lag and 0.05

Table 3.2 Model Parameters and Their Sources. All but N and g Are Adjustable in the Parameters Panel

Symbol	Meaning	Value	Source
n_{Fe}	iron in the zone	10 153 kmol	6.0 h of feed
N	cells	80	refinement
k_2, E_2	wüstite step	1.0 h^{-1}, 70 kJ/mol	teaching
k_1, E_1	higher oxides	6 h^{-1}, 60 kJ/mol	teaching
A_2, B_2	wüstite–iron	1.0836, 1899.0 K	JANAF fit [18]
A_1, B_1	magnetite–wüstite	7.2464, 7354.3 K	JANAF fit
ΔH_1, ΔH_2	reaction heats	34, 16 MJ/kmol O	JANAF, 1100 K
c_s	solids heat capacity	65 kJ/(kmol Fe K)	teaching
c_g	gas heat capacity	31 kJ/(kmol K)	teaching
UA	gas–solid, whole bed	4×10^7 kJ/(h K)	teaching
g	pellet grade	67.5 % Fe	DR grade
x_R, x_W	reducing, steam share at bustle	0.95, 0.025	teaching
p	purge share	0.02	teaching
P_{ref}, $n_{L,0}$	top pressure, loop gas	3 bar, 400 kmol	teaching
P_{rv}	relief	3.6 bar	teaching
C_{max}, $\dot{H}_{mu,max}$	compressor, makeup valve	9000, 6000 kmol/h	teaching
τ_h, T_{max}	heater lag, limit	0.1 h, 1000 °C	teaching
P_{wind}, $P_{el,max}$	wind farm, electrolyzer	600, 350 MW	teaching
e_{el}, ρ_{el}	specific energy, ramp	52 kWh/kg, 2100 MW/h	teaching
B_{max}, floor	buffer	30 t, 20 %	teaching
T_{st}	sticking, metallized	950 °C	level [15]
T_{lab}, θ_{lab}, σ_{lab}	laboratory	2, 0.5 h, 0.3 points	teaching

points of noise, the top pressure (0.005 bar), the bed pressure drop, a burden thermocouple at the bustle reading the bottom cell's solids (3 K), the buffer as a level (0.2 %) and the electrolyzer and makeup meters (0.5 %). The laboratory draws a spot sample of the discharge every 2 h from $t = 0$, reports it 0.5 h later with 0.3 points of noise and holds it until the next report. The report shown before the run is a noisy sample of the lined-out product dated hour −2. Every sample is stored with the product's true metallization at the moment it was drawn, which the scorecard and the challenge judge and no controller reads, and a source-level test of the application rejects any controller or estimator that names a truth field. Where a controller needs a plant constant, namely the buffer size, the electrolyzer and wind farm ratings, the specific energy, the holdup, the bustle composition or the sticking temperature, it takes the shipped value, and a second test checks that a parameter edit leaves every controller output unchanged at fixed measurements. The CSV export of Section 1.6.4 writes one row per simulated minute with the soft sensor before and after its bias, the laboratory report and its sampling hour, the gas ratio requested, commanded and applied, the bustle and top-gas measurements, the sticking margin, the buffer, feed, makeup and wind, and, under a `truth_` prefix, the product, the margin, the front and the cluster index.

3.3.7 Parameters, Initialization and Numerics

Table 3.2 lists the parameters with their defaults. Each case starts lined out. A pseudo-transient Euler integration at 0.004 h finds the bed's steady state for the case's stations, and the makeup is set to the difference between the bustle gas and the returning recycle. The solve is deterministic and cached by its inputs, so Reset replays exactly. The framework's engine (Section 1.6.1) integrates the whole plant with explicit Euler at $\Delta t = 5$ s, clamps the solid fractions to the simplex, runs the controllers every minute and logs a row every minute. Halving Δt moves case 1's metallization at hour 12 from 97.5069 to 97.5121 %. Doubling the cells to 160 moves the design metallization by 0.385 points, from 94.540 to 94.925 %, moves no case's measured lesson by more than 0.41 points and changes the verdict of none of 15 challenge policies on five seeds. The suggested challenge day takes 17 287 steps in 0.24 s headless, about 14 µs a step, and Chrome needs 24.0 s for the whole 24 h day at 1 h/s.

3.3.8 Assumptions

Hydrogen is the only reductant, with 5 % inert and 2.5 % steam at the bustle, so the natural-gas routes, carbon monoxide and carbon in the product are out of scope. The bed has one dimension and uniform pellets, with no radial gradients, channeling or gas distribution, which a real shaft's cone and bustle do not achieve. Hematite and magnetite are lumped, wüstite is FeO in the balances and $Fe_{0.947}O$ in the equilibrium, and the equilibrium fits are extrapolated below 570 °C, where wüstite is not stable but little reduces. The kinetics omit pore diffusion, gas-film resistance, swelling and disintegration. Heat capacities are constant, the walls lose nothing, and one conductance couples gas and solids. The bustle composition is declared rather than computed from the recycle and the dry makeup, and the makeup is the total-gas difference $G - R_c$, so at design it is 77 kmol/h, 3.0 %, above an exact species balance. Of the 53.4 kg of makeup hydrogen per tonne of DRI, 48.9 are consumed, 2.9 purged and 1.5 are this closure, and the condenser removes 45.2 t/h of water against 44.1 t/h made. The electrolyzer has no minimum load, no stack degradation and no grid purchase.

3.4 Dynamic Character

3.4.1 The Design Point and the Heat That Sets the Flow

At the shipped stations, 1690 Nm^3 of bustle gas per tonne of pellets at 900 °C and 3 bar, the bed lines out at 94.54 % metallization and makes 100.87 t/h of DRI with 93.69 % total iron, 6.58 % FeO, 4.85 % gangue and no carbon, discharged at 899 °C. The top gas leaves at 440 °C with 27.0 % steam against a 36.9 % ceiling at the bustle, and 74.7 % of the bustle gas is recycle. The bustle carries 3.99 times the hydrogen the reduction consumes (4.10 counting its steam), and on the basis of the pure-hydrogen study, hydrogen fed over the hydrogen that reduces all the iron fed, the design is 3.85 against the study's 3659.3/976.9 = 3.75 [15], reported there as 3.8. That study's adiabatic balance gives a 384 °C top gas at a 900 °C inlet against 440 °C here. The reduction consumes 52.2 kg of hydrogen per tonne of iron, against 54 kg for full reduction of hematite [1].

The equilibrium does not set this flow. The upper shaft, cells 9 to 40, averages 729 °C in the solids, where the wall is near 30.8 %, so the gas leaves even the coolest part of the bed below its local wall, as Figure 3.3d shows for the whole range of ratios. At 900 °C the top-gas utilization peaks at 27.3 % near 1600 Nm^3/t and falls on both sides, at lower ratios because the cold upper bed reduces nothing and at higher ratios because there is more gas than oxygen. No ratio approaches the ceiling. The compressor caps the applied ratio at 1887 Nm^3/t for 140 t/h at 900 °C. Below about 1600 Nm^3/t the metallization falls steeply, to 64.7 % at 1400 and 13.3 % at 1000, and a hotter bustle shifts the whole curve left, as shown in Figure 3.3c.

🕮 **Pause.** Try Exercise 3.2 before reading on.

3.4.2 A Fast Gas and a Slow Product

Case 1 raises the ratio from 1690 to 1800 Nm^3/t at hour 1. Within three minutes the top gas warms from 440.4 to 450.8 °C and its steam share falls from 27.04 to 26.45 %, as shown in Figure 3.3a. The product rises from 94.54 to 98.52 % by hour 30, reaching 10, 50, 63 and 90 % of that move 2.37, 6.73, 8.62 and 16.75 h after the step, as shown in Figure 3.3b. The gas reaches every pellet at once, but a pellet leaving the bottom has spent hours under the old gas, and the upper bed must warm before its reduction speeds up. The pellets charged at the step reach the bustle one residence later, at hour 7, 83.3 % of the way down at hour 6, as the application draws them in Figure 3.5.

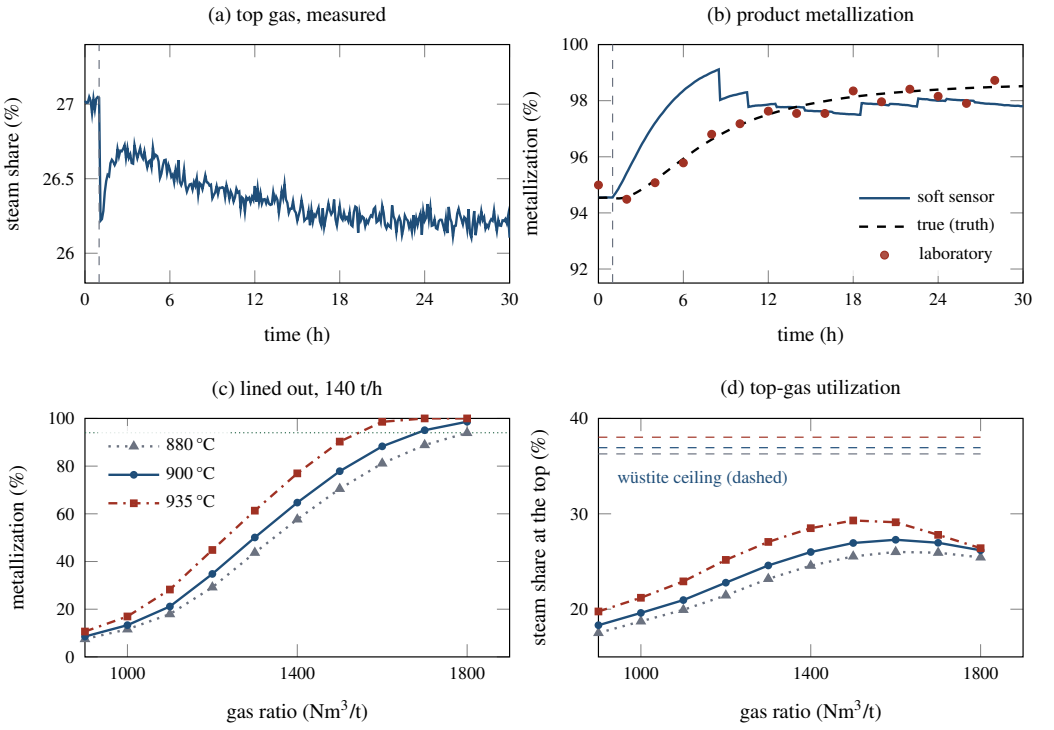


Figure 3.3 Two clocks and one operating map, computed by the application's engine. (a), (b) Case 1 in Manual with the gas ratio raised from 1690 to 1800 Nm^3/t at hour 1 (dashed line). The measured top-gas steam share steps within minutes, while the true product turns over the residence time and the soft sensor leads it until the laboratory reports, drawn at their sampling instants, move its bias from hour 8.5. (c), (d) The lined-out bed at 140 t/h against the applied gas ratio at three bustle temperatures, with the 94 % specification (dotted) and the wüstite ceiling at each bustle temperature (dashed, higher for a hotter bustle). The curves of the coolest bustle are dotted with triangles, those at 900 °C solid with circles and those at 935 °C dash-dotted with squares. In (b) the soft sensor is solid, the true product dashed and the laboratory samples dots. Points above 1887 Nm^3/t at 900 °C are not reachable because the compressor caps the flow.

TRY IT IN THE APP 1 · Raise the Gas: The Top Answers Now, the Product Later

Load the scenario, press Run and, after the start, raise Gas ratio to 1800 Nm^3/t. Watch the top-gas trend: its steam share falls from 27.04 to 26.45 % within three minutes. On the metallization trend the true product needs 6.73 h to make half its move, while the soft sensor gets there in 2.18 h.

A pellet lot shows the residence alone. In case 2 a lot with one point more iron enters at hour 1. The DRI rate rises at once from 100.87 to 102.41 t/h, because the bed's iron holdup is fixed, so the discharge iron rate rises at once with the richer feed while the old gangue-to-iron ratio is still leaving the bed. It falls back to 100.56 t/h at hour 7.5, half its drop at hour 6.90, when the lot's lower gangue per iron reaches the discharge. The DRI's total iron makes 10, 50 and 90 % of its step from 93.69 to 95.37 % at 4.38, 5.92 and 6.18 h. A change of the feed rate is not a lot. Cutting the feed from 140 to 120 t/h speeds or slows the whole bed at once and moves the product 10 % of its way in 0.58 h.

3.5 Implemented Control System

The controllers execute every minute on the measurement packet. A header switch selects Manual or Auto, and each loop has its own Auto and Man with a station that is live only in Man. Layer 1 is the gas-ratio station, $G_{sp} = R\dot{F}_{p,m}/22.414$ on the weigher, the heater setpoint and the top pressure loop on the makeup. Layer 2 is the metallization loop on the corrected soft sensor and the temperature push on the sticking margin. Layer 3 is the turndown supervisor on the buffer. The 94 % specification, the 20 % buffer floor and the sticking curve are teaching values. The real metallization target comes from the arc furnace's energy and slag balance and the buyer's specification, the real sticking limit from laboratory tests on the pellets in use [15] and the real buffer floor from the storage vessel's minimum pressure and the trip logic.

3.5.1 The Soft Sensor and Its Laboratory Correction

The soft sensor closes the oxygen balance on the top gas against the oxygen the pellets bring, filters it with the residence and adds a bias that the laboratory moves, as shown in Equation (3.7).

$$\begin{aligned} \hat{X} &= \frac{G_m x_R\,(s_{top} - x_W/x_R)}{1.5\,\dot{F}_{Fe,m}}, \qquad \hat{M} = 1 - 1.5\,(1-\hat{X}) \\ \hat{M}_{f,k} &= \hat{M}_{f,k-1} + \left(1 - \mathrm{e}^{-\Delta t_c/\tau_f}\right)(\hat{M}_k - \hat{M}_{f,k-1}) \\ \tau_f &= n_{Fe}/\dot{F}_{Fe,m}, \qquad \hat{M}_c = \mathrm{clip}(\hat{M}_f + b,\ 0,\ 1) \\ b &\leftarrow b + \lambda\left(M_{lab} - \hat{M}_f(t_s) - b\right) \end{aligned} \tag{3.7}$$

where s_{top} is the analyzer's steam share of the reducing gas, $\dot{F}_{Fe,m}$ the iron fed from the weigher and the pellet analysis, $\Delta t_c = 1\,\mathrm{min}$, $\tau_f \le 24\,\mathrm{h}$, t_s the sample's drawing time and $\lambda = 0.5$ with $|b| \le 8$ points. The update is the framework's bias updater, the exponentially weighted form of Equation (1.6) applied at each accepted report, as the hearth soft sensor of Chapter 2 applies it at each cast. In a steady bed with the right analysis the balance equals the product to 0.005 points. A report moves b only when the gas ratio and the feed applied over the 6 h before the sample stayed within 3 % and 3 t/h and the report is below 98.5 %, because a balance read during a transient charges the bed's stored oxygen to the bias.

Pause. Try Exercise 3.3 before reading on.

The clamp corrects an earlier build of the application. The balance counts the oxygen the bed gives up against the oxygen fed now, so after a feed cut or with a hotter bustle it passes 100 %, to 103.21 % at hour 17.5 of case 6 for 4.8 h and to 107.74 % at 965 °C in case 4, while the true product sits at 100 %. The corrected inference is now held within 0 to 100 %, the unclamped balance is logged as `met_raw_pct`, and the bias is computed on it.

3.5.2 Three PI Loops in the Velocity Form

The metallization loop, the temperature push and the pressure loop use the framework's velocity-form proportional-integral (PI) controller of Section 1.6.2 without derivative action, executed every Δt_c. In the symbols of Equations (1.13) to (1.16), this chapter writes the action a as σ, the output bias u_b as u_0 and the gain as K_c, and the tracking time T_t is τ_I. The metallization loop reads $\hat{M}_c$ and moves the ratio within 900 to 2600 $\mathrm{Nm^3/t}$ with $\sigma = +1$, reverse acting in the convention of Section 1.3 because more gas raises the metallization, $K_c = 10\,\mathrm{Nm^3/t}$ per point and $\tau_I = 3\,\mathrm{h}$. While the feed command has moved more than 3 t/h in the last 2 h it tracks the applied ratio instead of integrating, because the balance then reads the bed's transient. The push reads the measured margin $\Delta = T_{stick}(0.9) - T_{probe}$, the curve's metallized end less the burden thermocouple, and moves the heater setpoint within 800 to 1000 °C with $\sigma = -1$, direct acting because a hotter bustle closes the margin, $K_c = 0.5$ and $\tau_I = 0.5\,\mathrm{h}$, so it heats while the margin exceeds its setpoint. This is the constraint push that the spray dryer of Volume 5, Chapter 10 applies to its inlet temperature against a sticky margin. The push never reads

the curve at the soft sensor's value, which an early build did and which let a transient talk the push into sticking. The pressure loop moves the makeup within 0 to 6000 kmol/h with $\sigma = +1$, reverse acting, $K_c = 2000$ kmol/h per bar and $\tau_I = 0.25\,\mathrm{h}$. In Man each loop tracks its station, so Man to Auto is bumpless. The application also applies a station entry made since the last scan before the loop goes to Auto. In an earlier build without that rule, revealing the suggested settings in mid-run set the push to Auto and the bustle station to 935 °C in one instant, the push started from the old 900 °C, and the sample of hour 2 read 93.81 %. The mid-run failure was a lost station entry, not the bed. A grade feedforward, when switched on, scales the ratio by the pellet analysis over its value at the switch, so that the gas follows the iron fed.

3.5.3 The Turndown Override

When the buffer runs low, the supervisor caps either the pellet feed (priority Tons) or the gas ratio (priority Quality) with a low selector whose reset is fed back from the applied value, as shown in Equation (3.8).

$$\begin{aligned} \hat{B}_{+2h} &= B_m + \frac{2\,\mathrm{h}\,(\dot{H}_{el,m} - \dot{H}_{mu,m})}{B_{max}} \\ u_B &= \bar{u} + K_{c,B}\,\Delta u\,\big(\hat{B}_{+2h} - B_{SP}\big) \\ u &= \min(u_{req},\, u_B) \\ \bar{u}_k &= \bar{u}_{k-1} + \big(1 - \mathrm{e}^{-\Delta t_c/\tau_{I,B}}\big)\,(u_{applied} - \bar{u}_{k-1}) \end{aligned} \tag{3.8}$$

where Δu is the span of the capped station, 40 to 170 t/h or 900 to 2600 $\mathrm{Nm^3/t}$, $K_{c,B} = 1$, $\tau_{I,B} = 0.5\,\mathrm{h}$ and $B_{SP} = 40\,\%$ as shipped. The cap sits above the request while the projected buffer is above its setpoint, takes over at the request's value when it reaches it, integrates while selected and hands back when the buffer recovers. Nothing latches. The feed moves at most 30 t/h per hour whoever asks, and while the override holds the ratio the metallization loop tracks the applied value, so it does not wind up. The supervisor reads the measured imbalance projected two hours ahead and has no forecast.

3.5.4 The Optional Buffer MPC

The Advanced turndown method "MPC" replaces the override with a receding-horizon quadratic program on the framework's linear MPC (Section 1.6.3), as shown in Equation (3.9).

$$\begin{aligned} \min_{u}\ & -\sum_{k} u_k + 0.02 \sum_{j} \Delta u_j^2 + 2000\,\varepsilon^2 \\ & + 200\,\varepsilon \\ \text{s.t. } & B_{k+1} = B_k - T_s \hat{h}\, u_k + T_s\, \hat{w}_k \\ & B_{SP} B_{max} - \varepsilon \le B_k \le B_{max}, \qquad \varepsilon \ge 0 \\ & 40 \le u_k \le \dot{F}_{p,SP}, \qquad |\Delta u_j| \le 15\,\mathrm{t/h} \end{aligned} \tag{3.9}$$

where u is the pellet feed, B the buffer in tonnes of hydrogen, $T_s = 0.5\,\mathrm{h}$ over 48 steps in 12 move blocks, $\hat{w}_k = \min(P_{el,max}, f(t_k) P_{wind})/e_{el}$ the hydrogen the declared forecast promises, and $\hat{h}$ the makeup per tonne of pellets learned from the meters through a 1 h filter. The weights are those declared to the module. The electrolyzer capacity, wind farm, specific energy and buffer size in $\hat{w}_k$ and in the bounds are the shipped values, so a change to those parameters is model mismatch that the MPC meets in the measured buffer. The program is re-solved every half hour from the measured buffer and its first move caps the feed. It never sees the actual wind's deviation and it cuts only tons, while the metallization loop and the push hold the quality. "One plan" is the same program held to a single move block. The MPC therefore has information the override does not, the wind forecast, much as the scenario-tree program of Volume 9, Chapter 7 plans its releases on an ensemble rain forecast, and the comparison in Section 3.6 states that difference rather than credit it to the method. Over the challenge day the MPC and One plan made 48 solves without a failure.

Table 3.3 The Eight Cases and the Challenge

#	Case and events	Teaching point
1	Manual, raise the ratio after the start	two clocks
2	Manual, lot of 68.5 % Fe at 1 h	the residence
3	Auto, the same lot at 2 h, analysis still 67.5 %	soft sensor, laboratory
4	Auto, loop in Man, bustle by hand or the push	rate against sticking
5	Auto, lot at 1 h, feedforward off or on	measured grade
6	Auto, forecast day, buffer at 90 %, supervisor on	tons or quality, infeasible lull
7	Auto, compressor trip at 2 h for 0.5 h	pressure and bed
8	Manual, Map view, ratio changed at 0.25 h	the equilibrium wall
9	Challenge: 24 h on the forecast, lot at 3 h	the whole problem

Table 3.4 Measured Results, Default Seed. Metallization Is the True Product

Case and setting	Result
1, ratio 1690 → 1800	top gas in 3 min, product half in 6.73 h, 98.52 % at 30 h
1, soft sensor	half in 2.18 h, 99.01 against 96.61 % at 8 h
2, lot	DRI iron 50 % of its step at 5.92 h
3, correction on	low 93.37 % at 13.7 h, 94.44 % at 30 h, 10 of 16 short
3, correction off	92.55 % at 30 h and falling, 13 of 16 short
4, 965 °C by hand	sticking from 1.21 h, ΔP 0.99 → 1.17 bar
4, push at 15 / 5 K	935.4 / 945.1 °C, 99.41 / 99.86 %, no sticking
5, feedforward off / on	92.85 / 94.16 % at 24 h, ratio 1690 → 1715
6, Tons	1991 t, worst 94.53 %, buffer low 27.9 %
6, Quality	2518 t, worst 43.40 % (6 of 13), 42.6 kg/t
6, supervisor off	buffer dry at 15.7 h, worst 77.34 %
7, trip	3.69 bar, below 94 % from 2.33 to 9.33 h, low 86.74 %
8, ratio 1400	5 cells at the wall at 10.27 h, 67.71 % at 24 h
8, ratio 1000	no cell at the wall, 13.48 % at 24 h

The application keeps three models apart. The plant is the model of Section 3.3. The identified model is a learner's fit to case 1's step. The controllers' models are the residence filter, the declared holdup and pellet analysis inside the soft sensor, and the one-state buffer model of the MPC.

3.6 Nine Cases, Measured

The cases are listed in Table 3.3. Each value in this section is the shaft's own output under Node on seed 20261027, with five-second steps identical to the page's, a learner's change made after the run started where a row says so, and the challenge repeated on four further seeds. Loading a case first, picking it after another case and pressing Reset give the same nine runs, with Reset and the reload arrow as Section 1.6.4 describes them. The results are collected in Table 3.4.

> ✦ **DID YOU KNOW?** **Half an Hour Without Gas, Seven Hours Off Grade**
>
> From the simulator: in Table 3.4, a recycle compressor trip of half an hour at hour 2 leaves the product below 94 % from hour 2.33 to hour 9.33, with a low of 86.74 %. The burden kept moving while the bustle had no gas.

3.6.1 An Inferred Product and a Late Laboratory

The soft sensor is faster than the product because it reads the bed, not the discharge. In case 1 it makes half the product's move in 2.18 h against the product's 6.73 h. The first report it accepts after the step is from the sample of hour 8, drawn after six hours of steady burden, and the bias moves to −1.10 points at hour 8.5 and to −1.90 by hour 16.5, as the saw-tooth of Figure 3.3b shows. The gate's six hours are shorter than the product's 8.62 h to 63 %, so the bias absorbs part of the transient, and at hour 30 the soft sensor reads 97.81 % against a true 98.52 % with an analysis that is correct.

Pause. Try Exercise 3.1 before reading on.

Case 3 charges the same rich lot at hour 2 while the analysis still reads 67.5 %. The pellets then bring 1.5 % more oxygen than the balance counts, the soft sensor reads about two points high, and the metallization loop holds the ratio where the product slides. Without the correction the product falls to 92.55 % by hour 30 and is still falling, with 13 of 16 samples short. With it, the bias reaches −2.56 points, the loop raises the ratio from 1689 to 1730 Nm^3/t, and the product turns at 93.37 % at hour 13.7 and recovers to 94.44 % at hour 30, with 10 of 16 samples short. The correction is slow because each report describes the product of hours ago and the product answers the ratio over the residence. Case 5 shows the cheaper answer. A feedforward on the measured grade raises the ratio from 1690 to 1715 Nm^3/t at once and holds the product at 94.16 %, against 92.85 % at hour 24 without it. The ratio scaled by the iron alone leaves the product short because the bed holds a fixed amount of iron. The richer feed passes more iron through it, so each pellet spends less time under the gas. Raising the feed by the same share of iron, with no new lot, costs the same metallization.

3.6.2 Temperature, Sticking and a Trip

Hotter gas reduces faster and carries more heat. In case 4 a bustle set by hand to 965 °C at hour 1 lifts the product to 99.67 % by hour 5, but the metallized pellets at the bustle follow the gas within a quarter of an hour, sticking starts at hour 1.21, and the cluster index reaches 0.19 with the bed pressure drop rising from 0.99 to 1.17 bar by hour 10. The push in Auto at a 15 K margin settles the heater at 935.4 °C with the product at 99.41 % and no sticking, and at 5 K it reaches 945.1 °C. In every heating case the measured margin stayed within 10.3 K, 3.4 standard deviations of its thermocouple noise, of the true margin of the cell closest to the curve. Case 7 trips the recycle compressor for half an hour at hour 2. The top pressure rises to 3.69 bar at the relief, the bustle flow falls to nothing by hour 2.5, and the bed pressure drop from 0.98 to 0.04 bar. The burden keeps moving, so the product is below 94 % from hour 2.33 to hour 9.33 with a low of 86.74 % at hour 3.35, and three of the first six samples are short. A half-hour upset costs seven hours of off-specification iron.

3.6.3 The Wind, the Buffer and an Infeasible Lull

Case 6 runs the forecast day with the buffer starting at 90 %. With priority Tons the supervisor cuts the pellets from 140 to 40 t/h through the lull, the buffer bottoms at 27.9 % at hour 16.4, every sample stays at or above 94.53 % and the day makes 1991 t at 54.2 kg of makeup hydrogen per tonne. With priority Quality the tons stay, 2518 t at 42.6 kg/t, but the ratio falls to its 900 Nm^3/t floor and the worst sample is 43.40 % with 6 of 13 short. With the supervisor off the buffer runs dry at hour 15.7 and the worst sample is 77.34 %. A tonne of iron needs its oxygen removed, and hydrogen saved at the same tonnage is iron left unreduced.

> **TRY IT IN THE APP** **6 · The Electrolyzer Follows the Wind: Tons or Quality**
>
> Run the day with Scarce hydrogen costs at Tons and watch the pellet feed on the Gas loop view: it falls from 140 to 40 t/h through the lull, and every sample stays at or above 94.53 %. Then set it to Quality and press Reset. The day makes 2518 t instead of 1991 t, but the worst sample is 43.40 %.

The case also holds the honestly infeasible variant, full tons and every sample at 94 % through the six-hour lull with the shipped buffer. The oxygen alone needs 29.5 t of hydrogen over the lull at 140 t/h, while the buffer above its floor holds 24.0 t and the electrolyzer makes 4.2 t in the lull. The bed can shift some reduction around the lull, so the count is not a proof. A sweep of 24 constant settings of ratio, bustle temperature and a hand-set makeup station, and then of 960 schedules that pre-reduce the bed before the lull at up to 2600 Nm^3/t and 940 °C, starve it in the lull and repay it after, found none that kept both the buffer at 20 % and every sample at 94 %. The closest kept the buffer at 19.0 % with a sample at 80.85 %, and the best buffer with every sample at specification was 0.3 %.

Table 3.5 The Challenge Over 24 h on Five Seeds

Configuration	Pass	Tons	Worst (%)	Buffer low (%)
shipped	0/5	2443–2459	70.7–80.6	0
suggested	5/5	2085–2123	97.6	22.7–29.5
laboratory correction off	5/5	2083–2123	97.6	22.1–29.5
no push, 900 °C	0/5	2101–2145	88.7	15.6–22.2
no push, 140 t/h	0/5	1910–1962	94.3	30.2–36.6
fixed ratio 1690	5/5	2074–2111	98.7	20.4–27.5
supervisor off	0/5	2798–2817	54.5–61.5	0
priority Quality	0/5	2893–2912	25.1–32.8	13.0–20.0
override at 30 %	0/5	2183–2231	97.6	10.4–14.5
MPC at 30 %	5/5	2225–2297	97.6–97.8	33.8–36.0
One plan at 30 %	0/5	2159–2192	99.1	0–4.5
MPC at 50 %	5/5	2090–2160	97.9–98.0	53.9–56.8

LP bound with the forecast known and the design's 38.4 kg of hydrogen per tonne of pellets: 2417 t above a 20 % buffer, 2361 above 30 and 2248 above 50 %.

3.6.4 The Challenge

The challenge runs 24 h on the forecast with a richer lot at hour 3. It is met when the true metallization of every laboratory sample is at least 94 %, the true buffer never falls below 20 %, no cell rises above the sticking curve and the day makes at least 2050 t, with every parameter as shipped. Table 3.5 gives the configurations on five seeds. As shipped, with a fixed ratio of 1690 Nm^3/t, 900 °C, 140 t/h and the supervisor off, the day makes 2443 to 2459 t but runs the buffer dry at hour 15.7, and nine of 13 samples fall short. The suggested settings put the metallization loop in Auto at 95.5 %, set the bustle station to 935 °C and the push in Auto, raise the feed to 160 t/h and turn the supervisor on at Tons from 50 % with a gain of 3. They pass on five of five seeds with 2085 to 2123 t, a worst sample of 97.6 % and the buffer bottoming at 22.7 to 29.5 %. Revealed in mid-run at hour 0.75 they still meet the target with 2082.6 t, and at hour 3 they miss it on tons with 2048.4 t.

The table separates what the day needs from what it does not. The hotter bustle is what lets the compressor carry 160 t/h. Without the push the same feed leaves six or seven samples short at 88.7 %, and at 140 t/h the day falls short of the tons. The metallization loop and the laboratory correction are not needed, since a fixed ratio of 1690 Nm^3/t passes with 56.65 kg of hydrogen per tonne against 55.2 kg/t, so the loop's value on this day is the hydrogen it saves. The MPC plans the feed on the wind forecast, which the override does not have. At the suggested 50 % setpoint it makes 2129 t on the default seed against the override's 2096 t, and with the setpoint at 30 % it makes 2225 to 2297 t with the buffer at 33.8 % or more, where the override lets it fall to 10.4 to 14.5 % and One plan, with the same forecast, to 0 to 4.5 %. The advantage over the override is information, while the advantage over One plan is the freedom to shape the feed. A linear program that knows the forecast and holds the design's hydrogen per tonne bounds the day at 2361 t above a 30 % buffer, so the MPC at that floor leaves 64 to 136 t unmade. Figure 3.4 draws the table's ranges against the three limits of the challenge.

3.6.5 Verification Against the Ledgers and the Literature

The ledgers test the code and the literature tests the numbers, and the two are reported separately. On the code side, over a disturbed 12 h run with a trip, a ratio step, a temperature step and a lot, the oxygen fed less the oxygen discharged, the oxygen removed by the gas and the change of inventory closes to 3.8×10^{-9} of the oxygen removed, the buffer to 10^{-9} kmol, and the solids energy to below 5×10^{-10} of the reaction heat in three disturbed runs. Over 1.85 million reducing cell-steps no cell's gas passed its wüstite wall, the largest excursion being 1.1×10^{-16}. The design point was reproduced by hand in every printed digit. The suite of the application passes 45 of 45 checks and 10 of 10 completion checks. Against the literature, the agreement is of order only. The design carries 3.85 times the hydrogen on the pure-hydrogen study's basis against its 3.75 [15], a top gas 56 K hotter

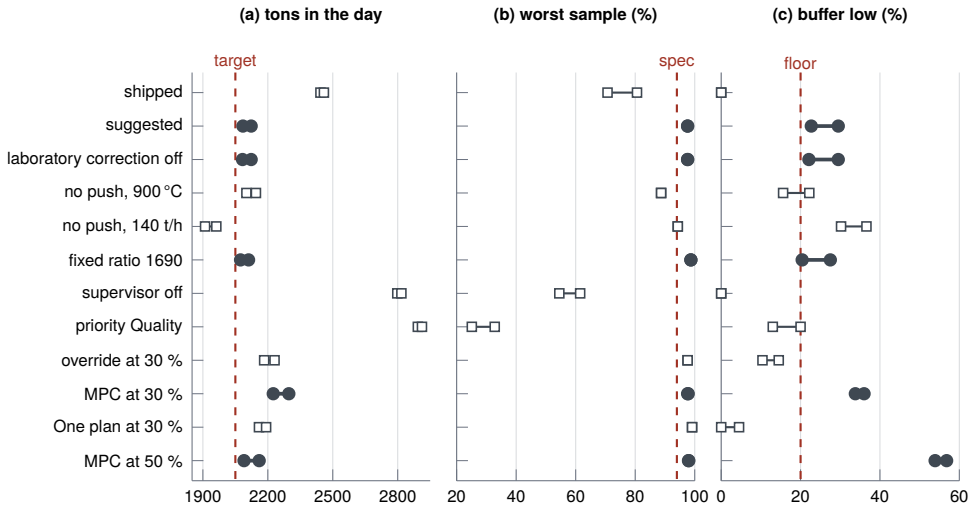


Figure 3.4 The challenge of Table 3.5 redrawn: for each configuration, the range over five seeds of (a) the tons made in the day, (b) the worst laboratory sample and (c) the lowest buffer, against the 2050 t target, the 94 % specification and the 20 % floor (dashed). Filled circles passed on five of five seeds, open squares on none. Notice that the settings that make the most tons fail on quality or on the buffer, that removing the push fails in both of its rows, and that at a 30 % setpoint the MPC keeps the buffer at 33.8 % or more where the override lets it fall to 10.4 %.

than its adiabatic balance, and 52.2 kg of hydrogen consumed per tonne of iron against 54 kg for full reduction [1]. A pellet reaching 95 % in 2.29 h matches the order of a zone a few meters tall in the pure-hydrogen simulations [6], but the rate constants were chosen, not fitted.

The shaft's product is the charge of the arc furnace application of Chapter 5, which declared its DRI before this model existed. The arc furnace declares 94 % metallization, no carbon, 92 % total iron, 4.5 % gangue and a cold charge at 25 °C, with about 600 °C for a hot charge. At design this shaft makes 94.5 %, no carbon, 93.7 % total iron and 4.85 % gangue at 899 °C. It matches in metallization and carbon, and because the model has no cooling zone its product would need to be cooled or charged hot. Neither application reads the other.

3.7 Technology in Practice

Figure 3.6 draws the loops of Section 3.5 with their hardware.

What Measures It

The product's metallization has no on-line instrument in this model. It is inferred from the top gas and checked by the laboratory. The model's top-gas analyzer, a lag of two minutes with 0.05 points of noise, stands for an on-line analyzer of hydrogen and steam. One common family of such analyzers, the process gas chromatograph, injects a measured sample into a column at intervals and tells the species apart by the time each takes to leave it, so it answers once per analysis [19]. The model's laboratory, a spot sample every two hours reported half an hour late with 0.3 points of noise, stands for an assay of the kind described below.

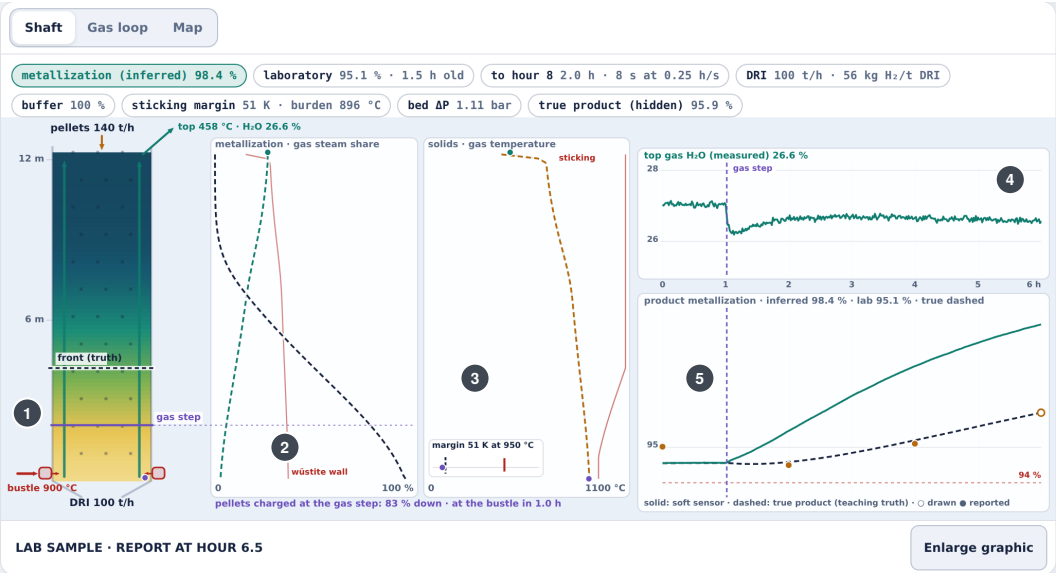


Figure 3.5 Case 1 at hour 6, five hours after the gas ratio was raised from 1690 to 1800 Nm^3/t on the Shaft view in Manual, cropped to the plant view. The bed is shaded by metallization with the reduction front marked as teaching truth and the pellets charged at the step drawn 83 % of the way down, one hour from the bustle (1). Beside it, the gas's steam share against the wüstite wall (2) and the temperatures against the sticking curve with a 51 K margin at the burden thermocouple's 896 °C (3). The top-gas plot has stepped to 26.6 % at the marker (4), while the product plot shows the soft sensor at 98.4 %, the true product at 95.9 % (dashed) and the laboratory's 95.1 % from the sample of hour 4 (5).

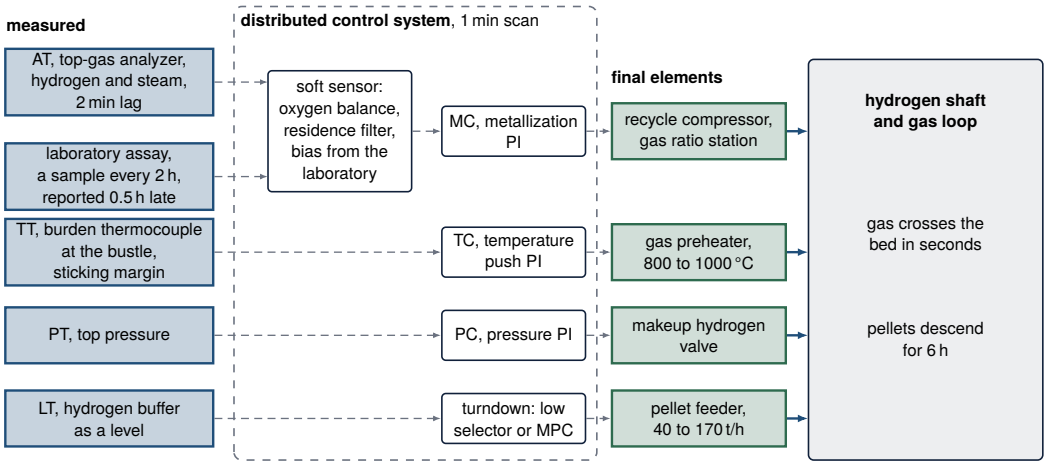


Figure 3.6 The loops of the simulated shaft drawn with their hardware. The metallization controller MC reads the soft sensor, which balances the oxygen on the top-gas analysis and takes its bias from the late laboratory assay, and sets the gas ratio that the recycle compressor delivers. The push TC holds the sticking margin on the gas preheater, PC holds the top pressure on the makeup hydrogen, and the turndown caps the pellet feed (or, at priority Quality, the gas ratio) when the buffer runs low.

TECHNOLOGY Sensor: The Laboratory Assay of Metallization

At the first pilot shaft in the United States run on hydrogen, pellet samples were discharged into sealed metal bins after cooling to room temperature in nitrogen. A sample of one kilogram was crushed and pulverized, and 100 mg of the powder was analyzed, the metallic iron by an iron(III) chloride titration under ISO 16878 and the total iron by a titanium(III) chloride titration under ISO 2597-2. Metallization is the moles of metallic iron over the moles of all the iron in the product [10]. Such an assay describes the pellets that left at the hour of sampling, and it reaches the control room only after the sample has been cooled, ground and titrated.

What Moves It

In the model the compressor delivers the requested bustle gas up to its capacity, the heater and the feeder follow their commands through lags of 0.1 h, and the makeup hydrogen closes the pressure.

TECHNOLOGY Final Element: The Recycle Compressor and the Gas Preheater

In the same pilot, the top gas of hydrogen and steam passed a filter for dust and fines and a condenser for its moisture, a compressor returned the remaining hydrogen to the feed, and excess hydrogen was burned in a thermal oxidizer. A preheater brought the gas to its inlet temperature before a bustle fed it into the shaft. A recycle of 400 liters a minute with 870 °C of external heating gave 91.8 % average metallization, against 67.8 % at 100 liters a minute and 850 °C [10]. A centrifugal compressor surges, violently and possibly destructively, when its flow falls too low against too high a pressure difference, and a recycle valve from discharge to suction prevents it [19].

Where the Controller Runs

Control loops in a plant are often nested or cascaded, with cycle times from milliseconds to minutes, and a distributed control system uses a supervisory loop to mediate a group of localized controllers [20]. In the application the plant is integrated every five seconds and the controllers execute every minute. On such a platform the metallization loop, the temperature push and the pressure loop are PI function blocks of the regulatory layer, the soft sensor is a calculation block whose bias is moved by each accepted laboratory report, and the turndown is a low selector with its reset fed back from the applied value. The buffer MPC, re-solved every half hour on the wind forecast, sits in the supervisory layer and writes a cap on the pellet feed.

IN PRACTICE Keeping a Soft Sensor Honest in Service

A soft sensor in service must be maintained and tuned regularly, because drifts of the process and of its instruments make its performance deteriorate. Most soft sensors have no automated maintenance, so their quality control is manual and a significant cost, and whether a model still works is often judged by an operator's visual comparison of its predictions with the target values. A drift of an instrument calls for recalibration or adaptation of the soft sensor, not for action on the process. Laboratory values arrive far less often than the automatic measurements, and where material takes time to pass through a process, measurements taken at the same moment in different places must be synchronized by that delay [8].

What Is Modernizing It

The first development is the pilot shaft run on hydrogen. A finite element model of the first such pilot in the United States predicted its average metallization and its internal temperatures, and was used extensively in commissioning the pilot and in predicting its startup for given operating parameters [10]. The second is the detailed model of the industrial shaft. A two-dimensional model of the reduction zone with a pellet-scale kinetic model shows that complete metallization can be reached with hydrogen in a reactor smaller than the shafts that now run on gas made from natural gas [1].

3.8 Limitations

The model is a reduction zone without a transition or cooling zone, so its DRI leaves at 899 °C and cannot be compared with a cold product without a model the application does not have. The kinetics are teaching values with a chemical-control shape, the bed is one-dimensional with uniform pellets, and the walls lose no heat, so profiles are plausible rather than predicted for a particular shaft. The bustle composition is declared, which puts the makeup 3 % above an exact species balance. The sticking curve's level has a laboratory source but its shape does not, and the cluster index and the pressure-drop law are teaching relations.

The soft sensor's residence filter is first order while the bed is close to plug flow, and its six-hour gate admits reports before the product has caught up, which leaves a 0.7-point error after case 1's step with a correct analysis. The MPC models only the buffer and plans on a forecast whose error it never sees, the override has no forecast at all, and the electrolyzer has no minimum load. Every limit of the scorecard is a teaching value, and passing a case in simulation qualifies no real plant's controls.

OPEN CHALLENGE **Running a Shaft on Wind Hydrogen**

A steelworks usually asks for hydrogen at a constant rate, while wind and sun deliver power intermittently, and storage must turn the one into the other. Alkaline electrolyzers, the most mature type, must run between 20 and 100 % of their rated power, which limits how closely they can follow the wind [16]. One pilot program designs its shaft to run on hydrogen or on a mixture with natural gas, so that production continues when hydrogen is short [10]. Progress would be a shaft, buffer and electrolyzer run as one plant.

3.9 Conclusion

A hydrogen shaft furnace is a plant whose quality is inferred, delayed and fed from a variable supply, and the simulator of this chapter logs the inference, the laboratory and the true product side by side on a moving bed whose ledgers close to 10^{-9}. Raising the gas moves the top gas within three minutes and the product half its way only after 6.73 h, while the soft sensor gets there in 2.18 h and overstates the product by 2.4 points at hour 8. The heat the hydrogen must carry, not the 36.9 % wüstite ceiling, sets a design flow of four times the hydrogen consumed and a 27.0 % top gas. A laboratory correction recovers a wrong pellet analysis over a day, a temperature push buys rate up to a sticking margin, and a six-hour lull with the shipped buffer cannot keep both full tons and full quality on any of 984 settings and schedules. On the challenge day the suggested settings pass on five of five seeds, and a forecast-based MPC adds tons because it knows the wind. The model omits the cooling zone, carbon and a computed loop composition. A fuller model would add a cooling zone that delivers a cold product, a soft sensor with an inventory state of the bed, measured kinetics and an electrolyzer with a minimum load. The shaft is a compact case for learning that the measurement a controller trusts may describe the plant now while the product describes the plant six hours ago.

KEY POINTS

- The gas crosses the shaft in seconds and the iron takes six hours, so the top gas answers within minutes and the product over hours.
- The soft sensor reads the whole bed, not the discharge, and overstates the product by 2.4 points at hour 8.
- The heat the hydrogen must carry, not the 36.9 % wüstite ceiling, sets a design flow of four times the hydrogen consumed.
- Through a six-hour lull the shipped buffer cannot keep both full tons and full quality.

Where else this applies. Volume 5, Chapter 4 grades a charge that descends for hours, and Volume 5, Chapter 10 corrects a soft sensor against late samples.

References

[1] F. Patisson and O. Mirgaux, "Hydrogen ironmaking: How it works," *Metals*, vol. 10, no. 7, p. 922, 2020, open copy: HAL hal-02968422.

[2] V. Vogl, M. Åhman, and L. J. Nilsson, "Assessment of hydrogen direct reduction for fossil-free steelmaking," *Journal of Cleaner Production*, vol. 203, pp. 736–745, 2018.

[3] H. Hamadeh, O. Mirgaux, and F. Patisson, "Detailed modeling of the direct reduction of iron ore in a shaft furnace," *Materials*, vol. 11, no. 10, p. 1865, 2018.

[4] D. R. Parisi and M. A. Laborde, "Modeling of counter current moving bed gas-solid reactor used in direct reduction of iron ore," *Chemical Engineering Journal*, vol. 104, no. 1–3, pp. 35–43, 2004.

[5] R. Béchara, H. Hamadeh, O. Mirgaux, and F. Patisson, "Optimization of the iron ore direct reduction process through multiscale process modeling," *Materials*, vol. 11, no. 7, p. 1094, 2018.

[6] A. Ranzani da Costa, D. Wagner, F. Patisson, and D. Ablitzer, "Modélisation d'un four à cuve de réduction directe du minerai de fer par l'hydrogène pur," *Revue de Métallurgie*, vol. 106, no. 10, pp. 434–439, 2009, preprint: arXiv:0911.4949.

[7] A. Ranzani da Costa, D. Wagner, and F. Patisson, "Modelling a new, low CO_2 emissions, hydrogen steelmaking process," *Journal of Cleaner Production*, vol. 46, pp. 27–35, 2013, preprint: arXiv:1402.1715.

[8] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

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

[10] A. Meshram, Y. Korobeinikov, D. D. Nogare, A. Zugliano, J. Govro, R. J. OMalley, and S. Sridhar, "Modeling the first hydrogen direct reduction pilot reactor for ironmaking in the USA using finite element analysis and its validation using pilot plant trial data," *Processes*, vol. 11, no. 12, p. 3346, 2023.

[11] A. Szemalikowska, "Advancements in transition to green steel in Europe through direct reduced iron production," *Materials Science and Welding Technologies*, vol. 68, no. 2, pp. 1–9, 2024.

[12] A. Rist and N. Meysson, "A dual graphic representation of the blast-furnace mass and heat balances," *JOM*, vol. 19, no. 4, pp. 50–59, 1967.

[13] A. Pineau, N. Kanari, and I. Gaballah, "Kinetics of reduction of iron oxides by H_2: Part I: Low temperature reduction of hematite," *Thermochimica Acta*, vol. 447, no. 1, pp. 89–100, 2006.

[14] D. Spreitzer and J. Schenk, "Reduction of iron oxides with hydrogen—A review," *steel research international*, vol. 90, no. 10, p. 1900108, 2019.

[15] A. Ranzani da Costa, "La réduction du minerai de fer par l'hydrogène : étude cinétique, phénomène de collage et modélisation," Ph.D. dissertation, Institut National Polytechnique de Lorraine, Nancy, France, 2011, hAL tel-01204934.

[16] F. Superchi, A. Mati, C. Carcasci, and A. Bianchini, "Techno-economic analysis of wind-powered green hydrogen production to facilitate the decarbonization of hard-to-abate sectors: A case study on steelmaking," *Applied Energy*, vol. 342, p. 121198, 2023.

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

[18] M. W. Chase, Jr., *NIST-JANAF Thermochemical Tables*, 4th ed., ser. Journal of Physical and Chemical Reference Data, Monograph No. 9. Washington, DC and Woodbury, NY: American Chemical Society and American Institute of Physics for the National Institute of Standards and Technology, 1998.

[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.

3.10 Exercises

3.1 Concept check · 2 min

In case 1 the gas ratio is raised from 1690 to 1800 $\mathrm{Nm^3/t}$ at hour 1. At hour 8 the corrected soft sensor reads 99.0 % while the true product is 96.6 %. Why does the soft sensor run ahead of the product?

(a) The top-gas analyzer's two-minute lag lets the balance overshoot its final value
(b) The first laboratory reports after the step have raised the bias by 2.4 points
(c) The soft sensor's residence filter is set shorter than the residence of the bed
(d) The balance counts the oxygen the whole bed loses now, not that of the pellets leaving

3.2 Concept check · 1 min

At the design point the top gas leaves with 27.0 % steam against a wüstite ceiling of 36.9 %. What sets the bustle flow at about four times the hydrogen the reduction consumes?

(a) The heat the hydrogen must carry to warm the pellets and drive an endothermic reduction
(b) The wüstite equilibrium, which the top gas would cross at any lower flow
(c) The recycle compressor, whose capacity caps the ratio at the design value
(d) The purge, which removes two percent of the loop gas on every pass

3.3 Concept check · 2 min

The laboratory correction uses a report only if the gas ratio and the feed stayed within 3 % and 3 t/h over the 6 h before the sample was drawn. What is this gate for?

(a) The laboratory's 0.3 points of noise grow too large while the bed is moving
(b) A balance read in a transient counts oxygen the bed stores or releases as a model error
(c) The metallization loop tracks the applied ratio during a feed move and cannot use a new bias
(d) Reports of samples drawn during a transient arrive later than the usual half hour

3.4 Theory and calculation · 6 min

The shaft holds $n_{Fe} = 10\,153$ kmol of iron and is fed 140 t/h of pellets. Take $M_{Fe} = 55.845\,\mathrm{kg/kmol}$.

(a) Compute the iron fed, $\dot{F}_{Fe}$, at 67.5 % Fe and the residence $n_{Fe}/\dot{F}_{Fe}$. *Checkpoint:* 1692.2 kmol/h; 6.000 h.

(b) The lot of cases 2 and 5 carries 68.5 % Fe at the same 140 t/h. Compute the new iron feed and residence, and the change of the residence in percent. *Checkpoint:* 1717.3 kmol/h; 5.912 h, 1.5 % shorter.

(c) The grade feedforward scales the ratio by the analysis. Compute the new ratio from 1690 $\mathrm{Nm^3/t}$, and the pellet feed, to the nearest 1 t/h, that would bring the iron fed back to its design value. *Checkpoint:* 1715 $\mathrm{Nm^3/t}$; 138 t/h.

3.5 Theory and calculation · 6 min

The wüstite step has $K_2 = \mathrm{e}^{A_2 - B_2/T}$ with $A_2 = 1.0836$ and $B_2 = 1899$ K, and the steam share of the reducing gas over wüstite cannot exceed $s_{eq} = K_2/(1+K_2)$.

(a) Compute K_2 and s_{eq} at 800 and 900 °C. Give s_{eq} to ± 0.1 points. *Checkpoint:* 0.504 and 33.5 %; 0.586 and 36.9 %.

(b) The bustle gas is 95 % reducing with $x_W = 0.025$, so the steam share of its reducing gas is $x_W/0.95$. Each mole of reducing gas that leaves at a steam share s_{top} has taken $s_{top} - x_W/0.95$ moles of oxygen. Compute the moles of reducing gas, and of hydrogen, per mole of oxygen at the design top gas of 27.04 %. *Checkpoint:* 4.10 and 3.99.

(c) Repeat (b) with the top gas at the 900 °C ceiling, and say how many times that minimum the design carries. *Checkpoint:* 2.92 and 2.84; about 1.4 times.

3.6 Theory and calculation · 7 min

Case 6 asks whether the shaft can keep 140 t/h and 94 % through the six-hour lull. Use $\dot{F}_{Fe} = 1692.2$ kmol/h, hematite at 1.5 oxygen per iron, a product at $m = 0.94$ whose remaining oxygen is wüstite, $M_{H_2} = 2.016\,\mathrm{kg/kmol}$, the 30 t buffer with its 20 % floor, the 600 MW wind farm at a capacity factor of 0.06 through the lull and $e_{el} = 52\,\mathrm{kWh/kg}$.

(a) Compute the hydrogen consumed by the oxygen alone over the six hours. *Checkpoint:* 29.5 t.

(b) Compute the hydrogen the buffer can give above its floor and what the electrolyzer makes in the lull, and the shortfall. *Checkpoint:* 24.0 t and 4.2 t; 1.3 t short.

(c) At what constant pellet feed would this count balance? Say why the count alone does not prove the lull feasible or infeasible. *Checkpoint:* about 134 t/h.

3.7 App challenge · 8 min

Load scenario 1, "Raise the gas: the top answers now, the product later," at the default seed in Manual. Run at 0.1 h/s and press Pause when the clock reads 1.0 h (anywhere from 0.95 to 1.05 h). Set Gas ratio to 1800 $\mathrm{Nm^3/t}$, press Run again and let the case pause itself at hour 8. Export the CSV file and read the rows against `Time (h)`. Give metallizations to ±0.1 points.

(a) From `top_h2o_pct`, give the average steam share of the top gas over the half hour before the step and over the half hour after it, each to ±0.1 points. *Checkpoint:* about 27.0 %, then about 26.3 %.

(b) At 8.0 h, read the corrected soft sensor `met_inf_pct` and the true product `truth_met_pct`. *Checkpoint:* 99.0 % and 96.6 %.

(c) The product lines out at 98.52 % from 94.54 %. What share of that move has each reading made at hour 8? *Checkpoint:* about 0.52 and 1.12.

3.8 App challenge · 8 min

Load scenario 3, "The soft sensor against the laboratory," at the default seed with no changes, so the header is in Auto and Laboratory correction is On. Run at 1 h/s until the case pauses at hour 30. Then switch Laboratory correction Off, press Reset and run again to hour 30. Read the scorecard's DRI tile after each run, and the CSV columns `truth_met_pct`, `met_bias_pts` and `ratio_applied` at 30.0 h. Give metallizations to ±0.1 points and the ratio to ±2 $\mathrm{Nm^3/t}$.

(a) With the correction on: the true product, the bias and the applied ratio at hour 30, and the samples below 94 % on the tile. *Checkpoint:* 94.4 %, −2.56 points, 1730 $\mathrm{Nm^3/t}$; 10 of 16.

(b) With it off: the true product at hour 30 and the samples below 94 %. *Checkpoint:* 92.6 %; 13 of 16.

(c) Explain why the correction recovers the product so slowly, and name a cheaper remedy the chapter offers for a known change of the pellets. *Checkpoint:* each report describes the product of hours ago; the grade feedforward of case 5.

3.9 App challenge · 9 min

Load the challenge, "a day on the wind forecast," at the default seed. Run at 1 h/s. The day pauses at hour 24 with the verdict held on the Challenge target tile. Report tons to ±2 t, the worst sample to ±0.1 points and the buffer low to ±1 %.

(a) Run as shipped. Read the tons, the worst sample, the samples below 94 % and the buffer low. *Checkpoint:* 2452 t, 74.6 %, 9 of 13, 0 %.

(b) Reveal the suggested settings, press Reset and run the day again. Read the same four values and the verdict. *Checkpoint:* 2096 t, 97.6 %, 0 of 13, 28 %; met.

(c) Keep the suggested settings, set Turndown method to MPC in the Advanced settings of the production card, press Reset and run again. Read the tons and the buffer low, and explain where the extra tons come from. *Checkpoint:* 2128 t, 56 %; the MPC plans the feed on the wind forecast.

3.10 Application area · 3 min

Case 7 trips the recycle compressor for half an hour, and the product stays below 94 % from hour 2.33 to hour 9.33 while the burden keeps moving.

(a) At the design DRI rate of 100.87 t/h, about how much off-specification iron does the trip cost? *Checkpoint:* about 706 t.

(b) The feed may move at most 30 t/h per hour. How long would cutting it from 140 to 40 t/h take, and what do plants do instead when the gas is lost? *Checkpoint:* 3.3 h; they stop the burden.

3.11 Application area · 4 min

The model has no cooling zone, so its DRI leaves at 899 °C. The arc furnace of Chapter 5 declares a cold charge at 25 °C, with about 600 °C for a hot charge. Use the model's solids heat capacity, 65 kJ per kmol of iron per kelvin, and the product's 93.69 % total iron.

(a) How much heat per tonne of DRI must a cooling zone remove to deliver a cold charge? *Checkpoint:* about 953 MJ/t.

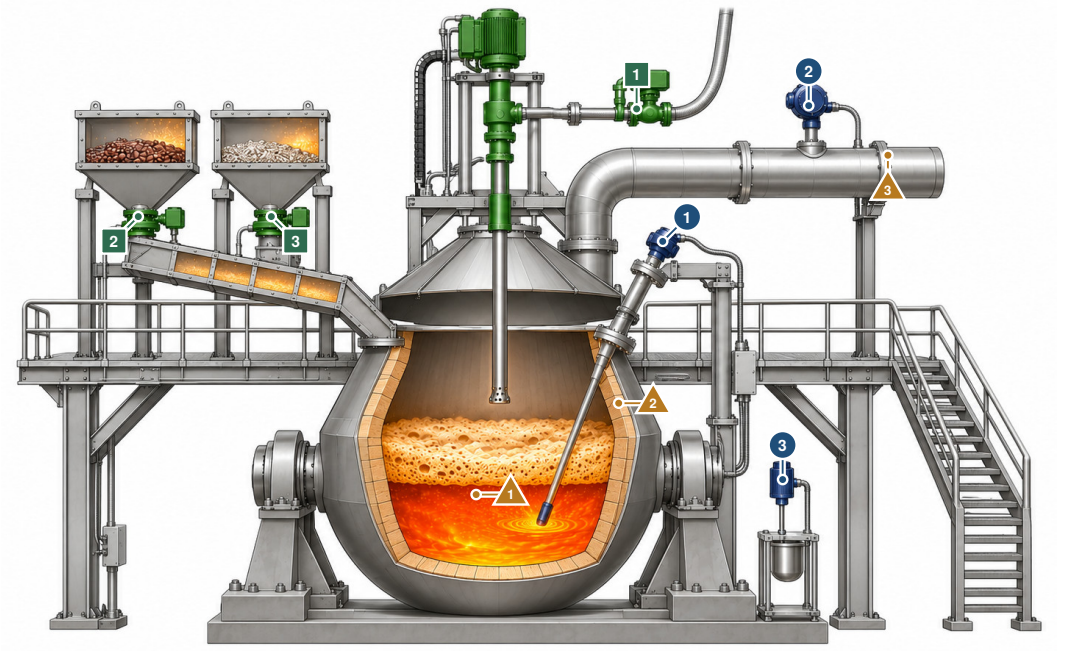
(b) How much of it would a hot charge at 600 °C keep, and why might a plant still cool the product? *Checkpoint:* about 627 MJ/t kept; storage and handling.

3.12 Review the agent · 5 min

An AI assistant reviewing the design writes: "The top gas leaves at 27.0 % steam, but the wüstite ceiling at 900 °C is 36.9 %. A quarter of the hydrogen does no work. Cut the gas ratio by the same share, from 1690 to 1235 Nm^3/t. The top gas will then run at the ceiling, the metallization will hold and the plant will save about 27 % of its hydrogen." Find the flaw, and say how you would check it in the app.

CHAPTER 4

Basic Oxygen Furnace Simulator for Endpoint Control With One Late Measurement



The controller moves 1 oxygen blown through the lance 2 iron ore added to cool the bath 3 ferrosilicon added to warm a heat that would end cold

The controller reads 1 sublance probe: carbon and temperature late in the blow 2 off-gas analyzer: how fast carbon leaves as gas 3 carbon and temperature tested at turndown

What disturbs it 1 hot metal and scrap that vary from heat to heat 2 a lining that loses more heat each heat 3 an error in the off-gas flow reading

Notice that the lance blows oxygen from above the bath while the sublance dips into it once, late in the blow, after which ore can cool the heat but only more oxygen, burning the carbon below its aim, can warm it. Illustration, not to scale.

IN PLAIN TERMS **One Blow, Judged at the End**

The main idea. A converter makes steel by blowing pure oxygen onto liquid iron. The oxygen burns out the extra carbon, and the burning heats each batch, called a heat, to the temperature the steel needs.

The challenge. Oxygen once blown cannot be taken back, and a probe checks the bath only once, near the end. After that the heat can be cooled with iron ore, but warming it would burn out too much carbon.

Why it matters. A heat that misses must be blown again, or it wastes good iron in the slag on top.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the fact box of Section 4.1, the history box of Section 4.2, the try-it boxes and the history box of Section 4.4, the fact box of Section 4.5, and the key points
Study	35 min	The whole chapter with its pause points and the twelve exercises
Practice	10 min	The first page, Section 4.5, Section 4.6, Section 4.7 and the open challenge after it

The control question A converter heat is blown once and judged at its end, and neither the carbon nor the temperature can be taken back. What can a charge model plan, what can one late measurement correct when the only correction left can cool but not heat, and what can only learning between heats follow?

Learning objectives

- Derive the decarburization rate, its time constant below the critical carbon and the iron burned per kilogram of carbon.
- Explain why a static charge model hits the carbon more often than the temperature.
- Size a hot margin for a one-way correction after the sublance.
- Compare the error of an integrated off-gas estimate with that of the same signal read as a rate.
- Compute the lag of run-to-run learning on a drifting heat loss and choose its weight.

Prerequisites Section 1.4.4 (charge model, rate endpoint, one-way correction and EWMA lag), 1.6.1 and 1.6.4 (engine, scenarios and the scorecard).

The app https://apmonitor.com/apps/bof-converter/

See also Volume 8, Chapter 9 for EWMA run-to-run correction from late results, Volume 8, Chapter 10 for learning between repetitions and Volume 8, Chapter 4 for a batch graded only at its end. Volume 5, Chapter 13 for an integrated estimate reset by late fixes and Volume 5, Chapter 4 for acting ahead of a late analyzer.

4.1 Introduction

A batch has no steady state to regulate around. What it delivers is decided by the state it reaches at the end, so batch control is largely a matter of planning the batch from what is known of its start, correcting the plan once or twice on the way, and using the result of each batch to improve the plan for the next. Run-to-run and batch-to-batch control formalize the last of these steps, and surveys have tied them to iterative learning and repetitive control as one family of methods for processes that repeat [1, 2]. Learning between the layers of a build in Volume 8, Chapter 10 is a member of the same family.

The basic oxygen furnace, or converter, is an unforgiving example. The steel sector is responsible for about 7 % of anthropogenic carbon dioxide emissions, and the blast furnace route that dominates its iron production [3] makes its steel in the converter. Each heat is charged with hot metal and scrap and blown with oxygen until the carbon falls to its aim, and the heat of oxidation brings the bath to tapping temperature on the way. One actuator, the oxygen volume, sets both endpoints. Too little oxygen leaves carbon that must be removed by a reblow, and too much oxidizes iron, because below a critical carbon the reaction becomes limited by carbon transport in the bath rather than by oxygen supply [4, 5]. A block diagram with one input and two outputs hides what makes the problem hard, namely that oxygen once blown cannot be withdrawn, that the two endpoints are coupled through the heat of oxidation, and that the only in-blow measurement arrives with 15 % of the blow left.

✦ DID YOU KNOW? Twenty-Eight Percent of American Steel

The United States made about 81 million metric tons of raw steel in 2024, 28 % of it in basic oxygen furnaces and the rest in electric arc furnaces [6], the subject of Chapter 5. A modern converter can make steel at 500 tons an hour [7]. The simulated heat of this chapter is 300 t, blown in about 21 min.

The cases of this chapter are built around the belief that a continuous measurement is always worth more than a single sample, and show where it fails. An off-gas analyzer reports the decarburization rate every second of the blow, while a sublance probe measures the bath once, at 85 % of the planned oxygen. Integrated from the start of the heat, the continuous signal accumulates every percent of flow meter error into the carbon estimate, and the single late sample does not. The broader class is batch endpoint control under a terminal specification, which appears in plasma etch, batch polymerization, fermentation, heat treatment and others, wherever a recipe must be fixed before the batch and

corrected from few measurements during it [2, 8]. The shape of the problem is the same in all of them. A model maps the measured initial state to a recipe, a late measurement re-plans the remainder, the correction available after that measurement is usually one-directional, and slow drift in the equipment is visible only across batches.

This chapter works with a converter simulator in which the static charge model, the sublance re-plan, the off-gas integral, the rate-based endpoint, the temperature margin and run-to-run learning can be switched on one at a time and compared over the same campaign of heats. All of these strategies are established practice. What is particular to this plant is a heat and mass balance of one heat, with the charge variation, instrument errors and controller logic written out completely enough to rebuild. The ten scenarios separate what feedforward from the charge can decide, what one late measurement can correct, and what only learning between heats can follow. Section 1.4.4 introduces these ideas on a smaller converter, and this chapter measures them on a full one. Section 4.2 reviews converter control and Section 4.3 gives the model of one heat, its oxygen distribution, heat balance and heat-to-heat variation. Section 4.4 describes the strategies as implemented, from the static charge model through the sublance re-plan and the two uses of the off-gas signal to reblowing and run-to-run learning. Section 4.5 reports the measured campaigns, from the fixed recipe through the late measurement, the integral against the rate, drift and learning and the low carbon aim to the graded challenge. Section 4.6 describes the real instruments, final elements and control platforms, Section 4.7 states the limitations, Section 4.8 concludes, and Section 4.9 holds the exercises.

4.2 Endpoint Control in Oxygen Steelmaking

Blowing pure oxygen onto a bath of hot metal became an industrial steelmaking process in the early 1950s, and the oxygen converter marked its fiftieth year of production in 2002 [9]. Without a model or an in-blow measurement, the decision to stop rests on a sample taken after turndown and a reblow when the sample is wrong.

> ⧖ HISTORY **From Open Hearth to Oxygen (1954 to the 1980s)**
>
> The basic oxygen furnace was introduced in North America in 1954. The basic open hearth it displaced had made steel there since well before 1900, but it consumed huge tonnages of refractories, needed many hours of labor per ton and was slow, and as labor and energy costs rose the change became inevitable. By 1969 the two processes each made 60 million tons a year. The oxygen process runs on its own heat with 70 percent or more hot metal in the charge. Static and dynamic models and sensors for the turndown have occupied steelmaking technologists since the 1960s. The ladle furnaces of the 1980s eased the temperature problem, because they allowed the temperature to be adjusted after tapping [7].

Computer control of the basic oxygen process was reported by 1964 [10]. The static part of such control is the charge calculation, a heat and mass balance run before the blow that turns the measured hot metal weight, composition and temperature and the scrap weight into an oxygen volume and a coolant addition. These models inherit every error in the charge data and in their own constants. Dynamic control, which corrects the blow from information gathered while it is in progress, was in operation within a few years [11], and by 1983 the control of oxygen steelmaking was treated as a subject of its own in reviews of iron and steelmaking control [12].

The in-blow information comes from two sources. A sublance lowers a disposable probe into the bath near the end of the blow and returns a carbon and a temperature, which the model uses to re-plan the remaining oxygen and coolant. Off-gas analysis measures the flow and composition of the gas leaving the vessel, from which the decarburization rate follows continuously. Model-based estimation of the carbon content from that rate was developed toward closed-loop control in the late 1990s [13, 14], and the gradient of the relative decarburization rate has been used more recently to predict the endpoint [15]. Static models are still refined, for example by combining heat and mass balances with neural networks trained on tens of thousands of heats [16].

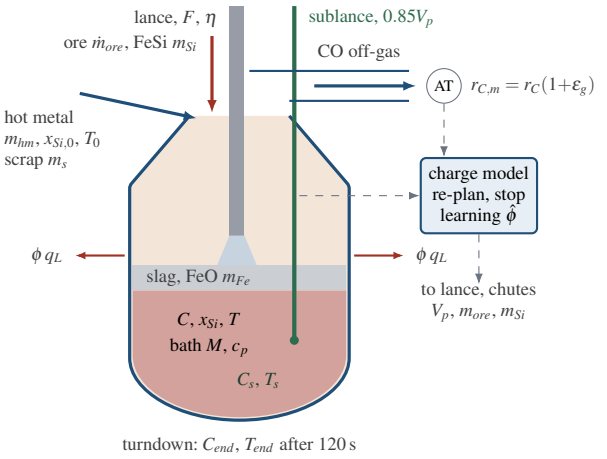


Figure 4.1 Basic oxygen furnace heat as modeled. Hot metal of mass m_{hm}, silicon $x_{Si,0}$ and temperature T_0 and scrap of mass m_s form a bath of mass M with carbon C, silicon x_{Si} and temperature T. The lance blows oxygen at flow F with efficiency η, iron ore and ferrosilicon are added as coolant and fuel, iron oxidized to FeO collects in the slag, and the vessel loses heat ϕq_L. The sublance measures C_s and T_s at 85 % of the planned oxygen V_p, the off-gas analyzer AT reports the decarburization rate with a flow error ε_g, and carbon and temperature are measured again at turndown.

The physics underneath has been modeled in detail. Multi-zone kinetic models treat the jet impact zone, the emulsion of metal droplets in slag and the bulk bath separately [5, 17], reaction models based on Gibbs free energy minimization follow slag and metal chemistry through the blow [18], and the change in decarburization mechanism at a critical carbon content has a theoretical interpretation of its own [4]. Learning between heats, the third element of the simulator, comes from a different industry. The exponentially weighted moving average (EWMA) controller of semiconductor manufacturing updates a model offset from each run's error, and its stability and its sensitivity to drift are well characterized [1, 8].

4.3 Model of One Heat

A nominal heat is charged with 250 t of hot metal at 1350 °C carrying 4.5 % carbon and 0.40 % silicon, and 50 t of scrap. Oxygen is blown at 750 $\mathrm{Nm^3/min}$ and the blow lasts about 21 min. The aim is 0.05 % carbon and 1650 °C, within ±0.01 % and ±10 °C. The vessel and its measurements are shown in Figure 4.1.

Table 4.1 classifies the variables. During a heat the model carries six states: carbon C and silicon x_{Si} in mass percent of the bath, temperature T, the iron oxidized into the slag m_{Fe}, the oxygen blown V and the blowing time t_b. The model heat-loss factor $\hat{\phi}$ changes only between heats.

4.3.1 Oxygen Distribution

Silicon burns first. Its oxidation is first order with time constant τ_{Si}, and it takes the oxygen it needs before carbon does, as shown in Equation (4.1).

$$r_{Si} = \min\left(\frac{x_{Si} M}{100\,\tau_{Si}},\, k_{Si} F\right), \qquad F_r = F - \frac{r_{Si}}{k_{Si}} \tag{4.1}$$

where r_{Si} is in kg/s, $M = m_{hm} + m_s$ is the bath mass in kilograms, F is the oxygen flow in $\mathrm{Nm^3/s}$, F_r is the oxygen left for carbon and $k_{Si} = 28/22.4 = 1.25\,\mathrm{kg/Nm^3}$ for $\mathrm{Si + O_2 \rightarrow SiO_2}$. At the nominal charge the silicon initially takes about 8 $\mathrm{Nm^3/s}$ of the 12.5 $\mathrm{Nm^3/s}$ supplied.

Table 4.1 Variables of One Heat and of the Campaign

Symbol	Role	Nominal	Range
C_{end}	CV, turndown carbon	0.05 %	±0.01 %
T_{end}	CV, turndown temperature	1650 °C	±10 °C
V_p	MV, planned oxygen	15 500 Nm^3	12 000–19 000
m_{ore}	MV, iron ore coolant	2.7 t	0–15
m_{Si}	MV, ferrosilicon silicon	0 kg	≥ 0
C, x_{Si}	states, bath carbon, silicon	4.5, 0.40 %	0.02–4.8
T	state, bath temperature	1350 °C	1250–1900
m_{Fe}	state, iron to slag	0	0–20 t
V	state, oxygen blown	0	0–19 000 Nm^3
t_b	state, blowing time	0	0–1500 s
$\hat{\phi}$	per heat, model loss factor	1	0.5–2
s	DV, charge variation scale	1	0–2
g	DV, off-gas error (1σ)	3 %	0–8
d	DV, heat loss growth	0	0–5 %/heat

Carbon takes the remaining oxygen with efficiency η while it is plentiful. Below the critical carbon C^* the rate is limited by carbon transport to the reaction zone and falls in proportion to the carbon left above the equilibrium value C_{eq}, Equation (4.2), the first form of Equation (1.4) in this chapter's symbols.

$$\begin{aligned} r_C &= \eta F_r k_C f(C) \\ f(C) &= \min\left(1,\ \max\left(0, \frac{C - C_{eq}}{C^* - C_{eq}}\right)\right) \end{aligned} \tag{4.2}$$

with $k_C = 24/22.4 = 1.071\,\mathrm{kg/Nm^3}$ for $\mathrm{C} + \frac{1}{2}\mathrm{O_2} \rightarrow \mathrm{CO}$. The share $1-\eta$ is taken to burn CO to CO_2 in the vessel and is not counted further. The oxygen that carbon could have taken below C^* and cannot oxidizes iron, Equation (4.3).

$$r_{Fe} = \eta F_r k_{Fe} \left(1 - f(C)\right) \tag{4.3}$$

where the iron oxidized per unit of oxygen for $2\mathrm{Fe} + \mathrm{O_2} \rightarrow 2\mathrm{FeO}$ is $k_{Fe} = 2(55.85)/22.4$, which is 4.99 kg/Nm^3. The composition and inventory states then follow from the rates, as shown in Equation (4.4).

$$\begin{aligned} \frac{\mathrm{d}C}{\mathrm{d}t} &= -\frac{100\, r_C}{M}, & \frac{\mathrm{d}x_{Si}}{\mathrm{d}t} &= -\frac{100\, r_{Si}}{M} \\ \frac{\mathrm{d}m_{Fe}}{\mathrm{d}t} &= r_{Fe}, & \frac{\mathrm{d}V}{\mathrm{d}t} &= F \end{aligned} \tag{4.4}$$

and the blowing time t_b advances at one second per second while oxygen is on.

The two regimes are drawn in Figure 4.2. Above C^* the plateau rate is $\eta F k_C = 12.7\,\mathrm{kg/s}$, about 0.25 % of the bath per minute. Below it, carbon decays exponentially with time constant $M(C^* - C_{eq})/(100\,\eta F k_C) \approx 78\,\mathrm{s}$, and the iron burned per kilogram of carbon removed is $(k_{Fe}/k_C)(1-f)/f$, about 47 kg at the 0.05 % aim and about 150 kg at 0.03 %. This is the reason the endpoint must be hit rather than overshot.

4.3.2 Heat Balance

The bath is one well-mixed volume whose temperature follows the heats of oxidation, less the scrap, the ore and the vessel losses, Equation (4.5).

$$\begin{aligned} M c_p \frac{\mathrm{d}T}{\mathrm{d}t} &= h_C r_C + h_{Si} r_{Si} + h_{Fe} r_{Fe} \\ &\quad - q_s - h_{ore} \dot{m}_{ore} - \phi\, q_L \end{aligned} \tag{4.5}$$

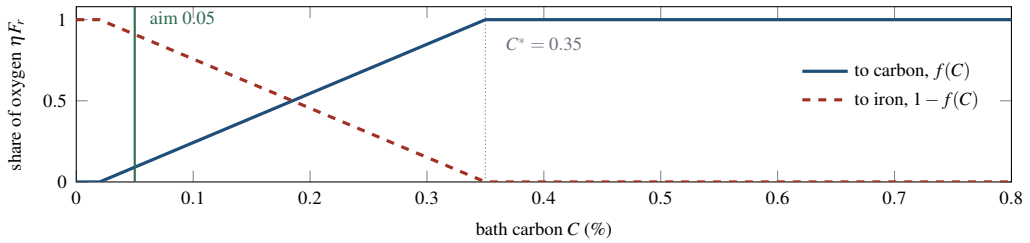


Figure 4.2 Division of the oxygen left after silicon between carbon and iron, from Equations (4.2) and (4.3), the share to carbon solid and the share to iron dashed. Above the critical carbon (dotted) all of it reaches carbon. Below it the share taken by carbon falls linearly to zero at $C_{eq} = 0.02\,\%$, and at the aim about 91 % of it oxidizes iron.

Table 4.2 Measurements Available to the Controller, With Their Errors (1σ)

Measurement	When	Error
bath mass M, scrap m_s	charge	$1.0s$, $0.7s$ t
hot metal silicon $x_{Si,0}$	charge	$0.03s$ %
hot metal temperature T_0	charge	$5s$ K
hot metal carbon C_0	—	assumed 4.5 %
sublance C_s, T_s	$V = 0.85V_p$	3 % of C, 4 °C
off-gas rate $r_{C,m}$	every second	ε_g, fixed per heat
turndown C_{end}, T_{end}	120 s after stop	2 % of C, 3 °C

in MJ/s, with Mc_p in MJ/K (about 246 MJ/K for 300 t), $h_C = 10$, $h_{Si} = 30$ and $h_{Fe} = 4.8\,\mathrm{MJ/kg}$. The scrap absorbs $h_s = 1.3\,\mathrm{MJ/kg}$ at a constant rate $q_s = m_s h_s/600$ over the first 600 s of blowing, the ore absorbs $h_{ore} = 4.0\,\mathrm{MJ/kg}$ as it is added, and the vessel loses $q_L = 42\,\mathrm{MJ/s}$ times the heat's loss factor ϕ while blowing. Temperature is frozen when there is neither blowing nor ore addition. One tonne of ore is worth about 16 °C, and one percent of heat loss over a 21 min blow about 2 °C. Below C^* the iron term dominates, because at the aim the iron burns at about 54 kg/s and heats the bath by about 1 K/s, so an overblown heat is also a hot one.

4.3.3 Heat-to-Heat Variation and Instruments

Every heat $k = 0, 1, \ldots$ draws its true charge and constants from independent standard normal deviates ξ, scaled by the charge variation s, Equation (4.6).

$$\begin{aligned} m_{hm} &= 250 + 2s\xi, \quad m_s = 50 + 1.5s\xi \quad \text{(t)} \\ C_0 &= 4.5 + 0.08s\xi, \quad x_{Si,0} = \max(0.1,\ 0.40 + 0.08s\xi) \\ T_0 &= 1350 + 12s\xi, \quad \varepsilon_g = g\xi/100 \\ \eta &= \mathrm{clamp}(0.95 + 0.012s\xi,\ 0.85,\ 1) \\ \phi &= 1 + dk/100 + 0.04s\xi \end{aligned} \tag{4.6}$$

The lining drift d is the growth of the heat loss in percent per heat. The charge model is told a measured charge rather than the true one, and the instrument errors are listed in Table 4.2. The hot metal carbon is not measured and is taken as 4.5 % on every heat.

The parameters are collected in Table 4.3. The heat starts at C_0, $x_{Si,0}$ plus any ferrosilicon silicon, and T_0, with the iron, oxygen and time states at zero. Physics advances in 1 s steps of the engine's classical fourth-order Runge-Kutta method (Section 1.6.1).

Table 4.3 Heat Model Parameters and Their Status in the Interface, Where Visible Means Adjustable

Symbol	Meaning	Value	Class
F	oxygen flow	750 Nm^3/min	visible
η	oxygen efficiency, nominal	0.95	visible
C^*	critical carbon	0.35 %	visible
C_{eq}	equilibrium carbon	0.02 %	hidden
τ_{Si}	silicon oxidation time	120 s	visible
c_p	bath heat capacity	0.82 kJ/(kg K)	hidden
q_L	heat loss, nominal	42 MJ/s	visible
h_{ore}	ore cooling	4.0 MJ/kg	visible
h_s	scrap melting heat	1.30 MJ/kg	visible
h_C, h_{Si}, h_{Fe}	heats of oxidation	10, 30, 4.8 MJ/kg	fixed
C_{aim}, T_{aim}	aim carbon, temperature	0.05 %, 1650 °C	visible
—	windows	±0.01 %, ±10 °C	visible

4.3.4 Assumptions

The bath is one well-mixed volume, and the scrap and hot metal are mixed into it from the start for composition. Real converters carry gradients of carbon and temperature, and the sublance samples one point in them [5].

Decarburization has two regimes separated by a fixed critical carbon. In practice the critical carbon depends on the blowing and stirring practice, and droplets ejected into the slag emulsion contribute a separate reaction zone [4, 17].

Slag chemistry, phosphorus, manganese and the burning of CO to CO_2 enter only through the fixed oxygen efficiency, and the iron oxidized does not reduce the bath mass. Following the slag and metal chemistry in full takes an equilibrium or kinetic reaction model [18].

Heat loss is proportional to blowing time. The vessel does not cool during turndown or between heats, and the scrap melts at a fixed rate regardless of its size or of the bath temperature.

The off-gas analyzer returns the true decarburization rate multiplied by one flow error per heat, without delay, dilution by air ingress or composition noise.

4.4 Endpoint Strategies as Implemented

The controller runs every 1 s of simulated time, the same step as the physics, and holds no proportional, integral or derivative terms. There is no continuous feedback loop, so the velocity-form algorithm of Section 1.6.2, with its anti-windup, derivative filtering and rate limits, does not apply. The manipulated variables are decided at discrete events: before the blow, at the sublance, at the stop, at turndown and between heats. More oxygen lowers the carbon and raises the temperature, and more ore lowers the temperature without changing the carbon. In the convention of Section 1.3, a loop of carbon on oxygen or of temperature on ore would be direct acting and a loop of temperature on oxygen reverse acting, although no such loop runs here.

4.4.1 Static Charge Model

The charge model simulates the heat from the measured charge with nominal constants, $\eta = 0.95$ and loss factor $\hat{\phi}$, using an explicit Euler step of 2 s. Forty bisection steps on the final oxygen volume over $[V, V + 30000]$ Nm^3 find the volume at which the simulated carbon reaches C_{aim}, which gives the end temperature with no coolant, $\hat{T}^0_{end}$. The temperature target is raised by the margin ΔT_h in automatic, and the ore follows from Equation (4.7).

$$m_{ore} = \max\left(0,\ \frac{\hat{M} c_p \left(\hat{T}^0_{end} - T_{aim} - \Delta T_h\right)}{h_{ore}}\right) \tag{4.7}$$

where $\hat{M}$ is the measured bath mass. The ore is added at a constant rate over the first 180 s, and the oxygen is bisected again with that ore schedule. A heat that would finish below target even with no ore is given ferrosilicon with the charge instead. A secant search of up to six steps, started from $(T_{aim} + \Delta T_h - \hat{T}^0_{end})\hat{M}c_p/h_{Si}$, finds the silicon that lands on target within 0.05 °C. In manual mode the operator's oxygen volume and ore mass replace the plan on every heat and no ferrosilicon is used.

4.4.2 Sublance Re-plan

When V first reaches $0.85V_p$ the sublance returns C_s and T_s. The model re-plans from these readings and the silicon it simulates for the oxygen blown, by bisection with no further ore, obtaining a new V_p and an end temperature $\hat{T}_{end}$, and if $\hat{T}_{end}$ exceeds T_{aim} it adds $\hat{M}c_p(\hat{T}_{end} - T_{aim})/h_{ore}$ of ore over the next 60 s. The re-plan aims at T_{aim} itself, without the margin. At 85 % of the plan the bath carbon is near C^*, between 0.26 and 0.49 % in case 3, and the re-plan leaves 1970 to 2770 Nm^3 of oxygen to blow. Ore after the sublance can remove heat, but the only way to add heat is more oxygen, which would take the carbon below its aim. This is the one-way correction of Equation (1.5), in which the margin ΔT_h plays the part of m.

> **TRY IT IN THE APP** **3 · The Sublance at 85 % of the Blow**
>
> Load the scenario and press Run, then watch the Oxygen, planned and blown trend: when the sublance measures the bath at 85 % of the blow, the planned oxygen changes. After ten heats all ten carbons are in the window and seven temperatures are, and the three misses all finish cold, between 1612 and 1628 °C.

4.4.3 Off-Gas Integral and Rate Endpoint

The measured rate $r_{C,m} = r_C(1 + \varepsilon_g)$ is integrated from the assumed charge carbon, Equation (4.8).

$$\hat{C}_g(t) = 4.5 - \frac{100}{\hat{M}} \int_0^t r_{C,m}\,\mathrm{d}t \tag{4.8}$$

The off-gas strategy stops when $\hat{C}_g \le C_{aim}$, or at 115 % of the static plan if the estimate never gets there. The integral inherits both the flow error and the unmeasured charge carbon. A 3 % error on the 4.45 % of carbon removed is 0.13 %, and the charge carbon scatters by 0.08 %, both larger than the whole ±0.01 % window. An integrated estimate that only a fresh measurement can reset is also the problem of dead reckoning in Volume 5, Chapter 13.

The rate strategy re-plans at the sublance and then inverts Equation (4.2) below C^*, Equation (4.9).

$$\hat{C}_r = C_{eq} + (C^* - C_{eq})\,\frac{r_{C,m}}{0.95\,F\,k_C} \tag{4.9}$$

The estimate is formed only when $r_{C,m}$ is below 97 % of the nominal plateau rate, and the blow stops when $\hat{C}_r \le C_{aim}$ or at 110 % of the re-planned V_p. A relative error in the rate is now a relative error in $C - C_{eq}$, which is 0.03 % at the aim, so a 3 % meter error costs about 0.001 % of carbon. Equation (4.9) is the second form of Equation (1.4) with the nominal plateau rate in place of r_{max}. The same signal that fails as an integral succeeds as an instantaneous measurement. The static and sublance strategies stop when V reaches V_p.

> **TRY IT IN THE APP** **6 · The Sublance and the Decarburization Rate**
>
> After the sublance, watch the Carbon near the end trend: the blow stops when the carbon read from the rate reaches the aim, not when the planned oxygen is in. Over ten heats the case gives nine hits, all ten carbons in the window and no reblows, at 7.2 °C rms. Then load scenario 4, which integrates the same signal, and count its reblows.

4.4.4 Turndown, Reblow and Scoring

After the stop the vessel is turned down for 120 s and C_{end} and T_{end} are measured. A heat measured above $C_{aim} + 0.01\,\%$ is reblown in automatic, up to three times, with the oxygen re-planned from the turndown measurement and the model's silicon, and no further ore. A heat is a hit when its final measured carbon and temperature are both inside their windows, after any reblows. The scorecard reports hits, carbon and temperature hits separately, reblows, the root mean square (rms) end temperature error, the iron to slag per heat and the ferrosilicon silicon per heat.

4.4.5 Run-to-Run Learning of the Heat Loss

After each heat the controller re-simulates it from the measured charge with the oxygen and ore actually used and its current $\hat{\phi}$, obtaining the end temperature $\hat{T}_{end,k}$ it would have predicted. The factor is moved by a fraction λ of the correction that explains the measured error, Equation (4.10).

$$\hat{\phi}_{k+1} = \hat{\phi}_k - \lambda \, \frac{\hat{M} c_p \, (T_{end,k} - \hat{T}_{end,k})}{q_L \, t_{b,k}} \tag{4.10}$$

where $t_{b,k}$ is the blowing time of heat k, the result is clamped to the range 0.5 to 2, and the default weight is $\lambda = 0.5$. This is an EWMA update of one model constant [8], and every unmodeled source of temperature error, including the scatter in η and T_0, is attributed to the heat loss. For a factor that grows by δ per heat the EWMA estimate settles at a lag of δ/λ behind the heat it plans, as Equation (1.6) states, 0.06 for $\delta = 0.03$ and $\lambda = 0.5$, which is worth about 13 °C and is one reason the temperature margin is kept when learning is on. The same EWMA structure corrects each head of a vial filler from late checkweigher results in Volume 8, Chapter 9.

> ⧖ HISTORY **A Shop Model That Learned (1994 to 2006)**
>
> One converter shop ran its own static-dynamic model from 1994 to 2006. A charge calculation set the hot metal, scrap, lime and oxygen for a first blow to an intermediate stop near 0.22 % carbon, chosen because the critical carbon there was about 0.20 to 0.30 %. The dynamic part took the carbon to fall exponentially and the temperature to rise linearly with the oxygen blown. After each heat the errors of the heat and oxygen balances corrected the next heat, weighting the last three heats by 10, 30 and 60 %, the most recent most. Heats tapped without a reblow and within 15 °C of target rose from about 40 to 65 % of the total to about 85 % [19].

4.5 Results

The ten scenarios are listed in Table 4.4 with the teaching point of each and what completes it. Each case is credited when the learner performs the listed actions and the simulated campaign reaches the listed state, as Section 1.6.4 describes. Every value in Table 4.5 and in the text was measured on the simulation engine at the application's default seed, over ten heats unless stated, with the case selected and any setting changed before the first oxygen of the campaign. A control or disturbance changed at that point applies from the first heat, and the application gives the same campaign when the change is carried through Reset. Cases 5 and 8 therefore compare two margins or two learning weights on the same sequence of heats.

> ✦ DID YOU KNOW? **Two Hundredths of a Percent, Five Tonnes of Iron**
>
> From the simulator: in Table 4.5, lowering the carbon aim from 0.05 to 0.03 % on the rate endpoint still gives eight hits, but the iron burned into the slag rises from 6.85 to 11.64 t per heat. Near the aim most of the oxygen burns iron instead of carbon.

Figure 4.3 draws the hits, reblows and rms errors of Table 4.5 as a chart.

Table 4.4 Scenarios and Completion Requirements

#	Scenario, teaching point and kind	Required to complete
1	A fixed recipe (manual) *open loop on a varying batch* (demonstration)	Blow the first heat to turndown in manual, then finish the ten heats with the recipe held at 15 350 Nm^3 of oxygen and 1.9 t of ore.
2	The static charge model *feedforward from the charge* (experiment)	Blow one heat on the static plan in automatic and compare its oxygen with the fixed recipe, then finish ten heats and compare the carbon hits with the temperature hits.
3	The sublance at 85 % *one late measurement* (experiment)	Reach the sublance re-plan at 85 % of the planned oxygen, then finish ten heats and identify the heats that still miss.
4	Integrating the off-gas *an integral that drifts* (demonstration)	Blow one heat stopped on the off-gas integral, then finish ten heats and count the reblows and the heats blown past the aim.
5	Aim hot, trim cold *a one-way correction* (experiment)	Finish ten heats with the sublance and a 25 °C hot aim, then Reset, set the margin to 10 °C before the first oxygen and finish ten more.
6	Sublance and rate *a rate read as a level* (experiment)	Blow one heat to a stop on the decarburization rate after the sublance, then finish ten heats and compare hits, reblows and rms error with cases 4 and 5.
7	A lining that wears *drift across batches* (demonstration)	With the heat loss growing 3 % per heat, reach the sublance of heat 8 or later, then finish the ten heats on the rate endpoint.
8	Learning the heat loss *run-to-run control* (experiment)	Finish ten heats on the wearing lining with learning at $\lambda = 0.5$, then Reset, move λ by at least 0.1 (to 0.3 or 1.0) within the first 1300 s and finish ten more.
9	Aiming lower on carbon *the iron yield cost* (experiment)	With the carbon aim at 0.03 %, blow one heat to turndown, then finish ten heats and compare the iron to slag per heat with case 6.
10	Challenge, twelve heats *graded campaign* (challenge)	Finish twelve heats with at least eight hits in both windows, no reblows and an end temperature error of 15 °C rms or less, with the model parameters and disturbances as given.[a]

[a] With the shipped settings (static model, no margin, no learning) the campaign gives 2 hits, 4 reblows and 38.6 °C rms and fails all three criteria. With the suggested settings (sublance and rate, 30 °C margin, learning) it gives 8 hits, no reblows and 12.6 °C rms, a pass with no hit to spare and 2.4 °C of rms margin.

Table 4.5 Measured Campaigns. Hits Count Both Windows, C and T Count Each Window Alone, rms Is the End Temperature Error

Case and setting	Hits	C	T	Reblows	rms (°C)	Fe (t)
1, manual recipe	1	6	2	0	25.2	6.62
2, static	2	9	2	1	19.0	7.18
3, sublance	7	10	7	0	18.2	6.79
4, off-gas integral	3	8	3	9	65.0	8.79
5, sublance, $\Delta T_h = 10$	6	10	6	0	13.3	6.79
5, sublance, $\Delta T_h = 25$	7	10	7	0	7.2	6.79
6, rate, $\Delta T_h = 25$	9	10	9	0	7.2	6.85
7, as 6, $d = 3$	5	10	5	0	23.7	6.85
8, as 7, $\lambda = 0.3$	8	10	8	0	12.5	6.85
8, as 7, $\lambda = 0.5$	7	10	7	0	12.2	6.85
8, as 7, $\lambda = 1.0$	7	10	7	0	14.5	6.85
9, as 6, $C_{aim} = 0.03$	8	10	8	0	7.6	11.64

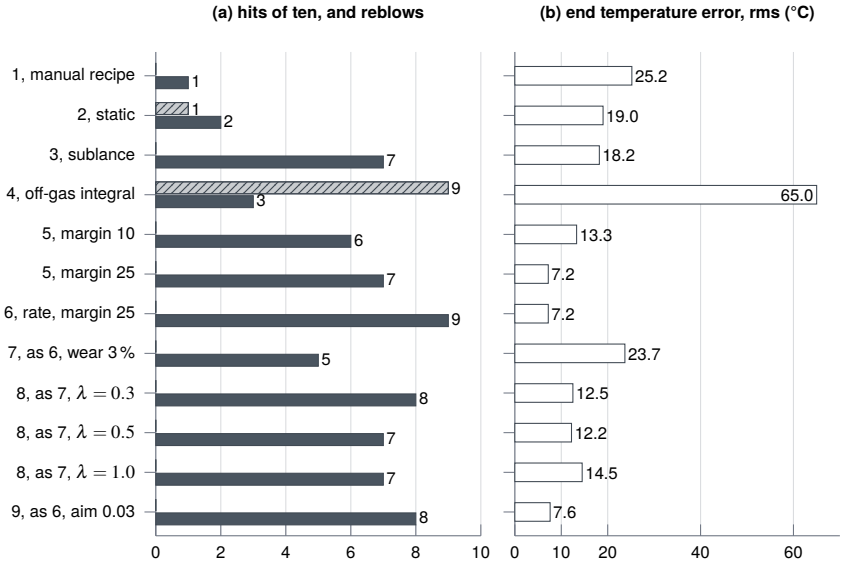


Figure 4.3 The measured campaigns of Table 4.5 as a chart: hits of ten solid and reblows hatched in (a), the rms end temperature error as open bars in (b). Notice that the off-gas integral (case 4) stands out with nine reblows and a 65.0 °C error, that the same signal read as a rate (case 6) gives the most hits, and that a wearing lining (case 7) undoes it until learning (case 8) recovers part of the loss.

4.5.1 Open Loop and the Static Model

A fixed recipe of 15 350 Nm^3 and 1.9 t of ore hits one heat in ten, with the end carbon spread from 0.036 to 0.081 % and the end temperature from 1615 to 1690 °C. The recipe is right for an average heat and no heat is average. The static charge model lands nine of ten carbons but only two temperatures, with an rms error of 19.0 °C. Carbon depends mostly on the oxygen blown, which the model controls. Temperature integrates every heat source and sink the model does not see, namely the efficiency, the heat loss and the errors in the charge measurements. With the charge variation and the meter error set to zero, so that every heat is nominal, the static model hits all ten heats at 3.4 °C rms (3.3 °C in the application's built-in test of the same condition, which draws other turndown errors), so the misses are due to what the model is not told rather than to the model structure.

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

4.5.2 One Late Measurement

The sublance re-plan puts all ten carbons in the window with no reblows and raises the temperature hits to seven. All three misses are cold, at 1612, 1615 and 1628 °C, while no heat finishes above 1655 °C. Ore after the sublance can take heat out, and nothing can put it back without burning carbon below its aim. Aiming the charge model 25 °C hot moves the static error to the side the sublance can correct, and the rms error falls from 18.2 to 7.2 °C. A 10 °C margin gives 13.3 °C. The hit count stays at seven with the larger margin because two heats still finish at 1637 °C, just outside the window, and the re-planned remainder of a third heat, trimmed with ore at the sublance, ends just above it.

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

4.5.3 An Integral That Drifts, a Rate That Does Not

Stopping on the integrated off-gas estimate gives three hits and nine reblows. Heats whose estimate falls too fast, because the meter reads high or the charge carbon exceeds the assumed 4.5 %, stop early and are reblown, and heats whose estimate never reaches the aim blow on to the 115 % limit.

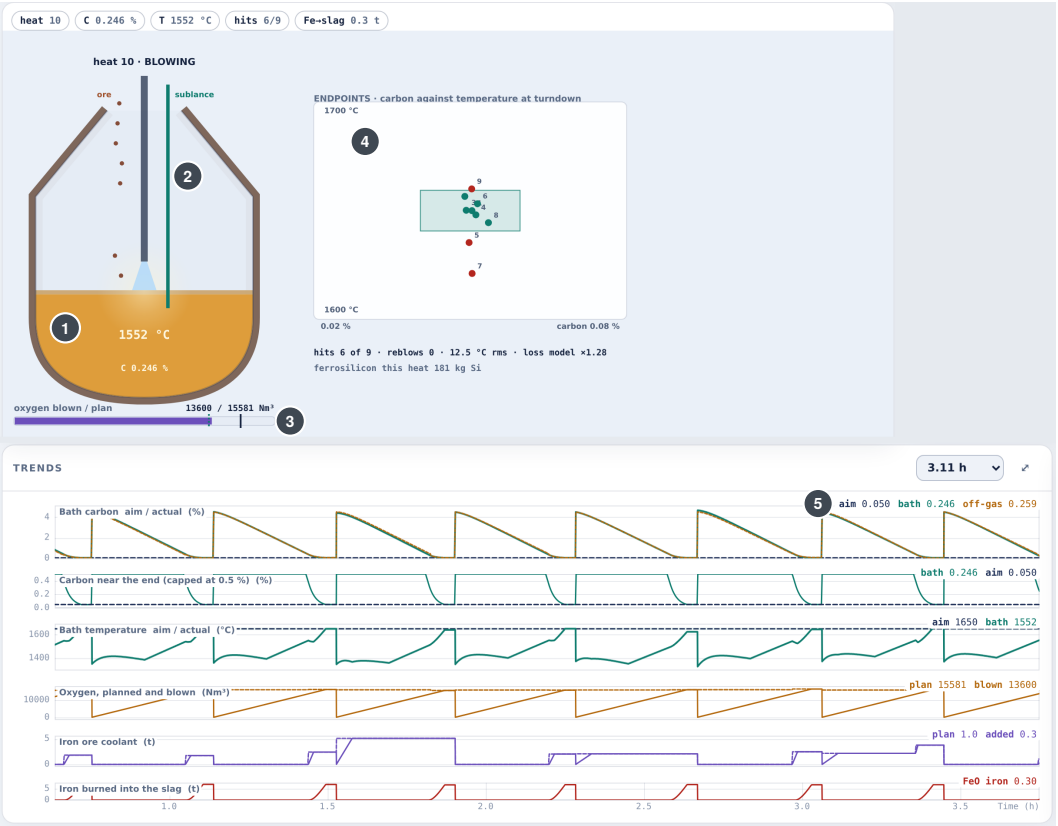


Figure 4.4 Scenario 8, learning the heat loss on a wearing lining with the sublance and rate endpoint and a 25 °C margin, at $t = 3.75$ h, cropped to the plant view and the trends. Heat 10 is blowing at 0.246 % carbon and 1552 °C (1), with 13 600 of a planned 15 581 Nm^3 blown (3), the sublance in the bath (2) and ore falling from the chute. The endpoint map (4) shows heats 1 to 9 against the target window, six inside it, heats 5 and 7 cold and heat 9 hot, with the model's loss factor at 1.28 and 181 kg of ferrosilicon silicon charged to this heat. The trends over 3.11 h show the carbon sawtooth of each heat with the off-gas estimate dashed (5), the bath temperature against its aim, the oxygen plan and volume blown, the ore planned and added, and the iron burned into the slag near the end of every blow.

Two heats end at 0.024 and 0.023 % carbon and at 1792 and 1794 °C, heated by the iron they burned, and the iron to slag averages 8.79 t per heat against 7.18 t for the static model. Setting the meter error to zero for the whole campaign does not rescue the integral, which then gives one hit with five reblows, because the scatter of the unmeasured charge carbon, 0.08 % (1σ), is by itself eight times the half-width of the window.

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

The rate endpoint uses the same signal after the sublance. With the 25 °C margin it hits nine heats, all ten carbons, with no reblows and 7.2 °C rms. Raising the meter error to 8 % still leaves all ten carbons in the window, with seven hits and 8.1 °C rms.

4.5.4 Drift and Learning

A lining whose heat loss grows 3 % per heat defeats the fixed margin. The first four heats still hit, the last four finish between 1601 and 1633 °C, and the campaign gives five hits at 23.7 °C rms. Carbon

Table 4.6 Challenge: Twelve Heats, Charge Variation 1.2, Drift 3 %/Heat

Configuration	Hits	Reblows	rms (°C)
as loaded, static	2	4	38.6
rate, $\Delta T_h = 30$, no learning	5	0	29.2
rate, $\Delta T_h = 50$, no learning	8	0	17.4
rate, $\Delta T_h = 30$, learning	8	0	12.6

is still right on every heat, because the sublance and the rate endpoint do not depend on the heat balance. With learning at $\lambda = 0.5$ the model factor climbs to 1.24 after the tenth heat, against a true value of 1.27 plus that heat's scatter, and the model charges ferrosilicon to the heats it now expects to be short, 122 kg of silicon per heat on average. The hits return to seven and the rms error to 12.2 °C over the campaign. During the tenth heat the model factor reads 1.28 with six hits in nine, as shown in Figure 4.4. A weight of 0.3 gives eight hits at 12.5 °C, and a weight of 1.0 gives seven at 14.5 °C, because each update then copies more of one heat's random error into the next plan [8].

4.5.5 The Iron Cost of a Low Carbon Aim

Lowering the aim from 0.05 to 0.03 % with the rate endpoint still gives eight hits, but the oxygen per heat rises from 15 643 to 16 695 Nm^3 and the iron to slag from 6.85 to 11.64 t. Both agree with the transport-limited regime of Equation (4.2). Taking $C - C_{eq}$ from 0.03 to 0.01 % at a time constant of 78 s takes $78 \ln 3 \approx 86$ s of blowing, about 1070 Nm^3, while the iron burns at 54 to 58 kg/s, about 4.8 t in total for 60 kg of carbon.

4.5.6 The Graded Challenge

The challenge runs twelve heats with the charge variation at 1.2 and a lining drift of 3 % per heat, starting from the static model with no margin and no learning. It is met with at least eight hits, no reblows and an rms error of 15 °C or less. The challenge is judged on the case as given, and a change to a model parameter, the charge variation, the meter error or the lining drift voids the attempt. Table 4.6 lists the measured outcomes.

As loaded the campaign gives two hits, four reblows and 38.6 °C rms, with one late heat at 1563 °C. The rate endpoint with a 30 °C margin fixes the carbon and the reblows but leaves five hits at 29.2 °C, because the drift eats the margin. A 50 °C margin buys eight hits at 17.4 °C, still outside the rms limit, because the late heats still finish cold, down to 1605 °C. Learning with the 30 °C margin meets all three targets, eight hits, no reblows and 12.6 °C, a pass with no hit to spare and 2.4 °C of rms margin, and charges 235 kg of ferrosilicon silicon per heat on average against 86 kg without learning. The rate endpoint, the margin and the learning each answer one of the campaign's three problems, the charge variation, the one-way correction and the drift.

4.6 Technology in Practice

Figure 4.5 draws the decisions of Section 4.4 with the hardware that measures and carries them out.

What Measures It

A shop measures the bath in the blow at a fixed share of the oxygen, by sublance, by a drop sensor or after turning the vessel down [19]. The temperature comes from a thermocouple, which is less accurate than a resistance thermometer but more rugged and with a wider span. It most often fails open, which reads as wild noise unless the instrument has a burnout mode that drives it high or low [20]. The model's sublance at 85 % of the planned oxygen, with a scatter of 3 % of the carbon and 4 °C, stands for this in-blow measurement, and the turndown reading 120 s after the stop for the test made by hand.

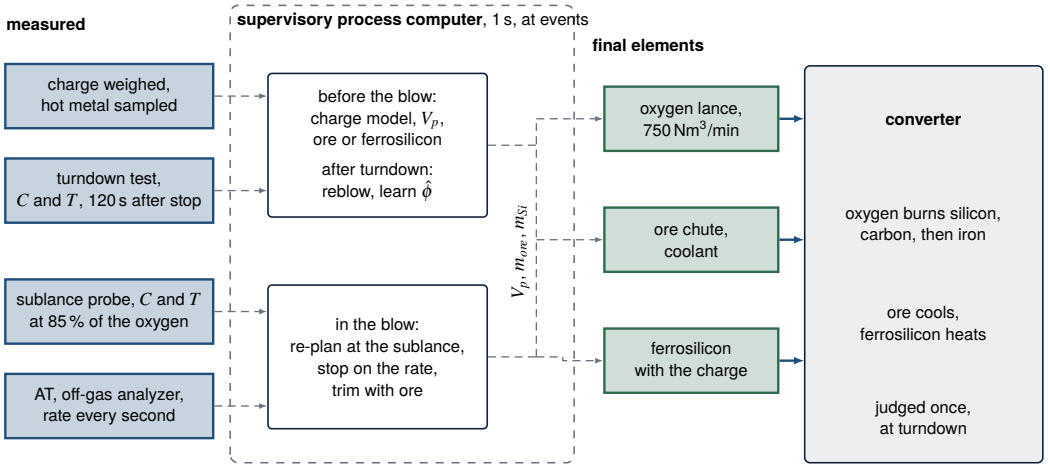


Figure 4.5 The decisions of the simulated converter drawn with their hardware. The charge model plans the oxygen volume V_p and the ore or ferrosilicon from the measured charge, the sublance and the off-gas analyzer AT correct and end the blow, and the turndown test decides a reblow and updates the heat-loss factor $\hat{\phi}$ for the next heat. Notice that every decision reaches the same three final elements, and that no continuous loop runs from a sensor to the lance.

> **TECHNOLOGY** **Sensor: The Off-Gas Analyzer**
>
> The gas leaving the converter carries the carbon removed from the bath as carbon monoxide, which makes up a steady 80 to 90 % of the gas until decarburization slows near the end of the blow. Plants commonly analyze it with infrared or mass spectrometric instruments, which need a clean sample at room temperature drawn from the duct far downstream of the furnace. The delay this causes, and the change in composition from continued reaction and air drawn into the gas, make the signal very difficult to use for dynamic control [21]. An infrared analyzer works because most gases of interest absorb infrared light at their own wavelengths [20].

The model's analyzer has neither the delay nor the air drawn in. It returns the true rate times one flow error per heat.

What Moves It

The lance delivers the one actuator that sets both endpoints.

TECHNOLOGY Final Element: The Oxygen Lance

The oxygen reaches the bath through a top-blown lance whose tip carries Laval nozzles, convergent and divergent passages that turn the pressure of the supply into the speed of a jet. Each nozzle has a design Mach number, pressure and flow rate and is usually run near them. Away from that point, shock or expansion waves form and the jet loses energy. Lances with several nozzles soften the blow to reduce spitting, dust and skull, but jets that merge before they reach the bath strike almost as hard as one, and lance design has been settled almost empirically [22].

The model blows a fixed 750 Nm^3/min with an efficiency η drawn for each heat, so the planned volume V_p is the decision and the blowing time follows from it. Iron ore is added through a chute at a constant rate over the first 180 s and over 60 s after the sublance, and ferrosilicon goes in with the charge.

Where the Controller Runs

Control loops in a plant are often nested, with cycle times from milliseconds to minutes, and a distributed control system uses a supervisory loop to mediate a group of localized controllers [23]. One shop's endpoint model ran its calculations for three adjacent heats at once: the heat in preparation, the heat in progress and the heat just finished [19]. The app executes its controller every 1 s, the step of the physics, but decides only at events. On a plant platform the charge model, the sublance re-plan, the stop rule and the learning are calculations of the supervisory process computer, which writes the oxygen volume and the additions to the oxygen flow loop and the feeders of the basic layer.

IN PRACTICE How a Shop Measures the Endpoint

Standard practice measures the temperature and carbon after about 90 % of the oxygen has been blown, and the temperature and dissolved oxygen at the end of the blow, taking a metal sample each time. The readings are taken by hand after turning the converter down, or with the vessel upright by a sublance or an in-blow drop sensor. In an industry association's 2020 data, 47 of 71 member shops, close to 70 %, used sublances or drop sensors, 37 used off-gas analysis, and only five relied on turndown tests alone [19].

What Is Modernizing It

The first development is optical sensing of the gas at the mouth. A mid-infrared tunable diode laser beam sent through the gas directly above the vessel was tested in four long-term field trials, and the infrared emission of the dusty gas correlated with the turndown carbon from below 0.03 to at least 0.15 % with a precision of ±0.007 %. The same emission signal predicted slopping of the foaming slag for 90 % of 100 heats, over 80 % with at least 30 s of warning [21]. The second is prediction from the records of past heats. A case-based model retrieves the stored heats most similar to a new one, learns as new heats are added and predicted the end temperature of a 120 t converter within 10 °C [24].

4.7 Limitations

The bath is treated as one perfectly mixed volume, and its critical carbon is fixed. Real decarburization runs in several zones whose share changes through the blow [5, 17].

The slag is a tally of oxidized iron, without phosphorus and manganese removal, lime dissolution, foaming or post-combustion [18].

The vessel loses heat only while blowing, so the waiting time between heats, during which a real vessel cools, is absent.

The off-gas analyzer has one flow error per heat and no transport delay or dilution by air drawn in at the hood [14].

The learning adapts one constant and blames every unmodeled heat error on the heat loss, which would mislead it if, for example, the oxygen efficiency drifted instead.

OPEN CHALLENGE **A Complete Model of the Converter**

The transport phenomena in a converter are complex and hard to observe directly, no complete intelligent model of the process had been published in open sources by 2022, and most models predict a single quantity such as the end temperature [24]. Rapid, accurate sensors for the hostile vessel are lacking, and gas analyzed downstream arrives late and altered [21]. A data-driven predictor for a 260 t converter put 90 % of carbons within 0.005 % and 94 % of temperatures within 10 °C, 84 % both, on one plant's records [25]. Progress would be a model of carbon, temperature and slag together that carries from one shop to another.

4.8 Conclusion

The converter combines in one batch the four elements of endpoint control, namely planning from the measured initial state, one late measurement, a correction that works in one direction and learning between batches. A fixed recipe hits one heat in ten, and the static charge model hits nine carbons but only two temperatures at 19.0 °C rms. The sublance at 85 % of the blow fixes every carbon and leaves only cold misses, and a 25 °C hot aim takes the rms error to 7.2 °C. The off-gas signal integrated from the charge gives nine reblows, while the same signal read as a rate below the critical carbon gives nine hits with none. A lining that loses 3 % more heat per heat brings that strategy down to five hits, and learning the heat loss brings it back to seven at 12.2 °C. Lowering the carbon aim by 0.02 % costs 4.8 t of iron per heat. A fuller model would add slag chemistry, the time between heats and lance height as a second actuator. The simulator is a compact case of controlling a batch that is blown once and judged at its end.

KEY POINTS

- A converter heat is blown once and judged at its end, and oxygen once blown cannot be taken back.
- The static charge model hits nine carbons but only two temperatures.
- The sublance fixes every carbon and leaves only cold misses, and a 25 °C hot aim takes the rms error to 7.2 °C.
- The off-gas signal integrated from the charge gives nine reblows, and read as a rate below the critical carbon it gives nine hits with none.
- Learning the heat loss brings a wearing lining back from five hits to seven.

Where else this applies. Volume 8, Chapter 9 applies the same EWMA correction from late results, and Volume 5, Chapter 13 resets an integrated estimate with late fixes.

References

[1] E. Del Castillo and A. M. Hurwitz, "Run-to-Run Process Control: Literature Review and Extensions," *Journal of Quality Technology*, vol. 29, no. 2, pp. 184–196, 1997.

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

[3] L. Holappa, "A General Vision for Reduction of Energy Consumption and CO2 Emissions from the Steel Industry," *Metals*, vol. 10, no. 9, p. 1117, 2020.

[4] J. H. Zong and J. K. Yoon, "Theoretical interpretation of the decarburization mechanism in convective oxygen steelmaking," *Metallurgical Transactions B*, vol. 21, no. 1, pp. 49–57, 1990.

[5] N. Dogan, G. A. Brooks, and M. A. Rhamdhani, "Comprehensive Model of Oxygen Steelmaking Part 1: Model Development and Validation," *ISIJ International*, vol. 51, no. 7, pp. 1086–1092, 2011.

[6] U.S. Geological Survey, "Mineral commodity summaries 2025: Iron and steel," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[7] J. Stubbles, "Energy use in the U.S. steel industry: A historical perspective and future opportunities," U.S. Department of Energy, Office of Industrial Technologies, Tech. Rep., 2000.

[8] A. Ingolfsson and E. Sachs, "Stability and Sensitivity of an EWMA Controller," *Journal of Quality Technology*, vol. 25, no. 4, pp. 271–287, 1993.

[9] F. Wallner and E. Fritz, “Fifty years of oxygen-converter steelmaking,” *Revue de Métallurgie*, vol. 99, no. 10, pp. 825–837, 2002.

[10] K. Katsura, K. Isobe, and T. Itaoka, “Computer control of the basic oxygen process,” *JOM*, vol. 16, no. 4, pp. 340–345, 1964.

[11] J. A. Glasgow, W. F. Porter, and J. Morrill, “Development and operation of BOF dynamic control,” *JOM*, vol. 19, no. 8, pp. 81–87, 1967.

[12] R. Williams, “Control of oxygen steelmaking,” in *Control and Analysis in Iron and Steelmaking*. Elsevier, 1983, pp. 147–176.

[13] D. Widlund, A. Medvedev, and R. Gyllenram, “Towards Model-Based Closed-Loop Control of the Basic Oxygen Steelmaking Process,” *IFAC Proceedings Volumes*, vol. 31, no. 23, pp. 69–74, 1998.

[14] A. Johansson, A. Medvedev, and D. Widlund, “Model-based estimation of decarburization rate and carbon content in the basic oxygen steelmaking process,” in *Conference Record of the 2000 IEEE Industry Applications Conference. Thirty-Fifth IAS Annual Meeting and World Conference on Industrial Applications of Electrical Energy (Cat. No.00CH37129)*, vol. 4, 2000, pp. 2578–2583.

[15] J. Terpak, M. Laciak, J. Kacur, M. Durdan, P. Flegner, and G. Trefa, “Endpoint Prediction of Basic Oxygen Furnace Steelmaking Based on Gradient of Relative Decarburization Rate,” in *2020 21th International Carpathian Control Conference (ICCC)*, 2020, pp. 1–4.

[16] R. Wang, I. Mohanty, A. Srivastava, T. K. Roy, P. Gupta, and K. Chattopadhyay, “Hybrid Method for Endpoint Prediction in a Basic Oxygen Furnace,” *Metals*, vol. 12, no. 5, p. 801, 2022.

[17] B. K. Rout, G. Brooks, M. Akbar Rhamdhani, Z. Li, F. N. H. Schrama, and A. Overbosch, “Dynamic Model of Basic Oxygen Steelmaking Process Based on Multizone Reaction Kinetics: Modeling of Decarburization,” *Metallurgical and Materials Transactions B*, vol. 49, no. 3, pp. 1022–1033, 2018.

[18] R. Sarkar, P. Gupta, S. Basu, and N. B. Ballal, “Dynamic Modeling of LD Converter Steelmaking: Reaction Modeling Using Gibbs’ Free Energy Minimization,” *Metallurgical and Materials Transactions B*, vol. 46, no. 2, pp. 961–976, 2015.

[19] T. D. Bradarić, Z. M. Slović, G. J. Stepanoski, and S. Kosanović, “Static-dynamic model for endpoint control used in Smederevo’s BOF shop,” *Metallurgical and Materials Engineering*, vol. 27, no. 4, pp. 449–456, 2021.

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

[21] S. Allendorf, D. Ottesen, and D. Hardesty, “AISI/DOE advanced process control program vol. 1 of 6: Optical sensors and controls for improved basic oxygen furnace operations,” American Iron and Steel Institute, Tech. Rep. DOE/ID/13205, 2002.

[22] K.-i. Naito, Y. Ogawa, T. Inomoto, S.-y. Kitamura, and M. Yano, “Characteristics of jets from top-blown lance in converter,” *ISIJ International*, vol. 40, no. 1, pp. 23–30, 2000.

[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] S.-y. Song, J. Li, and W. Yan, “Intelligent case-based hybrid model for process and endpoint prediction of converter via data mining technique,” *ISIJ International*, vol. 62, no. 8, pp. 1639–1648, 2022.

[25] M. Wang, C. Gao, X. Ai, B. Zhai, and S. Li, “Whale optimization end-point control model for 260 tons BOF steelmaking,” *ISIJ International*, vol. 62, no. 8, pp. 1684–1693, 2022.

4.9 Exercises

4.1 ○ Concept check · 2 min

In scenario 2, The static charge model, the model plans every heat from its measured charge. Over ten heats it lands nine carbons in the window but only two temperatures. Which explanation is right?

(a) The bisection searches the oxygen for the carbon aim, while the ore is set by a cruder formula

(b) The turndown temperature reading is less accurate than the carbon analysis, so hits are scored as misses

(c) Carbon follows the oxygen the model controls, while temperature sums heat terms the model is not told

(d) The temperature window is narrower relative to its aim, so the same model error misses it more often

4.2 ○ Concept check · 1 min

In scenario 3, The sublance at 85 % of the blow, every carbon lands in the window and seven temperatures do. The three misses finish at 1612, 1615 and 1628 °C. Why are all three cold?

(a) The sublance temperature reads high on average, so the re-plan removes heat the bath needed

(b) The vessel heat loss grows from heat to heat, and one reading per heat cannot see that drift

(c) The sublance measures before the scrap has melted, so its temperature misses the scrap's cooling

(d) Ore can take heat out after the sublance, but heat can be added only by blowing the carbon too low

4.3 ○ Concept check · 2 min

In scenario 4, Integrating the off-gas, the off-gas flow meter error is set to zero before Run, so the rate is measured exactly on every heat. What does the chapter report for the ten heats, and why?

(a) One hit and five reblows, because the integral starts from an assumed charge carbon that scatters widely

(b) Ten hits, because with an exact rate the integral equals the carbon removed and the estimate is exact

(c) Nine hits and no reblows, because the integral and the rate endpoint now read the same exact signal

(d) Three hits and nine reblows, as before, because the meter error was never a source of the misses

4.4 ▭ Theory and calculation · 5 min

A nominal heat has a bath of $M = 300\,\mathrm{t}$ with 0.40 % silicon. Oxygen is blown at $750\,\mathrm{Nm^3/min}$ with $\eta = 0.95$, $k_C = 24/22.4$ and $k_{Si} = 28/22.4$ in $\mathrm{kg/Nm^3}$, $\tau_{Si} = 120\,\mathrm{s}$, $C^* = 0.35\,\%$ and $C_{eq} = 0.02\,\%$.

(a) Find the plateau decarburization rate above C^* once the silicon is gone, in kg/s and in percent of the bath per minute. *Checkpoint:* 12.7 kg/s and 0.25 % per minute.

(b) Find the oxygen the silicon takes at the start of the blow. *Checkpoint:* $8.0\,\mathrm{Nm^3/s}$ of the $12.5\,\mathrm{Nm^3/s}$ supplied.

(c) Find the time constant with which the carbon decays below C^*. *Checkpoint:* About 78 s.

4.5 ▭ Theory and calculation · 8 min

Scenario 9 lowers the carbon aim from 0.05 to 0.03 %. Use the constants of the previous exercise, $k_{Fe} = 2(55.85)/22.4$ in $\mathrm{kg/Nm^3}$ and a time constant of 78 s below C^*.

(a) At each aim, find the share of the oxygen left after silicon that oxidizes iron and the iron burned per kilogram of carbon removed. *Checkpoint:* 0.909 and 47 kg at 0.05 %; 0.970 and 149 kg at 0.03 %.

(b) Find the blowing time and the oxygen needed to go from the higher aim to the lower one. *Checkpoint:* About 86 s and $1070\,\mathrm{Nm^3}$.

(c) Estimate the iron burned in that interval, within 0.2 t, and compare it with the measured rise in iron to slag between scenarios 6 and 9. *Checkpoint:* About 4.8 t, against 6.85 to 11.64 t measured.

4.6 Theory and calculation · 5 min

The bath has $\hat{M}c_p = 246\,\mathrm{MJ/K}$, ore cools by $h_{ore} = 4.0\,\mathrm{MJ/kg}$ and the vessel loses $q_L = 42\,\mathrm{MJ/s}$ times the loss factor while blowing.

(a) Find the change in end temperature worth one tonne of ore. *Checkpoint:* 16.3 °C.

(b) For a nominal heat the charge model plans 2.72 t of ore with no margin. Find the ore it plans with $\Delta T_h = 25\,^{\circ}\mathrm{C}$. *Checkpoint:* 1.18 t.

(c) The loss factor grows by 0.03 per heat and is learned with $\lambda = 0.5$. Find the steady lag of the estimate and the end temperature error it causes over a blow of 21 min. Repeat for $\lambda = 0.3$. *Checkpoint:* 0.06 and about 13 °C; 0.10 and about 22 °C.

4.7 App challenge · 8 min

Load scenario 2, The static charge model, change no setting, and run the ten heats at speed 300 with the default seed (the run stops itself at the end). Then load scenario 3, The sublance at 85 % of the blow, and do the same. Export the CSV file of the second run.

(a) For each run, read the Carbon hits, Temperature hits, Reblows and End T error, rms tiles of the scorecard. *Checkpoint:* Scenario 2: 9, 2, 1 and 19.0 °C. Scenario 3: 10, 7, 0 and 18.2 °C.

(b) In the file of scenario 3, read the column `end_T_C` in the rows where `metric_heats` steps up, and list the heats that miss the temperature window with their end temperatures, within 0.1 °C. *Checkpoint:* Heats 1, 5 and 7, at 1627.6, 1615.1 and 1611.5 °C.

Explain why the temperature hits rise from two to seven while the rms error barely moves.

4.8 App challenge · 9 min

Load scenario 5, Aim hot, trim cold, with Aim hot by at its shipped 25 °C. Use the default seed and speed 300, and let each campaign of ten heats run to its end.

(a) Run the scenario as shipped. Read Hit rate and End T error, rms, and find the end temperature of heat 5 in column `end_T_C` of the exported file. *Checkpoint:* 7 of 10, 7.2 °C, heat 5 at 1637.3 °C.

(b) Press Reset, set Aim hot by to 10 °C before Run, and read the same two tiles. *Checkpoint:* 6 of 10, 13.3 °C.

(c) Press Reset, set Aim hot by to 40 °C before Run, and read the same two tiles and the end temperature of heat 5. *Checkpoint:* 8 of 10, 6.4 °C, heat 5 at 1637.3 °C.

Read the end temperatures within 0.1 °C. Use Equation (1.5) to explain why heat 5 does not change between the last two runs.

4.9 App challenge · 8 min

Use the default seed and speed 300, and let each campaign of ten heats run to its end.

(a) Load scenario 7, A lining that wears, change no setting, and read Hit rate and End T error, rms. *Checkpoint:* 5 of 10, 23.7 °C.

(b) Load scenario 8, Learning the heat loss from heat to heat, change no setting, and read the same two tiles, the Ferrosilicon tile and the loss model factor on the learning card at the end. *Checkpoint:* 7 of 10, 12.2 °C, 122 kg of silicon per heat, $\times 1.24$.

(c) Press Reset, set Learning weight to 1.00 before Run, and read the same tiles and the factor. *Checkpoint:* 7 of 10, 14.5 °C, 178 kg, $\times 1.21$.

Explain why the weight of 1.00 ends with a lower factor than the weight of 0.50, although its lag behind a steady drift is smaller.

4.10 Application area · 4 min

The model's off-gas analyzer reports the decarburization rate without delay. Suppose a real analyzer reports it 10 s late, and the rate endpoint stops the blow when the delayed reading gives 0.05 % through Equation (4.9). Below C^* the carbon decays with the time constant of 78 s.

(a) Find the true carbon at the stop, and say whether it is inside the window. *Checkpoint:* 0.046 %, inside the window.

(b) Estimate the extra iron burned and the extra heating in the 10 s, and state the stopping reading that would compensate. *Checkpoint:* About 0.5 t and 10 °C; stop at a reading of 0.054 %.

4.11 Application area · 3 min

At the sublance of scenario 3 the bath carbon lies between 0.26 and 0.49 %, and the re-plan leaves 1970 to 2770 Nm^3 of oxygen to blow at 750 Nm^3/min.

(a) Find the blowing time left after the sublance. *Checkpoint:* 2.6 to 3.7 min.

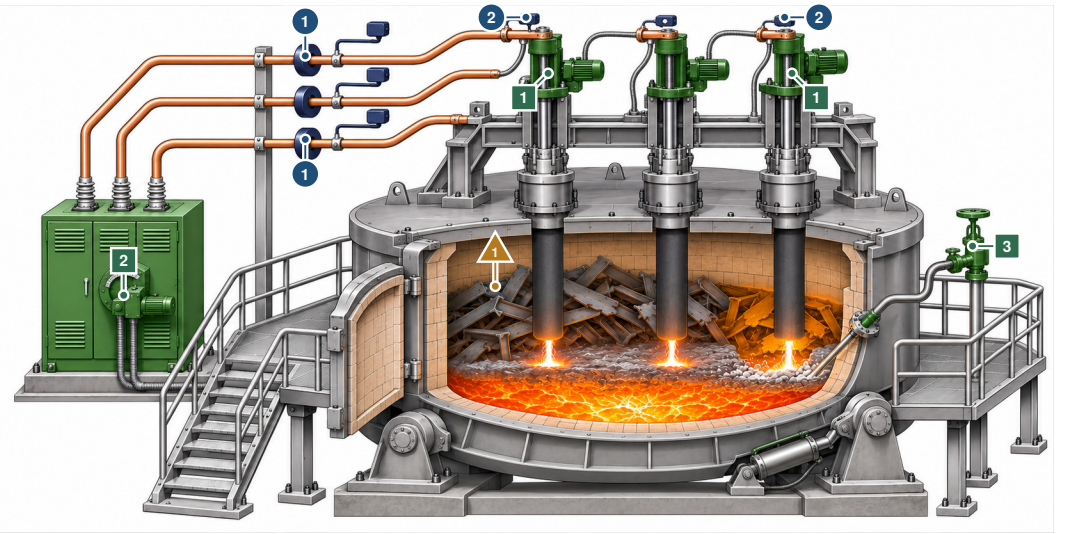
(b) Give one reason against taking the probe much later in the blow and one against taking it much earlier. *Checkpoint:* No numeric answer.

4.12 Review the agent · 5 min

An AI assistant is asked to pass scenario 10, the challenge, and replies as follows. "Use Sublance + rate for the carbon and raise Aim hot by to 50 °C. Leave learning off. The lining drift adds at most 33 % to the heat loss by heat 12, and a fixed margin covers it without the risk of a model fooled by noise. Expect at least eight hits, no reblows and an rms error well under 15 °C." Find the flaw in the proposal and say how it could be checked.

CHAPTER 5

Three Arcs on One Floating Neutral: Electrode Regulation in an Arc Furnace



The controller moves 1 the three electrode masts 2 the transformer tap changer 3 the carbon injection that foams the slag

The controller reads 1 the current in each phase 2 the voltage from each electrode arm

What disturbs it 1 scrap that caves in under an electrode

Notice that the three arcs all return their current through the one steel bath, so moving one electrode's mast changes the current read on all three cables. Illustration, not to scale.

IN PLAIN TERMS **Three Arcs That Share One Pool**

The main idea. An arc furnace melts scrap steel with three huge electric arcs, glowing columns of hot gas that carry current from carbon rods down to the metal. Each rod moves up and down to keep its arc at the right length.

The challenge. All three arcs send their current back through the same pool of melted steel, so moving one rod changes the current in all three. Each rod's controller must therefore watch a reading that the other two rods barely disturb.

Why it matters. That choice decides how much power a melt uses, how long it takes and how fast the furnace walls wear.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the fact box of Section 5.1, the history boxes of Section 5.2, the fact and try-it boxes of Section 5.6, 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 5.6, Section 5.8, Section 5.9 and the open challenge after it

The control question Three arcs return their currents through a steel bath that no wire ties to the transformer, so moving one electrode moves all three currents. Which variable should each phase's regulator hold, and which power program and slag practice bring a heat to tap within its energy, time and wear limits?

Learning objectives

- Explain from the floating neutral why lowering one electrode raises its neighbors' currents unequally.
- Compute the arc's extinction length and the circle-diagram maximum on a given tap.
- Evaluate the relative gain of current and impedance regulation from a circulant sensitivity matrix.
- Compare current and impedance regulators on the same step and through bore-in.
- Choose the regulator, power program and foamy slag that pass the graded heat, and say why one change alone does not.

Prerequisites Sections 1.6.2 (loop action and the velocity form), 1.4.7 (slow and fast parts joined) and 1.6.4 (scenarios and the scorecard).

The app https://apmonitor.com/apps/electric-arc-furnace/

See also Volume 6, Chapter 2 for loop pairing by the relative gain and the interaction it leaves. Volume 4, Chapter 2 for a protective override that acts only near its limit.

5.1 Introduction

At the top tap of the simulated 100 t furnace the transformer delivers 1000 V line to line, and three arcs 336 mm long each carry 60.1 kA and 26.6 MW. Lowering electrode 1 by 30 mm raises its own current to 63.6 kA, the current of phase 3 to 63.1 kA and that of phase 2 only to 60.8 kA. No electrode is wired to another, yet each moves the other two, because the three arc currents leave through the molten steel and the steel is not connected back to the transformer. A secondary circuit of this kind has a star point whose potential floats so that the three phasor currents add to zero, and an 80 MVA furnace modeled from plant data shows that point wandering inside the triangle of the phase voltages, reaching its edge when an arc goes out [1].

Arc furnaces that recycle scrap made about one third of the world's steel in 2011, at around 450 kWh/t of electricity [1], and a proposed fossil-free route reduces iron ore with hydrogen and makes steel from the product, with a variable share of scrap, for 3.48 MWh of electricity per tonne of liquid steel [2]. Every heat is melted under an electrode regulator, which moves three graphite columns on hydraulic masts to hold each arc between a short circuit and extinction while the scrap collapses under the tips. The problem has the shape of every decentralized multiloop system in which the loops share one node, such as parallel inverters on a common bus, zones of a heated slab or the tray loops of a column, and the relative gain of such a pairing measures how much one loop is changed by the others closing [3]. Volume 6, Chapter 2 pairs the loops of a reverse osmosis train by the same measure. Interaction between the phases of an arc furnace is documented as a source of unproductive electrode movement that industrial systems answer with dead zones [4], the floating-neutral circuit has been solved for proportional current control [5], and the current control of a three-phase submerged-arc furnace is a sibling problem [6].

> ✦ **DID YOU KNOW?** **Seven Tenths of American Steel**
>
> The United States made 81 million metric tons of raw steel in 2024, and 72 % of it came from electric arc furnaces, run by 49 companies at 104 minimills [7]. The converters of Chapter 4 made the rest. The simulated furnace of this chapter taps 100.45 t per heat.

The cases of this chapter are aimed at a belief that block diagrams invite, namely that three identical single-loop regulators on three identical phases act as three independent loops. In this model a current regulator tuned for phase 1 alone overshoots a 6 % step by 8 to 16 % and lies inside a quarter

of the step after 0.6 to 0.7 s, but with the other two phases also on automatic the same step overshoots 26 to 40 % and enters that band only after 2.0 to 2.6 s. The sign of the coupling is easy to guess wrong. The planning document from which the application was built expected a raised electrode to raise its neighbors' currents. A raised electrode lowers all three at the design point, and the second neighbor's sensitivity changes sign at shorter arcs.

The published treatments sit at either end of the problem. Furnace models are validated on plant data for operator training and optimization [1, 8], the arc's voltage characteristic is measured on pilot furnaces [9], and the control papers treat the electrode system without the heat around it [5, 10]. A student does not see, in one place, the circuit that couples the phases, a regulator choice that removes most of the coupling, and a heat whose energy, time and refractory depend on that choice. This chapter shows, from one heat of a simulated 100 t three-phase furnace in the browser, how current and impedance regulation, the power program and foamy slag decide the arcs' stability, the energy per tonne and the wear of the walls. The circuit, the arc law, the circle diagram and the relative gain are standard, and the analysis of Section 5.4, which places the change in sign of the neighbors' sensitivity where the neutral's displacement leaves a 30° cone, is elementary circuit algebra applied to the teaching model. The furnace's three currents add to zero to 10^{-6} of the short-circuit current at every step, a hand solve of the circuit agrees with the application to 0.1 %, and a measurement boundary keeps arc lengths, the bath between dips and the wear index away from the controllers. Nine cases run from one electrode moved by hand to a graded heat, each with a completion rule.

Section 5.2 follows the furnace and its regulators from ultra-high-power operation to hydrogen-reduced iron. Section 5.3 gives the model: the floating star, the arc and its extinction, the masts, the charge, bath, slag and wear, and the instruments. Section 5.4 derives the sign of the coupling, the relative gains of current and impedance regulation and the circle diagram that bounds the operating window. Section 5.5 describes the regulators, their overrides, the tap changer and the power program. Section 5.6 measures the nine cases, from one electrode moved by hand through the same step under both regulators, bore-in, melting, the flat bath, a tap step and a broken tip to the graded heat and a heat that no setting passes without foam. Section 5.7 separates code verification from agreement with published data, Section 5.8 describes the real instruments, final elements and controller platforms, Section 5.9 states the limitations, Section 5.10 concludes, and Section 5.11 holds the exercises.

5.2 Arc Furnaces and Their Electrode Regulators

How the electrode regulator grew from the push to ultra-high power into three loops on the arc resistance is told in the history box below. Two quantities can serve as the regulated variable, the arc current, which a current transformer gives directly, and the arc impedance or resistance, which needs the phase voltage measured against the furnace as well.

> ⧖ HISTORY **From Ultra-High Power to Three Loops (1968 to 2011)**
>
> The electrode regulator became the center of furnace automation when furnaces were pushed to the power their transformers could give. By 1968 an article in the Japanese electric-furnace literature was already titled as a trend toward ultra-high-power operation together with the automation of arc furnaces [11]. The analytical treatment followed within a decade. Billings and Nicholson modeled a three-phase furnace in 1977 to compare its control strategies [10], and a 1979 paper in *Automatica* combined the modeling and control of the electric arc furnace [12]. A model of an 80 MVA furnace published in 2011 controls the three electrode hydraulics through three univariable proportional integral derivative (PID) controllers, each on the relative error between a reference and the actual arc resistance [1]. That is the arrangement the simulator's impedance regulator follows.

Regulation then acquired a supervisory layer. A neural network electrode position system ran full time for five months in 1992, learning the relation between electrode position and the stability of furnace operation, and was reported to lower power consumption and furnace and electrode wear [13].

Table 5.1 Variables by Role, With Values at the Flat-Bath Design Point

Symbol	Role	Design	Range
$\lvert I_k \rvert$	CV of a current regulator, measured	60.1 kA	0–105
Z_k	CV of an impedance regulator, measured	7.48 mΩ	0.41–12
$v_{c,k}$	MV, mast speed command	0	±120 mm/s
V_{tap}	MV (discrete), tap voltage line to line	1000 V	600–1000 in 9
h_s^*	MV (on/off), slag foaming by injection	450 mm	80 or 450
—	DVs, cave-in, broken tip, tap steps	events	
λ_c	DV, cave-in frequency multiplier	1	0–3
z_k, v_k	states, tip height and mast speed	—	0.3–5.2 m
w_k, s_k	states, burn-back and surface under tip	0, h_b	—
m_s, m_b	states, solid charge and bath	0, 115.5 t	—
T_b, h_s	states, bath and slag height	1560 °C, 450 mm	—
L_k	truth, arc length	336 mm	$\leq L_{ext}$
V_N	truth, bath potential	0	—
W	truth, wear index	—	$\leq$ 0.30 kRWI·min
E, T_{dip}	measured, energy meter and dip	—	dip 10 s late

Dynamic models that describe melting, chemical reactions and energy additions were built so that the heat could be optimized as a whole [8], and heat-transfer models divided the arc's power among conduction to the steel, radiation to the surfaces and losses to the gas and the electrodes [14]. Foaming slag became the means of running long arcs in the flat bath. In a high-productivity furnace the foam height is usually limited by the amount of slag present, and it falls toward the end of the process as the carbon monoxide generation falls [15].

> ⧖ HISTORY **Minimills and the Race for Minutes (1970s to 2000)**
>
> The electric arc furnace was at first confined to specialty steel. In the 1970s, as capital costs rose and the profits of the integrated plants fell, entrepreneurs saw it, coupled with a billet caster and a bar or section mill, as a low-cost way to make steel. These minimills melted scrap, made long products such as reinforcing bar and broke into flat-rolled products in the 1990s. Key developments of the furnace included long arcs for electrical efficiency with foamy slags to protect the walls, computer controls and ever more powerful transformers. By 2000 many shops were racing to shave every possible minute from a heat, and the best had heat times of 40 to 45 minutes, with the power on for 30 to 35 of them [16].

The interaction between the phases has stayed a live control subject. A 2003 study solved the furnace circuit with a floating neutral and tried proportional current and power control on it [5], current control of a three-phase submerged-arc ferrosilicon furnace was reported in 2002 [6], and a 2023 study found that the mutual influence between phases causes unproductive electrode movements, built a matrix of mutual-influence coefficients that vary with the current increment, and proposed it as a replacement for the usual dead zone [4]. The newest change comes from the charge. The influence of direct-reduced iron on the furnace's energy balance has been quantified [17], and in a laboratory furnace an iron-oxide-rich slag from such pellets changed how the arc moved over the melt [18]. The simulator places these pieces in one heat and declares the iron it melts as the product of the hydrogen shaft furnace of Chapter 3.

5.3 The Furnace as Modeled

The simulated furnace has a 6 m shell, three electrodes of 0.6 m diameter on a 1.3 m pitch circle, nine transformer taps and one bucket of 85 t of scrap and 25 t of direct-reduced iron (DRI) charged onto a 15 t liquid heel at 1540 °C, as shown in Figure 5.1. The heat taps 100.45 t. The model is a declared teaching model assembled from public literature, not a port of any furnace supplier's or licensed simulation and not plant data. Table 5.1 classifies the signals.

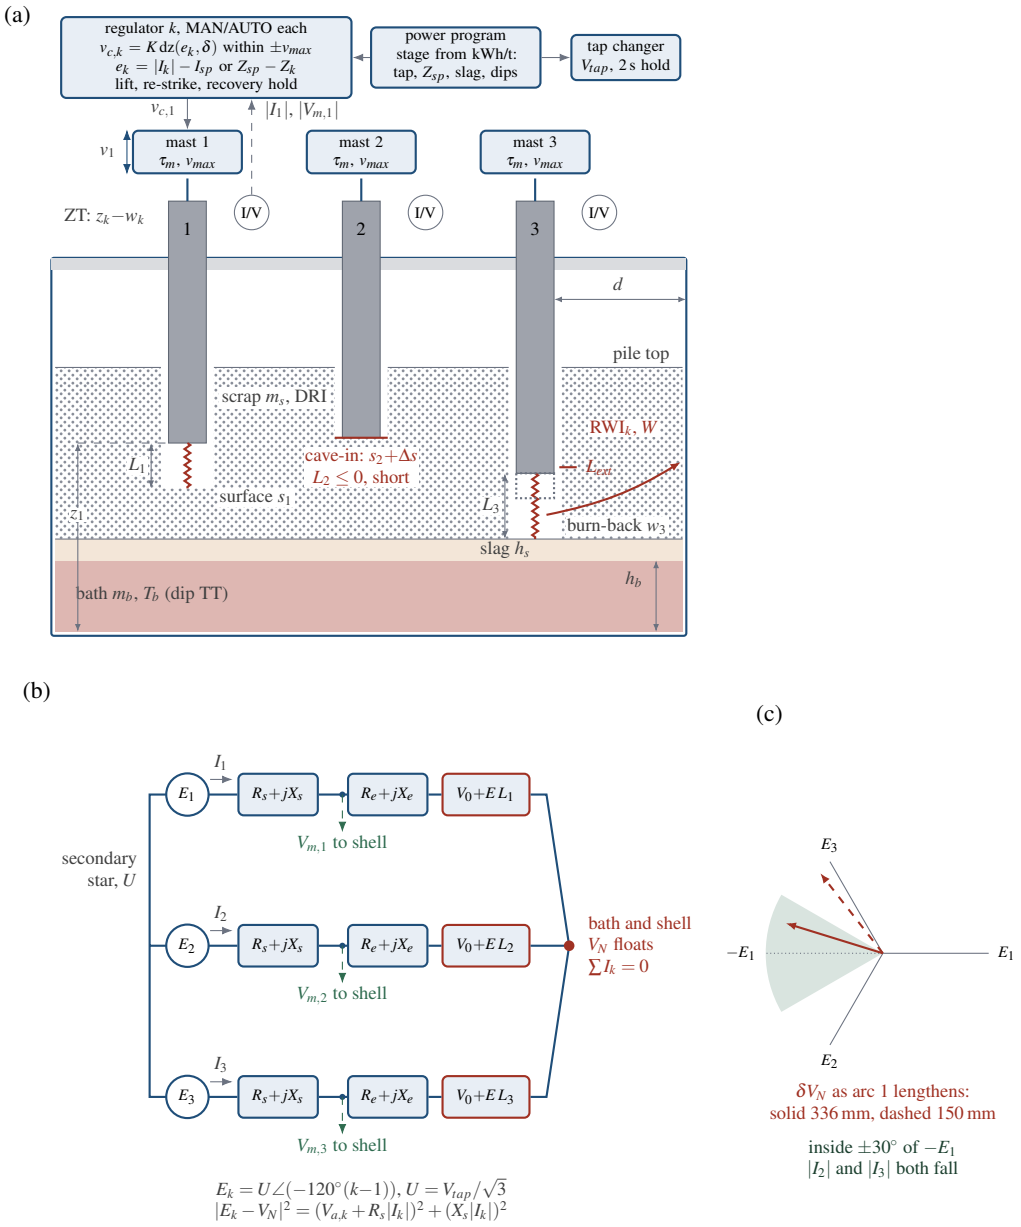


Figure 5.1 The furnace as modeled, after the application's Furnace view. (a) The section during bore-in. Masts of lag τ_m and rate limit v_{max} move tips at heights z_k (the encoders read $z_k - w_k$ before the burn-back w_k), each arc of length L_k ends on a surface s_k in its bore, a cave-in raises s_2 by Δs into a short, and arc 3 is drawn below its extinction length L_{ext}. The bath of mass m_b, level h_b and temperature T_b lies under a slag of height h_s, the bare part of an arc wears the wall at a distance d (RWI$_k$, integrated as W), and the regulator of each phase acts on its own measured current and voltage while the program sets taps and setpoints. (b) The secondary circuit, a star whose star point is the bath. Phase voltages E_k drive the series impedance $R_s + jX_s$, the column below the voltage tap $R_e + jX_e$ and an arc of voltage $V_0 + EL_k$, the regulator reads $V_{m,k}$ from the arm to the shell, and the bath potential V_N floats so that the currents add to zero. (c) The neutral's displacement when arc 1 lengthens, at the design length and at 150 mm, against the 30° cone about $-E_1$ inside which both neighbors' currents fall. Process model only. Not to scale.

5.3.1 A Star Whose Star Point Floats

The secondary is solved once per physics step as three root-mean-square phasors, a quasi-static description in which the arc's waveform within a cycle does not appear, as shown in Equation (5.1).

$$\begin{aligned} E_k &= U\angle\big(-120^\circ(k-1)\big), \qquad U = V_{tap}/\sqrt{3} \\ E_k - V_N &= (R_s + jX_s)\,I_k + V_{a,k}\,\frac{I_k}{|I_k|} \\ \sum_{k=1}^{3} I_k &= 0 \end{aligned} \tag{5.1}$$

where $V_{a,k}$ is the arc voltage of phase k, in phase with its current, and R_s and X_s lump the transformer, the flexible cables, the bus tubes and the electrode. The arc is a voltage set by its length rather than a resistance, so for a trial V_N each phase is explicit. Its magnitude obeys $(V_{a,k}+R_s|I_k|)^2+(X_s|I_k|)^2 = |E_k-V_N|^2$, a quadratic with one nonnegative root, and its angle follows from $I_k = (E_k-V_N)/(R_s + V_{a,k}/|I_k| + jX_s)$. A phase whose $|E_k-V_N|$ does not exceed $V_{a,k}$ carries nothing. The two real equations $\mathrm{Re},\mathrm{Im}\sum I_k(V_N) = 0$ are solved by a damped Newton iteration in the two components of V_N, started from the previous step's neutral, with a forward-difference Jacobian, a backtracking line search and a tolerance of 10^{-6} of the short-circuit current $U/|R_s+jX_s|$. Over whole heats it never needed more than four iterations, including steps with one phase out and another shorted. An extinguished phase carries no current, a short circuit ($L_k \le 0$) has $V_{a,k} = 0$, and with one or no phase able to conduct nothing flows.

The regulator measures each phase from its electrode arm to the furnace shell, which sits at the bath's potential, so the measured voltage includes the drop in the column below the measuring point, as shown in Equation (5.2).

$$V_{m,k} = V_{a,k}\frac{I_k}{|I_k|} + (R_e + jX_e)\,I_k, \qquad Z_k = \frac{|V_{m,k}|}{|I_k|} \tag{5.2}$$

A short circuit therefore reads $|R_e + jX_e| = 0.41\,\mathrm{m\Omega}$ rather than zero, which sets the lift's trip level.

5.3.2 The Arc, Its Extinction and Its Flicker

The arc voltage is linear in its length, with a seeded flicker on the length, as shown in Equation (5.3).

$$\begin{aligned} V_{a,k} &= V_0 + E\max(0, L_k + n_k), \qquad L_k = z_k - s_k \\ L_{ext} &= (k_{ext}U - V_0)/E \end{aligned} \tag{5.3}$$

where $V_0 = 40\,\mathrm{V}$ is the anode and cathode fall, $E = 1.2\,\mathrm{V/mm}$ the column gradient and n_k a colored noise of correlation time 0.1 s whose standard deviation is 14 mm while the arc burns on scrap and 3 mm on a flat bath, multiplied by $1 - 0.8\min(1, h_s/L_k)$ under foam. An arc longer than L_{ext} breaks and strikes again only when its tip comes within $L_{strike} = 25\,\mathrm{mm}$ of the charge. With $k_{ext} = 0.85$, L_{ext} is 212, 294 and 376 mm on taps 1, 5 and 9. The rule stands for an alternating arc that cannot reignite once the source voltage at current zero falls short of the arc's own, a teaching assumption without a measured source.

5.3.3 Masts, Surfaces and Electrode Consumption

Each mast is a hydraulic servo integrated exactly over a step for the speed command held through it, and the tip burns back in proportion to its current, as shown in Equation (5.4).

$$\begin{aligned} \tau_m\frac{\mathrm{d}v_k}{\mathrm{d}t} &= v_{c,k} - v_k, \qquad \frac{\mathrm{d}z_k}{\mathrm{d}t} = v_k + c_e|I_k| \\ \frac{\mathrm{d}w_k}{\mathrm{d}t} &= c_e|I_k|, \qquad P_k = V_{a,k}|I_k| \\ \frac{\mathrm{d}s_k}{\mathrm{d}t} &= -b\,P_k \quad \text{while } s_k > h_b \end{aligned} \tag{5.4}$$

where $\tau_m = 0.2\,\mathrm{s}$, $|v_{c,k}| \le v_{max} = 120\,\mathrm{mm/s}$, the stroke holds the tip between 0.3 and 5.2 m plus its burn-back, and the encoder reads $z_k - w_k$, the height a new electrode would have. The consumption c_e is declared from 1.5 kg of graphite per tonne of steel at 60 kA for 40 min, which burns 104.5 mm from each tip in a heat. An arc bores down through scrap at $b = 0.7\,\mathrm{mm/s}$ per megawatt of its power, and no surface stands above the pile top $h_b + m_s/(\rho A)$ with a bulk density $\rho = 1.2\,\mathrm{t/m^3}$ over $A = 28\,\mathrm{m^2}$, nor below the bath, whose level is $h_b = 0.145 + 0.007\,m_b$ in meters for m_b in tonnes.

A cave-in is a hard event in the post-step update. While more than 15 % of the charge is solid, an exponential clock with a mean of 30 s while any bore is open and 90 s once all are through picks an electrode at random and raises the surface under it by 100 to 300 mm, uniformly, never above the pile. A short arc is shorted by it, which is what the lift must clear. Arcs breaking and striking, shorts starting and ending, a broken tip, a tap step, a dip and the tap are also post-step events, each marked in the historian, so the derivative function stays pure.

5.3.4 Charge, Bath, Slag and Wear

The heat reaching the charge is the arcs' power weighted by how each arc is exposed, plus chemical heat, less the panels' losses, as shown in Equation (5.5).

$$\begin{aligned} Q &= \sum_k \eta_k P_k + P_{chem} - P_{loss} \\ \eta_k &= \eta_b + (\eta_{sc} - \eta_b)\,\chi_k \\ &\quad + (\eta_{sl} - \eta_b)\,c_k(1-\chi_k) \\ \frac{\mathrm{d}m_s}{\mathrm{d}t} &= -\frac{Q^+}{h_m}, \quad \frac{\mathrm{d}m_b}{\mathrm{d}t} = -Y\frac{\mathrm{d}m_s}{\mathrm{d}t} \quad (m_s > 0) \\ m_b c_p \frac{\mathrm{d}T_b}{\mathrm{d}t} &= Q \quad (m_s = 0), \qquad \tau_s \frac{\mathrm{d}h_s}{\mathrm{d}t} = h_s^* - h_s \end{aligned} \tag{5.5}$$

with the bath held at 1540 °C while $m_s > 0$, where $\chi_k = \mathrm{clip}_{[0,1]}\big((h_{top} - s_k)/L_k\big)$ is the share of the arc hidden by its bore's wall, $c_k = \min(1, h_s/L_k)$ the share under slag, η_{sc}, η_{sl} and η_b the efficiencies inside scrap, under foam and bare (0.90, 0.72 and 0.40), h_m the charge's melting energy (1.25 GJ/t for scrap and 1.42 GJ/t for DRI), Y its yield (0.92 and 0.89), $c_p = 0.82\,\mathrm{kJ/(kg\,K)}$ and $\tau_s = 60\,\mathrm{s}$. The efficiencies are a reduction by exposure of a heat-transfer model in which 15 to 20 % of the arc power reaches the steel by conduction and 75 to 85 % the furnace surfaces by radiation [14]. Burners and oxygen add $P_{chem} = 30\,\mathrm{MW}$ while scrap remains, up to a budget of 115 kWh/t of charge, so a slow heat cannot buy chemical energy by the minute, and $P_{loss} = 6 + 0.02(T_b - 1540)^+$ in megawatts. The slag foams toward $h_s^* = 450\,\mathrm{mm}$ when carbon is injected and less than 15 % of the charge is solid, and lies at 80 mm otherwise.

The wear of the walls is counted on the share of each arc that neither scrap nor slag covers, in the form of the refractory wear index named after Schwabe, as shown in Equation (5.6).

$$W = \int \max_k \Big(\frac{P_k V_{a,k}}{d^2}\,(1-\chi_k)(1-c_k) \Big)\,\mathrm{d}t \tag{5.6}$$

with P_k in megawatts, $d = 2.05\,\mathrm{m}$ from the electrode to the wall and W reported in kRWI·min. The energy meter integrates the arcs' power and the circuit's $R_s|I_k|^2$, which a high-current heat pays, and the energy per tonne is that meter divided by the tapped mass.

5.3.5 Instruments

Every controller works from one packet of measurements, which holds the three currents and the three arm-to-shell voltages, each with 0.3 % seeded noise and a 0.04 s transducer filter, $Z_k = V_k/I_k$ above 500 A, the metered power and energy, the three encoders, the tap indicator, the breaker, and the last thermocouple dip, which reads the true bath with $\sigma = 3\,^\circ\mathrm{C}$ and reports 10 s after it is taken. Arc lengths, the bath between dips, the scrap, the slag, the flicker and the wear are truth, logged in columns prefixed `truth_`, and the application's own tests fail if a controller file reads any of them.

Table 5.2 Model Parameters and Their Evidence. A Teaching Value Is Chosen for the Lesson, Not Measured on a Furnace

Symbol	Meaning	Value	Evidence
V_{tap}	taps, line to line	600–1000 V, 9	600–980 V in 18 [1]
R_s	circuit resistance	0.5 mΩ	0.37 mΩ + 0.12 mΩ [1]
X_s	operating reactance	5.5 mΩ	teaching value, cf. [19]
R_e, X_e	column below the voltage tap	0.1, 0.4 mΩ	teaching assumption
V_0	anode and cathode fall	40 V	common model value [9]
E	column gradient	1.2 V/mm	inside 0.5–3 V/mm [9]
k_{ext}	stability limit	0.85	teaching assumption
L_{strike}	re-strike gap	25 mm	teaching assumption
σ	flicker, scrap / flat	14 / 3 mm, 0.1 s	teaching value
m	scrap, DRI, heel	85, 25, 15 t	teaching value
f	supply frequency	60 Hz	declared (50 Hz in [1])
τ_m, v_{max}	mast lag, rate limit	0.2 s, 120 mm/s	teaching value (plan)
b	bore rate	0.7 mm/(s MW)	teaching value
—	cave-ins	30, 90 s, 100–300 mm	teaching value (plan)
c	electrode consumption	1.5 kg/t	teaching value (plan)
η_{sc}, η_{sl}, η_b	arc efficiency by exposure	0.90, 0.72, 0.40	reduction of [14]
P_{chem}, e_{chem}	chemical heat, budget	30 MW, 115 kWh/t	teaching value
P_{loss}	panels	$6 + 0.02\Delta T$ MW	teaching value
h_s^*, τ_s	foam height, time	450 mm, 60 s	teaching value
h_m	melting energy, scrap / DRI	1.25 / 1.42 GJ/t	teaching value (DRI table)
d	electrode to wall	2.05 m	geometry
W_{max}	wear budget	0.30 kRWI·min	teaching value (sweeps)

5.3.6 Parameters and the Declared Charge

Table 5.2 gives the parameters with the evidence for each. The circuit sits on a read source. The 80 MVA furnace of Logar et al. had a secondary of 600 to 980 V in 18 taps, rated secondary currents of 47 to 62 kA, arcs of 40 to 60 kA per phase, transformer and line resistances that add to 0.49 mΩ and a cathode drop of 30 to 40 V [1]. The arc law takes the 40 V that several furnace models use and a gradient inside the 500 to 3000 V/m that a non-linear model allows, while the pilot measurements themselves fit a non-linear relation [9]. The reactance is an operating value that includes the arc's own, a teaching value consistent with a circle-diagram analysis of a larger furnace whose operating reactance exceeds its short-circuit reactance [19]. The supply frequency is declared as 60 Hz, which enters only through the reactance and is 50 Hz in the read furnace.

The DRI is declared in the application's own file as the interface of the green steel line to a hydrogen shaft furnace. The table gives a metallization of 0.94, total iron 0.92, carbon 0, gangue 4.5 %, residual FeO 7.10 % from the metallization, a declared balance of 0.019 that closes the analysis, a charge temperature of 25 °C and a carbon demand of 12 kg/t to reduce the FeO. A review of the table found it plausible for hydrogen DRI, with a melting energy of 1.42 GJ/t against about 1.38 GJ/t by hand. The two apps are not coupled. Nothing is read from the shaft furnace, and a hot charge from a neighboring plant would be a mismatch for that app to report.

5.3.7 Initial State, Integration and Ledgers

A heat starts with a full bucket, the tips parked at 4.6 m over a pile whose top stands 3.5 m above the hearth, and every surface on that top. Other cases start lined out with the bores through and 60 % or 12 % of the charge left, or on a flat bath at 1560 °C, with balanced arcs of a stated length. The physics step is $\Delta t = 0.01$ s, the controller runs every 0.02 s and the historian logs at 5 Hz, the three clocks of Section 1.6.1. In place of the engine's Runge–Kutta step the application supplies its own update. Each step advances the flicker, solves the circuit at the step's start, holds the solution, integrates the masts exactly and the surfaces, charge, slag and ledgers explicitly, then applies the events. Six ledgers integrate the arc energy, the circuit losses, the net heat, the chemical heat, the panel losses and the wear.

5.3.8 Assumptions

Quasi-static phasors. The arc's waveform within a cycle, its harmonics and the flicker it puts on the supply network are absent, and the arc is a voltage in phase with its current. Intra-cycle arc models capture breakage and ignition that this description cannot [9].

A linear arc and a declared limit. The real voltage-length relation is non-linear and changes with the slag [9], and the stability limit k_{ext} and the re-strike gap are teaching values.

A declared charge. The pile, its bores and its cave-ins are a teaching model, not a melting simulation. The bath is held at 1540 °C while scrap remains, there is no chemistry, and the foam is a height that follows the injection, while the real foam height depends on the slag and its carbon monoxide supply [15].

One bucket, one heat. Charging, slag practice, sampling and the second bucket are out of scope, so tap to tap adds a declared 10 min to the power-on time.

5.4 How the Neutral Couples the Phases

5.4.1 The Sign of the Coupling

Raising electrode 1 lengthens arc 1, which lowers $|I_1|$ through its own equation and moves V_N. The neighbors' arc voltages do not change, and the magnitude relation under Equation (5.1) makes $|I_k|$ an increasing function of $|E_k - V_N|$ alone. To first order, therefore, a displacement δV_N changes phase k's current with the sign of

$$\delta|E_k - V_N| = -|\delta V_N| \cos\big(\arg \delta V_N - \arg E_k\big) \tag{5.7}$$

at a balanced point, where $V_N = 0$. Both neighbors' currents fall when δV_N lies within 90° of E_2 at −120° and of E_3 at +120° at once, which is the cone of ±30° about $-E_1$ drawn in Figure 5.1(c). A longer arc 1 drops more voltage, so its own current falls and the bath potential is pushed away from E_1. The push is not exactly opposite E_1, because the change in phase 1's current follows its own lag, 34.9° behind E_1 at the design point. At 336 mm the neutral moves 0.556 V per millimeter at 17.2° from $-E_1$ toward E_3, and $|E_k - V_N|$ changes by +0.531, −0.124 and −0.408 V/mm for phases 1, 2 and 3. All three currents fall, by 0.127, 0.035 and 0.114 kA/mm. Phase 3, which leads phase 1 by 120°, falls 0.90 as much as phase 1 and phase 2 only 0.27 as much. A star whose neutral were tied to the transformer would leave both neighbors unchanged. Over balanced arc lengths from 40 mm to extinction on taps 5 and 9 the push stays between 10 and 72° of $-E_1$, so the two neighbors never rise together, and the plan's expectation that a raised electrode raises the other two currents was a sign error, corrected during the build.

🕮 **Pause.** Try Exercise 5.1 before reading on.

5.4.2 Where the Second Neighbor Changes Sign

A shorter arc carries more current at a lower power factor, so its current lags further and the push rotates further from $-E_1$. At tap 9 the lag is 48.9 and 62.5° at 250 and 150 mm, the push sits 33.5 and 52.0° from $-E_1$, and phase 2's current rises by 6.6 and 31.6 A/mm when electrode 1 is raised, as shown in Figure 5.2(b). The sign of phase 2's sensitivity and the 30° crossing of the push were located independently on the application's solver. Both fall at 268.6 mm on tap 9 and 208.2 mm on tap 5, as Equation (5.7) requires. The program's stage points straddle the change. Melting at 317 mm and the flat bath at 336 mm on tap 9 lie on the side where every current falls, while the bore-in point, 200 mm on tap 5, lies on the other side, where phase 2 rises by 3.7 A/mm as electrode 1 goes up.

5.4.3 Interaction Measured as a Relative Gain

The sensitivity matrices of the three currents and the three measured impedances to the three tip heights are circulant at a balanced point. At the design point their diagonals are −0.1271 kA/mm and 0.0355 mΩ/mm, the current matrix's off-diagonal ratios are 0.90 and 0.27 and the impedance matrix's

0.39 and 0.12. The diagonal relative gain [3] of current regulation is 0.745 and that of impedance regulation 1.034. At the melting point the two are 0.639 and 1.007, and at the bore-in point 0.440 and 0.975. Divided by the diagonal, the current matrix has eigenvalues 2.17 and $0.68\angle\pm53^\circ$, and the impedance matrix 1.51 and $0.78\angle\pm18^\circ$. The complex pair belongs to the two modes in which the phases move against each other, and its 53° rotation is what turns three proportional current loops into a loop that rings. A phase's impedance is mostly its own arc's voltage over its own current, so the impedance matrix is nearly diagonal in the sense of the relative gain, although its off-diagonal ratio of 0.39 is not small.

5.4.4 The Circle Diagram and the Stable Window

For one balanced phase on one tap the arc power is $P(I) = I\big(\sqrt{U^2 - X_s^2 I^2} - R_s I\big)$. It is zero for an arc too long to carry current and for a short circuit at $U/|R_s + jX_s|$, and its maximum lies where the arc's resistance equals $|Z_s|$, as shown in Equation (5.8).

$$\begin{aligned} P^* &= \frac{U^2}{2\big(R_s + |Z_s|\big)} \le \frac{U^2}{2X_s} \\ I^* &= \frac{U}{\sqrt{2|Z_s|\big(R_s + |Z_s|\big)}}, \qquad Z_s = R_s + jX_s \end{aligned} \tag{5.8}$$

With $R_s = 0$ this is the circle of radius $U^2/(2X_s)$ with its maximum at a power factor of 0.707, the form used in a modified circle-diagram analysis of an AC furnace [19]. The per-phase maxima are 9.96, 17.71 and 27.67 MW at 42.47, 56.63 and 70.79 kA on taps 1, 5 and 9, as shown in Figure 5.2(a), and the circuit's resistance puts the top tap's maximum 8.7 % below the 30.30 MW of the lossless circle. The plan's bound $V_{tap}^2/(2X_s)$, 90.9 MW at the top tap, is the three-phase total of the same circle and is never exceeded. On tap 9 the arc is stable above 47.61 kA, where it reaches L_{ext}, so the program's setpoints of 60 and 65 kA sit on the long-arc side of the maximum, 40 and 58 mm short of extinction. The plan's 400 mm extinction length at the lowest tap would need an arc voltage of 520 V against a phase voltage of 346 V, so the limit was redefined as in Equation (5.3).

5.5 Electrode Regulation as Implemented

The controller runs every 0.02 s from the measurement packet. Its regulators are written in the application itself, not on the shared velocity-form PID of Section 1.6.2, because the regulated output is a mast speed and the mast already integrates it.

5.5.1 Current and Impedance Regulators

Each phase has its own regulator and its own MAN/AUTO switch, as shown in Equation (5.9).

$$\begin{aligned} v_{c,k} &= \mathrm{sat}_{v_{max}}\Big(K\,\mathrm{dz}(e_k,\delta) + \frac{K}{T_i}\int \mathrm{dz}(e_k,\delta)\,\mathrm{d}t\Big) \\ e_k &= |I_k| - I_{sp} \ \text{(kA, current)} \\ e_k &= Z_{sp} - Z_k \ \text{(m}\Omega\text{, impedance)} \\ \mathrm{dz}(e,\delta) &= e - \delta\,\mathrm{sgn}\,e \ \text{for } |e| > \delta, \ \text{else } 0 \end{aligned} \tag{5.9}$$

where K is K_I in mm/s per kA or K_Z in mm/s per mΩ, a positive e_k raises the electrode, the dead zone dz leaves errors inside $\pm\delta$ unanswered, the optional integral ($T_i = 0$ removes it) accumulates only while the output is not saturated, and it restarts from zero on a switch of regulator type, whose units differ, and on MAN. In MAN a mast drives toward its position station at 4/s times the gap between the station and the encoder reading $z_k - w_k$, clipped at the rate limit, and the station follows the encoder in AUTO, so transfers are bumpless. The header's MANUAL/AUTO throws all three phases. A raised electrode lowers its current and raises its impedance, so in the convention of Section 1.6.2

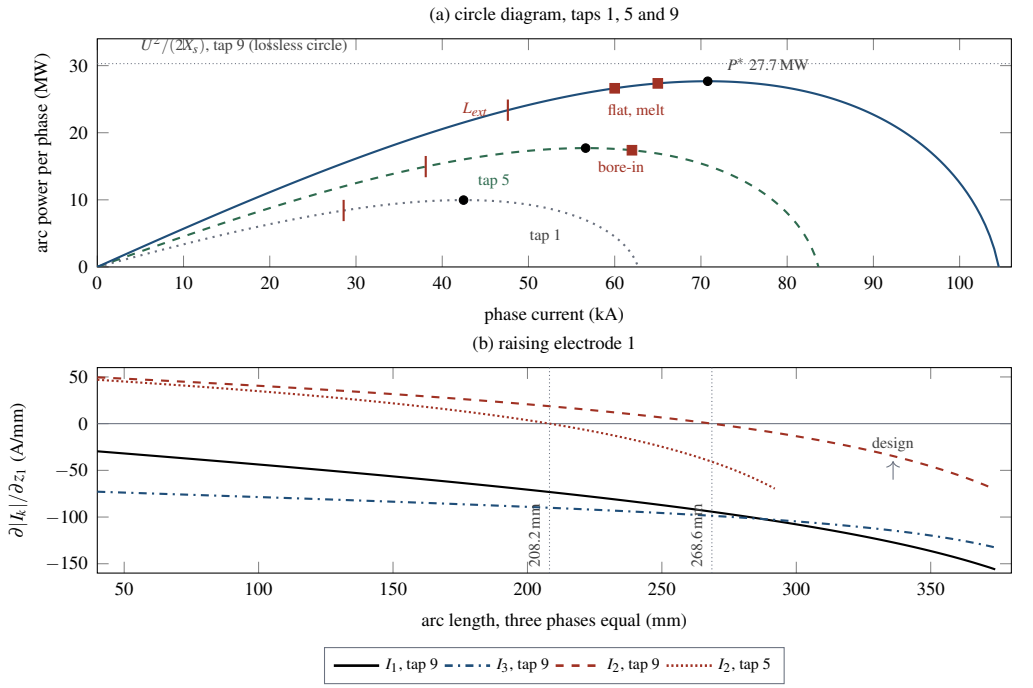


Figure 5.2 The operating window and the coupling, from the application's circuit functions. (a) Arc power per phase against current on taps 1 (dotted), 5 (dashed) and 9 (solid), with the maxima (dots), the stability limits at L_{ext} (bars) and the program's stage points (squares). (b) The change in each current per millimeter that electrode 1 is raised, with the three arcs of equal length: I_1 solid, I_3 dash-dotted, I_2 dashed on tap 9 and dotted on tap 5. Phase 2 changes sign at 268.6 mm on tap 9 and 208.2 mm on tap 5, where the neutral's push leaves the 30° cone.

the current regulator is direct acting, $a = -1$ with $e_k = a(I_{sp} - |I_k|)$, and the impedance regulator is reverse acting, $a = +1$. The mast integrates the speed command, so a steady burn-back r leaves an error of only r/K, 0.0011 mΩ for 0.044 mm/s at $K_Z = 40$.

5.5.2 Overrides Without Latches

Three overrides act only while their condition holds, like the protective override of Volume 4, Chapter 2, whose output rests on its lower clamp throughout healthy operation. The lift raises at v_{max} while a phase reads below $Z_{trip} = 1.0\,\text{m}\Omega$. The impedance rather than the current detects a short, because with a floating neutral a shorted phase's neighbor can carry more current than the shorted phase. On tap 4 with arc 1 shorted and the others at 200 mm the hand solve gives 69.3 kA in the short, 76.8 kA in the phase that leads it and 48.7 kA in the third. The re-strike lowers at v_{max} while a phase carries less than 1 kA. With the re-strike on, a recovery hold keeps a healthy phase still while another phase is out, shorted, or below half its own impedance setpoint, because two arcs in series across a line voltage read a high impedance that would drive them onto the charge just as the third strikes. A phase that is itself below half its setpoint never holds.

🕮 **Pause.** Try Exercise 5.3 before reading on.

That last rule corrects one of two controller defects that a review of the application found and fixed. The hold had latched, since a phase held whenever any other phase read low, including when it was itself low, so two phases recovering together held each other. On a flat bath with two tips broken at once all three arcs stayed at 61 to 75 mm, 100 kA and 35 to 38 MW for good. At the start of a heat the three fresh strikes held each other short, which had let the impedance regulator on the top tap

alone pass 3 of the build's six seeds, with a wear index of 0.18 to 0.42. With the latch removed it passes none of the six and none of 21 seeds, at 0.36 to 1.05 over the 21. The same latch had produced an 86 s short in a bore-in at $K_Z = 10$ with the lift off, now 2.04 s. The second defect was the tap changer's hold, which returned a zero speed before the lift was tested, so a cave-in during a tap step stayed shorted until the tap had stepped. On the challenge's seed this was a 2.06 s short. With the lift acting through the hold it lasts 0.53 s, and 1.00 s at most over 21 seeds.

5.5.3 Tap Changer and Power Program

The tap changer makes one step of 50 V per operation. The regulators hold the electrodes for 2 s, the voltage steps, and the next operation may start 3 s later. The program switches its stage on the metered energy per tonne of charge, and each stage names a current at a tap, from which the impedance setpoint is that of the balanced operating point on the declared circuit, so both regulators are asked for the same arcs. Bore-in runs tap 5 at 62 kA with 200 mm arcs to 45 kWh/t, melting tap 9 at 65 kA with 317 mm arcs to 300 kWh/t, and the flat bath tap 9 at 60 kA under foam, or tap 6 at 66 kA with 214 mm arcs if the slag is left flat. Dips start at 340 kWh/t and repeat every minute. Between dips the program carries the last reading forward by the metered energy at a declared net heating efficiency, 0.62 under foam and 0.32 bare, into the declared bath, and calls the tap when that estimate reaches 1630 °C once a dip has shown the bath above 1560 °C. A Fixed tap mode keeps the stages' setpoints on one tap.

5.5.4 Three Models, and Why No Optimizer

The simulation model is Section 5.3. The application fits no model to its own data, and the identified models a student builds are the sensitivity matrices of Section 5.4, measured by moving one electrode at a time. The controller's internal models are the declared circuit that turns a stage's current into an impedance setpoint and the program's declared bath mass, heat capacity and efficiencies. An operating-point optimizer, optional in the application's plan, was not built, since the wear term it would trade against power needs the slag height, which nothing measures, and on the declared values it would return the stage table.

5.6 Nine Cases on the Running Furnace

The cases appear in Table 5.3, and each carries the briefing and completion rule of Section 1.6.4. Each number in this section was read off a headless copy of the engine, advanced in the same 10 ms steps that the page uses, on the seed that the case declares, and a change that a sentence mentions is made after Run whenever the case's own text asks for it then. The review script that measures every case, rerun for this chapter, reproduced its earlier record line for line. Cases 2 and 3 are quoted as ranges over 31 step instants from 2 to 8 s after Run, because the arcs' flicker at the instant of the step moves the result, and Figures 5.3 and 5.4 quote one instant, 2.0 s. The first eight cases are summarized in Table 5.4.

> ✦ DID YOU KNOW? **A Wider Deadband That Moved No Less**
>
> From the simulator: in case 5 of Table 5.4, a 0.5 mΩ deadband cuts the masts' travel from 7.3 to 1.4 mm/s. Doubling it to 1.0 mΩ leaves the travel at 1.5 mm/s, and one arc wanders out to 376 mm and breaks.

5.6.1 One Electrode by Hand

In case 1 the currents read 60.0, 60.1 and 60.0 kA at 5 s, and with Phase 1 position lowered 30 mm they read 63.6, 60.8 and 63.1 kA at 10 s, with the bath potential 16 V off center. Raised 30 mm above its start, the same electrode takes all three down, to 56.1, 58.7 and 56.5 kA. The independent hand solve of the same states gives 63.80, 61.02 and 63.57 kA lowered and 56.16, 58.93 and 56.71 kA raised, without the noise and flicker of the running case.

Table 5.3 The Nine Cases

#	Start and settings	Teaching point
1	flat bath, 336 mm, all three on MAN	the coupling by hand
2	flat bath, current, $K_I = 20$, phase 1 offset +6 %	three loops ring
3	flat bath, impedance, $K_Z = 80$, offset −6 %	the same step settles
4	full bucket, stage taps, impedance $K_Z = 40$ (seed 4)	cave-ins, the lift
5	60 % of charge left, tap 9, 317 mm, deadband	extinction margin
6	13 t left, slag off, stage taps, Slag on	foam hides a long arc
7	mid-melt, tap 8, $Z_{sp} = 6.01$ mΩ, Tap up	a discrete step, held
8	flat bath, impedance, a 200 mm broken tip	burn-back, re-strike
9	full heat, current $K_I = 30$, Fixed tap 9 (seed 5)	graded heat

Table 5.4 Measured Results of Cases 1 to 8 on Their Own Seeds. Ranges in Cases 2 and 3 Span 31 Step Instants

#	Configuration	Measured
1	Phase 1 position −30 mm	60.0, 60.1, 60.0 → 63.6, 60.8, 63.1 kA, $\lvert V_N \rvert$ 16 V
2	$K_I = 20$, three on AUTO	overshoot 26–40 %, in at 2.0–2.6 s, I_3 swing 2.9–3.8 kA
	$K_I = 20$, phase 1 alone	8–16 %, in at 0.6–0.7 s
	$K_I = 30$, three on AUTO	41–222 %, in at 4.2–8.0 s, 0–2 arcs broken
3	$K_Z = 80$, three on AUTO	9–23 %, in at 0.5–0.7 s, I_3 swing 1.2–1.9 kA
	$K_Z = 80$, phase 1 alone	13–25 %, in at 0.5–0.7 s
4	impedance, $K_Z = 40$	10 cave-ins, 6 shorts (0.76 s), no arc broken, melting at 318 s
	current, $K_I = 30$	3 arcs broken, 7.7 s out, melting at 326 s
5	deadband 0, 0.5, 1.0 mΩ	masts 7.3, 1.4, 1.5 mm/s, longest arc 344, 360, 376 mm, 0, 0, 1 broken
6	slag off, 10 min	tap 6, 59 MW, wear 6.38
	Slag on at 2.5 min	TAP at 1629 °C after 390 s, wear 0.79
7	Tap up, tap 8 to 9	step 2.01 s after the button, 74.2 → 82.9 MW
8	200 mm tip broken	struck at 4.48 s, within 5 % 3.50 s later

5.6.2 Current Against Impedance on the Same Step

Case 2 asks phase 1 alone for 6 % more current, 3.6 kA, with each phase on a proportional current regulator at $K_I = 20$, the gain at which phase 1 by itself overshoots 8 to 16 % and is inside a quarter of the step at 0.6 to 0.7 s. With all three on automatic the step overshoots 26 to 40 %, settles inside a quarter of the step only after 2.0 to 2.6 s, and swings phase 3's current by 2.9 to 3.8 kA, as shown in Figure 5.3. The Furnace view two seconds after such a step is shown in Figure 5.4. At $K_I = 30$ the three loops go unstable at most instants, with overshoots of 41 to 222 %, phase 3 reaching 91.4 kA and up to two arcs broken. At $K_I = 10$ they ring less, 17 to 24 % with settling at 1.2 to 1.6 s. Case 3 makes the same request of impedance regulators at $K_Z = 80$, a 0.45 mΩ shorter arc. All three on automatic overshoot 9 to 23 % and settle at 0.5 to 0.7 s, the same as phase 1 alone (13 to 25 %, 0.5 to 0.7 s), and phase 3's current moves only 1.2 to 1.9 kA. Phase 3's current still changes when electrode 1 moves, but its impedance hardly does, so its regulator moves electrode 3 a little rather than passing the error round the phases.

🕮 **Pause.** Try Exercise 5.2 before reading on.

> **TRY IT IN THE APP** **2 · Current Regulation on All Three Phases**
>
> Load the scenario and press Run, then set Phase 1 offset to +6 %. Watch I3 on the Phase currents trend: phase 3's current swings by 2.9 to 3.8 kA although only phase 1 was asked to move, and the step overshoots 26 to 40 %. Then load scenario 3 and make the same request of the impedance regulators.

5.6.3 Bore-In, Melting and the Flat Bath

Case 4 bores a full bucket on tap 5. At the pause, when the program steps to melting, the impedance regulator had ridden 10 cave-ins and 6 short circuits, the longest 0.76 s, with no arc broken, and the program stepped up at 318 s. Current regulation at $K_I = 30$ on the same cave-ins broke 3 arcs, the

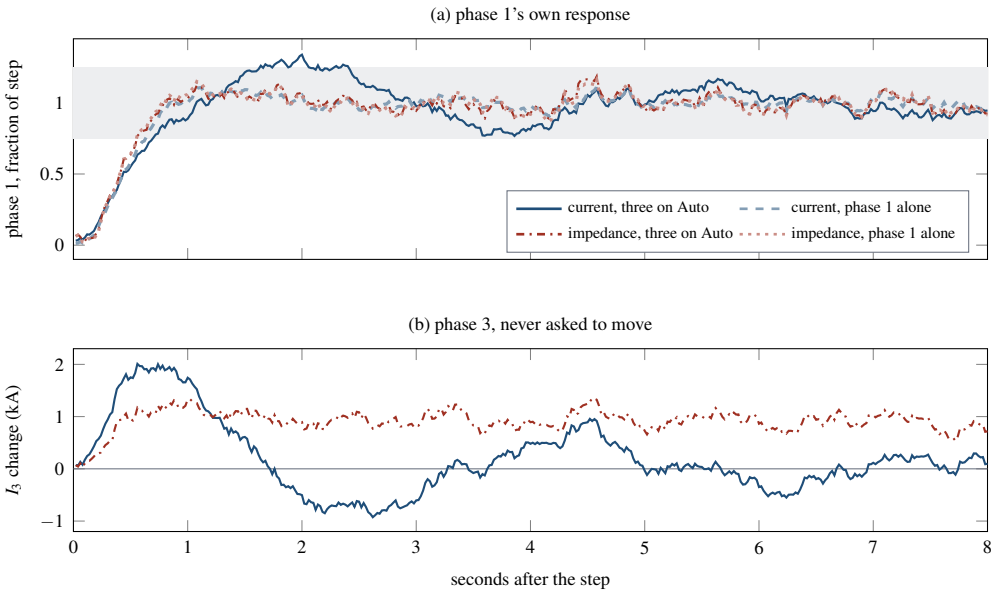


Figure 5.3 Phase 1 offset stepped 2.0 s after Run, measured signals from the running engine. (a) Phase 1's current (case 2, +6 %, $K_I = 20$) or impedance (case 3, −6 %, $K_Z = 80$) as a fraction of the step, with phases 2 and 3 on AUTO (current solid, impedance dash-dotted) or held on MAN from the start (current dashed, impedance dotted), and the band of a quarter of the step shaded. The current regulators overshoot 31.5 % together against 9.0 % alone and enter the band at 2.43 against 0.66 s. The impedance regulators overshoot 16.3 against 13.6 % and both enter at 0.60 s. (b) Phase 3's current with all three on AUTO, which swings 2.94 kA under current regulation (solid) and 1.41 kA under impedance regulation (dash-dotted).

longest out 2.6 s and 7.7 s in all, and paused at 326 s. Over six seeds the current regulator broke 2 to 5 arcs where the impedance regulator broke none. A short in one phase raises a neighbor's current by tens of kiloamperes, which a current regulator reads as its own error and answers by pulling its arc toward extinction. At $K_Z = 10$ the impedance regulator no longer reaches the mast rate on its own, and with the lift off as well the longest short grows from 0.98 to 2.04 s and three arcs break instead of one.

Case 5 melts on tap 9 at 65 kA, arcs of 317 mm against 376 mm to extinction, with 66.0 t of scrap left. With no deadband the masts travel 7.3 mm/s on average over 5 min and the longest arc reaches 344 mm. A 0.5 mΩ deadband, the size of the 14 mm flicker in impedance, cuts the travel to 1.4 mm/s while the longest arc rises only to 360 mm. At 1.0 mΩ the travel is no lower, 1.5 mm/s, the arcs wander to the band's edges and one reaches 376 mm and breaks.

Case 6 starts with 13.2 t of scrap in a 103 t bath and the slag flat, so the program runs tap 6 at 66 kA with 214 mm arcs and 59 MW. The wear index reads 0.15 at 2.5 min and 6.38 at 10 min, 21 times the challenge's budget. With Slag on at 2.5 min the program steps to tap 9 at 162 s, the bare long arcs wear the walls to 0.77 at 3.5 min while the foam is thin, and from 4 min, under 367 mm of foam rising to 443 mm, the index stops moving. The heat taps at 1629 °C after 390 s with the index at 0.79, and at 0.23 if the slag is started at 1 min.

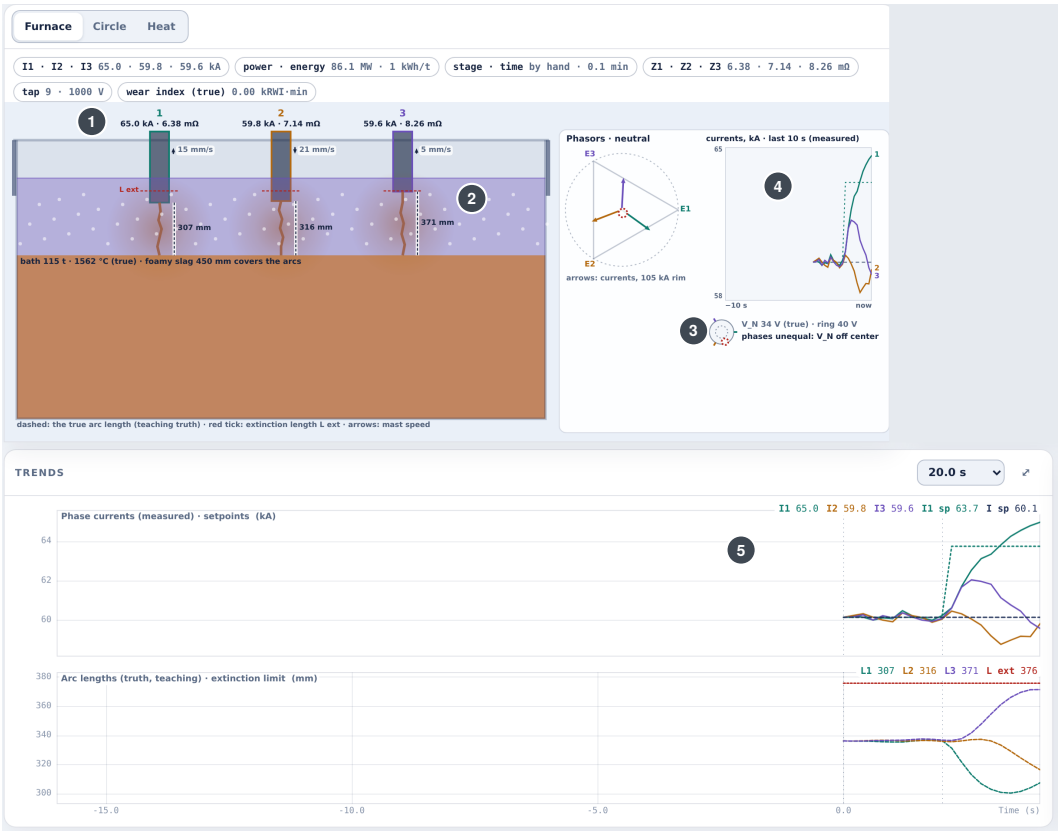


Figure 5.4 Case 2 on the Furnace view, 2.0 s after Phase 1 offset was set to +6 % with all three current regulators on AUTO ($t = 4.0$ s), cropped to the plant view and the trends. The badges (1) read 65.0, 59.8 and 59.6 kA and 6.38, 7.14 and 8.26 mΩ against a phase 1 setpoint of 63.7 kA and 60.1 kA for the others. Arc 1 has shortened to 307 mm, arc 3 (2) has drawn out to 371 mm, 5 mm from its extinction tick, the mast arrows show 15, 21 and 5 mm/s, and the bath potential sits 34 V off center on the magnified ring (3). The 10 s trace (4) and the trend (5) show I_1 past its dashed setpoint while I_3 overshoots and turns back.

> **TRY IT IN THE APP** **6 · The Flat Bath: Foamy Slag Hides a Long Arc**
>
> Load the scenario, press Run and watch the Wear index (truth) tile. With the slag flat it reads 0.15 at 2.5 min. Press Slag on then: the program steps to tap 9, the index climbs to 0.77 while the foam is thin and stops moving from about 4 min, and the heat taps at 1629 °C with the index at 0.79.

5.6.4 A Tap Step and a Broken Tip

Case 7 runs tap 8 at a fixed 6.01 mΩ, with currents of 62.7, 63.5 and 62.3 kA, arcs of 283 to 295 mm and 74.2 MW. Tap up steps the voltage 2.01 s after the button with the electrodes held. The currents jump to about 73 kA, and the regulators raise the electrodes until the impedance reads 6.01 mΩ again. Twenty seconds after the button the currents are 67.6, 67.5 and 65.9 kA, the arcs 303 to 309 mm and the power 82.9 MW. The same impedance on a higher tap is a longer arc at more current. In case 8 the tips burn 5.2 mm in 120 s and the encoders go down 5.2 mm while arc 1 stays at 336 mm. A 200 mm broken tip lengthens arc 1 to 536 mm, past extinction, the re-strike brings it down to strike 4.48 s later while the other two hold, and 3.50 s after the strike all three are within 5 % of 7.48 mΩ.

Table 5.5 The Graded Heat. Seed 5 in Detail, Then Passes Over the Build's Six Seeds and Over 21 Seeds

Configuration	T2T (min)	kWh/t	Wear	Out (s)	6	21
shipped, current $K_I = 30$, tap 9	40.3	411	0.64	195	0	0
impedance $K_Z = 40$, tap 9	38.6	412	0.63	0	0	0
current $K_I = 30$, stage taps	41.2	413	0.303	85	0	1
impedance, tap 7	43.6	405	0.40	0	4	12
current $K_I = 10$, stage taps	40.9	413	0.25	54	5	16
suggested, impedance, stage taps	40.5	414	0.23	0	6	19

Every configuration taps at 1630 °C on seed 5. T2T is tap to tap, Out the time with an arc out.

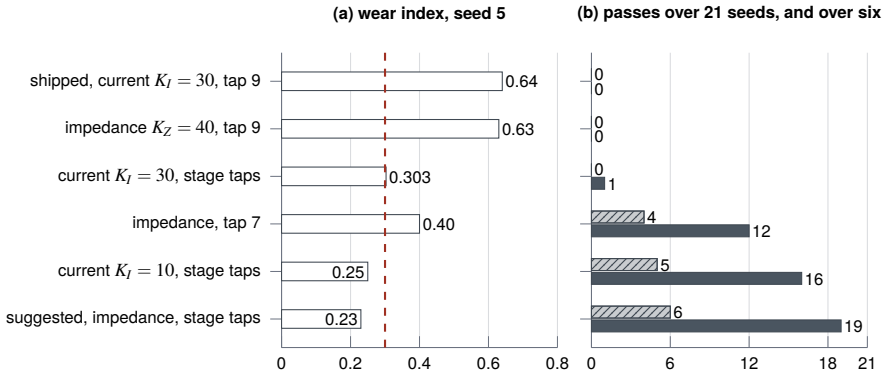


Figure 5.5 The graded heat of Table 5.5 as a chart: the wear index on seed 5 as open bars against the limit of 0.30 (dashed) in (a), and the passes over 21 seeds (solid) and over the build's six (hatched) in (b). Notice that either single change from the shipped settings, the impedance regulator or the stage taps, still misses the wear limit on seed 5 and passes at most once, while the two changes together pass 19 of 21.

5.6.5 The Graded Heat

The challenge runs one heat from a full bucket to the tap and asks for a tap at 1620 °C or hotter on the true bath, a tap-to-tap time of 48 min or less (power-on time plus the declared 10 min), at most 420 kWh/t on the meter, a wear index of at most 0.30, no arc out for more than 5 s at a time and at most 30 s of arcs out over the heat. The last criterion was added because a re-strike takes about 3.3 s, so the plan's 5 s per episode never binds, while a heat that breaks an arc every minute is not acceptable practice. Every limit is a teaching value. Real tap temperatures come from the ladle route and the caster, energy targets from a shop's benchmarks, wear limits from its refractory practice and time with arcs out from its supply contract. Table 5.5 gives the configurations on the case's seed and over 21 seeds, and Figure 5.5 draws their wear and passes.

As shipped the heat taps hot enough, after 30.3 min of power and at 411 kWh/t, but breaks 60 arcs, leaves arcs out for 195 s and wears the walls to 0.64. Each single change fails on this heat, the impedance regulator on the top tap on wear (0.63, its long arcs striking on top of a fresh pile) and the stage taps with the shipped current regulator on arcs out (85 s, 28 arcs broken). Together they tap at 1630 °C after 30.5 min of power, at 414 kWh/t with a wear index of 0.23 and no arc out. In that heat the program steps to melting at 5.29 min and to the flat bath at 24.42 min, takes three dips and taps at 30.47 min, through 43 cave-ins and 4 shorts, the longest 0.53 s. The claim is not universal. Over 21 seeds the suggested settings pass 19 times. Both misses are on wear, one from two cave-ins a second apart that left an arc burning above the pile at the top tap for half a minute, the other from wear in the first two minutes. The stage taps alone pass once and the shipped settings and the impedance regulator alone never. A current regulator detuned to $K_I = 10$ passes 16 of 21 with the stage taps,

Table 5.6 Code Verification and Agreement With Published Data

Check	Result
Code verification	
$\lvert\sum I_k\rvert$, suggested heat	$\le 1.00 \times 10^{-6}$ of I_{sc} on 182 843 steps
independent hand solve, Python	design point, case 1, short, open phase within 0.1 %
Newton against the plan's fixed point	agree to 0.1 A
circle maximum by sweep	Eq. (5.8) to 6×10^{-7} at every tap
melting and flat-bath energy	closed to 10^{-9}
whole heat, first law	0.036 % gap (6 MW before strike)
energy meter against trapezoid	0.006 %
halving Δt	suggested -0.001, shipped $+0.089$ %, verdicts kept on 8 seeds
load path (first, list, Reset)	9 of 9 cases identical
Agreement with published values	
currents at the program's points	60–66 kA at about 104 MVA, against arcs of 40–60 kA and a rated 47–62 kA at 80 MVA [1]
arc voltage at 100, 200, 336 mm	160, 280, 443 V, on the 40 V intercept and inside 0.5–3 V/mm [9]
energy per tonne	414 kWh/t against about 450 typical [1]
circle diagram, $R_s = 0$	radius $U^2/(2X_s)$, pf 0.707 [19]

while the impedance regulator with the stage taps passes all six of the build's seeds at gains of 20, 40, 80 and 120. The current regulator has to be detuned to live with the interaction, and the impedance regulator does not.

5.6.6 An Honestly Infeasible Heat

Without foamy slag no setting passes. Over 333 settings on three seeds, every tap program, both regulators and hand-set taps and arcs for the whole heat or from the flat bath on, and 286 more around the closest, none meets the challenge even when tapped by an oracle the instant the true bath reaches 1620 °C. An arc short enough to hide in the 80 mm slag runs at high current toward the short-circuit end of the circle, where the metered loss grows with the square of the current, and an arc long enough to carry the power wears the walls. The closest, 80 mm arcs on tap 5 from the flat bath on, misses 420 kWh/t by 0.36 kWh/t, 0.09 %, unchanged at half the time step. A margin that thin is stated beside the result, and a heat tapped on dips would miss by more. The plan's own infeasible variant, a current regulator at the shipped gain unable to ride the cave-ins, holds on the challenge's seed (none of 128 settings at $K_I = 30$ passes) but not on others (8 of 128 on one seed, 1 on another), so it is kept as a comparison and not offered as the proof.

5.7 Code Verification and Agreement With Published Data

Verification and agreement are kept apart in Table 5.6. The first block checks that the code solves its own equations. The second compares the teaching model's numbers with published ones, which supports plausibility and not validation, since no furnace was fitted.

The first-law gap of the whole heat is a declared consequence of the model, not an integration error. During the first 9 s, before the arcs strike, the panels' 6 MW is booked in the net-heat ledger while the bath, held at 1540 °C while scrap remains, does not pay it. In the suggested heat the arcs deliver 38.55 MWh and the circuit loses 3.06 MWh, a meter reading of 41.61 MWh or 414.2 kWh/t on 100.45 t, while the chemical heat spends 101.7 kWh of its 115 kWh per tonne of charge, at a mean arc power of 75.9 MW and an arc efficiency to the charge of 0.872. A heat costs 2.7 µs per physics step in Node, 0.50 s for the suggested heat, and in the browser the heat runs at the full 200× speed, 8.7 to 10.8 s of wall time in two separate sets of measurements.

5.8 Technology in Practice

Figure 5.6 draws the regulators of Section 5.5 with the hardware that measures and moves them, and the power program above them.

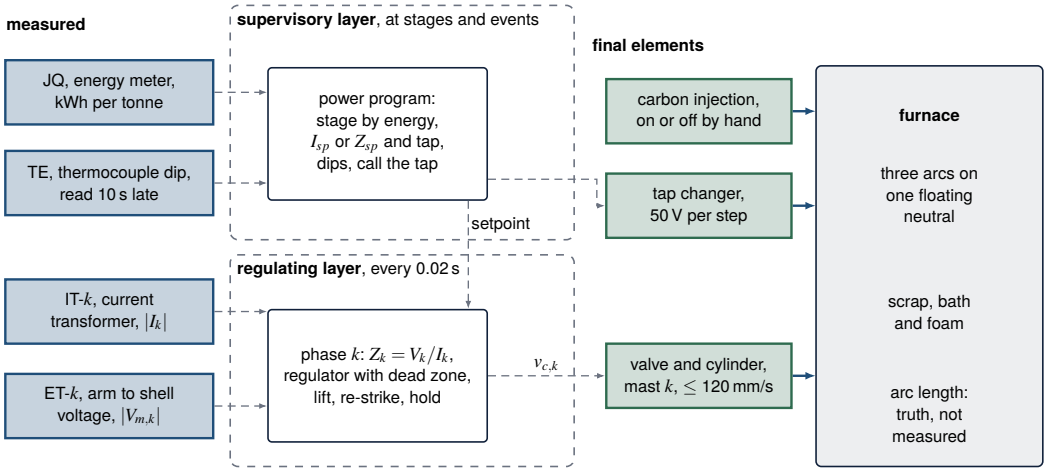


Figure 5.6 The control of the simulated furnace drawn with its hardware, for one phase k of three. The current transformer IT-k and the arm-to-shell voltage ET-k feed the phase's regulator, which commands the speed $v_{c,k}$ of the hydraulic mast. The power program reads the energy meter JQ and the dips TE and sets the regulators' setpoints and the tap, while the carbon injection that foams the slag is switched by hand. Notice that the regulator closes its loop on electrical signals alone: the arc length it aims at is never measured.

What Measures It

Both regulators of this chapter work from two signals per phase, the current in the phase's conductor and the voltage from its electrode arm to the furnace shell. In plant practice instrument transformers bring both down to levels a panel instrument can take and isolate the instrument from the high voltage of the power circuit, a current transformer for the current and a potential transformer for the voltage [20]. The model's 0.3 % noise and 0.04 s transducer filter on each current and voltage stand for this chain.

> **TECHNOLOGY Sensor: Current and Voltage Transformers**
>
> A current transformer is an iron ring around the power conductor, which passes through it as a primary of one turn, while many turns of secondary wire around the ring step the line current down to a standard 5 A for the instrument. A potential transformer steps the voltage down in the same way, with standard accuracy classes of 0.3, 0.6 and 1.2 % of its rated ratio. Driven through too large a burden, a current transformer's core saturates and its signal no longer follows the shape and size of the line current, and an open secondary develops a dangerous voltage [20]. The arc length itself is difficult to measure in an industrial furnace, so it is controlled through its dependence on the voltage [9].

The arc length that both regulators aim at is logged as truth in the model and read by no controller, as in a real furnace.

What Moves It

Each electrode hangs from an arm on a mast that a hydraulic cylinder raises and lowers. In the 80 MVA furnace that the chapter's circuit is drawn from, three pumps keep an accumulator of 7 m^3 at about 64 bar for the three cylinders [1]. The model's mast, with a lag of 0.2 s and a rate limit of 120 mm/s, stands for this servo without the valve's hysteresis or its leakage.

TECHNOLOGY Final Element: The Hydraulic Electrode Mast

In one 80 MVA furnace the regulator drives three control valves. Each sets the flow of oil from a pressurized accumulator to the cylinder that lifts one electrode assembly of arm, graphite column, cables and cooling water, about 17 000 kg. A command of ± 10 V takes the valve, through a hydraulic pre-amplifier, from full flow out of the cylinder to full flow in. The valve shows hysteresis, is non-linear within ± 20 % of opening and leaks about 0.5 % of its rated flow. The actuators must accelerate the assembly to 100 mm/s in about 100 to 500 ms [1].

The slower final elements are the tap changer, which steps the transformer 50 V at a time while the regulators hold the electrodes, and the carbon injection that foams the slag.

Where the Controller Runs

Control loops in a plant are often nested, and supervisory and lower-level loops run with cycle times from milliseconds to minutes [21]. In the 80 MVA furnace the three PID regulators send their outputs through an input and output card as voltage signals to the valves, and the operating data were logged over a one-second window [1]. The app runs its regulators every 0.02 s, a regulating layer that holds the three regulators, the lift, the re-strike and the recovery hold. The power program, the tap changer and the dips form a slower supervisory layer that sets the regulators' setpoints at each stage.

IN PRACTICE How a Shop Foams Its Slag

Slag foaming is critical for a productive arc furnace. The foam shields the refractories from the intense arc, which allows a high power input. Carbon is injected into the slag, where it reacts with iron oxide to form carbon monoxide, and the decarburization of the metal adds more gas. Foam made by reaction has small bubbles and can last several minutes after the gas stops, while gas injection alone gives larger bubbles and a less stable foam. Simulated furnace slags in the laboratory foamed highest at about 25 % iron oxide, although the gas generation kept rising at higher contents [15].

What Is Modernizing It

The first development is machine vision of the arc itself. On a pilot-scale alternating-current furnace the arc length was measured by image analysis and set against the furnace's electrical data, and the relation between arc length and voltage was found to be non-linear and dependent on the slag composition [9], the relation for which the linear arc law of Equation (5.3) stands in. The second is the validated furnace model. The 80 MVA furnace model was fitted to operating data so that it could serve control design, the optimization of energy consumption and an operator-training simulator [1].

5.9 Limitations

The circuit is quasi-static at 60 Hz, without the arc waveform, harmonics, supply flicker or the series reactor, and a direct-current furnace, whose single electrode has no phase coupling, is not modeled. The arc's voltage is linear in its length, while measured arcs are non-linear and slag dependent [9], and the extinction rule, the re-strike gap, the flicker, the cave-ins, the bore rate and the exposure efficiencies are teaching values. One bucket is melted, the bath is one temperature, there is no chemistry, and the foam is a height, so an iron-oxide-rich slag from hydrogen-reduced iron, which changed the arc's movement in a laboratory furnace [18], has no effect here. The DRI table is a declaration, not a product of the shaft furnace. The recovery hold is the build's design rather than a cited practice, and the regulators read perfect transducers apart from noise and a filter. Every limit of the challenge is a teaching value. The verdicts of single-change policies flip on knife-edge seeds, which is why the counts are given over 21 seeds, and the suggested settings miss 2 of them.

OPEN CHALLENGE **Melting Iron Reduced by Hydrogen**

Pellets reduced with hydrogen carry no carbon, unlike pellets reduced with natural gas, so carbon must be added in the furnace both to finish reducing the iron and to foam the slag. A changing slag composition and a larger amount of slag are among the identified challenges of the hydrogen route, and more slag dissolves more refractory and raises the power demand. Optical emission spectroscopy, already proven in pilot and industrial furnaces, offers a way to follow the arc, the residual hydrogen and the slag while such pellets melt [18]. Progress would be a slag and power practice for carbon-free charges set from measurements inside the furnace rather than from experience with scrap.

5.10 Conclusion

The arc furnace is a plant whose three loops share one floating node, and the simulator logs that node beside the currents, impedances and true arc lengths of one heat. Lowering one electrode 30 mm raises all three currents, phase 3 0.90 as much as phase 1 and phase 2 0.27 as much, and the second neighbor changes sign exactly where the neutral's push leaves a 30° cone, at 268.6 mm on the top tap. The relative gain is 0.745 for current regulation and 1.034 for impedance regulation, and on the same 6 % step three current regulators overshoot 26 to 40 % where three impedance regulators overshoot 9 to 23 % and settle as fast as one. Through bore-in the current regulator breaks 2 to 5 arcs on six seeds and the impedance regulator none, foam is the difference between a wear index of 6.38 in ten minutes and one that stops, and without foam no setting meets the heat, the closest by 0.09 %. The suggested impedance regulator with the stage taps passes the graded heat at 414 kWh/t on 19 of 21 seeds. The model is quasi-static and its charge declared. A fuller model would add the arc's waveform and its flicker on the supply, a non-linear arc law fitted to measurements, slag chemistry with a foam that follows it, a second bucket and the shaft furnace's product read rather than declared. The furnace shows why the choice of what a loop measures can matter more than how it is tuned.

KEY POINTS

- Three arcs close their circuit through one steel bath that no wire ties back to the transformer, so moving one electrode moves all three currents.
- Lowering one electrode raises all three currents, phase 3 0.90 as much as phase 1 and phase 2 0.27 as much.
- The relative gain is 0.745 for current regulation and 1.034 for impedance regulation, and on the same step three impedance regulators settle as fast as one.
- Without foamy slag no setting meets the graded heat, and the suggested settings pass it on 19 of 21 seeds.
- The choice of what a loop measures can matter more than how it is tuned.

Where else this applies. Volume 6, Chapter 2 pairs its loops by the same relative gain, and Volume 4, Chapter 2 has a protective override that acts only near its limit.

References

[1] V. Logar, D. Dovžan, and I. Škrjanc, "Mathematical modeling and experimental validation of an electric arc furnace," *ISIJ International*, vol. 51, no. 3, pp. 382–391, 2011.

[2] V. Vogl, M. Åhman, and L. J. Nilsson, "Assessment of hydrogen direct reduction for fossil-free steelmaking," *Journal of Cleaner Production*, vol. 203, pp. 736–745, 2018.

[3] E. Bristol, "On a new measure of interaction for multivariable process control," *IEEE Transactions on Automatic Control*, vol. 11, no. 1, pp. 133–134, 1966.

[4] J. Kozyra, A. Lozynskyy, Z. Łukasik, A. Kuśmińska-Fijałkowska, A. Kutsyk, and L. Kasha, "Increasing the level of autonomy of control of the electric arc furnace by weakening interphase interactions," *Energies*, vol. 16, no. 24, p. 8114, 2023.

[5] B. Boulet, G. Lalli, and M. Ajersch, "Modeling and control of an electric arc furnace," in *Proceedings of the 2003 American Control Conference*, vol. 4, 2003, pp. 3060–3064.

[6] A. S. Hauksdóttir, A. Gestsson, and A. Vésteinsson, "Current control of a three-phase submerged arc ferrosilicon furnace," *Control Engineering Practice*, vol. 10, no. 4, pp. 457–463, 2002.

[7] U.S. Geological Survey, "Mineral commodity summaries 2025: Iron and steel," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[8] R. D. M. MacRosty and C. L. E. Swartz, "Dynamic modeling of an industrial electric arc furnace," *Industrial & Engineering Chemistry Research*, vol. 44, no. 21, pp. 8067–8083, 2005.

[9] H. Pauna, T. Willms, M. Aula, T. Echterhof, M. Huttula, and T. Fabritius, "Electric arc length-voltage and conductivity characteristics in a pilot-scale AC electric arc furnace," *Metallurgical and Materials Transactions B*, vol. 51, no. 4, pp. 1646–1655, 2020.

[10] S. A. Billings and H. Nicholson, "Modelling a three-phase electric arc furnace: a comparative study of control strategies," *Applied Mathematical Modelling*, vol. 1, no. 7, pp. 355–361, 1977.

[11] K. Ogawa, "A trend of ultra high power operation and automization of electric arc furnaces," *Denki-Seiko [Electric Furnace Steel]*, vol. 39, no. 1, pp. 62–72, 1968.

[12] S. A. Billings, F. M. Boland, and H. Nicholson, "Electric arc furnace modelling and control," *Automatica*, vol. 15, no. 2, pp. 137–148, 1979.

[13] W. E. Staib and R. B. Staib, "The intelligent arc furnace controller: a neural network electrode position optimization system for the electric arc furnace," in *Proceedings of the 1992 International Joint Conference on Neural Networks (IJCNN)*, vol. 3, 1992, pp. 1–9.

[14] V. Logar, D. Dovžan, and I. Škrjanc, "Modeling and validation of an electric arc furnace: Part 1, heat and mass transfer," *ISIJ International*, vol. 52, no. 3, pp. 402–412, 2012.

[15] H. Matsuura and R. J. Fruehan, "Slag foaming in an electric arc furnace," *ISIJ International*, vol. 49, no. 10, pp. 1530–1535, 2009.

[16] J. Stubbles, "Energy use in the U.S. steel industry: A historical perspective and future opportunities," U.S. Department of Energy, Office of Industrial Technologies, Tech. Rep., 2000.

[17] M. Kirschen, K. Badr, and H. Pfeifer, "Influence of direct reduced iron on the energy balance of the electric arc furnace in steel industry," *Energy*, vol. 36, no. 10, pp. 6146–6155, 2011.

[18] H. Pauna, T. Kokkonen, P. Cavaliere, M. Bayat, A. Mirowska, T. Alatarvas, M. Huttula, and T. Fabritius, "Hydrogen direct reduced iron melting in an electric arc furnace: Benefits of in situ monitoring," *Journal of Sustainable Metallurgy*, vol. 11, no. 4, pp. 4655–4667, 2025.

[19] J. L. García-Muñoz, G. Rojo, I. Unamuno, and J. J. Laraudogoitia, "Analysis of AC EAF electrical magnitudes for different process conditions using a modified circular diagram perspective," *Archives of Metallurgy and Materials*, vol. 53, pp. 607–612, 2008.

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

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

5.11 Exercises

5.1 ⌕ Concept check · 2 min

In case 1 all three phases are on MAN on a flat bath with arcs of 336 mm, and Phase 1 position is lowered by 30 mm. What happens to the currents of phases 2 and 3?

(a) Neither changes, because no wire connects one electrode to another

(b) Both rise, phase 3 nearly as much as phase 1 and phase 2 only a little

(c) Both fall, because the three currents must add to zero when phase 1 rises

(d) Both rise by as much as phase 1, because the bath potential stays at zero

5.2 ⌕ Concept check · 2 min

In case 2 three current regulators ring on a step of phase 1's setpoint, while in case 3 three impedance regulators settle the same kind of step as fast as phase 1 alone. What explains the difference?

(a) The impedance regulators run at four times the gain, so they act before the neighbors move
(b) The impedance regulators use a deadband that ignores the neighbors' small errors
(c) Each phase's impedance is mostly its own arc's voltage over its own current
(d) The impedance step is too small to move the neighbors' currents at all

5.3 Concept check · 2 min

The short-circuit lift raises an electrode at full rate while its phase reads an impedance below $Z_{trip} = 1.0\,\mathrm{m\Omega}$. Why is a short detected by a low impedance and not by a high current?

(a) The current transformers saturate in a short, so the current reading is lost
(b) A shorted phase reads exactly zero impedance, so any positive trip level is safe
(c) A short circuit draws less current than a normal arc on the top tap
(d) With a floating neutral a neighbor of the shorted phase can carry more current than the shorted phase

5.4 Theory and calculation · 5 min

The arc of Equation (5.3) has $V_0 = 40\,\mathrm{V}$, $E = 1.2\,\mathrm{V/mm}$ and $k_{ext} = 0.85$. Taps 1, 5 and 9 give 600, 800 and 1000 V line to line.

(a) Find the arc voltage of a 336 mm arc. *Checkpoint:* 443 V.

(b) Find the phase voltage U and the extinction length L_{ext} on taps 1, 5 and 9, to the nearest millimeter. *Checkpoint:* 212, 294 and 376 mm.

(c) The application's plan set a 400 mm extinction length on the lowest tap. Show that no arc of that length can burn there. *Checkpoint:* 520 V against a phase voltage of 346 V.

5.5 Theory and calculation · 6 min

One balanced phase on tap 9 has $U = 577.4\,\mathrm{V}$, $R_s = 0.5\,\mathrm{m\Omega}$ and $X_s = 5.5\,\mathrm{m\Omega}$. Use Equation (5.8).

(a) Find $|Z_s|$, the maximum arc power P^* per phase and the current I^* at which it occurs. *Checkpoint:* 5.523 mΩ, 27.67 MW, 70.79 kA.

(b) Find the lossless bound $U^2/(2X_s)$ and how far below it P^* lies, in percent of the bound. *Checkpoint:* 30.30 MW, 8.7 %.

(c) Find the short-circuit current $U/|Z_s|$. The arc is stable above 47.61 kA. On which side of the maximum do the program's setpoints of 60 and 65 kA lie, and why is that the side to run on? *Checkpoint:* 104.5 kA, the long-arc side.

5.6 Theory and calculation · 7 min

At a balanced point the sensitivity of the three currents to the three tip heights, divided by its diagonal, is the circulant matrix

$$G = \begin{bmatrix} 1 & a & b \\ b & 1 & a \\ a & b & 1 \end{bmatrix},$$

with $a = 0.90$ and $b = 0.27$ for the currents and $a = 0.39$ and $b = 0.12$ for the impedances. The eigenvalues of a circulant matrix are $\mu_m = 1 + a\,\omega^m + b\,\omega^{2m}$ with $\omega = e^{j2\pi/3}$ and $m = 0, 1, 2$, and every diagonal element of its inverse is $\frac{1}{3}\sum_m 1/\mu_m$.

(a) Find the eigenvalues of the current matrix as magnitude and angle. *Checkpoint:* 2.17 and $0.685\angle \pm 52.7°$.

(b) Find the diagonal relative gain $\lambda = g_{11}\,(G^{-1})_{11}$ of current regulation, to three decimals. *Checkpoint:* 0.742.

(c) Repeat (a) and (b) for impedance regulation and say what the two results mean for three loops tuned one phase at a time. The chapter's values, from unrounded sensitivities, are 0.745 and 1.034. *Checkpoint:* 1.51 and $0.781\angle\pm17.4°$, then 1.035.

5.7 App challenge · 8 min

Load scenario 2, "Current regulation on all three phases", at the default seed with no changes, and set the speed to $1\times$. Press Run and, at any time between 2 and 8 s, set Phase 1 offset to $+6\,\%$. Run at least 10 s more and read the Phase 1 step response tile. Then load scenario 3, "Impedance regulation on the same step", and repeat with an offset of $-6\,\%$.

(a) Give the overshoot, the time after which the response stayed inside a quarter of the step, and the I3 swing for each case. Because the arcs' flicker moves them with the instant of the step, give each as a range that holds across the window. *Checkpoint:* Case 2: 26 to 40 %, 2.0 to 2.6 s, 2.9 to 3.8 kA. Case 3: 9 to 23 %, 0.5 to 0.7 s, 1.2 to 1.9 kA.

(b) Phase 1 alone, with phases 2 and 3 on Man, settles inside a quarter of the step in 0.6 to 0.7 s at $K_I = 20$ and 0.5 to 0.7 s at $K_Z = 80$. What has the interaction cost each regulator? *Checkpoint:* Current: about three to four times the settling time. Impedance: nothing measurable.

5.8 App challenge · 7 min

Load scenario 5, "Melting at the top tap: extinction and the deadband", at the default seed and set the speed to $20\times$. Run until the case pauses at 5 min and read the Mast travel tile (average mast speed and longest arc) and the Arcs broken · longest out tile. Set Impedance deadband to 0.50 mΩ, press Reset and run to the pause again. Repeat with 1.00 mΩ. Give speeds to ±0.1 mm/s and lengths to ±1 mm.

(a) Tabulate the three runs. *Checkpoint:* 7.3, 1.4 and 1.5 mm/s, then 344, 360 and 376 mm, then 0, 0 and 1 broken.

(b) Express each deadband in millimeters of arc with the diagonal sensitivity of 0.0355 mΩ/mm, and explain why the larger deadband buys no less travel. *Checkpoint:* About 14 and 28 mm.

5.9 App challenge · 9 min

Load the challenge, "Challenge · the full heat" (case 9), at its own seed with no changes, set the speed to $200\times$ and run until the heat pauses at TAP (about 10 s). Read the verdict, the Heat tile, the Wear index (truth) tile and the Arcs broken · longest out tile. Repeat three times after Reset: with Regulator set to Impedance alone, with Power program set to Stage taps alone, and with both.

(a) Tabulate the energy per tonne, the tap-to-tap time, the wear index, the arcs broken and the verdict for the four heats. *Checkpoint:* Only both changes together pass, at 414 kWh/t and a wear of 0.23.

(b) Say what each single change fixes and what it leaves, in terms of the chapter's cases 4 and 6. *Checkpoint:* Impedance fixes the arcs out. Stage taps hold the wear near the budget.

5.10 Application area · 3 min

In case 6 the last 13.2 t of scrap melts with the slag left flat. The program then runs tap 6 at 66 kA with 214 mm arcs and 59 MW, and the wear index reaches 6.38 at 10 min. With Slag on at 2.5 min the program moves to tap 9 and the heat taps at 1629 °C after 390 s with the index at 0.79. Why does a foamy slag allow both more power and less wall wear, and why can no setting pass the challenge without it?

5.11 Application area · 3 min

Case 7 runs tap 8 at a fixed impedance setpoint of 6.01 mΩ with about 63 kA, arcs of 283 to 295 mm and 74.2 MW. Tap up steps the voltage 2.01 s after the button with the electrodes held, and 20 s after the button the furnace runs at about 67 kA, arcs of 303 to 309 mm and 82.9 MW.

(a) Explain why the same impedance on a higher tap gives a longer arc at more current and more power. *Checkpoint:* The tap moves the circle. The impedance fixes the ratio $|V_m|/|I|$.

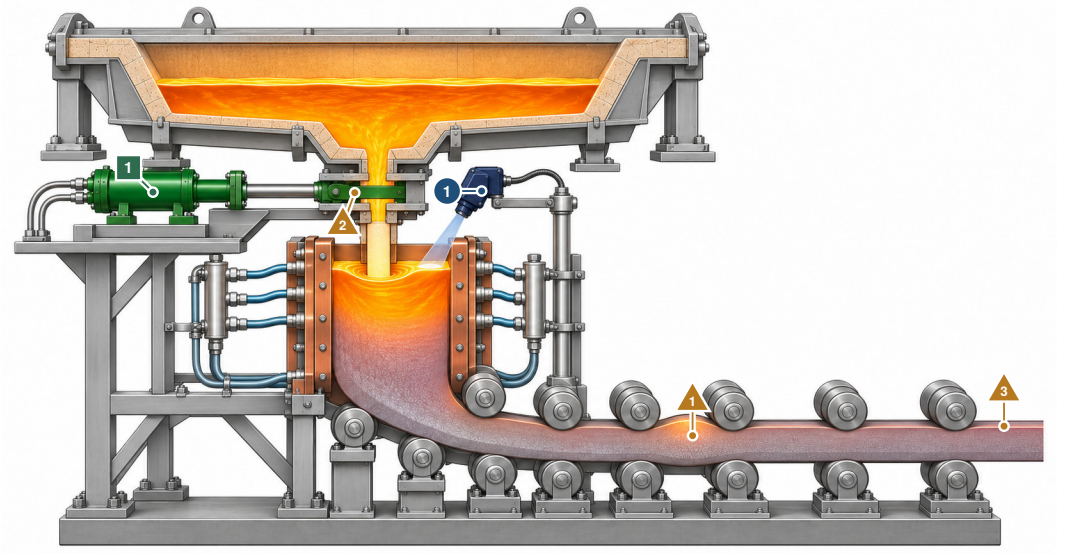
(b) The tap changer holds the regulators for 2 s. Why must the short-circuit lift still act during that hold? *Checkpoint:* A cave-in during a tap step stayed shorted 2.06 s. With the lift acting, 0.53 s.

5.12 Review the agent · 5 min

An AI assistant tuning the current regulators writes: "With phases 2 and 3 on Man, phase 1 at $K_I = 20$ overshoots a 6 % step by only 8 to 16 % and settles in under 0.7 s. The three phases are identical, so the same tuning holds for all three. To settle faster, raise all three to $K_I = 30$, which the single-phase test shows is still well damped." Find the flaw, and say how you would check it in the app before anyone uses the gain.

CHAPTER 6

Mold Level Simulator for Telling a Sticking Gate From Bulging and Overtuning



The controller moves 1 the slide gate, moved by its hydraulic cylinder

The controller reads 1 the level of the steel in the mold

What disturbs it 1 the strand bulging between its rolls 2 gate plates that stick 3 a change in casting speed

Notice that the strand leaves at a rate set by the casting speed, not by the level, so every gap between the flow through the gate and the steel withdrawn makes the level climb or fall until the gate moves. Illustration, not to scale.

IN PLAIN TERMS **Three Reasons a Steel Level Swings**

The main idea. Melted steel pours into a copper box with no bottom. It freezes there into a slab that is pulled out below. A sliding gate above the box sets how fast steel flows in, and a controller moves the gate to keep the level steady.

The challenge. The level can swing because the gate sticks, because the slab bulges between its rollers, or because the controller reacts too hard. All three draw almost the same wavy line, and each fix works only on its own cause.

Why it matters. The steel surface in the box decides the quality of the slab, and the wrong fix wears out the gate and leaves the swing.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the fact box of the Introduction, the history box of Section 6.2 and the one of Section 6.4, the fact and try-it boxes of Section 6.5, and the key points
Study	35 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 6.5, Section 6.6, Section 6.7 with its two open challenges

The control question The mold of a continuous caster is an integrator fed through a slide gate, and its level can oscillate because the gate plates stick, the strand bulges between its rolls or the controller carries too much gain. On a trend the three look alike. Which tests tell them apart, and what does each remedy cost when it is applied to the wrong cause?

Learning objectives

- Derive the integrating gain of the mold from the gate's lens area and the volume balance, and compute SIMC settings and the ultimate gain from it.
- Explain why a sticking gate on an integrating plant cycles even without integral action.
- Apply the open-loop test, the speed-change test and the signature plot to name the cause of an oscillation.
- Compute the bulging period from the roll pitch and casting speed, and the gate motion a canceller needs.
- Compare the knocker, the canceller and retuning by level variation and gate travel, and choose the order of repairs.

Prerequisites Sections 1.6.2 (velocity form and anti-windup), 1.4.5 (a disturbance frequency set by speed) and 1.6.4 (scenarios, tiles and the CSV file).

The app https://apmonitor.com/apps/continuous-caster/

See also Volume 2, Chapter 5 for an integrating level tuned by SIMC and Volume 2, Chapter 4 for a sticking steam valve. Volume 1, Chapter 13 for loop-performance tests on recorded loops, one of them this caster. Volume 7, Chapter 11 for a disturbance whose frequency the speed sets and Volume 7, Chapter 13 for two limits that look alike but need different remedies.

6.1 Introduction

An oscillating control loop invites retuning, and retuning is often the wrong response. Surveys of industrial loop performance have found that a large share of oscillating loops oscillate because of the final control element or an external periodic disturbance rather than the controller settings, and that the diagnosis has to separate those causes before any remedy is chosen [1, 2]. A sticking valve responds to tuning, just never enough. A periodic disturbance responds to tuning in either direction depending on where its frequency sits against the loop bandwidth. Only a genuinely overtuned loop is cured by less gain.

The steel continuous caster is an unusually clean home for that problem. The mold is an integrator, because the strand withdraws steel at a rate set by the casting speed and not by the level, and on an integrating plant a sticking valve produces a limit cycle at any nonzero gain, even without integral action [3]. The strand bulges between its support rolls and pumps the liquid core into and out of the mold at a frequency that is known exactly from the casting speed and the roll pitch [4]. The measurement delay and the hydraulic actuator put a hard ceiling on the gain. All three occur on real casters, often at once, and all three produce a level trend that looks roughly the same.

DID YOU KNOW? Nearly All Steel Is Cast Continuously

By 2000, barring specific situations, continuous casting had entirely replaced ingot casting in the United States, and it improved yield dramatically [5]. The country then made about 95 million tons of cast products a year. The slab caster of this chapter withdraws about 52 kg/s of steel from one strand.

The simulator is built against one misconception, that an oscillation is a tuning problem until proven otherwise. The broader class is loop performance diagnosis, which appears wherever a control valve can stick and a disturbance can be periodic: paper machine headboxes, flotation cells, drum levels, and others [1, 6]. The shape of the diagnosis is the same in all of them. Opening the loop separates an internal oscillation from an external one. Moving the frequency of a suspected external source confirms or rules it out. Plotting the controller output against the valve position exposes stiction as a parallelogram rather than a line [7].

This chapter works with a mold level simulator in which all three causes can be switched on alone or together, the diagnostic tests can be run as they are on a plant, and each remedy can be applied to each cause and priced in level variation and gate wear. Stick-slip models, stiction compensation by pulse injection and bulging cancellation are all established methods. The mold, gate, strand and sensor model is written out completely enough to be rebuilt, the diagnostic instruments are drawn in the interface, and the scenarios are ordered so that diagnosis comes before treatment. Section 6.2 gives the domain history. Section 6.3 sets out the model: the mold as an integrator, the slide gate, two-parameter stiction, bulging, clogging and the sensor. Section 6.4 describes the level controller, the knocker, the bulging canceller and the diagnostic instruments. Section 6.5 reports the measured results: the healthy reference, the three signatures, combined faults with the order of repairs, and the graded challenge. Section 6.6 turns to the real sensor, gate and controller platform, Section 6.7 states the limitations and Section 6.8 draws the conclusions. The exercises follow.

6.2 Mold Level Control on Continuous Casters

Continuous casting displaced ingot casting for most of the world's steel over the second half of the twentieth century, because it removes the reheating and primary rolling of ingots and raises yield [8]. Liquid steel runs from a ladle into a tundish, a refractory buffer vessel about a meter deep, and from the tundish through a flow control device and a submerged entry nozzle into a bottomless water-cooled copper mold. A shell freezes against the copper and is withdrawn continuously while support rolls contain it until the core has solidified. Mathematical modeling of the fluid flow, heat transfer and solidification in this process has been reviewed at length [9].

The meniscus is where surface quality is decided. Mold flux floats on it and the first shell forms against it, so level fluctuations either entrain flux into the solidifying steel or disturb the initial shell, and both appear later as surface and subsurface defects. Instrumenting the mold to understand and prevent those defects was a central theme of caster research in the 1980s [10]. Level measurement moved from floats and radioactive sources to eddy-current and electromagnetic sensors, and flow control from stopper rods alone to stopper rods and slide gates, the latter being two refractory plates with round bores slid across each other by a hydraulic cylinder.

> ⧖ HISTORY **Automatic Mold Level Control (1955 to 1984)**
>
> The first continuous casting machine in Japan was installed in 1955, and by 1982 almost all common steels there were cast this way [11]. A paper of that year called the development of a mold level sensor one of the most important and difficult tasks of the process. Sensors using gamma rays, thermocouples and eddy currents had been tried, and only a few could work beside the electromagnetic stirrers then being added for high grade steel. By 1984 one slab caster held its mold level within 3 mm by moving the tundish slide gate continuously with a hydraulic unit [12]. Casting then started at the push of a button, and the mold operator was freed from the lasting tension of controlling the level by hand. The same note tied level fluctuation to entrapped scum and to longitudinal cracks.

Once the level was measured continuously it was put under closed-loop control, and the history of the loop since is largely a history of its disturbances. Robust controllers were designed against the variation in nozzle gain and the measurement delay [13], and supervisory schemes, including fuzzy supervision of the level controller, were used to handle changing operating conditions [14]. As casting speeds rose and shells thinned, unsteady bulging became a recognized source of periodic level fluctuation, and its mechanism was measured and simulated in detail [4]. Because its period is the roll pitch divided by the casting speed, dedicated cancellation schemes that exploit the known frequency were developed and validated on plants [15].

Stiction in slide gates is a maintenance reality with a separate literature of its own. The data-driven two-parameter model of a sticking valve gives a stick band and a slip jump that reproduce the characteristic trends of industrial loops [7], a friction compensator that adds short pulses to

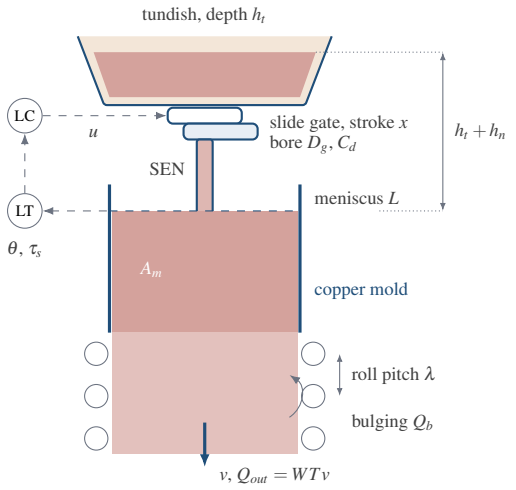


Figure 6.1 Mold level loop on a slab caster. Steel from a tundish of depth h_t passes a slide gate of bore D_g at stroke x and a submerged entry nozzle (SEN) into a mold of free-surface area A_m. The strand of width W and thickness T leaves at casting speed v between support rolls of pitch λ, and bulging between the rolls adds a periodic flow Q_b. The level transmitter LT reads the meniscus L after a delay θ and a filter τ_s, and the controller LC sets the gate command u.

Table 6.1 Classification of Variables at the Design Operating Point

Symbol	Role	Nominal	Range
L	CV, true meniscus level	100 mm	0–190 mm (lip)
L_m	measured level	100 mm	0–200 mm span
u	MV, gate command	57.0 %	0–100
x	state, gate position	57.0 %	0–100
c	state, clog fraction	0	0–0.6
ϕ	state, bulging phase	—	—
v	DV, casting speed	1.20 m/min	0.6–2.0
h_t	DV, tundish depth	1.00 m	0.5–1.3
S, J	DV, gate stiction	0, 0 %	0–8, 0–4
a_b	DV, bulging amplitude	0 mm	0–10
r_c	DV, clogging rate	0 %/min	0–3
n	DV, level noise (1σ)	0.2 mm	0–2

the valve command, the knocker, was proposed for pneumatic control valves [16], and methods to detect stiction from routine data, including in integrating loops where the signature differs, have been compared extensively [2, 3, 6]. Volume 1, Chapter 13 applies such loop-performance tests to recorded loops, one of them this caster's overtuned case. The level loop on a caster therefore sits at the intersection of three mature literatures that are rarely taught together.

6.3 Process Model

The simulated caster makes a slab 1.60 m wide and 0.23 m thick at 1.2 m/min. The strand withdraws 7.36 L/s of steel, about 52 kg/s, from a free surface of 0.353 m^2 once the nozzle is subtracted, and the slide gate balances that withdrawal at 57 % of its stroke. The arrangement is shown in Figure 6.1.

The variables are classified in Table 6.1. The level is measured from the zero of a 200 mm sensor span, with the mold lip 190 mm above that zero and the setpoint at 100 mm. Four states are integrated: the level L, the gate position x, the clog fraction c and the bulging phase ϕ.

6.3.1 The Mold as an Integrator

A volume balance on the steel above the strand gives the level derivative, as shown in Equation (6.1).

$$A_m \frac{\mathrm{d}L}{\mathrm{d}t} = Q_{in}(x, c, h_t) - Q_{out}(v) + Q_b(t) \tag{6.1}$$

where $A_m = WT - \pi D_{SEN}^2/4$ is the slab section less the nozzle, and the withdrawal is $Q_{out} = WTv$. No term on the right depends on L, so the mold has no self-regulation. Every imbalance ramps the level at a constant rate, and the only thing that stops the ramp is the gate.

6.3.2 The Slide Gate

Two plates with equal round bores slide across each other, so the open area is the lens where two circles of radius $D_g/2$ overlap. With the stroke $s = x/100$ and the center offset $d = 2(1-s)$ in bore radii, the open fraction of the bore is given by Equation (6.2).

$$f(s) = \frac{1}{\pi}\left[2\cos^{-1}\frac{d}{2} - \frac{d}{2}\sqrt{4-d^2}\right] \tag{6.2}$$

The gate flow is an orifice driven by the tundish head, reduced by the clog, as shown in Equation (6.3).

$$Q_{in} = C_d \frac{\pi D_g^2}{4}\sqrt{2g\,(h_t + h_n)}\, f(s)\,(1-c) \tag{6.3}$$

The lens opens slowly from closed and fastest near full open, so the process gain depends on where the gate sits. At the design point each percent of stroke is worth 0.18 L/s and the integrating gain is $K_I = 0.51\,\mathrm{mm/s}$ per percent of stroke.

The hydraulic cylinder follows the position delivered by the stiction model, x_s, as a first-order lag with a speed limit, Equation (6.4).

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \mathrm{clamp}\left(\frac{x_s - x}{\tau_a}, -R_a, R_a\right) \tag{6.4}$$

with $\tau_a = 0.15\,\mathrm{s}$ and $R_a = 15\,\%/\mathrm{s}$.

6.3.3 Two-Parameter Stiction

The plates stick according to the two-parameter model of stiction, a stick band plus dead band S and a slip jump J [7], written in the simulator as a force balance on the plate evaluated at every controller scan. With the static band $f_s = (S+J)/2$ and the kinetic band $f_d = (S-J)/2$, a stuck plate at x_s stays put until the command u satisfies $|u - x_s| > f_s$, slips to $x_s = u - \mathrm{sgn}(u - x_s) f_d$, and then tracks at the kinetic distance until the command stops or reverses, Equation (6.5).

$$x_s \leftarrow \begin{cases} u - \mathrm{sgn}(u - x_s)\, f_d, & |u - x_s| > f_s, \text{ or moving} \\ & \text{and } |u - x_s| > f_d \\ x_s \text{ (stuck)}, & \text{otherwise} \end{cases} \tag{6.5}$$

A reversal therefore costs S of dead command travel, and every breakaway ends in a jump of J. With $S = 0$ the plate follows the command exactly. The stiction scenarios use $S = 5\,\%$ and $J = 1.5\,\%$, so the plate needs 3.25 % of command pressure to break away.

6.3.4 Bulging, Clogging and the Sensor

Bulging is one sinusoid at the roll-passing frequency, Equation (6.6).

$$Q_b = A_m a_b \omega \cos\phi, \qquad \frac{\mathrm{d}\phi}{\mathrm{d}t} = \omega = \frac{2\pi v}{\lambda} \tag{6.6}$$

Table 6.2 Model Parameters. Class marks what the application shows: visible parameters are adjustable in the parameter panel, hidden ones are not

Symbol	Meaning	Value	Class
W, T	slab width, thickness	1.60, 0.23 m	visible
D_{SEN}	nozzle outside diameter	0.14 m	visible
λ	support roll pitch	0.25 m	visible
D_g	slide-gate bore	70 mm	visible
C_d	discharge coefficient	0.75	visible
h_n	tundish floor to meniscus	0.50 m	visible
τ_a	actuator lag	0.15 s	visible
R_a	actuator speed limit	15 %/s	visible
θ	measurement and flow delay	0.30 s	visible
τ_s	sensor filter	0.10 s	visible
—	slag-entrainment band	±5 mm	visible
—	level alarm	±20 mm	visible
—	casting speed ramp rate	0.30 m/min/min	hidden
g	gravity	9.81 m/s^2	fixed

where a_b is the level swing the bulging would cause with the gate held still. At 1.2 m/min on 0.25 m rolls the period is 12.5 s. The frequency is set by a length and a speed, the case of Section 1.4.5, as a road feature under a moving car sets the frequency in Volume 7, Chapter 11. The clog grows at $\mathrm{d}c/\mathrm{d}t = r_c/6000$ per second while $c < 0.6$, and a clog release event resets it to zero.

The eddy-current sensor sees the level after a transport and measurement delay $\theta = 0.30$ s, adds colored noise with a 0.25 s time constant, and passes the sum through a first-order filter $\tau_s = 0.10$ s before clamping it to the 200 mm span. The level itself is clamped at the lip, where an overflow is flagged, and a level alarm is raised at 20 mm from setpoint.

The parameters are listed in Table 6.2. The casting speed does not jump when its setting changes, because it is ramped at 0.30 m/min per minute in the post-step correction.

6.3.5 Assumptions

The meniscus is flat and single-valued. Standing waves, jet impingement from the nozzle ports and argon bubbling are represented only as band-limited sensor noise, although on a real caster they are a large part of the measured level signal [9].

The gate flow is an orifice with a constant discharge coefficient through the geometric overlap of two bores. Erosion of the bore, argon injection and flow separation in the nozzle are not modeled.

Solidification, mold oscillation and mold flux are outside the model, and the withdrawal is exactly slab section times speed. Shell shrinkage and the density difference between liquid and solid steel are neglected.

Bulging is a single sinusoid at a single roll pitch with an amplitude that does not depend on speed. Real unsteady bulging combines several pitches and grows with speed and shell temperature [4].

Stiction is evaluated at the 50 ms controller scan with the two-parameter model. Velocity-dependent friction and backlash are not included. Clogging is uniform and is removed only in one piece.

6.4 Implemented Controller

6.4.1 Velocity Form PI Loop

The level controller is the framework's velocity-form PID of Section 1.6.2, which differentiates the measurement rather than the error and executes every $\Delta t = 0.05$ s. The algorithm as coded for this

loop is given by Equations (6.7) to (6.10).

$$e_k = L_{SP} - L^f_{m,k} \tag{6.7}$$

$$P_k = K_c e_k, \qquad D_k = -K_c \tau_D \frac{L^f_{m,k} - L^f_{m,k-1}}{\Delta t} \tag{6.8}$$

$$\Delta u^I_k = \frac{K_c}{\tau_I} e_k \Delta t \tag{6.9}$$

$$\begin{aligned} s_k &= s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ c_k &= u_{ff,k} + s_k + I_{k-1} \\ u_k &= \text{clamp}(c_k, 0, 100) \\ I_k &= I_{k-1} + \Delta u^I_k + \frac{\Delta t}{T_t}(u_k - c_k) \end{aligned} \tag{6.10}$$

where L^f_m is the measured level after the optional level filter τ_f, s_k is the running sum of the proportional and derivative increments, c_k is the demanded output, u_{ff} is the canceller feedforward of Section 6.4.3, and $T_t = \tau_I$ is the tracking time. In the symbols of Equations (1.14) and (1.15), s_k is S_k, c_k is v_k and u_{ff} takes the place of the bias u_b. Anti-windup is the back-calculation term of Equation (6.10), which removes $\Delta t/\tau_I = 1/92$ of the excess per scan at the default $\tau_I = 4.6\,\text{s}$ and runs only when integral action is present. The knocker pulse is added after the clamp, $u_{cmd,k} = \text{clamp}(u_k + u_{kn,k}, 0, 100)$, and only u_{cmd} reaches the gate plates, so the integral never sees the pulses. Physics advances by a fourth-order Runge-Kutta step of 0.01 s. The process gain from stroke to level is positive, so a low level gives a positive error and opens the gate. The loop is therefore reverse acting, $a = +1$ in the convention of Section 1.6.2. The derivative time defaults to zero and the optional level filter τ_f to zero, because the measured level is full of surface waves. Transfer between manual and automatic is bumpless in both directions, as Section 1.6.2 describes.

The default tuning follows the SIMC rule for an integrating plant with dead time [17]. The apparent dead time is the measurement delay, the sensor filter, the actuator lag and half a scan, $\theta_a \approx 0.30 + 0.10 + 0.15 + 0.025 = 0.58\,\text{s}$. With the closed-loop time constant $\tau_c = \theta_a$ the rule gives Equation (6.11).

$$\begin{aligned} K_c &= \frac{1}{K_I(\tau_c + \theta_a)} \approx 1.7\,\%/\text{mm} \\ \tau_I &= 4(\tau_c + \theta_a) \approx 4.6\,\text{s} \end{aligned} \tag{6.11}$$

For the same plant the ultimate gain is about $\pi/(2K_I\theta_a) \approx 5.3\,\%/\text{mm}$ and the ultimate period about $4\theta_a \approx 2.3\,\text{s}$. The simulated loop tolerates a little more for small upsets, which still decay at 5.5 %/mm and grow at 6.0 %/mm, because the sensor filter and the actuator lag also lower the loop gain at that frequency, which the dead-time approximation ignores. The top of the integral time slider, 60 s, means no integral action. Volume 2, Chapter 5 identifies and tunes another integrating level, a cooling tower basin, by the same rule.

6.4.2 The Knocker

The knocker adds short pulses to the command in the direction the controller output has been moving [16]. Every period $h_k = 0.5\,\text{s}$ the sign of the change in controller output since the previous pulse is recorded, ignoring changes under 0.005 %, and for the pulse width $\tau_k = 0.1\,\text{s}$ a pulse u_{kn} of amplitude a_k and that sign is added to the command. The default amplitude is 2.0 %, and about half the stick band is the usual starting point. Each pulse carries the plate across part of its stick band, so the gate moves when the controller asks rather than after the error has grown enough to break it free.

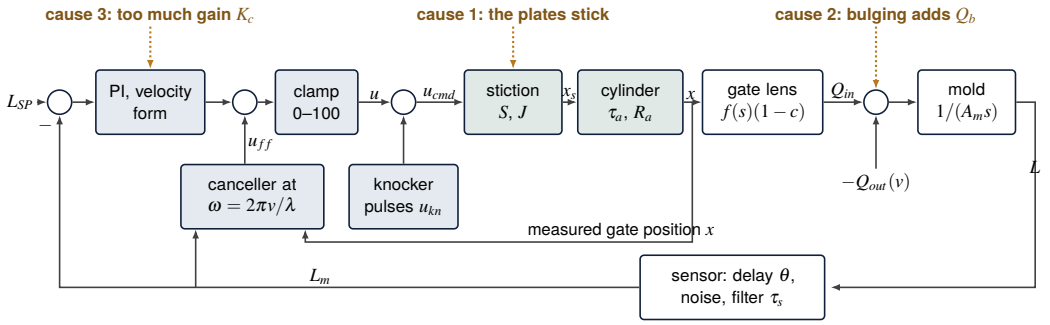


Figure 6.2 The control structure as a block diagram. The PI output and the canceller's feedforward u_{ff} are summed before the clamp, and the knocker's pulses are added after it. The three causes of the chapter enter at three different blocks: the plates' stiction sits between command and gate, bulging adds flow at the mold, and too much gain lives in the PI block. Notice that the mold integrates every flow imbalance, so each cause can make the level swing, and only a test that separates the blocks tells them apart.

> ⧖ HISTORY **The Slide Gate (1885 to 1986)**
>
> The idea of the sliding gate was patented in the United States in 1885, but it was first put into practice at a West German steelworks in 1964 [18]. Its use spread internationally from 1968, and by 1973 it served 115 plants of 80 firms in Japan alone. With no stopper rod in the ladle, pouring could be run by remote control, and its rate was easier to control automatically. Sticking soon showed up as a control problem. A 1986 note reported large level swings caused by solidified steel adhering to the gate and delaying its motion [19]. Its remedy superposed a vibratory signal on the gate command, so that the gate vibrated all the time, and the level fluctuation fell by about 80 percent. The knocker of this chapter pursues the same aim with pulses in one direction only.

6.4.3 The Bulging Canceller

The canceller exploits the fact that the bulging frequency $\omega = 2\pi v/\lambda$ is known. Two least-squares tone estimators track the phasors of the measured level and the measured gate position at that frequency, after a 20 s high-pass, with an averaging time $T_a = 12$ s. The level phasor is moved back through the sensor filter and forward through the delay to the mold, and a flow balance at the one frequency, $j\omega A_m \hat{L} = (\partial Q_{in}/\partial x)\hat{X} + \hat{Q}_b$, turns the two phasors into the phasor of the bulging flow. The gate motion that cancels it is advanced by $\omega\tau_a$ for the actuator lag and corrected for the phase lead of the high-pass, then capped at 15 % of stroke and added to the command as the feedforward u_{ff} in Equation (6.10), in automatic mode only. It ramps in over 5 s when switched on. The estimators run in both modes, so the canceller is warm when it is engaged. Figure 6.2 assembles the level controller, the canceller and the knocker into one block diagram.

6.4.4 Diagnostic Instruments

The interface computes the dominant oscillation period of the last minute of measured deviation once per second. Whether the level is oscillating is decided by the height of the second peak of the autocorrelation, and the period is the mean spacing of upward zero crossings with hysteresis at 30 % of the amplitude, which is $\sqrt{2}$ times the standard deviation of the same minute. It also shows the RMS of the true level and the gate travel per minute, both exponentially weighted with a 30 s time constant, the share of the run outside the slag-entrainment band, the gate reversal rate, and a plot of the controller output against the gate position over the last 40 s, with the healthy diagonal dashed. A thin line along the diagonal is a healthy gate, and a parallelogram whose width is the stick band is stiction [7].

6.5 Results

The nine teaching scenarios and the challenge are listed in Table 6.3 in the order the application offers them, with the teaching point of each and what completes it. A case is credited when the learner performs the listed actions and the simulation reaches the listed state, as Section 1.6.4 describes. The challenge is judged on the case as given, and changing a model parameter or a disturbance voids the attempt. All reported values come from the running application at its default seed, with each setting made before the run starts unless a time is given. Every case with stiction starts with the plate stuck at 57.6 % of stroke, 0.6 % open of balance, because a plate stuck exactly at balance never has a reason to move.

> **+ DID YOU KNOW? A Remedy That Only Adds Wear**
>
> From the simulator: with bulging present and a healthy gate, switching the knocker on moves the level only from 1.93 to 1.78 mm RMS, while the gate travels 2.1 times as far (Table 6.4, case 8). The pulses meant for a sticking gate find nothing to free.

6.5.1 The Integrator and the Healthy Reference

In scenario 1 the gate starts 0.4 % of stroke too far open, a surplus of 0.07 L/s, and the level rises about 0.19 mm/s. Setting the gate to the exact balance stops it, and a clog of half a percent per minute then turns the perfect balance into a 3.2 mm error in one minute (1.9 mm with the stroke rounded to 0.1 %, as the application's balance helper sets it). The manual case makes the absence of self-regulation something the student has to fight rather than read about.

Scenario 2 is the reference. With 0.2 mm of sensor noise the true level holds to 0.13 mm RMS, a 3 mm setpoint step overshoots by 0.99 mm, and the gate still travels 47 % of its stroke per minute chasing noise, which is the baseline wear on a healthy loop.

6.5.2 Three Signatures

The signatures of the three causes are summarized in Table 6.4. With $S = 5\,\%$ and $J = 1.5\,\%$ the level cycles with an RMS of 1.11 mm and a mean period of 13.7 s, although single cycles range from 9.8 to 18.2 s. The gate trace becomes a staircase under a sawtooth command, and the signature plot opens into a parallelogram, as shown in Figure 6.3. Removing integral action does not stop it (1.02 mm RMS in slower, irregular ramps), because the plate breaks free only when the error reaches $f_s/K_c \approx 1.9$ mm and the mold integrates the leftover imbalance in between. Half the gain makes the cycle slower and larger, 3.1 mm in amplitude at a mean period of 24 s, and double the gain makes it smaller, 0.67 mm RMS, while the gate travels 71 % of its stroke per minute against 23 %. The knocker at its default 2.0 % takes the RMS to 0.36 mm, and at 2.5 %, half the stick band, to 0.22 mm, at the price of 83 % per minute of gate travel.

🕮 **Pause.** Try Exercise 6.2 before reading on.

> **TRY IT IN THE APP 3 · A Sticking Slide Gate**
>
> Load scenario 3, A sticking slide gate, and press Run. Watch the gate signature plot in the scene for half a minute. A healthy gate would draw a thin line along the dashed diagonal. This one opens into a parallelogram about one stick band wide, while the gate position on the lower trend climbs in steps under a sawtooth command.

Bulging of 4 mm open-loop amplitude at 12.5 s appears as a 2.7 mm amplitude in automatic. Opening the loop with the gate held at the balance stroke increases it to 3.9 mm, which identifies an external disturbance, and raising the speed to 1.5 m/min shortens the period to 10.0 s, exactly roll pitch over the new speed, which identifies bulging. More gain shrinks this oscillation too, which is why the gain alone cannot tell it from stiction. Double the gain gives 1.3 mm, and half the gain

Table 6.3 Scenarios and Completion Requirements

#	Scenario, teaching point and kind	Required to complete
1	Hold the meniscus by hand *the integrator, manual* (experiment)	In Manual, let the level drift at least 1.5 mm from 100 mm, move the gate command to bring it back, then hold it within 1 mm of 100 mm and moving less than 0.05 mm/s for 5 s, and leave the gate alone in Manual for 60 s.
2	A healthy loop *reference trend* (experiment)	After 30 s in Auto, raise the setpoint from 100 to 103 mm and hold the level within 0.5 mm of 103 mm for 10 s.
3	A sticking slide gate *stiction signature* (experiment)	Watch 30 s of the stick-slip cycle at the shipped tuning, run 60 s proportional only ($\tau_I = 60$ s), then 30 s with $\tau_I = 4.6$ s restored and the knocker on.
4	Bulging between the support rolls *open the loop, move ω* (experiment)	Run in Auto until an oscillation period is measured, open the loop with the gate at the balance stroke for 60 s without a level alarm, return to Auto and raise the speed by 0.3 m/min, then run 30 s at 1.5 m/min with the canceller on.
5	Too much gain *the delay ceiling* (demonstration)	Run to 100 s at $K_c = 5.0\,\%/\text{mm}$ through the setpoint step at 20 s, then 30 s at $K_c = 2.5\,\%/\text{mm}$, then 20 s at $K_c = 5.5\,\%/\text{mm}$.
6	All three at once *the order of repairs* (experiment)	Watch 30 s as loaded, then 30 s each at $K_c = 1.7\,\%/\text{mm}$, with the canceller added, and with the knocker added, in that order.
7	Two faults, one trend *the speed-change test* (experiment)	Read the oscillation period before the speed rise at 90 s, run 60 s at 1.6 m/min with both add-ons off, then 30 s with the canceller on and 30 s more with the knocker added.
8	The knocker on the wrong fault *the price of a wrong remedy* (experiment)	Run 60 s with the knocker on, Reset and switch it off, run 60 s without it, then 30 s with the canceller on instead.
9	A clogging nozzle *a leading indicator* (demonstration)	Run until the gate passes 62 % with the level within 1 mm of setpoint, through the clog release at 240 s with the level more than 2 mm high, until it holds within 0.5 mm of 100 mm for 10 s.
10	Challenge *graded run* (challenge)	Keep the true level within ± 1.5 mm of setpoint from 200 to 300 s, through the speed rise from 1.2 to 1.5 m/min at 60 s, with the model parameters and disturbances as given.[a]

[a] With the shipped settings ($K_c = 3.0\,\%/\text{mm}$, no add-ons) the largest deviation over the scored window is 3.34 mm, more than twice the band. With the suggested settings ($K_c = 1.7\,\%/\text{mm}$, $\tau_I = 4.6$ s, knocker at 2.5 %, canceller on) it is 0.77 mm, inside the band with 0.73 mm to spare. The knocker and the canceller together are enough at the shipped gain (1.24 mm with the default 2.0 % knocker), as the challenge text states.

gives 4.8 mm, larger than with the loop open because the loop then amplifies at this frequency. The canceller takes the level back to 0.14 mm RMS, the noise floor of the healthy loop (0.13 mm).

Pause. Try Exercise 6.1 before reading on.

TRY IT IN THE APP 4 · Bulging Between the Support Rolls

Load scenario 4 and run until the Oscillation period tile shows a value near the 12.5 s of the Bulging period at this speed tile. Switch to Manual and press Gate to balance: the swing grows from 2.7 to about 3.9 mm, so the cause lies outside the loop. Back in Auto, press Speed +0.3 m/min, and the period falls to 10.0 s.

Tripling the gain to 5.0 %/mm produces the third signature. A 3 mm setpoint step overshoots by 3.8 mm, and the gate spends most of the next 12 s at its 15 %/s speed limit, so the first swings are 3.8 s apart and shorten to 2.3 s as they decay. The loop then keeps ringing on the sensor noise with a mean period of 2.7 s, far faster than either of the other two, and the gate travels 380 % of its stroke

Table 6.4 Measured Oscillation Signatures

Case and change	Level	Other
3, stiction, SIMC	1.11 mm RMS	period 13.7 s
3, proportional only	1.02 mm RMS	irregular
3, half gain	3.1 mm amplitude	period 24 s
3, double gain	0.67 mm RMS	travel ×3.1
3, knocker 2.0 %	0.36 mm RMS	travel ×1.5
3, knocker 2.5 %	0.22 mm RMS	travel ×3.6
4, bulging, automatic	2.7 mm amplitude	period 12.5 s
4, loop opened[a]	3.9 mm amplitude	external
4, speed 1.5 m/min	—	period 10.0 s
4, double gain	1.3 mm amplitude	—
4, half gain	4.8 mm amplitude	amplified
4, canceller	0.14 mm RMS	noise floor
5, $K_c = 5.0$[b]	overshoot 3.8 mm	ring 2.7 s
5, $K_c = 5.5$[b]	runs away	saturation
8, knocker, no stiction	1.93 to 1.78 mm RMS	travel ×2.1
9, clog release[b]	4.4 mm spike	gate 57.0 to 64.4%

Each change is made before the run. RMS and amplitude ($\sqrt{2}$ times the standard deviation) are of the true level, and they, the mean period and the gate travel ratio are taken from 120 to 480 s unless marked. [a] One minute in manual at the balance stroke from 36 s. [b] Scripted events as loaded (setpoint step at 20 s, clog release at 240 s).

per minute over the minute after the step. Half the gain removes the ringing, with a 1.6 mm overshoot and 0.19 mm RMS at 120 s. The speed limit also turns the stability margin into a cliff. With the gain at 5.1 %/mm or more when the step arrives, the swings grow instead of decaying, and at 5.5 %/mm the gate saturates 9.8 s after the step and the level alarm follows at 37 s. The same 5.5 %/mm set after the step has passed rings with swings of up to 2.4 mm over the next four minutes but does not run away, so where the cliff lies depends on the size of the upset as well as on the gain.

6.5.3 Combined Faults and the Order of Repairs

Scenario 6 combines all three at $K_c = 4.5\,\%/\text{mm}$. The level RMS is 0.97 mm and the gate travels 205 % of its stroke per minute, more than four times the healthy loop. Bringing the gain back to 1.7 %/mm first makes the RMS jump to 3.1 mm, because the high gain was holding the bulging down. The canceller then takes it to 1.29 mm and the knocker to 0.49 mm, with the gate traveling 87 % per minute. Adding the knocker while the gain is still 4.5 %/mm, with or without the canceller, pushes the loop over its stability cliff (29 mm RMS with both, the gate saturated 67 % of the time), so the order of the repairs matters as well as their choice.

🕮 **Pause.** Try Exercise 6.3 before reading on.

Scenario 7 has stiction and bulging together with nearly the same period at 1.2 m/min, which is the case the trend alone cannot resolve. The speed starts to rise toward 1.6 m/min at 90 s and reaches it at 170 s. The bulging then moves to 9.3 s (9.4 s computed), while stiction on its own cycles at 13.7 s at 1.2 m/min and 13.4 s at 1.6 m/min. The canceller alone takes the RMS from 3.7 mm to 1.3 mm over two minutes. It needs 5.0 % of gate motion either way at the higher speed, and that motion has to get through plates that stick by 5 %, so adding the knocker brings the RMS to 0.50 mm.

A remedy applied to a cause that is not present costs wear and buys nothing. The knocker on a healthy gate with bulging present doubles the gate travel, from 104 to 220 % per minute, while the level RMS moves only from 1.93 mm to 1.78 mm. The canceller in its place gives 0.14 mm with 86 % per minute of travel. A clog of 4 % per minute walks the gate from 57.0 % to 64.4 % while the true level stays within 0.41 mm of setpoint, and the process gain drifts from 0.51 to 0.44 mm/s per percent. The release at 240 s spikes the level 4.4 mm within 2 s, and the loop holds it within 0.5 mm again from 248 s. The gate position is the leading indicator of the clog, not the level.

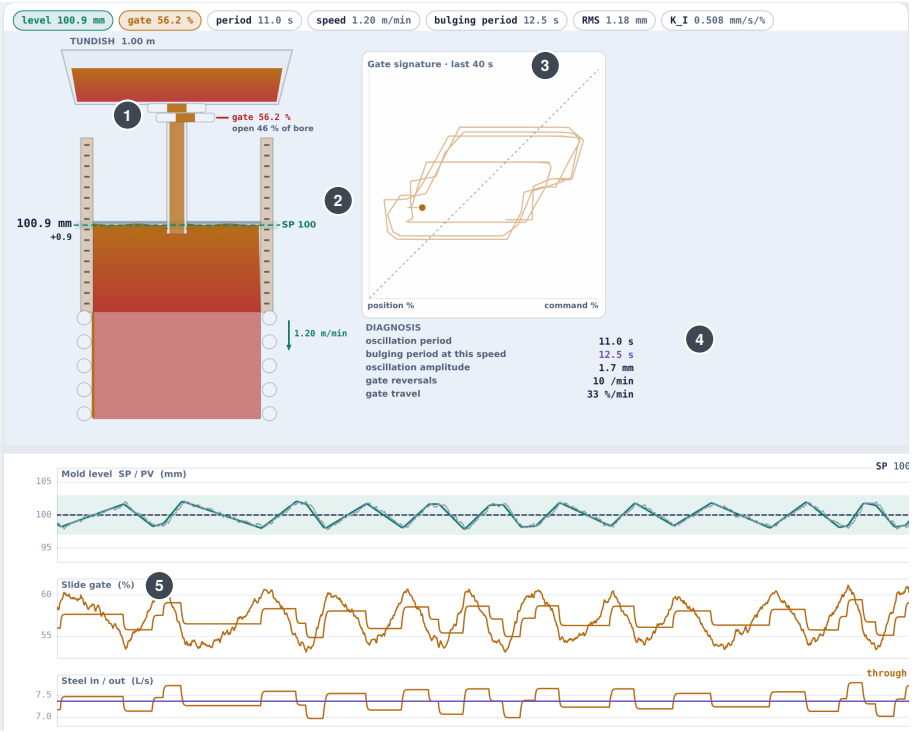


Figure 6.3 Mold level loop with a sticking slide gate, scenario 3, at 190 s with the SIMC tuning in automatic, cropped to the caster view and the trends. The caster section on the left shows the tundish, the slide gate (1) at 56.2 % stroke, the mold with the setpoint and the slag-entrainment band (2), and the strand between its rolls. The gate signature (3) plots controller output against gate position over the last 40 s and has opened into a parallelogram about one stick band wide, and the diagnosis readouts (4) sit below it. The trends show the true level cycling about 4 mm peak to peak with the sensor reading dotted, the gate position (5) moving as a staircase under a sawtooth command, and the gate flow stepping about the constant withdrawal of 7.36 L/s. The level RMS reads 1.18 mm at this instant.

6.5.4 The Graded Challenge

The challenge combines stiction, 4 mm bulging and a gain of 3.0 %/mm, with a speed rise from 1.2 to 1.5 m/min starting at 60 s. It is met when the true level stays within 1.5 mm of setpoint from 200 to 300 s, read from the 10 Hz record. Table 6.5 reports the largest deviation in that window for each combination of remedies, and Figure 6.4 draws the same numbers as a chart.

Retuning alone makes the peak worse, 5.31 against 3.34 mm, for the same reason as scenario 6. Each remedy on its own removes only its own cause, so neither the knocker (3.86 mm) nor the canceller (2.22 mm) meets the target after the retuning or at the shipped gain, and neither passes on its own at any gain tried from 1.7 to 4.0 %/mm. The three together give 0.77 mm with a level RMS of 0.26 mm at the end of the run. The two add-ons also pass without the retuning, at 0.87 mm, because 3.0 %/mm is well below the gain at which the loop of scenario 5 tips over. The retuning buys margin, 0.77 against 0.87 mm, rather than the pass itself, and the application's challenge text says so.

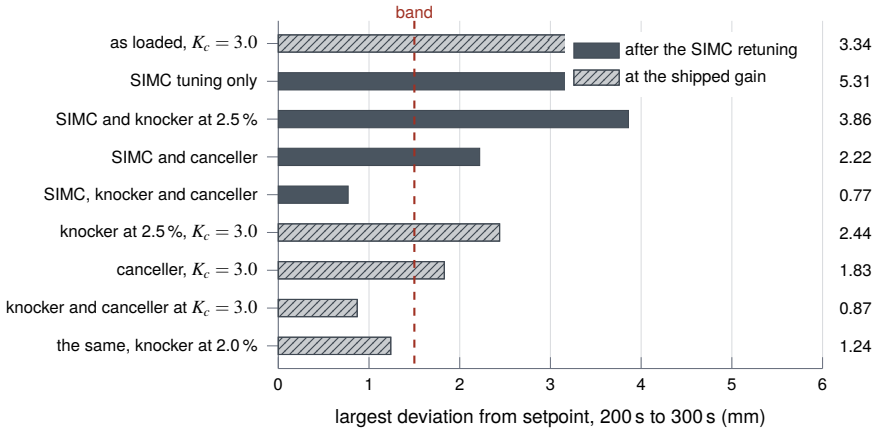


Figure 6.4 The graded challenge of Table 6.5 as a chart: the largest deviation of the true level in the scored window for each combination of remedies, after the SIMC retuning (solid) and at the shipped gain (hatched), against the band of 1.5 mm (dashed). Values are in the column at the right. Notice that retuning alone makes the peak worse than as loaded, that no single remedy passes, and that the knocker and the canceller together pass at either gain.

Table 6.5 Challenge: Largest True Level Deviation From 200 to 300 s

Configuration	Deviation (mm)
as loaded, $K_c = 3.0$	3.34
SIMC tuning only	5.31
SIMC and knocker at 2.5 %	3.86
SIMC and canceller	2.22
SIMC, knocker and canceller	0.77
knocker at 2.5 %, $K_c = 3.0$	2.44
canceller, $K_c = 3.0$	1.83
knocker and canceller at $K_c = 3.0$	0.87
the same, knocker at 2.0 %	1.24

6.6 Technology in Practice

What Measures It

The controlled variable of this chapter is the meniscus level, read by a sensor above the mold. Not every mold level sensor is an eddy-current coil. An electrode-type sensor of 1982 kept its electrode at a constant electrical resistance in the slag layer with a servo drive, over a range of 500 mm with a response time of 1 s [11].

TECHNOLOGY Sensor: The Eddy-Current Mold Level Meter

An eddy-current meter hangs a coil a short distance above the meniscus. The coil carries a high-frequency alternating current that induces eddy currents in nearby metal, and the closer the metal, the stronger the currents and the heavier the load on the oscillator that drives the coil [20]. The mold powder conducts too once it melts, but at a low frequency such as 10 kHz the coil impedance is almost unaffected by the thickness of the molten layer [21]. A non-contact meter of this kind described in 1984 read the steel level through the powder to within 1 mm with a quick response, and a differential coil made it more stable against the position of the mold wall [22]. A detector 30 mm across fitted billet molds.

In the model of Section 6.3 the sensor is the 0.30 s delay, the colored noise and the 0.10 s filter, clamped to a 200 mm span. A real meter also sees the surface waves and jets that the model represents only as band-limited noise.

What Moves It

The flow from the tundish into the mold is controlled by a stopper rod or by a slide gate [23]. Where a stopper is used, the most common measure of clogging in a plant is the rise of the stopper rod position. It climbs sharply when the clog is severe, and the casting speed must then be reduced and the nozzle replaced. The final element is thus also the plant's clog gauge, as the gate position is in scenario 9.

TECHNOLOGY Final Element: The Slide Gate

A slide gate controls the flow by sliding one refractory plate across another, so that their holes overlap more or less [18]. The sliding faces are polished very precisely so that steel cannot leak between them, and the plates must keep their strength when hot and resist abrasion and corrosion, and a hydraulic cylinder moves the plate. A plate that sticks jumps to a new position only when enough force has built up, and on an integrating process the measurement then changes slope at every slip [20]. Deposits in the nozzle below act on the same flow: slight clogging may at first raise the flow rate, but heavy clogging restricts the channel [23].

The simulated gate is the lens of two 70 mm bores, moved by a cylinder with a 0.15 s lag and a 15 %/s speed limit, and its sticking is the stick band S and slip jump J of Equation (6.5). Figure 6.5 draws the loop with these devices in place.

Where the Controller Runs

A loop this fast runs in a local controller. Programmable logic controllers serve as local controllers within the supervisory scheme of a distributed control system and carry proportional integral derivative (PID) control among their functions, and control loops run with cycle times from milliseconds to minutes [24]. The simulated controller executes every 0.05 s. Its velocity-form PI block, the knocker added after the clamp and the canceller feedforward would share one fast task. Records kept for analysis are often slower than the loop: one forecasting study used plant data sampled at 1 Hz [25].

IN PRACTICE Check the Gate Before You Retune

Before any tuning, a field handbook advises checking instrument health with small up-and-down steps of the output in manual, after asking operations how far the output may move and the measurement may drift [20]. The steps show how much stiction the final element has. The same text warns that a sticky valve cannot be fully overcome by tuning: on an integrating process, less proportional gain lowers the frequency of the cycle but does not end it. A study of ten refinery loops adds the cost argument [1]. Special tests cost labor and lost production, and where the signature points to a valve nonlinearity there is little benefit in retuning. The valve is better examined at a scheduled shutdown and replaced if necessary.

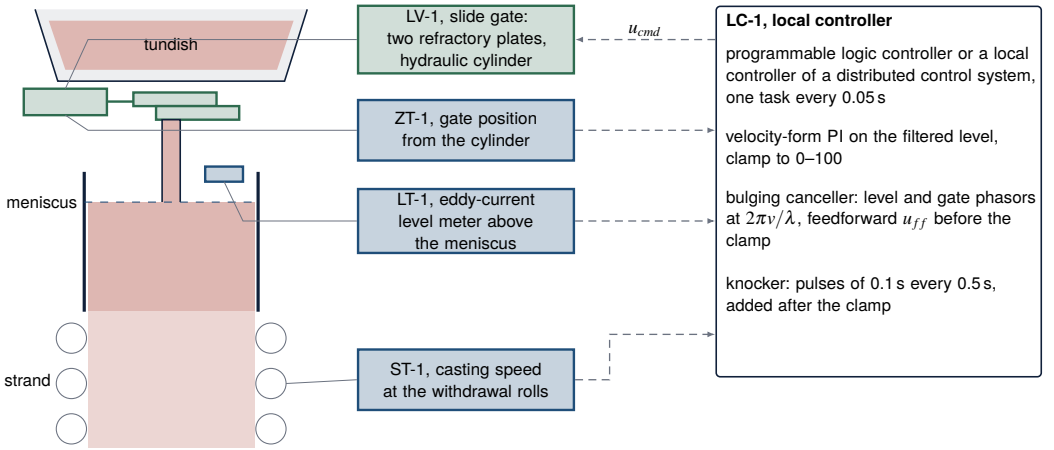


Figure 6.5 The mold level loop drawn with its hardware. The eddy-current meter LT-1 reads the meniscus, ZT-1 the gate position and ST-1 the casting speed, and the local controller LC-1 sends one gate command to the slide gate LV-1. Notice that the canceller and the knocker need no extra hardware: both live in the controller task, and the one gate command carries the sum of all three.

What Is Modernizing It

Loop performance assessment computes figures of merit from routine data without taking loops off line, and in 1997 it was becoming widely used in the petrochemical, chemical and paper industries [1]. Its oscillation signatures point the engineer to the special test that is most likely to confirm the cause, the approach of Volume 1, Chapter 13. Learning-based forecasting is a second development. A sequence model trained on plant records from 13 casts, 2.2 million points in all, predicts the next ten one-second samples of the mold level fluctuation, and in its tests the mean absolute error fell from 0.517 to 0.150 against a benchmark of the same kind trained with a plain squared-error loss [25]. A third acts on the flow itself. Superimposed alternating and direct magnetic fields in the mold roughly halved the surface defects on steel sheet in trials on an industrial caster, compared with a direct field alone [26]. The present model has no electromagnetic brake, as Section 6.7 states.

6.7 Limitations

The flat meniscus leaves out the surface waves and jet effects that dominate real mold level signals [9]. Their absence makes the signature plot cleaner than it is on a plant, where the same test is run on a noisier signal with more averaging.

Bulging is one tone. Real unsteady bulging carries several frequencies from roll segments of different pitch, and a canceller built for one of them leaves the others [4, 15].

> ⚠ **OPEN CHALLENGE** **Bulging at High Casting Speed**
>
> Mold level hunching from unsteady bulging has been called a main obstacle to raising the casting speed of thin slab casters, which cast strands 50 to 100 mm thick at about three times the speed of conventional slab casters [27]. In 2002 its origin and detailed mechanism were still not clearly known. Casters suppress it with more secondary cooling water and roll pitches that do not repeat [4], and the roll gap on a running caster is hard to measure directly. Progress would be a roll layout and a cooling practice that keep the bulging below what the level loop can reject at the highest speed.

The stiction model is data-driven rather than physical, so it reproduces the signature of a sticking gate but not the effect of temperature and refractory wear on the stick band [7].

The oscillation period reading wanders on an irregular stick-slip cycle, between 10.9 and 16.6 s over the first eight minutes of scenario 3, because it averages the three to five upward crossings of the last minute while single cycles range from 9.8 to 18.2 s. The level RMS and the gate travel are steadier, and the stick-slip periods quoted here are means over six minutes.

There is no stopper rod, no argon control and no electromagnetic braking, all of which interact with the level loop on modern casters.

OPEN CHALLENGE **Clogging Nozzles in New Steel Grades**

Alumina clogging was still a major issue in slab casting in 2000 [5], and a 2024 review finds it unsolved for two growing families of steel [23]. High aluminum steels for lighter cars change the inclusions and the reaction with the nozzle refractory, and rare earth treated steels form solid inclusions that are hard to remove. Most remedies in use touch the steel, such as nozzle materials and coatings, and can affect its cleanliness. The review points to non-contact treatment by external magnetic or electric fields, and to several methods coordinated together.

6.8 Conclusion

A mold level loop that oscillates can be oscillating for three unrelated reasons, and the simulator of this chapter makes each of them identifiable with the test that plant engineers use. A sticking gate with a 5 % stick band cycles at 1.11 mm RMS with integral action and 1.02 mm without it and draws a parallelogram on the command against position plot, and a knocker at half the stick band takes it to 0.22 mm. Bulging survives opening the loop, growing from 2.7 mm to 3.9 mm, and moves with the casting speed exactly as roll pitch over speed predicts, and a canceller at that one frequency returns the level to the 0.14 mm noise floor. Too much gain rings at about 2.7 s and is cured by halving it. Applied to the wrong cause, each remedy costs gate travel and fixes nothing, and in the challenge no single remedy meets the target while the three together bring the largest deviation from 3.34 mm to 0.77 mm. The model leaves out surface waves, multiple bulging frequencies and the physics behind the stick band. A fuller model would add a stopper rod, multiple roll pitches, electromagnetic braking and an automated stiction test. The caster is an accessible browser case for diagnosing an oscillation before treating it.

KEY POINTS

- The level in the mold can swing because the gate sticks, the strand bulges or the controller pushes too hard, and the three swings look alike.
- A sticking gate cycles even without integral action, and a knocker at half the stick band takes the level to 0.22 mm RMS.
- Bulging survives opening the loop and moves with the casting speed, and a canceller at its frequency returns the level to the 0.14 mm noise floor.
- A remedy on the wrong cause costs gate travel and fixes nothing. In the challenge only all three remedies together bring the peak from 3.34 to 0.77 mm.

Where else this applies. A sticking steam valve in Volume 2, Chapter 4 and a disturbance whose frequency the speed sets in Volume 7, Chapter 11 call for the same diagnosis before treatment.

References

[1] N. Thornhill and T. Hägglund, “Detection and diagnosis of oscillation in control loops,” *Control Engineering Practice*, vol. 5, no. 10, pp. 1343–1354, 1997.

[2] M. Jelali and B. Huang, Eds., *Detection and Diagnosis of Stiction in Control Loops*. Springer London, 2010.

[3] A. Horch and A. J. Isaksson, “Detection of valve stiction in integrating processes,” in *2001 European Control Conference (ECC)*, 2001, pp. 1327–1332.

[4] H. Ohno, Y. Miki, and Y. Nishizawa, "Generation Mechanism of Unsteady Bulging in Continuous Casting-1 - Development of Method for Measurement of Unsteady Bulging in Continuous Casting-," *ISIJ International*, vol. 56, no. 10, pp. 1758–1763, 2016.

[5] J. Stubbles, "Energy use in the U.S. steel industry: A historical perspective and future opportunities," U.S. Department of Energy, Office of Industrial Technologies, Tech. Rep., 2000.

[6] A. Horch, "A simple method for detection of stiction in control valves," *Control Engineering Practice*, vol. 7, no. 10, pp. 1221–1231, 1999.

[7] M. Shoukat Choudhury, N. Thornhill, and S. Shah, "Modelling valve stiction," *Control Engineering Practice*, vol. 13, no. 5, pp. 641–658, 2005.

[8] R. I. L. Guthrie and M. M. Isac, "Continuous Casting Practices for Steel: Past, Present and Future," *Metals*, vol. 12, no. 5, p. 862, 2022.

[9] B. G. Thomas, "Modeling of the continuous casting of steel—past, present, and future," *Metallurgical and Materials Transactions B*, vol. 33, no. 6, pp. 795–812, 2002.

[10] J. K. Brimacombe, "Empowerment with knowledge— toward the intelligent mold for the continuous casting of steel billets," *Metallurgical Transactions B*, vol. 24, no. 6, pp. 917–935, 1993.

[11] H. Imada and Y. Nimura, "A new electrode-type mold level sensor for continuous casting machine," *Transactions of the Iron and Steel Institute of Japan*, vol. 22, no. 12, pp. 991–996, 1982.

[12] "Technology of automatic mold bath level control," *Transactions of the Iron and Steel Institute of Japan*, vol. 24, no. 9, p. 777, 1984, new Technology note.

[13] T. Kurokawa, T. Kondo, T. Mita, K. Liu, and M. Sampei, "Development of mold level control in continuous casting by H_∞ control theory," in *Proceedings of IEEE International Conference on Control and Applications*, 1993, pp. 865–871.

[14] M. Dussud, S. Galichet, and L. Foulloy, "Fuzzy Supervision for Continuous Casting Mold Level Control," *IFAC Proceedings Volumes*, vol. 33, no. 17, pp. 1137–1142, 2000.

[15] K. Jabri, E. Godoy, D. Dumur, A. Mouchette, and B. Bèle, "Cancellation of Bulging Effect on Mold Level in Continuous Casting: Experimental Validation," *IFAC Proceedings Volumes*, vol. 42, no. 23, pp. 191–196, 2009.

[16] T. Hägglund, "A friction compensator for pneumatic control valves," *Journal of Process Control*, vol. 12, no. 8, pp. 897–904, 2002.

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

[18] M. Ohnishi, T. Nishio, and T. Terashima, "Utilization of sliding gate nozzle in Japan," *Transactions of the Iron and Steel Institute of Japan*, vol. 16, no. 3, pp. 164–167, 1976.

[19] "Minimizing fluctuation of mold bath level," *Transactions of the Iron and Steel Institute of Japan*, vol. 26, no. 11, p. 1010, 1986, new Technology note.

[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] Y. Nakamori, O. Ichiko, Y. Mimura, K. Yamaguchi, and M. Ota, "Measurement of molten powder thickness in continuous casting mold by eddy current method," *Transactions of the Iron and Steel Institute of Japan*, vol. 25, no. 1, pp. 50–53, 1985.

[22] K. Sano, S. Ando, Y. Kawase, I. Takahashi, M. Takenaka, and K. Tsuboi, "The eddy-current type continuous casting mold level meter," *Tetsu-to-Hagané*, vol. 70, no. 9, pp. 1160–1166, 1984, in Japanese with an English synopsis.

[23] W. Yang, L. Zhang, Y. Ren, W. Chen, and F. Liu, "Formation and prevention of nozzle clogging during the continuous casting of steels: A review," *ISIJ International*, vol. 64, no. 1, pp. 1–20, 2024.

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

[25] X. Liu, D. Han, Z. Li, H. Wang, W. Li, M. Fu, and X. Wang, "Time series forecasting model of mold level fluctuations integrating multi-scale feature decomposition and trend flexible alignment," *ISIJ International*, vol. 66, no. 7, pp. 867–875, 2026.

[26] Y. Miki, A. Matsui, H. Ohno, and N. Aramaki, "New electromagnetic flow control system for optimization of molten steel flow in continuous casting mold," *ISIJ International*, vol. 65, no. 8, pp. 1162–1171, 2025.

[27] U.-S. Yoon, I.-W. Bang, J. H. Rhee, S.-Y. Kim, J.-D. Lee, and K. H. Oh, "Analysis of mold level hunching by unsteady bulging during thin slab casting," *ISIJ International*, vol. 42, no. 10, pp. 1103–1111, 2002.

6.9 Exercises

6.1 Concept check · 2 min

In scenario 4, Bulging between the support rolls, the true level swings with an amplitude of 2.7 mm in automatic at the SIMC tuning. With the loop switched to Manual and the gate held at the balance stroke, the swing grows to 3.9 mm. What does this test show?

(a) The loop makes the oscillation itself, and opening it removed the damping the controller supplied

(b) The oscillation enters from outside the loop, and the controller had been taking part of it off

(c) The gate plates stick, because a held gate cannot move and the mold integrates whatever imbalance is left

(d) The sensor delay sets the period, so the swing in Manual is the true one and the swing in Auto is filtered

6.2 Concept check · 1 min

In scenario 3, A sticking slide gate, the integral time is set to 60 s, which removes integral action. The level still cycles, at 1.02 mm RMS against 1.11 mm with the SIMC tuning. Why does the cycle survive?

(a) The knocker keeps adding pulses that break the plate free at regular intervals

(b) The proportional gain of 1.7 %/mm lies above the ultimate gain of the loop

(c) The mold integrates the imbalance while the plate is stuck, until the error breaks it free

(d) The setting of 60 s still leaves a slow integral that winds up between the slips

6.3 Concept check · 2 min

Scenario 6, All three at once, ships at $K_c = 4.5\,\%/\mathrm{mm}$ with a sticking gate and 4 mm of bulging, and the level RMS is 0.97 mm. Lowering the gain to 1.7 %/mm first raises the RMS to 3.1 mm. What is the right reading of that rise?

(a) The SIMC tuning does not suit this plant, so the gain should go back to 4.5 %/mm

(b) The lower gain has enlarged the stiction cycle, which shows that stiction is the main cause

(c) The lower gain has brought the loop to its stability limit, where it amplifies every upset

(d) The high gain was holding the bulging down, and the retuning has unmasked it for the canceller

6.4 Theory and calculation · 5 min

The mold has the integrating gain $K_I = 0.51\,\mathrm{mm/s}$ per percent of stroke. The measurement delay is 0.30 s, the sensor filter 0.10 s, the actuator lag 0.15 s and the controller scan 0.05 s.

(a) Find the apparent dead time θ_a, counting half a scan, to two decimals. *Checkpoint:* 0.58 s.

(b) With $\tau_c = \theta_a$, find the SIMC gain and integral time of Equation (6.11). *Checkpoint:* $K_c \approx 1.7\,\%/\mathrm{mm}$, $\tau_I \approx 4.6\,\mathrm{s}$.

(c) Derive the ultimate gain and period of an integrator with dead time θ_a, evaluate them, and find the gain margin of the SIMC loop. *Checkpoint:* 5.3 %/mm, 2.3 s, a margin of about 3.1.

6.5 Theory and calculation · 6 min

The gate of scenario 3 has $S = 5\,\%$ and $J = 1.5\,\%$ and starts stuck at 57.6 % of stroke, where the balance stroke is 57.0 %. Each percent of stroke is worth 0.18 L/s, the free surface is $A_m = 0.353\,\mathrm{m}^2$ and $K_c = 1.7\,\%/\mathrm{mm}$.

(a) Find the static and kinetic bands f_s and f_d, and show that every breakaway ends in a jump of J. *Checkpoint:* $f_s = 3.25\,\%$, $f_d = 1.75\,\%$, jump 1.5 %.

(b) Find the surplus flow while the plate stays stuck and the rate at which the level rises. *Checkpoint:* 0.11 L/s and 0.31 mm/s.

(c) With proportional action only and the level starting at setpoint, find the measured error at which the plate breaks free and estimate when that happens. *Checkpoint:* 1.9 mm, after about 6 s plus the sensor delay.

6.6 Theory and calculation · 5 min

Bulging of open-loop amplitude $a_b = 4$ mm acts on a mold with $A_m = 0.353\,\mathrm{m}^2$ between rolls of pitch $\lambda = 0.25$ m. At 1.2 m/min the gate gives 0.180 L/s per percent of stroke, and at 1.6 m/min, where it sits further open, 0.190 L/s per percent.

(a) Find the bulging period and angular frequency at each speed. *Checkpoint:* 12.5 s and 0.503 rad/s; 9.4 s and 0.670 rad/s.

(b) Find the amplitude of the bulging flow Q_b at each speed. *Checkpoint:* 0.71 L/s and 0.95 L/s.

(c) Find the gate motion a canceller must make to cancel it at each speed, and compare it with the stick band of the gate in scenario 7. *Checkpoint:* 3.9 % and 5.0 %, against $S = 5\,\%$.

6.7 App challenge · 9 min

Load scenario 3, A sticking slide gate, at the default seed. Make three runs to 480 s at speed 10 (48 s of wall time each), reloading the case with the reload button before each and making the change before Run: (i) as loaded, (ii) Knocker on, (iii) Knocker on with Knocker amplitude at 2.5 % (Knocker settings). Export the CSV file after each run. Use the rows from 120 to 480 s.

(a) Compute the RMS of `level_error_mm`, the true level minus setpoint, for each run, within 0.02 mm. *Checkpoint:* 1.11, 0.36 and 0.22 mm.

(b) Compute the gate travel per minute as the sum of the absolute changes of `gate_pos_pct` divided by 6 min, within 2 %/min, and its ratio to run (i). *Checkpoint:* 23, 34 and 83 %/min; ratios 1.5 and 3.6.

Which run would you keep on a plant until the plates are changed, and why is 2.5 % called half the stick band?

6.8 App challenge · 8 min

Load scenario 4, Bulging between the support rolls, at the default seed, and run at speed 5. Read each tile within 0.2 s.

(a) Run to about 120 s as loaded and read the Oscillation period and the Bulging period at this speed tiles. *Checkpoint:* 12.5 s and 12.5 s.

(b) Reload the case, pause at about 60 s, press Speed +0.3 m/min, run on to about 180 s and read the same two tiles. *Checkpoint:* 10.0 s and 10.0 s.

(c) Reload the case, set Proportional gain to 3.4 %/mm before Run, run to 480 s at speed 10 and export the CSV file. Compute the amplitude, $\sqrt{2}$ times the standard deviation of `level_error_mm` from 120 to 480 s, within 0.05 mm. Repeat at 0.85 %/mm. *Checkpoint:* 1.27 mm and 4.81 mm.

Which result names the cause, and why can the gain alone not tell bulging from stiction?

6.9 App challenge · 9 min

Load the challenge, A speed change on a bad day, at the default seed. Run each attempt the full 300 s at speed 10 and read the result the app gives at 300 s, the worst level error from 200 to 300 s (the largest $|$`level_error_mm`$|$ from 200 s in the CSV file gives the same value), within 0.02 mm.

(a) Run the case as loaded. *Checkpoint:* 3.34 mm, a fail.

(b) Press Reset, switch the Knocker and the Bulging canceller on with Proportional gain left at 3.0 %/mm, and run again. *Checkpoint:* 1.24 mm, a pass.

(c) Press the reload button, set Proportional gain to 1.7 %/mm with both add-ons off, and run again. *Checkpoint:* 5.31 mm, a fail.

Explain why the retuning alone makes the result worse, and what the retuning adds once both add-ons are on.

6.10 Application area · 4 min

In scenario 9, A clogging nozzle, the clog grows at 4 % per minute until it breaks away at 240 s. The withdrawal is 7.36 L/s and $A_m = 0.353\,\mathrm{m}^2$.

(a) Find the clog fraction at the release and the share of extra steel the same gate stroke passes once the clog is gone. *Checkpoint:* 0.16 and 19 %.

(b) Find the extra flow and the rate at which the level starts to rise, and compare the measured 4.4 mm spike with the ± 5 mm slag-entrainment band. *Checkpoint:* 1.40 L/s, 4.0 mm/s; inside the band by 0.6 mm.

What should the operator trend to see the clog coming, and why does the level trend not show it?

6.11 Application area · 3 min

The model's bulging is one tone at one roll pitch. Suppose a strand has segments of 0.25 m and 0.30 m pitch.

(a) Find the two bulging periods at 1.2 m/min and at 1.5 m/min. *Checkpoint:* 12.5 and 15 s; 10 and 12 s.

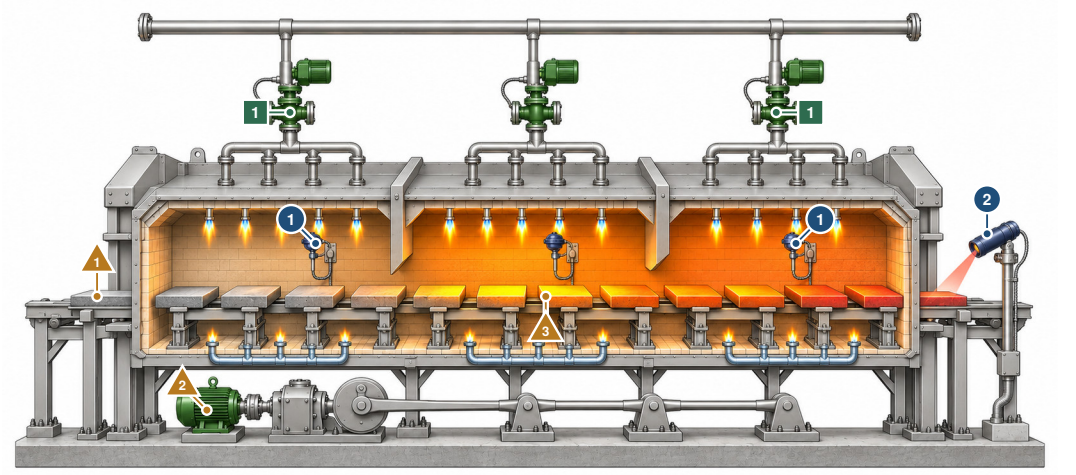
(b) What would the speed-change test show, and what would a canceller locked to the 0.25 m frequency leave in the level? *Checkpoint:* Both periods shorten by the same ratio. The 0.30 m component is left.

6.12 Review the agent · 5 min

A process engineer asks an AI assistant about scenario 4. The assistant replies: "The level swings 2.7 mm at $K_c = 1.7\,\%/\mathrm{mm}$, and doubling the gain cut it to 1.3 mm. So the bulging is a tuning problem. Raise K_c to 4.5 %/mm and the swing will fall below 1 mm with no new equipment." Find the flaw, and say how you would check the claim in the app before acting on it.

CHAPTER 7

Walking-Beam Reheat Furnace Simulator for Slab Tracking and Level 2 Setpoint Control



The controller moves 1 fuel and air valves of the three zones

The controller reads 1 zone thermocouples 2 pyrometer at the discharge door

What disturbs it 1 temperature of the charged slabs 2 mill pace and mill delays, through the walking-beam drive 3 a fall in heat transfer to the slabs

Slabs walk from the charge door at the left to the mill at the right while each zone fires from above and below. Notice that the sensors read the zone gas and the slab surface, never the inside of a slab. Illustration, not to scale.

IN PLAIN TERMS **Heating Steel That Nobody Can See Inside**

The main idea. A reheat furnace heats thick slabs of steel for a rolling mill, and each slab must leave hot all the way through. A computer model follows every slab through three fired zones and sets how hot each zone runs.

The challenge. Heat soaks into a slab slowly, so a slab that leaves sooner leaves colder even when every zone holds its heat. The only instrument aimed at a slab sees its skin, never its middle, and that reading must keep the model honest.

Why it matters. A cold slab rolls badly, and extra heat wastes fuel in a furnace that burns around the clock.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 7.2 and Section 7.4, the fact box of the Introduction, the fact and try-it boxes of Section 7.5, and the key points
Study	40 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 7.5, Section 7.6, Section 7.7 with its two open challenges

The control question A reheat furnace must deliver every slab at a mean temperature that no instrument reads, while the mill changes its pace, stops for a roll change and sends hot slabs. Which source of zone setpoints puts the slabs on target, and what keeps a slab-tracking model honest when its heat transfer is wrong?

Learning objectives

- Compute the conduction time, Biot number, residence time and throughput of the furnace from its parameters.
- Explain why zone loops on setpoint and a loop on the slab surface both discharge cold slabs after a pace change.
- Describe how a slab-tracking model turns each slab's predicted discharge mean into zone setpoints and uses the charge temperature as feedforward.
- Calculate a correction of the exchange factor from a pyrometer residual, and weigh the correction gain against the pyrometer noise.
- Compare the slab model with a joint predictive profile on accuracy, fuel and a mill delay.

Prerequisites Sections 1.4.2 (slab conduction and the three clocks), 1.4.5 (material tracking), 1.6.2 (velocity form) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive option.

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

See also Volume 5, Chapter 11 for a multi-zone kiln whose brick core is never measured, and Volume 5, Chapter 5 for a moving ribbon that carries its temperature through the half thickness. Volume 9, Chapter 4 for an offset-free predictive controller whose estimator carries a bias.

7.1 Introduction

Feedback acts on the variable it measures. When the measured variable and the product quality are the same thing, that is enough, but on many plants the quality lives inside the product and the sensors see only its surroundings or its skin. A kiln measures its gas and not the clinker, a dryer its air and not the moisture at the center of the web, and a batch reactor its jacket and not the conversion. The multi-zone kiln of Volume 5, Chapter 11 has the same difficulty with the core of a brick. The general answer is a model that runs alongside the plant, driven by the measured inputs and corrected by whatever the plant does measure, which is the structure formalized in the linear filtering problem [1] and extended to unknown model parameters by treating them as additional states [2].

The slab reheat furnace in front of a hot strip or plate mill, such as the mill of Chapter 8, is an unusually clear case. Slabs are charged cold or warm, walked through several fired zones for a couple of hours, and discharged to the mill at a mean temperature that decides whether they roll well, and neither the mean nor the core is measured on line. The furnace is fired around the clock, so every degree of overheating has a price in fuel, and the energy consumption of reheating furnaces under different charging conditions and heating patterns remains an active research subject [3, 4]. A transfer function from zone setpoint to slab temperature hides the part that matters, because the slab response depends on how long each slab stays and where it is when a change occurs, and those quantities move with the mill schedule.

> ✦ **DID YOU KNOW? A Tenth of a Steelworks' Energy**
>
> In Japan around 1980, reheating furnaces burned nearly 10 percent of the total energy of an integrated steelworks and 30 to 70 percent at works without ironmaking [5]. Continuous slab reheating furnaces then ran at about 350 thousand kilocalories per tonne of steel.

Three misconceptions shaped the simulator. One is that zone temperatures held on setpoint deliver slabs on target, when a faster mill shortens every slab's stay. A second treats the surface temperature read by the discharge pyrometer as a measure of the slab, when a hotter furnace widens the gap between surface and center. A third holds that an estimator which agrees with itself is right, when

a change in radiant exchange leaves a slab-tracking model confident and wrong. The same pitfalls appear wherever a moving product carries an internal state, and the answer is the general state and parameter estimator [1, 2].

This chapter works with a walking-beam furnace simulator in which the zone setpoints come from the operator, from a pyrometer trim loop, from a slab-tracking model that can be corrected from the pyrometer, or from a joint predictive profile, with each choice judged on the discharged slabs and the fuel. Slab tracking, model-based setpoint calculation and exchange factor correction are established methods, cited below, and the chapter specifies the furnace model, the tracking model with its sensitivity shadow and the setpoint calculation in enough detail to be rebuilt. The scenarios separate what Level 1 holds, what a proxy measurement can steer, what a model can plan and what only the measurement can correct. Section 7.2 traces furnace control from zone loops to slab tracking. Section 7.3 gives the furnace and slab model, and Section 7.4 the two control levels, the tracking model, its correction and the predictive option. Section 7.5 reports what each scenario measured, Section 7.6 describes the real instruments and control platforms, Section 7.7 states the limitations, and Section 7.8 concludes.

7.2 From Zone Loops to Slab Tracking

A continuous reheating furnace is a long refractory chamber through which slabs move from a charging door to a discharging door. In a walking-beam furnace movable beams lift the slabs, carry them one step and set them down on fixed beams, which allows heating from below as well as above. Burners are grouped into zones along the length, conventionally a preheating zone near the charge door, one or more heating zones, and a soaking zone near the discharge door in which the surface is held while the center catches up. Detailed heat transfer models of reheating furnaces resolve the radiative exchange inside the chamber and the transient conduction inside each slab [6, 7].

> ⧖ HISTORY **The Walking Beam Replaces the Pusher (1967)**
>
> Until the mid-1960s the furnaces that heated large steel for hot rolling were, with few exceptions, multi-zone pusher furnaces, and walking hearths served only special or small furnaces [8]. A large pusher needs a strong pushing mechanism to drive the slabs along the skids. The slabs sliding on skids and hearth were scored on their undersides, the deeper scores survived as defects in the rolled product, and piling and buckling of slabs limited the length of the furnace. In June 1967 a Japanese steelworks began to run four large walking-beam furnaces. With no soaking hearth, each slab was heated over its whole surface from charge to discharge, and a walking-beam furnace was said to need three quarters of the length of a pusher of the same capacity.

The zone temperature loop is an ordinary loop from a thermocouple to the fuel and air of one zone, and it leaves open the question of what temperature each zone should hold. By 1970 that question was being posed as an optimization of the heating of a slab through a reheating furnace [9]. Two decades later integrated control systems for the reheating furnace of a hot strip mill were being reported [10], and the optimal heating of slabs was formulated for control purposes [11]. Because the quantity to be optimized is the unmeasured slab temperature, estimation arrived with it, and the extended Kalman filter was applied to the average slab temperatures of a simulated reheat furnace by 1990 [12].

Model predictive control followed the same path from zone temperatures to slab temperatures. Identification and predictive control of a slab reheating furnace were reported in 2002 [13], a nonlinear predictive controller was developed for a continuous slab reheating furnace [14], and a nonlinear predictive controller based on a numerical model was applied to a walking-beam furnace [15]. The operating events that such controllers exist to handle have their own literature. Mixed loading and mill delays are treated as a rolling-horizon optimization in walking-beam furnaces [16], and the effect of charging conditions on energy consumption and slab heating quality has been assessed on an industrial-scale furnace [4].

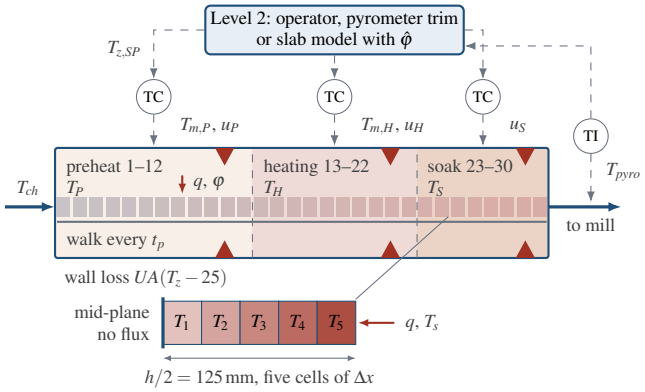


Figure 7.1 Walking-beam reheat furnace as modeled. Thirty slabs charged at T_{ch} walk one position every t_p through the preheat, heating and soak zones at temperatures T_P, T_H and T_S, fired from above and below at u_z percent of capacity and losing heat through the walls. Each slab receives the radiant flux q with exchange factor φ, and half of its thickness is five conduction cells T_1 to T_5 from the mid-plane to the surface temperature T_s (inset). The zone controllers TC read the filtered thermocouples $T_{m,z}$ and follow setpoints $T_{z,SP}$ from Level 2, which also receives the pyrometer reading T_{pyro} of each discharged slab.

Table 7.1 Variables, Their Roles and Default Values

Symbol	Role	Default	Range
$\bar{T}_{dis}$	CV, discharge mean, unmeasured	1200 °C	SP 1150–1260
T_z	CV (Level 1), zone temperature	1150, 1320, 1230	SP 900–1380
$T_{m,z}$	measured, zone thermocouple	T_z	noise 2 °C
T_{pyro}	measured, discharge surface	—	noise 5 °C
u_z	MV, zone firing	holding	0–100 %
$T_{z,SP}$	MV of Level 2, zone setpoint	nominal	800–1350
$T_{i,j}$	states, slab cells (30×5)	periodic	—
E	state, fuel fired	0	—
t_p	DV, mill pace	300 s	180–600
T_{ch}	DV, charge temperature	25 °C	25–800
$\Delta\varphi$	DV, exchange factor change	0 %	±30
—	DV, mill delay length	60 min	5–120

The weakest number in any of these models is the radiant exchange between the furnace and the slab. It is commonly lumped into a total heat exchange factor and estimated, either by solving an inverse heat conduction problem from measured temperatures [17] or from trials in which a slab instrumented with thermocouples travels through the furnace [18]. A recent review of reheating furnaces covers the mathematical models, heating patterns, control systems and energy analysis together [3]. The simulator follows that lineage in miniature, with a one-dimensional slab, one exchange factor, a tracking model and a correction from the discharge pyrometer.

7.3 Furnace and Slab Model

Thirty slab positions pass through three fired zones, preheat (positions 1 to 12), heating (13 to 22) and soak (23 to 30), as drawn in Figure 7.1. Every slab is 10 m long, 1.5 m wide and 250 mm thick, 29.25 t of steel, heated over 30 m^2 of top and bottom faces. At the nominal pace of one walk every $t_p = 300$ s the furnace discharges 351 t an hour and each slab stays 2.5 h, 60 min in preheat, 50 min in heating and 40 min in soak. The mill wants a discharge mean of 1200 °C.

Table 7.1 gives the role and range of each variable. The plant carries 154 states, five cell temperatures for each of the thirty positions, three zone temperatures and a fuel accumulator. The walking

beam, the instruments and the two tracking models run in the controller at every step.

7.3.1 Conduction Through the Slab

Each slab is heated alike from above and below, so half of its thickness h is modeled with $N_c = 5$ cells of width $\Delta x = h/(2N_c) = 25\,\mathrm{mm}$. Cell j exchanges heat with its neighbors by conduction, as shown in Equation (7.1).

$$\rho c_p \Delta x \frac{\mathrm{d}T_j}{\mathrm{d}t} = \frac{k}{\Delta x}\left(T_{j-1} - T_j\right) + \frac{k}{\Delta x}\left(T_{j+1} - T_j\right) \tag{7.1}$$

where $k = 30\,\mathrm{W/(m\,K)}$ and $\rho c_p = 5.07\,\mathrm{MJ/(m^3\,K)}$. The first term is zero for the cell at the mid-plane ($j = 1$), by symmetry, and for the surface cell ($j = 5$) the second term is replaced by the radiant flux q of Equation (7.2).

$$\begin{aligned} q &= \varphi\sigma\left[(T_z + 273.15)^4 - (T_s + 273.15)^4\right] \\ T_s &= T_5 + \frac{q\Delta x}{2k} \end{aligned} \tag{7.2}$$

where $\sigma = 5.670 \times 10^{-8}\,\mathrm{W/(m^2\,K^4)}$, T_z is the temperature of the zone over the slab and φ is the overall exchange factor, 0.6 by default. The surface lies half a cell outside the center of cell 5, and the implicit pair is solved by five fixed-point iterations starting from $T_s = T_5$. The slab mean is the average of the five cells and the core is T_1.

Two dimensionless numbers set the character of the slab. The radiant heat transfer coefficient $4\varphi\sigma T^3$ of Equation (1.2), where Chapter 1 writes σ_{SB} for σ and ℓ for $h/2$, is about $460\,\mathrm{W/(m^2\,K)}$ at 1230 °C, so the Biot number on the half thickness is about 1.9 and the surface runs well ahead of the center. The conduction time $(h/2)^2\rho c_p/k$ is about 44 min, comparable to the time a slab spends in any one zone, so a slab moves on well before it equilibrates with any zone.

7.3.2 Zone Heat Balance and Fuel

Each zone is a single lumped temperature, fired at u_z percent of its burner capacity P_z and loaded by the slabs under it, as shown in Equation (7.3).

$$C_z \frac{\mathrm{d}T_z}{\mathrm{d}t} = \eta(T_z)\,P_z \frac{u_z}{100} - A\sum_{i \in z} q_i - UA\,(T_z - 25) \tag{7.3}$$

with $C_z = 100\,\mathrm{MJ/K}$ for every zone, $P_z = 100$, 90 and 30 MW for preheat, heating and soak, $A = 30\,\mathrm{m^2}$ per slab, $UA = 0.003\,\mathrm{MW/K}$ and the flux sum converted to megawatts. The available-heat fraction of Equation (7.4) falls linearly as the zone runs hotter, standing in for the larger flue gas loss of a hotter zone.

$$\eta(T_z) = \mathrm{clamp}\left[0.85 - 0.0004\,(T_z - 900),\ 0.4,\ 0.9\right] \tag{7.4}$$

The fuel accumulator integrates the gross firing, $\dot{E} = \sum_z P_z u_z/100$, and the fuel per tonne is E divided by the tonnes discharged. Linearizing Equation (7.3) about the nominal state with the slabs frozen gives zone time constants of 10.1, 10.0 and 16.5 min for preheat, heating and soak. At that state the zones fire 61.5, 46.9 and 9.6 MW, 117.9 MW in all or about 1.21 GJ/t, while heating steel from 25 to 1199 °C takes 0.763 GJ/t at the model's heat capacity, a thermal efficiency of about 63 %.

7.3.3 Walking Beam, Instruments and Events

A walk clock advances by the integration step except during a mill delay. When it reaches t_p, the slab in position 30 is discharged and recorded (mean, core, true surface and pyrometer reading), every slab moves one position, and a new slab enters position 1 with all five cells at T_{ch}. The pyrometer reads $T_{pyro} = T_s + \sigma_p\varepsilon$ with $\sigma_p = 5\,^\circ\mathrm{C}$ and ε a standard normal deviate. The zone thermocouples add noise of $\sigma_{tc} = 2\,^\circ\mathrm{C}$ and are smoothed by a first-order lag with $\tau_{tc} = 30\,\mathrm{s}$, updated as $T_{m,z} \leftarrow T_{m,z} + (T_z + \sigma_{tc}\varepsilon - T_{m,z})\,\Delta t/(\tau_{tc} + \Delta t)$. The true exchange factor is $\varphi(1 + \Delta\varphi/100)$, and $\Delta\varphi$ is the

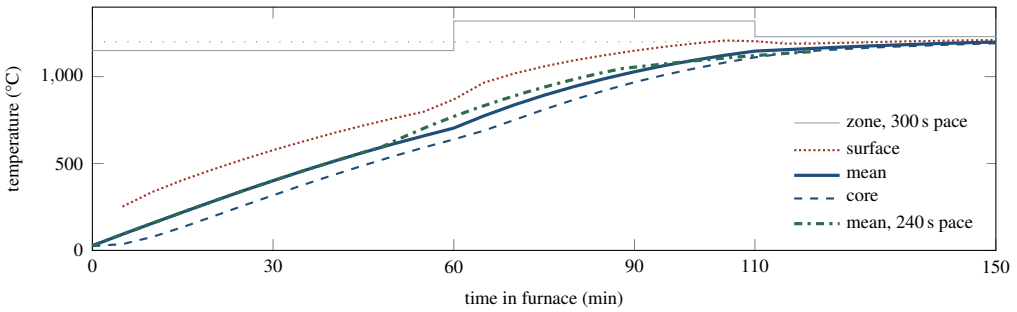


Figure 7.2 Slab heating curves at the operator setpoints of 1150, 1320 and 1230 °C, from the true furnace at its periodic steady state with thermocouple noise off: zone temperature (thin gray steps), surface (dotted), mean (thick solid) and core (dashed) at the 300 s pace, and the mean at the 240 s pace (dash-dotted), with the target as a thin loosely dotted line. At the 300 s pace the slab leaves with a mean of 1199 °C, a core of 1192 °C and a surface of 1211 °C. At 240 s the same zones deliver a mean of 1145 °C after 2 h.

Table 7.2 Furnace and Slab Parameters. Visible Parameters Have Sliders in the Show Model Dialog, Hidden Ones Are Kept From the Student, and Fixed Ones Are Built Into the Source

Symbol	Meaning	Value	Access
φ	exchange factor, true	0.6	visible
k	steel conductivity	30 W/(m K)	hidden
ρc_p	steel heat capacity	5.07 MJ/(m^3 K)	hidden
h	slab thickness	250 mm	visible
N_c	cells per half thickness	5	fixed
A	heated area per slab	30 m^2	fixed
—	slab mass	29.25 t	fixed
C_z	zone heat capacity	100 MJ/K	visible
P_P, P_H, P_S	burner capacity	100, 90, 30 MW	visible
UA	zone wall loss	0.003 MW/K	hidden
T_{max}	zone setpoint limit	1350 °C	visible
η	available heat, Eq. (7.4)	0.85 at 900 °C	fixed
τ_{tc}	thermocouple filter	30 s	fixed
σ_{tc}, σ_p	thermocouple, pyrometer noise	2, 5 °C	slider
N	slab positions (12, 10, 8)	30	fixed

disturbance the tracking model does not know. The events available to the scenarios are a mill delay that stops the walk clock for the delay length, a pace change to 240 s, hot charge at 450 °C and a 15 % fall in φ.

At reset, including a restart that keeps the learner's settings (Section 1.6.1), the true furnace and both tracking models are filled with the periodic steady state of the operator setpoints, held below T_{max}, and the pace and charge temperature then in force, found by marching one slab through the thirty positions with 10 s explicit Euler steps. The zones start at those setpoints. In automatic each zone loop starts from the firing that holds its zone, computed from Equation (7.3) with the derivative set to zero, and in manual from the firing stations, which default to 62.5, 55 and 31.5 %, close to that firing at the nominal state. A 10 s fourth-order Runge-Kutta step advances the plant, and the controller executes at that period too.

The consequence of the conduction time is drawn in Figure 7.2. The surface climbs to about 1208 °C late in the heating zone, falls as the slab enters the cooler soak zone and then holds while the core closes on it, so the surface-to-core gap shrinks from 127 °C at 105 min to 19 °C at discharge. At the faster pace the same zone temperatures act for 30 min less and the mean reaches only 1145 °C.

Table 7.2 collects the parameters.

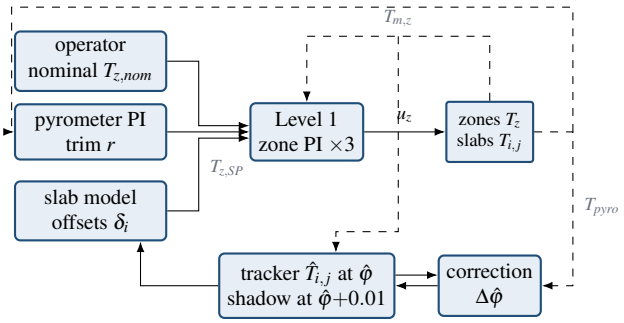


Figure 7.3 The two control levels. One of three Level 2 sources supplies the zone setpoints $T_{z,SP}$ to three Level 1 PI loops. The slab model plans from a tracker driven by the measured zone temperatures, and the correction compares each pyrometer reading with the tracker's predicted surface, using the shadow tracker for sensitivities, and updates $\hat{\varphi}$ and every tracked slab.

7.3.4 Assumptions

Heat flows through the thickness only. Real slabs are colder where they rest on the water-cooled beams and at their edges, effects that a one-dimensional slab cannot represent.

Each zone has one temperature and one exchange factor applies to every slab. The gas and wall temperatures of a real zone vary along its length and differ from the thermocouple, and the total heat exchange factor of a real furnace is estimated from measured temperatures rather than assumed [17, 18].

Steel properties are constant. The heat capacity of carbon steel peaks at its phase transformation and its conductivity changes with temperature, so the mid-furnace temperatures are the least accurate part of the heating curve.

Every slab has the same size, grade and target, and the burners respond instantly to the firing command. The available-heat fraction is a linear function of zone temperature standing in for the combustion and flue gas balance, and scale growth is not modeled.

7.4 Two Levels of Control

The control system is hierarchical, as drawn in Figure 7.3. Level 1 holds the zones. Level 2 decides what they hold, recomputing the three setpoints every 120 s and at reset.

7.4.1 Level 1: Zone Loops

Each zone loop is the velocity-form PID of Section 1.6.2, here without derivative action and with no output bias, executed every $\Delta t = 10$ s on the filtered thermocouple. In the symbols of Equations (1.13) to (1.16), the setpoint is $T_{z,SP}$, the measurement is $T_{m,z}$ and the output u_z is limited to 0 to 100 %. The integral I is seeded with the holding firing at reset and is the only term the limit acts on, by back-calculation with a tracking time equal to τ_I. The loops are reverse acting in the convention of Section 1.6.2, with action $a = +1$, since the zone temperature rises with the firing and a zone below its setpoint raises its firing. The defaults are $K_c = 0.6$ percent per °C for all three zones and $\tau_I = 10$, 10 and 15 min. There is no rate limit on the firing and no feedforward from the slab load. In manual mode each loop tracks the operator's firing so that the transfer to automatic is bumpless.

The loops are fast next to the slabs. With noise off, a 50 °C step in the operator's soak setpoint, which reaches the loop at the next Level 2 calculation, brings the soak zone within 10 °C in 16.0 min and first within 2 °C in 22.8 min, after which it rises to a peak 5.5 °C above the new setpoint at 43 min. This separation is what allows Level 2 to treat a zone setpoint as though it were the zone temperature.

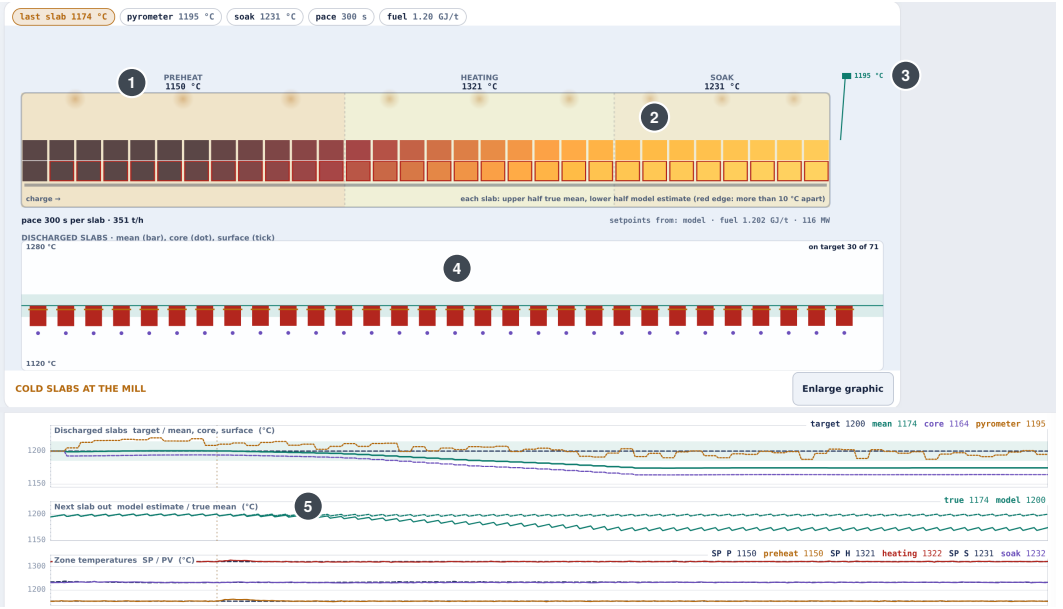


Figure 7.4 A model that is wrong, case 7, six hours into the run with the slab model setting the zones and correction off, cropped to the furnace view and the first three trends. The furnace section shows the zones at 1150, 1321 and 1231 °C, and each slab drawn in two halves, the upper colored by its true mean and the lower by the model's estimate, with a red outline around the lower half wherever the two differ by more than 10 °C. The discharged-slab strip below it shows every recent mean below the target band and 30 of 71 slabs on target, with the status reporting cold slabs at the mill. In the trends the discharged mean falls from 1200 °C to 1174 °C after the fall in heat transfer at 1 h, the pyrometer reads 1195 °C, the model's estimate of the slab about to leave is 1200 °C while its true mean is 1174 °C, and the zone setpoints barely move. Callouts: (1) the zones and their temperatures, (2) the slabs in two halves with the red outlines, (3) the pyrometer reading at the door, (4) the discharged-slab strip with its target band, (5) the trend of the model's estimate against the true mean of the next slab out.

7.4.2 Level 2: Operator and Pyrometer Trim

With operator setpoints the zones hold the nominal values $T_{z,nom}$ of 1150, 1320 and 1230 °C, limited to T_{max}. The pyrometer strategy adds a trim r to the heating and soak setpoints, $T_{H,SP} = \mathrm{clamp}(1320 + r, 800, T_{max})$ and likewise for soak, while preheat stays at nominal. The trim comes from a second velocity-form PI of the same algorithm, also reverse acting, executed once per discharged slab with $\Delta t = t_p$, on the error $e = (\bar{T}_{SP} + 12) - T_{pyro}$. The offset of 12 °C is the assumed surface excess over the mean at the nominal pace. The defaults are $K_c = 0.5$ and $\tau_I = 30$ min, and the trim is limited to ± 120 °C. The T_{max} clamp on the heating setpoint acts after the trim and is not seen by its back-calculation, so a trim above 30 °C leaves the heating zone at 1350 °C while the integral keeps raising the soak setpoint. When another strategy is active the trim loop tracks its last output.

7.4.3 Level 2: The Slab-Tracking Model

The tracker holds its own copy $\hat{T}_{i,j}$ of every slab, integrated by explicit Euler at 10 s from Equations (7.1) and (7.2) with the measured zone temperatures $T_{m,z}$ in place of T_z and its own exchange factor $\hat{\varphi}$, which starts at 0.6. It shifts and charges with the true furnace at every walk, so it knows each slab's position, charge temperature and elapsed time exactly. With $\hat{\varphi}$ correct its estimate of the

next slab out differs from the truth by at most 0.34 °C over the eight hours of case 4, pace change included.

Every 120 s each tracked slab i asks for the uniform offset δ_i on the nominal profile that brings its predicted discharge mean to the target, Equation (7.5).

$$\bar{\hat{T}}_i(t_{dis};\ \mathrm{clamp}\,[T_{z,nom}+\delta_i,\ 800,\ T_{max}]) = \bar{T}_{SP} \tag{7.5}$$

The prediction marches the slab's cells forward in explicit Euler steps of at most 40 s, or 0.8 of the stability limit $\rho c_p \Delta x^2/2k$ where that is shorter (52.8 s at 250 mm), through the time left at its present position and then t_p at each later position, under the offset zone temperatures and at the present pace. Equation (7.5) is solved for $\delta_i \in [-250, 150]$ by a bracketed secant (false position) iteration of at most seven steps, stopping when the predicted mean is within 0.3 °C of target. A slab that cannot reach the target at the upper bound, or overshoots it at the lower one, returns that bound. The zone setpoint is the nominal plus the average offset of the slabs under it, Equation (7.6).

$$T_{z,SP} = \mathrm{clamp}\left[T_{z,nom} + \frac{1}{n_z}\sum_{i\in z}\delta_i,\ 800,\ T_{max}\right] \tag{7.6}$$

Each δ_i is computed as if it applied to every zone the slab has yet to cross, while the zones receive averages, so the calculation is a heuristic rather than an optimization. The prediction assumes the present pace continues. During a mill delay the walk clock is frozen, the time left at the present position stays constant, and the model responds to the delay through the extra heating it sees in the tracked slabs rather than through knowledge of the delay length. A stopped door slab is not an imminent discharge. At each two-minute Level 2 update during a stop, the soak setpoint is held at the product mean-temperature target, while the preheat and heating zones retain their model calculations. Model planning resumes when the beam moves. The controller uses the current stopped state without forecasting the end of the delay.

> ⧖ HISTORY **Computers Plan the Heating (1985 to 1987)**
>
> By the mid-1980s computer control of a reheating furnace was built around a calculation of the temperature inside the steel, and because it ran on line it had to be fast [19]. Most systems then solved difference equations, since the thermal properties of steel are nonlinear, and a 1985 paper proposed an analytical solution to cut the computing time. A system reported in 1987 computed an optimum temperature profile for the steel in real time to minimize fuel, predicted the temperature of each piece along the furnace with a difference equation, and set each zone's temperature as a weighted average of that profile [20]. In commercial operation it saved at least 10 percent of energy against manual operation and held the discharge temperature within 15 °C of its aim.

7.4.4 Correcting the Model From the Pyrometer

When correction is on, each discharged slab gives one residual, the pyrometer reading minus the surface $\hat{T}_s$ the tracker predicts for that slab from its cells, the measured soak temperature and $\hat{\varphi}$. A shadow tracker runs the same equations at $\hat{\varphi}+\varepsilon$ with $\varepsilon = 0.01$, so the sensitivity of every tracked cell to the exchange factor is a finite difference, $S_k = (\hat{T}'_k - \hat{T}_k)/\varepsilon$. The update is a damped Gauss-Newton step on the scalar $\hat{\varphi}$, Equation (7.7).

$$\Delta\hat{\varphi} = \frac{g\,\mathrm{clamp}\left(T_{pyro}-\hat{T}_s,\ -60,\ 60\right)}{\max\left(20,\ \partial\hat{T}_s/\partial\hat{\varphi}\right)} \tag{7.7}$$

with gain $g = 0.3$ by default and $\hat{\varphi}$ limited to $[0.3, 1.0]$. The same step is then applied to every tracked cell and its shadow, $\hat{T}_k \leftarrow \hat{T}_k + S_k\Delta\hat{\varphi}$, so that the slabs still in the furnace take on the history they would have had under the new exchange factor at once. In the wrong-model scenario the surface sensitivity is about 135 °C per unit of φ, so a 1 °C residual moves $\hat{\varphi}$ by about 0.002 at the default gain.

7.4.5 A Joint Predictive Zone Profile

A receding-horizon alternative jointly chooses the three zone setpoints, rather than averaging independently requested offsets. It is the app's predictive option of Section 1.6.3, built here on a slab model of its own. Predictive reheating formulations motivate this comparison [14, 16, 21]. Its separate 150-cell tracker uses the conduction and radiation balances of Equations (7.1) and (7.2) at declared nominal values: 250 mm, $k = 30\,\mathrm{W/(m\,K)}$, $\rho c_p = 5.07\,\mathrm{MJ/(m^3\,K)}$ and $\varphi = 0.6$. Measured zone temperatures drive ten-second observer steps. The measured charge temperature and current walking pace determine slab entry and shifts. Neither true slab cells, changed thermal coefficients nor a future event schedule is supplied.

When correction is enabled, each discharged pyrometer reading updates a surface-output bias as $b^+ = \mathrm{clip}[b + 0.2(T_{pyro} - \hat{T}_s - b), -40, 40]$. The same bias corrects predicted mean and core outputs. This approximation does not identify the exchange factor or its influence on the internal profile. Its restrictions differ from the parameter correction of the slab model, and no general offset-free guarantee is assigned to it [22]. Volume 9, Chapter 4 treats an offset-free predictive controller whose estimator carries a disturbance bias.

Every 120 s, six cohorts at positions 5, 10, 15, 20, 25 and 30 are predicted to discharge at the present pace. At 300 s per slab, the longest remaining interval is at most 130 min. There is one constant future zone profile, replanned each decision, rather than a time-varying sequence of future setpoints. Prediction uses 20 s Euler cell steps and approximates future zone tracking as instantaneous, as the slab model's predictor does. During a stop, only the two soak cohorts constrain a stationary 30 min forecast. No stop-end time is read.

Central perturbations of ± 5 K produce a local 12-output, three-input map $\hat{y} \simeq c + Hu$, whose outputs are six discharge means and six cores. For u the zone setpoints and u^- their previous values, the five-variable quadratic program minimizes

$$J = \sum_i \frac{(\hat{\hat{T}}_{i,dis} - T_{SP})^2}{225} + 0.0003\|u - u^-\|^2 + 0.0001\sum_z u_z + 0.1\sigma_m^2 + \sigma_c^2. \tag{7.8}$$

Each predicted mean has band $T_{SP} \pm 15$ K enlarged by shared slack $\sigma_m \geq 0$, and each core has floor $T_{SP} - 40 - \sigma_c$ K. Setpoints satisfy $800 \leq u_z \leq T_{max}$ and move at most 30 K from the previous profile clipped into this box. A 10^{-8} ridge is added to the three setpoint diagonals. The linear temperature price is a proxy, not a calculation of fuel cost. In a stop, the objective and output bands apply only to the two soak cohorts.

Independent implicit-surface calculations differ from the five-substitution surface routine by at most 0.002 95 K at three temperatures. Repeating the map with ± 2.5 K changes sensitivities by at most 3.08×10^{-6} K/K. Across the measured runs, the largest nonlinear versus local-map output error at the chosen profile is 0.552 K. A failed or invalid solve retains the last valid profile under the zone PIs, clamped to the current setpoint ceiling even on failed startup. The next valid profile resumes planning. The other three strategies do not change when this option is present.

7.5 Results

Table 7.3 lists the eleven scenarios, and Table 7.4 gives what each one measured. Every value was measured on the application's engine at its default seed, which the application uses on every run (Section 1.6.1), and the spread over other noise seeds is stated where it matters. A slab is on target when its discharge mean is within 15 °C of the target, and the rms error is over every slab discharged in the run. An alternative setting such as the operator setpoints in case 5 is chosen before the run starts, and the application gives one run whether it is chosen then or carried through a restart. Figure 7.5 draws the main outcomes as a chart.

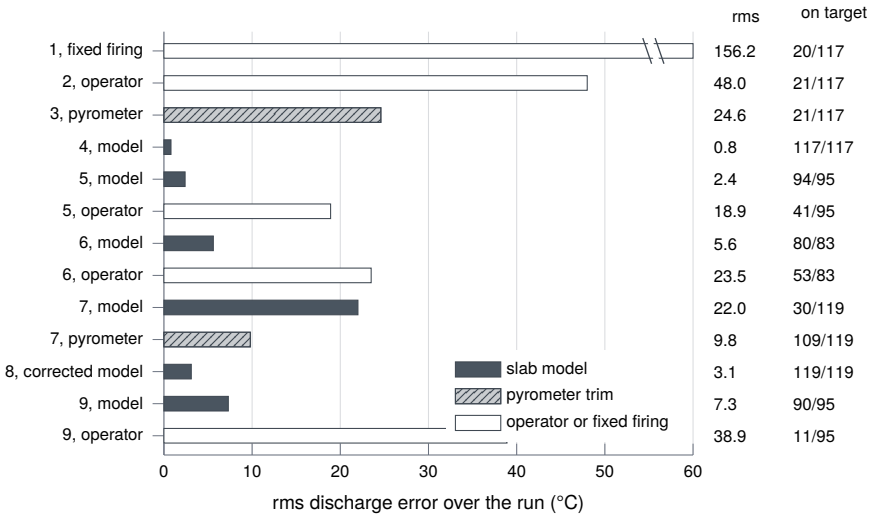


Figure 7.5 The scenario outcomes of Table 7.4 as a chart: the rms discharge error of each case and setpoint source (bars and first column, in °C) and the slabs on target (second column). The case 1 bar is cut at 60 °C. Notice that wherever a case is run with two setpoint sources, the solid bar of the slab model is the shorter, except in case 7, where the model is wrong and uncorrected and the hatched pyrometer bar beats it.

Table 7.3 Scenario Sequence. Events at 1 h Unless Stated

#	Scenario	Setpoints	Event
1	fixed firing (manual)	—	pace 240 s
2	holding the zones	operator	pace 240 s
3	pyrometer trim	pyrometer	pace 240 s
4	slab tracking	model	pace 240 s
5	hot charge	model	$T_{ch} = 450\,°C$
6	mill delay	model	60 min at 2 h
7	a wrong model	model	φ −15 %
8	model corrected	model, $g = 0.3$	φ −15 %
9	harder grade	model	target 1240 °C
10	predictive profile	MPC, bias correction	pace 240 s
11	challenge, 12 h	operator	all four

7.5.1 Holding the Firing and Holding the Zones

In case 1 the firing stays at 61, 52 and 33 % when the mill speeds up from 300 to 240 s per slab. The faster slabs arrive colder and draw more heat, the zones sag to 1131, 1226 and 1073 °C after eight hours, and the last slabs leave at about 1004 °C. The zones are not fast in open loop either. Raising the soak firing from 33 to 43 % at the start of case 1 lifts the soak zone above an otherwise identical run by 22.7 °C in 30 min, 32.5 °C in 1 h and 37.4 °C in 2 h, still rising because the slabs beneath it warm and take less heat.

Case 2 closes the zone loops at the operator setpoints. The zones stay within 7 °C of their setpoints, the largest excursion being a 6.8 °C dip of the soak zone at two hours as the faster slabs reach it. The slabs leave at a mean of about 1145 °C with a core of 1128 °C, the heating curve of Figure 7.2 at the faster pace. Only 21 of 117 slabs are on target, all of them discharged in the first 1.6 h. Level 1 has done its job and the product is 55 °C cold.

🕮 **Pause.** Try Exercise 7.1 before reading on.

Table 7.4 Measured Scenario Outcomes. Cases 1 to 6 and 9 Over 8 h, 7 and 8 Over 10 h. Late Mean Is the Average of the Last Ten Slabs

Case and setpoints	on target	rms (°C)	late mean	GJ/t
1, fixed firing	20/117	156.2	1004	0.991
2, operator	21/117	48.0	1145	1.134
3, pyrometer	21/117	24.6	1177	1.169
4, model	117/117	0.8	1200	1.214
5, model	94/95	2.4	1200	0.923
5, operator	41/95	18.9	1224	0.955
6, model	80/83	5.6	1200	1.246
6, operator	53/83	23.5	1199	1.260
7, model	30/119	22.0	1174	1.199
7, pyrometer	109/119	9.8	1192	1.218
8, corrected model	119/119	3.1	1203	1.235
9, model	90/95	7.3	1240	1.299
9, operator	11/95	38.9	1199	1.227

Table 7.5 Tuning Studies. Pyrometer Trim Gain in Case 3 Over 8 h, and Correction Gain and Pyrometer Noise in Case 8 Over 10 h

Case	Setting	on target	rms (°C)	late mean
3	$K_c = 0.2$	21/117	30.3	1174
3	$K_c = 0.5$ (default)	21/117	24.6	1177
3	$K_c = 1.0$	21/117	22.4	1177
8	$g = 0.1$, $\sigma_p = 5$	119/119	4.4	1202
8	$g = 0.3$, $\sigma_p = 5$ (default)	119/119	3.1	1203
8	$g = 1.0$, $\sigma_p = 5$	119/119	3.5	1205
8	$g = 0.1$, $\sigma_p = 15$	115/119	6.9	1206
8	$g = 0.3$, $\sigma_p = 15$	113/119	7.8	1209
8	$g = 1.0$, $\sigma_p = 15$	102/119	9.6	1214

TRY IT IN THE APP **2 · Holding the Zones, Not the Slabs**

Load scenario 2 and press Run. Watch the discharged-slab strip under the furnace after the mill speeds up at one hour. The zones stay within 7 °C of their setpoints, yet the slabs drop below the target band and stay there, about 55 °C cold. Only the slabs that left in the first 1.6 h are on target.

7.5.2 A Loop on the Surface

Case 3 trims the heating and soak setpoints to hold the pyrometer at 1212 °C. The loop succeeds, and the last ten slabs have a surface of 1212 °C (1210 °C as read by the noisy pyrometer) with a mean of 1177 °C and a core of 1159 °C. The heating setpoint reaches its 1350 °C limit at 4.3 h. At the faster pace the surface-to-core gap widens to about 53 °C against 19 °C at the nominal pace. The shorter stay causes most of the widening, since the operator setpoints of case 2 leave nearly the same gap at the same pace, and the hotter zones add little to it. The fixed 12 °C offset therefore no longer relates the surface to the mean, and 96 of 117 slabs still leave more than 15 °C cold. The rms error halves, from 48.0 to 24.6 °C.

No tuning of the trim loop repairs this, as reported in Table 7.5. A gain of 0.2 settles the mean at about 1174 °C and a gain of 1.0 at about 1177 °C, the same as the default, and all three put the same 21 slabs on target. The loop is controlling the wrong variable, and a faster loop only reaches the wrong answer sooner.

Pause. Try Exercise 7.2 before reading on.

7.5.3 Planning With the Slab Model

Case 4 takes the setpoints from the slab model. At the pace change every tracked slab has less time left, and the Level 2 calculation on the scan of the change raises the setpoints from nominal to 1192, 1347 and 1244 °C. They end at about 1195, 1350 and 1277 °C, with the heating zone on its limit from 1.1 h. All 117 slabs leave within 15 °C of target, 0.8 °C rms. The surface-to-core gap still reaches 53 °C, about as much as under the operator setpoints at the same pace, and the calculation does not constrain it. The price is 1.214 GJ/t of fuel against 1.134 GJ/t for the operator setpoints, which deliver cold slabs.

The charge temperature enters the tracker as a known initial condition, so hot charging acts as feedforward. In case 5, slabs charged at 450 °C lower each zone's setpoint as they reach it, to minima of 1126, 1294 and 1161 °C, and the soak setpoint first falls 110 min after the first hot slab enters, when that slab reaches the soak zone. Of 95 slabs, 94 are on target at 0.923 GJ/t. With the operator setpoints the hot slabs leave at about 1224 °C, 54 of 95 are more than 15 °C hot, and the fuel is 0.955 GJ/t, about 3.5 % more.

In case 6 the mill stops for an hour at two hours, on the scan at which a walk falls due, so the slab at the door waits out the delay. The tracked slabs keep heating while the beams stand still, and the model cuts the preheat and heating setpoints to 1117 and 1264 °C by the end of the delay. The soak setpoint holds 1200 °C during the delay and falls to 1133 °C at 3.6 h as the slabs that waited under the heating zone reach it. The furnace burns 249 GJ during the delay against 295 GJ with the operator setpoints, and the first fourteen slabs after the restart leave at 1194 to 1216 °C against 1227 to 1257 °C. Over the run 80 of 83 slabs are on target with the model and 53 with the operator setpoints.

> **DID YOU KNOW? Better and Cheaper With Hot Slabs**
>
> From the simulator: with slabs charged at 450 °C in case 5, the slab model puts 94 of 95 on target at 0.923 GJ/t, while the operator setpoints put 41 of 95 on target at 0.955 GJ/t (Table 7.4). Here accuracy and fuel do not trade against each other: using the charge temperature wins both.

Case 9 raises the target to 1240 °C. The soak setpoint rises to about 1334 °C, but the slabs near the door have too little time left, and the five discharged from the change to 20 min after it leave at 1200 to 1225 °C. From 25 min on every slab is within 15 °C of the new target, and 90 of 95 are on target over eight hours. The operator setpoints, unchanged, put 11 of 95 on target.

7.5.4 A Wrong Model and Its Correction

Case 7 lowers the true exchange factor by 15 % at one hour. The tracker still computes with 0.6, so every tracked slab heats faster in the model than in the furnace. From four hours on, its estimate of the next slab out reads 1196 to 1200 °C while the true mean is 1167 to 1174 °C, the setpoints stay within 6.2 °C of nominal, and the slabs leave at about 1174 °C, as shown in Figure 7.4. Only 30 of 119 are on target. Over the last ten slabs the pyrometer reads 17.0 °C below the predicted surface while the true mean is 25.8 °C below the predicted mean, so the residual is available and nothing uses it. The pyrometer strategy on the same furnace puts 109 of 119 on target, because it at least sees the disturbance.

Pause. Try Exercise 7.3 before reading on.

> **TRY IT IN THE APP 7 · A Model That Is Wrong**
>
> Load scenario 7 and run past four hours. Each slab is drawn in two halves, the upper colored by its true mean and the lower by the model's estimate, with a red edge where they differ by more than 10 °C. The model's next slab out reads 1196 to 1200 °C while its true mean is 1167 to 1174 °C.

Case 8 switches the correction on. The exchange factor estimate first falls below 0.53 at 3.2 h, and between four and ten hours it averages 0.508 with a standard deviation of 0.008 and a range of 0.491

Table 7.6 Challenge Outcomes Over 12 h. Worst Is the Largest Discharge Error in the Last 4 h

Setpoints from	on target	rms (°C)	worst (°C)	GJ/t
operator (as loaded)	22/159	59.7	97.8	1.011
pyrometer	59/159	35.9	46.5	1.040
model	66/159	23.2	50.7	1.040
model, corrected	150/159	8.6	37.0	1.066
same, $K_{c,S} = 1.2$	151/159	7.7	33.1	1.066

to 0.535, against a true value of 0.51. All 119 slabs are on target, 3.1 °C rms, with means between 1192 and 1205 °C. With the pyrometer noise set to zero the estimate stays within 0.507 to 0.510 over the same interval and the error falls to 1.5 °C rms, so the residual scatter of $\hat{\varphi}$ is the pyrometer noise passed through Equation (7.7).

The correction gain trades speed against that noise, as reported in Table 7.5. At 5 °C noise a gain of 1.0 brings $\hat{\varphi}$ below 0.53 by 1.6 h but spreads it over 0.462 to 0.565, while a gain of 0.1 takes until 3.5 h and holds it within 0.499 to 0.518. At 15 °C noise every gain loses slabs, and the largest gain loses the most, 17 of 119 off target against 4 at a gain of 0.1. That ranking holds on average over ten noise seeds (14.9 against 6.4 slabs off target) but not at every seed.

7.5.5 The Graded Challenge

The twelve-hour challenge opens with a one-hour mill delay at one hour, the pace change at three, the fall in exchange factor at five and hot charge at seven, starting on the operator setpoints. It is met when at least 90 % of slabs are on target and none is more than 40 °C off in the last four hours. The outcomes appear in Table 7.6.

As loaded, 22 of 159 slabs (13.8 %) are on target, the worst late error is 97.8 °C, both criteria fail, and slabs leave between 1101 and 1257 °C. The uncorrected model answers the delay, the pace and the hot charge, and the pyrometer the change in heat transfer, but each alone puts fewer than half the slabs on target. Only the corrected model meets both criteria, with 150 of 159 slabs on target and a worst late error of 37.0 °C. That worst slab is the first hot-charged slab to reach the door, 37.0 °C hot at 9.05 h, while the cold-charged slab ahead of it leaves 16.5 °C cold, because the soak zone receives the average offset of both kinds. The corrected model also burns the most fuel, 1.066 GJ/t, because the other configurations discharge 73 to 109 slabs cold. Its margins are narrow, six slabs and 3.0 °C, and on sixty additional noise seeds (1 through 60) it passes at 47. The suggested settings therefore also raise the soak loop gain from 0.6 to 1.2, so the soak zone follows its fast-moving setpoint more closely (17.3 °C rms deviation at logged samples against 22.7 °C). They put 151 slabs on target (95.0 %) with a worst late error of 33.1 °C, again the first hot-charged slab, passing by seven slabs and 6.9 °C, and pass at 58 of those 60 additional seeds. The run is the same whether the settings are chosen before Run or carried through a restart. Changing a model parameter or a disturbance voids the attempt.

7.5.6 Joint Planning and Its Counterfactuals

Case 10 runs the predictive profile with both corrections on, the same seed, sensors and pace step at 1 h as case 4. Table 7.7 measures 8 h, except 10 h for the wrong-model/correction pair and 12 h for the challenge. In case 10 the rms error falls from 3.15 to 1.52 °C, with all 117 slabs on target, while fuel rises from 1.215 to 1.221 GJ/t. Setpoint travel falls from 2236 to 520 °C, as shown in Figure 7.6.

This nominal gain is not evidence that joint optimization is required. With correction off under both methods, the slab model achieves 0.756 °C rms against MPC's 1.348 °C. The smoother bias avoids part of the noise-driven parameter adjustment. A six-setting hindsight slab-model grid has best rms 2.83 °C with correction on, and establishes no optimum over all classical tunings. Half-step integration preserves the paired result, 3.153 versus 1.520 °C.

Table 7.7 Matched Slab-Model and Predictive Profiles

Case, law	On target	rms (°C)	GJ/t
Case 10, Slab	117/117	3.15	1.215
Case 10, MPC	117/117	1.52	1.221
Pace, no bias, Slab	117/117	0.76	1.214
Pace, no bias, MPC	117/117	1.35	1.222
Hot charge, Slab	94/95	2.43	0.923
Hot charge, MPC	94/95	3.14	0.871
Mill delay, Slab	80/83	5.55	1.246
Mill delay, MPC	61/83	13.91	1.387
Wrong model, Slab	30/119	21.99	1.199
Wrong model, MPC	29/119	22.67	1.194
Correction on, Slab	119/119	3.14	1.235
Correction on, MPC	114/119	8.67	1.208
Grade step, Slab	90/95	7.27	1.299
Grade step, MPC	90/95	7.91	1.255
Challenge, Slab	151/159	7.66	1.066
Challenge, MPC	125/159	13.15	1.074

On hot charge, MPC uses 0.871 versus 0.923 GJ/t, with 94 of 95 slabs on target under each law. The final slab-plus-zone stored heat is 830.6 rather than 896.2 GJ, against 2421.5 rather than 2564.4 GJ of firing. The 65.6 GJ lower inventory qualifies the 142.9 GJ firing reduction, so GJ/t alone is not a steady-operation saving. On the mill delay it loses, 61 rather than 80 of 83 slabs, with 13.91 rather than 5.55 °C rms and higher fuel. A stationary half-hour forecast and one constant profile do not reproduce the slab model's delay rule. Output bias also corrects changed heat transfer less effectively than parameter adaptation, 114 versus 119 of 119 slabs on target. With the challenge's suggested settings the MPC fails at 125 of 159 slabs, while the slab model passes at 151 of 159. Its worst late error is 32.96 °C, so the failed criterion is the fraction of on-target slabs.

Core prediction is not core measurement. On the grade step, four MPC slabs rather than three under the slab model have true cores below their target minus 40 K. Actual zone temperatures can also overshoot bounded setpoints: in the challenge with the suggested settings the heating zone peaks at 1354.5 °C under MPC and 1357.6 °C under the slab model. Soft output bands and hard setpoint boxes are consequently not physical temperature guarantees. Chrome mean solve times are 1.60–1.76 ms over six 60-plan captures, at most 0.0880 % core at 60 s/s playback and 5.28 % at the top 3600 s/s setting. Ordinary comparisons have no rejected plans. Eight injected invalid solver results are rejected, and two independent measurement rows reproduce exactly.

7.6 Technology in Practice

What Measures It

The variable the mill cares about, the discharge mean, has no sensor. Two instruments stand in for it. Each zone loop reads a thermocouple through the furnace wall, and a radiation pyrometer reads the surface of each slab as it leaves. The furnace temperature is thus a reading at one place. A 1973 study noted that steelmakers did not define the representative gas temperature of a furnace in the same way, and examined how the position of the reading changes the exchange factor estimated from it [23].

> **TECHNOLOGY Sensor: The Radiation Pyrometer**
>
> A radiation pyrometer concentrates the infrared light from a hot surface onto a small sensing element with a lens or a mirror, and the radiant energy grows with the fourth power of absolute temperature [24]. That steep law limits a pyrometer to a narrow range wherever good accuracy is needed. Its calibration does not depend on distance, because a pyrometer farther away sees a larger patch of the same surface. It reads only the surface, never the temperature inside an object. Its main weakness is emittance, which depends on the material, the finish and the shape and is known for certain only by testing the object at a known temperature.

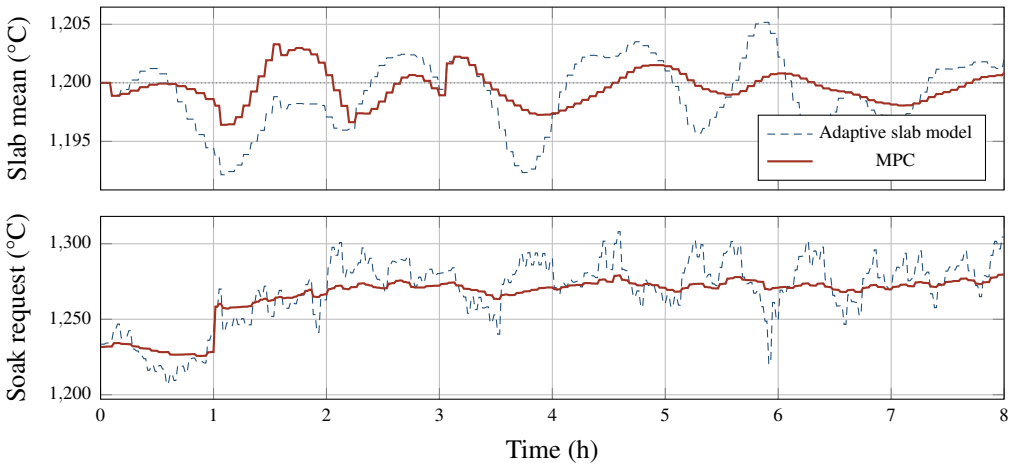


Figure 7.6 Discharged means and soak setpoints in case 10 with both corrections enabled, under the adaptive slab model (dashed) and the MPC (thick solid), with the 1200 °C target dotted. The pace changes after one hour.

In the model of Section 7.3 the thermocouple is the 2 °C noise and the 30 s filter, and the pyrometer is the true surface plus 5 °C of noise. A real pyrometer also carries the error of an uncertain emittance, which the model does not represent.

What Moves It

The zone loop's output u_z is a firing command in percent of burner capacity. On a real furnace it becomes setpoints for the fuel and air flows of the zone's burners, fired from above and below the slabs.

> **TECHNOLOGY** **Final Element: Cross-Limited Fuel and Air**
>
> A zone's firing command is carried out by two flow loops, one on the fuel and one on the combustion air of its burners [24]. With a plain ratio, a step in the command can put fuel ahead of air for a moment, a rich and dangerous mixture. Cross-limiting prevents it with two selectors. The air controller takes the greater of the command and the measured fuel flow, and the fuel controller the lesser of the command and the measured air flow, so air leads on a rise, fuel leads on a fall, and every transient runs lean. A rate limiter on the command keeps a lean transient from blowing the flame out, and a multiplier sets the air to fuel ratio.

The simulated burners respond instantly to the command, with no flow loops, no air ratio and no rate limit, as Section 7.3 states. Figure 7.7 draws the loops with the real devices in place.

Where the Controller Runs

The chapter's two levels map onto the control hierarchy of a plant. In a distributed control system a supervisory controller sends set points to the distributed field controllers and requests data from them, and the field controllers move their actuators on those commands and on sensor feedback [25]. Supervisory loops and lower loops run with cycle times from milliseconds to minutes. The three zone PI blocks of Level 1 would run in such a field controller, and the simulator executes them every 10 s. The slab-tracking model, its correction and the predictive option belong to the supervisory Level 2 computer, which in the simulator recomputes the setpoints every 120 s.

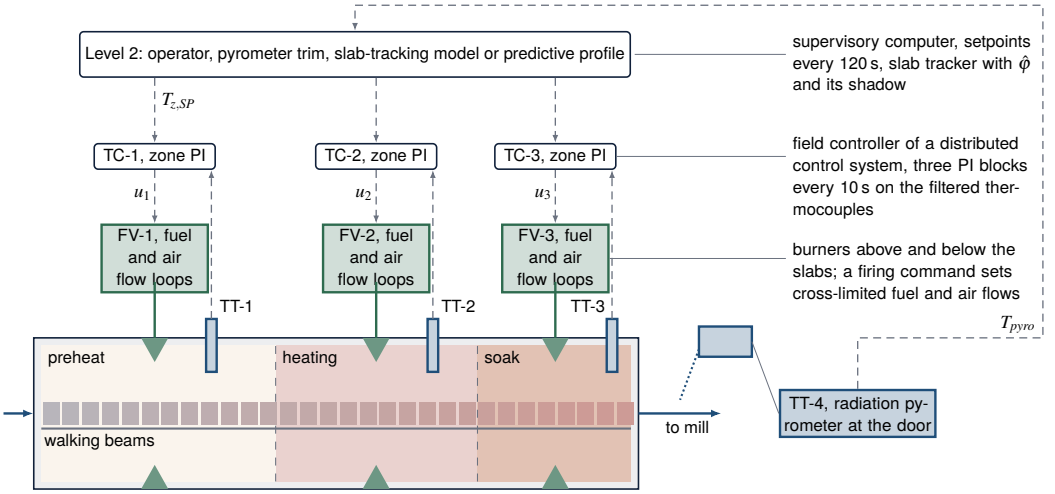


Figure 7.7 The reheat furnace loops drawn with their hardware. Each zone controller TC-1 to TC-3 reads a thermocouple through the roof, TT-1 to TT-3, and sends a firing command to the fuel and air flow loops FV-1 to FV-3 of its burners. The radiation pyrometer TT-4 reads each slab's surface at the door and reports only to Level 2. Notice that no device reads the slab mean: Level 2 turns the pyrometer and the zone readings into the three setpoints.

> **IN PRACTICE** **Decide Which Way a Thermocouple Fails**
>
> The most common failure of a thermocouple is to fail open, known as burning out [24]. An open input lets electrical noise from power lines, motors and drives reach the instrument, where it reads as a wildly varying temperature. Instruments therefore carry a burnout mode that sends the reading to a set end of the scale, and the choice is a safety decision. For a furnace in which too high a temperature is the greater danger, the field text advises high burnout: a burnt-out thermocouple then reports a false high temperature, and the controller lowers the firing. In the reverse-acting zone loops of this chapter, a reading pinned high would cut the firing of that zone.

What Is Modernizing It

Predictive control is reaching the events that the slab model of this chapter handles by rule. A strategy for mixed loading and delays in a walking-beam furnace solves its optimization by a rolling horizon with weights adjusted in real time, and in simulation it raised the control accuracy from 74.79 to 99.17 percent over three typical delay cases [16], the direction taken by the predictive option of Section 7.4.5.

Hot charging changes the furnace's work. Slabs charged directly from the caster need less reheating than slabs cooled to air temperature, which take longer and use more energy [26], as case 5 shows when the model lowers its setpoints for slabs charged at 450 °C.

Combustion itself is changing. A United States energy study lists a next-generation system for scale-free reheating, which uses preheated or oxygen-enriched air to control the flue gas, improves the quality and yield of the steel, and is estimated to save 0.2 GJ/t during reheating [26]. Electric heating, which would remove the flame altogether, is still an open challenge.

7.7 Limitations

The slab is one-dimensional with constant properties, so skid marks, edge effects and the heat capacity peak at the phase transformation are absent.

One exchange factor applies to the whole furnace and changes as a step. In a real furnace it varies by zone and drifts with scale and atmosphere, and a single pyrometer residual per slab cannot separate the zones.

> **OPEN CHALLENGE** **An Exchange Factor That Will Not Stay Put**
>
> Every slab-tracking model rests on a total heat exchange factor, and the factor belongs to the operation as much as to the furnace. A 1973 study estimated it at 0.64 for heavy oil firing and 0.3 for gas firing, both with the gas temperature read by a thermocouple through the furnace wall [23]. A 2021 model of a regenerative furnace found that it rises as fuel consumption or the hot charging temperature falls, and as productivity rises [18]. Its authors offer the result for dynamic compensation of the factor in on-line control models. Progress would be a factor estimated zone by zone, while the furnace runs, from measurements a plant already has.

The slab-model calculation is a per-slab offset averaged by zone, not an optimization over the schedule. It has no fuel term, no bound on the surface-to-core gap, no knowledge of the delay length and no look-ahead to announced grade or pace changes. Predictive formulations that optimize the zone temperatures over a horizon, including through delay operations, are the established answer to these omissions [14, 16].

The predictive option adds soft core and mean bands for six cohorts, but uses one future profile, nominal properties and a scalar surface bias. These restrictions explain the delay and changed-transfer losses. Its temperature price is not a fuel objective.

The fuel per tonne does not credit the heat stored in the slabs still in the furnace, so it compares fairly only runs that end with similar furnace contents.

> **OPEN CHALLENGE** **Heating Slabs With Electricity**
>
> Slabs are reheated to about 1250 °C for rolling, and 1 to 3 percent of the steel is lost to oxidation on the way [27]. Electric resistance heating needs no combustion atmosphere, so it could remove that loss and cut carbon dioxide emissions, and it offers more control of zone temperature than gas burners. For large slabs it is not yet possible. They need many heating elements close to the slab surface, which means elements across the roof, and the usual molybdenum silicide element creeps badly: it serves hung on the side walls, but mounted across the roof it must run well below its maximum temperature. Progress would be an element that keeps its shape when mounted horizontally at furnace temperature.

7.8 Conclusion

This chapter studied a reheat furnace, which heats a moving product whose one important temperature is never measured. In the simulator, zone loops that hold their setpoints through a pace change deliver slabs 55 °C cold, and a pyrometer loop that holds the surface leaves the mean between 1174 and 1177 °C for every trim gain tried. A slab-tracking model that plans the zone setpoints puts all 117 slabs on target through the same pace change, uses the charge temperature of hot slabs as feedforward, and cuts the fuel burned in a one-hour delay from 295 to 249 GJ. When the exchange factor falls 15 % the same model discharges 30 of 119 slabs on target while reading on target, and a Gauss-Newton correction from the pyrometer restores all 119 with $\hat{\varphi}$ averaging 0.508 against a true 0.51. The Level 2 calculation is a heuristic without fuel or gap terms, and the slab is one-dimensional. The joint predictive option saves hot-charge fuel, but loses to the slab model's delay rule and fails the challenge with the suggested settings. A fuller model would add per-zone exchange factors, mixed slab sizes and announced delays. The furnace is an accessible case of estimating, planning with and correcting the state of a product nobody measures.

KEY POINTS

- The furnace must deliver slabs hot all the way through, and no instrument reads the inside of a slab.
- Zone loops that hold their setpoints through a pace change deliver slabs 55 °C cold, and a pyrometer loop on the surface leaves the mean between 1174 and 1177 °C for every trim gain tried.
- A slab-tracking model that plans the zone setpoints puts all 117 slabs on target through the same pace change and uses the charge temperature of hot slabs as feedforward.
- When the exchange factor falls 15 % the model reads on target while 30 of 119 slabs are, and a correction from the pyrometer restores all 119.

Where else this applies. A brick core that is never measured in Volume 5, Chapter 11 and a moving ribbon in Volume 5, Chapter 5 call for the same model run beside the plant.

References

[1] R. E. Kalman, “A New Approach to Linear Filtering and Prediction Problems,” *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[2] L. Ljung, “Asymptotic behavior of the extended Kalman filter as a parameter estimator for linear systems,” *IEEE Transactions on Automatic Control*, vol. 24, no. 1, pp. 36–50, 1979.

[3] F. Qi, J. Xu, Y. Hou, and B. Li, “Digitization and greening of industrial reheating furnaces: A review of mathematical model, heating pattern, control system and energy analysis in the reheating furnace,” *Applied Thermal Engineering*, vol. 263, p. 125320, 2025.

[4] J. Xu, B. Li, and F. Qi, “Effect assessment of charging conditions on energy consumption and slab heating quality in an industrial-scale reheating furnace,” *Applied Thermal Engineering*, vol. 265, p. 125565, 2025.

[5] Y. Wada, “On the minimum specific heat consumption of the reheating furnace in the iron and steel industry,” *Transactions of the Iron and Steel Institute of Japan*, vol. 20, no. 3, pp. 192–199, 1980.

[6] M. Y. Kim, “A heat transfer model for the analysis of transient heating of the slab in a direct-fired walking beam type reheating furnace,” *International Journal of Heat and Mass Transfer*, vol. 50, no. 19-20, pp. 3740–3748, 2007.

[7] A. Steinboeck, D. Wild, T. Kiefer, and A. Kugi, “A mathematical model of a slab reheating furnace with radiative heat transfer and non-participating gaseous media,” *International Journal of Heat and Mass Transfer*, vol. 53, no. 25-26, pp. 5933–5946, 2010.

[8] K. Kono, M. Fujiura, K. Goto, and J. Yamasaki, “Walking beam slab reheating furnace,” *Tetsu-to-Hagane*, vol. 55, no. 14, pp. 1392–1398, 1969, in Japanese.

[9] H. Pike and S. Citron, “Optimization studies of a slab reheating furnace,” *Automatica*, vol. 6, no. 1, pp. 41–50, 1970.

[10] K. Yorifuji, K. Morita, H. Miura, and J. Ichihara, “A new integrated control system for the reheating furnace of a hot strip mill,” in *Conference Record of the IEEE Industry Applications Society Annual Meeting*, 1989, pp. 1467–1471.

[11] N. Yoshitani, Y. Naganuma, and T. Yanai, “Optimal Slab Heating Control for Reheating Furnaces,” in *1991 American Control Conference*, 1991, pp. 3030–3035.

[12] H. J. Wick, “Estimation of Average Slab Temperatures of a Simulated Reheat Furnace using the Extended Kalman Filter,” in *1990 American Control Conference*, 1990, pp. 2953–2954.

[13] G. van Ditzhuijzen, D. Staalman, and A. Koorn, “Identification and model predictive control of a slab reheating furnace,” in *Proceedings of the International Conference on Control Applications*, vol. 1, 2002, pp. 361–366.

[14] A. Steinboeck, D. Wild, and A. Kugi, “Nonlinear model predictive control of a continuous slab reheating furnace,” *Control Engineering Practice*, vol. 21, no. 4, pp. 495–508, 2013.

[15] X. M. Nguyen, F. Lawayeb, P. Rodriguez-Ayerbe, D. Dumur, and A. Mouchette, “Nonlinear model predictive control of steel slab walking-beam reheating furnace based on a numerical model,” in *2014 IEEE Conference on Control Applications (CCA)*, 2014, pp. 191–196.

[16] Z. Yang, X. Luo, and J. Qiao, “A combined adaptive entropy-TOPSIS and model predictive control strategy for mixed loading and delay operations in the reheating furnace,” *Optimal Control Applications and Methods*, vol. 44, no. 5, pp. 2775–2795, 2023.

[17] X. Luo and Z. Yang, “A new approach for estimation of total heat exchange factor in reheating furnace by solving an inverse heat conduction problem,” *International Journal of Heat and Mass Transfer*, vol. 112, pp. 1062–1071, 2017.

[18] W. Ji, G. Li, L. Wei, and Z. Yi, "Modeling and determination of total heat exchange factor of regenerative reheating furnace based on instrumented slab trials," *Case Studies in Thermal Engineering*, vol. 24, p. 100838, 2021.

[19] T. Endo and Y. Funahashi, "A prediction algorithm of temperature in slab in reheating furnace by use of analytical solution," *Tetsu-to-Hagane*, vol. 71, no. 10, pp. 1335–1342, 1985, in Japanese with an English synopsis.

[20] "Computer control system for reheating furnace," *Transactions of the Iron and Steel Institute of Japan*, vol. 27, no. 7, p. 604, 1987, new Technology note.

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

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

[23] S. Matsunaga, "On the total heat exchange factor of the slab reheating furnace basing on the slab reheating experiment," *Tetsu-to-Hagane*, vol. 59, no. 2, pp. 301–307, 1973, in Japanese with an English synopsis.

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

[25] 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] K. Jamison, C. Kramer, S. Brueske, and A. Fisher, "Bandwidth study on energy use and potential energy saving opportunities in U.S. iron and steel manufacturing," U.S. Department of Energy, Advanced Manufacturing Office, Tech. Rep. DOE/EE-1231, 2015, oSTI 1248755.

[27] M. D. Mulholland, K. K. Bawane, and E. K. Putihprayogi Pek, "Creep performance of novel composite heating elements for reheat furnace electrification," TMS Energy Materials 2026, Idaho National Laboratory report INL/CON-26-93505, 2026, oSTI 3382474.

7.9 Exercises

7.1 ⌕ Concept check · 2 min

In case 2 the mill pace rises from 300 to 240 s per slab at 1 h. The three zone loops then hold their zones within 7 °C of the operator setpoints, yet the slabs leave about 55 °C cold. Which explanation is right?

(a) Each slab now stays 2 h instead of 2.5 h, so the same zones put less heat into it

(b) The zone loops are too slow for the pace change, so the zones sag below their setpoints for hours

(c) The faster slabs reach the soak zone hotter, so the soak loop cuts its firing and the mean falls

(d) The soak loop winds up against its lower firing limit after the pace change and stays there

7.2 ⌕ Concept check · 2 min

In case 3 a PI loop trims the heating and soak setpoints until the pyrometer reads the target plus 12 °C, 1212 °C. It succeeds, yet after the pace change the mean of the last ten slabs is 1177 °C. Which explanation is right?

(a) The 5 °C pyrometer noise biases the trim integral toward a surface that is too low

(b) At the faster pace the surface runs further ahead of the mean, so the fixed offset no longer fits

(c) The heating setpoint sits at its 1350 °C limit, and a higher trim gain would reach target

(d) The trim leaves the preheat zone at nominal, and that cold zone sets the mean the slabs reach

7.3 ⌕ Concept check · 1 min

In case 7 the true exchange factor falls by 15 % and the slab model is not corrected. Its estimate of the next slab out reads about 1200 °C while the true mean is about 1174 °C. Which signal shows that the model is wrong?

(a) The zone thermocouples, which read lower than the model expects them to read
(b) The model's estimate of the next slab out, which drifts below the target
(c) The pyrometer reading minus the surface the model predicts for the same slab
(d) The fuel per tonne, which climbs above its value with the right exchange factor

7.4 Theory and calculation · 6 min

The slab is 250 mm thick with $k = 30\,\mathrm{W/(m\,K)}$, $\rho c_p = 5.07\,\mathrm{MJ/(m^3\,K)}$, $N_c = 5$ cells per half thickness and an exchange factor $\varphi = 0.6$.

(a) Find the cell width Δx and the conduction time $(h/2)^2 \rho c_p / k$. *Checkpoint:* 25 mm and 44.0 min.

(b) Find the radiant coefficient $4\varphi\sigma T^3$ at 1230 °C and the Biot number on the half thickness. *Checkpoint:* 462 W/(m^2 K) and 1.93.

(c) The predictor of the slab model takes explicit Euler steps of 40 s, shortened to 0.8 of the stability limit $\rho c_p \Delta x^2 / 2k$ where that is shorter. Find the step at 250 mm and at the 180 mm minimum of the thickness slider. *Checkpoint:* 40 s (limit 52.8 s) and 21.9 s (limit 27.4 s).

7.5 Theory and calculation · 6 min

Each slab is 10 m long, 1.5 m wide and 250 mm thick and weighs 29.25 t. The furnace holds thirty slabs, twelve in preheat, ten in heating and eight in soak, and the model's steel has $\rho c_p = 5.07\,\mathrm{MJ/(m^3\,K)}$.

(a) Find the throughput and the residence time at the nominal pace of 300 s and at 240 s, and the time spent in the soak zone at each pace. *Checkpoint:* 351 t/h, 2.5 h, 40 min. 438.75 t/h, 2.0 h, 32 min.

(b) Find the density implied by the slab mass, the specific heat of the model's steel and the heat needed per tonne to raise the steel from 25 to 1199 °C. *Checkpoint:* 7800 kg/m^3, 650 J/(kg K), 0.763 GJ/t.

(c) At the nominal state the zones fire 117.9 MW in all. Find the gross fuel per tonne and the thermal efficiency. *Checkpoint:* 1.21 GJ/t and 63 %.

7.6 Theory and calculation · 6 min

The correction of Equation (7.7) moves $\hat{\varphi}$ by $g\,(T_{pyro} - \hat{T}_s)$ divided by the surface sensitivity. In the wrong-model case the sensitivity is about 135 °C per unit of φ and the default gain is $g = 0.3$.

(a) The true exchange factor falls by 15 % from 0.6. Find its new value. *Checkpoint:* 0.51.

(b) Over the last ten slabs of case 7 the pyrometer reads 17.0 °C below the predicted surface. Find the step in $\hat{\varphi}$ that one such residual gives at the default gain. *Checkpoint:* -0.038.

(c) A single reading carries pyrometer noise with a standard deviation of 5 °C. Find the step that this alone gives at $g = 0.3$ and at $g = 1.0$, and say what the difference means for the estimate. *Checkpoint:* 0.011 and 0.037.

7.7 App challenge · 8 min

Load scenario 2, Holding the zones, not the slabs, and change no setting. The mill speeds up to a slab every 240 s at 1 h. Run 8 h at the default seed, at the 3600 speed setting.

(a) Read the On target, Discharge error rms, Fuel and Worst surface-core gap tiles of the scorecard. *Checkpoint:* 21 of 117, 48.0 °C, 1.134 GJ/t, 53 °C.

(b) Set Zone setpoints from to Slab model, press Reset and run 8 h again. Read the same four tiles. *Checkpoint:* 117 of 117, 0.8 °C, 1.214 GJ/t, 53 °C.

(c) Find the extra fuel per tonne of the second run as a percentage of the first, and say what the two gap tiles show about the cause of the wide gap. Accept the tile values exactly, and the percentage within 0.2 points. *Checkpoint:* 7.1 %.

7.8 App challenge · 9 min

The true exchange factor falls by 15 % at 1 h in both scenarios below. Use the default seed and the 3600 speed setting, and run each for 10 h.

(a) Load scenario 7, A model that is wrong, and run it as given. Read the On target, Discharge error rms and Model exchange factor tiles. *Checkpoint:* 30 of 119, 22.0 °C, 0.600.

(b) Load scenario 8, Correcting the model from the pyrometer, and run it as given. Read the same tiles, then export the CSV file and find the first row after 1 h at which the column `phi_hat` is below 0.53. The time column is in hours. Accept a time within 0.05 h. *Checkpoint:* 119 of 119, 3.1 °C, 0.535, and 3.18 h.

(c) The tile of part (b) reads 0.535 against a true 0.51. Is that evidence of a bias in the estimate? Say how the file answers the question.

7.9 App challenge · 8 min

Load scenario 6, A mill delay. The mill stops for 60 min at 2 h. Use the default seed and the 3600 speed setting.

(a) Run 8 h as given, with the Slab model. Export the CSV file and read the column `fuel_GJ` in the rows at 2.0 and 3.0 h, and the soak setpoint `SP_S_C` during the delay. Read the On target tile at 8 h. *Checkpoint:* 854.2 and 1103.6 GJ, a soak setpoint of 1200 °C, and 80 of 83.

(b) Press Reset, set Zone setpoints from to Operator before Run, and run 8 h again. Read the same rows of `fuel_GJ` and the On target tile. *Checkpoint:* 852.3 and 1147.6 GJ, and 53 of 83.

(c) Find the fuel burned during the delay in each run. Accept each within 1 GJ. *Checkpoint:* 249.4 GJ and 295.2 GJ.

7.10 Application area · 4 min

Case 5 charges slabs at 450 °C from 1 h at the nominal pace of 300 s. Over 8 h the slab model uses 0.923 GJ/t and the operator setpoints 0.955 GJ/t.

(a) Find the fuel the model saves per hour and per day of operation at this pace, taking the run's average per tonne as representative. *Checkpoint:* 11.2 GJ/h and 270 GJ per day.

(b) The predictive profile of case 10 fires 142.9 GJ less than the slab model over the same hot-charge run, but ends with 65.6 GJ less heat stored in the slabs and zones. Bound the saving that steady operation could keep, and say why the bound is not reached. *Checkpoint:* Less than 77.3 GJ over the run.

7.11 Application area · 3 min

In a real walking-beam furnace scale grows mostly on slabs in the preheat zone, so the exchange factor there falls while the soak zone changes little. The app's model has one exchange factor for the whole furnace and one pyrometer at the door.

(a) What does the correction of case 8 do with the residuals such a furnace produces, and what is still wrong afterward?

(b) Name one measurement that would let a plant estimate a factor for each zone.

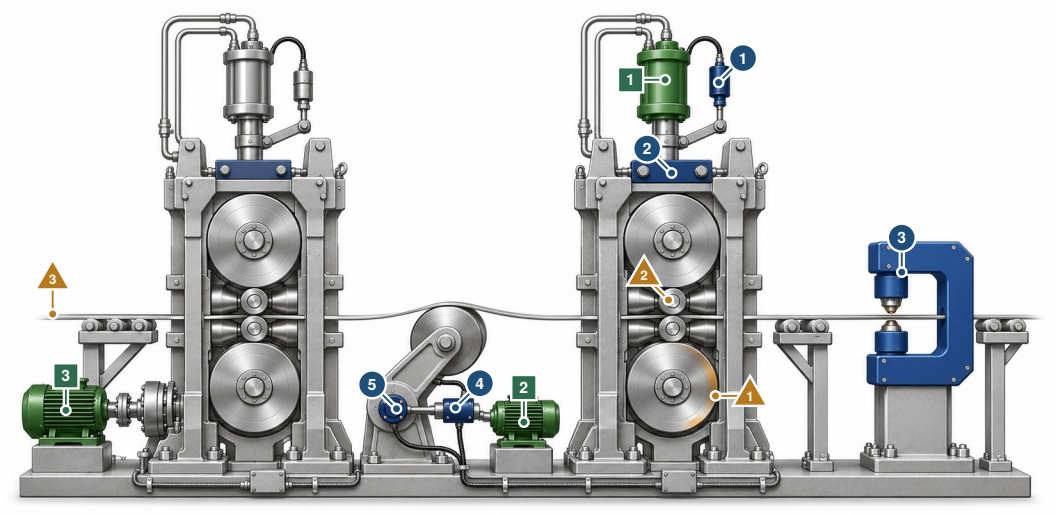
7.12 Review the agent · 5 min

An AI assistant reviews case 10 and writes: "With both corrections on, the joint predictive profile cuts the rms discharge error from 3.15 to 1.52 °C with all 117 slabs on target, for only 0.5 % more

fuel. Joint optimization of the three zones is clearly better than averaging per-slab offsets. Replace the slab model with the MPC for every shift, including the twelve-hour challenge.” Find the flaw and say how you would check it.

CHAPTER 8

Hot Strip Mill Simulator for Inferential Gauge Control With a Sign-Inverted Error



The controller moves 1 capsule that sets the F7 roll gap 2 looper motor that holds the strip tension 3 F6 drive, trimmed to hold the looper angle

The controller reads 1 capsule position 2 roll force 3 X-ray gauge, 5 m downstream 4 strip tension 5 looper angle

What disturbs it 1 back-up roll out of round 2 work rolls growing as they heat 3 incoming strip, thick patches and cold bands

Notice that capsule position and roll force are read at the stand at once, while the X-ray gauge sees the strip only half a second later. Illustration, not to scale.

IN PLAIN TERMS **Rolling Hot Steel to an Exact Thickness**

The main idea. A hot strip mill squeezes red-hot steel between heavy rolls into a thin sheet. The squeeze stretches the frame that holds the rolls, so the set gap is not the thickness that comes out.

The challenge. The thickness can be worked out at once from how hard the rolls push, or read half a second late by a meter down the line. The quick answer is right when the steel changes but points the wrong way when the rolls warm and grow, so the late meter must correct it.

Why it matters. Strip that is too thick or too thin sells for less or is rolled again.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	The first page, the two history boxes of Section 8.2, the fact box of Section 8.3, the fact and try-it boxes of Section 8.7, 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 8.7, Section 8.8, Section 8.9 and its open challenge

The control question A finishing stand can infer the strip thickness at once from its roll gap and roll force, or measure it half a second later with an X-ray gauge downstream. For which disturbances is the fast inferred gauge exact, for which is it wrong the other way, and what must sit underneath it?

Learning objectives

- Derive the gap gain and the deadbeat gain of a stand from its mill modulus and the plastic modulus of the pass.
- Sort a disturbance by the gaugemeter error identity into one the inferred gauge reads exactly and one it inverts or biases.
- Explain from the length-indexed dead time why a loop on the X-ray gauge cannot carry the gain of a loop on the gaugemeter.
- Set up a monitor trim on the X-ray gauge, with its action and its speed relative to the gauge loop.
- Choose a gauge control architecture for a coil by asking which measurement can see each disturbance.

Prerequisites Sections 1.4.6 (mill spring and gaugemeter), 1.4.5 (delay as a length), 1.3 (loop action), 1.6.2 and 1.6.4.

The app https://apmonitor.com/apps/hot-strip-mill/

See also Volume 5, Chapter 3 for a loop closed through a slow analyzer, Volume 5, Chapter 10 for a soft sensor trimmed by a delayed laboratory and Volume 5, Chapter 13 for an estimate corrected by late fixes.

8.1 Introduction

A property that a plant is paid for is often measured somewhere other than where it is made. A composition leaves the column before the analyzer sees it, a coating weight is read a frame downstream of the applicator, and a thickness is read after the metal has left the rolls. The standard answer is to compute the property from instruments that sit at the point of manufacture, run the loop on that computed number, and correct it slowly from the honest measurement when the honest measurement finally arrives [1–3]. The arrangement is old, it is everywhere, and the usual way of describing its weakness is that the inferred value drifts.

Drift is the wrong word for what goes wrong. An inferred measurement is built out of a model, and the model is exact for whatever it accounts for and blind to whatever it does not. For the blind class the error is not a slow random walk. It is a deterministic function of the disturbance, and it can carry the opposite sign to the change in the true property and be several times larger. A loop closed on such a measurement then amplifies that disturbance while holding its own measurement exactly on setpoint, and the operator sees a controller working perfectly.

A hot strip mill finishing stand is the clearest instance of this in industry, because both halves of the argument are arithmetic rather than statistics. The stand is a spring that deforms under the load it applies, so the delivered thickness is the commanded roll gap plus the roll force divided by the mill modulus. At the design pass of this chapter's simulator the gap is 1.000 mm, the stretch is 2.000 mm and two thirds of the product is machine deflection. Read the other way the same relation is a thickness measurement, computed from a position transducer and a load cell that answer in milliseconds, which is what the British Iron and Steel Research Association gaugemeter has been since the 1950s [4, 5]. The X-ray gauge that measures the strip itself stands 5.0 m past the stand, which at 10.00 m/s is 0.500 s of pure dead time, and on an entry thickness wave the two measurements at identical tuning give 11.6 µm and 173.8 µm of root-mean-square deviation.

The inversion is what the simulator is built to show. Every term of the gaugemeter except the modulus is measured at the stand, so any disturbance that arrives with the metal, a cold band, a thick patch on the transfer bar or a change in interstand tension, appears in the roll force and the inferred value matches the truth once the force has settled, straying by at most 1.5 µm while it moves. Anything that moves the roll surface without moving the position transducer does the opposite. Work rolls growing thermally, back-up rolls that are not round and a modulus calibration taken on

a different roll set all close the gap, raise the force, and are read as a thicker strip when the strip is getting thinner. The error is the whole of the gap offset, while the true gauge moves by only 0.2104 of it, so the error is 4.75 times the size of the real change and points the other way. Measured open loop, 200 µm of thermal crown moves the delivered gauge −41.9 µm and the gaugemeter reading +156.7 µm, an error of +198.6 µm. With the gaugemeter loop closed at the tuning that wins every other comparison in this chapter, the delivered gauge ends 198.6 µm thin, which is the entire crown.

The class generalizes past steel. Any plant that runs a fast estimator under a slow reference measurement inherits the same division of labor and the same exposure, including analyzer-trimmed property models, laboratory-corrected soft sensors and inertial platforms corrected by satellite fixes [2, 3]. On a moving-material line it inherits a second property, because every transport delay is a length divided by a speed and every periodic disturbance is a wavelength in delivered material rather than a period in seconds [6]. However, a student handed a soft sensor as a transfer function block never meets an upset whose estimated error points the wrong way. This chapter works with a hot strip mill simulator in which the same inferred gauge is exact for one class of disturbance and inverted for another, so that the choice of measurement appears as a hundred micrometers of product rather than as a remark about model mismatch. The mill spring, the roll force model and monitor gauge control are all standard practice. What the chapter adds for the reader is an exact error identity for the inferred gauge, used to order four control architectures against six disturbances, an ordering in which the architecture decides the outcome and the tuning barely moves it. Section 1.4.6 introduced the mill spring and the gaugemeter error on the running example W-1, and this chapter measures them on a full stand.

Section 8.2 traces gauge control on a finishing train from the screwdown to monitor gauge control. Section 8.3 specifies the thirteen-state two-stand model down to its implicit stand solution, its instrument chain and its length-indexed transport delays, with every constant and initial value, and Section 8.3.8 states where each simplification fails. Section 8.4 derives the error identity and sets the inferred gauge against the delayed X-ray gauge. Section 8.5 describes the five loops, the architecture switch and the eccentricity estimator. Section 8.6 lists the fifteen scenarios, which open on a roll gap moved by hand and close on one coil carrying five disturbances and a speed change. Section 8.7 reports what each scenario measures, and Section 8.8 sets the loops in their real hardware. Sections 8.9, 8.10 and 8.11 give the limitations, the lessons that transfer outside steel and a summary, and Section 8.12 holds the exercises.

8.2 Gauge Control on a Finishing Train

Strip was rolled to thickness for half a century before anything on a mill measured thickness. A rolling schedule fixed a screwdown position for each stand from the incoming bar, the ordered gauge and experience of the grade, the operator adjusted it by hand between coils, and the product was checked with a micrometer on the cold coil hours later. That arrangement needs a large margin, because strip below the ordered thickness cannot be shipped. The theoretical apparatus arrived first, with roll pressure through the arc of contact put on a defensible footing by the slab analysis of the early 1940s [7] and reduced to a form a mill engineer could evaluate for a hot pass a decade later [8]. Both give the force. Neither gives the thickness, because the force deflects the machine.

The step that turned the deflection from a nuisance into an instrument was taken at the British Iron and Steel Research Association, and the two history boxes of this section tell how the gaugemeter and the hydraulic capsule came about.

The second measurement arrived in parallel. Radiation gauges, first isotope and then X-ray, made an absolute traceable thickness available on the runout table, and because the reading describes strip that has already been made they were put to work correcting the inferred gauge rather than closing the fast loop. Monitor gauge control, in which the fast gaugemeter carries the regulation and the slow absolute gauge trims its bias, was standard on hot mills through the 1960s and remains the structure a

modern finishing train runs [5, 9]. The two measurements are complementary in exactly the way the frequency content of the disturbances requires, and no later technique has displaced the arrangement.

> ⧖ HISTORY **The Gaugemeter and Lock-On (1950s)**
>
> Before 1960 the British Iron and Steel Research Association developed automatic gauge control based on the load cell, and with it a gauge meter formula that would drive the later progress of gauge control [10]. In its first form the formula computed the exit thickness at the head end of each strip from the roll gap and the roll force, and the result was stored, a step called lock-on. The roll gap was then adjusted so that the thickness of the rest of the strip agreed with the stored value, which removed the unevenness along the rolling direction. The target was the head end's own thickness. Close study of the mill modulus in the 1960s then opened the way to the absolute mode of automatic gauge control [10].

Three refinements were added to that skeleton and all three are in the simulator. Mass-flow control uses the fact that the volume rate through each stand must match, so the thickness ratio across two stands is the speed ratio their drives have to run at, which keeps the looper from being emptied every time a gauge controller moves a gap. Looper and interstand tension regulation became a subject of its own, because the looper angle is a pure integral of the mass-flow mismatch and the tension it sets feeds back into the roll force and therefore into the gauge [6]. Roll eccentricity compensation answered the one disturbance neither gauge loop can reach, by demodulating the force signal against the measured back-up roll angle and subtracting the once-per-revolution component before the force is divided by the modulus [11]. A century of rolling theory is reviewed in [10], and the modeling detail behind the schedules, the mill spring calibration and the automatic gauge control architectures is collected in [12, 13].

What did not change is the exposure. The gaugemeter is a measurement of the mill, not of the strip, and it is only as good as the mill's calibration was on the day it was taken. Stand stiffness governs achievable rolling accuracy and is quoted as a chord over the working range, because the force against stretch curve is not linear at low load [14]. A modulus measured on one roll set, one width and one oil temperature is used on all of them until somebody retakes it.

> ⧖ HISTORY **The Hydraulic Screwdown (Late 1960s)**
>
> A rolling mill with a hydraulic screw-down device was developed at the end of the 1960s. It could change the mill modulus, and it entered tandem strip mills in the 1970s to raise the precision of gauge control [10]. Its control used a mechanical servo at first and later an electro-hydraulic servo, which moved the screw-down eight times faster than the conventional one. The speed brought a new problem. A fast roll gap loop and the interstand tension loop now responded at different rates and disturbed each other, and that interference opened the way to non-interference control of gauge and tension [10]. The capsule and servo valve of this chapter's F7 stand are of that kind.

8.3 Process Model

The simulated plant is stands F6 and F7 of a seven-stand finishing train, the looper between them, and the runout table as far as the X-ray gauge. 1200 mm of C-Mn strip enters F6 at 4.600 mm, leaves it at 3.600 mm, and leaves F7 at 3.000 mm and 880 °C, running 10.00 m/s at the F7 exit and 8.333 m/s into it. F6 takes 13.697 MN on a 1.400 mm gap and F7 takes 12.000 MN on a 1.000 mm gap, against mill moduli of 6.226 and 6.000 MN/mm. Interstand tension is 10.0 MPa, held by 13.535 kN m of looper torque at a 25° arm angle. The gauge tolerance is ±25 µm and the mill is paid by the meter inside it. The machine, its instruments and each symbol appearing below are drawn in Figure 8.1.

What the five loops reach, and what they do not, is sorted in Table 8.1. Units are those of a rolling mill, so thickness, gap and contact length appear in mm, roll force in MN, mill modulus in MN/mm, flow stress and strip tension in MPa, the strip path and the line speed in m and m/s, looper torque in kN m and temperature in °C. The clock is in seconds, because the transport delay sets the time scale

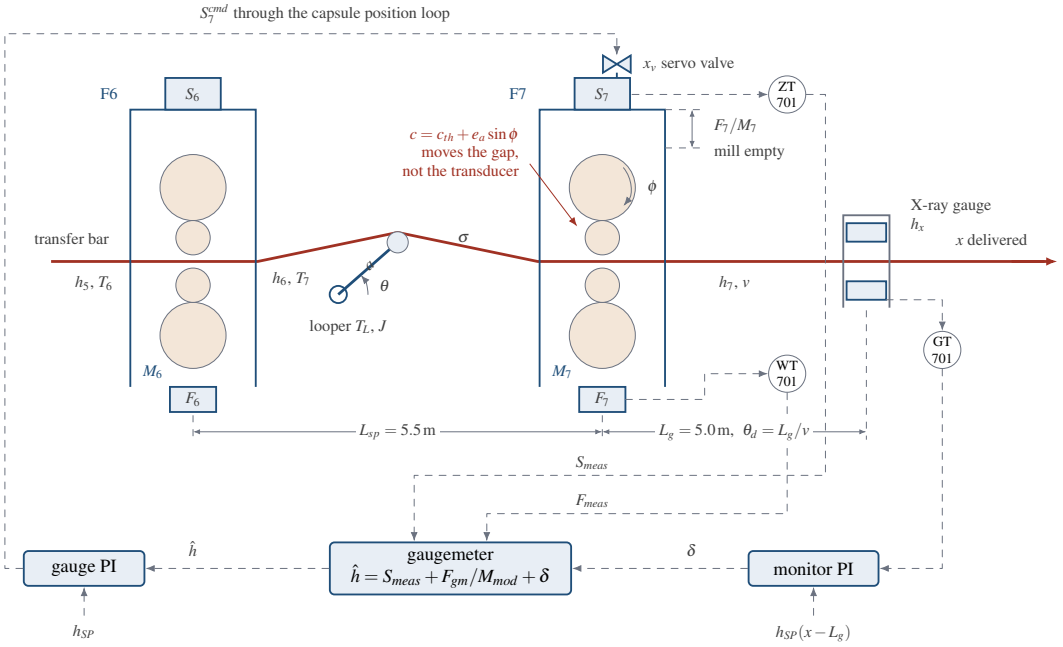


Figure 8.1 Stands F6 and F7 of a hot strip finishing train with the two rival thickness measurements. Each stand is a spring of modulus M_i standing F_i/M_i taller under load, so the delivered gauge h_7 is the commanded gap S_7 plus that stretch plus the gap offset c made by thermal crown c_{th} and runout $e_a \sin\phi$. Transducer ZT-701 and load cell WT-701 answer in milliseconds and the gaugemeter combines them into $\hat{h}$, while gauge GT-701 reads the strip itself L_g downstream for $\theta_d = L_g/v$ of transport delay. The looper sets the interstand tension σ, which subtracts from the roll force.

and the interesting events last tenths of a second. Thirteen states are integrated and every conversion between a meter and a millimeter is made once, at the point of use.

8.3.1 The Mill Spring and the Two Curves

A stand is an elastic structure loaded by the pass it is making. The housing, the screws, the chocks, the bearings and the oil column in the hydraulic capsule stretch together, and on a light hot gauge they stretch by more than the strip is thick. The delivered thickness is therefore the unloaded roll gap plus that stretch, and the roll force that sets the stretch is itself a function of the thickness, so the two have to be solved together as shown in Equation (8.1).

$$h_i = S_i + c_i + \frac{F_i(h_i)}{M_i}, \qquad i \in \{6,7\} \tag{8.1}$$

where h_i is the exit thickness in mm, S_i is the unloaded roll gap the position transducer reads, M_i is the mill modulus in MN/mm and c_i is everything that moves the roll surface without moving the transducer. F6 carries no such offset in this model, and F7 carries the one given in Equation (8.2).

$$c_7 = c_{th} + e_a \sin\phi \tag{8.2}$$

where c_{th} is roll thermal crown as a gap offset in mm, negative when the rolls grow into the strip, e_a is the back-up roll runout amplitude and ϕ is the back-up roll angle. Both terms are the subject of Section 8.4.

Equation (8.1) is the intersection of the elastic line of the stand with the plastic curve of the pass [4], and it cannot be marched. The plastic curve is nearly four times the steeper of the two, so

Table 8.1 Classification of Variables, With the Design Line-Up Every Scenario Opens On.

Symbol	Role	Nominal	Range
h_7	CV, unmeasured truth	3.000 mm	2.4–3.6 mm SP
$\hat{h}$, h_x	CV as inferred and as delayed	3.000 mm	—
S_7 (ZT-701)	CV of the capsule loop	1.000 mm	0.05–2.60 mm
σ	CV, measured	10.0 MPa	4–22 MPa SP
θ	CV, measured	25°	12–38° SP
S_7^{cmd}	MV of the gauge loop	1.000 mm	0.05–2.60 mm
x_v	MV of the capsule loop	0 %	±100 %
δ	MV of the monitor loop	0 mm	±0.25 mm
T_L	MV of the tension loop	13.535 kN m	0–32 kN m
r_6	MV of the angle loop	0 %	±5 %
v_{SP}	DV, operator set	10.0 m/s	4–14 m/s
S_6^{cmd}	DV, operator set	1.400 mm	0.6–2.4 mm
A_e, λ_e	DV, entry wave and its wavelength	0 %, 30 m	0–6 %, 8–120 m
ΔT_s, λ_s	DV, skid depth and pitch	0 °C, 125 m	0–60 °C, 40–300 m
e_a	DV, back-up roll runout	0 mm	0–0.080 mm
c_{tgt}	DV, thermal crown target	0 mm	−0.30–0.15 mm
ε_M	DV, modulus calibration error	0 %	±8 %
T_{bias}	DV, bulk temperature bias	0 °C	−50–40 °C
σ_x, σ_F	DV, gauge and load cell noise	0, 0	0–12 µm, 0–0.30 MN
x_v, S_7, S_6	states, valve and two capsules	0, 1.000, 1.400 mm	±1, 0–4 mm
v, r_6	states, the two drives	10.00 m/s, 0	0.3–14 m/s, ±0.05
x, ϕ	states, strip and roll angle	0, 0	m, 0–2π
θ, $\dot{\theta}$	states, looper arm	25°, 0	3–45°
σ, T_L	states, tension and torque	10.0, 13.535	0–60 MPa, 0–32
c_{th}, T_b	states, the two slow drifts	0, 0	as their disturbances

substituting the force back into the right-hand side amplifies an error by that factor on every pass and successive substitution diverges immediately. The application solves it by Newton iteration on the residual $r(h) = h - S - c - F(h)/M$, with the slope re-evaluated by a forward difference at every pass. The derivative $r'(h) = 1 + Q/M \approx 4.8$ is bounded away from zero and never changes sign, and the iteration stops at a residual of 10^{-13} mm. The solve is called eight times per physics step, four Runge-Kutta stages on two stands, and from the previous step's answer as a warm start it converges in two or three iterations.

> **DID YOU KNOW? Most of the Stretch Is in the Rolls**
>
> Of the elastic deformation of a rolling mill under load, the rolls take 40 to 70 %, the housing 10 to 16 % and the screw-down device 4 to 20 %. Bending and flattening of the rolls is by far the largest share [10].

8.3.2 Roll Force and Flow Stress

The separating force follows the slab treatment of the roll bite [7, 8], with the mean pressure taken as the plane-strain flow stress lifted by a Coulomb friction hill and lowered by both strip tensions, as shown in Equation (8.3).

$$F = w L_c \left[1.1547\, k_f - \tfrac{1}{2}(\sigma_b + \sigma_f)\right] Q_p, \quad L_c = \sqrt{R' \Delta h} \tag{8.3}$$

$$Q_p = \frac{h_m}{\mu L_c}\left(\mathrm{e}^{\mu L_c / h_m} - 1\right) \tag{8.4}$$

where w is the strip width in mm, L_c is the contact arc length, $R' = R\,f_H$ is the work roll radius after Hitchcock flattening, Δh is the draft, h_m is the mean of the entry and exit thicknesses, μ is the roll-strip friction coefficient, and σ_b and σ_f are the back and front tensions. Millimeters times millimeters times N/mm^2 is newtons, so a factor of 10^{-6} converts the product to MN. At the design pass the arc is 16.72 mm against a mean thickness of 3.30 mm, which puts Q_p at 2.350, so the friction hill is most of the reason a 3 mm pass costs twelve meganewtons. The bracket in (8.3) is floored at 15 MPa so that a large tension cannot drive the computed pressure negative.

Flow stress is the reason a hot strip is a different material every second, as shown in Equation (8.5).

$$k_f = k_0\,\mathrm{e}^{-\beta(T-T_0)}\left(\frac{\max(1,\dot{\varepsilon})}{\dot{\varepsilon}_0}\right)^{m}, \quad \dot{\varepsilon} = \frac{\bar{v}}{L_c}\ln\frac{h_{in}}{h_{out}} \tag{8.5}$$

where k_0 is the flow stress at the reference temperature T_0 and reference strain rate $\dot{\varepsilon}_0$, β is the softening coefficient, m is the strain rate exponent, and $\bar{v}$ is the mean of the entry and exit strip speeds through the arc, converted to mm/s at the point of use. The temperature term carries the disturbance that matters. A skid mark 35 °C deep raises the flow stress by 16 %, and with both gaps frozen it puts 134.0 µm of extra thickness on the product.

8.3.3 The Two Plant Gains

Differentiating (8.3) at the operating point gives the plastic modulus, and combining it with (8.1) gives the two numbers every gauge argument in this chapter turns on, as shown in Equation (8.6).

$$Q = -\frac{\partial F}{\partial h}, \qquad \frac{\partial h}{\partial S} = \frac{M}{M+Q}, \qquad \frac{\partial h_{out}}{\partial h_{in}} = \frac{\partial F/\partial h_{in}}{M+Q} \tag{8.6}$$

At the design pass $Q = 22.513\,\mathrm{MN/mm}$ against a mill modulus of 6.000 MN/mm, so a gap move arrives at the product attenuated to 0.2104 of itself and an entry thickness variation walks through at 0.6534, both checked numerically by the application's verification suite. The reciprocal of the first, $(M+Q)/M = 4.752$, is the deadbeat gain, the gap move that cancels a gauge error in one step, and the shipped gauge loop is tuned just below it at 4.00. A mill is a soft spring against a stiff strip, which is why no hot mill has ever been run on gap position alone.

8.3.4 Actuators, Drives and the Looper

The F7 capsule is a velocity device. A servo valve spool follows the valve command through a short lag and the capsule integrates the spool position, so the position loop closes around an integrator and needs almost no reset action of its own, as shown in Equations (8.7) and (8.8).

$$\tau_v\frac{\mathrm{d}x_v}{\mathrm{d}t} = \frac{\mathrm{clamp}(u_v,\pm 100)}{100} - x_v \tag{8.7}$$

$$\frac{\mathrm{d}S_7}{\mathrm{d}t} = K_q\,x_v, \qquad \tau_{g6}\frac{\mathrm{d}S_6}{\mathrm{d}t} = S_6^{cmd} - S_6 \tag{8.8}$$

where x_v is the normalized spool position, u_v is the valve command in percent and K_q is the capsule slew rate at full valve in mm/s. F6 has no gauge control in this application, so its capsule is represented as a closed position loop with a single time constant. The drives, the delivered strip coordinate and the back-up roll angle follow Equation (8.9).

$$\begin{aligned} \tau_d\frac{\mathrm{d}v}{\mathrm{d}t} &= v_{SP} - v, & \tau_{s6}\frac{\mathrm{d}r_6}{\mathrm{d}t} &= \frac{r_6^{cmd}}{100} - r_6 \\ \frac{\mathrm{d}x}{\mathrm{d}t} &= v, & \frac{\mathrm{d}\phi}{\mathrm{d}t} &= \frac{v}{(1+f_s)R_b} \end{aligned} \tag{8.9}$$

where x is the strip delivered past F7 in meters, R_b is the back-up roll radius in meters and f_s is the forward slip at the roll exit. At 10.00 m/s the back-up roll turns at 13.87 rad/s, which is 2.207 Hz, and that frequency moves with the line speed while the runout amplitude does not.

The looper is a torque balance about its pivot and the interstand span is an elastic member, as shown in Equation (8.10).

$$\begin{aligned} J\,\ddot{\theta} &= 10^{3}\,(T_L - T_{grav} - T_{strip}) - c_L\,\dot{\theta} \\ \tau_T\,\frac{\mathrm{d}T_L}{\mathrm{d}t} &= T_L^{cmd} - T_L \\ \frac{\mathrm{d}\sigma}{\mathrm{d}t} &= K_s\left(v_{in,7} - v_6 + \frac{\mathrm{d}L}{\mathrm{d}\theta}\,\dot{\theta}\right) \end{aligned} \tag{8.10}$$

where J is the arm inertia in kg m^2, c_L is the hydraulic damping, $T_{grav} = W\cos\theta$ is the arm weight torque and T_{strip} is the torque the strip pull applies. The factor 10^3 converts the torques from kN m to N m for the inertia. The geometry is two straight runs from each roll bite up to a looper roll sitting $a\sin\theta$ above the pass line, as shown in Equation (8.11).

$$\begin{aligned} \ell &= \sqrt{a^2\sin^2\theta + (L_{sp}/2)^2}, \qquad L = 2\ell \\ T_{strip} &= \frac{2\,\sigma\,h_6\,w\,(a\sin\theta/\ell)\,a\cos\theta}{10^{3}} \end{aligned} \tag{8.11}$$

where K_s is the span stiffness in MPa/m of length mismatch, $v_{in,7} = v\,h_7/h_{in,7}$ is the strip speed into F7, with $h_{in,7}$ the F7 entry thickness of Equation (8.15), and $v_6 = v\,\rho_{mf}(1 + r_6)$ is the F6 exit speed. Between its 3° and 45° stops the arm stores 115 mm of strip, of which only 41 mm lies below the 25° design angle. At 10.00 m/s that is one hundredth of a second of production, which is why the looper angle is never left to an operator and why the F6 drive is slaved to the thickness ratio through the mass-flow lag ρ_{mf} of Equation (8.12).

$$\tau_{mf}\,\frac{\mathrm{d}\rho_{mf}}{\mathrm{d}t} = \mathrm{clamp}\left(\frac{h_7}{h_{in,7}},\,0.4,\,1.0\right) - \rho_{mf} \tag{8.12}$$

8.3.5 What Arrives, and What Drifts

Two disturbances arrive with the metal and both are written as functions of delivered strip length rather than of time, because that is the coordinate the X-ray gauge, the coiler and the customer's complaint all live in, as shown in Equation (8.13).

$$\begin{aligned} h_5(x) &= h_5^{0}\left[1 + 0.01\,A_e\sin\left(\frac{2\pi x}{\lambda_e}\right)\right] \\ T_7(x) &= T_0 + T_b - \Delta T_s\,\mathrm{e}^{\,\kappa[\cos(2\pi x/\lambda_s)-1]} \end{aligned} \tag{8.13}$$

where A_e is the transfer bar thickness variation in percent of 4.600 mm, λ_e and λ_s are wavelengths in delivered meters, ΔT_s is the skid mark depth and $\kappa = 8$ makes the wrapped Gaussian about three times narrower than its own pitch while staying smooth enough for a fixed-step integrator. F6 rolls material that will be delivered L_{st} meters later, so its entry thickness is evaluated at $x + L_{st}$ and its temperature is the F7 profile at the same place plus ΔT_{67}. A skid pitch of 1.5 m on the slab becomes 125 m on 3 mm strip, which arrives every 12.5 s at ten meters a second and is therefore far slower than roll runout and separable from it.

The two slow drifts are first-order approaches to their targets, as shown in Equation (8.14).

$$\tau_c\,\frac{\mathrm{d}c_{th}}{\mathrm{d}t} = c_{tgt} - c_{th}, \qquad \tau_b\,\frac{\mathrm{d}T_b}{\mathrm{d}t} = T_{bias} - T_b \tag{8.14}$$

The work roll thermal time constant τ_c of 40 s is what makes the drift slow enough for a slow loop to catch, which is the whole design of monitor gauge control.

8.3.6 Instruments and the Length-Indexed Delay

Four measurements feed the controllers. Each is a first-order lag stepped with the exact discrete update $z_k = \alpha z_{k-1} + (1-\alpha)u_k$, where $\alpha = \mathrm{e}^{-\Delta t_c/\tau}$, so no scan period can destabilize it. The load cell carries $\tau_{load} = 0.004\,\mathrm{s}$, the position transducer $\tau_{pos} = 0.002\,\mathrm{s}$, the X-ray gauge $\tau_{xray} = 0.030\,\mathrm{s}$ and the tensiometer $\tau_{sig} = 0.020\,\mathrm{s}$. The position transducer reads the capsule, which is above the whole roll stack, and that single fact is why roll runout is invisible to it. Gauge and load cell noise enter as first-order filtered Gaussian sequences, scaled to put the stationary standard deviation on the slider value. Their generator is seeded at 20260914, which makes a noisy run repeat bit for bit.

Two transport delays exist in the model, the X-ray gauge's and the span between the stands, both keyed to delivered strip length, as shown in (8.15).

$$h_x = h_7(x - L_g)\,, \quad \theta_d = \frac{L_g}{v}, \quad h_{in,7} = h_6(x - L_{st}) \tag{8.15}$$

Each delay line is a ring buffer keyed on x with linear interpolation between samples, and the span line starts full of strip rolled at the starting conditions. Indexed on time it would be right at 10.00 m/s and quietly wrong at every other speed the mill runs, whereas indexed on length the dead time follows every acceleration for free. Chapter 10 measures the same distinction for a predictor on a coating line. On the gap step of scenario 1 the application measures 0.525 s of lag at 10 m/s and 0.860 s at 6 m/s between the true and reported gauges crossing half of the settled step. That is L_g/v plus about 25 ms, which the 0.030 s averaging time of the gauge and the 0.004 s scan at which the delay line is written and read add to the transport delay.

8.3.7 Parameters, Initialization and Numerics

Table 8.2 holds every constant. The model dialog shows and adjusts all but seven of them, and those seven carry a hidden mark, meaning their values are never displayed and have to be worked out from behavior instead.

The mill modulus deserves a sentence of its own, because 6.00 MN/mm looks small beside the figures a rolling engineer carries in the hundreds. Those figures are the same number in different units, since 6.00 MN/mm is 612 tf/mm of force per millimeter of stretch, and the check that matters is the stretch the value produces. At 12.000 MN the stand stands 2.000 mm taller than when it is empty, which is two thirds of a 3.000 mm product and is why the commanded gap is 1.000 mm. A stiffer stand would pass a smaller fraction of the entry variation and would need a smaller deadbeat gain [14]. F6 carries 6.2261 MN/mm and stretches 2.200 mm under 13.697 MN.

No scenario starts away from equilibrium. Bisection on Equation (8.1) at the scheduled exit thicknesses gives the two roll gaps, returning 1.4000 mm at F6 and 1.0000 mm at F7, and the looper torque is initialized to the 13.535 kN m that balances the arm weight and the strip pull at 25°, of which 5.891 kN m lifts the arm itself. The application's verification suite re-solves those three constants and fails if they have drifted. States advance on a 0.002 s fixed step under classical Runge-Kutta of fourth order, controllers execute on a 0.004 s period, and the historian keeps 50 samples for each simulated second. The engine, its control period and its speed setting are those of Section 1.6.1, and a different speed rescales simulated time against the wall clock and nothing else.

Four conditions raise an annunciation without stopping the run, and each clears only after the condition has receded by a margin. Interstand tension below 1.5 MPa is a loop-out, above 28 MPa is necking, roll force above 20.0 MN is past the stand rating, and a roll gap driven to zero leaves the capsule with nothing to give. Tension is clamped to 0 and 60 MPa and the looper angle to its two stops, with the angular rate zeroed at a stop so the integrator cannot pass through it.

8.3.8 Assumptions

The model is detailed along the strip and absent across it, and each simplification fails somewhere.

One dimension. There is no strip width profile, no roll crown across the barrel, no flatness and no work roll bending, shifting or pair-crossing. On a real mill the shape across the strip takes at least as

Table 8.2 Model Parameters, in Two Panels. Hidden Constants Are Never Displayed and Have to Be Inferred From the Plant's Behavior.

Symbol	Meaning	Value	Unit	Shown
w	strip width	1200	mm	adjustable
R	work roll radius	340	mm	adjustable
f_H	roll flattening factor	1.37	—	adjustable
μ	roll-strip friction	0.30	—	adjustable
R_b	back-up roll radius	0.70	m	adjustable
F_{max}	stand force limit	20.0	MN	adjustable
L_g	F7 to the X-ray gauge	5.0	m	adjustable
L_{sp}	F6 to F7 stand centers	5.5	m	adjustable
h_5^0	F6 entry thickness	4.600	mm	adjustable
h_6^0	F6 exit thickness	3.600	mm	adjustable
h_7^0	F7 exit thickness	3.000	mm	adjustable
T_0	strip temperature at F7	880	°C	adjustable
ΔT_{67}	F6 hotter than F7	25	°C	adjustable
k_0	flow stress, T_0 and $\dot{\varepsilon}_0$	224.79	MPa	adjustable
β	flow stress softening	0.0045	1/°C	adjustable
m	strain rate exponent	0.13	—	adjustable
K_s	span tension stiffness	300	MPa/m	adjustable
M_7	F7 mill modulus	6.00	MN/mm	adjustable
M_6	F6 mill modulus	6.2261	MN/mm	adjustable
τ_v	servo valve response	0.006	s	adjustable
K_q	capsule slew, full valve	9.0	mm/s	adjustable
τ_{g6}	F6 capsule response	0.030	s	adjustable
τ_d	main drive response	0.50	s	adjustable
τ_{s6}	F6 speed trim response	0.15	s	adjustable
τ_{mf}	mass-flow lag	0.25	s	adjustable
τ_T	looper torque response	0.08	s	adjustable
τ_c	roll thermal constant	40	s	adjustable
τ_b	bulk temperature lag	6.0	s	adjustable
J	looper inertia	600	kg m^2	adjustable
a	looper arm length	0.80	m	adjustable
c_L	looper damping	5000	N m s	adjustable
W	looper gravity torque	6.50	kN m	adjustable
τ_{load}	load cell response	0.004	s	adjustable
τ_{pos}	transducer response	0.002	s	adjustable
τ_{xray}	X-ray averaging time	0.030	s	adjustable
b	gauge tolerance	0.025	mm	adjustable
b_σ	tension tolerance	2.0	MPa	adjustable
σ_{lo}	strip loops out below	1.5	MPa	adjustable
σ_{hi}	strip necks above	28	MPa	adjustable
f_s	forward slip at the exit	0.03	—	hidden
L_{st}	F6 to F7, delivered	6.6	m	hidden
$\dot{\varepsilon}_0$	reference strain rate	100	1/s	hidden
σ_{56}	F5 to F6 back tension	10.0	MPa	hidden
θ_{min}	looper bottom stop	3	°	hidden
θ_{max}	looper top stop	45	°	hidden
τ_{sig}	tensiometer response	0.020	s	hidden

much control effort as the thickness along it, and the two interact through the roll stack, because a gauge correction changes the load, which changes the roll deflection, which changes the crown [13].

The mill spring is one linear modulus per stand. A real force against stretch curve is nonlinear at low load where the chocks are still bedding in, which is why mill modulus is quoted as a chord over the working range [14]. The modulus error disturbance makes the point that the modulus is a calibration, but the linearity is an idealization.

Hitchcock flattening is frozen at 1.37. Real flattening grows with load, so the model understates the plastic modulus far from the design point, although within the working range the two gains of (8.6) are accurate.

The strip is one temperature. Through the thickness and across the width the strip is isothermal, where in reality the surface is far below the core after descaling and the deformation reheats it. The profile is written in delivered meters rather than emerging from a heat balance down the mill.

No mill resonance. Stands, drives and the roll stack are first-order lags, so the third-octave chatter mode near 100 Hz to 200 Hz that a real gauge controller must stay clear of does not exist, and neither does torsional coupling between the two stands through the strip.

The mass-flow link is idealized. F6's speed follows the true thickness ratio through a 0.25 s lag, where a real scheme forms that ratio from the X-ray gauge and the upstream gaugemeters. Feeding the compensation a biased gaugemeter turns a gauge bias into a permanent speed error and parks the looper on a stop, which is not what the application is about.

The eccentricity estimator is handed a perfect roll angle. It assumes one sinusoid on one roll, where a real stand has four rolls at three frequencies plus harmonics and the estimator has to acquire the phase rather than being given it [11].

Two stands, not seven. There is no threading, no head end, no tail out and no coiler, so the transient a real gauge control system spends most of its effort on, the first meters of a coil, is absent.

Three disturbances are stylized. A thermal crown arriving on a clean first-order approach to exactly 200 µm is not what a real roll set does, but it isolates the gap-side disturbance and makes the steady state of Equation (8.17) readable off a trend. A modulus error applied as a step is equally artificial, because a calibration goes stale between roll changes rather than during a coil, although it separates a pure measurement bias from every dynamic effect in the plant. The skid marks are one wrapped Gaussian of a single pitch and depth, and that regularity is what makes their 12.5 s arrival interval measurable.

Repeatability itself is unrealistic. Freezing the step and the scan period yields an identical trajectory on any machine, whereas a distributed control system in a mill carries timing jitter and an execution order that shifts.

8.4 The Two Measurements

8.4.1 An Exact Error Identity

The gaugemeter rearranges Equation (8.1) into a thickness, using the position transducer, the conditioned load cell signal and a modulus that is a calibration rather than a constant of nature, as shown in Equation (8.16).

$$\hat{h} = S_{meas} + \frac{F_{gm}}{M_{mod}} + \delta, \qquad M_{mod} = M_7 \left(1 + \frac{\varepsilon_M}{100}\right) \tag{8.16}$$

where ε_M is the modulus calibration error in percent and δ is the monitor trim of Section 8.5. Subtracting the truth of (8.1) from (8.16) leaves the whole argument of the chapter in one line, as shown in Equation (8.17).

$$\hat{h} - h_7 = -c_7 + F_7 \left(\frac{1}{M_{mod}} - \frac{1}{M_7}\right) + \delta \tag{8.17}$$

Three properties of (8.17) carry the application. Nothing on the right depends on the strip, so every disturbance that arrives with the metal, a thick patch on the transfer bar, a skid mark, a cold slab or a tension change, leaves the error unmoved and the inferred gauge is exact. Measured open loop

Table 8.3 Open-Loop Response to Six Disturbances, Roll Gap Held. The Rule That Separates the Two Blocks Is Whether the Disturbance Reaches the Strip or the Roll Gap. The Skid and Cold-Strip Errors Are Read at the Band Bottom and in Steady State, and the Entry-Wave Error Is the Lag of the Force Filters.

Disturbance	Δh_7 (μm)	$\hat{h} - h_7$ (μm)	Sign
Arrives with the metal			
Skid mark, 35 °C deep	+134.0	−0.04	—
Entry wave, 3 % of 4.6 mm	±51.2	±1.6	—
Cold strip, −30 °C	+115.5	0.00	—
Moves the roll gap			
Thermal crown, −200 μm	−41.9	+198.6	inverted
Runout, 40 μm amplitude	±8.6	∓39.6	inverted
Modulus error, +5 %	0.0	−95.2	bias

at the bottom of a 134.0 μm skid mark the gaugemeter error is 0.04 μm, and on a whole strip 30 °C cold, which moves the product 115.5 μm, it is zero to the digits shown once the strip has settled. Only while the force is moving does the 14 ms of load cell and gaugemeter filtering open a gap, at most 1.6 μm on any strip-side row of Table 8.3. Second, a gap offset enters at full size with a minus sign while the true gauge moves by only $M/(M+Q)$ of it, so the ratio of error to real change is $-(M+Q)/M = -4.75$. Third, a modulus error is a pure bias proportional to the force, invisible at the stand and constant for the coil, which is what the monitor trim δ exists to cancel and to display.

Table 8.3 separates the two halves on the running model. Every row is open loop with the roll gap held, so the entries are properties of the plant and the instruments rather than of any controller. Each row is the scenario that carries the disturbance with Gauge control switched to manual before the run, and a steady disturbance is read once it has settled.

The crown row is the one worth checking against the algebra. Equation (8.17) says the error should be exactly the negative of the gap offset, and at 200 s after the step the crown state has reached −198.65 μm and the measured error is 198.65 μm. The true gauge has moved −41.91 μm, so the measured ratio is -4.741 against the -4.752 that (8.6) predicts, the small difference being the rise in Q as the pass gets heavier. The modulus row is (8.17) evaluated directly, since $12\,\text{MN}\,(1/6.30 - 1/6.00) = -95.2\,\mu\text{m}$.

🕮 **Pause.** Try Exercise 8.2 before reading on.

8.4.2 Where the Two Curves Cross

Figure 8.2 draws Equation (8.1) the way a mill engineer reads it, as the intersection of the elastic line of the stand with the plastic curve of the pass, and puts the crown disturbance on the same axes.

Reading the figure explains the mechanism in one sentence. Crown slides the mill line left, the intersection moves down the plastic curve to a thinner strip and a higher force, and the gaugemeter converts that higher force back into a thickness along the line the mill no longer sits on. A controller that believes the reported 3.1586 mm shuts the gap until the reported value is 3.000 mm again, which by Equation (8.17) is precisely the condition $h_7 = h_{SP} + c_7$. The steady state of a gaugemeter loop under a gap offset is therefore a delivered gauge low by the entire offset, exactly, and the measured result at 210 s is −198.6 μm against a crown of −198.65 μm.

🕮 **Pause.** Try Exercise 8.1 before reading on.

8.4.3 What the Five Meters Cost the Other Loop

The X-ray gauge has no such defect. It measures the product, it is absolute and traceable, and its error under every disturbance in Table 8.3 is zero once the strip reaches it. What it has instead

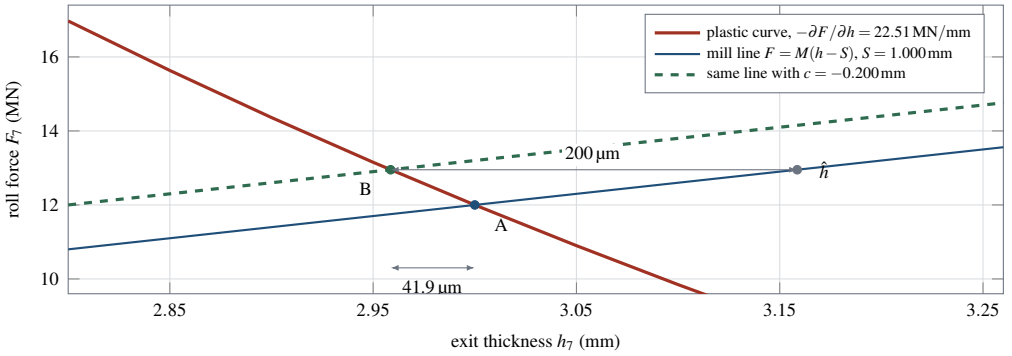


Figure 8.2 Mill line and plastic curve at the F7 design pass, and what 200 µm of thermal crown does to both. The plastic curve is the thick solid line, the mill line the thin solid line and the mill line under crown the dashed line. Crown displaces the mill line bodily to the left without moving the position transducer, so the intersection moves from A to B, 41.9 µm thinner at 12.951 MN, while the gaugemeter reads that force back along the undisplaced line and reports 3.1586 mm. The gap between B and $\hat{h}$ is the whole of the crown, the wrong way up.

is 0.500 s of transport delay at line speed, and dead time is the one plant property no tuning rule removes [15, 16]. A pure delay θ_d in series with the static gap gain of (8.6) reaches −180° of phase at a period of $2\theta_d = 1.0$ s where the loop gain is the reciprocal of the gap gain, so the ultimate gain is $1/0.2104 = 4.75$. Adding the capsule, the X-ray averaging filter and reset action moves both numbers, and at $K_c = 4.00$ and $\tau_I = 0.25$ s the delayed loop cycles at 1.374 s rather than 1.0 s, or at 2.09 s when 6 m/s makes θ_d 0.833 s. Neither is a property of the controller. Both are five meters of runout table divided by a speed.

That is the whole of the trade. One measurement is exact for the strip, blind to the mill and available in four milliseconds, and the other is true about everything and half a second late. Monitor gauge control runs both and splits the job by frequency, with the gaugemeter above about one radian a second where it is the only measurement that exists and the X-ray below it where it is the only measurement that is right.

8.5 Controllers as Implemented

8.5.1 The Shared Algorithm

All five loops run the velocity-form PID of Section 1.6.2, Equations (1.13) to (1.16), with back-calculation on the integral alone and the tracking constant left at its default of τ_I. The action a takes the value +1 on a reverse-acting loop and −1 on a direct-acting one, in the convention of Section 1.3. Here the scan is $\Delta t_c = 0.004$ s and the output bias is zero throughout. Every loop sets $\tau_D = 0$, which idles the derivative path and its measurement filter. Nor does any loop impose a rate limit, since the capsule slew K_q, the valve lag τ_v and the torque lag τ_T already supply the real ones from the machinery. Selecting manual solves for the integral against the operator's actuator value, which makes both transfers bumpless.

8.5.2 The Five Loops and Their Directions

Table 8.4 lists what each loop reads, what it writes and which way it acts. Two of the directions are worth stating in words, because they are the two that get reversed in practice.

Opening the gap thickens the strip, so the gauge loop is reverse acting. Its gain of 4.00 is set just below the deadbeat value of 4.752 from (8.6), with the remainder held back for the phase the capsule and the load cell filter cost. The monitor loop is direct acting, and this is the sign that gets flipped. Its output is a bias added to the gaugemeter's reading, and raising a measurement drives the loop

Table 8.4 The Five Implemented Loops, With the Tuning Every Scenario Except 4, 14 and 15 Opens On.

CV	MV	K_c	τ_I (s)	Limits	a
$\hat{h}$ or h_x	S_7^{cmd}	4.00	0.25	0.05–2.60	+1
h_x	δ	0.60	2.0	±0.25	−1
S_7	x_v	500	1.00	±100	+1
σ	T_L	0.300	0.80	0–32	+1
θ	r_6	0.120	4.0	±5	+1

underneath it to shut the gap and roll thinner, so the manipulated variable moves the process variable down. Getting that one backwards does not oscillate. It walks the gauge quietly off specification over half a coil, which is worse, and the application's verification suite tests the direction of all five loops by perturbing each and checking the sign of the response.

🕮 **Pause.** Try Exercise 8.3 before reading on.

The capsule loop closes around the integrator of (8.8), so it is nearly a pure proportional loop, and 500 %/mm against $K_q = 9.0\,\mathrm{mm/s}$ at full valve places the closed capsule at 45/s. It is in service in manual mode as well, because it is the servo that makes the capsule a position device at all and what manual changes is only who writes its setpoint. The looper angle loop is in service in both modes for the same reason, since the angle is a pure integral of the mass-flow mismatch and one percent of F6 speed swings the arm 27 °/s. Manual mode releases only the tension loop and the gauge loop.

8.5.3 The Architecture Switch

The gauge loop reads one of three numbers, selected by a four-position control that is the whole argument of the application. In *manual* the gap is whatever the operator typed, in *X-ray only* the loop reads h_x, in *gaugemeter* it reads (8.16) without the trim, and in *gaugemeter with monitor* it reads (8.16) in full with the monitor loop in service. The monitor setpoint is h_{SP} delayed as in (8.15) and averaged as the X-ray gauge is. When the monitor loop is out of service its integral tracks a zero output, so switching it in is bumpless.

8.5.4 The Synchronous Eccentricity Estimator

Neither gauge loop can reach 2.207 Hz, which is fourteen radians a second against a delayed loop that stops at one and a fast loop that would have to chase a signal of the wrong sign. What the mill has instead is independent knowledge of where the back-up roll is, from a shaft encoder, and two integrators in the rotating frame turn that into an estimate of the once-per-revolution force component, as shown in Equation (8.18).

$$
\begin{aligned}
F_{ecc} &= e_a^c \cos\phi + e_b^s \sin\phi \\
\frac{\mathrm{d}e_a^c}{\mathrm{d}t} &= 2K_{ecc}\,(F_{sig} - \bar{F} - F_{ecc})\cos\phi \\
\frac{\mathrm{d}e_b^s}{\mathrm{d}t} &= 2K_{ecc}\,(F_{sig} - \bar{F} - F_{ecc})\sin\phi \\
F_{gm} &= F_{sig} - F_{ecc}
\end{aligned}
\tag{8.18}
$$

where F_{sig} is the load cell signal after the adjustable force filter τ_F, $\bar{F}$ is a first-order average with a 1.0 s time constant, removed first so the estimator locks onto the ripple rather than spending its first second chasing twelve meganewtons of steady load, and K_{ecc} is a convergence rate in reciprocal seconds. The factor of two undoes the mean of $\cos^2$, which is what makes K_{ecc} a rate rather than an arbitrary number. Setting K_{ecc} to zero switches the estimator off and zeroes both states. At $K_{ecc} = 3/\mathrm{s}$ in scenario 12 the estimate reaches 90 % of its final amplitude 0.19 s after it is switched in and settles at 0.177 MN, against the 0.189 MN the runout is actually worth in the open-loop force.

The alternative offered alongside it is the force filter τ_F, and the comparison fails instructively, because a low-pass slow enough to remove 2.207 Hz also removes the entry thickness wave at 0.333 Hz that the whole scheme exists to catch. With $\tau_F = 0.20$ s in place of the estimator, a 3 % entry wave fired on top of the runout of scenario 12 reaches 63.6 µm of gauge error, against 43.9 µm with the estimator in service. The estimator removes energy at one frequency the mill has independent knowledge of and leaves the rest of the spectrum alone. The phasor canceller of Chapter 6 removes strand bulging in the same way, at a frequency computed from the casting speed and the roll pitch.

8.5.5 Three Models Kept Apart

Equations (8.1) to (8.15) are the simulated plant, thirteen integrated states carrying two implicit algebraic solves. Whatever model the controller holds is far smaller, because the gauge loop carries nothing beyond a gain and a reset time and its tuning encodes one plant number, the gap gain $M/(M+Q) = 0.2104$, whose reciprocal is the deadbeat gain. The gaugemeter is a third thing again, a measurement model consisting of one constant M_{mod} whose disagreement with the plant's M_7 is the modulus term in (8.17). Confusing the three is easy here, because the symbol M appears in the plant, in the measurement and in the tuning rule and means something different each time. The mill modulus M_7 is a property of the machine, M_{mod} is a number in a calculation, and $K_c = 4.00$ is a controller gain in millimeters of gap per millimeter of gauge error.

8.6 Scenarios

Table 8.5 lists fifteen scenarios ordered the way a laboratory session would run them. A scenario pins the opening mode, the architecture, the five tunings, every disturbance setting and any timed event, then puts a question before the run and a measurement after it. Three tuning sets recur. The shipped set is $K_c = 4.00$ and $\tau_I = 0.25$ s on the gauge loop, the detuned set that a delayed measurement forces is $K_c = 1.2$ and $\tau_I = 1.5$ s, and the deliberately poor set that opens the challenge is $K_c = 0.8$, $\tau_I = 2.5$ s, a monitor trim too slow to converge inside a coil and a force filter of 0.20 s.

Each scenario also carries a completion rule, condensed in the last column of Table 8.5, and credit, Reset and the challenge verdict work as Section 1.6.4 describes. A run that needs a different architecture or tuning is repeated from Reset. The challenge is judged on the coil as given, so changing a model parameter or a disturbance voids the attempt. The scenario names in the table are the ones the application shows, and all but scenarios 5 and 15 are experiments.

8.7 Results

Each figure below comes off the running application. The exported record sets the true gauge beside both measurements, the capsule position beside its command, the live plastic modulus and gap gain, and the running yield and deviation metrics, which is enough to redo any measurement here. Root-mean-square gauge deviation and meters off gauge accumulate with the length rolled, so both are reported together with the interval over which they were taken. Peak deviations, capsule positions and limit cycle periods carry no such dependence. Each run loads its scenario, makes any listed change before Run and runs from $t = 0$, and peaks, crossings and periods are read at the 2 ms physics step.

8.7.1 Two Thirds of the Product Is Mill Spring

Scenario 1 settles the question with no controller in the picture. Opening the F7 gap from 1.000 to 1.250 mm moves the delivered gauge from 3.0000 to 3.0546 mm, which is 54.6 µm of product for 250 µm of gap. Closing it to 0.750 mm gives 2.9474 mm, another 107.2 µm for 500 µm. The ratio is 0.218 on the opening step and 0.214 on the closing one, both a little above the linearized 0.2104, which describes only small moves, and the arithmetic is on the force trace, because opening the gap 250 µm lightens the pass, the force falls from 12.000 to 10.828 MN, the stretch falls from 2.000 to

Table 8.5 Scenario Exercises and Completion Requirements. Each Scenario Is Named as the Application Shows It, With the Mechanism It Exercises in Italics.

#	Scenario	Required to complete
1	The gap you set is not the gauge you get. *the mill spring, a gap gain of 0.21*	Run in Manual through the gap moves at 6, 16 and 26 s, then move the gap at least 5 µm off 1.000 mm and hold it for 3 s.
2	Measuring the mill modulus. *force against stretch, not gauge against gap*	Run in Manual through the three gap moves and hold the gap at 1.000 mm for 5 s after the strip arrives 30 °C colder at 26 s.
3	What the roughing mill sends you. *entry walk-through at 0.65*	Run 9 s in Manual on the 3 % wave, three of its periods, then 9 s more with the wave cut to 1 %.
4	Feedback from five meters downstream. *what half a second of dead time permits*	Run 30 s in Auto on the X-ray gauge with $K_c \leq 1.3$ and the 3 % wave.
5	Pushing the delayed loop. *the limit cycle at $2\theta_d$, a looper on its stops* (demonstration)	On the X-ray gauge at $K_c \geq 3.9$ let the gauge error pass ±100 µm, lower K_c to 3.0 or less, and hold the error within ±40 µm for 9 s.
6	The gauge you cannot measure, inferred. *an inferred measurement with no delay*	With gaugemeter and monitor, hold the true gauge within 5 µm of 2.950 mm for 2 s after the step at 8 s, and of 3.000 mm for 2 s after the step back at 20 s.
7	Skid marks. *the disturbance the gaugemeter is exact for*	Run 25 s of skid marks on each of gaugemeter with monitor at $K_c \geq 3.9$, no gauge control, and the X-ray gauge at $K_c = 1.2$ and $\tau_I = 1.5$ s.
8	The rolls grow into the strip. *the sign inversion, driven by the loop*	Keep the gaugemeter loop without trim in Auto for 110 s after the crown starts at 10 s.
9	The monitor trim. *splitting one CV by frequency*	Run 100 s past the crown onset with the trim at 0.60 and 2.0 s, then 60 s each with the trim gain at 0 and at 1.5 with $\tau_I = 0.3$ s.
10	A calibration that is a season old. *a bias no instrument at the stand can find*	Run the gaugemeter loop with the +5 % error to at least 18 s, switch the monitor in until the true gauge holds within 5 µm for 3 s, then set the error to −5 % and see the gauge error pass −20 µm.
11	Once per revolution. *amplification at 2.207 Hz*	Run 5 s each with gaugemeter and monitor (estimator off), no gauge control, the gaugemeter alone at $K_c \geq 3.9$, and the X-ray gauge at $K_c = 1.2$ and $\tau_I = 1.5$ s.
12	Telling the roll from the strip. *synchronous demodulation against a shaft angle*	Run to 20 s with the estimator in at $K_{ecc} \geq 2.9/\text{s}$ from 12 s, then 3 s with $K_{ecc} = 0$ and $\tau_F \geq 0.15$ s, then fire the 3 % entry wave and run 6 s more.
13	The looper, the tension and why the gauge cares. *the looper, mass flow and why the gauge cares*	Run 10 s past the F6 gap opening at 42 s, then Reset, open the gauge loop before 8 s and run 5 s past the 16 MPa tension step.
14	Slow down, and watch the dead time grow. *a plant that changes under a fixed tuning*	Run 15 s at 6 m/s on the X-ray gauge with $K_c \leq 1.3$, the same on gaugemeter and monitor, and, from the reloaded case, on the X-ray gauge at $K_c = 4.0$ and $\tau_I = 0.25$ s.
15	Challenge: One coil, everything at once. *five disturbances and a speed change at once* (challenge)	Hold the true delivered gauge within ±25 µm of 3.000 mm over the last third of the 90 s coil, with the setpoint, model parameters and disturbances as given.[a]

[a] With the shipped configuration (X-ray gauge, $K_c = 0.8$, $\tau_I = 2.5$ s) the worst error over the last third is 61.2 µm, 36.2 µm outside the band, at 65.4 % prime yield. With the suggested configuration (gaugemeter and monitor, $K_c = 7.0$ and $\tau_I = 0.15$ s, trim 1.2 and 1.2 s, $K_{ecc} = 2/\text{s}$) it is 20.3 µm, inside with 4.7 µm to spare, at 98.8 % prime yield.

1.805 mm, and 195 µm of the 250 µm given away comes straight back as a housing no longer standing as tall.

> **TRY IT IN THE APP 1 · The Gap You Set Is Not the Gauge You Get**
>
> Load scenario 1 and press Run in Manual. At 6 s the F7 roll gap opens by 250 µm. Watch the delivered gauge: it moves from 3.0000 to 3.0546 mm, only 54.6 µm, because the roll force falls from 12.000 to 10.828 MN and the stand stretches less.

Scenario 2 turns the same experiment into the calibration a mill crew takes on the head end of a coil. The modulus is not the slope of gauge against gap. It is the slope of force against stretch, and stretch is $h - S$, which is readable from the trends. At $+100\,\mu$m the force is 11.530 MN and the stretch 1.9216 mm, so $\Delta F/\Delta\delta = 0.470/0.0784 = 6.00\,\text{MN/mm}$, and at $-200\,\mu$m the same ratio is $0.472/0.0787$. The cold strip that follows moves the gauge 115.5 µm with the gap untouched and the modulus unchanged, because what changed was the load.

Scenario 3 runs a 3 % transfer bar variation, which is 138 µm of amplitude on the 4.600 mm entering F6, against no gauge control at all. F6 passes 57 % of what it is given and F7 passes 65 %, so the amplitude is 78 µm at the F6 exit and 51 µm at the product against a 25 µm tolerance. Read at 60 s the root-mean-square deviation is 36.2 µm and the prime yield is 32.5 %. A hot mill without gauge control does not make slightly worse strip. It makes mostly seconds.

8.7.2 What Five Meters Cost

Scenario 4 closes the loop on the X-ray gauge with the loop detuned to $K_c = 1.2$ and $\tau_I = 1.5\,$s to suit the delay. Read at 60 s on the same disturbance it buys almost nothing, with 34.2 µm of root-mean-square deviation against 36.2 µm open loop and the prime yield moving from 32.5 to 35.1 %. The loop is working, and it is working a wavelength late, and a correction applied late to a periodic disturbance is as likely to add to it as to subtract from it.

Scenario 5 answers the obvious objection by restoring $K_c = 4.00$ and $\tau_I = 0.25\,$s, the tuning the gaugemeter uses successfully in scenario 6, on a disturbance cut to a third. The loop limit-cycles at 1.374 s with a peak deviation of 328 µm and 164 µm of root-mean-square deviation read at 60 s, against 12.0 µm for doing nothing at all. The part that matters on a real mill is below the gauge trace, because the gap swings 2.31 mm, the looper cannot store the mass-flow swing that follows, the arm reaches both its 3° and 45° stops and the tension swings between 0.0 and 34.8 MPa. Loose strip between two stands at ten meters a second is a cobble and a wrecked roll set. Neither 4.75 nor 1.374 s is a tuning constant.

8.7.3 The Same Gain on a Measurement That Is Not Late

Scenario 6 closes the loop on $\hat{h}$ at the identical 4.00 and 0.25 s, with the monitor trim in service. A 50 µm setpoint step is inside 5 µm of the new setpoint 0.54 s after it takes effect, crosses it at 0.65 s and overshoots by 4.5 µm at 1.03 s. With the gaugemeter alone the same step settles inside 5 µm in 0.99 s with no overshoot, and the difference belongs to the trim. The step is no error to its delayed setpoint, but when the thinner strip reaches the X-ray gauge 0.51 s after the step the true gauge is still 11 µm short of the new setpoint, so the trim rises to 17.5 µm at 0.57 s and the gauge loop rolls that much thinner. The bare $S + F/M$ parts from the truth by at most 25.3 µm, 18 ms into the capsule's move, and by more than 5 µm for 60 ms, which is the load cell and its filter fourteen milliseconds behind a moving capsule. Twelve seconds after the step the inferred and true thicknesses agree to 0.20 µm, and that residue is the trim δ, because the uncorrected $S + F/M$ agrees with the truth to 0.002 µm, as Equation (8.17) requires with the mill unchanged.

Table 8.6 puts all four architectures on one disturbance at identical gains, which is the comparison the application exists to make.

Reading Table 8.6 across gives the result the chapter is built on. The delayed loop at the gain the fast loop carries comfortably is seven times worse than doing nothing, at 173.8 µm of root-mean-

Table 8.6 Four Architectures on the Two Strip-Side Disturbances, at Identical Gauge Loop Tuning Except Where Noted. The Entry Wave Block Is Scenario 3 With the Wave Cut to 2 %, Read From the Record Over $15 \le t \le 45$ s, Which Is 300 m of Strip, and the Skid Block Is Scenario 7 Read From the Scorecard at $t = 80$ s, Over 800 m.

Architecture	RMS (µm)	Peak (µm)	Prime (%)	Off gauge (m)
Entry wave 2 %, $K_c = 4.00$ and $\tau_I = 0.25$ s				
Manual, gap held	24.1	34.4	52.3	143
X-ray only	173.8	340.3	6.1	282
Gaugemeter	11.6	16.4	100.0	0
Gaugemeter + monitor	11.5	16.4	100.0	0
Skid marks 35 °C, X-ray detuned to 1.2 and 1.5 s				
Manual, gap held	41.3	134.0	80	159
X-ray only, detuned	32.2	107.8	81	151
Gaugemeter	11.6	32.7	90	82
Gaugemeter + monitor	10.2	32.7	95	42

square deviation against 24.1 µm, while the inferred loop holds 11.6 µm and puts no strip outside the band. Nothing about the two controllers differs, because the plant, the disturbance, the gain and the reset time are identical and the whole difference is the 0.500 s between the roll bite and the gauge. The monitor trim changes the result by less than 0.1 µm here, which is correct, because an entry wave leaves the gaugemeter no bias to remove.

> **DID YOU KNOW? Same Gain, Opposite Outcomes**
>
> From the simulator: on the 2 % entry wave at identical tuning, the loop on the X-ray gauge puts 282 m of 300 m off gauge and the loop on the gaugemeter puts none, while a held gap leaves 143 m off gauge (Table 8.6). Half a second of delay turns the best controller into the worst.

8.7.4 The Disturbance the Gaugemeter Was Invented For

Scenario 7 puts cool bands on the strip. The slab rested on water-cooled skid rails in the reheat furnace, so it carries bands 1.5 m apart that the mill has stretched more than eighty times, arriving 125 m apart and therefore every 12.5 s. Thirty-five degrees is 16 % more flow stress, F6 feels it first and its exit thickens by 99.3 µm, F7 meets that strip and the cold band together 0.66 s later, and 134.0 µm of extra gauge reaches the coiler with nothing controlling it. The lower block of Table 8.6 runs the four architectures on that one disturbance.

The gaugemeter is right for this disturbance and not merely fast. A skid mark is a change in the strip, every effect of a change in the strip appears in the roll force, and the roll force is one of the two numbers the gaugemeter is built from, so Equation (8.17) does not move and the inferred error at the bottom of a 35 °C band is 0.04 µm against a 134.0 µm disturbance. The measurement does not have to be told about temperature or about the furnace.

8.7.5 The Three It Gets Backwards

Scenarios 8 to 12 are the other half of Table 8.3, now with the loops closed.

Scenario 8 is the headline. The gaugemeter loop holds its own measurement dead on setpoint and leaves the delivered gauge 198.6 µm thin, against 41.9 µm with the loop open, so closing the loop made the disturbance 4.7 times worse. At 210 s the capsule is commanded at 0.197 mm instead of 1.000 mm with 16.82 MN on a stand scheduled for 12.000 MN, and the prime yield over the run is 7.4 %. The controller worked perfectly and destroyed the coil.

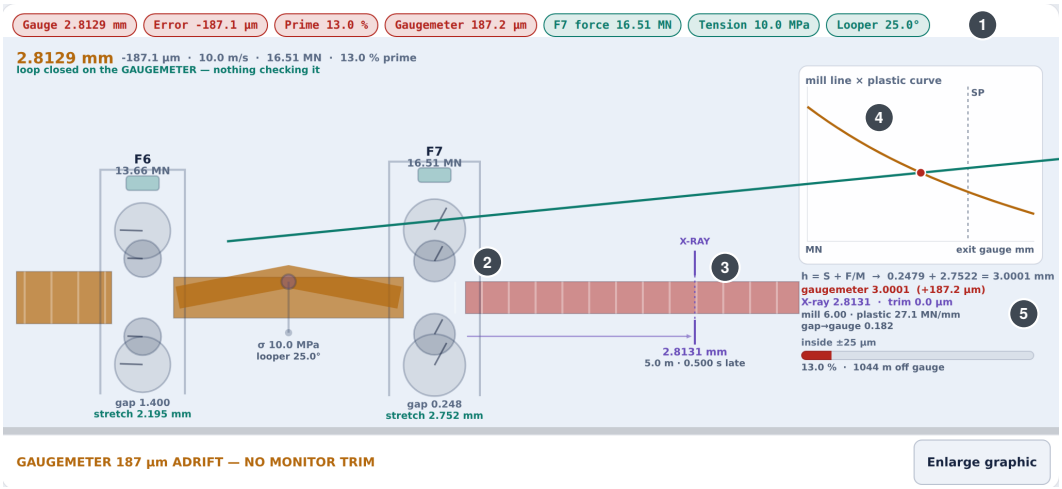


Figure 8.3 Scenario 8 at $t = 120$ s, with 200 µm of thermal crown 110 s into a 40 s time constant and the gauge loop closed on the gaugemeter alone, cropped to the plant view. The badge row reads the gaugemeter 187.2 µm adrift against a true error of −187.1 µm, positive by nearly the whole crown while the truth is negative. The capsule has been driven to 0.2479 mm against a scheduled 1.000 mm and the force to 16.51 MN against 12.000 MN. Callouts: (1) the badge row, (2) stand F7 with its force, gap and stretch, (3) the X-ray gauge and its late reading, (4) the mill line and plastic curve at the present pass, (5) the gaugemeter reading against the X-ray reading.

> **TRY IT IN THE APP** **8 · The Rolls Grow Into the Strip**
>
> Load scenario 8, leave Gauge control at Gaugemeter and run past two minutes. The work rolls start to grow at 10 s. The inferred gauge stays on 3.000 mm while the delivered gauge walks thin. By 210 s the capsule is commanded at 0.197 mm and the strip is 198.6 µm thin.

Scenario 9 adds one slow loop and nothing else. Comparing the X-ray gauge with the setpoint and adding the result as a bias to the gaugemeter's reading, at $K_c = 0.60$ and $\tau_I = 2.0$ s, roughly an eighth the speed of the loop it corrects, holds the peak gauge error to 12.5 µm and the steady error to 0.13 µm against 198.6 µm without it. The trim walks to −0.1985 mm, the capsule settles at 1.198 mm and the force returns to 12.00 MN, which is where the stand should have been all along. A faster trim, at 1.5 and 1.0 s, holds the peak to 3.1 µm, but at 1.5 and 0.3 s, or 3.0 and 1.0 s, the trim has become the delayed loop again and cycles at up to 201 µm.

Scenario 10 is the second job a monitor loop does. A new roll set went in and nobody retook the modulus, so the arithmetic divides by 6.30 MN/mm and the stand has 6.00 MN/mm. Equation (8.17) predicts −95.2 µm at a frozen force and the mill delivers +81.5 µm, because the force falls as the loop opens the gap. Either number is three times the tolerance band and constant for the coil, and no instrument at the stand can find it, because the position is right, the force is right and the arithmetic is right while the constant is wrong. Switching the monitor loop in at 30 s brings the error inside 5 µm 14.9 s later, inside 1 µm after 23.6 s, and parks the trim at +0.0952 mm, which is the calibration error displayed in millimeters. A trim that walks a long way is a maintenance request.

Scenarios 11 and 12 are the disturbance neither loop can reach. 40 µm of runout on a 0.70 m back-up roll at ten meters a second breathes the gap at 2.207 Hz while the position transducer, above the whole roll stack, never moves. The measured gauge ripple, half its peak-to-peak value, is 8.6 µm with no gauge control, 8.0 µm with the X-ray loop at its detuned 1.2 and 1.5 s and 21.0 µm with the gaugemeter, so the fast loop is two and a half times worse than doing nothing. Equation (8.6)

predicts the trends to within 0.5 µm, since $0.2104 \times 40\,\mu\text{m} = 8.4\,\mu\text{m}$ of true ripple runs against $-0.7896 \times 40\,\mu\text{m} = -31.6\,\mu\text{m}$ of reading, measured open loop as 8.6 µm and 31.1 µm. At the shipped 4.00 and 0.25 s the X-ray loop would only limit-cycle, as in scenario 5, which is why scenario 11 runs it detuned. The X-ray loop is untroubled only because 2.207 Hz is fourteen radians a second and the loop stops at one, so being slow is not always a defect. Switching the estimator of (8.18) in at $K_{ecc} = 3/\text{s}$ in scenario 12 takes the ripple from 20.3 µm, with gaugemeter and monitor, to 8.2 µm, just under the no-control value, while losing none of the gaugemeter's speed elsewhere.

8.7.6 The Looper, and a Plant That Changes

Scenario 13 shows what the machinery underneath is for. Tension subtracts from the pressure the rolls apply at 0.024 MN/MPa at frozen thickness, which is 0.83 µm of gauge per megapascal at F7 alone. With the gauge loop open the step from 10 to 16 MPa moves the delivered gauge −10.6 µm, or 1.77 µm per megapascal, because the same tension also pulls the F6 exit 8.6 µm thinner. Closed, the effect appears in the actuator instead of the product, because the gap moves 49 µm through the same step while the delivered gauge strays 1.3 µm and settles within 0.03 µm. The gaugemeter takes a tension change out without being told, because it is one more thing that arrives through the roll force. Opening the F6 gap by 80 µm thickens the F6 exit by only 22 µm and is also a mass-flow change that the looper absorbs and the F6 speed trim walks back, with the delivered gauge moving 9.0 µm at worst.

Scenario 14 changes the plant rather than the disturbance. The line drops from 10 to 6 m/s for the tail out, so the dead time grows from 0.500 to 0.833 s while the gauge has not moved. The slowdown is also a disturbance in its own right, because the strain rate and with it the flow stress fall with the speed, and with the gap held it leaves the strip 56.5 µm thin. The detuned loop removes that over about fifteen seconds with a worst excursion of 67.4 µm, and then survives the new dead time, with the worst excursion going from 34 to 39 µm and the root-mean-square deviation from 22.9 to 22.5 µm, read over the ten seconds before the change and over 25 to 60 s. What has gone is the room for integral action, because the integral gain a delayed loop can carry falls as the reciprocal of the delay, while the proportional limit set by the gap gain barely moves. The scenario 5 tuning, whose short integral time leans on that room, cycles in this case at 2.09 s with 381 µm peaks in the fifteen seconds after the change, against 1.374 s at full speed in scenario 5. What does not change is the disturbance wavelength, because the entry wave is 30 m long and arrives every five seconds instead of every three, slower in time and identical in strip. On the gaugemeter at its own 4.00 and 0.25 s the slowdown costs a 6.4 µm transient and no margin at all, because there is no length in the measurement to divide by anything. Left at the case's detuned 1.2 and 1.5 s, the gaugemeter removes the slowdown itself barely faster than the X-ray loop does, peaking at 50.2 µm.

8.7.7 The Graded Challenge

The challenge is ninety seconds of one coil carrying a 1.5 % transfer bar, 16 °C skid marks, 45 µm of runout, a modulus 4 % adrift, 150 µm of thermal crown arriving over the run and a speed change to 7 m/s, starting on the X-ray gauge with the poor tuning. Six disturbances, and each is defeated by a different line in the scheme. Table 8.7 measures what each line is worth. Figure 8.4 draws its five rows as a chart.

Table 8.7 summarizes the whole application. The two configurations that stay on the delayed measurement differ in every tuning, a gauge loop at 0.8 and 2.5 s with weak inner loops against 1.2 and 1.5 s with sound ones, and by almost nothing in result, at 61.2 against 60.9 µm of peak deviation, so tuning is not the variable. The configuration that runs the fast loop without a slow one behind it is the worst of the five at 65.9 µm and 29.9 % prime yield, because the crown and the stale modulus go straight into the product. Adding the monitor trim takes the peak to 28.6 µm and the prime yield to 89.5 %, and adding the estimator takes them to 20.3 µm and 98.8 % with 10 m off gauge in 767 m. Get the architecture right and almost any sensible tuning passes. Get it wrong and no tuning does.

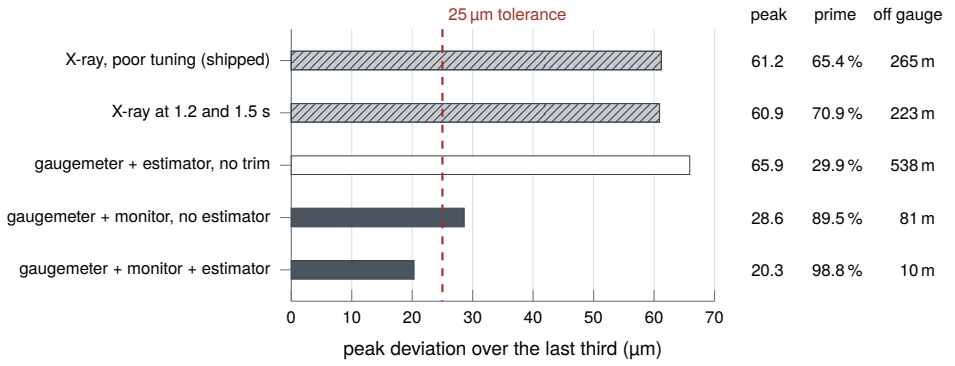


Figure 8.4 The five configurations of Table 8.7 as a chart: the peak deviation over the last third of the coil (bars and first column, in µm), with prime yield and meters off gauge in the other two columns. Hatched bars read the X-ray gauge, the white bar the gaugemeter without a trim, solid bars the gaugemeter with the monitor trim. Notice that the two hatched bars, whose tunings differ in every loop, are almost equal, and only the full scheme falls inside the dashed tolerance line.

Table 8.7 Five Configurations on the Same Ninety-Second Coil, Each After the First Built From the Suggested Settings With One Change. Peak Deviation Is the Application's Score, the Worst Logged Error Over the Last Third ($t \geq 59.4$ s), and RMS, Prime Yield and Meters off Gauge Are Cumulative Over the Whole Run of 767 m.

Configuration	Peak (µm)	RMS (µm)	Prime (%)	Off gauge (m)
Shipped: X-ray, poor tuning	61.2	27.0	65.4	265
X-ray at 1.2 and 1.5 s	60.9	24.0	70.9	223
Gaugemeter + estimator, no monitor trim	65.9	47.0	29.9	538
Gaugemeter + monitor, no estimator	28.6	17.1	89.5	81
Gaugemeter + monitor + estimator (suggested)	20.3	9.9	98.8	10

8.8 Technology in Practice

The scheme of Section 8.5 runs on three instruments, one hydraulic actuator and a layered control system. Figure 8.5 draws the gauge loops with that hardware in place of the schematic's symbols.

What Measures It

TECHNOLOGY **Sensor: The X-Ray Thickness Gauge**

An X-ray thickness gauge sends a beam through the strip from a source on one side to a detector on the other. The transmitted intensity falls exponentially with the thickness crossed, following Beer's law [17]. At low photon energies the attenuation grows with about the fourth to fifth power of the atomic number, so the reading depends on the alloy as well as on the thickness. Gauges of this family have been moving from radioactive sources to radiation generators [17]. Inside a finishing train there is no room for the gauge and the environment is hostile, so it stands only behind the final stand [18]. In this chapter it is GT-701, a 0.030 s averaging lag at the end of 5.0 m of runout table.

The gaugemeter is built from two instruments at the stand. The position transducer reads the capsule, and the load cell under the roll stack reads the roll force. A load cell is a steel element of precisely known modulus with strain gauges bonded to it, and the gauges change resistance as the element compresses or relaxes on a microscopic scale [19]. Stray electric currents through a load cell are read as a change of load, and the currents of arc welding equipment can damage it for good. In

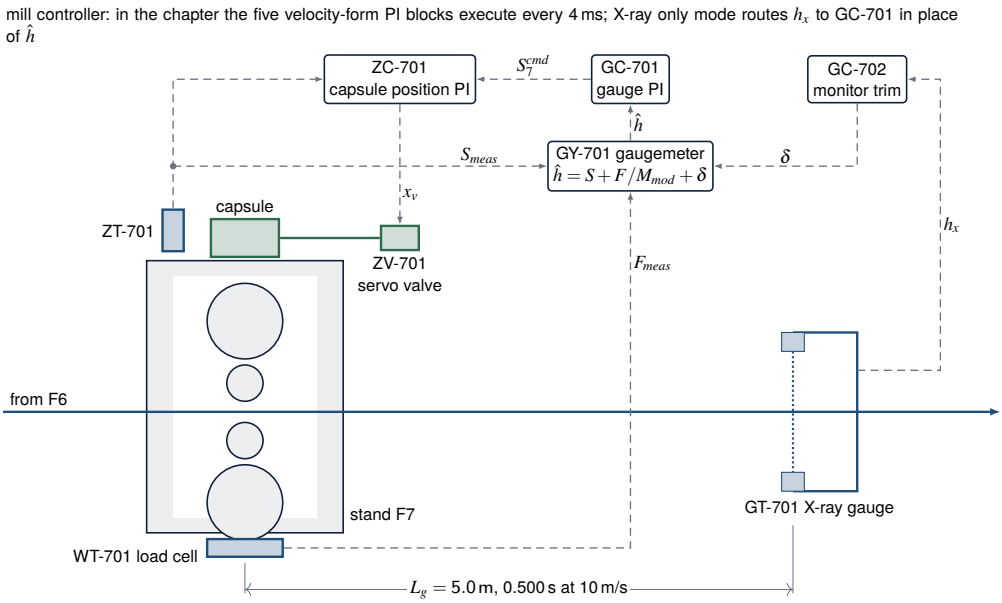


Figure 8.5 The F7 gauge loops drawn with their hardware. The position transducer ZT-701 on the capsule and the load cell WT-701 under the roll stack feed the gaugemeter GY-701 at once, and the gauge controller GC-701 sends the gap command to the capsule position loop ZC-701, which drives the servo valve ZV-701. The X-ray gauge GT-701 reads the strip itself 5.0 m downstream and feeds only the monitor trim GC-702, whose output δ biases the gaugemeter. Notice that no instrument at the stand reads the strip: the fast loop sees the mill, and only the slow loop sees the product.

the chapter's model the load cell WT-701 is a 0.004 s lag and the transducer ZT-701 a 0.002 s lag, which is why the inferred gauge is available at once.

What Moves It

The gap is moved by the capsule under the servo valve, inside a position loop that the chapter models as a proportional loop around the capsule's integrator. That loop, ZC-701 in the figure, is in service in manual and automatic alike, and the gauge loop only writes its setpoint.

> **TECHNOLOGY** **Final Element: The Hydraulic Capsule**
>
> A hydraulic gap control system moves the rolls and creates the rolling force with high-response servo valves, hydraulic cylinders and a precise measurement of the cylinder's movement, and accumulators hold its supply pressures constant [20]. Hydraulic oil hardly compresses, so a cylinder positions firmly and a modest piston generates a very large force. The oil must be purged of gas bubbles, the lines kept free of leaks, and the system stores almost no energy of its own [19]. A model of such a system carries the mill housing stretch and the compression of the oil as parameters [20]. Here the servo valve is a 0.006 s lag and the capsule an integrator of 9.0 mm/s at full valve.

Where the Controller Runs

A published model of a finishing mill stacks the work in layers: a setup calculation that fixes the gaps, speeds and tensions of each coil, a monitor thickness control that corrects from the X-ray gauge, the automatic gauge control, and the hydraulic gap control that moves the rolls [20]. That model's gap control samples every 1 ms. A more recent design names the programmable logic controller as the platform for a multivariable looper and gauge law [21]. Control loops of this kind run with cycle times from milliseconds to minutes [22]. The chapter's five loops all execute every 4 ms. The capsule

loop belongs to the gap control layer, the gauge loop with its architecture switch to the gauge control layer, and the monitor trim to the monitor layer above it.

> **IN PRACTICE** **A Gain Held Below the Limit**
>
> Inside the finishing train the thickness cannot be measured, so the roll force automatic gauge control works from the roll force alone, and its gain is usually set below the value a stable loop would allow. A smaller gain keeps the loop stable through the large variation of the mill modulus and the uncertainty that roll eccentricity adds, at the price of variation left uncorrected. The monitor control trims the roll gaps from an X-ray gauge that stands only behind the final stand. Skid marks from the reheating furnace are the major part of the entry thickness variation [18]. This chapter's shipped gain of 4.00 likewise sits below the deadbeat value of 4.752.

What Is Modernizing It

Three developments touch this chapter's loops. Adaptive feedforward gauge control extracts the long-term variation of the previous stand's exit gauge, mostly skid marks, with an adaptive filter and feeds it to the next stand. It runs on the conventional control system and was shown on real coil data [18]. Machine vision now measures the dimension this chapter leaves out. A shape meter at the finishing exit projects a pattern of light-emitting diode dots onto the strip and reads it with an area camera. Its angle error against a standard target stays within 0.45° at two sigma, it succeeds on more than 98.5 % of each coil, and it feeds flatness control through the work roll bender and leveling [23]. Diagnosis has moved onto the exit gauge itself: comparing the spectra of thickness defects relates each one to a roll, which supports a policy for changing and maintaining rolls, and a least-squares estimate shortens the eccentricity tests [24].

8.9 Limitations

What the application demonstrates soundly is inferential control and the way it fails, and what it is not is a rolling model. One dimension, a linear mill spring, frozen Hitchcock flattening and an absent shape controller between them forbid any statement about the schedule, crown, flatness or chatter of a real mill. The absent second dimension is the larger omission, because work roll bending, shifting and pair-crossing, the roll crown across the barrel and the flatness the coiler shape meter reads all interact with thickness through the roll stack, and on many mills coordinating gauge against shape is a harder problem than either alone [13]. The thermal crown here is a scalar gap offset standing in for a profile.

Three further boundaries matter for anyone reading the numbers as mill data. The strip is isothermal, so the flow stress of (8.5) is an average no instrument on a real mill measures and the skid depth is a modeling input rather than a temperature anybody reads. The interstand span is one lumped elastic state whose stiffness is dominated by roll-bite slip, and slip is not modeled explicitly. There is no threading, no head end and no tail out, so the transient that consumes most of a real gauge control system's effort, the first meters of a coil before the X-ray gauge has anything to read, is outside the model entirely [9]. Five features were kept, the mill spring, the two plant gains, the length-indexed delay, the gap-side offset and the modulus calibration, and each was kept because a gauge controller must be designed against it. Section 8.3.8 states where each simplification fails.

> **OPEN CHALLENGE** **Gauge, Looper and Tension as One Problem**
>
> Gauge, looper angle and strip tension interact, because every gap move made by the gauge control changes the mass flow the looper must absorb. Practice still runs separate proportional-integral loops that do not take these dynamic relations into account [21]. A multivariable gauge and tension control based on optimal regulator theory was proved in simulation and experiments in 1993 [25]. An integrated design of 2019 removed the interaction in simulation and was described as a prospect for use on site [21]. Progress would be an integrated controller proved on a working mill across strip sizes.

8.10 Discussion

The measured result that transfers outside steel is the ordering in Table 8.7. Three of the five configurations differ only in architecture and span 20.3 to 65.9 µm of peak deviation, while the two that differ only in tuning span 60.9 to 61.2 µm. A loop built on an inferred measurement is therefore not usefully described by its tuning until the disturbance classes have been enumerated and each one assigned to a part of the scheme. Equation (8.17) is the tool for that enumeration, and its analogue exists wherever an estimator is used in a loop. A soft sensor built from temperatures and flows is exact for whatever those measurements carry and structurally wrong for whatever changes the relationship between them [2, 3].

The frequency split is the other transferable piece. The crossover between the fast inferred measurement and the slow true one is set by the dead time of the true one and by nothing else, which is why it moves with the line speed here and why a monitor trim tuned at 10 m/s is carrying less margin at six. The same reasoning fixes where an analyzer trim, a laboratory correction or a satellite fix belongs in a cascade, and the reset time of the slow loop is the design variable that places it. Volume 5, Chapter 10 corrects a moisture soft sensor from a delayed laboratory in this way, and Volume 5, Chapter 13 corrects a dead-reckoned position from late acoustic fixes. The application exposes that choice directly, because the monitor trim reset time is a slider and the two failure modes at its ends are both reachable inside a scenario.

Yield is the economic weight behind all of it. A hot strip mill is paid by the meter inside ±25 µm, and the configurations in Table 8.7 put between 10 and 538 m of one 767 m coil outside that band. Strip rolled off gauge is downgraded or rerolled and the energy spent reheating the slab is spent either way, and scenario studies of the European iron and steel industry put process efficiency and yield alongside fuel switching among the levers available to it [26].

Methods other than the five loops fit the same written model. Capsule travel, a stand force rating and two looper stops make the thirteen-state plant a constrained optimization, and the constraint that binds in scenario 8 is the roll force rather than the gauge. An estimator running on these equations over a horizon could tell the gap offset c_7 apart from the modulus error ε_M instead of lumping both into one trim, the distinction a maintenance crew needs and a single bias cannot make.

8.11 Conclusion

The last two stands of a hot strip finishing train make a compact laboratory for inferential control. Thirteen states are integrated across the two stands, the looper and the runout table, with each stand’s exit thickness solved implicitly against its plastic curve at every integration stage, two transport delays indexed on delivered strip length and two slow drifts that make an inferred gauge wrong. At the design pass the stand stretches 2.000 mm under 12.000 MN, so a 1.000 mm gap delivers 3.000 mm of strip and only 0.2104 of a gap move reaches the product. Rearranged, the same relation is a thickness measured at the stand, and its difference from the truth is exactly the negative of the unseen gap offset plus a term in the modulus calibration error, which divides every disturbance in two. On a 2 % entry wave at identical gains the inferred loop holds 11.6 µm of root-mean-square deviation with no strip off gauge while the delayed loop gives 173.8 µm and 282 m off gauge in 300 m, and on 200 µm of thermal crown the same inferred loop leaves the gauge 198.6 µm thin against 41.9 µm with the loop open. A slow X-ray trim at an eighth the speed of the loop it corrects takes that to a 12.5 µm peak and 0.13 µm at 210 s, and a synchronous estimator locked to the back-up roll angle takes the 2.207 Hz ripple from 20.3 to 8.2 µm. A fuller model would add the cross-strip dimension, a head end and a tail out, several eccentric rolls at different frequencies and a mill resonance to stay clear of. Within its limits the simulator shows where an inferred measurement is exact, where it is inverted, and what to put underneath it.

KEY POINTS

- A finishing stand stretches under the force of rolling, so the gap it is set to is not the thickness it delivers.
- The gaugemeter is exact for every disturbance that arrives with the metal and inverted for anything that moves the roll surface without moving the position transducer.
- On a 2 % entry wave at identical gains the inferred loop holds 11.6 µm of root-mean-square deviation while the delayed loop gives 173.8 µm.
- A slow X-ray trim takes the crown error from 198.6 µm to a 12.5 µm peak, and the architecture, not the tuning, decides the challenge coil.

Where else this applies. A fast estimate trimmed by a late true measurement also corrects the moisture soft sensor of Volume 5, Chapter 10 and the dead-reckoned position of Volume 5, Chapter 13, and the roll-angle canceller has a twin in the strand bulging canceller of Chapter 6.

References

[1] B. Joseph and C. B. Brosilow, "Inferential control of processes: Part I. steady state analysis and design," *AIChE Journal*, vol. 24, no. 3, pp. 485–492, 1978.

[2] C. Brosilow and M. Tong, "Inferential control of processes: Part II. the structure and dynamics of inferential control systems," *AIChE Journal*, vol. 24, no. 3, pp. 492–500, 1978.

[3] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

[4] W. C. F. Hessenberg and R. B. Sims, "Principles of continuous gauge control in sheet and strip rolling," *Proceedings of the Institution of Mechanical Engineers*, vol. 166, no. 1, pp. 75–90, 1952.

[5] G. W. Alderton and W. C. F. Hessenberg, "The continuous measurement and control of gauge in strip rolling," *International Materials Reviews*, vol. 1, no. 1, pp. 239–269, 1956.

[6] I. S. Choi, J. A. Rossiter, and P. J. Fleming, "Looper and tension control in hot rolling mills: A survey," *Journal of Process Control*, vol. 17, no. 6, pp. 509–521, 2007.

[7] E. Orowan, "The calculation of roll pressure in hot and cold flat rolling," *Proceedings of the Institution of Mechanical Engineers*, vol. 150, no. 1, pp. 140–167, 1943.

[8] R. B. Sims, "The calculation of roll force and torque in hot rolling mills," *Proceedings of the Institution of Mechanical Engineers*, vol. 168, no. 1, pp. 191–200, 1954.

[9] T. S. Bilkhu, "Hot strip mill gauge control," in *UKACC International Conference on CONTROL 2010*, 2010, pp. 150–155.

[10] M. Ataka, "Rolling technology and theory for the last 100 years: The contribution of theory to innovation in strip rolling technology," *ISIJ International*, vol. 55, no. 1, pp. 89–102, 2015.

[11] A. Kugi, W. Haas, K. Schlacher, K. Aistleitner, H. M. Frank, and G. W. Rigler, "Active compensation of roll eccentricity in rolling mills," *IEEE Transactions on Industry Applications*, vol. 36, no. 2, pp. 625–632, 2000.

[12] W. L. Roberts, *Hot Rolling of Steel*. New York, NY: Marcel Dekker, 1983.

[13] V. B. Ginzburg, *Steel-Rolling Technology: Theory and Practice*. New York, NY: Marcel Dekker, 1989.

[14] K. Wiedemer, "Significance of mill stand stiffness in relation to rolling accuracy," *Metals Technology*, vol. 1, no. 1, pp. 181–185, 1974.

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

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

[17] International Atomic Energy Agency, "Technical data on nucleonic gauges," International Atomic Energy Agency, Vienna, IAEA-TECDOC 1459, 2005, iSBN 92-0-107805-6.

[18] Y. K. Lee, Y. J. Jang, and S. W. Kim, "Adaptive feed-forward automatic gauge control in hot strip finishing mill," *ISIJ International*, vol. 47, no. 10, pp. 1444–1451, 2007.

[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] P. Bhowal and S. K. Mukherjee, "Modeling and simulation of hydraulic gap control system in a hot strip mill," *ISIJ International*, vol. 36, no. 5, pp. 553–562, 1996.

[21] J. Sun, Y. Hu, F. Yin, Y. Hu, W. Peng, and D. Zhang, "Looper-gauge integrated control in hot strip finishing mill using inverse linear quadratic theory," *ISIJ International*, vol. 59, no. 9, pp. 1562–1572, 2019.

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

[23] Y. Isei, T. Kato, M. Osugi, and T. Ohta, "Development of a shape meter employing the LED dot pattern projection method for a hot strip finishing mill," *ISIJ International*, vol. 60, no. 12, pp. 2876–2885, 2020.

[24] D. F. Garcia, J. M. Lopez, F. J. Suarez, J. Garcia, F. Obeso, and J. A. Gonzalez, "A novel real-time fuzzy-based diagnostic system of roll eccentricity influence in finishing hot strip mills," *IEEE Transactions on Industry Applications*, vol. 34, no. 6, pp. 1342–1350, 1998.

[25] K. Kimura, S. Nakagawa, A. Haraguchi, and H. Miura, "Advanced gauge control technology for hot strip mill," *Tetsu-to-Hagané*, vol. 79, no. 3, pp. 352–359, 1993, in Japanese with an English synopsis.

[26] N. Pardo and J. A. Moya, "Prospective scenarios on energy efficiency and CO_2 emissions in the European iron and steel industry," *Energy*, vol. 54, pp. 113–128, 2013.

8.12 Exercises

8.1 Concept check · 2 min

In scenario 8 the gauge loop is closed on the gaugemeter alone, with no monitor trim, at $K_c = 4.00$ and $\tau_I = 0.25\,\text{s}$. The work rolls grow 200 µm into the strip on a 40 s time constant. Long after the crown has settled, where is the delivered gauge?

(a) On the setpoint, because the integral action removes any steady error

(b) About 42 µm thin, the open-loop change, because the loop cannot see the crown

(c) About 200 µm thick, because the gaugemeter reads the higher force as thicker strip

(d) About 200 µm thin, because the loop holds the gaugemeter reading on setpoint

8.2 Concept check · 1 min

Which of these disturbances does the gaugemeter read exactly once the roll force has settled?

(a) A strip that arrives 30 °C colder than the last one

(b) Work rolls that grow 200 µm into the strip as they heat

(c) A back-up roll that is out of round by 40 µm

(d) A mill modulus calibration that is 5 % too high

8.3 Concept check · 2 min

The monitor trim compares the X-ray gauge with the delayed setpoint and adds its output δ to the gaugemeter reading. Why does it run with the action $a = -1$, direct acting in the convention of Section 1.3?

(a) The X-ray gauge sits downstream, and a transport delay reverses the sign of a gain

(b) A larger trim raises the reading the gauge loop holds, so the gap shuts and the strip thins

(c) A trim is a bias on a measurement, and every bias loop in the apps runs with -1

(d) A thermal crown makes the gaugemeter error positive, so the trim has to come out negative

8.4 Theory and calculation · 6 min

The F7 stand has a mill modulus $M_7 = 6.00\,\text{MN/mm}$ and rolls 3.000 mm strip under 12.000 MN, with no gap offset. The plastic modulus of the pass is $Q = 22.513\,\text{MN/mm}$.

(a) Find the mill stretch and the unloaded roll gap. *Checkpoint:* 2.000 mm and 1.000 mm.

(b) Find the gap gain $\partial h/\partial S$ and the deadbeat gain of the gauge loop. *Checkpoint:* 0.2104 and 4.752.

(c) A back-up roll out of round by 40 µm moves the roll surface without moving the position transducer. Find the amplitude of the ripple in the true gauge and in the gaugemeter reading, and say whether the two move together. *Checkpoint:* 8.4 µm and 31.6 µm, in opposite directions.

8.5 Theory and calculation · 5 min

The gaugemeter divides the F7 force by $M_{mod} = 6.30\,\mathrm{MN/mm}$, a calibration 5 % high, while the stand has $M_7 = 6.00\,\mathrm{MN/mm}$. There is no gap offset.

(a) Find the gaugemeter error at the design force of 12.000 MN. *Checkpoint:* −95.2 µm.

(b) The gauge loop on the gaugemeter alone opens the gap until the force has fallen to 10.27 MN. Find the steady error of the delivered gauge. *Checkpoint:* +81.5 µm.

(c) The monitor trim is switched in and returns the force to 12.000 MN. Find where the trim parks, and where it would park for a calibration 5 % low, $M_{mod} = 5.70\,\mathrm{MN/mm}$. *Checkpoint:* +0.0952 mm and −0.1053 mm.

8.6 Theory and calculation · 6 min

The X-ray gauge stands $L_g = 5.0\,\mathrm{m}$ past F7. Model the loop closed on it as the gap gain 0.2104 in series with a pure delay θ_d and a proportional controller.

(a) Find θ_d at 10 and at 6 m/s. *Checkpoint:* 0.500 s and 0.833 s.

(b) Find the ultimate gain and the period of oscillation at the ultimate gain at each speed. *Checkpoint:* 4.75 at both speeds, with periods of 1.0 s and 1.67 s.

(c) The app limit-cycles at $K_c = 4.00$ and $\tau_I = 0.25\,\mathrm{s}$, with periods of 1.374 s at 10 m/s and 2.09 s at 6 m/s. Name what the simple model leaves out, and say whether the simple model predicts a limit cycle at this gain. *Checkpoint:* The capsule, the X-ray averaging filter and the reset action. It does not.

8.7 App challenge · 8 min

Load scenario 1, The gap you set is not the gauge you get, and change no setting. The F7 roll gap command steps from 1.000 to 1.250 mm at 6 s and to 0.750 mm at 16 s. Run at a speed of 1 with the default seed and pause once between 12 and 15.9 s and once between 22 and 25.9 s, where the readings are steady, or read the rows at 15 s and 25 s in the exported CSV file (columns `gauge_mm`, `F7_MN` and `stretch7_mm`).

(a) In the first pause, read the Gauge error, Roll force and Mill stretch tiles, and find the gap gain of the opening step. *Checkpoint:* 54.6 µm, 10.83 MN, 1.805 mm, and a gap gain of 0.218.

(b) In the second pause, read the same tiles and find the gap gain of the closing step from 1.250 to 0.750 mm. *Checkpoint:* −52.6 µm, 13.18 MN, 2.197 mm, and a gap gain of 0.214.

(c) From the readings of (a) and (b), estimate the mill modulus. *Checkpoint:* 6.00 MN/mm, within 0.05.

Accept tile readings within one unit of their last digit.

8.8 App challenge · 9 min

Load scenario 8, The rolls grow into the strip. Its crown event starts at 10 s. Use the default seed and a speed of 2.

(a) Run the scenario as shipped, with Gauge control at Gaugemeter, to 120 s. Read the Gauge error, Gaugemeter error and Roll force tiles. *Checkpoint:* −187.1 µm, 187.2 µm and 16.51 MN.

(b) Reload the scenario, set Gauge control to Manual before Run, and run to 120 s again. Read the same three tiles. *Checkpoint:* −39.5 µm, 187.2 µm and 12.89 MN.

(c) Find the ratio of the two gauge errors, and explain why the gaugemeter error is the same in both runs. *Checkpoint:* 4.7.

Accept each reading within 0.2 µm or 0.02 MN.

8.9 App challenge · 9 min

Load the challenge, One coil, everything at once. Each run covers the whole 90 s coil at a speed of 2 with the default seed. The app pauses by itself at the end of the coil, when the clock reads 90.0 s. For each run, read the challenge result, the worst delivered gauge error over the last third of the coil, and the Prime yield tile at the end.

(a) Run the challenge as shipped, on the X-ray gauge with the poor tuning. *Checkpoint:* 61.2 µm and 65.4 %.

(b) Reload, press Reveal suggested tuning, and run again. *Checkpoint:* 20.3 µm and 98.8 %.

(c) Open the advanced controls, set Eccentricity estimator rate to 0, press Reset and run again. Name the disturbance that now spoils the coil. *Checkpoint:* 28.6 µm and 89.5 %, the back-up roll runout.

Accept each value within 0.2 µm or 0.2 %.

8.10 Application area · 4 min

A mill builder offers a stiffer F7 housing with $M_7 = 12.0\,\mathrm{MN/mm}$. Keep the pass of the chapter: 12.000 MN, $Q = 22.513\,\mathrm{MN/mm}$, and the entry sensitivity $\partial F/\partial h_{in}$ implied by the walk-through factor of 0.6534 on the present stand.

(a) Find the stretch, the gap gain and the deadbeat gain of the stiffer stand. *Checkpoint:* 1.000 mm, 0.348 and 2.88.

(b) Find the fraction of an entry thickness variation that reaches the product. *Checkpoint:* 0.540, against 0.6534.

(c) Find the gaugemeter error for 200 µm of crown, and its ratio to the true change. *Checkpoint:* 200 µm, and 2.88 times the true change.

8.11 Application area · 3 min

The slab rested on water-cooled skids 1.5 m apart in the reheat furnace, and the cool bands arrive 125 m apart on the 3.000 mm strip. Neglect any change in width.

(a) Find the elongation from slab to strip and the slab thickness it implies. *Checkpoint:* 83.3 and 250 mm.

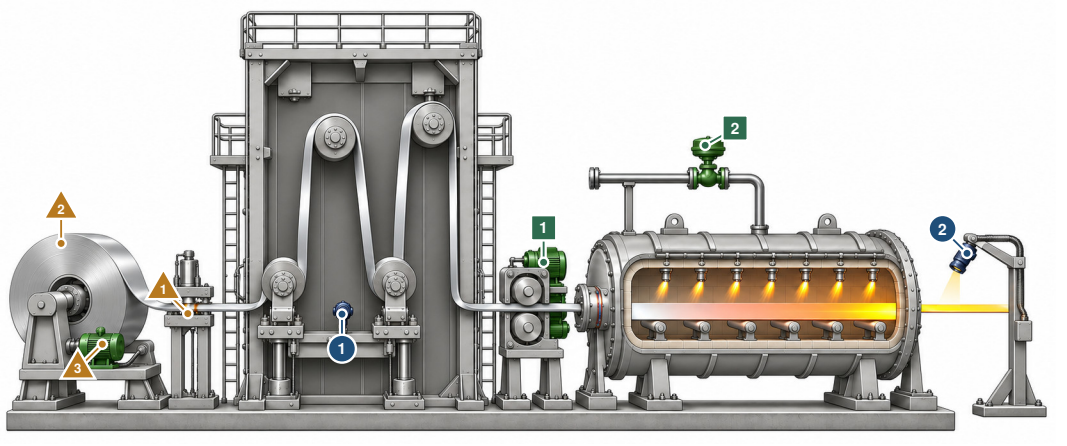
(b) Find the time between bands at 10 and at 7 m/s, the tail-out speed of the challenge. *Checkpoint:* 12.5 s and 17.9 s.

8.12 Review the agent · 5 min

An AI assistant is asked how to deal with the back-up roll runout and replies as follows. “Skip the eccentricity estimator and set the force signal filter to $\tau_F = 0.20\,\mathrm{s}$. Its corner is at 0.80 Hz, so it cuts the 2.207 Hz runout by a factor of 2.9, while the 0.333 Hz entry wave keeps 92 % of its amplitude. The filter needs no shaft encoder and no tuning, so it is the more robust choice and costs nothing at the entry wave.” Find the flaw and say how you would check it.

CHAPTER 9

Annealing Line Entry Looper Simulator for Averaging Level Control and Speed Feedforward



The controller moves 1 speed of the drive rolls that pull strip out of the looper 2 firing of the furnace

The controller reads 1 strip stored in the looper 2 strip temperature at the furnace exit

What disturbs it 1 weld stops, some longer than the looper can carry 2 shorter or thicker coils 3 a slower entry

Notice that the looper stands between the welder and the furnace: the strip it stores lets the furnace run on while the entry stops for each weld. Illustration, not to scale.

IN PLAIN TERMS **Keeping a Furnace Fed While the Entry Stops**

The main idea. An annealing line heats rolled steel strip in a long furnace that must keep running. Every new coil is welded to the last one while the entry stands still, so a tall store of strip in loops feeds the furnace meanwhile.

The challenge. If the speed is pulled down to keep that store full, every weld slows the furnace, and the strip comes out at the wrong temperature. The store can only smooth the supply of strip, and a stop longer than it holds must be seen coming and met by slowing early.

Why it matters. Strip heated to the wrong temperature leaves the line off specification.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	The first page, the history boxes of Sections 9.1 and 9.2, the fact box of Section 9.3, the try-it and fact boxes of Section 9.5, 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 9.5, Section 9.6, Section 9.7 and its open challenge

The control question A looper stores strip between an entry that stops for every weld and a furnace whose strip temperature depends on the line speed. How should the looper loop be tuned so that the welds stay out of the speed, and what must the speed and the firing do when the supply, a stop or the strip thickness changes?

Learning objectives

- Compute the strip one weld draws from the looper and the average speed an interrupted entry can supply.
- Explain why a looper loop resting on its upper clamp leaves the speed untouched when its setpoint lies below the lowest level of a weld.
- Distinguish what a looper absorbs from what only a schedule cap or look-ahead through a stop can fix.
- Scale the line speed with strip thickness from the product $v_p d$ and pair each speed change with a firing feedforward.
- Compare a predictive looper controller with the capped loop over a stated window and name the window that decides the verdict.

Prerequisites Sections 1.4.5 (weld draw and average supply) and 1.4.7 (strip exit temperature on vd), 1.6.2 and 1.6.4 (the PI loops and the scenarios), and 1.6.3 for the looper MPC only. Chapter 6 gives a first integrating level loop.

The app https://apmonitor.com/apps/annealing-line/

See also Volume 6, Chapter 5 for a line whose product answers to its speed and Volume 6, Chapter 7 for level loops inside a recycle. Volume 3, Chapter 6 for a computed feedforward with the feedback left as a slow trim and Volume 3, Chapter 2 for choosing a structure by where the disturbance enters.

9.1 Introduction

An interrupted operation joined to a continuous one needs an inventory between them. Surge drums sit between batch and continuous units, buffer stocks decouple the stations of automatic transfer lines [1], and accumulators store strip and web so that coil changes do not stop a processing line [2, 3]. The level of such an inventory is a means rather than an end. The loop that moves its outflow is judged by how steady it keeps that outflow while the level stays inside its limits, which is the problem posed as averaging level control in the process literature [4, 5].

The misconception this chapter corrects is that every level loop should hold its setpoint. A block diagram of an integrator under PI control rewards tight tuning, because tight tuning gives the smallest level error. The continuous annealing line shows what that error costs downstream. The strip temperature at the furnace exit depends on the line speed, the strip thickness and the furnace temperature [6], so a loop that holds the looper by moving the process speed passes every weld stop into the product. In the simulator the same entry and the same furnace put 71.3 % of the strip outside its temperature window with a tight looper loop and none with an averaging one. The line carries two further lessons that a transfer function does not show. An inventory can smooth a supply but cannot raise its average, and no averaging tuning can carry a stop longer than the inventory holds at full speed.

The line also has two time scales. The strip temperature at the furnace exit answers changes of line speed, strip thickness and furnace temperature with different response times and amplitudes [6]. In the simulator speed acts within one residence time, while the furnace follows its firing over minutes, the two clocks of Section 1.4.7. The freedom the looper gives the process speed is what allows speed to be used as a fast manipulated variable when the strip thickness changes at a weld, and the same speed changes then call for a feedforward to the slow firing. This shape, a fast manipulated variable borrowed from an inventory and a slow one that must be coordinated with it, recurs wherever a buffer feeds a rate-sensitive unit, such as web processing lines [3] like the coating line of Volume 6, Chapter 5, plants whose recycle loops pass through surge tanks [5] like the reactor and column of Volume 6, Chapter 7, galvanizing and coating lines such as Chapter 10, and others.

This chapter compares averaging and tight looper tuning, a schedule cap on speed, look-ahead through entry stops and two feedforwards on the same sequence of coils. Each of these strategies is standard on strip lines. The model covers the entry sequence, the looper, the drive, the strip heating and the furnace, with the controllers written as coded, including a property of the looper loop that follows from its anti-windup and its upper clamp. Twelve scenarios separate what the looper absorbs, what only the schedule can fix, what advance knowledge of a stop can save and what speed can do that the furnace cannot. Section 1.4.5 introduces the weld draw and the average supply on a smaller line, and this chapter measures them on a full one. Section 9.2 reviews the line and its control. Section 9.3 gives the model of the looper, the entry sequence, the drive, the strip heating and the furnace. Section 9.4 describes the looper and strip temperature loops as implemented, their options, the optional looper MPC and manual mode. Section 9.5 reports the measured scenarios, from tight and averaging tuning through short coils, a long weld, firing for the speed and thickness changes to the MPC and the graded challenge. Section 9.6 sets the loops in their real hardware, Section 9.7 states the limitations, Section 9.8 concludes, and Section 9.9 holds the exercises.

> ⧖ HISTORY **From 25 to 600 Meters a Minute (1936 to the 1970s)**
>
> The first vertical continuous annealing line went into operation in 1936. It annealed black plate for tinplate at 25 m/min and heated the strip by electric resistance. Lines spread from the 1940s, and by the first half of the 1950s they ran at 250 to 300 m/min. From the late 1950s through the 1960s the speed jumped to 450 to 600 m/min. Gas jet cooling made a great contribution: it cut the number of strand passes by half and replaced cooling by dipping the strip in water. A review of 1985 judged that the technology had approached a marginal line speed of 600 m/min [7].

9.2 Continuous Annealing and Its Control

Annealing restores the formability of cold-rolled strip, and drawing grades depend on it [8]. Analysis and control of a continuous annealing line were already the subject of a 1964 article in an electrical engineering transactions [9], and a continuous annealing process for drawing-quality strip was reported in the Transactions of the Iron and Steel Institute of Japan in 1977 [8]. Running the strip through the furnace one coil after another is what makes the entry looper necessary. The entry cuts and welds coils at rest, and only the looper allows the furnace section to keep moving while it does.

> ⧖ HISTORY **Sheet for Forming Waits Until 1972**
>
> Continuous annealing of cold-rolled sheet for forming was studied from the 1930s, yet the first commercial vertical line for such sheet started only in October 1972. The obstacle was metallurgical: there was no way to make the carbon dissolved in the steel precipitate in a short time [7]. Earlier work had needed an over-aging treatment of 30 minutes to an hour, too long for a running strip. A process reported in 1977 quenched the strip in water first and so finished the over-aging in less than a minute [8].

Strip temperature control became a named subject as these lines spread. A system for a continuous annealing and processing line was described in 1984 [10], another for a continuous annealing line in 1993 [11], and measurement of the strip temperature with a contact measuring roll was reported for such a line in 1992 [12]. The furnace itself has been modeled in detail. Mathematical models have been published for direct-fired [13] and radiant tube [6] furnaces. The radiant tube model computes the strip temperature distribution from a heat balance and identifies strip thickness, line speed and tube temperature as the dominant influences, with strip width of little effect [6]. The heating of the strip inside an annealing furnace has also been modeled mathematically [14], and low-order identified models now support predictive control of the heating section with an adaptive observer [15]. Energy use is part of the motivation, and furnace models have been applied to improve the energy efficiency of deep-drawing sheet production [16].

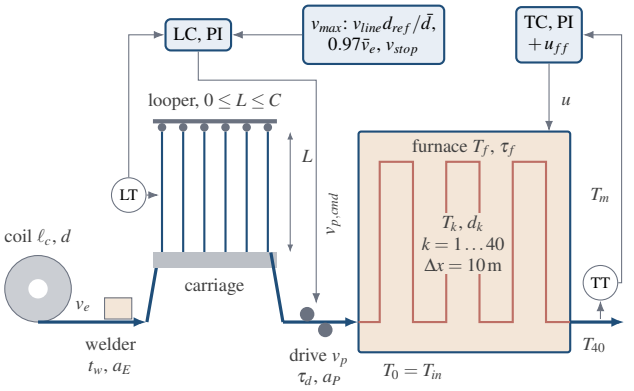


Figure 9.1 Continuous annealing line as modeled. The entry pays out coils of length ℓ_c and thickness d at speed v_e and stops for a weld of duration t_w. The looper stores the strip length L up to its capacity C, the process drive runs the strip at v_p, and the furnace at temperature T_f heats the strip in 40 cells of $\Delta x = 10\,\mathrm{m}$ from T_{in} to the exit temperature T_{40}, measured by the pyrometer TT as T_m. The looper controller LC sets the process speed command within the limit v_{max}, and the temperature controller TC sets the firing u.

Looper control grew up alongside. A strip accumulator is a carriage that moves to lengthen or shorten a set of strands, and its dynamics and control in continuous strip processing lines have been analyzed as a subject of their own [2]. A later model includes the carriage dynamics, the average tension in the accumulator spans and the driven rollers on the process and exit sides, and a controller with a tension observer was compared with the industrial controller of a continuous processing line [3]. The inventory question, how much the stored strip should be allowed to vary and how gently the process speed should respond, is the same one that chemical engineers posed for surge tanks, where averaging level control has been formulated as an optimal control problem for multi-tank systems [4] and systematic PI tuning rules are still being developed for plants with recycle [5].

9.3 Model of the Line

The simulated line pays out coils of 2100 m at up to 420 m/min, stands for 60 s at every weld, stores up to 600 m of strip in the looper and runs the furnace at 300 m/min for 1.0 mm strip, which must leave at 800 °C ±10 °C. Each symbol of the equations that follow is placed on the layout of the line in Figure 9.1.

> **+ DID YOU KNOW? Ten Meters Between Top and Bottom Rolls**
>
> In a vertical annealing furnace the strip runs up and down between rolls more than 10 m apart, and as line speeds rose the span grew to 15 to 20 m. A line started in 1982 was built for more than a million tonnes a year at 450 m/min [7]. This chapter's furnace holds 400 m of strip.

Table 9.1 classifies the variables. The model carries 47 states: the stored strip L, the entry speed v_e, the process speed v_p, the furnace temperature T_f, forty strip cell temperatures T_1 to T_{40}, and three accumulators for strip entered, strip processed and strip off specification. The entry speed is a state only for bookkeeping, because it is set each step by the entry sequence.

9.3.1 Looper and Entry Sequence

The looper is a pure store of strip, Equation (9.1).

$$\frac{\mathrm{d}L}{\mathrm{d}t} = v_e - v_p, \qquad 0 \le L \le C \tag{9.1}$$

Table 9.1 Variables of the Line

Symbol	Type and meaning	Initial	Range
$\ell = 100L/C$	CV, looper level	90 % at start	0–100 %
T_{40}	CV, strip exit temperature	800 °C	SP 720–860
$v_{p,cmd}$	MV, process speed command	300 m/min	60–v_{max}
u	MV, furnace firing	91.5 %	0–100 %
v_e	entry speed (sequence)	300 m/min	0–540
L	state, stored strip	540 m	0–C
v_p	state, process speed	5 m/s	≥ 0
T_f	state, furnace temperature	957.6 °C	—
T_k	states, strip cells 1–40	steady profile	—
ℓ_c	DV, coil length	2100 m	500–4000
t_w	DV, weld stop	60 s	20–240
$v_{e,max}$	DV, entry top speed	420 m/min	180–540
d_{next}	DV, thickness of next coils	1.0 mm	0.5–2.0
σ	DV, pyrometer noise (1σ)	1 °C	0–5

where L is clipped to its bounds after every step. The entry is not a loop but a fixed sequence executed every 0.5 s. While a coil runs, the entry accelerates or decelerates at a_E toward $v_{e,max}$, or toward $\min(v_{e,max}, v_p)$ once $L \geq 0.95C$, and it never exceeds v_p with the looper full. It begins to decelerate when the length R left on the coil satisfies $R \leq v_e^2/(2a_E) + v_e\Delta t$, stops, stands for the weld time t_w, and then starts the next coil. The length of each new coil is drawn as $R = \ell_c(1 + 0.1\xi)$ from a standard normal deviate ξ of the seeded generator, and the first coil starts with $0.55\ell_c$ left.

Two relations follow from the sequence. They are the two halves of Equation (1.8) with this line's symbols. The strip drawn from the looper by one weld, when the entry was running at the process speed before it, is

$$\Delta L_w = v_p \left(t_w + \frac{v_p}{a_E} \right) \tag{9.2}$$

because the entry spends v_p/a_E below the process speed in its deceleration and acceleration together. For the default line this is 350 m, and the measured drop of the level at each weld in case 3 is 347.5 m. A 180 s weld draws 950 m, more than the looper holds. The average rate the entry can supply over a coil cycle is Equation (9.3).

$$\bar{v}_e = \frac{\ell_c}{\ell_c/v_{e,max} + t_w + v_{e,max}/a_E} \tag{9.3}$$

It is 337 m/min for 2100 m coils, 277 m/min for 1000 m coils and 237 m/min for 1000 m coils with the entry derated to 330 m/min. The last two are below the 300 m/min the furnace takes.

9.3.2 Process Drive

The process section follows its speed command through a lag, with a tighter limit on acceleration than on deceleration, Equation (9.4).

$$\frac{\mathrm{d}v_p}{\mathrm{d}t} = \mathrm{clip}\left(\frac{v_{p,cmd} - v_p}{\tau_d},\ -4a_P,\ a_P \right) \tag{9.4}$$

With $a_P = 0.05\,\mathrm{m/s^2}$ the line gains at most 3 m/min per second and loses at most 12 m/min per second. When the looper is empty the process speed is set equal to the entry speed after the step, and the controller limits its command to v_e whenever $L \leq 0.5\,\mathrm{m}$. Each entry into that condition counts as one line stop.

9.3.3 Strip Heating

The 400 m strip path is cut into $N = 40$ cells of length Δx. Each cell is advected by the process speed with first-order upwinding and heated toward the furnace temperature with a time constant proportional to its thickness, Equation (9.5).

$$\frac{\mathrm{d}T_k}{\mathrm{d}t} = -\frac{v_p}{\Delta x}\,(T_k - T_{k-1}) + \frac{T_f - T_k}{\tau_{ref}\, d_k / d_{ref}} \tag{9.5}$$

for $k = 1 \ldots 40$, with $T_0 = T_{in} = 25\,^{\circ}\mathrm{C}$, $\tau_{ref} = 44\,\mathrm{s}$ and $d_{ref} = 1.0\,\mathrm{mm}$. The thickness d_k of each cell is looked up from a list of coil records. When a coil is welded on, its record starts at the strip coordinate $P + L$, where $P = \int v_p\,dt$ is the strip processed, so the new coil reaches the furnace after the strip stored ahead of it. Cell k carries the thickness at coordinate $P - (k - \frac{1}{2})\Delta x$, and $\bar{d}$ is the mean over the forty cells.

At steady state Equation (9.5) gives

$$T_{40} = T_f + (T_{in} - T_f)\,g, \qquad g = \left(\frac{v_p\,\tau_{ref}\,d/d_{ref}}{v_p\,\tau_{ref}\,d/d_{ref} + \Delta x} \right)^N \tag{9.6}$$

so the exit temperature depends on speed and thickness only through their product $v_p d$, as the plug-flow form of Equation (1.12) does. At 300 m/min and 1.0 mm, $g = 0.169$, against $\exp(-80/44) = 0.162$ for exact plug flow with the same 80 s residence time. The furnace temperature that holds 800 °C is 957.6 °C, and one degree of furnace temperature moves the strip by 0.83 °C. With the furnace held, slowing the line by 30 m/min raises the strip to 827.6 °C, and a 1.3 mm strip at unchanged speed leaves at 721.7 °C.

9.3.4 Furnace, Pyrometer and Quality

The furnace is a first-order lag on the firing, loaded by the strip mass rate, Equation (9.7).

$$\tau_f \frac{\mathrm{d}T_f}{\mathrm{d}t} = 500 + 5u - 30\left(v_p \bar{d} - 5\right) - T_f \tag{9.7}$$

with $\tau_f = 240\,\mathrm{s}$, u in percent, v_p in m/s, $\bar{d}$ in mm and temperatures in °C. Full firing at the nominal load gives 1000 °C. A 1.3 mm strip at 300 m/min would need 1062 °C, while full firing at that load reaches only 955 °C. The pyrometer adds Gaussian noise of standard deviation σ to T_{40} and passes it through a first-order filter of 5 s, updated as $T_m \leftarrow T_m + (T_{raw} - T_m)\,\Delta t/(5 + \Delta t)$. Strip leaving the furnace with $|T_{40} - T_{SP}| > 10\,^{\circ}\mathrm{C}$ accumulates at the rate v_p into the off-specification length, which is judged on the true exit temperature rather than the measurement.

The two time scales appear in open-loop step tests in manual mode, drawn in Figure 9.2. Slowing the process from 300 to 270 m/min raises the strip temperature by 40.5 °C, with 63 % of the change after 75 s, because the strip heats longer at once and the furnace, less loaded, then rises by 15 °C. Raising the firing by 5 % raises the strip by 20.8 °C, with 63 % after 270 s. Speed is the faster handle by more than a factor of three.

Table 9.2 lists the parameters. The line starts with $L = 0.9C$, both speeds at 5 m/s, the furnace at the temperature that puts the strip at 800 °C, the strip cells on the steady profile of Equation (9.6) and the firing at 91.5 %, the value that holds that furnace temperature. All states advance together under the fourth-order Runge-Kutta method, stepped every 0.5 s (Section 1.6.1).

9.3.5 Assumptions

The looper is a store of strip without carriage inertia or tension. Real accumulators couple carriage motion to strip tension and to the speeds of the rolls on either side, which limits how fast they fill and empty [3].

The entry sequence is ideal, with one acceleration, a weld time that changes only by event, and coil lengths scattered by 10 %. Coil preparation, threading and variable weld times are absent.

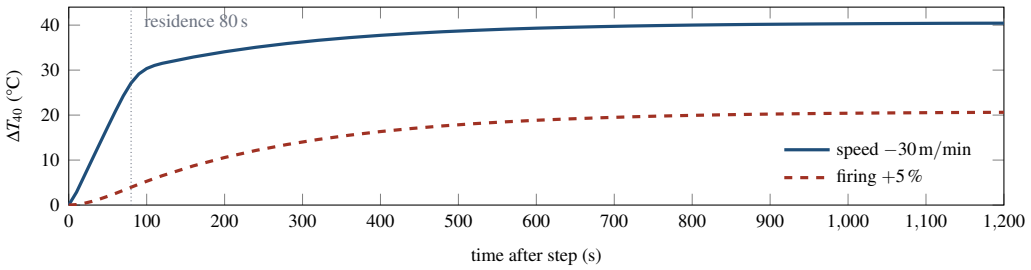


Figure 9.2 Strip exit temperature after a step in process speed and after a step in firing, from 300 m/min, 1.0 mm strip and 800 °C in manual mode, as simulated. The speed response (solid) reaches most of its change within one residence time and then creeps up as the furnace temperature follows the lighter load. The firing response (dashed) follows the 240 s furnace lag.

Table 9.2 Parameters and Their Status in the Interface. Disturbances Are Listed in Table 9.1

Symbol	Meaning	Value	Class
C	looper capacity	600 m	adjustable
a_E	entry acceleration	0.5 m/s^2	adjustable
a_P	process acceleration limit	0.05 m/s^2	adjustable
τ_d	process drive lag	2 s	adjustable
τ_{ref}	strip heating time at 1.0 mm	44 s	adjustable
τ_f	furnace time constant	240 s	adjustable
T_{in}	strip entry temperature	25 °C	hidden
—	strip temperature window	±10 °C	adjustable
N, Δx	cells, cell length	40, 10 m	fixed
d_{ref}	reference thickness	1.0 mm	fixed
—	coefficients of (9.7)	500, 5, 30	fixed
—	pyrometer filter	5 s	fixed
—	coil length scatter (1σ)	10 %	fixed

The strip heats toward one furnace temperature. Real furnaces have heating, soaking and cooling sections with several zones, and radiant heat exchange that is nonlinear in temperature [6, 13].

Strip width is constant, so thickness alone sets the mass per meter, which the radiant tube study supports for the strip temperature itself [6]. The upwind cells add numerical dispersion to what is physically a pure transport, which rounds the front of a thickness change.

Quality is judged by exit temperature alone. Annealing quality also depends on the soaking time and the cooling rate, which are not modeled.

9.4 Controllers as Implemented

Both loops execute every 0.5 s, the same step as the physics, with the velocity-form PID of Section 1.6.2. The derivative time is zero in both loops and no output slew limit is set, so each scan reduces to the PI recursion of Equation (9.8).

$$
\begin{aligned}
e_k &= \alpha\,(r_k - y_k), \qquad P_k = K_c e_k \\
s_k &= s_{k-1} + (P_k - P_{k-1}) \\
c_k &= u_{b,k} + s_k + I_{k-1} \\
u_k &= \mathrm{clip}\,(c_k,\ u_{min},\ u_{max}) \\
I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t}\right]\Delta t
\end{aligned}
\tag{9.8}
$$

where $\alpha = \pm 1$ is the action, $T_t = \tau_I$ and $\Delta t = 0.5\,\text{s}$. These are Equations (1.13) to (1.16) without the derivative, with α for a, s_k for S_k and c_k for v_k, since v is a speed in this chapter. The sum s_k telescopes to P_k and is never clipped, and the clip acts on the emitted output alone. Anti-windup is back-calculation with a tracking time equal to the integral time [17], so at a limit only I is corrected, toward the value that reproduces the applied output.

9.4.1 Looper Loop

The looper loop measures the level ℓ in percent and sets the process speed command. Its action is direct ($\alpha = -1$) in the convention of Section 1.6.2, because a higher level must raise the outflow, so $e = \ell - \ell_{SP}$. The bias u_b is zero, the integral is initialized to the starting speed of 5 m/s, and the output is clipped between 1 m/s and a time-varying upper limit v_{max}, Equation (9.9).

$$v_{max} = \min\left(v_{line}\frac{d_{ref}}{\bar{d}},\ 0.97\,\bar{v}_e,\ v_{stop} \right) \tag{9.9}$$

The line speed v_{line} is 300 m/min. The thickness factor $d_{ref}/\bar{d}$ is applied when the thickness feedforward is on, the schedule cap $0.97\bar{v}_e$ from Equation (9.3) when the schedule option is on, and the look-ahead limit v_{stop} only during a deceleration or a weld with look-ahead on. With the thickness feedforward off the first term is v_{line} itself, and the other two terms drop out when their options are off. The loop can therefore only slow the line below the limit, never speed it up. The default tuning is $K_c = 0.75\,\text{m/min}$ per percent, $\tau_I = 30\,\text{min}$ and $\ell_{SP} = 30\,\%$.

The upper clamp and the back-calculation together give the loop a property that decides its behavior. While the output rests on v_{max}, the last line of Equation (9.8) with $u_b = 0$, $s_k = K_c e_k$ and $T_t = \tau_I$ reduces to

$$I_k - I_{k-1} = \frac{\Delta t}{\tau_I}\left(v_{max} - I_{k-1}\right) \tag{9.10}$$

in which the error has cancelled. The integral relaxes toward the limit, and once it is there the command stays at v_{max} for as long as $\ell > \ell_{SP}$. Below the setpoint the command is close to $v_{max} - K_c(\ell_{SP} - \ell)$, with integral action too slow to matter within one weld. The speed dip at a weld is therefore about K_c times the depth to which the level falls below the setpoint, and a setpoint below the lowest level of a normal weld leaves the process speed untouched whatever the gain.

The look-ahead limit uses the known duration of the entry stop, Equation (9.11).

$$v_{stop} = \max\left(0.5\,\text{m/s},\ \frac{L - 0.05C}{\max(1\,\text{s},\ t_{left})} \right) \tag{9.11}$$

where t_{left} is the weld time remaining plus $v_p/(2a_E)$ during a weld, and the next weld time plus v_e/a_E plus $v_p/(2a_E)$ during the deceleration before it. The term $v_p/(2a_E)$ is the standstill equivalent to the entry's acceleration back to the process speed. The stored strip less a reserve of 5 % of capacity then lasts until the entry has caught up.

9.4.2 Strip Temperature Loop

The temperature loop measures T_m and sets the firing with reverse action ($\alpha = +1$) in the same convention, output limits 0 and 100 %, $K_c = 0.25\,\%$ per °C and $\tau_I = 4\,\text{min}$. Its integral is initialized to the starting firing. With the firing feedforward on, the bias is the steady-state firing for the present speed and mean thickness, measured from the base case, Equation (9.12).

$$\begin{aligned} u_{ff} &= u^*(v_p, \bar{d}) - u^*(5\,\text{m/s}, 1\,\text{mm}) \\ u^*(v,d) &= \frac{T_f^*(v,d) - 500 + 30(vd - 5)}{5} \end{aligned} \tag{9.12}$$

where $T_f^* = (T_{SP} - T_{in}g)/(1-g)$ inverts Equation (9.6) and g is evaluated with the uniform thickness $\bar{d}$. When the option is switched, the integral is moved by the change in u_{ff} so that the firing does not

jump. The feedforward is exact at steady state for uniform strip, because it inverts the same equations the plant uses, and the pyrometer loop is left to trim what it misses, as the settled water measurement trims the computed coagulant dose in Volume 3, Chapter 6. For 1.0 mm strip it asks for 55.4 % firing at 174 m/min and 81.5 % at 268 m/min, against 91.5 % at 300 m/min.

The thickness feedforward on speed is exact in the same sense. Because the exit temperature of Equation (9.6) and the furnace load of Equation (9.7) both depend on $v_p\bar{d}$, scaling the speed by $d_{ref}/\bar{d}$ leaves both unchanged, and 1.3 mm strip at 231 m/min needs the same 957.6 °C furnace as 1.0 mm strip at 300 m/min.

9.4.3 Looper Model Predictive Controller

A model predictive controller (MPC) can replace the looper loop as an option, set up as Section 1.6.3 describes. The looper PI stays the default, in the application and in all scenarios but scenario 11, and it is the fallback. Every $T_s = 4\,\text{s}$, or eight scans, it solves one quadratic program for the process speed command over $N_p = 120$ steps, or 480 s, in 16 move blocks. The first four span one step each and the remaining twelve span 2, 2, 3, 4, 5, 6, 8, 10, 12, 15, 20 and 29 steps. The first move alone is applied and held until the next decision, the receding-horizon form [18, 19].

The controller's model is written from Equations (9.1) and (9.4) rather than identified, because a step test would return the same integrator behind the same lag with fitting error added. It is Equation (9.13).

$$\frac{\mathrm{d}L}{\mathrm{d}t} = v_e - v_h - v_p + q, \quad \frac{\mathrm{d}v_p}{\mathrm{d}t} = \frac{v_{p,cmd} - v_p}{\tau_d}, \quad \frac{\mathrm{d}q}{\mathrm{d}t} = 0 \tag{9.13}$$

Here q is an unexplained net inflow supplied by the estimator, and $v_h \geq 0$ is the rate at which the entry holds strip back while the looper is full, a second decision because the full-looper rule is not linear. The lag takes its declared value $\tau_d = 2\,\text{s}$, copied when the application loads, so an edit of the drive changes the plant and not the model. With a zero-order hold the model becomes Equation (9.14).

$$\begin{aligned} L_{k+1} &= L_k - \tau_d(1-\phi)\,v_{p,k} + T_s\left(v_{e,k} - v_{h,k} + q_k\right) \\ &\quad - \left[T_s - \tau_d(1-\phi)\right] v_{p,cmd,k} \\ v_{p,k+1} &= \phi\, v_{p,k} + (1-\phi)\, v_{p,cmd,k} \end{aligned} \tag{9.14}$$

where $\phi = \mathrm{e}^{-T_s/\tau_d} = 0.135$. The application's sampled matrices equal this closed form to 3×10^{-15}. Run open loop for 120 s after a 30 m/min step of the command in manual mode and fed the measured entry speed, the model stays within 0.111 m/min of the plant's process speed and 0.025 m of its stored strip, a residual caused by the drive's deceleration clip in the first second of the step.

The decisions are the moves of the command, measured from the command applied now, and the held-back rate, both constant within each block. Each plan minimizes Equation (9.15), in which $\hat{L}_k$ and $\hat{v}_{p,k}$ are predictions, j counts the blocks and n_j is the length of block j in steps, with $n_0 = 1$ for the step from the applied command.

$$\begin{aligned} \mathscr{J} &= \sum_{k=1}^{N_p} \left(\frac{\hat{v}_{p,k} - v_{max}}{0.1\,\text{m/s}}\right)^2 + \lambda \sum_{j=1}^{16} \left(\frac{\Delta v_{p,cmd,j}}{0.05\,\text{m/s}}\right)^2 \\ &\quad + w_h \sum_{k=0}^{N_p-1} v_{h,k}^2 + \kappa\,\varepsilon_1^2 + \nu\,\varepsilon_1 + \nu_{SP}\,\varepsilon_2 \\ \text{s.t. } & 0.05C - \varepsilon_1 \leq \hat{L}_k \leq 0.97C + \varepsilon_1, \quad \hat{L}_k \geq \frac{\ell_{SP}C}{100} - \varepsilon_2 \\ & 1\,\text{m/s} \leq v_{p,cmd,j} \leq v_{max}, \quad 0 \leq v_{h,j} \leq v_{e,max}, \quad \varepsilon_o \geq 0 \\ & -4a_P T_s n_{j-1} \leq \Delta v_{p,cmd,j} \leq a_P T_s n_{j-1} \end{aligned} \tag{9.15}$$

The speed term prices the shortfall from the present line-speed limit v_{max}, because the furnace is fired for that speed. With $\lambda = 1$ a move of 3 m/min costs as much as 6 m/min of shortfall for one step, and

$w_h = 0.01$ per $(\text{m/s})^2$. The level limits are soft with one slack per output, so ε_1 is the deepest predicted crossing of the 5 % reserve that look-ahead also keeps or of a top above the entry's 95 % fill rule, with $\kappa = 10^3$ per m^2 and $\nu = 10^3$ per m. The level setpoint is a cheap second floor, $v_{SP} = 10$ per m, kept against what the plan is not told. A limit already crossed is re-entered along a funnel with a 20 s time constant. The command box is the looper PI's own clamp, so the schedule cap and look-ahead bound the MPC when they are on, and the rate rows are the drive's 12 m/min up and 48 m/min down per decision. A dual active-set method [20] solves the program, and any plan the solver does not report optimal, or whose values are not finite or whose command leaves its box, is rejected.

The entry speed over the horizon comes from a preview with three settings. Sequence, the default, runs the entry sequence forward at 0.5 s with its own rules from its present phase and speed, the strip R left on the coil being paid out, known from the entry's coil tracking, and the weld time left or the next weld's time, 60 s until the welder announces a longer one, and then the scheduled coils. It never sees an event before its announcement or a coil before it is drawn, which tests on every case with events confirm. Schedule gives only what the cap and look-ahead know, the rate $\bar{v}_e$ of Equation (9.3) while the entry runs and, once it decelerates for a weld, the time until it runs again, and Off holds the entry speed measured now. Neither the cap nor look-ahead reads what Sequence reads, so Schedule is the like-for-like setting.

A steady-state Kalman filter on (L, v_p, q) runs every 4 s under both looper options. It predicts with the mean command and entry speed of the last period and corrects with the measured stored strip and process speed, with measurement variances of $0.25\ \text{m}^2$ and $2.5 \times 10^{-5}\ \text{m}^2/\text{s}^2$ and process variances per decision of $0.25\ \text{m}^2$, $4 \times 10^{-4}\ \text{m}^2/\text{s}^2$ and $\rho \times 10^{-4}\ \text{m}^2/\text{s}^2$ for q, $\rho = 0.01$. The random-walk inflow makes the plan offset-free [21], although with a model this close to the plant ρ changes nothing measurable.

A first design, $T_s = 1$ s with 180 steps, solved but left 21.20 % of scenario 11 and 13.74 % of the short coils of case 4 with the firing feedforward off specification. A 4 s period at 120 steps, inside the solver's documented limit, is a horizon design, the shortest period that spans a long weld from its announcement through the stop, about 410 s, and two short-coil cycles, 434 s by Equation (9.3). A 2 s period scored better on scenario 11, 17.22 against 17.46 %, and was not chosen. The setpoint floor was added before tuning, after a probe outside the comparison case, and it lowers scenario 11's share from 18.05 to 17.46 %. One pass over nine values of λ and five of ρ on scenario 11 then scored the speed shortfall and the strip temperature IAE relative to the looper PI with the schedule cap and look-ahead. The pass was to admit only settings that move the command no more than that structure, but none did, 339 against 205.5 m/min at best, so the condition was dropped, and the largest λ within 10 % of the best score is $\lambda = 1$.

A rejected plan hands the command to the looper PI, which tracks the applied command every scan, and the next valid plan takes it back. A forced rejection with the pyrometer noise off moves the command by 0.279 m/min against the PI's own move of 0.283 m/min on that scan, and no plan was rejected in any measured run. The entry sequence with its full-looper rule, the empty-looper limit and the temperature loop behave under the MPC as they do under the PI, and a weld extended by 360 s without notice in case 3 empties the looper under both, where the empty-looper rule stops the line. The controller reads the measured stored strip and speeds, the entry sequence's state, the schedule and its declared model, and a test that replays one plant's measurements into a second plant with other drive, strip and furnace parameters finds identical commands and plans over 3000 steps. A plan takes 0.87 to 0.89 ms in Chrome, 6.5 to 6.7 % of a core at 300 times real time, at 59.9 frames per second.

9.4.4 Manual Mode

In manual mode the operator sets the process speed command, clipped to 0 to 420 m/min, and the firing directly. Both PI loops track the applied values, the looper loop its command limited to v_{max}, by setting $I_k = u_k - u_{b,k} - s_k$ in Equation (9.8), so the transfer to automatic is bumpless, as Section 1.6.2 describes. The entry sequence and the empty-looper limit apply in both modes.

9.5 Results

Table 9.3 gives the twelve scenarios in order. All numbers in Table 9.4 and in this section come from runs of the application engine at its default seed unless ten seeds are stated, each over the window given with it, with any setting changed before the run starts, and they agree with the figures quoted in the application's own scenario notes. At the default speed of 60 times real time a one-hour case runs in one minute. Figure 9.3 draws the off-specification shares and line stops of Table 9.4 as a chart.

Table 9.3 Scenarios in Order

#	Scenario	Event and setting
1	Fixed process speed (manual)	1000 m coils at 30 min
2	Holding the looper tightly	ℓ_{SP} 75 %, K_c 2, τ_I 5 min
3	Averaging level control	ℓ_{SP} 30 %, K_c 0.75, τ_I 30 min
4	Short coils	1000 m coils at 10 min
5	Schedule cap	as 4, cap $0.97\bar{v}_e$
6	A weld longer than the looper	180 s weld near 10 min
7	Look-ahead	as 6, v_{stop} on
8	Firing for the speed	as 7, u_{ff} on
9	Thicker strip at a weld	1.3 mm coils at 5 min
10	Speed feedforward	as 9, $d_{ref}/\bar{d}$ on
11	Fill the looper or cap the speed	1000 m coils, MPC, long weld
12	Challenge, two hours	four events, tight tuning

Table 9.4 Measured Scenarios. Off Spec Is the Share of Strip Processed Outside ±10 °C, Speed Is the Time Mean and Standard Deviation

# (min)	Stops	Off spec	Speed (m/min)	Level (%)	T_{40} (°C)
1 (90)	5	18.8 %	287.6 ± 49.9	0–96	800–953
2 (60)	0	71.3 %	287.2 ± 22.4	43–96	782–845
3 (60)	0	0.0 %	300.0 ± 0.0	38–96	800–800
4 (90)	7	53.1 %	278.9 ± 57.4	0–96	753–910
5 (90)	0	6.6 %	271.9 ± 9.9	33–96	800–828
6 (40)	1	30.9 %	285.0 ± 60.0	0–96	736–947
7 (40)	0	31.3 %	288.8 ± 34.3	5–96	749–903
8 (40)	0	18.5 %	288.8 ± 34.3	5–96	732–861
9 (60)	0	77.5 %	300.0 ± 0.0	38–96	718–800
10 (60)	0	2.5 %	245.8 ± 28.2	38–97	778–821
11 (90)	0	7.6 %	274.9 ± 21.5	5–97	744–843
12 (120)	1	76.0 %	251.1 ± 40.8	0–96	716–939

9.5.1 Fixed Speed, Tight and Averaging Tuning

With the process fixed at 300 m/min and 2100 m coils, the looper swings between 38 and 96 % and nothing else moves. After the switch to 1000 m coils the entry supplies 277 m/min on average, the looper first runs empty at 42.5 min, and the line stops five times in 90 min. The tight loop of case 2, set at 75 %, cuts the speed at every weld, to as low as 223 m/min, and the strip runs between 782 and 845 °C with 71.3 % off specification, although the entry could supply the furnace at full speed throughout. The averaging loop of case 3 lets the level fall 347.5 m at each weld without reaching its setpoint, the speed stays at 300 m/min with zero standard deviation, and no strip is off specification.

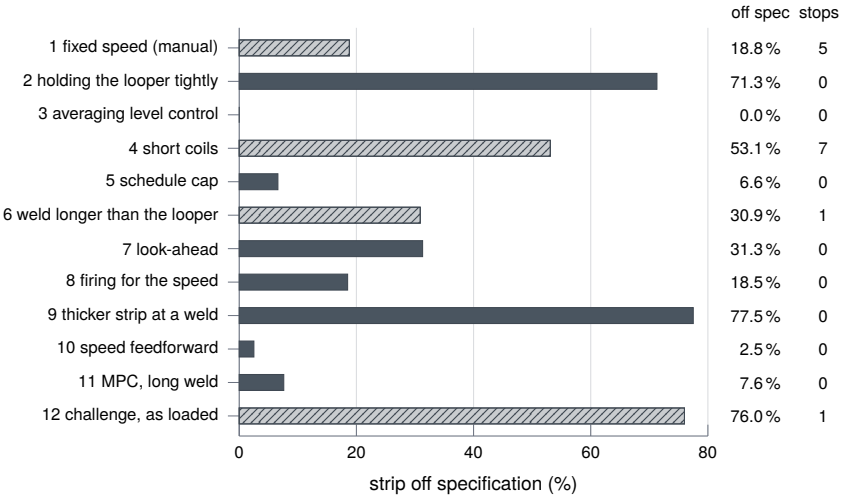


Figure 9.3 The twelve scenarios of Table 9.4 as a chart: the share of strip off specification (bars and first column) and the line stops (second column). Hatched bars mark runs in which the looper ran empty and stopped the line at least once, solid bars runs without a stop. Notice two pairs: tight against averaging tuning (rows 2 and 3), and a speed held against a speed scaled by thickness (rows 9 and 10). Each pair runs on the same entry and the same furnace and goes from one of the longest bars of the chart to one of the shortest.

> **TRY IT IN THE APP** **3 · Averaging Level Control**
>
> Load scenario 3 and press Run. Watch the looper level and the process speed on the trends. At every weld the level falls by about 347.5 m of strip, to no lower than 38 %, while the process speed stays at 300 m/min for the whole hour and no strip leaves off specification.

Table 9.5 separates the setpoint from the gain, which cases 2 and 3 change together. The tight gain and integral time of case 2 with the 30 % setpoint give exactly the result of case 3, because the level never falls below 38 % and Equation (9.10) keeps the command on its limit. With the 75 % setpoint the dip scales with the gain, 31, 77 and 11 m/min for gains of 0.75, 2 and 0.25. These are somewhat above the proportional estimate $K_c(\ell_{SP} - \ell_{min})$ of 27, 65 and 9 m/min, because the integral adds to the cut. A setpoint of 40 %, two points above the lowest level, already gives a dip of 2 m/min. In this loop the setpoint decides whether the welds reach the speed at all, and the gain decides by how much.

Pause. Try Exercise 9.1 before reading on.

Table 9.5 Setpoint Versus Gain in Case 3, One Hour. Gain in m/min per Percent, Dip Is the Lowest Process Speed Below 300 m/min

ℓ_{SP} (%)	K_c	τ_I (min)	Dip (m/min)	Off spec
30	0.75	30	0	0.0 %
30	2	5	0	0.0 %
40	0.75	30	2	0.0 %
50	0.75	30	10	0.0 %
75	0.25	30	11	0.0 %
75	0.75	30	31	17.5 %
75	2	5	77	71.3 %

9.5.2 A Supply That Is Too Small

Short coils defeat the averaging loop. At 30 % the looper holds 180 m, less than one weld draws, so by the time the level is low enough for the loop to act the next weld empties it. Case 4 stops the line seven times, first at 22 min, with 53.1 % of the strip off specification. The schedule cap of case 5 brings the speed to 268 m/min at the schedule change, $0.97 \times 277\,\mathrm{m/min}$, the looper never falls below 33 %, and the line never stops. The single speed change overheats the strip to 828 °C while the furnace follows, leaving 6.6 % off specification, and the firing feedforward reduces that to 2.3 %.

Pause. Try Exercise 9.2 before reading on.

9.5.3 A Stop Longer Than the Looper

A 180 s weld draws about 950 m by Equation (9.2). In case 6 the looper runs empty at 12.4 min, the line stands with the entry for 70 s, and strip that waited in the furnace leaves at up to 947 °C, with 30.9 % off specification over 40 min. Doubling the capacity to 1200 m carries the same weld with no stop and no strip off specification. Look-ahead in case 7 lowers the speed limit as soon as the entry decelerates for the weld and holds it at 174 m/min while the entry stands. The looper bottoms at 5 %, and the line never stops. The furnace is still fired for 300 m/min, so the slow strip peaks at 903 °C and 31.3 % is off specification, about as much as with the stop. After the weld the process accelerates at its limit to 268 m/min and then creeps, reaching 290 m/min about 80 s after the entry restarts, because the level is far below its setpoint and the integral has relaxed toward the look-ahead limit while the entry stood still.

> **DID YOU KNOW? Saving the Stop Did Not Save the Strip**
>
> From the simulator: look-ahead keeps the line running through the 180 s weld that stopped it in case 6, yet the share of strip off specification rises from 30.9 to 31.3 % (Table 9.4). The furnace is still fired for the full speed while the line runs slow.

9.5.4 Firing for the Speed

Adding the firing feedforward in case 8 cuts the firing to a low of 35.1 % during the weld, against 54.4 % with the PI loop alone, and the strip peaks at 861 °C instead of 903 °C. The furnace cannot return as fast as the line, so after the weld the strip falls to 732 °C, below the 749 °C of case 7, while the firing sits at 100 % for 133 s. Over 40 min the off-specification share falls from 31.3 to 18.5 %, with both the hot and the cold strip reduced, from 826 to 616 m and from 2786 to 1521 m. The moment of Figure 9.4 is 30 min into this run, when all of the off-specification strip has already been made. The same feedforward applied to the tight loop of case 2 lowers its off-specification share from 71.3 to 55.1 %, which shows that a feedforward softens speed changes but does not excuse creating them.

9.5.5 Thickness Changes at a Weld

With the speed held, 1.3 mm strip needs a furnace hotter than full firing gives. In case 9 the firing is saturated for 46.3 min, the strip falls to 718 °C and 77.5 % of it is off specification, and adding the firing feedforward changes nothing (77.5 %). Scaling the speed by $d_{ref}/\bar{d}$ in case 10 ramps the process from 300 to 231 m/min as the weld travels through the furnace, the looper absorbs the difference between entry and process, and the strip stays between 778 and 821 °C with 2.5 % off specification. All of that strip leaves in the 100 s while the weld is inside the furnace and $\bar{d}$ lies between 1.16 and 1.30 mm, because the speed set from the mean thickness is exact only when the furnace holds one thickness.

Pause. Try Exercise 9.3 before reading on.



Figure 9.4 Scenario 8, firing for the speed through a 180 s weld with look-ahead and averaging tuning, at $t = 29.83$ min with a 30 min trend window, cropped to the plant view and the trends. The looper is back at 96 % (574 of 600 m), the entry and process run at 300 m/min, the strip leaves at 802 °C, and the furnace is at 960 °C with 92 % firing. The trends show four weld stops of the entry. During the long one, from about 10 to 13.5 min, the dashed line limit steps down and the process speed follows it, the level falls to about 5 %, the firing is cut and then saturates at 100 % as the line speeds up again, and the strip temperature rises above and then falls below its setpoint while the furnace temperature dips. The badge row and the status line show no line stops and 2136 m (25.1 %) of strip off specification. Callouts: (1) the badge row, (2) the looper and its stored strip, (3) the furnace with its firing and the exit temperature, (4) the entry and process speed bars above the status line, (5) the long weld on the trends, where the line limit and the process speed step down.

TRY IT IN THE APP **10 · Speed Feedforward on Thickness**

Load scenario 10 and run past the weld at five minutes, where the coils turn 1.3 mm thick. As the weld travels through the furnace the process speed ramps from 300 to 231 m/min and the looper takes up the difference. The strip stays between 778 and 821 °C, and the 2.5 % off specification all leaves while the weld is inside the furnace.

9.5.6 Predictive Control Through an Announced Long Weld

Scenario 11 runs 1000 m coils from the start with the averaging tuning and the firing feedforward of case 8, and at 680 s the welder announces that the next weld will take 180 s, a stop that the looper cannot carry at 300 m/min. Table 9.6 compares the MPC with the looper PI under the schedule cap and look-ahead, the strongest structure without the MPC, called the capped loop below, at the same seed, noise and duration, with no case retuned and 600 plans in 2400 s, none rejected. Two windows are needed, because the case starts at 300 m/min with the furnace fired for that speed, so a window from $t = 0$ holds a start-up term unrelated to the weld.

Under the MPC the line runs at 288 m/min with the looper at 46.3 % when the weld is announced. The plan then holds the line between 194 and 285 m/min, so the entry, still at 420 m/min, fills the

Table 9.6 Scenario 11 Under Five Looper Structures, Default Seed and Mean of Ten Seeds. Speed SD and Command Travel in m/min Over 0 to 2400 s

	0–2400 s		680–3600 s				
Structure	Off (%)	km	Off (m)	Off (%)	km	SD	Travel
PI, cap, look-ahead	21.60	10.427	1690	13.25	12.752	25.1	205
MPC, Sequence	17.46	10.752	1877	14.43	13.006	30.6	428
MPC, Schedule	40.48	10.767	5291	40.71	12.997	42.6	1071
MPC, Off	87.69	8.448	9598	94.17	10.192	99.6	3840
PI, hindsight[a]	16.17	10.667	1724	13.49	12.781	34.3	472
Mean of ten seeds							
PI, cap, look-ahead	21.85	10.41	1713	13.45	12.74	—	—
MPC, Sequence	19.42	10.77	2268	17.40	13.03	—	—
MPC, Schedule	31.78	10.79	4951	38.05	13.02	—	—
PI, hindsight[a]	17.77	10.53	—	14.84	12.62	—	—

[a] The best of 1080 fixed settings of the looper PI chosen after the run: look-ahead, no cap, $v_{line} = 290$ m/min, $\ell_{SP} = 90$ %, $K_c = 0.25$ m/min per percent, $\tau_I = 30$ min.

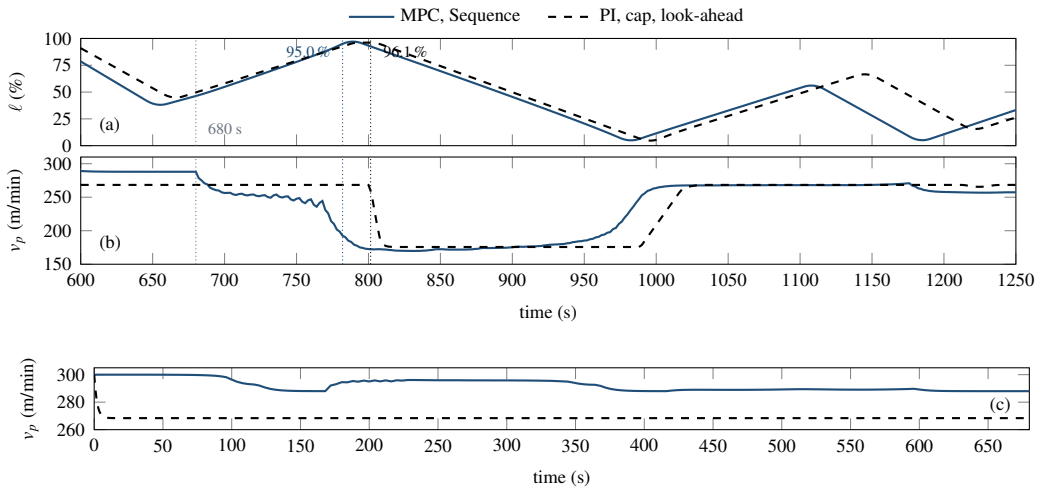


Figure 9.5 Scenario 11 under the MPC with the Sequence preview (blue, solid) and the capped loop (black, dashed). (a) Looper level and (b) process speed from 600 to 1250 s, with dotted marks at the announcement (680 s) and at the entry's deceleration for the long weld under each structure, 782 s at 95.0 % for the MPC and 801.5 s at 96.1 % for the capped loop. (c) Process speed from 0 to 680 s, where the schedule cap steps the line from 300 to 268 m/min at $t = 0$.

looper to 95.0 % by the time it decelerates for the weld at 782 s. The capped loop runs at 268 m/min throughout and arrives as full, 96.1 % at 801.5 s, as shown in Figure 9.5, and neither stops. The fill is itself a slowdown, so the strip leaves its window at 720.5 s under the MPC against 819.5 s, and from 680 to 1200 s the MPC puts 1685 against 1301 m off specification.

Over the first 2400 s the MPC puts 17.46 % of the strip off specification against 21.60 % and processes 10.752 against 10.427 km. That lead comes from the start of the run. The cap steps the speed limit to 268 m/min at $t = 0$, which puts 563 m of strip off specification before the announcement, while the MPC slows gradually to 288 m/min on the strip stored at the start and puts none, as Figure 9.5c shows. From the announcement to one hour, 680 to 3600 s, the MPC puts more strip off specification, 1877 against 1690 m (14.43 against 13.25 %), while carrying 2.0 % more strip. Over ten seeds it puts more off specification in that window at all ten, 17.40 against 13.45 % on average with 2.3 % more strip, and less over the first 2400 s at all ten, 19.42 against 21.85 %.

A fixed setting of the looper PI chosen with hindsight beats the MPC on its own headline. Of 1080 fixed settings of the loop's options and tuning, 587 run without a stop, and the best over the first 2400 s puts 16.17 % off specification with 10.667 km, lower than the MPC at 9 of 10 seeds in both windows. It is a bound rather than a design, chosen after the run was seen.

The fill is the preview's doing. Told only what the cap and look-ahead know, the Schedule preview reaches the weld at 78.6 %, slows to 127 m/min and puts 40.48 % off specification, and told nothing, the Off preview puts 87.69 %. With the level limits and the setpoint floor unpriced, the Sequence plan stops the line three times.

The MPC loses elsewhere too. It moves the speed more, a standard deviation of 30.6 against 25.1 m/min and 428 against 205 m/min of command travel. On steady short coils with the firing feedforward, case 5 over 90 min, the cap's single speed change suits the strip better, 2.29 against 6.15 % off specification, although the MPC carries 2.8 % more strip. A long weld announced 40 s into an ordinary weld stop halts the line under the MPC in 3 of 3 runs and never under the capped loop or look-ahead alone, because Equation (9.11) lets look-ahead slow the line to 30 m/min, below the MPC's 60 m/min floor. On the long coils of cases 6 to 8 the looper is at 95.6 % at the long weld under every structure, and the MPC and look-ahead tie, 30.34 against 31.29 %. Under the MPC the looper of case 4 never empties, 0 stops against 7, so the application's instructions run cases 1 to 10 on the looper PI.

9.5.7 The Graded Challenge

The challenge runs two hours with 1000 m coils from 10 min, 1.2 mm strip from 30 min, a 180 s weld at 60 min and the entry derated to 330 m/min from 80 min, starting from the tight tuning of case 2 with every option off. It is met with no line stops, at most 15 % of the strip off specification and the strip within ± 20 °C over the last 40 min, the last row of Table 9.7, and the outcomes are collected in Table 9.8.

As loaded the line stops once and puts 76.0 % of the strip off specification. The suggested configuration meets all three targets, with no stops, 10.1 % off specification and the strip between 797 and 815 °C over the last 40 min. Removing any one element breaks a target. Without look-ahead the long weld stops the line once, and without either feedforward or the schedule cap the off-specification share rises above 15 %. The schedule cap has a price in throughput, 29.87 against 30.77 km of strip processed without it. Keeping the tight tuning with every option on gives 57.5 %, so the feedforwards cannot compensate for a looper loop that passes the welds into the speed.

With the rules unchanged, the MPC alone still fails the challenge at 43.5 % off specification, because the tight 75 % setpoint becomes its floor and both feedforwards stay off. With the suggested settings and the MPC in place of the looper PI the target is met at 10.2 %, and with the cap and look-ahead also off at 8.0 % with 2.4 % more strip, the last rows of Table 9.8. Under the suggested settings the difficult weld is announced while a weld is already in progress. The event then extends that weld alone by 120 s, and the following weld keeps its scheduled duration. An earlier version of the app also lengthened the following weld, which raised the suggested configuration's off-specification share from 10.1 to 14.2 %.

9.6 Technology in Practice

The two loops of Section 9.4 run on a level measurement, a pyrometer, a speed drive and the firing of the furnace. Figure 9.6 draws them with that hardware in place of the schematic's symbols.

What Measures It

On a real line the strip temperature is measured in more places than the exit. In one furnace of five zones the strip readings, taken together during operation, were judged indispensable for the control of each zone and of the line as a whole [22]. The chapter has one reading at the exit, TT-902. The speed of a roll can be measured by a tachogenerator, a small generator whose voltage is proportional

Table 9.7 Scenarios and Completion Requirements

#	Scenario	Kind	Required to complete
1	Fixed process speed	Demonstration	Run in manual mode with the speed at 300 m/min and the firing at 91.5 % through two welds, then on to 60 min, past the switch to 1000 m coils at 30 min.
2	Holding the looper tightly	Demonstration	With the tight tuning ($\ell_{SP} = 75\,\%$, $K_c = 2$ m/min per %, $\tau_I = 5$ min), run until a weld cuts the process speed below 270 m/min, then 30 min more with the tuning unchanged.
3	Averaging level control	Experiment	With the averaging tuning ($\ell_{SP} = 30\,\%$, $K_c = 0.75$ m/min per %, $\tau_I = 30$ min), hold the process speed at or above 299.4 m/min for 30 min, then run to 60 min with the same tuning.
4	Short coils	Demonstration	With the averaging tuning and no schedule cap, run through three welds of short coils after the switch at 10 min, then on to 60 min.
5	Schedule cap	Experiment	With the schedule cap on, run until the speed limit falls below 276 m/min after the switch at 10 min, then on to 60 min.
6	A weld longer than the looper	Demonstration	With the averaging tuning, run until the 180 s weld empties the looper and stops the line, then 5 min more with the tuning unchanged.
7	Look-ahead	Experiment	With look-ahead on, run until the speed limit falls below 240 m/min at the long weld, then 10 min more with look-ahead on.
8	Firing for the speed	Experiment	With the firing feedforward on, run until the process speed is below 270 m/min with the firing below 70 %, then 10 min more with the feedforward on.
9	Thicker strip at a weld	Demonstration	With the thickness feedforward off, run until the mean thickness in the furnace exceeds 1.29 mm, then 15 min more with it off.
10	Speed feedforward	Experiment	With the thickness feedforward on, run until the process speed falls below 252 m/min, then 15 min more with it on.
11	Fill the looper or cap the speed	Experiment	With the MPC setting the process speed, run in automatic mode to 1100 s, past the weld announced at 680 s and the stop it causes. Restart with the looper PI, the schedule cap and look-ahead, run to 1100 s again, and confirm a comparison of the looper level at the stop, the strip processed and the strip off specification.
12	Challenge, two hours	Challenge	At 2 h, no line stops, at most 15 % of the strip off specification, and $\lvert T_{40} - T_{SP}\rvert \le 20\,^{\circ}\mathrm{C}$ over the last 40 min, with the model parameters and disturbances as given.[a]

[a] With the shipped settings (tight tuning, every option off) the run has one line stop, 76.0 % of the strip off specification and a worst deviation of 46.7 °C over the last 40 min, so all three criteria fail. With the suggested settings (averaging tuning, schedule cap, look-ahead and both feedforwards) it has no stop, 10.1 % off specification and a worst deviation of 15.0 °C, inside the limits by 4.9 percentage points and 5.0 °C.

Table 9.8 Challenge, Two Hours. Suggested Is Averaging Tuning With the Schedule Cap, Look-Ahead and Both Feedforwards. MPC Replaces the Looper PI

Configuration	Stops	Off spec	Last 40 min (°C)
as loaded	1	76.0 %	753–843
suggested	0	10.1 %	797–815
suggested, no schedule cap	0	17.1 %	781–817
suggested, no look-ahead	1	10.2 %	797–815
suggested, no thickness ff	0	16.5 %	794–829
suggested, no firing ff	0	22.1 %	799–820
all options, tight tuning	0	57.5 %	786–824
as loaded, MPC	0	43.5 %	773–822
suggested, MPC	0	10.2 %	797–815
suggested, MPC, no cap or look-ahead	0	8.0 %	792–806

to speed, or by an optical encoder [23]. In the model the level LT-901 is the stored length itself, with no noise and no lag.

line controller: in the chapter both PI blocks execute every 0.5 s, the step of the physics; the optional MPC replaces LC-901 and plans every 4 s

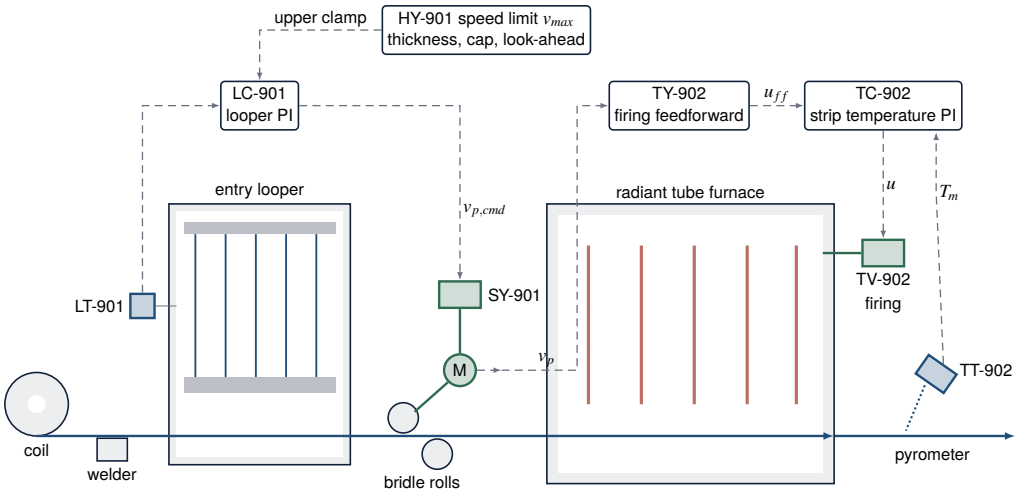


Figure 9.6 The two loops of the line drawn with their hardware. The looper level LT-901 feeds the looper controller LC-901, whose speed command goes to the process drive SY-901, a motor and the bridle rolls that pull the strip out of the looper. The limit HY-901 is the loop's upper clamp, set by the thickness feedforward, the schedule cap and look-ahead. The pyrometer TT-902 reads the strip at the furnace exit and feeds the strip temperature controller TC-902, which sets the firing TV-902 of the radiant tubes, and the feedforward TY-902 adds the firing that the present speed and mean thickness need. Notice that the only path from the looper to the strip temperature runs through the process speed.

TECHNOLOGY **Sensor: A Pyrometer Looking at Bright Strip**

A radiation pyrometer infers a surface temperature from the light the surface emits, so it must know the surface's emittance [23]. Cold-rolled strip in a reducing atmosphere is bright: one study measured an average emissivity of 0.22. The strip is also colder than the furnace walls, so wall radiation reflected off the strip adds to the reading. In one test this back radiation was about 14 °C. A water-cooled shielding flange near the strip, a gas-purged optical filter and a contact thermometer for intermittent correction brought the estimated error below 1 %. A weld has a rougher surface and a higher emissivity, so it shows on the trend as a short pulse of false high temperature [22]. The chapter's pyrometer adds noise and a 5 s filter and has no bias of this kind.

What Moves It

In the heating section of a vertical furnace the strip is heated by radiant tubes arranged in several zones, and the furnace temperature of each zone is controlled independently [7]. Measurements on a working furnace found the line speed the most effective variable in controlling the strip temperature [22], the fast handle of Figure 9.2. The chapter lumps the zones into one furnace temperature with a 240 s lag.

TECHNOLOGY Final Element: The Bridle Roll Drive

The bridle roll drive plays the main role in strip speed control on an annealing line. Its motor, gears and rolls form a torsional spring, and backlash in reducers and couplings lowers the natural frequency of that spring until the speed loop can vibrate. Direct-current motors needed new brushes and commutator parts every two to three years, and for helper rolls inverter-fed alternating-current motors were expected to replace them [7]. Such a variable frequency drive sets the speed by the frequency of the power it feeds the motor [23]. The chapter's drive is a 2 s lag with an acceleration limit of 0.05 m/s^2.

Where the Controller Runs

A line of this kind is run from a layered control system. A supervisor sends set points to distributed controllers, which drive the actuators from sensor feedback, and loops at both levels run with cycle times from milliseconds to minutes [24]. On an annealing line an operation optimization layer sits between the production scheduling layer and the process control layer and sets the process control variables [25]. The chapter's PI blocks execute every 0.5 s, the step of the physics, and the optional MPC plans every 4 s. Both loops, the limit HY-901 and the feedforward TY-902 belong to the process control layer. The schedule cap reads the coil schedule, which belongs to the layer above.

IN PRACTICE Watching the Strip Track

Operators of continuous annealing furnaces know that poor tracking often ends in a broken strip. A temperature difference across the width deforms the strip and changes how it touches the rolls. The rolls are given a crown or a thick middle so that the strip centers itself, but too much crown buckles the center and too strong a centering force wrinkles the strip. The roll profile and the strip tension must therefore be set to suit each operating condition [7]. The model of this chapter has neither tension nor tracking.

What Is Modernizing It

Three developments touch the loops of this chapter. Predictive control of the heating section has moved to linear low-order models identified from the furnace, with parameter estimation and an adaptive observer standing in for quantities that have no sensor. In simulation near actual operating conditions it was compared with a PI controller and with nonlinear predictive control [15]. Data-driven models predict the strip temperature from the operating data of a large furnace: a neural network whose parameters are tuned by an improved swarm optimization was verified on data collected at the production site [26]. Operation optimization now predicts the range of the next disturbances from a time series and adjusts how robust the settings must be, which in tests on actual production data improved strip quality and capacity utilization and reduced the total energy consumption [25].

9.7 Limitations

The looper has no carriage dynamics or tension loops, which in practice limit the rate at which strip can be stored and released [3].

The furnace is one lumped temperature with a linear firing gain. Multi-zone furnaces with radiant exchange respond differently in each zone and nonlinearly in temperature [6, 13], and practical strip temperature control uses their models directly [15].

There is no exit looper and no exit section, so the shears and the coiler at the exit never constrain the process speed.

The feedforwards use the same equations as the plant and so have no model error. The look-ahead knows the weld duration exactly and the schedule cap knows the coil length on average.

Quality is judged on exit temperature alone, and the soaking and cooling that decide the final properties are absent.

The MPC's model omits the strip temperature and the furnace, so its only route to the strip is the price on the speed shortfall, and it also omits the drive's acceleration clip and the scatter of the

coils still to come. Its Sequence preview takes the coil tracking and the announced weld time as exact. A long weld announced while the entry already stands finds the plan unable to slow below 60 m/min where look-ahead may go to 30 m/min. The tuning pass and the headline comparison both use scenario 11, whose start at 300 m/min favors a gradual slowdown in any window that begins at $t = 0$. Manual mode entered during an empty-looper stop starts from the loop's request (277 m/min under the looper PI), not from the limited 0 m/min.

OPEN CHALLENGE A Strip Temperature Model That Holds On Line

The strip temperature in an annealing furnace depends on tube and furnace temperatures, strip width and thickness, line speed and material. The process has a large time delay and a strongly coupled nonlinearity, which limits controllers designed on a mechanism model, and data-driven models trained on production data are now a research focus [26]. Detailed nonlinear models cost much computation in a large furnace, and some quantities have no sensor [15]. Progress would be a model that stays accurate across sizes and speeds and runs inside the controller on a working line.

9.8 Conclusion

The annealing line places a surge inventory between an entry that stops for every weld and a furnace whose product temperature depends on speed. Holding the looper tightly puts 71.3 % of the strip off specification, and letting it swing through its range puts none. Because the loop can only slow the line and its integral relaxes to the limit, the setpoint relative to the lowest level of a weld decides whether any weld reaches the speed. When short coils lower the average supply to 277 m/min the averaging loop lets the looper empty seven times, and a cap at 97 % of the supply removes every stop. A 180 s weld empties the 600 m looper unless look-ahead slows the line to 174 m/min, and firing for the speed then lowers the off-specification share from 31.3 to 18.5 %. A change to 1.3 mm strip saturates the firing and leaves 77.5 % of the strip cold, against 2.5 % with the speed scaled by thickness. A looper MPC fills the looper before an announced long weld, but from the announcement on it leaves 14.4 against 13.3 % of the strip off specification. Carriage and tension dynamics, a multi-zone furnace, an exit section and model error in the feedforwards would make the line more realistic. The line as it stands is a lesson in using an inventory for what it is for.

KEY POINTS

- A looper is a store of strip that lets the furnace keep running while the entry stops for each weld.
- Holding the looper tightly puts 71.3 % of the strip off specification, and letting it swing through its range puts none.
- The setpoint relative to the lowest level of a weld decides whether any weld reaches the speed, because the loop can only slow the line.
- An inventory can smooth a supply but cannot raise its average.
- A change to 1.3 mm strip saturates the firing and leaves 77.5 % of the strip cold, against 2.5 % with the speed scaled by thickness.

Where else this applies. Averaging level control also governs the surge levels inside the recycle of Volume 6, Chapter 7, and a store between an interrupted entry and a rate-sensitive process recurs in the coating line of Chapter 10.

References

[1] J. A. Buzacott, "Automatic Transfer Lines with Buffer Stocks," *International Journal of Production Research*, vol. 5, no. 3, pp. 183–200, 1967.

[2] P. Pagilla, S. Garimella, L. Dreinhoefer, and E. King, "Dynamics and control of accumulators in continuous strip processing lines," *IEEE Transactions on Industry Applications*, vol. 37, no. 3, pp. 934–940, 2001.

[3] P. R. Pagilla, I. Singh, and R. V. Dwivedula, "A Study on Control of Accumulators in Web Processing Lines," *Journal of Dynamic Systems, Measurement, and Control*, vol. 126, no. 3, pp. 453–461, 2004.

[4] N. Ye, T. McAvoy, M. Piovoso, and K. Kosanovich, "Optimal Averaging Level Control for Multi-Tank Systems," *IFAC Proceedings Volumes*, vol. 27, no. 2, pp. 413–418, 1994.

[5] A. Gupta, P. Srivastava, and N. Kaistha, "Systematic Tuning of PI Averaging Level Control for Recycle Systems," *IFAC-PapersOnLine*, vol. 59, no. 6, pp. 421–426, 2025.

[6] F. He, Z.-X. Wang, G. Liu, and X.-L. Wu, "Calculation Model, Influencing Factors, and Dynamic Characteristics of Strip Temperature in a Radiant Tube Furnace during Continuous Annealing Process," *Metals*, vol. 12, no. 8, p. 1256, 2022.

[7] T. Fukushima, "Recent Technological Progress in High Speed Continuous Annealing," *Transactions of the Iron and Steel Institute of Japan*, vol. 25, no. 4, pp. 278–293, 1985.

[8] H. Kubotera, K. Nakaoka, K. Araki, K. Watanabe, A. Nishimoto, and K. Iwase, "Development of a Continuous Annealing Process for the Drawing Quality Steel Strip," *Transactions of the Iron and Steel Institute of Japan*, vol. 17, no. 11, pp. 663–670, 1977.

[9] E. Y. Kung, "Analysis and control of a continuous annealing line," *IEEE Transactions on Communication and Electronics*, vol. 83, no. 70, pp. 121–127, 1964.

[10] K. Yui, T. Sato, H. Ikegami, and K. Suehiro, "Development of Strip Temperature Control for the Continuous Annealing and Processing Line," *IFAC Proceedings Volumes*, vol. 17, no. 2, pp. 1725–1730, 1984.

[11] K. Yahiro, H. Shigemori, K. Hirohata, T. Ooi, M. Haruna, and K. Nakanishi, "Development of strip temperature control system for a continuous annealing line," in *Proceedings of IECON '93 - 19th Annual Conference of IEEE Industrial Electronics*, 1993, pp. 481–486.

[12] H. Seriu, Y. Mashino, Y. Nakajima, and H. Ogawa, "Strip Temperature Measurement by Temperature Measuring Roll for Continuous Annealing Line," *IFAC Proceedings Volumes*, vol. 25, no. 19, pp. 253–256, 1992.

[13] S. Strommer, M. Niederer, A. Steinboeck, and A. Kugi, "A mathematical model of a direct-fired continuous strip annealing furnace," *International Journal of Heat and Mass Transfer*, vol. 69, pp. 375–389, 2014.

[14] S. W. Taylor and S. Wang, "Modelling steel strip heating within an annealing furnace," *Pacific Journal of Mathematics for Industry*, vol. 9, no. 1, p. 5, 2017.

[15] M. Seo, J. Ban, M. Cho, B. Y. Koo, and S. W. Kim, "Low-Order Model Identification and Adaptive Observer-Based Predictive Control for Strip Temperature of Heating Section in Annealing Furnace," *IEEE Access*, vol. 9, pp. 53 720–53 734, 2021.

[16] H. Meunier and M. Cambier, "Use of Furnace Modelling to Improve Energy Efficiency in the Deepdrawing Steel Sheet Industry," in *Energy Efficiency in Process Technology*. Springer Netherlands, 1993, pp. 815–824.

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

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

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

[20] 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.

[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. Iuchi, J. Ohno, and R. Kusaka, "Temperature Measurement System of Steel Strips in a Continuous Annealing Furnace," *Transactions of the Iron and Steel Institute of Japan*, vol. 16, no. 4, pp. 195–203, 1976.

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

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

[25] Y. Wang, X. Wang, Z. Dong, and Z. Wang, "A Data-Driven Multiobjective Dynamic Robust Modeling and Operation Optimization for Continuous Annealing Production Process," *ISIJ International*, vol. 60, no. 6, pp. 1225–1236, 2020.

[26] H. Ding, H. Shen, and Q. Xie, "Predictive Modeling of Strip Temperature in Continuous Annealing Furnace: An Improved Optimization Algorithm," *ISIJ International*, vol. 64, no. 5, pp. 808–816, 2024.

9.9 Exercises

9.1 Concept check · 2 min

In scenario 3 the looper level falls from about 96 to about 38 % at every weld, yet the process speed stays at 300 m/min with zero standard deviation for the whole hour. Which explanation is right?

(a) The gain of 0.75 m/min per percent is so small that each speed dip is below the resolution of the tile

(b) The integral time of 30 min is long against a weld, so the loop cannot move the speed within one

(c) The level never falls below the 30 % setpoint, so the output rests on its upper limit whatever the gain

(d) The schedule cap holds the speed command at 97 % of the average rate the entry can supply

9.2 Concept check · 1 min

After the switch to 1000 m coils in scenario 4 the entry can supply about 277 m/min on average while the furnace section takes 300 m/min. Which statement about the looper is right?

(a) A higher looper gain keeps the process at 300 m/min, because the loop reacts before the level is low

(b) A looper of twice the capacity removes the line stops for good, because it can carry two welds at once

(c) A looper setpoint of 90 % removes the line stops, because more strip is held back before each weld

(d) The mean process speed must fall to the mean supply, because a looper smooths a supply but cannot raise it

9.3 Concept check · 2 min

From the weld after 5 min in scenario 9 the coils are 1.3 mm thick, the line stays at 300 m/min and 77.5 % of the strip leaves off specification, with or without the firing feedforward. Scenario 10 scales the speed by $d_{ref}/\bar{d}$ and leaves 2.5 % off specification. Why does the speed succeed where the firing fails?

(a) The exit temperature and the furnace load depend on $v_p d$ only, and at 300 m/min the strip needs a furnace above full firing

(b) The firing feedforward is switched off whenever the thickness feedforward is off, so scenario 9 runs on the PI loop alone

(c) The pyrometer reads the thicker strip low, so the firing loop acts on a measurement that is wrong by 30 %

(d) The furnace lag of 240 s is too long for the firing to reach a new furnace temperature within the hour

9.4 Theory and calculation · 6 min

The default line runs the process at $v_p = 5\,\mathrm{m/s}$ (300 m/min). The entry accelerates and decelerates at $a_E = 0.5\,\mathrm{m/s^2}$, a normal weld stands $t_w = 60\,\mathrm{s}$ and the looper holds $C = 600\,\mathrm{m}$. In scenario 3 the looper holds 574 m when the entry begins to slow for a weld.

(a) Use Equation (9.2) to find the strip a normal weld draws and the strip a 180 s weld draws. *Checkpoint:* 350 m and 950 m.

(b) Find the lowest stored strip after a normal weld, in meters and in percent of capacity, and say which looper setpoints keep the welds out of the process speed. *Checkpoint:* 224 m, 37.3 %; any setpoint below about 37 %.

(c) Find the longest weld the looper can carry from 574 m at 300 m/min without slowing the line. *Checkpoint:* 104.8 s.

9.5 Theory and calculation · 6 min

The entry runs at up to $v_{e,max} = 420\,\mathrm{m/min}$, a weld stands $t_w = 60\,\mathrm{s}$ and $a_E = 0.5\,\mathrm{m/s^2}$.

(a) Use Equation (9.3) to find the average supply for coils of 2100 m and of 1000 m, in m/min. *Checkpoint:* 336.9 m/min and 276.7 m/min.

(b) The schedule cap is $0.97\,\bar{v}_e$. Find the cap for 1000 m coils, and the coil length below which the entry cannot feed 300 m/min on average. *Checkpoint:* 268.4 m/min; 1295 m.

(c) From 80 min of the challenge the entry is limited to 330 m/min with 1000 m coils. Find the new average supply and the cap. *Checkpoint:* 237.3 m/min and 230.2 m/min.

9.6 Theory and calculation · 7 min

Use Equation (9.6) with $N = 40$, $\Delta x = 10\,\mathrm{m}$, $\tau_{ref} = 44\,\mathrm{s}$, $d_{ref} = 1.0\,\mathrm{mm}$ and $T_{in} = 25\,^\circ\mathrm{C}$, and Equation (9.7) at steady state.

(a) Find g and the furnace temperature T_f^* that gives 800 °C at 300 m/min and 1.0 mm. *Checkpoint:* $g = 0.169$, $T_f^* = 957.6\,^\circ\mathrm{C}$.

(b) With the furnace held there, find the exit temperature of 1.3 mm strip at 300 m/min, and the speed at which 1.3 mm strip leaves at 800 °C. *Checkpoint:* 721.7 °C; 230.8 m/min.

(c) Find the furnace temperature that 1.3 mm strip at 300 m/min needs and the firing that holds it. Compare with the firing range. *Checkpoint:* 1062.4 °C; 121.5 %, above the 100 % limit.

9.7 App challenge · 9 min

Run the app at the default seed and the default speed of 60×, at which one simulated hour takes one minute. Read the scorecard tiles at 60 min.

(a) Load scenario 2, *Holding the looper tightly*, and run 60 min. Read the *Off-spec strip* and *Process speed* tiles, and the lowest process speed on the *Entry and process speeds* trend or as the minimum of `vp_mpm` in the exported CSV, to within 2 m/min. *Checkpoint:* 71.3 %; $287 \pm 22.4\,\mathrm{m/min}$; lowest about 224 m/min.

(b) Load scenario 3, *Averaging level control*, and run 60 min. Read the same two tiles and the *Looper range used* tile. *Checkpoint:* 0.0 %; $300 \pm 0.0\,\mathrm{m/min}$; 38 to 96 %.

(c) Load scenario 3 again, set *Looper level setpoint* to 75 % before Run, keep the gain and integral time as loaded, and run 60 min. Read the *Off-spec strip* tile to within 0.2 % and the lowest process speed to within 2 m/min, and explain the result with Equation (9.10). *Checkpoint:* 17.5 %; lowest about 269 m/min.

9.8 App challenge · 8 min

Run scenarios 4, *Short coils: an inventory cannot fix a supply*, and 5, *Capping the speed at what the entry can supply*, at the default seed for 90 min each, at the default speed of 60× (90 s each).

(a) Run scenario 4 to 90 min and read the *Line stops* and *Off-spec strip* tiles. *Checkpoint:* 7 stops; 53.1 %.

(b) Run scenario 5 to 90 min. After the switch at 10 min, read on the *Line speed* card what the entry can supply and the speed limit. At 90 min read *Line stops*, *Off-spec strip*, *Looper range used* and *Strip temperature range*. Compare the card with your answer to Exercise 9.5. *Checkpoint:* 277 m/min and 268 m/min; 0 stops; 6.6 %; 33 to 96 %; 800 to 828 °C.

(c) Load scenario 5 again, set *Fire for the speed (feedforward)* On before Run, run to 90 min and read *Off-spec strip* to within 0.2 %. Say where the off-specification strip of (b) came from. *Checkpoint:* 2.3 %.

9.9 App challenge · 8 min

Run scenarios 6, *A weld longer than the looper*, 7, *Looking ahead through the stop*, and 8, *Firing for the speed*, at the default seed for 40 min each, at the default speed of 60× (40 s each). Read the scorecard tiles at 40 min.

(a) Scenario 6: read *Line stops*, *Off-spec strip* and *Strip temperature range*. *Checkpoint:* 1; 30.9 %; 736 to 947 °C.

(b) Scenario 7: read the same tiles, *Looper range used*, and the lowest process speed during the long weld on the *Entry and process speeds* trend, to within 2 m/min. *Checkpoint:* 0; 31.3 %; 749 to 903 °C; 5 to 96 %; 174 m/min.

(c) Scenario 8: read *Off-spec strip* and *Strip temperature range*. Explain why scenario 7 leaves as much strip off specification as scenario 6 and what the firing feedforward changes. *Checkpoint:* 18.5 %; 732 to 861 °C.

9.10 Application area · 4 min

(a) Doubling the looper to 1200 m carries the 180 s weld of scenario 6 with no stop and no strip off specification. Explain with Equation (9.2) why, and say whether the larger looper would also remove the line stops of scenario 4. *Checkpoint:* 950 m fits; no, it only postpones them.

(b) The briefing of scenario 1 notes that a line stop can break the strip and cost an hour of rethreading. Give two reasons from the chapter why a line prefers the slowdown of scenario 7 to the stop of scenario 6, although both leave about 31 % of the strip off specification. *Checkpoint:* Strip that stands in the furnace overheats most, and only a slowdown can be paired with a firing feedforward.

9.11 Application area · 3 min

The Limitations section notes that the feedforwards use the plant's own equations and that look-ahead knows the weld duration exactly.

(a) Look-ahead keeps a reserve of 5 % of the 600 m looper. At the look-ahead speed of 174 m/min, how long an overrun of the announced weld does that reserve cover, and what happens after it? *Checkpoint:* About 10 s; the looper empties and the line stops.

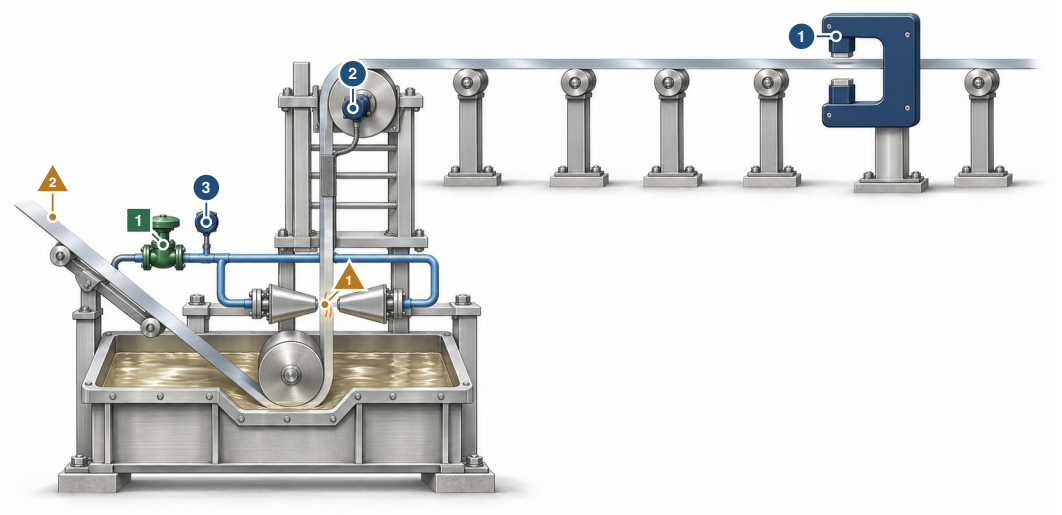
(b) On a real line the thickness and firing feedforwards would carry model error. Which loop removes what they miss, and on what time scale? *Checkpoint:* The strip temperature loop on the firing, over minutes.

9.12 Review the agent · 5 min

An AI assistant, shown Table 9.6, writes: "Over 40 min of scenario 11 the MPC puts 17.46 % of the strip off specification against 21.60 % for the looper PI with the schedule cap and look-ahead, and it processes 3 % more strip. The MPC is the better looper controller, and the line should run every case on it." Find the flaw and say how you would check it with the app.

CHAPTER 10

Galvanizing Line Simulator for Coating Weight Control With a Delay Fixed in Length



The controller moves 1 air pressure to the knives

The controller reads 1 coating gauge, far down the strip 2 line speed 3 air pressure at the knives

What disturbs it 1 coil shape and flutter at the knives 2 a change of line speed

Notice that the knives act just above the zinc, while the gauge sees the coating only after a long run of strip, so every correction is late. Illustration, not to scale.

IN PLAIN TERMS **Holding a Zinc Coating With a Late Gauge**

The main idea. A galvanizing line pulls steel strip out of melted zinc, and two flat jets of air blow the extra zinc back. Harder air leaves a thinner coat, so the controller sets the air pressure to hold the coat on target.

The challenge. The coat can be measured only after it cools, far down the line, so the gauge always reports old news. When the line slows, that wait grows, and a controller that counts it in seconds instead of in length of strip starts to swing.

Why it matters. Every gram of zinc kept as a safety margin costs money, so a steadier coat saves zinc.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history and fact boxes of the Introduction, the history box of Section 10.2, the try-it and fact boxes of Section 10.6, and the key points
Study	35 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 10.6, Section 10.7, Section 10.8 and its open challenges

The control question The coating gauge of a galvanizing line sits a fixed length of strip after the air knives, so its delay changes with every line speed, and the coating's sensitivity to knife pressure changes with every grade. Which loop holds the coating through both, and what is each gram of spread worth in zinc?

Learning objectives

- Compute the knife pressure, process gain and gauge delay of the wiping line at a given speed and coating.
- Explain why more gain on a loop with a long delay widens the coating spread.
- Show why a Smith predictor must keep a transport delay in meters, and find the speeds at which one kept in seconds fails.
- Use speed feedforward to keep the loop gain constant through a slowdown.
- Price coating variance as target margin and zinc, and say why a heavy lower tail needs more margin than the normal formula.

Prerequisites Sections 1.4.5 and 1.4.6 (delay as a length, wiping law and coating gain), 1.6.2 and 1.6.4 (the PI algorithm, scenarios and scorecard).

The app https://apmonitor.com/apps/galvanizing/

See also Volume 5, Chapter 5 for a Smith predictor on a moving ribbon, Volume 5, Chapter 4 for acting ahead of a late analyzer and Volume 5, Chapter 3 for tuning through a slow measurement. Volume 3, Chapter 6 for feedforward from a measured disturbance with a late trim, and Volume 3, Chapter 7 for a ratio paced by throughput.

10.1 Introduction

Many measurements in the process industries are taken downstream of the point where the product is made. Basis weight on a paper machine, film thickness after a die, the composition of a blend at the end of a pipeline and the gauge of rolled strip are all read after the material has traveled some distance, and in each case the dead time the controller faces is that distance divided by a throughput, as Section 1.4.5 sets out. Dead-time compensation began with the Smith predictor and has grown into a family of methods, and the stability of these compensators when their model of the delay is wrong has been analyzed since 1980 [1, 2]. A dead time also sets a floor under the variance that any feedback controller can reach, because nothing the controller does reaches the measurement before it, which is the basis of minimum-variance performance assessment [3, 4].

> ⧖ HISTORY **A Predictor for Dead Time (1957)**
>
> The Smith predictor was proposed in 1957. It splits the model of a process into a fast part without delay and a pure delay, and from them it builds a signal that anticipates the process output. The controller acts on that signal, and the delay drops out of the characteristic equation of the closed loop, but only while the model matches the plant. The original design was meant for plants that are stable in open loop, and later work extended it to integrating and unstable processes. A review of 2022 calls it the best known and most widely used method of dead-time compensation for a single loop, in teaching and in industry [1].

Hot-dip galvanizing is a clear example of this class. Zinc coatings protect steel sheet against corrosion, and the metallurgy of zinc-coated steel is reviewed in [5]. On a continuous line the strip leaves a bath of molten zinc vertically, and two gas jets, the air knives, blow the excess back into the bath [6, 7]. The coating is measured only after it has frozen, here 100 m of strip later. At 150 m/min that is 40 s, and at 90 m/min it is 67 s. A transfer function with a fixed dead time hides this, and it also hides that the sensitivity of the coating to knife pressure moves with the operating point. The simulator is built to correct two beliefs. The first is that a dead time is a number of seconds, which leads to a Smith predictor that works at one speed and oscillates at another. The second is that a tighter loop gives a narrower product, which a long delay makes false.

The economic side is equally direct. The coating target must sit above the specification minimum by a margin that covers the spread, so a reduction in variance is worth the zinc saved by moving the target down, and economic assessments of advanced control value variance reduction in this way [8]. On the 1.5 m strip of the simulator one gram per square meter on both sides at 150 m/min is 27 kg of zinc, or about $81, every hour.

> ✦ DID YOU KNOW? **Most Zinc Ends Up on Steel**
>
> World mine production of zinc was about 12 million tonnes in 2024, an estimate. In the United States galvanized steel was the leading use of refined zinc, and aluminum and plastics are the substitutes for galvanized sheet in automobiles [9].

This chapter works through a galvanizing line simulator in which feedback on the late gauge, speed feedforward and a Smith predictor with its delay kept in seconds or in meters are compared on the same coils, speed changes and grade changes, with every result priced in zinc. All of these strategies are established ones. Section 10.2 reviews coating weight control. Section 10.3 gives a model complete enough to rebuild, including the strip-coordinate history that makes the gauge late. Section 10.4 describes the loop and the two forms of the predictor, and Section 10.5 analyzes the linearized loop at each operating point. Section 10.6 reports eleven scenarios that separate what a delay fixed in length does to feedback, to a predictor and to the economics of the coating target. Section 10.7 describes the hardware of real lines, Section 10.8 states the limitations of the model and Section 10.9 sums up.

10.2 Coating Weight Control on Galvanizing Lines

The history of zinc-coated steel, and the metallurgy of the coatings made on continuous strip lines, are reviewed in [5, 10]. On a continuous line the coating is no longer set by the dipping of each piece but by what is removed from a moving strip as it leaves the bath, and gas jet finishing became the means of removal. An analytical description of jet finishing for hot-dip coatings on strip appeared in 1976 [6], and an analysis of jet stripping of liquid coatings followed in 1984 [11]. The strip speed, the nozzle pressure and the standoff distance from the nozzle to the strip are the process parameters of jet wiping [7], and they are the three quantities of the correlation used here.

Closed-loop control of coating weight arrived in the same decade, and its first steps are told in the history box that follows. The coating gauge of such lines is necessarily downstream of the knives, because the coating must freeze before it can be measured, so every one of these loops faces a transport delay.

> ⧖ HISTORY **Coating Weight Under Computer Control (1975 to 1992)**
>
> Computer control of zinc coating weight on a galvanizing line was reported in 1975 [12], and the regulation of coating thickness in continuous galvanizing of sheet was described in 1977 [13]. A review of 2006 dates the idea of coating weight control on continuous lines to 1976. The analytical wiping model of the time was too complicated for an on-line computer, so a simple exponential model in line speed, knife distance and pressure took its place in the controllers. In 1992 an adaptive controller was built on a linearized form of that model [14]. The wiping law of Equation (10.1) has the same exponential form.

Physical modeling of the wiping has continued since. An engineering model of jet wiping predicts the final coating thickness in agreement with experiments, and identifies splashing of the film as the limit on how hard the jets can wipe [7]. A coating weight model for the continuous hot-dip process was published in the same year [15]. Automatic coating weight control systems for continuous galvanizing lines, and coating weight control on an industrial line at the Florange works, were reported in 2008 and 2010 [16, 17].

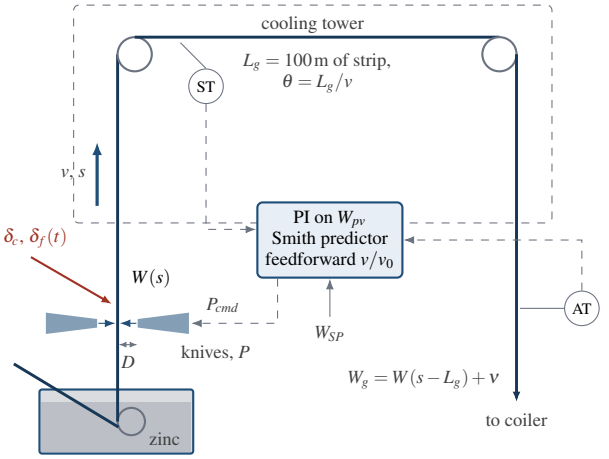


Figure 10.1 Continuous galvanizing line as modeled. Strip enters the zinc pot, turns under the sink roll and rises at speed v between two air knives at pressure P and distance D from the strip, which leave a coating $W(s)$ at strip coordinate s. The coil shape offset δ_c and the flutter δ_f move D unmeasured. The strip travels $L_g = 100$ m over the cooling tower to the gauge AT, which reads the coating made at $s - L_g$ plus noise ν. The controller receives the gauge, the speed from ST and the target W_{SP}, and sends the pressure command P_{cmd} to the knives.

Table 10.1 Variables of the Coating Loop

Symbol	Role	Nominal	Range
W_g	CV, gauge coating per side	76 g/m^2	60–160
W	coating at the knives	76 g/m^2	—
P_{cmd}	MV, knife pressure command	26.2 kPa	2–70
P	state, knife pressure	26.2 kPa	2–70
v	state, line speed	2.5 m/s	1–3.5
s	state, strip coordinate	0	≥ 0
A, A_b	states, area coated, below W_{min}	0	m^2
Z, Z_x	states, zinc used, above W_{min}	0	kg
v_{SP}	DV, line speed setting	150 m/min	60–210
δ_c	DV, coil shape offset (1σ)	4 %	0–8
δ_f	DV, flutter (1σ)	1 %	0–6
L_c	DV, coil length	2000 m	500–5000
ν	DV, gauge noise (1σ)	1 g/m^2	0–5
W_{min}	DV, specification minimum	70 g/m^2	30–150

10.3 Model of the Wiping Line

The line runs a strip 1.5 m wide at a reference speed of 150 m/min. The specification requires at least 70 g/m^2 of zinc on each side, and the default target is 76 g/m^2. The strip, the knives, the cooling tower and the gauge are shown in Figure 10.1.

Table 10.1 classifies the variables. The controlled variable is the gauge reading, not the coating at the knives, which no instrument sees. The model carries seven states, the knife pressure P, the line speed v, the strip coordinate s and four accumulators for the scorecard. The knife distance D and the gauge are algebraic, updated once per controller step.

10.3.1 Wiping, Knives and Line

The coating left on each side follows a power law in speed, distance and pressure, Equation (10.1), which is the wiping law of Equation (1.11) with this line's constants.

$$W = W_0 \left(\frac{v}{v_0}\right)^{0.6} \left(\frac{D}{D_0}\right) \left(\frac{P}{P_0}\right)^{-0.6} \tag{10.1}$$

where $W_0 = 70\,\mathrm{g/m^2}$ is the coating at the reference speed $v_0 = 2.5\,\mathrm{m/s}$, distance $D_0 = 10\,\mathrm{mm}$ and pressure $P_0 = 30\,\mathrm{kPa}$. The code bounds v below at 0.05 m/s and P at 0.5 kPa inside the power law. A faster strip drags out more zinc, knives farther away wipe less, and a higher pressure wipes more, which are the directions of the jet stripping analyses [6, 11]. The pressure that gives a coating W follows by inversion, and it gives $P = 26.16\,\mathrm{kPa}$ for $76\,\mathrm{g/m^2}$ at the reference condition.

The sensitivity of the coating to pressure is the process gain of the loop, Equation (10.2).

$$K_p = \frac{\partial W}{\partial P} = -0.6\,\frac{W}{P} \tag{10.2}$$

It is $-1.74\,\mathrm{g/m^2}$ per kPa at the reference condition. Holding $76\,\mathrm{g/m^2}$ at 90 m/min needs $P = 15.7\,\mathrm{kPa}$ and gives $K_p = -2.91$, and a heavy grade of $145\,\mathrm{g/m^2}$ at 150 m/min needs 8.91 kPa and gives $K_p = -9.76$, which is 5.6 times the nominal value. The gain rises whenever the coating rises or the pressure falls, and both happen on a heavy grade.

The knife pressure follows its command through a first-order lag with the command clipped to the actuator range, and the line speed ramps toward the speed the line sets, Equation (10.3).

$$\begin{aligned} \tau_P \frac{\mathrm{d}P}{\mathrm{d}t} &= \mathrm{clamp}(P_{cmd}, P_{min}, P_{max}) - P \\ \frac{\mathrm{d}v}{\mathrm{d}t} &= \mathrm{clamp}\left(\frac{v_{SP} - v}{2\,\mathrm{s}},\ -a_v,\ a_v\right), \qquad \frac{\mathrm{d}s}{\mathrm{d}t} = v \end{aligned} \tag{10.3}$$

with $\tau_P = 1\,\mathrm{s}$, $P_{min} = 2\,\mathrm{kPa}$, $P_{max} = 70\,\mathrm{kPa}$ and $a_v = 0.02\,\mathrm{m/s^2}$. A slowdown from 150 to 90 m/min therefore takes about 50 s.

10.3.2 Strip Shape

The knife distance is not measured and not manipulated. It carries a step at every coil and a flutter, Equation (10.4).

$$D = D_0\left(1 + \delta_c + \delta_f(t)\right) \tag{10.4}$$

Each time a further coil length L_c of strip has passed the knives, δ_c is redrawn as $\sigma_c \xi$ with ξ standard normal and $\sigma_c = 4\,\%$. The first change comes after $0.6L_c$, 8 min into a run at 150 m/min, and later coils last 13.3 min at that speed. The flutter is a first-order filtered Gaussian sequence with a correlation time of 60 s, scaled to a stationary standard deviation $\sigma_f = 1\,\%$, Equation (10.5).

$$x_k = a x_{k-1} + \sqrt{1 - a^2}\,\xi_k, \quad a = \mathrm{e}^{-\Delta t/60\,\mathrm{s}}, \quad \delta_f = \sigma_f x_k \tag{10.5}$$

started from $x_0 = 0$. Because W is proportional to D, a coil step of one standard deviation moves the coating by $3.0\,\mathrm{g/m^2}$ and the flutter by $0.76\,\mathrm{g/m^2}$ at the target.

10.3.3 The Gauge and the Strip Coordinate

Every controller step records the coating made at the knives against both the time t and the strip coordinate s in a ring buffer of 6000 entries, 20 min at the 0.2 s step. The gauge reads the entry for the strip that is now L_g past the knives, Equation (10.6).

$$W_g(t) = W\left(s(t) - L_g\right) + \nu(t), \qquad \theta(t) = \frac{L_g}{v(t)} \tag{10.6}$$

Table 10.2 Parameters of the Line, Adjustable in the Parameter Panel Unless Marked Fixed

Symbol	Meaning	Value	Class
W_0	reference coating	70 g/m^2	adjustable
v_0	reference speed	150 m/min	adjustable
P_0	reference pressure	30 kPa	adjustable
D_0	reference knife distance	10 mm	adjustable
—	speed, pressure exponents	0.6, −0.6	fixed
τ_P	knife pressure lag	1.0 s	adjustable
P_{min}, P_{max}	pressure limits	2, 70 kPa	adjustable
L_g	knives to gauge	100 m	adjustable
a_v	line speed ramp	0.02 m/s^2	adjustable
—	speed approach time	2 s	fixed
b	strip width	1.5 m	adjustable
p_{Zn}	zinc price	\$3.0/kg	adjustable
W_{min}	minimum per side	70 g/m^2	adjustable
—	flutter, noise correlation	60, 2 s	fixed

The lookup interpolates linearly between the two entries that bracket $s - L_g$. Equation (10.6) is the delay defined by a length of Equation (1.7), and the X-ray gauge of Chapter 8 is keyed to strip length in the same way. The noise ν is a second filtered Gaussian sequence of the form of Equation (10.5) with a 2 s correlation time and a standard deviation of 1 g/m^2. The delay θ is 40 s at 150 m/min, 67 s at 90 m/min and 33 s at 180 m/min. During a speed change the gauge reads strip that was coated at whatever speed it had when it passed the knives, which Equation (10.6) handles without a separate delay calculation.

10.3.4 Zinc and the Scorecard

The accumulators integrate the strip area, the zinc used and the zinc above the minimum on both sides, Equation (10.7).

$$\begin{aligned} \frac{\mathrm{d}A}{\mathrm{d}t} &= vb, \qquad \frac{\mathrm{d}Z}{\mathrm{d}t} = \frac{2Wvb}{1000} \\ \frac{\mathrm{d}Z_x}{\mathrm{d}t} &= \frac{2\max(0, W - W_{min})\,vb}{1000} \end{aligned} \tag{10.7}$$

in m^2/s and kg/s, with $b = 1.5$ m, and A_b grows at vb while $W < W_{min}$. The scorecard reports the mean $\bar{W}$ and standard deviation σ_W of the coating at the knives weighted by strip length since the start of the run, the strip below the minimum $100A_b/A$, the zinc cost $p_{Zn}Z_x/t$ at $p_{Zn} = \$3$ per kilogram averaged over the run, and a target this spread allows, $W_{min} + 2.33\sigma_W$, which would put 1 % of the strip below the minimum if the coating were normally distributed. One gram per square meter on both sides costs $2vbp_{Zn}$, which is \$81 an hour at 150 m/min and \$49 at 90 m/min.

The parameters are collected in Table 10.2, and the disturbance sliders and their ranges in Table 10.1, where the gauge noise is shown only in the advanced view. The line starts at $v = v_{SP}$ with the pressure that puts the coating on target in automatic, or at the manual pressure, and $s = 0$. The history is filled for 120 s behind the knives with that initial coating, so the gauge agrees with the knives at the start. The states advance in 0.2 s steps of the fourth-order Runge-Kutta engine of Section 1.6.1, with D and P_{cmd} held over each step, and a scenario replays exactly at its seed.

10.3.5 Assumptions

The wiping correlation has fixed exponents and a coating proportional to the knife distance. Jet wiping models also depend on the nozzle geometry, the tilt of the jet and the film properties, and the jets cannot wipe beyond the onset of splashing [7].

Both sides and the whole width are one coating. Real lines coat the two sides with separate knives and see crossbow and edge effects across the width, so the distance and the coating vary across the strip.

The knife distance is disturbed but never moved. On real lines the knife position is a second actuator used for large changes, and electromagnetic or mechanical stabilization reduces the flutter at its source.

The gauge is a point measurement with light filtered noise. A traversing gauge scans the width and mixes time with position, which lengthens the effective delay.

The strip does not stretch between the knives and the gauge, so the delay is exactly L_g/v, and the zinc price, the width and the minimum are constant apart from the grade change of two scenarios.

10.4 The Coating Loop as Implemented

One loop moves the knife pressure to hold the coating. It runs every 0.2 s, the same step as the physics, before each integration step. In the convention of Section 1.6.2 the loop is direct acting, $a = -1$, because more pressure leaves less zinc, so a coating above target raises the pressure. The app's Show Model dialog uses the same name.

10.4.1 Discrete PI Algorithm

The loop uses the velocity-form controller of Section 1.6.2 with the derivative term at zero. Equations (10.8) to (10.10) are Equations (1.13) to (1.16) written for this loop, with the output limits divided by the feedforward gain.

$$e_k = a\,(W_{SP} - W_{pv,k}), \quad a = -1, \quad p_k = K_c e_k \tag{10.8}$$

$$S_k = S_{k-1} + p_k - p_{k-1}, \quad c_k = S_k + I_{k-1} \tag{10.9}$$

$$u_k = \text{clamp}\left(c_k, \frac{P_{min}}{g_k}, \frac{P_{max}}{g_k}\right)$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{\tau_I}\right] \Delta t \tag{10.10}$$

where $\Delta t = 0.2\,\text{s}$, W_{pv} is the loop measurement of the next subsection, and g_k is the feedforward gain of Section 10.4.3, equal to one when feedforward is off. The integral is unwound by back-calculation with a tracking time equal to τ_I. There is no derivative term, no rate limit and no output filter, so the pressure lag τ_P is the only smoothing between the controller and the knives. The command is $P_{cmd,k} = \text{clamp}(g_k u_k, P_{min}, P_{max})$. At reset the integral is seeded with the initial pressure divided by g, and in manual the same expression is solved for I_k against the operator's pressure, which makes both transfers bumpless. The default tuning is $K_c = 0.12\,\text{kPa per g/m}^2$ and $\tau_I = 15\,\text{s}$, and the predictor scenarios use $K_c = 0.3$ and $\tau_I = 4\,\text{s}$.

10.4.2 Three Loop Measurements

The measurement W_{pv} is chosen in the app. The first choice is the gauge itself, $W_{pv} = W_g$. The other two are Smith predictors built on the nominal wiping model, Equation (10.11).

$$\hat{W}(t) = W_0 \left(\frac{v}{v_0}\right)^{0.6} \left(\frac{P}{P_0}\right)^{-0.6} \tag{10.11}$$

which is Equation (10.1) with the distance at D_0, evaluated with the measured speed and the actual knife pressure, and recorded in the same history as W. The predictor in meters compares the gauge with the model's value for the strip the gauge is reading, and the predictor in seconds looks the model up a fixed time $\theta_0 = L_g/v_d$ back, Equation (10.12). The design speed $v_d = 150\,\text{m/min}$ is fixed in the

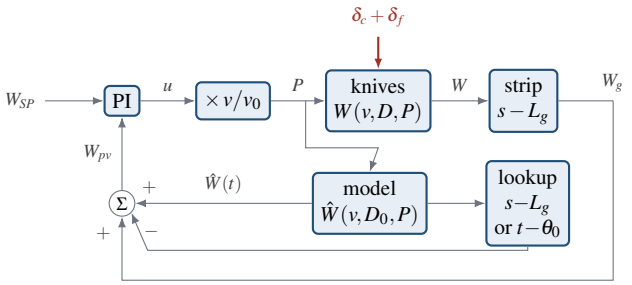


Figure 10.2 Coating loop with the Smith predictor. The model runs on the measured speed and pressure at the nominal distance, its value now enters the loop measurement with the gauge, and its value for the strip the gauge reads is subtracted. The lookup is by strip coordinate in one option and by a fixed 40 s in the other. With the gauge option the model path is absent, and with feedforward off the multiplier is one.

code and does not follow the reference speed v_0 of the wiping correlation, so $\theta_0 = 40\,\text{s}$ for the 100 m gauge distance.

$$\begin{aligned} W_{pv}^m &= \hat{W}(t) + W_g(t) - \hat{W}\left(s - L_g\right) \\ W_{pv}^s &= \hat{W}(t) + W_g(t) - \hat{W}\left(t - \theta_0\right) \end{aligned} \tag{10.12}$$

The loop structure is drawn in Figure 10.2. Substituting Equation (10.6) into the predictor in seconds separates what it sees, Equation (10.13).

$$\begin{aligned} W_{pv}^s = \hat{W}(t) &+ \left[W - \hat{W}\right]_{s-L_g} + \nu \\ &+ \left[\hat{W}(s - L_g) - \hat{W}(t - \theta_0)\right] \end{aligned} \tag{10.13}$$

The first bracket is the model error for one piece of strip, which is the correction a Smith predictor is meant to apply. The second bracket is zero only while $\theta = \theta_0$. At 90 m/min it carries the negative of every change in the model coating for the 27 s between the time the predictor assumes the change reached the gauge and the time it really does, and the loop reads that as a disturbance. The predictor in meters has no second bracket at any speed. Volume 5, Chapter 5 applies a Smith predictor on another moving product, a glass ribbon whose gauge reads glass formed long before.

10.4.3 Speed Feedforward

Equation (10.1) holds the coating constant when P is proportional to v. With feedforward on, the controller output u is the pressure the loop would use at the reference speed, and the knives receive $g_k u_k$ with $g_k = v_k/v_0$, which gives the coating of Equation (10.14).

$$W = W_0 \left(\frac{D}{D_0}\right) \left(\frac{u}{P_0}\right)^{-0.6} \tag{10.14}$$

The speed has disappeared, so a speed change no longer moves the coating, and the gain seen by the controller, $\partial W/\partial u = -0.6W/u$, is the nominal $-1.74\,\text{g/m}^2$ per kPa at every speed for the $76\,\text{g/m}^2$ target. The limits on u are divided by g_k so that the anti-windup acts on the pressure actually available. The feedforward still sees the ramp through the 1 s pressure lag. Scaling a manipulated variable by a measured throughput is the speed feedforward of Section 1.4.6. Chapter 9 feeds the line speed forward to the firing of its furnace, and Volume 3, Chapter 6 computes a coagulant dose from a measured raw-water disturbance and leaves the late measurement a slow trim.

Table 10.3 Linearized PI Loop at the Scenario Operating Points. GM Is the Gain Margin and T_u the Period at Phase Crossover

Operating point	K_c, τ_I	P	$\lvert K_p \rvert$	θ (s)	GM	T_u (s)
150 m/min, case 2	0.12, 15	26.2	1.74	40	3.05	114
150 m/min, case 3	0.25, 10	26.2	1.74	40	1.02	127
90 m/min, case 4	0.12, 15	15.7	2.91	67	1.16	214
90 m/min, feedforward	0.12, 15	26.2[a]	1.74	67	1.93	214
145 g/m^2, case 10	0.12, 15	8.9	9.76	40	0.55	114

[a] controller output u, the pressure at the reference speed

10.4.4 Three Models Kept Apart

The simulation model is Equations (10.1) to (10.7) with the true, disturbed distance. The controller's internal model is Equation (10.11), which shares the plant's parameters and exponents and differs only in the distance, so its error is exactly the coating change caused by $\delta_c + \delta_f$. No model is identified in the application. The tunings are supplied with the scenarios, and the control model used to analyze them in Section 10.5 is the linearization of the plant at each operating point.

10.5 The Linearized Loop

Around an operating point the coating responds to the pressure command as a first-order lag with dead time, $G(s) = K_p \mathrm{e}^{-\theta s}/(\tau_P s + 1)$, where K_p comes from Equation (10.2), θ from Equation (10.6) and $\tau_P = 1$ s. With the PI controller the open loop is $K_c(1 + 1/\tau_I s)G(s)$, and its gain margin and the period at its phase crossover were computed for the operating points of the scenarios. The results are listed in Table 10.3.

The default tuning has a gain margin of 3.05 at 150 m/min. Doubling the gain and shortening the integral time leaves a margin of 1.02, so case 3 sits on its stability limit. A slowdown to 90 m/min multiplies the gain by 1.67 and the delay by 1.67 at once and leaves 1.16. Feedforward restores the nominal gain, and the margin rises to 1.93 against the longer delay alone. On the heavy grade the margin is 0.55 and the linear loop is unstable. The simulated pressure rings near these periods at the 0.2 s step. It crosses its mean upward every 121 to 155 s in the bursts of case 3, its peaks come 191 to 236 s apart after the slowdown of case 4, and case 10 settles into a saturated limit cycle of 144 s.

The delay also bounds what the loop can remove, because the best achievable variance is set by the disturbance content faster than θ [3, 4]. At 150 m/min the coil shape steps last twenty delay times and the flutter correlation time is 1.5 delay times, so the loop can remove the first and barely touch the second. A Smith predictor removes the delay from the stability of the loop, but the gauge still reports each piece of strip θ late, so the predictor cannot lower this floor.

10.6 Scenarios and Measured Results

Table 10.4 gives the eleven scenarios, which follow the pattern of Section 1.6.4. All values below come from runs of the engine at its default seed, and the statistics are those of the coating at the knives weighted by strip length. Unless stated otherwise, the standard deviation, range and strip below the minimum exclude the first 5 min, and the zinc cost is the scorecard's average over the whole run. The scorecard itself accumulates from $t = 0$, so it reads slightly different values, for example 1.39 g/m^2 in case 2 and 4.58 g/m^2 with 9.4 % below the minimum in case 4. The coil shapes are random, so a result that rests on a few coil changes moves with the seed, and the spread over ten other seeds is given where a comparison depends on it. Cases 9 and 10 compare two loop measurements on the same coils by pressing Reset, which restarts the case from $t = 0$ and keeps every setting the learner changed. Figure 10.4 draws the rows of Table 10.5 as a chart.



Figure 10.3 Case 7, the Smith predictor with its delay in meters, at $t = 15$ min after a slowdown from 150 to 90 m/min at 5 min, cropped to the badges, the plant view and the trends. The badges (1) read 76.9 g/m^2 at the gauge, 75.1 g/m^2 at the knives, 14.9 kPa, 90 m/min and a 67 s delay. The air knives (2) sit just above the zinc, the strip path over the cooling tower (3) carries the 100 m that make the delay, and the gauge (4) reads the coating at its end. The histogram of the last 10 min sits well above the 70 g/m^2 minimum and around the 76 g/m^2 target, and the line under it reads 1.17 g/m^2 and $336 an hour of zinc above the minimum. The trends over 15 min show the coating at the knives between 73.4 and 79.0 g/m^2 through the slowdown, the knife pressure falling from about 26 to 15 kPa as the speed ramps down and the delay rises from 40 to 67 s, the unmeasured knife distance stepping 2.9 % closer at a coil change at 9.7 min (5), after which the coating at the knives dips to 71.6 g/m^2, and the average zinc cost.

Table 10.4 Scenarios and What Each Teaches

#	Scenario	Lesson
1	A fixed knife pressure (manual)	open loop, coils and speed
2	A PI loop on the gauge	feedback on a late signal
3	More gain, more variance	the best gain is finite
4	The line slows down	delay and gain change together
5	Speed feedforward	a measured disturbance
6	Smith predictor, delay in s	a delay in the wrong unit
7	Smith predictor, delay in m	a delay in the right unit
8	What a gram is worth	the target and the spread
9	A strip that flutters	variance control cannot remove
10	A heavier coating grade	a gain that grows 5.6 times
11	Challenge, one hour	grade and two speed changes

Table 10.5 Measured Scenarios. σ_W and Range in g/m^2, Below Is the Strip Under W_{min}, Cost Is Zinc Above W_{min} in \$/h

#	Configuration and window	σ_W	Range	Below	Cost
1	manual, 1–20 min	3.45	73.6–83.9	0.0 %	785
1	manual, 1–30 min	10.32	53.7–83.9	22.8 %	531
2	PI, 5–60 min	1.43	70.6–82.8	0.0 %	482
3	PI, K_c 0.25, τ_I 10	4.49	62.9–94.4	7.8 %	504
4	PI, slowdown, 5–40 min	5.08	56.9–89.0	11.6 %	311
5	PI and feedforward	1.25	72.5–78.9	0.0 %	310
6	Smith (s), 1–5 min	0.89	73.2–77.5	0.0 %	—
6	Smith (s), 5–40 min	53.35	28.7–218.2	40.1 %	1312
7	Smith (m), 5–40 min	1.24	71.5–79.2	0.0 %	308
8	PI, W_{SP} 73.5	1.41	68.3–80.1	2.0 %	281
9	PI, flutter 3 %	2.66	67.6–83.3	1.9 %	484
9	Smith (m), flutter 3 %	2.61	67.7–83.5	1.8 %	485
10	PI, heavy, 10–30 min	114.87	43.6–389.5	64.6 %	2915
10	Smith (m), heavy	2.94	75.6[a]–150.1	0.65 %	572

[a] the minimum is the coating at the instant of the grade change

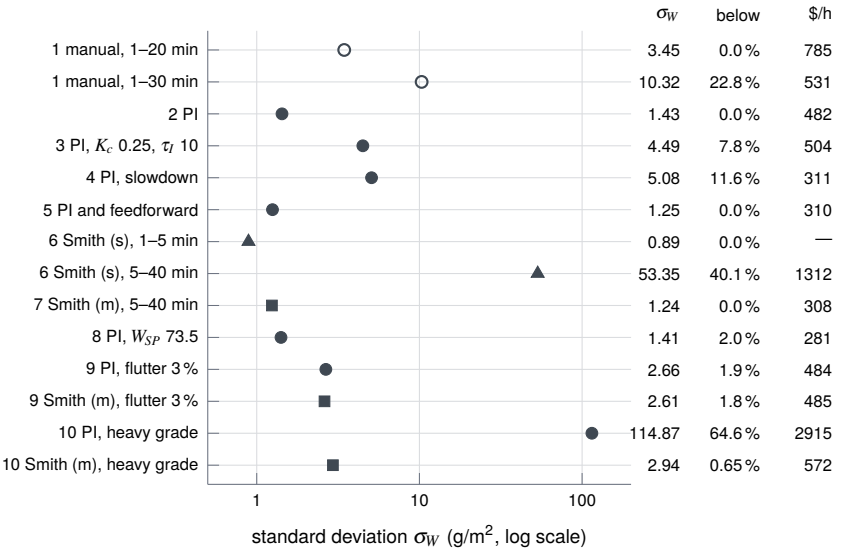


Figure 10.4 The fourteen rows of Table 10.5 as a chart: the standard deviation of the coating on a log scale, with its value, the strip below the minimum and the zinc cost above the minimum in the three columns at the right. Open circles mark the manual runs, filled circles the PI loop on the gauge, triangles the predictor with its delay in seconds and squares the predictor with its delay in meters. Notice that the two points near the right edge, the predictor in seconds after the slowdown and the PI loop on the heavy grade, sit more than a decade to the right of the predictor in meters facing the same slowdown and the same grade.

10.6.1 A Fixed Pressure and a Loop on the Gauge

Held at 26.2 kPa, the coating wanders between 73.6 and 83.9 g/m^2 with the coil shapes over the first 20 min, averaging 79.9 g/m^2 because the first coil change, at 8 min, moves the strip 0.93 mm farther from the knives, and zinc above the minimum costs \$785 an hour. When the line slows to 90 m/min at 20 min, less zinc is dragged out, the coating falls to between 54 and 57 g/m^2, and 22.8 % of the strip

in the half hour is below the minimum. The PI loop on the gauge holds 76.0 g/m^2 with a standard deviation of 1.43 g/m^2 and none of the strip below the minimum, at $482 an hour.

> **TRY IT IN THE APP** **1 · A Fixed Knife Pressure (MANUAL)**
>
> Load scenario 1 and press Run. Watch the coating at the knives on the top trend. For the first 20 min it wanders between 73.6 and 83.9 g/m^2 as the coils change. When the line slows to 90 m/min at 20 min, less zinc is dragged out, and the coating falls to between 54 and 57 g/m^2, below the 70 g/m^2 minimum.

The residual is not mainly flutter. With the flutter set to zero the standard deviation falls only to 1.25 g/m^2, while with the coil shape changes set to zero it falls to 0.87 g/m^2, and removing the gauge noise gives 1.40 g/m^2. Most of the variance left by the loop is the transient after each coil step, which the loop cannot see for 40 s and then removes over a couple of minutes. The default gain is near the best one for this delay. At $\tau_I = 15$ s the standard deviation is 1.68, 1.52, 1.43, 1.40, 1.40 and 1.42 g/m^2 for gains of 0.06, 0.09, 0.12, 0.15, 0.18 and 0.20. In case 3 the loop is at the edge of stability, as Table 10.3 predicts, and rings in bursts after some coil changes. The standard deviation rises to 4.49 g/m^2 and 7.8 % of the strip falls below the minimum. With the coil changes switched off the same tuning still gives 3.30 g/m^2, so the ringing is sustained by the flutter and the noise as well.

Pause. Try Exercise 10.1 before reading on.

10.6.2 A Slowdown, Feedforward and Two Predictors

At 5 min in cases 4 to 7 the line slows to 90 m/min. With the PI loop alone the coating first drops to about 57 g/m^2, because the speed step reaches the knives at once and the gauge 67 s later, and the loop then rings with every coil change. Over 40 min the standard deviation is 5.08 g/m^2 and 11.6 % of the strip is below the minimum. With the speed set to 90 m/min before Run in case 2, so that the line runs at 90 m/min from the start, the same loop still gives 3.36 g/m^2 and 3.9 % below from 5 to 60 min, more than twice the 1.43 g/m^2 at 150 m/min, so the ringing belongs to the operating point and not to the transition. Speed feedforward keeps the coating above 74.6 g/m^2 through the slowdown and lowers the standard deviation to 1.25 g/m^2 with none of the strip below the minimum, no more than the 1.44 g/m^2 of case 2 at 150 m/min over the same window. At 90 m/min from the start feedforward lowers the standard deviation from 3.36 to 2.29 g/m^2, because it restores the nominal loop gain. Over ten other seeds case 4 gives 4.50 to 7.85 g/m^2 and case 5 gives 0.93 to 1.62 g/m^2. At 90 m/min from the start the ten seeds give 1.57 to 2.70 g/m^2 without feedforward, below the 3.36 of the shipped seed, and 1.12 to 1.76 g/m^2 with it, lower on every seed, so the benefit of feedforward does not depend on the coil shapes drawn.

The Smith predictor with its delay in seconds and the tighter tuning holds 0.89 g/m^2 over the first 5 min at 150 m/min. After the slowdown the second bracket of Equation (10.13) is no longer zero, and the loop oscillates with a period of about 45 s and a growing amplitude until the pressure command reaches its upper limit at 1065 s. The coating swings between 28.7 and 218 g/m^2 and 40.1 % of the strip is below the minimum. The same predictor keeping its delay as 100 m of strip stays stable with the same tuning, at 1.24 g/m^2 with none of the strip below the minimum. The wiping model includes the speed, so this predictor sees the slowdown at once, as the feedforward does, and adding feedforward to it changes the result by little (1.20 g/m^2).

> **TRY IT IN THE APP** **6 · A Smith Predictor With Its Delay in Seconds**
>
> Load scenario 6 and run past the slowdown at 5 min. The loop oscillates with a period of about 45 s and a growing amplitude until the pressure command reaches its upper limit at 1065 s, and the coating swings between 28.7 and 218 g/m^2. Then load scenario 7, the same predictor with its delay kept as 100 m of strip. It stays stable at 1.24 g/m^2.

The time-based predictor was also run at constant speeds, as case 2 with the predictor and tuning of case 6 and the line speed set before Run, and measured from 10 min to 60 min. It is stable at 122.5, 125, 130, 135, 140, 160, 180, 200 and 210 m/min, with standard deviations of 1.08 to 1.79 g/m^2. At 120 m/min, where the true delay is 50 s against the assumed 40 s, the standard deviation grows from 1.91 g/m^2 between 10 and 30 min to 10.22 g/m^2 between 40 and 60 min, and at 105 m/min it reaches 68.2 g/m^2. A delay 9 s too short is tolerated and one 10 s too short is not, which the linearized loop of the predictor confirms, while a delay 11 s too long, at 210 m/min, is still stable. Robustness studies of dead-time compensators address exactly this sensitivity [1, 2].

Case 7 ten minutes after the slowdown is shown in Figure 10.3. No strip falls below the minimum over the 40 min, although the coil change at 9.7 min brings the strip 2.9 % closer to the knives and the coating dips to 71.6 g/m^2 before the loop sees it.

10.6.3 The Price of the Spread

Lowering the target of case 2 from 76 to 73.5 g/m^2 cuts the zinc cost from \$482 to \$281 an hour, about \$200 an hour or \$1.75 million over a year of continuous running. The strip below the minimum rises from none to 2.0 %, although the scorecard's normal-theory target for this spread is 73.2 g/m^2. At that target 3.0 % of the strip is below the minimum. A sweep of the target gives 1.39, 1.18, 1.04, 0.92 and 0.45 % below at 74.0, 74.5, 74.7, 74.8 and 75.0 g/m^2, so 1 % below needs a margin of about 4.75 g/m^2, 3.3 standard deviations rather than 2.33. The coil steps give the coating a heavy lower tail, and a target set from the standard deviation alone understates the margin. Each gram per square meter of standard deviation still costs at least $2.33 \times \$81 = \$189$ an hour in target margin at 150 m/min.

> **DID YOU KNOW?** **Twice the Gain, Three Times the Spread**
>
> From the simulator: raising the loop gain from 0.12 to 0.25 and shortening the integral time from 15 to 10 s widens the spread of the coating from 1.43 to 4.49 g/m^2, and the strip below the minimum rises from none to 7.8 % (Table 10.5). On a late gauge a tighter loop gives a wider product.

Tripling the flutter to 3 % raises the standard deviation to 2.66 g/m^2 with the PI loop and 2.61 g/m^2 with the tightly tuned predictor in meters, and the target this spread allows to about 76 g/m^2. The predictor does no better, and over ten other seeds the two differ by at most 0.3 g/m^2, because the gauge still reports the flutter 40 s late. Variance of this kind has to be removed at the strip, not in the controller.

Pause. Try Exercise 10.2 before reading on.

10.6.4 A Heavier Grade and the Challenge

At 10 min in case 10 the minimum rises to 137.5 g/m^2 and the target to 145 g/m^2. The knives drop toward 9 kPa, the gain rises 5.6 times, and the PI loop tuned for the light grade goes into a limit cycle between the pressure limits, with the coating between 43.6 and 389 g/m^2 and 64.6 % of the strip below the new minimum. The predictor in meters carries the gain in its model, and its inner loop sees the model rather than the delayed gauge, so the same tuning settles on the new target at 2.94 g/m^2 with 0.65 % of the strip below the minimum, all of it in the first 8 s after the change, while the coating climbs past the new minimum.

Pause. Try Exercise 10.3 before reading on.

Table 10.6 Challenge: Heavy Grade at 10 min, 90 m/min at 30 min, 180 m/min at 45 min. Strip Below W_{min} Over the Hour

Configuration	Below
as loaded, PI 0.12, 15 s, no feedforward	50.5 %
PI 0.03, 30 s, no feedforward	14.9 %
PI 0.03, 30 s, feedforward	12.2 %
Smith (s) 0.3, 4 s, feedforward	24.2 %
Smith (m) 0.3, 4 s, feedforward (suggested)	0.22 %

The challenge runs one hour with the grade change of case 10, a slowdown to 90 m/min at 30 min and a speed-up to 180 m/min at 45 min. It is met with at most 2 % of the strip below the minimum over the hour and the gauge within $\pm 20\,\mathrm{g/m^2}$ of the target over the last 20 min. The measured configurations are listed in Table 10.6. As loaded the loop goes unstable at the grade change and 50.5 % of the strip is below the minimum. The app's verdict reads 50.45 % and a worst gauge error over the last 20 min of 260.9 g/m^2, so both clauses fail. Detuning the PI loop four times keeps it stable on the heavy grade, but the coating then takes minutes to climb to the new minimum, and 34.7 % of the strip between 10 and 30 min is below it. Without feedforward the slowdown adds 16.4 % of the strip between 30 and 45 min, and with it none. The predictor with its delay in seconds fails at the two speed changes. The predictor in meters with feedforward meets the target with 0.22 % below, all of it in the first 8 s after the grade change, and over the last 20 min, which include the speed-up, the coating at the knives stays between 140.4 and 152.0 g/m^2 with a standard deviation of 1.83 g/m^2, while the gauge stays within 8.0 g/m^2 of the target. The verdict therefore passes 1.78 points inside the limit on the strip below the minimum and 12.0 g/m^2 inside the band, and the app gives this one run whether the suggested settings are revealed before Run or carried through Reset. The same predictor without feedforward gives the same result, because its model already carries the speed. The grade change needs the model's gain, and the speed changes need the delay in meters. Over ten other seeds the suggested settings leave 0.25 to 2.6 % of the strip below the minimum, and two of the ten exceed 2 %, so the margin of the graded run depends on the coil shapes it draws.

10.7 Technology in Practice

The loop of Section 10.4 runs on a coating gauge, a speed measurement and the pressure of two air knives. Figure 10.5 draws it with that hardware in place of the schematic's symbols.

What Measures It

On a working line the coating gauge stands where the chapter puts it, behind the cooling tower, and the air knife settings are adjusted from the deviation between the preset and the measured coating [18]. The gauge measures the top and the bottom coating, and the controller averages each side over one scan of the head [14]. The chapter has one fixed gauge, AT-1001, and treats the two sides as one coating. The predictor reads two more signals, the line speed ST-1002 and the knife pressure PT-1003, because its wiping model needs both.

> **TECHNOLOGY** **Sensor: An X-Ray Gauge Past the Cooling Tower**
>
> An X-ray fluorescence gauge measures a coating by the fluorescent X-rays that its atoms emit when irradiated, or by how much the coating absorbs the X-rays coming back from the steel beneath. The method serves coatings up to about 20 µm thick, galvanized strip among them [19]. The coating must be solid first, so the gauge stands behind the cooling tower, 100 to 200 m of strip past the knives [18]. Its head often traverses the width, so the points it measures trace a dashed diagonal line on the moving strip [14]. The chapter's gauge AT is a fixed point 100 m past the knives with 1 g/m^2 of filtered noise.

line controller: AC-1001 with AY-1001 and SY-1002 executes every 0.2 s, the step of the physics, against a gauge delay of 40 s at 150 m/min and 67 s at 90 m/min

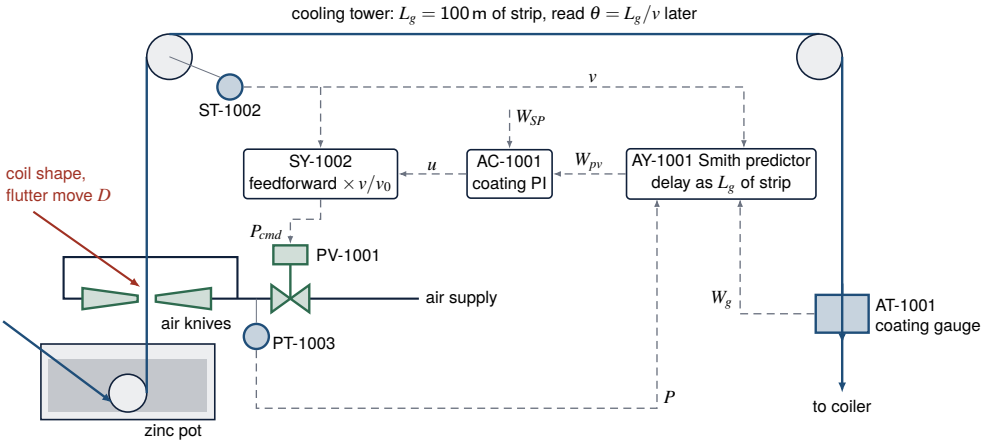


Figure 10.5 The coating loop drawn with its hardware. The air knives blow both sides of the strip just above the zinc pot, fed through the knife pressure PV-1001. The strip travels 100 m over the cooling tower to the coating gauge AT-1001, an X-ray gauge in practice. The Smith predictor AY-1001 combines the gauge with its wiping model, which reads the line speed ST-1002 and the knife pressure PT-1003, and looks the model up by strip length. The PI block AC-1001 acts on that signal, and the feedforward SY-1002 scales its output by the speed. Notice how far apart the knives that act and the gauge that sees are along the strip.

What Moves It

Real lines blow each side of the strip with its own knife. The controller computes a pressure for the top and the bottom knife, and it corrects a difference between the two sides by moving the knife distances [14]. The chapter has one pressure for both knives, PV-1001, and a knife distance that only the disturbances move.

> **TECHNOLOGY** **Final Element: The Air Knife**
>
> An air knife is a slot nozzle that blows a flat gas jet at the strip just above the bath. The pressure gradient and the shear stress of the jet on the liquid zinc set the final film thickness, and most of the zinc runs back into the pot [20]. The handle is the pressure in the knife chamber [14]. It cannot be raised without end. A harder jet brings tonal noise, splashing and an uneven coating, which is why lines hold their speed down instead [20], and splash leaves first from the strip edges [21]. A pressure difference between the two knives also shakes the strip [22]. The chapter's knives follow the pressure command through a 1 s lag between 2 and 70 kPa.

Where the Controller Runs

A line of this kind is run from a layered control system. A supervisor sends set points to distributed controllers, which drive the actuators from sensor feedback, and loops at both levels run with cycle times from milliseconds to minutes [23]. On a galvanizing line the coating control has a preset part that acts when a weld passes the knives and a feedback part that trims along the coil [14]. The chapter keeps only the feedback part. Its PI block AC-1001 executes every 0.2 s, the step of the physics and far shorter than the 40 s delay at 150 m/min. The predictor AY-1001 and the speed feedforward SY-1002 belong to the same block, and the coating target is the set point that a supervisor sends to it.

IN PRACTICE **How a Coating Loop Is Run on a Line**

On a working line the coating is set in two ways. When the weld between two coils passes the knives, a new pressure and knife distance are preset from a model that learns from the coating measured on earlier coils, because no feedback can reach the head of a coil in time. Feedback then trims through a Smith predictor whose delay is taken from the line speed and the distance from the knives to the gauge. It ignores errors within $\pm 1\,\mathrm{g/m^2}$, a dead band, and a feedforward meets the changes of line speed that the annealing furnace ahead of the pot calls for [14].

What Is Modernizing It

Three developments touch the loop of this chapter. Electromagnetic actuators now act on the strip without touching it. One system of five pairs of actuators and controllers, designed on a reduced finite element model of the strip, was shown effective for vibration and flatness control in simulation and in experiments under various tensions [22]. Data-driven models predict the coating from production data. Models trained for three bath chemistries found the knife pressure, the strip speed and the knife distance to be the main factors, and the sensitivity of the coating to each of them differed with the coating type [18], which is the moving gain of Equation (10.2) seen from the data. The knives themselves are changing too. A multi-slot knife with a weak auxiliary jet beside the main jet wiped a thinner coating than a single slot under similar main jet conditions in laboratory tests [20].

10.8 Limitations

The wiping correlation is a power law with fixed exponents, and it ignores the nozzle geometry, the jet tilt, the zinc temperature and the splashing limit [7].

OPEN CHALLENGE **Thinner Coatings Without Splashing**

Zinc resources are limited, and saving zinc means thinner coatings. Gas wiping meets two limits: the liquid film splashes as the jet pressure rises, and the wiping capacity of the equipment runs out. Hot-dip sheet with less than $30\,\mathrm{g/m^2}$ is difficult to make as a standard product, and zinc that begins to solidify under the cooling jet appears to resist wiping [21]. For now lines keep the strip speed low to stay clear of these troubles [20]. Progress would be a knife that wipes thinner at full speed, and a loop that holds its target closer to the minimum.

The two sides and the full width are one coating, so crossbow, edge build-up and side-to-side control are absent. The controller's model shares every parameter of the plant except the knife distance. A model with the wrong reference coating or exponent would add a steady model error that the gauge corrects slowly, which none of the scenarios tests.

The knife distance is only a disturbance. Real lines also move the knives, and they stabilize the strip mechanically or electromagnetically, which is the remedy case 9 points to.

OPEN CHALLENGE A Strip That Will Not Hold Still

Curvature and vibration of the strip at the knives make the coating uneven. More tension calms the strip but wears out the bushes and sleeves of the sink roll, and rolls that squeeze the strip can scratch it [22]. Active control is hard. The strip vibrates in many modes at once, only a non-contact position sensor can be used on a moving strip, and a controller built on the full state risks spill-over into modes it does not model [22]. Vibration also grows harder to reduce as the strip speeds up [24]. Progress would be a strip held still at full speed, which would remove the variance that case 9 shows the coating loop cannot.

The gauge is a fixed point sensor, without the scan period of a traversing gauge.

10.9 Conclusion

The galvanizing line combines a dead time fixed in meters, a process gain that moves with the coating and the speed, and a product specification that turns variance into zinc. A PI loop on the gauge holds 1.43 g/m^2, most of it the transient after each coil step, and doubling its gain triples the spread. A slowdown to 90 m/min raises the gain and the delay together and leaves 11.6 % of the strip below the minimum, and speed feedforward removes that loss by restoring the nominal gain. A Smith predictor that keeps its delay in seconds oscillates once the true delay exceeds its own by about 10 s, while the same predictor keeping the delay as 100 m of strip stays at 1.24 g/m^2 and also carries the 5.6-fold gain of a heavy grade. Lowering the target by 2.5 g/m^2 saves about $200 an hour, but the lopsided spread needs a margin of 3.3 standard deviations for 1 % of the strip below the minimum. The model treats both sides and the width as one coating with fixed wiping exponents. A fuller model would add two-sided knives, knife position as a second actuator and a traversing gauge. The simulator offers an accessible way to practice the control of a process whose dead time is set by a distance.

KEY POINTS

- The coating can be measured only after it has frozen, a fixed length of strip past the knives, so the gauge reads it late, and later still when the line slows.
- A PI loop on the gauge holds 1.43 g/m^2, and doubling its gain triples the spread.
- A slowdown raises the gain and the delay together, and speed feedforward removes the loss by restoring the nominal gain.
- A Smith predictor that keeps its delay in seconds oscillates once the true delay exceeds its own by about 10 s, while the same predictor keeping the delay as 100 m of strip stays stable.

Where else this applies. A delay set by a distance and a throughput also governs the glass ribbon of Volume 5, Chapter 5 and the strip gauge of Chapter 8.

References

[1] J. E. Normey-Rico, T. L. M. Santos, R. C. C. Flesch, and B. C. Torrico, “Control of dead-time process: From the Smith predictor to general multi-input multi-output dead-time compensators,” *Frontiers in Control Engineering*, vol. 3, p. 953768, 2022.

[2] Z. Palmor, “Stability properties of Smith dead-time compensator controllers,” *International Journal of Control*, vol. 32, no. 6, pp. 937–949, 1980.

[3] T. J. Harris, “Assessment of control loop performance,” *The Canadian Journal of Chemical Engineering*, vol. 67, no. 5, pp. 856–861, 1989.

[4] L. Desborough and T. Harris, “Performance assessment measures for univariate feedforward/feedback control,” *The Canadian Journal of Chemical Engineering*, vol. 71, no. 4, pp. 605–616, 1993.

[5] A. Marder, “The metallurgy of zinc-coated steel,” *Progress in Materials Science*, vol. 45, no. 3, pp. 191–271, 2000.

[6] J. A. Thornton and H. F. Graff, “An analytical description of the jet finishing process for hot-dip metallic coatings on strip,” *Metallurgical Transactions B*, vol. 7, no. 4, pp. 607–618, 1976.

[7] A. Gosset and J.-M. Buchlin, "Jet Wiping in Hot-Dip Galvanization," *Journal of Fluids Engineering*, vol. 129, no. 4, pp. 466–475, 2007.

[8] M. Bauer and I. K. Craig, "Economic assessment of advanced process control – A survey and framework," *Journal of Process Control*, vol. 18, no. 1, pp. 2–18, 2008.

[9] U.S. Geological Survey, "Mineral commodity summaries 2025: Zinc," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[10] A. R. Marder and F. E. Goodwin, "History of zinc-coated steel," in *The Metallurgy of Zinc Coated Steels*. Elsevier, 2023, pp. 1–14.

[11] C. H. Ellen and C. V. Tu, "An Analysis of Jet Stripping of Liquid Coatings," *Journal of Fluids Engineering*, vol. 106, no. 4, pp. 399–404, 1984.

[12] T. Noe and T. Schuerger, "Computer Control of Zinc-Coating Weight on the Fairfield No. 4 Galvanizing Line," *IFAC Proceedings Volumes*, vol. 8, no. 1, pp. 404–411, 1975.

[13] D. L. Grinberg, "Regulating coating thickness in continuous galvanizing of steel sheet," *Metallurgist*, vol. 21, no. 3, pp. 185–188, 1977.

[14] K.-T. Shin, H.-D. Park, and W.-K. Chung, "Synthesis method for the modelling and robust control of coating weight at galvanizing process," *ISIJ International*, vol. 46, no. 10, pp. 1442–1451, 2006.

[15] E. Elsaadawy, G. Hanumanth, A. Balthazaar, J. McDermid, A. Hrymak, and J. Forbes, "Coating Weight Model for the Continuous Hot-Dip Galvanizing Process," *Metallurgical and Materials Transactions B*, vol. 38, no. 3, pp. 413–424, 2007.

[16] Y.-T. Kim, "An automatic coating weight control for continuous galvanizing line," in *2008 International Conference on Control, Automation and Systems*, 2008, pp. 1715–1720.

[17] N. Guelton and A. Lerouge, "Coating weight control on ArcelorMittal's galvanizing line at Florange Works," *Control Engineering Practice*, vol. 18, no. 10, pp. 1220–1229, 2010.

[18] H. Wang, K. Zha, K. Yang, W. Zhang, S. Zhu, F. Xu, and Q. Xie, "Data-driven prediction and process parameter analysis of zinc coating thickness in hot-dip galvanizing," *ISIJ International*, vol. 66, no. 7, pp. 857–866, 2026.

[19] International Atomic Energy Agency, "Technical data on nucleonic gauges," International Atomic Energy Agency, Vienna, IAEA-TECDOC 1459, 2005, iSBN 92-0-107805-6.

[20] A. Yahyaee Soufiani, J. R. McDermid, A. N. Hrymak, and F. E. Goodwin, "A parametric study of a multi-slot air knife for coating thickness reduction," *ISIJ International*, vol. 60, no. 5, pp. 1040–1051, 2020.

[21] H. Kobayashi, G. Takeda, K. Katoh, and T. Wakimoto, "Coating weight reduction technology in gas wiping of hot-dip galvanizing on steel strip," *ISIJ International*, vol. 64, no. 11, pp. 1716–1722, 2024.

[22] W. Jung, Y. J. Jang, S. Lim, and S. Won, "Active vibration and flatness control of a strip in a continuous galvanizing line using positive position feedback control," *ISIJ International*, vol. 53, no. 5, pp. 854–865, 2013.

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

[24] J. Li, Y.-H. Yan, X.-H. Guo, and Y.-Q. Wang, "Research on vibration control method of steel strip for a continuous hot-dip galvanizing line," *ISIJ International*, vol. 52, no. 6, pp. 1072–1079, 2012.

10.10 Exercises

10.1 ⌕ Concept check · 2 min

In scenario 3, More gain, more variance, the PI loop on the gauge runs with $K_c = 0.25$ and $\tau_I = 10\,\text{s}$ instead of 0.12 and 15 s. Measured from 5 to 60 min, the standard deviation of the coating rises from 1.43 to 4.49 g/m^2. Which explanation is right?

(a) The knife pressure now reaches its limits after each coil change, so the clipped loop cannot finish its corrections

(b) The coating becomes less sensitive to pressure as the loop raises it, so the loop is detuned just when it is needed

(c) The gain margin falls from 3.05 to 1.02, so a loop that sees each correction 40 s late rings near its stability limit

(d) The loop now follows the gauge noise with its 2 s correlation time, and the knives copy that noise onto the strip

10.2 Concept check · 1 min

In scenario 9, A strip that flutters, the flutter is tripled to 3 %. The PI loop leaves 2.66 g/m^2 and the tightly tuned Smith predictor in meters 2.61 g/m^2. Why does the predictor do no better?

(a) Its tighter tuning makes the loop ring, which cancels what it gains on the flutter

(b) The flutter is too small against the coil steps to change the spread of either loop

(c) The flutter changes the line speed, which the predictor handles no better than the gauge

(d) Its model runs at the nominal distance, so the flutter reaches it only through the late gauge

10.3 Concept check · 2 min

In scenario 10, A heavier coating grade, the target rises from 76 to 145 g/m^2 at 10 min. The PI loop tuned for the light grade goes into a limit cycle, while the Smith predictor in meters, with a higher gain of 0.3, settles on the new target. What explains the difference?

(a) The process gain grows 5.6 times at the low pressure, and the predictor's inner loop sees it without the delay

(b) The heavy grade lengthens the gauge delay, which the predictor in meters tracks and the gauge loop does not

(c) The higher minimum leaves less room below the target, so the same spread puts more strip below the minimum

(d) The heavy grade needs a pressure below the knives' lower limit, which only a predictor can work around

10.4 Theory and calculation · 6 min

Use the wiping correlation of Equation (10.1) with $W_0 = 70\,\mathrm{g/m^2}$, $v_0 = 150\,\mathrm{m/min}$, $P_0 = 30\,\mathrm{kPa}$, $D = D_0$ and the gauge 100 m after the knives.

(a) Find the pressure that gives 76 g/m^2 at 150 m/min, and the process gain K_p there. *Checkpoint:* 26.16 kPa and −1.74 g/m^2 per kPa.

(b) Repeat at 90 m/min, and give the gauge delay at both speeds. *Checkpoint:* 15.7 kPa, −2.91 g/m^2 per kPa, 40 s and 66.7 s.

(c) Repeat for 145 g/m^2 at 150 m/min, and show that the ratio of the two gains at 150 m/min is $(145/76)^{1+1/0.6}$. *Checkpoint:* 8.91 kPa, −9.76 g/m^2 per kPa, ratio 5.6.

10.5 Theory and calculation · 6 min

At 90 m/min the linearized loop of Section 10.5 has $G(s) = K_p \mathrm{e}^{-\theta s}/(\tau_P s + 1)$ with $\theta = 66.7\,\mathrm{s}$ and $\tau_P = 1\,\mathrm{s}$, under the default PI tuning $K_c = 0.12$ and $\tau_I = 15\,\mathrm{s}$. Table 10.3 gives the period at phase crossover as 214 s.

(a) Check that the open-loop phase is -180° at $\omega = 2\pi/214$, adding the phases of the integral term, the delay and the lag. *Checkpoint:* $66.2^\circ + 112.1^\circ + 1.7^\circ = 180.0^\circ$.

(b) Find the gain margin with $|K_p| = 2.91$, the gauge loop of case 4. *Checkpoint:* 1.16.

(c) Find the gain margin with speed feedforward, which makes the gain seen by the controller 1.74 g/m^2 per kPa, and say why the period does not change. *Checkpoint:* 1.93.

10.6 Theory and calculation · 6 min

The Smith predictor in seconds looks its model up a fixed $\theta_0 = 40$ s back, while the gauge reads strip $L_g = 100$ m back.

(a) Find the true delay and the mismatch $\theta - \theta_0$ at 120, 130 and 210 m/min. *Checkpoint:* 50 s and +10 s, 46.2 s and +6.2 s, 28.6 s and −11.4 s.

(b) The chapter reports that a delay 9 s too short is tolerated and one 10 s too short is not. Find the speed at which the true delay is $\theta_0 + 9$ s. *Checkpoint:* 122.4 m/min.

(c) At a constant 120 m/min the model coating $\hat{W}$ steps by $\Delta\hat{W}$ at $t = 0$ and the model is otherwise exact. Using Equation (10.13), give the second bracket as a function of time. *Checkpoint:* $-\Delta\hat{W}$ from 40 to 50 s, zero otherwise.

10.7 App challenge · 8 min

Use the default seed and speed 120, and change no setting. Near the end of each run lower the speed to 5 and press Pause when the clock reads 60 min.

(a) Load scenario 2, A PI loop on the gauge, run it, and read the Coating, mean $\pm$ σ, Below the minimum and Zinc above the minimum tiles. *Checkpoint:* 76.0 ± 1.39 g/m^2, 0.00 %, \$482 an hour.

(b) Load scenario 3, More gain, more variance, and read the same tiles. *Checkpoint:* 76.0 ± 4.3 g/m^2, about 7.1 %, about \$500 an hour.

Read the standard deviation within 0.2 g/m^2 and the strip below within 0.2 points. Explain why the zinc cost hardly changes while the spread triples, and what the spread would cost if the target were set from it.

10.8 App challenge · 8 min

Use the default seed and speed 120, change no setting, and pause each run when the clock reads 40 min. Read Coating, mean $\pm$ σ, Below the minimum and Gauge delay.

(a) Scenario 4, The line slows down. *Checkpoint:* 75.4 ± 4.58 g/m^2, 9.3 % to 9.4 %, 67 s.

(b) Scenario 5, Speed feedforward. *Checkpoint:* 75.9 ± 1.18 g/m^2, 0.00 %, 67 s.

(c) Scenario 7, A Smith predictor with its delay in meters. *Checkpoint:* 75.8 ± 1.18 g/m^2, 0.00 %, 67 s.

Read the standard deviations within 0.05 g/m^2. Explain why two different remedies give the same result after the slowdown.

10.9 App challenge · 8 min

Load scenario 2, A PI loop on the gauge. Before Run, set Loop measurement to Smith, delay in s, Coating loop gain to 0.30, Coating loop integral time to 4 s and Line speed to 120 m/min. Run at speed 120 with the default seed and pause when the clock reads 60 min. Then press Reset, set Line speed to 130 m/min, keep the other settings and run the same way.

(a) Read Coating, mean $\pm$ σ, Below the minimum and Gauge delay at 120 m/min, the standard deviation within 0.2 g/m^2. *Checkpoint:* 76.0 ± 6.2 g/m^2, about 12.5 %, 50 s.

(b) Read the same tiles at 130 m/min, the standard deviation within 0.05 g/m^2. *Checkpoint:* 76.0 ± 1.12 g/m^2, 0.00 %, 46 s.

Using the mismatch of each run, explain why 10 m/min decides between a stable and an unstable loop.

10.10 Application area · 4 min

A line holds the coating with the PI loop of case 2 at 150 m/min and 1.43 g/m^2. The chapter measures 1.25 g/m^2 with the flutter set to zero and 0.87 g/m^2 with the coil shape changes set to zero, and 2.66 g/m^2 when the flutter is tripled to 3 %. Each gram per square meter of standard deviation costs at least $189 an hour in target margin.

(a) An electromagnetic stabilizer would remove the flutter, and better coil shape or knife position control would remove the coil steps. Estimate the margin each saves on the line as it is. *Checkpoint:* About $34 and $106 an hour.

(b) Estimate what the stabilizer saves on a line whose strip flutters at 3 %, if it brings the flutter back to 1 %. *Checkpoint:* About $232 an hour.

Say why these estimates are lower bounds for the coil steps in particular.

10.11 Application area · 3 min

The coating gauge of the simulated line is 100 m of strip after the knives, beyond the cooling tower.

(a) Say why the gauge cannot simply be placed at the knives. *Checkpoint:* No numeric answer.

(b) A redesigned cooling section would allow a gauge 60 m after the knives. Find its delay at 150 and 90 m/min. *Checkpoint:* 24 s and 40 s.

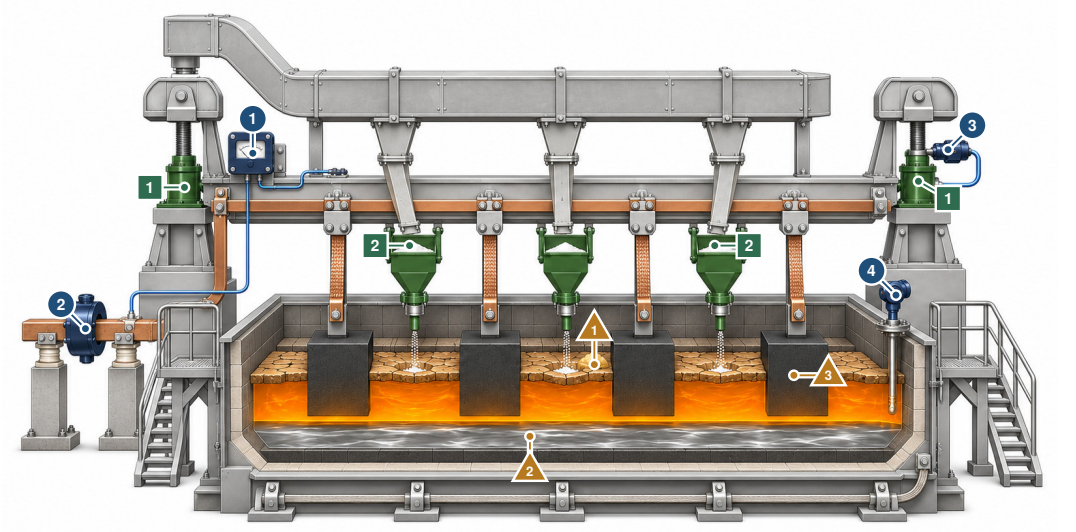
Say what the shorter delay would buy and what it would not change.

10.12 Review the agent · 5 min

An AI assistant is asked to pass the challenge, A heavy grade at three speeds, and replies as follows. “Keep the gauge as the loop measurement, because a predictor adds a model that can be wrong. Table 10.3 gives the default PI loop a gain margin of 0.55 on the heavy grade, so cut the gain four times, to 0.03 kPa per g/m^2, with an integral time of 30 s. That restores a margin of about 3. Turn speed feedforward on, so the slowdown and the speed-up do not move the loop gain. Both clauses will pass.” Find the flaw in the proposal and say how it could be checked.

CHAPTER 11

Alumina Feed Control Through a Non-Monotone Resistance in a Smelting Cell Simulator



The controller moves 1 the jacks that raise and lower the anode beam 2 the feeders that drop alumina into the bath

The controller reads 1 the voltage across the cell 2 the current through the cell 3 the height of the anode beam 4 a daily dip of the bath

What disturbs it 1 a wet lot of alumina that dissolves slowly 2 the rolling of the metal pad 3 an anode change

Notice that no device reads the alumina dissolved in the bath, so the controller judges it from the voltage and the current alone. Illustration, not to scale.

IN PLAIN TERMS **Feeding a Smelting Cell It Cannot See Into**

The main idea. Aluminum is made by sending a huge electric current through hot melted salt that holds a white powder called alumina. Feeders drop small doses of the powder into the bath, and a computer decides when by watching the voltage.

The challenge. The voltage is lowest in the middle, so one reading can mean too little powder or too much. If the powder runs low, the voltage jumps many times higher and the cell gives off strong greenhouse gases.

Why it matters. Feeding well saves power and keeps those gases out of the air.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the fact box of the Introduction, the history box of Section 11.2, the try-it and fact boxes of Section 11.6, 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 11.6, Section 11.7, Section 11.8 and its open challenge

The control question A smelting cell is fed through a pseudo-resistance that has a minimum, so one reading cannot say on which side of it the alumina lies. How does a controller keep the alumina on the steep side, a few tenths of a percent from the anode effect, without ever measuring it, and what does each part of the controller buy in alumina, superheat, energy and anode-effect time?

Learning objectives

- Explain why a reading with a minimum cannot close a simple loop, and compute the slope of the curve from its three voltage terms.
- Describe the tracking cycle of demand feeding and show why it absorbs a constant error in the feeders' dose.
- Place a slope threshold and filter window between the pad's noise and the anode-effect wall with the half-window rule.
- Relate the target resistance to the heat balance, the ledge and the sludge, and judge a daily-dip inference for bias.
- Price an anode effect in minutes and CO_2-equivalent, and name the disturbances that only a margin can absorb.

Prerequisites Sections 1.4.6 (soft sensors corrected by a slower measurement), 1.4.4 (an endpoint read from a rate), 1.4.2 (thermal inertia), 1.6.2 and 1.6.4. Chapter 4 and Chapter 8 help but are not required.

The app https://apmonitor.com/apps/aluminum-smelting/

See also Volume 7, Chapter 4 for regulation at the peak of a plant characteristic, where the gain changes sign. Volume 5, Chapter 12 for a sensor that rises on both sides of the right dose and Volume 5, Chapter 3 for a loop closed on a slow analyzer. Volume 9, Chapter 8 for an operating point chosen against a moving optimum and a constraint.

11.1 Introduction

Primary aluminum is made by the Hall–Héroult process, in which alumina dissolved in an electrolyte of molten cryolite and aluminum fluoride is reduced on carbon anodes, and modern potlines run at about 300 to 600 kA [1]. The average alumina concentration is usually kept within 2 to 4 wt%, the most modern smelters need close to 13 kWh for a kilogram of aluminum against a world average near 14, and a cell whose alumina runs too low has an anode effect, a cell voltage of perhaps 30 to 40 V with CF_4 and C_2F_6 in the anode gas [1]. Over 100 years CF_4 carries a global warming potential of 6630 against CO_2 and C_2F_6 one of 11 100 [2]. Cell control feeds the alumina by watching how the cell resistance changes with it [3], through a pseudo-resistance computed from the cell voltage, a quantity easily confused with the cell's true ohmic resistance [4], and measured against bath samples that pseudo-resistance falls to a minimum and climbs again toward the anode effect [5].

✦ DID YOU KNOW? Seventy-Two Million Tons a Year

The world's smelters made about 72 million metric tons of primary aluminum in 2024, an estimate, 43 million of them in China. Their capacity at the end of the year was about 78 million tons [6].

The minimum defeats the obvious ways of feeding a cell. A resistance looks like a measure of alumina, yet in this cell 8.68 µΩ stands for 2.5 % or 5.14 % alumina, and while the alumina falls from 4.1 to 2.7 % the filtered reading wanders only between 8.595 and 8.641 µΩ. Feeding what Faraday's law says the cell consumes looks safe, yet feeders that deliver 2.8 % less than they are counted for take the cell from 3.50 % into the anode effect at 32.5 h, while until then the controller's tally of shots never strays more than 1.8 kg from balance and the true shortfall reaches 165 kg. A lower slope threshold looks like an earlier warning, yet at 0.03 µΩ/h the metal pad's noise ends the underfeeds early, and the alumina climbs past the minimum to 4.03 % and spends 31 min above its band, on the side where an underfeed lowers the resistance. Demand feeding at the default 0.10 µΩ/h keeps the alumina between 2.14 and 3.07 % for two days after its first underfeed, with no anode effect and

no measurement of alumina. A block diagram with one gain from feed to resistance hides all of it, because the sign of that gain flips at the minimum.

The difficulty belongs to a class. Whenever a measured output has an extremum in the variable that matters, a reading cannot say which side of the extremum the plant is on, and the controller has to move the input on purpose to learn the sign of the local slope. Extremum seeking control does this with a deliberate perturbation, and its first stability proof allowed a general nonlinear dynamic plant whose equilibrium map has a maximum [7]. A photovoltaic array's power against its operating voltage has such a maximum, and at least 19 distinct methods have been introduced to track it [8]. A burner whose efficiency peaks at one air ratio has the same shape, and so has a braking tire's force against its slip, which Volume 7, Chapter 4 regulates at the peak. The tailpipe sensor of Volume 5, Chapter 12 rises on both sides of the right urea dose, so that a high reading may mean too little urea or too much. The smelting cell is the sharper teaching case for two reasons. Its useful operating point is not the extremum but a band on the steep side, a few tenths of a percent from a failure, and the plant moves its own measurement for reasons that are not alumina, an anode change, a tap, anodes that burn and metal that rises, which a slope detector must not mistake for an answer.

The curves that make a cell hard to feed were measured on industrial cells by sampling the bath through a feeding cycle [5], and the feeding, resistance regulation and heat-balance control of potlines are documented as practice [3, 9]. A real cell does not let the alumina be watched while the controller cannot see it, an anode effect be provoked on purpose, or the same noise be replayed against another threshold. This chapter does all three on a simulator of one 300 kA prebake cell whose controller cannot invert its only continuous measurement, and it measures how much alumina, superheat, energy and anode-effect time each part of a cell controller is worth. Demand feeding, resistance regulation and a daily correction of the heat balance follow the practice the literature reviews. To them the chapter adds a half-window rule that places the end of an underfeed, from the curve's slope, the drift of the anode–cathode distance and the fit's window, within 0.03 % of alumina of the engine at the default window. The model has 11 states and 17 ledgers, its thresholds are declared teaching values, and its controller sees only the voltage, the current, the beam's encoder, the crew's signals and a sampled bath dip, so that its view and the truth are logged side by side. Nine cases and a challenge separate the curve, the threshold, the anode effect, two crew operations, the heat balance and an infeasible wet lot.

Section 11.2 traces the domain from the first cells to the process computer. Section 11.3 builds the cell: the alumina between a shot and the anode, the voltage and the control room's pseudo-resistance, the bath, ledge and superheat, the beam and the crew, the anode effect and the metal pad's noise. Section 11.4 shows why the curve has a minimum, what slope an underfeed produces, where an underfeed ends and how the beam moves resistance and heat together. Section 11.5 describes the controller loop by loop, and Section 11.6 reports the measured runs of every case and of the challenge. Section 11.7 describes the real sensors, final elements and control platform, and Section 11.8 lists what the model leaves out, and Section 11.9 sums up.

11.2 Pots, Point Feeders and Process Computers

> ⧖ **HISTORY** **Invented Twice in One Year (1886)**
>
> Charles Martin Hall in the United States and Paul Héroult in France were both born in 1863, and each independently electrolyzed alumina dissolved in molten cryolite. Héroult bought a dynamo of 400 A and 30 V for his experiments, and his French patent dates from 23 April 1886. Hall applied for his U.S. patent on 9 July 1886, and his carbon anode was consumed in the overall reaction. The first commercial company started on 18 September 1888 with steel cells lined with baked carbon and cylindrical carbon anodes 3 inches across. The cells were called pots and the electrolyte bath, a terminology that has survived, and apart from scale, hooding and fume recovery a prebake cell is still built the same way. World production grew from 4682 t in 1900 to 49.7 million metric tons in 2013 [10].

The process dates from 1886, as the history box above tells, and the cell has changed since in its anode, its feeding and its control. The anode changed first. The Søderberg anode, a continuous and self-baking anode, has an electrical resistivity about 30 % higher than a prebaked one and makes its baking fumes hard to collect, and Søderberg cells are being replaced by cells whose anodes are baked at 1150 to 1200 °C in a separate furnace [1]. Feeding changed next. Older designs received large and infrequent additions of alumina, while modern cells carry point feeders that dose it automatically, which leaves anode changing as the most labor-intensive manual routine [1]. A cell of 16 to 40 prebaked anodes changes about one anode a day, and the change is the largest operating disturbance a prebake cell sees [1].

Computers came to the cell in the early 1960s, first to set its energy input by moving the anode beam and to feed alumina by watching how the cell resistance changed with the alumina concentration [3]. A paper of 1966 presented computer control of reduction cells as the maximum production for a given power input and, with automatic feeding added, the possibility of a completely automated potline [11]. Point breaking and feeding in prebake cells, combined with alternating overfeeding and underfeeding, then brought the alumina concentration to 2 to 3 %, which has remained an optimum for cell performance, and AlF_3 came under computer control through automatic feeding [3]. A review of automatic cell control lists its objectives as alumina feeding, resistance regulation through the anode–cathode distance (ACD), bath chemistry and the heat balance, with the routine operations on the cells accounted for [9]. Pseudo-resistance curves measured against bath samples on five types of cell put their minimum at 5.0 to 5.5 % alumina and their anode effect at 1.6 to 2.2 %, 100 to 300 mV above the minimum, and showed a fall of about 100 mV within 30 min of a switch from overfeeding to underfeeding, which the study suggested may come from sludge dissolving [5]. Cell voltages fell for 50 years and have leveled off at 4.1 to 4.2 V in cells above 200 kA [12]. For the heat balance, plant data in which the energy disturbance dominated led to a strategy of nearly constant AlF_3 addition with the energy manipulated through the resistance reference [13].

The anode effect became an environmental problem in the 1990s, when the high global warming potential of perfluorocarbons raised public concern and anode effects were held to be their only significant source [2]. By 1994 the industry was voluntarily exploring how to reduce the number and the duration of anode effects [14]. A measurement program sponsored by the U.S. Environmental Protection Agency at seven smelters found a clear trend toward lower emissions with fewer and shorter anode effects, with a scatter that pointed to other operational factors [15]. Anode effects are triggered by low alumina near the anode surface, start on a single anode, spread until the whole cell is on anode effect, and end only when at least part of the anode–cathode gap is short-circuited [16]. The voluntary programs cut perfluorocarbon emissions in ten countries by 31 to 78 % against 1990, 46 % on average [2], prediction and practical elimination of anode effects reduced them to a minimum [3], and a modern prebake cell can run for more than a week, even for months, without one [1]. The simulator puts these pieces into one cell and keeps the true alumina beside every decision its controller makes.

11.3 One Prebake Cell at 300 kA

The simulated cell has 24 prebaked anodes on a line current of 300 kA, modeled as one well-mixed bath over one metal pad, as shown in Figure 11.1. Its 11 integrated states are the dissolved alumina m_a, the alumina dispersing from dry and from wet shots, m_s and m_w, the sludge m_k, the bath solvent m_b, the excess aluminum fluoride m_f, the temperature T shared by bath and metal, the ledge thickness L, the metal pad's depth h_m, the beam position b and the anodes' burn-off z. Seventeen ledgers integrated with them record what was fed, consumed, made, tapped, burned, heated, lost and emitted, so that each balance closes from the state alone. The feeder shots, the start and end of an anode effect, the metal pad's noise and the timed disturbances are applied between integration steps and stay out of the derivative function. Table 11.1 classifies the signals.

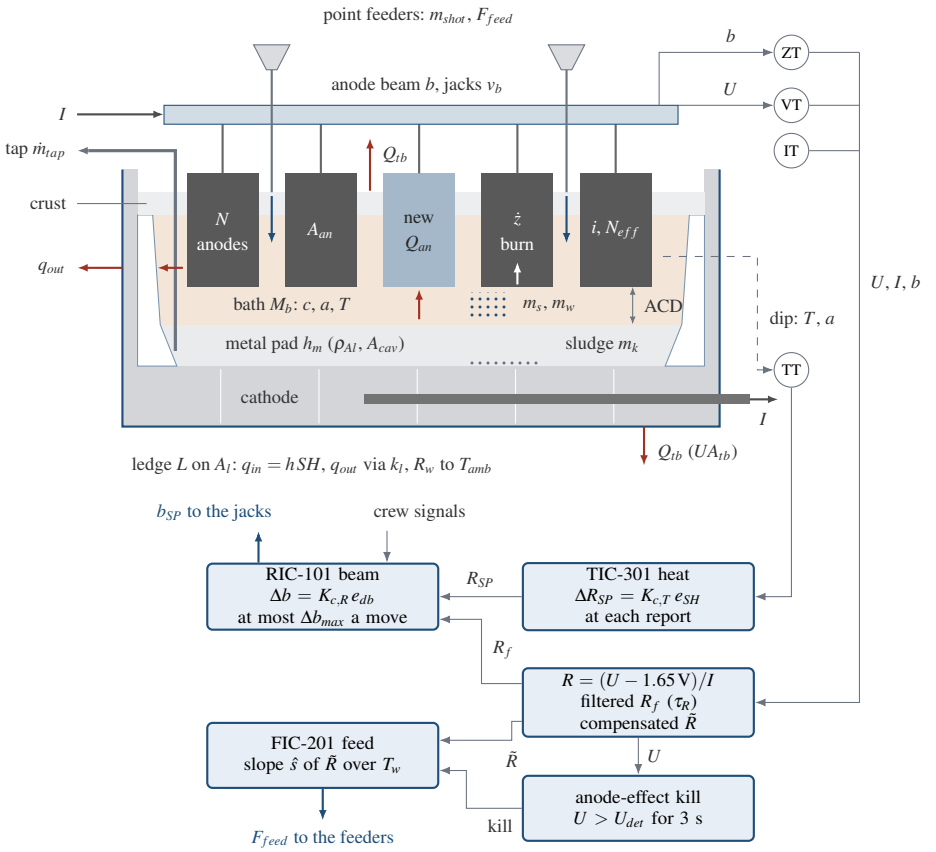


Figure 11.1 One prebake cell in longitudinal section, not to scale, with five of its 24 anodes, and its controller. The beam b, moved by jacks at v_b, carries anodes of area A_{an} that burn upward at $\dot{z}$ over a bath M_b with alumina c, excess AlF_3 a and temperature T, the ACD above a metal pad h_m on A_{cav}, and a ledge L on the sidewall area A_l. Point feeders drop shots m_{shot} at the rate F_{feed} into dispersing pools, m_s for dry alumina and m_w for a wet lot, or into sludge m_k. A new anode draws Q_{an}, so N_{eff} anodes carry I at the density i, and a tap siphons $\dot{m}_{tap}$. The controller reads U (VT), I (IT), b (ZT), the crew's signals and a daily dip of T and a (TT), and moves b_{SP}, F_{feed} and its own target R_{SP}, with e_{db} the error outside the dead band and $e_{SH} = SH_{SP} - SH_{est}$.

11.3.1 Alumina Between a Shot and the Anode

Each shot of a working feeder delivers $m_{shot} = 1.75\,\mathrm{kg}$, while the controller counts 1.80 kg, a declared error of 2.8 %. A share φ of the shot sinks to the metal as sludge and the rest joins a dispersing pool, m_w while a wet lot is on the feeders and m_s otherwise. The pools release alumina to solution as shown in Equation (11.1).

$$\begin{aligned}
&\frac{\mathrm{d}m_s}{\mathrm{d}t} = -\frac{m_s}{\tau_{dis}}, \qquad \frac{\mathrm{d}m_w}{\mathrm{d}t} = -\frac{m_w}{\tau_{wet}}, \qquad \frac{\mathrm{d}m_k}{\mathrm{d}t} = -\frac{m_k}{\tau_k} \\
&\frac{\mathrm{d}m_a}{\mathrm{d}t} = \frac{m_s}{\tau_{dis}} + \frac{m_w}{\tau_{wet}} + \frac{m_k}{\tau_k} - r_{cons} \\
&r_{cons} = \frac{\eta_{CE}\, I\, M_{Al_2O_3}}{6F}, \qquad c = \frac{100\, m_a}{m_a + m_b + m_f}
\end{aligned} \tag{11.1}$$

Table 11.1 Variables of the Smelting Cell, Their Roles and Defaults

Symbol	Role	Default	Range
R_f	CV of RIC-101, filtered pseudo-resistance	8.63 µΩ SP	7.4–10.0
SH	CV of TIC-301, superheat (inferred daily)	10 K SP	3–20
c	truth, alumina, never measured	3.5 % start	band 2–4
b	MV, beam position from its encoder	200 mm	0–400
F_{feed}	MV, alumina feed rate	180.70 kg/h	0–450
R_{SP}	MV of TIC-301, target resistance	8.63 µΩ	7.4–10.0
F_{AlF_3}	MV, AlF_3 addition	1.6 kg/h	0–6
I	DV, line current, measured	300 kA	250–320
w	DV, slower dissolution of a lot	1, wet $\frac{10}{3}$	1–5
μ	DV, pad noise multiplier	1	0–3
—	DVs, events	anode change, tap, feeder stuck	
—	DVs, events	pad instability, current drop	
m_a, m_s, m_w, m_k	states, alumina pools (kg)	315, 8.8, 0, 21.7	—
m_b, m_f	states, bath solvent, excess AlF_3 (kg)	7560, 1125	—
T, L	states, temperature, ledge	959.8 °C, 12.6 cm	—
h_m, b, z	states, pad, beam, burn-off	20, 20, 0 cm	—
U, I	measured every 1 s	4.24 V, 300 kA	—
T_{dip}, a_{dip}	measured daily, 1 h late, ±1.5 K, ±0.2 %	06:00	—

where η_{CE} is the current efficiency, $M_{Al_2O_3}$ the molar mass of alumina, F Faraday's constant and c the alumina in wt%. The excess AlF_3 is $a = 100\,m_f/(m_a+m_b+m_f)$, its mass grows with the addition F_{AlF_3} and falls with a loss r_{AlF_3} to the lining and the fumes, and the crew holds the bath level, so $\mathrm{d}m_b/\mathrm{d}t = 0$. At 300 kA and 95 % the cell consumes 180.70 kg/h of alumina and makes 95.64 kg/h of aluminum, 2.295 t a day, and the base feed is one shot every 36 s. In the 9000 kg bath one percent of alumina is 90 kg, half an hour of consumption. The three times and the sinking share depend on the superheat SH, the lot's factor w and the alumina, as shown in Equation (11.2).

$$\begin{aligned} g &= \frac{13}{\max(0, SH)+3}, \qquad \chi = \mathrm{clip}_{[0.05,2]}\left(\frac{c_{sat}-c}{c_{sat}-3.5}\right) \\ \tau_{dis} &= \frac{\tau_{dis,0}\,g}{\chi}, \qquad \tau_{wet} = w\,\tau_{dis}, \qquad \tau_k = \frac{\tau_{k,0}\,g}{\chi} \\ \varphi &= \mathrm{clip}_{[0,0.9]}\left(\varphi_0 + \varphi_{wet}[w>1] + 0.25\,\mathrm{clip}_{[0,1]}\left(\tfrac{5-SH}{5}\right)\right) \end{aligned} \tag{11.2}$$

where $\mathrm{clip}_{[p,q]}(x) = \min(q, \max(p,x))$, g is 1 at 10 K of superheat and 4.3 at the liquidus, χ slows dissolution toward the solubility c_{sat}, and $[w>1]$ is 1 under a wet lot. The lot of the cases has $w = 10/3$, so its shots disperse in 10 min instead of 3 min and a further 30 % of each sinks. Below 5 K of superheat a growing share of every shot sinks too, which is how a cold bath starves its own solution.

11.3.2 Voltage and the Control Room's Number

The cell voltage is built from a reversible part, two anode overvoltages, a cathode overvoltage and two ohmic drops, as shown in Equation (11.3).

$$\begin{aligned} U &= E_{rev} + \eta_{rxn} + \eta_{conc} + \eta_c + IR_{ext} \\ &\quad + \frac{I\,(ACD + d_{bub})}{\kappa N_{eff} A_{an} f_{fr}} \\ \eta_{rxn} &= \frac{R_g T_K}{\alpha F} \ln \frac{i}{i_0\,(c/3.5)^m} \\ \eta_{conc} &= \frac{R_g T_K}{2F} \ln \frac{c}{\max(c - c_{crit}, 0.01)} \\ i &= \frac{I}{N_{eff} A_{an}}, \qquad c_{crit} = c_{crit,0}\,\frac{i}{i_{nom}} \\ R &= \frac{U - 1.65\,\mathrm{V}}{I} \end{aligned} \tag{11.3}$$

where R_g is the gas constant, T_K the absolute temperature, α the anode's transfer coefficient, i_0 its exchange current density at 3.5 % alumina and m the exponent on alumina, c_{crit} the critical alumina of the moment, which scales with the current density against its design value i_{nom}, d_{bub} the bubble layer counted as the bath it replaces, f_{fr} the fringing of current around the anodes and R_{ext} the anodes, cathode and busbars. The concentration overvoltage has a pole at c_{crit}, the mass-transfer limit where the anode effect starts, and is held at its value 0.01 % above it. The conductivity follows the correlation that the application takes from the moving-electrode measurements of [17], as shown in Equation (11.4), with the alumina, AlF_3 and CaF_2 in wt% and κ in S/cm.

$$\ln \kappa = 1.977 - 0.0200\,c - 0.0131\,a - 0.0060\,w_{CaF_2} - \frac{1204.3}{T_K} \tag{11.4}$$

At the design point, 3.5 % alumina, 959.8 °C, an ACD of 4.0 cm and 0.847 A/cm^2, the 4.237 V across the cell divide into 1.20 V reversible, 0.499 and 0.038 V at the anode, 0.10 V at the cathode, 1.550 V in the bath and 0.850 V in the conductors, with a conductivity of 2.089 S/cm. The control room's pseudo-resistance R subtracts 1.65 V, an intercept of the voltage–current line that lies below the true back electromotive force, so R is not the ohmic resistance of the cell [4]. It reads 8.625 μΩ at the design point, where the cell voltage is a little above the 4.1 to 4.2 V at which cells above 200 kA have settled [12], and one millimeter of ACD moves it 0.117 μΩ.

11.3.3 Bath, Ledge and Superheat

The liquidus is the empirical equation of [18] without its LiF, MgF_2 and KF terms, as shown in Equation (11.5), in °C with a, w_{CaF_2} and c in wt%.

$$\begin{aligned} T_{liq} &= 1011 + 0.50a - 0.13a^{2.2} - \frac{3.45\,w_{CaF_2}}{1 + 0.0173\,w_{CaF_2}} \\ &\quad + 0.124\,w_{CaF_2} a - 0.00542\,(w_{CaF_2} a)^{1.5} \\ &\quad - \frac{7.93\,c}{1 + 0.0936\,c - 0.0017\,c^2 - 0.0023\,a\,c} \end{aligned} \tag{11.5}$$

It is 949.8 °C at the design point and falls 6.0 K for each percent of alumina near 2.6 %, so the superheat $SH = T - T_{liq}$ moves with the feeding cycle. The bath and the metal share one heat capacity

C, and the ledge freezes or melts against the superheat, as shown in Equation (11.6).

$$\begin{aligned} C\frac{\mathrm{d}T}{\mathrm{d}t} &= I\,(U_s - E_{tn}) + \delta_{AE} f_{AE} I\,(U_{AE} - U_s) \\ &\quad - q_A\, r_{dis} - Q_{side} - UA_{tb}(T - T_{amb}) - Q_{an} \\ \rho_l \Delta H_f \frac{\mathrm{d}L}{\mathrm{d}t} &= q_{out} - q_{in}, \qquad q_{in} = h\,SH \\ q_{out} &= \frac{T_{liq} - T_{amb}}{L/k_l + R_w} \end{aligned} \tag{11.6}$$

where U_s is the voltage of Equation (11.3) at the smooth ACD, without the pad's noise, E_{tn} the reaction enthalpy as a voltage with the alumina already hot, δ_{AE} is 1 during an anode effect, f_{AE} the share of its extra power that heats the bath, q_A the heat that warms and dissolves alumina at the rate r_{dis} at which it dissolves, $Q_{side} = A_l\, q_{in}$ the heat into the ledge, ρ_l and ΔH_f the frozen bath's density and latent heat, k_l the ledge's conductivity and R_w the resistance of the sidewall and shell. When the ledge has melted to L_{min} and $q_{in} > q_{out}$, the bath heats a bare wall, $Q_{side} = A_l(T - T_{amb})/(1/h + L_{min}/k_l + R_w)$, and L stops changing. At the design point the 1271 kW the cell draws are 570 kW of reaction, 60 kW for the feed, 112 kW into a ledge 12.6 cm thick and 529 kW through the top and bottom.

11.3.4 Beam, Anodes, Metal and the Crew

The anodes hang from the beam, they burn from below, and the metal rises under them, so the ACD follows Equation (11.7).

$$\begin{aligned} ACD &= ACD_0 + (b - b_0) + z - (h_m - h_{m,0}) \\ \frac{\mathrm{d}z}{\mathrm{d}t} &= \frac{b_C\, \dot{m}_{Al}}{\rho_C N A_{an}}, \qquad \frac{\mathrm{d}h_m}{\mathrm{d}t} = \frac{\dot{m}_{Al} - \dot{m}_{tap}}{\rho_{Al} A_{cav}}, \qquad \frac{\mathrm{d}b}{\mathrm{d}t} = v_b \end{aligned} \tag{11.7}$$

where $\dot{m}_{Al} = \eta_{CE} I M_{Al}/(3F)$ is the aluminum made, b_C the net carbon burned per kilogram of aluminum and ρ_C the anode density. The anodes burn 0.723 mm/h while the metal rises 0.904 mm/h, so the ACD closes 0.181 mm/h unless the beam moves. A tap siphons $m_{tap} = 2300$ kg at 12 000 kg/h, for 11.5 min, and opens the ACD by 21.7 mm. An anode change takes n anodes out for $t_{out} = 10$ min and sets cold ones, so the cell has $N_{eff} = N - n$ anodes carrying current while they are out and $N_{eff} = N - n\,\mathrm{e}^{-\theta/\tau_{pick}}$ at an age θ after their return, when the new anodes draw $Q_{an} = nQ_{ac}\,\mathrm{e}^{-\theta/\tau_{an}}$ from the bath. A cell of 24 anodes that last about 30 days changes one every 30 h or so, close to the one a day that cells of 16 to 40 anodes see [1], so the cases fire changes as events rather than on a schedule.

11.3.5 The Anode Effect and Its Emissions

The anode effect is a hybrid state, as shown in Equation (11.8).

$$\begin{aligned} &\text{starts when } c \le c_{crit} \\ &\text{ends when } c \ge c_{crit} + \Delta c_{self} \\ &\quad \text{or when } ACD \le ACD_{kill} \text{ and } c \ge c_{crit} + \Delta c_{kill} \\ &U = U_{AE} + \sigma_{fl}\, \nu_k, \qquad \dot{m}_{CF_4} = 60\, S_{CF_4} P_{day} \\ &\dot{m}_{C_2F_6} = F_{C_2F_6} \dot{m}_{CF_4} \end{aligned} \tag{11.8}$$

where a new standard normal draw ν_k arrives with every step, $\sigma_{fl} = 2.5$ V a flicker, P_{day} the aluminum made in a day in metric tons, and the emission rates are in kilograms per hour. The second ending is the kill. With the alumina 0.2 % above critical, pumping the anodes down to an ACD of 15 mm sweeps the gas film away, which stands for the short circuit of the anode–cathode gap that a termination requires [16]. The emission coefficients follow the slope method, kilograms of CF_4 per metric ton of

aluminum for each anode-effect minute per cell-day, so every anode-effect minute of this cell vents 0.328 kg of CF_4, which with the GWP-100 values of 6630 and 11 100 [2] is 2.62 t of CO_2-equivalent. During an effect the cell runs at 9.0 MW instead of 1.27 MW.

11.3.6 The Metal Pad's Noise

The ACD the anodes see carries a noise w_{pad} that stands for the metal pad's roll and heave, as shown in Equation (11.9), updated once per step Δt in seconds.

$$\begin{aligned} \dot{x}_{k+1} &= \dot{x}_k - (2\zeta\omega\dot{x}_k + \omega^2 x_k)\,\Delta t + \sigma_x\sqrt{\Delta t}\;\nu_{1,k} \\ x_{k+1} &= x_k + \dot{x}_{k+1}\Delta t \\ y_{k+1} &= \beta y_k + \sqrt{1-\beta^2}\,\sigma_{slow}\nu_{2,k} \\ w_{pad} &= \lambda\,\mu\,(x+y) \end{aligned} \tag{11.9}$$

where x is a lightly damped oscillator of frequency $\omega = 2\pi/T_{pad}$ and damping ratio ζ whose stationary standard deviation is σ_{pad}, with $\sigma_x = \sigma_{pad}\sqrt{4\zeta\omega^3}$, y a first-order heave with $\beta = \mathrm{e}^{-\Delta t/\tau_{slow}}$, μ the multiplier the operator sets and λ is 4 during a pad instability and 1 otherwise. The noise enters the voltage the controller measures and not the heat balance, and it moves the pseudo-resistance by 0.033 μΩ, one standard deviation from second to second.

11.3.7 Start and Numerics

Every case starts at 00:00 in thermal balance. The bath, AlF_3 and alumina masses follow from $M_{b,0}$, a_0 and c_0, the superheat solves $I(U - E_{tn}) - q_A r_{cons} - A_l h SH - UA_{tb}(T - T_{amb}) = 0$ by bisection, the ledge is the thickness at which $q_{out} = q_{in}$, and the pools hold what the base feed would leave in transit, $r_{cons}(1-\varphi)\tau_{dis}$ dispersing and $r_{cons}\varphi\tau_k$ as sludge. At the design point the superheat is 10.01 K, the ledge 12.6 cm and the alumina in transit 30.4 kg. Table 11.2 lists the parameters. The integrator is fourth-order Runge–Kutta with a fixed step of 1 s, the controller scans every second, the historian logs every 10 s, and the speed selector runs from 0.01 to 4 simulated hours per second, so a two-day case takes 192 s at the default 0.25 h per second. The order of work in a step, the Reset that keeps the learner's settings and the reload that restores the case's own are those of every framework app (Section 1.6.1). The pad's noise and the dip's errors draw on one generator seeded at 20 261 011. The ledgers close every balance to a relative 1×10^{-9} through whole runs, and halving the step with the pad's noise off moves the lowest alumina of a day of demand feeding by under 0.01 %, the bath temperature by under 0.05 K and the count of triggers by none, two of the 41 checks the application passes.

11.3.8 Assumptions

One well-mixed bath. Real anode effects are triggered by low alumina near an anode surface and start on a single anode before they spread [16]. Here one concentration stands for the whole bath, so no local anode effect can start while the mean is above the wall.

A teaching curve. The Tafel, mass-transfer and ohmic constants are placed so that the minimum sits at 3.5 % and the anode effect at 1.80 % at the design current density. Measured curves had their minimum at 5.0 to 5.5 % and their anode effect at 1.6 to 2.2 %, 100 to 300 mV above the minimum [5], and this one rises 137 mV from its minimum to 1.85 %, inside that range.

A threshold for the anode effect. The effect is a threshold on the mean alumina, and its 30 V, the share of its power that heats the bath and the kill condition are teaching values, as are the slope coefficients of its emissions. Real alumina feeding and anode-effect prediction belong to each plant's computer control [3, 9], and real emissions are measured on the potlines that make them [15].

A bath level held by the crew. Crushed bath goes in as the ledge grows and is ladled out as it melts, so the ledge moves heat but not the bath's composition. A growing ledge would otherwise concentrate the AlF_3 and the alumina and lower the liquidus.

Table 11.2 Model Parameters, With M for a Show Model Slider and C for a Constant Fixed in the Code

Symbol	Meaning	Value	Class
I_{nom}, N	design current, anodes	300 kA, 24	M
A_{an}	area of one anode, 1.80 m × 0.82 m	1.476 m^2	M
η_{CE}	current efficiency	95 %	M
$M_{b,0}$	liquid bath at the start	9000 kg	M
c_0	alumina at the start (case 1: 5.0)	3.5 %	M
a_0, w_{CaF_2}	excess AlF_3 at the start, CaF_2	12.5, 5 %	M
ACD_0, b_0	ACD and beam at the start	4.0 cm, 200 mm	M, C
E_{rev}	reversible voltage	1.20 V	M
α, m	transfer coefficient, exponent on alumina	1.0, 0.492	M
i_0	exchange current density at 3.5 %	0.0077 A/cm^2	M
$c_{crit,0}$	critical alumina at i_{nom}	1.80 %	M
η_c	cathode overvoltage	0.10 V	M
f_{fr}, d_{bub}	fringing, bubble layer	1.159, 0.43 cm	M
R_{ext}	anodes, cathode and busbars	2.833 μΩ	M
U_{AE}, σ_{fl}	voltage of an anode effect, flicker	30 V, 2.5 V	M, C
—	intercept of R	1.65 V	C
A_{cav}, $h_{m,0}$	metal pad area, depth at the start	46 m^2, 20 cm	M
ρ_{Al}	liquid aluminum density	2300 kg/m^3	M
b_C, ρ_C	net carbon, anode density	0.415, 1550 kg/m^3	M
t_{out}, τ_{pick}	anode out of the cell, current pickup	10 min, 8 h	M
Q_{ac}, τ_{an}	heat a new anode draws, heating time	60 kW, 6 h	M
m_{tap}, $\dot{m}_{tap}$	metal of one tap, siphon rate	2300 kg, 12 000 kg/h	M
t_{dip}	hour of the daily dip	06:00	M
C	heat capacity of bath and metal	55 MJ/K	M
E_{tn}	reaction enthalpy as a voltage	1.90 V	M
q_A	heat to warm and dissolve alumina	1.2 MJ/kg	M
A_l, h	ledge area, bath-ledge coefficient	16 m^2, 700 W/(m^2 K)	M
k_l, R_w	ledge conductivity, wall resistance	1.8 W/(m K), 0.06 m^2 K/W	M
T_{amb}	potroom air at the shell	40 °C	M
ρ_l, ΔH_f	frozen bath density, latent heat	2100 kg/m^3, 530 kJ/kg	M
UA_{tb}	top and bottom loss coefficient	575 W/K	M
L_{min}	thinnest ledge, bare wall below	1.0 cm	M
f_{AE}	anode-effect power to the bath	0.10	M
r_{AlF_3}	AlF_3 lost to lining and fumes	1.6 kg/h	M
m_{shot}	true dose (counted 1.80)	1.75 kg	M
$\tau_{dis,0}$, $\tau_{k,0}$	dispersing, sludge times at 10 K	3 min, 4 h	M
φ_0, φ_{wet}	sinking share, extra from a wet lot	0.03, 0.30	M
c_{sat}	alumina solubility	7.5 %	M
Δc_{kill}, Δc_{self}	above critical to kill, to end alone	0.2, 1.0 %	M
ACD_{kill}	ACD at which a pump sweeps the film	15 mm	M
S_{CF_4}	slope coefficient (per AE-min/day)	0.143	M
$F_{C_2F_6}$	C_2F_6 per kilogram of CF_4	0.121	M
—	GWP-100 of CF_4, C_2F_6	6630, 11 100	M
σ_{pad}, T_{pad}, ζ	pad roll: deviation, period, damping	0.25 mm, 45 s, 0.15	M
σ_{slow}, τ_{slow}	pad heave: deviation, correlation time	0.10 mm, 120 s	M
λ	roll factor in a pad instability	4	C

The constants of Equation (11.2) (13, 3, 0.25, 5 K and the clip limits) and the floors of Equation (11.3) (0.05 % alumina and 0.01 % above critical) are class C. Constants: $F = 96485.33\,\mathrm{C/mol}$, $R_g = 8.314462\,\mathrm{J/(mol\,K)}$, $M_{Al_2O_3} = 0.101961\,\mathrm{kg/mol}$ and $M_{Al} = 0.0269815\,\mathrm{kg/mol}$.

Constants where the cell varies. The current efficiency is a constant 95 %, the pad's noise is a declared spectrum that cannot grow by itself, dissolution is first order in three pools, and the bath, metal and walls are lumped, with one ledge on 16 m^2 and one loss coefficient for the top and bottom.

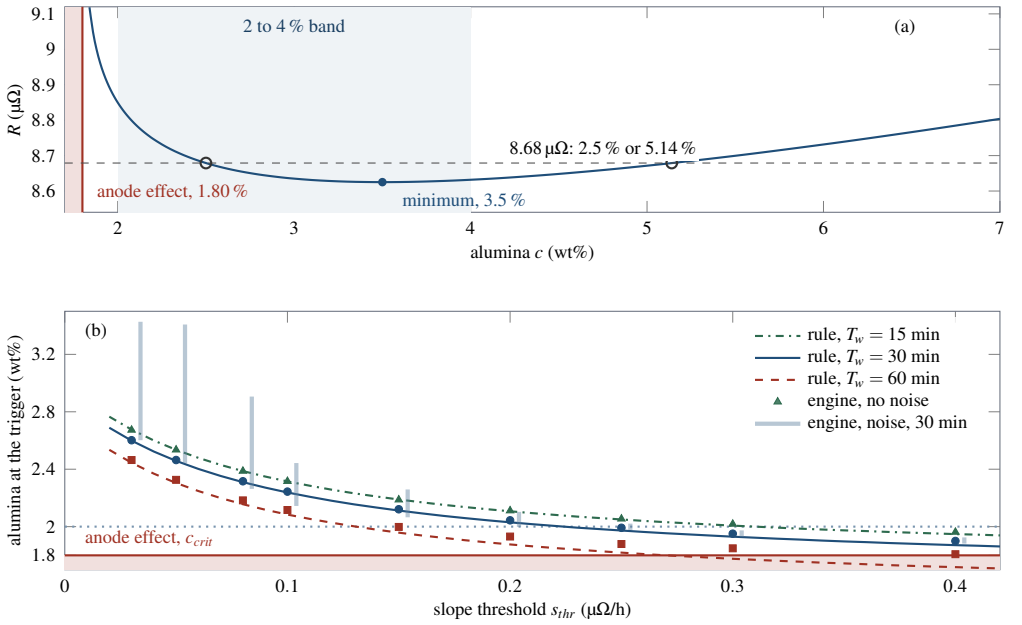


Figure 11.2 The curve the controller reads and where its underfeeds end. (a) The pseudo-resistance against alumina at the design point, 959.8 °C, an ACD of 4.0 cm and 300 kA. (b) The alumina at which an underfeed ends against the slope threshold, from the half-window rule of Equation (11.11) (dash-dotted, solid and dashed lines for windows of 15, 30 and 60 min) and the engine's first trigger in case 3 with the pad noise off (triangles, circles and squares for the same windows), with the spread of every trigger in 24 h at the default noise and the 30 min window (shaded bars).

11.4 A Reading With Two Answers

11.4.1 Why the Curve Has a Minimum

The minimum is where three slopes cancel. Differentiating Equation (11.3) at fixed temperature, ACD and current gives Equation (11.10), per wt% of alumina.

$$I\frac{\partial R}{\partial c} = -\frac{R_g T_K}{\alpha F}\frac{m}{c} + \frac{R_g T_K}{2F}\left(\frac{1}{c} - \frac{1}{c - c_{crit}}\right) + 0.0200\,U_{bath} \tag{11.10}$$

where U_{bath} is the ohmic drop in the bath. At 3.5 % the reaction term is −14.9 mV per percent, the concentration term −16.1 mV and the ohmic term +31.0 mV, 2 % of the 1.550 V drop, and they cancel. At 2.5 % the anode terms grow to −20.9 and −54.6 mV while the ohmic term stays at 30.4 mV, so the curve falls 0.151 μΩ per percent there, 0.341 μΩ at 2.2 % and 0.784 μΩ at 2.0 %. Figure 11.2(a) shows the result. A reading of 8.679 μΩ means 2.5 % or 5.14 %, 2.2 % reads as 6.24 % and 2.0 % as 7.58 %, while above the minimum the curve rises only 0.172 μΩ from 4 to 7 %. A reading says how far the cell is from the minimum and nothing about the side.

🕮 **Pause.** Try Exercise 11.1 before reading on.

11.4.2 The Slope an Underfeed Produces

An underfeed reads the side through a time derivative, as the rate endpoint of Chapter 4 reads a carbon that a level cannot give (Section 1.4.4). During a 70 % underfeed the alumina falls 0.615 %/h in case 3 (0.609 to 0.617 over its 21 underfeeds), faster than the 0.602 %/h of the counted dose because every shot is 2.8 % short, and slower than the 0.641 %/h of the true dose because the dispersing pools drain as the feed drops. The alumina-driven slope $s_c = -\dot{c}\,\partial R/\partial c$ is then 0.027 μΩ/h at 3.0 %, 0.093 at

2.5 %, 0.210 at 2.2 % and 0.482 at 2.0 %. Two things stand between that slope and the controller. The closing ACD lowers every slope by $d = 0.181\,\text{mm/h} \times 0.117\,\mu\Omega/\text{mm} = 0.021\,\mu\Omega/\text{h}$, which the beam corrects only in steps of its own. The pad's noise leaves a standard deviation in the least-squares slope of 0.071 μΩ/h over a 15 min window, 0.027 over 30 min and 0.010 over 60 min, measured at the flat bottom with the beam held, and 0.109 μΩ/h over 30 min in a pad instability. The noise falls about as the window to the power -1.4, near the $-3/2$ of white noise, and the price of a longer window is lag.

11.4.3 Where an Underfeed Ends

A least-squares line through the last T_w of a steady descent estimates the slope near the middle of its window, half a window behind the present alumina, so an underfeed ends near the alumina of Equation (11.11).

$$c_{trig} \approx c^*(s_{thr} + d) - |\dot{c}|\,\frac{T_w}{2}, \qquad s_c\big(c^*(s)\big) = s \tag{11.11}$$

where $c^*(s)$ is the alumina at which the alumina-driven slope equals s and s_{thr} is the threshold. With the noise switched off, the engine's first trigger lies within 0.013 % of the rule at 15 min, 0.028 % at 30 min and 0.090 % at 60 min over thresholds of 0.03 to 0.40 μΩ/h, and above it at the longer windows, where the curve's convexity across the window makes the fit lead its midpoint, as shown in Figure 11.2(b). At the default 0.10 μΩ/h and 30 min the rule gives 2.238 % and the engine 2.244 %. The rule sets the two edges of a usable threshold. The lower edge is the noise. A threshold within two or three standard deviations of the flat-bottom slope ends underfeeds on the roll, and with the default noise the triggers of a day spread over 2.60 to 3.43 % at 0.03 μΩ/h and 2.26 to 2.91 % at 0.08, against 2.14 to 2.44 % at 0.10. The upper edge is the wall, because c_{trig} must stay above c_{crit} with room for the overfeed's alumina to dissolve. At 60 min the rule reaches 1.80 % at 0.27 μΩ/h, and the engine meets the wall at 0.40 with the noise off.

11.4.4 One Handle for Resistance and Heat

The beam moves the resistance and the heat together. One millimeter of ACD is 0.117 μΩ, every 0.1 μΩ of target is 9.0 kW of heat at 300 kA, and a kelvin of superheat drives 11.2 kW into the ledge against 0.575 kW through the top and bottom. Solving the heat balance at 3.5 % for a given target shows the consequence. The steady superheat rises 0.76 K for each 0.1 μΩ, and a target of 7.90 μΩ instead of 8.625 settles at 4.47 K with an ACD of 3.37 cm. The bath alone answers in $C/(A_l h + UA_{tb}) = 1.30\,\text{h}$, while the ledge integrates an imbalance for days, the clocks of thermal inertia in Section 1.4.2. A beam loop that holds the resistance therefore holds the heat as well, unless its target is wrong, and the only measurement of the heat is a dip a day.

11.5 The Cell Controller, Loop by Loop

The controller reads five inputs and nothing else, the cell voltage every second with the pad's noise in it, the line current, the beam's encoder, the crew's signals of an anode change or a tap with their start times, and a daily bath dip. It never reads the alumina, the superheat, the ACD or the ledge. It forms the pseudo-resistance R every scan and filters it with a first-order lag of $\tau_R = 2\,\text{min}$, held while an anode effect is detected. Table 11.3 lists the settings.

11.5.1 Resistance Regulation on the Shared PID

Every $T_{reg} = 5\,\text{min}$ the beam loop forms the error $e = R_{SP} + R_{ac} - R_f$, removes a dead band, $e_{db} = \text{sgn}(e)\max(0, |e| - \delta_R)$ with $\delta_R = 0.08\,\mu\Omega$, and passes e_{db} as the setpoint of the framework's shared PID with a process value of zero and $\Delta t = T_{reg}$. That PID is the velocity form of Section 1.6.2, Equations (1.13) to (1.16), with the action $a = +1$, reverse acting because a higher beam raises the resistance. Its clamp keeps the beam within its 0 to 400 mm of travel, its slew holds a move to $\Delta b_{max} = 3\,\text{mm}$ per interval, it has no derivative action ($\tau_D = 0$), and its integral is back-calculated with a tracking time equal to τ_I. Because τ_I equals the interval, the back-calculation returns the

Table 11.3 Controller Settings and Constants, Defaults and Ranges

Symbol	Setting	Default	Range
RIC-101, resistance to beam			
R_{SP}	target resistance	8.63 µΩ	7.4–10.0
$K_{c,R}$	gain, mm/µΩ	6	0–20
$\tau_{I,R}$, T_{reg}	integral time, interval	5, 5 min	1–60, 1–30
δ_R	dead band	0.08 µΩ	0–0.4
Δb_{max}	move limit	3 mm	0.5–30
τ_R	resistance filter	2 min	0.5–15
—	anode-change, tap compensation	off, off	on
s_R	assumed beam sensitivity	0.12 µΩ/mm	C
ΔR_{an}, $\hat{\tau}_{pick}$	assumed rise per anode, pickup	0.22 µΩ, 8 h	C
FIC-201, alumina feed			
—	feeding	demand	manual, fixed
s_{thr}	slope threshold	0.10 µΩ/h	0.01–0.6
u_{under}	underfeed	70 %	30–100
u_{over}	overfeed	150 %	100–300
t_{over}, t_{base}	overfeed, base period	45, 30 min	5–180, 0–180
t_{arm}	arming time	20 min	2–90
$t_{under,max}$	longest underfeed	4 h	0.5–12
—	slope filter	least squares	two lags
T_w	slope window	30 min	5–120
t_{hold}	hold after an event	1 h	5–240 min
u_{fixed}	fixed feed	100 %	50–150
—	counted dose, $\hat{\eta}_{CE}$	1.80 kg, 0.95	C
—	feeders, shortest gap	4, 20 s	C
Anode-effect detection and kill			
U_{det}	detection, held 3 s	8 V	5–20
m_{kill}	kill burst	36 kg	0–90
Δb_{pump}	pump stroke	30 mm	5–60
—	automatic termination	on	off
t_{crew}	crew response, termination off	8 min	1–60
TIC-301, superheat to target resistance			
—	heat loop mode	manual	auto
SH_{SP}	superheat target	10 K	3–20
$K_{c,T}$	gain, µΩ/K	0.06	0–0.3
$\tau_{I,T}$	integral time	24 h	6–240
c_{as}	alumina the inference assumes	2.6 %	1.5–6
—	step limit, target range	0.3 µΩ, 7.4–10.0	C

integral to the applied output at every call, and each move is exactly $\Delta b = K_{c,R}\, e_{db}$ within $\pm\Delta b_{max}$. With $K_{c,R} = 6\,\mathrm{mm/\mu\Omega}$ a move removes $6 \times 0.117 = 0.70$ of the error outside the band. The jacks run the beam to its setpoint at 10 mm/min, and a resistance above its target lowers the beam, which the application's direction check confirms. Regulation pauses while an anode effect is detected and for 10 min after it.

Two options feed a known cause forward. Tap compensation lowers the beam setpoint with the metal the siphon takes, $\dot{m}_{tap}\Delta t/(\rho_{Al}A_{cav})$ per scan at the controller's assumed density and area, and carries the move in the PID's integral so that the loop owns it. Anode-change compensation adds to the target the resistance the controller expects the missing current to add, $\sum_j \Delta R_{an}\, n_j\, \mathrm{e}^{-\max(0,\,\theta_j - t_{out})/\hat{\tau}_{pick}}$, with $\Delta R_{an} = 0.22\,\mu\Omega$ against the plant's step of 0.25 and $\hat{\tau}_{pick} = 8\,\mathrm{h}$, passed through the same first-order lag τ_R as the reading, so that R_{ac} and R_f step together at a change.

11.5.2 Feeding Modes and the Tracking Cycle

The feed loop's base rate F_b is the consumption r_{cons} of Equation (11.1) at the measured current and an assumed current efficiency $\hat{\eta}_{CE}$ of 95 %, 180.70 kg/h at 300 kA, and the loop turns any rate into shots of 1.80 kg through an accumulator that fires the next of four feeders, each at most every 20 s. Fixed feeding asks for $u_{fixed}F_b$. Demand feeding runs the tracking cycle and decides every 30 s. It underfeeds at $u_{under} = 70\,\%$ of F_b until, after an arming time of 20 min, the slope of

Equation (11.12) passes $s_{thr} = 0.10\,\mu\Omega/\mathrm{h}$, or until the underfeed has lasted 4 h. It then overfeeds at 150 % for $t_{over} = 45\,\mathrm{min}$, feeds F_b for 30 min and underfeeds again.

$$\tilde{R}_j = \bar{R}_j - s_R\,(b_j - b_{ref}), \qquad \hat{s} = \frac{\sum_j (t_j - \bar{t})(\tilde{R}_j - \bar{\tilde{R}})}{\sum_j (t_j - \bar{t})^2} \tag{11.12}$$

where $\bar{R}_j$ are 30-s block means of the unfiltered resistance at block times t_j, $s_R = 0.12\,\mu\Omega/\mathrm{mm}$ is the beam sensitivity the controller assumes against the plant's 0.117, b_{ref} the beam at the start of the underfeed, and the sums run over the blocks of the last $T_w = 30\,\mathrm{min}$. A second estimator, two first-order lags of $T_w/4$ on the resistance and on its block-to-block change, is an option. An anode change, a tap or an anode effect puts a hold on the cycle, because its step in resistance is not alumina. The estimator forgets its blocks, an underfeed in progress drops to the base rate, and a fresh underfeed starts once the hold, $t_{hold} = 1\,\mathrm{h}$ plus a tap's duration, has passed.

11.5.3 Anode-Effect Detection and Kill

The detector declares an anode effect when the voltage stays above $U_{det} = 8\,\mathrm{V}$ for 3 s and releases it after 10 s below. With automatic termination on the kill starts at once, and with it off the crew starts it $t_{crew} = 8\,\mathrm{min}$ later. The kill fires a burst of 36 kg, raises every feeding mode to 200 % of the base rate until the effect is over, and pumps the beam 30 mm down at 30 mm/min, holds it 10 s and returns it, first 30 s after the kill starts and then a minute after each return while the effect lasts. Demand feeding then runs an overfeed, a base period and a fresh underfeed.

11.5.4 The Daily Dip and the Target Resistance

The heat loop runs on the framework's shared sampled analyzer. A dip drawn at 06:00 every 24 h reports the bath temperature with 1.5 K of noise and the excess AlF_3 with 0.2 % an hour later, at 07:00, and the loop infers the superheat as $SH_{est} = T_{dip} - T_{liq}(c_{as}, a_{dip})$ from Equation (11.5) at an assumed alumina $c_{as} = 2.6\,\%$ and 5 % CaF_2. In AUTO each report moves the target through the shared PID with $K_c = 0.06\,\mu\Omega/\mathrm{K}$ and $\tau_I = 24\,\mathrm{h}$, equal to the report interval, so each step is $K_{c,T}(SH_{SP} - SH_{est})$, limited to 0.3 µΩ and to 7.4 to 10.0 µΩ, and a superheat below its target raises the target, so the loop is reverse acting ($a = +1$). In MANUAL the target is the operator's, 8.63 µΩ by default. The inference inherits the liquidus's 6.0 K per percent of alumina, so an error of 0.5 % in c_{as} is 3 K of superheat that no dip reveals. The inference is a soft sensor of the kind Section 1.4.6 describes, exact for what its model contains and biased by the alumina it assumes.

11.5.5 Manual Mode and Three Kinds of Model

The master MANUAL and AUTO switch throws all three stations, and each station has a mode of its own. The transfer is bumpless both ways, as in every framework app (Section 1.6.2). Equations (11.1) to (11.9) are the simulation model. No model is fitted to the plant's data. The controller's internal models are the base rate at 95 % and 1.80 kg, the beam sensitivity s_R, the anode-change term with its pickup, the metal density and pad area of the tap compensation, and the liquidus at an assumed alumina, and each differs from the plant by a declared amount.

11.6 Measured Runs of the Cell

Table 11.4 lists the cases, and Table 11.5 collects their scorecards, which the application keeps on the truth at every step. The engine produced each number below on its own, at the default seed, in the window that its row or sentence names, with a variant's settings changed after the case is chosen and before it runs, the same run that a reset with those settings gives. Each case starts from the reference line of Table 11.3, demand feeding at 0.10 µΩ/h with the heat loop in MANUAL at 8.63 µΩ, and each case alters that line only where its row says so. Figure 11.3 draws the alumina range of every row against the band and the anode-effect wall.

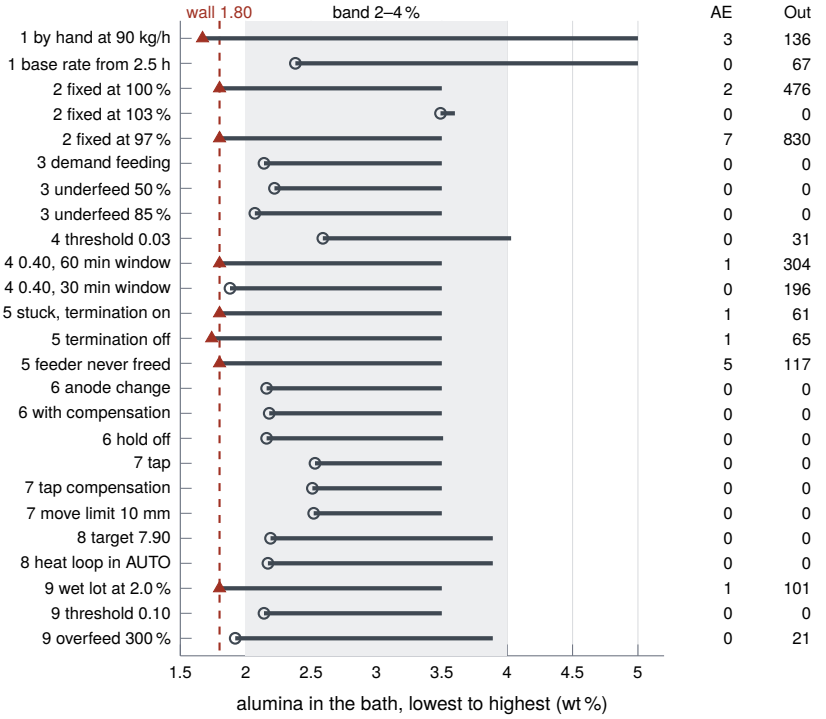


Figure 11.3 The 25 rows of Table 11.5 as a chart: each bar runs from the lowest to the highest alumina of its run, the shaded band is 2 to 4 % and the dashed line is the anode-effect wall at 1.80 %. A filled triangle marks a run with an anode effect, an open circle one without, and the columns at the right give the anode effects (AE) and the minutes outside the band (Out). Notice that every run that reaches the wall is fed on the clock, has a stuck feeder or a wet lot, or ends its underfeeds behind a 60 min window at 0.40 μΩ/h, while the runs of cases 3, 6, 7 and 8 stay inside the band.

Table 11.4 Cases of the Cell and What Each Teaches

#	Change from the reference line	Teaching point
1	$c_0 = 5.0\,\%$, 90 kg/h by hand, beam held, termination off	the curve by hand
2	fixed feeding at 100 % of F_b	open-loop drift
3	none	the tracking cycle
4	$s_{thr} = 0.03\,\mu\Omega/\mathrm{h}$	threshold against noise
5	a feeder stuck from 1 to 5.5 h	anode effect and kill
6	an anode change at 1 h	current redistribution
7	a tap at 1 h	a rate-limited beam
8	$R_{SP} = 7.90\,\mu\Omega$	the heat balance
9	$s_{thr} = 0.25\,\mu\Omega/\mathrm{h}$, a wet lot at 2.0 %	honest infeasibility
10	fixed feeding, $R_{SP} = 8.30\,\mu\Omega$, anode change, tap, wet lot	graded two days

11.6.1 A Reading That Hardly Moves

Case 1 starts the cell at 5.0 % alumina, feeds it 90 kg/h by hand, about half its consumption, holds the beam and leaves anode effects to the crew. The alumina falls 0.98 %/h between 1 and 3 h. The filtered resistance starts at 8.692 μΩ, reaches its lowest, 8.595 μΩ, at 1.37 h and 3.76 %, and stays between 8.595 and 8.641 μΩ while the alumina falls from 4.1 to 2.7 %, a band the size of the pad's noise. It reads 8.703 μΩ at 3 h and 2.15 %, 8.857 μΩ at 3.25 h and 1.90 %, and 9.167 μΩ just before 3.35 h, when the alumina reaches 1.80 % and the anode effect starts. The voltage jumps to about 30 V,

Table 11.5 Truth Scorecard of Each Case and Variant. AE Counts Anode Effects and Out the Minutes Outside 2–4 % Alumina

Case and configuration	Window (h)	AE	Alumina (%)	Out (min)	Superheat (K)	*SE* (kWh/kg)	Other
1 by hand at 90 kg/h (shipped)	5	3	1.67–5.00	136	6.5–29.3	23.43	98.4 t CO_2-eq
1 base rate from 2.5 h	8	0	2.38–5.00	67	7.1–10.9	13.27	
2 fixed at 100 % (shipped)	48	2	1.80–3.50	476	9.0–14.5	13.49	first at 32.53 h
2 fixed at 103 %	48	0	3.49–3.60	0	9.4–10.0	13.23	
2 fixed at 97 %	48	7	1.80–3.50	830	8.8–14.3	14.12	first at 15.82 h
3 demand feeding (shipped)	48	0	2.14–3.50	0	7.9–10.4	13.24	21 triggers
3 underfeed 50 %	48	0	2.22–3.50	0	7.5–10.6	13.24	25 triggers
3 underfeed 85 %	48	0	2.07–3.50	0	8.4–10.5	13.25	14 triggers, 1 timeout
4 threshold 0.03 μΩ/h (shipped)	24	0	2.59–4.03	31	7.6–10.8	13.23	triggers at 2.60–3.43 %
4 0.40 μΩ/h, 60 min window	24	1	1.80–3.50	304	7.8–13.1	13.51	11.65 t CO_2-eq
4 0.40 μΩ/h, 30 min window	24	0	1.88–3.50	196	7.8–10.5	13.27	
5 stuck feeder, termination on (shipped)	12	1	1.80–3.50	61	7.7–13.0	13.76	4.45 min, 11.65 t CO_2-eq
5 termination off	12	1	1.74–3.50	65	7.7–19.3	14.65	12.45 min, 32.58 t CO_2-eq
5 feeder never freed	12	5	1.80–3.50	117	7.7–13.9	15.75	58.6 t CO_2-eq
6 anode change (shipped)	12	0	2.16–3.50	0	5.1–10.0	13.27	ledge to 17.2 cm
6 with compensation	12	0	2.18–3.50	0	6.3–10.0	13.37	ledge to 15.8 cm
6 hold off	12	0	2.16–3.51	0	4.7–10.0	13.27	false trigger at 2.91 %
7 tap (shipped)	3	0	2.53–3.50	0	8.7–11.9	13.52	R to 10.38 μΩ
7 tap compensation	3	0	2.51–3.50	0	8.5–10.0	13.28	R to 8.64 μΩ
7 move limit 10 mm	3	0	2.52–3.50	0	8.7–10.4	13.42	R to 10.02 μΩ
8 target 7.90 μΩ (shipped)	48	0	2.19–3.89	0	2.2–10.0	12.57	sludge to 129 kg
8 heat loop in AUTO	48	0	2.17–3.89	0	2.9–10.0	12.91	sludge to 70 kg
9 wet lot at 2.0 % (shipped)	12	1	1.80–3.50	101	7.6–14.2	14.12	19.9 t CO_2-eq
9 threshold 0.10 μΩ/h	12	0	2.14–3.50	0	7.9–10.4	13.26	the lot never comes
9 overfeed 300 %	12	0	1.92–3.89	21	7.6–12.7	13.26	sludge to 312 kg

38.2 V at its flickering peak, the superheat goes from 6.7 to 18.2 K, and in the 12.5 min before the crew's kill the cell vents 4.11 kg of CF_4, 32.8 t of CO_2-equivalent. With the feed unchanged, effects follow at 4.00 and 4.64 h, and the five hours cost 98.4 t of CO_2-equivalent at 23.4 kWh/kg. Fed the base rate from 2.5 h, when the resistance has barely begun to rise, the cell has no anode effect and settles near 2.38 %.

> **TRY IT IN THE APP** **1 · Feeding by Hand: Find the Minimum, Cross It**
>
> Load scenario 1 and press Run. Watch the resistance R in the badge row while the hidden alumina falls. Between 4.1 and 2.7 % alumina the reading stays between 8.595 and 8.641 µΩ, a band the size of the pad's noise. Near 3.35 h the alumina reaches 1.80 %, and the voltage jumps to about 30 V.

11.6.2 A Tally That Balances While the Bath Runs Down

Case 2 feeds F_b, one shot every 36 s, without looking at the resistance. The shots are 2.8 % short, so the cell loses about 5 kg of alumina an hour while the controller's tally stays within 1.8 kg of balance until the first anode effect. At 12 h the alumina is 2.86 % and the filtered resistance 8.555 µΩ, at 24 h 2.24 % and 8.567 µΩ, because the bottom of the curve is flat and the beam loop holds the resistance to its target. At 32.53 h the alumina reaches the wall with a true shortfall of 165 kg, and the kill ends the effect in 4.4 min. A second effect follows at 42.19 h, and the alumina spends 476 min outside its band. At 103 % of F_b the drift reverses, to 3.60 % after two days, and at 97 % the first effect comes at 15.82 h and six more follow about every 4.7 h. Open-loop feeding would need the dose and the current efficiency to a fraction of a percent, held for days.

> **DID YOU KNOW?** **Three Percent Either Way**
>
> From the simulator: fixed feeding at 103 % of the Faradaic rate keeps the alumina between 3.49 and 3.60 % for two days with no anode effect, while 97 % brings the first anode effect at 15.82 h and seven in the two days (Table 11.5). A feeder can be counted, but its dose cannot be trusted that closely.

11.6.3 Riding the Low Side

Case 3 is the reference line. The first underfeed, from 3.50 %, runs 2.10 h and ends at 2.24 % on a slope of 0.103 µΩ/h. After it the cycle repeats every 2.29 h on average, 21 triggers in two days, each at 2.14 to 2.45 % alumina and a slope of 0.100 to 0.112 µΩ/h. From the first trigger on, the alumina stays between 2.14 and 3.07 %, the superheat between 7.9 and 10.4 K, and no underfeed times out, at 13.24 kWh/kg with 67 beam moves that total 5.9 mm. By 48 h the controller has counted 132 kg more than its base rate, while the cell received 113 kg less than it consumed, having started at the minimum. The cycle absorbs a constant feed error because it ends each underfeed on the curve and not on the clock, as shown in Figure 11.4. A 50 % underfeed triggers 25 times at 2.24 to 2.69 %, every 1.88 h, and an 85 % one 14 times at 2.07 to 2.29 %, every 3.14 h, with one underfeed timing out.

Pause. Try Exercise 11.2 before reading on.

11.6.4 Threshold, Window and Noise

Case 4 lowers the threshold to 0.03 µΩ/h, less than two standard deviations of the noise above the flat-bottom slope of −0.021 µΩ/h. The ten triggers of the day come at 2.60 to 3.43 % alumina, half of them or more where the noise-free resistance rose more slowly than the threshold, as the application's own check confirms, so the roll ends the underfeeds. Each early end still buys a full overfeed, the alumina ratchets past the minimum to 4.03 %, and 31 min are spent above the band. At 0.40 µΩ/h the 30 min fit ends its underfeeds at 1.88 to 1.92 %, a few hundredths of a percent from the wall. Behind a 60 min window every underfeed of the day ends between 1.805 and 1.810 %. After the second, which ends at 5.108 h, the overfeed's alumina dissolves too slowly, and the anode effect starts at 5.123 h, about a minute after the trigger and the only one in 24 h. At 0.10 µΩ/h the alumina stays between 2.14 and 3.50 %, and a pad instability fired at 3.36 h, just after the second underfeed of that run begins, ends it at 3.68 h and 2.68 % instead of at 4.44 h and 2.21 %.

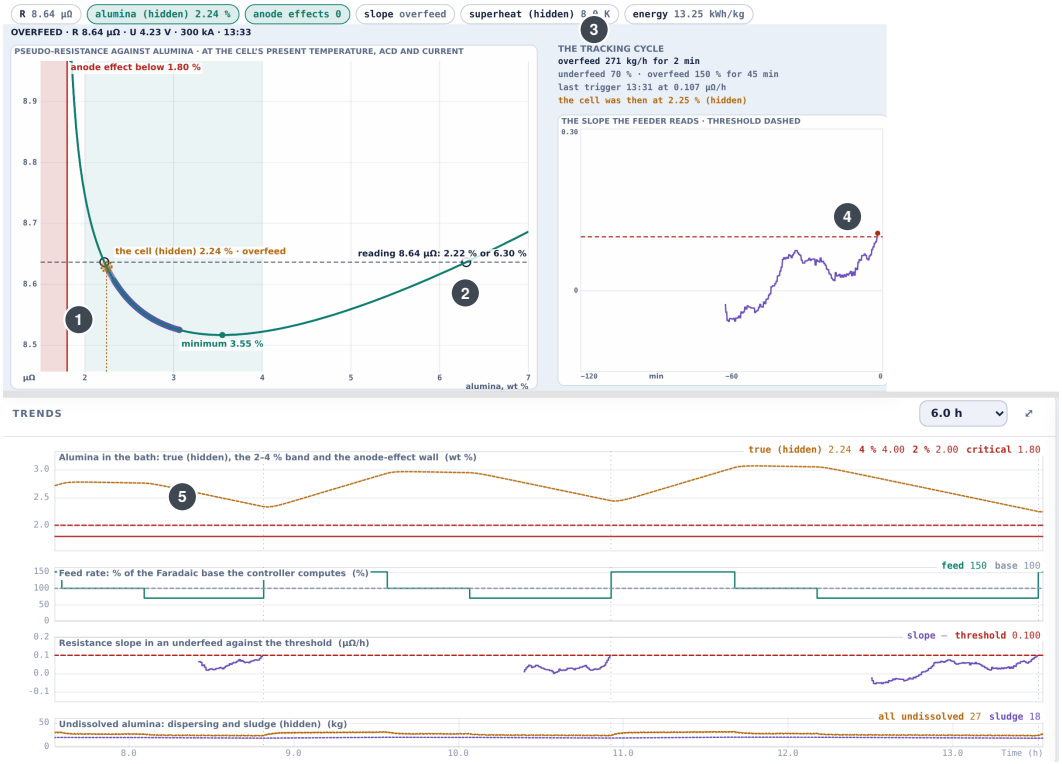


Figure 11.4 The cell in case 3 (Demand feeding rides the slope) at $t = 13.55$ h (13:33) on the Curve view with the 6 h trend window and every control as shipped. The header reads OVERFEED, R 8.64 µΩ, U 4.23 V and 300 kA, and the badges give the hidden alumina at 2.24 %, no anode effect, a hidden superheat of 8.0 K and 13.25 kWh/kg. The curve, drawn at the present 964.3 °C and ACD of 3.91 cm, has its minimum at 3.55 % and its wall at 1.80 %, and the dashed reading of 8.64 µΩ crosses it at 2.22 % and 6.30 %. The tracking panel records the trigger at 13:31 at 0.107 µΩ/h with the cell then at 2.25 %, the sixth of the run, followed by 2 min of overfeed at 271 kg/h, 150 % of 181 kg/h, and the slope panel shows that underfeed's slope climbing to the dashed 0.10 µΩ/h. Over the window the trends show three underfeeds with the alumina between 2.24 and 3.07 %. The callouts mark (1) the hidden alumina where the cell sits on the curve, (2) the reading crossing the curve on both sides of its minimum, (3) the tracking panel, (4) the slope reaching its dashed threshold and (5) the alumina trend above the dashed 2 % line of its band.

11.6.5 A Stuck Feeder and the Kill

Case 5 sticks one of the four feeders from 1 to 5.5 h. It still fires and is still counted, but nothing falls, so the 150 % overfeed delivers 109.8 % of the consumption and the underfeed 50.1 %, and the cycle cannot put back what its underfeeds take. At 4.447 h the alumina reaches 1.80 %. Automatic termination ends the effect in 4.45 min with 1.46 kg of CF_4 and 0.177 kg of C_2F_6, 11.65 t of CO_2-equivalent, while the superheat rises from 8.76 to 12.96 K. Left for the crew, the effect lasts 12.45 min, vents 4.09 kg of CF_4, 32.58 t of CO_2-equivalent, heats the bath to 19.30 K of superheat and raises the 12 h specific energy from 13.76 to 14.65 kWh/kg. Never freed, the feeder brings effects at 4.45, 6.03, 7.60, 9.17 and 10.74 h.

> **TRY IT IN THE APP** **5 · An Anode Effect and Its Termination**
>
> Load scenario 5 and run 12 h. The stuck feeder takes the alumina to 1.80 % at 4.447 h, and automatic termination ends the effect in 4.45 min for 11.65 t of CO_2-equivalent. Turn automatic termination off and run again. The crew's kill comes later, the effect lasts 12.45 min and vents 32.58 t.

11.6.6 Current That Moves Between Anodes

Case 6 changes one anode at 1 h. While it is out the other 23 carry the current, the noise-free resistance steps from 8.619 to 8.868 μΩ and the critical alumina rises from 1.800 to 1.878 %. The new anode draws 52.2 kW an hour after the change, and the cell has 23.38 of its 24 shares at 5 h. The beam loop reads the step as a long ACD and lowers the beam 1.4 mm, taking heat away when the cold anode is drawing it, so the superheat falls from 8.71 to 5.13 K at 5.25 h and the ledge thickens from 12.7 to 17.2 cm. With the compensation the target and the filtered reading step together, the beam setpoint stays within 0.02 mm of where it was until 6 h, and the superheat bottoms at 6.27 K at 5.32 h and the ledge at 15.8 cm. With the hold off, the step reads as a slope, ends an underfeed a minute after the change at 2.91 % alumina and starts an overfeed the cell did not need, and the superheat bottoms at 4.70 K.

Pause. Try Exercise 11.3 before reading on.

11.6.7 A Tap Under a Rate-Limited Beam

Case 7 taps 2300 kg at 1 h, faster than the move limit lets the beam follow. The filtered resistance peaks at 10.38 μΩ at 1.25 h, 1.75 μΩ above target, with the ACD opened to 5.57 cm, and returns to its dead band at 1.78 h, 47 min after the tap began, after 10 moves that total 20.9 mm. The extra voltage heats the bath from 8.7 to 11.9 K of superheat and costs 64 kWh more over the three hours than the same run without a tap. With tap compensation the beam follows the metal, the resistance never exceeds 8.64 μΩ, and the energy stays within 5 kWh of the untapped run. A move limit of 10 mm returns in 32 min for 35 kWh but still peaks at 10.02 μΩ.

11.6.8 A Target That Starves the Bath of Heat

Case 8 sets the target to 7.90 μΩ, 66 kW less heat than the 8.63 μΩ that holds the cell in balance, and the beam loop closes the ACD as told. Its first underfeed ends at the first armed decision, at 0.33 h and 3.32 %, because the beam's 4.9 mm descent is taken out of the slope at the controller's 0.12 μΩ/mm instead of the plant's 0.117 (with 0.117 in a copy of the code it runs to 1.69 h), and the overfeed that follows lifts the alumina to 3.89 % at 1.22 h. The superheat reaches 2.85 K at 3.9 h, where the next underfeed ends, and then swings about 1.6 K either side of a mean of 3.83 K from 6 to 48 h, lowest at the end of each underfeed and 2.22 K at 14.5 h. That mean lies 0.64 K under the 4.47 K of the static balance, as case 3 lies 0.70 K under its own. A cold bath dissolves alumina slowly and sinks more of it, so the undissolved alumina reaches 106 kg at 12 h and 129 kg at its most, against 32 kg in case 3, while the ledge grows from 12.6 to 22.6 cm by 12 h and to 37.4 cm by 48 h. The feed logic reads the dissolved alumina running low as an underfed cell and feeds 109 kg more than case 3 in the two days, most of it into sludge. The energy is cheaper, 12.57 kWh/kg against 13.24, and the ledge and the sludge are the bill. With the heat loop in AUTO the dip of 06:00, reported at 07:00, estimates 1.89 K where the truth was 5.04 K, because the alumina was 3.03 % against the assumed 2.6 and the thermocouple read 0.84 K low. The loop raises the target by its 0.3 μΩ limit to 8.2, and to 8.5 μΩ at the next report, and the superheat bottoms at 2.85 K at 3.9 h, before the first report, recovers to 6.71 K by 12 h and ends at 7.64 K, with 70 kg of sludge at most and 12.91 kWh/kg.

11.6.9 A Wet Lot Below Two Percent

Case 9 raises the threshold to 0.25 μΩ/h, so that underfeeds run deep, and lets a wet lot reach the feeders the moment the alumina falls through 2.0 %, at 2.497 h in the first underfeed. The slope passes the threshold 2.20 min later at 0.257 μΩ/h and the overfeed starts, yet the alumina keeps falling,

Table 11.6 The Two-Day Challenge in Eight Configurations

Configuration	AE	Alumina (%)	Out (min)	Superheat (K)	*SE* (kWh/kg)	
as shipped	5	1.80–3.50	594	3.40–16.21	14.14	fails
demand feeding	0	2.02–3.50	0	2.80–10.01	12.96	fails
+ heat loop	0	1.91–3.50	31.65	4.33–10.72	13.20	fails
+ compensation	0	1.92–3.53	30.1	5.48–11.17	13.23	fails
suggested[a]	0	2.25–3.56	0	5.32–12.50	13.32	passes
no compensation	0	2.26–3.55	0	4.68–12.15	13.29	fails
heat loop manual	0	2.25–3.50	0	3.32–10.12	12.99	fails
fixed feeding	5	1.80–3.50	439	6.61–17.82	14.33	fails

[a] demand feeding at 0.08 μΩ/h with a 60 min overfeed, heat loop in AUTO, anode-change compensation on.

because the lot's shots disperse in 10 min and lose a further 30 % to sludge, so that at 150 % they dissolve about what the cell consumes. The anode effect starts at 2.769 h, 16.4 min after the lot, and vents 19.9 t of CO_2-equivalent. The limit is one of information, not tuning. An oracle that meets the lot at 2.0 % with an immediate 45 min overfeed, on a cell that had been underfeeding at 126 kg/h, still reaches the wall 75.2 min after the lot at 150 %, half an hour after its overfeed ends, and 218 min after it at 200 %. Only 300 % holds the cell, which bottoms at 1.966 %. In case 9 itself a 300 % overfeed bottoms at 1.92 % with 312 kg of sludge, while on dry alumina at the default 0.10 μΩ/h the same overfeed carries the cycle up to 6.60 % in two days. The defense is margin, as it is for the inferred limit of Volume 15, Chapter 8. At thresholds of 0.15 and 0.10 μΩ/h the underfeeds end above 2 %, at 2.06 and 2.14 % at the lowest, and the lot never meets a cell that low.

11.6.10 Two Days at 300 kA

The challenge runs 48 h through an anode change at 8 h, a tap at 20 h and a wet lot from 30 h, and asks for no anode effect, alumina between 2 and 4 % for all but 30 min, superheat between 5 and 15 K and at most 13.5 kWh/kg, scored on the truth. The challenge is judged on the case as given, so that a changed model parameter or disturbance voids the attempt. As shipped the cell is fed F_b with a target of 8.30 μΩ left from an energy drive. Table 11.6 compares the configurations. As shipped the superheat fails first, below 5 K at 8.99 h, an hour after the anode change, and 3.40 K at its lowest. The drifting alumina leaves the band at 27.07 h, the first anode effect starts at 30.08 h, five minutes after the wet lot, four more follow, and the two days end at 14.14 kWh/kg and 107 t of CO_2-equivalent.

Each element repairs one line of the target. Demand feeding alone removes the anode effects and lets the low target cool the bath to 2.80 K. Adding the heat loop lifts the target at the first dip, but the anode change still takes the superheat to 4.33 K and the default cycle dips to 1.91 % on the wet lot. Adding the anode-change compensation keeps the superheat above 5 K, at 5.48 K at its lowest, but leaves the alumina 30.07 min outside the band, four seconds more than allowed, a margin within the noise, since it passes on three of five other seeds and the suggested configuration on all six. The suggested configuration rides higher and keeps the alumina between 2.25 and 3.56 % with no minute out of the band and the superheat between 5.32 and 12.50 K, at 13.32 kWh/kg, margins of 0.32 K on the superheat and 0.18 kWh/kg on the energy. Without its compensation it bottoms at 4.68 K, with the heat loop in MANUAL at 3.32 K, and on fixed feeding it has five anode effects. Its success also leans on an inference error. Both dips read the superheat low, 3.30 against 8.07 K and 5.75 against 10.31 K, because the cell held 3.31 and 3.23 % alumina against the assumed 2.6, so the loop keeps raising the target, to 8.86 μΩ by the end.

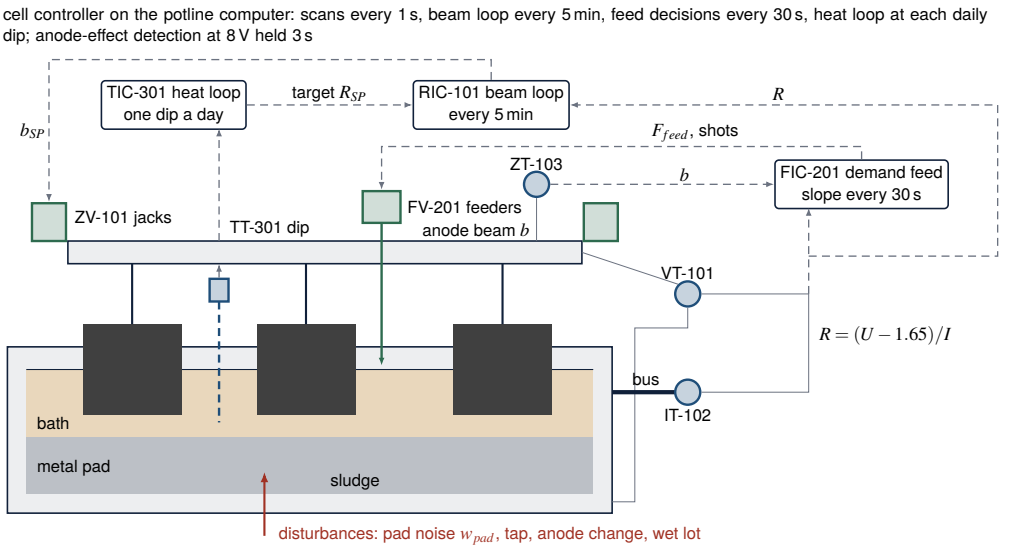


Figure 11.5 The cell's three loops drawn with their hardware. The cell voltage VT-101 and the line current IT-102 give the pseudo-resistance, and the beam encoder ZT-103 lets the controller remove its own beam moves from the slope. The beam loop RIC-101 moves the anode beam on its jacks ZV-101, the demand feed loop FIC-201 fires the point feeders FV-201, and the heat loop TIC-301 turns a daily dip TT-301 into the beam loop's target. Notice that no sensor reaches the alumina in the bath: every loop reads it through the voltage.

11.7 Technology in Practice

The controller of Section 11.5 reads two electrical signals, the beam position and a daily dip, and it moves the beam and the feeders. Figure 11.5 draws its loops with that hardware in place of the schematic's symbols.

What Measures It

The chapter's controller never reads the dissolved alumina. Its pseudo-resistance is computed from the cell voltage VT-101 and the line current IT-102, and the beam encoder ZT-103 lets the slope estimator take the beam's own moves out of the reading. The heat balance is seen once a day, through the dip TT-301 of the bath temperature and the excess aluminum fluoride.

> **TECHNOLOGY Sensor: Cell Voltage and Line Current**
>
> Cell control rests on two electrical readings. The voltage of each cell is read directly, and one potline system gathered the voltages of hundreds of pots over a shared cable to a central control computer [19]. A cell normally runs at 4.0 to 4.5 V, and an anode effect raises it to perhaps 30 to 40 V [1]. The line current is the same for every cell in series and is kept constant [1]. It can be read without contact from the magnetic field around the conductor, which Hall-effect sensors resolve even for direct current. A shared cable is a weak point, because a cut loses every reading it carries at once [19]. The chapter's VT and IT are read every second, with the pad's noise on the voltage.

What Moves It

The second handle is the anode beam. The anodes burn away, so the beam has to be lowered step by step to keep the anode–cathode distance, while regular tapping holds the metal surface at about the same height. When the beam reaches its lower stop, a beam-raising machine carried by the overhead

crane lifts it again [1]. The chapter's beam ZV-101 travels 400 mm on jacks that run at 10 mm/min, with a move limit of 3 mm in each interval, and the feeders FV-201 turn a requested rate into shots.

> **TECHNOLOGY** **Final Element: The Point Feeder**
>
> A point feeder keeps a hole open in the crust and drops a measured dose of alumina into the bath. The alumina comes from a bin in the cell's superstructure, and two to six volumetric feeders in turn add about 1 kg every minute or so, doses small enough to dissolve and disperse fast. The holes lie along the center line between the two rows of anodes, and the aluminum fluoride goes in through them too. Too much alumina forms sludge, which settles under the metal and raises the cell's resistance [1]. The chapter's four feeders drop 1.75 kg and count 1.80 kg a shot, and case 5 shows a feeder that fires and is counted while nothing falls.

Where the Controller Runs

The potline current is held constant, and each cell has its own voltage adjustment for its heat balance and its condition [1]. Cell voltages are gathered to a central control computer [19]. Control systems of this kind run a supervisory layer over localized controllers, with cycle times from milliseconds to minutes [20]. The chapter's cell controller scans every second. Its beam loop RIC-101 acts every 5 min through the shared velocity-form PID, the feed loop FIC-201 decides every 30 s, and the heat loop TIC-301 moves the target of RIC-101 once a day, a cascade whose outer loop is the slowest of the three.

> **IN PRACTICE** **Anode Effects Are Counted and Reported**
>
> Across the industry a high-voltage anode effect is typically a cell voltage above 8 V for 3 s, the rule of the chapter's detector. Many point-fed smelters end it automatically by moving the anodes and feeding alumina, mostly with no manual step [21]. In the United States a smelter reports its anode-effect frequency, duration and minutes per cell-day, computed monthly, and the method that measures them. Its slope coefficients must be measured again at least every ten years, or sooner when a change in the control algorithm alters the mix of anode effects or the termination routine [22].

What Is Modernizing It

Three developments touch the loops of this chapter. The first is power modulation. Smelters have tried running at a varying power since the 1970s, and poor control of the heat balance then froze cryolite in the bottom of cells. Recent work optimizes the line current and the anode–cathode distance together under limits on the heat balance, so that a potline can follow a power price or a renewable supply, and offers this as a foundation for online control policies [23]. Case 8 shows what an unwatched heat balance costs. The second is online gas measurement. Instruments at the smelter can measure perfluorocarbons continuously and separate the emissions of high-voltage and low-voltage anode effects, though neither continuous nor sampled measurement is routine in the industry yet [21]. The third is the inert anode, first suggested in 1886. It would give off oxygen instead of carbon dioxide, and an anode that lasted the life of the cell would end the anode change of case 6, but every material dissolves to some extent in the corrosive cryolite [1].

11.8 Limitations

One well-mixed bath stands for a cell whose alumina differs between feeders and far corners, so an anode effect cannot start on one anode while the mean is above the wall, as real ones do [16]. The curve, every threshold, the bands of 2 to 4 % and 5 to 15 K, the emission coefficients and the challenge's energy ceiling are teaching values, and their real counterparts come from each plant's cell control and measurement campaigns [9, 15], as Section 11.3.8 records. Each disturbance arrives as a step, which a potroom would rarely allow, but a step keeps one cause apart from the others. The current efficiency is constant, so a hot, squeezed or unstable cell loses energy here but not metal. The

crew's bath level removes the ledge's effect on the bath's composition, the pad's noise cannot grow by itself, and the bath, metal and walls are lumped. The controller's assumptions differ from the plant only by declared amounts, and a dip a day is its only reading of the heat.

OPEN CHALLENGE Anode Effects the Voltage Does Not Show

Perfluorocarbons can also form without the cell voltage rising. Such low-voltage anode effects have been found since the late 2000s in high-amperage cells with many large anodes, and they come from the same upset on a smaller, local scale. Higher current density, a smaller anode–cathode distance, larger anodes with less bath to dissolve alumina, more anodes for each feeder and longer cells all make them more likely [21], and larger cells with fewer feeders for their current have bath that differs more from place to place [23]. Progress would be a controller that finds a local shortage of alumina before the gas forms, which one voltage across the whole cell cannot.

11.9 Conclusion

Feeding a Hall–Héroult cell means reading a measurement that cannot say on which side of its minimum the alumina lies, and the cases put a price on that blindness and on each remedy. On the cell's curve 8.68 µΩ means 2.5 % or 5.14 %, and the filtered reading moves only 0.046 µΩ while the alumina falls from 4.1 to 2.7 %. Feeding the Faradaic rate with shots 2.8 % short reaches the anode effect at 32.5 h with the controller's tally within 1.8 kg of balance, while demand feeding at 0.10 µΩ/h holds 2.14 to 3.07 % for two days without measuring alumina. A half-window rule places the end of each underfeed within 0.03 % of the engine at the default window and bounds the threshold between the noise, which ends underfeeds early at 0.03 µΩ/h, and the wall, which a 60 min window meets at 0.40. A stuck feeder's anode effect costs 11.65 t of CO_2-equivalent when it is killed at once and 32.58 t when the crew arrives after 8 min, and a target 0.73 µΩ too low runs at 12.57 kWh/kg while the ledge grows to 37.4 cm. Everything rests on one lumped cell and on teaching values. A fuller model would add alumina gradients and local anode effects, a current efficiency that varies, ledge-driven changes of the bath's composition and the neighboring cells of a potline. The cell is a place to learn control through a measurement with an extremum.

KEY POINTS

- The controller reads a resistance that is lowest in the middle, so one reading cannot say whether the alumina is too low or too high.
- Feeding the Faradaic rate with shots 2.8 % short reaches the anode effect at 32.5 h while the controller's tally stays within 1.8 kg of balance.
- Demand feeding at 0.10 µΩ/h holds 2.14 to 3.07 % for two days without measuring alumina, because it ends each underfeed on the curve and not on the clock.
- A stuck feeder's anode effect costs 11.65 t of CO_2-equivalent when it is killed at once and 32.58 t when the crew arrives after 8 min.

Where else this applies. Regulation near the peak of a plant characteristic also governs the braking tire of Volume 7, Chapter 4, and a sensor that rises on both sides of the right dose the urea loop of Volume 5, Chapter 12.

References

[1] H. Kvande and P. A. Drabløs, "The aluminum smelting process and innovative alternative technologies," *Journal of Occupational and Environmental Medicine*, vol. 56, no. 5S, pp. S23–S32, 2014.

[2] A. T. Tabereaux and D. S. Wong, "Awakening of the aluminum industry to PFC emissions and global warming," in *Light Metals 2021*, ser. The Minerals, Metals & Materials Series, L. Perander, Ed. Springer, 2021, pp. 554–564.

[3] V. Potocnik and M. Reverdy, "History of computer control of aluminum reduction cells," in *Light Metals 2021*, ser. The Minerals, Metals & Materials Series, L. Perander, Ed. Springer, 2021, pp. 591–599.

[4] W. Haupin, "Interpreting the components of cell voltage," in *Essential Readings in Light Metals: Volume 2, Aluminum Reduction Technology*, G. Bearne, M. Dupuis, and G. Tarcy, Eds. Springer, 2016, pp. 153–159.

[5] H. Kvande, B. P. Moxnes, J. Skaar, and P. A. Solli, "Pseudo resistance curves for aluminium cell control - alumina dissolution and cell dynamics," in *Essential Readings in Light Metals: Volume 2, Aluminum Reduction Technology*, G. Bearne, M. Dupuis, and G. Tarcy, Eds. Springer, 2016, pp. 760–766.

[6] U.S. Geological Survey, "Mineral commodity summaries 2025: Aluminum," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[7] M. Krstić and H.-H. Wang, "Stability of extremum seeking feedback for general nonlinear dynamic systems," *Automatica*, vol. 36, no. 4, pp. 595–601, 2000.

[8] T. Esram and P. L. Chapman, "Comparison of photovoltaic array maximum power point tracking techniques," *IEEE Transactions on Energy Conversion*, vol. 22, no. 2, pp. 439–449, 2007.

[9] P. Homsi, J.-M. Peyneau, and M. Reverdy, "Overview of process control in reduction cells and potlines," in *Essential Readings in Light Metals: Volume 2, Aluminum Reduction Technology*, G. Bearne, M. Dupuis, and G. Tarcy, Eds. Springer, 2016, pp. 739–746.

[10] T. R. Beck, "ECS classics: Hall and Héroult and the discovery of aluminum electrolysis," *The Electrochemical Society Interface*, vol. 23, no. 2, pp. 36–37, 2014.

[11] T. K. McMahon and G. P. Dirth, "Computer control of aluminum reduction cells," *JOM*, vol. 18, no. 3, pp. 317–319, 1966.

[12] H. Kvande and W. Haupin, "Cell voltage in aluminum electrolysis: A practical approach," *JOM*, vol. 52, no. 2, pp. 31–37, 2000.

[13] T. Drengstig, D. Ljungquist, and B. A. Foss, "On the AlF_3 and temperature control of an aluminum electrolysis cell," *IEEE Transactions on Control Systems Technology*, vol. 6, no. 2, pp. 157–171, 1998.

[14] A. T. Tabereaux, "Anode effects, PFCs, global warming, and the aluminum industry," *JOM*, vol. 46, no. 11, pp. 30–34, 1994.

[15] B. P. Leber, A. T. Tabereaux, J. Marks, B. Lamb, T. Howard, R. Kantamaneni, M. Gibbs, V. Bakshi, and E. J. Dolin, "Perfluorocarbon (PFC) generation at primary aluminum smelters," in *Essential Readings in Light Metals: Volume 2, Aluminum Reduction Technology*, G. Bearne, M. Dupuis, and G. Tarcy, Eds. Springer, 2016, pp. 1015–1023.

[16] G. Tarcy and A. Tabereaux, "The initiation, propagation and termination of anode effects in Hall-Heroult cells," in *Light Metals 2011*, S. J. Lindsay, Ed. Wiley, 2011, pp. 329–332.

[17] X. Wang, R. D. Peterson, and A. T. Tabereaux, "Electrical conductivity of cryolitic melts," in *Essential Readings in Light Metals: Volume 2, Aluminum Reduction Technology*, G. Bearne, M. Dupuis, and G. Tarcy, Eds. Springer, 2016, pp. 57–64.

[18] A. Solheim, S. Rolseth, E. Skybakmoen, L. Støen, Å. Sterten, and T. Støre, "Liquidus temperatures for primary crystallization of cryolite in molten salt systems of interest for aluminum electrolysis," *Metallurgical and Materials Transactions B*, vol. 27, no. 5, pp. 739–744, 1996.

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

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

[21] Intergovernmental Panel on Climate Change, "Chapter 4: Metal industry emissions," in *2019 Refinement to the 2006 IPCC Guidelines for National Greenhouse Gas Inventories, Volume 3: Industrial Processes and Product Use*. Geneva: IPCC, 2019. [Online]. Available: https://www.ipcc-nggip.iges.or.jp/public/2019rf/pdf/3_Volume3/19R_V3_Ch04_Metal_Industry.pdf

[22] U.S. Environmental Protection Agency, "Mandatory greenhouse gas reporting, subpart F: Aluminum production, 40 CFR 98.60 to 98.68," Code of Federal Regulations, Title 40, Volume 23, revised as of July 1, 2024, U.S. Government Publishing Office, 2024. [Online]. Available: https://www.govinfo.gov/app/collection/cfr/2024/title40

[23] C.-J. Wong, A. A. Larkin, J. Bao, M. Skyllas-Kazacos, B. J. Welch, N. Ahli, M. Faraj, and M. Mahmoud, "Optimisation of power modulation for Hall-Héroult cells: Process operability and constraints as virtual energy storage," arXiv preprint 2511.07893, 2025.

11.10 Exercises

11.1 Concept check · 2 min

On the curve of Figure 11.2(a), drawn at the design point, the control room reads a pseudo-resistance of 8.68 μΩ. What does that reading alone say about the alumina?

(a) It is near 3.5 %, because the curve is flattest at its minimum
(b) It is near 2.5 or near 5.1 %, and the reading cannot say which
(c) It is near 2.5 %, because only the low side can raise the reading
(d) Nothing at all, because the metal pad's noise hides every reading of R

11.2 Concept check · 1 min

Scenario 2 feeds the Faradaic rate with feeders 2.8 % short and reaches the anode effect at 32.5 h. Scenario 3, Demand feeding rides the slope, has the same feeders and no anode effect in two days. Why does demand feeding not drift?

(a) It reads the alumina in the daily dip and trims the feed rate to it
(b) Its underfeeds and overfeeds average exactly to the Faradaic rate
(c) It ends each underfeed on the slope of the curve, not on the clock
(d) It corrects its count of shots to the true dose of 1.75 kg

11.3 Concept check · 2 min

In scenario 6, The anode change, one of 24 anodes is out of the cell at 1 h. The noise-free resistance steps from 8.619 to 8.868 μΩ while the ACD and the alumina have not moved. With anode-change compensation off, what does the beam loop do?

(a) Holds the beam, because the step lies inside its dead band of 0.08 μΩ
(b) Raises the beam, reading the step as a short ACD, and so heats the bath
(c) Waits out the hold after the event, so the beam stays where it was
(d) Lowers the beam, reading the step as a long ACD, and so cools the bath

11.4 Theory and calculation · 5 min

The cell runs at $I = 300\,\mathrm{kA}$ with a current efficiency of 95 %. Use $F = 96\,485.33\,\mathrm{C/mol}$ and $M_{Al_2O_3} = 0.101\,961\,\mathrm{kg/mol}$ in Equation (11.1).

(a) Find the alumina consumption in kg/h and the time between shots of the counted 1.80 kg. *Checkpoint:* 180.70 kg/h and about 36 s.

(b) Each shot really delivers 1.75 kg. On fixed feeding at 100 %, how much alumina does the cell lose per hour? *Checkpoint:* About 5.0 kg/h.

(c) The bath starts with 315 kg of dissolved alumina in 9000 kg in all, 3.50 %, and the other 8685 kg stay constant. Find the dissolved alumina at the wall of 1.80 % and estimate when the cell reaches it. *Checkpoint:* About 159 kg, a loss of 156 kg, in about 31 h.

11.5 Theory and calculation · 8 min

At the design point, $T = 959.8\,^{\circ}\mathrm{C}$, $\alpha = 1$, $m = 0.492$, $c_{crit} = 1.80\,\%$, the bath's ohmic drop is 1.550 V and $I = 300\,\mathrm{kA}$. Use $R_g = 8.314\,462\,\mathrm{J/(mol\,K)}$ and $F = 96\,485.33\,\mathrm{C/mol}$ in Equation (11.10).

(a) Evaluate the reaction, concentration and ohmic terms at 3.5 % alumina, in mV per percent. *Checkpoint:* -14.9, -16.1 and $+31.0$, a sum near zero.

(b) Repeat the two anode terms at 2.5 %, take the ohmic term there as 30.4 mV per percent, and convert the sum to μΩ per percent. *Checkpoint:* -20.9 and -54.6; about -0.150 μΩ per percent.

(c) A 70 % underfeed lowers the alumina 0.615 % per hour, and the closing ACD lowers R by 0.021 µΩ/h. What slope does the controller see at 2.5 %, and how much alumina does a 30 min least-squares window lag behind? *Checkpoint:* About 0.072 µΩ/h, below the threshold of 0.10; a lag of about 0.15 %.

11.6 Theory and calculation · 5 min

At 300 kA the beam loop holds a target resistance. The ledge area is 16 m^2 with a bath-to-ledge coefficient of 700 W/(m^2 K), the top and bottom lose 575 W/K, and bath and metal hold $C = 55\,\mathrm{MJ/K}$.

(a) How much heat does each 0.1 µΩ of target add? *Checkpoint:* 9.0 kW.

(b) With the liquidus fixed, by how much does the steady superheat rise per 0.1 µΩ? Starting from 10.01 K at 8.625 µΩ, estimate the steady superheat at a target of 7.90 µΩ. *Checkpoint:* 0.76 K, and about 4.47 K.

(c) Find the time constant of the bath alone. *Checkpoint:* 1.30 h.

11.7 App challenge · 8 min

Load scenario 2, Fixed feeding at the Faradaic rate, at the default seed and a speed of 4 h per second. For each run below, let the clock pass 48 h, pause, export the CSV file and read the row at `Time (h)` = 48 and the rows named.

(a) Run as shipped. Find the first row where `ae_count` is 1, and in the row at 32.5 h read `feed_balance_kg` and the difference `fed_true_kg` minus `consumed_kg`. Give times within 0.02 h and masses within 2 kg. *Checkpoint:* 32.53 h; about −1 kg counted against −164 kg real.

(b) Press Reset, set the Fixed feed rate to 97 % before Run, and find the first anode effect and the count at 48 h. *Checkpoint:* 15.82 h; 7 anode effects.

(c) Press Reset, set 103 %, and read `ae_count` and `c_true_pct` at 48 h, the latter within 0.01 %. *Checkpoint:* 0 anode effects; 3.60 %.

11.8 App challenge · 9 min

Load scenario 4, The slope threshold against the noise, at the default seed and a speed of 4 h per second. For each run below, let the clock pass 24 h, pause, export the CSV file and read the row at `Time (h)` = 24.

(a) Run as shipped, with the Slope threshold (low side) at 0.03 µΩ/h. Read `metric_c_min_pct`, `metric_c_max_pct`, `metric_c_out_min` and `triggers`, the alumina within 0.01 %. *Checkpoint:* 2.59 to 4.03 %, 31 min out of the band; 10 triggers.

(b) Press Reset, set the Slope threshold to 0.40 µΩ/h, switch to Advanced and set the Slope filter window to 60 min before Run. Read `ae_count`, `co2e_t`, `metric_c_min_pct` and `SE_kWhkg`, the energy within 0.01 kWh/kg. *Checkpoint:* 1 anode effect, 11.65 t CO_2-eq; 1.80 %; 13.51 kWh/kg.

(c) The slope fitted over 30 min at the flat bottom has a standard deviation of 0.027 µΩ/h, and the closing ACD lowers every slope by 0.021 µΩ/h. How many standard deviations above the flat-bottom slope does a threshold of 0.03 sit? *Checkpoint:* About 1.9.

Explain each run with the two edges of a usable threshold.

11.9 App challenge · 9 min

Load scenario 8, The heat balance: a target set too low, at the default seed and a speed of 4 h per second, with the Target resistance at its shipped 7.90 µΩ.

(a) Run with the Heat loop in Manual until the clock passes 48 h, pause and export the CSV file. In the row at `Time (h)` = 48 read `metric_SH_min_K`, `metric_SH_max_K`, `SE_kWhkg`, `metric_sludge_max_kg` and `ledge_cm`, the ledge within 0.1 cm. *Checkpoint:* 2.2 to 10.0 K; 12.57 kWh/kg; 129 kg at most; 37.4 cm.

(b) Press Reset, set the Heat loop to Auto on the TIC-301 card, and run past 7 h. Read the dip on the TIC-301 card and the target it wrote, and read `SH_true_K` in the row at 6.0 h. *Checkpoint:* Superheat 1.9 K from the dip of 06:00, target 8.20 µΩ; true superheat 5.0 K.

(c) Run on past 48 h, export the file again and read the first four columns of (a) in the row at 48 h. *Checkpoint:* 2.9 to 10.0 K; 12.91 kWh/kg; 70 kg at most.

Explain why the superheat still bottoms below 3 K with the loop in Auto.

11.10 Application area · 4 min

An anode effect vents CF_4 at a slope coefficient of 0.143 kilograms per metric ton of aluminum for each anode-effect minute per cell-day, and C_2F_6 at 0.121 of the CF_4 by mass. The cell makes 2.295 t of aluminum a day, and the GWP-100 values are 6630 for CF_4 and 11 100 for C_2F_6.

(a) Find the CF_4, the C_2F_6 and the CO_2-equivalent of one anode-effect minute. *Checkpoint:* 0.328 kg, 0.040 kg and 2.62 t.

(b) In scenario 5 the automatic kill ends the effect in 4.45 min and the crew, arriving late, in 12.45 min. What does the wait for the crew cost in CO_2-equivalent? *Checkpoint:* About 21 t.

Name one cost of the late kill besides the emissions.

11.11 Application area · 3 min

The model has one well-mixed bath and a constant current efficiency.

(a) Real anode effects start on one anode where the alumina is low. Where in a cell would that be, and what does it mean for the band a plant keeps above the wall? *Checkpoint:* Far from the feeders; a margin above the mean wall.

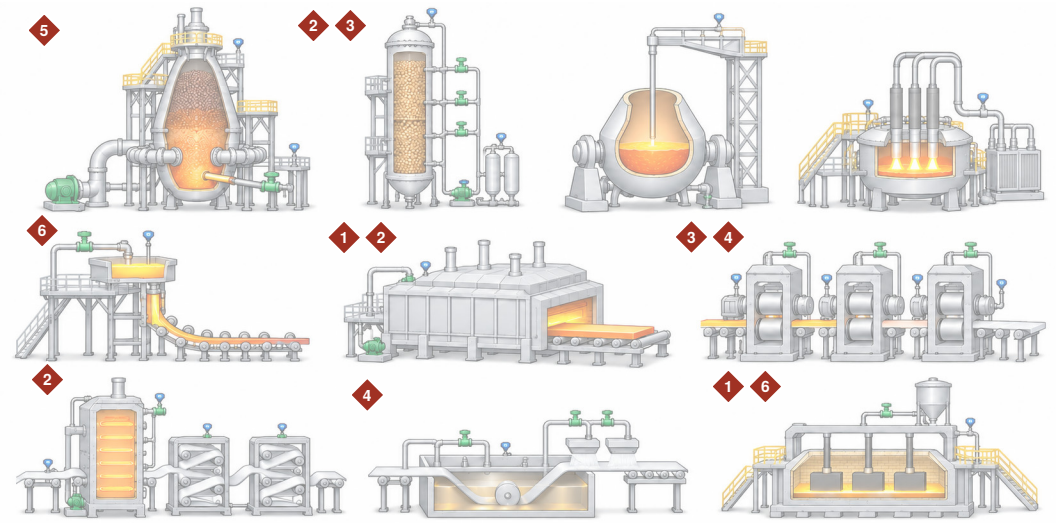
(b) Scenario 8 runs at 12.57 kWh/kg against 13.24 for the reference line of scenario 3. Why might a real cell not show that saving? *Checkpoint:* A squeezed, cold cell would also lose metal.

11.12 Review the agent · 5 min

A learner runs the challenge with the suggested settings: Demand feeding, Slope threshold 0.08 µΩ/h, Overfeed duration 60 min, Heat loop Auto and Anode-change compensation On. It passes. An AI assistant writes: *The heat loop is what saves the bath. Its dips report a superheat of 3.3 K on day 1 and 5.7 K on day 2, so the bath really was cold, and the loop rightly raised the target to 8.86 µΩ. A more accurate liquidus estimate would only make the run pass with more margin.* Find the flaw, and say how to check it in the app.

CHAPTER 12

Summary and Outlook



Open gaps 1 state inside 2 schedules 3 blind side 4 delay as length 5 two clocks 6 diagnosis

Notice that the shaft furnace, the reheat furnace, the strip mill and the smelting cell each carry two gaps, and that no gap marks the oxygen converter or the arc furnace. Illustration, not to scale.

IN PLAIN TERMS **What Ten Metal Plants Taught**

The main idea. Each chapter of this book ran one unit of a steel or aluminum plant in a simulator, from the blast furnace to the coating line. Every unit had a slow part, such as a furnace full of hot metal, and a fast part, such as a valve or a motor.

The challenge. A fast loop can hold its own reading steady while the slow part drifts away.

Why it matters. The units that did best planned the slow part with a model, corrected it from the few true readings, and gave the fast part room to move. Many of the open problems lie between the plan and the loops.

Ten chapters followed the metal from a blast furnace to coated strip and added an aluminum smelting cell. Each joined a slow part to a fast one. Table 1.1 of Chapter 1 listed, for each chapter, what is slow, what is fast and what joins them, and Section 1.4.7 named four ways of joining them: a setpoint computed by a model of the slow part, a feedforward from the fast part's measured changes, an inventory that lets the two run at different rates, and a slow loop that returns a fast actuator to its range. 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.

12.1 What the Chapters Showed

Tables 12.1 and 12.2 compare the chapters. The second column condenses Table 1.1. The other columns name an arrangement that worked, one that failed, and the measured result with the table or section of the chapter that reports it.

Table 12.1 How the Slow and the Fast Parts Were Joined, Chapters 2 to 6

Chapter	Slow part and fast part	What worked	What failed	Measured result
2, blast furnace	Coke after a 5 h descent. Coal within the hour but bounded. The hearth dipped every 2 h.	A slow mid-ranging loop on coke. A bias update of the soft sensor at each cast. A gentler coal loop.	The coal loop alone past its limit. A fast mid-ranging loop. A loop tuned as if the hearth were read continuously.	Last 12 h at 1477 against 1504 °C with mid-ranging (Table 2.4). At $K_v = 1.0$, 7.0 h cold (Table 2.6). Challenge's coldest hearth 1337 against 1459 °C (Table 2.7).
3, hydrogen shaft	Product over a 6 h residence. Top gas within minutes. Hydrogen from a wind-fed buffer.	A hotter bustle under the sticking push. The supervisor on tons. Feedforward on the measured grade.	The supervisor off. Priority on quality through the lull. The soft sensor trusted while the product turns.	Challenge passed on 5 of 5 seeds against 0 of 5 shipped. Supervisor off, buffer dry and worst sample 54.5 to 61.5 % (Table 3.5). Feedforward 94.16 against 92.85 % (Table 3.4).
4, oxygen converter	Lining heat loss drifting heat to heat. One blow of about 21 min and one sublance reading.	Charge model, sublance, the rate endpoint, a hot aim, learning the heat loss between heats.	A fixed recipe. The off-gas integral from the charge. A fixed margin on a wearing lining.	Hits in ten heats 1, 2 and 9 for recipe, static model and rate endpoint. Integral 3 hits, 9 reblows. Wearing lining 5 hits, 7 with learning (Table 4.5).
5, arc furnace	Bath heat and taps in 50 V steps. Three electrodes every 0.02 s on one floating neutral.	Impedance regulation with the staged tap program and foamy slag.	Current regulation, which reads a neighbor's arc as its own error. The top tap on a fresh pile.	Step overshoot 26 to 40 against 9 to 23 % (Table 5.4). Graded heat passed on 19 of 21 seeds against none shipped (Table 5.5).
6, caster mold	Casting speed sets the bulging frequency. Mold level integrates, sensor 0.3 s late.	Diagnosis first, then a canceller at the roll-passing frequency, a knocker and a lower gain.	Any single remedy. Retuning first, which removes the gain that held the bulging down.	Period 12.5 to 10.0 s at 1.5 m/min, canceller to the 0.14 mm floor (Table 6.4). Challenge 3.34 as loaded, 5.31 retuned, 0.77 mm with all three (Table 6.5).

The columns make five points, and Figure 12.1 draws one pair of results from each row.

A fast loop holds what it measures. The most repeated result of the volume is a loop that did its job while the slow part went elsewhere. In Chapter 7 the zone loops held their setpoints within 7 °C through a pace change and the slabs left 55 °C cold. A pyrometer trim held the surface at 1212 °C and left the mean at 1174 to 1177 °C at every trim gain tried (Table 7.5). In Chapter 8 the gaugemeter loop held its own reading on setpoint against thermal crown and left the strip 198.6 µm thin, against 41.9 µm with the loop open. In Chapter 2 a drifting soft sensor read 1514 °C while the hearth sat at 1464 °C (Table 2.4). In Chapter 11 the feeder's tally stayed within 1.8 kg of balance while shots 2.8 % short ran the bath down to an anode effect. Each fast loop worked. Each measured a variable that the slow part could leave behind.

The slow measurement corrects the fast one, and slowly. Where a true measurement exists it arrives late, and its correction is a filter. The hot strip mill's X-ray trim, at roughly an eighth the speed of the loop it corrects, held the crown's peak error to 12.5 µm. Faster trims became the delayed loop again and cycled at up to 201 µm (Section 8.7). The blast furnace's bias update at each cast held the hearth at 1492 °C against 1464 uncorrected, and against a drift of 3 °C per hour a weight of 1.0 beat one of 0.2, 5.1 against 9.6 °C rms, because the drift outweighed the dip noise. The reheat furnace's correction from the pyrometer put all 119 slabs on target where the uncorrected model put 30, and

Table 12.2 How the Slow and the Fast Parts Were Joined, Chapters 7 to 11

Chapter	Slow part and fast part	What worked	What failed	Measured result
7, reheat furnace	A slab 2.5 h inside whose mean is never measured. Mill pace and zone loops.	Slab-tracking setpoints every 120 s, the charge temperature as feedforward, a correction from the pyrometer.	Zone loops on fixed setpoints. A pyrometer trim at any gain. The model uncorrected after the exchange factor fell.	On target 21, 21 and 117 of 117 for zones, pyrometer and model. After the fall, 30 and 119 of 119 uncorrected and corrected (Table 7.4). Challenge 22 against 151 of 159 (Table 7.6).
8, hot strip mill	X-ray gauge 0.5 s late, crown and modulus drifting. Gaugemeter at the stand on a 4 ms scan.	The gaugemeter, an X-ray trim at an eighth of its speed and a synchronous eccentricity estimator.	The same gains on the X-ray gauge. The gaugemeter alone against crown.	Entry wave 11.6 against 173.8 µm rms (Table 8.6). Challenge peak 61.2 against 20.3 µm, prime yield 65.4 against 98.8 % (Table 8.7).
9, annealing line	Furnace with a 240 s lag. Line speed and an 80 s residence. An entry that stops for every weld.	A looper allowed to swing, a speed cap, look-ahead, speed scaled by thickness, firing for the speed.	A tight looper loop. The speed held through a thicker coil.	Off specification 71.3 tight against 0.0 % averaging, 77.5 against 2.5 % on 1.3 mm strip (Table 9.4). Challenge 76.0 against 10.1 % (Table 9.8).
10, galvanizing line	Gauge 100 m downstream. Knife pressure fast, line speed in ramps, a heavy grade's gain 5.6 times.	Speed feedforward. A Smith predictor that keeps its delay in meters of strip.	The predictor with its delay in seconds. A PI tuned for the light grade on the heavy one.	Below the minimum after a slowdown 11.6 against 0.0 % with feedforward. Predictor 40.1 in seconds, 0.0 in meters (Table 10.5). Challenge 50.5 against 0.22 % (Table 10.6).
11, smelting cell	Heat and ledge, one dip a day. Resistance read every second, with a minimum it cannot see past.	Demand feeding on the resistance slope, a heat loop on the daily dip, compensation of anode changes.	Feeding on a tally. A low resistance target left from an energy drive.	Fixed feeding's first anode effect at 32.53 h, none in two days on demand (Table 11.5). Challenge five anode effects as shipped, none suggested (Table 11.6).

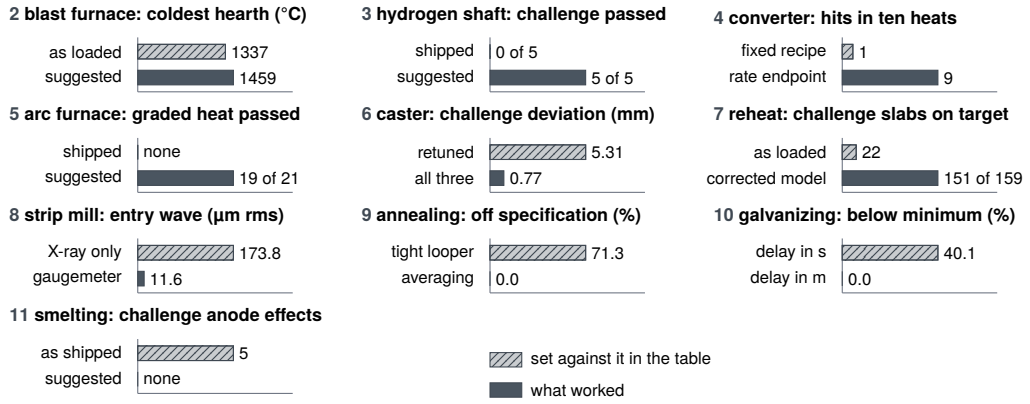


Figure 12.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 12.1 and 12.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 2, 3, 4, 5 and 7, and a smaller one in the others. The blast-furnace panel is the coldest hearth of the challenge, and on a scale from zero its two bars look close. Notice that in most rows the better result came from a change of structure or measurement, not of tuning alone.

with noisier readings the largest correction gain lost the most slabs (Table 7.5). The converter learned its lining's heat loss between heats. A weight of 0.5 restored seven hits, and a weight of 1.0 gave seven at 14.5 against 12.2 °C rms, because each update copied more of one heat's random error into the next plan (Table 4.5). The weight of a slow correction is set by the drift it must follow and the noise it must not pass.

Speed is the fast handle, and an inventory lets it move. The exit temperature of a continuous furnace depends on speed and thickness only through *vd* (Equation (1.12)). The annealing line that held its speed through a change to 1.3 mm strip left 77.5 % of it outside the window, and the line that scaled its speed by thickness left 2.5 %. The looper made the speed change possible, and only when allowed to swing. Held tightly at 75 %, it passed every weld into the speed and put 71.3 % of the strip off specification. Allowed to fall from 96 to 38 %, it held 300 m/min and put none off (Table 9.4). In Chapter 10 a slowdown raised the gain and the delay together by 1.67, and speed feedforward restored the gain, lifting the gain margin from 1.16 to 1.93 against 3.05 at full speed, but not the delay (Table 10.3). In Chapter 6 the speed set the frequency of the disturbance. Bulging moved to a 10.0 s period at 1.5 m/min, exactly roll pitch over speed, and a canceller at that frequency took the level to the noise floor. The hydrogen shaft's buffer stood between the wind and the furnace as the looper stood between the welds and the strip. Through a six-hour lull no setting or schedule among 984 kept both the full tons and every sample at specification (Section 3.6).

A slow, wide actuator returns a fast, limited one to its range. The blast furnace's coal acts within the hour and stops at 200 kg/tHM. Two burden changes asked for about 67 kg/tHM of extra coal against 50 of room, so the coal sat on its limit for about 26 h. A mid-ranging loop moved the coke, which waits five hours, and lifted the last half day from 1477 to 1504 °C for $1.73 against $1.68 million of fuel a day (Table 2.4). It had to be slow. At $K_v = 0.5$ the hearth overshot to 1541 °C, and at 1.0 the loops fought through the coke delay, the coke reached its 420 kg/tHM limit and the hearth spent 7.0 h cold (Table 2.6). Neither arrangement prevented the first dip, to 1448 °C in both, because the burden reached the hearth before either fuel could. The resting point of the fast actuator is an economic choice. A coal target of 190 kg/tHM saved about $33,000 a day against 150, and it left the coal 10 kg/tHM of room where a burden change asked for about 21 (Section 2.5). The arc furnace divides its work the same way, between electrodes regulated every 0.02 s and a transformer tap that moves in 50 V steps.

The structure decided more than the tuning. The clearest measurement is the hot strip mill's challenge. Two configurations on the delayed X-ray gauge differed in every tuning and by almost nothing in result, 61.2 against 60.9 µm of peak deviation, while three that differed in structure spanned 20.3 to 65.9 µm (Table 8.7). The arc furnace showed the same with what a loop measures. Over 21 seeds the impedance regulator with the staged taps passed the graded heat 19 times, and a current regulator passed 16 times only when detuned to $K_I = 10$ (Table 5.5). In Chapter 6 no single remedy met the challenge, and retuning first made the peak worse. On the converter the same off-gas signal gave nine reblows integrated from the charge and nine hits with none read as a rate after the sublance. A tuning study asks how well a structure can do. It cannot rescue a structure that measures the wrong variable.

12.2 Choosing a Method

The chapters suggest an order of questions for a unit that is new to the reader. Each question names the chapter where its answer decided the result, and Figure 12.2 draws the list as a chart.

1. *Which clocks does the unit run on?* Estimate the zone time constant, the product's conduction or residence time and the measurement's delay first (Section 1.4.2, Chapter 7, Chapter 3).

2. *Is the product measured, and where?* If not, plan the slow part with a model and correct it from the slowest true measurement (Chapter 7, Chapter 2). For an inferred measurement, write its error identity, as Equation (1.10) does, and give each disturbance to the fast loop or the slow trim (Chapter 8).
3. *Is the delay a time or a length?* Keep a predictor's delay in meters of strip and feed the speed forward to the gain (Chapter 10). Check what a speed change does to the gain, the delay and the disturbance frequency together (Chapter 6).
4. *Can the fast actuator carry the worst load?* Apply the capacity test of Equation (1.3). If it fails, close a slow loop on the wide actuator, slow in gain as well as integral time against that actuator's delay, and choose the resting point from the fuel price and the disturbances (Chapter 2).
5. *Is the product a batch judged at its end?* Plan from the measured charge, aim to the side that the late correction can reach, and learn a drifting constant with a weight that passes little noise (Chapter 4).
6. *Is the line fed in pieces?* Size the inventory against the longest stop with Equation (1.8), let it swing, cap the speed below the average supply and slow ahead of a stop that the inventory cannot carry (Chapter 9).
7. *Has an oscillation more than one cause, or a reading more than one side?* Diagnose before treating. Open the loop, change the speed and read the signature (Chapter 6). Work on the side of an extremum that the reading can tell (Chapter 11).

A predictive controller paid where it held information the loops did not have. On the hydrogen shaft it planned the feed on the wind forecast. At a 30 % buffer setpoint it made 2225 to 2297 t on five seeds with the buffer at 33.8 % or more, where the override let the buffer fall to 10.4 to 14.5 % (Table 3.5). On the blast furnace it scheduled the coke on its five-hour delay, kept the coal off its limit in scenario 10 and cut the IAE from 1016 to 644 °C h for 395 t more coke over three days. A coal loop retuned in hindsight reached 512 by moving the coal three times as far (Table 2.8). Prediction lost where a simpler structure already carried the timing. Through the reheat furnace's mill delay it put 61 of 83 slabs on target against 80 for the slab model's delay rule, and it failed the challenge at 125 of 159 where the slab model passed at 151 (Table 7.7). On the annealing line it filled the looper before an announced long weld, and from the announcement to one hour it left 14.43 against 13.25 % of the strip off specification, more on all ten seeds (Table 9.6). Before reaching for prediction, try the model of the slow part, the feedforward and the inventory. Volume 9 treats that choice in full.

12.3 Where Practice Is Heading

Hydrogen, and the electricity behind it. A design study of fossil-free steelmaking by direct reduction with hydrogen found that it needs 3.48 MWh of electricity per tonne of liquid steel, mainly for the electrolyzers. It named flexibility in production and in electricity demand as a key feature, with grid balancing through storage of hydrogen and of hot-briquetted iron or through the share of scrap [1]. A study that joined a long-term energy system planning tool to a process model of the same route sized the storage it needs. For the United Kingdom's industry the optimal system held 600 GWh of hydrogen storage, and 140 GWh if a fifth of the electricity came from dispatchable generation [2]. The shaft of Chapter 3 meets this problem over one day. Its shipped buffer could not carry a six-hour lull at full tons and full quality. In the editor's view the size of the buffer and the supervisory layer that trades tons against quality will become part of the control design of such plants, not a fixed input to it.

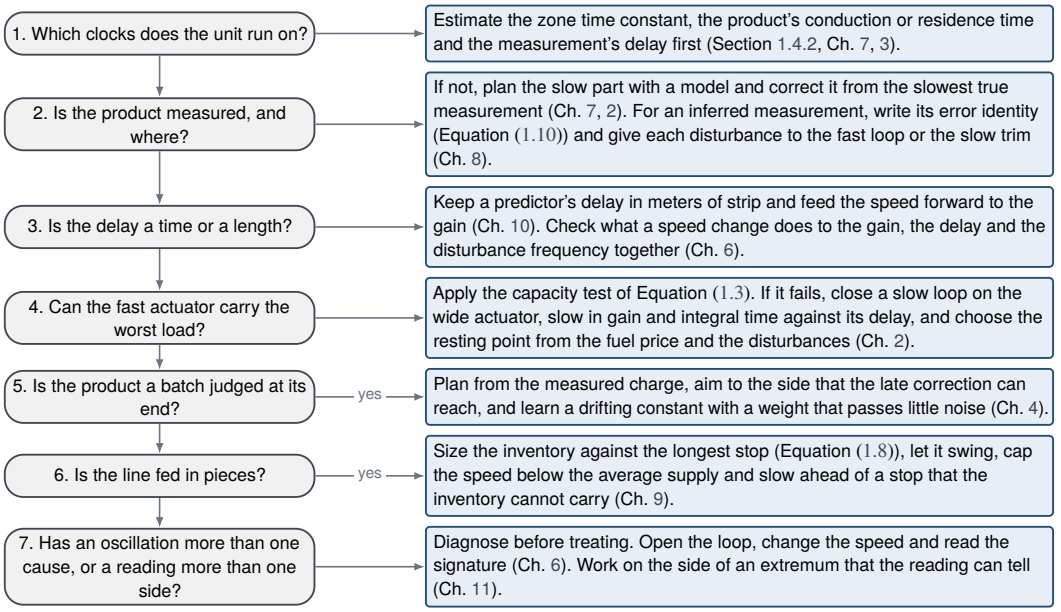


Figure 12.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the first three questions ask about clocks, measurements and delays, and are answered before any loop is tuned.

> ⧖ HISTORY **2004 to 2026: Learning From the Same Readings**
>
> The second history box of Chapter 1 ended in 1984 with a mold level held within 3 mm. Twenty years later a study of blast furnace hot metal temperature still described two or three measurements by cast, taken in the runner at the exit of the tap hole [3]. In 2025 a model predicted the silicon of the next cast of one furnace from the past shift's operating conditions, with an accuracy of 91 % over 200 predictions, because the hostile environment inside precludes a direct temperature measurement [4]. In 2026 a model forecast the mold level of a slab caster from plant data sampled once a second, 2.2 million points over 13 casts [5]. In the editor's reading the models have moved from physics to data, and they still learn from the dip and the level sensor of the earlier boxes.

Electricity prices in the melt shop and the potroom. Production schedules have been optimized together with the purchase of electricity. A continuous-time formulation that adds energy awareness to the daily schedule was applied to the melt shop of a stainless steel plant, and optimizing the schedule jointly with the net electricity cost mitigated the effect of high day-ahead prices [6]. Aluminum smelters have investigated modulating their power to offer services to the grid, driven by energy prices made volatile by renewable sources [7], and shell heat exchangers have been studied as a way to widen the window over which a pot's power can be modulated [8]. The cell of Chapter 11 shows what such modulation asks of control. A resistance target 0.73 μΩ too low ran at 12.57 kWh/kg while the ledge grew to 37.4 cm (Section 11.9). In the editor's view a modulated cell needs an estimate of its heat between the dips, because each change of power is a deliberate upset of the slowest balance in the volume.

The caster and the mill planned together. Planning the continuous caster and the hot strip mill as one problem has been done through "virtual" slabs, possible forms of steel not yet cast, which pass information between the two plans. The approach was solved by mathematical programming with

heuristics and was implemented at several mills [9]. Such a plan sends slabs to the reheat furnace hot. In Chapter 7 the slab model took the charge temperature as a feedforward and burned 0.923 GJ/t against 0.955 on the operator setpoints with 94 against 41 of 95 slabs on target (Table 7.4). In the editor's view the plan and the furnace model will share more, such as the charge temperature of each slab and the length of each announced delay.

Endpoints predicted from data. Data-driven models have been built to predict the endpoint temperature and the phosphorus, manganese, sulfur and carbon of a converter heat. A model on static features was compared with one that adds time series from the blow, which give a better estimate of the endpoint and allow the process to be tuned, and simple feature extraction improved the prediction of phosphorus, manganese and temperature [10]. The static charge model of Chapter 4 hit only two temperatures in ten at 19.0 °C rms, and all ten at 3.4 °C with every heat nominal (Section 4.5), so its misses came from what it was not told. In the editor's view a better prediction will narrow the hot margin, while the one-way correction and the learning between heats will remain.

Seeing inside the vessel and on the strip. New instruments shrink the part of the plant that a model must infer. A European project has studied muon radiography, which uses the muons of cosmic rays that pass through the furnace, to image the cohesive zone of a blast furnace, where the descending burden begins to soften and melt [11]. On the mill, a shape meter that projects a pattern of LED dots onto the strip measured flatness over more than 98.5 % of every coil, and the meter and the flatness control built on it have been implemented in a hot strip finishing mill [12]. In the editor's view the blast furnace of Chapter 2, whose hearth is read only at the cast, has the most to gain from such an instrument.

A mold with a magnetic actuator. A 2025 study superimposed alternating and direct magnetic fields on the molten steel in a slab mold, in simulation and at an industrial caster, and found that surface defects on the sheets could be roughly halved against a direct field alone [13]. The caster of Chapter 6 lists electromagnetic braking among what a fuller model would add (Section 6.8). In the editor's view a field that shapes the flow becomes one more input to the level loop's plant, and the diagnosis of that chapter will have to tell its effect apart from bulging and a sticking gate.

Hydrogen iron in the arc furnace. Iron reduced by hydrogen carries none of the carbon that natural gas leaves in the pellets, so the carbon must be adjusted in the arc furnace to metallize the iron and to foam the slag. A changing and growing slag has been named one of the challenges of melting it, with more refractory dissolution and a higher power demand. Optical emission spectroscopy, proven in pilot and industrial furnaces, was proposed to watch the arc, the residual hydrogen and the slag during melting [14]. In the editor's view the shaft of Chapter 3 and the arc furnace of Chapter 5 then become one control problem, because the product of the first sets the slag under which the electrode regulators of the second work. Figure 12.3 sets developments of this section and of the volume's chapters on three lanes.

12.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [20]. 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.

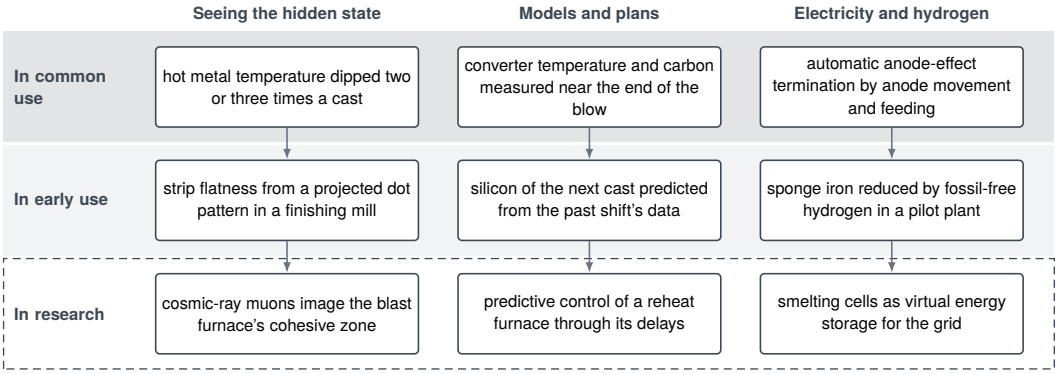


Figure 12.3 Developments in three threads. The top lane holds what the cited source calls standard or common practice, or in use at many plants, the middle lane what it reports in service at some sites or on plant data, and the bottom lane what it reports from simulation, an optimization study or as a project, in the editor's reading. Sources by thread: seeing the hidden state [3, 11, 12], models and plans [4, 15, 16] and electricity and hydrogen [17–19]. Notice that the first thread moves from a reading taken at the cast toward an image of the inside of the vessel.

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 [21]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

These five are the professional form of the habit introduced in Section 1.7: check before trusting, and know enough to check.

12.4.1 How This Volume Was Made

Each of the ten apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace each number to a run. All ten referee reports are self-reviews by the agent or stage that wrote the paper, and three carry a second pass on the predictive-control material added later. Eight of the ten papers were measured again in a consistency round and an app-fix round, and the other two were written after them. Defects went into a log for each app, 105 entries with a severity in all. Of these, 18 are wrong behavior of an app or its model, 65 are documentation, 15 cosmetic or display, and 7 something else, such as a failing test. 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. Five cases show what went wrong and what caught it.

A furnace that skipped its own delay. In the blast furnace app a coke rate, coal rate, blast or burden chosen before Run acted at once. A coke rate of 360 kg/tHM warmed the hearth 5.7 °C in the first hour and 14.7 °C by the fourth, as if the coke had no five-hour descent, and a coal rate showed no inverse dip. The model started each thermal deviation at zero while its delay line held the new inputs for all earlier times, so the furnace began at the base temperature with a history that said the new rates had always been in force. It was found in an app-fix round, when two of the app's own checks failed after a shared engine change made a setting before Run a starting condition. The furnace now starts steady at the settings in force, and no shipped run moved.

A sampling effect that was a tuning effect. Case 3 of the same app was meant to show what the two-hour cast costs a loop tuned for a continuous measurement. Writing the paper, the agent ran the tuning on an ideal continuous measurement, where at the shipped gain of 1.2 the hearth swung from 1401 to 1562 °C too. The case moved to a gain of 0.8, and a consistency round found 0.8 still past the continuous limit: the loop settled at 0.65 but not 0.70 on the continuous measurement and at 0.45 but not 0.50 on the casts. The 48 h minimum-temperature check of the first fix could not show a cycle that grows over a day. The chapter now states both limits and the factor of about 1.4 by which the hold lowers the stable gain.

A span with no length. The hot strip mill's second stand rolled the first stand's present exit, with no transport delay across the span between them. The first stand reads its incoming strip 6.6 m ahead, so a skid band thickened the strip at the second stand 0.66 s before the cold band itself arrived there, and the open-loop skid peak was 119.6 against the 133.5 µm that a steady solve of both stands gives. An agent fixing other defects logged it. The span is now keyed on the delivered coordinate of the steel, so its delay follows every speed change. Every run moved, and the challenge's suggested estimator rate had to fall from 3 to 2/s to pass again. A volume whose lesson is that delay is a distance had a model that left one out.

Controllers that read the plant's constants. In the hydrogen shaft app the predictive controller, the buffer override, the soft sensor's residence filter and the sticking push took the plant's own parameters at each execution. An edit of the bed holdup from 10153 to 12000 kmol moved the

residence filter from 6.03 to 7.12 h, and an edit of the sticking temperature from 950 to 910 °C moved the push's measured margin from 48.8 to 8.8 K, with every measurement unchanged. A truth audit found it. The controllers now read the shipped design values, and a test edits 32 parameters and checks that 15 controller outputs stay fixed at fixed measurements. No number of the paper changed, because no case pins a parameter.

A comparison that rested on its window. The annealing line's predictive controller was compared with the capped looper loop on one scenario. Over the first 2400 s it put 17.46 against 21.60 % of the strip off specification. The referee pass on the new material found that 563 m of the capped loop's off-specification strip came from its speed step at the start of a case that begins at full speed, before any weld was announced. The paper and the chapter now report both windows, and from the announcement to one hour the predictive controller loses on all ten seeds. The same pass left one shipped behavior for the owner. A difficult weld announced during a weld also lengthened the next weld, and correcting it took the suggested challenge configuration from 14.2 to 10.1 % off specification.

The book round. The ten chapter records list 143 rows of numeric verification, each a stated value or group of values set beside its reproduction. Seven records report no drift. The other three each name one printed value that a run does not reproduce, among them an annealing table cell printed 10.297 where the run gave 12.997. The paper and the chapter now print 12.997, and the other two stand as printed. Seven explanations in five chapters were contradicted or found incomplete by runs and were corrected without a new number. The shaft paper blamed the feedforward's shortfall on the heat a richer lot needs, when the cause is the residence. A fixed holdup passes the richer feed's iron faster, so each pellet spends less time under the gas. The paper, the app's explanation and its hints now name the residence. The mill paper said a slowed mill loses margin because the ultimate gain falls with the delay, and a bisection showed that the proportional limit barely moves while the room for integral action shrinks. The other five chapters tested their causal claims by run, and the claims held. Six statements in four chapters named a loop's action opposite to the convention of Section 1.6.2, and one app's model dialog, the galvanizing line's, still does. Chapter 1 was written from the converted chapters, and five of its statements were corrected from the chapter reports. It had said, for one, that speed feedforward removes both effects of a slowdown, when it removes the gain change and not the longer delay.

Every error above was caught by a check that could fail: a step timed against its delay, a comparison on an ideal measurement, a hidden value moved while the measurements were held, a second window, a cause removed and the run repeated. None was caught by rereading the text that contained it. Four working rules follow. Time every delay the model claims, from a start that is already steady. Give a controller only what the plant measures. Report a comparison in every window that changes its sign. Test a cause by removing it.

12.5 Gaps in Metals Process Control

Six gaps belong to this volume's topic, and the gaps common to agent-era practice follow them in Section 12.5.1. In the editor's view none is closed. A sentence on difficulty or progress that carries no citation is the editor's view. Figure 12.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the units in the picture that opens this chapter.

1. *Knowing the product's state inside the furnace.* The reheat model of Chapter 7 read its slabs on target while 30 of 119 were, until a correction from the pyrometer found the exchange factor. The smelting cell's heat loop in Chapter 11 passed its challenge partly because both dips read the superheat low, 3.30 against 8.07 K at the first. Trials with instrumented slabs in a regenerative

reheating furnace found that its total heat exchange factor rises as the fuel consumption or the hot charging temperature falls and as the productivity rises, and offered the result for the dynamic compensation of online control models [22]. In a blast furnace the hostile environment precludes a direct temperature measurement, and the silicon of the hot metal serves as a proxy that arrives with significant delays [4]. The gap is hard because the state that matters lies inside a slab, a hearth or a bath where no instrument lives, and a model's error shows only at the next true reading. An estimate that states its own uncertainty between true readings is not routine.

2. *Schedules that reach the control layer.* The reheat furnace's Level 2 model in Chapter 7 is not told how long a delay will last. The annealing line's predictive controller in Chapter 9, told only what the cap and look-ahead know, put 40.48 against 17.46 % of the strip off specification, and the controller of the shaft in Chapter 3 gained tons because it read the wind forecast. Standard ways to pass delays, weld lengths, coil thicknesses and supply forecasts from planning to control, with a measure of what each is worth, are missing. The Open challenge box after this list takes the gap further.

3. *Inferred measurements that name their blind side.* The gaugemeter's error identity divides every disturbance of the mill in Chapter 8 into one it sees and one it inverts. The shaft's soft sensor in Chapter 3 overstated the product by 2.4 points at hour 8 while the product turned (Table 3.4). In a hot strip finishing mill the gauge cannot be measured at the stand, for lack of space and a hostile environment, so the rolling force is the only measurement of the gauge controller, whose gain is usually set low to stay stable as the mill modulus varies [23]. A review of soft sensors noted that material takes a given run time through a process, so readings taken at the same moment at different places should not be related directly [24]. The gap is hard because an inferred value is checked only against a late or rare true reading, and its blind side shows only when a disturbance of that kind arrives. Most soft sensors in metals plants have no error identity, and a method that derives one from the sensor's model and places the slow trim's crossover from it is not available.

4. *Delays kept in material coordinates.* The galvanizing predictor of Chapter 10 that held its delay in seconds failed once the true delay exceeded its own by about 10 s, and the one that held it in meters carried a 5.6-fold gain as well. The model of Chapter 8 first had no delay across its span. The Smith predictor, one of the most popular dead-time compensators, separates the process model into a fast model free of delay and a delay model [25]. On a galvanizing line the coating gauge stands where 100 to 200 m of strip has left the zinc pot, and the knives are adjusted from its deviation [26]. The gap is hard because most formulations write the delay model as a time, while the plant's delay is a length divided by a speed that the line keeps changing. Controller blocks still take a dead time in seconds, and a block that takes a length and a measured speed would remove a class of error.

5. *Two actuators on two clocks without a fight.* The mid-ranging loop of Chapter 2 at $K_v = 1.0$ fought the coal loop through the coke delay, and the model has no anti-windup coupling between the two loops (Section 2.7). In 1982 one blast furnace held its hot metal temperature by the injection rate of heavy fuel oil [27], and the trend in coal injection rates is toward a long-term average of 200 kg/tHM, 40 % of the total fuel rate [28]. The gap is hard because the fast actuator now carries a large share of the fuel and runs near its limit, while the actuator that relieves it acts hours later. Tuning rules for mid-ranging across a delay of hours, with a guarantee that the slow loop cannot drive the fast one onto a limit, are thin.

6. *Diagnosis before treatment.* The three causes of oscillation in the caster of Chapter 6 were separated only by opening the loop and changing the speed, and each remedy applied to the wrong cause cost gate travel. The smelting cell's reading in Chapter 11 cannot tell on which side of its minimum the alumina lies. A 1997 study of oscillating control loops noted that a full diagnosis typically requires a special test for which the loop is taken out of routine operation [29],

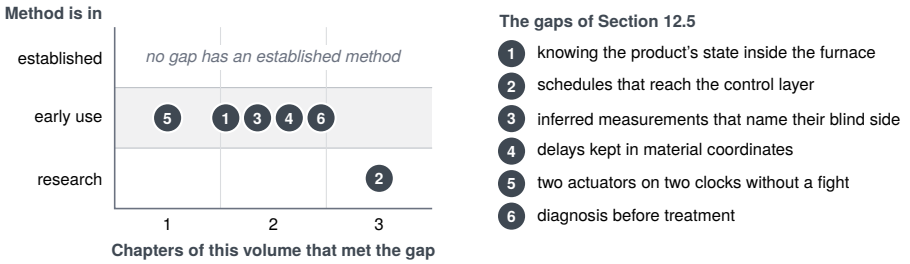


Figure 12.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: two for gap 1 (Chapters 7 and 11), three for gap 2 (3, 7 and 9), two for gap 3 (3 and 8), two for gap 4 (8 and 10), one for gap 5 (2) and two for gap 6 (6 and 11). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the gap most chapters met has its method only in research.

and a study of thin slab casting called mold level hunching from unsteady bulging a main obstacle to higher casting speed, whose origin and detailed mechanisms were not then known clearly [30]. The gap is hard because the causes share one band of frequencies and a running mold cannot be opened at will. The caster chapter lists an automated stiction test for a fuller model (Section 6.8), and a test that names the cause of an oscillation in a running mold is not standard.

OPEN CHALLENGE Schedules That Reach the Loops

Gap 2 is the gap most chapters met. A study of continuous annealing placed the optimization of process settings between the production scheduling layer and the process control layer [31]. A predictive strategy for a reheat furnace handled three typical delays and raised its control accuracy from 74.79 to 99.17 %, in simulation [16]. Hydrogen made from intermittent renewable power meets a steel plant's steady demand with possible mismatches between the two [18]. The gap is hard because a schedule belongs to planners and a loop to control engineers, and neither owns the message between them. Progress would be a standard message that carries delays, weld lengths, thicknesses and forecasts to the control layer, with a test of what each is worth.

12.5.1 Gaps Common to Agent-Era Practice

The final chapter of each volume lists gaps specific to its topic. Seven more are common to every topic in the series. In the editor's view none of them is closed.

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [32]. 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 [21]. 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 [33]. 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 [34]. 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 [35]. 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 holds what it measures. In several chapters the slow part that it did not measure went elsewhere.
- A slow, true measurement corrects a fast or inferred one, and its weight is set by the drift it must follow and the noise it must not pass.
- Speed is the fast handle of a continuous line, and an inventory lets it move.
- A slow, wide actuator returns a fast, limited one to its range, and it must be slow against its own delay.
- The structure of the loops decided more than their tuning.

References

[1] V. Vogl, M. Åhman, and L. J. Nilsson, "Assessment of hydrogen direct reduction for fossil-free steelmaking," *Journal of Cleaner Production*, vol. 203, pp. 736–745, 2018.

[2] A. J. Pimm, T. T. Cockerill, and W. F. Gale, "Energy system requirements of fossil-free steelmaking using hydrogen direct reduction," *Journal of Cleaner Production*, vol. 312, p. 127665, 2021.

[3] J. Jiménez, J. Mochón, J. S. d. Ayala, and F. Obeso, "Blast Furnace Hot Metal Temperature Prediction through Neural Networks-Based Models," *ISIJ International*, vol. 44, no. 3, pp. 573–580, 2004.

[4] O. Raza, N. Walla, T. Okosun, K. Leontaras, J. Entwistle, and C. Zhou, "Prediction of silicon content in a blast furnace via machine learning: A comprehensive processing and modeling pipeline," *Materials*, vol. 18, no. 3, p. 632, 2025.

[5] X. Liu, D. Han, Z. Li, H. Wang, W. Li, M. Fu, and X. Wang, "Time series forecasting model of mold level fluctuations integrating multi-scale feature decomposition and trend flexible alignment," *ISIJ International*, vol. 66, no. 7, pp. 867–875, 2026.

[6] H. Hadera, I. Harjunkoski, G. Sand, I. E. Grossmann, and S. Engell, "Optimization of steel production scheduling with complex time-sensitive electricity cost," *Computers & Chemical Engineering*, vol. 76, pp. 117–136, 2015.

[7] N. Depree, R. Düssel, P. Patel, and T. Reek, "The "virtual battery": Operating an aluminium smelter with flexible energy input," in *Light Metals 2016*. Wiley, 2016, pp. 571–576.

[8] P. Lavoie, S. Namboothiri, M. Dorreen, J. J. J. Chen, D. P. Zeigler, and M. P. Taylor, "Increasing the power modulation window of aluminium smelter pots with shell heat exchanger technology," in *Light Metals 2011*. Wiley, 2011, pp. 367–374.

[9] P. Cowling and W. Rezig, "Integration of continuous caster and hot strip mill planning for steel production," *Journal of Scheduling*, vol. 3, no. 4, pp. 185–208, 2000.

[10] D. A. Sala, A. Jalalvand, A. Van Yperen-De Deyne, and E. Mannens, "Multivariate time series for data-driven endpoint prediction in the basic oxygen furnace," in *Proc. 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2018, pp. 1419–1426.

[11] L. Bonechi *et al.*, "BLEMAB european project: Muon imaging technique applied to blast furnaces," arXiv preprint arXiv:2110.10327, 2022.

[12] Y. Isei, T. Kato, M. Osugi, and T. Ohta, "Development of a shape meter employing the LED dot pattern projection method for a hot strip finishing mill," *ISIJ International*, vol. 60, no. 12, pp. 2876–2885, 2020.

[13] Y. Miki, A. Matsui, H. Ohno, and N. Aramaki, "New electromagnetic flow control system for optimization of molten steel flow in continuous casting mold," *ISIJ International*, vol. 65, no. 8, pp. 1162–1171, 2025.

[14] H. Pauna, T. Kokkonen, P. Cavaliere, M. Bayat, A. Mirowska, T. Alatarvas, M. Huttula, and T. Fabritius, "Hydrogen direct reduced iron melting in an electric arc furnace: Benefits of in situ monitoring," *Journal of Sustainable Metallurgy*, vol. 11, no. 4, pp. 4655–4667, 2025.

[15] T. D. Bradarić, Z. M. Slović, G. J. Stepanoski, and S. Kosanović, "Static-dynamic model for endpoint control used in Smederevo's BOF shop," *Metallurgical and Materials Engineering*, vol. 27, no. 4, pp. 449–456, 2021.

[16] Z. Yang, X. Luo, and J. Qiao, "A combined adaptive entropy-TOPSIS and model predictive control strategy for mixed loading and delay operations in the reheating furnace," *Optimal Control Applications and Methods*, vol. 44, no. 5, pp. 2775–2795, 2023.

[17] Intergovernmental Panel on Climate Change, "Chapter 4: Metal industry emissions," in *2019 Refinement to the 2006 IPCC Guidelines for National Greenhouse Gas Inventories, Volume 3: Industrial Processes and Product Use*. Geneva: IPCC, 2019. [Online]. Available: https://www.ipcc-nggip.iges.or.jp/public/2019rf/pdf/3_Volume3/19R_V3_Ch04_Metal_Industry.pdf

[18] F. Superchi, A. Mati, C. Carcasci, and A. Bianchini, "Techno-economic analysis of wind-powered green hydrogen production to facilitate the decarbonization of hard-to-abate sectors: A case study on steelmaking," *Applied Energy*, vol. 342, p. 121198, 2023.

[19] C.-J. Wong, A. A. Larkin, J. Bao, M. Skyllas-Kazacos, B. J. Welch, N. Ahli, M. Faraj, and M. Mahmoud, "Optimisation of power modulation for Hall-Héroult cells: Process operability and constraints as virtual energy storage," arXiv preprint 2511.07893, 2025.

[20] 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.

[21] 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.

[22] W. Ji, G. Li, L. Wei, and Z. Yi, "Modeling and determination of total heat exchange factor of regenerative reheating furnace based on instrumented slab trials," *Case Studies in Thermal Engineering*, vol. 24, p. 100838, 2021.

[23] Y. K. Lee, Y. J. Jang, and S. W. Kim, "Adaptive feed-forward automatic gauge control in hot strip finishing mill," *ISIJ International*, vol. 47, no. 10, pp. 1444–1451, 2007.

[24] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

[25] J. E. Normey-Rico, T. L. M. Santos, R. C. C. Flesch, and B. C. Torrico, "Control of dead-time process: From the Smith predictor to general multi-input multi-output dead-time compensators," *Frontiers in Control Engineering*, vol. 3, p. 953768, 2022.

[26] H. Wang, K. Zha, K. Yang, W. Zhang, S. Zhu, F. Xu, and Q. Xie, "Data-driven prediction and process parameter analysis of zinc coating thickness in hot-dip galvanizing," *ISIJ International*, vol. 66, no. 7, pp. 857–866, 2026.

[27] M. Hatano, Y. Matoba, K. Otsuka, M. Yoshiki, and T. Miyagi, "Automatic control of hot metal temperature of blast furnace," *Transactions of the Iron and Steel Institute of Japan*, vol. 22, no. 7, pp. 534–542, 1982.

[28] C. Zhou, “Computational fluid dynamics (CFD) modeling for high rate pulverized coal injection (PCI) into the blast furnace,” Purdue University Calumet, final report to the U.S. Department of Energy, Hammond, IN, Tech. Rep., 2008.

[29] N. Thornhill and T. Hägglund, “Detection and diagnosis of oscillation in control loops,” *Control Engineering Practice*, vol. 5, no. 10, pp. 1343–1354, 1997.

[30] U.-S. Yoon, I.-W. Bang, J. H. Rhee, S.-Y. Kim, J.-D. Lee, and K. H. Oh, “Analysis of mold level hunching by unsteady bulging during thin slab casting,” *ISIJ International*, vol. 42, no. 10, pp. 1103–1111, 2002.

[31] Y. Wang, X. Wang, Z. Dong, and Z. Wang, “A Data-Driven Multiobjective Dynamic Robust Modeling and Operation Optimization for Continuous Annealing Production Process,” *ISIJ International*, vol. 60, no. 6, pp. 1225–1236, 2020.

[32] 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.

[33] 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.

[34] 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.

[35] L. Bainbridge, “Ironies of automation,” *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

12.6 Exercises

12.1 Discussion · 8 min

After the pace change of Chapter 7 an agent proposes to double the pyrometer trim gain “so that the slabs leave hot enough”. Using Table 7.5 and Section 1.4.2, explain why this cannot work, which structure does, and what that structure still needs.

12.2 Discussion · 8 min

On the blast furnace of Chapter 2 an agent proposes to raise the mid-ranging gain from 0.2 to 1.0, at the same 12 h integral time, “to return the coal to its target faster after a burden change”. Using Table 2.6 and Section 1.4.3, say what the proposal will do, why, and what you would try instead to reduce the first cold dip.

12.3 Discussion · 8 min

On scenario 8 of Chapter 8, with thermal crown, an agent reports that the gaugemeter loop “held the gauge on setpoint through the whole coil” and recommends removing the X-ray monitor trim to save its maintenance. Using Equation (1.10) and Table 8.7, explain what the report missed and what evidence would settle it.

12.4 Discussion · 8 min

A vendor’s Smith predictor block for the coating loop of Chapter 10 takes its dead time in seconds. An agent sets it to the 40 s of the 150 m/min design speed. The line runs between 90 and 180 m/min. Using Table 10.5 and Section 1.4.5, say what will happen, and what the block needs.

12.5 Discussion · 7 min

For the twelve-heat challenge of Chapter 4 an agent proposes a 50 °C hot aim with no learning, “because learning adds a model that can be wrong, while a margin cannot”. Using Table 4.6 and Section 1.4.4, weigh the proposal.

12.6 Discussion · 9 min

An agent reports that the looper MPC of Chapter 9 cuts the strip off specification in scenario 11 by 19 percent, from 21.60 to 17.46 %, against the capped loop. Using Table 9.6 and Section 12.4.1, list the questions to ask before accepting it, and the check from the record that would have caught its weakness.

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) 12 min. (b) 46.9 min, Bi = 1.88.

1.3 (a) 30 °C, leaving the hearth 15 °C cold. (b) 15 kg/t, with the fast fuel back at its resting point.

1.5 (a) 0.083 %, eight times the half-window. (b) 0.0004 %, 208 times smaller.

1.7 (a) 150 m, 220 s. (b) Yes, the entry supplies 2.66 m/s on average.

1.9 (a) 3.20 mm, 0.20. (b) The strip thickens by 10 µm, and the gaugemeter reads 50 µm below the truth. (c) 3.25 mm.

Chapter 2. Blast Furnace Mid-Ranging of Coal and Coke

2.1 (a). The raceway path of Equation (2.3) has no dead time and a negative gain.

2.3 (c). More coke lets the coal loop take coal out, so the measurement falls when the output rises and the loop has $a = -1$.

2.5 (a) 1.2 and 2.4 h. (b) $K_c \approx 0.26$, $\tau_I = 1.5$ h. (c) $K_c \approx 0.45$, a ratio of about 1.7, where the stability boundaries give about 1.4.

2.7 (a) 1500.0 °C at 6 h and 1519.0 °C at 16 h. (b) 1516.1 °C at 21.0 h, a dip of 3.6 °C. (c) 1532.0 °C. Coke about 2.0 and coal about 1.2 °C per kg/tHM.

2.9 (a) MPC: 375.0 and 157.1 kg/tHM. Coal loop: 357.2 and 198.9 kg/tHM. (b) 192.0 kg/tHM and no time at the limit, against the limit from 21.75 h. (c) The plan starts the coke as soon as its disturbance estimate shows the hearth cooling, for 395 t more coke over three days.

2.11 (a) $0.25 \times 80 = 20\,^{\circ}\mathrm{C}$, offset by $20/1.2 = 16.7\,\mathrm{kg/tHM}$. (b) The coal cools the raceway at once and warms the hearth only after one to three hours, so the first hour dips deeper and the return overshoots. (c) Remove the extra coal ahead of the planned return, by about the coal path's delay and lag.

Chapter 3. Hydrogen Direct Reduction Shaft Furnace

3.1 (d). The balance reads the bed, the product is the bottom cell.

3.3 (b). The bias must describe the model's error, not the bed's transient.

3.5 (a) 0.504, 33.5 %; 0.586, 36.9 %. (b) 4.10 and 3.99. (c) 2.92 and 2.84; the design carries about 1.4 times the minimum.

3.7 (a) About 27.0 then 26.3 %. (b) 99.0 % and 96.6 %. (c) About 0.52 of the move for the product and 1.12 for the soft sensor.

3.9 (a) 2452 t, worst 74.6 %, 9 of 13 short, buffer low 0 %. (b) 2096 t, 97.6 %, none short, 28 %, met. (c) 2128 t with the buffer low at 56 %. The MPC knows the forecast, which the override does not.

3.11 (a) About 953 MJ/t. (b) About 627 MJ/t, about 0.17 MWh/t. Cold DRI can be stored and shipped, hot DRI must be charged soon.

Chapter 4. Basic Oxygen Furnace Endpoint Control

4.1 (c). The end carbon is set mostly by the oxygen blown, and the end temperature by every heat source and sink, including the efficiency, the heat loss and the charge errors that the model does not see.

4.3 (a). The integral starts from 4.5 % carbon on every heat, while the true charge carbon scatters by 0.08 % (1σ), eight times the half-width of the window.

4.5 (a) 0.909 and 47 kg; 0.970 and 149 kg. (b) About 86 s and 1070 Nm^3. (c) About 4.8 t, close to the measured rise of 4.79 t per heat.

4.7 (a) 9, 2, 1, 19.0 °C; 10, 7, 0, 18.2 °C. (b) Heats 1, 5 and 7 at 1627.6, 1615.1 and 1611.5 °C. The sublance trims the hot heats into the window, while the cold ones keep their full error.

4.9 (a) 5 of 10, 23.7 °C. (b) 7 of 10, 12.2 °C, 122 kg, $\times 1.24$. (c) 7 of 10, 14.5 °C, 178 kg, $\times 1.21$. With a weight of 1 the factor is set by the last heat alone, including its random error.

4.11 (a) 2.6 to 3.7 min. (b) Later leaves too little time and oxygen to correct, and the ore added over 60 s might not be in before the stop. Earlier leaves more of the blow unmeasured, so the re-plan's own error grows.

Chapter 5. Electrode Regulation of an Arc Furnace

5.1 (b). The bath floats, so the shorter arc 1 moves its potential and raises the voltage left for the other two phases unequally.

5.3 (d). On tap 4 with arc 1 shorted the hand solve gives 69.3 kA in the short and 76.8 kA in the phase that leads it.

5.5 (a) 5.523 mΩ, 27.67 MW, 70.79 kA. (b) 30.30 MW, 8.7 % below. (c) 104.5 kA. The setpoints lie between the stability limit and I^*, on the long-arc side of the maximum.

5.7 (a) Case 2: overshoot 26 to 40 %, inside from 2.0 to 2.6 s, I3 swing 2.9 to 3.8 kA. Case 3: 9 to 23 %, 0.5 to 0.7 s, 1.2 to 1.9 kA. (b) The current loops take three to four times as long as one loop alone. The impedance loops settle as fast as one.

5.9 (a) Shipped: 411 kWh/t, 40.3 min, wear 0.64, 60 arcs broken (195 s out), not met. Impedance: 412, 38.6 min, 0.63, none, not met (wear). Stage taps: 413, 41.2 min, 0.30 (0.303), 28 broken (85 s out), not met. Both: 414, 40.5 min, 0.23, none, met. (b) See the solution.

5.11 (a) More voltage across the same measured ratio $|V_m|/|I|$ means both a larger voltage and a larger current: the arc lengthens to take the voltage and the power rises. (b) A cave-in can short an arc during a tap step. With the hold blocking the lift the short lasted 2.06 s.

Chapter 6. Caster Mold Level Oscillation Diagnosis

6.1 (b). An oscillation the loop makes stops when the loop is opened. This one grows, so it is a disturbance, and the controller was attenuating it.

6.3 (d). The canceller then takes the RMS to 1.29 mm and the knocker to 0.49 mm.

6.5 (a) 3.25 %, 1.75 %, jump 1.5 %. (b) 0.11 L/s, 0.31 mm/s. (c) 1.9 mm, about 6 s plus the sensor delay.

6.7 (a) 1.11, 0.36, 0.22 mm. (b) 23, 34, 83 %/min (ratios 1.5 and 3.6).

6.9 (a) 3.34 mm. (b) 1.24 mm, pass. (c) 5.31 mm.

6.11 (a) 12.5 and 15 s at 1.2 m/min; 10 and 12 s at 1.5 m/min. (b) Both periods fall by the ratio of the speeds, which marks both as bulging, and the canceller leaves the 0.30 m component.

Chapter 7. Reheat Furnace Slab-Tracking Control

7.1 (a). The zones are held, but the time each slab spends under them falls by a fifth.

7.3 (c). Over the last ten slabs the pyrometer reads 17.0 °C below the predicted surface.

7.5 (a) 351 t/h, 2.5 h, 40 min at 300 s. 438.75 t/h, 2.0 h, 32 min at 240 s. (b) 7800 kg/m^3, 650 J/(kg K), 0.763 GJ/t. (c) 1.21 GJ/t, 63 %.

7.7 (a) 21 of 117, 48.0 °C, 1.134 GJ/t, 53 °C. (b) 117 of 117, 0.8 °C, 1.214 GJ/t, 53 °C. (c) 7.1 % more fuel per tonne. The gap is the same under both strategies, so the faster pace widens it, not the model.

7.9 (a) 854.2 and 1103.6 GJ, soak setpoint 1200 °C, 80 of 83. (b) 852.3 and 1147.6 GJ, 53 of 83. (c) 249.4 GJ with the model and 295.2 GJ with the operator setpoints, 46 GJ less in one hour.

7.11 (a) It lowers the single $\hat{\varphi}$ until the predicted discharge surface matches on average, so the predicted profile inside the furnace is wrong zone by zone. (b) A trial slab instrumented with thermocouples that travels through the furnace.

Chapter 8. Hot Strip Mill Inferential Gauge Control

8.1 (d). By Equation (8.17) the loop that holds $\hat{h}$ on setpoint leaves $h_7 = h_{SP} + c_7$, low by the whole crown.

8.3 (b). The X-ray reading falls when δ rises, so the process gain seen by the trim is negative and the loop is direct acting.

8.5 (a) −95.2 µm. (b) +81.5 µm, thick. (c) +0.0952 mm, and −0.1053 mm for the low calibration.

8.7 (a) 54.6 µm, 10.83 MN, 1.805 mm, gap gain 0.218. (b) −52.6 µm, 13.18 MN, 2.197 mm, gap gain 0.214. (c) 6.00 MN/mm.

8.9 (a) 61.2 µm, 65.4 %, a fail. (b) 20.3 µm, 98.8 %, a pass with 4.7 µm to spare. (c) 28.6 µm, 89.5 %, a fail caused by the runout.

8.11 (a) An elongation of 83.3, a slab of 250 mm. (b) 12.5 s and 17.9 s.

Chapter 9. Annealing Line Looper and Speed Feedforward

9.1 (c). With the level above the setpoint the integral relaxes to the limit and the command stays on v_{max}.

9.3 (a). At 231 m/min the 1.3 mm strip sees the same $v_p d$ as 1.0 mm strip at 300 m/min.

9.5 (a) 336.9 m/min, 276.7 m/min. (b) 268.4 m/min, 1295 m. (c) 237.3 m/min, 230.2 m/min.

9.7 (a) 71.3 %, 287 ± 22.4 m/min, about 224 m/min. (b) 0.0 %, 300 ± 0.0 m/min, 38 to 96 %. (c) 17.5 %, about 269 m/min.

9.9 (a) 1 stop, 30.9 %, 736 to 947 °C. (b) 0 stops, 31.3 %, 749 to 903 °C, 5 to 96 %, 174 m/min. (c) 18.5 %, 732 to 861 °C.

9.11 (a) About 10 s. After it the looper empties and the line stops with the entry. (b) The pyrometer PI on the firing, over several minutes, set by the 240 s furnace lag and the 4 min integral time.

Chapter 10. Galvanizing Line Coating Weight Control

10.1 (c). The tighter tuning leaves almost no gain margin against the 40 s delay, and the loop rings in bursts.

10.3 (a). At 8.91 kPa the gain is -9.76 against -1.74, and the gauge loop's gain margin falls to 0.55.

10.5 (a) $66.2° + 112.1° + 1.7° = 180.0°$. (b) 1.16. (c) 1.93. The phase does not depend on the gain, so the crossover frequency is the same.

10.7 (a) 76.0 ± 1.39 g/m^2, 0.00 %, \$482 an hour. (b) 76.0 ± 4.32 g/m^2, 7.11 %, \$504 an hour. The mean stays at the target, so the zinc used barely moves, but the target the spread allows rises from 73.2 to 80.1 g/m^2.

10.9 (a) 76.0 ± 6.21 g/m^2, 12.46 %, 50 s. (b) 76.0 ± 1.12 g/m^2, 0.00 %, 46 s. The true delay is 10 s longer than the predictor's at 120 m/min and 6 s longer at 130 m/min, either side of the tolerated 9 s.

10.11 (a) The coating must freeze before it can be measured. (b) 24 s and 40 s. The coil steps would be removed sooner and the loop could take more gain, but the delay would still change with speed and the flutter would still be mostly too fast.

Chapter 11. Alumina Feed Control in a Smelting Cell

11.1 (b). The curve has a minimum, so every reading above it has one alumina on each side.

11.3 (d). The loop sees a resistance above target and closes the ACD, just when the cold anode is drawing heat from the bath.

11.5 (a) -14.9, -16.1, $+31.0$ mV per percent. (b) -20.9 and -54.6; -0.150 µΩ per percent. (c) About 0.072 µΩ/h; 0.15 %.

11.7 (a) 32.53 h; −1 kg counted, −164 kg real. (b) 15.82 h, 7 anode effects. (c) None, 3.60 %.

11.9 (a) 2.2 to 10.0 K, 12.57 kWh/kg, 129 kg, 37.4 cm. (b) 1.9 K, 8.20 µΩ, true 5.0 K. (c) 2.9 to 10.0 K, 12.91 kWh/kg, 70 kg. The lowest point comes at 3.9 h, before the first dip is reported.

11.11 (a) Far from the point feeders, so the mean alumina must stay well above the wall. (b) In the model a squeezed or cold cell loses energy but not metal; a real one would lose current efficiency and grow ledge and sludge.

Chapter 12. Summary and Outlook

12.1 A good answer notes that the trim holds the surface at 1212 °C at every gain and leaves the mean at 1174 to 1177 °C, with the same 21 of 117 slabs on target at gains of 0.2, 0.5 and 1.0. The fixed surface-to-mean offset holds only at the nominal pace, because a shorter stay widens the surface-to-core gap from about 19 to 53 °C. The loop controls the wrong variable, and a faster loop reaches the wrong answer sooner. A slab-tracking model that plans the zone setpoints puts all 117 on target, and it needs a correction from the pyrometer when the heat transfer changes, 30 against 119 of 119 slabs.

12.3 A good answer notes that the loop held the inferred gauge, not the strip. Thermal crown is a gap-side offset, which the gaugemeter's error identity reads with the opposite sign, so closing the loop left the strip 198.6 µm thin against 41.9 µm open, 4.7 times worse, with a prime yield of 7.4 % and the roll force at 16.82 against 12.000 MN scheduled. The X-ray trim at an eighth the loop's speed held the peak error to 12.5 µm. In the challenge the configuration without the trim had a peak of 65.9 µm and 29.9 % prime yield against 20.3 and 98.8 %. The evidence is the true gauge from the X-ray or the export, and the force against its schedule.

12.5 A good answer notes that the 50 °C margin gives eight hits and no reblows but 17.4 °C rms, which fails the 15 °C limit, because the lining's drift still lets the late heats finish cold, down to 1605 °C. A larger margin trades rms error for hits. Learning with a 30 °C margin gives eight hits, no reblows and 12.6 °C, a pass with no hit to spare. Learning answers the drift, the margin answers the one-way correction. The concern about learning is fair: its weight sets how much heat-to-heat noise it copies (14.5 °C rms at a weight of 1.0 against 12.2 at 0.5), and its learned factor should be checked against the run.

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.

air knife A slot that blows a sheet of air across the strip leaving the zinc bath and wipes the excess coating back. *Chapter 10*.

alumina feed control Dosing alumina into a smelting cell so that its dissolved share stays in a narrow band that no instrument reads. *Chapter 11*.

anode beam The bar that carries the carbon anodes of a smelting cell and is lowered step by step as they burn away. *Chapter 11*.

anode effect An upset of a smelting cell short of alumina, in which the anodes make perfluorocarbons, strong greenhouse gases. *Chapter 11*.

automatic gauge control (AGC) A loop that moves the roll gap of a rolling stand to hold the thickness of the strip leaving it. *Chapter 8*.

averaging level control Level control tuned to let an inventory swing within its limits so that the flow it passes on stays steady. *Chapter 9*.

back-calculation A way to stop windup: when a controller output is limited, its integral is pulled toward the output actually applied. *Chapter 2*.

basic oxygen furnace (BOF) A tilting vessel, also called a converter, in which oxygen blown onto molten iron burns out carbon to make steel. *Chapter 1*.

bias update A correction added to an inferred value each time a true measurement arrives, so that the estimate follows the plant. *Chapter 2*.

Biot number A ratio that tells how far the surface of a heated body runs ahead of its center. *Chapter 1*.

blast furnace A tall shaft furnace in which burning coke reduces iron ore and melts it into liquid iron called hot metal. *Chapter 1*.

bridle rolls A set of driven rolls that grip the strip and set its speed and tension on a strip line. *Chapter 9*.

burden The layers of ore, pellets and coke charged into the top of a blast furnace. *Chapter 2*.

cast (blast furnace) One draining of liquid iron and slag through the taphole; the hot metal is dipped for temperature once per cast. *Chapter 2*.

charge model A calculation from mass and heat balances that plans the oxygen and coolant of a heat from its measured charge. *Chapter 1*.

coating gauge An instrument past the cooling tower that reads the zinc on the strip many meters after the air knives. *Chapter 10*.

coating weight The mass of zinc per square meter on each side of a galvanized strip. *Chapter 1*.

continuous annealing Heat treating cold-rolled strip as it runs through a furnace, instead of as stacked coils. *Chapter 1*.

continuous casting Pouring steel into a water-cooled mold from which a strand with a solid skin is withdrawn without stopping. *Chapter 6*.

controller action Whether a controller raises or lowers its output when its measurement rises, called reverse or direct acting. *Chapter 1*.

critical carbon The carbon content below which decarburization slows, because carbon rather than oxygen then limits it. *Chapter 1*.

current transformer An instrument transformer that brings a large current down to a small one that an instrument can read safely. *Chapter 5*.

dead time The time between a change at a plant input and the first sign of it at the measurement. *Chapter 6*.

decarburization Removing carbon from molten iron with oxygen, which turns it into carbon monoxide gas. *Chapter 1*.

direct reduced iron (DRI) Iron ore turned to metal by a reducing gas without melting, leaving solid, porous pellets. *Chapter 3*.

electric arc furnace (EAF) A furnace that melts scrap or reduced iron with electric arcs struck from three graphite electrodes. *Chapter 5*.

electrode regulation Moving each electrode of an arc furnace up or down to hold the current or impedance of its arc. *Chapter 5*.

electrolyzer A device that splits water into hydrogen and oxygen with an electric current. *Chapter 3*.

endpoint control Deciding when a batch stops, such as the carbon and temperature at which a converter heat ends. *Chapter 1*.

entry looper A store of strip at the entry of a line that feeds the process while the entry stops to weld coils. *Chapter 1*.

EWMA Exponentially weighted moving average: an update that moves an estimate a fixed fraction of the way toward each new result. *Chapter 1*.

feedforward control Moving a handle as soon as a disturbance or a change of plan is known, before the controlled variable has moved. *Chapter 1.*

final control element The device a controller moves to act on the plant, such as a valve, a burner, a cylinder or a drive. *Chapter 1.*

floating neutral The star point of a three-phase arc circuit, which is not grounded and moves as the arcs change. *Chapter 5.*

foamy slag Slag puffed up by gas so that it covers the arcs and shields the furnace walls from their heat. *Chapter 5.*

galvanizing Coating steel with zinc against rust, here by running strip through a bath of molten zinc. *Chapter 1.*

gaugemeter A strip thickness computed at the stand from the roll gap and the roll force, through the stretch of the mill. *Chapter 1.*

Hall–Héroult cell The electrolytic cell that makes aluminum from alumina dissolved in molten salt, using carbon anodes. *Chapter 11.*

hearth (blast furnace) The bottom of a blast furnace, where liquid iron and slag collect before each cast. *Chapter 2.*

hot charging Charging slabs into a reheat furnace while they are still warm, which saves fuel. *Chapter 7.*

hot metal silicon The silicon dissolved in the iron of a blast furnace, which rises with the furnace heat and is analyzed with each cast. *Chapter 2.*

hot strip mill A train of rolling stands that rolls a reheated slab into strip a few millimeters thick. *Chapter 1.*

hydraulic capsule The hydraulic cylinder that sets the roll gap of a rolling stand, moved by a fast servo valve. *Chapter 8.*

inferential control Controlling a variable that no instrument reads by acting on an estimate computed from what is measured. *Chapter 1.*

integrating process A process, such as the steel level in a mold, that keeps drifting while inflow and outflow differ. *Chapter 6.*

interstand tension The pull in the strip between two rolling stands, held by the looper. *Chapter 8.*

inverse response A response that first moves the wrong way before it settles in the right direction. *Chapter 2.*

Kalman filter An estimator that blends a model's prediction with noisy measurements, each weighted by how much it is trusted. *Chapter 2.*

ledge The frozen layer of electrolyte on the side walls of a smelting cell, which grows or melts with the cell's heat. *Chapter 11.*

load cell A steel element with strain gauges bonded to it that reads the force passing through it. *Chapter 1.*

looper (rolling mill) An arm with a roll between two stands that pushes up on the strip to hold its tension. *Chapter 8.*

material tracking Keeping a record of each slab, coil or length of strip by its position as it moves through a plant. *Chapter 1.*

meniscus The top surface of the liquid steel in a caster mold, whose level the mold loop holds. *Chapter 6.*

metal pad The layer of molten aluminum under the bath of a smelting cell, whose waves add noise to the cell voltage. *Chapter 11.*

metallization The share of the iron in a reduced pellet that has become metal rather than oxide. *Chapter 3.*

mid-ranging control A slow loop that moves a wide actuator so that a fast, limited one returns to the middle of its range. *Chapter 1.*

mill modulus The stiffness of a rolling stand: the force that stretches it by one millimeter. *Chapter 1.*

minimill A steel plant that melts scrap in an electric arc furnace and casts and rolls the steel on one site. *Chapter 5.*

model predictive control (MPC) Control that plans future moves with a model, applies the first move and plans again at the next step. *Chapter 2.*

mold flux A powder layer that floats on the liquid steel in a caster mold; level swings can trap it in the steel. *Chapter 6.*

mold level control Holding the steel level in a caster mold by moving the gate or stopper that feeds it. *Chapter 6.*

monitor gauge control A slow trim from a thickness gauge after the stand that removes the bias of the gaugemeter. *Chapter 1.*

nozzle clogging Solids building up inside the nozzle that feeds a caster mold, which narrows the flow. *Chapter 6.*

off-gas analysis Measuring the flow and makeup of the gas leaving a converter to follow how fast its carbon burns. *Chapter 1.*

one-way correction A late correction that can move a result in one direction only, such as cooling a converter heat that is too hot. *Chapter 1.*

oxygen lance A water-cooled pipe lowered into a converter that blows oxygen onto the bath. *Chapter 4.*

plastic modulus How much the roll force rises for each millimeter the strip is made thinner. *Chapter 1.*

point feeder A feeder that drops a fixed shot of alumina into the bath of a smelting cell each time the controller asks. *Chapter 11.*

power program The schedule of tap and arc settings that an arc furnace steps through as energy per tonne accumulates. *Chapter 5.*

prebake anode A carbon anode baked in a separate furnace before it is hung in a smelting cell. *Chapter 11.*

pseudo-resistance A resistance computed from the voltage and current of a smelting cell, which passes through a minimum as alumina changes. *Chapter 11.*

pulverized coal injection (PCI) Blowing finely ground coal into a blast furnace through its tuyeres to replace part of the coke. *Chapter 2.*

pyrometer An instrument that reads the temperature of a surface from the radiation it gives off, without touching it. *Chapter 1.*

reblow A further blow of oxygen given to a converter heat whose carbon was found above its window. *Chapter 4.*

reheat furnace A furnace that brings slabs back to rolling temperature before the hot strip mill. *Chapter 1.*

relative gain array (RGA) A table that measures how strongly the loops of a plant with several inputs and outputs interact. *Chapter 5.*

run-to-run control Correcting a model or recipe between batches from the result of the batch just finished. *Chapter 1.*

selector A function that passes the higher or the lower of several controller outputs to one actuator. *Chapter 3.*

slide gate A plate valve under a ladle or tundish that slides across the nozzle to set the flow of steel. *Chapter 6.*

Smith predictor A controller that acts on a model of the delay-free response, so dead time drops out while the model is right. *Chapter 10.*

soaking zone The last zone of a reheat furnace, where the slab surface is held while its center catches up. *Chapter 7.*

soft sensor A model that estimates an unmeasured variable from variables that are measured. *Chapter 1.*

sublance A probe dipped into a converter bath near the end of the blow to measure its carbon and temperature. *Chapter 1.*

superheat How far a melt lies above the temperature at which it starts to freeze. *Chapter 11.*

tap changer A switch on a furnace transformer that changes its voltage in steps. *Chapter 5.*

thermal crown The swelling of hot rolls, which changes the gap the strip sees without moving the gap transducer. *Chapter 1.*

thermal inertia The heat a furnace, a slab or a cell stores, which makes its temperature slow to follow a change of firing. *Chapter 1.*

thermocouple A temperature sensor in which two joined metals make a small voltage that grows with temperature. *Chapter 1.*

total heat exchange factor One factor that lumps the radiant exchange between a furnace and a slab, estimated from measured temperatures. *Chapter 7.*

transport delay The time material takes to travel from an actuator to a measurement; on a line, a length divided by the speed. *Chapter 1.*

tundish The vessel between the ladle and the mold that holds steel and feeds the mold through a nozzle. *Chapter 6.*

turndown Tilting the converter after the blow so that the bath can be sampled and measured. *Chapter 4.*

tuyere A water-cooled nozzle through which hot air and injected fuel enter a blast furnace. *Chapter 2.*

unsteady bulging Swelling of the cast strand between support rolls, which pumps the mold level at a period set by roll pitch and speed. *Chapter 6.*

walking-beam furnace A reheat furnace whose moving beams lift the slabs and carry them forward step by step. *Chapter 7.*

weld stop The pause of a line's entry while the tail of one coil is welded to the head of the next. *Chapter 9.*

X-ray gauge An instrument that reads strip thickness from the share of an X-ray beam that passes through the strip. *Chapter 8.*

zinc giveaway Zinc coated above the specified minimum, paid for but not required. *Chapter 10.*

Index of Control Topics

actuator saturation, 35
AI agents
  causal claims, 267
  engineering practice, 265
  open problems, 270
  reproduction, 268
  truth audit, 267
  verification record, 267
  while studying, 18
alumina feed control, 232
anode effect
  low-voltage, 254
  reporting, 253
automatic gauge control, 159
  in service, 181
averaging level control, 187, 188
  setpoint below the swing, 194

back-calculation, 29, 194
batch control, 8, 72
bias update, 34, 59
  weight, 260
Biot number, 140
bumpless transfer, 16, 29, 60, 196

coating weight control, 11, 211
  practice, 226
commissioning
  hot metal temperature control, 41
constraint control, 59
control layers, 13
control loop performance, 117
control structure
  against tuning, 262
controller action, 171
current regulation, 101
current transformer, 109

data export, 17
data-driven models
  history, 264
dead band, 101, 243
dead time
  apparent, 122
  variable, 211
deadbeat gain, 165
decarburization, 8, 76
decentralized control, 92
disturbance cancellation
  known frequency, 123

eccentricity compensation, 162
economics of variance, 211
endpoint control, 8, 72
  measurement practice, 86
EWMA, 8, 75
expert system, 25
extremum seeking control, 234

feedforward
  from the initial state, 74
  measured charge temperature, 148
  speed, 211
feedforward control, 36, 60, 187, 244
  from a schedule, 194
  look-ahead, 194
  model-based, 194
final control element, 12
flow scheme, 4
friction
  stick-slip, 120

gas analysis
  off-gas, 85
gas wiping
  thinning limits, 226
gaugemeter, 10
gaugemeter error identity, 169
Gauss-Newton method, 144

hierarchical control, 136
  Level 1 and Level 2, 142

impedance regulation, 101
infeasibility, 62, 108
inferential control, 10, 26, 49, 159, 238
  non-monotone measurement, 232
  wrong variable, 260
inferential measurement, 160
integrated estimate, drift of, 79
integrating process, 120
  gain, 120

inventory
  surge, 10, 187, 188
inverse response, 27

Kalman filter, 17, 31
  extended, 138
  input disturbance, 196

laboratory measurement, 56
least-squares slope, 243
limit cycle, 124
  from excess loop gain, 125
load cell, 13
loop interaction, 92
looper-gauge interaction, 181

manipulated variable
  as a health indicator, 126
mass-flow control, 162
material tracking, 9, 143
  history, 144
measurement
  history, 3
measurement selection, 172
metallization
  laboratory assay, 66
metals production
  control design procedure, 262
  open problems, 268
mid-ranging control, 6, 23, 262
  loop action, 30
  tuning, 7
minimum-variance assessment, 219
model predictive control, 195
  against simpler control, 263
  in the apps, 16
  move blocking, 31, 60, 195
  offset-free, 31
  preview of a disturbance, 60, 196
  soft constraints, 196
  zone profile, 145
monitor gauge control, 10, 161
moving horizon estimation, 17

offset-free control, 196
one-way correction, 8, 79
oscillation detection, 123
oscillation diagnosis, 116
override control, 48
  external reset, 60
  protective, 102

parameter estimation, 144
periodic disturbance, 120
PID control
  anti-windup, 16, 122, 217
  velocity form, 16, 121, 142, 193, 217, 243
PID tuning
  SIMC rule, 29, 122
process gain
  operating point, 215
  sign reversal, 242
pyrometer, 13

rate limit
  of an actuator, 125, 243
ratio control
  cross-limited, 151
relative gain array, 100
residence time, 52
run-to-run control, 8, 73
  heat-loss learning, 80
  in a steel shop, 80

sampled-data control, 23, 245
  gain ceiling, 32
  hold as dead time, 29
scenarios, 17
  completion rules, 17
scheduling
  and control, 270
scorecard, 17
selector
  low, 60
self-regulation
  absent, 120
simulation engine, 15
simulator framework, 14
slab heating, 5, 140
Smith predictor, 211
  delay mismatch, 212
  history, 212
soft sensor, 11, 24, 48
  bias, 169, 245
  laboratory correction, 59
  maintenance, 66
state estimation
  slab temperature, 143
stiction compensation
  knocker, 122
strip temperature
  modeling, 206

strip tracking, 205
strip vibration, 227
synchronous demodulation, 172

temperature measurement
  immersion thermocouple, 40
thermal inertia, 5
thermocouple, 12
  burnout, 152
time-scale separation, 12, 192, 262
transport delay, 211
  indexed on length, 167
  length-indexed, 9

ultimate gain, 122

valve position control, 6, 24
valve stiction, 118
  field test, 129
  signature plot, 123
  two-parameter model, 120

Index of Applications

air knife, 211
    as final element, 225
aluminum smelting, 232
anode beam, 239
anode change, 239
anode effect, 235
    termination, 245
arc extinction, 97
arc voltage, 97

basic oxygen furnace, 8, 72
    endpoint prediction, 87, 265
    history, 74
bath dip, 245
blast furnace, 6, 23
    muon imaging, 265
bore-in, 104
bridle rolls, 205
burden, 25
burden descent, 27

cast (blast furnace), 25
cell voltage
    measurement, 252
charge model, 74
charge variation, 77
circle diagram, 101
coating gauge, 211
    X-ray fluorescence, 224
coating target, 223
coating weight control
    history, 213
coil shape, 215
compressor trip, 62
continuous annealing, 12, 187
    history, 189
    of sheet, 189
continuous casting, 116
    mold, 118
critical carbon, 76

demand feeding, 244
direct reduced iron, 48
direct reduction, 50
    history, 50
direct-reduced iron, 99
    melting, 111

electric arc furnace, 92
    DRI charge, 64
electrode consumption, 98
electrode mast, 97
    hydraulic, 110
electrode regulation, 92
    history, 94
electrolyzer, 51
    intermittent supply, 67
electromagnetic flow control, 265
entry looper, 10, 187, 190

finishing mill, 159
floating neutral, 93
flow stress, 165
foamy slag, 95
    practice, 110
friction hill, 164

galvanizing, 11
galvanizing line, 211
gaugemeter, 161
    history, 162

Hall–Héroult cell, 232
    history, 234
hearth (blast furnace), 25
heat balance, 76
hot charging, 148, 264
hot metal silicon, 28
    prediction, 42
hot metal temperature, 23
    automatic control, 25
hot stove, 36
hot strip mill, 10, 159
hot-dip galvanizing, 211
hydraulic capsule, 161
    history, 162
    servo valve, 180
hydrogen buffer, 51
hydrogen direct reduction, 263
    arc furnace melting, 265
hydrogen ironmaking
    history, 51

instrument, 12
integrated steel works, 4
interstand tension, 162
iron yield, 76

jet wiping, 213

ledge, 239
level sensor
    eddy-current, 118, 129
line stop, 191
lining wear, 80
liquidus, 238
looper, 162, 262

meniscus, 118
metal pad, 240
metallization, 55
mill delay, 148
mill modulus, 162
mill spring, 161
mill stretch, 163
minimill, 95
mold flux, 118
mold level control, 116
    history, 118
moving bed, 51

nozzle clogging, 121
    open problems, 131

off-gas analysis, 74
oxygen lance, 86
oxygen steelmaking, 72

pellet grade, 58
perfluorocarbons, 235
plastic modulus, 165
point feeder, 235
    as final element, 253
power modulation, 264
power program, 103
prebake anode, 235
prime yield, 173
pseudo-resistance, 233
pulverized coal injection, 24
    history, 26
    metering, 41
pyrometer, 140, 192
    in an annealing furnace, 204
    radiation, 150

raceway, 27
radiant heat transfer, 140
radiant tube furnace, 189
reblow, 80
recycle compressor, 55, 66
refractory wear index, 98
reheat furnace, 5, 136
    electrification, 153
    slab temperature, 137
roll eccentricity, 162
roll force, 164

scrap cave-in, 98
secondary circuit, 97
servo valve, 165
shaft furnace (direct reduction), 48
shrinking-core model, 53
silicon oxidation, 75
skid marks, 166
slide gate, 118
    history, 123
    refractory plates, 129
sludge, 236
soaking zone, 138
steel slab, 136
steel strip, 187
steelmaking, 1
    history, 3
    milestones, 2
sticking (DRI pellets), 55
strip accumulator, 190
strip flutter, 215
strip temperature, 189
sublance, 74
submerged entry nozzle, 118
superheat, 238
Søderberg anode, 235

tap changer, 103
tap-to-tap time, 107
tapping, 239
thermal crown, 166
total heat exchange factor, 139
    open problems, 153
transfer bar, 166
tundish, 118
turndown, 80
tuyere, 25

unsteady bulging, 118
    at high speed, 130

walking-beam furnace, 138
history, 138
weld stop, 190
wüstite equilibrium, 53

X-ray gauge, 161
principle, 179

zinc giveaway, 216

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.