INTERACTIVE CONTROL ENGINEERING

Control in Mining and Mineral Processing

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 15

Control in Mining and Mineral Processing

Variable ore, late assays and equipment at its limit

John D. Hedengren
Editor

Advanced Process Solutions, LLC

Interactive Control Engineering: Principles and Applications
Volume 15: Control in Mining and Mineral Processing
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/v15-mining-minerals/

How to cite this volume. John D. Hedengren, Ed., *Control in Mining and Mineral Processing*, vol. 15 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 ore from the feeders to finished product. Chapters 2 to 6 cover the concentrator: blending on a belt, grinding, classification, flotation and thickening. Chapters 7 to 10 cover leaching, electrorefining, solvent extraction and a lithium refinery. Chapter 11 covers a cement kiln.

The ore changes by the hour, the assay arrives late, and the equipment runs at its limit. Override, cascade and inferential control are the working tools, and each chapter names the technology volume that teaches its main technique.

John D. Hedengren
2026

Contents

Preface v

Notation x

1 Introduction: Variable Ore and Late Assays Along the Circuit 1

1.1 The Question This Volume Answers . . . 1
1.2 A Short History . . . 2
1.3 Notation and Conventions . . . 4
1.4 From Ore to Product . . . 4
1.5 Instruments and Final Elements in This Volume . . . 11
1.6 The Simulator Framework . . . 13
1.7 Working With AI Agents While Studying . . . 17
1.8 Chapter Roadmap . . . 17
1.9 Exercises . . . 20

2 Three Feeders and One Belt Scale: Ore Blending Through Speed-Dependent Delays 22

2.1 Introduction . . . 23
2.2 Belt Scales, Raw-Mix Proportioning and Delayed Feed . . . 24
2.3 The Conveyor as Modeled . . . 25
2.4 Dynamic Character . . . 27
2.5 Implemented Control System . . . 29
2.6 Nine Cases, Measured . . . 32
2.7 Technology in Practice . . . 36
2.8 Limitations . . . 37
2.9 Conclusion . . . 38
2.10 Exercises . . . 40

3 Semi-Autogenous Mill Simulator for Impact-Zone Control Against a Moving Constraint 43

3.1 Introduction . . . 44
3.2 Background . . . 45
3.3 Process Model . . . 46
3.4 Nonlinear Characteristics and Dynamic Response . . . 52
3.5 Controller Design . . . 54
3.6 Results . . . 57
3.7 Technology in Practice . . . 64
3.8 Limitations . . . 66
3.9 Discussion . . . 66
3.10 Conclusion . . . 67
3.11 Exercises . . . 69

4 Hydrocyclone Classification Simulator for Cascade Control and Hidden Loop Coupling 72

4.1 Introduction . . . 73
4.2 Background . . . 74
4.3 Process Model . . . 76
4.4 Controller Design . . . 80
4.5 Results . . . 82
4.6 Technology in Practice . . . 87
4.7 Limitations . . . 89

4.8 Discussion . . . 90
4.9 Conclusion . . . 90
4.10 Exercises . . . 91

5 A Browser-Based Froth Flotation Simulator for Economic Setpoint Selection and Loop Interaction 95
5.1 Introduction . . . 96
5.2 Background . . . 97
5.3 Process Model . . . 98
5.4 Dynamic Analysis . . . 100
5.5 Controller Design . . . 101
5.6 Results . . . 103
5.7 Technology in Practice . . . 109
5.8 Limitations . . . 111
5.9 Discussion . . . 112
5.10 Conclusion . . . 112
5.11 Exercises . . . 114

6 A Mining Thickener Simulator for Teaching Override Selector Control 117
6.1 Introduction . . . 118
6.2 Process Model . . . 120
6.3 Controller Design . . . 123
6.4 Results . . . 127
6.5 Technology in Practice . . . 130
6.6 Limitations . . . 132
6.7 Discussion . . . 133
6.8 Conclusion . . . 133
6.9 Exercises . . . 135

7 Soft Sensing and Inferential Control in a Gold Carbon-in-Leach Simulator 138
7.1 Introduction . . . 139
7.2 Background . . . 140
7.3 Process Model . . . 141
7.4 Effect of the Reagent Setpoints on Leach Performance . . . 143
7.5 Controller Design . . . 144
7.6 Results . . . 148
7.7 Technology in Practice . . . 158
7.8 Limitations . . . 160
7.9 Discussion . . . 161
7.10 Conclusion . . . 161
7.11 Exercises . . . 164

8 A Copper Electrorefining Simulator for Inferential Constraint Control 167
8.1 Introduction . . . 168
8.2 Background . . . 169
8.3 Process Model . . . 169
8.4 Controller Design . . . 172
8.5 Results . . . 173
8.6 Technology in Practice . . . 179
8.7 Limitations . . . 181
8.8 Discussion . . . 181
8.9 Conclusion . . . 182

8.10 Exercises . . . 184

9 Two Time Scales and a Late Laboratory in Rare Earth Solvent Extraction 187
9.1 Introduction . . . 188
9.2 From Ion Exchange to Counter-Current Cascades . . . 189
9.3 The Cascade as Modeled . . . 190
9.4 One Stage, Seventy Stages and Two Speeds . . . 194
9.5 Control as Implemented . . . 197
9.6 Ten Cases on the Running Cascade . . . 199
9.7 Technology in Practice . . . 203
9.8 Limitations . . . 205
9.9 Conclusion . . . 206
9.10 Exercises . . . 207

10 A Lithium Hydroxide Refinery Simulator for Interacting Loops on Three Units in Series 211
10.1 Introduction . . . 212
10.2 Background . . . 213
10.3 Process Model . . . 214
10.4 Dynamic Analysis . . . 217
10.5 Controller Design . . . 218
10.6 Results . . . 219
10.7 Technology in Practice . . . 224
10.8 Limitations . . . 226
10.9 Discussion . . . 227
10.10 Conclusion . . . 228
10.11 Exercises . . . 229

11 Rotary Cement Kiln Simulator for Ratio Feedforward and Dead-Time-Limited Quality Control 232
11.1 Introduction . . . 233
11.2 Background . . . 234
11.3 Process Model . . . 235
11.4 Controller Design . . . 240
11.5 Results . . . 244
11.6 Technology in Practice . . . 251
11.7 Limitations . . . 253
11.8 Discussion . . . 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 . . . 266
12.5 Gaps in Mineral Process Control . . . 268
12.6 Exercises . . . 272

Answers to Selected Exercises 274

Glossary 278

Index of Control Topics 281

Index of Applications **284**

About the Editor **286**

Volumes in This Series **287**

Notation

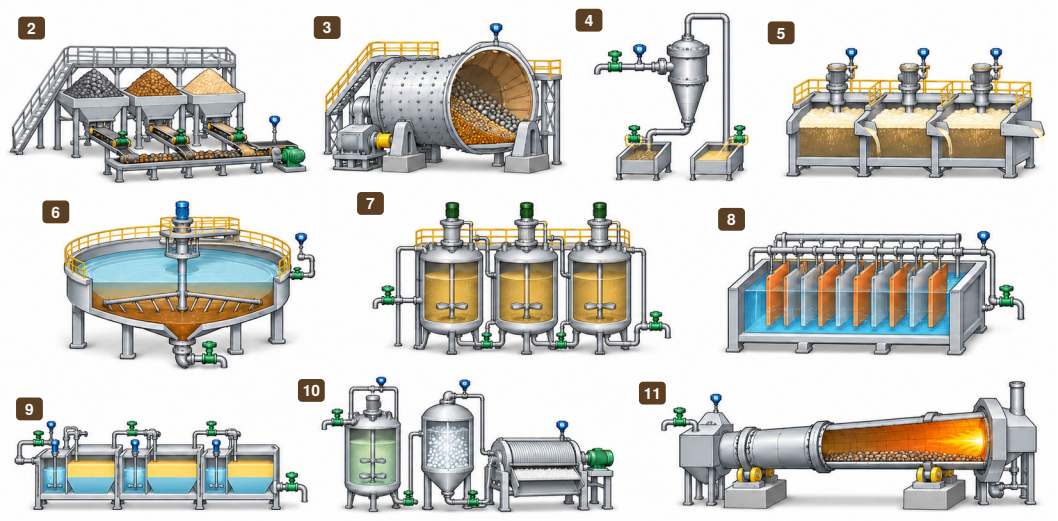
The chapters share the symbols below. A symbol used by one chapter only is defined where it first appears, and each chapter states its own units.

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: Variable Ore and Late Assays Along the Circuit



The chapters 2 feeders and belt 3 semi-autogenous mill 4 hydrocyclone 5 flotation cells 6 tailings thickener 7 carbon-in-leach tanks 8 electrorefining cell 9 solvent extraction 10 lithium refinery 11 cement kiln

Notice that the top row grinds, classifies and floats the ore, the middle row thickens the tailings, leaches gold and refines copper, and the bottom row separates rare earths, refines lithium and burns cement. Illustration, not to scale.

IN PLAIN TERMS **Running a Mineral Plant on Late Answers**

The main idea. A mineral plant grinds rock to powder and separates the small share that holds the metal from the waste. Some plants then dissolve the metal with chemicals and recover it in pure form.

The challenge. The rock changes from hour to hour in hardness and in metal content, and nothing measures that as it arrives. A laboratory often tells how the plant did. Its answer comes hours after the sample was taken.

Why it matters. Good control keeps metal coming from changing rock while it waits for those late answers. It must not break the machines that do the work.

THREE WAYS TO READ THIS CHAPTER

Explore	8 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, 1.4.2 and 1.4.5, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, working each number of the example C-1, then Exercises 1 to 10.
Practice	10 min	This page, Section 1.4.6, Section 1.4.7 with Table 1.1, Section 1.5 with Figures 1.7 and 1.9, and Section 1.8.

1.1 The Question This Volume Answers

A mine sends a concentrator rock whose hardness, grade and mineralogy change from one hour to the next, and nothing weighs those properties as the rock arrives. The plant grinds it, separates the valuable mineral from the waste, dewaters what is left and extracts the metal by chemistry, in units that hold minutes, hours or weeks of material. The product is judged by assays, chemical analyses of samples that a laboratory reports hours after the sample was cut, or by on-stream analyzers that

cycle through a few streams in minutes. This volume answers one question: how is a circuit of this kind controlled when its feed varies without being measured and its quality is known late, and which techniques recur from the belt to the cathode?

Seven control problems define the industry, and every chapter meets several of them [1, 2].

- *The ore changes and is not measured.* A belt scale weighs the total tonnage and nothing else. Hardness, grade and leachability show up later as a change in the plant's behavior.
- *Quality is measured late and on a sample.* A tailings or residue assay describes the plant of several hours before, and some variables, such as a circulating load or the gold left in a leach tail, have no instrument at all.
- *Recycle stretches and amplifies.* A mill closed by classifiers returns most of its coarse material, so a small change in breakage becomes a large change in the load the circuit carries.
- *Inventories are large and slow.* A thickener bed turns over in hours, a leach train holds most of a day and a solvent extraction cascade takes weeks to move its inventory.
- *Equipment limits bind and move.* Mill load, rake torque, cyclone roping, belt rating and current density are hard limits, and several of them move with the ore.
- *The objective is a trade, not a setpoint.* Grade against recovery, density against torque and purity against recovery each have an economic optimum that is often flat and must be chosen.
- *Inferred variables look right while they are wrong.* A mill toe computed from the charge, a soft sensor and a limiting current density come from models, and a loop that holds them can hold the plant in the wrong place.

Ten chapters answer the question on ten plants, each with an app that the reader runs in a browser. Chapters 2 to 6 follow the concentrator: feeders and a belt scale, a semi-autogenous mill, hydrocyclone classification, froth flotation and a tailings thickener. Chapters 7 to 9 extract and refine: gold by carbon-in-leach, copper by electrorefining and rare earths by solvent extraction. Chapters 10 and 11 add two pyroprocessing plants, a lithium hydroxide refinery with a rotary calciner and a rotary cement kiln.

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

- Write the mass balance of a closed grinding circuit and say how a change of ore or of cut size moves its circulating load and its product.
- Compute a recovery and a circulating load from assays, and say which assay error matters most.
- Choose an operating point on a grade and recovery trade from an economic margin, and say how flat the optimum is.
- Say how old an assay is when it is used, and design the bias update that joins it to a soft sensor.
- Wire an override selector in the safe direction, say which disturbances a cascade removes, and size the margin an inferred limit needs.

This chapter develops each of these on one small example and names the chapter that uses it. PID forms and tuning are taken as known from Volume 2, ratio, feedforward and cascade from Volume 3, selectors and override from Volume 4, dead time, soft sensors and estimation from Volume 5, interaction and the relative gain array from Volume 6 and predictive and economic control from Volume 9.

1.2 A Short History

The milestones of Figure 1.1 follow the plants of this volume, from copper refined by current to the first look inside a turning mill.

Copper came first. It could be refined by electrolysis on a commercial scale only once the dynamo made large currents cheap, and the earliest serious attempt was made under a first patent dated 1865. A hundred or more troughs were coupled in series, and an encyclopedia of 1910 judged the process identical in principle with the practice of its own day [3]. A patent filed in 1890 whirled water under

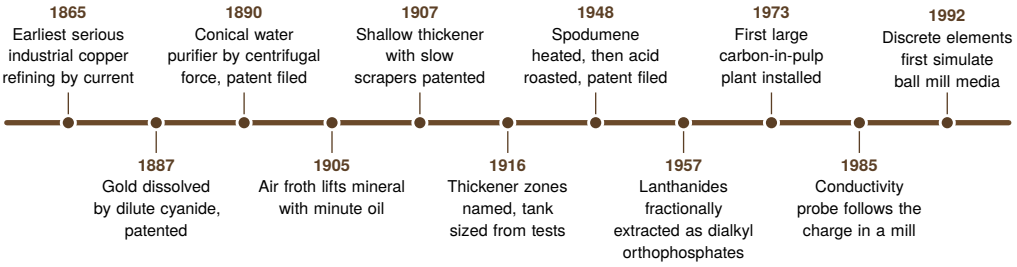


Figure 1.1 Milestones of mineral processing and its measurement. The dates are from [3] (1865), [4] (1887), [5] (1890), [6] (1905), [7] (1907), [8] (1916), [9] (1948), [10] (1957), [11] (1973), [12] (1985) and [13] (1992). A United States patent is dated by its first filing. Notice that every unit of the volume was in service long before an instrument could see inside it. Events are evenly spaced, not to a time scale.

pressure into an inverted cone, so that impurities driven to the wall left through an opening at its point [5]. That opening is the apex of Chapter 4.

> ⧖ HISTORY **1887 to 1916: Wet Methods Join the Mill**
>
> A British patent of 1887 dissolved gold from powdered ore with a cyanide solution kept extremely dilute, because then it takes the gold and silver in preference to the base metals [4]. Twenty years later a patent for the slimes of the cyanide process described a shallow tank, 35 feet across in its example, in which scrapers on slowly turning arms moved the thickened solids gently to a central discharge [7]. Flotation came between them. By March 1905 three metallurgists had cut the oil to a fraction of one percent of the ore and beat air into the pulp, and a froth rose carrying 70 to 80 percent of the mineral. By 1916 it was in extensive use in the United States, Australia, Sweden, Chile and Cuba [6]. In the same year the continuous thickener was divided into four zones and sized from laboratory settling tests [8].

Cement changed its kiln in the same years. Rotating steel cylinders replaced the fixed kilns, petroleum fired them in the United States from about 1895, and by 1910 powdered coal blown in by air was almost universal [14]. The middle of the century added chemistry and theory. A patent filed in 1948 heated spodumene to make it leachable before roasting it with acid [9], the route of Chapter 10. The rare earths were extracted as their dialkyl orthophosphates in 1957 [10], and the low throughput and rising cost of ion exchange later moved their industrial separation to solvent extraction [15].

> ⧖ HISTORY **1973 to 2012: Measuring What the Plant Hides**
>
> In 1973 an analysis of conveyor belt weighers found limits on their performance that still kept them from their widest use, and placed the most serious errors in the element that measures the weigh length [16]. In 1985 a conductivity probe set through the liner bolts of a laboratory mill followed the toe and the shoulder of the tumbling charge. In 1999 and 2000 accelerometers on the shell of an industrial semi-autogenous mill gave a vibration signal sensitive to its operating conditions, tested against speed, feed rate, ball load and discharge density [12]. Work began in 2004 to extend an online free cyanide analyzer to the cyanide bound to metals, and by 2012 its readings from 5 to 155 parts per million agreed with laboratory methods within 10 percent across oxide, sulfide and copper-rich ores [17].

Gold followed. In 1950 an elution process let loaded carbon be stripped and reused, and the first large carbon-in-pulp plant was installed in South Dakota in 1973 [11]. In 1992 the discrete element method first simulated the media of a ball mill, though by 2017 the method had not matured enough to predict a mill's product size and throughput [13].

1.3 Notation and Conventions

Solids flows are in tonnes per hour and assays in percent of the element by mass. In the grinding circuit F is the fresh feed, φ its coarse share, b the share of the coarse that the mill breaks on one pass, E_c and E_f the shares of the coarse and the fine that the cyclones send to the underflow U, and g the gain of the recycle. In flotation f, c and t are the copper assays of feed, concentrate and tailings, and R the recovery. A leach train has N tanks of residence τ_t and a rate constant k. An assay is sampled every T_s and reported θ later, and λ is the weight of a bias update. The framework's h is its integration step and its a the controller action. Neither is used for anything else in this chapter; Chapter 8 uses a for the nodule height and writes the action as a sign. 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 loop that adds steam when a leach residue assay is high is direct acting, because more heat lowers the residue.

The chapters keep their own symbols, and several collide. β is the neodymium and praseodymium separation factor in Chapter 9, the bias update gain in Chapter 7 and the pressure exponent of the cut size in Chapter 4. R_f is the froth recovery in Chapter 5 and the water recovery to underflow in Chapter 4, and R is also the air to fuel ratio in Chapter 11 and the gas constant in Chapter 10. In Chapter 4, g is the share of the coarse product ground per pass and g_r the grinding effectiveness. σ is the roping factor in Chapter 4, the ore settleability in Chapter 6 and the loop sign in Chapter 9. δ is a boundary layer in Chapter 8, an apex exponent in Chapter 4 and an impact offset in Chapter 3. τ_c is the compression time constant in Chapter 6 and a closed-loop time constant elsewhere, and u is the current density utilization in Chapter 8. In Chapter 8, g, θ, β and σv have that chapter's own meanings, defined where they first appear. In Chapter 5, g is the concentrate grade (the c of this chapter), θ_c the collector coverage, T_s the predictive controller's decision period and R also its input weight. In Chapter 11, λ is the excess air ratio (the bias update weight here), β the burnability of the raw mix and N the kiln speed.

1.4 From Ore to Product

1.4.1 The Flowsheet

Figure 1.2 places each chapter's plant on the path from ore to product [18]. Feeders draw ore from stockpiles onto a conveyor that carries it to a semi-autogenous mill, and hydrocyclones split the ground product into a fine overflow that goes on and a coarse underflow that returns for more grinding. Comminution uses more of a concentrator's energy than any other step [19]. A copper ore then goes to flotation, which lifts the sulfide mineral into a froth concentrate, and the concentrate goes to a smelter whose anodes are refined electrolytically. A gold ore goes from the cyclones to a carbon-in-leach train. Both tailings streams go to a thickener that returns water to the mill. Other plants start from a concentrate made elsewhere: a rare earth leach liquor that solvent extraction separates, and a spodumene concentrate that a lithium refinery calcines, leaches and crystallizes. The cement kiln shares two things with them, a raw mix blended from feeders and a rotary kiln.

> **+ DID YOU KNOW? Billions of Tonnes a Year**
>
> A report of 2005 put the ore treated by flotation at nearly 2 billion tonnes a year worldwide [20]. The world made an estimated 4.0 billion metric tons of cement in 2024 [21], each tonne of it burned in a kiln like the one of Chapter 11.

The techniques the chapters share are developed below on one running example, a copper circuit called C-1, drawn in Figure 1.3. A feeder puts 450 t/h of ore on a belt 150 m ahead of a ball mill. The mill is closed by cyclones, and their overflow feeds a rougher flotation bank of six cells. The tailings go to a thickener, and the concentrate is leached in four tanks in series. Solvent extraction transfers the dissolved copper to a clean electrolyte, and an electrowinning tankhouse plates it as cathode. C-1

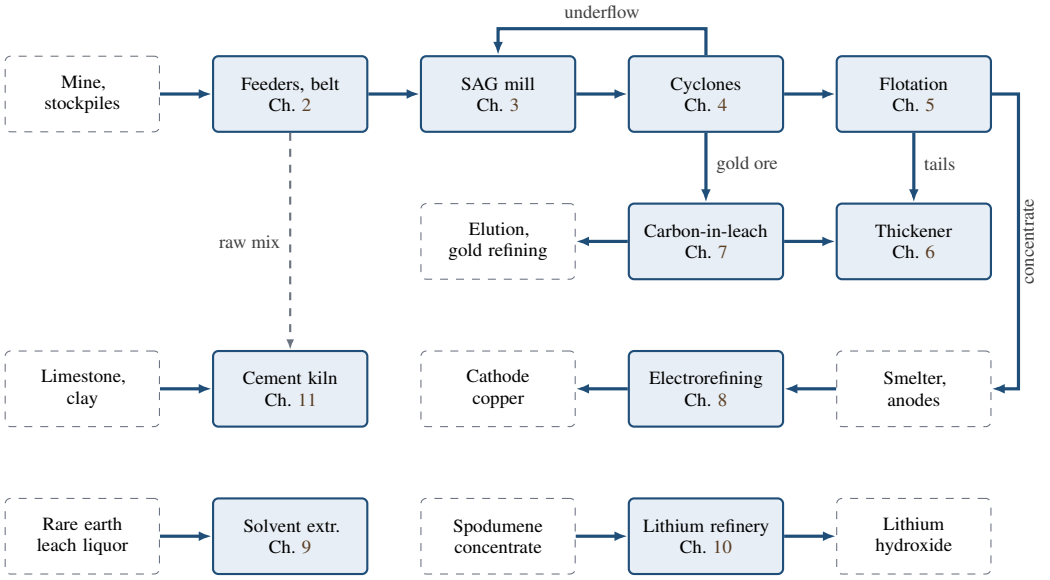


Figure 1.2 The Plants of the Volume on the Path From Ore to Product. Filled Blocks Are the Plants of Chapters 2 to 11, Dashed Blocks Lie Outside the Volume. The Top Two Rows Are a Concentrator, With Thickener Water Returning to the Mill (Not Drawn). The Dashed Arrow Marks That Feeders and a Belt Scale Also Build a Cement Raw Mix.

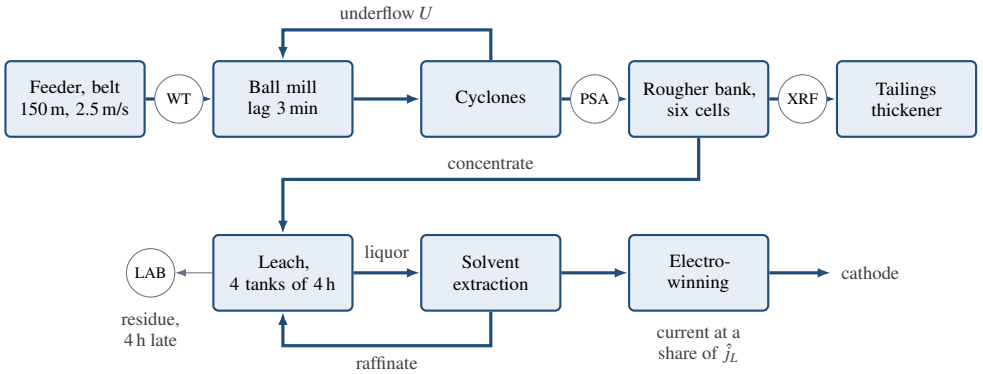


Figure 1.3 The Running Example C-1. Circles Mark the Measurements on the Product Path: the Belt Scale (WT), a Particle Size Analyzer (PSA) on the Cyclone Overflow, an On-Stream X-Ray Fluorescence Analyzer (XRF) on the Flotation Streams, and a Laboratory Assay (LAB) of the Leach Residue, Sampled Every 4 h and Reported 4 h Later. The Tankhouse Runs at a Share of a Limiting Current Density $\hat{j}_L$ That Is Computed, Not Measured.

leaches its concentrate on site instead of selling it to a smelter, a route some plants take, so that one example runs from ore to cathode [22]. Every number of C-1 is declared here or computed from these declarations by the script kept with this chapter's record. The numbers are chosen for clear arithmetic and are not taken from any one chapter.

1.4.2 Feed, Comminution and the Circulating Load

Ore on a belt is a delay set by distance, not by time. C-1's belt carries $\lambda = 50\,\mathrm{kg/m}$ at $v = 2.5\,\mathrm{m/s}$, and a belt scale reports the product $3.6\,v\lambda = 450\,\mathrm{t/h}$. Ore placed 150 m upstream arrives after $L/v = 60\,\mathrm{s}$.

If the belt slows to 2.0 m/s, the scale reads 360 t/h at once although no feeder has moved, and the delay becomes 75 s. Chapter 2 measures both effects on three feeders and one scale, where slowing the loaded belt from 2.0 to 1.5 m/s cuts the reading from 900 to 675 t/h. It also shows the deeper limit. A scale weighs the total, so it cannot say which feeder is short, and a total on its setpoint does not prove the blend. The blend decides how hard the ore is, and the mill feels that before any instrument.

C-1's mill is closed by cyclones, and two size classes are enough to see why the circuit behaves as it does. Call material coarser than the product size coarse. The fresh ore is $\varphi = 0.80$ coarse, the mill breaks $b = 1/3$ of the coarse it receives on each pass, and the cyclones send $E_c = 0.90$ of the coarse and $E_f = 0.25$ of the fine to the underflow, which returns to the mill [18, 23]. The coarse that reaches the mill is the fresh coarse plus what the cyclones return, so a balance at the mill feed gives Equation (1.1).

$$M_c = \frac{\varphi F}{1-g}, \qquad g = E_c\,(1-b), \qquad \frac{O_c}{F} = \frac{\varphi\,(1-E_c)(1-b)}{1-g} \tag{1.1}$$

Here M_c is the coarse in the mill feed and O_c the coarse that escapes in the overflow. The recycle gain is $g = 0.60$, so the mill sees 2.5 times the fresh coarse, 900 t/h, and the product carries 60 t/h of coarse, 13.3 % of it. The underflow returns 540 t/h of coarse and 130 t/h of fine, a circulating load of $U/F = 149\,\%$, and the mill passes 1120 t/h. The recycle also stretches time. If the mill delays its discharge by a first-order lag of $\tau_m = 3$ min, the coarse in the circuit answers a change with the time constant $\tau_m/(1-g) = 7.5$ min, and a simulation of the loop returns the same value.

> **+ DID YOU KNOW? One Percent of the Energy**
>
> A 2006 report for the United States government put the energy that semi-autogenous and ball mills spend on size reduction at about 99 trillion British thermal units a year. Grinding in such mills uses only about one percent of the input energy, and the mills also consume tons of steel balls and liners [24].

A harder ore arrives. With $b = 1/4$ the recycle gain rises to 0.675, the circulating load to 193 % and the product coarse to 18.5 %, while the belt scale reads 450 t/h throughout. Figure 1.4 shows the choices that remain. A finer cut, $E_c = 0.933$, restores the product, but the circulating load climbs to 216 % and the mill must pass 1420 t/h. If the mill can take 1300 t/h, the circuit can hold the product only by cutting the fresh feed to 412 t/h, 8.5 % less. Keeping 450 t/h at the mill limit leaves 19.5 % coarse in the product. Fineness, throughput and the mill limit cannot all be held, and which one gives way is a control decision.

No instrument measures a circulating load. Plants compute it from size assays of three streams around the cyclones. If d, u and o are the coarse shares of the cyclone feed, underflow and overflow, a balance of total solids and of coarse at the cyclones gives Equation (1.2).

$$\frac{U}{O} = \frac{d-o}{u-d} \tag{1.2}$$

C-1's streams assay 0.536, 0.806 and 0.133, which return 149 %. The difference in the denominator is small, so a feed assay 0.01 high reports 159 %, an error of 6.4 % from one point of one assay.

Chapter 4 builds its circuit on the same recycle, with a coarse recirculation gain of $0.918 \times 0.68 = 0.624$, and turns a 3 min transport lag into a quarter hour of circuit dynamics. When grinding weakens there, the circulating load climbs from 288 to 386 % while the level and pressure loops hold their setpoints. Chapter 3 adds the limit that moves. Its power has an interior maximum in the mill filling, so past it more feed returns less grinding, and the toe of the charge that the media should strike moves with the ore, the ball charge and the feed [25]. Chapter 10 shows a quieter form of the same lesson: a lower ore grade returns every controlled variable to its setpoint while the plant's output falls by about a quarter.

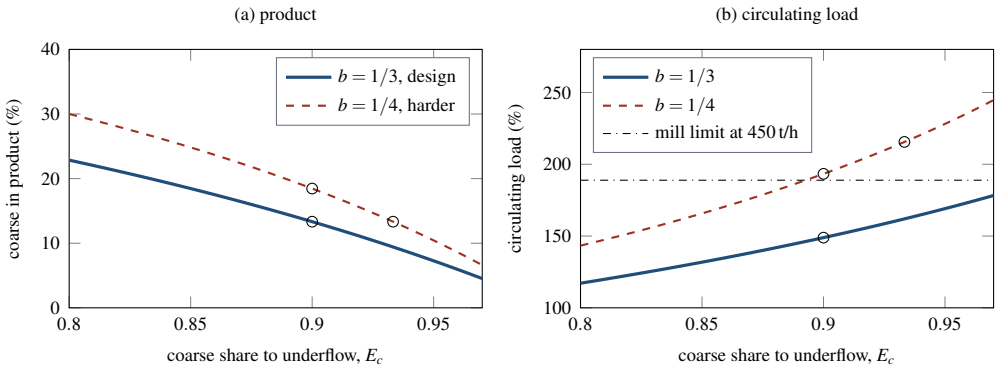


Figure 1.4 C-1's Closed Grinding Circuit at 450 t/h of Fresh Ore. (a) The Coarse Share of the Product and (b) the Circulating Load Against the Share of Coarse the Cyclones Return, for the Design Ore (Thick Solid) and a Harder Ore (Dashed). Circles Mark the Design, the Harder Ore at the Design Cut, and the Harder Ore With the Product Restored. The Dash-Dotted Line Is the Circulating Load at Which the Mill Reaches 1300 t/h.

1.4.3 Flotation and Thickening

C-1's cyclone overflow carries 1.00 % copper as chalcopyrite, which holds 34.6 % copper, so 2.89 % of the ore is the mineral. Each of the six cells is treated as a perfectly mixed tank in which the mineral floats at a first-order rate and the gangue reports to the froth by a much slower rate, mostly by entrainment. For a bank of n equal cells the recovery of each class is Equation (1.3).

$$R = 1 - \left(1 + \frac{p k \tau}{n}\right)^{-n} \tag{1.3}$$

where k is the rate constant, τ the residence of the bank and p a pull factor that lumps the air rate and the froth depth, one at the design. With $k\tau/n = 0.5$ for the mineral and 0.002 for the gangue, the bank recovers 91.2 % of the copper in a concentrate of 24.0 % copper, and the tailings assay 0.091 %. Pulling harder recovers more mineral and much more gangue, so grade falls as recovery rises, as Figure 1.5a shows.

Which point on the curve is best is an economic question. If a tonne of copper is worth $V = \$8000$ and each tonne of concentrate costs $W = \$250$ to ship and treat, the margin is V times the copper recovered less W times the concentrate made. It is greatest where the last tonne of concentrate pulled carries just enough copper to pay for itself, as Equation (1.4) states.

$$\frac{\partial\,(\text{copper})}{\partial\,(\text{concentrate})} = \frac{W}{V} \tag{1.4}$$

Here that incremental grade is 3.13 %. At the design the last tonne pulled carries 10.6 %, so the bank should pull harder. The optimum lies at 97.6 % recovery and 20.3 % grade and earns \$1158 an hour more. It is flat. The margin stays within 1 % of its best from 95.0 to 99.0 % recovery, as Figure 1.5b shows, so the operating point is a standing decision that a single shift of data cannot locate. Chapter 5 finds the same shape on a mechanical cell whose handle is the froth depth. Its optimum lies near 0.53 m against a design of 0.30 m, and its margin stays within 1 % of the best between 0.47 and 0.61 m [26, 27].

Recovery is not measured either. It is computed from three assays by the two-product formula of Equation (1.5), which follows from the balances of solids and of copper across the bank [18].

$$R = \frac{c\,(f - t)}{f\,(c - t)} \tag{1.5}$$

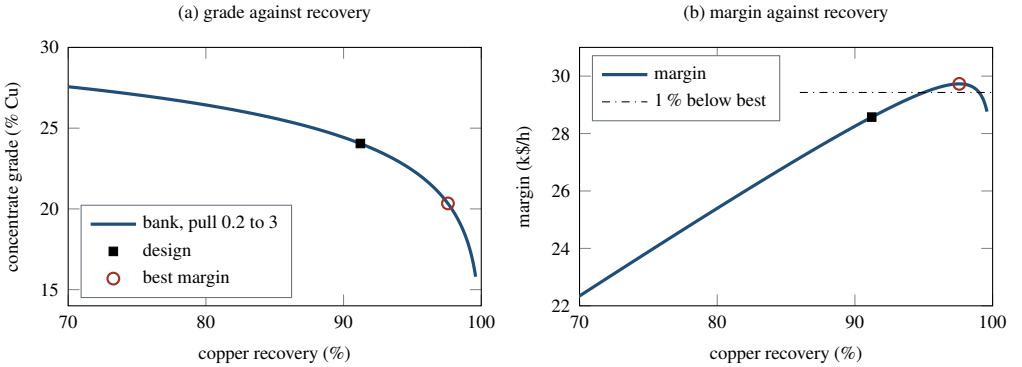


Figure 1.5 C-1's Rougher Bank as the Pull Factor Moves From 0.2 to 3. (a) Concentrate Grade Against Copper Recovery. (b) The Margin, Copper at $8000/t Less $250 per Tonne of Concentrate. The Square Marks the Design and the Circle the Best Margin. Every Point Above the Dash-Dotted Line Is Within 1 % of the Best.

With C-1's assays it returns 91.2 %. The assays matter unequally. An error of 0.01 % copper in the tailings assay moves the computed recovery by 0.97 points, the same error in the feed assay by 0.09 points, and an error of half a point in the concentrate assay hardly at all. A rougher's tailings assay is the one to get right, and it is also the lowest number the analyzer must read. Chapter 5 reports 80.5 % at the design of its cell, and the formula returns 80.6 % from its rounded feed, concentrate and tailings assays of 0.80, 5.48 and 0.176 %.

The tailings thickener recovers water, and its product is a dense underflow. The density at which solids leave depends on how long they have been compressed under the bed above them, so a denser underflow needs a slower pump, a deeper bed and more rake torque [28, 29]. In Chapter 6 the torque rises from 54.9 % of rating at 62 % solids to 85.4 % at 66, and the 92 % trip is reached at 66.54 % solids, so a density setpoint above that value is infeasible rather than ambitious. The economic optimum is pressed against a mechanical limit, which is the condition for the override of Section 1.4.6.

1.4.4 Leaching and Solvent Extraction

C-1 leaches its concentrate in $N = 4$ agitated tanks of $\tau_t = 4\,\mathrm{h}$ each. If the copper dissolves at a rate first order in what remains, with $k = 0.5\,\mathrm{h}^{-1}$ at 90 °C, each tank divides the undissolved copper by $1 + k\tau_t$, and the fraction left in the residue is Equation (1.6) [30].

$$\frac{x_N}{x_0} = (1 + k\,\tau_t)^{-N} \tag{1.6}$$

At $k\tau_t = 2$ that is 3^{-4}, 1.23 % left and 98.8 % extracted. One tank of the same 16 h would leave 11.1 %, which is why leach and adsorption circuits are trains. The rate depends on the ore and on the conditions. Here k doubles for every 10 K, so a train 5 K cooler leaves 2.94 %, and an ore that leaches 20 % more slowly leaves 2.19 %. The train answers either change with a delay set by its residence, and the residue reaches 63 % of its new value 5.1 h after the temperature drop. Chapter 7 has five tanks of 3.14 h with a rate that depends on cyanide, oxygen and pH. When a copper parcel arrives, its head tank stays within 10.2 mg/L of setpoint while the free cyanide in tank five falls from 89.7 to 53.2 mg/L. Chapter 10 leaches a calcine between a kiln and a crystallizer with no storage between them.

Solvent extraction moves the copper from the leach liquor into an organic phase and back into a clean electrolyte. In each stage the copper divides between the phases by a distribution ratio D, and with organic and aqueous flows O and A the extraction factor is $\mathscr{E} = DO/A$. Counter-current

stages compound it, and the fraction left in the raffinate after N ideal stages is the Kremser result of Equation (1.7).

$$\frac{a_N}{a_0} = \frac{\mathscr{E} - 1}{\mathscr{E}^{N+1} - 1} \tag{1.7}$$

At $\mathscr{E} = 2$ one stage leaves 33.3 %, two stages 14.3 % and three 6.7 %. C-1 uses two, because its raffinate returns to the leach and nothing it leaves is lost.

Separating two similar metals is harder than extracting one. If the extractant prefers one metal by a separation factor β, a stage at total reflux raises the log-odds $\ln[x/(1-x)]$ of the product's purity by at most $\ln\beta$, the Fenske result for distillation [31, 32]. Neodymium over praseodymium has $\beta = 1.5$, so each stage adds at most 0.405, and raising a feed of 75 % neodymium to a purity of 99.5 % takes 10.4 stages even at total reflux. Chapter 9 measures 0.167 per scrub stage at its design reflux, against 0.405 at total reflux, and its 20 scrub stages make 99.67 %. A cascade this long has two speeds. A change of reflux redistributes the inventory already there, and Chapter 9 measures the purity move 63 % of its way in 25.3 h. A change of feed grade must rebuild the inventory, which takes roughly the inventory change divided by the imbalance that drives it, $37.7/0.0525 = 718$ h, and the measured purity takes 877 h to cover 63 % of its fall. The same estimate, inventory over flow, gives the thickener bed's hours and the leach train's day.

1.4.5 Sampled Assays

An assay is a sample cut at one instant and reported later. If samples are cut every T_s and reported θ after cutting, the value in hand describes the plant of between θ and $\theta + T_s$ ago, on average $\theta + T_s/2$. C-1's leach residue is sampled every 4 h and reported 4 h later, so it is 4 to 8 h old, 6 h on average. The chapters range from minutes to most of a day.

- Chapter 7 samples its tailings every 4 h and reports them 6 h later, so a result is 6 to 10 h old.
- Chapter 9 analyzes every 4 h and reports an hour later, 1 to 5 h, and offers an X-ray fluorescence option every 30 min.
- Chapter 11 adds 15 min of free lime turnaround to a sampling cycle of up to 10 min, an effective dead time of 15.9 to 25.9 min against a burning zone whose own time constant is 13 min.
- Chapter 5, Chapter 4 and Chapter 3 use on-stream analyzers with dead times of 2 to 3.5 min.

A loop on a raw assay can act only when a report arrives, and it then corrects a plant that has already moved. In Chapter 7 the loop on the assay moves the cyanide setpoint past 122 mg/L 10.26 h after a feed grade step, against 0.26 h for a loop on a soft sensor that reads the feed grade, and it overshoots its final setpoint of 140 mg/L to 150 mg/L. Delay and sampling in general are the subject of Volume 5.

A soft sensor computes the product continuously from measured inputs [33]. C-1's runs a copy of the four tank balances driven by the measured temperature, with the design leachability. It cannot see a change in the ore, so an assay corrects it through a bias $\hat{b}$ updated at each report by Equation (1.8).

$$\hat{b} \leftarrow \hat{b} + \lambda \left[y_{lab}(t_i) - \hat{y}(t_i) - \hat{b} \right], \qquad \hat{y}_{est}(t) = \hat{y}(t) + \hat{b} \tag{1.8}$$

where $y_{lab}(t_i)$ is the assay of the sample cut at t_i and $\hat{y}(t_i)$ is the soft sensor's value at that same instant, stored when the sample was cut. The update is an exponential filter on the assay residual with the time constant $-T_s/\ln(1-\lambda)$, 9.3 h at $\lambda = 0.35$ and $T_s = 4$ h. Against white assay noise of spread σ the bias carries the variance $\lambda\sigma^2/(2-\lambda)$, so $\lambda = 0.35$ passes 46 % of the noise spread and $\lambda = 1$ all of it. A run of C-1 with assay noise of 0.10 % returns spreads of 0.049 and 0.103 against the predicted 0.046 and 0.100.

Figure 1.6 runs C-1's train through two disturbances. The cooler train is measured, the soft sensor follows it exactly, and the bias does not move. The first assay that shows it arrives at 16 h. If the bias compared that assay with the soft sensor's present value instead of its value at the sampling instant, it would read the plant's own response as a model error and pull the estimate 0.38 % low, as the dashed curve of panel (b) shows. The slower ore is not measured, so the soft sensor stays at 2.94 % while the residue climbs toward 4.85 %. The estimate first moves at 76 h, when the first assay that contains the

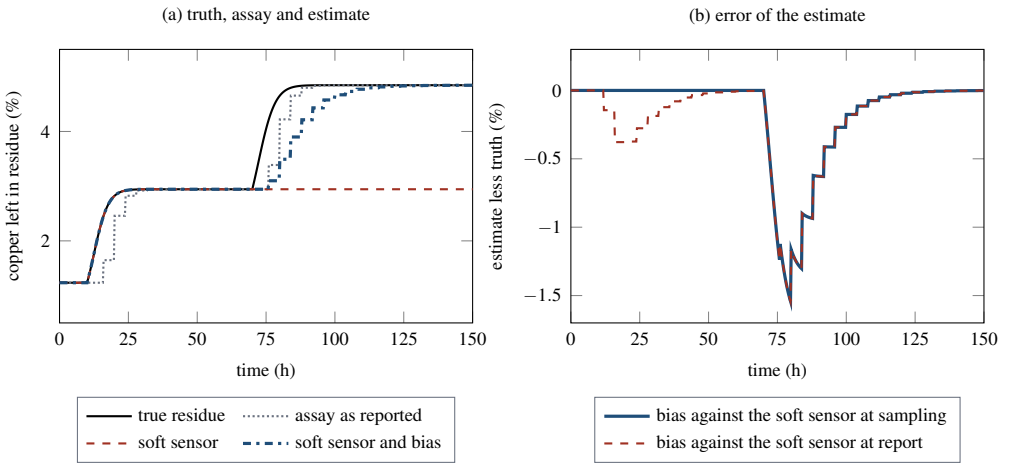


Figure 1.6 C-1's Leach Residue With the Train 5 K Cooler From 10 h (Measured) and an Ore That Leaches 20 % More Slowly From 70 h (Not Measured). (a) The True Residue (Solid), the Assay Held From Each Report (Dotted), the Soft Sensor (Dashed) and the Soft Sensor With the Bias Update at $\lambda = 0.35$ (Dash-Dotted). (b) The Error of the Estimate When the Assay Is Compared With the Soft Sensor at the Sampling Instant (Solid) and at the Report (Dashed). No Assay Noise.

change is reported, and stays more than 0.05 % in error until 42 h after the change. For a disturbance the soft sensor cannot read, it is no faster than the assay. Chapter 7 measures exactly this: a lower leachability first moves its cyanide setpoint at 20.1 h with the loop on the soft sensor, as it does with the loop on the assay. With assay noise of 0.02 g/t, a bias gain of one raises its cyanide travel by 14 to 23 % over a gain of 0.35 across seven noise seeds, and on every seed gives a higher peak tailings grade than a gain of 0.15.

> ✦ **DID YOU KNOW? The Cyanide Nobody Measured**
>
> In 1995 a change to copper-rich ore at a mine in Australia raised the cyanide bound to metals in its tailings dam to 350 parts per million, and about 2700 birds died. The mine measured only free cyanide, so the rise went undetected [17].

1.4.6 Override, Cascade and Inferential Control as Practiced

Override. When a protective controller and a quality controller share one actuator, a selector passes one of their demands. The rule that decides its direction is simple. The selector must pass the demand that moves the actuator in the direction that relieves the limit. In Chapter 6 the density, torque and bed controllers share the underflow pump, and more pump relieves both torque and bed, so the selector is a high select. Wired as a low select, which looks the same on a drawing, it blocks the protective demand when more pump is needed and drives the rake past its trip. C-1's grinding circuit needs the opposite. A mill-load or circulating-load controller protects the mill by reducing fresh feed, so it meets the feed controller in a low select. With the product held, the throughput then falls to 412 t/h on the harder ore of Section 1.4.2. Each losing controller must have its integral fed back toward the selected output, the external reset of Volume 4, or it winds up while it waits. Chapter 6 shows that a correct reset does not repair a reversed selector, and that a protective selector keeps its place beside a predictive controller.

Cascade. An inner loop removes a disturbance only if the disturbance enters before the inner measurement [34]. On C-1's cyclones, a worn pump impeller lowers the feed pressure and a pressure loop on pump speed restores it, while a harder ore leaves the pressure alone and only the outer size

loop can find it. Chapter 4 measures the integrated product error falling by a factor of 49 for impeller wear and not at all for a coarser feed. It also finds two disturbances that do enter the pressure, worn liners and an extra cyclone, which the cascade makes about three times worse, because restoring the pressure draws more dilution water and thins the feed. Chapter 10 finds the same split on its kiln's fuel. A header pressure loss falls in integral error by a factor of 52 to 93 under cascade, depending on when the loss lands in the controller scan, and a loss of heating value not at all, because the inner flow loop sees the gas flow and not its energy.

Ratio. A ratio station moves one flow with another in the same scan, before any analyzer can report. Chapter 2 sets each feeder through its calibration in the recipe's proportion, where speeds in the recipe's own proportion deliver the wrong blend. Chapter 11 ratios the combustion air to the coal command and removes a carbon monoxide excursion that oxygen feedback alone cannot prevent, because the analyzer sees the firing move only after its transport delay [35]. Chapter 9 keys its scrub and organic flows to the metered feed and trims them from the late laboratory.

Inferential control. Several chapters control a variable that no instrument reads. An electrowinning or electrorefining cell is limited by the limiting current density $j_L = nFDC/\delta$, computed from a handbook diffusivity D, an assayed concentration C and a correlation for the boundary layer δ [22]. If the estimate may overstate j_L by a fraction ε, a controller that must keep the true utilization below η_{lim} has to aim at the setpoint of Equation (1.9).

$$\eta_{SP} = \frac{\eta_{lim}}{1+\varepsilon} \tag{1.9}$$

C-1's tankhouse has a limit of $\eta_{lim} = 0.50$, and allowing for a 10 % overstatement sets the target at 0.455 and gives up 9.1 % of the current. Chapter 8 shows why the margin is worth it. A diffusivity 20 % low overstates j_L by 1.25, so a cell held at a believed 59.0 % runs at 73.7 % and shorts on nodules under flat trends. The penalty is asymmetric. An optimistic estimate ends in a short, and a pessimistic one costs a linear 10.3 % of copper. Chapter 3 holds an impact angle measured by a shell microphone, with the toe computed from the charge, and Chapter 7 a tailings grade from a soft sensor, and each shows a case in which the inferred variable sits on setpoint while the plant does not.

1.4.7 Variable Ore and Late Assays Along the Circuit

A circuit is controlled in layers. Fast loops hold flows, levels, pressures and temperatures on continuous instruments. Ratio stations and feedforward move flows with a measured load. Selectors keep equipment inside its limits. Above them, slow loops on analyzers, laboratories or soft sensors move the setpoints, and an economic layer chooses where the trades sit [2, 36]. Table 1.1 lists, for each chapter, how the ore or the feed varies, what is measured late or not at all, and the structure that handles it. Chapter 12 adds the measured results and compares how variable ore and late assays are handled along the circuit.

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 level or a bed depth is read in four chapters, a pressure, a bearing load or a rake torque in three, and a pH or oxygen probe in four. The product itself is read by an on-stream analyzer in five chapters and by a laboratory assay in four, and three chapters control a variable that only a model computes. On the other side, a valve or a flow loop is moved in six chapters, and a variable-speed drive, a solids feeder and a reagent dosing pump in four each.

The flow of ore is weighed on the belt. A weighfeeder is a conveyor belt with a section supported by rollers on one or more load cells, and a speed sensor measures the belt, so that the mass flow is the product of the two. The belt loading fluctuates in normal operation and the calculation has its own latency [37], so the trend is noisy and a little late. Figure 1.8 draws the load cell, and Chapter 2 models the scale as a mean loading times the measured speed.

Table 1.1 For Each Chapter, How the Ore or Feed Varies, What Is Measured Late or Not at All, and the Structure That Handles It

Ch.	The ore or feed varies	Late or not measured	Structure
2	Feeder calibration, belt speed, an empty silo	Blend at the scale, never	Ratio through calibrations, master PI, bump test, supervisor
3	Hardness, coarse feed, ball charge, liner wear	Toe of the charge, inferred, size 2 min late	Impact angle on speed, load on feed, toe tracking
4	Coarseness, grindability, pump wear	Product size 2 min late, circulating load never	Pressure cascade under a size loop, sump level
5	Feed grade, throughput, blower capacity	Grades 2.0 min and 3.5 min late	Froth depth, grade and tailings loops, economic setpoint
6	Feed solids, settleability	Settleability never, density, torque and bed in minutes	High select of three demands, external reset
7	Grade, leachability, copper and sulfide parcels	Tailings assay 6 to 10 h old	Cyanide cascade on a soft sensor with a bias update
8	Purification upset, passivation, dirty electrolyte	Limiting current, never, copper 2 h late	Utilization on an inferred limit, voltage override
9	Feed grade and concentration	Purity and recovery 1 to 5 h old	Ratio stations, two sampled PI loops, feed feedforward
10	Ore grade, fuel header, heating value	Product strength two units downstream	Fuel cascade, three units without storage
11	Raw meal, coal heating value, false air	Free lime 15.9 to 25.9 min late	Air ratioed to fuel with oxygen trim, free lime loop

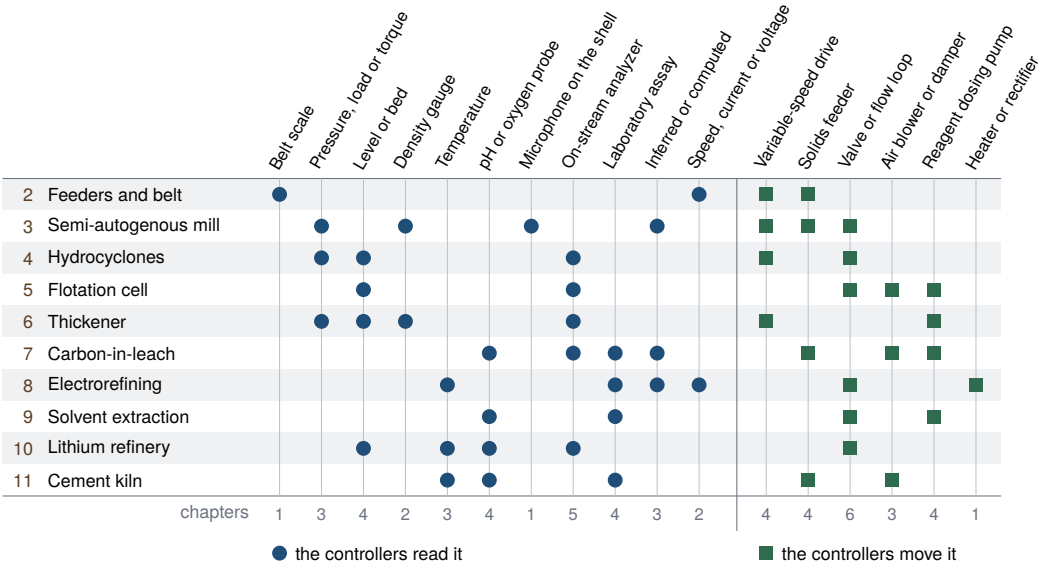


Figure 1.7 Sensors and final elements by chapter. A round mark means that the chapter's controllers read the device, a square mark that they move it. A mark is set only where the chapter's variables table, control structure or technology section names the device, and a flow station counts as a valve or flow loop. A blank says that the model leaves the device out. Notice that nine of the ten chapters read an analyzer, a laboratory assay or a computed value, and only the feeders read nothing but weight and speed.

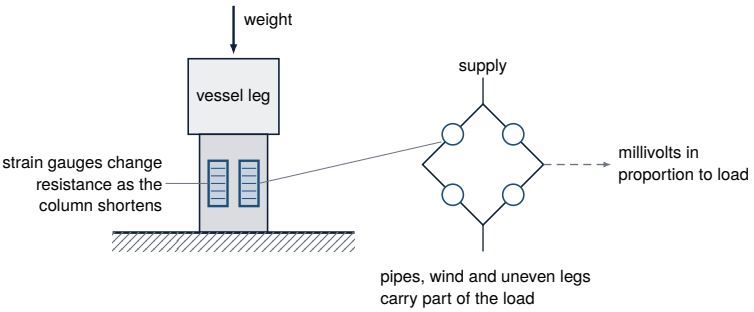


Figure 1.8 Principle of a strain-gauge load cell, drawn under the leg of a vessel. Notice that the cell reads only the load that passes through it. On a belt scale the cells carry the rollers of the weighing section, so they weigh whatever lies on that section, from whichever feeder it came [37].

Layer	What it does in this volume	Where it runs	How often
Economic and predictive control	chooses where the trades sit and moves several handles against several limits	**Supervisory application** on a computer that writes setpoints to the loops below	every 2 s to 0.4 h in the apps
Assay and soft sensor loops	moves a setpoint from a late assay or from an estimate corrected by one		once per assay; every 0.4 h to 4 h in the apps
Ratio, selectors and overrides	keys flows to a measured load, passes the safe demand to a shared actuator	**Distributed control system** or programmable controller, as function blocks	at each scan of the loops below
Regulatory control	PI loops on flows, levels, pressures, temperatures, speed and pH		milliseconds to minutes; 1 s to 5 min in the apps
Field devices	scales, gauges, probes, analyzers, drives, feeders, valves and pumps	on wires or a bus to the controller	continuously

Figure 1.9 The control layers of this volume, the platform that hosts each and how often each acts. The range of loop periods is that of [38], and the ranges in the apps are those of the chapters' technology sections. Notice that the assay layer acts a few times a day, the loops below it thousands of times between two reports.

Most of the handles of the concentrator turn rather than throttle. A variable-speed drive sets the speed of a feeder, a mill or a slurry pump. Speed control is more energy-efficient than throttling, but a variable-speed machine often cannot change a flow as fast as a control valve and offers limited options for a fail-safe state. Pumps in slurry service wear quickly [37], and Chapter 4 and Chapter 6 model that wear as a pump that delivers less than its speed promises.

Figure 1.9 places the layers of Section 1.4.7 on their platforms. Plant control loops run with cycle times from milliseconds to minutes, and a distributed control system uses a supervisory loop to mediate a group of localized controllers [38]. The apps run their regulatory loops every 1 s in Chapter 2 up to every 5 min in Chapter 9, their predictive controllers every 2 s in Chapter 2 up to every 0.4 h in Chapter 7, and their loops on assays once per report, every 4 h in Chapter 9.

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

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.

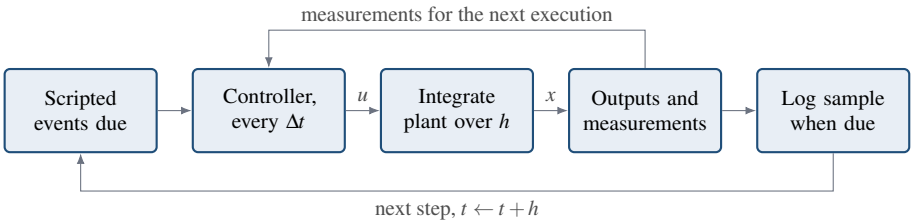


Figure 1.10 Work of the Engine in One Integration Step

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.

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

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

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

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

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

where $a = \pm 1$ is the controller action, PV^f is the measurement after an optional first-order filter, u_b is the output bias, S is the running sum of proportional and derivative increments, v is the output before limiting, u is the output sent to the actuator, I is the integral term and T_t is the tracking time constant, equal to τ_I unless a loop sets another value. The function sat clamps v_k to the actuator range and then applies a rate limit if the loop declares one.

The limit acts on the emitted output u_k only. The sum S_k is never clamped, so it equals $P_k + D_k$ and the output before limiting is the position form $u_b + K_c e_k + D_k + I_{k-1}$. Anti-windup touches the integral alone. While the output is limited, the last term of Equation (1.13) drives I toward the value that reproduces the output applied, which is back-calculation [39]. With no integral action, Equation (1.13) is skipped. Clamping the running sum would instead discard each increment that crosses a limit and keep each one that returns, so symmetric noise would push the output off the limit even in a proportional-only loop.

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

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

1.6.3 The Predictive Control and Estimator Options

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

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

where w is a measured or scheduled disturbance and c is a constant left by linearization. The model comes from a step test at the design point, from the app's balances linearized there or again at every

decision, or from the model the reader fits. It never reads the true states or parameters of the run, so a parameter the reader changes in Show Model is model mismatch, as on a real plant.

At each decision the controller minimizes the objective of Equation (1.15) over a horizon of N steps.

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

where r is the setpoint, Q weights the tracking error, R is the move suppression on each input change Δu_j and ρ_i is the price of a slack $s_i \geq 0$ that lets an output band be broken when it cannot be held. Actuator limits are hard constraints. The quadratic program is solved in the browser by a dual active-set method [40]. 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 [41]. 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 [42].

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.

Table 1.3 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Feeders and a belt scale	A total on target does not prove the blend.	introductory	2 h
3	Semi-autogenous mill	Holding the impact angle is not aiming at the toe.	intermediate	3 h
4	Hydrocyclone cluster	A cascade helps only with what enters its inner loop.	introductory	2.5 h
5	Rougher flotation cell	The best froth depth is a flat economic choice.	intermediate	2.5 h
6	Tailings thickener	The selector's direction decides the protection.	introductory	2.5 h
7	Carbon-in-leach train	A soft sensor is only as fast as what it reads.	advanced	3 h
8	Electrorefining cell	An inferred limit needs a margin for model error.	intermediate	2.5 h
9	Solvent extraction cascade	Two time scales and a late laboratory set the loops.	advanced	2.5 h
10	Lithium hydroxide refinery	One disturbance reaches three loops at three times.	intermediate	2.5 h
11	Rotary cement kiln	Ratio for the fast path, a patient loop for the slow assay.	advanced	2.5 h

Chapters 2, 4 and 6 form a first course: a blend that one total cannot see, a cascade sorted by where each disturbance enters, and a selector wired in the safe direction. They rest on Sections 1.4.2 and 1.4.6. Chapters 3, 5, 8 and 10 add a moving constraint, an economic setpoint, an inferred limit and units in series. Chapters 7, 9 and 11 combine late assays with slow dynamics and are best left until Section 1.4.5 is familiar. A reader who follows the ore can read the concentrator chapters in their order, which is the order of Figure 1.2.

KEY POINTS

- A mineral circuit is controlled while its feed varies without being measured and its quality is known late.
- Fineness, throughput and the mill limit cannot all be held, and which one gives way is a control decision.
- The best point on the grade and recovery curve is economic, and it is flat, so a single shift of data cannot locate it.
- An assay describes the plant of hours before, and for a disturbance that the soft sensor cannot read, the soft sensor is no faster than the assay.
- A selector must pass the demand that relieves the limit, and a controller on an inferred limit must aim below it.

Where else this applies. Selectors and override return in Volume 4, dead time, soft sensors and estimation in Volume 5 and economic and predictive control in Volume 9.

References

[1] D. Hodouin, S.-L. Jämsä-Jounela, M. T. Carvalho, and L. Bergh, "State of the art and challenges in mineral processing control," *Control Engineering Practice*, vol. 9, no. 9, pp. 995–1005, 2001.

[2] D. Hodouin, "Methods for automatic control, observation, and optimization in mineral processing plants," *Journal of Process Control*, vol. 21, no. 2, pp. 211–225, 2011.

[3] "Copper," in *Encyclopædia Britannica*, 11th ed. Cambridge: Cambridge University Press, 1910, vol. 7, pp. 102–110, read in the Wikisource transcription (en.wikisource.org, 1911 Encyclopædia Britannica/Copper), accessed November 19, 2026.

[4] J. S. MacArthur, R. W. Forrest, and W. Forrest, "Process of obtaining gold and silver from ores," U.S. Patent 403,202, application filed November 9, 1887, 1889, patented May 14, 1889.

[5] E. Bretney, "Water-purifier," U.S. Patent 453,105, 1891, application filed November 3, 1890; patented May 26, 1891.

[6] Supreme Court of the United States, "Minerals Separation, Ltd. v. Hyde, 242 U.S. 261," Opinion of the court, decided December 11, 1916; Legal Information Institute, law.cornell.edu/supremecourt/text/242/261, 1916.

[7] J. V. N. Dorr, "Decanting apparatus," U.S. Patent 867,958, filed February 26, 1907, issued October 15, 1907, 1907.

[8] F. Concha and A. Barrientos, “A critical review of thickener design methods,” *KONA Powder and Particle Journal*, vol. 11, pp. 79–104, 1993.

[9] R. B. Ellestad and K. M. Leute, “Method of extracting lithium values from spodumene ores,” U.S. Patent 2,516,109, application filed September 16, 1948, 1950, patented July 25, 1950.

[10] D. F. Peppard, G. W. Mason, J. L. Maier, and W. J. Driscoll, “Fractional extraction of the lanthanides as their di-alkyl orthophosphates,” *Journal of Inorganic and Nuclear Chemistry*, vol. 4, no. 5–6, pp. 334–343, 1957.

[11] N. Tsuchida, “Studies on the mechanism of gold adsorption and elution in the carbon-in-pulp process,” Ph.D. dissertation, Murdoch University, Perth, Western Australia, 1984.

[12] S. B. Donkor, “On-line sensors for measuring the total ball and charge level in tumbling mills,” MSc (Eng) thesis, University of Cape Town, Centre for Minerals Research, 2014, hdl.handle.net/11427/13225.

[13] L. M. Tavares, “A review of advanced ball mill modelling,” *KONA Powder and Particle Journal*, vol. 34, pp. 106–124, 2017.

[14] “Cement,” in *Encyclopædia Britannica*, 11th ed. Cambridge: Cambridge University Press, 1910, vol. 5, pp. 653–659, read in the Wikisource transcription (en.wikisource.org, 1911 Encyclopædia Britannica/Cement), accessed November 19, 2026.

[15] K. L. Lyon, “Separation of adjacent rare earth elements using solvent extraction,” M.S. thesis, University of Idaho, Moscow, ID, 2016. [Online]. Available: https://www.lib.uidaho.edu/digital/etd/items/lyon_idaho_0089n_10815.html

[16] A. E. Hidden, “Errors in conveyor belt weigher systems,” *Measurement and Control*, vol. 6, no. 6, pp. 245–252, 1973.

[17] W. A. M. van der Merwe, “The development of an online amperometric technique to measure free and WAD cyanide in gold plant leach liquors and effluent streams,” M.Eng. dissertation, University of Cape Town, 2012.

[18] B. A. Wills and J. A. Finch, *Wills’ Mineral Processing Technology: An Introduction to the Practical Aspects of Ore Treatment and Mineral Recovery*, 8th ed. Oxford: Butterworth-Heinemann, 2016.

[19] J. Jeswiet and A. Szekeres, “Energy consumption in mining comminution,” *Procedia CIRP*, vol. 48, pp. 140–145, 2016.

[20] C. R. Boeree, “Up-scaling of froth flotation equipment,” MSc thesis, report AES/RE/14-08, Delft University of Technology, 2014, resolver.tudelft.nl/uuid:660e58d7-a82d-46d9-a8ff-1bfba42b2b05.

[21] U.S. Geological Survey, “Mineral commodity summaries 2025: Cement,” U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[22] M. E. Schlesinger, M. J. King, K. C. Sole, and W. G. Davenport, *Extractive Metallurgy of Copper*, 5th ed. Oxford: Elsevier, 2011.

[23] K. Nageswararao, D. M. Wiseman, and T. J. Napier-Munn, “Two empirical hydrocyclone models revisited,” *Minerals Engineering*, vol. 17, no. 5, pp. 671–687, 2004.

[24] R. Rajamani, J. Delgadillo, and V. Duriseti, “Online SAG mill pulse measurement and optimization,” University of Utah, Metallurgical Engineering Department, Semi-annual technical progress report DE-FC26-04NT42088, 2006, u.S. Department of Energy, Office of Scientific and Technical Information.

[25] D. Wei and I. K. Craig, “Grinding mill circuits: A survey of control and economic concerns,” *International Journal of Mineral Processing*, vol. 90, no. 1–4, pp. 56–66, 2009.

[26] B. J. Shean and J. J. Cilliers, “A review of froth flotation control,” *International Journal of Mineral Processing*, vol. 100, no. 3–4, pp. 57–71, 2011.

[27] L. G. Bergh and J. B. Yianatos, “The long way toward multivariate predictive control of flotation processes,” *Journal of Process Control*, vol. 21, no. 2, pp. 226–234, 2011.

[28] G. J. Kynch, “A theory of sedimentation,” *Transactions of the Faraday Society*, vol. 48, pp. 166–176, 1952.

[29] S. P. Usher and P. J. Scales, “Steady state thickener modelling from the compressive yield stress and hindered settling function,” *Chemical Engineering Journal*, vol. 111, no. 2–3, pp. 253–261, 2005.

[30] J. McLaughlin and G. E. Agar, “Development and application of a first order rate equation for modelling the dissolution of gold in cyanide solution,” *Minerals Engineering*, vol. 4, no. 12, pp. 1305–1314, 1991.

[31] M. R. Fenske, “Fractionation of straight-run Pennsylvania gasoline,” *Industrial & Engineering Chemistry*, vol. 24, no. 5, pp. 482–485, 1932.

[32] F. Xie, T. A. Zhang, D. Dreisinger, and F. Doyle, “A critical review on solvent extraction of rare earths from aqueous solutions,” *Minerals Engineering*, vol. 56, pp. 10–28, 2014.

[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] P. R. Krishnaswamy, G. P. Rangaiah, R. K. Jha, and P. B. Deshpande, "When to use cascade control," *Industrial & Engineering Chemistry Research*, vol. 29, no. 10, pp. 2163–2166, 1990.

[35] K. S. Stadler, J. Poland, and E. Gallestey, "Model predictive control of a rotary cement kiln," *Control Engineering Practice*, vol. 19, no. 1, pp. 1–9, 2011.

[36] D. Sbárbaro and R. del Villar, Eds., *Advanced Control and Supervision of Mineral Processing Plants*, ser. Advances in Industrial Control. London: Springer, 2010.

[37] 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.

[38] 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.

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

[40] 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.

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

[42] 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 feeder 84 m upstream of a belt scale puts 600 t/h of ore on a belt that runs at 2.0 m/s.

(a) Find the belt loading in kg/m and the transport delay from the feeder to the scale. *Checkpoint:* 83.3 kg/m and 42 s.

(b) The belt slows to 1.6 m/s with the feeder unchanged. What does the scale read at once, and what is the new delay? *Checkpoint:* 480 t/h and 52.5 s.

1.2 Concept check · 2 min

Three feeders in a fixed recipe share one belt scale, and a master PI loop on the scale's total sets all three speeds in proportion. One feeder loses 20 % of its delivery per unit speed, and the loop brings the total back to its setpoint. What is true afterward?

(a) the blend is restored, because every feeder was raised in proportion

(b) the weak feeder can be found from the scale's delay pattern in steady state

(c) the total is on setpoint and the blend is still wrong

(d) the total is wrong, because the scale reads load times speed

1.3 Theory and calculation · 5 min

A ball mill closed by cyclones receives $F = 300\,\mathrm{t/h}$ of fresh ore with a coarse share $\varphi = 0.75$. The mill breaks $b = 0.30$ of the coarse it receives on each pass, and the cyclones return $E_c = 0.85$ of the coarse to the mill.

(a) Find the recycle gain and the coarse in the mill feed. *Checkpoint:* $g = 0.595$, 555.6 t/h.

(b) Find the coarse in the product and its share. *Checkpoint:* 58.3 t/h, 19.4 %.

(c) The mill is a lag of 4 min. Find the circuit's coarse time constant. *Checkpoint:* 9.9 min.

1.4 Theory and calculation · 3 min

Size assays around a cyclone cluster give coarse shares of 0.48 in the cyclone feed, 0.70 in the underflow and 0.15 in the overflow.

(a) Find the ratio of underflow to overflow solids. *Checkpoint:* 1.500.

(b) The feed assay is in fact 0.49. Find the ratio again. *Checkpoint:* 1.619.

1.5 Theory and calculation · 4 min

A rougher bank's assays are 0.60 % copper in the feed, 22.0 % in the concentrate and 0.08 % in the tailings.

(a) Find the copper recovery and the mass pull, the share of the feed that reports to the concentrate. *Checkpoint:* 86.98 % and 2.37 %.

(b) The true tailings assay is 0.10 %. Find the recovery. *Checkpoint:* 83.71 %.

1.6 Concept check · 2 min

At C-1's design the last tonne of concentrate the bank pulls carries 10.6 % copper, while the ratio of the cost of a tonne of concentrate to the value of a tonne of copper is 3.13 %. What should the operating point do?

(a) pull harder, because the last tonne pulled earns more than it costs

(b) pull less, because a higher grade always earns a higher margin

(c) stay, because the design point maximizes grade times recovery

(d) pull less, because recovery above 90 % carries no value

1.7 Theory and calculation · 3 min

A copper concentrate leaches at a rate first order in the copper that remains, with $k = 0.3\,\mathrm{h}^{-1}$. Find the copper left in the residue for one tank of 15 h and for three tanks of 5 h each, all perfectly mixed.

1.8 Theory and calculation · 5 min

(a) A counter-current solvent extraction section has three ideal stages and an extraction factor of 2.5. Find the fraction of the metal left in the raffinate. *Checkpoint:* 3.94 %.

(b) Two metals are separated with a separation factor of 2. Find the least number of ideal stages that raises a feed of 50 % to a purity of 99 % at total reflux. *Checkpoint:* 6.63 stages.

1.9 Theory and calculation · 5 min

A laboratory samples a tailings stream every 8 h and reports each result 2 h after the sample was cut.

(a) Between which ages does the result in hand describe the plant, and what is its average age? *Checkpoint:* 2 to 10 h, 6 h.

(b) A bias update of the form of Equation (1.8) runs at $\lambda = 0.5$. Find its time constant and the share of the assay's noise spread it passes to the estimate. *Checkpoint:* 11.5 h, 57.7 %.

1.10 Concept check · 3 min

A throughput controller and a mill-load controller both set C-1's fresh feed, and the mill-load controller asks for less feed when the load is high. Which selector belongs between them?

(a) a high select, so the protective controller wins when its demand is higher

(b) a high select, because the throughput controller is the primary loop

(c) either selector, provided the losing controller has external reset

(d) a low select, because the protective action is less feed

CHAPTER 2

Three Feeders and One Belt Scale: Ore Blending Through Speed-Dependent Delays



The controller moves 1 the speed of each feeder, set through its drive

The controller reads 1 a switch that trips when a silo is empty 2 the belt scale's load cells 3 the belt's speed sensor 4 each feeder drive's actual speed

What disturbs it 1 a feeder that delivers less ore per turn 2 a silo running empty 3 the belt slowed or speeded up

Notice that the one scale sits after all three feeders, so it weighs only their sum, and each feeder's ore reaches it after its own trip along the belt. Illustration, not to scale.

IN PLAIN TERMS **Three Feeders, One Scale, One Moving Belt**

The main idea. Three bins drop different kinds of ore onto one long moving belt, and a scale near its end weighs everything that passes. A controller sets the speed of each feeder so that the mill gets the right total and the right mix.

The challenge. Ore from the far feeder takes about a minute to reach the scale, and ore from the near one only seconds, so every correction shows up late and at different times. The scale weighs only the total, so it cannot tell which feeder has fallen short.

Why it matters. A total that looks right can still send the wrong mix of ores to the mill.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 2.2, the fact boxes of the Introduction and of Section 2.6, the try-it boxes of Section 2.4, 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 2.6, Section 2.7, Section 2.8 with its open challenge

The control question Three feeders drop ore onto one belt, and one scale downstream weighs the combined stream after a delay that each feeder's distance and the belt speed set. What can a loop on that one total achieve, and what does a total on target leave unproven?

Learning objectives

- Compute each feeder's delay and the flow a belt scale reads after a belt speed change.
- Allocate a total through the calibrations, and explain why speeds in the recipe's proportion miss it.
- Explain why a master PI on the total must wait for its slowest path.
- Find a drifted feeder gain with a bump test, and judge its result.
- Compare three structures on an infeasible shift and say where the MPC earns its place.

Prerequisites Sections 1.4.2, 1.4.6, 1.6.2, 1.6.3 and 1.6.4.

The app https://apmonitor.com/apps/ore-feed/

See also Volume 5, Chapter 6 (feeder ratio control, tracked material). Volume 3, Chapter 6 (a dose computed ahead of a slow measurement). Volume 9, Chapter 6 (MPC on a delay that follows the flow).

2.1 Introduction

A grinding circuit is fed from stockpiles or bins through feeders that drop ore onto a collecting conveyor, and the feed is weighed by a belt scale somewhere along that conveyor before it reaches the mill. In the generic plant simulated here three silos stand over feeders at 0, 60 and 100 m along a belt that runs at 2.0 m/s, and the scale sits at 120 m. The plant is asked for 900 t/h at a 20:50:30 mass blend, which the feeders deliver at speeds of 45, 56.25 and 45 % because their gains differ (4, 8 and 6 t/h per percent). A 10 % speed step first moves the scale 10 t/h after 62.0 s from feeder A, 30.9 s from B and 11.1 s from C. Slowing the loaded belt from 2.0 to 1.5 m/s cuts the flow the scale reads from 900 to 675 t/h before any feeder has moved.

Plants with this shape feed the mill of Chapter 3 and many of the minerals processing circuits whose control methods have been surveyed [1, 2], and the same arrangement builds cement raw mixes, where proportioning several raw materials has been a control problem since at least the 1970s [3–5]. The class is a multivariable dead-time process whose input-output pairs carry different delays [6], with the extra feature that the delays are set by the integral of a conveyor speed that is itself adjusted to save drive energy [7, 8], which makes them time-varying input delays [9]. A block diagram with three fixed dead times hides two consequences that the simulator makes visible. The first is that a belt scale multiplies the load on its weighbridge by the belt speed [10], so a speed change is also an immediate flow disturbance from ore that no controller can recall. The second is that a total measured downstream of three feeders in fixed proportion cannot say which feeder is short. A student who trusts a block diagram tends to assume that a total-flow loop at its setpoint is also delivering the recipe, and the simulator is built to show otherwise.

> ✦ **DID YOU KNOW? Two Belts, Two Delays**
>
> A textbook water plant has the same problem. Lime and flocculant ride belts with dead times of 60 and 5 s, so a feedforward that moves both at once delivers the lime 55 s too soon and the pH rises [11]. The galvanizing line of Volume 16, Chapter 10 weighs its coating 100 m of strip downstream, a delay that changes with the line speed.

Teaching material on blending treats the composition analyzer and the proportioning loop of a cement plant [3, 5], or a dead-time compensator with constant delays [12], and rarely puts a speed-dependent transport, an unobservable blend and an infeasible target in one plant. This chapter describes, with every equation needed to rebuild it, a three-feeder ore conveyor simulated in a web browser, and measures, on identical measurements and disturbances, what ratio feedforward, a master

PI and an MPC each achieve on it. The transport law, the ratio allocation, the PI and the MPC are established methods. The simulator tracks the ore source by source on a grid fixed to the belt, so arrivals follow the integrated travel exactly and mass closes to round-off through stops, refills and spills. Its MPC condenses the prediction in closed form from the estimated belt, and every controller reads only a packet of declared measurements. Nine cases with completion rules carry the lessons, and the MPC's measured standing against ratio feedforward is reported as found, within 0.5 pp of on-spec share except where belt loading binds.

Section 2.2 reviews belt weighers, raw-mix proportioning and delayed feed. Section 2.3 sets out the conveyor, its drives, the belt grid and the scale. Section 2.4 measures the three delays, the flow change that a belt speed change causes, and what one total cannot say. Section 2.5 describes the supervisor, ratio feedforward, the master PI, the estimator with its bump test, and the MPC. Section 2.6 reports the nine cases, the controller comparison and the graded challenge, Section 2.7 sets the loop in its real hardware, Section 2.8 states the limitations and Section 2.9 draws the conclusions. The exercises at the end of the chapter run the app's cases.

2.2 Belt Scales, Raw-Mix Proportioning and Delayed Feed

The errors of belt weighers and the international recommendation that now governs them are traced in the history box below [10, 13]. Later work showed that the length of the conveyor itself affects belt weigher accuracy [14]. The simulated scale follows the recommendation's principle directly, since it reads the load on a 1 m weighbridge times the measured speed.

> ⧖ HISTORY **Weighing Ore in Motion (1973 to 2014)**
>
> In 1973 an analysis of continuous belt weighers found their performance limited in ways that still held back their widest adoption. It divided a weigher into four essential elements, found the most serious errors in the weigh-length element, with the measurement of belt travel a possible further source, and stressed several weighing idlers, a long idler pitch and a low belt tension [13]. The international recommendation of 2014 defines the instrument as one that weighs a bulk product continuously on a conveyor belt without stopping the belt. Its totalizing device integrates the load per unit length times the belt speed, it covers single-speed and variable-speed conveyors, and its four accuracy classes allow in-service errors from 0.2 to 2 % of a totalized load [10].

Cement plants proportion several raw materials into a raw meal, and the blending loop there has a control literature of its own. Self-tuning adaptive control was applied to cement raw material blending by 1978 [3], stochastic models and self-tuning control of a continuous raw-material mixing system followed in 1980 [4], and robust gain-scheduled blending control of raw-mix quality was reported in 2011 [5]. Those studies act on the quality of the mix. The simulated ore feed is closed on tonnage alone, with the blend set by proportioning, and the wider methods for observing and controlling mineral processing circuits are the subject of two reviews [1, 2].

Transport delay entered control practice through dead-time compensation, whose origin and extension to a varying delay the box below traces [6, 9, 12].

> ⧖ HISTORY **From Smith's Predictor to a Moving Delay (1957 to 2010)**
>
> In 1957 O. J. M. Smith proposed a predictor for processes with dead time. It splits the process model into a fast, delay-free part and a delay, and its primary controller acts on a signal that anticipates the output, which removes the delay from the closed loop's characteristic equation when the model is exact. It was proposed for open-loop stable plants, most dead-time compensators since derive from it, and by 1973 it had been extended to plants with several inputs and outputs [6]. A conveyor's delay changes with its speed. An explicit predictor feedback for a time-varying input delay was published in 1991, and in 2010 its exponential stability was proved through a transport equation whose propagation speed is not constant [9].

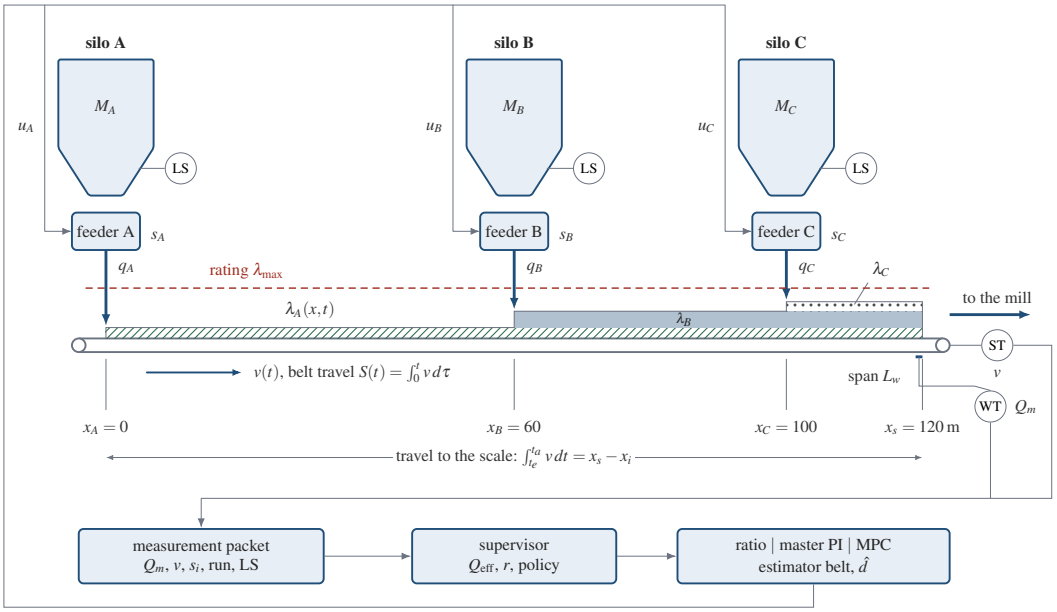


Figure 2.1 Three-feeder ore conveyor as modeled, positions to scale along the belt. Silo inventories M_i feed drives at actual speed s_i, which deliver q_i onto the belt at x_i. Source-resolved loadings $\lambda_i(x,t)$ ride the belt at speed $v(t)$ to the weighbridge of span L_w at x_s, where the weight transmitter (WT) and the speed sensor (ST) give the reading Q_m. Empty switches (LS), drive feedback and the scale form the measurement packet, the supervisor sets the effective target Q_{eff} and recipe r, and the selected structure returns the commands u_i. The dashed line is the 200 kg/m loading rating $\lambda_{\max}$. Not a construction drawing.

A conveyor adds a further twist, because its delay is a distance divided by a speed that changes, and belt speed has itself become a manipulated variable, adjusted to the material flow to save drive energy [7, 8]. Model predictive control is widely applied in industry to multivariable plants with delays and constraints [15], and it is the natural frame here because the belt's loading limit and the ore already in transit are both constraints on future decisions.

2.3 The Conveyor as Modeled

The simulated plant is generic and has no proprietary source, as shown in Figure 2.1. Section 1.4.2 introduces its two transport effects on the running example of Chapter 1. Silos A, B and C hold inventories M_i over feeder drives whose actual speeds s_i follow commands u_i. Feeder i delivers q_i onto the belt at x_i (0, 60 and 100 m), the belt moves at a uniform speed $v(t)$, and a belt scale at $x_s = 120$ m weighs the combined stream over a weighbridge of span L_w. Each silo carries an optional empty switch (LS). Ore that crosses the scale is delivered to the mill, which is outside the model.

Each variable carries a role in Table 2.1. Eight quantities are integrated: the three drive speeds, the belt speed, the three silo inventories and the scale transmitter's filtered reading. The belt's load is not an ordinary state vector but a grid of masses per source, described below. Units are seconds, meters, metric tonnes, t/h, m/s and kg/m, and a blend is an as-fed mass fraction.

2.3.1 Drives, Silos and Delivery

The controller's command reaches each drive as an applied value $\tilde{u}_i$, which is the command clamped to 0–100 %, or zero while the belt is stopped or the drive is tripped. The drive follows it through an exact first-order lag of time constant $\tau_d = 2$ s whose step is capped by the slew limit, as shown in

Table 2.1 Variables, Their Roles and Design-Point Values

Variable	Role	Design	Range
Q_m	CV, scale reading	900 t/h	SP 0–1800 t/h
f_i	CV (truth only), arriving blend	20:50:30	recipe r
u_A, u_B, u_C	MV, speed commands	45, 56.25, 45 %	0–100 %
D	MV of the master, total demand	900 t/h	0–$D_{\max}$
v	DV (measured), belt speed	2.0 m/s	1.0–3.0, or 0
g_i	DV (hidden), true-gain factor	1	0.5–1.2
b, σ	DV, scale bias and noise	0, 2 t/h	±5 %, 0–300 %
s_i	states, drive speeds	45, 56.25, 45 %	0–100 %
M_i	states, silo inventories	100, 200, 150 t	≥ 0
$\lambda_i(\sigma)$	grid, loading per source	25, 62.5, 37.5 kg/m	spill above 240
q_i	output, delivered flows	180, 450, 270 t/h	$\leq 100 k_i g_i$
Q	output (truth), flow at scale	900 t/h	—

Equation (2.1).

$$\begin{aligned} \Delta s_i &= \mathrm{clamp}\big((\tilde{u}_i - s_i^n)(1 - \mathrm{e}^{-\Delta t/\tau_d}),\ \pm \dot{s}\,\Delta t\big) \\ s_i^{n+1} &= \mathrm{clamp}\big(s_i^n + \Delta s_i,\ 0,\ 100\big) \end{aligned} \tag{2.1}$$

where $\dot{s} = 5\,\%/\mathrm{s}$ while the drive may run and 20 %/s when the interlock stops it. Each feeder delivers the true gain times the step's mean speed, and the silo gives up at most what it holds, as shown in Equation (2.2).

$$\begin{aligned} q_i &= g_i\, k_i\, \tfrac{1}{2}\big(s_i^n + s_i^{n+1}\big) \\ m_i &= \min\Big(\frac{q_i\,\Delta t}{3600},\ M_i^n\Big), \qquad M_i^{n+1} = M_i^n - m_i \end{aligned} \tag{2.2}$$

where k_i is the nameplate gain (4, 8 and 6 t/h per percent, a hidden model parameter that no controller reads), g_i a hidden factor that the calibration drift and the true-gain sliders change, and m_i the tonnes withdrawn in the step. An empty silo stops delivery at once while its drive may run on, so the inventory never becomes negative.

2.3.2 The Belt in Its Own Coordinate

The continuum model of each source's loading λ_i (kg/m) is the transport law of Equation (2.3).

$$\frac{\partial \lambda_i}{\partial t} + v(t)\,\frac{\partial \lambda_i}{\partial x} = \frac{q_i(t)}{3.6}\,\delta(x - x_i) \tag{2.3}$$

Because the whole belt surface moves at one speed and ore never moves relative to it, the law becomes trivial in the belt's own coordinate $\sigma = x - S(t)$, where $S(t) = \int_0^t v\,d\tau$ is the belt's travel. Along σ every loading is constant in time, each feeder deposits at the moving point $\sigma_i = x_i - S(t)$, and the scale removes whatever passes $\sigma_s = x_s - S(t)$. The simulator therefore keeps cells of $h = 0.25\,\mathrm{m}$ fixed in σ, each holding one mass per source in a ring buffer long enough for the 120 m belt plus 20 m of travel per step. The belt speed follows its setting at $0.1\ \mathrm{m/s^2}$ ($0.25\ \mathrm{m/s^2}$ when stopping), and the travel advances by the trapezoid $S^{n+1} = S^n + \frac{1}{2}(v^n + v^{n+1})\Delta t$, which is exact for a ramp. Over one step, feeder i spreads its $1000\,m_i$ kg over the interval $[x_i - S^{n+1},\, x_i - S^n]$ in proportion to each cell's overlap, and on a stopped belt the whole deposit lands in the cell under the chute, which is how a rundown piles up. The crossing point then sweeps from $x_s - S^n$ to $x_s - S^{n+1}$, removing the crossed fraction of each cell it enters, so every kilogram is either on the belt or counted at the scale. Ore above $\lambda_{\mathrm{spill}} = 240\,\mathrm{kg/m}$ in the cells a chute has just swept falls off and is counted as spill. The only approximation is the uniformity of a cell, which smears an arrival over at most one cell (0.125 s at 2 m/s).

The true flow at the scale is the crossed mass per step, $Q = 3.6\sum_i \Delta m_{i,\mathrm{crossed}}/\Delta t$ in t/h with the mass in kilograms. Two relations follow from the coordinate change and are used to check the code. Ore

discharged at t_e arrives when the belt has traveled the distance, and while the belt is running a source's arriving flow is its discharge flow scaled by the ratio of the speeds, as shown in Equation (2.4).

$$\int_{t_e}^{t_a} v(t)\,dt = x_s - x_i, \qquad q_{i,s}(t) = q_i(t_e)\,\frac{v(t)}{v(t_e)} \tag{2.4}$$

At constant speed the first relation reduces to distance over speed (60, 30 and 10 s at 2 m/s). The second explains why a slowdown cuts the flow at once. The belt laid at 2.0 m/s carries 125 kg/m, and at 1.5 m/s it crosses the scale at $3.6 \times 1.5 \times 125 = 675$ t/h. A delay recomputed as distance over the current speed would relocate the ore already on the belt, which this grid never does.

Pause. Try Exercise 2.1 before reading on.

2.3.3 The Belt Scale

The scale reads the load on its weighbridge times the measured speed, with a bias and seeded noise, through a transmitter filter, as shown in Equation (2.5).

$$\begin{aligned} Q_{\text{raw}} &= \max\bigl(0,\ 3.6\,\bar{\lambda}_{L_w}\,v\,(1+b) + \gamma\,\sigma_s z\bigr) \\ Q_m^{n+1} &= Q_m^n + \bigl(Q_{\text{raw}} - Q_m^n\bigr)\bigl(1 - \mathrm{e}^{-\Delta t/\tau_s}\bigr) \end{aligned} \tag{2.5}$$

where $\bar{\lambda}_{L_w}$ is the mean loading over the 1 m span ahead of 120 m, b the bias, $\sigma_s = 2$ t/h the noise level at $\gamma = 1$ (the slider's 100 %), z a unit colored noise with a 2 s correlation time seeded per case, and $\tau_s = 2$ s. The noise level sits inside the 1 % in-service error that a class 1 belt weigher may show on a totalized load [10], although that error is defined on totals, not on an instantaneous rate.

2.3.4 Parameters, Initialization and Numerics

The parameters appear in Table 2.2. Every value is a teaching value, except the belt weigher principle, which comes from the international recommendation [10]. A case starts lined out, with the belt filled as a long steady run would have left it, source i at $k_i g_i s_i/(3.6v)$ from x_i to the scale. At the design point that is 25, 62.5 and 37.5 kg/m, 125 kg/m at the scale and 7.5 t of ore in transit. The physics step is $\Delta t = 0.1$ s, the controllers and the supervisor run every 1 s, the MPC decides every 2 s, and the historian logs at 1 Hz. Halving Δt moves case 6's delivered tonnes under the master PI by 0.0002 % (134.981 against 134.981 t) and the arrival of a step on A from 60.40 to 60.30 s, and halving the cell moves a variable-speed run by less than 0.000 05 %.

2.3.5 Assumptions

The belt is one rigid body at one speed with no stretch, slip or tracking, and ore does not move relative to it, which holds for a troughed belt below its surcharge angle but not for a steep incline. Delivery is linear in speed through a hidden gain, with no dead band, flooding, bridging or rat-holing, and the only starvation is an empty silo. Blend is an as-fed mass fraction, so assay, moisture, particle size and segregation are absent. The empty switch trips when the cone is empty and a refill is instantaneous. The rating and the spill loading are teaching values, where a real rating comes from the belt width, troughing angle, speed and bulk density through conveyor design standards. Each assumption breaks where its named effect matters, such as a sticky ore for linear delivery or a long incline for the rigid belt.

2.4 Dynamic Character

2.4.1 One Measurement, Three Delays

The scale sees each feeder through its own path. With the noise off, a 10 % step on A first moves the reading 10 t/h after 61.7 s, which matches the 61.67 s obtained by convolving the 60 s travel with the 2 s drive lag, the half-second the ore takes to fill the 1 m span and the 2 s filter. With the shipped noise the steps on A, B and C are first seen after 62.0, 30.9 and 11.1 s, and the equal speed steps add

Table 2.2 Model and Controller Parameters. H: Hidden From Learner and Controllers, V: Visible, A: Adjustable

Symbol	Meaning	Value	
k_A, k_B, k_C	nameplate (true) feeder gains	4, 8, 6 t/h/%	H
$\hat{k}_i$	controllers' calibrations	4, 8, 6 t/h/%	A
τ_d, $\dot{s}$	drive lag, slew limit	2 s, 5 %/s	A
$\dot{s}_{\text{stop}}$	drive ramp under the interlock	20 %/s	A
x_i, x_s	discharge points, scale	0, 60, 100, 120 m	V
h	belt grid cell	0.25 m	V
a, a_{stop}	belt acceleration, stop	0.1, 0.25 m/s^2	A
$\lambda_{\max}$	belt loading rating	200 kg/m	A
λ_{spill}	spill loading	240 kg/m	A
L_w, τ_s	weighbridge span, scale filter	1 m, 2 s	A
σ_s	scale noise (1σ, 2 s colored)	2 t/h	A
M_{sw}	empty switch trip (clears above 5 t)	0 t	A
$Q_{\min}$	least flow with a defined blend	90 t/h	V
m	constraint margin on the rating	2 %	A
K_c, T_i	master PI (case 3 ships 1.5, 15 s)	0.7, 30 s	A
w_Q, w_B, $w_{\Delta u}$	MPC throughput, blend, move weights	1, 3, 1	A
H, T_{mpc}	MPC horizon, decision period	240, 2 s	A
ρ	loading-slack price	1000 per (kg/m)^2	A
$\tau_{\hat{d}}$, B	correction filter, bound	20 s, 25 %	V

40, 80 and 60 t/h (901 to 942 to 1019 to 1082 t/h), so one scale sees three plants with different gains and delays. A master loop on the total therefore faces a dead time that is a weighted mixture of three, about 13 s for 30 % of the response, 33 s for 50 % and 63 s for the last 20 %.

2.4.2 A Speed Change Is a Flow Disturbance

In case 6 the belt slows from 2.0 to 1.5 m/s over 60–65 s and speeds up to 2.5 m/s over 300–310 s. Under ratio feedforward the feeder commands do not move at all (largest change 0), so the true flow at the scale is the transport's own response, and it agrees with the closed form of Equation (2.4) to 0.79 t/h away from arrival fronts and to 1.2 t/h within a second of one, as shown in Figure 2.2 (ratio feedforward, thick solid line). The flow falls to 675 t/h as the belt slows and climbs in three stairs as ore laid at the new speed arrives, C's first and A's last. Ore that A laid at 60 s arrives at 139.2 s, after 79.2 s of travel, where a fixed delay at the entry speed would predict 60 s. Back at 900 t/h the belt carries 166.7 kg/m, and the speed-up to 2.5 m/s then delivers a surge to 1497 t/h ($3.6 \times 2.5 \times 166.7$) for the same reason. A master PI (dashed line) reads the dip as a shortfall and raises every feeder, so the heavier ore then overshoots to 1007 t/h during the slow stretch.

> **TRY IT IN THE APP** **6 · The Belt Changes Speed: Master PI Against MPC**
>
> Load the scenario and press Run on the Plant view. Watch the scale as the belt slows from 2.0 to 1.5 m/s at 60 s: the flow falls toward 675 t/h before any feeder has moved. The master PI reads the dip as a shortfall and raises every feeder, so the heavier ore later overshoots to about 1007 t/h on the slow belt.

2.4.3 What One Total Cannot Say

With three feeders in fixed proportion, a shortfall of a given size moves the total identically whichever feeder causes it, so the total cannot separate the three gains. In case 7 B's true gain falls 20 % over a minute while its drive runs on. The master PI restores the total to 896 t/h by 300 s by raising all three feeders by $1/0.9$, which leaves a true blend of 22.2:44.4:33.4 against the 22.2:44.4:33.3 that the arithmetic $200:400:300$ predicts, while the controllers' own estimate still reads 20.0:50.0:30.1. The blend becomes observable only when one feeder moves alone, which is the purpose of the bump test of Section 2.5.3. Delay differences let a separately excited feeder be identified, but they do not make the blend continuously observable in fixed-ratio operation. The loss-in-weight feeders of Volume 12, Chapter 12 meet the same limit, where a held weighed flow is not a held product concentration.

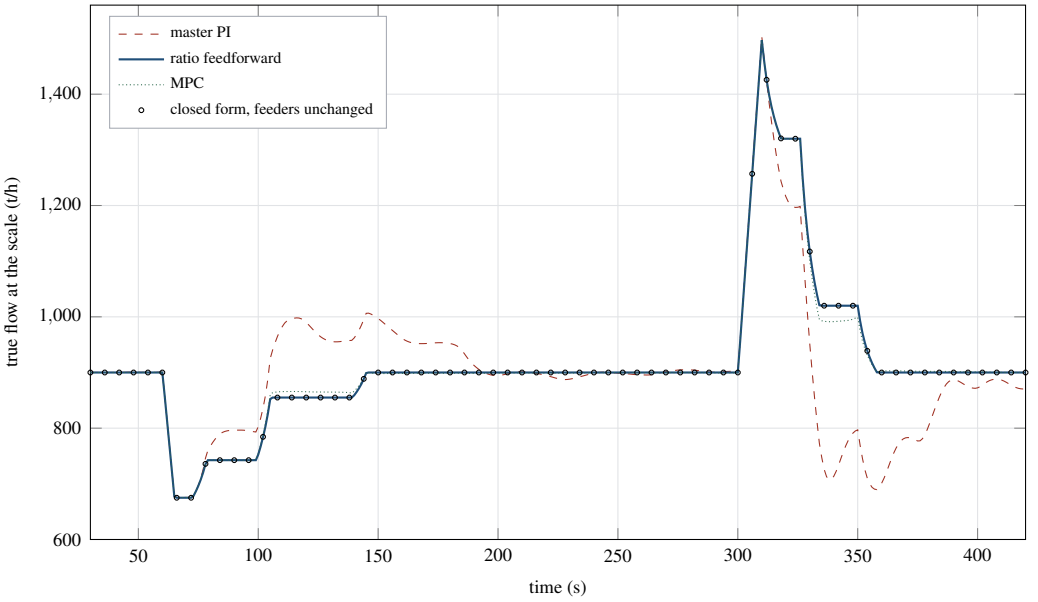


Figure 2.2 True flow at the scale in case 6, belt slowed from 2.0 to 1.5 m/s at 60 s and speeded to 2.5 m/s at 300 s. Master PI dashed, ratio feedforward thick solid, MPC dotted. The circles are Equation (2.4) with the feeders held at their design flows. Ratio feedforward and MPC lie on or near the closed form, while the master PI chases the dip.

> **TRY IT IN THE APP** **7 · A Feeder's Calibration Drifts**
>
> Press Run on the Blend view and compare the estimated and true arriving bars. By 300 s the master PI has the total back at 896 t/h, yet the true blend is 22.2:44.4:33.4 while the estimate still reads 20.0:50.0:30.1. The warning BLEND UNCONFIRMED latches at 147 s and stays, although the total looks right.

2.5 Implemented Control System

Every structure reads one measurement packet, published once per physics step: the scale reading Q_m, the belt speed and its run status, each drive's actual speed, run permissive and trip, and the empty switches where installed. Known geometry, the operator's targets, the configured limits, the drive lag and slew, and the calibrations $\hat{k}_i$ are also available. The true gains, the per-source flows at the scale, the silo inventories, the belt's true load and the scenario's event schedule are not. A test greps the four controller files for truth reads, and a behavioral check changes C's hidden gain from 100 to 90 % at 120 s. The scale reading first differs at 129.4 s and the commands of the master PI and the MPC at 130.1 s, the next scan, while ratio feedforward never changes. A different hidden inventory leaves every command identical for 300 s.

2.5.1 Supervisor

A supervisor shared by every structure sets the target and the recipe before any controller acts. It normalizes the recipe weights to fractions r_i, rejects an all-zero recipe and keeps the last valid one, and computes the effective target of Equation (2.6).

$$Q_{\text{eff}} = \min\Big(Q_{\text{SP}},\ \min_{r_i>0} \frac{100\,\hat{k}_i}{r_i},\ 3.6\,v\,\lambda_{\max}(1-m)\Big) \tag{2.6}$$

At 2 m/s the recipe limit is 1600 t/h (B's 800 t/h is half of it) and the belt limit is $1440 \times 0.98 = 1411.2\,\text{t/h}$, and at 1.2 m/s the belt limit is $864 \times 0.98 = 846.7\,\text{t/h}$. With the supervisor Off the request

passes untouched. A source is unavailable when its drive may not run or its switch has tripped. Under the preserve-recipe policy, the default, a missing required source sets the target to zero, because no positive flow satisfies the recipe, while the temporary-recipe policy renormalizes the available fractions (20:50:30 without B is 40:0:60) and labels the result. When the source returns, the original recipe comes back on a 60 s ramp. A stopped belt interlocks every feeder in every mode, including Manual. On restart the target ramps up at 30 t/h per second, and for one belt transit after the ramp the master's feedback waits, because the ore the stop retained arrives first and a PI reading it would hold the feeders back. An unconfirmed feed loss freezes the target where it was.

2.5.2 Ratio Feedforward and the Master PI

Ratio feedforward, the ratio structure of Section 1.4.6, allocates $u_i = r_i Q_{\text{eff}}/\hat{k}_i$ through the calibrated inverse, which gives 45, 56.25 and 45 % at the design point. The master PI is one loop on the total that sets a demand D, limited before allocation and split by the recipe, as shown in Equation (2.7).

$$D = \mathrm{PI}(Q_{\text{eff}} - Q_m) \in [0, D_{\max}], \qquad u_i = \frac{r_i D}{\hat{k}_i} \tag{2.7}$$

where $D_{\max}$ is the least of 1800 t/h, the recipe's calibrated capacity and the belt's rating at the measured speed $3.6v\lambda_{\max}$ with the supervisor On, and the feeders' summed capacity $\sum 100\hat{k}_i$ with it Off. The PI is the framework's velocity-form controller of Section 1.6.2, reverse acting in that section's convention (action $a = +1$, so the demand rises while the scale reads below the target) on the error $Q_{\text{eff}} - Q_m$, with no derivative and a scan of 1 s. Its back-calculation anti-windup acts on the integral only, with tracking time T_i, so the emitted demand is clamped and the integral follows the limited value. Limiting the one demand keeps the recipe, where clipping three speeds independently would change it silently. In Manual, ratio and MPC operation the PI tracks $\sum_i \hat{k}_i u_i$, the demand the applied commands represent, so every transfer is bumpless (largest command change 0.000 % in the two seconds after each of five switches at 900 t/h). Three PIs acting on the same scale error would fight, which is why a single master sets one demand.

2.5.3 Estimator, Blend Confidence and Bump Test

The controllers keep their own belt, a separate instance of the same grid filled every second from each drive's measured speed times its calibration (zero while the drive may not run or its switch has tripped) and advanced at the measured belt speed, the material tracking that the tableting line of Volume 5, Chapter 6 applies to its powder. From it come the predicted arrivals $\hat{a}_i$, their sum $\hat{Q}$, the same sum through the scale's filter $\hat{Q}_m$, and the residual $Q_m - \hat{Q}_m$. A bounded aggregate correction follows the residual, as shown in Equation (2.8).

$$\hat{d} \leftarrow \hat{d} + \left(\mathrm{clamp}(Q_m - \hat{Q}_m,\ \pm B) - \hat{d}\right)\left(1 - \mathrm{e}^{-\Delta t/\tau_{\hat{d}}}\right) \tag{2.8}$$

where B is 25 % of $\hat{Q}_m$ and at least 50 t/h and $\tau_{\hat{d}} = 20$ s. The correction freezes while the belt is stopped, the predicted flow is under 90 t/h, a feed loss is unconfirmed or a bump test runs. It is added to the MPC's predicted total and never assigned to a feeder. A warning, BLEND UNCONFIRMED, latches when $|\hat{d}|$ exceeds 3 % of the predicted flow for 20 s, and a total that matches again does not clear it. Only a completed bump test or a change of calibration does. Without empty switches, a measured flow 20 % (at least 60 t/h) below the model for 3 s raises an unconfirmed feed-loss alarm that names no source.

The bump test holds the master and first lets the longest belt transit plus 30 s of settling pass. It records a 60 s baseline, steps one drive by 10 %, waits that feeder's transport delay plus 30 s, and records a second 60 s window. It reports $\hat{k}_i = (\bar{Q}_{\text{after}} - \bar{Q}_{\text{before}})/(\bar{s}_{\text{after}} - \bar{s}_{\text{before}})$. This prevents earlier master corrections arriving on the belt from being mistaken for the test step. Reset clears the momentary test selection. The operator applies the result by typing it into the calibration, which also rescales the estimated ore of that source already in transit.

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

2.5.4 Model Predictive Control

Every 2 s the MPC plans the three commands over $H = 240\,\mathrm{s}$ ($N = 120$ steps), which covers the 120 s farthest travel at the minimum running speed of 1 m/s plus settling. The moves are blocked in twelve blocks that lengthen with distance into the future, four single steps first and a final block of 75 steps. The measured belt speed is held over the horizon, and no speed change, demand change or empty silo is previewed. With v held, every predicted quantity is a known linear function of the decisions with fractional delays, as shown in Equation (2.9).

$$\begin{aligned} s_i(\tau) &= s_i(0)\,\mathrm{e}^{-\tau/\tau_d} + \sum_j u_{ij}\left[g(\tau-\tau_j^0) - g(\tau-\tau_j^1)\right] \\ \lambda_i(x,t) &= \begin{cases} \hat{k}_i\, s_i(t-\delta)/(3.6v), & t \ge \delta = (x-x_i)/v \\ \hat{\lambda}_i(x-vt), & t < \delta \end{cases} \\ \hat{Q}(t) &= 3.6\,v \sum_i \lambda_i(x_s,t) + \hat{d} \end{aligned} \tag{2.9}$$

where $g(x) = 1 - \mathrm{e}^{-x/\tau_d}$ for $x > 0$ and zero otherwise, block j spans $[\tau_j^0, \tau_j^1)$, $\hat{\lambda}_i$ is the estimator's belt, and a feeder whose switch has tripped has a zero gain. The ore already on the belt is the free response. Nothing is relocated when v changes, because the next decision rebuilds the map from the estimated belt at the new speed. The objective weighs the scaled total error, the scaled blend mismatch and the moves, as shown in Equation (2.10).

$$\begin{aligned} J = \sum_{k=1}^{N} \Bigg[& w_Q \Big(\frac{\hat{Q}_k - Q_{\mathrm{eff}}}{Q_n} \Big)^2 \\ & + w_B \sum_i \Big(\frac{\hat{a}_{i,k} - r_i \sum_j \hat{a}_{j,k}}{Q_n} \Big)^2 \Bigg] \\ & + w_{\Delta u} \sum_{\text{blocks}} \sum_i \Big(\frac{\Delta u_i}{100} \Big)^2 + \rho \sum_{p=1}^{3} \varepsilon_p^2 \end{aligned} \tag{2.10}$$

with $Q_n = 900\,\mathrm{t/h}$. The blend term uses the source-flow mismatch rather than a fraction, so it stays well behaved near zero flow, which is also why the throughput metrics, not the blend term, distinguish a stopped plant from a good one. The constraints are $0 \le u_{ij} \le 100$ (zero for an unavailable feeder), a move per block of at most $\dot{s}\,\tau_d = 10\,\%$ (which keeps a single move inside the drive's slew, although two successive moves in the same direction can still meet it, a nonlinearity the linear map leaves out), and a loading bound at every step just past each of the three discharge points, $\lambda(x_p,t_k) \le \max(\lambda_{\max}(1-m), \lambda_{\mathrm{free}}(x_p,t_k)) + \varepsilon_p$ with $\varepsilon_p \ge 0$. The loading anywhere on the belt is the loading at the last discharge point upstream, carried downstream unchanged, so three points cover the conveyor. The free response λ_{free}, every command at zero from now, is the least loading any plan can reach, so ore already over the cap after a stop or a slowdown bounds only itself and is reported as unavoidable rather than consuming the slack. The condensed problem has 39 variables (36 block moves and three slacks) and up to 507 inequality rows, and it is solved by the shared dense QP, a dual method for strictly convex quadratic programs [16], within 600 iterations, with a ridge of 10^{-9} scaled by the decision part of the Hessian only. Move blocking follows the usual receding-horizon practice [17].

A plan is applied only when the QP returns optimal within the iteration budget, every planned command is finite and inside its bounds, and the first move is inside the slew cap. Anything else is a fallback to the master PI under the same supervisor, held for at least 10 s and marked in the log. The prediction is built in the app rather than on the shared state-space MPC of Section 1.6.3 because a belt with these delays would need roughly 50 to 100 delay states per decision. The district heating network of Volume 9, Chapter 6 faces the same choice, with a delay that follows the manipulated flow.

Table 2.3 The Nine Cases

#	Setup and events	Teaching point
1	Manual, +10 % steps on A, B, C	three delays, three gains
2	Manual at 27.25:68.25:41 % speeds	speed ratio is not mass ratio
3	Master PI 1.5, 15 s, target 900 to 1080 t/h at 30 s	delayed overcorrection, retuning
4	Recipe 20:50:30 to 30:40:30 at 30 s	unequal arrivals in a blend change
5	Silo B at 37.5 t, refill 100 t at 480 s	empty-silo policies, detection
6	Belt 2.0 to 1.5 at 60 s, 2.5 m/s at 300 s	speed factor, variable delay
7	B's true gain −20 % over 60–120 s	total right, blend wrong, bump test
8	Target 1800 t/h at 30 s, stop 300–360 s	feasibility, anti-windup, restart
9	1000 t/h at 60 s, belt 1.2 m/s at 300 s, B at 75 t, refill at 780 s, supervisor Off	an infeasible shift, graded on policy

Three models are kept apart: the plant of Section 2.3, the estimator's belt filled from calibrations, and the MPC's linear map, which shares the plant's drive lag and slew but not its true gains.

2.6 Nine Cases, Measured

The nine cases appear in Table 2.3. All start from the design point with the master PI at 0.7 and 30 s unless the row says otherwise. All figures quoted here were measured headless, with the engine of Section 1.6.1 advanced through the same steps the browser takes, from a fresh load at the shipped seed. The exported CSV file pairs the measured signals with `truth_` columns that no controller reads (Section 1.6.4). On-spec tonnes are true tonnes that crossed the scale above 90 t/h with every fraction within 2 pp of the requested recipe. The total-flow integrated absolute error (IAE) is $\int |Q_m - Q_{\mathrm{SP}}|\, dt$ in tonnes over the stated window, including any shortfall the physics makes unavoidable.

2.6.1 Code Verification

The transport is checked against its analytic limits. A 1 s pulse arrives after distance over speed to 1 ms from every feeder at 1, 2 and 3 m/s, and under a speed ramp at 73.870 s against the integrated travel's 73.900 s, where a fixed delay at the entry speed would say 60 s. A pulse's centroid is unchanged to 10^{-9} m across a speed change. A spike run with a stop during which the feeders keep running, a restart to 2.5 m/s, an empty silo and a refill closes the belt balance to 3.7×10^{-16} of the handled mass. Closed independently from the silo states, the belt inventory and the score's delivered tonnes, the balance leaves residuals of 3×10^{-15} to 4×10^{-14} relative over case 5 with its refill, case 8 with the supervisor Off and the shipped challenge, whose 8.065 t of spill is counted. Lined out, the true flow and drive speeds stay put to 3.8×10^{-10} for 29 s.

The MPC's prediction reproduces a closed-form step response on an empty belt to 1.5×10^{-6} t/h at 2.0 and 1.2 m/s. Played forward without feedback on a duplicate of the belt it planned from, at 0.1 s with the drives' lag and slew, the accepted plan of case 4 at 32 s reproduces the predicted scale flow with a median error of 0.000 t/h and a worst error of 15.64 t/h, at the edges of the estimator's one-second deposits, and the loading past C within 0.10 kg/m. With 250 kg/m of A's ore placed on the belt, the plan stays optimal and reports 54.00 kg/m as unavoidable. A forced solver failure at 60 s hands the plant to the PI with no command change, and the MPC takes back at 71.1 s. Over the suggested challenge run headless the median solve takes 0.61 ms and the slowest 2.21 ms over 600 solves. In a desktop browser the 1200 s challenge plays in 20.0 s of wall time at 60× at 1366×768 and on a 390×844 phone, with the slowest solve between 8.7 and 14.1 ms, a wall-clock figure that varies from run to run.

2.6.2 Allocation and Tuning

In case 2 speeds in the recipe's own proportion, 27.25:68.25:41 %, deliver 902 t/h at a true blend of 12.1:60.6:27.3, with B 10.6 pp over its share because it has the largest gain. Setting the speeds through the calibrations (45:56.25:45 %) brings the blend to 20.0:50.0:30.0, and the 60 s steady band of 2 % and 2 pp is first completed at 183 s, because the new mix needs a full belt transit to replace

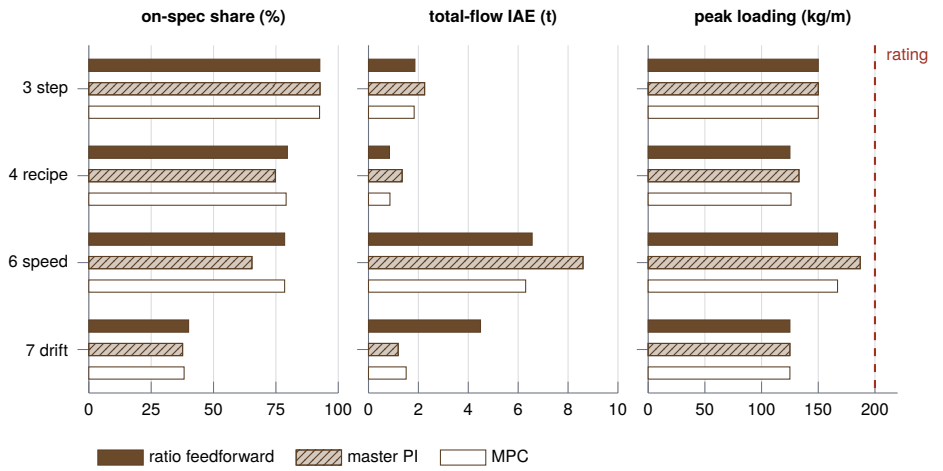


Figure 2.3 Three columns of Table 2.4, means over three paired seeds: ratio feedforward solid, master PI hatched, MPC open. The dashed line is the belt's 200 kg/m loading rating. Notice that the master PI loses on-spec share in cases 4 and 6, where it chases dips the transport makes, yet has the lowest IAE in case 7, where no structure gets past 40 % on spec.

the old one before the 60 s hold can start. In case 3 the shipped master PI at $K_c = 1.5$ and $T_i = 15$ s integrates during the wait for the slow paths, peaks at 1245 t/h on a 1080 t/h target and is still outside 2 % 594 s after the step (IAE 5.47 t to 630 s). At 0.7 and 30 s it peaks at 1081 t/h and stays inside 2 % from 61 s after the step (IAE 2.25 t). A sweep of 56 tunings ranks $K_c = 1.0$, $T_i = 30$ s lowest in IAE (1.81 t) at the price of a 24 t/h overshoot and 105 s to settle, so the shipped tuning trades a little IAE for a 1 t/h overshoot.

2.6.3 Controller Comparison

Table 2.4 compares the three structures on identical measurements, noise and events, and Figure 2.3 draws three of its columns, with no preview for any, as means over three paired seeds whose on-spec shares agree within 1 pp. The question is whether the MPC shows a reproducible benefit on the blend-change and speed-change cases. Against the tuned master PI it does. In case 4 the on-spec share rises from 74.8 to 79.1 % and the IAE falls from 1.35 to 0.86 t, and in case 6 from 65.4 to 78.5 % and from 8.60 to 6.30 t, with no higher peak loading (126 against 133 kg/m, 167 against 187 kg/m). Against ratio feedforward it does not. In case 4 ratio feedforward reaches 79.6 % and 0.84 t, slightly better than the MPC, and in case 6 both reach 78.5 % while the MPC's IAE is 4 % lower (6.30 against 6.56 t). With exact calibrations the PI's feedback is the handicap, because it chases dips that the transport makes and nobody can remove. The total dips to about 810 t/h between 70 and 90 s in case 4 under every structure, because B's cut arrives after 30 s and A's increase only after a minute. Where the calibration is wrong (case 7, to 300 s) the PI holds the total better than the MPC (IAE 1.19 against 1.51 t), likely because the MPC sees the shortfall only through $\hat{d}$ and its 20 s filter, ratio feedforward leaves the total 91 t/h short (4.49 t), and no structure fixes the blend (37.6 to 40.0 % on spec).

Table 2.4 Structures Compared, Means Over Three Paired Seeds

Case (window)	Structure	On-spec %	Blend pp	IAE t	Peak kg/m	Travel %
3 step (630 s)	ratio FF	92.6	0.24	1.86	150	29
	master PI	92.7	0.27	2.25	150	75
	MPC	92.5	0.26	1.83	150	38
4 recipe (300 s)	ratio FF	79.6	1.80	0.84	125	34
	master PI	74.8	2.31	1.35	133	84
	MPC	79.1	1.87	0.86	126	41
6 speed (540 s)	ratio FF	78.5	1.39	6.56	167	0
	master PI	65.4	3.06	8.60	187	270
	MPC	78.5	1.46	6.30	167	18
7 drift (300 s)	ratio FF	40.0	3.19	4.49	125	0
	master PI	37.6	3.30	1.19	125	28
	MPC	38.2	3.30	1.51	125	17

Blend is the mass-weighted worst-source error against the requested recipe, peak the highest true loading anywhere on the belt, and travel the summed command movement of the three feeders.

+ DID YOU KNOW? Feedback That Makes the Blend Worse

From the simulator: in Table 2.4, adding the master PI in case 6 lowers the on-spec share from 78.5 % under ratio feedforward to 65.4 %, and the feeders' summed command travel rises from 0 to 270 %. The PI chases a dip that only the transport made.

2.6.4 Calibration, Empty Silos and Capacity

In case 7 the aggregate correction settles at $\hat{d} = -100\,\mathrm{t/h}$ and BLEND UNCONFIRMED latches at 147 s. Starting the B test at 300 s produces its result at 570 s: a 64.51 t/h change for a 10 % speed step, or $\hat{k}_B = 6.451\,\mathrm{t/h}$ per percent against a true 6.40. Applying 6.45 on the calibration grid reaches the sustained blend and flow band at 724.5 s under PI and 666.2 s under MPC. Starts from 295 to 305 s give 6.429 to 6.479. Starts immediately after the warning, from 148 to 152 s, give 6.443 to 6.450. The clearing and settling stage makes the early test usable too, at the cost of about 270 s with the master held.

In case 5 silo B empties at 300.1 s (37.5 t at 450 t/h is 300 s) and its switch is read 1.0 s later. B's last ore keeps arriving for half a minute. Under preserve recipe production stops at 301.1 s, restarts on the ramp when B feeds again at 481.2 s, and the plant delivers 131.2 t in 720 s, 88.1 % on spec and 48.8 t short. Under temporary recipe the feeders go to 40:0:60 (A 360 t/h, C 540 t/h), and the plant delivers 179.7 t, 3.1 t short, of which 58.8 % is on the original recipe. The three structures agree within 2.7 pp on both policies. Without the switches the feed-loss alarm comes at 334.1 s, 34.0 s after the silo emptied, and cannot name the silo.

In case 8 the supervisor holds the 1800 t/h request at 1411 t/h, the belt limit less the margin, with the recipe at 19.9:50.0:30.1 at 200 s, a peak of 196 kg/m and no time over the rating through the stop and restart. With the supervisor Off the demand runs to 1800 t/h, B saturates at 100 % while A and C reach 90 %, and the plant delivers 1700 t/h at 21.2:47.1:31.8 with the belt at 236 kg/m, over its rating for 499 s. The integral tracks the limited demand. In a test that holds the demand at its 1800 t/h limit for 270 s and then returns the target to 900 t/h, the demand leaves the limit at the next scan and the scale is back within 2 % 169.8 s later.

2.6.5 The Challenge

The shift of case 9 asks for 1000 t/h from 60 s, slows the belt to 1.2 m/s at 300 s, where its rating carries 864 t/h, and lets B run dry before a truck brings 100 t at 780 s. The target is infeasible by design, and in ten runs across structures, supervisor settings and master tunings none holds 880 t/h for 120 s inside the rating (best 862.6 t/h). The verdict therefore grades recognition and policy over 1200 s: no spill and at most 30 s above the rating, a correct empty-silo policy (at most 3 t fed while B

Table 2.5 Challenge Verdicts Over 1200 s

Structure, supervisor	Policy	Over s	Spill t	Useful t	Verdict
master PI, Off (shipped)	preserve	516	8.07	193	fail
master PI, On	preserve	52	0	198	fail
master PI, On	temporary	172	0.25	182	fail
ratio FF, On	preserve	18	0	197	pass
ratio FF, On	temporary	148	0.43	210	fail
MPC, On	preserve	0	0	196	pass
MPC, On (suggested)	temporary	0	0	216	pass

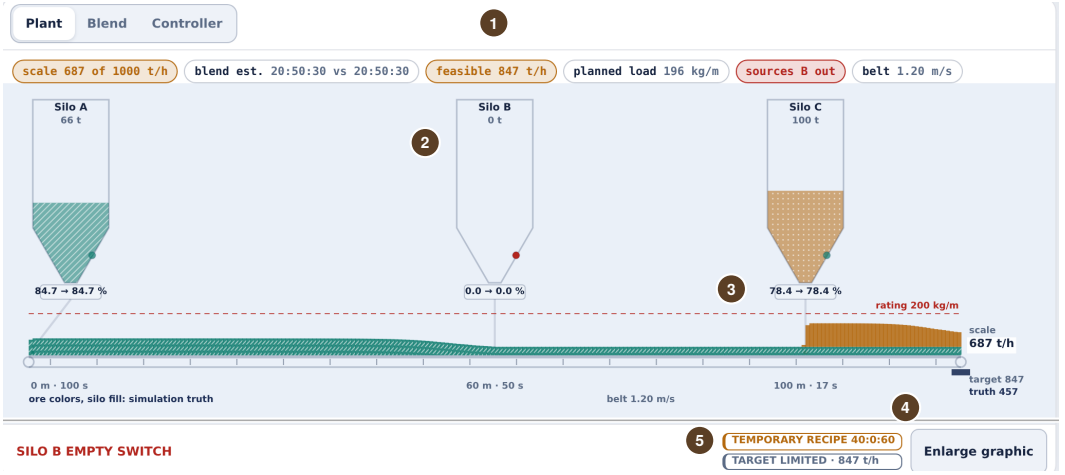


Figure 2.4 The challenge at $t = 640$ s on the Plant view with the suggested settings (MPC, supervisor On, temporary recipe 40:0:60 labeled), cropped to the plant view and its status line, with the belt at 1.20 m/s. (1) The effective target is 847 t/h against 1000 t/h asked, and the MPC's planned loading is 196 kg/m. (2) Silo B is empty. (3) A and C run at 84.7 and 78.4 % with silos at 66 and 100 t. (4) B's last ore has just crossed the scale, so the true flow is 457 t/h while the filtered reading is 687 t/h. (5) The status line labels the temporary recipe and the limited target.

is out, or at least 80 % of the fed tonnes on the labeled temporary recipe), at least 180 t of useful ore, and the last 120 s within 2 % of the effective target and 2 pp of 20:50:30. Table 2.5 gives the rows, and three paired seeds return identical verdicts.

The MPC is the only structure that keeps the belt inside its rating under the temporary recipe. When B empties at 589.4 s on the slow belt, the MPC raises A at once but holds C at or below 42.4 %, because C's extra ore would land on B's last ore still passing C's chute, and it starts raising C at 622 s, when that ore has passed, in moves of 10 % per decision to 78.4 %. The peak loading stays at 199 kg/m. Under the same policy the master PI and ratio feedforward raise C at once and spend 172 and 148 s above the rating in the shift. Under preserve recipe ratio feedforward also passes with 18 s over the rating, so the MPC's clear advantage in this plant is constraint handling on ore already placed, not tracking. Holding C back has a cost, as shown in Figure 2.4. B's last ore crosses the scale at 640 s while A's increase is still about 50 s away on the belt, so the true flow drops to 457 t/h (the filtered reading shows 687 t/h) and recovers to 816 t/h by 700 s as C's and then A's extra ore arrive.

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

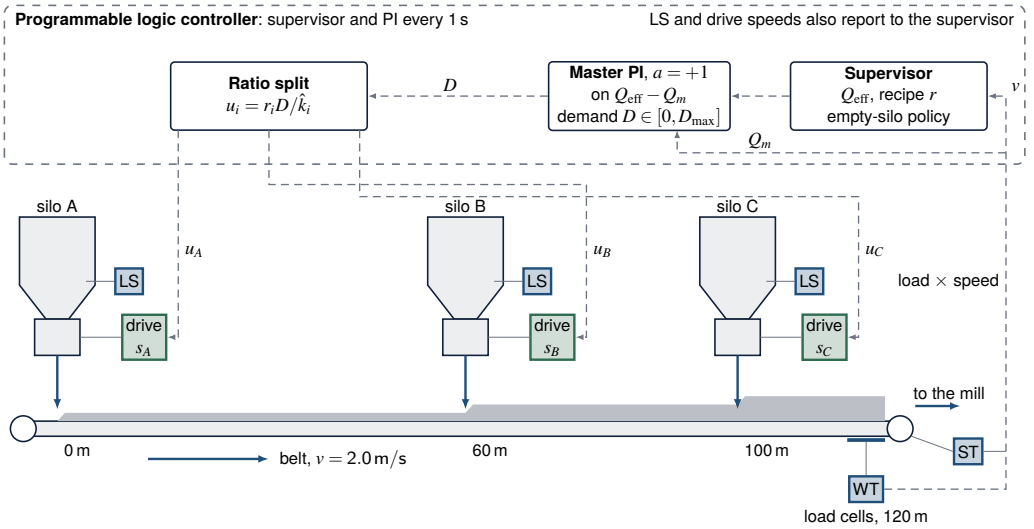


Figure 2.5 The master loop of the simulated conveyor drawn with its hardware. The belt scale's load cells WT and its speed sensor ST give the flow Q_m, the empty switches LS and the drives' own speed feedback complete the measurement packet, and the supervisor and the master PI, reverse acting with $a = +1$, set one demand that the recipe and the calibrations split into three speed commands for the variable-frequency drives. Notice that the only measurement of mass sits 120 m from feeder A, after every feeder.

2.7 Technology in Practice

The structures of Section 2.5 run on one belt scale, three variable-speed feeder drives, the silos' empty switches and one controller. Figure 2.5 draws the master loop with that hardware.

What Measures It

The model of Section 2.3 reduces the scale to the mean loading over a 1 m span times the measured speed, with a bias, a colored noise of 2 t/h and a 2 s transmitter filter.

> **TECHNOLOGY** **Sensor: A Belt Scale**
>
> A belt scale weighs a fixed length of moving belt. The rollers under that length rest on one or more load cells, which read the weight of ore per unit length, and a speed sensor reads the belt's speed. The product of the two is the mass flow [11]. The international recommendation calls the speed sensor a displacement transducer and drives it from the side of the belt that carries no product, so that slip cannot affect it. It sets the empty-belt zero over a whole number of belt revolutions [10]. A span error scales every reading by the same fraction, which the bias b of the model stands for. A zero error adds the same flow to every reading, even on an empty belt.

The international recommendation gives a weigher's errors on totalized loads, not on an instantaneous rate, so the model's noise and filter are teaching values. The empty switches are read once a second, and the drives report their actual speeds, run permissives and trips.

What Moves It

Each drive is modeled as an exact first-order lag of 2 s whose step is capped at 5 %/s, and the feeder delivers its true gain times the speed, with no dead band, flooding or bridging.

TECHNOLOGY Final Element: A Variable-Frequency Drive on Each Feeder

Each feeder is turned by an alternating-current motor whose speed a variable-frequency drive sets. The drive rectifies the line power to direct current, smooths it, and an inverter switches it back to alternating current at the frequency and voltage the commanded speed needs. A drive spins the motor only as fast as needed, ramps the speed up and down to spare the mechanism, and can brake the load [11]. A feeder that delivers less ore per turn, as in case 7, shows nothing on its drive's speed signal, which keeps reading the commanded speed.

The model's lag and slew stand for the drive's ramp, and its hidden gain factor g_i stands for the ore a feeder delivers per turn, which the drive cannot see.

Where the Controller Runs

Programmable logic controllers were designed first for discrete tasks such as conveyor belt management, and they now run PID loops as readily, with variable-speed drives among their outputs [11]. Such a controller works as a local controller under a supervisory scheme or as the primary controller of a smaller system, and supervisory and lower-level loops run with cycle times from milliseconds to minutes [18]. The app integrates every 0.1 s, runs the supervisor and the controllers every 1 s and the MPC every 2 s.

IN PRACTICE How a Belt Weigher Is Verified

OIML R 50 tests a belt weigher in place, fully assembled where it will work. Before a product test the conveyor runs, preferably loaded, for at least 30 min at its nominal speed. A control instrument kept near the weigher finds the true mass of each test load, weighed before or after it crosses the belt, to at least a third of the permitted error. For tests without the conveyor, standard weights and a device that simulates the belt's travel stand in for the ore and the belt [10].

In such a controller the chapter's structure is a supervisor that limits the target and keeps the recipe, a ratio block that divides the demand through the calibrations, one PI block whose output is limited before allocation and whose integral tracks the applied demand, the estimator's belt, and the MPC above them.

What Is Modernizing It

The first development is a belt whose speed is itself controlled. Speed control regulates a conveyor's speed to match a varying feed rate, since most conveyors run partly loaded. On a simulated terminal conveyor it raised the belt's utilization from 45 to 69 % and cut its electricity by over 10 %, while too steep a speed ramp raises the belt tension that breaks belts at their splices [8]. Every such change is the flow disturbance of case 6, and it moves each feeder's delay. The second is analysis on line. In a pilot plant an online X-ray fluorescence analyzer read the copper in a suspension with a relative error of 6 % [19], the kind of measurement that would make the blend of Section 2.4 observable. The third is model predictive control, widely applied to multivariable plants with delays and constraints [15], whose clearest gain in this chapter is the loading constraint of the challenge.

2.8 Limitations

The geometry, gains, rating and noise make one self-consistent teaching design point and are not fitted to a plant. The belt is rigid and its ore never slides, delivery is linear in speed, the blend is a mass fraction with no assay, and the silos are bins with an instantaneous refill and a switch that trips only when the cone is empty. The scale has no dropout event, no zero drift and no belt-tension error, although the weigh-length element is the largest error source of a real belt weigher [13].

OPEN CHALLENGE **Measuring the Blend, Not Only the Total**

Devices for online elemental analysis are widely available, yet mineral processing still controls mostly from laboratory samples, which are slow, expensive and error prone [19]. A 2026 review of ore blending names the barriers to its industrial use: data from many sources that are hard to join, models too slow for real-time control, and models that do not carry over from one ore to the next [20]. In this chapter the arriving blend is seen only by the truth log. Progress would be an analysis on the belt fast and accurate enough for a blend loop to close on.

The controllers know the drive lag and slew exactly, which leaves the gains, the scale bias and the events as their only sources of mismatch. The estimator fills its belt once a second, which leaves prediction errors of up to 15.6 t/h at the edges of a deposit (median zero), and the true loading in the challenge peaks at 199 kg/m against the MPC's cap of 196 kg/m. The bump test carries about 2 % of noise error (6.45 against 6.40 for a 300 s start), and its result still varies with the moment the test starts. The MPC holds the measured belt speed over its horizon, is told of no future event, and is not shown to beat ratio feedforward when the calibrations are right. Its advantage over the tuned PI in cases 4 and 6 comes from not reacting to dips it can predict, and its one clear win is constraint handling in the challenge. Every safety and capacity limit is a teaching value, and a simulated pass is not a conveyor design or a weighing certification. Case 4 is stylized, since a real recipe change would be scheduled against the transport, but it isolates the off-spec ore that unequal arrivals cause.

2.9 Conclusion

The ore feed simulator puts three feeders behind one belt scale and tracks their ore on a grid fixed to the belt, so arrivals follow the integrated travel and a slowdown from 2.0 to 1.5 m/s cuts the scale's flow from 900 to 675 t/h with no feeder moving. Speeds in a recipe's proportion deliver 12.1:60.6:27.3 instead of 20:50:30, a master PI tuned to its slowest path reaches a step with almost no overshoot, and a feeder that loses 20 % of its gain leaves the total at 896 t/h and the blend wrong until a bump test finds 6.45 against 6.40. On identical inputs the MPC beats the tuned master PI through a recipe change (79.1 against 74.8 % on spec) and a speed change (78.5 against 65.4 %), comes no closer to the recipe than ratio feedforward in either, and is the only structure that keeps the belt inside its rating through the infeasible shift under a temporary recipe. The model is lumped and its parameters are teaching values. Further work would add a cross-belt analyzer, scale dropout and drift, a belt-speed schedule shared with the controllers, and stockpile blending upstream of the silos. The simulator shows that a total on target is a measurement of the total and nothing more.

KEY POINTS

- The scale sees each feeder only after its ore has traveled to it, about a minute from A and ten seconds from C.
- A belt scale multiplies its load by the belt speed, so a slowdown to 1.5 m/s cuts the flow it reads from 900 to 675 t/h at once.
- Speeds in a recipe's proportion deliver 12.1:60.6:27.3 instead of 20:50:30, because the feeders' gains differ.
- A feeder that loses 20 % of its gain leaves the total at 896 t/h and the blend wrong until a bump test finds it.
- The MPC's clear win is keeping the belt inside its rating, not tracking.

Where else this applies. Feeder ratio control on tracked material in Volume 5, Chapter 6, a delay fixed in meters in Volume 16, Chapter 10, and MPC on a delay that follows the flow in Volume 9, Chapter 6.

References

[1] D. Hodouin, S.-L. Jämsä-Jounela, M. T. Carvalho, and L. Bergh, "State of the art and challenges in mineral processing control," *Control Engineering Practice*, vol. 9, no. 9, pp. 995–1005, 2001.

[2] D. Hodouin, "Methods for automatic control, observation, and optimization in mineral processing plants," *Journal of Process Control*, vol. 21, no. 2, pp. 211–225, 2011.

[3] L. Keviczky, J. Hetthéssy, M. Hilger, and J. Kolostori, "Self-tuning adaptive control of cement raw material blending," *Automatica*, vol. 14, no. 6, pp. 525–532, 1978.

[4] T. Westerlund, H. T. Toivonen, and K.-E. Nyman, "Stochastic modelling and self tuning control of a continuous cement raw material mixing system," *Modeling, Identification and Control*, vol. 1, no. 1, pp. 17–37, 1980.

[5] V. de Oliveira, M. Amrhein, and A. Karimi, "Robust gain-scheduled blending control of raw-mix quality in cement industries," in *IEEE Conference on Decision and Control and European Control Conference*, 2011, pp. 4189–4194.

[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] S. Zhang and X. Xia, "Modeling and energy efficiency optimization of belt conveyors," *Applied Energy*, vol. 88, no. 9, pp. 3061–3071, 2011.

[8] D. He, Y. Pang, and G. Lodewijks, "Green operations of belt conveyors by means of speed control," *Applied Energy*, vol. 188, pp. 330–341, 2017.

[9] M. Krstic, "Lyapunov stability of linear predictor feedback for time-varying input delay," *IEEE Transactions on Automatic Control*, vol. 55, no. 2, pp. 554–559, 2010.

[10] International Organization of Legal Metrology, "Continuous totalizing automatic weighing instruments (belt weighers). Part 1: Metrological and technical requirements," International Organization of Legal Metrology, Paris, International Recommendation OIML R 50-1, Edition 2014 (E), 2014. [Online]. Available: https://www.oiml.org/en/files/pdf_r/r050-1-e14.pdf

[11] 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.

[12] J. E. Normey-Rico and E. F. Camacho, *Control of Dead-time Processes*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2007.

[13] A. E. Hidden, "Errors in conveyor belt weigher systems," *Measurement and Control*, vol. 6, no. 6, pp. 245–252, 1973.

[14] V. K. Donis, A. E. Rachkovskii, and V. M. Sin, "How the conveyor belt length affects belt weigher accuracy," *Measurement Techniques*, vol. 47, no. 2, pp. 163–167, 2004.

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

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

[17] R. Cagienard, P. Grieder, E. C. Kerrigan, and M. Morari, "Move blocking strategies in receding horizon control," *Journal of Process Control*, vol. 17, no. 6, pp. 563–570, 2007.

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

[19] L. Rybok, T. Kowoll, A. Thomas, A. Nad, and C. Bachmann, "Online elemental analysis for process control in the mineral processing industry," *TOS Forum*, vol. 2022, no. 11, p. 131, 2022.

[20] Z. Cai, H. Luo, and W. Yang, "The research progress and application status of ore blending technology: A review," *Physicochemical Problems of Mineral Processing*, vol. 62, no. 4, p. 226165, 2026.

2.10 Exercises

2.1 Concept check · 2 min

The belt of Figure 2.1 is lined out at 900 t/h and 2.0 m/s. With every feeder held where it is, the belt is slowed to 1.5 m/s. What does the true flow at the scale do in the first seconds?

(a) It stays at 900 t/h until ore laid at the new speed reaches the scale

(b) It rises, because ore laid at the lower speed is heavier per meter of belt

(c) It falls only after feeder C's delay, when C's ore laid at the lower speed arrives

(d) It falls at once to about 675 t/h, because an unchanged loading passes the scale more slowly

2.2 Concept check · 1 min

In case 7 feeder B's true gain has fallen 20 %. By 300 s the master PI has the total back at 896 t/h and the controllers' estimated blend reads 20.0:50.0:30.1. Which statement is right?

(a) The total cannot show which feeder is short, and only a feeder moved alone can reveal B's gain

(b) The blend is right, because the estimator's belt is filled from the measured drive speeds

(c) The blend is wrong, and the aggregate correction $\hat{d}$ assigns the missing flow to feeder B

(d) The blend is right once $\hat{d}$ has settled, because a settled correction means the model fits

2.3 Concept check · 2 min

In the challenge with the MPC and the temporary recipe, silo B empties at 589.4 s on the 1.2 m/s belt. The MPC raises A at once but holds C at or below 42.4 % until about 622 s. Why does it hold C back?

(a) C's drive is slew limited, so the MPC lets A's drive settle before moving C

(b) B's last ore is still passing C's chute, and C's extra ore on it would load the belt above its rating

(c) The supervisor brings the temporary recipe in on a ramp, so C's share is still rising

(d) The MPC does not learn that B is empty until B's last ore has crossed the scale

2.4 Theory and calculation · 6 min

Use the design point of Table 2.1: feeders at 0, 60 and 100 m, the scale at 120 m, the belt at 2.0 m/s, and feeder flows of 180, 450 and 270 t/h.

(a) Compute each feeder's travel time to the scale and its loading on the belt between its chute and the scale, to 0.1 kg/m. *Checkpoint:* 60, 30, 10 s; 25.0, 62.5, 37.5 kg/m

(b) The belt slows to 1.5 m/s with the feeders unchanged. Compute the loading of the ore laid at the new speed, to 0.1 kg/m, and the flow at the scale, to 1 t/h, when that ore is under the scale and the belt is then speeded up to 2.5 m/s. *Checkpoint:* 166.7 kg/m; 1500 t/h

(c) In case 6 the belt decelerates at 0.1 m/s^2 from 2.0 m/s at 60 s. Compute when ore that feeder A laid at 60 s reaches the scale, to 0.1 s. *Checkpoint:* 139.2 s

2.5 Theory and calculation · 6 min

The recipe is 20:50:30 by mass and the calibrations are 4, 8 and 6 t/h per percent, equal to the true gains.

(a) Compute the three speeds that deliver 900 t/h at the recipe. *Checkpoint:* 45, 56.25, 45 %

(b) Case 2 sets the speeds in the recipe's own proportion, 27.25, 68.25 and 41 %. Compute the total, to 1 t/h, and the blend, to 0.1 % of the total for each feeder. *Checkpoint:* 901 t/h; 12.1:60.6:27.3

(c) With the supervisor On and a 2 % margin, compute the effective target of Equation (2.6), to 0.1 t/h, for 1000 t/h asked on a 1.2 m/s belt, first at 20:50:30 and then under the temporary recipe 40:0:60. *Checkpoint:* 846.7 t/h both

2.6 Theory and calculation · 8 min

In case 7 feeder B's true gain falls from 8 to 6.4 t/h per percent while its calibration stays at 8. The target is 900 t/h at 20:50:30 and the noise is ignored.

(a) Under ratio feedforward, compute the true total at the scale, the true blend to 0.1 %, and the value at which the aggregate correction $\hat{d}$ settles. *Checkpoint:* 810 t/h; 22.2:44.4:33.3; −90 t/h

(b) Under the master PI, compute the demand D that brings the true total back to 900 t/h, the three speed commands, the true blend and the settled $\hat{d}$. *Checkpoint:* 1000 t/h; 50, 62.5, 50 %; 22.2:44.4:33.3; −100 t/h

(c) A bump test steps B's speed by 10 %. Compute the true change at the scale and the gain the test should report. The run reads 64.51 t/h for that step. Compute the error of the reported gain, to 0.1 %. *Checkpoint:* 64 t/h; 6.40 t/h per percent; 0.8 %

2.7 App challenge · 8 min

Load scenario 1, "Three paths, one scale", at the default seed. Before Run, set Scale noise to 0 % (an Advanced disturbance). Run at 5×. At about 10 s raise Feeder A speed from 45 to 55 %, at about 100 s raise Feeder B speed to 66.25 %, and at about 160 s raise Feeder C speed to 55 %. Run to the pause at 300 s and export the CSV, whose time column is `Time (min)`.

(a) For each step, find the first row in which the command column (`cmd_A_pct` and so on) shows the new value, and the first later row in which `scale_tph` is at least 10 t/h above its value in the row before the step. Give the time between them in seconds, to 1 s. *Checkpoint:* 61, 30, 11 s

(b) Read `scale_tph` in the row before each later step and in the last row, to 1 t/h, and give each feeder's change in t/h per percent. *Checkpoint:* 940, 1020, 1080 t/h; 4, 8, 6

(c) Explain why each time in (a) is longer than the travel distance over the belt speed. *Checkpoint:* drive lag, weighbridge fill, filter

2.8 App challenge · 8 min

Load scenario 7, "A feeder's calibration drifts", at the default seed with the master PI as shipped. Run at 5× and export the CSV once you pass 300 s. Its time column is `Time (min)`.

(a) Give the time, in seconds to 1 s, of the first row in which `blend_low` is 1. *Checkpoint:* 148 s

(b) Pause between 295 and 305 s. Read $\hat{d}$ on the Feeder calibration card, to ± 1 t/h. Choose Bump feeder B as the Bump test, run on, and read the result on the same card when it appears, about 270 s later, to 0.01. Type its nearest value on the calibration grid into Feeder B calibration. *Checkpoint:* −100 t/h; 6.43 to 6.48 t/h per percent

(c) Set Feeder B calibration back to its shipped 8.00, press Reset, and check that the Bump test has returned to Off. Run again, start the bump test on B between 148 and 152 s, just after the warning, and read the result. Explain why a test started this early still finds the drifted gain. *Checkpoint:* 6.44 to 6.45 t/h per percent

2.9 App challenge · 8 min

Load the challenge, “Deliver the shift blend” (case 9), at the default seed. Each run lasts 1200 s at 60×. Read the Challenge target tile at the end of each run, and the target on the Total flow card after 300 s.

(a) Run as shipped (master PI, supervisor Off, Preserve recipe). Give the spill, the seconds over the rating, the useful tonnes and the verdict. *Checkpoint:* 8.07 t; 516 s; 193 t; not met

(b) Set the Feasibility supervisor On, the Empty-silo policy to Temporary recipe and the Control structure to MPC, press Reset and run. Give the effective target after 300 s, the seconds over the rating, the useful tonnes and the verdict. *Checkpoint:* 847 t/h; 0 s; 216 t; met

(c) Keep the supervisor On, choose Ratio and Preserve recipe, press Reset and run. Give the seconds over the rating and the verdict, and say why ratio feedforward with Temporary recipe fails. *Checkpoint:* 18 s; met

2.10 Application area · 4 min

A concentrator has three feeders in fixed ratio and one belt scale, and the mill’s grind changes when the blend does. The engineer plans to check the feeder calibrations with bump tests like the app’s. Name three conditions that make a bump test result trustworthy, and say what the test costs the plant while it runs.

2.11 Application area · 3 min

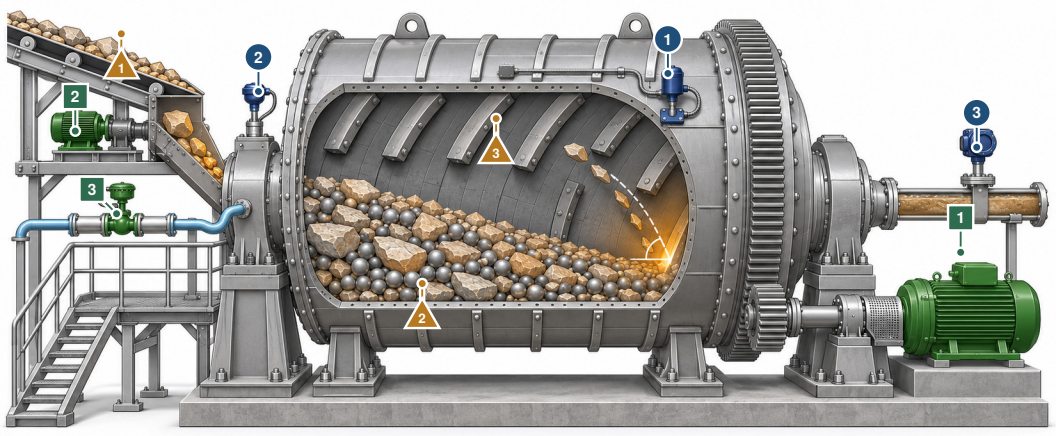
Silo B of a 20:50:30 mill feed runs empty for three minutes before a truck arrives. Case 5 measures two policies. Under preserve recipe the plant delivers 131.2 t in 720 s, 88.1 % on spec and 48.8 t short. Under temporary recipe it delivers 179.7 t, 3.1 t short, 58.8 % on the original recipe. Which would you choose, and what decides it?

2.12 Review the agent · 5 min

An AI assistant reviews case 7 and writes: “After B’s drift the master PI holds the scale within 1 % of the 900 t/h target, so the plant is back on recipe. The BLEND UNCONFIRMED warning is a nuisance alarm. Clear it automatically once the total has stayed within 1 % of its target for 60 s.” Find the flaw and say how you would check it in the app.

CHAPTER 3

Semi-Autogenous Mill Simulator for Impact-Zone Control Against a Moving Constraint



The controller moves 1 the shell speed, through the drive motor 2 the ore feed, through the feeder drive 3 the dilution water, through its valve

The controller reads 1 a sound pickup on the shell, from which the impact angle is inferred 2 the trunnion bearing pressure, which weighs balls, rock and water together 3 the slurry density at the discharge

What disturbs it 1 harder or coarser ore 2 steel balls lost from the charge 3 lifters wearing down

Notice the falling ore landing on the lower edge of the tumbling pile, where it breaks rock best, while no instrument reads where that edge is. Illustration, not to scale.

IN PLAIN TERMS Aiming Falling Steel at the Rock

The main idea. A huge turning drum grinds ore by lifting steel balls and rock and letting them fall. The speed of the drum decides where they land. They break the most rock when they land on the lower edge of the tumbling pile.

The challenge. No instrument can see that edge. The automatic control holds a landing point worked out from the sound of the drum. When some steel balls are lost, the pile grows and the falling balls land short of its edge, while every reading still looks right.

Why it matters. A mill that looks fine on every screen can still grind the ore far coarser than it should.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	9 min	The first page, the fact box of the Introduction, the history boxes of Section 3.2, the try-it box of Section 3.4, the fact box of Section 3.6.7, and the key points
🕮	**Study**	60 min	The whole chapter with its pause points and the twelve exercises
⛰	**Practice**	25 min	The first page, Section 3.6, Section 3.7, Section 3.8 with its open challenge

The control question A semi-autogenous mill grinds best when its media land on the toe of the charge, but the loop can hold only an inferred impact angle while the toe moves with the filling, the ball charge and the liner. When does holding the angle meet the grinding objective, and what do a toe-tracking setpoint and a predictive controller add?

Learning objectives

- Compute the impact angle and its local gain from the detachment condition and the landing point, and explain why the gain grows with speed.
- Locate the power maximum in the filling and explain the overload runaway past it.
- Identify the impact-angle gain from step tests at two speeds and relate a tuning to the ultimate gain.
- Explain why a lost ball charge lifts the toe with every controlled variable on setpoint, and why toe tracking from the bearing pressure aims the wrong way.
- Judge a predictive controller against the loops under a common load limit, and name what its disturbance estimate contributes.

Prerequisites Sections 1.4.2 (grinding circuits), 1.6.2 (velocity form and loop action), 1.4.6 (inferential control and back-off) and 1.4.5 (late analyzers), and Chapter 2 for the feed belt. Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/sag-mill/

See also Volume 4, Chapter 4 for throughput pushed against limits measured late and Volume 4, Chapter 2 for a limit at the peak of a characteristic. Volume 5, Chapter 10 for a constraint that is inferred. Volume 7, Chapter 3 for a gain set by the operating point and Volume 7, Chapter 4 for regulation at the peak of a curve.

3.1 Introduction

Comminution consumes a larger share of the electricity used in a concentrator than any other operation, and the semi-autogenous grinding mill is where most of it is spent [1]. The mill is a rotating drum in which ore breaks itself with the help of a partial charge of steel balls. Above a certain fraction of critical speed the outermost media leave the shell on a parabola and land somewhere inside the mill, and where they land decides what the machine does. Short of the toe of the charge they bury themselves in slurry and grind by attrition. On the toe they strike rock and break it. Past the toe they strike bare liner, and a liner campaign is a capital item measured in months.

DID YOU KNOW? One Percent of the Energy

A 2006 report for the United States government put the energy that semi-autogenous and ball mills spend on size reduction at about 99 trillion British thermal units a year. Grinding mills use only about one percent of that input energy, and they also consume tons of steel balls and liners [2].

The application therefore controls an impact angle ψ, measured from the bottom of the mill toward the toe, using mill speed as the manipulated variable. Considered on its own the loop is an easy one. The gain is large, the open-loop response is a pair of first-order lags with no dead time worth reporting, and a competently tuned controller brings a 5° setpoint step inside half a degree of the new target in 0.32 min.

None of that is what makes a semi-autogenous mill difficult, and three properties of the process do. The first is that the process gain is a function of the operating point, because the detachment condition is $\cos\theta_d = \phi^2$ in the fraction of critical speed ϕ and a cosine is not a straight line. The second is that the mill power has an interior maximum in the volumetric filling, so beyond a certain load more feed returns less grinding, which returns more load. The third, and the reason for this chapter, is that the target itself moves. The quantity that governs breakage is the impact margin of Equation (3.1),

$$m = \psi - \theta_{\text{toe}}, \tag{3.1}$$

in which θ_{toe} is the angular position of the toe of the charge. The toe depends on the filling, which depends on the ore, on the ball charge and on the feed. A controller holding ψ exactly is holding one term of (3.1), and Section 3.6 measures three separate disturbances that move the other term while every trend on the operator's display remains correct.

The scenarios built on that structure teach a specific list of process dynamics and control principles. Operating-point-dependent process gain, the identification of a gain and a time constant from an open-loop step, the tuning that follows from them, time-scale separation between three loops on one unit, an inferential measurement carrying bias and noise, derivative action on a noisy signal, actuator saturation with anti-windup, dead time from a conveyor and from a sampled size analyzer, supervisory setpoint computation, and the distinction between a controlled variable and a control objective. Each of these is standard material in a first course [3–5], and the mill supplies a setting in which the last of them cannot be avoided. The physical fundamentals underneath the model are equally standard, because the plant is three mass balances on the coarse solids, the fine solids and the water, a force balance that decides where a ball leaves the shell, an empirical power correlation and an energy-to-size relation. Industrially the behavior matters because a grinding circuit that is aimed badly consumes its energy without reducing the ore, and because the economic loss appears on the product size and the liner rather than on any loop trend [6].

This chapter works with a semi-autogenous mill simulator in which the controlled variable is one term of a difference, and it quantifies what that structure costs a regulatory loop that is otherwise well tuned. The mill sits downstream of the feeders and belt of Chapter 2, and its feed conveyor dead time appears in the load loop below. Section 3.2 places the model against grinding-mill control practice. Section 3.3 gives the equations and constants needed to rebuild the simulator. Section 3.4 analyzes the two nonlinearities, the gain against speed and the power against filling, and the open-loop response times. Section 3.5 describes the three loops, the supervisory setpoint and the optional predictive controller. Section 3.6 reports the measured scenarios, from the mill run by hand to the graded shift and the predictive controller against the load limit. Section 3.7 describes the real instruments, drives and control platforms, and Sections 3.8 to 3.10 state the limitations, the three results that generalize past this mill and the conclusions.

3.2 Background

Autogenous and semi-autogenous milling displaced multi-stage crushing in large concentrators because a single mill accepts primary crusher product and delivers a feed the ball mills can finish. The change moved the plant's principal energy consumer into a unit whose internal state is almost entirely unmeasured. Standard practice controls the mill through the variables that can be instrumented, which are the ore feed rate, the dilution water rate and, on a variable-speed drive, the shell speed, and it infers everything else. Mill load is inferred from trunnion bearing pressure or from load cells, discharge density from a nuclear or Coriolis gauge, and charge motion from shell-mounted microphones or accelerometers. The standard reference works on comminution circuits describe both the equipment and the reasons the measurements are indirect [7, 8].

> ⧖ HISTORY **Listening to the Charge (1985 to 2014)**
>
> In 1985 a conductivity probe set through the liner bolts of a laboratory mill was first used to follow the toe and shoulder of the charge. Later sensors fall into two groups. On-shell sensors turn with the mill and see every point of the shell. Off-shell devices stand beside it: load cells, the bearing pressure, and acoustic sensors such as accelerometers or an array of microphones. In 1999 and 2000 accelerometers on the shell of an industrial semi-autogenous mill gave a vibration signal sensitive to its operating conditions, and the same method was later used to estimate the toe and the shoulder of the charge. By 2014 such sensors gave both angles online on plant and pilot mills [9].

Control practice on grinding circuits has been surveyed repeatedly, and the surveys agree on two points. Most concentrators run regulatory loops with a supervisory layer above them rather than a single multivariable controller, and the economic objective is a trade between throughput, product size and the cost of grinding media and liners rather than a setpoint [6]. Methods for automatic control, observation and optimization in mineral processing plants have been reviewed with the same conclusion about the dominance of inferential measurement [10]. Model structures suitable for control have been identified and validated on run-of-mine grinding circuits, with hold-ups of rock, solids, water and steel as the states [11], and observers have been built on those structures to reconstruct the hold-ups that no instrument reports [12]. The application follows the same state selection at a reduced size classification.

The ballistic description of charge motion is older than any of that work. A single ball riding the shell leaves it where the shell can no longer supply the centripetal acceleration, follows a parabola and lands on the charge, which gives the impact angle in closed form. Discrete element methods resolve what that description omits, including a distribution of radii, collisions in flight and the lifter geometry that changes as the liner wears [13]. Mill power is treated separately and empirically, because it depends on the mass and the shape of the charge rather than on individual trajectories. The correlation used here is Austin's, which is the form most often quoted for semi-autogenous mills [14]. Published measurements show that slurry filling affects the net power draw independently of total filling, which the single correlation cannot represent [15].

> ⧖ HISTORY **From Energy Laws to Simulated Charges (1867 to 1992)**
>
> Attempts to relate the energy spent in grinding to the size reduction it buys began with laws of 1867 and 1885 and ran through Bond's work of 1952. Each describes grinding by one number, the specific energy, which is the power divided by the throughput. Bond's work index is still measured in a standard laboratory ball mill 30 cm across and 30 cm long. In 1992 the discrete element method was first used to simulate the motion of the media in a ball mill, in two dimensions. Software now follows the media in three, as a quantitative mill simulation needs, but the method has not yet matured enough to predict a mill's product size and throughput [16].

3.3 Process Model

The simulated mill, its three loops and the two inferred measurements they run on are shown in Figure 3.1. The clock is in minutes throughout, while solids and water flows are quoted per hour in the units a concentrator uses, so every rate divides by 60 on its way into a derivative.

3.3.1 The Design Point

The mill is 10.4 m in diameter by 5.1 m long, 433.2 m^3 internal, with a critical speed $N_c = (60/2\pi)\sqrt{g/R} = 13.116$ rpm. At the design point it runs at 78.8 % of critical, which is 10.34 rpm, and holds 241.7 t of balls, 102.7 t of rock and 18.2 t of water for a total volumetric filling of 0.2482. It draws 13.00 MW against a 13.0 MW rating, takes 1600 t/h of ore at 100 mm F_{80} and a work index of 15 kWh/t, adds 640 m^3/h of dilution water for 69.9 % solids in the discharge, and produces 1901 µm

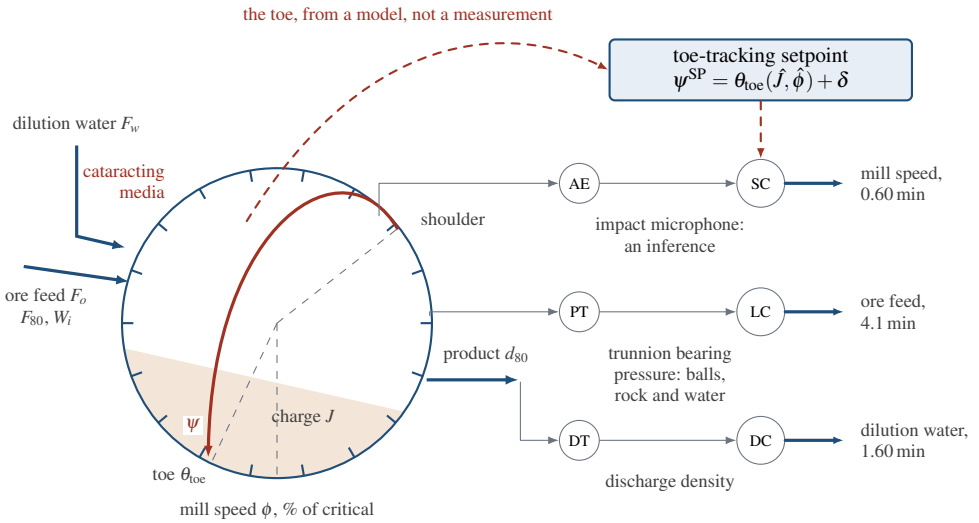


Figure 3.1 The simulated mill and its three loops. The impact-angle controller (SC) moves mill speed against an angle inferred from a shell-mounted microphone. The load controller (LC) moves ore feed against trunnion bearing pressure, which reads the mass of balls, rock and water together. The density controller (DC) moves dilution water. Each output arrow carries the measured open-loop time to sixty-three percent of the response, so the separation between the three loops is visible in the figure. The dashed path (red) is the optional supervisory layer that recomputes the impact-angle setpoint from a toe position estimated through the bearing pressure, which comes from a model rather than from any instrument.

at a specific energy of 8.13 kWh/t. The impact angle is 25.16° and the toe is at 25.14°, so the margin of (3.1) is 0.02° and the media land on the toe.

3.3.2 Charge Geometry and the Position of the Toe

The total filling J counts balls and rock, because the water occupies the interstices, as shown in Equation (3.2).

$$J = \frac{J_b}{100} + \frac{M_c + M_f}{\rho_r V}, \qquad V = \frac{\pi}{4} D^2 L \tag{3.2}$$

where ρ_r is the rock bulk density and V is the mill volume. The charge surface at rest is a horizontal chord cutting off a segment of area $J\pi R^2$, so the half-angle α at which that surface meets the shell solves Equation (3.3), which the model solves by Newton iteration.

$$\alpha - \sin\alpha\cos\alpha = \pi J \tag{3.3}$$

Rotation tilts the chord by a dynamic angle of repose that steepens with speed, which lifts the shoulder and drops the toe, as shown in Equation (3.4).

$$\theta_{toe} = \alpha(J) - [\beta_0 + \beta_1(\phi - 0.75)] \tag{3.4}$$

Equation (3.4) is the term in (3.1) that no instrument on the mill reports, and its dependence on J is the subject of Section 3.6.5. At the design filling, $\alpha = 65.98°$ and the repose angle is 40.84°, which places the toe at 25.14°.

3.3.3 Ballistics of the Cataracting Media

A ball riding the shell leaves it where the centripetal requirement exceeds the radial component of gravity, that is where $\cos\theta_d = \phi^2$ with θ_d measured from the top of the mill and ϕ the fraction of

Table 3.1 Controlled, manipulated and disturbance variables, with the states that carry them. Neither controlled variable in the first two loops is measured directly.

Symbol	Meaning	Unit	Range
Controlled variables			
ψ	impact angle, inferred	°	−20 to 90
p_{brg}	trunnion bearing pressure	kPa	4200–6600
ρ_d	discharge density	% solids	58–78
Manipulated variables			
ϕ	mill speed	% critical	50–100
F_o	ore feed rate	t/h	0–3200
F_w	dilution water rate	m^3/h	100–1500
Disturbances			
W_i	ore work index	kWh/t	10–22
F_{80}	feed size, 80 % passing	mm	40–180
J_b	ball filling	%	6–18
a_f	feed availability	%	0–100
g	grate open area	—	0.5–1.3
w	liner wear	%	0–100
States			
M_c	coarse solids hold-up	t	—
M_f	fine solids hold-up	t	—
M_w	water hold-up	t	—
ψ	impact angle	°	—
ϕ	mill speed, after the drive lag	—	0.20–1.10
F_o	ore feed on the belt	t/h	—
d_{80}	product size	µm	—
w	liner wear	%	0–100

critical speed. It then flies a parabola, as shown in Equations (3.5) and (3.6).

$$x = R\sin\theta_d - R\omega\cos\theta_d\, t, \tag{3.5}$$

$$y = R\cos\theta_d + R\omega\sin\theta_d\, t - \tfrac{1}{2}gt^2 \tag{3.6}$$

Imposing $x^2 + y^2 = R^2$ on the flight cancels the leading terms identically, which is what the detachment condition buys, and leaves the flight time in closed form as $t_f = 4\sin\theta_d\cos\theta_d/\omega$. The landing point follows without integrating anything, as shown in Equation (3.7).

$$\frac{x_f}{R} = \sin\theta_d\,(1 - 4\cos^2\theta_d), \qquad \frac{y_f}{R} = \cos\theta_d\,(1 - 4\sin^2\theta_d) \tag{3.7}$$

The pair lies on the circle of radius R to machine precision, which the application's verification suite confirms at every speed the drive can reach. The impact angle is then $\psi = \mathrm{atan2}(-x_f, -y_f)$, positive toward the toe and negative when the media roll off short of the bottom of the mill.

Worn lifter bars do not grip, so the charge turns more slowly than the drive does. The model applies the slip as a reduction in the effective speed that enters (3.4) and (3.7), as shown in Equation (3.8).

$$\phi_{\mathrm{eff}} = \phi\left(1 - \frac{s_w}{100}\cdot\frac{w}{100}\right) \tag{3.8}$$

where $s_w = 10\,\%$ is the slip at a fully worn liner and w is the wear in percent. At 65 % wear the charge slips 6.5 %, so a drive commanded to 78.8 % of critical throws media as though it were at 73.67 %.

3.3.4 Mill Power

Power comes from Austin's empirical relation for semi-autogenous mills [14], as shown in Equation (3.9).

$$P = K_P D^{3.5} L\rho_c\,\phi_{\mathrm{eff}} J(1 - 0.937J)\left[1 - \frac{0.1}{2^{9-10\phi_{\mathrm{eff}}}}\right] \tag{3.9}$$

where $\rho_c = (M_b + M_c + M_f + M_w)/(JV)$ is the charge bulk density, M_b is the ball mass and K_P is a calibration constant that places the design duty on the nameplate rating. The bracket rolls the power off as the mill approaches critical speed. The product $\rho_c J$ is the charge mass divided by the mill volume, so the filling dependence is really $(M_b + M_c + M_f + M_w)(1 - 0.937J)$, which rises, reaches a maximum and then falls. That interior maximum is the source of the overload runaway analyzed in Section 3.4.

3.3.5 Breakage, Product Size and Liner Wear

Impact energy landing on the toe breaks rock, and the application makes the efficiency a Gaussian in the margin with an attrition floor, because even a cascading mill grinds something, as shown in Equation (3.10).

$$\eta_I = \exp\left[-\tfrac{1}{2}\left(\frac{m}{\sigma_\psi}\right)^2\right], \qquad \eta_B = \kappa + (1-\kappa)\,\eta_I \tag{3.10}$$

where σ_ψ is the width of the impact zone and κ is the attrition floor. Impacts landing *beyond* the toe reach bare shell, so the liner wear rate depends on the positive part of the margin alone and on how hard the mill is being driven, as shown in Equation (3.11).

$$\frac{\mathrm{d}w}{\mathrm{d}t} = \frac{1}{60}\left[k_{w0} + k_{w1}\frac{P}{P_{\text{rated}}}\left(\frac{\max(0,\,m)}{10}\right)^2\right] \tag{3.11}$$

The asymmetry between (3.10), which is symmetric in the margin, and (3.11), which is one sided, is the reason a supervisory setpoint is worth having and also the reason it is dangerous, both of which Section 3.6.6 measures.

Product size comes from Bond's third theory applied to the fraction of the specific energy that actually did size reduction, as shown in Equation (3.12).

$$\frac{1}{\sqrt{d_{80}}} = \frac{\eta_0\,\eta_B\,P/F_o}{10\,W_i} + \frac{1}{\sqrt{F_{80}}} \tag{3.12}$$

where η_0 is the fraction of applied energy credited to Bond-efficient breakage and sizes are in microns. A mill that is well aimed and a mill that is badly aimed therefore differ in product size at the same power draw, which is exactly what a power-based controller cannot see.

3.3.6 Inventory Balances

The three mass balances close on the coarse solids, the fine solids and the water, as shown in Equations (3.13) to (3.15). Solids leave through the grate as a first-order discharge, and fines and water leave together because the discharge carries the mill's own internal density.

$$\frac{\mathrm{d}M_c}{\mathrm{d}t} = \frac{1}{60}\left[F_o(1-f_f) - r_b\right] \tag{3.13}$$

$$\frac{\mathrm{d}M_f}{\mathrm{d}t} = \frac{1}{60}\left[F_o f_f + r_b - k_s g M_f\right] \tag{3.14}$$

$$\frac{\mathrm{d}M_w}{\mathrm{d}t} = \frac{1}{60}\left[F_w + F_o u_m - k_s g M_w\right] \tag{3.15}$$

where f_f is the fines fraction already in the feed, u_m is the feed moisture fraction, k_s is the grate discharge rate constant in reciprocal hours and g is the grate open area multiplier. The breakage rate of the coarse class is first order in the coarse hold-up and scales with power, with how well that power is aimed, and inversely with the work index, as shown in Equation (3.16).

$$r_b = k_b\,\frac{P}{P_{\text{rated}}}\,\eta_B\,\frac{W_i^{\text{ref}}}{W_i}\,M_c \tag{3.16}$$

Four first-order lags carry the rest of the plant, as shown in Equation (3.17). The impact angle reaches the ballistic value with the time the charge needs to rearrange itself, the drive reaches its commanded speed with the drive lag, the belt reaches its commanded rate with the feeder lag, and the product size reaches the value of (3.12) with the transport lag of the discharge and the cyclone overflow.

$$\begin{aligned} \frac{\mathrm{d}\psi}{\mathrm{d}t} &= \frac{\psi_{ss}(\phi_{\mathrm{eff}}) - \psi}{\tau_\psi}, & \frac{\mathrm{d}\phi}{\mathrm{d}t} &= \frac{\phi_{\mathrm{cmd}} - \phi}{\tau_\phi}, \\ \frac{\mathrm{d}F_o}{\mathrm{d}t} &= \frac{F_o^{\mathrm{app}} - F_o}{\tau_F}, & \frac{\mathrm{d}d_{80}}{\mathrm{d}t} &= \frac{d_{80,ss} - d_{80}}{\tau_d} \end{aligned} \tag{3.17}$$

The applied feed is the command delayed by one belt transit and scaled by availability, $F_o^{\mathrm{app}} = a_f F_o^{\mathrm{cmd}}(t - \theta_F)/100$, so the conveyor dead time appears in the load loop without entering the integrator. The belt starts loaded with the feed the mill started on, so a command issued at $t = 0$ also arrives one transit later. After every step the three hold-ups are floored at zero, the wear is clipped to 0 % to 100 % and the speed state is clipped to 0.20 to 1.10 of critical.

3.3.7 Measurements

Neither of the first two controlled variables is measured. The impact angle comes from a shell-mounted microphone, which the application models as the true angle plus an operator-set bias and band-limited noise, passed through a first-order filter of 0.40 min, as shown in Equation (3.18).

$$\tau_m \frac{\mathrm{d}\hat{\psi}}{\mathrm{d}t} = \psi + b_\psi + \sigma_n n_\psi(t) - \hat{\psi} \tag{3.18}$$

where n_ψ is a zero-mean noise source of unit variance with a 0.25 min correlation time and σ_n is the noise amplitude, zero by default, 2.5° in scenario 10 and 1° in the graded challenge. The filter and the noise are updated at each controller scan. The mill load comes from trunnion bearing pressure, which is affine in the mass of balls, rock and water together, as shown in Equation (3.19).

$$p_{\mathrm{brg}} = p_0 + k_{\mathrm{brg}}\left(M_b + M_c + M_f + M_w\right) \tag{3.19}$$

Adding dilution water therefore raises the load reading with no additional rock in the mill, and removing steel lowers it with no rock removed at all. The discharge density $\rho_d = 100 M_f/(M_f + M_w)$ is measured through a 0.20 min filter with its own noise, and the product size is reported through a pure dead time of 2.0 min representing the sampling cycle of an on-stream size analyzer. State observers exist for exactly the hold-up problem this creates [12]. The three loops deliberately do not use one, the optional predictive controller of Section 3.5.4 carries its own, and the toe-tracking setpoint of Section 3.5.3 reads that filter's estimate whichever structure is selected.

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

3.3.8 Parameters

Table 3.2 lists every constant needed to rebuild the simulator. The application exposes most of them for adjustment, and marks the two that are calibration artifacts rather than physical properties as hidden.

3.3.9 Initialization and Numerics

Initialization solves for a genuine equilibrium rather than a plausible one. The fine and water hold-ups follow from (3.14) and (3.15) in one line each once the discharge constant is known. The coarse hold-up does not, because the breakage rate of (3.16) depends on the power, the power of (3.9) depends on the filling, and the filling of (3.2) depends on the coarse hold-up. The application iterates the substitution to machine precision, so every scenario opens with the plant derivatives at zero for its initial inputs, except scenarios 4 and 9, which deliberately start away from equilibrium with the

Table 3.2 Model constants with their default values. Hidden constants are not offered for adjustment because they calibrate the correlations rather than describe the mill.

Symbol	Meaning	Value	Unit
D	mill diameter	10.4	m
L	mill length inside liners	5.1	m
P_{rated}	motor rating	13000	kW
$P_{\max}$	motor overload limit	14800	kW
K_P	power model constant (hidden)	1.45518	—
ρ_b	ball charge bulk density	4.65	t/m^3
ρ_r	rock charge bulk density	1.85	t/m^3
β_0	repose angle at 75 % critical	40	°
β_1	repose steepening with speed	22	° per unit ϕ
σ_ψ	impact-zone width	12	°
κ	attrition floor	0.35	—
k_b	coarse breakage rate constant	23.8	1/h
k_s	grate discharge rate constant	37.9	1/h
η_0	Bond energy utilization	0.365	—
W_i^{ref}	reference work index (hidden)	15	kWh/t
f_f	fines already in the feed	0.10	—
u_m	feed moisture	0.03	—
τ_ψ	charge rearrangement lag	0.30	min
τ_ϕ	drive speed lag	0.25	min
τ_F	feeder response lag	0.35	min
θ_F	feed conveyor dead time	0.80	min
τ_d	product size transport lag	2.5	min
θ_{PSI}	size analyzer dead time	2.0	min
τ_m	impact meter averaging	0.40	min
k_{w0}	baseline liner wear	0.020	%/h
k_{w1}	wear per squared overthrow	0.200	%/h
s_w	charge slip at a fully worn liner	10	%
p_0	bearing pressure at empty mill	1000	kPa
k_{brg}	bearing pressure per tonne	12.37	kPa/t
J_{trip}	overload alarm filling	0.45	—
Defaults for the operator-set inputs			
ψ^{SP}	impact angle setpoint	25.0	°
$p_{\mathrm{brg}}^{\mathrm{SP}}$	mill load setpoint	5485	kPa
ρ_d^{SP}	discharge density setpoint	70.0	%
ϕ	manual mill speed	78.8	% critical
F_o	manual ore feed	1600	t/h
F_w	manual dilution water	640	m^3/h
σ_{brg}	bearing pressure noise	8	kPa
σ_ρ	density gauge noise	0.3	%

impact angle at 5° and with 245 t of rock in the mill. Where the stations chosen before a run have no equilibrium below the overload band, such as a feed the mill cannot grind at the chosen speed, the run starts from the scenario's own stations and the chosen speed and feed arrive as steps. A manual-mode calculation bisects (3.7) against (3.4) to return the speed that lands the media exactly on the toe at the present filling, which is 78.80 % of critical at the design charge.

Integration is fixed-step fourth-order Runge-Kutta at 0.002 min (0.12 s), on the engine of Section 1.6.1. The three controllers execute on a 0.02 min scan (1.2 s), and the historian records twenty samples per simulated minute.

Table 3.3 Impact angle and its local gain against mill speed, computed from (3.7). A negative ψ means the media land short of the bottom of the mill and are cascading rather than cataracting.

ϕ (% crit.)	Shoulder θ_d (°)	ψ (°)	$\partial\psi/\partial\phi$ (°/%)
60	68.90	−26.70	2.211
66	64.18	−12.53	2.521
70	60.66	−1.98	2.761
78.8	51.61	25.16	3.456
86	42.30	53.09	4.393
90	35.90	72.29	5.276

3.3.10 Assumptions

The model carries two solid size classes and one water inventory rather than a population balance, so d_{80} is computed from an energy relation and not from a breakage and selection description. The ballistics are those of a single ball at the shell radius on a frictionless parabola, with no distribution of radii, no collisions in flight and no explicit lifter profile. The charge surface is a plane chord tilted by an angle of repose that depends on speed alone, so there is no slurry pool and no separate slurry filling. Power comes from one empirical correlation whose constant is fitted to place the design duty on the nameplate. Breakage is first order in the coarse hold-up. Ball wear and ball consumption are absent, so the ball charge changes only when a disturbance changes it. The liner wear rate is an index rather than a rate in millimeters. The actuators are ideal within their limits, so there is no stiction, no hysteresis and no stroking time on the feeders or the drive. Section 3.8 states where each of these assumptions fails.

3.4 Nonlinear Characteristics and Dynamic Response

3.4.1 Process Gain Against Mill Speed

The detachment condition places θ_d through an inverse cosine, and (3.7) then maps θ_d to the impact angle, so the process gain $\partial\psi/\partial\phi$ is a function of the operating point. Table 3.3 evaluates both at six speeds.

The gain in Table 3.3 rises by a factor of 2.09 from 66 % of critical to 90, and it never changes sign. A controller tuned at the design point is therefore more aggressive above that speed and more sluggish below it, which is the classic case for gain scheduling and the reason scenario 3 asks for the gain twice. Volume 7, Chapter 3 asks the same question of a plant whose stiffness changes with the operating point, where a linearization at one point promises little about another. A secant measured with the 3 % open-loop step of scenario 3 reads 3.61 °/% at the design point against the local derivative of 3.456, because the step traverses a curve rather than a line.

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

3.4.2 Power Against Filling

The second nonlinearity is not mild. Table 3.4 evaluates (3.9) at the design speed while holding the water inventory at its design value of 18.2 t, so the only variable is how much rock is in the mill.

The maximum in Table 3.4 is at $J = 0.432$ and 14.18 MW. Its position follows from (3.9) in one line, because the charge mass is affine in the filling, $M = M_0 + \rho_r V(J - J_b/100)$ with M_0 the ball and water mass, so maximizing $M(1 - 0.937J)$ gives Equation (3.20).

$$J^{*} = \frac{1}{2}\left(\frac{1}{0.937} - \frac{M_0}{\rho_r V} + \frac{J_b}{100}\right) \tag{3.20}$$

Substituting $M_0 = 259.9\,\mathrm{t}$, $\rho_r V = 801.5\,\mathrm{t}$ and $J_b = 12\,\%$ returns $J^{*} = 0.4315$. The last column of Table 3.4 shows that the geometric term alone peaks much later, at $J = 1/(2 \times 0.937) = 0.534$, so it is

Table 3.4 Mill power against total filling at the design speed, with the water hold-up fixed at 18.2 t. The motor overload limit is 14.8 MW and the filling alarm is at 0.45.

J	Rock hold-up (t)	P (MW)	$J(1-0.937J)$
0.20	64	12.300	0.1625
0.28	128	13.375	0.2065
0.35	184	13.947	0.2352
0.42	240	14.175	0.2547
0.432	250	**14.180**	0.2571
0.48	289	14.097	0.2641
0.55	345	13.687	0.2666

Table 3.5 Open-loop steady-state gain and response times at the design point, for steps of 3 % of critical speed, 50 m³/h of water and 100 t/h of ore. The ore feed rows include the 0.80 min conveyor dead time.

Pairing	Gain	t_{63} (min)	t_{95} (min)
Mill speed → impact angle	3.61 °/%	0.60	1.31
Dilution water → density	–0.0299 % per m³/h	1.60	5.24
Ore feed → bearing pressure	0.76 kPa per t/h	4.13	9.04
Ore feed → product size	1.87 µm per t/h	3.41	7.82

the mass in the charge and not the shape factor that places the maximum. Past J^{*} adding load reduces power, which reduces breakage, which adds load. That positive feedback is the semi-autogenous overload runaway, it is a consequence of (3.9) rather than something added on top of it, and the operator sees it as a load loop that has stopped responding while the power trend falls. Here the peak is a limit to stay below, while Volume 7, Chapter 4 regulates at the peak of a friction curve, where the process gain passes through zero.

> **TRY IT IN THE APP** **9 · Overload: The Runaway on the Far Side of the Power Peak**
>
> Load the scenario and press Run in Manual at 20×. The mill starts with 245 t of rock and the feed left at 1600 t/h. Watch the filling climb until the overload alarm fires at 22.7 min. Then reload the scenario and press Aim at the toe before Run: it sets 83.2 % of critical, and the mill never reaches the alarm.

3.4.3 Open-Loop Response Times

Table 3.5 reports the open-loop response of each controlled variable to a step in its own manipulated variable, measured at the design point of scenario 3 with all three loops in manual. Each response is the stepped run minus an otherwise identical unstepped run, timed from the scan at which the step took effect.

The impact loop is the fast one, and it is fast because its dynamics are the drive lag of 0.25 min and the charge rearrangement lag of 0.30 min in series, with no dead time. The impact meter adds a further 0.40 min of averaging, which raises the time to sixty-three percent of the measured response to 1.04 min. The impact-angle gain is read 10 min after the step, because a mill thrown past the toe wears its liner faster, the slip of (3.8) grows, and the angle drifts down by 0.06° over the following hour. The load loop is the slow one at 4.13 min, a factor of 6.9 on the impact loop, and it also carries the only dead time in the three loops. The separation is what allows the three loops to be tuned one at a time in the order the scenarios use.

3.5 Controller Design

3.5.1 Velocity Form PID Algorithm

All three loops run the velocity-form PID of Section 1.6.2 [17], with the derivative on the filtered measurement, so that moving a setpoint produces no derivative kick. In this chapter's symbols the error and the three increments are given by Equations (3.21) and (3.22).

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k) \tag{3.21}$$

$$\begin{aligned} \Delta u_k^P &= K_c\,(e_k - e_{k-1}) \\ \Delta u_k^I &= \frac{K_c}{\tau_I}\,e_k\,\Delta t_c \\ \Delta u_k^D &= -\,a\,K_c\,\tau_D\,\frac{\mathrm{PV}_k^f - 2\mathrm{PV}_{k-1}^f + \mathrm{PV}_{k-2}^f}{\Delta t_c} \end{aligned} \tag{3.22}$$

where $\Delta t_c = 0.02\,\mathrm{min}$ (1.2 s) is the scan period common to the three loops and a is the action. It is $+1$, reverse acting in the convention of Section 1.6.2, on the impact and load loops. It is -1, direct acting, on the density loop, because more dilution water lowers the discharge density, and the app's hint on the density loop gain uses the same name. The term PV^f is the measurement passed through a first-order filter, $\mathrm{PV}_k^f = \gamma \mathrm{PV}_{k-1}^f + (1-\gamma)\mathrm{PV}_k$ with $\gamma = \exp(-\Delta t_c/\tau_f)$ and $\tau_f = 0.10\,\mathrm{min}$, so that derivative action does not amplify the noise on an inferred angle. The increments are combined, the demand is clipped to the actuator span, and the integral is corrected by back-calculation, as shown in Equation (3.23).

$$\begin{aligned} z_k &= z_{k-1} + \Delta u_k^P + \Delta u_k^D \\ c_k &= z_k + I_{k-1} \\ u_k &= \min(u_{\max},\ \max(u_{\min},\ c_k)) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\,(u_k - c_k) \end{aligned} \tag{3.23}$$

Here z_k is the running sum S_k of Equation (1.11), c_k is the output before limiting, v_k of Equation (1.12), and I_k carries the integral contribution together with the starting actuator position, which seeds it when the controller is initialized. The tracking time is $T_t = \tau_I$, and Section 1.6.2 shows why the

integral alone is corrected and the running sum is never truncated. The ore feed command sits at its 3200 t/h limit through the crusher trip of scenario 7, where back-calculation with the 3.0 min integral time of the load loop holds the integral near the limit, so the command leaves it within a minute of the supply returning at 48 min and is down to 1787 t/h four minutes later. Transfer between manual and automatic is bumpless in either direction, as Section 1.6.2 describes. A loop switched to automatic therefore inherits the integral it was tracking, whereas a loop that starts in automatic is initialized with the integral at the current actuator position and applies the whole proportional term to the error it finds, a difference that Section 3.6.6 measures. The actuator spans are 50 % to 100 % of critical on the drive, 0 t/h to 3200 t/h on the ore feed and 100 m^3/h to 1500 m^3/h on the dilution water, and no output carries a rate limit.

3.5.2 The Three Loops and Their Tunings

Impact angle goes to mill speed, mill load to ore feed and discharge density to dilution water. The pairing follows the response times of Table 3.5, and the loops are commissioned in that order in the scenarios. Continuous cycling of the impact loop under proportional control alone puts the ultimate gain at $K_u = 2.4\,\%$ of critical per degree with an ultimate period of $P_u = 1.12$ min, and the integral time at its 30 min maximum, the nearest setting to proportional control the application offers, gives the same two values, so the Ziegler-Nichols proportional-integral settings are $K_c = 1.08$ and $\tau_I = 0.93$ min [18]. The tuning the application recommends backs off from there to $K_c = 0.75$ and $\tau_I = 1.8$ min, because the measurement is inferred. The cautious tuning the scenarios load with is $K_c = 0.15$ and $\tau_I = 8$ min. Load control uses $K_c = 1.2$ t/h per kPa with $\tau_I = 3.0$ min against a plant gain of 0.76 kPa per t/h, and density control uses $K_c = 40\,m^3/h$ per percent with $\tau_I = 2.0$ min, both of which place the loop gain near unity and neither of which is retuned in any scenario.

3.5.3 The Supervisory Setpoint

The optional supervisory layer replaces the operator's fixed impact-angle setpoint with an estimate of the toe position plus an operator-settable offset, as shown in Equation (3.24).

$$\psi^{\mathrm{SP}} = \theta_{\mathrm{toe}}\left(\hat{J}, \hat{\phi}\right) + \delta \tag{3.24}$$

No instrument reports the toe, the filling, the ball charge or the liner wear, so Equation (3.24) is evaluated every scan on the state of the Kalman filter of Section 3.5.4, which runs under every structure. The filling $\hat{J}$ is the declared 12 % ball filling plus the rock $\hat{M}_c + \hat{M}_f$ that the filter infers from the bearing pressure, the impact meter and the density, and $\hat{\phi}$ is the filter's speed state, the commanded speed through the drive lag with no liner slip. The offset is $\delta = 0$ by default. This is feedforward on a constraint rather than on a disturbance. Section 1.4.6 treats inferential control and the back-off from an inferred limit, and Volume 5, Chapter 10 holds a dryer against a sticky margin that is inferred in the same way. It follows the rock that the bearing pressure sees, nothing corrects its ball charge or its liner, and Section 3.6 measures both what it buys and what it costs.

3.5.4 Model Predictive Control

An optional fourth controller replaces the three loops with one model predictive controller (MPC) that moves the mill speed, the ore feed and the dilution water together. It is the last control structure offered, never the default of a scenario, and it acts only in automatic mode. A decision is made every $T_s = 0.5$ min (25 scans) over $N = 60$ steps (30 min) in 12 move blocks of $[1,1,1,1,2,2,3,4,6,9,15,15]$ steps, 36 input decisions per quadratic program. The plant receives only the first move, which it keeps until the next plan arrives, and planning starts at $t = 0$. Section 1.6.3 describes the receding-horizon arrangement, which dominates industrial predictive control [19], and predictive control has been applied to ball mill grinding circuits [20] and studied in robust nonlinear form on a run-of-mine ore milling circuit [21].

The controller's model is the process model of Section 3.3 with the declared values of Table 3.2 and the design ore and charge ($W_i = 15\,\mathrm{kWh/t}$, $F_{80} = 100$ mm, $J_b = 12\,\%$, $g = 1$, $w = 0$), never the

values of the run. Its seven states are M_c, M_f, M_w, ϕ, ψ, the meter output $\hat{\psi}$ and the feed entering the mill, and the belt is its own record of the feed it commanded. Three estimated disturbances enter as shown in Equation (3.25).

$$\begin{aligned} \frac{\mathrm{d}M_c}{\mathrm{d}t} &= \tfrac{1}{60}\left[F_o(1-f_f) - r_b + d_B\right] \\ \frac{\mathrm{d}M_f}{\mathrm{d}t} &= \tfrac{1}{60}\left[F_o f_f + r_b - d_B - k_s M_f\right] \\ \tau_m \frac{\mathrm{d}\hat{\psi}}{\mathrm{d}t} &= \psi + d_\psi - \hat{\psi}, \qquad y_\rho = \rho_d + d_\rho \end{aligned} \tag{3.25}$$

Here d_B is a shortfall in coarse breakage in t/h, which a harder ore or a starved belt produces, d_ψ is an offset on the impact angle, which a worn liner and a biased meter produce alike, d_ρ is an offset on the density reading y_ρ, and (3.10) is evaluated at $\psi + d_\psi$. The outputs are the meter reading, the bearing pressure of (3.19), y_ρ, the power of (3.9), which no instrument measures, and the margin over the controller's own toe, which (3.4) gives from the estimated rock and the design ball charge and which (3.24) also uses. At the design point the linearized static gains equal the perturbed plant steady state within 0.004 % on all fifteen pairs, for example 3.456 °/% from speed to impact angle and 0.747 kPa per t/h from feed to bearing pressure. Second-order fits to open-loop plant steps leave residuals of 0.00 to 1.67 % of the response and agree with the model's gains within 0.03 to 2.2 %, except from speed to bearing pressure (−8.92 against −6.32 kPa/%), where the efficiency peaks with the media on the toe and a step either way lowers it. Because the speed-to-angle gain reaches 4.78 °/% at 88 % of critical, the model is relinearized at every decision by central differences and discretized with an exact hold. The belt becomes two delay states, so a new feed decision reaches the model's feeder 1.0 min later against 0.80 min on the plant.

The quadratic program minimizes Equation (3.26).

$$\begin{aligned} J = \sum_{k=1}^{N} \Big[& Q_\psi \left(\hat{\psi}_k - \psi^{\mathrm{SP}}\right)^2 + Q_\rho \left(y_{\rho,k} - \rho_d^{\mathrm{SP}}\right)^2 - c_F F_k^{\mathrm{cmd}} \Big] \\ & + \sum_{j=1}^{12} \Delta v_j^{\top} S \, \Delta v_j + \sum_{i=1}^{4} \left(\mu_i s_i^2 + \nu_i s_i\right) \end{aligned} \tag{3.26}$$

Commands are scaled to percent of critical and hundreds of t/h and m³/h, Δv_j is the change in the three commands between successive blocks (the first from the commands applied now), $Q_\psi = 1$ per °², $Q_\rho = 0.5$ per squared percent solids, $c_F = 0.2$ per hundred t/h per step and $S = \mathrm{diag}(8, 1, 0.05)$. With toe tracking the first term holds the margin over the controller's own toe at the offset δ. Each slack $s_i \geq 0$ serves one soft limit over the whole horizon at quadratic and linear prices (μ_i, ν_i): the bearing pressure at or under the mill load setpoint, read as a ceiling, at $(10^4, 10^3)$ per hundred kPa, the predicted power at or under the 14.8 MW overload limit at $(10^5, 10^4)$ per MW, and the meter (or margin) within ±3° of its target and the density within ±2 % solids, both at $(100, 10)$. Where the response with no new moves already lies past a limit, the limit follows it for two steps, while new feed is still on the belt, and then closes on it as $1 - \mathrm{e}^{-(t_k - 1\,\mathrm{min})/5\,\mathrm{min}}$. Every slack was zero in each check made, including a load ceiling set 86 kPa under the bearing pressure, so the soft limits act mainly through these funnels. The hard limits are the loops' own spans, and the overload alarms are unchanged.

Offset is removed by estimating the three disturbances [22] with a steady-state Kalman filter whose gain solves the discrete Riccati equation for the model augmented with three random walks, linearized at the design point and discretized at the 0.02 min scan. Every scan the nonlinear model advances on the commands sent and is corrected by the gain times the residuals of the three instruments, and the slowest error mode is 8.84 min. The filter operates whichever structure is selected, and neither structure receives a preview. A solution is discarded unless it is optimal, finite and within its bounds,

and the three PIDs, kept aligned with each applied command, resume control bumplessly until a valid plan returns, although no plan was rejected in any case measured. The plant runs on the actual ore, charge and liner, the instruments impose their filters, bias and noise, and the controller sees nothing but the three instruments, its own commands, the setpoints, its tuning and the motor limit. A decision takes 1.2 to 1.7 ms on average in Node.js.

3.6 Results

Every number in this section was produced by driving the application's own model at its shipped constants and seed, and each metric is defined where it is first used. Settings are changed at $t = 0$ before the run starts unless the text says otherwise, and the application gives the same run when a setting is carried through Reset instead. Integral error is the integral of the absolute deviation of the true impact angle from its setpoint, and speed travel is the accumulated absolute change in the drive speed expressed in percent of critical, the quantity the application reports as speed travel.

3.6.1 The Optimum and the Overload by Hand

Scenario 2 is a search for an optimum rather than a setpoint. On a new liner the settled product size runs out to 5875 µm at 75 % of critical and 4676 µm at 85, against 1901 µm at the optimum. The fast side is bounded twice, by a liner wear rate that reaches 0.84 %/h at 85 % and by the 14.8 MW motor limit, which the settled mill crosses near 87.5 %. The slow side has no limit of its own and fills the mill.

Scenario 9 opens in manual on the far side of the power peak of Section 3.4, with the feed left at 1600 t/h, and the overload alarm fires at 22.7 min. The manual calculation of Section 3.3.9 returns 83.2 % of critical for the charge the scenario starts with, and with that speed set before the run the mill never reaches the alarm.

3.6.2 A Setpoint Step Under Two Tunings

The two tunings of scenarios 4 and 5 are compared on the same 5° impact-angle setpoint step, made in scenario 5 at 10 min, once the loop has settled on its 25.0° setpoint, with the cautious gains entered at the instant of the step and the response followed for 120 min. The cautious controller, $K_c = 0.15$ and $\tau_I = 8$ min, reaches the new setpoint within half a degree after 42.2 min without overshoot, for an integral error of 73.7 ° min and 1.60 % of speed travel. The recommended controller, $K_c = 0.75$ and $\tau_I = 1.8$ min, is inside the same half degree after 0.32 min, for an integral error of 4.82 ° min, a factor of 15.3, at the cost of 7.70 % of speed travel against 1.60. The recommended tuning overshoots to 33.0° on a step whose target is 30.0, which is 60 % of the step, and it stays inside the half-degree band only after 3.43 min. Scenario 4 poses the same contrast as a recovery from an impact angle of 5°, where the cautious loop needs 47.2 min to come within 2° of setpoint and 63.2 min to come within 1°, while the recommended tuning holds within 1° from 3.8 min.

That is a textbook result on a textbook plant, and it is the least informative measurement in this chapter.

3.6.3 A Measurement That Is an Inference

Scenario 10 turns on 2.5° of impact-microphone noise with derivative action set to $\tau_D = 0.5$ min, and the comparisons here use the thirty-minute runs the scenario prescribes. With a clean signal the controller makes essentially no speed moves at steady state, at 0.14 % of travel. With the noise it travels 172 % of critical speed over the same thirty minutes, the peak liner wear rate rises from 0.020 to 0.112 %/h, and the wear consumed over the run rises from 10.0 to 12.3 m%. Removing derivative action alone cuts the travel to 72.5 %, while the peak wear rate of 0.093 %/h, the wear consumed of 12.8 m% and the integral error of the true angle, 56.9 ° min against 54.8, barely move. Behind the drive and charge lags the true angle follows only the slower part of the noise, which the proportional and integral terms chase with or without derivative action. Over seven noise seeds the travel spans 146 % to 172 % with derivative action and 56 % to 73 % without, while the peak wear rate spans

Table 3.6 Sixty-five percent liner wear, three configurations, after 120 min.

	ϕ_{cmd} (%)	ϕ_{eff} (%)	ψ (°)	Margin (°)	d_{80} (µm)
Manual at 78.8 %	78.80	73.67	8.62	−29.38	8416
Auto, fixed setpoint	84.23	78.76	25.00	−0.17	1906
Auto, toe tracking	83.95	78.50	24.11	−1.13	1915

Table 3.7 A two-point drop in ball filling, with the impact-angle loop holding its 25° setpoint throughout.

t (min)	J	θ_{toe} (°)	ψ (°)	Margin (°)	d_{80} (µm)
39	0.2506	25.37	25.59	+0.22	2367
50	0.2788	28.37	25.10	−3.27	4136
70	0.2721	27.69	24.98	−2.71	3485

about 0.05 %/h to 0.13 %/h in both cases, so the travel comparison is robust and the wear comparison lies within the seed spread.

Nothing has happened to the mill. The controller is chasing an instrument, the actuator is a variable-speed drive on a 13 MW motor, and the cost appears on the drive and the liner rather than on any trend the loop can see.

3.6.4 A Worn Liner Moves the Plant and Not the Measurement

Scenario 8 runs the mill four months into a liner campaign at 65 % lifter wear, where (3.8) slips the charge 6.5 % against the shell.

Table 3.6 is the case the loop handles. Left in manual at the speed that was correct on new liners, the mill slides into cascading, the impact angle collapses to 8.62° and the product coarsens from 1901 to 8416 µm, a factor of 4.4. In automatic the loop commands 5.4 points more speed and the product is within 0.3 % of design. The plant moved and the controller absorbed it, which is what controllers are for.

Two observations survive that success. The actuator has spent 5.4 of the 50 points of its range on a disturbance that will keep growing for another two months. Toe tracking coarsens the product slightly here, 1915 against 1906 µm, because its toe is computed for a new liner. The filter's speed state is the drive speed of 83.95 % rather than the charge speed of 78.50 %, which overstates the repose angle by 1.2° and places the tracked toe 1.13° below the true one.

3.6.5 A Constraint That Moves While the Setpoint Does Not

The instructive case is the one in which the toe moves. In the graded scenario the ball filling is reduced from 12 to 10 % at 40 min, which is the ordinary consequence of a missed ball top-up. The scenario runs with its cautious tuning and a fixed setpoint, and its graded run ends at 70 min.

Two percentage points of ball filling are 40.3 t of steel, and (3.19) does not distinguish steel from rock. The bearing pressure falls at once from 5487 to 4989 kPa, the load controller reads that as an empty mill and raises the ore feed from 1431 to a peak of 2243 t/h on the belt at 43 min, and the rock hold-up climbs from 104.7 to 147.4 t at 47 min before settling near 138 t once the bearing is back on setpoint. The filling rises with it, and (3.4) lifts the toe by three degrees, as shown in Table 3.7. The impact angle is held within 0.6° of its setpoint for the remaining half hour of the graded run, and the margin goes from just positive to −3.3° at 50 min, with its worst value of −4.0° at 47 min. The media are now landing inside the slurry, (3.10) falls, and the product coarsens from 2367 to 4136 µm, which is 75 %.

The three controlled variables are all on setpoint while that happens. The impact angle is on setpoint, the mill load is on setpoint, the discharge density is on setpoint, and the mill power has

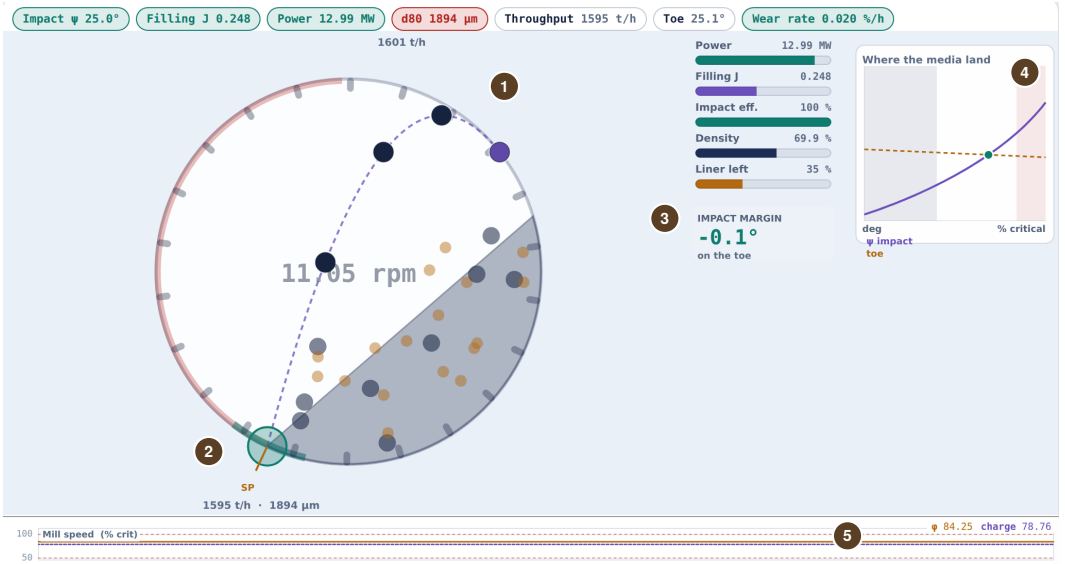


Figure 3.2 Scenario 8 at $t = 800$ min (13.33 h on the application clock) in automatic, with the liner 65 % worn at the start, cropped to the plant view and the mill speed strip. The mill is drawn in end view with the charge, the cataracting trajectory (1) from (3.7) and the impact point (2) marked against the toe. The impact angle is on its 25.0° setpoint and the margin (3) reads −0.1°, on the toe. The loop commands 84.25 % of critical instead of the 78.8 % that was correct on new liners, the charge turns at 78.76 % (the dashed line on the speed strip, 5), and the product is 1894 µm. The inset (4) plots the impact angle and the toe against speed with the operating point on both.

moved by 2.9 %. The only manipulated variable that moved is the ore feed, and it moved because the load loop was defending a mass rather than an inventory of rock. The one instrument that reports the problem is the size analyzer, which is two minutes of dead time and a 2.5 min lag behind the mill (Section 1.4.5), and the only remedy available to the impact loop is to move its setpoint.

Pause. Try Exercise 3.3 before reading on.

3.6.6 Tracking the Toe on Three Disturbances

Equation (3.24) follows the toe whenever the rock in the mill moves, because the bearing pressure sees rock, and the application measures it on three such disturbances. The ball-charge drop of Table 3.7 is not one of them, and Section 3.6.7 returns to it.

On a harder ore the two setpoint modes are indistinguishable, because the filling barely moves. On a starved conveyor, where a crusher trip empties the mill and the toe falls with it, the fixed setpoint leaves the impacts beyond the toe and the peak liner wear rate reaches 0.0566 %/h, while toe tracking follows the toe down and holds 0.0206 %/h, a factor of 2.7.

The overload scenario shows the cost of that benefit. The scenario opens in manual with the mill already carrying 245 t of rock at $J = 0.4257$, which puts the charge surface high and the toe at 43.9° while the drive is still at 72 % of critical and the media are landing at 3.7°. The filter starts from the bearing pressure and finds that filling, so toe tracking asks for 43.9°, 18.9° above the fixed setpoint. Switched to automatic before the run starts, the controller starts with its integral at the 72 % speed and applies the whole proportional term to the 40.3° error. The command saturates at 100 % of critical, the shell reaches 92.7 %, and the impact angle overshoots to 64.8°, which is 25.2° past the toe and onto bare liner. The peak liner wear rate reaches 1.51 %/h, 75 times the baseline of 0.0200, and the power reaches 15.41 MW. The fixed setpoint also crosses the motor limit briefly, at 14.89 MW,

Table 3.8 Fixed setpoint against toe tracking on three disturbances, measured over 150 min with the recommended impact loop tuning in scenarios 6, 7 and 9. The overload case is switched to automatic before the run starts.

Scenario	Setpoint	Peak wear (%/h)	Peak P (MW)	Peak J
Harder ore	fixed	0.0207	13.069	0.2586
	toe	0.0201	13.114	0.2584
Starved feed	fixed	0.0566	13.810	0.3122
	toe	**0.0206**	13.995	0.3096
Overload	fixed	0.0200	14.892	0.4257
	toe	**1.5065**	**15.414**	0.4257

Table 3.9 The graded shift over its 70 min, with the worst deviations of the true impact angle from the setpoint and from the toe. Only the cautious tuning with the load loop slowed to $K_{c,L} = 0.6$ and $\tau_{I,L} = 6$ min, the suggested configuration, meets all five clauses. The others fail the toe clause, and the sharp tuning and the cautious tuning with toe tracking also fail the setpoint clause.

Configuration	SP err (°)	Toe err (°)	IAE (° min)	Travel (%)	Peak wear (%/h)	d_{80} (µm)
Cautious, $K_c = 0.15$	0.60	3.98	15.0	9.8	0.0277	3485
plus toe tracking	4.24	3.86	26.9	10.7	0.0308	3505
slower load loop	0.60	2.42	15.0	9.8	0.0266	3264
Recommended, $K_c = 0.75$	2.51	5.45	49.6	64.9	0.0302	3482
plus toe tracking	2.45	6.36	47.8	59.1	0.0302	3501
Sharp, $K_c = 1.6$	45.5	47.1	1236	2650	4.489	7058

but never leaves the baseline wear rate. Once the load loop has emptied the mill both modes settle identically, and there toe tracking is the better of the two, holding the media on a toe that has come down to 32.6° and delivering 4699 µm against 5280 µm for the fixed setpoint.

The distinction is not subtle once it is stated. Toe tracking is correct while the filling is where the plant wants it and the ball charge and liner are the declared ones. In an overload it aims at a toe that sits 19° above the fixed setpoint and is about to fall by 11°, so it demands a speed move of about 12 points of critical, and with the whole of that move reaching the drive at once it ends on bare liner. The model that computes the setpoint cannot tell the difference, because it is open loop on the filling, which is the variable moving fastest. A supervisory feedforward needs an interlock against the constraint it is tracking, and the application ships without one so that the failure can be staged deliberately [10].

3.6.7 The Graded Shift and a Tuning That Makes It Worse

The challenge runs four unannounced disturbances over seventy minutes with 1° of impact-meter noise, starting from the cautious tuning and a fixed setpoint, and judges the whole run at 70 min. It asks for the impact angle within ±4° of its setpoint and within ±3° of the moving toe, the mill power under 14.8 MW, the filling under 0.40 and a final liner wear rate under 0.06 %/h.

Table 3.9, drawn in Figure 3.3, contains a result that contradicts the obvious advice. The tuning the application recommends, and that Section 3.6.2 shows to be 15 times better on integral error after a setpoint step, is 3.3 times worse here with a fixed setpoint and uses 6.6 times the actuator travel. Every disturbance in this scenario is slower than the loop already was, and the measurement carries a degree of noise, so the extra gain buys nothing and amplifies what the meter is reporting. Neither tuning meets the target on a fixed setpoint with the shipped load loop, because the ball-charge drop of Table 3.7 moves the toe away from the setpoint both of them hold. The cautious loop misses the toe by 3.98°, 0.98° beyond the limit, and the recommended loop by 5.45°. As shipped, the cautious

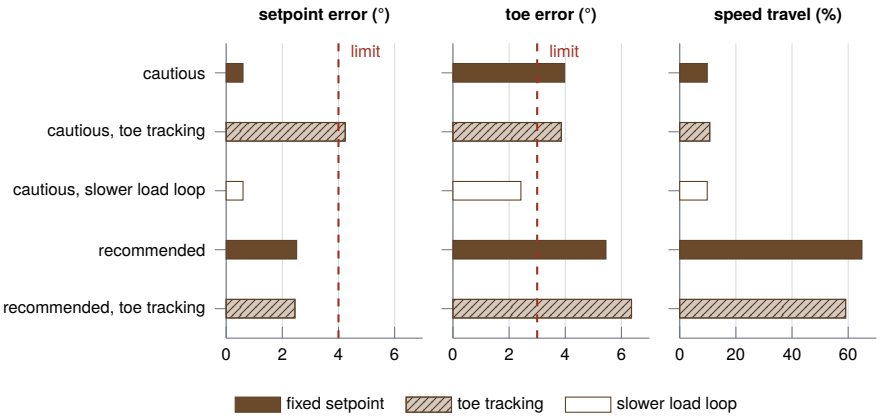


Figure 3.3 The graded shift of Table 3.9 for the five configurations on the loops' scale: the worst error from the setpoint, the worst error from the moving toe and the speed travel over 70 min. The dashed lines are the challenge's clauses of 4° and 3°. The sharp tuning, far off every scale, stays in the table. Notice that only the slower load loop (open bar) clears the toe clause, and that the recommended tuning spends 6.6 times the travel of the cautious one.

run fails on the toe clause alone, with a setpoint error of 0.60°, a peak power of 13.06 MW, a peak filling of 0.284 and a final wear rate of 0.020 %/h.

Pushing the gain further makes the loop unstable. At $K_c = 1.6$ with $\tau_I = 0.6$ min, a gain two thirds of the ultimate gain, the impact angle is outside the ±4° band for 48.5 min of 70, the drive travels 2650 % of critical speed, the peak liner wear rate reaches 4.49 %/h and the product moves from 3485 to 7058 µm.

> ✦ DID YOU KNOW? **A Better Tuning That Loses the Shift**
>
> From the simulator: in Table 3.9 the recommended tuning, $K_c = 0.75$, ends the graded shift with an integral error of 49.6 ° min and 64.9 % of speed travel, against 15.0 and 9.8 % for the cautious $K_c = 0.15$, and it misses the toe by 5.45 against 3.98°.

Toe tracking does not help here, because the toe it follows is inferred from the bearing pressure. When the bearing pressure falls with the steel at 40 min, the filter reads the fall as rock leaving and its toe drops from 25.4 to 19.6° within a minute, while the true toe first falls 2.2° with the lost steel and then climbs to 28.9° at 46.9 min as the load loop fills the mill with rock. With the supervisory setpoint on, the cautious loop misses the toe by 3.86° and its own moving setpoint by 4.24°, and the recommended loop misses the toe by 6.36°. Of 90 impact tunings, K_c from 0.05 to 1.2 and τ_I from 0.6 to 30 min, none passes with either setpoint mode.

What does help is the load loop. The toe climbs out of the band because the load loop overshoots while it replaces the steel, carrying the rock to 147.4 t before it settles near 138 t. With the load loop at $K_{c,L} = 0.6$ t/h per kPa and $\tau_{I,L} = 6$ min, a quarter of its shipped integral action, the rock peaks at 135.1 t at 55.1 min and the toe at 27.32°, and the cautious loop on its fixed setpoint stays within 2.42° of the toe. This is the challenge's suggested configuration. It passes with a margin of 0.58° on the toe clause, a setpoint error of 0.60°, a peak power of 13.09 MW, a peak filling of 0.269 and a final wear rate of 0.020 %/h, and the settings give the same run entered before the run or carried through Reset. It passes on all eight noise seeds tried, with worst toe errors of 2.26° to 2.91°, while toe tracking passes on none, at 3.79° to 4.41° with the cautious tuning and 5.00° to 7.07° with the recommended one. The slower load loop costs 1.9 % of the ore over the run (1898.9 against 1934.9 t).

Table 3.10 The three loops (L) and the MPC (M) over 70 min at each scenario's own load setpoint, with R the loops at the recommended impact tuning. Peak is the largest true bearing pressure over the setpoint, reached in the first minute in scenarios 4 and 9, which start away from it. Peak power above 14.8 MW raised a motor overload. The MPC made 140 decisions in every case with no rejected plan.

Sc.		Ore (t)	Peak (kPa)	Feed travel (t/h)	Density IAE (% min)	Peak P (MW)
6	L	1733.5	84.2	8208	12.7	13.069
	M	1733.2	**24.2**	2247	19.5	13.012
11	L	1845.2	84.2	8277	16.4	13.090
	M	1844.0	**24.2**	2636	23.3	13.059
7	L	1840.8	683.5	9082	41.2	13.810
	M	**1750.3**	560.3	6335	66.2	13.637
4	L	1742.6	1332.3	8540	24.0	13.535
	R	1847.5	1332.3	8961	41.8	14.709
	M	1860.9	1363.1	3521	24.3	**14.801**
9	L	2376.8	895.5	9515	23.7	**14.892**
	M	2341.0	895.4	3852	24.8	14.709
12	L	1898.9	94.3	4277	17.6	13.095
	M	1924.8	83.6	3860	32.3	13.035

3.6.8 Model Predictive Control Against the Three Loops

Scenario 11 hands the tuned mill of scenario 5 to the MPC of Section 3.5.4 and changes the work index from 15 to 19.5 kWh/t at 8 min, to 12 at 30 min and back to 15 at 50 min, and the challenge is now scenario 12. Every comparison runs 70 min on each scenario's own seed, noise and events, with the same instruments, no preview for either structure, and the MPC at the settings of scenario 11 throughout. Scenario 9 ships in manual, where neither structure acts, so its rows are the case switched to automatic before the run, as in Table 3.8, and the scenario 12 rows use the challenge's suggested settings. Measures are taken at every integration step, and ore is the solids the mill discharges. Feed travel is the measure the application reports, summed between logged samples at twenty per minute (summed at every scan, the loops' value in scenario 11 is 13 645 t/h rather than 8277, while the MPC holds each move for half a minute).

The harder ore at 8 min cuts the breakage rate of (3.16) by 23 %, the ratio $15/19.5$. By 10 min the estimator reads a shortfall of 371 t/h, the MPC lowers the feed command from 1555 at 8 min to 1349 t/h at 12 min, and the bearing pressure peaks 24.2 kPa over the setpoint at 10.83 min. The load loop answers the bearing pressure itself, which peaks 84.2 kPa over at 12.82 min, and after the softer ore arrives it lets the bearing fall 143.8 kPa under the setpoint at 34.70 min, against 76.9 kPa at 33.50 min under the MPC, as shown in Figure 3.4. Over the hour the two make the same ore (1844.0 against 1845.2 t) because both hold nearly the same average load (5479.9 against 5484.1 kPa, the MPC treating the setpoint as a ceiling), and scenario 6 gives the same tie and the same peaks, as listed in Table 3.10. In scenarios 5, 8 and 10 and in the challenge as given the ore of the two structures differs by at most 0.53 %.

The tie becomes a difference when the load must respect a limit. Table 3.11 gives each structure the highest setpoint, found in hindsight by bisection to 0.1 kPa, at which the true bearing pressure never exceeds 5485 kPa after the first minute. The MPC then makes 4.1 % more ore than the loops in both scenarios (1812.7 against 1741.2 t in scenario 11). Ore made is almost a function of the setpoint alone, 1.2 to 1.3 t per kPa for every structure, so the gain is the overshoot the MPC avoids. On three other noise seeds, searched on a 5 kPa grid, the gain is 3.9, 4.2 and 8.3 %, the last because a noise excursion holds the loops at 5355 kPa. Volume 4, Chapter 4 makes the same trade of throughput against the overshoot of a limit on a machine whose limits are measured late or not at all.

The gain belongs to the breakage-shortfall estimate rather than to the optimization. Without d_B

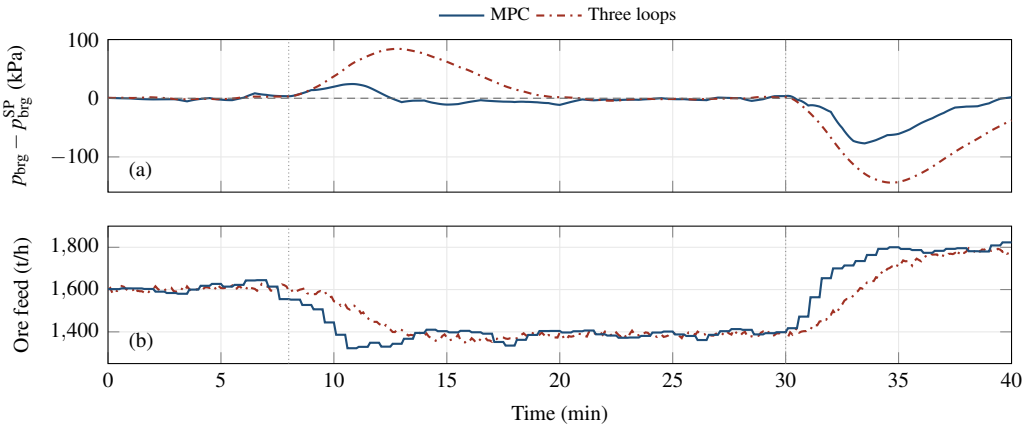


Figure 3.4 Scenario 11 under the MPC and under the three loops, both at the 5485 kPa load setpoint. (a) True bearing pressure over the setpoint. (b) Ore feed command. The MPC is the solid line (blue) and the three loops the dash-dotted line (red). The ore hardens at 8 min and softens at 30 min (vertical dotted lines).

Table 3.11 Ore over 70 min with the bearing pressure held at or under 5485 kPa after the first minute, each structure at its own hindsight setpoint, which is the same in both scenarios. Feed travel is for scenario 11.

Structure	Setpoint (kPa)	Sc. 6 (t)	Sc. 11 (t)	Feed travel (t/h)
Three loops, as shipped	5403.9	1636.1	1741.2	8266
MPC, as shipped	5460.5	**1703.8**	**1812.7**	2687
Loops with $\hat{d}_B$ feedforward	5451.1	1694.2	1805.4	12689
Load loop retuned for ore	5432.3	1669.2	1778.1	16524
MPC, load tracked, no reward	5465.9	1711.6	1823.7	4211

the same MPC crosses the setpoint by 108 kPa in scenario 11, while an MPC that tracks the load setpoint from both sides with no feed reward does slightly better than the shipped one. The decisive comparator gives the shipped loops the MPC's own estimate, which the filter computes under every structure, as a feedforward that subtracts $g\hat{d}_B/(1-f_f)$ from the load loop's output, where $g = 1$ cancels the shortfall in (3.13). The best of $g = 0.25$, 0.5, 0.75, 1 and 1.5 is 0.5, and the loops then come within 0.4 % of the MPC in scenario 11 and 0.6 % in scenario 6. The load loop retuned for ore, the best of six gains and five integral times ($K_c = 2.4$ t/h per kPa, $\tau_I = 2$ min), comes within 1.9 %, and 1804.8 t with the feedforward added. What remains particular to the MPC is the feeder, which travels a third to a sixth as far as the three loop structures listed, and the third seed, on which the feedforward gains 0.18 % over the loops against the MPC's 8.3 %.

In scenario 9 the MPC raises no motor overload where the loops raise one for 0.25 min, and makes 1.5 % less ore. Its slow speed moves (a speed weight of 8 in S) are responsible rather than its power limit, because removing the limit gives the identical run, while with a speed weight of 0.5 it crosses 14.8 MW only without the limit.

The MPC loses four comparisons. On the starved belt of scenario 7 it makes 4.92 % less ore. During the crusher trip the estimator reads the missing ore as faster breakage ($\hat{d}_B = -2178$ t/h at 47 min), so its model expects added feed to fill the mill quickly, and the feed command reaches only 2363 t/h by 40 min where the load loop is at its 3200 t/h stop from 38.85 min. In scenario 4 it raises one motor overload, at 14.801 MW for 0.10 min, where neither the cautious loops nor the recommended impact tuning do, because its power limit acts on a prediction with no meter behind it.

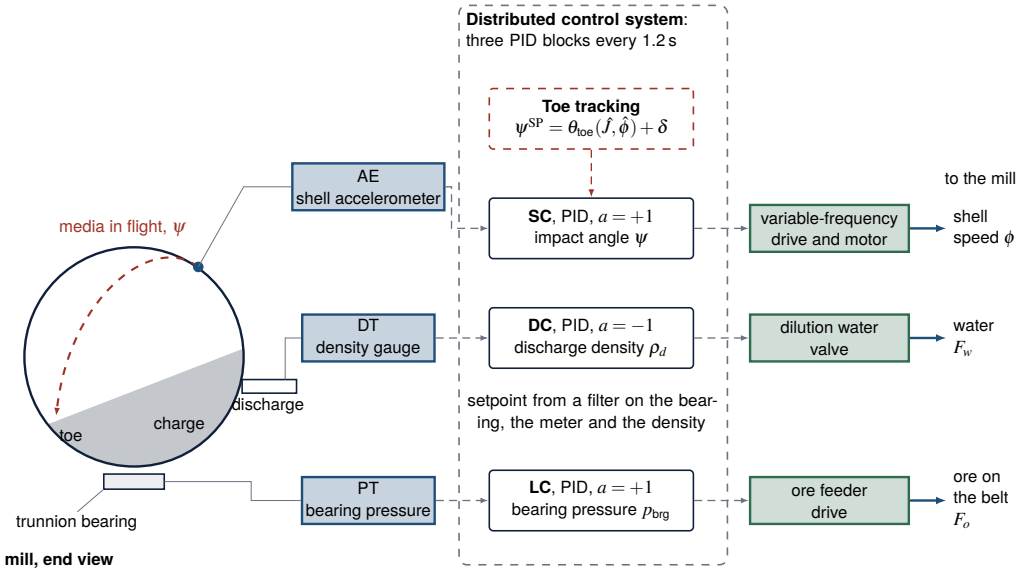


Figure 3.5 The three loops of the simulated mill drawn with their hardware. The shell accelerometer AE, the trunnion bearing pressure PT and the discharge density gauge DT feed three PID blocks, SC and LC reverse acting with $a = +1$ and DC direct acting with $a = -1$, which move the variable-frequency drive of the shell, the dilution water valve and the ore feeder drive. The dashed block writes the impact-angle setpoint from a filter's toe. Notice that no device on the figure reads the toe, the filling or the ball charge.

It holds the density less tightly in every case but scenario 8, because the feed reward and the bearing ceiling trade water for rock, and in scenario 11 the slurry averages 70.19 against 70.03 % solids after 10 min. With the challenge's suggested settings the loops pass at 2.42° from the toe and the MPC, which holds the same fixed setpoint and has no load loop to slow, fails at 3.97°. With toe tracking both structures aim at the filter's toe, which comes from the rock estimated from the bearing pressure with a 12 % ball filling, so when the filling falls to 10 % at 40 min the filter reads the falling bearing as rock leaving and lowers the toe while the true toe rises. The MPC then misses the toe by 5.68° and the loops at the recommended tuning by 6.36°. Aimed at the true toe, which no instrument reports, the loops would miss by 2.31° and the MPC by 3.17°, the difference coming from its slow speed moves.

3.7 Technology in Practice

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

What Measures It

The model of Section 3.3 reduces the impact meter to the true angle plus a bias and band-limited noise behind a 0.40 min filter, the bearing pressure to an affine function of the charge mass, the density gauge to a noisy filtered reading, and the size analyzer to a 2.0 min dead time. The bearing pressure and load cells are off-shell devices [9], and each weighs the charge as a whole.

TECHNOLOGY Sensor: An Acoustic Pickup on the Mill Shell

The tumbling charge shakes the shell with an intense vibration. An accelerometer, a piezoelectric crystal on the shell, picks it up, and the vibration where the charge strikes the shell gives the toe and shoulder angles. Microphones beside the mill listen to the same events without touching it [9]. The meter infers an angle from sound, so a bias moves the true angle while its own trend sits on setpoint, and noise makes the speed loop chase the instrument, as scenario 10 shows.

No instrument on the mill reports the toe, the filling or the ball charge, and the common way to measure the filling is with the mill stationary [9].

IN PRACTICE Crash Stop and Grind-Out

On a plant the filling is still measured with the mill stopped. A crash stop cuts the feed and de-energizes the motor, and the inside diameter, the chord across the charge surface and the height from the roof to that surface are measured. To find the ball load alone, the mill first runs without feed until all the rock is ground out and discharged, and the chord and height are measured again. The method costs production time, which is why online sensors are sought [9]. In this chapter no instrument reports the ball charge at all.

What Moves It

The model's drive is a first-order lag of 0.25 min over a span of 50 to 100 % of critical with no rate limit, and the ore feed reaches the mill through a 0.35 min feeder lag and a 0.80 min belt.

TECHNOLOGY Final Element: A Variable-Frequency Drive on the Mill Motor

The shell speed is set by a variable-frequency drive on the mill motor. The drive rectifies the line power to direct current, smooths it, and an inverter switches it back to alternating current at the frequency and voltage the commanded speed needs. It lets a fixed-speed motor deliver power over a wide range of speeds, ramps the speed up and down to spare the mechanism, and can brake the load [23]. Worn lifters let the charge turn more slowly than the shell, and the drive's own speed signal cannot show it.

The slip of Equation (3.8) stands for that lost grip: the drive holds its commanded speed while the charge turns at ϕ_{eff}.

Where the Controller Runs

Supervisory and lower-level loops run with cycle times from milliseconds to minutes, and a distributed control system uses a supervisory loop that sends set points to a group of localized controllers [24]. Programmable logic controllers among them run PID loops and drive variable-speed motors [23]. The app integrates every 0.002 min, runs the three PIDs on a 0.02 min (1.2 s) scan and the MPC every 0.5 min. Semi-autogenous mills have traditionally been run by expert control systems built on predefined rules. An upper layer watches the bearing pressure and the motor power against their limits and suggests setpoints, and a regulatory layer moves the tonnage, the solids percentage and the mill speed [25]. In such a platform the chapter's structure is three velocity-form PID blocks with back-calculation, the toe-tracking calculation as a supervisory block that writes the impact setpoint, and the MPC above them.

What Is Modernizing It

The first development is sensing the charge. Acoustic and vibration sensors on or beside the shell now give the toe and shoulder angles online, and the filling estimated from them agreed with the filling measured on the stopped mill [9]. A load sensor package mounted on the lifter bolts outside the shell replaced an instrumented ball that the plant mill destroyed, and it worked well in a pilot mill [2]. The second is the digital twin. A model trained on 68 h of plant data from a mill under expert control predicted its behavior 2.5 min ahead at 30 s intervals with errors under 5 % [25]. The third is model

predictive control, which has been studied in robust nonlinear form on a run-of-mine ore milling circuit [21] and which Section 3.6.8 tests against the three loops here.

3.8 Limitations

The charge is two size classes and a water inventory. A real mill needs a population balance over many size intervals to say anything about the product distribution, and the single d_{80} reported here follows from Bond's law applied to an efficiency-weighted specific energy rather than from any breakage and selection description [8]. The application can say that a badly aimed mill grinds coarser and cannot say what the size distribution looks like.

The ballistics are those of a single ball at the shell radius on a frictionless parabola. Real charge motion involves a distribution of radii, a slip velocity that varies through the charge, collisions in flight and lifter geometry that changes as the liner wears, all of which discrete element methods resolve and this model does not [13]. The impact angle here is the trajectory of an idealized outermost ball, and it should be read as an index rather than as a physical measurement.

Ball wear and ball consumption are absent, so the ball charge changes only when a scenario changes it. Slurry pooling, grate blinding and pebble recycle are absent, and pebble handling is a large part of how a real semi-autogenous circuit responds to a hard ore [7]. Power comes from a single empirical correlation with no allowance for slurry filling separately from total filling [15], which is the largest simplification in (3.9), and the shape of Table 3.4 therefore places the power maximum more confidently than any real mill would support.

The liner wear rate of (3.11) is an index driven by the positive part of the margin and by the power. Real wear depends on impact energy, media size, liner alloy and slurry abrasivity, so the numbers in Tables 3.8 and 3.9 should be read as ratios rather than as rates. The actuators are ideal within their limits, so stiction, hysteresis and stroking time are absent, and the noise sources are stationary where a real impact meter drifts.

> **OPEN CHALLENGE** **Seeing Inside a Tumbling Mill**
>
> A monitoring device inside the shell of a running mill is close to impossible, because the tumbling charge is so violent. Plants read the charge indirectly, through the motor power, load cells and, more recently, the mill's sound. An instrumented grinding ball that transmitted the impacts it felt worked in a plant mill but could not withstand its impacts [2]. In this chapter the toe, the filling and the ball charge are reported by no instrument. Progress would be a sensor that survives inside the charge and reports where the media land against the toe.

The predictive controller's model knows the declared parameters and the design ore and nothing else. Ore hardness, feed size, ball filling, liner wear, grate open area and feed availability are model error, absorbed in part by the three estimated disturbances, and two absorptions point the wrong way. A starved belt reads as faster breakage (scenario 7), and a lower ball charge reads as rock leaving, so the filter's toe moves opposite to the true toe. The power it limits is a prediction with no meter to correct it (scenario 4). The belt is rounded to two decisions, the density filter is omitted, the product size is not an output, and the shared solver prices one slack per output across the whole horizon. The toe-tracking setpoint of (3.24) aims at the same filter's toe under either structure, so a ball-charge loss or a worn liner moves the true toe and not the tracked one, and toe tracking fails the challenge under both structures (6.36 and 5.68° from the toe).

3.9 Discussion

Three results generalize past this mill.

The first is that when the operating objective is a difference, controlling one term of it is not control. Equation (3.1) makes that explicit here, and Table 3.7 measures the consequence at 75 % of product coarsening with the impact angle held to 0.6° throughout. The same structure appears

wherever a loop holds a temperature against a dew point, a pressure against a bubble point, or a flow against a demand that nobody measures. The diagnosis is to write the objective down and check whether the measurement appears in it alone.

The second is that supervisory feedforward on a constraint is not free. Table 3.8 has one supervisory layer cutting the peak wear rate on one disturbance by a factor of 2.7 and, on another, multiplying the peak wear rate by 75 and driving the motor 0.61 MW past its limit, using exactly the same equation, because the constraint it tracks is itself the variable that is moving fastest. On a third, the ball-charge drop of the graded shift, it aims the wrong way, because the measurement its toe is inferred from cannot tell steel from rock. Any scheme that computes a setpoint from a model of the plant inherits the plant's failure modes, and the standard remedy is an interlock on the variable the model assumes is well behaved [10].

The third is that a faster tuning is not always a better one. Table 3.9 shows the recommended tuning multiplying the integral error by 3.3 and the actuator travel by 6.6 on a plant where the measurement is an inference and every disturbance is slower than the loop. A tuning study of this mill therefore ends not with the question of how tight the loop can be made but with the questions of what the disturbances look like and whether the measurement deserves to be believed [5, 17].

3.10 Conclusion

A semi-autogenous mill is aimed rather than regulated. The application computes where the media land from Davis ballistics in closed form, places the landing point on the shell to machine precision at every speed, and shows that the loop holding that angle is straightforward, at an integral error 15 times lower from the recommended tuning than from a cautious one on a 5° setpoint step. The remaining scenarios establish that holding the angle is not the problem the mill poses. A worn liner costs 5.4 points of critical speed and nothing else. A two-point drop in ball filling removes 40 t of steel, the load controller replaces it with rock, the toe rises three degrees, and the impact margin goes from +0.2 to −3.3° with the product coarsening by 75 % while all three controlled variables stay on setpoint. Tracking a toe inferred from the bearing pressure cannot see that drop and aims the wrong way. It cuts the peak liner wear rate by a factor of 2.7 on a starved conveyor, while in an overload the identical equation multiplies the wear rate by 75 and drives the motor past its limit, because the constraint it follows is the variable moving fastest. Over the graded shift the recommended loop tuning has 3.3 times the integral error of the tuning the scenario loads with, and only a slower load loop, which replaces the lost steel with rock without overshoot, keeps the media within the 3° the target allows from the toe. An optional predictive controller on the same balances crosses the load setpoint by 24 rather than 84 kPa when the ore hardens and feeds 4.1 % more ore under a common load limit, but the loops given its breakage estimate as a feedforward come within 0.4 %, and it loses 4.9 % on a starved belt. The mill rewards knowing where the target is rather than how tightly the setpoint is held.

KEY POINTS

- The mill grinds best when its falling media land on the toe of the charge.
- The impact-angle gain rises by a factor of 2.09 from 66 to 90 % of critical.
- Past the power maximum at $J = 0.432$ more load gives less power: the overload runaway.
- A two-point drop in ball filling coarsens the product by 75 % while all three controlled variables stay on setpoint.
- Toe tracking cuts peak wear 2.7 times on a starved conveyor and raises it 75 times in an overload.

Where else this applies. A limit at a peak in Volume 4, Chapter 2, regulation at a peak in Volume 7, Chapter 4, an inferred constraint in Volume 5, Chapter 10.

References

[1] J. Jeswiet and A. Szekeres, "Energy consumption in mining comminution," *Procedia CIRP*, vol. 48, pp. 140–145, 2016.

[2] R. Rajamani, J. Delgadillo, and V. Duriseti, "Online SAG mill pulse measurement and optimization," University of Utah, Metallurgical Engineering Department, Semi-annual technical progress report DE-FC26-04NT42088, 2006, u.S. Department of Energy, Office of Scientific and Technical Information.

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

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

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

[6] D. Wei and I. K. Craig, "Grinding mill circuits—a survey of control and economic concerns," *International Journal of Mineral Processing*, vol. 90, no. 1–4, pp. 56–66, 2009.

[7] B. A. Wills and J. A. Finch, *Wills' Mineral Processing Technology: An Introduction to the Practical Aspects of Ore Treatment and Mineral Recovery*, 8th ed. Oxford: Butterworth-Heinemann, 2015.

[8] T. J. Napier-Munn, S. Morrell, R. D. Morrison, and T. Kojovic, *Mineral Comminution Circuits: Their Operation and Optimisation*. Brisbane: Julius Kruttschnitt Mineral Research Centre, University of Queensland, 1996.

[9] S. B. Donkor, "On-line sensors for measuring the total ball and charge level in tumbling mills," MSc (Eng) thesis, University of Cape Town, Centre for Minerals Research, 2014, hdl.handle.net/11427/13225.

[10] D. Hodouin, "Methods for automatic control, observation, and optimization in mineral processing plants," *Journal of Process Control*, vol. 21, no. 2, pp. 211–225, 2011.

[11] J. D. le Roux, I. K. Craig, D. G. Hulbert, and A. L. Hinde, "Analysis and validation of a run-of-mine ore grinding mill circuit model for process control," *Minerals Engineering*, vol. 43–44, pp. 121–134, 2013.

[12] J. D. le Roux, A. Steinboeck, A. Kugi, and I. K. Craig, "An EKF observer to estimate semi-autogenous grinding mill hold-ups," *Journal of Process Control*, vol. 51, pp. 27–41, 2017.

[13] B. K. Mishra and R. K. Rajamani, "Numerical simulation of charge motion in ball mills: Lifter bar effect," *Mining, Metallurgy & Exploration*, vol. 10, no. 2, pp. 86–90, 1993.

[14] L. G. Austin, "A mill power equation for SAG mills," *Mining, Metallurgy & Exploration*, vol. 7, no. 1, pp. 57–63, 1990.

[15] F. K. Mulenga and M. H. Moys, "Effects of slurry filling and mill speed on the net power draw of a tumbling ball mill," *Minerals Engineering*, vol. 56, pp. 45–56, 2014.

[16] L. M. Tavares, "A review of advanced ball mill modelling," *KONA Powder and Particle Journal*, vol. 34, pp. 106–124, 2017.

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

[18] J. G. Ziegler and N. B. Nichols, "Optimum settings for automatic controllers," *Transactions of the ASME*, vol. 64, no. 8, pp. 759–765, 1942.

[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] X.-s. Chen, J.-y. Zhai, S.-h. Li, and Q. Li, "Application of model predictive control in ball mill grinding circuit," *Minerals Engineering*, vol. 20, no. 11, pp. 1099–1108, 2007.

[21] L. C. Coetzee, I. K. Craig, and E. C. Kerrigan, "Robust nonlinear model predictive control of a run-of-mine ore milling circuit," *IEEE Transactions on Control Systems Technology*, vol. 18, no. 1, pp. 222–229, 2010.

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

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

[24] 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] P. Quintanilla, F. Fernández, C. Mancilla, M. Rojas, and D. Navia, "Digital twin with automatic disturbance detection for an expert-controlled SAG mill," *Minerals Engineering*, vol. 220, p. 109076, 2025.

3.11 Exercises

3.1 Concept check · 1 min

The load loop of the mill holds one measured variable at its setpoint. What does that measurement respond to?

(a) The coarse rock hold-up alone, because the steel and the water stay fixed between ball additions

(b) The total volumetric filling J, because the gauge reads the volume that balls and rock occupy

(c) The mass of balls, rock and water together, through the trunnion bearing pressure

(d) The ore hardness, through the torque that the bearing transmits to the drive

3.2 Concept check · 2 min

Scenario 3 steps the mill speed by 3 % of critical, first from 78.8 % and later from 88 %. How does the impact-angle gain of the second step compare with that of the first?

(a) It is the same, because one parabola maps every speed to the angle at which the media land

(b) It is smaller, because the charge rearrangement lag absorbs more of a step at a higher speed

(c) It is smaller, because the media already land past the toe and the angle saturates there

(d) It is larger, because the shoulder angle set by $\cos\theta_d = \phi^2$ falls faster as speed rises

3.3 Concept check · 2 min

In the graded shift the ball filling falls from 12 to 10 % at 40 min. With the cautious tuning and a fixed setpoint, the impact angle, the mill load and the density all stay on setpoint, yet the product coarsens by 75 %. Which explanation is right?

(a) The load loop replaces the lost steel with rock, the filling rises, the toe rises and the fixed angle lands short of it

(b) The impact meter reads high once the steel leaves, so the speed loop holds a true angle well below its setpoint

(c) The mill power falls by a third with the steel, so less energy reaches the ore at the same speed

(d) The density loop adds water to hold the slurry, and the water buries the media in the charge

3.4 Theory and calculation · 6 min

A ball leaves the shell where $\cos\theta_d = \phi^2$, with θ_d measured from the top of the mill and ϕ the fraction of critical speed, and lands at the point of Equation (3.7), with $\psi = \operatorname{atan2}(-x_f, -y_f)$. The mill radius is $R = 5.2\,\mathrm{m}$ and $g = 9.81\,\mathrm{m/s^2}$. Give angles to 0.01°.

(a) Find θ_d and ψ at $\phi = 0.86$. *Checkpoint:* $\theta_d = 42.30°$ and $\psi = 53.09°$.

(b) The liner is 65 % worn and the drive runs at 78.8 % of critical. Find the effective speed of Equation (3.8) with $s_w = 10\,\%$, and the impact angle at that speed. *Checkpoint:* $\phi_{\text{eff}} = 0.7368$ and $\psi = 8.63^\circ$.

(c) Find the flight time $t_f = 4 \sin\theta_d \cos\theta_d / \omega$ at 78.8 % of critical on a new liner, with $\omega = \phi\sqrt{g/R}$. *Checkpoint:* 1.80 s.

3.5 Theory and calculation · 5 min

At fixed speed Equation (3.9) makes the mill power proportional to $M(1 - 0.937J)$, where the charge mass is $M = M_0 + \rho_r V(J - J_b/100)$ and M_0 is the mass of balls and water. Use $M_0 = 259.9\,\text{t}$, $\rho_r V = 801.5\,\text{t}$, $J_b = 12\,\%$ and $k_{\text{brg}} = 12.37\,\text{kPa/t}$.

(a) Show that the power is largest at the J^* of Equation (3.20), evaluate it and find the rock hold-up there. *Checkpoint:* $J^* = 0.4315$ with 249.7 t of rock.

(b) Find where the geometric term $J(1 - 0.937J)$ alone peaks, and say what places the true maximum earlier. *Checkpoint:* $J = 0.534$.

(c) The overload alarm is at $J = 0.45$. Find the rock between the power peak and the alarm and the bearing pressure that rock adds. *Checkpoint:* 14.8 t and 183 kPa.

3.6 Theory and calculation · 6 min

Continuous cycling of the impact loop under proportional control gives $K_u = 2.4\,\%$ of critical per degree and $P_u = 1.12\,\text{min}$ at the design point, where the local gain is 3.456 °/%. At 90 % of critical it is 5.276 °/% (Table 3.3). The lags and the meter do not change with speed.

(a) Find the Ziegler–Nichols PI settings, $K_c = 0.45K_u$ and $\tau_I = P_u/1.2$. *Checkpoint:* $K_c = 1.08$ and $\tau_I = 0.93\,\text{min}$.

(b) The recommended tuning has $K_c = 0.75$. Find the dimensionless loop gain, K_c times the local gain, at the design point and at 90 %, and the K_c that keeps the design loop gain at 90 %. *Checkpoint:* 2.59 and 3.96, and $K_c = 0.49$.

(c) Estimate the ultimate gain at 90 % and say what the sharp tuning, $K_c = 1.6$, would do there. *Checkpoint:* About 1.57, below 1.6.

3.7 App challenge · 8 min

Load scenario 3, Step test: the gain is not a constant (MANUAL), with the default seed. The mill starts at 78.8 % of critical. Press Run at 5×, set Mill speed to 81.8 % at 1.0 min, to 88.0 % at 4.0 min and to 91.0 % at 7.0 min, pause at 10.0 min and export the CSV file. A step made a few tenths of a minute later gives the same readings three minutes after it.

(a) Read the true impact angle, column `impact_deg`, at 1.0, 4.0, 7.0 and 10.0 min. *Checkpoint:* 25.16, 35.99, 62.20 and 77.53°, each within 0.05°.

(b) Compute the secant gain of the first and of the last step in degrees per percent of critical, and compare each with the local gains of Table 3.3. *Checkpoint:* 3.61 and 5.11 °/%.

(c) Read the meter column `impact_meas_deg` and the true angle at 2.0 min, and explain the difference. *Checkpoint:* 31.69° against 34.64°, within 0.05°.

3.8 App challenge · 8 min

Load scenario 4, First try in AUTO (weak tuning), with the default seed. The mill starts at 72 % of critical with the impact angle at 5° and its setpoint at 25°, under $K_c = 0.15$ and $\tau_I = 8\,\text{min}$.

(a) Run 50 min at 20× and export the CSV file. Find the first row at which `impact_deg` is within 2° of `impact_sp_deg`, and read the angle at 30 min. *Checkpoint:* 47.25 min, within 0.1 min, and 20.8°.

(b) Press Reset, set Impact loop gain to 0.75 and Impact loop integral time to 1.8 min before Run, and run 10 min. Find the largest true angle and the last row at which the angle is 1° or more from setpoint. *Checkpoint:* 35.4° and 3.75 min, each within one row.

3.9 App challenge · 9 min

Load the Challenge, Hold the impact zone through a shift change, with the default seed. Each run lasts 70 min, which takes 3.5 s at 20×. When the run ends, the challenge target in the progress checklist shows a result line with the worst setpoint error and the worst toe error. The exported CSV file gives time in hours for this run, and the rock hold-up is `state_coarse_t` plus `state_fine_t`.

(a) Run the case as shipped. Read the two errors, and find the peak rock hold-up and its time. *Checkpoint:* 0.6° and 4.0°, and 147.4 t at 0.783 h.

(b) Open Advanced, set Load loop gain to 0.60 and Load loop integral time to 6.0 min, press Reset and run again. Read the same three values and the last value of `metric_tonnes` in both runs. *Checkpoint:* 0.6°, 2.4° and 135.1 t at 0.918 h, with 1934.9 and 1898.9 t of ore, each within 0.1 t.

Explain why slowing the load loop passes a clause on the impact angle.

3.10 Application area · 4 min

Worn lifters slip the charge by $s_w = 10\,\%$ at a fully worn liner (Equation (3.8)). The fixed 25.0° setpoint is met with the charge turning at 78.76 % of critical (Table 3.6), and the drive spans 50 to 100 %.

(a) Find the drive command that holds that charge speed at 65 and at 100 % wear, and the points of drive span used at the end of the campaign beyond the new-liner 78.8 %. *Checkpoint:* 84.24 and 87.51 %, each within 0.02 %, and 8.7 points.

(b) A new-liner mill settled at 87.5 % of critical draws 14.8 MW. Does the motor limit bind at the end of the liner campaign in this model? Use Equation (3.9). *Checkpoint:* No.

(c) Name the signal an engineer would trend to follow the slip, which no instrument reports.

3.11 Application area · 4 min

A crew adds two percentage points of ball filling to the mill at its design point ($J = 0.2482$) while the load loop holds 5485 kPa. The ball charge has a bulk density of 4.65 t/m^3, the mill volume is 433.2 m^3, $\rho_r V = 801.5$ t and the bearing reads 12.37 kPa/t.

(a) Find the steel added and the rise in bearing pressure before the load loop acts. *Checkpoint:* 40.3 t and 498 kPa.

(b) The load loop returns the bearing to its setpoint by removing rock. Find the new filling and say which way the toe moves. *Checkpoint:* $J = 0.218$. The toe falls.

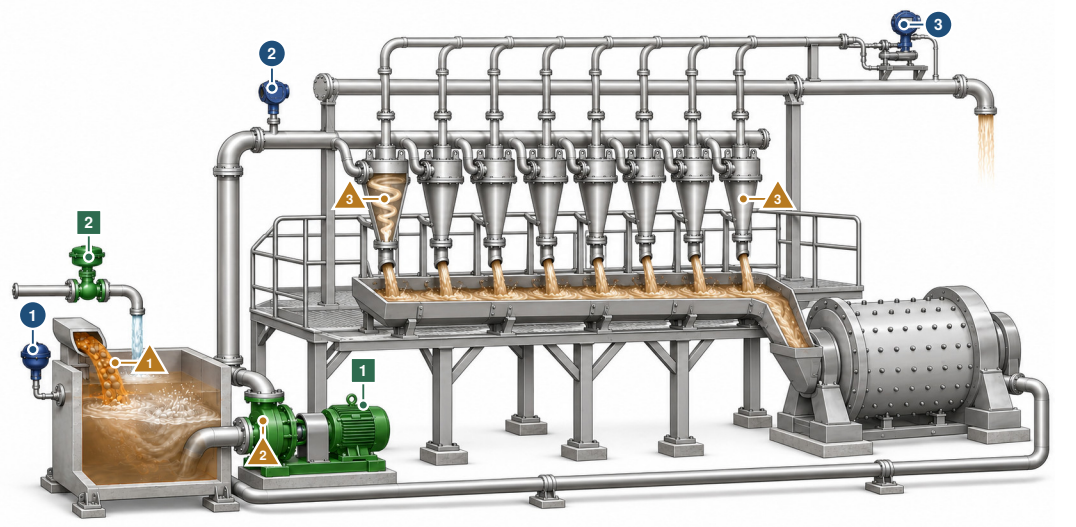
(c) Say what the operator should do with the load setpoint during the addition.

3.12 Review the agent · 5 min

An AI assistant reviews the graded shift and writes: “The ball charge drop at 40 min moves the toe, and a fixed 25° setpoint cannot follow it. Switch Impact setpoint to Track the toe. It recomputes the toe from the charge every scan, so the setpoint follows the toe up as the filling rises, and the run passes the ±3° toe clause with the cautious tuning unchanged.” Find the flaw and say how you would check it.

CHAPTER 4

Hydrocyclone Classification Simulator for Cascade Control and Hidden Loop Coupling



The controller moves 1 the pump speed, through the feed pump's motor drive
2 the dilution water into the sump, through its valve

The controller reads 1 the sump level 2 the pressure of the cyclone feed
3 the fineness of the overflow, from a size analyzer on a sample line

What disturbs it 1 coarser new ore 2 a worn pump impeller 3 worn cyclone liners

Notice the water falling into the sump: the level loop that sets it also sets how thick the cyclone feed is, and with it most of the product size. Illustration, not to scale.

IN PLAIN TERMS **Sorting Ground Ore in a Spinning Cone**

The main idea. Ground ore and water are pumped into cones that spin them. Coarse grains go back to the mill, and fine grains leave as product.

The challenge. Pumping harder makes the product a little finer, but it also draws more water into the tank in front of the pump. That thins the mix, and a thinner mix sorts much finer, so the loop that only keeps the tank level steady quietly sets most of the product size.

Why it matters. A fix that steadies one reading can push the product off target through a loop that nobody watches for quality.

THREE WAYS TO READ THIS CHAPTER

Explore	8 min	The first page, the history and fact boxes of Section 4.2, the fact box and the try-it box of Section 4.5, 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 4.5, Section 4.6, Section 4.7 with its open challenge

The control question A size loop on the cyclone overflow sets the setpoint of a fast pressure loop on the feed pump, and the pressure is a weak handle on the cut. Which disturbances does that cascade take out of the product, and where does the outer loop's authority over the cut size actually come from?

Learning objectives

- Sort a disturbance by whether it enters the pressure measurement, and predict whether the cascade helps, hurts or does nothing.
- Split a change in cut size into a pressure part and a feed density part with the cut-size correlation.
- Compute the sump's integrating gain and hold-up and the circuit's recycle balance from the model.
- Explain why a tighter level loop holds the level better and the product no better.
- Commission the three loops from the fastest outward and judge the result against the challenge's four bands.

Prerequisites Sections 1.4.2, 1.4.6, 1.6.2 and 1.6.4.

The app https://apmonitor.com/apps/hydrocyclone/

See also Volume 3, Chapter 2 (disturbances sorted by where they enter). Volume 3, Chapter 3 (how fast an inner loop must be). Volume 6, Chapter 7 (a recycle that amplifies a small change).

4.1 Introduction

The product specification on a grinding circuit is written on the cyclone overflow, and the actuator nearest to it is the cyclone feed pump. Between the two sits an empirical correlation whose exponent on feed pressure is –0.225, so a 60 % increase in pressure returns a 10 % reduction in cut size. That is a weak handle, and the pump power required to obtain it rises with the cube of speed. The exponent alone does not justify placing the product size controller on the pump, yet that is the standard configuration in concentrators.

The justification is in the sump rather than in the correlation. The vessel in front of the pump holds 0.26 min of slurry at design flow, and its only spare volume is dilution water, so the level controller owns the single degree of freedom in the water balance. A larger draw on that sump lowers the level, the level controller adds water, the cyclone feed thins, and the cut-size correlation is exponential in feed solids. Stepping the pressure setpoint of the application from 90.20 kPa to 110.0 kPa moves the pump from 65.01 % to 72.32 %, the dilution water from 653.8 m^3/h to 753.0 m^3/h, the feed from 44.82 % to 43.13 % solids, and the cut size from 151.4 µm to 134.3 µm. Decomposing the logarithm of that cut-size ratio assigns 37.1 % to the pressure term and 62.9 % to the density term. Two thirds of the outer loop's steady-state authority therefore passes through the level controller, which appears on no control diagram as a product quality loop, and it arrives within half a minute rather than over the quarter hour the circulating load requires. The feed density covers 63 % of its final change 0.43 min after the step.

The first subject of this chapter is more conventional and is the reason the circuit is useful as a teaching plant. The feed pressure responds to pump speed with a gain of 2.59 kPa/% and a time constant of 7.5 s, while the mill recycle has a time constant of 5.1 min and reaches 95 % of its change in 15.8 min. A separation of 41 to 1 in dominant time constant, and of more than 100 to 1 in 95 % response time against the closed pressure loop, satisfies the standard condition for a cascade [1, 2] (Section 1.4.6), and the application quantifies the benefit. An impeller worn to 86 % of design delivery, applied identically under both architectures, produces 154.9 % min of integrated fineness error and 33.3 min off specification when the fineness controller drives the pump through a fixed characteristic, and 3.2 % min with 0.1 min off specification when a pressure loop is placed between them. A coarser new feed produces 53.7 % min under the single loop and 53.1 % min under the cascade, because ore coarseness does not appear in a pressure transmitter.

Classifying disturbances into those two groups before any tuning is performed is the analysis the plant is built to teach, and Volume 3, Chapter 2 makes the same sort on a level loop. On this

circuit the classification is not resolved by the block diagram alone. Worn cyclone liners appear in the pressure measurement and belong to the first group by the usual test, yet the cascade responds to them by moving the pump, which moves the density, which moves the product by 5.94 % of fineness at the peak against 1.97 % when the pump is left where a static characteristic put it. An inner loop that returns its own measurement to setpoint has rejected only the part of the disturbance that lived in that measurement.

The experiments the application supports are the standard content of a process dynamics and control course. The sump is an integrating process with 0.26 min of hold-up, the pressure loop is a first-order plant with a sampled measurement, the size analyzer contributes 2 min of dead time to the outer loop, the mill recycle is a positive feedback path that slows the circuit to a quarter hour and drives the size controller to its limit when the mill grinds poorly, and the three loops interact through a variable that is not measured in either of them. Cascade structure, integrating level control, anti-windup on a saturating actuator, dead time in a sampled loop, loop interaction and commissioning order are all exercised on one plant. The physical fundamentals are equally standard. The model is a set of conservation statements on water and on two solids classes in a perfectly mixed sump, closed by an empirical statement of the centrifugal separation and by a transport lag around the mill, and the time-scale separation between an algebraic cyclone, a sump that turns over in seconds and a quarter-hour recycle is what makes the reduction legitimate [1, 3].

This chapter works with a browser-hosted grinding circuit classifier, an eight-state circuit with a genuine recycle, three architectures and eleven scenarios, each with a completion requirement. It quantifies the cascade against six disturbances sorted by whether they enter the inner measurement, and it decomposes the outer loop's gain into the designed path and the path through the level controller. Two measured outcomes carry the analysis. The cascade divides the integrated fineness error by 49 for pump wear and by 1.01 for feed coarseness, and 63 % of the outer loop's steady-state authority is carried by the level controller. A third result is a commissioning order whose justification is a moving plant gain rather than a convention. The cut-size correlation, the partition curve and the cascade arithmetic are all standard.

Section 4.2 reviews how the cyclone displaced the classifiers it replaced and why its control rests on an instrument with a cycle time of minutes. Section 4.3 sets out the sump, pump, cyclone and mill model and its design point. Section 4.4 describes the three loops, the two architectures for the size controller and their tuning. Section 4.5 reports the disturbance comparison, the decomposition of the outer loop's gain, the effect of level tuning, the two unmeasured constraints and the graded challenge. Section 4.6 sets the loops on real instruments and platforms, Section 4.7 states what the model leaves out, Section 4.8 draws the general lesson for cascades, Section 4.9 summarizes, and the exercises in Section 4.10 work through the scenarios.

4.2 Background

Wet classification decides the size at which ore leaves a grinding circuit, and before the hydrocyclone it was performed by rake and spiral classifiers, which are large, mechanically complex and limited to relatively coarse separations. The cyclone displaced them because it has no moving parts, occupies a fraction of the floor area, and separates finer. A modern circuit runs a cluster of them in parallel on a common manifold, with individual cyclones valved in and out to match the throughput, and the assembly is described by a correlation rather than by a mechanism [4, 5].

The correlations have a long history in the mineral processing literature. Empirical forms relating the cut size to feed pressure, spigot and vortex finder diameters and feed solids concentration were fitted to plant and pilot data through the 1970s and remain the working tools of circuit design, and their coefficients have been refitted and compared repeatedly as the data sets have grown [6]. A parallel line of work builds the separation from the fluid mechanics of the swirling flow and the settling of individual particles in it, which gives a mechanistic account of the shape of the correlation

but does not replace it for design [7]. Both traditions agree on the exponents that matter here, which are small on pressure and large on the geometry and the feed density.

> ⧖ HISTORY **A Cone That Purified Water (1891)**
>
> A United States patent of 26 May 1891, filed in November 1890 by Eugene Bretney of Indianapolis, describes a water purifier that frees water of impurities by centrifugal force. Water under pressure enters the top of a tank shaped as an inverted cone through a pipe set at one side, so that the inflow is tangential and the water whirls. The impurities driven to the wall descend to an adjustable annular opening at the point of the cone, and the tapering form, the patent explains, leads them there in a nearly straight line. The purified water leaves through a pipe that starts on the axis just below the cover and passes up through it [8]. The opening at the point is what this chapter calls the apex.

The physical content of the model is elementary. Mass is conserved separately on water and on the solids in each size class, the sump is an integrator because its outflow is set by a pump rather than by its own inventory, and the driving force for separation is the centrifugal field that the pressure drop creates inside a cyclone with no moving parts. Textbook treatments of process dynamics develop exactly these elements, integrating inventory, first-order actuator dynamics, recycle and the conditions under which a fast loop may be treated as an algebraic gain by a slower one [1, 3, 9]. What the circuit adds is that all of them appear together in a plant small enough to run in a browser.

Comminution accounts for a large share of the energy used in a concentrator. Estimates place the comminution share of total mine site energy in the tens of percent, with the associated emissions following [10]. Because the classifier sets the size at which material leaves the circuit, it also sets how much of that energy is spent regrinding material that was already fine enough, and it sets the size distribution presented to flotation or leaching. That is the industrial reason the classification loops are among the most valuable regulatory loops in the plant.

> ✦ DID YOU KNOW? **Dense Grains Ground Twice**
>
> Plant sampling of an iron ore grinding circuit, reported to the United States Department of Energy in 2005, found that the largest single source of overgrinding was the return of dense material to the mills by the hydrocyclones. Simulations indicated about 50 percent more throughput from the same number of mills once the circulating load was reduced [11].

The control literature on grinding circuits reports the difficulty rather than the value. Surveys of milling circuit control find that most industrial circuits run on single-loop regulatory control with an operator supplying the supervisory layer, and that the economic case for advanced control is routinely made and less routinely realized [12]. The measurement problem is a large part of the difficulty. Particle size is measured on stream by instruments with cycle times of minutes, ore properties are not measured at all, and the states a model-based controller would use are not observable [13]. Model development for these circuits has concentrated on obtaining a representation simple enough to be identified from plant data and rich enough to hold the recycle, and validation against operating data remains the hard part [14]. Supervisory schemes that manipulate sump water and cyclone feed rate against inferred size have been reported for ball mill circuits since the 1990s [15], and feedforward from measured feed size distribution is an active line of work for exactly the disturbance a cascade cannot see [16].

However, a cascade presented in a textbook is a pair of transfer functions and a rule about time-scale separation. That presentation does not show a cascade that is correctly structured, correctly tuned, and delivering most of its steady-state authority through a variable that appears in neither of its two loops. The application of this chapter provides a plant on which that distinction is measurable.

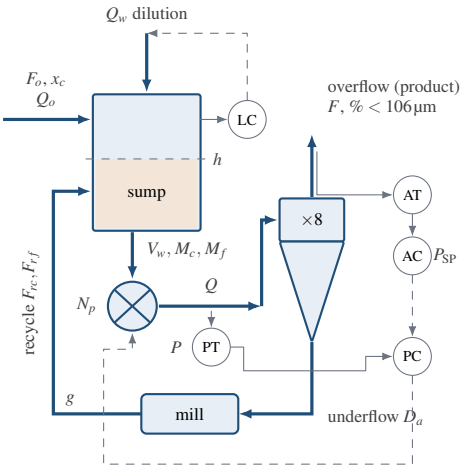


Figure 4.1 Closed grinding circuit as implemented. Sump inventory V_w, M_c, M_f under level control LC manipulating dilution water Q_w, a variable-speed pump N_p delivering Q to eight cyclones, a pressure transmitter PT and pressure controller PC forming the inner loop, and a particle size analyzer AT with controller AC writing the pressure setpoint P_{SP}. Underflow of apex diameter D_a returns to the mill, where a fraction g of the coarse is ground, and re-enters the sump as F_{rc} and F_{rf}.

4.3 Process Model

The plant is a closed grinding circuit reduced to the two units that decide the product. New ore and water enter a sump. A variable-speed centrifugal pump draws from the sump and feeds a manifold of eight identical hydrocyclones. The cyclone overflow is the product and leaves the circuit. The underflow returns to a ball mill, is ground, and comes back to the sump. Dilution water is added to the sump under level control. The arrangement, the instruments and the symbols used in the equations are drawn in Figure 4.1.

4.3.1 Working Units and the Design Point

The clock is in minutes, solids in t/h, slurry in m^3/h, pressure in kPa and size in µm. At the design condition the circuit takes 250 t/h of new feed at 85 % coarse with 110 m^3/h of water, the pump runs at 65 % and delivers 1380 m^3/h to eight cyclones at 90.19 kPa, the level controller holds 50 % with 653.8 m^3/h of dilution water, and the cyclone feed is 862 t/h of solids at 44.82 % by weight, which is 23.13 % by volume. The cut size is 151.4 µm, the overflow is 81.44 % passing 106 µm, the underflow is 67.3 % solids and the circulating load is 245 % of new feed.

4.3.2 Sump Inventory

The sump is treated as perfectly mixed in three quantities, water volume and two solids classes:

$$\frac{\mathrm{d}V_w}{\mathrm{d}t} = \tfrac{1}{60}\left(Q_w + Q_u + Q_o - Q\,\tfrac{V_w}{V}\right) \tag{4.1}$$

$$\frac{\mathrm{d}M_c}{\mathrm{d}t} = \tfrac{1}{60}\left(F_o x_c + F_{rc} - \tfrac{60 M_c}{\tau_s}\right) \tag{4.2}$$

$$\frac{\mathrm{d}M_f}{\mathrm{d}t} = \tfrac{1}{60}\left(F_o(1-x_c) + F_{rf} - \tfrac{60 M_f}{\tau_s}\right) \tag{4.3}$$

with total volume, hold-up time and level

$$V = V_w + \frac{1000\,(M_c + M_f)}{\rho_s}, \quad \tau_s = \frac{V}{Q/60}, \quad h = \frac{100\,V}{A_s H_s}. \tag{4.4}$$

Table 4.1 Controlled, Manipulated and Disturbance Variables

Symbol	Meaning	Unit	Range
Controlled variables			
F	overflow passing 106 µm	%	60–95
h	sump level	%	0–100
P	cyclone feed pressure	kPa	0–190
Manipulated variables			
N_p	feed pump speed	%	25–95
Q_w	dilution water	m^3/h	50–1800
P_{SP}	pressure setpoint (cascade output)	kPa	40–170
Disturbance variables			
F_o	new feed to the circuit	t/h	120–420
x_c	coarse fraction of the new feed	–	0.65–0.96
Q_o	water with the new feed	m^3/h	30–320
n	cyclones open	–	5–12
w_p	pump delivery multiplier	–	0.70–1.05
w_c	cyclone liner condition	–	0.65–1.10
r_o	ore classification multiplier	–	0.80–1.30
g_r	mill grinding effectiveness	–	0.40–1.30
States			
V_w	sump water volume	m^3	—
M_c, M_f	sump coarse and fine solids	t	—
F_{rc}, F_{rf}	mill recycle coarse and fine	t/h	—
N_p, Q_w, D_a	actuator positions	—	—

Here Q_u is the water returning with the mill discharge and ρ_s is the solids density. Equations (4.1) to (4.3) are statements of conservation of mass on three components that share one vessel, and each carries the factor $\frac{1}{60}$ because the clock is in minutes while the streams are in hours. The pump is a linear characteristic in speed,

$$Q = k_p N_p w_p, \tag{4.5}$$

which is the relation that the impeller wear disturbance degrades and that the single-loop architecture inverts to obtain a speed from a pressure setpoint.

4.3.3 Cyclone Pressure and Cut Size

Feed pressure follows the square of the flow through each individual cyclone and scales with the slurry density,

$$P = k_h \left(\frac{Q}{n} \right)^2 \frac{\rho_{sl}}{1395\,\mathrm{kg/m^3}} w_c, \tag{4.6}$$

so opening a cyclone lowers the pressure at fixed total flow and worn liners lower it at fixed flow and density. The cut size uses a reduced form of the established empirical correlation [6],

$$d_{50} = d_{50,0} \left(\frac{P}{P_0} \right)^{-\beta} \left(\frac{D_a}{D_{a,0}} \right)^{-\delta} \mathrm{e}^{\gamma(\phi_v - \phi_{v,0})} r_o, \tag{4.7}$$

with $\beta = 0.225$ on pressure, $\delta = 0.71$ on apex diameter, $\gamma = 0.063/\%$ on the volumetric feed solids ϕ_v and r_o an ore classification multiplier that is unity at the design condition. All three exponents

carry the same sign convention. Raising the pressure, opening the apex or thinning the feed each produces a finer cut. The magnitudes are unequal by construction. Pressure is weak, apex diameter is a little over three times stronger, and the feed solids term is exponential rather than a power law, which is why the 1.20 % change in ϕ_v of Table 4.5 moves the cut size further than the 22 % change in pressure that produced it.

The split of each size class between overflow and underflow is a corrected partition curve of logistic form with the water recovery added as a bypass,

$$E_i = R_f + (1 - R_f)\,\frac{1}{1 + (d_{50}/d_i)^m}, \tag{4.8}$$

evaluated at a representative coarse size $d_c = 300\,\mu\text{m}$ and fine size $d_f = 55\,\mu\text{m}$, with m the sharpness. The corrected term is one half at $d_i = d_{50}$, which is the definition of the cut size. The water recovery to underflow follows the apex,

$$R_f = R_{f,0}\left(\frac{D_a}{D_{a,0}}\right)^{a_{Rf}}, \tag{4.9}$$

so closing the apex sends less water to the underflow and thickens it.

Roping is represented by blending the partition of (4.8) toward the total solids mass split S_u as the underflow density rises,

$$E_i^* = \sigma E_i + (1 - \sigma) S_u, \quad \sigma = \text{sat}\left(\frac{\rho_{rope} - \rho_u}{\rho_{rope} - \rho_{soft}}, 0, 1\right), \tag{4.10}$$

where ρ_u is the underflow percent solids and σ runs from unity at a healthy apex to zero at a rope. At $\sigma = 0$ every size class reports to the underflow in the proportion in which it arrived, which is the definition of no separation, so (4.10) is the mechanism by which a roped cyclone stops classifying rather than classifying differently. Because S_u and ρ_u depend on the split they determine, (4.8) to (4.10) are solved by fixed-point iteration at each evaluation.

4.3.4 Circulating Load and Mill Recycle

Underflow solids go to the mill, where a fraction is ground from coarse to fine and the whole stream returns after a transport lag:

$$\frac{\mathrm{d}F_{rc}}{\mathrm{d}t} = \frac{U_c\,(1 - g_e) - F_{rc}}{\tau_m}, \qquad \frac{\mathrm{d}F_{rf}}{\mathrm{d}t} = \frac{U_f + U_c g_e - F_{rf}}{\tau_m}, \tag{4.11}$$

$$g_e = \min\left(1,\ \frac{F_{mill,\max}}{U_c + U_f}\right) g\, g_r. \tag{4.12}$$

Equations (4.11) make the circulating load a genuine positive feedback path with loop gain $E_c(1 - g_e)$ rather than a fitted transfer function. This is the recycle gain g of Equation (1.1), with the share g_e ground per pass in the place of Chapter 1's b, and the same positive feedback loads a reactor's recycle in Volume 6, Chapter 7. At the design condition $E_c = 0.918$ and $g_e = 0.32$, so the coarse recirculation gain is 0.624 and the total solids recirculation ratio is 0.71. The pressure setpoint step of Table 4.5, which raises the cyclone feed rate by 47 t/h, moves the recycle streams of (4.11) with a 63 % response time of 5.1 min and a 95 % response time of 15.8 min, so a 3 min transport lag becomes a quarter hour of circuit dynamics. Lowering the mill grinding effectiveness g_r raises that loop gain toward unity, where the circuit would have no stable operating point. Within the range the application offers it stops short of unity. At $g_r = 0.4$ with the size controller saturated the gain settles at 0.80 and the circulating load at 412 %.

Table 4.2 Model Parameters. Those Marked A Are Adjustable

Symbol	Meaning	Value	Unit
A_s	sump cross-section	4.0	m^2
H_s	sump depth	3.0	m
k_p	pump capacity per unit speed	21.23	$m^3/(h\,\%)$
τ_{pu}	drive lag	0.10	min
τ_w	dilution water valve lag	0.15	min
τ_a	apex actuator travel time	1.2	min
n	cyclones open	8	—
k_h	cyclone resistance coefficient	3.035×10^{-3}	$kPa\,h^2/m^6$
$d_{50,0}$	reference cut size (A)	150	µm
P_0	reference pressure	90	kPa
β	exponent on pressure (A)	0.225	—
δ	exponent on apex diameter (A)	0.71	—
γ	sensitivity to feed solids (A)	0.063	1/%
$\phi_{v,0}$	reference feed solids	22.97	% vol
$D_{a,0}$	reference apex diameter	100	mm
m	partition sharpness (A)	3.0	—
$R_{f,0}$	water recovery at reference apex	0.28	—
a_{Rf}	water split exponent on apex	1.8	—
d_c	representative coarse size	300	µm
d_f	representative fine size	55	µm
d_{spec}	product specification mesh	106	µm
g	coarse ground to fine per pass (A)	0.32	—
τ_m	mill transport lag (A)	3.0	min
$F_{mill,\max}$	mill throughput limit	1400	t/h
ρ_{rope}	underflow density at roping	80	% solids
ρ_{soft}	density at loss of sharpness	72	% solids
$P_{\max}$	feed pressure limit	190	kPa
τ_{psa}	analyzer averaging	0.4	min
θ_{psa}	analyzer dead time (A)	2.0	min
τ_{fh}	level transmitter filter	0.03	min
τ_{fP}	pressure transmitter filter	0.02	min
Δt_c	controller scan	0.02	min
ρ_s	solids density	2700	kg/m^3

4.3.5 Initialization

The application is initialized from a converged steady state rather than from an assumed one. At reset, the dilution water that closes the sump water balance at the operator's pump speed is found by direct calculation, the recycle streams are converged by damped substitution of (4.11) until neither moves by more than 1×10^{-11}, and the sump inventory is split between water and the two solids classes at the cyclone feed density implied by that converged state. The residual on every derivative at $t = 0$ is below 1×10^{-3} in relative terms, the remaining error being the difference between the exactly balanced water rate, 653.759 m^3/h, and the rounded value carried as the operator input, 653.8 m^3/h,

which moves the level by less than a fifth of a percent in half an hour. The pressure transmitter and the size analyzer start on the pressure and fineness of that state, so a case that starts off the design point, or a setting chosen before the run, is measured correctly from the first scan.

4.3.6 Assumptions

1. **Two size classes.** Solids are coarse or fine at two representative sizes rather than a distribution. The specification, percent passing 106 µm, is computed from the split of those two classes. A real circuit is described by a full distribution and the specification is one point on it, so the application cannot show a bimodal feed or the shape change that a roping cyclone actually produces.

2. **The cyclone is algebraic.** Its residence time is a fraction of a second against a 0.26 min sump, so (4.6) through (4.10) are evaluated instantaneously. This breaks down only for transients faster than the pressure loop itself.

3. **The correlation is reduced.** Vortex finder diameter, inlet area, cone angle and cyclone diameter are fixed and absorbed into $d_{50,0}$ and k_h. Fitted constants for a real installation differ from one another by more than the range this application explores [6].

4. **Roping is smooth.** The sharpness degrades continuously with underflow density. Real roping is abrupt, hysteretic and often intermittent between cyclones on the same manifold, and an operator observes it as a change in the shape of the discharge spray rather than as a number.

5. **The mill is one efficiency and one lag.** There is no ball charge, no power draw and no relationship between the distribution entering the mill and the one leaving it. Mill dynamics are the subject of Chapter 3.

6. **The sump is perfectly mixed.** Real sumps stratify, and a nuclear level gauge reads a density-weighted level that moves when the density moves, which would couple the level measurement to the density path directly.

7. **No rheology.** Water and solids balances only. Slurry viscosity is one of the physical reasons a dense cyclone feed classifies badly and it is absent here, so the exponential density term of (4.7) carries the whole of that effect.

8. **Fixed-step integration.** Classical fourth-order Runge–Kutta at 0.005 min with a 0.02 min controller scan (Section 1.6.1). Every run reproduces bit for bit on any machine, which a grinding circuit does not.

4.4 Controller Design

Cascade structure and the conditions under which it pays are standard material [1, 2] (Section 1.4.6), and the arrangement here is the conventional one. The application implements three loops and offers three architectures. The level and pressure loops are always available. The fineness controller can be switched off, wired to the pump through a fixed characteristic, or cascaded onto the pressure setpoint, and the comparison between the last two is the subject of Section 4.5.

4.4.1 Velocity Form PID Algorithm

All three loops run the velocity-form algorithm of Section 1.6.2 [17], Equations (1.10) to (1.13), with the output bias carried in the integral term, which starts at the actuator position. The scan period $\Delta t_c = 0.02$ min (1.2 s) is common to the three loops. The pressure and fineness loops are proportional-integral. The level loop also carries a derivative term, zero by default, taken on the measurement after a first-order filter of $\tau_f = 0.05$ min, so that derivative action added to the level loop does not amplify transmitter noise. The limit acts on the output alone, and back-calculation with

tracking time $T_t = \tau_I$ corrects the integral, so each scan removes $\Delta t_c/\tau_I$ of the clipped excess, one three-hundredth for the fineness loop at $\tau_I = 6\,\text{min}$. Manual operation, and a loop that the chosen architecture leaves idle, track the actuator value in force, and in automatic mode the manual actuator values follow the controller outputs after every scan, so transfer in either direction is bumpless. The actuator spans are 25 % to 95 % on the pump, 50 m^3/h to 1800 m^3/h on the dilution water and 40 kPa to 170 kPa on the pressure setpoint written by the cascade, and no output carries a rate limit. All three loops are reverse acting in the ISA sense, action $a = +1$ in Chapter 1's convention, because more pump speed raises the pressure, more pressure raises the fineness and more dilution water raises the level, so a measurement above setpoint has to reduce the output. With the error written as SP − PV, that is what a positive K_c delivers.

4.4.2 Level Loop and Its Coupling to Feed Density

The level controller manipulates dilution water against a sump holding 0.26 min. The plant is integrating, and the volume-balance gain follows directly from (4.4):

$$\left.\frac{\mathrm{d}h}{\mathrm{d}t}\right|_{Q_w} = \frac{100}{60 A_s H_s} = 0.139\,\%/\text{min per m}^3/\text{h}. \tag{4.13}$$

Equation (4.13) is exact at the instant of a step in Q_w, because the composition terms in (4.1) to (4.4) cancel and leave the total volume responding only to the difference between the volumetric inflow and outflow. The pump removes a fixed total volume whatever the composition of the sump, so the relation holds closely after the step as well. The slope measured on the application in manual mode between 0.35 and 0.85 min after a step of 50 m^3/h is 0.141 %/min per m^3/h, which is 1.4 % larger, and it rises to 0.145 between 0.85 and 1.5 min as the underflow returns part of the extra water to the sump. For tuning, the sump is an integrator with the gain of (4.13). Default tuning is $K_c = 4\,(\text{m}^3/\text{h})/\%$ and $\tau_I = 8\,\text{min}$, and the tuning the application suggests after commissioning is $K_c = 18\,(\text{m}^3/\text{h})/\%$ with $\tau_I = 1.5\,\text{min}$.

4.4.3 Inner Pressure Loop and SIMC Tuning

The inner loop manipulates pump speed against feed pressure. Perturbing (4.5) and (4.6) about the design point, with the water balance closed, gives 2.59 kPa/%. The 63 % response time of the transmitted pressure to a manual pump step from 65 % to 66 % is 7.5 s, set by the 0.10 min drive lag and the 0.02 min transmitter filter with no dead time worth reporting. Applying the SIMC rules of Skogestad [18] to a first-order plant with a closed-loop time constant set equal to the open-loop one, $\tau_c = \tau_p = 0.125\,\text{min}$ and $\theta_p = 0$, gives

$$K_c = \frac{\tau_p}{K_p\,(\tau_c + \theta_p)} = \frac{1}{K_p} = 0.39\,\%/\text{kPa}, \qquad \tau_I = \tau_p. \tag{4.14}$$

The shipped default is $K_c = 0.15\,\%/\text{kPa}$ with $\tau_I = 2\,\text{min}$, deliberately conservative so that the effect of tightening the loop can be measured against a slack starting point, and the tuning the application suggests is $K_c = 0.8\,\%/\text{kPa}$ with $\tau_I = 0.15\,\text{min}$. How much faster than the outer loop an inner loop must be is measured directly in Volume 3, Chapter 3, where the lags lie within a factor of two of one another.

4.4.4 Outer Fineness Loop and the Two Architectures

The single-loop architecture inverts a calibrated pump characteristic. Given a demanded pressure it returns the speed that produced that pressure on a clean pump, at nominal density, with eight cyclones open:

$$N_p = \mathcal{N}^{-1}(P_{\text{SP}})_{\,w_p=1,\; w_c=1,\; n=8,\; \phi_v=\phi_{v,0}}\,. \tag{4.15}$$

Every disturbance that invalidates one of those four subscripts passes straight to the product. The cascade replaces (4.15) with the pressure loop, so the fineness controller writes P_{SP} and the inner loop finds whatever speed that requires.

Table 4.3 The Eleven Scenarios and the Mechanism Each Isolates

#	Title	Mechanism
1	Sump hold-up (manual)	0.26 min of inventory sets the fastest loop
2	What moves the cut size	the three exponents of (4.7)
3	Level control	feed rate raises level, water is removed, feed thickens
4	Inner loop alone	2.59 kPa/%, 7.5 s
5	Fineness without a cascade	pump wear through a static characteristic
6	The same wear, cascaded	pump wear inside the inner measurement
7	Where the authority is	fineness target step, split between the two paths
8	Circulating load runaway	recycle gain $E_c(1-g_e)$ raised
9	Roping	closed apex, then more feed, collapses the partition
10	Liners and a ninth cyclone	disturbances the inner loop hides in part
11	Challenge	four disturbances and noisy instruments over 130 min

Manual mode offers two helper buttons that perform these calculations by hand. Balance the sump returns the dilution water that closes the sump water balance at the present pump speed, which is the operation the level controller performs continuously and which makes the density coupling explicit. Hold this pressure returns the pump speed and dilution water that produce the pressure setpoint at steady state for the pump, cyclones and feed as they stand. It plays the role of (4.15) without its fixed calibration, and a wear disturbance applied afterward shows why such a characteristic ages.

The outer loop's plant gain is not the pressure exponent of (4.7). Measured on the closed-loop circuit with the level controller in service it is 0.275 %/kPa of fineness at the design point, falling to 0.231 %/kPa as a secant across the 19.8 kPa step of Table 4.5, and 63 % of it arrives through the density path. That path is also the fast one. After the step of Table 4.5 the feed density covers 63 % of its final change in 0.43 min, against the 5.1 min of the recycle, and it overshoots, reaching 41.80 % solids one minute after the step before the returning solids bring it back to 43.13 %. The shipped default tuning is $K_c = 0.3\,\mathrm{kPa/\%}$ with $\tau_I = 40\,\mathrm{min}$, and the scenarios that measure performance use $K_c = 1.2\,\mathrm{kPa/\%}$ with $\tau_I = 6\,\mathrm{min}$, obtained by treating the density path rather than the recycle as the dominant dynamic. The challenge scenario suggests a tighter $K_c = 2.0\,\mathrm{kPa/\%}$ with $\tau_I = 4\,\mathrm{min}$. A tuning derived from the recycle, $K_c = 7\,\mathrm{kPa/\%}$, drives the output between its limits of 40 kPa and 170 kPa with a period of 6.0 min in scenario 7.

4.5 Results

Table 4.3 lists the eleven scenarios and the mechanism each isolates. They run on the framework of Section 1.6.4, with a briefing, a completion rule, a scorecard and a CSV export for each, and every IAE quoted below is the scorecard's, integrated from the start of the run. Scenarios 1 and 2 operate the sump and the cyclone by hand, scenarios 3 and 4 close the level and pressure loops one at a time, which is the commissioning order the plant later justifies, scenarios 5 to 7 make the cascade argument and its decomposition, scenarios 8 and 9 are the two constraints that no instrument reports,

Table 4.4 Integrated Absolute Fineness Error Over 120 min and Peak Deviation, Same Disturbance Applied Under Both Architectures

	single loop		cascade	
Disturbance	IAE % min	peak %	IAE % min	peak %
Impeller wear to 0.86	154.9	13.08	**3.2**	2.81
Liner condition to 0.80	25.1	1.97	73.1	5.94
Ninth cyclone opened	26.6	2.08	76.8	6.20
Coarse feed 0.85 to 0.91	53.7	1.61	53.1	1.63
New feed 250 to 320 t/h	199.0	7.13	210.1	8.46
Grinding effectiveness 0.6	432.5	7.88	433.8	8.03

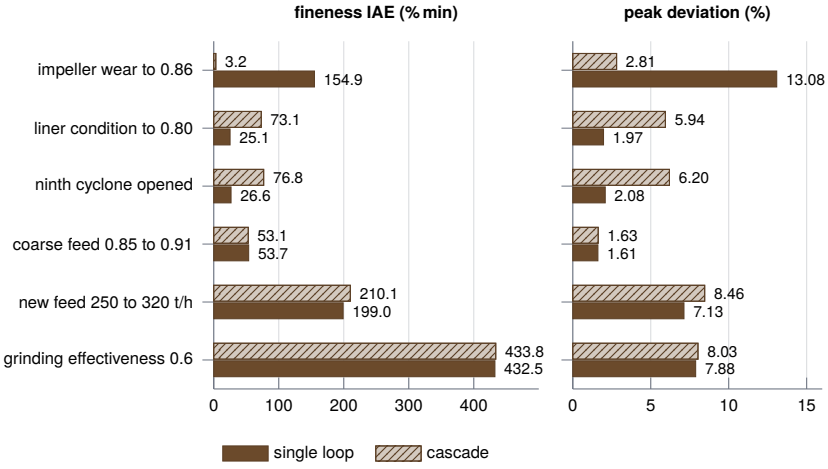


Figure 4.2 Table 4.4 as bars, single loop solid and cascade hatched. Notice the three groups: the cascade wins only on impeller wear, loses on the liners and the ninth cyclone, and changes nothing on the three disturbances that never reach the pressure transmitter.

scenario 10 shows a cascade making two resistance disturbances worse, and scenario 11 grades the whole commissioning sequence against four bands at once.

4.5.1 Disturbance Rejection Under Both Architectures

Table 4.4 applies six disturbances at $t = 6$ min under both architectures with identical outer-loop tuning and reports the integrated absolute fineness error over 120 min, read from the application's scorecard. Each run starts from scenario 6, with the architecture chosen before the run and the disturbance moved at 6 min. The ordering differs from the ordering a block diagram predicts.

Three groups appear, and Figure 4.2 draws them. Impeller wear is the case the textbook describes. It enters the inner measurement, the inner loop returns the pressure to within 2 kPa of setpoint in 10 s, and the integrated error falls by a factor of 49. Feed coarseness, feed rate and grinding effectiveness are invisible to a pressure transmitter, and the two architectures are within 6 % of one another on all three.

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

Table 4.5 Pressure Setpoint Step From 90.20 kPa to 110.0 kPa With the Level Loop in Service

Quantity	before	after	change
Pump speed N_p (%)	65.01	72.32	+7.31
Dilution water Q_w (m³/h)	653.8	753.0	+99.2
Feed solids (% wt)	44.82	43.13	−1.69
Feed solids ϕ_v (% vol)	23.13	21.93	−1.20
Cut size d_{50} (µm)	151.43	134.28	−17.15
Fineness F (%)	81.44	86.02	+4.58
Decomposition of the change in $\ln d_{50}$			
Pressure factor $(P/P')^{-\beta}$		0.95633	37.1 %
Density factor $e^{\gamma\Delta\phi_v}$		0.92721	62.9 %

The middle group carries the result of interest. Worn liners and an extra cyclone both lower the resistance and therefore both appear in the pressure measurement, which places them in the first group by the usual test. The cascade makes them worse by a factor of three in integrated error. The reason is the density path. A single loop leaves the pump speed where the static characteristic put it, so the flow does not change, the density barely changes, and the cut size moves only through the weak pressure term, from 90.20 kPa to 71.86 kPa and the overflow from 81.44 % to 80.16 % passing. The cascade restores the pressure by raising the pump speed from 65.0 % to 73.4 %, which draws 118 m³/h more dilution water into the sump, thins the feed from 44.82 % to 42.46 % solids and takes the overflow to 85.49 % passing, an excursion three times the size of the one it prevented. The excursion is on the fine side and the outer loop brings the fineness back within 0.5 % of target 39 min after the wear, so no product is lost, but the disturbance the inner loop rejected was carried into the product by the correction itself.

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

4.5.2 Decomposition of the Outer Loop Gain

Table 4.5 reports the steady states before and after a pressure setpoint step from 90.20 kPa to 110.0 kPa, made at 5 min in scenario 6 with the size controller switched off, which is the step scenario 4 scripts at 4 min. The run reaches the tabulated values within 60 min, and they agree with the steady-state solution of the model to the digits shown. The decomposition uses (4.7) directly, since the pressure and density terms are separable in the logarithm, and Figure 4.3 draws the two paths it separates. Scenario 7 reaches the same move from the outer loop, stepping the fineness target from 81.4 % to 86 % under the cascade. After the hour its completion rule asks for, the pressure has risen to 107.3 kPa and the same decomposition assigns 36.6 % to pressure and 63.4 % to density.

Two consequences follow. The outer loop's plant gain is not the one an engineer reading (4.7) alone would compute, and is larger by a factor of 2.7, so a tuning derived from the correlation puts more gain around the loop than intended, by the same factor, and a loop more aggressive than its design. The level controller, which has no product variable anywhere in its configuration, is also a quality controller, and widening its proportional band on the grounds that it is only a level loop detunes the product size.

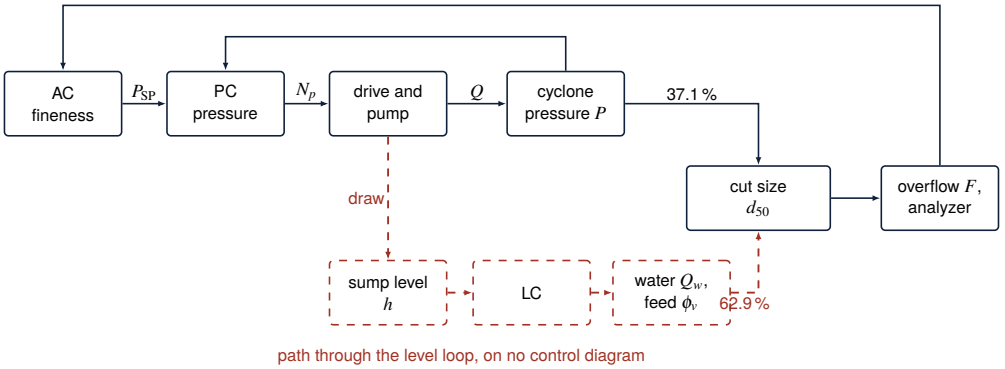


Figure 4.3 The two paths from the size controller AC to the cut size. The designed path runs through the pressure loop PC. The dashed path runs from the pump's draw on the sump through the level loop LC and the dilution water to the feed density. Notice that the dashed path carries the larger share of the steady-state change, 62.9 % in Table 4.5.

Table 4.6 Challenge Scenario Over 130 min With Three Level Tunings and the Suggested Cascade Tuning, Everything Else Unchanged

K_c ((m^3/h)/%)	τ_I (min)	level IAE (% min)	water travel (m^3/h)	fineness IAE (% min)
4	8	524.7	3465	138.9
8	4	168.7	6919	142.1
18	1.5	66.9	15682	149.4

✦ DID YOU KNOW? The Level Loop Moves the Cut

From the simulator: in Table 4.5 a pressure step from 90.20 to 110.0 kPa moves the cut size from 151.43 to 134.28 µm. Of that change, 62.9 % comes through the feed density that the level loop sets and only 37.1 % through the pressure itself.

The coupling also runs in the other direction. Scenario 3 raises the new feed from 250 t/h to 320 t/h at 3 min with the level and pressure loops closed and the size controller off, and returns it at 18 min. The level rises by at most 1.6 %, the controller removes water from 653.8 m^3/h to 603.8 m^3/h, the feed thickens from 44.82 % to 49.50 % solids, and by 18 min the overflow has fallen from 81.44 % to 70.71 % passing, within 0.03 % of its new steady state. A level loop holding its setpoint and a pressure loop holding 90.2 kPa have together moved the product by 10.7 % of fineness.

TRY IT IN THE APP **3 · Level Control on Dilution Water**

Load the scenario and press Run. The level and pressure loops are closed and the size loop is off. At 3 min the new feed rises from 250 to 320 t/h. Watch the level barely move while the dilution water falls, then watch the fineness slide from 81.44 to 70.71 % passing by 18 min, with the pressure on its setpoint throughout.

4.5.3 Effect of Level Loop Tuning

Because the level controller is coupled to the product, a natural question is whether tightening it improves the product. Table 4.6 answers that question on the challenge scenario, and the improvement does not appear.

Level control improves by a factor of 7.8 and costs a factor of 4.5 in valve travel, while the fineness error rises by 7.6 %. Only the tightest tuning keeps the level inside the challenge's ± 4 % band, with largest deviations of 10.8 %, 7.2 % and 3.4 % for the three rows. The challenge carries transmitter noise, and the rise in fineness error with the tight tuning holds on each of six noise seeds, by 3 % to 12 %. The tighter loop passes level noise and every level deviation into the dilution water, and the feed density it delivers travels three times as far over the run. The product error in this scenario is set by the two feed rate steps, which the cascade cannot see, and no amount of level tuning addresses them. The coupling established in Table 4.5 is a steady-state authority rather than a disturbance rejection path, and the distinction between the two is what this measurement establishes.

4.5.4 Unmeasured Constraints: Roping and Circulating Load

Two limits bind in this circuit and neither is measured on most plants. Scenario 9 closes the apex from 100 mm to 68 mm at 5 min and raises the new feed from 250 t/h to 420 t/h at 40 min. The closed apex alone takes the underflow from 67.3 % to 73.4 % solids, with σ in (4.10) at 0.83. The feed step takes it to 77.9 % against the 80 % at which σ reaches zero, the roping alarm ($\sigma < 0.35$) sounds 0.77 min after the step, and the circuit settles with the underflow at 77.0 % and $\sigma = 0.38$. The coarse and fine recoveries converge from 0.918 and 0.313 to 0.725 and 0.455, the partition curve flattens, and the overflow drops to 53.9 % passing, with 74.4 min of the 80 min run off specification. The outer loop drives the pressure setpoint to its 170 kPa limit by 42.8 min, and the pressure loop holds it with an integrated error of 7.40 kPa min over the run. Scenario 8 lowers the grinding effectiveness to 0.55 and asks for 92 % fineness. The recycle gain rises, the circulating load climbs from 288 % to 386 % and the cyclone feed from 971 t/h to 1215 t/h, the outer loop holds its output at the 170 kPa limit for 79.7 % of the 90 min run, and the fineness settles at 76.5 %.

TRY IT IN THE APP **9 · A Roping Apex Stops Classifying**

Load the scenario and run it for 80 min. The apex closes to 68 mm at 5 min and the new feed rises to 420 t/h at 40 min. Watch the underflow density climb, the roping alarm sound and the partition curve flatten while the pressure and level loops stay on their setpoints. The overflow ends at 53.9 % passing.

In both cases the level and pressure loops hold their setpoints while the size controller sits at its limit. The variables that failed, underflow density and circulating load, are computed rather than measured, and the operator's cue for the first of them is the shape of the underflow discharge.

Pause. Try Exercise 4.1 before reading on.

Figure 4.4 shows the circuit at the end of the roping scenario.

4.5.5 Graded Challenge Scenario

Scenario 11 applies four disturbances over two hours: new feed to 305 t/h at 8 min, impeller wear to 0.88 at 30 min, coarse feed to 0.91 at 55 min and feed back to 255 t/h at 85 min, with noise of 0.5 %, 1.2 kPa and 0.5 % (one standard deviation) on the level transmitter, the pressure transmitter and the size analyzer. It opens with the fineness controller off and the cautious tunings of scenario 1. The case

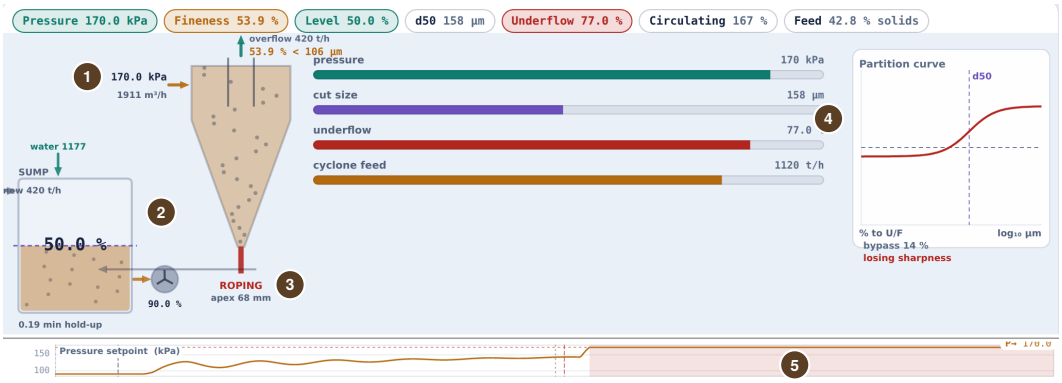


Figure 4.4 Scenario 9 at $t = 80$ min (1.33 h on the application clock), cropped to the plant view and the pressure setpoint strip, 40 min after the new feed was raised from 250 t/h to 420 t/h with the apex already closed to 68 mm. The pressure (1) and level (2) loops are holding 170.0 kPa and 50.0 %, while the underflow is at 77.0 % solids with the roping alarm on (3), the partition curve has flattened (4) and the overflow has fallen to 53.9 % passing 106 µm. The size controller's output is at its 170 kPa limit (5) and the product is off specification.

is passed when, over the full 130 min, the feed pressure stays within ±4 kPa of its setpoint outside the first minute after each upset, the level stays within ±4 %, the fineness returns within ±1.5 % of target after each of the four upsets, and the underflow never exceeds 76 % solids. Left as shipped it fails, with pressure deviations to 13.3 kPa, level deviations to 13.6 %, the fineness back after none of the upsets, 964.9 % min of fineness error and 121.9 min off specification. Commissioned from the fastest loop outward with the cascade in service and the tunings the application suggests, level $K_c = 18\,(\mathrm{m}^3/\mathrm{h})/\%$ with $\tau_I = 1.5$ min, pressure $K_c = 0.8\,\%/\mathrm{kPa}$ with $\tau_I = 0.15$ min and fineness $K_c = 2.0\,\mathrm{kPa}/\%$ with $\tau_I = 4$ min, the same run passes with 3.1 kPa, 3.4 %, four recoveries out of four and 69.5 % solids, and gives 149.4 % min and 14.9 min, a factor of 6.5 in integrated error. The verdict is the same on six noise seeds, with 149.4 % min to 164.9 % min for the suggested tunings. The scenarios' fineness tuning of $K_c = 1.2\,\mathrm{kPa}/\%$ with $\tau_I = 6$ min gives 284.0 % min and fails, because the fineness is still 2.7 % below target when the impeller wears at 30 min. Of the remaining 149.4 % min, 108.6 % min follows the two feed rate steps, which reach the product through the feed density and which no architecture in this application can detect before the analyzer reports them.

4.6 Technology in Practice

Figure 4.5 draws the three loops of Section 4.4 with their hardware.

What Measures It

The model reduces the pressure transmitter to a 0.02 min filter and the level transmitter to a 0.03 min filter, both with optional noise, and those filters stand for the sensing element and its seal.

> **TECHNOLOGY** **Sensor: A Pressure Transmitter With a Remote Seal**
>
> In most modern pressure transmitters a flexible metal isolating diaphragm passes the process pressure through a fill fluid to a silicon diaphragm. Its strain gauges unbalance a bridge circuit, and the voltage becomes a 4 to 20 mA or fieldbus signal. Where sludge would clog an impulse line, the diaphragm is set at the pipe as a remote seal joined to the transmitter by a filled capillary. The fill must stay free of gas, since bubbles absorb the diaphragm's motion and ruin the accuracy, and a height difference between seal and transmitter adds a constant offset to the reading [19].

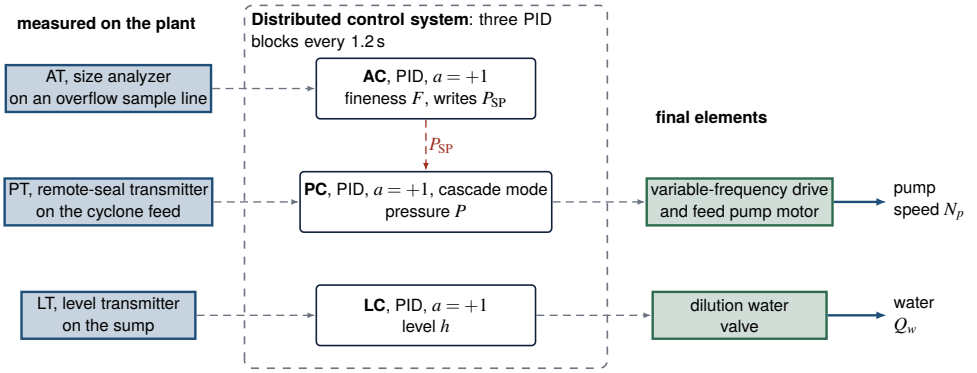


Figure 4.5 The three loops of the simulated circuit drawn with their hardware. The level transmitter LT, the pressure transmitter PT and the size analyzer AT feed three PID blocks, all reverse acting with $a = +1$. LC moves the dilution water valve, PC moves the variable-frequency drive of the feed pump, and AC writes the setpoint of PC, which runs in cascade mode. Notice that no block on the figure reads the feed density, which carries most of AC's authority.

The product measurement is slower. Particle size is measured on stream by instruments with cycle times of minutes [13], which the model represents as a 2.0 min dead time followed by a 0.4 min average. The feed density, the underflow density and the circulating load have no instrument at all.

What Moves It

The model's pump drive is a first-order lag of 0.10 min over a span of 25 to 95 % with no rate limit, the dilution water valve a lag of 0.15 min, and the apex actuator a travel time of 1.2 min.

> **TECHNOLOGY Final Element: A Variable-Frequency Drive on the Feed Pump**
>
> The cyclone feed pump is centrifugal, the most common pump design in the process industries, and its discharge pressure rises as its flow falls. Its speed is set by a variable-frequency drive, which rectifies the line power to direct current, smooths it, and inverts it back to alternating current at the frequency the commanded speed needs. A fixed-speed motor then delivers useful power over a wide range of speeds, and the drive can ramp the speed to spare the mechanism [19]. A worn impeller shows no fault at the drive: the speed is right and the flow is short.

The delivery multiplier w_p of Equation (4.5) stands for that short flow at the right speed.

Where the Controller Runs

Control loops in a plant are often nested or cascaded, the setpoint of one loop being based on the variable measured by another, and supervisory and lower-level loops run with cycle times from milliseconds to minutes. A distributed control system uses a supervisory loop to mediate a group of localized controllers [20], and programmable logic controllers run PID loops as often as they switch discrete devices [19]. The app integrates every 0.005 min and runs its three PID blocks on a 0.02 min (1.2 s) scan. On such a platform the chapter's structure is three velocity-form PID blocks with back-calculation, the output of the fineness block wired to the remote setpoint of the pressure block.

IN PRACTICE **Commissioning a Cascade From the Inside Out**

The secondary controller of a cascade is tuned first. It does not depend on the primary, while the primary can exert no control at all if the secondary sits in manual. The secondary then has three modes: manual, automatic with a local setpoint, and cascade, in which the primary writes its setpoint. A secondary left in automatic defeats the cascade as surely as a controller left in manual [19]. In this circuit the pressure loop is that secondary, and the level loop comes before both.

What Is Modernizing It

The first development is seeing inside the cyclone. Its internal flow is difficult or impossible to observe directly, because of the reinforced walls and the opaque slurry, and electrical impedance tomography estimates the conductivity field inside from electrodes outside. It collects several hundred frames a second, enough for real-time monitoring, and Bayesian methods now carry the numerical error of the reconstruction into its uncertainty [21]. The second is control above the cascade. Model-based supervision of the sump water and the cyclone feed rate [15] and feedforward from the measured feed size distribution [16], both met in Section 4.2, reach disturbances that no pressure transmitter sees. The third is a classification model that accounts for density. A Department of Energy project on iron ore circuits set out to model the mechanisms by which hydrocyclones retain fine, dense particles until they are overground, and to improve existing models to account for that behavior [11].

4.7 Limitations

The two size classes are the largest simplification. A real classifier is judged on a distribution, and the specification used here is a scalar computed from the split of two representative sizes. That is adequate for the control questions the application asks and inadequate for any statement about product distribution, because it cannot represent the shape change that a degraded partition curve actually produces.

OPEN CHALLENGE **Grains That Differ in Density**

A hydrocyclone sorts by density as well as by size. In an iron ore circuit, a cut set so that light quartz leaves only below its liberation size sends the dense magnetite to an even finer size, and the cyclone returns it to the mill to be overground. Earlier models could not handle such mixed-density feeds, nor the fish-hook in the measured efficiency curve, where some finer particles are more likely to report to the underflow than coarser ones [22]. This chapter has one solids density. Progress would be a classification model simple enough for control that keeps a partition curve for each mineral.

The cut-size correlation is reduced to three variables and its constants are fitted rather than measured. Vortex finder diameter, inlet geometry and cone angle are absorbed into two constants, and published comparisons of these correlations show that fitted coefficients differ between installations by more than the ranges explored here [6]. Nothing in this application supports a numerical claim about a particular cyclone. Mechanistic models built from the swirling flow field give a physical account of the same behavior and are correspondingly harder to identify [7].

The mill is a single grinding efficiency and a transport lag, with no power draw, no charge and no size distribution. That is sufficient to make the circulating load a genuine recycle with the right time constant and insufficient to say anything about grinding. Slurry rheology is absent entirely, so the whole effect of a dense feed on classification is carried by the exponential term of (4.7). Roping is smooth where a real apex is abrupt and hysteretic. The particle size analyzer is a continuous dead time equal to its nominal cycle time followed by a 0.4 min average, with optional noise. Unlike the sampled assays of Section 1.4.5, it does not hold its reading between cycles, and it has none of the calibration drift, sample line plugging or distribution-dependent bias that make on-stream sizing difficult in practice.

4.8 Discussion

The instruction to commission a control system from the fastest loop outward is usually given as a convention. This circuit supplies a reason. Until the level loop is holding, the cyclone feed density wanders, so through (4.6) the pressure loop's plant gain wanders, so the fineness loop's plant is not stationary. Tuning any loop against a moving plant produces a tuning valid at one condition, and the order in which the loops are closed is what makes each plant stationary in turn.

The same reasoning identifies what a cascade can and cannot be expected to do, and Table 4.4 shows that the usual test is incomplete. The question is not only whether a disturbance appears in the inner measurement, but also whether the inner loop's correction moves a variable that the outer loop cares about. Worn liners satisfy the first condition and fail the second. On this circuit that is a benign failure, because the excursion is on the fine side, but the general form of it is not benign. Any cascade whose inner loop manipulates a flow into a vessel under level control is coupled to the composition of that vessel, and the coupling is invisible in both loops.

Feedforward is the remaining option for the disturbances that no cascade reaches. Feed size distribution is measurable, and feedforward from it has been demonstrated for grinding circuits [16]. The application does not implement it, and the coarse feed row of Table 4.4, 53.1 % min of integrated error that the cascade does not remove, is the available improvement.

4.9 Conclusion

A grinding circuit classifier holds two time scales a factor of 41 apart in one manifold, and the resulting cascade is worth a factor of 49 in integrated product error against the disturbance it was designed for and nothing at all against three of the six disturbances tested. The third group in between is the result this chapter is built around, where an inner loop rejects a disturbance in its own measurement and carries it into the product through the level controller, and where the outer loop's steady-state authority is 63 % density and 37 % pressure. Both effects run through a loop that appears on no control diagram as a quality loop, and both are measurable in the browser in under two minutes. The commissioning order that follows, level first and outward, is then a statement about which plant gains are stationary when rather than a convention. On a plant with a genuine recycle, the application lets a reader measure cascade structure, loop interaction and commissioning order.

KEY POINTS

- The product size depends more on how thick the cyclone feed is than on the feed pressure, and the level loop sets that thickness.
- The cascade divides the integrated fineness error by 49 for pump wear and by 1.01 for feed coarseness.
- Worn liners and an extra cyclone appear in the pressure, yet the cascade makes them worse by a factor of three in integrated error.
- A tighter level loop holds the level 7.8 times better and the product no better.
- Commissioning from the level loop outward makes each loop's plant stationary in turn.

Where else this applies. The same disturbance sort on a level loop in Volume 3, Chapter 2, an inner loop's speed in Volume 3, Chapter 3, a recycle that amplifies in Volume 6, Chapter 7.

References

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

[2] F. G. Shinskey, *Process Control Systems: Application, Design, and Tuning*, 4th ed. New York: McGraw-Hill, 1996.

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

[4] B. A. Wills and J. A. Finch, *Wills' Mineral Processing Technology: An Introduction to the Practical Aspects of Ore Treatment and Mineral Recovery*, 8th ed. Oxford: Butterworth-Heinemann, 2015.

[5] T. J. Napier-Munn, S. Morrell, R. D. Morrison, and T. Kojovic, *Mineral Comminution Circuits: Their Operation and Optimisation*. Brisbane: Julius Kruttschnitt Mineral Research Centre, University of Queensland, 1996.

[6] K. Nageswararao, D. M. Wiseman, and T. J. Napier-Munn, "Two empirical hydrocyclone models revisited," *Minerals Engineering*, vol. 17, no. 5, pp. 671–687, 2004.

[7] K. T. Hsieh and R. K. Rajamani, "Mathematical model of the hydrocyclone based on physics of fluid flow," *AIChE Journal*, vol. 37, no. 5, pp. 735–746, 1991.

[8] E. Bretney, "Water-purifier," U.S. Patent 453,105, 1891, application filed November 3, 1890; patented May 26, 1891.

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

[10] J. Jeswiet and A. Szekeres, "Energy consumption in mining comminution," *Procedia CIRP*, vol. 48, pp. 140–145, 2016.

[11] S. K. Kawatra, T. C. Eisele, T. Weldum, D. Larsen, R. Mariani, and J. Pletka, "Optimization of comminution circuit throughput and product size distribution by simulation and control," Michigan Technological University, Quarterly Technical Progress Report, U.S. DOE award DE-FC26-01NT41062, 2005.

[12] D. Wei and I. K. Craig, "Grinding mill circuits—a survey of control and economic concerns," *International Journal of Mineral Processing*, vol. 90, no. 1–4, pp. 56–66, 2009.

[13] D. Hodouin, "Methods for automatic control, observation, and optimization in mineral processing plants," *Journal of Process Control*, vol. 21, no. 2, pp. 211–225, 2011.

[14] J. D. le Roux, I. K. Craig, D. G. Hulbert, and A. L. Hinde, "Analysis and validation of a run-of-mine ore grinding mill circuit model for process control," *Minerals Engineering*, vol. 43–44, pp. 121–134, 2013.

[15] V. R. Radhakrishnan, "Model based supervisory control of a ball mill grinding circuit," *Journal of Process Control*, vol. 9, no. 3, pp. 195–211, 1999.

[16] M. G. Maritz, J. D. le Roux, and I. K. Craig, "Feed size distribution feedforward control for a grinding mill circuit," *IFAC-PapersOnLine*, vol. 52, no. 14, pp. 7–12, 2019.

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

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

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

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

[21] C. J. Oates, J. Cockayne, R. G. Aykroyd, and M. Girolami, "Bayesian probabilistic numerical methods in time-dependent state estimation for industrial hydrocyclone equipment," *Journal of the American Statistical Association*, vol. 114, no. 528, pp. 1518–1531, 2019.

[22] H. J. Walqui, T. C. Eisele, and S. K. Kawatra, "Optimization of comminution circuit throughput and product size distribution by simulation and control," Michigan Technological University, Quarterly Technical Progress Report, U.S. DOE award DE-FC26-01NT41062, 2003.

4.10 Exercises

4.1 ○ Concept check · 1 min

In scenarios 8 and 9 the level and pressure loops hold their setpoints while the product fails. Which of these quantities does the circuit of Figure 4.1 measure with an instrument?

(a) The circulating load returned to the mill

(b) The underflow density at the apex

(c) The cyclone feed pressure

(d) The mill grinding effectiveness g_r

4.2 Concept check · 2 min

The cyclone liners wear, and the liner condition w_c falls from 1.00 to 0.80 at a fixed pump speed. Which statement agrees with Table 4.4?

(a) The cascade restores the pressure in seconds, so the fineness barely moves
(b) Neither architecture can see the wear, so both give the same integrated error
(c) The single loop is worse, because its pressure falls and nothing restores it
(d) The cascade is about three times worse, because restoring the pressure thins the feed

4.3 Concept check · 2 min

A coarser new feed, x_c from 0.85 to 0.91, gives 53.7 % min of integrated fineness error under the single loop and 53.1 % min under the cascade. Why does the cascade help so little?

(a) Ore coarseness does not appear in the pressure transmitter, so the inner loop never acts on it
(b) The analyzer's dead time is longer under the cascade, which cancels what the inner loop gains
(c) The size controller saturates at 170 kPa under both architectures
(d) The inner loop's correction thins the feed by as much as it restores the pressure

4.4 Theory and calculation · 6 min

Use the cut-size correlation, Equation (4.7), with $\beta = 0.225$ and $\gamma = 0.063/\%$.

(a) At a fixed apex and feed density, by what factor must the pressure rise for a cut 10 % finer? Give it to two decimals. *Checkpoint:* 1.60

(b) The pressure step of Table 4.5 takes the pressure from 90.20 to 110.0 kPa and the feed solids ϕ_v from 23.13 to 21.93 % by volume. Compute the pressure factor and the density factor to four decimals, and the share of the change in $\ln d_{50}$ that each carries, to 0.1 %. *Checkpoint:* 0.9563 and 0.9272; 37.1 % and 62.9 %

(c) An engineer computes the outer loop's plant gain from the pressure term alone and tunes the fineness controller for it. By what factor is the real gain larger, and is the tuned loop slower or faster than intended? *Checkpoint:* 2.7; faster

4.5 Theory and calculation · 6 min

The sump of Table 4.2 has a cross-section of 4.0 m^2 and a depth of 3.0 m. At the design point it runs at 50 % level and the pump draws 1380 m^3/h.

(a) Compute the slurry volume and the hold-up time, to 0.01 min. *Checkpoint:* 6.0 m^3; 0.26 min

(b) Compute the integrating gain of the level on the dilution water, in %/min per m^3/h, to three decimals. *Checkpoint:* 0.139

(c) Scenario 1 cuts the dilution water by 200 m^3/h in Manual. Ignoring the valve lag, estimate how long the level takes to fall from 50 to the 12 % suction alarm, to 0.01 min. *Checkpoint:* 1.37 min

4.6 Theory and calculation · 7 min

At the design point the new feed is 250 t/h at 85 % coarse, the cyclones send $E_c = 0.918$ of the coarse and $E_f = 0.313$ of the fine to the underflow, and the mill grinds $g_e = 0.32$ of the coarse it receives. Treat Equations (4.11) at steady state.

(a) Compute the coarse recirculation gain to four decimals and the coarse solids in the cyclone feed, to 0.1 t/h. *Checkpoint:* 0.6242; 565.5 t/h

(b) Compute the fine solids in the cyclone feed, to 0.1 t/h, and the circulating load, the underflow solids as a percentage of the new feed, to 1 %. *Checkpoint:* 296.4 t/h; 245 %

(c) With the grinding effectiveness at its lowest setting the gain settles at 0.80. By what factor does the coarse recycle now amplify the new coarse feed, and how many times the design amplification is that, to one decimal? *Checkpoint:* 5; 1.9

4.7 App challenge · 8 min

Open scenario 6, "The same disturbance, with the cascade", at the default seed.

(a) Run at 1 min per second and, when the clock reads between 5.5 and 7 min, press *Pump impeller wears (86 %)*. Switch to 20 min per second and run to 120 min. Read the fineness loop's IAE on the scorecard, within 0.1 % min, and the Off-spec time tile. *Checkpoint:* 3.2 % min; 0.1 min

(b) Reload the scenario with the reload button, set Size control to Single loop before Run, and repeat. Read the same two tiles, the IAE within 0.5 % min. *Checkpoint:* 154.9 % min; 33.3 min

(c) Compute the ratio of the two IAE values and explain it from where the disturbance enters. *Checkpoint:* about 49

4.8 App challenge · 9 min

Open scenario 7, "Where the outer loop's authority really comes from", as shipped at the default seed. The fineness target steps from 81.4 to 86 % at 5 min.

(a) Run at 20 min per second to at least 90 min (1.5 h), pause and export the CSV. The file's time column is in hours. From the last row before 5 min and the row nearest 1.000 h, read `press_kPa`, `feed_vol_pct` and `d50_um`, to 0.01 kPa, 0.001 % and 0.1 µm. *Checkpoint:* 90.15, 23.133, 151.5; 107.27, 22.059, 136.1

(b) Split the change in $\ln d_{50}$ into a pressure part and a density part with Equation (4.7), within 0.3 %, and check that the two factors reproduce the cut size at 1 h. *Checkpoint:* 36.6 % pressure, 63.4 % density

(c) Read `fine_pct` at 1.000 h and `press_kPa` at 1.500 h, within 0.1. Is the circuit at steady state after the hour, and how does the split compare with Table 4.5? *Checkpoint:* 85.5 %; 108.9 kPa

4.9 App challenge · 8 min

Open scenario 11, the challenge, at the default seed, and run each part for the full 130 min at 20 min per second.

(a) Run it as shipped. Read the challenge verdict and the fineness loop's IAE, within 1 % min. *Checkpoint:* fails: 13.3 kPa, 13.6 %, 0 of 4; 964.9 % min

(b) Reload, press Reveal suggested tuning and run again. Read the verdict, the IAE within 0.5 % min and the Off-spec time tile. *Checkpoint:* passes: 3.1 kPa, 3.4 %, 4 of 4, 69.5 %; 149.4 % min; 14.9 min

(c) Reload, reveal the suggested tuning, then set the fineness loop to $K_c = 1.20\,\mathrm{kPa/\%}$ and $\tau_I = 6\,\mathrm{min}$ before Run. Read the verdict and the IAE within 1 % min, and say which upset the fineness misses. *Checkpoint:* fails, 3 of 4; 284.0 % min; the feed step at 8 min

4.10 Application area · 4 min

Roping, Equation (4.10), sets the sharpness factor σ from the underflow density, with $\rho_{rope} = 80\,\%$ and $\rho_{soft} = 72\,\%$ solids.

(a) Compute σ at 73.4 and 77.0 % solids, to three decimals, and the underflow density above which the alarm at $\sigma < 0.35$ sounds, to 0.1 %. *Checkpoint:* 0.825, 0.375; 77.2 %

(b) On a real circuit nothing reads the underflow density. What does the operator watch, and name two remedies once the apex ropes. *Checkpoint:* the shape of the discharge; open the apex, cut the solids, open a cyclone

4.11 Application area · 4 min

The Introduction says the pressure is a weak handle on the cut and that pump power rises with the cube of speed. Hydraulic power is $P\,Q$.

(a) At a fixed feed density and eight cyclones, $P \propto Q^2$. By what factor must the flow and the hydraulic power rise for a cut 10 % finer by pressure alone? Give two decimals. *Checkpoint:* 1.26; 2.02

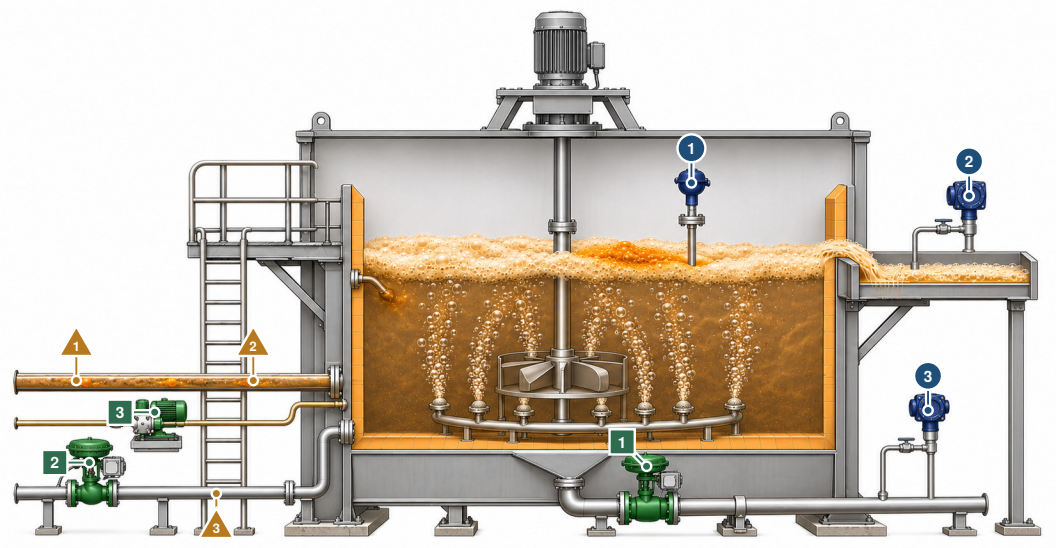
(b) In Table 4.5 the pump goes from 65.01 to 72.32 % with the pressure from 90.20 to 110.0 kPa, and Q is proportional to speed. Compute the rise in hydraulic power and the fall in cut size, both in percent, to one decimal. *Checkpoint:* 35.7 %; 11.3 %

4.12 Review the agent · 5 min

An AI assistant reviewing the challenge writes: "Table 4.5 shows that 63 % of the outer loop's authority passes through the level controller, so the level loop is the real quality loop. Raising the level gain from 4 to 18 $(m^3/h)/\%$ and cutting its integral time from 8 to 1.5 min will therefore cut the challenge's fineness error substantially." Find the flaw and say how to check it on the app.

CHAPTER 5

A Browser-Based Froth Flotation Simulator for Economic Setpoint Selection and Loop Interaction



The controller moves 1 the tailings valve that drains the pulp 2 the air blown into the cell 3 the collector dose

The controller reads 1 the froth depth 2 the copper grade of the concentrate 3 the copper grade of the tailings

What disturbs it 1 leaner ore in the feed 2 more feed, or a thicker one 3 a blower that delivers less air

Notice that the air pipe and the tailings valve both move the froth depth, the air within seconds and the valve over minutes. Illustration, not to scale.

IN PLAIN TERMS **Skimming a Metal Off a Froth of Bubbles**

The main idea. A flotation cell blows air through a stirred mix of ground ore and water. Grains of the copper mineral stick to the bubbles and rise into a foam on top, which spills over the edge as the product.

The challenge. A deeper foam gives a richer product but lets more copper fall back, so the best depth is a choice about money, not a fixed rule. The air also changes the foam depth within seconds, far faster than the outlet valve that is meant to hold it.

Why it matters. Holding the wrong depth well can cost more than holding the right depth badly.

THREE WAYS TO READ THIS CHAPTER

Explore	8 min	The first page, the history box of Section 5.1, the fact box of Section 5.2, the try-it box of Section 5.6.1, the fact box of Section 5.6.6, and the key points
Study	50 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 5.6, Section 5.7, Section 5.8 and the Discussion with their open challenges

The control question A flotation cell is paid for its margin, and no froth depth improves both the concentrate grade and the recovery. Where should the depth sit, and what do the three loops that hold the cell do to each other when the air reaches the froth depth seventeen times faster than the valve assigned to control it?

Learning objectives

- Compute the margin of an operating point from its grades with the two-product formula, and explain why the best froth depth gives up recovery.
- Calculate the gas holdup, the interface displacement and the bubble surface area flux that an air step produces.
- Explain from measured response times why a level loop cannot keep an air move out of the froth depth, and recognize an over-tuned grade loop by the swings it spreads to its neighbors.
- Judge a configuration of the graded scenario by its setpoints and its tuning, and price each decision.
- Separate an economic MPC's gain into the operating point, the dropped tailings setpoint and the online optimization.

Prerequisites Sections 1.4.3 (grade, recovery, the economic optimum and the two-product formula), 1.4.5 (late assays), 1.6.2 (velocity form and loop action) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/froth-flotation/

See also Volume 9, Chapter 8 for a setpoint chosen on cost and an MPC gain that does not come from the horizon, and Volume 9, Chapter 10 for where a supervisory setpoint belongs. Volume 6, Chapter 2 for loops that interact although the relative gain says they should not, and Volume 6, Chapter 3 for pairing two quality loops on one product.

5.1 Introduction

A flotation cell separates a valuable mineral from gangue by rendering the valuable particles hydrophobic with a collector reagent, dispersing air through the slurry, and overflowing the loaded froth into a concentrate launder [1]. Two numbers describe the outcome. Grade is the mass fraction of valuable mineral in the concentrate, and recovery is the fraction of the valuable mineral in the feed that reports to the concentrate. No operating point improves both, because every action that carries more valuable mineral over the lip also carries more gangue, so the concentrate stream grows and its composition moves toward the feed composition [2]. Comminution and concentration account for a large share of the energy consumed by a mine [3], so the position chosen on the grade-recovery curve carries a direct economic consequence.

The trade-off makes flotation a useful teaching process, because regulation and optimization appear in the same experiment and cannot be separated. The cell exercises the material and component balances behind a self-regulating level, time-scale separation between actuators, measurement dead time from composition analyzers, actuator saturation with anti-windup, and the difference between holding a setpoint and choosing one [4–6]. As many have noted in assessments of process control education, there is a need to give students exercises in which the control objective is economic rather than purely regulatory [7, 8]. However, most course simulations supply a setpoint and grade the tracking error, which omits the setpoint decision that dominates flotation economics.

This chapter works with a simulated flotation cell in which the economic objective and the regulatory loops share one experiment. The cell's margin-optimal froth depth lies near 0.55 m against a 0.30 m design value, and its air rate reaches the froth depth seventeen times faster than the tailings valve assigned to control it. The chapter also measures an optional economic model predictive controller (MPC) on the air and the collector against the three loops, and separates the share of its gain that a fixed setting from the same model already captures. Section 5.2 reviews the grade-recovery trade-off and flotation control practice. Section 5.3 sets out the model, Section 5.4 measures how fast each actuator reaches each variable, and Section 5.5 gives the three loops and the MPC. Section 5.6 reports the scenario results, from an open-loop air step through the limit cycle of an aggressive grade

Table 5.1 Steady-state economics against froth depth with the level controller holding each depth. Net value is the concentrate value less reagent and blower costs.

h_f (m)	R_f	Grade (%)	Recovery (%)	Valve (%)	Net ($/h)
0.20	0.698	4.18	84.51	50.4	6882
0.30	0.583	5.48	80.54	54.8	8963
0.40	0.487	6.94	76.29	58.1	10028
0.50	0.407	8.50	71.71	60.7	10428
0.55	0.372	9.31	69.30	61.8	**10447**
0.60	0.340	10.12	66.81	62.9	10372
0.70	0.284	11.72	61.60	64.8	9998
0.80	0.237	13.24	56.16	66.7	9400

loop to the graded scenario and the MPC comparison. Section 5.7 describes the real instruments, final elements and control platforms of such a cell. Sections 5.8 to 5.10 give the limitations, what generalizes past this cell and the conclusion, and the Exercises follow.

> ⧖ HISTORY **A Froth Lifted by Air (1905)**
>
> In 1916 the Supreme Court of the United States upheld a patent issued in 1906 to three London metallurgists, Henry Livingston Sulman, Hugh Fitzalis Kirkpatrick-Picard and John Ballot. Earlier oil processes used so much oil that they were too costly for lean ores. By March 1905 the three had cut the oil to a fraction of one percent of the ore and beat air into the pulp by agitation. A coherent froth rose to the surface carrying between 70 and 80 percent of the mineral, and the court found that the lifting force came chiefly from the air bubbles, not from the oil. By 1916 the process was in extensive use in the United States, Australia, Sweden, Chile and Cuba [9]. The froth that this chapter's cell overflows is that froth.

5.2 Background

5.2.1 The Grade-Recovery Trade-off

What reaches the launder is the product of what the pulp collects and what the froth passes. The froth phase returns particles to the pulp by drainage and coalescence, and both loss mechanisms grow with froth depth, while water drains faster than an attached particle settles. A deeper froth therefore loses some valuable mineral and much more water, so the concentrate is smaller and richer. The grade-recovery curve traced by the froth depth is the fundamental performance object in mineral processing [1, 2]. Section 1.4.3 derives the same trade for a bank of cells and places the economic optimum where the last tonne of concentrate pulled just pays for itself, Equation (1.4). This chapter's handle is the froth depth of one mechanical cell.

Table 5.1 lists the steady-state economics of the simulated cell as the froth depth setpoint is moved. Deepening the froth from the 0.30 m design value to 0.55 m raises the net value by $1484 /h (17%) and costs 11.2 points of recovery, which is the figure reported in plant metal accounting. A finer sweep at the same air rate puts the maximum at 0.53 m and $10 452 /h, so the 0.55 m row is within $5 /h of it. The optimum is also shallow. The margin is within 1% of its maximum anywhere between 0.47 and 0.61 m, so the optimum cannot be located from a single shift of routine operating data, and the operating depth is in practice a standing decision rather than a calculation, the same flatness that Figure 1.5b shows for the bank of Chapter 1. Letting the air move as well, the application's optimum search (the *Go to the model optimum* helper), on a grid of 0.05 m and 2 m^3/min, finds 0.60 m at 18 m^3/min (1.50 cm/s), worth $10 617 /h at 71.0 % recovery and 9.11 % copper.

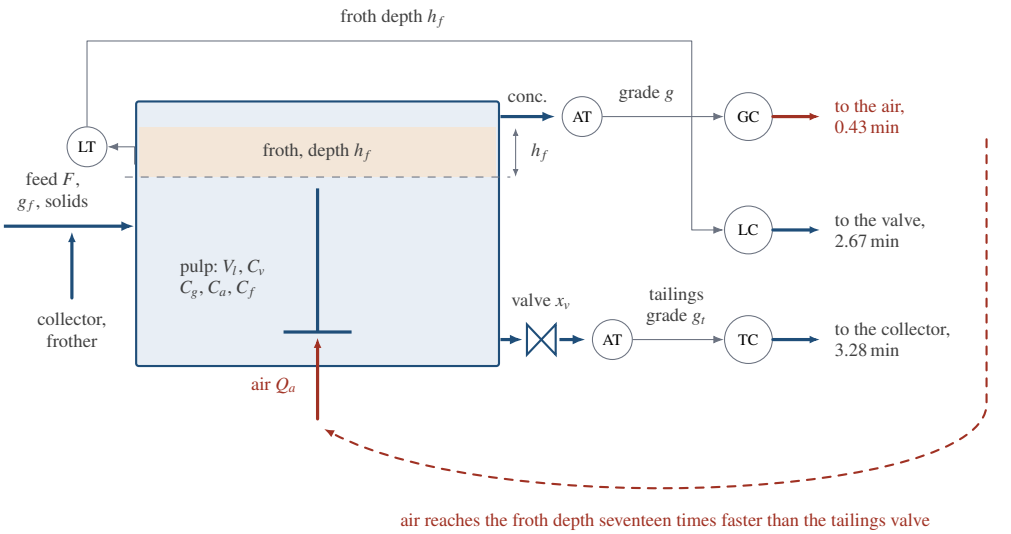


Figure 5.1 The simulated cell. The level controller (LC) holds froth depth on the tailings valve, the grade controller (GC) holds concentrate grade on the air rate, and the tailings-grade controller (TC) holds the tailings assay on the collector dose. Each output arrow carries the measured open-loop time to sixty-three percent for its own variable from Table 5.2. The dashed path, drawn red, marks the interaction studied in Section 5.6, in which the air moves the froth depth in 0.16 min while the valve needs 2.67 min.

5.2.2 Flotation Control Practice

Industrial flotation control is organized in layers, with base-level loops stabilizing pulp levels, air rates, and reagent flows, and a supervisory layer moving their setpoints toward an economic objective [10, 11]. The pairing implemented here, froth depth on the tailings valve, concentrate grade on the air rate, and tailings grade on the collector dose, is the conventional one [1, 11]. Multivariable and predictive control have a long but only partially adopted record in flotation, because the process is nonlinear, the assays arrive with dead time, and the objective is a position on a curve rather than a fixed setpoint [10, 12]. The simulated cell reproduces these features at a scale a single experiment can hold.

> ✦ **DID YOU KNOW? Two Billion Tonnes a Year**
>
> A 2014 thesis on scaling up flotation equipment cites a 2005 report that nearly 2 billion tonnes of ore a year were treated by flotation worldwide. Its compilation of the largest mechanical cells on the market runs from 3.0 m^3 in 1969 to 660 m^3 in 2012 [13].

5.3 Process Model

Figure 5.1 shows the cell, its three loops, and the two composition analyzers. The model integrates seven states, the pulp liquid volume, the valuable and gangue solids concentrations, the collector and frother concentrations, and the two actuator lags. The engine that integrates them and executes the loops is the framework of Section 1.6.1.

5.3.1 Design Point

The model time unit is minutes. The cell takes 12 m^3/min of feed at 0.80 % copper and 33 % solids. At design the tailings valve sits at 54.8 %, the air at 14.4 m^3/min, and the collector at 30 g/t, which gives a froth depth of 0.300 m, a superficial gas velocity of 1.20 cm/s, a gas holdup of 12.0 %, a

Sauter mean bubble diameter of 1.496 mm, and a bubble surface area flux of 48.1/s. The concentrate is 1.20 m^3/min at 5.48 % copper, the tailings assay is 0.176 %, the recovery is 80.5 %, and the margin is \$8963 /h. Section 1.4.3 checks this recovery against the two-product formula, Equation (1.5).

5.3.2 Cell Geometry and Gas Holdup

The liquid inventory obeys a total volume balance, with the tailings flow set by the valve and the head, and the concentrate flow set by the water the froth carries over the lip. The gas holdup is essentially algebraic in the air rate,

$$\varepsilon_g = k_\varepsilon \frac{Q_a}{A}, \qquad h_f = H_{\text{lip}} - \frac{V_l}{A\,(1-\varepsilon_g)} \tag{5.1}$$

where k_ε (min/m) is the holdup per unit superficial gas velocity, A is the cell cross section, V_l is the liquid volume, and H_{lip} is the lip height. A change in air therefore displaces the pulp interface within seconds by expanding the same liquid over a taller column, before any change in flotation has occurred.

Concentrate flow is a weir relation on the amount by which the interface exceeds the lip, which makes flooding a physical event rather than an alarm condition. When the level controller loses the cell, the froth depth reaches zero and the launder receives whole pulp instead of concentrate, so the grade collapses toward the feed grade. The graded scenario of Section 5.6.6 fails any run in which the froth depth falls to the 0.06 m collapse depth.

5.3.3 Collection Rate From Bubble Surface Area Flux

Flotation rate is proportional to the rate at which bubble surface passes through the pulp, which is the standard reduction of the two gas variables to one [14–16],

$$S_b = \frac{6\,j_g}{d_{32}}, \qquad j_g = \frac{Q_a}{A} \tag{5.2}$$

where j_g is the superficial gas velocity and d_{32} is the Sauter mean bubble diameter set by the frother concentration. The design value of 48.1/s sits within the range measured on industrial cells [14]. The valuable and gangue rate constants scale with S_b, and the valuable rate scales additionally with the fraction of particle surface the collector has covered,

$$\theta_c = \frac{C_a}{C_a + K_a} \tag{5.3}$$

which is a Langmuir isotherm and therefore saturates. Gangue is recovered without the collector, so collector added past the point where (5.3) has flattened adds gangue and little valuable mineral. Section 5.6.4 measures the cost of that overdose.

5.3.4 Froth Recovery

The froth phase returns particles to the pulp through drainage and coalescence, and both losses grow exponentially with depth,

$$R_f = \mathrm{e}^{-\lambda_p h_f}, \qquad R_w = \mathrm{e}^{-\lambda_w h_f} \tag{5.4}$$

with $\lambda_p = 1.8/\text{m}$ for particles and $\lambda_w = 4.5/\text{m}$ for water, because water drains faster than an attached particle settles. That single inequality produces the grade-recovery trade-off, since a deeper froth loses some valuable mineral and much more water. Table 5.1 is (5.4) evaluated against a copper price. Above a critical superficial gas velocity of 1.4 cm/s a lumped destabilization term reduces the froth recovery and raises entrainment, which bounds the benefit of additional air.

Pause. Try Exercise 5.2 before reading on.

Table 5.2 Measured open-loop time to sixty-three percent of the largest excursion after a unit step from the design point. The air and the valve reach the froth depth seventeen times apart.

Actuator	Variable	τ_{63} (min)
Air rate	froth depth h_f	0.16
Air rate	concentrate grade	0.43
Tailings valve	froth depth h_f	2.67
Tailings valve	concentrate grade	3.37
Collector dose	concentrate grade	4.74
Collector dose	tailings grade	3.28

5.3.5 Measurements, Initialization, and Assumptions

The froth depth measurement carries a 0.08 min transmitter filter. The concentrate and tailings analyzers are modeled with 2.0 and 3.5 min of dead time and first-order lags of 0.5 and 0.8 min, which is what limits the two quality loops. They read continuously, unlike the sampled assays of Section 1.4.5, so the age of a reading is a fixed dead time rather than a range. Both analyzers and the transmitter are sampled at the controller scan, and the dead times are applied exactly because they are whole multiples of it. Initialization solves the seven-state steady state directly, so the application starts at an exact equilibrium for any actuator setting, and a manual-mode helper, *Hold this froth depth*, returns the valve position that holds the froth depth setpoint at the present air rate.

🕮 **Pause.** Try Exercise 5.1 before reading on.

The pulp is one perfectly mixed zone with two lumped mineral classes and no particle size distribution, and the froth is a recovery factor rather than a transported phase. Integration is fixed-step fourth-order Runge-Kutta at 0.005 min with a 3 s controller scan, on the engine of Section 1.6.1.

5.4 Dynamic Analysis

The cell has three actuators. The tailings valve sets how much slurry leaves the bottom, the air rate sets how much gas is dispersed through the pulp, and the collector dose sets how hydrophobic the valuable mineral is.

Two of these actuators move the froth depth at very different speeds. The valve changes the liquid inventory, a volume that settles with a time constant of minutes because the tailings flow rises with the head. Air acts through the algebraic holdup of (5.1) and displaces the interface within seconds. The two effects of air on froth depth also run in opposite directions. The immediate holdup effect raises the interface and thins the froth, while at a fixed valve position the added air carries more water over the lip, drains the liquid inventory, and leaves the steady-state froth deeper.

Table 5.2 lists the measured open-loop response times and is the source of the interaction measured in Section 5.6. Each entry steps one actuator by one unit (1 % of valve, 1 m^3/min of air, 1 g/t of collector) in manual mode and times the true variable from the scan at which the move reached the plant. For the air the largest excursion is the fast one, a 0.018 m thinning of the froth at 0.53 min and a fall of 0.33 points of grade at 1.43 min, both of which partly reverse as the drained inventory deepens the froth again. The steady-state gain matrix of this cell is unremarkable, and a relative gain array computed from it would not flag the pairing [4]. Volume 6, Chapter 2 meets the same warning from another plant, a relative gain at the textbook value for no interaction on two loops that interact anyway. What matters is that one actuator reaches a shared variable seventeen times faster (0.16 versus 2.67 min) than the actuator assigned to control it. A grade controller acting on the air can therefore disturb the froth depth faster than the level controller can correct it, and the level controller has no recourse except to move faster on a valve that cannot. Section 5.6.2 measures the consequence for the level loop, and Section 5.6.3 what follows when the grade loop is tuned aggressively.

5.5 Controller Design

The three loops use the velocity-form PID of Section 1.6.2, Equations (1.10) to (1.13), executed every $\Delta t = 0.05\,\text{min}$ (3 s) with the derivative on the measurement. More air lowers the grade and more collector lowers the tailings assay, so both quality loops have negative process gains and run with action $a = -1$, direct acting in the convention of Section 1.6.2, as the application's hints also call them. The level loop runs with $a = +1$, reverse acting, because opening the tailings valve drains the pulp and deepens the froth. For the air the saturation function of Equation (1.12) applies a rate limit r after the clamp to the actuator range.

Anti-windup is back-calculation with the tracking time $T_t = \tau_I$. While the output is clipped or rate-limited, the last term of Equation (1.13) drives the integral toward the value that reproduces the applied output, so the integral alone is corrected and the running sum is never truncated. Each scan removes $\Delta t/\tau_I$ of the excess, $1/20$ for the aggressive grade loop of Section 5.6.3 and $1/80$ for the recommended one. In manual mode, or with a quality loop switched out, each controller sets its integral so that its output equals the manual station, which makes the transfer to automatic bumpless.

Froth depth is paired with the tailings valve, concentrate grade with the air rate, and tailings grade with the collector dose, tuned by the SIMC rules [17] applied to measured step responses. The three loops run as PI, and the level loop carries an optional derivative with a 0.06 min filter. The air command is limited to 3 to 40 m^3/min and rate-limited at $r = 7\,\text{m}^3/\text{min}$ per minute, so the damper travels its full range in about 5.3 min. The valve and the collector are limited to 0 to 100 % and 0 to 95 g/t with no rate limit. The air that reaches the cell is the command multiplied by the blower delivery, so a blower fault appears to the grade loop as a loss of gain rather than as a limit. The application allows the tailings-grade loop, or both quality loops, to be switched out, which is how a large fraction of industrial rougher banks run, and Section 5.6.5 measures why.

The level loop is fast by necessity rather than by choice. Its own plant is slow, at 2.67 min, but the disturbance it must reject arrives in 0.16 min from an actuator inside the same control system. The scenarios ship $K_c = 180\,\%/\text{m}$ and $\tau_I = 1.2\,\text{min}$, and the graded scenario suggests $K_c = 300\,\%/\text{m}$, both aggressive for a level loop, for the reasons measured in Section 5.6.6.

5.5.1 Economic Model Predictive Control

Section 1.6.3 describes the MPC option the framework apps share. This section gives what is specific to this cell. A fourth structure hands the two quality handles to an economic MPC, offered as an option and never the default of a scenario. It sets the air command Q_a and the collector dose D_c together, with no tailings setpoint, while the froth depth PI keeps the tailings valve. A decision is made every $T_s = 1\,\text{min}$ (20 controller scans) over $N = 60$ steps grouped into 12 move blocks of $[1,1,1,1,2,2,3,4,6,9,15,15]$ steps, so each quadratic program carries 24 decisions and one slack. The first move goes to the plant and stays in force until the following decision, and the first plan is made at $t = 0$. This is the receding-horizon structure of industrial linear MPC [18], with the plant value in place of a setpoint as the objective [19].

The controller's model is distinct from the simulated plant. Its static part is the steady state of the balances of Section 5.3 evaluated with the declared parameter values, the design feed (0.80 % copper, $12\,\text{m}^3/\text{min}$, 33 % solids), and the froth depth held at its setpoint h_f^{sp}. A held depth fixes the liquid volume through (5.1), so both grades follow in closed form as $y_{ss} = [G_{ss}(Q_a, D_c),\, T_{ss}(Q_a, D_c)]$. Its dynamic part is a set of input lags identified from steps of $\pm 2\,\text{m}^3/\text{min}$ of air and $\pm 5\,\text{g/t}$ of collector with the depth loop closed, placed ahead of the static map as shown in Equation (5.5).

$$\begin{aligned} \tau_A \dot{\xi}_A &= Q_a - \xi_A, \quad \tau_1 \dot{\xi}_1 = D_c - \xi_1, \quad \tau_2 \dot{\xi}_2 = \xi_1 - \xi_2 \\ \hat{y} &= y_{ss}(u_0) + b + K(u_0)\,[\xi_A - Q_{a,0},\ \xi_2 - D_{c,0}]^{\top} \end{aligned} \tag{5.5}$$

Here ξ_A, ξ_1 and ξ_2 are lagged commands with $\tau_A = 1.2\,\text{min}$, $\tau_1 = 2.38\,\text{min}$ and $\tau_2 = 0.92\,\text{min}$, $u_0 = [Q_{a,0}, D_{c,0}]$ is the command applied at the decision, K is the Jacobian of y_{ss} there, b is the

estimated output bias, and $\hat{y} = [\hat{G}, \hat{T}]$ is the predicted true concentrate and tailings grade. The lags are discretized exactly at T_s. The first-order fit from air to concentrate grade leaves a root-mean-square residual of 6.2% of the response, and the second-order fit from collector to tailings grade leaves 0.07%. Collector to concentrate grade fits no single linear shape (21.6% residual), because its gain changes sign at 29.3 g/t, where the collector completes a monolayer and the selectivity loss behind (5.3) begins. The static map is therefore kept nonlinear and relinearized at every decision, a Hammerstein structure. At the design commands the gains are −0.243 and −0.0085 percent copper per m^3/min of air on the concentrate and tailings grades, and −0.0069 and −0.0013 per g/t of collector, equal to five digits to central differences of the full balance.

The objective is the net value of the cell. The concentrate and the copper in it follow from the two grades by the two-product formula at the design feed grade, as shown in Equation (5.6).

$$\begin{aligned} \Pi &= 60\Big[\dot{m}_c\,(g\,\phi\,p_{Cu} - c_c) - \tfrac{F_s}{10^6}\,(c_a D_c + c_f D_f) - c_q Q_a\Big] \\ \dot{m}_c &= F_s\,\frac{f-t}{g-t} \end{aligned} \tag{5.6}$$

Here Π is in \$/h, g, t and f are the concentrate, tailings and feed grades as mass fractions (Chapter 1 writes the concentrate assay c), and the second line is the two-product formula of Equation (1.5) written for the concentrate mass. In it $F_s = 4999\,\mathrm{kg/min}$ is the design solids feed, D_f is the frother dose, $\phi = 0.965$ is the payable fraction, $p_{Cu} = \$9.0/\mathrm{kg}$, the concentrate handling cost is $c_c = \$0.22/\mathrm{kg}$, the collector and frother costs are $c_a = 2.8$ and $c_f = \$3.5/\mathrm{kg}$, and the blower cost is $c_q = \$0.02$ per m^3/min per minute. The quadratic program minimizes Equation (5.7).

$$\begin{aligned} J = \sum_{k=1}^{N} \Big(\ell_y^{\top} \hat{y}_k + \ell_u^{\top} v_k + v_k^{\top} R\, v_k \Big) \\ + \sum_{j=1}^{12} \Delta v_j^{\top} S\, \Delta v_j + \rho\, s^2 + \sigma s \end{aligned} \tag{5.7}$$

Here J is in thousands of dollars per hour, $v_k = u_k - u_0$, ℓ_y and ℓ_u are -10^{-3} times the gradients of Π with respect to the grades and to the commands at fixed grades, and the input weight

$$R = 10^{-3}\,\mathrm{diag}\Big(\max\big(-\tfrac{1}{2}\Pi_{QQ}, 1\big),\ \max\big(-\tfrac{1}{2}\Pi_{DD}, 0.05\big) \Big)$$

uses the second derivatives Π_{QQ} and Π_{DD} of the steady-state value in the two commands. The move weight $S = \mathrm{diag}(0.005, 0.0005)$ penalizes the change Δv_j between successive blocks, the first measured from u_0. Because $\hat{y}_k$ approaches $y_{ss}(u_0) + b + K v_k$, the sum is close to N times the second-order expansion of $-10^{-3}\Pi$ about the present commands, so each plan heads for the point where the steady-state value stops rising. On the design ore that point is 8.34 m^3/min of air and 36.9 g/t of collector, 7.43 % copper, a 0.245 % tailings grade, 71.8 % recovery and \$9796 /h, against \$8963 /h at the design point. The horizon has no measurable effect, since 10 and 120 min both earn \$8213 /h in scenario 11.

The air and collector commands are held by hard bounds of 3 to 40 m^3/min and 0 to 95 g/t, the clamps of the shipped loops, and the air may move at most 7 m^3/min per decision, the grade loop's slew, which is also applied on every scan. The predicted grade carries a soft floor $\hat{G}_k \geq \underline{g}_k - s$, $s \geq 0$, with $\underline{g} = \max(G^{sp}, 4.5\,\%)$, prices $\rho = 100$ and $\sigma = 10$ in the units of J with s in percent copper, and one slack for the whole horizon. Where the predicted free response lies below the floor, the floor at step k closes on it as $1 - \mathrm{e}^{-kT_s/3\,\mathrm{min}}$, which keeps the slack at zero whenever the floor is reachable. An unreachable 12 % floor gives a slack of 1.57 % copper and no fallback.

Offset is removed with an output bias on each grade [20]. On every scan the model grades at the lagged commands are evaluated at the measured froth depth, delayed by the analyzer dead times of

2.0 and 3.5 min, and filtered by their lags of 0.5 and 0.8 min. The analyzer reading y^m then updates $b \leftarrow b + \lambda\,(y^m - \hat{y}^m - b)$ with $\lambda = 1 - \mathrm{e}^{-\Delta t/3\,\mathrm{min}}$, so each residual is aligned to the time its sample was taken, the alignment that Equation (1.8) recommends for a late assay. After a 2 m^3/min air step the largest grade bias is 0.023 % copper, against 0.113 in a copy that compares each reading with the current model grade. The estimator runs under every structure, and no preview is given. A plan is used only if it is optimal, finite and inside its bounds. Otherwise the grade and tailings PIs, which track every applied command, take over without a bump, the trend is marked, and the next valid plan resumes. A forced non-finite plan at 30 min hands over with steps of 8×10^{-6} m^3/min and 8×10^{-5} g/t, and no plan was rejected in any scenario. The plant runs on the actual feed, blower delivery and reagent strength, the analyzers add dead time, lag and noise, and the controller reads only the analyzers, the depth transmitter, the commands it sent, the setpoints and the 4.5 % limit. A copy fed only those readings reproduces all 60 decisions of a run on a plant with changed hidden parameters. The mean solve takes 0.11 to 0.30 ms in Node.js.

5.6 Results

The application carries twelve scenarios, from manual operation of the cell through mapping the grade-recovery curve and measuring the loop interaction to the economic MPC and the graded scenario. Each is run, credited and exported as Section 1.6.4 describes. Reset restarts a run from $t = 0$ with every setting the learner has changed, initialized at the steady state of the current manual stations, while reloading the scenario restores its own settings. A repeated open-loop step therefore needs the reload, because Reset would start the cell already settled at the stepped manual station.

5.6.1 Open-Loop Response to the Air Rate

Scenario 3 steps the air from 14.4 to 20 m^3/min in manual mode with the valve fixed. The added holdup thins the froth from 0.300 to 0.201 m within 0.48 min, and the thinned froth briefly discharges what it held, so the displayed recovery, which is the copper overflow rate divided by the copper feed rate, peaks at 112% at 0.35 min. The grade bottoms at 3.71 % at 1.1 min. The froth then deepens past its starting point, because the added air carries more water over the lip and drains the liquid inventory, and the cell settles at 0.362 m, 82.69 % recovery, 4.99 % grade, and $8422 /h. A steady-state sweep at the same valve position shows the whole range. Raising the air from 9 to 26 m^3/min takes the bubble surface area flux from 30.0 to 87.4/s and the grade from 6.08 to 4.10 %. Recovery rises from 75.74 % through a maximum of 82.7 % near 21 m^3/min, where the froth destabilization term overtakes the added bubble surface, and ends at 81.3 %. The froth depth moves from 0.225 to 0.457 m without any valve change, and the margin falls from 9150 to $6411 /h across the range.

One actuator therefore moves the recovery, the grade, and a variable that has its own controller. This coupling drives the closed-loop behavior measured in the remaining scenarios.

> **TRY IT IN THE APP** **3 · Step the Air, One Handle, Three Effects (MANUAL)**
>
> Load the scenario, press Run and let the trends settle, then raise the air flow from 14.4 to 20 m^3/min in one move. Watch the froth depth thin from 0.300 to 0.201 m within 0.48 min, then deepen past where it started and settle at 0.362 m, while the concentrate grade falls and settles at 4.99 %.

5.6.2 Level Loop Disturbance Rejection

Scenario 4 runs the level loop alone and steps the air from 14.4 to 22 m^3/min at 8 min. The level controller is in automatic and correctly tuned, at $K_c = 180\,\%/\mathrm{m}$, and the froth depth still departs by 0.111 m within 0.31 min of the step, which is more than twice the ±0.05 m band the graded scenario allows. Detuning the loop to 30 %/m widens the excursion only to 0.127 m, while the lowest grade falls to 3.21 % against 3.67 %.

The loop itself is not at fault. The disturbance arrives seventeen times faster than the actuator can answer, so the peak excursion is set by the ratio of the two response times rather than by the tuning.

Table 5.3 Scenario 8 over 240 min with a well-tuned grade loop and with an aggressive one. Recovery is the copper recovered over the copper fed for the whole run.

	Grade range (%)	Valve travel (%)	Recovery (%)	Mean net ($/h)
Tuned	4.84–5.52	8	76.3	7270
Aggressive	1.05–8.56	2257	76.3	3305

The available remedies are to constrain how fast the air is allowed to move or to feed the air command forward to the valve. The application implements the first as the fixed rate limit of Section 5.5, which acts on the grade loop's command rather than on a manual step such as this one, and omits the second deliberately, so that the limitation remains measurable.

5.6.3 Limit Cycle From an Aggressive Grade Loop

Scenario 8 closes the grade loop on the air with an aggressive tuning of the kind chosen to get the grade back faster, $K_c = 12$ m^3/min per percent copper and $\tau_I = 1.0$ min against the recommended 2.4 and 4.0 min, and lets a modest feed-grade disturbance, from 0.80 to 0.68 % copper at 12 min, start the loop moving.

Table 5.3 and Figure 5.2 are the central measurement of this chapter. The aggressive grade loop oscillates against the dead time of its own analyzer, and it cycles even with the froth depth held exactly at its setpoint. Every air move it makes also moves the froth depth in about ten seconds, the level loop chases it with a valve that needs nearly three minutes, and the froth depth and the grade swing with the air. Tightening the level loop does not remove the cycle, because the cycle belongs to the grade loop. The cell settles into a sustained oscillation with an 11.3 min period in which the air command falls to its 3 m^3/min floor on every cycle and climbs at its rate limit to about 34 m^3/min, short of the 40 m^3/min ceiling. The cell value first goes negative at 20.5 min and reaches about $–21 300 /h on the worst half-cycles.

The concentrate grade swings between 1.05 and 8.56 %, which is below the cleaner circuit's 4.5 % minimum for half of every cycle (49.5 % of the time after 30 min). The oscillation costs no copper. Both tunings recover 76.3 % of the copper fed over the four hours, but the aggressive loop ships 186 t of concentrate at a mean grade of 3.37 % against 115 t at 5.48 %, and the concentrate handling cost rises from 6305 to $10 251 /h. The tailings valve travels 2257 % of its range against 8 %, and the margin falls from 7270 to $3305 /h, a loss of 55% of the cell's income from one tuning decision on one loop.

> **TRY IT IN THE APP** **8 · An Aggressive Grade Loop Wrecks Its Neighbors**
>
> Load it and press Run in Auto. After the feed grade dips at 12 min, watch the air flow fall to its 3 m^3/min floor on every cycle and climb at its rate limit, the froth depth and the valve swing with it, and the cell value turn negative. Then set the grade loop back to a gain of 2.4 and an integral time of 4.0 min, and run again.

5.6.4 Collector Dosage and Selectivity

Scenario 9 steps the collector from 30 to 55 to 90 g/t. Surface coverage follows (5.3) from 70.4 to 87.7 %, so recovery rises from 80.5 to 83.8 %, an improvement of about three points. Grade falls from 5.48 to 3.89 %, and the margin falls from 8963 to $5962 /h.

The response has the same shape as the froth depth response and a different lesson. The added collector pays for a benefit that (5.3) has already saturated, while the gangue it activates has not saturated, so each increment buys less recovery and costs more grade. Reagent overdose is one of the few plant conditions in which increasing a beneficial input produces a clear economic loss that is visible only in the assay [21].



Figure 5.2 Scenario 8 at 2.50 h, a sustained limit cycle with the air command falling to its 3 m^3/min floor on every cycle, cropped to the plant view and the first four trends. (1) The froth layer in the cell, with the instantaneous recovery of 34.4 % beneath it for a cell value of $3618 /h. (2) The operating point, which lies below the inset's 35 % recovery axis, so it is drawn hollow on the frame, well off the steady-state grade-recovery curve. (3) The alarm line, which reports the concentrate below the cleaner minimum. (4) The froth depth, oscillating about its 0.30 m setpoint. Below it are the tailings valve and the concentrate grade, which crosses the 4.5 % cleaner minimum twice per cycle. (5) The air flow, falling to its floor on every cycle.

5.6.5 Two Quality Loops on Two Independent Actuators

Section 5.6.6 contains the measurement. Adding the tailings-grade controller on the collector dose, so that both quality loops run, takes the graded scenario from 9890 to $7072 /h and the recovery over the run from 68.7 to 50.1 %. The loop is not badly tuned. It asks the collector to hold a tailings assay whose achievable value is set by the froth depth the other quality loop has chosen, and it closes the gap by overdosing into the selectivity loss of Section 5.6.4. The collector reaches its 95 g/t stop at 15.5 min and stays there, the tailings assay still ends at 0.405 % against a 0.18 % setpoint, and the grade loop cuts the air to between 3.0 and 6.4 m^3/min to hold its own setpoint against the overdosed gangue.

5.6.6 Graded Scenario

The graded scenario runs 140 min against a feed grade drop from 0.80 to 0.68 % copper at 25 min, a throughput increase from 12 to 14 m^3/min at 60 min, and a solids increase from 33 to 36 % at 95 min, with 0.006 m of noise on the level transmitter and 0.18 % copper on the concentrate assay, with a smaller share of it on the tailings assay. It is judged at 140 min. The froth depth must stay within

Table 5.4 The graded scenario over 140 min. Each row after the second changes one decision of the suggested configuration. Worst error is the largest froth depth error from 25 min on, and recovery is the copper recovered over the copper fed for the run.

Configuration	Worst error (m)	Mean net ($/h)	Min grade (%)	Recovery (%)
As loaded	0.118	8282	4.16	78.0
Suggested	**0.039**	**9890**	5.48	68.7
as-found setpoints	0.037	8580	4.88	77.5
depth 0.80 m	0.053	8987	5.10	62.5
both quality loops	0.039	7072	5.48	50.1
as-found level loop	0.142	9787	5.48	68.2

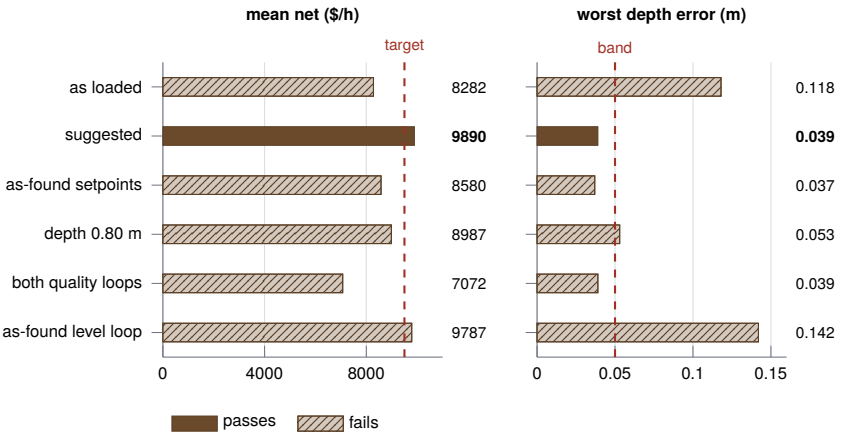


Figure 5.3 The graded scenario of Table 5.4: mean net value against the $9500 /h target, and worst froth depth error against the 0.05 m band. Only the suggested configuration (solid) clears both. Notice that the as-found setpoints hold the depth best of all and still fall short of the target, while the as-found level loop nearly reaches the target and breaks the band.

±0.05 m of its setpoint from 25 min on, the average value must exceed $9500 /h, the concentrate must spend less than 0.05 min below the 4.5 % cleaner minimum, and the froth must never collapse. The suggested configuration is the grade loop alone at 0.55 m and 8.4 % copper, with the collector left at 30 g/t and the level loop at $K_c = 300\,\%/\mathrm{m}$ and $\tau_I = 1.2\,\mathrm{min}$. The attempt is judged on the case as given, so changing a model parameter or a disturbance setting, or firing an extra disturbance, voids it. The application gives the same run whether the suggested settings are applied before Run or carried through Reset.

Every row of Table 5.4 after the second fails, and each fails for a different reason, which is the property a scoring scenario should have. Figure 5.3 draws the two measures that decide most of the verdicts.

The as-loaded configuration fails all three clauses, with a worst depth error of 0.118 m at the throughput step, a mean of $8282 /h, and 7.0 min below the cleaner minimum. Its froth depth setpoint is the design value, which Table 5.1 prices at about $1500 /h below the maximum, and its level loop runs at the as-found $K_c = 30\,\%/\mathrm{m}$ and $\tau_I = 5\,\mathrm{min}$. Returning the suggested configuration to the as-found setpoints of 0.30 m and 5.5 % copper fails on the margin at $8580 /h, and it does so with the best froth depth tracking in the table, a worst error of 0.037 m. Accurate regulation at the wrong setpoint costs $1309 /h.

> ✦ **DID YOU KNOW? The Tightest Loop Fails**
>
> From the simulator: of the six configurations in Table 5.4, the one with the as-found setpoints holds the froth depth best, with a worst error of 0.037 m. It still fails the graded scenario, earning \$8580 /h against the \$9500 /h target, because it holds the wrong depth.

Excessive depth fails differently. At 0.80 m the grade loop must pull a richer froth down to 8.4 % copper, so it runs the air at 20 to 30 m^3/min, well past the 16.8 m^3/min at which the froth turns churn-turbulent, and the mean net of \$8987 /h is below the target. Its worst depth error of 0.053 m is within the noise, because eight noise seeds give 0.047 to 0.061 m, while the margin fails on every seed.

Enabling the second quality loop fails on the margin by a wide gap, at \$7072 /h and 50.1 % recovery. Returning the level loop to its as-found tuning fails only on the band, at 0.142 m during the throughput step, for the reason established in Section 5.4, while it costs just \$103 /h. The 180 %/m the other scenarios ship is marginal here, with worst errors of 0.046 to 0.057 m over the eight seeds, which is why the suggestion is 300 %/m. This level loop is fast because the disturbance it rejects originates inside the same control system at seventeen times its own speed.

🕮 **Pause.** Try Exercise 5.3 before reading on.

5.6.7 Economic MPC Against the Three Loops

Scenario 11 loads the cell of scenario 5 with the MPC of Section 5.5.1 on the two quality handles, raises the grade request from 5.5 to 7.0 % copper at 10 min, and lowers the feed grade from 0.80 to 0.55 % copper between 50 and 100 min. Every comparison runs 140 min on one seed with identical instruments, noise and scripted events, and no structure receives a preview. The MPC keeps the settings of scenario 11 in every case and is not retuned. Settings that differ from a scenario's own are applied just after the first integration step, so every run starts from the same state. Two comparators locate the source of the result. The first enters the MPC model's own optimum, 8.34 m^3/min and 36.9 g/t, as fixed stations with both quality loops open, so it has neither an optimizer nor quality feedback. The second gives the shipped loops the same optimum as setpoints, 7.43 % copper with the collector fixed at 36.9 g/t for the grade loop alone, and 7.43 % with a 0.245 % tailings grade for both loops.

Against the grade and tailings loops as shipped, the MPC earns 812, 562 and \$1050 /h more in scenarios 5, 6 and 11, and it never lets the true grade fall below 4.5 % copper, as listed in Table 5.5. The fixed stations earn within \$11 /h of it in all three cases (\$2 /h more in scenario 5 and \$11 /h less in scenarios 6 and 11), and the grade loop alone at the optimum comes within \$27 /h. The gain is therefore the operating point that the model selects rather than the optimization over the horizon. Volume 9, Chapter 8 also finds an MPC gain on its data hall that does not come from the horizon. At a froth depth of 0.30 m the value of the optimum changes little with these disturbances. On the 0.55 % feed the optimum lies at 6.26 m^3/min and 36.9 g/t and earns \$5606 /h at steady state, only \$172 /h more than the design optimum held there, and the lean feed lasts 50 of the 140 min. The pairing lesson of Section 5.6.5 survives. Both loops given the optimum's own setpoints earn 1403, 278 and \$1250 /h less than the MPC, and in scenarios 5 and 11 the tailings loop again drives the collector to its 95 g/t stop, at 10.8 min. What settles the pairing argument is the absence of a tailings setpoint, which the MPC shares with the grade loop alone. In scenario 11 the grade loop alone as shipped already recovers \$933 /h of the MPC's \$1050 /h. Holding the MPC's collector at 30 g/t costs \$68 /h.

The lean feed separates the structures most clearly, as shown in Figure 5.4. At 70 min the MPC holds 6.99 % copper on 5.12 m^3/min of air and 37.0 g/t of collector for \$5482 /h. The grade and tailings loops hold 6.42 % with the air at its 3 m^3/min floor and the collector at its stop since 39.2 min, for \$4422 /h, and the fixed stations give 5.44 % for \$5434 /h. Over the lean period the three average 5560, 4472 and \$5545 /h. The MPC recovers the grade more slowly than the grade loop alone. Its grade falls to 5.56 % at 54 min and rises without overshoot to stay above 6.9 % from 67.2 min, while

Table 5.5 Cell value and grade violations over 140 min in the pairing case (scenario 5), the feed grade drop (scenario 6), and the MPC case (scenario 11). Value is the run average. Below 4.5 counts minutes of true grade under the cleaner minimum, and below request counts minutes more than 0.1 % copper under the larger of the grade setpoint in force and 4.5 %. The MPC made 140 decisions in each case with no fallback.

Structure	Value ($/h)	Below 4.5 (min)	Below req. (min)
Scenario 5, the pairing case			
Grade and tailings loops, as shipped	8945	0	0
Grade loop alone, as shipped	8979	0	0
Economic MPC	**9758**	0	0
Model optimum as fixed stations	9760	0	0
Grade loop alone at the optimum	9757	0	12.1
Grade and tailings loops at the optimum	8354	0	71.8
Scenario 6, the feed grade drop			
Grade and tailings loops, as shipped	7651	3.6	16.5
Grade loop alone, as shipped	7708	3.4	17.3
Economic MPC	**8213**	0	0
Model optimum as fixed stations	8202	0	0
Grade loop alone at the optimum	8187	0	37.8
Grade and tailings loops at the optimum	7935	3.3	36.8
Scenario 11, the MPC case			
Grade and tailings loops, as shipped	7163	0	100.8
Grade loop alone, as shipped	8097	0	27.7
Economic MPC	**8213**	0	16.7
Model optimum as fixed stations	8202	0	51.2
Grade loop alone at the optimum	8186	0	39.7
Grade and tailings loops at the optimum	6963	0	95.0

the grade loop first passes 6.9 % at 56.4 min and then swings between 6.56 and 7.30 % until 73.7 min. The bias sees the lean feed only through 2 min of analyzer dead time and a 3 min filter. Told the true feed grade, the same MPC never falls below 6.59 %. The same delay acts when the feed recovers at 100 min. The MPC is still on its lean-feed air of about 5 m^3/min, so the grade rises to 9.25 % at 105 min before the air returns to about 8.3 m^3/min, while the grade and tailings loops swing about the request.

The MPC loses two cases to the best fixed setting of the grade loop alone, found by a search that knows the whole case and is therefore a bound rather than a controller. In scenario 7 it earns $13 777 /h against $13 865 /h, a loss of $88 /h. After the step to 38 % solids at 55 min its design-feed model keeps dosing 36.9 g/t where that ore's optimum is 28.2 g/t. Holding the collector at 28.2 g/t from 55 min gives $13 877 /h, and telling the model the true feed gives $13 887 /h, both above the bound. In scenario 10 it earns $9715 /h against $9739 /h, a loss of $24 /h, because the model takes the air command as delivered while the blower supplies 62 % of it. At 50 min it commands 12.57 m^3/min (7.79 m^3/min delivered) against 13.51 m^3/min, and told the delivery it earns $9755 /h. In scenario 8 the MPC replaces the aggressive grade loop and earns $7864 /h against $3625 /h, with no time below 4.5 % against 61.3 min.

In the graded scenario the MPC does not move the froth depth. With the shipped settings it

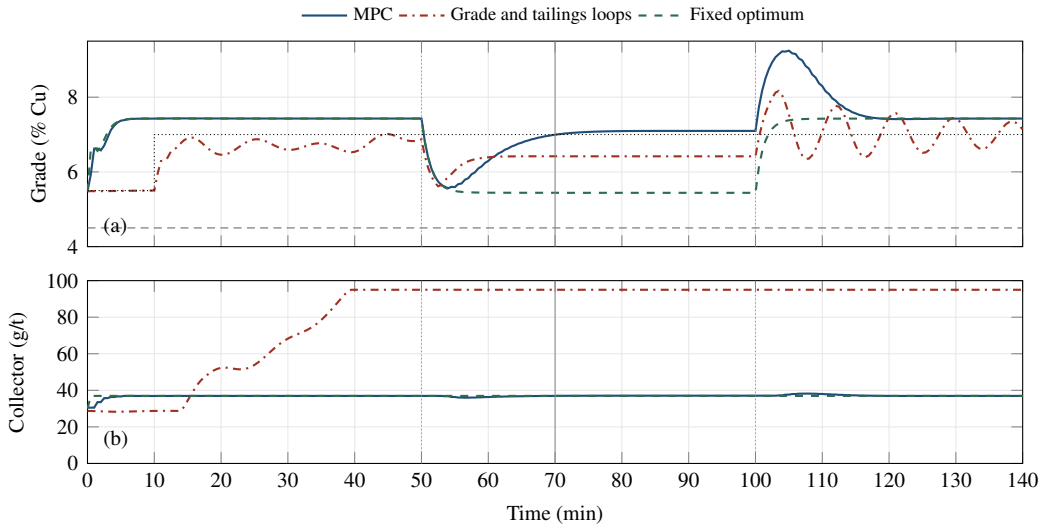


Figure 5.4 Scenario 11 under the MPC (solid), the shipped grade and tailings loops (dash-dotted), and the MPC model's optimum as fixed stations (thick dashes). (a) True concentrate grade with the request (dotted) and the 4.5 % minimum (thin gray dashes). (b) Collector dose. The lean feed lasts from 50 to 100 min (dotted), and the solid line marks 70 min.

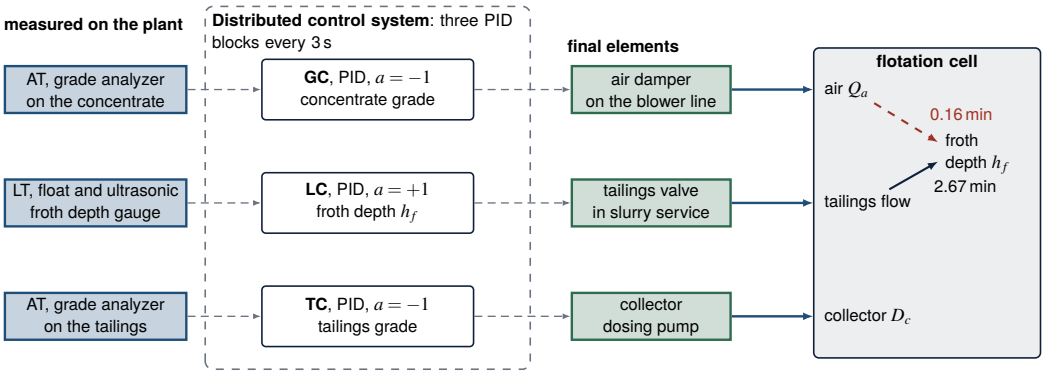


Figure 5.5 The three loops of the simulated cell drawn with their hardware. The froth depth gauge LT and the two grade analyzers AT feed three PID blocks. LC is reverse acting and moves the tailings valve, GC and TC are direct acting and move the air damper and the collector dosing pump. The dashed path is the one that matters: the air reaches the froth depth in 0.16 min, while the valve that LC owns needs 2.67 min.

averages \$9032 /h and still fails both the \$9500 /h target and the band, with a worst depth error of 0.139 m from the as-found level loop. With the suggested settings it earns \$9917 /h against \$9890 /h for the suggested grade loop, with a worst depth error of 0.041 m. Both pass, and on five other noise seeds the MPC leads by 20 to \$26 /h.

5.7 Technology in Practice

Figure 5.5 draws the three loops of Section 5.5 with their hardware.

What Measures It

The model reduces the froth depth gauge to a 0.08 min transmitter filter, which stands for the float, the ultrasonic sensor and their electronics.

> **TECHNOLOGY Sensor: A Float and an Ultrasonic Gauge for Froth Depth**
>
> An ultrasonic level gauge times a sound pulse from a piezoelectric transducer to a surface and back, and the time of flight gives the distance. Foam is a poor target: its top returns a weak echo and the foam scatters the echo from the liquid below, and the speed of sound in the air above shifts with its temperature [22]. On industrial flotation cells the froth depth is therefore measured with a float coupled to an ultrasonic sensor, and because the overflow lip is fixed, the froth depth and the pulp level add to a constant. A thick slurry can stick the float and make the reading wrong [23]. A stuck float reads a depth that no longer follows the pulp, and the level loop then moves the valve against it.

The two grades are measured by on-stream analyzers, which the model represents as 2.0 and 3.5 min of dead time with lags of 0.5 and 0.8 min. On a plant the concentrate grade can be estimated on line by an X-ray fluorescence analyzer, an estimate that demands maintenance and calibration [24]. The bubble size, the gas holdup and the collector coverage have no instrument in the model.

What Moves It

The model limits the tailings valve to 0 to 100 % and the collector to 0 to 95 g/t, with no rate limit on either. The air command passes a clamp of 3 to 40 m^3/min and a rate limit of 7 m^3/min per minute, and the air that reaches the cell is the command multiplied by the blower delivery.

> **TECHNOLOGY Final Element: A Slurry Valve on the Tailings**
>
> The tailings valve throttles a slurry of ground ore and water. Valves that narrow the flow path with a flexible sheet, much as a hand pinches a hose, suit flows that carry solids, although the elasticity of the sheet makes precise throttling difficult. Hard particles also erode the trim and body of a valve in slurry service, so such a valve is a fast-wearing part that is rebuilt or replaced at regular intervals. The air is throttled by a different device, a damper, which handles large flows of air at low pressure [22]. A worn tailings valve passes more slurry at the same opening, so the output that holds a given froth depth drifts as the trim wears.

A blower that cannot deliver, scenario 10, appears to the grade loop as a loss of gain rather than as a limit.

Where the Controller Runs

Control loops in a plant are often nested or cascaded, and supervisory and lower-level loops run with cycle times from milliseconds to minutes. A distributed control system uses a supervisory loop to mediate a group of localized controllers [25], and programmable logic controllers run PID loops as often as they switch discrete devices [22]. The app integrates every 0.005 min and runs its three PID blocks on a 0.05 min (3 s) scan. On such a platform the chapter's structure is three velocity-form PID blocks with back-calculation, and the economic MPC is a supervisory application that writes the air and collector commands once a minute while the depth block keeps the valve.

> **IN PRACTICE The Analyzer Behind the Grade Loop**
>
> A review of flotation control counts the available instrumentation among the biggest challenges at industrial scale. On-line estimation of concentrate grade with an X-ray fluorescence analyzer generally demands a significant amount of maintenance and calibration, and plant grade and recovery data usually vary widely [24]. Where instruments fall short, operators run the cells by eye: in sulphur flotation, control has relied mainly on operators who watch the froth's appearance from a spot away from the control room [23]. In this chapter the grade loop sees the cell only through its analyzer, 2.0 min late.

What Is Modernizing It

The first development is machine vision on the froth. Image analysis at the top of a cell measures the number, size, shape, speed and stability of the bubbles at the surface, although estimating the mineral content of the froth from images has not yet been fully implemented at industrial scale [24]. One study estimates the froth depth itself from froth images and the measured air, tailings and feed flows, in place of a float that sticks [23]. The second is model predictive control, whose partial adoption in flotation Section 5.2 describes and whose gain on this cell Section 5.6.7 traces to the operating point. The third is searching for the best air rate on line: extremum-seeking schemes on real cells exploit the peak in air recovery [26], a maximum that this chapter's model represents only with a lumped term.

5.8 Limitations

The pulp is one perfectly mixed zone with two lumped mineral classes. Real flotation is size-dependent in a way that dominates plant behavior. Fine particles float slowly because they collide with bubbles rarely, coarse particles detach, and the recovery-by-size curve has an interior maximum that sets what a circuit can achieve. Predicting grade and recovery from liberation and size data is an active subject [27], and this model does not substitute for it.

The froth is a recovery factor rather than a transported phase. The lumped destabilization term above 1.4 cm/s of superficial gas velocity gives the steady-state recovery an interior maximum against the air rate (82.7 % near 21 m^3/min at the design valve position), but the froth transport behavior behind peak air recovery (residence time, bubble bursting, and surface mobility) is not represented. Extremum-seeking schemes that exploit peak air recovery on real cells [26] would find this surrogate maximum rather than one produced by froth physics.

> ⚠ **OPEN CHALLENGE** **Froth Physics in a Predictive Controller**
>
> Little evidence of successful model predictive control of flotation at industrial scale is available. The models used for it so far describe the pulp with kinetic rate equations and often treat the froth the same way, although froth performance is driven by bubble coalescence, bubble bursting and the motion of liquid and solids [24]. This chapter's froth is a recovery factor of the same simple kind. Progress would be froth models that carry this physics and are still simple enough to run inside a predictive controller.

There is one cell. A real rougher bank is eight to twelve cells in series with a cleaner and a scavenger around it, the grade-recovery curve of the circuit is not the curve of any cell in it [1, 2], and the recycle streams between the units are where the control problem becomes multivariable in the conventional sense.

Reagent chemistry is one collector with a Langmuir coverage and one frother with a bubble-size effect. Real collectors are selective by mineral and by pulp potential, depressants and activators are absent, and the difference between two collectors at the same dose can exceed the difference between two doses of one [21]. Entrainment of fine gangue with the recovered water is folded into the gangue rate constant rather than modeled as a separate mechanism, which understates how much of the grade problem is hydraulic rather than chemical.

The MPC's model is the declared balance at the design feed and the setpoint depth. It does not know the feed grade, rate or solids, the blower delivery, the reagent strength, or any parameter or price changed in a run. The output biases absorb their effect on the grades but not on the location of the optimum, which is the cause of the 88 and $24 /h losses in scenarios 7 and 10. One lag shape serves both grades, so the slower answer of the gangue to the collector (3.95 min from the balance, against 0.86 min for the valuable mineral) is absent. Recovery is not measured and enters the objective only through the two-product formula at the design feed grade. The objective and the cell value both charge the reagent pumped, of which only the fraction set by the reagent strength acts in the pulp. With a strength of 0.7 from the start of scenario 5, the MPC still settles at 36.9 g/t of collector and earns $9492 /h, against $9758 /h at full strength. The grade floor carries one slack for the whole

horizon, the froth depth is not a decision of the MPC, and the MPC does not coordinate the air with the tailings valve, so the fast path of Section 5.4 stays with the level loop. Because the optimum moves little under the scripted disturbances, this cell credits the MPC mainly with the operating point, which a fixed setting from the same model also reaches.

5.9 Discussion

Three results generalize past this cell.

The first is that when the objective is a maximum rather than a target, the setpoint decision dominates the tuning decision. Table 5.4 contains a configuration with the best regulatory performance in the study, a worst error of 0.037 m, that earns \$1309 /h less than the suggested one because its setpoint is wrong. The worst regulation in the table, nearly four times larger at 0.142 m, costs \$103 /h. A course that teaches only regulation prepares a student for the smaller part of this problem [7]. Volume 9, Chapter 10 asks the same question of an activated sludge plant, where the choice is where a supervisory setpoint belongs.

The second is that interaction is a statement about bandwidth as much as about gain. The steady-state gain matrix of this cell is unremarkable, and the air nevertheless reaches the froth depth seventeen times faster than the valve that holds it. Every air move of the grade loop therefore arrives at the level loop as a disturbance it cannot stop. An aggressive grade loop, which cycles even with the froth depth held fixed, carries its own limit cycle into the froth depth, the valve and the margin. The diagnostic question is not only which actuator affects which variable, but how fast, compared with the loop that owns that variable [4].

> ⚠ **OPEN CHALLENGE** **Air in the Level Model**
>
> The air rate directly affects the pulp level of an aerated tank, yet the hydraulic model behind a published level control strategy for a bank of six cells left that effect out. A model that includes it predicts pulp level changes better, but it has been validated only at laboratory scale and has never been tested as part of a level control strategy [24]. This chapter's cell carries the effect through the gas holdup. Progress would be level controllers on industrial banks that take every air move into account.

The third is that a second loop is not automatically an improvement. The tailings-grade controller in Table 5.4 is well posed and well tuned, and it costs \$2818 /h, because it holds a variable whose achievable value is set by a decision another loop has already made. Loops that contest the same economic trade-off are better replaced by one loop and an explicit optimizer [6, 10, 12].

5.10 Conclusion

A flotation cell is judged on a maximum rather than a setpoint, and the application locates that maximum near 0.55 m of froth depth, 9.31 % grade, and 69.30 % recovery, worth \$10 447 /h against \$8963 /h at the design point. Reaching it costs 11.2 points of the recovery figure reported in plant metal accounting, and the margin is within 1% of its maximum over a 0.14 m span of depth, so the setpoint is a standing decision rather than a calculation. Beneath the economics sits an interaction carried by bandwidth. Air reaches the froth depth in 0.16 min while the tailings valve needs 2.67 min, so every move of the grade loop disturbs the level loop faster than the valve can answer. An aggressively tuned grade controller carries its limit cycle into the froth depth, takes the air to its floor on every cycle, travels the valve through 2257 % of its range in four hours, and removes 55% of the margin without losing any copper. The graded scenario is won by selecting the correct depth and lost, in four measured ways, by four otherwise reasonable configuration choices. An optional economic MPC on the air and the collector earns 562 to \$1050 /h more than the three loops in scenarios 5, 6 and 11, but its own model's optimum entered as fixed settings comes within \$11 /h of it, so on this cell its value lies in the operating point and in having no tailings setpoint. Economic setpoint selection and bandwidth interaction can both be measured on the one cell, with the scenarios of Section 5.6.

KEY POINTS

- A flotation cell is judged on its margin, and the best froth depth gives up recovery to earn more.
- Deepening the froth from 0.30 m to 0.55 m raises the net value by 17% and costs 11.2 points of recovery.
- Air reaches the froth depth in 0.16 min while the tailings valve needs 2.67 min, so the level loop cannot stop an air move.
- An aggressively tuned grade loop cycles against its own analyzer and removes 55% of the margin without losing any copper.
- The economic MPC's gain lies in the operating point and in having no tailings setpoint, which fixed settings from its model nearly match.

Where else this applies. A setpoint chosen on cost in Volume 9, Chapter 8, loops that interact against the relative gain in Volume 6, Chapter 2, and two quality loops on one product in Volume 6, Chapter 3.

References

[1] B. A. Wills and J. A. Finch, *Wills' Mineral Processing Technology: An Introduction to the Practical Aspects of Ore Treatment and Mineral Recovery*, 8th ed. Oxford: Butterworth-Heinemann, 2015.

[2] J. Drzymała, P. B. Kowalczuk, M. Oteng-Peprah, D. Foszcz, A. Muszer, T. Henc, and A. Łuszczkiewicz, "Application of the grade-recovery curve in the batch flotation of Polish copper ore," *Minerals Engineering*, vol. 49, pp. 17–23, 2013.

[3] J. Jeswiet and A. Szekeres, "Energy consumption in mining comminution," *Procedia CIRP*, vol. 48, pp. 140–145, 2016.

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

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

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

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

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

[9] Supreme Court of the United States, "Minerals Separation, Ltd. v. Hyde, 242 U.S. 261," Opinion of the court, decided December 11, 1916; Legal Information Institute, law.cornell.edu/supremecourt/text/242/261, 1916.

[10] D. Hodouin, "Methods for automatic control, observation, and optimization in mineral processing plants," *Journal of Process Control*, vol. 21, no. 2, pp. 211–225, 2011.

[11] B. J. Shean and J. J. Cilliers, "A review of froth flotation control," *International Journal of Mineral Processing*, vol. 100, no. 3–4, pp. 57–71, 2011.

[12] L. G. Bergh and J. B. Yianatos, "The long way toward multivariate predictive control of flotation processes," *Journal of Process Control*, vol. 21, no. 2, pp. 226–234, 2011.

[13] C. R. Boeree, "Up-scaling of froth flotation equipment," MSc thesis, report AES/RE/14-08, Delft University of Technology, 2014, resolver.tudelft.nl/uuid:660e58d7-a82d-46d9-a8ff-1bfba42b2b05.

[14] B. K. Gorain, J.-P. Franzidis, and E. V. Manlapig, "Studies on impeller type, impeller speed and air flow rate in an industrial scale flotation cell. Part 4: Effect of bubble surface area flux on flotation performance," *Minerals Engineering*, vol. 10, no. 4, pp. 367–379, 1997.

[15] D. A. Deglon, F. Sawyerr, and C. T. O'Connor, "A model to relate the flotation rate constant and the bubble surface area flux in mechanical flotation cells," *Minerals Engineering*, vol. 12, no. 6, pp. 599–608, 1999.

[16] K. Heiskanen, "On the relationship between flotation rate and bubble surface area flux," *Minerals Engineering*, vol. 13, no. 2, pp. 141–149, 2000.

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

[19] M. Ellis, H. Durand, and P. D. Christofides, "A tutorial review of economic model predictive control methods," *Journal of Process Control*, vol. 24, no. 8, pp. 1156–1178, 2014.

[20] 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.

[21] S. J. Adkins and M. J. Pearse, "The influences of collector chemistry on kinetics and selectivity in base-metal sulphide flotation," *Minerals Engineering*, vol. 5, no. 3–5, pp. 295–310, 1992.

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

[23] M. He and B. Sun, "On-line froth depth estimation for sulphur flotation process with multiple working conditions," *IEEE Access*, vol. 7, pp. 124 774–124 784, 2019.

[24] P. Quintanilla, S. J. Neethling, and P. R. Brito-Parada, "Modelling for froth flotation control: A review," *Minerals Engineering*, vol. 162, p. 106718, 2021.

[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] D. A. Wepener, J. D. le Roux, and I. K. Craig, "Extremum seeking control to optimize mineral recovery of a flotation circuit using peak air recovery," *Journal of Process Control*, vol. 129, p. 103033, 2023.

[27] M. Gupta, K. Huang, and R.-H. Yoon, "Predicting the recovery and grade of a rougher flotation circuit from liberation data," *Minerals Engineering*, vol. 188, p. 107853, 2022.

5.11 Exercises

5.1 Concept check · 1 min

Which of these quantities reaches a controller of Figure 5.1 from an instrument?

(a) The Sauter mean bubble diameter d_{32}
(b) The copper grade of the concentrate
(c) The collector coverage θ_c of the particles
(d) The bubble surface area flux S_b

5.2 Concept check · 2 min

In Equation (5.4) water drains from the froth with $\lambda_w = 4.5/\text{m}$ and attached particles with $\lambda_p = 1.8/\text{m}$. What does the inequality $\lambda_w > \lambda_p$ produce as the froth deepens?

(a) A recovery that rises, because particles have longer to attach
(b) A flooded cell at any depth under the 0.06 m collapse depth
(c) A concentrate that becomes smaller and richer in copper
(d) A bubble surface area flux that falls with the froth recovery

5.3 Concept check · 2 min

What does the graded scenario of Section 5.6.6 reward?

(a) The most copper recovered, at any grade above the cleaner minimum
(b) The tightest froth depth regulation at the 0.30 m design depth
(c) All three loops closed, each holding its setpoint through the upsets
(d) The highest average cell value, with the froth depth held inside its band

5.4 Theory and calculation · 8 min

At the design point the solids feed is $F_s = 4999\,\text{kg/min}$ at 0.80 % copper, the concentrate assays 5.48 % and the tailings 0.176 %. Copper is worth \$9.0/kg with a payable fraction of 0.965, and concentrate handling costs \$0.22/kg. The collector (30 g/t) costs \$2.8/kg, the frother (20 g/t) \$3.5/kg, and the blower \$0.02 per m^3/min per minute at 14.4 m^3/min.

(a) Compute the concentrate solids from the two-product formula, to 1 kg/min. *Checkpoint:* 588 kg/min

(b) Compute the net value with Equation (5.6), to \$10 /h, and compare it with Table 5.1. *Checkpoint:* \$8970 /h against 8963

(c) At 0.55 m the table gives 9.31 % copper and 69.30 % recovery at the same air and doses. Find the tailings assay to 0.001 %, then the net value to \$10 /h. *Checkpoint:* 0.261 %; \$10 450 /h

5.5 Theory and calculation · 6 min

At design the cell takes 14.4 m^3/min of air, which gives a superficial gas velocity of 1.20 cm/s, a gas holdup of 12.0 % and a Sauter mean diameter of 1.496 mm. The overflow lip is 3.0 m above the bottom and the froth is 0.300 m deep.

(a) From Equations (5.1) and (5.2), find the cross section A, the holdup coefficient k_ε to three decimals, and S_b to 0.1/s. *Checkpoint:* 20 m^2; 0.167 min/m; 48.1/s

(b) Scenario 3 steps the air to 20 m^3/min. With the liquid volume unchanged, find the new holdup and the froth depth, to 0.001 m. *Checkpoint:* 16.7 %; 0.149 m

(c) The run of Section 5.6.1 thins the froth only to 0.201 m. Give two reasons from the model. *Checkpoint:* damper lag; the inventory drains meanwhile

5.6 Theory and calculation · 5 min

The grade loop runs the velocity-form PID of Section 1.6.2 with action $a = -1$, $K_c = 2.4\,m^3/min$ per percent copper and $\tau_I = 4.0$ min, every $\Delta t = 0.05$ min, with the air limited to 3 to 40 m^3/min and rate-limited at 7 m^3/min per minute.

(a) The analyzer reading jumps from setpoint to 0.5 % copper above it. Find the proportional move and the integral move per scan, with their signs, and say whether the loop is direct or reverse acting in Chapter 1's convention. *Checkpoint:* +1.2 m^3/min; +0.015 m^3/min per scan; direct

(b) How far may the air move in one scan, how long does the damper take to cross its range, and how long does the rate limit take to deliver the proportional move of scenario 8's $K_c = 12$ for the same error? *Checkpoint:* 0.35 m^3/min; 5.3 min; 0.86 min

(c) What fraction of the excess does back-calculation remove from the integral per scan at $\tau_I = 4$ and at $\tau_I = 1$ min? *Checkpoint:* 1/80; 1/20

5.7 App challenge · 8 min

Open scenario 4, "Level control alone (quality loops off)", at the default seed. The air steps from 14.4 to 22 m^3/min at 8 min.

(a) Run it as shipped, with the froth depth gain at 180 %/m, at 5 min per second to 30 min, pause and export the CSV. Its time column is in minutes. Find the most negative `froth_error_m` and the lowest `grade_pct` between 8 and 25 min, within 0.002 m and 0.02 %, and read the froth depth loop's IAE tile. *Checkpoint:* −0.111 m at 8.30 min; 3.67 %; IAE 0.294

(b) Press Reset, set Froth depth gain to 30 %/m, and run and read the same three values. *Checkpoint:* −0.127 m at 8.40 min; 3.21 %; IAE 0.727

(c) The gain fell by a factor of six. By what factor did the peak error grow, and why so little? *Checkpoint:* about 1.14

5.8 App challenge · 9 min

Open scenario 8, "An aggressive grade loop wrecks its neighbors", at the default seed. Run each part for 240 min at 20 min per second, then read the Average value tile, within \$5 /h, and the froth depth loop's IAE tile, within 0.01, and MV travel tile.

(a) Run it as shipped, with the grade loop at $K_c = 12$ and $\tau_I = 1.0$ min. *Checkpoint:* \$3305 /h; IAE 9.805; travel 2257 %

(b) Reload, set Grade loop gain to 2.4 and Grade loop integral time to 4.0 min before Run, and run. *Checkpoint:* \$7270 /h; IAE 0.036; travel 8.21 %

(c) Reload, leave the grade loop as shipped, set Froth depth gain to its 600 %/m maximum before Run, and run. Does a tighter level loop cure the cell? *Checkpoint:* \$1325 /h; IAE 2.986; travel 3410 %; no

5.9 App challenge · 9 min

Open the challenge scenario, "find the money on the grade–recovery curve", at its default seed. Run each part for the full 140 min at 20 min per second and read the challenge result in the progress panel: the worst depth error, within 0.002 m, the average value, within \$5 /h, and the minutes below 4.5 % copper.

(a) Run it as shipped. *Checkpoint:* fails: 0.118 m, \$8282 /h, 7.0 min

(b) Reload the scenario with its reload button, press Reveal suggested tuning and run. *Checkpoint:* passes: 0.039 m, \$9890 /h, 0.0 min

(c) Reload, press Reveal suggested tuning, then set Froth depth setpoint to 0.30 m and Concentrate grade setpoint to 5.5 % copper before Run. Read the result, and compute what the wrong setpoints cost against (b). *Checkpoint:* fails: 0.037 m, \$8580 /h; about \$1310 /h

5.10 Application area · 4 min

A plant manager's monthly report leads with copper recovery, and the move from 0.30 m to 0.55 m of froth in Table 5.1 cuts it from 80.54 to 69.30 %. The feed carries 2400 kg of copper an hour, and the concentrate falls from 588 to 298 kg/min.

(a) Price the copper given up (payable at 0.965 of \$9.0/kg) and the handling saved (at \$0.22/kg), each to \$10 /h. *Checkpoint:* \$2340 /h; \$3830 /h

(b) Give two reasons from the chapter why the plant should not simply move to the optimum and stay there. *Checkpoint:* flat optimum; the cleaners, contract and ore decide it

5.11 Application area · 4 min

Scenario 9 steps the collector from 30 to 90 g/t with the quality loops open. Recovery rises from 80.5 to 83.8 %, the grade falls from 5.48 to 3.89 %, and the cell value falls from 8963 to \$5962 /h. The solids feed is 4999 kg/min and the collector costs \$2.8/kg.

(a) Compute the loss per extra g/t and the cost of the extra collector itself, each to \$1 /h. *Checkpoint:* \$50 /h per g/t; \$50 /h

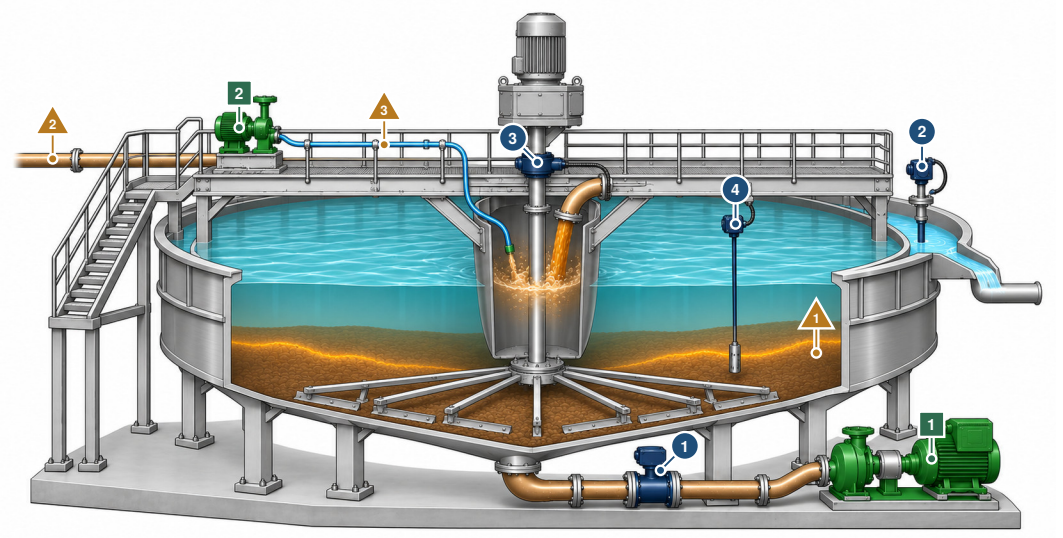
(b) An operator doses collector to hold the tailings assay. What does the tailings trend show in this scenario, and what should set the dose instead? *Checkpoint:* it improves at every step; the margin

5.12 Review the agent · 5 min

An AI assistant reads Table 5.5 and writes: "In scenario 11 the economic MPC earns \$8213 /h against \$7163 /h for the shipped loops, \$1050 /h more. That gain is the value of online optimization over a one-hour horizon, so the plant should buy the optimizer." Find the flaw and say how to check it on the app.

CHAPTER 6

A Mining Thickener Simulator for Teaching Override Selector Control



The controller moves 1 the speed of the underflow pump 2 the flocculant dose

The controller reads 1 the density of the underflow 2 the clarity of the overflow water 3 the load on the rake 4 the depth of the settled bed

What disturbs it 1 finer ore that settles more slowly 2 more feed, or a thicker one 3 weaker flocculant from the make-up

Notice that the one pump at the bottom sets the underflow density, the bed depth and the load on the rake together. Illustration, not to scale.

IN PLAIN TERMS **Squeezing Water Out of Mine Waste**

The main idea. A thickener is a huge round tank where ground rock settles out of the water used to process it. Clear water spills over the rim to be used again, and a pump at the bottom draws off the thick mud.

The challenge. Slowing the pump keeps the mud in the tank longer, so it packs denser and less water is lost. But a deeper, heavier layer is harder for the slow rake to drag, so two safety controllers must be able to speed the pump up first.

Why it matters. Water that leaves with the mud must be replaced, and a stalled rake in a full tank takes days to recover.

THREE WAYS TO READ THIS CHAPTER

Explore	9 min	The first page, the two history boxes of Section 6.1, the fact box of Section 6.2.2, the try-it box of Section 6.4.1, the fact box of Section 6.4.4, 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 6.4, Section 6.5 and Section 6.6 with its open challenge

The control question Slowing the underflow pump is the only way to make a denser underflow, and it is also what deepens the bed the rake has to drag. How should a density controller and two protective controllers share that one pump, and what happens when the selector between them is wired the wrong way?

Learning objectives

- Compute the bed residence time, underflow density, bed depth and rake torque at an operating point from the compression and torque laws.
- Choose a selector's direction from the action of each controller and the direction that relieves each limit.
- Explain what external reset does to a losing controller and predict when a protection takes the pump.
- Compare the selector with the predictive controller on torque, density error and pump travel, and say what each protection costs.
- Choose a density setpoint the thickener can hold through a shift of disturbances.

Prerequisites Sections 1.4.3 (density against torque), 1.4.6 (selector direction and external reset), 1.6.2, 1.6.3 and 1.6.4.

The app https://apmonitor.com/apps/thickener/

See also Volume 4, Chapter 2 and Volume 4, Chapter 5 for protective overrides on other machines, and Volume 4, Chapter 4 for a low selector where less feed is the safe direction. Volume 2, Chapter 3 for back-calculation at an actuator bound.

6.1 Introduction

A mining thickener recovers process water by letting solids settle. The feed enters a large shallow tank through a central feedwell, the solids form a bed, the clarified water overflows a peripheral launder and is returned to the plant, and a slowly rotating rake drags the thickened solids to a central discharge cone where a variable-speed pump removes them.

> **HISTORY A Shallow Tank with Slow Scrapers (1907)**
>
> In 1907 John V. N. Dorr, of Lead, South Dakota, patented a decanting apparatus for the slimes of the cyanide process. The slimes had to be thickened until they carried one to three parts of liquid to one part of solids. In the cones used for this, solids collected on the slopes and then slid down suddenly, choking the discharge. Dorr's tank was of large diameter but shallow, 35 feet across and no more than 12 feet deep in the example he gave, and scrapers on slowly turning arms moved the thickened material gently and continuously to a central discharge without stirring the liquid above. The arms could be raised and lowered [1]. The vessel of this chapter, with its rake and central cone, keeps that form.

The economic objective is the densest underflow the unit can produce, because every cubic meter of water that leaves with the tailings has to be replaced. The limit on that density is mechanical rather than economic. The two collide directly because only one quantity sets the underflow density at steady state. Solids leave the bed at the concentration the compression zone has produced, and that concentration depends on how long the solids have been under the weight of the bed above them, as shown in Equation (6.1).

$$\begin{aligned} \tau_b &= \frac{M_s}{F_{\text{in}}}, \qquad X_u^* = X_{\min} + \Delta X \left(1 - \mathrm{e}^{-\tau_b/\tau_c}\right) \\ X_u &= X_u^* \, \sigma^{0.15} \max\left(0.80,\ 1 - 0.06 \max\left(0, \frac{D_f}{D_0} - 2.2\right)\right) \end{aligned} \tag{6.1}$$

where M_s is the solids inventory in tonnes, F_{in} the solids feed rate in tonnes per hour, τ_b the bed residence time in hours, X_u the underflow solids concentration in kg/m^3, $X_{\min}$ and ΔX the uncompressed density and the compression span, τ_c the compression time constant, σ the ore settleability as a fraction of the design value, D_f the flocculant dose and D_0 the reference dose. The second line is the overdose penalty, because a floc built at more than 2.2 times the reference dose

Table 6.1 Steady-state cost of underflow density at design feed. The rake trip is at 92 %.

Density (%)	Pump (%)	τ_b (h)	Bed (m)	Torque (%)
55	66.47	2.60	1.093	31.4
60	58.00	3.87	1.419	45.9
62	54.99	4.59	1.597	54.9
64	52.17	5.55	1.831	67.2
66	49.52	6.93	2.170	85.4
68	47.03	9.36	2.783	117.5

holds water and yields a thinner underflow at the same residence time. Every symbol carries the value listed in Table 6.4. Equation (6.1) is an inventory whose product quality depends on holdup, which is a structure that recurs wherever a residence time sets a conversion [2]. Slowing the underflow pump raises the inventory, which raises the residence time, which raises the density, and no second actuator has that effect. The inventory is also what the rake has to drag through, so the torque of Equation (6.2) rises with the depth and the density together.

$$T_q = T_{q0} + k_\tau h_b \left(\frac{X_u}{X_{\mathrm{ref}}} \right)^n \tag{6.2}$$

Here T_q is the rake torque as a percentage of the rated value, T_{q0} the torque with no bed, k_τ the torque per meter of bed, h_b the bed depth in meters, X_{ref} a reference density and n the density exponent.

Table 6.1 states the design problem in full. The relationship is convex, because the step from 66 to 68 % solids costs 32.1 points of rated torque while the whole span from 55 to 62 % costs 23.5. The 92 % trip is reached at 66.54 % solids, so a setpoint above that value is not aggressive but infeasible.

That structure is what override control exists for. The application measures four properties of it: whether it works, what it costs, what external reset contributes, and what happens when the selector is wired backwards. The same structure, a protective controller and a quality controller sharing one actuator, protects a compressor near surge in Volume 4, Chapter 2 and the rotor and engine of a combine harvester in Volume 4, Chapter 4, where less feed is the safe direction and the selector is a low select.

> ⧖ HISTORY **Sizing a Thickener by Experiment (1916)**
>
> The first equation for the capacity of a thickener, a mass balance, was published by Mishler in 1912. In 1916 Coe and Clevenger divided a continuous thickener into four zones, from clear liquid at the top to a compression zone at the bottom, and sized the tank from laboratory settling tests. Where water has to be pressed out of the pulp, they held, the tank must keep the compressing pulp long enough to reach the required density. Used with a safety factor, their method was still the most popular in the mining industry in 1993, and the theory of sedimentation that Kynch published in 1952 opened a new field of research [3]. The compression law of Equation (6.1) rests on the same residence time.

The scenarios built on the six states cover a specific list of process dynamics and control principles, namely an integrating inventory whose holdup sets product quality, a slow dominant time constant against fast measurement filters and transport delays, actuator saturation at both ends of the pump and the reagent, override selection between competing controllers, external reset as the anti-windup mechanism of a selector, and a strongly nonlinear gain that steepens toward a constraint [2, 4, 5]. The physical fundamentals are a solids mass balance on the vessel, a compression rate law driven by residence time, a hindered settling velocity set by feed concentration and reagent dose, and a separation of time scales between a bed that turns over in four and a half hours and instruments that respond in minutes [6, 7]. Water recovery is the economic reason the industry cares, because

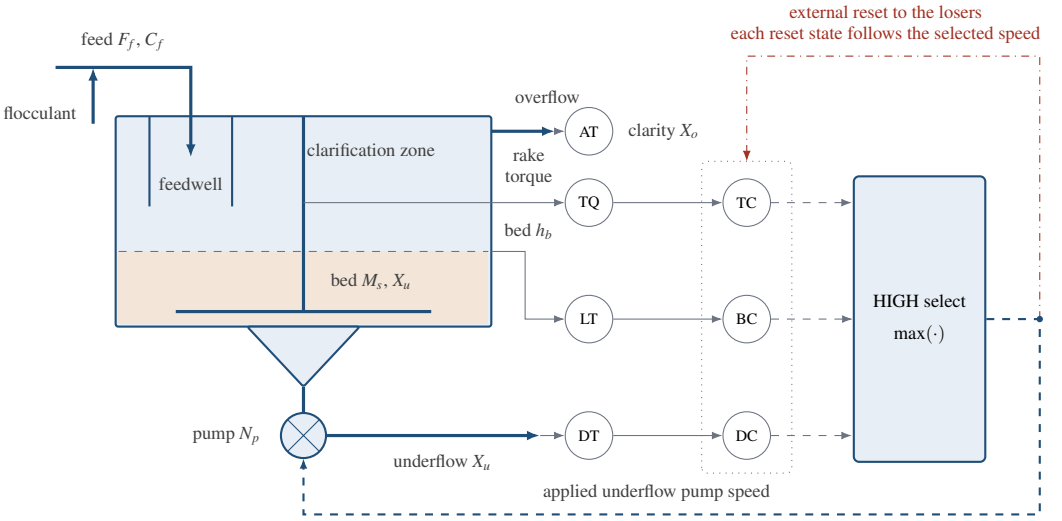


Figure 6.1 The simulated thickener. Only the underflow pump sets the underflow density at steady state, and it does so by setting the bed inventory, which is also what the rake has to drag. Three controllers ask for pump speed at once, the density controller for the plant's benefit and the torque and bed controllers for the machine's, and a high selector passes whichever wants the most pump. The dash-dotted path (red) is the external reset, which feeds the selected speed back to losing loops over their tracking times, and the dashed lines carry the three demands and the applied speed.

a concentrator that loses water to its tailings pays for the replacement, and a tripped rake in a full bed takes days to recover [8]. Section 1.4.3 places the thickener in the circuit of Chapter 1, and Section 1.4.6 states the rule for a selector's direction.

Section 6.2 gives the model: compression, clarity and flocculation, the bound on the inventory, the six state equations and every parameter. Section 6.3 describes the four PID loops, the high selector, its external reset and the optional predictive controller. Section 6.4 reports a pump move by hand, every scenario under the selector and under the MPC, the thickening-feed case in detail and the graded shift. Section 6.5 describes the real instruments, pump and control platform, Section 6.6 states what the model leaves out, and Sections 6.7 and 6.8 draw the lessons.

6.2 Process Model

Figure 6.1 shows the vessel, the three controllers that compete for the underflow pump and the selector between them.

6.2.1 Variables and Signal Roles

Table 6.2 classifies the quantities the operator and the four controllers act on. Six of them are integrated states and the rest are algebraic functions of those states, the parameters and the disturbances.

🕮 **Pause.** Try Exercise 6.2 before reading on.

6.2.2 Design Operating Point

The simulation clock runs in minutes. The vessel is 30 m in diameter, an area of 707 m^2, with the feedwell 3.4 m above the floor. At design it takes 1200 m^3/h of feed at 150 kg/m^3, which is 180 t/h of solids, with the pump at 55 % and the flocculant at 19.7 g/t. That holds 827 t of solids in a 1.597 m bed at a residence time of 4.59 h, discharges 177 m^3/h of underflow at 1017 kg/m^3, which is 61.99 % solids, returns 1023 m^3/h of water at 30.0 mg/L, and loads the rake to 54.9 %. The 19.7 g/t dose is

Table 6.2 Signal roles. Nominal values are the design steady state of Section 6.2.2. The range column gives the setpoint range for a controlled variable and the adjustable range for everything else.

Symbol	Role	Nominal	Range
X_u	CV, state	61.99 % solids	48–70 %
X_o	CV, state	30.0 mg/L	10–200 mg/L
T_q	override CV, state	54.9 %	55–90 %
h_b	override CV	1.597 m	1.00–3.20 m
N_p	MV, state	55.0 %	10–100 %
D_f	MV, state	19.7 g/t	0–70 g/t
M_s	state	827 t	—
F_f	DV	1200 m^3/h	500–2100 m^3/h
C_f	DV	150 kg/m^3	70–280 kg/m^3
σ	DV	1.00	0.35–1.60
s_f	DV	1.00	0.15–1.30
w_p	DV	1.00	0.60–1.10

the value the clarity loop settles at against its 30 mg/L setpoint, and the 100 mg/L recycle limit is therefore a factor of 3.3 above the design overflow.

> **DID YOU KNOW? Twelve Thickeners Behind One Plant**
>
> A 2017 survey of tailings thickeners lists a copper plant treating 32.5 million tonnes of ore a year with twelve paste thickeners, each 24 m across, and another of 14 million tonnes a year with two high-rate thickeners 45 m across [9]. The vessel of this chapter is 30 m across.

6.2.3 Overflow Clarity and Flocculation

The overflow carries whatever does not settle faster than the water rises. The application makes that an exponential in the ratio of the two,

$$v_r = \frac{F_f - F_u}{A}, \qquad X_o = X_{o,0}\exp\!\left(g\,\frac{v_r}{v_s}\right) \tag{6.3}$$

where v_r is the rise rate in m/h, F_f the feed rate, F_u the underflow rate, A the vessel area, X_o the overflow solids concentration, $X_{o,0}$ a clarity scale factor and g the clarity sensitivity. The exponential is flat while the ratio is small and steepens sharply as the rise rate approaches the settling velocity. At design the ratio is 0.257 and the overflow is 30.0 mg/L, and the ratio is clamped to 4 in the implementation so the exponential cannot overflow numerically.

Settling velocity depends on the flocculant through a Langmuir curve, so the reagent saturates,

$$v_s = v_0\,\sigma\,\frac{D_f/(D_f + K_d)}{D_0/(D_0 + K_d)}\,\mathrm{e}^{-k_h C_f} \tag{6.4}$$

where v_0 is the free settling velocity of flocculated solids, σ the ore settleability, D_f the flocculant dose, D_0 the reference dose, K_d the half-saturation dose, k_h the hindered settling coefficient and C_f the feed solids concentration in t/m^3. Table 6.3 gives the resulting operating map, and the graded shift of Section 6.4.4 shows what happens when the reagent runs out of authority.

The saturation is what limits the reagent as a handle. Doubling the dose from 30 to 60 g/t raises the settling velocity from 6.459 to 7.536 m/h, an increase of 17 %, and cuts the overflow from 22 to only 17 mg/L, while the rake torque rises from 54.9 to 56.7 % because the overdose term of (6.1) forces a 1.664 m bed to hold the same underflow density that a 1.597 m bed holds at the design dose. Cutting the dose is the expensive direction, because 5 g/t puts the overflow at 396 mg/L, four times the recycle limit, at a torque that is 0.4 % lower than the design value. Flocculant is therefore a trim rather than a lever.

Table 6.3 Flocculant dose against settling velocity and overflow at design feed. The dose is bounded at 70 g/t.

Dose (g/t)	v_s (m/h)	Overflow (mg/L)	Torque (%)
5	2.660	396	54.5
10	4.110	70	54.9
19.7	5.620	30	54.9
30	6.459	22	54.9
60	7.536	17	56.7

Pause. Try Exercise 6.3 before reading on.

6.2.4 Upper Bound on the Solids Inventory

If the pump cannot remove solids as fast as they arrive, the bed grows. It cannot grow forever, and the mechanism that stops it is the one that matters operationally. Once the bed reaches the feedwell there is no clarification zone left, the feed short-circuits to the launder, and the solids leave over the weir instead of through the pump. The application makes that explicit by blending (6.3) into a solids balance above $h_b = 2.9\,\mathrm{m}$, which both bounds the inventory and gives the steady-state solver a root to bisect. It is also what makes the failure mode in Section 6.4.4 so expensive, because an overflow of tens of thousands of milligrams per liter is not a clarity excursion but the tailings going to the process water tank.

6.2.5 State Equations

The six integrated states are the solids inventory, the underflow density, the pump speed, the flocculant dose, the overflow solids and the rake torque, as given by Equation (6.5).

$$\begin{aligned} \frac{\mathrm{d}M_s}{\mathrm{d}t} &= \frac{1}{60}\left[F_{\mathrm{in}} - \frac{F_u X_u}{1000} - \frac{F_o X_o}{1000}\right] \\ \tau_{\mathrm{comp}}\frac{\mathrm{d}X_u}{\mathrm{d}t} &= X_u^{\mathrm{ss}} - X_u, \qquad \tau_p \frac{\mathrm{d}N_p}{\mathrm{d}t} = N_p^{\mathrm{cmd}} w_p - N_p \\ \tau_f \frac{\mathrm{d}D_f}{\mathrm{d}t} &= D_f^{\mathrm{cmd}} s_f - D_f, \qquad \tau_{\mathrm{clar}}\frac{\mathrm{d}X_o}{\mathrm{d}t} = X_o^{\mathrm{ss}} - X_o \\ \tau_q \frac{\mathrm{d}T_q}{\mathrm{d}t} &= T_q^{\mathrm{ss}} - T_q \end{aligned} \tag{6.5}$$

with the algebraic relations of Equation (6.6).

$$\begin{aligned} F_u &= \min(k_p N_p, F_f), \qquad F_o = \max(0, F_f - F_u) \\ h_b &= \frac{1000\,M_s}{A\,\phi_b X_u} \end{aligned} \tag{6.6}$$

Here t is in minutes, M_s in tonnes, the flows in m^3/h, $F_{\mathrm{in}} = F_f C_f$ in t/h, and the factor of 60 converts the hourly solids terms to the minute clock. The superscript ss denotes the steady-state expressions of (6.1), (6.3) and (6.2), each of which the corresponding state approaches through a first-order lag. The mean bed concentration is a fixed fraction ϕ_b of the underflow concentration, which is what turns an inventory into a depth. Pump wear w_p and flocculant make-up strength s_f are operator-set disturbances that scale the commanded values, and both are unity at design.

6.2.6 Parameters, Initialization and Assumptions

Table 6.4 lists every constant with its value and unit.

Table 6.4 Model parameters and their design values. The pump is bounded to 10–100 % and the flocculant dose to 0–70 g/t.

Symbol	Meaning	Value	Unit
A	vessel area	707	m^2
k_p	pump capacity	3.218	m^3/h per %
ϕ_b	bed to underflow concentration ratio	0.72	—
$h_{b,\max}$	bed depth reaching the feedwell	3.4	m
$X_{\min}$	uncompressed underflow density	400	kg/m^3
ΔX	compression density span	850	kg/m^3
τ_c	compression time constant	3.55	h
τ_{comp}	underflow density lag	25	min
v_0	free settling velocity	12	m/h
k_h	hindered settling coefficient	4.5	m^3/t
K_d	flocculant half-saturation dose	12	g/t
D_0	reference flocculant dose	25	g/t
$X_{o,0}$	clarity scale factor	2.963	mg/L
g	clarity sensitivity to rise rate	9	—
τ_{clar}	clarification zone lag	12	min
T_{q0}	rake torque with no bed	12	%
k_τ	torque per meter of bed	26.875	%/m
n	torque exponent on density	2.2	—
X_{ref}	torque reference density	1017	kg/m^3
τ_q	rake torque lag	4	min
τ_p	pump lag	0.5	min
τ_f	feedwell mixing lag	1.5	min
θ_d	density gauge delay	1.2	min
θ_c	turbidity analyzer delay	2.5	min
τ_{bm}	bed level probe filter	3.0	min
τ_{qm}	rake torque signal filter	3.0	min
—	rake torque alarm and trip	75, 92	%
—	recycle water clarity limit	100	mg/L
—	integration step	0.02	min
—	controller scan	0.2	min

Initialization bisects the solids inventory at which the solids balance closes, which is monotone because a larger inventory means a denser underflow at the same volumetric rate. A manual-mode helper inverts the same balance to return the pump speed that would hold a chosen density.

The bed is one lumped inventory with an empirical compression law rather than a solved settling-velocity partial differential equation, which is the standard modeling framework for these units [6, 10] and is deliberately not used here. The mean bed concentration is assumed to be a fixed fraction of the underflow concentration, the rake torque is assumed algebraic in depth and density, and the four instruments are assumed to be first-order filters or pure delays with no drift. The engine of Section 1.6.1 integrates with fixed-step RK4 at 0.02 min and a 12 s controller scan.

6.3 Controller Design

All four loops run the velocity-form PID of Section 1.6.2, Equations (1.10) to (1.13), with the derivative taken on the measurement [11]. The scan is $\Delta t_c = 0.2$ min. The output is clipped to the actuator range, 10 to 100 % on the pump and 0 to 70 g/t on the flocculant. The bias is zero, and the integral starts at the output in force at $t = 0$. Anti-windup is back-calculation with a tracking time $T_t = \tau_I$, so while the output is clipped the integral moves toward the value that reproduces it by $\Delta t_c / \tau_I$ of the excess each scan ($1/1500$ on the density loop at $\tau_I = 300$ min), as on the blending tank of Volume 2, Chapter 3. The running sum of proportional and derivative increments is never truncated. The tracking time affects the reported results. In a counterfactual run in which the integral absorbs

Table 6.5 Controller tunings. The protection loops are the same in both sets. Derivative time is zero on every loop. The last four rows apply only with the pump on the predictive controller of Section 6.3.3.

Loop	Gain unit	Default	Faster
Density, K_c	% pump per % solids	3	6
Density, τ_I	min	600	300
Clarity, K_c	g/t per mg/L	0.25	0.90
Clarity, τ_I	min	60	15
Rake torque, K_c	% pump per % torque	2.2	
Rake torque, τ_I	min	120	
Bed level, K_c	% pump per m	150	
Bed level, τ_I	min	100	
MPC, T_s and horizon	min	10, 1200	
MPC move price, λ	—	30	
MPC filter intensity, ρ	—	1	
MPC limit margins	% torque, m bed	0.5, 0.02	

the whole excess in one scan, the backwards selector of scenario 11 peaks at 100.9 % rather than 109.9 %, and the recommended configuration of Section 6.4.4 misses its density band. The values measured with the shipped tracking time are in Tables 6.6 and 6.8.

All four loops set $a = -1$, because the process gain is negative in every case. A faster pump gives a thinner underflow, a shallower bed and less torque, and more flocculant gives a cleaner overflow, so a measurement above setpoint has to be answered with more of the manipulated variable. The controller output therefore rises as the measurement rises, which makes all four direct acting in the convention of Section 1.6.2, the one used in the process control literature [2, 11]. The application names them direct acting as well. It is the common sign across the three pump loops that makes a single selector correct for all of them.

Two tunings ship with the application and both appear in the results. Table 6.5 lists them. The default set is deliberately slow, which is what the manual and exploratory scenarios use, and a faster set drives the automatic scenarios of Section 6.4. The integral times follow from the SIMC rules [12] applied to gains measured on the application, and the protection loops keep one tuning throughout.

6.3.1 Override Selector

Three controllers compete for one pump,

$$N_p = \max\left(N_{\text{dens}},\, N_{\text{torque}},\, N_{\text{bed}}\right) \tag{6.7}$$

where N_{dens}, N_{torque} and N_{bed} are the three pump demands. The direction of (6.7) is the design decision that matters [2, 5]. Both protecting controllers relieve their constraint by speeding the pump up, so the safe action is more pump and the selector has to pass the largest demand. A switching margin of 0.75 % of pump speed is applied, so the incumbent keeps the actuator until another demand beats it by that margin, which stops three controllers whose outputs external reset draws toward the same speed from trading the pump back and forth on measurement noise. The application offers the same wiring as a low select in scenario 11, whose results are in Section 6.4.2.

🕮 **Pause.** Try Exercise 6.1 before reading on.

6.3.2 External Reset

After each loop's ordinary PID update, external reset adds a correction from its locally clamped request u_j to the selected command u_{sel}:

$$I_{j,k+1} = I^{local}_{j,k+1} + \frac{\Delta t_c}{T_{t,j}}\left(u_{sel,k} - u_{j,k}\right). \tag{6.8}$$

Here I_j is the integral term of controller j, the I of Equation (1.12), $I_{j,k+1}^{local}$ is its value after the ordinary update and Δt_c is the controller scan. The local update already includes the loop's own output-clamp back-calculation, so the additional term counts only the selector mismatch. The tracking time is $T_{t,j} = \tau_{I,j}$. A finite tracking time matters: assigning the entire selected command to every losing integral makes its next request jump by its proportional term, and alternating losers can ratchet the pump to its limit. With finite reset feedback, ownership and pump motion follow the controller dynamics rather than that scan artifact. The winner needs no selector correction. A switching margin of 0.75 % suppresses small noisy ownership changes. Turning reset off is a windup experiment. Manual transfer retains the command after protection, including when MPC supplies the competing request.

6.3.3 Model Predictive Controller

A model predictive controller (MPC) can take the density controller's place on the pump. The density controller remains the default in the application and in every scenario except scenario 12, and it is the fallback. Under the MPC the pump command is still the output of the high selector, as shown in Equation (6.9).

$$N_p^{\mathrm{cmd}} = \max\left(N_{\mathrm{MPC}},\, N_{\mathrm{torque}},\, N_{\mathrm{bed}}\right) \tag{6.9}$$

The torque and bed controllers keep their tunings, the switching margin and the external reset of Equation (6.8), so the override structure of Figure 6.1 stays in service above the plan as a backstop, and the density controller tracks the applied speed exactly on every scan so that a return to it is bumpless. Every $T_s = 10\,\mathrm{min}$, or 50 scans, one quadratic program plans the pump over $H = 120$ steps, a horizon of 20 h, in twelve move blocks, short near the present and long toward the end, the first four of one step each and the last of 75 steps. Only the first move is carried out, and it holds until the plan is renewed [13], which is the receding-horizon form of Section 1.6.3, used in industrial practice [14].

The controller's model is not identified. It is the balance of Equations (6.1) to (6.6) restricted to the four states the pump moves, evaluated with the declared values of Table 6.4 and the design feed, as shown in Equation (6.10).

$$\begin{aligned} \frac{\mathrm{d}M_s}{\mathrm{d}t} &= \frac{1}{60}\left[F_{\mathrm{in}}^0 + d_1 - \frac{F_u X_u}{1000} - \frac{F_o X_o}{1000}\right] \\ \tau_{\mathrm{comp}}\frac{\mathrm{d}X_u}{\mathrm{d}t} &= X_u^{\mathrm{ss}}(M_s) + d_2 - X_u \\ \tau_p\frac{\mathrm{d}N_p}{\mathrm{d}t} &= N_p^{\mathrm{cmd}} - N_p, \qquad \tau_q\frac{\mathrm{d}T_q}{\mathrm{d}t} = T_q^{\mathrm{ss}}(h_b, X_u) - T_q \end{aligned} \tag{6.10}$$

Here $F_{\mathrm{in}}^0 = 180\,\mathrm{t/h}$ is the design solids feed of $1200\,\mathrm{m^3/h}$ at $150\,\mathrm{kg/m^3}$, the settleability is one, the dose is the reference dose D_0, the pump delivers its command, and d_1 in t/h and d_2 in $\mathrm{kg/m^3}$ are disturbances the model cannot explain. A change to a plant parameter, the feed or the ore therefore reaches the plant and not the model, and appears to the controller as mismatch. The flocculant is left out, because the delivered dose is the command times a make-up strength s_f that the controller is never told, and the overdose penalty of Equation (6.1), when it occurs, lands in d_2. At the design point the model's steady state equals the plant's at 61.9945 % solids, 54.905 % torque and a 1.5969 m bed, and its steady gains for a pump step from 55 to 56 % are −0.6796 points of solids, −3.3519 points of torque and −0.065 23 m of bed per point of pump, against −0.6796, −3.3517 and −0.0652 measured on the plant. Its slow eigenvalue corresponds to a time constant of 903 min, the bed, and its fast one to 25.7 min, the compression lag.

At each decision the four states are linearized at the present estimate and the speed applied now by central differences and sampled with an exact zero-order hold that carries the affine remainder, so the steep dependence of torque on depth and density in Equation (6.2) enters at the local operating point. The decisions are block moves ΔN_j of the pump command, the first measured from the speed

applied now, and each plan minimizes Equation (6.11).

$$\begin{aligned} \mathscr{J} = &\sum_{k=1}^{H} \left(\frac{\hat{w}_k - w_{\mathrm{SP}}}{0.25\,\%} \right)^2 + \lambda \sum_{j=1}^{12} \left(\frac{\Delta N_j}{1\,\%} \right)^2 \\ &+ \kappa_T \varepsilon_T^2 + \nu_T \varepsilon_T + \kappa_h \varepsilon_h^2 + \nu_h \varepsilon_h \\ \text{s.t. } &\hat{T}_{q,k} \le \bar{T}_k + \varepsilon_T, \quad \hat{h}_{b,k} \le \bar{h}_k + \varepsilon_h, \quad \varepsilon_T, \varepsilon_h \ge 0 \\ &10\,\% \le N_j \le 100\,\% \end{aligned} \tag{6.11}$$

Here $\hat{w}_k$, $\hat{T}_{q,k}$ and $\hat{h}_{b,k}$ are the predicted underflow solids in percent by mass, rake torque and bed depth k steps ahead, w_{SP} is the density setpoint, j counts the blocks, N_j is the command in block j, and ε_T and ε_h are slacks. With $\lambda = 30$, a move of one point of pump at a block boundary costs as much as a density error of 1.37 points held for one step. The limits are $\bar{T} = T_{\mathrm{SP}} - 0.5\,\%$ and $\bar{h} = h_{\mathrm{SP}} - 0.02\,\mathrm{m}$, under the setpoints of the two override controllers and read where those controllers read them, so a plan that settles on a limit leaves the overrides idle. The bed limit is a ceiling because the bed controller only ever asks for more pump. Both limits are soft. Each output carries a single slack over the whole horizon, so each slack is the deepest predicted crossing, priced at $\kappa_T = 10^5$ per $\%^2$ and $\nu_T = 10^4$ per % of torque and $\kappa_h = 10^7$ per m^2 and $\nu_h = 10^6$ per m of bed. A limit the estimate has already crossed is re-entered along a 60 min funnel. The command range is the pump range every loop clamps to. A dual active-set method [15] solves the program, and a plan that is not reported optimal, is not finite or leaves its range is rejected. With one block, density alone and $\lambda = 30$, the solver's first move for a one-point setpoint step is −2.367 033 736 % against −2.367 033 739 % by hand.

A steady-state Kalman filter estimates ten states, the four of Equation (6.10), the density gauge with its 1.2 min delay written as a first-order lag of the same length, the torque and bed signal filters, and the random-walk disturbances d_1, d_2 and a torque offset d_3 added to the predicted torque. Its gain is computed once from the model linearized at the design point, with measurement variances of $(0.25\,\%)^2$, $(1.2\,\%)^2$ and $(0.03\,\mathrm{m})^2$, the noise levels of the graded shift, and process variances per decision of 4ρ $(\mathrm{t/h})^2$, 25ρ $(\mathrm{kg/m^3})^2$ and 0.04ρ $\%^2$ on the three disturbances with $\rho = 1$. Each decision it integrates the nonlinear model over the period with the mean command applied and corrects with the density gauge, the torque signal and the bed probe, the three signals the PID loops read. It runs under both pump options, and at $t = 0$ it anchors on the measurements with the plant taken to be at rest. The estimated disturbances make the plan offset-free [16].

The design first proposed was $T_s = 1\,\mathrm{min}$ with 120 steps. After a one-point pump step at the design point the plant takes 892, 850 and 832 min to reach 63 % of its density, torque and bed change, so a 2 h horizon sees 10.0 to 14.1 % of the change and a 20 h horizon 74.5 to 76.4 %. Before the reset correction, the weights were set once, on scenario 12, from 40 pairs of λ and ρ scored by the mean of the density error, torque excess and bed excess integrals relative to the selector's, admitting only pairs that moved the pump no more than the selector. Twelve were admitted, and the largest λ within 10 % of the best score is $\lambda = 30$ with $\rho = 1$. At $T_s = 1\,\mathrm{min}$ none was admitted, the least travel being 147.3 against 123.1 %, and without that condition the rule picks $\lambda = 30$ with $\rho = 0.1$. Section 6.4.3 reports both periods.

A rejected plan hands the pump to the density controller, which has been tracking it, and control returns at the next accepted plan. Three rejections forced at 300 min in scenario 12 move the pump by 0.0065 % on the hand-over scan, and the MPC is back at 330 min. Every measured run had zero rejected plans. The controller reads the three measured signals, the speed command, the setpoints and its declared model, and nothing of the plant's state, parameters or disturbances. A test that feeds the recorded signals of a plant with fifteen parameters, the feed, the ore, the reagent and the pump changed to a separate controller reproduces all of its 60 estimates and plans. A plan takes 384 to 481 µs in Chrome, 0.77 to 0.96 % of one core at 200 times real time, and the application runs at 59.6 frames per second under both pump options.

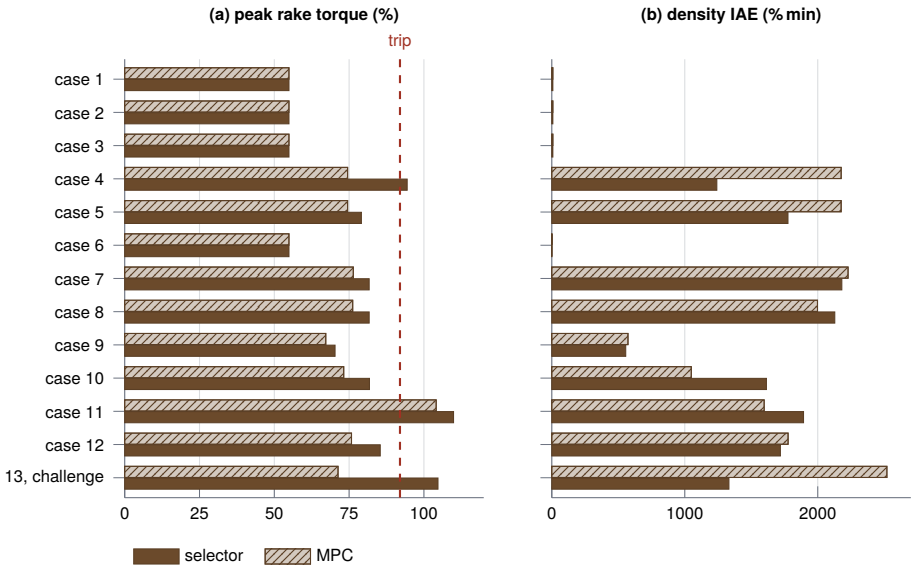


Figure 6.2 The selector and the MPC in every scenario, from Table 6.6: (a) peak rake torque against the 92 % trip (dashed line) and (b) density IAE. Solid bars are the selector and hatched bars the MPC. Notice cases 4, 12 and 13, where the MPC holds the rake lower at the cost of more density error, and case 11, where the wrong selector lets the rake pass its trip under both.

6.4 Results

6.4.1 A Pump Move by Hand

In scenario 1 the thickener runs in Manual for one minute before the pump is slowed from 55 to 51 %. Six hours later the bed has risen from 1.60 to 1.69 m, the density from 61.99 to 62.81 % and the torque from 54.9 to 59.6 %. The bed moves first, then the density and then the torque, and most of the change is still to come. Choosing 51 before Run instead starts the plant at the equilibrium of that speed and shows no response.

> **TRY IT IN THE APP** **1 · Operate the Thickener by Hand (MANUAL)**
>
> Load the scenario, press Run in Manual for one minute, then pause and slow the underflow pump from 55 to 51 %. Run on for six hours and watch the bed depth rise first, then the underflow density and then the rake torque. After six hours the bed has risen only from 1.60 to 1.69 m, and most of the change is still to come.

6.4.2 Selector and Predictive Comparisons

The results that follow use the external reset of Equation (6.8). All comparisons use the same seed, instruments, events and case settings. The selector is the structure each case loads with. Only the pump-control option changes for the MPC comparison. Each case runs for 1440 min, except the 900 min challenge. Density IAE is the integral of absolute true underflow-density error in percent-minutes. Torque is percent of rating. The complete-run measurements are listed in Table 6.6, and Figure 6.2 draws its torque and density columns as bars.

The low selector of case 11 blocks the protective demands when more pump is needed, under either pump option. A correct external reset does not repair a reversed selection. Turning reset off, as in case 10, instead lets a losing controller integrate an error while its command is unused. These are separate experiments from prediction. A protection must retain authority under both pump-control options.

Table 6.6 Selector and MPC in Every Scenario. Each Pair Is Selector / MPC

Case	Density IAE	Torque peak (%)	Bed peak (m)
1	8 / 8	54.9 / 54.9	1.597 / 1.597
2	8 / 8	54.9 / 54.9	1.597 / 1.597
3	8 / 8	54.9 / 54.9	1.597 / 1.597
4	1241 / 2176	94.4 / 74.5	2.338 / 1.966
5	1776 / 2176	79.1 / 74.5	2.055 / 1.966
6	1 / 1	54.9 / 54.9	1.598 / 1.597
7	2181 / 2228	81.7 / 76.4	2.560 / 2.391
8	2128 / 1998	81.7 / 76.2	2.497 / 2.335
9	556 / 573	70.3 / 67.2	1.894 / 1.831
10	1614 / 1049	81.8 / 73.2	2.062 / 1.898
11	1893 / 1597	109.9 / 104.1	2.635 / 2.529
12	1719 / 1777	85.4 / 75.8	2.451 / 2.239
13	1332 / 2520	104.7 / 71.3	3.280 / 2.590

Table 6.7 Case 12 Over 24 Hours

Configuration	IAE	Peak torque (%)	Torque override (min)
Selector	1719	85.45	416
MPC	1777	75.77	0
MPC, plan limits off	1459	81.50	241
MPC, selector off	1777	75.77	0
MPC, both off	1403	88.37	0
MPC, 1 min period	1182	75.42	296

Case 11 is deliberately extreme. Under the low selector the rake passes its trip and stays above it for the rest of a 600 min run. A physical machine would have stopped at the trip long before that, but the case shows what a selector wired in the wrong direction removes, which a real plant reveals only once.

6.4.3 Thickening Feed and the Cost of Protection

In case 12 the feed thickens from 150 to 200 kg/m^3 between 60 and 480 min, at a 64.5 % density request. The finite-reset selector peaks at 85.45 % torque, against 75.77 % for MPC. Over 24 h, density IAE is 1719 for the selector and 1777 for MPC, so the selector has slightly less density error while MPC has the smaller torque peak. Pump travel is 116.3 against 114.1 %. The selector's torque override owns the pump for 416 min, and the nominal MPC needs no torque takeover in this case. This does not establish that the protective selector is redundant: with MPC's plan limits left out, the override acts for 241 min and limits the torque peak to 81.50 %, while removing both plan limits and selector lets it reach 88.37 %. Table 6.7 separates these effects.

The declared MPC model, ten-minute period, twenty-hour horizon, move price and estimator have not been retuned. A one-minute period changes both the horizon and travel, so that row is a counterfactual rather than the shipped law. Compare torque, quality, clarity and movement together.

6.4.4 Graded Fifteen-Hour Shift

The challenge, case 13, runs 900 min through four scripted disturbances. It asks for the worst density error from 594 min within ±1.5 % solids, the rake torque below 92 %, no more than 90 min above 100 mg/L and the bed below the 3.4 m feedwell, with the model parameters and disturbances as given. The case loads with the overrides off, the slow default tuning and a 63 % setpoint. Its worst density error is then 3.02 % solids, twice the band, the torque passes the trip, and the bed rises past the 2.9 m depth at which the feed starts to short-circuit to the launder. The overflow of 81 800 mg/L that follows is stylized, and it is the clearest available statement of what a bed near the feedwell costs. The suggested configuration puts the high selector across torque and bed with external reset, the faster tuning, and 58.5 % density and 30 mg/L clarity setpoints. It meets all four requirements. The flocculant reaches its 70 g/t stop after the finer ore and stays there, so some time above the recycle

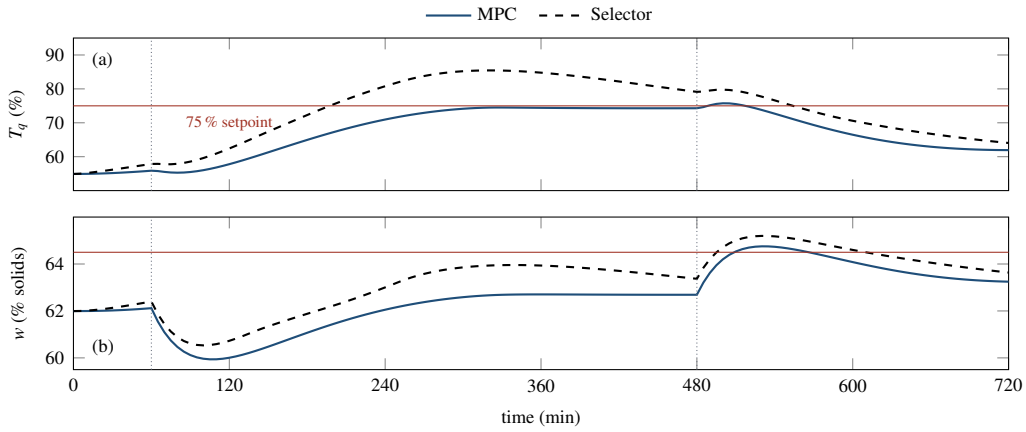


Figure 6.3 Case 12 over its first twelve hours: (a) rake torque and (b) underflow density under the selector (dashed black) and the MPC (solid blue). The thin horizontal lines (red) are the torque and density setpoints, and the dotted vertical lines mark the thick-feed interval. Both structures retain the same protective selector.

Table 6.8 Challenge Measurements

Configuration	Density IAE	Torque (%)	Bed (m)	Dirty (min)
As loaded	1332	104.7	3.280	350.8
As loaded, MPC	2520	71.3	2.590	0.0
Suggested	950	65.9	2.567	42.9
Suggested, MPC	1442	59.9	2.425	79.0

limit cannot be avoided. Table 6.8 gives complete-run measurements, rather than substituting density tracking for the separate torque, bed and recycle requirements.

With the suggested settings the MPC has more density error, 1442 against 950 percent-minutes, and more time above the recycle limit, 79.0 against 42.9 min. Selecting a feasible density setpoint, protecting the bed and maintaining clarity remain necessary, regardless of the pump controller. A transfer to Manual starts at the speed selected after protection and cannot release a protective override.

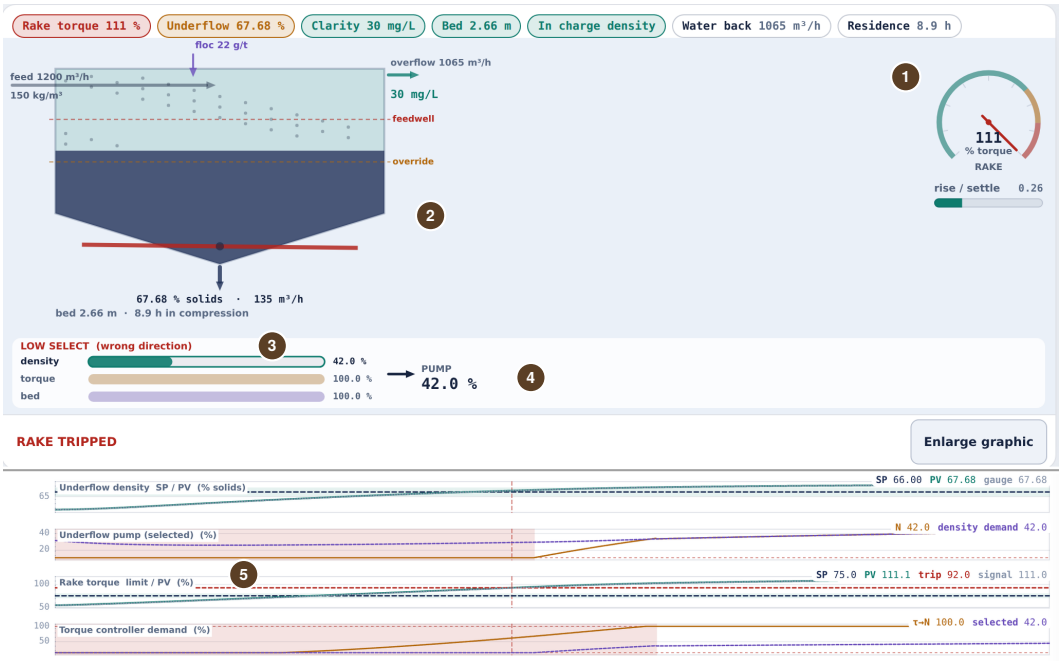


Figure 6.4 The application during scenario 11, under the low selector, with the plant view above the first four trends. (1) The rake gauge and the status badges read the torque past its trip, and the alarm line below the selector reports the rake tripped. (2) The vessel draws the bed, dark, under the clear zone, with dashed lines at the feedwell and at the bed override setpoint. (3) The selector panel names the wiring in force, a low select marked as the wrong direction, and draws the three pump demands. (4) The pump receives the smallest demand, the density controller's, while the torque and bed controllers ask for full speed. (5) The torque trend climbs past its limit and trip lines while the torque controller's demand rises and is never selected.

> ✦ **DID YOU KNOW? The Better Score Tripped the Rake**
>
> From the simulator: in the graded shift as loaded, with the overrides off, the PID run of Table 6.8 has less density error than the predictive controller, 1332 against 2520 percent-minutes. It also drives the rake to 104.7 %, past its trip, and keeps the overflow above the recycle limit for 350.8 min.

6.4.5 Application Access

The application runs in the browser with no installation and all computation is client-side. Manual and automatic modes are provided, and every run exports the full logged trajectory for offline analysis, as described in Section 1.6.4. The three competing pump demands and the selected output are displayed together with the wiring in force, which is what makes the comparison of Table 6.6 visible while it happens. Figure 6.4 shows that view during scenario 11. Manual mode supplies the inverse of the solids balance, returning the pump speed that would hold a chosen underflow density at steady state, so Table 6.1 can be built by hand and the ceiling found before a controller finds it. A second helper, Find the safe density, reports the densest underflow at which the steady torque reaches the torque override setpoint with the present feed.

6.5 Technology in Practice

Figure 6.5 draws the four loops of Section 6.3 with their hardware.

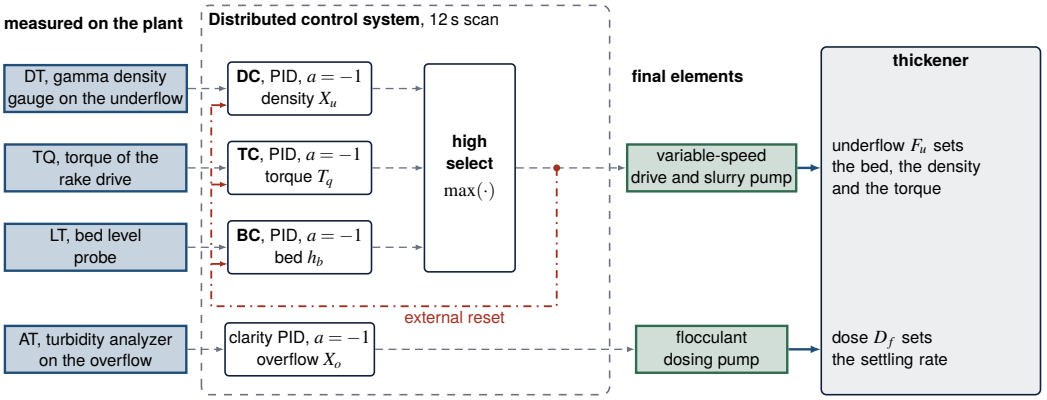


Figure 6.5 The four loops of the simulated thickener drawn with their hardware. The density gauge DT, the rake torque TQ and the bed level LT feed three PID blocks whose pump demands meet in a high select, and the selected speed goes to the drive of the underflow pump and back to all three blocks as their external reset (dash-dotted). The clarity loop on the turbidity analyzer AT owns the flocculant dosing pump alone.

What Measures It

The model reduces the underflow density gauge to a 1.2 min delay with no drift.

> **TECHNOLOGY** **Sensor: A Gamma Density Gauge on the Underflow**
>
> A radiation gauge sends gamma rays from a shielded source through the pipe or vessel to a detector on the far side. The rays are attenuated by the material they cross, so a denser slurry lets fewer reach the detector. A cesium-137 source, with a half-life of 30 years, drifts less in calibration than cobalt-60 at 5.3 years, and the source is made safe for maintenance by closing a lead shutter [17]. Gamma density meters on slurry pipelines are in wide industrial use [18]. The accuracy of radiation gauges varies with coating on the vessel wall, source decay and detector drift [17]. Such errors build slowly, and a density loop holds a wrong reading as faithfully as a right one.

The other three signals are the rake torque and the bed level, each filtered over 3.0 min, and an overflow turbidity analyzer 2.5 min late. A plant study lists rake torque, underflow density, overflow turbidity, the solids interface level and the solids inventory, read as a bed pressure, among the measured variables of a thickener [19], and full-scale profiles of a thickener bed have been taken with a gamma-ray gauge [20].

What Moves It

The model runs the pump from 10 to 100 % of speed through a 0.5 min lag, with the delivered speed scaled by the wear factor, and the flocculant from 0 to 70 g/t through a 1.5 min mixing lag.

TECHNOLOGY Final Element: A Variable-Speed Slurry Pump

The underflow is drawn off by a pump whose motor runs on a variable-frequency drive. Changing the speed of the machine that moves a fluid wastes less energy than throttling its outlet, but a variable-speed machine often cannot change the flow as quickly as a control valve, and it offers few options for a fail-safe state. Pumps in slurry service wear quickly and are rebuilt or replaced at regular intervals [17]. A worn pump delivers less at the same speed, which is the pump wear w_p of Table 6.2, so the speed that holds a given density creeps upward between overhauls.

The same plant study names the flocculant rate and the underflow rate as the two variables normally used to control a thickener, with the feed rate touched only in an emergency [19], the division of handles in Table 6.2.

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 [21]. Programmable logic controllers run PID loops as well [17]. The app runs its four PID blocks on a 0.2 min (12 s) scan, inside that range and short against a bed that turns over in four and a half hours. On such a platform the chapter's structure is three PID blocks feeding a high-select block whose output returns to each block as its external reset, and the predictive controller is a supervisory application that writes a pump request into the same selector every 10 min.

IN PRACTICE Watching the Bed Through Pressure

On one operating thickener, pressure sensors at several heights give the average bed density between each pair of points, and that profile is used to control the bed volume. A high pressure near the top of the wall raises an alarm for an over-expanded bed and, a little higher, drives the flocculant to its maximum. Comparing the bed density in the cone with the density actually pumped shows how well the rake delivers [22]. A pressure tap at the bottom of a settling vessel is kept clear of sediment by a small purge of clean water through a partly open valve [17].

What Is Modernizing It

The first development is predictive control with the rake inside the plan. Published work on paste thickeners has added a rake torque constraint to a predictive controller to prevent rake lifting and pumping problems, and has used a sampling interval that is shorter near the present than far ahead [19], as the move blocks of Section 6.3.3 are. The second is data-driven models: one study fitted regression models of the paste line pressure difference and of the rake torque to a year of one-minute data from 73 measurements [19]. The third is online rheology: an automated capillary rheometer on a side stream reports how the underflow flows without a laboratory sample [23], a property that this chapter's model does not carry.

6.6 Limitations

The bed is one lumped inventory with an exponential compression law. The established description of a continuous thickener is a conservation law with a flux function whose shape produces discontinuities, and the analysis of those solutions is a substantial literature in its own right [6, 10]. Nothing here reproduces a settling front, a rising interface or the discontinuous behavior at the feed level, and the regulator designs that follow from those solutions [24] have no counterpart in this model. The origin of a thickener's steady states is treated properly in that literature and not here.

Compression is a residence-time exponential rather than a compressive yield stress, so the application cannot say what the underflow will be for a particular material, only how it responds to a change in inventory [7].

⚠ OPEN CHALLENGE A Sedimentation Model That Predicts

Published sedimentation models of thickeners are not fully predictive. Each holds for one flocculation state, and none accounts for the shear that raking applies to the aggregates, so they substantially under-predict the underflow density. A twenty-year industrial research program on thickeners found its other sub-models well advanced and named the understanding of sedimentation and dewatering as the main barrier [20]. Progress would be a model that predicts the underflow from the unit's dimensions, its flows, the feed and the flocculation.

Rake torque is an algebraic function of depth and density. Real torque depends on rake geometry, on the yield stress of the bed at the depth the arms run, on whether the arms are raised, and on the presence of islands and rat-holes that form and collapse [25]. The trip in this application is an alarm rather than a mechanism, and the bed that a stalled rake sits in consolidates far more than this model represents.

Flocculant is a single dose with a single Langmuir response. Real flocculation depends on the mixing energy in the feedwell, the polymer molecular weight and charge, the dilution water and the shear history, and the difference between a good and a bad feedwell is larger than the difference between a good and a bad dose. Underflow rheology is absent, so the application cannot represent the transition from a pumpable to an unpumpable underflow that limits real deep cone units.

The predictive model uses the design-feed balance with the reagent at its reference dose and no clarification zone. Feed, ore, reagent strength and pump wear are represented only through three estimated disturbances. Its limits are soft, its ten-minute period delays its response between decisions, and its pump command can change discretely. The protection loops remain active under both pump options, so an apparent advantage of prediction in case 12 cannot be attributed to a delayed reset. The suggested challenge retains density and clarity losses under the MPC, as reported in Table 6.8.

Finally, the plant has one thickener. Real circuits have several in series or in parallel with recycle, and the water balance across them is where the economic argument actually lives [8].

6.7 Discussion

Selector direction determines whether protection can reach the actuator. A correct reset law cannot compensate for the wrong selector direction. The reset law also shapes the comparison with predictive control. Assigning the selected command to every losing integral at each scan ratchets the pump, and the finite reset feedback of Equation (6.8) removes that artifact and changes the quality tradeoff. Comparisons must use the same instruments and protections, and must include density and clarity alongside torque. A quality-only score can prefer a run that violates a mechanical constraint, which motivates separate physical criteria rather than one scalar score [26, 27].

6.8 Conclusion

A thickener makes denser underflow by retaining a heavier bed, so quality and mechanical protection compete for the same pump. The high selector passes protective demands in the safe direction. Its losing reset states follow the selected command, and Manual retains that command after protection. In the thickening-feed comparison, the selector and MPC torque peaks are 85.45 and 75.77 %, with density IAE 1719 and 1777 percent-minutes. Under the suggested challenge settings, MPC instead loses density tracking and clarity. The scenarios therefore distinguish physical protection, prediction and the selection of a feasible operating point without attributing an inherited reset defect to the optimizer.

KEY POINTS

- A thickener makes a denser underflow by holding a heavier bed, so the plant's water and the rake's safety compete for one pump.
- The steady torque reaches the 92 % trip at 66.54 % solids, so a higher setpoint is infeasible.
- Both protections relieve their limits with more pump, so the selector is a high select, and external reset cannot repair a reversed one.
- In case 12 the MPC peaks at 75.77 % torque against 85.45 % for the selector, with density IAE 1777 against 1719.
- A feasible setpoint, bed protection and clarity remain necessary whichever controller drives the pump.

Where else this applies. One actuator shared by a quality loop and a protection in Volume 4, Chapter 2 and Volume 4, Chapter 4.

References

[1] J. V. N. Dorr, "Decanting apparatus," U.S. Patent 867,958, filed February 26, 1907, issued October 15, 1907, 1907.

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

[3] F. Concha and A. Barrientos, "A critical review of thickener design methods," *KONA Powder and Particle Journal*, vol. 11, pp. 79–104, 1993.

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

[5] F. G. Shinskey, *Process Control Systems: Application, Design, and Tuning*, 4th ed. New York: McGraw-Hill, 1996.

[6] G. J. Kynch, "A theory of sedimentation," *Transactions of the Faraday Society*, vol. 48, pp. 166–176, 1952.

[7] S. P. Usher and P. J. Scales, "Steady state thickener modelling from the compressive yield stress and hindered settling function," *Chemical Engineering Journal*, vol. 111, no. 2–3, pp. 253–261, 2005.

[8] B. A. Wills and J. A. Finch, *Wills' Mineral Processing Technology: An Introduction to the Practical Aspects of Ore Treatment and Mineral Recovery*, 8th ed. Oxford: Butterworth-Heinemann, 2015.

[9] B. Pirouz, S. Javadi, and K. Seddon, "Thickener performance variability: underflow solids concentration and flowrate," in *Paste 2017: Proceedings of the 20th International Seminar on Paste and Thickened Tailings*. Beijing: University of Science and Technology Beijing, 2017, pp. 29–40.

[10] R. Bürger, K. H. Karlsen, and J. D. Towers, "A model of continuous sedimentation of flocculated suspensions in clarifier-thickener units," *SIAM Journal on Applied Mathematics*, vol. 65, no. 3, pp. 882–940, 2005.

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

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

[13] D. 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.

[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] 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.

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

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

[18] International Atomic Energy Agency, "Technical data on nucleonic gauges," International Atomic Energy Agency, Vienna, IAEA-TECDOC 1459, 2005.

[19] J. Ruuska, E. Ruhanen, J. Kauppi, S. Kauvosaari, and M. Kosonen, "Multivariate linear regression model of paste thickener," in *Proceedings of the 61st SIMS Conference on Simulation and Modelling*, ser. Linköping Electronic Conference Proceedings, vol. 176, 2021, pp. 160–164.

[20] P. D. Fawell, J. B. Farrow, A. R. Heath, T. V. Nguyen, A. T. Owen, D. Paterson, M. Rudman, P. J. Scales, K. Simic, D. W. Stephens, J. D. Swift, and S. P. Usher, "20 years of AMIRA P266 "Improving thickener technology": how has it changed the understanding of thickener performance?" in *Paste 2009: Proceedings of the Twelfth International Seminar on Paste and Thickened Tailings*. Perth: Australian Centre for Geomechanics, 2009, pp. 59–68.

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

[22] M. Cooks, D. Mostert, L. Coetzee, and D. J. van Schalkwyk, "Thickener experiences at Langer Heinrich mine," in *Paste 2012: Proceedings of the 15th International Seminar on Paste and Thickened Tailings*. Perth: Australian Centre for Geomechanics, 2012, pp. 141–153.

[23] A. G. Chryss, A. Mönch, and K. Constanti-Carey, "Online rheology monitoring of a thickener underflow," in *Paste 2019: Proceedings of the 22nd International Conference on Paste, Thickened and Filtered Tailings*. Perth: Australian Centre for Geomechanics, 2019, pp. 495–504.

[24] S. Diehl, "A regulator for continuous sedimentation in ideal clarifier-thickener units," *Journal of Engineering Mathematics*, vol. 60, no. 3–4, pp. 265–291, 2007.

[25] X. Chen, X. Jin, H. Jiao, Y. Yang, and J. Liu, "Pore connectivity and dewatering mechanism of tailings bed in raking deep-cone thickener process," *Minerals*, vol. 10, no. 4, p. 375, 2020.

[26] D. Hodouin, "Methods for automatic control, observation, and optimization in mineral processing plants," *Journal of Process Control*, vol. 21, no. 2, pp. 211–225, 2011.

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

6.9 Exercises

6.1 Concept check · 1 min

In Figure 6.1 the torque and bed controllers both relieve their limits by speeding the pump up. Which selector belongs between them and the density controller?

(a) A high select, because more pump is the safe direction for both limits

(b) A low select, because the two protecting controllers are direct acting

(c) A high select, but only while external reset is switched off

(d) Either direction, as long as external reset is switched on

6.2 Concept check · 2 min

Which of these quantities does no instrument on the thickener of Figure 6.1 measure, so that every loop sees it only through its effects?

(a) The underflow density X_u

(b) The ore settleability σ

(c) The rake torque T_q

(d) The overflow clarity X_o

6.3 Concept check · 2 min

At design feed, Table 6.3 doubles the flocculant dose from 30 to 60 g/t. Which statement agrees with the table?

(a) The underflow gets denser, because faster settling shortens compression

(b) The overflow falls by more than half, so the clarity loop gains authority

(c) The overflow falls only from 22 to 17 mg/L and the torque rises slightly

(d) Neither the overflow nor the torque moves, because the reagent is saturated

6.4 Theory and calculation · 6 min

At the design point the thickener holds $M_s = 827\,\mathrm{t}$ of solids and takes $1200\,\mathrm{m^3/h}$ of feed at $150\,\mathrm{kg/m^3}$. Use Equations (6.1), (6.6) and (6.2) with the values of Table 6.4, $\sigma = 1$ and no overdose penalty.

(a) Compute the solids feed, the bed residence time to 0.01 h and the compressed underflow concentration X_u^* to 1 kg/m^3. *Checkpoint:* 180 t/h; 4.59 h; 1017 kg/m^3

(b) Compute the bed depth h_b to 0.001 m and the rake torque to 0.1 %. *Checkpoint:* 1.597 m; 54.9 %

(c) An underflow of 66 % solids is $X_u = 1129.3\,\mathrm{kg/m^3}$. What residence time gives it, to 0.01 h, and what inventory, to 1 t? *Checkpoint:* 6.93 h; 1247 t

6.5 Theory and calculation · 5 min

Use Equation (6.2) and Table 6.1.

(a) At 66 % solids the bed is 2.170 m deep and $X_u = 1129.3\,\mathrm{kg/m^3}$. Compute the torque to 0.1 %. *Checkpoint:* 85.4 %

(b) Compute the torque cost per point of solids from 55 to 60 % and from 66 to 68 %, to 0.1, and their ratio to one decimal. *Checkpoint:* 2.9 and 16.1; 5.5

(c) Interpolate linearly between the 66 and 68 % rows for the density at which the torque reaches the 92 % trip, to 0.01 %. Is the estimate above or below the model's 66.54 %, and is that the safe side? *Checkpoint:* 66.41 %; below, the safe side

6.6 Theory and calculation · 6 min

The torque loop loses a high selector to a density demand of 55 % pump speed. Its measured torque is 76 %, setpoint 75 %, proportional gain 2.2, integral time and tracking time 120 min, scan 0.2 min, zero bias and no derivative. The integral carried into this scan is 50 %. Neither local output limit binds.

(a) Find the torque loop's emitted request on this scan and say whether it beats the incumbent by the 0.75 % switching margin. *Checkpoint:* 52.2 %; no.

(b) Include ordinary integration and the selector back-calculation of Equation (6.8). Find the next integral and, with unchanged measurements, the next request. *Checkpoint:* 50.00833 and 52.208 33 %.

(c) Explain why assigning $I = 55$ after every losing scan would give a different handover, and why a constant one-point torque violation can eventually take the pump with finite reset feedback. *Checkpoint:* Instant assignment jumps the next request to 57.2; finite feedback moves the integral toward the selection over its tracking time.

6.7 App challenge · 8 min

Open scenario 1, "Operate the thickener by hand (MANUAL)", at the default seed and a speed of 60 (60 min per second).

(a) Run for one minute, pause, set Underflow pump speed to 51.0 % and run again. Pause when the clock reads 361 min (6.02 h). Read the Underflow and Bed residence tiles, then export the CSV and read `bed_m` and `tq_pct` in its last row. Tolerances: 0.02 %, 0.01 h, 0.002 m and 0.1 %. *Checkpoint:* 62.81 %, 4.97 h; 1.691 m, 59.6 %

(b) Reload the scenario with the reload button, set the pump to 51.0 % before Run, run 30 min and read the same four values. *Checkpoint:* 64.87 %, 6.08 h; 1.960 m, 74.2 %, unchanged from $t = 0$

(c) Run (b) is the steady state at 51 %. What fraction of the steady change in density had arrived in (a) after six hours, to 0.01? Compare it with $1 - \mathrm{e}^{-6/15}$ for a 15 h time constant and explain the difference. *Checkpoint:* 0.28 against 0.33

6.8 App challenge · 8 min

Open scenario 11, "The wrong selector", at the default seed and a speed of 200 (200 min per second). The CSV time column of these runs is in hours.

(a) Run it as loaded, with Override selector on Low select, to 10 h and export the CSV. Find the first row with `tq_pct` at or above 92, within 0.02 h, the largest `tq_pct`, within 0.2, and the controller `override_who` names and the `pump_cmd_pct` during the first three hours. *Checkpoint:* 3.68 h; 109.9 %; bed, 10.0 %

(b) Press Reset, set Override selector to Torque + bed before Run, and run to 10 h. Read the largest `tq_pct`, within 0.2, and the first row in which `override_who` reads torque, within 0.02 h. *Checkpoint:* 79.1 %; 4.19 h

(c) Press Reset, set Override selector to None and run to 10 h. Read `tq_pct` in the row at 10.000 h, within 0.2. Rank the three wirings by peak torque and explain why the low select is worse than none. *Checkpoint:* 91.2 %; low, none, high

6.9 App challenge · 8 min

Open the challenge, "Challenge · sit on the torque limit for a shift", at the default seed and a speed of 200, and run each part for the full 900 min. Read the result line of the progress checklist, which gives the worst density error, the peak torque, the minutes above 100 mg/L and the largest bed.

(a) Run it as loaded. Read the four values and say which requirements fail. *Checkpoint:* 3.02 %, 104.7 %, 351 min, 3.28 m; density, torque and clarity fail

(b) Reload, press Reveal suggested tuning and run again. Read the four values. *Checkpoint:* 1.30 %, 65.9 %, 43 min, 2.57 m; passes

(c) Reload, press Reveal suggested tuning, set External reset to Off before Run and run again. Read the four values, say which requirements fail, and explain how the missing reset causes it. *Checkpoint:* 1.67 %, 71.5 %, 94 min, 2.72 m; density and clarity fail

6.10 Application area · 4 min

Table 6.1 runs the pump at 54.99 % for 62 % solids and at 49.52 % for 66 %. The pump capacity is $k_p = 3.218\,\mathrm{m^3/h}$ per %, and the application values recovered water at $0.45 per cubic meter.

(a) Compute the underflow at both densities and the extra water recovered at 66 %, to 0.1 $\mathrm{m^3/h}$, and its value per hour, to the cent. *Checkpoint:* 177.0 $\mathrm{m^3/h}$ and 159.4 $\mathrm{m^3/h}$; 17.6 $\mathrm{m^3/h}$; $7.92 per hour

(b) Find the safe density reports 64.96 % solids at 50.9 % pump for the 75 % torque setpoint. What share of the gain in (a) does that point keep, to 1 %, and what does the plant buy with the rest? *Checkpoint:* 75 %

6.11 Application area · 3 min

The rake trips at 92 % and the torque override ships with a 75 % setpoint. In the graded shift the torque signal carries noise with a standard deviation of 1.2 %.

(a) How far past the setpoint did the torque go under the selector in cases 5 and 12 of Table 6.6? *Checkpoint:* 4.1 and 10.4 points

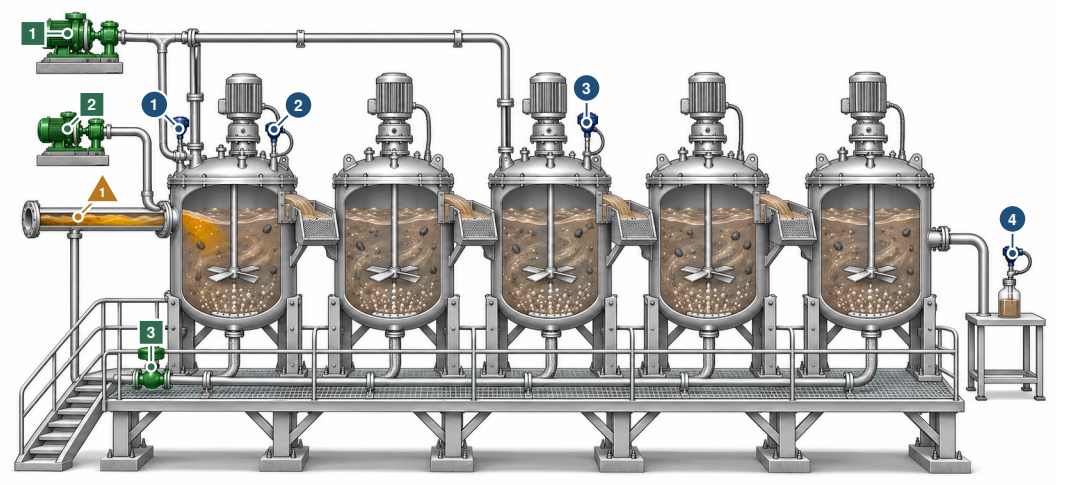
(b) Add three standard deviations of the noise to the larger overshoot. How many points are left below the trip, to one decimal, and why should an engineer not raise the setpoint to 88 % for a denser underflow? *Checkpoint:* 3.0 points

6.12 Review the agent · 5 min

An AI assistant reviewing Table 6.7 writes: "With the selector off the MPC peaks at 75.77 % torque, against 75.77 % with it, and the torque override held the pump for zero minutes. The overrides add nothing once the MPC is in charge, so remove them and save the commissioning effort." Find the flaw and say how to check it on the app.

CHAPTER 7

Soft Sensing and Inferential Control in a Gold Carbon-in-Leach Simulator



The controller moves 1 cyanide dosing pump, to tanks 1 and 3 2 lime dosing pump 3 air valve
The controller reads 1 free cyanide analyzer 2 pH probe 3 dissolved oxygen probe
4 tailings sample for the laboratory

What disturbs it 1 changes in the ore: grade, leaching rate, copper, sulfides

The gold left in the slurry leaving the last tank decides what the circuit earns, and the laboratory reports it hours after the slurry has gone. Illustration, not to scale.

IN PLAIN TERMS **Steering by a Report That Comes Late**

The main idea. Gold is taken out of crushed rock by stirring it with a weak cyanide solution in a row of five tanks. Grains of charcoal in the tanks soak up the dissolved gold, and more cyanide frees more gold but costs more.

The challenge. The plant is paid on the gold left in the waste, which a laboratory reports hours late. A quick estimate from a computer model helps, but it misses a change in the ore until the laboratory corrects it.

Why it matters. Too little cyanide throws gold away, and too much wastes reagent that then has to be destroyed.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history and fact boxes of Section 7.2, the history box of Section 7.3, the try-it boxes of Sections 7.6.2 and 7.6.4, the fact box of Section 7.6.5, and the key points
Study	70 min	The whole chapter with its pause points and the twelve exercises
Practice	30 min	The first page, Section 7.6, Section 7.7, Section 7.8, and the open challenges after it and in Section 7.9

The control question A carbon-in-leach circuit is paid on tailings gold that a laboratory reports six to ten hours late, on a train that takes most of a day to answer. Should an outer loop move the cyanide setpoint, and on the assay or on a soft sensor the assay corrects?

Learning objectives

- Explain why a loop on a late, sampled assay acts in steps, and when a soft sensor is faster.
- Compute the time constant and the noise of a bias update, and choose its gain.
- Show why holding the head-tank setpoints does not hold the tailings.
- Judge from a measured week whether a supervisory loop beats a fixed setpoint.
- Separate what an economic predictive controller gains from its target and from its prediction.

Prerequisites Sections 1.4.4, 1.4.5, 1.4.6 and 1.6.2, and Section 1.6.3 for the predictive controller and the estimator. Chapter 5 has an analyzer delay and an economic setpoint on a faster unit.

The app https://apmonitor.com/apps/gold-cil/

See also Volume 5, Chapter 10, Volume 5, Chapter 4 and Volume 5, Chapter 3 for late quality measurements. Volume 9, Chapter 10 and Volume 9, Chapter 4 for supervisory setpoints and offset removal.

7.1 Introduction

A carbon-in-leach circuit dissolves gold from ground ore with cyanide and adsorbs it onto activated carbon in the same tanks. The reagent decision is a cyanide addition rate. The variable the circuit is paid on is the gold left in the tailings, and the interval between the two is longer than in any other application in this series.

The plant contributes about half of that interval. A step in the sodium cyanide addition reaches sixty-three percent of its effect on the free cyanide concentration of the first tank in 2.68 h and on the tailings grade in 11.52 h, with ninety-five percent of the tailings response taking 20.46 h, because five stirred tanks in series each hold about 3.14 h of slurry and each is a first-order lag on the one before it.

The measurement contributes the remainder. Tailings grade is a fire assay taken every four hours and reported six hours later, so the most recent available number describes the circuit as it was between six and ten hours earlier (Section 1.4.5 explains the age of a sampled assay). No continuous instrument for that variable exists.

A controller closed on that measurement is not poorly tuned. It is structurally slow, and the standard industrial remedy is an inferential estimate. This chapter works with a simulated carbon-in-leach circuit that supplies such an estimate, shows three distinct ways in which it fails, and measures whether closing an outer loop around it improves operation at all.

The circuit shows process order through five first-order lags in series, dead time and sampled measurement through the laboratory assay, cascade structure and time-scale separation between the secondary and outer loops, inferential control through the soft sensor, discrete bias correction as a filter with its own dynamics, and actuator travel as a cost that tuning must account for. Section 7.2 places the problem in mineral processing practice. Section 7.3 gives the five-tank model, the soft sensor and its bias update, and Section 7.4 ranks the three reagent setpoints at steady state. Section 7.5 describes the secondary loops, the outer loop, the predictive controller and the moving horizon estimator. Section 7.6 closes the outer loop on the assay and on the soft sensor (Section 7.6.1), meets unmeasured ore changes (Section 7.6.2), varies the bias update gain (Section 7.6.3), passes copper and sulfide parcels (Section 7.6.4), runs the graded week (Section 7.6.5), and puts the predictive controller (Section 7.6.6) and the estimator (Section 7.6.7) in place of the loop and the update. Section 7.6.5 holds the central result, that on a week of difficult ore the recommended outer loop saves little reagent and costs most of a day above the tailings limit. Section 7.7 names the hardware of a real circuit, Sections 7.8 to 7.10 state the model's limits, what generalizes and what was found, and Section 7.11 holds the exercises.

7.2 Background

Cyanidation followed by adsorption onto activated carbon is the dominant hydrometallurgical route for gold recovery, and carbon-in-leach places the carbon in the same tanks as the leach so that dissolution and adsorption proceed together [1, 2]. Gold dissolution requires cyanide and dissolved oxygen together, and the alkalinity that keeps the cyanide in its ionic form is maintained with lime. Three reagents therefore set the leaching rate, and all three are measurable in the head tank on continuous instruments.

> ⧖ HISTORY **A Dilute Cyanide Solution (1887)**
>
> In 1887 John Stewart MacArthur and Robert and William Forrest, of Glasgow, patented in Britain a process that dissolves gold from powdered ore with a solution of a cyanide. Kept extremely dilute, the solution took the gold and silver in preference to the base metals, and stirring the ore with it in an agitated vessel sped the process. Their United States patent was granted in May 1889. A second patent of the same year dealt with ores holding partly oxidized pyrites, which absorbed much of the cyanide. The ore was first neutralized with an alkali or an alkaline earth, then leached, and the gold was precipitated from the solution on zinc [3, 4]. The lime addition of this chapter does the same work.

The variable that carries the economics is not measurable on those instruments. Mineral processing plants routinely regulate the fast intermediate variables while the grade or recovery that determines revenue arrives from a laboratory hours later, and the resulting mismatch between the control layer and the economic layer is a long-standing problem in the field [5]. Where an instrument is missing, the choices are to add one, to accept the delay, or to infer the variable from what is already measured, and the economics of that choice have been posed as a sensor placement problem [6].

> ✦ DID YOU KNOW? **The Cyanide That Was Not Measured (1995)**
>
> In 1995 a change to copper-rich ore at a mine in Australia raised the cyanide bound to metals, the weak-acid-dissociable cyanide, to 350 parts per million in its tailings dam, and about 2700 birds died. The mine measured only free cyanide, so the rise went undetected [7].

Inferential control takes the third route. A model computes an estimate of the unmeasured controlled variable from measurements that are available now, the controller closes on the estimate, and the delayed laboratory result is used only to correct the estimate's offset [8, 9]. Section 1.4.6 describes the structure as practiced, and Volume 5, Chapter 10 applies it to a dryer whose product moisture is inferred from a soft sensor that a sampled and delayed laboratory corrects. The structure is a cascade in the usual sense. The secondary loops reject fast disturbances in variables that respond in hours, and the outer loop trims their setpoints against a variable that responds in half a day [8, 10]. Cascade control earns its keep when the secondary loop is several times faster than the primary one, and the ratio of about four to one between the cyanide response and the tailings response places this circuit inside that condition.

Dead time is what makes the alternative unattractive. A proportional-integral controller closed directly on a measurement whose dead time is comparable to the dominant plant time constant must be detuned severely to remain stable, and the detuning removes the disturbance rejection the loop was installed to provide [11, 12]. Tuning rules that account for dead time explicitly, such as the internal model control family and the SIMC rules, make the trade-off visible in the gain itself [13, 14].

The physical fundamentals underneath the simulator are conservation laws applied to each tank. Component mass balances on the unleached gold in the solids, the dissolved gold, the gold loaded on carbon, the free cyanide and the alkalinity give five first-order ordinary differential equations per tank, with the flow terms setting the residence time and the reaction terms setting the driving force. Adsorption is driven by departure from a Freundlich equilibrium loading, and dissolution is driven by the unleached gold remaining above a refractory floor. Dissolved oxygen equilibrates in minutes

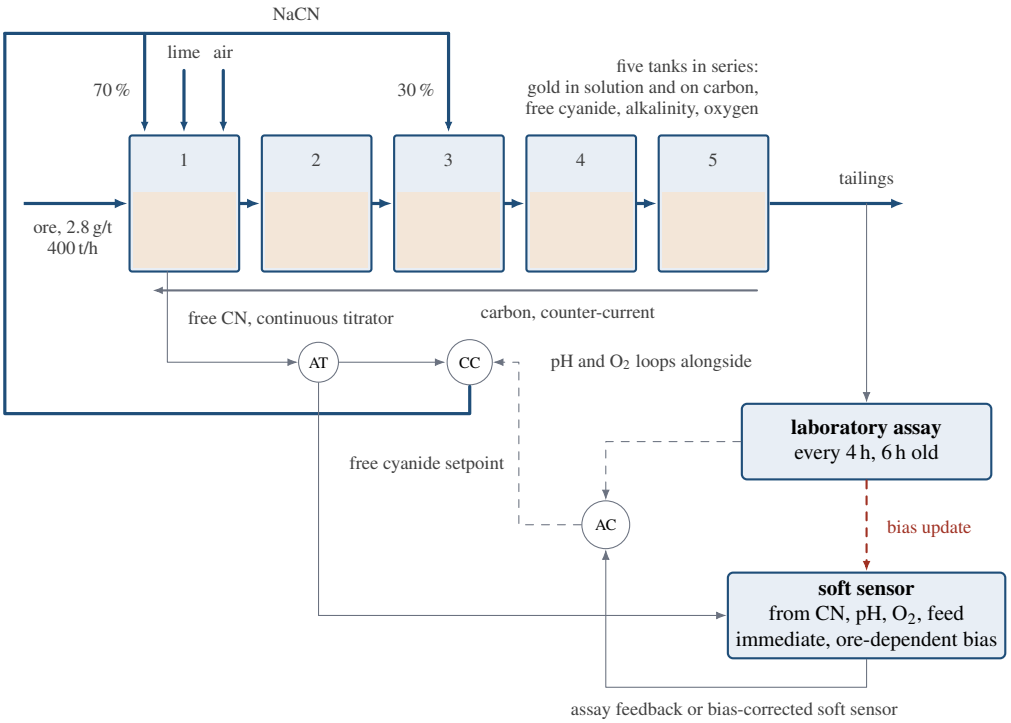


Figure 7.1 The simulated carbon-in-leach train. Three secondary loops hold free cyanide on the sodium cyanide addition, head-tank pH on the lime addition and dissolved oxygen on the air rate, all on continuous instruments. Seventy percent of the sodium cyanide enters tank one and thirty percent enters tank three. An outer controller (AC) moves the free cyanide setpoint against a tailings grade that is not measured, taking it either from the laboratory assay, which is sampled every four hours and reported six hours later, or from a soft sensor that is available immediately and carries a structural error whenever the ore changes. The dashed path from the laboratory assay to the soft sensor, drawn red, is the bias update that reconciles the two.

against tank residence times of hours, so time-scale separation permits it to be treated algebraically rather than as an additional state.

7.3 Process Model

Figure 7.1 shows the train, the three secondary loops and the two competing outer architectures.

7.3.1 Design Point

The clock is in hours. Five tanks of 2000 m^3 treat 400 t/h of ore at 2.8 g/t gold, giving a slurry residence time of 3.14 h per tank and 15.7 h over the train. At design the addition rates are 198.1 kg/h of sodium cyanide, 66.3 kg/h of lime and 100 % air, which hold 120 mg/L of free cyanide, pH 10.40 and 6.0 mg/L of dissolved oxygen in the head tank. The unleached gold in the solids falls from 1.099 g/t in tank one to 0.045 g/t in tank five and the dissolved gold from 0.031 mg/L to 0.0001 mg/L, so the tailings assay 0.295 g/t at 89.5 % recovery, and the loaded carbon leaves tank one at 4058 g/t. Sodium cyanide consumption at the design point is 0.495 kg/t. Seventy percent of the sodium cyanide is added to tank one and thirty percent to tank three, and all of the lime enters tank one, so the free cyanide falls from 120 mg/L in tank one to 89.7 mg/L in tank five and the pH from 10.40 to 10.05.

7.3.2 Leach and Adsorption Kinetics

Gold dissolution is first order in the unleached leachable gold with a rate constant that depends on all three reagents [15], and the cyanide and oxygen terms are saturating, as shown in Equation (7.1). The leach train of Section 1.4.4 has the same first-order dissolution at a fixed rate, and Equation (1.6) gives its residue.

$$r = k \frac{[CN^-]}{K_{CN} + [CN^-]} \cdot \frac{[O_2]}{K_O + [O_2]} \cdot f(\mathrm{pH}) \ (Au - Au_\infty) \tag{7.1}$$

where r is the specific dissolution rate in g/t per hour, k is the intrinsic rate constant of 0.947/h, K_{CN} and K_O are the half-rate concentrations of 60 mg/L and 3.0 mg/L, $f(\mathrm{pH})$ is a dimensionless factor with a maximum at pH 10.4, and Au_∞ is the refractory fraction. The refractory fraction is 0.25 g/t of the 2.8 g/t feed and is the floor that no reagent addition reaches. The pH factor falls away on both sides of its maximum, which is what produces the interior optimum of Table 7.1, and the underlying surface chemistry is more involved than either of the saturating factors [16].

Free cyanide is consumed by the dissolving gold, by cyanicides in the ore and by volatilization as hydrogen cyanide, the last of which depends on pH through the fraction of the free cyanide present in the protonated form, as shown in Equation (7.2).

$$f_{HCN} = \frac{1}{1 + 10^{\mathrm{pH} - 9.31}} \tag{7.2}$$

where 9.31 is the pK_a of hydrogen cyanide and f_{HCN} is dimensionless. At the design pH of 10.40 the protonated fraction is 7.5 %, and at the safety floor of pH 9.8 it is 24.5 %, so the pH loop protects the cyanide loop as well as the operators. Soluble copper raises the cyanicide term by more than a factor of two [17] and is the disturbance of Section 7.6.4.

Carbon adsorption is driven by the departure of the loading from a Freundlich equilibrium and proceeds at a finite film transfer rate, counter-current to the slurry, so tank five contacts the freshest carbon and holds the lowest solution grade.

> ⧖ HISTORY **Carbon Joins the Leach (1880 to 1973)**
>
> Wood charcoal was patented by Davis in 1880 to scavenge gold from chlorination solutions, and once cyanidation was established it was patented for cyanide solutions as well. At a mine in Western Australia in 1917 it showed only a very low capacity for gold, and the gold could be recovered only by burning the charcoal, so zinc precipitation prevailed. In 1950 Zadra eluted gold from activated carbon with a hot alkaline cyanide solution and electrowon it, which let the carbon be reused, and the first large carbon-in-pulp plant was installed in South Dakota in 1973 [18]. Carbon-in-leach puts the carbon in the leach tanks themselves and suits ores that leach fast or rob gold from the solution [19].

7.3.3 Tailings Grade Estimation

Three values of the tailings grade are computed and reported separately.

The first is the value computed by the model, which is available to the application and to no instrument on a real circuit. The second is the laboratory assay, which is that value sampled every four hours and reported six hours later. The third is the soft sensor, which evaluates (7.1) forward from the measured reagent concentrations and the feed grade and reports immediately. The soft sensor assumes that every tank sits at the head-tank cyanide, pH and oxygen, which the profile of Section 7.3.1 shows to be false, so it is structurally optimistic and its accuracy is set by its assumption about the ore. A bias term reconciles the estimate with the assay, as shown in Equation (7.3).

$$b \leftarrow b + \beta \left(Au_{\mathrm{lab}} - \hat{Au}_{\mathrm{soft}} - b \right) \tag{7.3}$$

where b is the bias in g/t, Au_{lab} is the assay just reported, $\hat{Au}_{\mathrm{soft}}$ is the soft sensor output recorded at the time the sample was cut rather than the current one, and β is a dimensionless operator setting.

Table 7.1 Steady state against each secondary setpoint, with the other two holding. 0.25 g/t of the feed is refractory and no reagent reaches it.

Setpoint	Value	Tailings (g/t)	Recovery (%)	NaCN (kg/t)	Margin ($/h)
Free cyanide	60 mg/L	0.3503	87.49	0.255	76007
	120 mg/L	0.2950	89.46	0.495	77423
	200 mg/L	0.2789	90.04	0.813	**77516**
	260 mg/L	0.2742	90.21	1.051	77357
pH	9.60	0.3967	85.83	0.511	74268
	10.40	0.2950	89.46	0.495	77423
	10.55	0.2922	89.57	0.494	**77505**
	11.20	0.3521	87.43	0.487	75604
Oxygen	3.0 mg/L	0.3446	87.69	0.493	75940
	6.0 mg/L	0.2950	89.46	0.495	77423
	7.2 mg/L	0.2883	89.70	0.496	**77502**

It is the update of Equation (1.8) with the weight λ written β, and the app labels it λ, the Soft sensor bias update gain. The update is applied once per assay, so (7.3) is an exponential filter on the assay innovation with a time constant of $-T_s/\ln(1-\beta)$ for a sampling interval T_s of 4 h, which approaches T_s/β for a small gain. The update runs only while the outer loop reads the soft sensor, under the soft sensor loop or the predictive controller of Section 7.5.2. The application supplies a default of $\beta = 0.40$ and the graded week uses 0.35. Sections 7.6.2 to 7.6.5 measure what happens at $\beta = 0$, 0.15, 0.35 and 1.0.

🕮 **Pause.** Try Exercise 7.2 before reading on.

7.3.4 Numerical Implementation and Assumptions

Initialization marches the five tanks forward from the feed at the given addition rates and sweeps the counter-current carbon backward until neither direction moves, so the returned state is an exact equilibrium. Helper routines return the lime addition for a chosen pH and the cyanide addition for a chosen free cyanide concentration, both by bisection.

Gold is one lumped leachable class plus a refractory floor rather than a particle size distribution, carbon is one loading per tank with no regeneration circuit, and cyanide destruction, elution and the gold room are outside the boundary. Integration is fixed-step fourth-order Runge-Kutta at 0.01 h with a 3 min controller scan, as the engine of Section 1.6.1 runs it. The free cyanide titrator reports the head-tank concentration 0.40 h late, the pH electrode is a first-order lag of 0.05 h, and the cyanide pump, lime feeder and air damper are first-order lags of 0.08 h, 0.12 h and 0.06 h.

7.4 Effect of the Reagent Setpoints on Leach Performance

Three reagent setpoints are available to an operator and their steady-state effects differ in shape.

Table 7.1 contains three different shapes. The cyanide response saturates, because raising the setpoint from 200 mg/L to 260 mg/L lowers the tailings grade by 0.0047 g/t and raises the reagent consumption by 0.238 kg/t, so the operating margin has a broad flat maximum near 200 mg/L. The oxygen response saturates more sharply and the blower reaches its 250 % ceiling at 7.2 mg/L. The pH response has an interior optimum at 10.55, because the leach rate falls above that value while the free cyanide is lost as hydrogen cyanide below it.

The pH optimum is shallow and sits close to the design value. It is worth $82 per hour against the design pH of 10.40 and $3237 per hour against pH 9.60, so the penalty for operating low is roughly

forty times the prize for finding the optimum. The lower bound on pH is a safety constraint at 9.8, below which hydrogen cyanide evolves in quantity. The upper bound is economic and no instrument on the circuit reports it.

None of these setpoints is the difficulty. The difficulty is that the tailings grade is the variable the mine is paid on, it responds in half a day, and it is measured four times a shift with a six-hour reporting delay.

7.5 Controller Design

Three secondary loops and one outer loop use the velocity-form PID law of Section 1.6.2, Equations (1.10) to (1.13) [10]. Each scan computes an increment to the reagent addition rather than the rate itself, and the derivative is taken on the measurement, so a setpoint move produces no derivative kick. The controller scan is $\Delta t = 0.05\,\mathrm{h}$, the measurement passes through a first-order filter of 0.2 h for the derivative alone, and only the cyanide loop offers a derivative time, which is zero as supplied. The output is clipped to the actuator span and, where a loop declares one, moved no more than $r\,\Delta t$ per scan by a rate limit of r per hour. Only the outer loop has a rate limit, $r = 8\,\mathrm{mg/L}$ per hour. Anti-windup is back-calculation on the integral alone with a tracking time $T_t = \tau_I$, the last term of Equation (1.13). While the output is clipped or rate limited it drives the integral toward the value that reproduces the emitted output, removing $\Delta t / \tau_I$ of the excess each scan, and the running sum of proportional and derivative increments is never truncated. With $T_t = \tau_I$ and no derivative, the integral relaxes toward the emitted output with time constant τ_I instead of integrating the error while the ramp delivers the proportional demand, so the integral does not wind up behind the rate limit. In manual, and in the outer loop whenever it is switched out, the integral tracks so that the output equals the operator's value exactly, which gives bumpless transfer, and in automatic the manual stations follow the controller outputs after every execution. The secondary loops are tuned by the SIMC rules [14]. Each has the action $a = +1$, reverse acting in the convention of Section 1.6.2, because more reagent raises its own measurement. Free cyanide is controlled by the sodium cyanide addition between 0 and 500 kg/h, head-tank pH by the lime addition between 0 and 600 kg/h and dissolved oxygen by the air rate between 20 % and 250 %, all three on continuous instruments.

7.5.1 Time-Scale Separation Between the Loops

The three secondary loops close on measurements that respond in hours and are available continuously. A step in sodium cyanide reaches sixty-three percent of its effect on the head-tank free cyanide in 2.68 h and ninety-five percent in 7.90 h, and a step in lime reaches sixty-three percent of its effect on head-tank pH in 3.18 h, both measured in the application from a ten percent step in manual. Both responses are close to first order, but the titrator adds 0.40 h of dead time to the cyanide measurement, so the cyanide loop sees its own step reach sixty-three percent only after 3.11 h.

The outer loop closes on a variable whose response to the same actuator takes 11.52 h to sixty-three percent and 20.46 h to ninety-five percent, and which is not measured. The ratio between the two responses is about 4.3 to one on the plant alone before the assay is considered, which places the arrangement inside the condition under which cascade control is normally recommended [8]. Satisfying that condition is not sufficient, and Section 7.6.5 reports why.

The outer controller moves the free cyanide setpoint against the tailings grade in the ordinary cascade arrangement, and four configurations are available: the outer loop switched out, the outer loop closed on the laboratory assay, the outer loop closed on the soft sensor with the bias update of (7.3), and the predictive controller of Section 7.5.2. The app's Outer loop selector names them Local, Assay trim, Soft sensor and MPC. The first three configurations are compared in Section 7.6, and the fourth in Section 7.6.6. The outer controller is direct acting ($a = -1$) in the convention of Section 1.6.2, as the app's hint also names it, because raising the cyanide setpoint raises the leach rate and therefore lowers the tailings grade, so a grade above target calls for more cyanide. Its output

is held between 40 mg/L and 400 mg/L and rate limited to 8 mg/L per hour so that a new assay does not step the secondary setpoint.

7.5.2 Model Predictive Control

The fourth outer-loop setting replaces the outer PI controller with an economic receding-horizon quadratic program [20, 21] that moves the free cyanide setpoint, which the cyanide loop follows as before. Section 1.6.3 describes the predictive option of the apps in general, and this section gives what is particular to this one. The pH and oxygen loops are unchanged. The three configurations above stay the application's default and the starting structure of every scenario but scenario 11, which opens on the predictive controller. While the plan acts, the outer PI controller tracks the setpoint it writes, so a hand-over in either direction carries no bump.

Controller model. The steady part of the model is the soft sensor itself. Its estimate $w(c)$, which is (7.1) evaluated with every tank at the head-tank conditions and with the soft sensor's structural error included, is computed at the measured pH, oxygen, feed grade and throughput from the soft sensor's own constants, which are the plant's declared values, and it is linearized about the measured free cyanide c_0 at every decision by central differences of ±5 mg/L, which give a slope w' and a curvature w''. Ahead of it sits the cyanide loop, identified from setpoint steps of ±10 mg/L at the tuning of scenario 11 with noise off as a first-order lag of 0.48 h behind 0.66 h of dead time (0.475 h and 0.645 h in an independent refit, with a residual of 1.05 % of the response), and carried as one decision period of delay followed by a lag $\tau_{CN} = 0.74$ h, whose mean delay of 1.14 h is close to the refit's 1.12 h. Behind it sit five equal lags, one per tank, at the residue time constant of the soft model's own tank balance, $\tau_r = \tau_{res}/(1+\tau_{res}k_s)$, where $\tau_{res} = 3.14$ h is the slurry residence per tank (the τ_t of Section 1.4.4) and k_s is the soft model's leach rate at the measured point, the rate constant k times the three factors of Equation (7.1), 0.421/h at design, so that $\tau_r = 1.353$ h. These lags reproduce a cyanide step from the soft model's output to the true tailings with a residual of 15.2 %, against 27.7 % for lags at the residence time, and refitting the gain alone, to 1.22, lowers the residual to 3.8 %, so what remains is a gain error. The model is given in Equation (7.4).

$$
\begin{aligned}
d_{k+1} &= u_k \\
\tau_{CN}\,\dot{c} &= d - c \\
\tau_r\,\dot{\ell}_1 &= w(c_0) + w'\,(c - c_0) - \ell_1 \\
\tau_r\,\dot{\ell}_i &= \ell_{i-1} - \ell_i, \qquad i = 2,\ldots,5 \\
\hat{A}u_k &= \ell_{5,k} + b
\end{aligned}
\tag{7.4}
$$

where u_k is the free cyanide setpoint decided at step k in mg/L, d is the setpoint decided one period earlier, c is the measured head-tank free cyanide, ℓ_i are the lags in g/t, b is the bias of (7.3), and $\hat{A}u_k$ is the predicted tailings grade. The continuous part is discretized exactly with the input held over each period and rebuilt at every decision. Between decisions the lags are driven at every controller scan by the soft sensor's output, whichever outer loop acts, so the last lag is the soft estimate delayed through the tanks. The bias is the soft sensor's own, updated by (7.3) with the same gain β and the same residual aligned to the time the sample was cut, which makes the prediction offset-free through an output disturbance model [22] that moves only when an assay is reported. Volume 9, Chapter 4 removes offset instead with a Kalman estimator that carries a fuel bias, a disturbance at the input. A result that arrives between two decisions moves b at once, and the next decision plans with it. The feed grade, throughput, pH and oxygen are held at their measured values over the horizon, because nothing in the circuit announces their future.

The soft sensor's structural error is therefore the controller model's. The steady slope of the true tailings against free cyanide is 1.25 times the model's at design, 1.84 times on the ore of scenario 11 at 200 mg/L and 2.0 times on the ore of scenario 6, because the model holds every tank at the head-tank

cyanide while the cyanide falls along the train. The bias corrects the level of the estimate and not its slope.

Objective and constraints. Every $T_d = 0.4\,\mathrm{h}$, the titrator cycle of eight controller scans, the plan covers $N = 60$ steps (24 h) in ten move blocks of $1, 1, 1, 2, 2, 3, 4, 6, 9$ and 31 steps, and the first move of the minimizer of Equation (7.5) is sent and held until the next decision.

$$\begin{aligned} J = \sum_{k=1}^{N} \left[SP_{Au}\hat{Au}_k + c_{CN}u_k + \tfrac{1}{2}SP_{Au}w''\,(u_k - c_0)^2 \right] T_d \\ + SP_{Au}T_{lag}\hat{Au}_N + s_m \sum_j \Delta u_j^2 + \sum_{k\in\mathscr{K}} \left(\rho_2 \varepsilon_k^2 + \rho_1 \varepsilon_k \right) \end{aligned} \tag{7.5}$$

Here S is the measured ore rate, 400 t/h at design, and P_{Au} is the declared gold price of \$78 per gram, so 0.001 g/t of tailings is worth \$31.2 per hour. The cyanide price c_{CN}, in dollars per hour per mg/L of setpoint, is the declared sodium cyanide price of \$2.40 per kilogram times the measured addition per mg/L in the head tank (the applied addition over the titrator reading, each filtered over 2 h, 1.651 kg/h per mg/L at design), plus the declared destruction price of \$3.50 per kilogram on the fraction 0.7471 of the head-tank cyanide that reaches tank five at the declared defaults, carried by 488.9 m³/h of solution, which totals \$5.24 per hour per mg/L at design. With the curvature term the unconstrained optimum is $u^* = c_0 - (SP_{Au}w' + c_{CN})/(SP_{Au}w'')$, the setpoint at which the next mg/L buys as much gold as it costs. The terminal term, with $T_{lag} = T_d + \tau_{CN} + 5\tau_r$ (7.9 h at design), credits the gold that a late move has still to deliver after the horizon, so that a move held for 24 h is credited 24.2 h of its gold rather than 16.6 h. The move price s_m is \$0.50 per (mg/L)² at each block boundary j.

The tailings target Au_{SP} becomes a soft ceiling, $\hat{Au}_k \le Au_{SP} + \varepsilon_k$, judged at every sixth step ($\mathscr{K} = \{1, 7, \ldots, 55\}$, one point every 2.4 h), each with a slack $\varepsilon_k \ge 0$ of its own priced at $\rho_2 = 10^{10}$ per (g/t)² and $\rho_1 = 10^7$ per g/t. An excess is therefore paid for at the step where it occurs, and an excess that the plant cannot avoid late in the day does not make the same excess free in the near hours. The linear price is 800 times the gold that 1 g/t of tailings over one step is worth at design. Judging every step gives the same net margin to within \$1.7 per hour on every case of Table 7.4 but costs a plan 1.5 ms in Chrome instead of about 0.2 ms. The setpoint is bounded hard below by the outer loop's 40 mg/L and above by the lesser of 400 mg/L and the reach of the pump, its 500 kg/h maximum divided by the measured addition per mg/L, about 303 mg/L at design, above which the head tank cannot follow the setpoint. The first move is bounded hard by the outer loop's own ramp, 3.2 mg/L per decision, and the later blocks are free. The ten inputs and the ten slacks form a quadratic program whose ridge is set to 10^{-6} explicitly, because the solver's default scales with the largest diagonal entry, 2×10^{10} from the ceiling's price, and the resulting ridge of 20 on the absolute setpoint pulls the end of the plan at hour 90 of scenario 11 from 161.8 mg/L to 128.6 mg/L against the economics. The horizon and the move price were set once on scenario 11, where horizons from 8 to 48 h fall within \$4.9 per hour of one another and move prices from 0.1 to 10 times the shipped value within \$1.5, and they are not retuned for any other case.

Acceptance, fallback and cost. The controller accepts a plan when the solver's status is optimal, the arrays have the expected dimensions, every planned setpoint lies within its bounds, the first planned setpoint equals the move sent and lies within the ramp, and the objective value, the slack and the predicted tailings are all finite. Otherwise the soft sensor loop, whose integral has tracked the applied setpoint, takes the setpoint, a fallback mark enters the historian, and the next accepted plan takes it back. The first plan is solved at $t = 0$. Failures forced at six instants of scenario 11 with noise off move the setpoint by at most 0.12 mg/L at the hand-over and by at most the 3.2 mg/L ramp at the take-back. In Chrome a plan takes 211 to 220 µs on average over the 600 plans of scenario 11, 0.055 % of one core at one plant hour per second.

Three models kept apart. The plant is the five-tank model of Section 7.3 with its true leachability, cyanicides, acid and oxygen demands and cyanide profile. The identified models are the cyanide loop's lag and dead time and the five residue lags. The controller's model is (7.4) on the soft sensor's linearization, fed only the measured free cyanide, pH, oxygen, feed grade and throughput, the applied setpoint and addition, the laboratory through the bias, the declared prices and the soft sensor's own constants, held at the plant's declared values, so that an edit of the plant's parameters reaches the controller only through its measurements. Both controllers hold their own copy of seven constants (the rate constant, the two half-rate concentrations, the pH optimum and width, the refractory gold and the ore's specific gravity) at the declared values, so a parameter changed in the simulated plant becomes a model mismatch that the laboratory has to reveal, as it would on an operating circuit. The application's tests confirm that it never reads the leachability, a parcel, the true tailings or the plant state, by source inspection, by seventeen deliberately altered copies of the code that the inspection rejects, and by raising the hidden tank-five residue by thirty percent at an assay instant, after which the bias, the plans and the setpoint stay bit-identical through the next two decisions.

7.5.3 Moving Horizon Estimation of the Bias

The bias b read by the soft sensor loop and by the predictive controller can instead come from a moving horizon estimator (MHE) [23], offered as a second laboratory update beside (7.3) and selected by the app's Lab update switch (Section 1.6.3 introduces the estimator options of the apps). It acts only under those two outer loops, because the other two use no bias. The update of (7.3) stays the application's default and the setting of every scenario but scenario 12, and the soft sensor, the predictive controller and the loops are unchanged, reading the estimate wherever they read the bias of the update. Correcting a leaching soft sensor from its residual against measured data, and choosing how that correction adapts, is an established problem in gold cyanidation [24], and the estimator offers a second answer to it beside the fixed-gain update.

The estimator advances one step per laboratory sample and places each sample at the instant it was cut rather than the instant its assay is reported. A result that arrives 6 h late therefore enters at the step it describes with no delay states, because an estimator that refits a batch of past measurements can place a delayed one where it belongs [25], and the bias in force is the estimate at the newest cut, since a random walk predicted across the six hours in transit stays where it is. Volume 5, Chapter 4 reconciles a residence-time predictor with an analyzer and a laboratory that report late in the same spirit. The model is given in Equation (7.6).

$$x_{k+1} = x_k + (r_{k+1} - r_k) + w_k, \qquad y_k = x_k + v_k \tag{7.6}$$

where x_k is the true tailings grade at cut k in g/t, y_k is the assay of that sample, v_k is the laboratory error, w_k is the change of the bias between consecutive samples, and r_k is the model's tailings grade at the cut, so that $b_k = x_k - r_k$. The reference r_k is not the soft sensor's value at the cut, which (7.3) uses, but that output passed through the five residue lags τ_r of (7.4), the lag ℓ_5 that the predictive controller's estimator integrates at every scan under every outer loop. It is what the model says the tanks had delivered when the sample was cut. Over the last $N = 12$ transitions, which hold thirteen results and span 48 h at the 4 h sampling interval, the estimate solves the quadratic program of Equation (7.7).

$$\begin{aligned} \min_{x_s, w_s, \ldots, w_{k-1}} \quad & \frac{(x_s - \bar{x}_s)^2}{\bar{P}} + \sum_{j=s}^{k-1} \frac{w_j^2}{\sigma_w^2} + \sum_{j=s}^{k} \frac{(y_j - x_j)^2}{\sigma_{lab}^2} \\ \text{subject to (7.6)}, \quad & x_j \geq A u_\infty, \quad \left| w_j \right| \leq 3\sigma_w \end{aligned} \tag{7.7}$$

Here $s = k - N$, and the first term is the arrival cost in its filtering form, which carries every older result through $\bar{x}_s$, the estimator's own earlier estimate moved one step through the model, and $\bar{P}$, the a-priori variance of the Kalman recursion run beside it [23]. The laboratory repeatability $\sigma_{lab} =$

0.015 g/t is declared, five percent of the 0.3 g/t design tailings, and the assay noise that a scenario sets is never read. The drift σ_w follows from a rule fixed before any run, that the steady Kalman gain of the random walk in (7.6) equal the default bias gain $\beta = 0.4$. That requires $\sigma_w^2/\sigma_{lab}^2 = \beta^2/(1-\beta)$, so $\sigma_w = 0.00775$ g/t per sample, and $\bar{P}$ starts at its steady value $\sigma_w^2/\beta = 1.5 \times 10^{-4}$ (g/t)^2, where it stays. The bounds are declared as well. The floor $Au_\infty = 0.25$ g/t is the refractory gold of (7.1), below which no tailings can fall, and $3\sigma_w = 0.0232$ g/t per sample is the three-sigma limit of the drift model. That limit shapes the fitted path rather than capping the move that one result causes, so an assay 0.15 g/t high on scenario 12 moves the estimate by 0.0535 g/t with no bound active. No upper bound is set, since the feed grade moves the tailings and a constant ceiling would never bind.

While no bound is active, a moving horizon estimate with this arrival cost equals the Kalman filter's for any window length, and an independent Kalman filter reproduces every update of scenarios 5 and 9 to 12 to within 4.4×10^{-16} g/t. With these weights the gain is 0.4 from the first sample, and the estimator then reduces to (7.3) at $\beta = 0.4$ with r_k in place of the soft sensor's static value at the cut. Handed that static value as its reference, it reproduces the bias update at $\beta = 0.4$ to within 1.2×10^{-15} g/t over 240 h of scenario 12. What the estimator adds to the update is therefore the lagged reference, the two bounds and the window, and Section 7.6.7 removes each in turn.

An update is used only when the solver reports it optimal and the estimate and the window are finite and on or above the floor. Otherwise the update of (7.3), at the operator's β, moves the bias for that result, a fallback mark enters the historian, and the next accepted update takes the bias back. Taking the bias over, on a switch to the estimator, at the start of a run or after a fallback, is ramped at the drift limit per sampling interval, 2.9×10^{-4} g/t per controller scan. An iteration limit, a non-finite estimate and an estimate below the floor, each forced at the 128 h result of scenario 12 with noise off, are rejected with the update's bias reproduced exactly, and the estimator retakes the bias over 26 scans while the cyanide setpoint moves at most 0.4 mg/L per scan, the outer loop's own ramp. The estimator reads the assay, the time its sample was cut, the model's tailings from its own history and its frozen constants, and never the true tailings, the ore or a plant parameter. Source inspection in the application's tests rejects sixteen deliberately altered copies that break this rule, and an edit of each of 42 plant parameters and 11 disturbance settings at 30 h, under both outer loops, never moves the estimate before it moves a measurement.

Each update is a quadratic program in thirteen unknowns, the first state of the window and twelve drifts, solved by the same dense active-set solver as the plan, once per 4 h of plant time. On the browser's frame loop at one plant hour per second an update takes 227 to 300 µs on average over two measurements, under 0.01 % of one core, against 0.11 % to 0.12 % for the plans of the predictive controller. At the top speed of 2000 plant hours per second the estimator takes 0.8 % to 1.7 % of a core and the plans 24.5 %, a load that predates the estimator and keeps the page below the requested speed with either update. Every timing was taken while other work loaded the host.

7.6 Results

The app presents thirteen scenarios in order: the train walked by hand, the three reagent handles ranked, the three secondary loops, the outer loop on the assay and then on the soft sensor, three ore disturbances, the two failure modes of the soft sensor, the predictive controller, the moving horizon estimator and a graded challenge. Each has a completion rule judged on the run, as Section 1.6.4 describes, and the challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt. Reset restarts a case from $t = 0$ with the reader's settings kept, while the reload button restores the scenario's own. Every result below is a scenario as it loads, with its own seed, events and noise, unless the text names a change.

7.6.1 Outer Loop on the Assay and on the Soft Sensor

Scenarios 4 and 5 apply the same disturbance to the same plant, a feed grade step from 2.8 to 3.2 g/t at 10 h, with the outer loop closed first on the assay and then on the soft sensor, and run for 200 h.

Both configurations hold the tailings grade. The peak is 0.3018 g/t on the assay and 0.2991 g/t on the soft sensor, and both return to the 0.295 g/t target at a recovery of 90.78 %. The soft sensor loop moves the cyanide setpoint past 122 mg/L 0.26 h after the step, because the feed grade is one of its inputs, and the assay loop 10.26 h after it, once the first assay containing the richer ore has been reported. The difference also appears in the actuator. The assay loop moves the cyanide addition through 428 kg/h of cumulative travel over the run and the soft sensor loop through 296, a factor of 1.4, because the assay loop corrects a six-hour-old error in four-hour steps and overshoots its final setpoint of 140 mg/L to 150 mg/L, while the soft sensor loop approaches it without overshoot.

That comparison supports treating an inferential estimate as a measurement and judging it on the same terms as one.

7.6.2 Unmeasured Ore Disturbances

Scenario 6 reduces the ore leachability to 0.82 at 12 h without changing any quantity the soft sensor reads. The estimate is computed from the reagent concentrations and the feed grade, all of which are held on setpoint, so it stays at 0.295 g/t while the true tailings grade reaches 0.320 g/t by 20 h and peaks at 0.3249 g/t at 26 h. The first assay that contains the change is reported at 20 h, and the cyanide setpoint first moves at 20.1 h, as it does in the same case with the outer loop on the assay. For this class of disturbance the soft sensor is no faster than the assay. The loop recovers at a cost of 745 kg/h of cyanide travel by 110 h, when the scenario's next disturbance arrives, two and a half times the travel of the whole feed grade case of Scenario 5.

> **TRY IT IN THE APP** **6 · The Disturbance the Soft Sensor Cannot See**
>
> Load the scenario and press Run with the outer loop on the soft sensor, as it loads. Watch the three tailings values after the harder ore arrives at 12 h. The soft sensor stays at 0.295 g/t while the true grade reaches 0.320 g/t by 20 h, and the cyanide setpoint does not move until the first assay that holds the change is reported.

Scenario 9 introduces the same effect deliberately, by biasing the soft sensor fourteen percent optimistic at 12 h. The estimate and the assay diverge by 0.0403 g/t at that instant, the bias update brings the true tailings grade back within 0.002 g/t of target by 79 h, and the cyanide travel over the 240 h run reaches 2193. The loop returns to setpoint only because the assay remains available to correct it. With $\beta = 0$ the tailings settle at 0.354 g/t, a recovery of 87.4 %, while the estimate, which the loop holds near its target, reads 0.297 g/t at the 140 h recalibration.

A soft sensor removes dead time from the loop and introduces a model error that the loop cannot distinguish from a process change. Whether the exchange is favorable depends on the bias update, and the bias update is limited by the assay interval.

Pause. Try Exercise 7.1 before reading on.

7.6.3 Effect of the Bias Update Gain

Scenario 10 sets $\beta = 1$, so that each assay replaces the bias outright, and applies the feed grade step of Scenario 5 at 12 h with 0.02 g/t of noise on the assay.

Table 7.2 quantifies the failure mode. Each assay is six hours old and carries the laboratory's error, so replacing the bias with it steps the estimate by that error and by whatever the plant did six hours earlier, and the outer controller passes each step on to the cyanide setpoint. With the scenario's noise the setpoint at $\beta = 1$ wanders over 101 mg/L against 53 mg/L at $\beta = 0.15$, the actuator travel is fourteen percent higher than at 0.35 and sixty-three percent higher than at 0.15, and the peak tailings grade is 0.0055 g/t and 0.0069 g/t worse. The reagent consumption does not change. Without assay noise a gain of one still doubles the travel with no measurable change in the tailings, so the dead time alone costs actuator wear and the noise costs control. Over six other noise seeds the travel at $\beta = 1$

Table 7.2 The feed grade step of Scenario 10 under three bias gains over 240 h, with the scenario's assay noise and without it. The setpoint range is taken from 24 h onward.

β	Assay noise (g/t)	CN travel (kg/h)	NaCN (kg/t)	CN setpoint (mg/L)	Peak tailings (g/t)
1.0	0.02	5472	0.512	78–179	0.3158
0.35	0.02	4803	0.519	93–171	0.3103
0.15	0.02	3356	0.517	101–154	0.3089
1.0	0	425	0.568	126–140	0.2990
0.35	0	220	0.568	130–140	0.2986

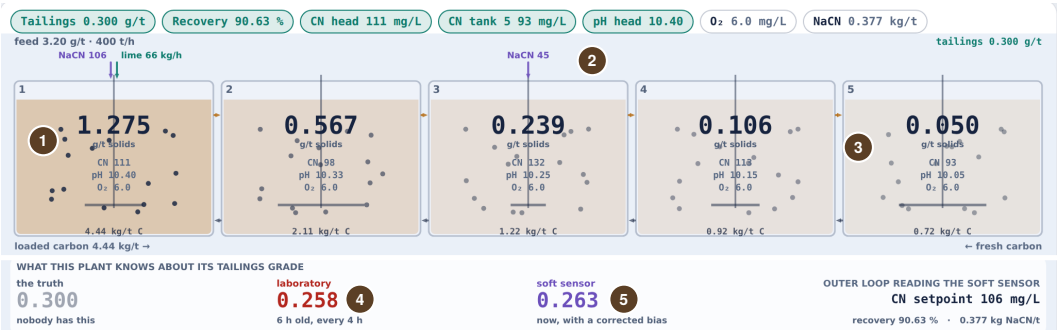


Figure 7.2 Scenario 10 at 170 h, cut to the plant view and the tailings panel, with the bias update replacing the bias outright at every noisy assay. (1) The head tank, where the sodium cyanide and the lime enter and the free cyanide is held. (2) The second cyanide addition, to tank three. (3) Tank five, the last stage of leaching. The three values of the tailings grade are displayed together: the value computed by the model at 0.300 g/t, (4) the laboratory assay at 0.258 g/t reported six hours after the sample was cut, and (5) the soft sensor at 0.263 g/t, 0.037 g/t below the truth. The free cyanide loop is holding its moving setpoint, now 106 mg/L.

exceeds that at 0.35 by 14 to 23 percent and that at 0.15 by 56 to 94 percent, and the peak is higher at $\beta = 1$ than at 0.15 on every seed.

The configuration the application supplies uses a bias gain near a third, which filters the assay across roughly three samples and gives the correction a time constant of about 9 h. Setting the update gain is therefore a filter design problem, and the sampling interval and the assay noise rather than the controller tuning set the answer.

7.6.4 Copper and Sulfide Parcels

Scenario 7 sends a copper parcel through the circuit from 12 h to 120 h, multiplying the cyanicide term by 2.2. The head-tank cyanide loop holds the free cyanide within 10.2 mg/L of its setpoint, but the free cyanide in tank five falls from 89.7 mg/L to 53.2 mg/L by 28 h, so the last stage of leaching slows while every setpoint is met and the tailings grade rises to 0.3084 g/t. The soft sensor loop recovers the tailings by raising the cyanide setpoint to 160 mg/L, which takes the addition to 311 kg/h and the reagent intensity from 0.495 kg/t to 0.777 kg/t. At steady state with the setpoint held at 120 mg/L the same parcel would cost $466 per hour of operating margin. Moving half of the addition to tank three before the run holds tank five at or above 76.3 mg/L through the parcel and the tailings at or below 0.299 g/t. Holding tank one on setpoint says nothing about tank five, which is where the last of the gold has to dissolve.

Table 7.3 Outer loop ablation over the graded week of 200 h. Every row after the first uses the suggested secondary tuning and setpoints (pH 10.5, 6.5 mg/L of oxygen). Above is the time the tailings spend above 0.40 g/t, CN error is the largest departure of the free cyanide from its setpoint, and the last column gives the application's verdict.

Outer loop	Above (h)	Peak (g/t)	NaCN (kg/t)	CN error (mg/L)	Met
As shipped, CN at 120 mg/L	80.9	0.4368	0.529	8.5	no
None, CN at 120 mg/L	**0.0**	0.3982	0.535	6.1	yes
None, CN at 125 mg/L	**0.0**	0.3929	0.557	6.2	yes
Soft sensor, target 0.36	21.6	0.4521	0.525	10.1	yes
Soft sensor, target 0.33	**0.0**	0.3911	0.716	18.2	no
Assay, target 0.36	52.7	0.5232	0.559	9.7	yes
Soft sensor, $\beta = 0$	172.0	0.5887	0.237	9.5	no

TRY IT IN THE APP 7 · A Copper Parcel Eats the Cyanide

Run the scenario and watch the dashed tank-five trace on the cyanide loop. The head tank stays within 10.2 mg/L of its setpoint, yet tank five falls from 89.7 to 53.2 mg/L by 28 h. Then reset, open Show Model, move the share of the cyanide added to the head tank from 0.70 to 0.50 and run again: tank five now stays at or above 76.3 mg/L.

Scenario 8 sends a sulfidic parcel at 12 h, which multiplies the oxygen demand by 2.4 and, 0.1 h later, the acid demand by three. The air command reaches its 250 % ceiling at 13.7 h and the dissolved oxygen settles at 5.32 mg/L against its 6.0 mg/L setpoint. The lime addition rises from 66.3 kg/h to 82.3 kg/h and holds the head tank within 0.04 of pH 10.40, but the pH in tank five falls below 9.8 at 15.0 h and reaches 9.50, where the alkalinity is exhausted, by 17.3 h. The tailings grade reaches 0.334 g/t at 29 h, and the outer loop recovers it only by raising the cyanide setpoint to 367 mg/L, with the sodium cyanide pump at its 500 kg/h limit from 115 h and 1.109 kg/t consumed over the 240 h run. The soft sensor reads the head-tank pH, which is held, and so reads 0.034 g/t optimistic at 28 h. A loop defends the point at which it measures, and where the variable matters along the whole train a single addition point is not a substitute for a tighter loop.

Pause. Try Exercise 7.3 before reading on.

7.6.5 Outer Loop Ablation Over the Graded Week

The graded challenge runs 200 h against harder ore at 20 h (leachability 0.80), a copper parcel at 60 h, an acid parcel at 105 h and a richer feed of 3.3 g/t at 150 h, with noise on the titrator, the pH electrode and the assay, starting on local control with cautious tunings. It is scored on the free cyanide within ±12 mg/L of its setpoint and the head-tank pH at or above 9.8 at every integration step, the tailings below 0.40 g/t immediately before 60 h, 105 h and 150 h and at 200 h, and the cumulative consumption under 1.05 kg/t at the end.

Table 7.3 is the central measurement of this chapter, with the hours above 0.40 g/t and the peak taken at the integration step. With the settings as shipped the circuit fails, because the tailings are 0.409 g/t before the richer feed and 0.433 g/t at the end, while the cyanide (8.5 mg/L), pH (10.32) and reagent (0.529 kg/t) clauses hold. Figure 7.3 draws the hours above the limit and the reagent of every row.

With the suggested secondary tuning, no outer loop and the free cyanide held at its design 120 mg/L, the tailings grade never exceeds 0.40 g/t and the circuit uses 0.535 kg/t of sodium cyanide. The soft sensor cascade at a target of 0.36 g/t also passes as scored, at 0.525 kg/t, two percent less reagent, but it spends 21.6 h above the limit after the harder ore and peaks at 0.4521 g/t, because the

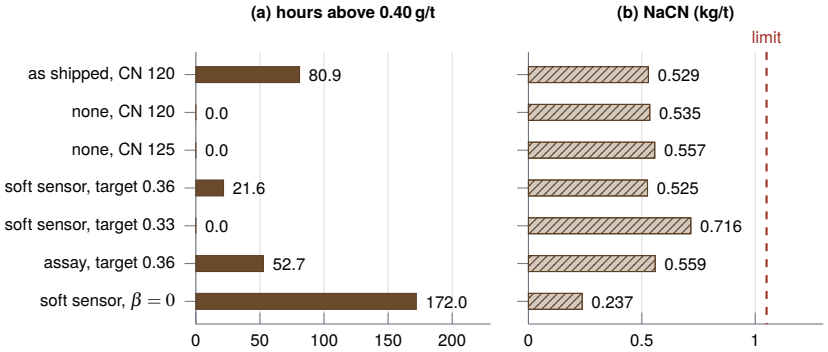


Figure 7.3 The graded week of Table 7.3 as a chart. A fixed cyanide setpoint with no outer loop spends no hours above the tailings limit, while the recommended soft sensor cascade saves a little reagent and spends 21.6 h above it. The loop with no bias update uses the least reagent of any row and the most hours above the limit.

challenge judges the tailings only at its checkpoints. As scored, its worst cyanide error is 10.14 mg/L, the worst checkpoint tailings 0.379 g/t, the lowest pH 10.40 and the consumption 0.525 kg/t, which passes with 1.9 mg/L and 0.021 g/t to spare, and the app gives the same run whether the suggested settings are revealed before Run or carried through Reset. Tightening the target to 0.33 g/t keeps the tailings under the limit throughout and takes the consumption to 0.716 kg/t, a third more than the fixed setpoint, and the richer feed then drives the sodium cyanide pump to its 500 kg/h limit for 7.6 h, so the free cyanide ends 18.2 mg/L below its setpoint and the cyanide clause fails. On the shipped noise seed no target from 0.300 to 0.380 g/t, in steps of 0.005, both keeps the tailings under the limit and passes. Targets up to 0.330 break the cyanide clause, and targets from 0.335 to 0.375 pass with 5.4 h to 43.1 h above the limit. Every row with the suggested tuning moves the cyanide addition through 10 031 kg/h to 12 670 kg/h, most of it in answer to titrator noise.

The comparison depends partly on the noise. Over six other noise seeds the fixed setpoint at 120 mg/L exceeds the limit once, for 5.1 h with a peak of 0.4004 g/t, while 125 mg/L stays at or below 0.395 g/t on every seed at 0.556 kg/t to 0.557 kg/t. The cascade at 0.36 spends 21 h to 38 h above the limit on every seed, and at 0.33 it both holds the limit and passes on four of the six. A fixed setpoint must also stay within 12 mg/L of the 120 mg/L it starts from, because the cyanide clause counts the error that a setpoint step itself creates, as the challenge text states.

The last row shows the behavior of the outer loop when its estimate is not corrected. At $\beta = 0$ the soft sensor remains optimistic through the ore change, the controller acts as though the tailings grade were acceptable, it reduces the cyanide consumption to 0.237 kg/t, and the circuit spends 172 h of the 200 h week above the tailings limit. That row is the argument for the bias update and against reading the reagent number in isolation.

> **+ DID YOU KNOW? The Cheapest Week Was the Worst**
>
> From the simulator: with no bias update, the soft sensor loop of Table 7.3 uses the least reagent of any row, 0.237 kg/t, and keeps the tailings above the 0.40 g/t limit for 172 h of the 200 h week.

The excursion of the recommended cascade has a specific cause. Its target lies 0.065 g/t above the tailings grade at the start, so at $K_c = 2500$ (mg/L)/(g/t) the proportional term alone demands a cyanide setpoint about 160 mg/L below the one the loop inherits, and the rate limit carries the setpoint down to 48 mg/L by 10 h and to its 40 mg/L floor by 18.6 h. The harder ore at 20 h then changes a property the soft sensor cannot see (Section 7.6.2), so the loop learns of it only from the assays, and the tailings pass 0.40 g/t at 25.8 h and return under it at 47.4 h. The setpoint stays low because

the error stays large, not because the integral winds up: at 18.6 h the integral part of the output is 76 mg/L and the proportional part cancels it on a soft sensor reading 0.03 g/t below the target. As counterfactual model runs, shortening the tracking time T_t of Equation (1.13) from 14 h to 2 h leaves 20.4 h above the limit, and only a one-scan tracking time removes the excursion on the shipped seed (peak 0.399 g/t, 0.491 kg/t), by discarding the proportional demand of every rate-limited move. That variant delays the Scenario 5 setpoint past 122 mg/L from 10.26 h to 12.91 h and still exceeds the limit for 22 h to 32 h on four of the six other seeds, so the application retains $T_t = \tau_I$ to preserve proportional response through rate-limited moves. A fixed setpoint at 120 mg/L instead runs the tailings near 0.29 g/t before the harder ore and stays under the limit through all four disturbances, with 0.002 g/t to spare after the richer feed. The value of an outer loop on an operating circuit comes from slow drift in the properties its estimator sees, and this week opens with a step in one that it does not see.

7.6.6 Model Predictive Control Against the Soft Sensor Loop

Table 7.4 runs scenarios 5 to 11 and the challenge, each from its own settings with its own seed, events, noise and duration, under four outer loops. The first is the structure the case ships with, which is the soft sensor loop, or local control on the challenge as shipped. The second is the predictive controller of Section 7.5.2 with the settings of scenario 11 and nothing retuned. The third gives the soft sensor loop the predictive controller's economics without its prediction, because every 0.4 h a setpoint optimizer moves the loop's tailings target to the steady optimum of the same soft model at the same prices, capped at the case's own target. The fourth is local control at the best fixed cyanide setpoint of a scan from 80 to 400 mg/L in steps of 10, chosen after the run and therefore a bound that no causal controller attains. No row receives a preview, and every row is scored on the application's net margin, gold at $78 per gram less sodium cyanide, lime, air and cyanide destruction.

The predictive controller earns more than the shipped structure on every case except scenario 6, from $10.3 per hour on scenario 8 to $188.6 on scenario 10 and $904.1 on the suggested challenge settings, and it keeps the true tailings under the target throughout scenarios 5, 9 and 10, where the loop spends 188.1, 74.3 and 200.1 h above it. The hours over the target favor the plan by construction, because a regulator that holds the target sits on it and exceeds it often by a small margin, at most 0.0041 g/t over 188.1 h on scenario 5, while the plan settles below it, so the depth is the fairer comparison. Almost all of the gain belongs to the economic target rather than to the prediction. The soft sensor loop with an optimized target comes within $12.9 per hour of the predictive controller on scenario 5, $13.9 on scenario 7, $1.6 on scenario 9, $1.9 on scenario 11 and $51.1 on the suggested challenge, and it earns $19.7 more on scenario 6, $2.9 more on scenario 8 and $1.7 more on scenario 10. The rule that a tailings target encodes is too lean in cyanide at the declared prices. After the richer feed of scenario 11 the soft sensor loop holds 0.295 g/t at 139 mg/L, while the plan settles at 162 mg/L with the tailings at 0.289 g/t, within 0.2 mg/L of the unconstrained optimum of 161.6 mg/L evaluated by hand from (7.5) at hour 90.

On scenario 8 the tailings ceiling cannot be met, and every judged step at which it is broken carries its own price. The plan therefore asks for all the cyanide the pump can deliver: the setpoint first reaches 300 mg/L at 44.9 h and stays between 309.8 and 310.5 mg/L from 60 h to the end, equal to the plan's own final setpoint, with the pump at its limit for 67.4 h against 124.9 h for the soft sensor loop, and the true tailings still end at 0.298 g/t. A single slack shared by the whole horizon, the form the controller used before, priced only the worst step there, held the applied setpoint between 267.6 and 275.3 mg/L from 60.1 h while its own final setpoint stood at least 35.5 mg/L higher, and earned $75.3 per hour more, because the cyanide that the ceiling demands buys less gold than it costs.

The fixed setpoint chosen in hindsight beats the predictive controller on every case, by $8.3 to $72.7 per hour, by $132.4 on scenario 8 and by $224.3 on the suggested challenge, where 230 mg/L fails the cyanide clause by its own initial step of 110 mg/L. The best fixed setpoints lie at 170 to 230 mg/L, above the plan's, because the plant's tailings respond to cyanide 1.25 to 2.0 times as strongly as the soft model says (Section 7.5.2), and the bias that corrects the model's level cannot

Table 7.4 Outer Loops Compared on Net Margin, Tailings and Cyanide. Over Is the Time the True Tailings Spend Above the Target in Force (0.36 g/t on the Challenge's Suggested Settings, 0.295 g/t Elsewhere) and Depth the Largest Excess, Discharge the Time Tank Five's Free Cyanide Spends Above 200 mg/L, and Travel the Summed Movement of the Cyanide Setpoint. Fixed Setpoints Are Chosen in Hindsight

Case	Outer loop	Net ($/h)	Over (h)	Depth (g/t)	NaCN (kg/t)	Discharge (h)	Travel (mg/L)	Plans, fallbacks
5, richer feed	Soft sensor (shipped)	89 170.7	188.1	0.0041	0.566	0.0	54	—
200 h	MPC	89 232.9	0.0	0	0.663	0.0	39	500, 0
	Soft sensor, optimized target	89 220.0	0.0	0	0.638	0.0	85	—
	Fixed 180 mg/L	89 241.2	0.0	0	0.739	0.0	0	—
6, harder ore	Soft sensor (shipped)	83 697.9	148.8	0.0299	0.852	29.5	389	—
240 h	MPC	83 679.5	94.3	0.0167	0.942	101.6	806	600, 0
	Soft sensor, optimized target	83 699.2	147.6	0.0221	0.891	29.3	355	—
	Fixed 200 mg/L	83 752.2	145.1	0.0137	0.816	0.0	0	—
7, copper parcel	Soft sensor (shipped)	77 324.5	84.7	0.0134	0.612	0.0	163	—
240 h	MPC	77 382.9	46.1	0.0034	0.688	0.0	415	600, 0
	Soft sensor, optimized target	77 369.0	103.4	0.0066	0.670	0.0	135	—
	Fixed 180 mg/L	77 405.1	0.0	0	0.796	0.0	0	—
8, sulfidic parcel	Soft sensor (shipped)	76 330.2	228.0	0.0390	1.109	141.0	337	—
240 h	MPC	76 340.5	224.1	0.0257	1.168	186.2	190	600, 0
	Soft sensor, optimized target	76 343.4	225.8	0.0333	1.112	139.4	336	—
	Fixed 220 mg/L	76 472.9	220.2	0.0127	0.895	0.0	0	—
9, structural error	Soft sensor (shipped)	77 405.9	74.3	0.0222	0.507	0.0	475	—
240 h	MPC	77 517.2	0.0	0	0.631	0.0	262	600, 0
	Soft sensor, optimized target	77 515.6	0.0	0	0.612	0.0	247	—
	Fixed 170 mg/L	77 537.9	0.0	0	0.695	0.0	0	—
10, $\beta = 1$	Soft sensor (shipped)	89 045.4	200.1	0.0208	0.512	0.0	1562	—
240 h	MPC	89 234.0	0.0	0	0.710	0.0	849	600, 0
	Soft sensor, optimized target	89 235.7	0.0	0	0.698	0.0	849	—
	Fixed 180 mg/L	89 243.0	0.0	0	0.739	0.0	0	—
11, richer feed,	Soft sensor	89 000.1	157.0	0.0245	0.695	0.0	303	—
then harder ore	MPC (shipped)	89 025.8	38.3	0.0168	0.783	28.3	452	600, 0
240 h	Soft sensor, optimized target	89 023.9	69.1	0.0179	0.749	0.0	278	—
	Fixed 190 mg/L	89 054.0	69.1	0.0101	0.778	0.0	0	—
Challenge,	Soft sensor (suggested)	79 144.1	106.0	0.0921	0.525	0.0	1299	—
suggested	MPC	80 048.2	16.0	0.0064	0.686	0.0	344	500, 0
settings,	Soft sensor, optimized target	79 997.1	26.3	0.0102	0.662	0.0	1128	—
200 h	Fixed 230 mg/L	80 272.5	0.0	0	1.012	0.0	0	—
Challenge,	Local (shipped)	78 994.9	180.0	0.1418	0.529	0.0	0	—
shipped	MPC	79 888.6	178.6	0.0590	1.042	0.0	370	500, 0
settings,	Soft sensor, optimized target	79 497.4	179.8	0.0923	0.668	0.0	245	—
200 h	Fixed 260 mg/L	79 907.9	173.9	0.0622	1.078	0.0	0	—

correct its slope. At hour 90 of scenario 11 the plant's own steady optimum is 178 mg/L against the plan's 161.8 mg/L. The soft sensor's structural error, inherited by the controller's model, is therefore what the hindsight setpoint exploits. In a counterfactual run that hands the plan the plant's true steady slope and curvature at every decision, information no operating circuit possesses, the predictive controller comes within $1.5 per hour of the hindsight setpoint on scenario 5, $7.4 on scenario 9, $10.2 on scenario 10 and $22.2 on scenario 11, and passes it on scenario 7 by $3.6 and on the suggested challenge by $14.2. Gaps of $65 and $126 per hour remain on scenarios 6 and 8, because the hindsight setpoint knows the unseen ore and the parcel from the first hour.

Scenario 11 applies the richer feed of scenario 5 at 12 h and a harder ore of leachability 0.85 from 100 h to 170 h, runs 240 h, and opens on the predictive controller, as shown in Figure 7.4. Neither structure sees the harder ore before the laboratory does. The loop first moves at 108.2 h and the

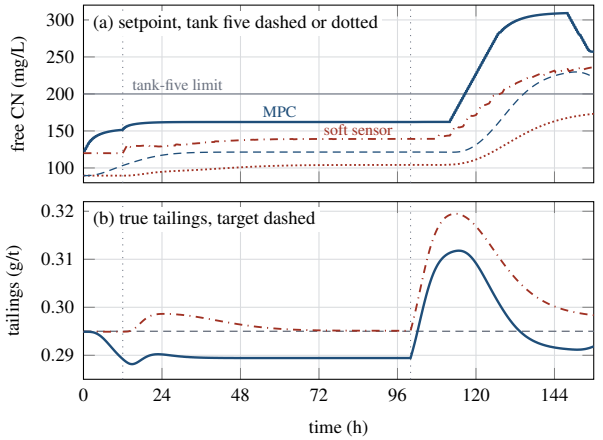


Figure 7.4 Scenario 11 from the start to 156 h under the predictive controller (solid and dashed lines, drawn blue) and the soft sensor loop (dash-dotted and dotted lines, drawn red). (a) The free cyanide setpoint (solid for the predictive controller, dash-dotted for the loop) and tank five's free cyanide (dashed and dotted) against its 200 mg/L discharge limit (thin gray line). (b) The true tailings (solid and dash-dotted) against the 0.295 g/t target (gray dashed). The vertical dotted lines mark the richer feed at 12 h and the harder ore at 100 h. The plan's setpoint peaks at 309 mg/L, the reach of the pump, at 147.7 h, and at 156 h it is back at 257 mg/L with the true tailings at 0.292 g/t, under the target since 133.5 h, while the loop's setpoint is 236 mg/L with the tailings at 0.298 g/t.

plan at 112.1 h, but once the bias grows the ceiling binds and the plan moves faster, past 200 mg/L at 116.5 h against 127.4 h for the loop, and further, to a peak of 309.2 mg/L at 147.7 h, at the reach of the pump, against the loop's 243.5 mg/L at 180.1 h. The true tailings are already back under the target from 133.5 h, while the soft estimate with its bias still reads 0.2953 g/t at 156 h, because the bias carries assays of samples cut six hours earlier and the model's slope on this ore is 1/1.84 of the plant's. Over the 240 h the plan earns \$89,025.8 per hour against \$89,000.1 on 0.783 kg/t of sodium cyanide against 0.695 kg/t, and within the harder ore, from 100 to 170 h, it earns \$89,158.3 against \$89,153.4 while the optimized target earns \$89,197.8.

The plan's price cannot be had by lowering the loop's target once. With the soft sensor loop's target lowered by hand to the 0.289 g/t at which the plan settles, scenario 11 earns \$89,003.5 per hour over the 240 h on 0.816 kg/t of sodium cyanide, against the plan's 0.783 kg/t, and so recovers \$3.4 of the plan's \$25.7. Local control at fixed setpoints of 160, 190 and 220 mg/L earns \$89,029.1, \$89,054.0 and \$89,023.1. The gain requires a target that moves with the conditions, as the optimized target of Table 7.4 does at \$89,023.9.

The predictive controller loses scenario 6 by \$18.4 per hour, \$83,679.5 against \$83,697.9, although its tailings spend 94.3 h above the target against 148.8 h at a depth of 0.0167 g/t against 0.0299 g/t. The harder ore at 12 h is invisible to the model until the assays report it, and once the bias grows the plan holds up to 309 mg/L from 44.0 to 74.2 h with the true tailings under the target, for the same two reasons as in scenario 11, the old samples in the bias and a slope half the plant's. After the scenario's richer feed at 110 h the plan runs at the reach of the pump, 309.8 mg/L at most, and over 110 to 160 h it earns \$88,982.8 per hour against the loop's \$89,044.1. The true slope does not remove the loss, since it raises scenario 6 by only \$8.1 per hour.

The plan also breaks a limit it does not model. The 200 mg/L discharge limit on tank five's free cyanide is an alarm that no shipped loop holds and the scorecard does not count, and the objective prices the cyanide that reaches tank five only through its destruction cost. Under the predictive controller tank five exceeds the limit for 28.3 h on scenario 11, 101.6 h on scenario 6 and 186.2 h on scenario 8, against 0, 29.5 and 141.0 h under the soft sensor loop, with peaks of 230, 233 and

211 mg/L, and on every other case neither structure reaches it. That discharge is a cost of the predictive controller as configured.

Under the suggested challenge settings the plan passes, with the free cyanide within 9.9 mg/L of its setpoint, the checkpoint tailings at or below 0.365 g/t, the pH at or above 10.40 and 0.69 kg/t of sodium cyanide, and it spends 16.0 h above its 0.36 g/t target against 106.0 h for the cascade of Table 7.3. With the shipped settings, whose cyanide loop keeps the cautious tuning, the plan moves the setpoint faster than that loop follows and the free cyanide strays 77.2 mg/L from it, so the challenge fails on the cyanide clause instead of the tailings clause on which local control fails.

7.6.7 Moving Horizon Estimation Against the Bias Update

Table 7.5 runs the soft sensor's two failure modes, the structural error of scenario 9 and the over-fast update of scenario 10, under the soft sensor loop and the predictive controller, each case from its own settings, seed, noise and 240 h duration. Row 1 is the bias update at the case's gain ($\beta = 0.4$ on scenario 9 and 1 on scenario 10). Row 2 is the estimator of Section 7.5.3 with the settings of scenario 12, retuned for no case. Row 3 is the bias update retuned once by the estimator's own rule, $\beta = 0.4$. The rows marked 4 remove one mechanism at a time: the bounds (4a), the window, cut to one transition, the smallest the solver accepts (4b), and the lagged reference, replaced by the soft sensor's value at the cut (4d). Row 4c is a bias update at $\beta = 0.4$, written without the estimator and fed the lagged reference, and row 4e gives the shipped update at the case's own gain that same reference.

No bound binds in any of these runs, so removing the bounds (4a), cutting the window to one transition (4b), or both, leaves each run unchanged to every printed digit, and the plain bias update at $\beta = 0.4$ on the lagged reference (4c) reproduces the estimator on every case. As shipped, the estimator is therefore a bias update at $\beta = 0.4$ against a lagged reference, and every difference from row 1 comes from that gain and that reference. Given the static reference instead (4d), it is row 3.

On scenario 10 under the soft sensor loop the estimator earns \$89,074.1 per hour against \$89,045.4 for the shipped update at $\beta = 1$, a gain of \$28.7, of which \$25.1 belongs to the gain, since the update at 0.4 earns \$89,070.5, and \$3.6 to the lagged reference. It also halves the soft sensor's error (0.0100 against 0.0209 g/t) and cuts the setpoint travel by a third (1062 against 1575 mg/L), while the shipped gain on the lagged reference (4e) stays close to row 1 at \$89,051.6. Under the predictive controller the estimator earns \$89,225.4 against \$89,234.0, with half the error (0.0106 against 0.0196 g/t) and less than half the travel (409 against 849 mg/L), and row 3 earns \$89,222.4, so a steadier bias is worth no margin to the plan on this case.

On scenario 9 the estimator loses. The soft sensor's fourteen percent error is the sensor's own and nothing in the plant moves when it appears, but the lagged reference treats it as though it had passed through the tanks and corrects it with their lag. Under the soft sensor loop the estimator earns \$77,386.3 per hour against \$77,405.9, the true tailings spend 127.7 h above the target against 74.4 h and peak at 0.3266 g/t against 0.3172 g/t, the highest of any outer loop on this case, and the error rises from 0.0119 to 0.0144 g/t. Under the predictive controller it earns \$77,507.7 against \$77,517.2, and tank five's free cyanide passes the 200 mg/L discharge limit for 2.1 h, with a peak of 200.5 mg/L, where no shipped structure reaches the limit on this case.

Judged as an estimator, the lagged reference helps where the soft sensor's own inputs move and hurts where its error is its own. Over the scenarios outside Table 7.5, the soft sensor's error under its own loop falls from 0.0008 to 0.0004 g/t on the feed grade step of scenario 5 and by at most 0.0002 g/t on scenarios 6, 7, 8 and 11, rises from 0.0162 to 0.0177 g/t on scenario 12, and falls from 0.0300 to 0.0210 g/t on the challenge's suggested settings. The estimator's own estimate of the tailings at each cut, which has the assays on both sides of it, is closer still, 0.0053 g/t on scenario 9 and 0.0109 g/t on scenario 10 under the soft sensor loop. None of these errors is ever fed back.

The window and the bounds act only with a noisier laboratory. With the assay noise raised to 0.05 g/t at $\beta = 0.4$, which no scenario ships, the floor is active in 28 of the 59 updates of scenario 9 and 17 of scenario 10 under the predictive controller. There it cuts the largest error of the soft sensor

Table 7.5 The Bias Update and the Moving Horizon Estimator on the Soft Sensor's Two Failure Modes, 240 h. Over Is the Time the True Tailings Spend Above the 0.295 g/t Target, Counted at Every Integration Step, Peak the Largest True Tailings, Tank Five the Largest Free Cyanide in Tank Five and Discharge the Time It Spends Above 200 mg/L, Error the RMS Error of the Soft Sensor With Its Bias Against the True Tailings After 12 h, and Travel the Summed Movement of the Cyanide Setpoint at Every Step. No Bound Is Active in Any of the 59 Updates of Any Row

Case	Loop	Row	Net ($/h)	Over (h)	Peak (g/t)	Tank five (mg/L)	Discharge (h)	Error (g/t)	Travel (mg/L)
9, structural error	Soft sensor	1 bias update, $\beta = 0.4$	77 405.9	74.4	0.3172	126.2	0.0	0.0119	480
		2 estimator	77 386.3	127.7	0.3266	147.1	0.0	0.0144	505
		3 bias update, $\beta = 0.4$	77 405.9	74.4	0.3172	126.2	0.0	0.0119	480
		4a no bounds	77 386.3	127.7	0.3266	147.1	0.0	0.0144	505
		4b one transition	77 386.3	127.7	0.3266	147.1	0.0	0.0144	505
		4c update on lagged reference	77 386.3	127.7	0.3266	147.1	0.0	0.0144	505
		4d static reference	77 405.9	74.4	0.3172	126.2	0.0	0.0119	480
	Predictive	1 bias update, $\beta = 0.4$	77 517.2	0.0	0.2950	173.0	0.0	0.0111	262
		2 estimator	77 507.7	0.0	0.2950	200.5	2.1	0.0143	327
		3 bias update, $\beta = 0.4$	77 517.2	0.0	0.2950	173.0	0.0	0.0111	262
		4a no bounds	77 507.7	0.0	0.2950	200.5	2.1	0.0143	327
		4b one transition	77 507.7	0.0	0.2950	200.5	2.1	0.0143	327
		4c update on lagged reference	77 507.7	0.0	0.2950	200.5	2.1	0.0143	327
		4d static reference	77 517.2	0.0	0.2950	173.0	0.0	0.0111	262
10, $\beta = 1$	Soft sensor	1 bias update, $\beta = 1$	89 045.4	200.1	0.3158	124.4	0.0	0.0209	1575
		2 estimator	89 074.1	195.1	0.3098	119.1	0.0	0.0100	1062
		3 bias update, $\beta = 0.4$	89 070.5	197.0	0.3102	120.5	0.0	0.0104	1070
		4a no bounds	89 074.1	195.1	0.3098	119.1	0.0	0.0100	1062
		4b one transition	89 074.1	195.1	0.3098	119.1	0.0	0.0100	1062
		4c update on lagged reference	89 074.1	195.1	0.3098	119.1	0.0	0.0100	1062
		4d static reference	89 070.5	197.0	0.3102	120.5	0.0	0.0104	1070
		4e $\beta = 1$ on lagged reference	89 051.6	203.5	0.3122	123.6	0.0	0.0198	1558
	Predictive	1 bias update, $\beta = 1$	89 234.0	0.0	0.2950	157.6	0.0	0.0196	849
		2 estimator	89 225.4	0.0	0.2950	159.4	0.0	0.0106	409
		3 bias update, $\beta = 0.4$	89 222.4	0.0	0.2950	176.2	0.0	0.0108	417
		4a no bounds	89 225.4	0.0	0.2950	159.4	0.0	0.0106	409
		4b one transition	89 225.4	0.0	0.2950	159.4	0.0	0.0106	409
		4c update on lagged reference	89 225.4	0.0	0.2950	159.4	0.0	0.0106	409
		4d static reference	89 222.4	0.0	0.2950	176.2	0.0	0.0108	417
		4e $\beta = 1$ on lagged reference	89 234.3	0.0	0.2950	157.6	0.0	0.0195	876

from 0.0656 to 0.0469 g/t on scenario 9 and from 0.0653 to 0.0469 g/t on scenario 10, against the same estimator without bounds, and the net margin does not change at the first decimal. Under the soft sensor loop the floor binds three times on scenario 9 and twice on scenario 10 and changes the rms error by less than 0.001 g/t.

Scenario 12 sets both effects in one run. The soft sensor loop runs at $\beta = 0.4$ with 0.02 g/t of assay noise, the feed grade rises from 2.8 to 3.2 g/t at 12 h, and the soft sensor reads fourteen percent optimistic from 120 h, as shown in Figure 7.5. Between 12 and 120 h the estimator's soft sensor is off by 0.0097 g/t rms against 0.0105 g/t and by more than 0.01 g/t for 38.2 h against 42.3 h. Between 120 and 240 h it is off by 0.0225 g/t against 0.0200 g/t, for 59.5 h against 54.6 h, and the true tailings peak at 0.3458 g/t against 0.3367 g/t, higher than under any shipped structure. Over the 240 h the loop earns $88,994.5 per hour on the estimator against $89,024.4 on the update, and no bound is active in 59 updates. The curves in Figure 7.5 are drawn from engine data exported at 0.2 h intervals, not from screen captures.

Across every existing scenario and the challenge on its suggested settings, each run under the four shipped outer loops and under the soft sensor loop and the predictive controller with the estimator, the head-tank pH minima are identical in every row. Beyond scenario 9, the estimator is less safe than every shipped structure on four cases. The challenge on its suggested settings under the soft

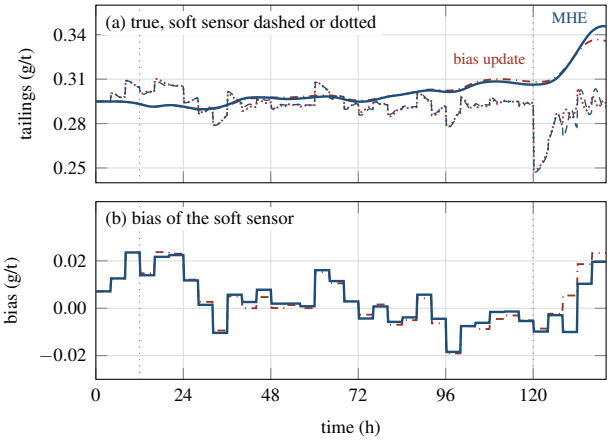


Figure 7.5 Scenario 12 from the start to 140 h with the moving horizon estimator (solid and dashed lines, drawn blue) and the bias update (dash-dotted and dotted lines, drawn red). (a) The true tailings (solid for the estimator, dash-dotted for the update) and the soft sensor with its bias (dashed and dotted). (b) The bias in force (solid and dash-dotted). The vertical dotted lines mark the richer feed at 12 h and the soft sensor's fourteen percent error from 120 h. At 140 h the true tailings are 0.346 g/t with the estimator and 0.336 g/t with the update, while the soft sensor reads 0.292 g/t and 0.293 g/t.

sensor loop spends 112.5 h above its target against 106.0 h, although its peak is lower (0.4294 against 0.4521 g/t) and so is its time above 0.40 g/t (17.7 against 21.6 h), and tank five's largest free cyanide under the predictive controller is higher by 0.2 mg/L to 1.8 mg/L on scenarios 6, 7 and 10 (232.9 against 232.7, 115.0 against 114.8 and 159.4 against 157.6 mg/L). Against the bias update under the same loop it also spends more hours above the target on scenario 7 (108.2 against 84.7 h under the soft sensor loop and 53.4 against 46.1 h under the predictive controller) and under the predictive controller on scenarios 6 and 11 (98.9 against 94.4 h, 41.4 against 38.3 h). It spends fewer hours above the discharge limit on scenarios 6, 8 and 11 (80.1 against 101.6 h, 130.7 against 141.0 h, 24.7 against 28.3 h).

7.7 Technology in Practice

Figure 7.6 draws the loops of Section 7.5 with their hardware.

What Measures It

The model reduces the free cyanide titrator to a reading of the head tank 0.40 h late, a cycle inside the range that online analyzers report.

> **TECHNOLOGY** **Sensor: An Online Free Cyanide Titrator**
>
> An online free cyanide analyzer filters a solids-free sample of solution from the leach pulp and titrates it with silver nitrate. Silver ions bind the free cyanide until an excess appears, and a silver electrode in the cell detects that end point as a change in potential. Such an analyzer reports every 5 to 30 minutes, against 2 to 12 hours for a sample taken and titrated by hand. Copper is its weakness: part of the cyanide bound to copper is titrated with the free cyanide and makes the end point hard to find and inaccurate [7]. On a copper-rich ore the reading can therefore stand above the true free cyanide, and the loop holds the wrong value.

The pH electrode is a first-order lag of 0.05 h and the oxygen probe reads the head tank continuously. The variable the circuit is paid on, the tailings grade, has no continuous instrument at all. The

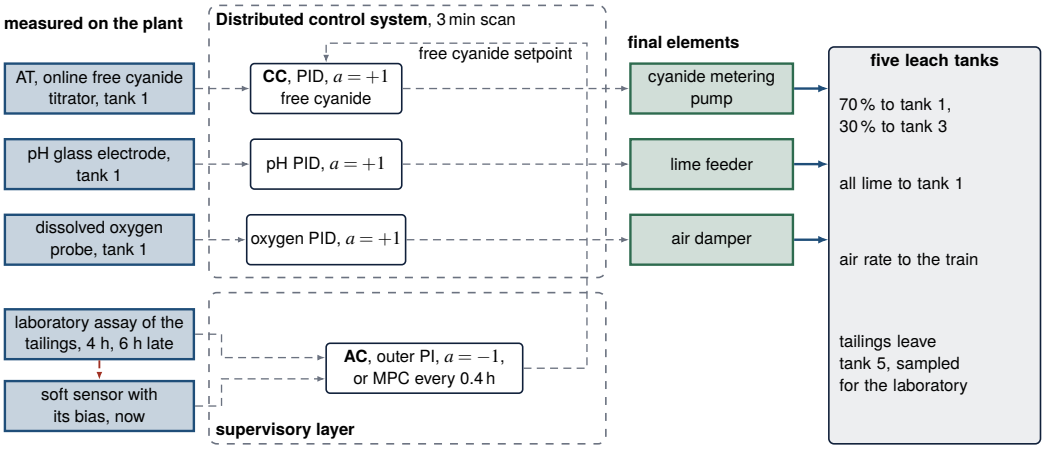


Figure 7.6 The loops of the simulated circuit drawn with their hardware. The online titrator AT, the pH electrode and the oxygen probe in tank one feed three PID blocks that move the cyanide metering pump, the lime feeder and the air damper. In the supervisory layer the outer controller AC, or the predictive controller in its place, reads the laboratory assay or the soft sensor whose bias the assay corrects (dashed), and writes the free cyanide setpoint of CC.

model takes it from a laboratory assay of a sample cut every four hours and reported six hours later, and from the soft sensor of Equation (7.3).

What Moves It

The model makes the cyanide pump, the lime feeder and the air damper first-order lags of 0.08 h, 0.12 h and 0.06 h, with the cyanide addition held between 0 and 500 kg/h and split seventy to thirty between tanks one and three.

> **TECHNOLOGY** **Final Element: A Cyanide Metering Pump**
>
> A metering pump is a pump mechanism, motor and drive electronics in one package, made to inject small, controlled flows of a chemical. A 4 to 20 mA signal sets its delivery from zero to full flow, and a knob on the pump sets what full flow is. Many metering pumps drive a plunger with a solenoid at a variable rate, so they deliver discrete pulses rather than a continuous stream [26]. A positive-displacement pump discharges a fixed volume with every turn and locks if its discharge is blocked [26]. A pump that has lost its delivery shows on a trend as a controller output that keeps rising while the free cyanide falls.

The sulfidic parcel and the graded week at the tighter target both drive the modeled pump to its 500 kg/h limit, the saturation that the full-flow setting of a real pump imposes in the same way.

Where the Controller Runs

Control loops in a plant are often nested or cascaded, with the setpoint of one loop based on the variable of another, and with cycle times from milliseconds to minutes. A distributed control system uses a supervisory loop to mediate a group of localized controllers [27], and programmable logic controllers run PID loops as well [26]. The app runs its PID blocks on a 3 min scan, short against the 2.68 h cyanide response and the 11.52 h tailings response. On such a platform the three secondary loops are PID function blocks, the soft sensor and its bias update are a calculation block that runs once per assay, and the outer loop or the predictive controller is a supervisory application that writes the free cyanide setpoint every 0.4 h.

> **IN PRACTICE Where the Cyanide Is Measured**
>
> In practice the free cyanide is measured in the first leach tank, by hand every 2 to 12 hours or with an online analyzer every 5 to 30 minutes. The consumption can change within two hours when the ore mineralogy changes, so the faster measurement holds the setpoint more tightly and sends less cyanide to the tailings [7]. The International Cyanide Management Code asks a mine to minimize its cyanide use and so limit the cyanide in its mill tailings (Standard of Practice 4.2), and to monitor the effects of cyanide on wildlife and on surface and ground water (Standard of Practice 4.9) [28].

What Is Modernizing It

The first development is the online analyzer. Control of the cyanide addition from an online measurement has been reported to cut cyanide consumption by at least five percent against manual control, and mines that sign the International Cyanide Management Code are required to measure the weak-acid-dissociable cyanide online as well [7]. The second is the data-driven soft sensor. Because an accurate mechanistic model of a gold leach is hard to build, data-driven modeling is used extensively in industrial plants, with methods from partial least squares to neural networks [29], where the soft sensor of this chapter evaluates the leach kinetics instead. The third is optimization of the setpoint itself, by real-time optimization of the leach under disturbance and model uncertainty [29] or by the economic predictive controller and moving horizon estimator of Sections 7.5.2 and 7.5.3.

7.8 Limitations

Gold is one leachable class and one refractory floor. Real ore has a distribution of liberation and grain size, so the leach curve is a spectrum of rate constants rather than a single one, and the refractory fraction itself depends on the grind [1]. The application represents that a slower ore raises the tailings grade and does not predict by how much for a particular deposit.

Carbon is one loading per tank with no regeneration circuit. Elution, acid washing, kiln regeneration and carbon attrition are all absent, and carbon management is a substantial part of the control task on an operating circuit [2]. Preg-robbing by carbonaceous ore, which is the failure mode that makes some deposits uneconomic, has no representation.

Cyanide destruction and the tailings storage facility are outside the boundary. The application reports a detoxification load and does not model it, and the environmental constraint on an operating circuit is at least as binding as the reagent cost.

The soft sensor is a steady-state evaluation of the same leach kinetics the plant model uses, with its constants fixed at the plant's declared values, differing in the ore parameters, in the assumption of a uniform profile along the train and in omitting the gold that leaves in solution. A soft sensor on an operating circuit is a fitted model of a different form, and its failure modes include extrapolation outside the data on which it was fitted, which cannot be staged here. What the application does represent is that an inference is no better than its most uncertain input, and that the most uncertain input is usually the ore [6].

The predictive controller inherits the soft sensor's model and its limits. Its slope of tailings against cyanide is 1/1.25 of the plant's at design and about half of it on a harder ore, no bias can correct a slope, and a fixed setpoint chosen in hindsight therefore beats it by \$8.3 to \$132.4 per hour. It sees an ore change only through the laboratory, which costs it scenario 6. Its ceiling carries a slack at every judged step, so a ceiling that cannot be met costs the plan at every step where it is broken, and on scenario 8 the plan runs the pump at its limit for 67.4 h, ends with the true tailings at 0.298 g/t and earns \$75.3 per hour less than the earlier form with one slack over the day, which postponed that cyanide. Honoring a ceiling the pump cannot reach is thus a design choice with a measured price. The plan holds the measured feed grade, throughput, pH and oxygen constant over its horizon, and it carries no constraint on tank five's free cyanide, which exceeds the 200 mg/L discharge limit for 28.3 h on scenario 11, 101.6 h on scenario 6 and 186.2 h on scenario 8.

The moving horizon estimator inherits the soft sensor's model through its reference and adds the tanks' lag to it. It treats every error as one that passed through the train, so an error at the sensor itself is corrected with that lag, which makes it less safe than every shipped structure on scenario 9. Its only unknown is an additive bias, so the soft sensor's understated slope remains. Its weights are declared rather than fitted, a laboratory repeatability of 0.015 g/t where scenarios 10 and 12 carry 0.02 g/t of assay noise, and because neither bound binds on a shipped case, the window and the bounds add nothing there that a bias update on the same reference lacks.

> **OPEN CHALLENGE** **Reaching the Plant's Own Optimum**
>
> A setpoint that is optimal on a model of a gold cyanidation leach is often not optimal on the plant. An accurate mechanistic model is hard to build for a long process with complex chemistry, and the plant drifts as its equipment ages and its ore changes. Published remedies fit local models from plant data near the current point and step the setpoint toward the plant's own optimum while the residue grade meets its limit, and they have been shown in simulation [29]. Progress would be such a method proved on an operating circuit.

7.9 Discussion

Three results generalize beyond this circuit.

The first is that a soft sensor is a measurement and should be evaluated as one. Sections 7.6.1 and 7.6.2 evaluate it that way, finding that it removes six hours of dead time and introduces a model error of up to 0.04 g/t, and that the second figure is tolerable only because an assay exists to correct it. A soft sensor without a reconciling measurement is a model running open loop rather than a faster measurement.

> **OPEN CHALLENGE** **Seeing the Ore Before It Leaches**
>
> Sensors on the ore feed can reveal trends in its grade and in how its gold occurs, and such readings could inform the cyanide addition. An operating plant still judges its recovery from samples, drawn in one study twice a day from every tank of a carbon-in-leach circuit for 209 days, and the variability of the feed and the uncertainty of sampling blur what they show [19]. Progress would be a measurement of how fast the incoming ore will leach, made before it reaches the first tank.

The second is that the reconciliation has dynamics of its own. Table 7.2 shows a bias update that is too fast passing the assay dead time and the laboratory's noise back into the loop, with the cyanide setpoint wandering over 101 mg/L against 53 mg/L at a gain of 0.15. The correct gain defines a filter length, and the sampling interval and the assay noise rather than the controller set that length. The moving horizon estimator of Table 7.5 leads to the same place, because with no bound active it is a filter of that gain and its result is decided by the gain and the reference it compares against, not by its window.

The third concerns when a supervisory layer is justified. An outer loop is worth closing when the plant drifts more slowly than the loop can act and through variables its estimate sees. On Table 7.3 the recommended cascade saves two percent of reagent and spends 21.6 h above the tailings limit, because its target spends the margin that a fixed setpoint keeps and the first disturbance is one its estimator cannot see. The secondary loops already reject what is available to be rejected. A supervisory layer that the disturbances do not justify is a common and costly choice [5, 12]. Volume 9, Chapter 10 asks the same question of an aeration plant, where a slow effluent ammonia loop sets a fast dissolved oxygen setpoint.

7.10 Conclusion

A carbon-in-leach circuit places 11.52 h of plant response and six hours of laboratory turnaround between a reagent decision and the evidence that judges it. The application integrates a five-tank leach and adsorption train, holds free cyanide, alkalinity and dissolved oxygen on continuous instruments,

and reports the tailings grade as a computed value, a delayed assay and a soft sensor together. The soft sensor removes the assay dead time and introduces an error of up to 0.04 g/t that only the laboratory detects, which makes the bias update the governing design decision: at a gain of one it passes the assay dead time and noise back into the loop and raises the actuator travel by sixty-three percent over a gain of 0.15, and at zero it starves the leach for 172 h of a 200 h week.

On the graded week the outer loop is hardly worth closing, because a fixed cyanide setpoint holds the tailings under 0.40 g/t throughout at 0.535 kg/t of reagent, while the recommended cascade saves two percent of the reagent and exceeds the limit for 21.6 h, and a target tight enough to stay under the limit costs a third more reagent and runs the cyanide pump into its limit. The excursion comes from the proportional response to the target itself rather than from integral windup. The result does not show that inferential control is inappropriate. It shows that whether to close a supervisory loop is a question that precedes the question of what to close it on.

An economic predictive controller on the soft sensor’s model earns $25.7 per hour more than the soft sensor loop on scenario 11, but a setpoint optimizer feeding that loop recovers all but $1.9 of it, a fixed setpoint chosen in hindsight earns $28.2 more than the plan, and the plan exceeds the tank-five discharge limit for 28.3 h where the loop does not. A moving horizon estimator of the bias, whose bounds never bind on a shipped case, is a bias update at a gain of 0.4 against a lagged model of the tanks, so its $28.7 per hour on the over-fast update belongs mostly to its gain, and it loses the structural error of scenario 9. Every result here except the setpoint optimizer’s is a shipped scenario at its own seed, so a reader can reproduce it in the app and change one setting at a time.

KEY POINTS

- The gold left in the tailings decides what the circuit earns, and a laboratory reports it six hours late.
- A soft sensor removes the assay dead time and adds an error of up to 0.04 g/t that only the laboratory detects.
- The bias update gain is a filter design, set by the sampling interval and the assay noise.
- On the graded week a fixed cyanide setpoint holds the tailings under 0.40 g/t, while the recommended cascade spends 21.6 h above the limit.
- Whether to close a supervisory loop precedes the question of what to close it on.

Where else this applies. A late, sampled quality measurement corrected into an inferential estimate also governs Volume 5, Chapter 10 and Volume 5, Chapter 4.

References

[1] J. O. Marsden and C. I. House, *The Chemistry of Gold Extraction*, 2nd ed. Littleton, CO: Society for Mining, Metallurgy, and Exploration, 2006.

[2] B. A. Wills and J. A. Finch, *Wills’ Mineral Processing Technology: An Introduction to the Practical Aspects of Ore Treatment and Mineral Recovery*, 8th ed. Oxford: Butterworth-Heinemann, 2015.

[3] J. S. MacArthur, R. W. Forrest, and W. Forrest, “Process of obtaining gold and silver from ores,” U.S. Patent 403,202, application filed November 9, 1887, 1889, patented May 14, 1889.

[4] ——, “Process of separating gold and silver from ore,” U.S. Patent 418,137, application filed April 4, 1889, 1889, patented December 24, 1889.

[5] D. Hodouin, “Methods for automatic control, observation, and optimization in mineral processing plants,” *Journal of Process Control*, vol. 21, no. 2, pp. 211–225, 2011.

[6] E. Mazzour, D. Hodouin, and S. Makni, “Optimal sensor implementation in metallurgical plants: an application to a generic mineral separation plant,” *International Journal of Mineral Processing*, vol. 69, no. 1–4, pp. 185–203, 2003.

[7] W. A. M. van der Merwe, “The development of an online amperometric technique to measure free and WAD cyanide in gold plant leach liquors and effluent streams,” M.Eng. dissertation, University of Cape Town, 2012.

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

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

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

[11] ——, "The future of PID control," *Control Engineering Practice*, vol. 9, no. 11, pp. 1163–1175, 2001.

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

[13] D. E. Rivera, M. Morari, and S. Skogestad, "Internal model control: PID controller design," *Industrial & Engineering Chemistry Process Design and Development*, vol. 25, no. 1, pp. 252–265, 1986.

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

[15] J. McLaughlin and G. E. Agar, "Development and application of a first order rate equation for modelling the dissolution of gold in cyanide solution," *Minerals Engineering*, vol. 4, no. 12, pp. 1305–1314, 1991.

[16] G. Senanayake, "Kinetics and reaction mechanism of gold cyanidation: Surface reaction model via Au(I)–OH–CN complexes," *Hydrometallurgy*, vol. 80, no. 1–2, pp. 1–12, 2005.

[17] M. Adams, R. Lawrence, and M. Bratty, "Biogenic sulphide for cyanide recycle and copper recovery in gold-copper ore processing," *Minerals Engineering*, vol. 21, no. 6, pp. 509–517, 2008.

[18] N. Tsuchida, "Studies on the mechanism of gold adsorption and elution in the carbon-in-pulp process," Ph.D. dissertation, Murdoch University, Perth, Western Australia, 1984.

[19] H. J. Glass, R. Abdulwahed, and S. Iyakwari, "Monitoring gold recovery by sampling a carbon-in-leach plant," in *Proceedings of the 57th International October Conference on Mining and Metallurgy*, Bor, Serbia, 2026, pp. 177–181.

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

[21] M. Ellis, H. Durand, and P. D. Christofides, "A tutorial review of economic model predictive control methods," *Journal of Process Control*, vol. 24, no. 8, pp. 1156–1178, 2014.

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

[23] 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.

[24] Q. Liu, Z. Wang, Y. Liang, Z. Wang, F. Peng, and D. He, "Soft sensor modeling and adaptive model correction strategy for gold cyanide leaching process," in *2019 Chinese Control And Decision Conference (CCDC)*. IEEE, 2019, pp. 6065–6070.

[25] R. Dubois, S. Bertrand, and A. Eudes, "Performance evaluation of a moving horizon estimator for multi-rate sensor fusion with time-delayed measurements," in *2018 22nd International Conference on System Theory, Control and Computing (ICSTCC)*. IEEE, 2018, pp. 664–669.

[26] 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.

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

[28] International Cyanide Management Institute, "The International Cyanide Management Code: principles and standards of practice for mining operations," cyanidecode.org/about-the-cyanide-code/the-cyanide-code, 2026, accessed November 19, 2026.

[29] Q. Liu, D. He, Z. Wang, and J. Shi, "A data-driven iterative optimization compensation method based on PJIT-PLS for gold cyanidation leaching process," *IEEE Access*, vol. 7, pp. 138 683–138 693, 2019.

7.11 Exercises

7.1 Concept check · 2 min

In scenario 6 the ore leachability falls to 0.82 at 12 h while the feed grade, the cyanide, the pH and the oxygen stay where they were. With the outer loop on the soft sensor, the cyanide setpoint first moves at 20.1 h, the same time as with the outer loop on the assay. Why is the soft sensor no faster here?

(a) The titrator's 0.40 h dead time delays the cyanide reading the soft sensor uses

(b) The outer loop's rate limit holds the setpoint until its error has grown past a threshold

(c) The soft sensor passes its estimate through the five tank lags before the loop reads it

(d) Leachability is not one of its inputs, so only an assay that contains the change can move it

7.2 Concept check · 1 min

Each assay updates the bias by Equation (7.3). Which value of the soft sensor is the assay compared with?

(a) Its value stored at the instant the sample was cut, six hours before the report

(b) Its value at the instant the assay is reported, the latest estimate available

(c) Its average over the four hours between the last two samples that were cut

(d) Its value at the report, passed through the five residue lags of the tanks

7.3 Concept check · 2 min

In scenario 7 a copper parcel multiplies the cyanicide term by 2.2. The head-tank cyanide loop stays within 10.2 mg/L of its setpoint, yet the tailings rise to 0.3084 g/t. Which explanation is right?

(a) The titrator's delay lets the head-tank cyanide wander far from its setpoint between readings

(b) The parcel consumes cyanide along the train, so tank five falls to 53.2 mg/L while tank one is held

(c) The copper acidifies the later tanks, and the head-tank pH electrode cannot see their pH fall

(d) The soft sensor stops reading the feed grade while the parcel passes through the train

7.4 Theory and calculation · 6 min

The tailings are sampled every $T_s = 4$ h and each assay is reported 6 h after its sample was cut. The bias update of Equation (7.3) runs at gain β.

(a) How old is the assay in hand at any instant, at least, at most and on average? *Checkpoint:* 6 h, 10 h, 8 h.

(b) Find the time constant of the bias update at $\beta = 0.15$, 0.35 and 0.40, to 0.1 h. *Checkpoint:* 24.6 h, 9.3 h, 7.8 h.

(c) The assay carries white noise of spread 0.02 g/t. Using Section 1.4.5, find the spread it leaves in the bias at $\beta = 0.15$ and 1. Then take a single assay 0.02 g/t high with nothing else changing: how far does the outer loop's proportional term, at $K_c = 2500\,\mathrm{mg/L}$ per g/t, move the cyanide setpoint at each gain, and how long does the 8 mg/L per hour ramp take to deliver the larger move? *Checkpoint:* 0.0057 g/t and 0.020 g/t; 7.5 mg/L and 50 mg/L; about 6.3 h.

7.5 Theory and calculation · 6 min

At design the circuit treats 400 t/h of ore, gold is worth \$78 per gram and sodium cyanide costs \$2.40 per kilogram. Holding one more mg/L of free cyanide in the head tank takes 1.651 kg/h of sodium cyanide, and a fraction 0.7471 of the head-tank cyanide reaches tank five and leaves in 488.9 $\mathrm{m^3/h}$ of solution, where its destruction costs \$3.50 per kilogram.

(a) What is 0.001 g/t of tailings worth per hour, to the dime? *Checkpoint:* \$31.2 per hour.

(b) What does one more mg/L of setpoint cost per hour, to the cent, and how much must it lower the tailings to pay for itself, to two significant figures? *Checkpoint:* \$5.24 per hour; about 0.17×10^{-3} g/t per mg/L.

(c) Use the cyanide rows of Table 7.1 to say whether the steps from 120 to 200 mg/L and from 200 to 260 mg/L pay. *Checkpoint:* The first pays, the second does not.

7.6 Theory and calculation · 5 min

The soft sensor assumes every tank sits at the head-tank conditions. At design the head tank holds 120 mg/L of free cyanide, 6.0 mg/L of oxygen and pH 10.40, where $f(\mathrm{pH}) = 1$. Take $k = 0.947/\mathrm{h}$, $K_{\mathrm{CN}} = 60\,\mathrm{mg/L}$, $K_{\mathrm{O}} = 3.0\,\mathrm{mg/L}$, a residence of 3.14 h per tank, five tanks, a feed of 2.8 g/t and 0.25 g/t of refractory gold.

(a) Find the soft model's leach rate k_s at these conditions, to three decimals. *Checkpoint:* 0.421/h.

(b) Find the tailings grade the soft sensor computes, to three decimals, with Equation (1.6) applied to the leachable gold, and compare it with the true 0.295 g/t. *Checkpoint:* 0.288 g/t, about 0.007 g/t optimistic.

(c) Find the residue time constant τ_r of one tank and the sum of the five, to 0.01 h and 0.1 h. *Checkpoint:* 1.35 h and 6.8 h.

7.7 App challenge · 8 min

Open scenario 4, "Close the loop on the laboratory assay", and then scenario 5, "Close it on a soft sensor instead", each as loaded and at the default seed. Both step the feed grade from 2.8 to 3.2 g/t at 10 h. Run each for 200 h at 500 h per second and export the CSV, whose time column is Time (h).

(a) Find the first row in which `cn_setpoint_mgL` exceeds 122 mg/L in each run, within 0.1 h, and how long after the feed step each one comes. *Checkpoint:* 20.3 h and 10.3 h; about 10.3 and 0.3 h after the step.

(b) Read the MV travel tile of the free cyanide loop on the scorecard at 200 h in each run, within 2 kg/h. *Checkpoint:* 428 kg/h and 296 kg/h.

(c) Find the largest cyanide setpoint in scenario 4 and when it occurs, within 1 mg/L and 0.5 h, and the setpoint at 200 h. Explain the overshoot. *Checkpoint:* About 150 mg/L at 40 h; about 140 mg/L.

7.8 App challenge · 9 min

Open scenario 10, "A bias update fast enough to bring the dead time back", at the default seed. It runs the soft sensor loop with Soft sensor bias update gain 1.0, a Tailings assay noise of 0.02 g/t and a feed grade step at 12 h. Run each part for 240 h at 500 h per second. Read the free cyanide loop's MV travel tile, within 2 kg/h, and the NaCN tile, and from the CSV the lowest and highest `cn_setpoint_mgL` at or after 24 h, within 1 mg/L.

(a) Run it as loaded. *Checkpoint:* 5472 kg/h; 0.512 kg/t; 78 to 179 mg/L.

(b) Press Reset, set Soft sensor bias update gain to 0.15 before Run, and run. *Checkpoint:* 3356 kg/h; 0.517 kg/t; 101 to 154 mg/L.

(c) Reload the scenario with its reload button, set Tailings assay noise (Advanced) to 0 before Run, keep the gain at 1.0, and run. Read the travel. What does the dead time alone cost, and what does the noise cost? *Checkpoint:* 425 kg/h.

7.9 App challenge · 9 min

Open the graded challenge, "a week of difficult ore", at its default seed. Run each part for the full 200 h at 500 h per second, and read the challenge result in the progress panel and the NaCN tile. Use the CSV to count the hours with `tails_gt` above 0.40 g/t, at 0.1 h a row, within 0.2 h.

(a) Run it as loaded. Which clause fails? *Checkpoint:* Not met: the tailings, 0.433 g/t; cyanide 8.5 mg/L, pH 10.32.

(b) Reload, press Reveal suggested tuning and run. Count the hours above 0.40 g/t and read the peak of `tails_gt`. *Checkpoint:* Met: 10.1 mg/L, 0.379 g/t, pH 10.40; 0.525 kg/t; 21.6 h; 0.452 g/t.

(c) Press Reset, set Outer loop to Local before Run, leaving the free cyanide setpoint at 120 mg/L, and run. Compare with (b). *Checkpoint:* Met: 6.1 mg/L, 0.396 g/t; 0.535 kg/t; no hour above 0.40 g/t.

7.10 Application area · 4 min

A plant meters all of its cyanide into the first tank and holds the head-tank free cyanide on a tight loop.

(a) Using scenario 7 of Section 7.6.4, say what happens in tank five when a copper-bearing ore arrives, and what changing the split of the addition between tanks one and three achieves. *Checkpoint:* Tank five falls to 53.2 mg/L; with half in tank three it stays at or above 76.3 mg/L.

(b) The predictive controller of Section 7.6.6 pushes tank five's free cyanide above its 200 mg/L discharge limit for 28.3 h on scenario 11. What would you add before running such a plan on a plant? *Checkpoint:* A constraint or override on tank five's cyanide.

7.11 Application area · 4 min

The head-tank pH is held by lime at a setpoint chosen by the operator.

(a) Use Equation (7.2) to find the share of the free cyanide present as hydrogen cyanide at pH 9.5. *Checkpoint:* 39 %.

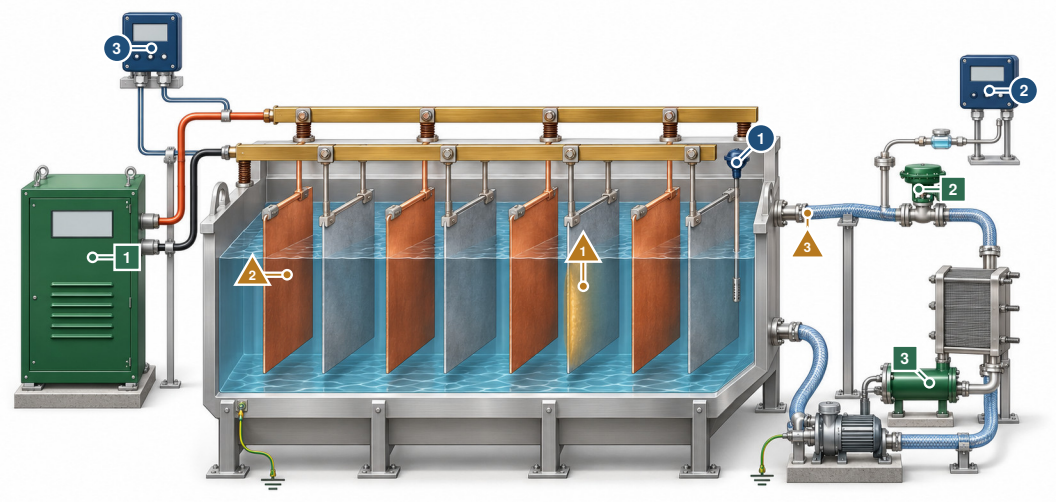
(b) The margin optimum lies at pH 10.55, \$82 per hour above the design pH of 10.40, and pH 9.60 costs \$3237 per hour. On which side of the optimum would you err, and what does the sulfidic parcel of scenario 8 add to the argument? *Checkpoint:* The high side; tank five can fall below the floor while the head tank is held.

7.12 Review the agent · 5 min

An AI assistant summarizes Table 7.5: "On scenario 10 the moving horizon estimator earns \$28.7 per hour more than the shipped bias update, because its 48-hour window and its bounds reject noisy assays that the update believes. Replace the bias update with the estimator on every circuit." Find the flaw and say how you would check it in the app.

CHAPTER 8

A Copper Electrorefining Simulator for Inferential Constraint Control



The controller moves 1 rectifier that sets the cell current 2 electrolyte flow valve 3 trim heater

The controller reads 1 electrolyte temperature 2 copper in the electrolyte, from a sample, two hours late 3 cell voltage

What disturbs it 1 solids and nodules at the cathode 2 anodes that passivate 3 leaner copper in the returning electrolyte

An electrorefining cell whose current is pushed toward a limit at the cathode surface that no instrument in the picture can read. Illustration, not to scale.

IN PLAIN TERMS **Pushing a Copper Cell Toward an Unseen Limit**

The main idea. A refining cell passes an electric current through an acid bath. The current carries copper from a rough slab onto a clean sheet, and more current makes more copper.

The challenge. Past a certain point, small bumps on the sheet grow faster the taller they get, and the sheet is spoiled. That point is worked out from book values and a late lab test, not read from a gauge. So a wrong guess goes unseen while every reading looks normal.

Why it matters. A safe margin costs a little copper, but crossing the true limit can spoil a whole crop.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history box of Section 8.2, the fact box of Section 8.3.2, the try-it and fact boxes of Sections 8.5.1 and 8.5.2, the try-it box of Section 8.5.5, the history box of Section 8.5.6, and the key points
Study	40 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 8.5, Section 8.6, Section 8.7 and the open challenge after it

The control question A tankhouse makes more copper the harder it drives the rectifier, until nodules on the cathode start to compound past a stability limit, and the limiting current density that sets that limit is computed, not measured. How close to an inferred constraint should a constraint controller run, and what does a wrong estimate cost while every loop reports success?

Learning objectives

- Compute the limiting current density, the utilization and the stability limit at an operating point from the mass-transfer, growth and leveler laws.
- Explain why the copper loop is direct acting and how closing its valve lowers the production ceiling.
- Predict the true utilization a constraint controller holds when its estimate of the limit is biased, and size a back-off for it.
- Weigh a linear gain in copper against a compounding loss in grade over a cathode cycle, and choose a defensible utilization setpoint.
- Distinguish a disturbance the constraint controller defends from one it cannot see, and choose the remedy for each.

Prerequisites Sections 1.4.6 (inferential control and the back-off from an inferred limit), 1.4.5, 1.6.2 and 1.6.4.

The app https://apmonitor.com/apps/copper-electrorefining/

See also Volume 5, Chapter 10 and Volume 5, Chapter 11 for control on an inferred variable that no instrument reaches, and Volume 4, Chapter 2 and Volume 4, Chapter 6 for a production rate held against a limit, with an override behind it.

8.1 Introduction

An electrorefining cell dissolves an impure copper anode and deposits pure copper onto a cathode by passing current between them through an acidic copper sulfate electrolyte [1]. Faraday's law makes the production rate exactly proportional to the current, so the economic instruction to a tankhouse is to raise the rectifier current. Mass transfer sets the limit: copper ions reach the cathode surface by diffusion across a concentration boundary layer, and as the applied current density approaches the limiting current density the deposit roughens and the cathode is downgraded [2]. A protrusion on the cathode reaches further into the boundary layer than its surroundings, so it is better supplied and grows faster the taller it is, and the leveling agent in the electrolyte is what holds that instability back [3, 4].

No instrument on an industrial cell measures the limiting current density. The value is computed from a handbook diffusivity, a boundary-layer correlation and the assayed bulk concentration, so the controlled variable of the production loop is an inferred quantity. Inferential measurements are common in the process industries wherever a controlled variable has no online sensor [5, 6], and the loop inherits every error in the inference while its own performance statistics report only the error it can see. Two other plants of this volume hold inferred variables: the impact angle of Chapter 3 and the tailings grade of Chapter 7.

The app follows the browser and laboratory formats of the same series [7]. Its scenarios teach constraint control against an estimated limit, inferential measurement, an instability that compounds past a threshold, override protection and actuator saturation, which are standard topics in process control texts [8, 9]. Surveys of process control education ask for problems in which the economic decision and the control structure interact [10, 11]. Industry requires the same judgment, because the margin held between a constraint setpoint and an estimated limit is the routine economic decision on a tankhouse.

Section 8.2 places the cell in a tankhouse and in its control hierarchy. Section 8.3 gives the model in enough detail to replicate it: the mass-transfer limit, a nodule that grows past a stability limit set by the leveling agent, the cell voltage and the design point. Section 8.4 describes the two regulatory loops, the constraint controller and the voltage override. Section 8.5 reports the measured results

of the twelve scenarios: where the stability limit lies and what it is worth, how the additive and the temperature move it, two plant upsets, two errors in the estimate, dirty electrolyte and the graded crop. Section 8.6 names the hardware of a real tankhouse, and Sections 8.7 and 8.8 state the limitations and three observations that carry beyond this cell. The inferred limit and the back-off it needs are introduced in Section 1.4.6.

8.2 Background

A commercial tankhouse runs hundreds of cells in electrical series at current densities near 300 A/m^2 to 400 A/m^2, producing Grade A cathode at specific energies around 250 kWh/t [1, 12]. Electrorefining is a small share of the energy a copper operation uses compared with comminution [13], so the operating pressure on a tankhouse is throughput rather than power. Deposit quality is managed with leveling additives such as thiourea and glue, dosed to a concentration in the electrolyte [4], and anode passivation raises the cell voltage when the anodes are pushed too hard [14, 15]. The circulating electrolyte is filtered because suspended slime and fines that settle on the cathode become nodules [1]. Mass transfer at the cathode sets the achievable current density, and flow enhancements such as jetting between the electrodes change that limit substantially [16]. Impurity deportment through the purification circuit decides whether the cathode meets specification at least as often as physical roughness does [17].

> ⧖ HISTORY **Refining Copper by Current (1865)**
>
> Copper could be refined by electrolysis on a commercial scale only once the dynamo made large currents cheap. The eleventh edition of the *Encyclopaedia Britannica* credits Elkington, whose first patent is dated 1865, with the earliest serious attempt. Crude copper was cast into anode plates, and sheets of deposited copper served as cathodes, hung alternately in troughs of copper sulfate solution. A hundred or more troughs were coupled in series, the electrolyte was circulated through them once a day, and the current density was 5 or 6 amperes per square foot. Silver and the other insoluble impurities settled to the bottom of each trough. The encyclopedia judged the process identical in principle with the practice of its own day [18], and the cell of this chapter still works the same way.

The control structure on such a plant is regulatory PID at the base with supervisory logic above it [8, 19]. Where the supervisory layer pushes against a constraint that has no sensor, the constraint value is estimated from available measurements and model parameters, a structure the minerals industry uses widely for grades, recoveries and other unmeasured quantities [6]. The failure mode documented in this chapter, a constraint controller holding an estimate while the plant crosses the true limit, is the standard exposure of that structure.

8.3 Process Model

Figure 8.1 shows the simulated cell, its two regulatory loops and the supervisory constraint controller above them. The model integrates nine states: cell temperature, bulk copper concentration, rectifier current, electrolyte flow, heater duty, leveler concentration, nodule height, deposited mass and suspended solids. The thermal, copper, leveler and solids balances are conservation equations on a single well-mixed electrolyte volume with first-order actuator lags, so the model separates a fast electrical time scale (a rectifier lag of 0.02 h) from a slow thermal one (about 3 h) [9].

8.3.1 Mass Transfer Limit

Copper ions reach the cathode surface by diffusing across a boundary layer, so there is a current density above which the surface would be starved, the limiting current density

$$j_L = \frac{nFDC}{\delta}, \qquad \delta = \delta_0 \left(\frac{F_{\mathrm{ref}}}{F_e} \right)^{n_\delta} \tag{8.1}$$

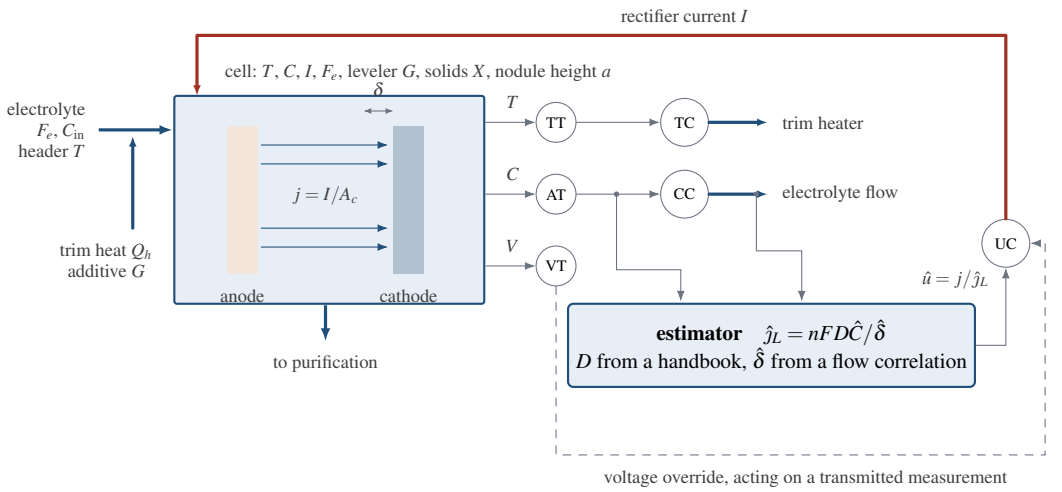


Figure 8.1 The simulated cell. Two regulatory loops hold cell temperature on a trim heater (TC) and bulk copper on the electrolyte flow (CC). A supervisory constraint controller (UC) holds current density utilization on the rectifier current, but the utilization it acts on is an estimate: the limiting current density is computed from a handbook diffusivity, a correlation boundary layer and the bulk assay, and no instrument on the cell measures it. Thick lines carry the manipulated variables (the heater and the flow valve on the right, the rectifier current across the top), thin arrows carry signals, and the dashed path is the cell voltage override, the only constraint in the structure that acts on a transmitted measurement.

where n is the number of electrons transferred, F the Faraday constant, D the diffusivity of the copper ion, C the bulk copper concentration, δ the boundary layer thickness, and δ_0, F_{ref} and n_δ correlate the layer to the electrolyte flow F_e. The controlled variable of the production loop is the fraction of that limit in use,

$$u = \frac{j}{j_L} = \frac{I/A_c}{j_L} \tag{8.2}$$

where j is the applied current density, I the rectifier current and A_c the cathode area. In the notation of Equation (1.9) u is the utilization η. Every term in the denominator of (8.2) is either a laboratory value or a correlation. The diffusivity comes from a handbook, the boundary layer from a flow correlation, and the surface concentration is not the bulk assay. No instrument on the cell measures j_L, and no instrument measures u. Figure 8.2 sets the path the controller closes beside the path the deposit follows.

🕮 **Pause.** Try Exercise 8.1 before reading on.

8.3.2 Deposit Quality

The cathode's quality is one physical quantity, the height a of a representative nodule, which follows a linearized morphological-stability law [3] written per unit of deposit thickness and multiplied by the plating velocity v,

$$\frac{da}{dt} = v\left[\frac{g\,(u-u_s)}{\delta}\,a + \beta X\right], \qquad v = \frac{\eta_c\, j\, M_{\text{Cu}}}{nF\rho_{\text{Cu}}} \tag{8.3}$$

where g is a tip amplification factor, u_s the stability limit, β a nucleation coefficient, X the suspended solids in the cell, η_c the cathode current efficiency, M_{Cu} the molar mass and ρ_{Cu} the density of copper. A bump that reaches a/δ into the boundary layer is better supplied by that fraction,

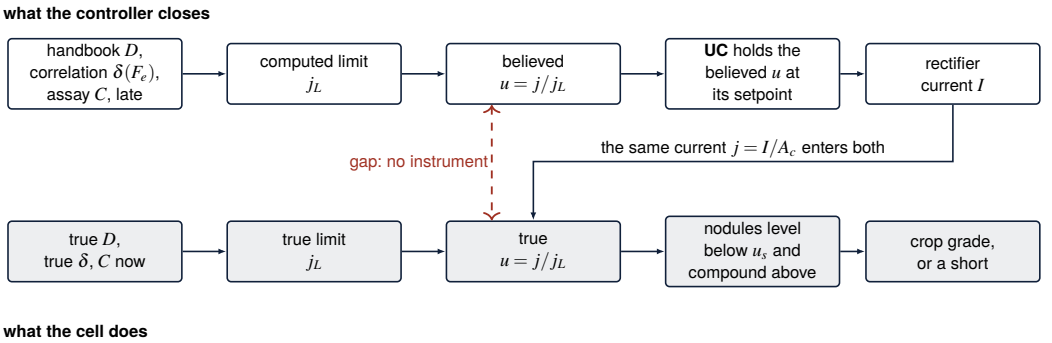


Figure 8.2 The two paths of the constraint loop. The controller closes the top path on a limit computed from a handbook value, a correlation and a late assay, and its integrated error measures only how well it holds that belief. The nodules follow the bottom path, set by the true limit through the same current, and the gap between the two rows reaches no instrument on the cell.

and with kinetics of first order in the surface concentration at a common electrode potential the local current answers its local limit with a sensitivity of exactly u, so the first term in the bracket is dimensionless, and $g = 2.2$ carries the focusing of the diffusion field onto a protrusion. Above u_s a nodule grows in proportion to its own height and its height compounds. Below u_s the same term is negative, the hollows fill faster than the bumps grow, and the plate levels itself. Particles that land on the cathode start new nodules at a rate proportional to X. Because every term carries $v \propto j$, nothing grows or levels while the current is off, whatever the temperature. A fresh blank starts at $a_0 = 5\,\mu\text{m}$, the plate advances 41 µm/h at design, a crop harvested with nodules above 2 mm is downgraded, and a nodule that reaches the 35 mm gap shorts the cell, which no leveling undoes.

> ✦ DID YOU KNOW? **A Week of Copper Is Under a Centimeter**
>
> Smooth copper plated at 350 A/m^2 for a week is only 7.8 mm thick, yet nodules grow across the 25 mm to 35 mm gap to the anode and short it. Such shorts are the largest cause of lost current efficiency, which runs at 93 % to 98 % in operation [20].

The stability limit is set by the leveling agent,

$$u_s = u_0 + \Delta u_L\,\theta - \Delta u_{\text{org}}\,\frac{(G - G_{\text{org}})_+}{G_{\text{org}}}, \qquad \theta = \frac{G}{G + G_{50}(T)} \tag{8.4}$$

where $u_0 = 0.414$ is what a bare surface tolerates, θ the Langmuir coverage of the leveler at concentration G, $\Delta u_L = 0.17$ the lift at full coverage, $(x)_+$ the positive part, and the last term the collapse of the limit once the leveler is dosed above $G_{\text{org}} = 4.5\,\text{g/m}^3$ and occluded in the deposit instead [4]. At the design dose of 2 g/m^3 the limit is exactly 55 %. Adsorption is exothermic, so G_{50} rises with temperature with an Arrhenius factor of 60 kJ/mol about the design temperature $T_0 = 62\,^\circ\text{C}$.

The leveler balance carries the hot-electrolyte damage,

$$V_{\text{cell}}\,\dot{G} = \dot{Q}_{\text{add}} - \left[F_e + V_{\text{cell}}\left(k_a\,\frac{I}{I_0} + k_h(T)\right)\right] G \tag{8.5}$$

where k_a is the consumption at the electrodes at the design current I_0 and k_h an Arrhenius hydrolysis rate of 0.5/h at T_0 with an activation energy of 100 kJ/mol. The dosing recipe $\dot{Q}_{\text{add}}$ replaces the flow, consumption and hydrolysis losses at T_0, and nobody measures G online, so a hotter cell loses leveler faster than the recipe replaces it. Temperature therefore reaches the deposit only where it physically acts, through the diffusivity inside u and through the leveler, and there is no separate penalty for cold

or hot electrolyte. The suspended solids obey $V_{\text{cell}}\dot{X} = F_e(X_{\text{in}} - X) + \dot{m}_{\text{sl}} I/I_0 - V_{\text{cell}} k_s X$, the filtered return plus the anode slime kept in suspension less what settles, which gives 4.1 mg/L at design.

8.3.3 Cell Voltage

Cell voltage is contact plus activation plus ohmic plus a concentration overpotential that diverges as the current approaches its limit,

$$V = V_c + V_a\,\pi + \frac{R_g T}{nF}\ln\frac{1}{1-u} + I\,\frac{d\,\pi}{\kappa A_c} \tag{8.6}$$

where V_c is the contact and busbar drop, V_a the activation overpotential, π a passivation multiplier, R_g the gas constant, T the absolute temperature, d the interelectrode gap (written d as in the app's Show Model, since g is the tip factor of Equation (8.3)), κ the electrolyte conductivity and A_c the cathode area. The logarithm raises the energy cost sharply as u approaches 1. The voltage is also the second constraint: the app trips a limit at 0.40 V, which Section 8.5.4 reaches through anode passivation rather than through current.

8.3.4 Design Point

The simulation clock is in hours. At design the rectifier delivers 29 440 A, which is 320 A/m^2 over the 92 m^2 of cathode area, into electrolyte at 62.0 °C carrying 45.0 kg/m^3 of copper at 1.2 m^3/h. The boundary layer is 140 µm and the limiting current density is 586 A/m^2, so the design utilization is 54.6 %. The cell plates 33.68 kg/h at 0.285 V and 249 kWh/t, with the leveling additive at 2 g/m^3, dosed at 8.3 g/h, so the cell sits a hair under its 55 % stability limit and the nodules neither grow nor level. That specific energy is in the range real tankhouses report [12], and refining is a small share of the energy a copper operation uses compared with comminution [13], which is part of why the pressure on a tankhouse is throughput rather than power.

8.3.5 Initialization and Assumptions

Initialization solves the thermal, copper, leveler and solids balances at the given actuator settings, so the initial condition is an exact equilibrium. A manual-mode helper, Hold this utilization, sets the current that produces the utilization setpoint in steady state on the *true* limiting current density at the present flow and heater settings, which is the calculation the plant cannot make.

The cell is one lumped tank with a single cathode area and no cell-to-cell variation, where a real tankhouse has hundreds of cells in series and its worst cell sets the crop. Impurity deportment, short-circuit detection and stripping are absent, and anode slime enters only as suspended solids that seed nodules. The cathodes are harvested, and the run stops, at the end of the 168 h cycle (120 h in the graded crop), where the crop is graded once on the nodules it carries. Integration is fixed-step RK4 at 0.005 h with a 72 s controller scan (Section 1.6.1). The copper assay reaches the controllers 2 h after the sample is taken, and the cell temperature reaches them through a thermowell lag of 0.1 h.

8.4 Controller Design

The three loops run the velocity-form PID of Section 1.6.2, Equations (1.10) to (1.13), with the derivative acting on the measurement [21], executed every $\Delta t = 0.02$ h. The integral is seeded at reset with the initial actuator value, and the tracking time of the back-calculation is $T_t = \tau_I$. With that choice the integral settles on the clipped output itself, so a loop held against a limit for several tracking times leaves it on the scan at which its error changes sign. Only the temperature loop carries a derivative term, and every scenario ships it with $\tau_D = 0$.

Cell temperature is controlled with a trim heater (action +1, reverse acting in the convention of Section 1.6.2), bulk copper with the electrolyte flow (action −1, direct acting, because more flow carries the anode's surplus away), and above them the constraint controller moves the rectifier current against the estimated utilization of (8.2) (action +1, reverse acting). The app's controller notes and loop hints use the same names. In this chapter the symbol a is the nodule height, so the action is written

Table 8.1 Settled state against utilization setpoint with all three loops in service, read at the end of the 168 h cathode cycle of scenario 4 with the setpoint changed before the run. σv is the nodule growth exponent, a the nodule height at harvest, Down the hour at which a crosses the 2 mm downgrade line, and Value the app's cycle value, with the crop graded at harvest. The 70 % row is limited by the 42 kA rectifier and reads 69.9 %. Rows from 62 % short the cell, at 141.8, 88.9 and 64.1 h. Bold marks the most valuable crop.

u (%)	I (kA)	V (V)	σv (/h)	Rate (kg/h)	a (mm)	Down (h)	Value ($)
45	22.50	0.234	−0.046	25.74	0.01	–	37920
55	29.94	0.289	0.000	34.25	0.09	–	49113
58	32.28	0.306	0.022	36.92	1.07	–	52584
58.5	32.67	0.309	0.026	37.37	1.82	–	**53167**
59	33.06	0.312	0.030	37.83	3.18	152.6	34846
60	33.86	0.317	0.039	38.73	10.6	124.9	35608
62	35.46	0.329	0.057	40.57	35	91.5	37140
66	38.72	0.353	0.101	44.29	35	59.4	40242
70	42.00	0.378	0.152	48.05	35	43.9	43392

as its sign. The tuned set of scenarios 3 to 10 is close to the SIMC settings [22] on the two regulatory loops, 40 %/°C and 2 h on temperature against 47.7 %/°C and 2.1 h, and 0.35 (m^3/h)/(kg/m^3) and 8 h on copper against 0.33 and 10.5 h at the 0.4 m^3/h flow stop, where the process gain is nine times its design value and the 2 h assay delay dominates. The utilization loop, at 120 A per percent and 1 h, is integral dominated with a closed-loop time constant near 5.5 h, slower than the temperature loop that moves its limit, though faster than the assay-limited copper loop. Scenarios 1 and 2 and the graded cycle load cautious gains of 8 %/°C and 8 h, 0.3 and 15 h, and 25 A per percent and 4 h.

The distinction the app is built on is that the first two loops act on measurements and the third acts on an estimate. All three look identical on a control drawing, all three report an integrated error, and only two of them can be verified against an instrument.

A voltage override sits beside the constraint controller and acts after it. When the current the controller commands would put the cell voltage above 0.40 V at the present temperature, concentration and flow, the override bisects the command down to the largest current that keeps the voltage under the limit. The constraint controller's integral is then reset so that its output before limiting equals that current, as the framework does in Manual (Section 1.6.2), and the controller resumes from it without windup. The override acts in manual mode as well and does nothing while the commanded current keeps the cell under the limit, which at the design point it does by 115 mV, so in most scenarios switching it off changes nothing. It is the one protection on the cell that acts on a measurable quantity rather than an inferred one, like the protective override on measured flow in Volume 4, Chapter 2.

8.5 Results

8.5.1 Locating the Stability Limit

Scenario 2 asks the reader to find the stability limit in manual mode. Table 8.1 shows why that is difficult: below the limit the growth exponent σv, the bracket of (8.3) times v, is negative or zero, the self-amplifying growth levels or stops, while particles still add height, and there is no other symptom. The cell voltage rises smoothly through the limit, from 0.289 to 0.306 V between 55 and 58 %, a change of six percent on an instrument with comparable noise. Above it the nodule trend sits on the floor for most of the crop and then takes off, because a nodule that doubles every day grows half of its final height in the final day. The first evidence is a rough plate at the end of a cycle.

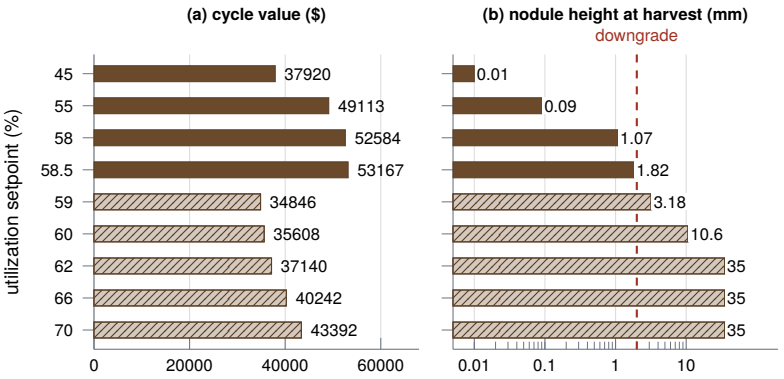


Figure 8.3 The utilization sweep of Table 8.1 as a chart. Solid bars are crops harvested below the downgrade line, hatched bars are downgraded. The value climbs to its peak at 58.5 % and then falls sharply at 59 %, the first setpoint whose nodules cross the 2 mm line, while every setpoint above it plates more copper and earns less.

Table 8.1 also prices the operating point over the 168 h cathode cycle. At 55 % the cell plates 5740 kg worth $49 113, at 58.5 % it plates 6213 kg worth $53 167 and is harvested with nodules of 1.82 mm, and at 66 % it plates 7250 kg worth $40 242. The 66 % cycle plates 17 % more copper than the 58.5 % cycle and earns 24 % less, because the nodules double every 6.9 h and cross the downgrade line at 59.4 h. A downgrade sells the whole crop at 65 % of the price, so the 59 % cycle, which crosses the line at 152.6 h, is worth $18 321 less than the 58.5 % cycle. The most valuable setpoint therefore sits 3.5 points above the stability limit, on a variable that is a ratio to an unmeasured quantity, and the sustainable ceiling over the cycle is 58.8 %, where a nodule doubles about once a day. Figure 8.3 draws the same sweep as a chart.

Scenarios 5 and 6 run the cycle at 59 and 66 %. The first is harvested at 3.18 mm, past the line it crosses at 152.6 h. The second is downgraded at 59.4 h and shorts at 88.9 h. Both crops sell downgraded, so the second earns $5396 more for its 969 kg of extra metal. The growth law also runs backwards. Lowering the 66 % setpoint to 55 %, the limit itself, stops the compounding, and the nodules then gain height only slowly, from the particles that keep seeding them, while lowering it to 50 % halves them every 25 h. Lowered at 59 h, with nodules of 1.92 mm, the crop is back under the line at 74.5 h, is harvested Grade A at 0.15 mm and is worth $50 524, $10 282 more than running on to the short. Lowered at 80 h, with 14.6 mm nodules, it is back under the line at 167.3 h, and at 85 h it is too late.

> **TRY IT IN THE APP** **6 · Push It to 66 %, the Cathode Is Scrap**
>
> Load the scenario and press Run. Watch the cathode nodule trend. It sits near the floor for days, then takes off, crosses the 2 mm downgrade line at 59.4 h and reaches the gap at 88.9 h, when the cell shorts. The crop plates more copper than at 58.5 % and still sells downgraded.

8.5.2 Additive Concentration

The leveling additive is the operator's defense against nodules, and its effect is not monotone.

Between 4.5 and 5 g/m^3 the stability limit in Table 8.2 falls eleven points, because the occlusion term of (8.4) overtakes the coverage term, and a crop that is harvested Grade A at 4.5 g/m^3 is downgraded at 61.4 h at 5 g/m^3. The coverage is still rising at 8 g/m^3 while the deposit is ruined, which is the behavior real tankhouses manage with thiourea and glue and is the reason additive is dosed to a concentration rather than to a rate [4]. Losing the additive is as costly in the other direction: with the pump stopped the limit falls toward 41.4 % and a crop at the design point is downgraded

Table 8.2 Additive setpoint at 58 % utilization over the 168 h cycle, read as in Table 8.1. θ is the leveler coverage and u_s the stability limit of (8.4). The setpoint slider stops at 8 g/m^3, and every downgraded row also shorts the cell.

G (g/m^3)	θ	u_s (%)	σv (/h)	a or Down	Value ($)
0.5	0.500	49.90	0.059	87.4 h	34104
2	0.800	55.00	0.022	1.07 mm	52584
4	0.889	56.51	0.011	0.28 mm	52563
4.5	0.900	56.70	**0.009**	0.24 mm	52558
5	0.909	45.74	0.089	61.4 h	34058
6	0.923	23.76	0.249	24.4 h	34048
8	0.941	−20.4	0.570	11.2 h	34028

within three days. At 58 % raising the dose from 2 to 4 g/m^3 halves the growth exponent and lowers the harvested nodule height from 1.07 to 0.28 mm, and changes the cycle value by $–21, the extra additive, because the crop is Grade A either way. The margin it buys pays only at a higher utilization or a dirtier electrolyte, which the graded cycle and scenario 11 use.

> **✦ DID YOU KNOW? The Most Leveler, the Fastest Ruin**
>
> From the simulator: in Table 8.2 the 8 g/m^3 dose gives the highest leveler coverage of any row, 0.941, yet its stability limit is −20.4 % and the crop is downgraded at 11.2 h. At 4.5 g/m^3 the same cell is harvested with nodules of 0.24 mm.

8.5.3 Effect of Temperature on the Limiting Current Density

Raising the cell temperature raises the diffusivity and therefore the ceiling of (8.1). At 58 % utilization over the 168 h cycle the limiting current density rises from 562.5 A/m^2 at 60 °C to 669.9 A/m^2 at 65 °C, and the rectifier follows it from 30.02 to 35.74 kA. The leveler balance (8.5) charges for it. The recipe was set at 62 °C, so the steady leveler falls from 2.21 to 1.71 g/m^3 over the same range, its coverage falls further because a warmer surface adsorbs less leveler at the same concentration, and the stability limit drops from 55.59 to 53.95 %, about 0.35 points per degree. The cycle value rises from 49 298 to $54 240 at 63 °C, where the crop is still harvested Grade A at 1.68 mm. One more degree reverses the result: at 64 °C the nodules double every 24 h instead of 27.5, the crop is downgraded at 156.5 h and the value falls to $36 223. At 68 °C, with the limit at 736.4 A/m^2 and the voltage still only 0.324 V, the crop is downgraded at 106.1 h and the cell shorts. Below the design point the heater runs out first, and the cell's own Joule heat holds it at 57.1 °C with the heater off against a 56 °C setpoint, where the recipe now runs the leveler rich and the nodules grow more slowly.

The temperature setpoint is therefore an optimum, with an edge that falls just above 63 °C at the design dose and moves with the dose, rather than a direction, and a supervisory optimizer built only on the limiting current would push past it. Lead anode corrosion and evaporation, which also accelerate with temperature, are outside the model.

8.5.4 Passivation and Purification Disturbances

Scenario 7 runs at 59 % and passivates the anodes at 12 h, which multiplies the activation and ohmic terms of (8.6) by 1.45 [14, 15]. The voltage steps to 0.419 V at the next physics step (0.402 V in the 0.05 h log), the override cuts the current from 31.8 to 29.6 kA at that scan, and the rectifier lag returns the voltage to the 0.400 V limit within about 0.1 h. Until the anodes are cleaned at 80 h the cell rides the limit at 29.5 kA and 350 kWh/t, against 265 kWh/t before, at a utilization of 54.4 % against its 59 % setpoint. Passivation enters neither the true limit nor its estimate, so the constraint controller

Table 8.3 Model errors with all three loops in service at a 59 % setpoint. Rows are read at 150 h in scenario 9 and at the end of each phase of scenario 10 (99.95 h and 168 h), whose estimate runs 12 % optimistic from 8 to 100 h and then 12 % pessimistic. Gap is true less believed, and σv the nodule growth exponent at the true utilization.

	Believed u (%)	True u (%)	Gap (pts)	σv (/h)
Diffusivity 20 % low	58.99	73.74	+14.75	0.140
Estimate 12 % optimistic	58.92	65.99	+7.07	0.099
Estimate 12 % pessimistic	59.13	52.03	−7.10	−0.019

is blind to it. With the interlock off the cell runs at up to 0.435 V and 380 kWh/t and finishes the 168 h cycle downgraded at 152.6 h, exactly as scenario 5 does, because passivation does not touch the deposit. The protected crop is harvested Grade A at 0.64 mm, because the override held the cell at 54.4 %, below its stability limit, from 12 to 80 h as a side effect.

Scenario 8 runs at 58 % and leans the purification circuit from 44 to 37 kg/m^3 between 10 and 70 h. The copper loop is direct acting, so an assay below setpoint cuts the electrolyte flow to retain the anode's dissolved surplus in the cell. The assay arrives 2 h late, the flow command reaches its 0.4 m^3/h minimum at 13.4 h, and the bulk copper settles at 39.0 kg/m^3 against a setpoint of 45.0 kg/m^3. The limiting current density falls from 590.5 to 365.4 A/m^2, a 38 % cut of which the leaner copper explains 13 points and the boundary layer thickened by the closed valve, from 139 to 195 µm, the rest, and the plating rate follows from 35.7 to 22.3 kg/h. The valve moves the limit faster than the utilization loop moves the rectifier, so the true utilization overshoots to 78.0 % at 13.4 h, where the nodules double every 6.7 h, while the belief gap peaks at only 4.34 points. The overshoot costs almost nothing, because (8.3) multiplies the nodules already on the plate and at 10 h they are 0.011 mm high: they reach 0.023 mm by 18 h. The same overshoot near the end of a crop run close to its ceiling would double nodules a millimeter high. The copper loop leaves its stop at 90.84 h, when the assay returns to setpoint, as the back-calculation of Equation (1.13) with $T_i = \tau_I$ predicts. The crop is harvested Grade A at 0.29 mm and finishes with 1288 kg less copper and $11 030 less value than the same case without the upset, a loss to a change upstream that no loop on the cell was configured to see.

🕮 **Pause.** Try Exercise 8.3 before reading on.

8.5.5 Model Error and the Belief Gap

Scenarios 9 and 10 are the central scenarios.

In the first row of Table 8.3 the true diffusivity is 20 % below the handbook value, so (8.1) overstates the ceiling by a factor of 1.25 and the constraint controller pushes the current until its *estimate* reads 59 %. The cell is at 73.74 %. The nodules at that utilization double every 4.9 h against every 23 h in the same case with the handbook right. The cathode is downgraded at 45.3 h, the cell shorts at 66.2 h when a nodule reaches the 35 mm gap, and the crop is worth $34 843, within $3 of scenario 5, which also sells downgraded.

Figure 8.4 shows the operator's view of that condition. At 150 h the temperature loop has an integrated absolute error of 0.243 °C h and the copper loop 3.737 kg h/m^3, both identical to the run with the handbook right, because the current, and with it every measured variable, follows the estimate. Every trend is flat while the cell is shorted by nodules. The app scores the utilization loop against the true utilization, which no plant can do, and that statistic alone records the error, at 2093 % h against 30 % h with the handbook right.

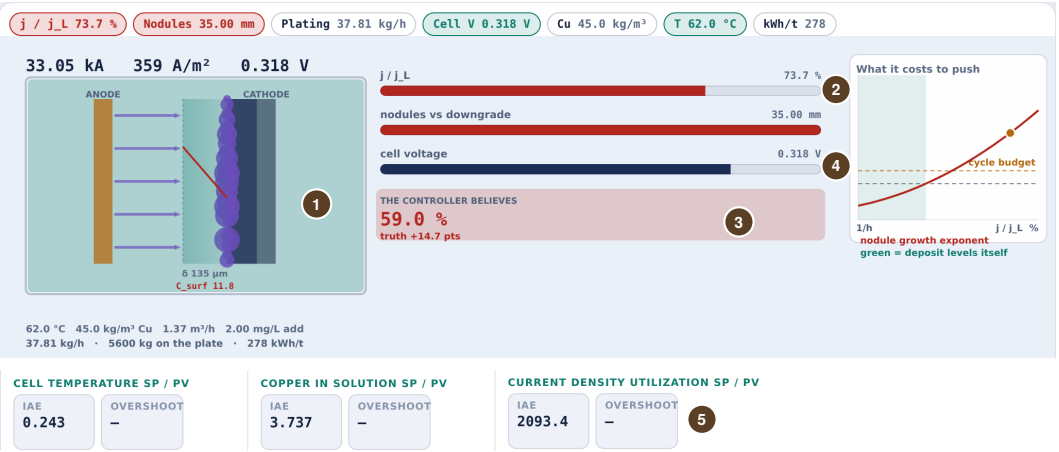


Figure 8.4 Scenario 9 at 150 h, with the true copper diffusivity 20 % below the handbook value the estimator uses, cut to the plant view and the integrated errors of the three loops. (1) The cathode, where a nodule has reached the 35 mm gap and shorted the cell. (2) The true utilization, 73.7 %, against (3) the believed utilization of 59.0 % that the constraint controller holds. (4) The cell voltage, which the override reads. (5) The integrated absolute errors: the temperature and copper loops report 0.243 and 3.737, and the utilization loop, scored on the true value, reports 2093.

> **TRY IT IN THE APP** **9 · The Model Is Wrong, and Nothing Says So**
>
> Run the scenario to 150 h and compare the Learn panel with the plant view. The temperature and copper loops report integrated errors of 0.243 and 3.737, the same as with the handbook right. The constraint controller believes 59.0 % while the cell runs at 73.7 %, and only the utilization score, taken on the true value, reads 2093.

The second and third rows are the two phases of scenario 10, in which the limiting-current estimate is biased 12 % optimistic for 92 h and then 12 % pessimistic. The optimistic phase runs the cell at a true 65.99 % while the estimate reads 58.92 %, the nodules double every 7 h, the cathode is downgraded at 63.2 h, and the cell shorts at 92.5 h, before the bias reverses. The pessimistic phase leaves the cell at 52.03 % while the controller believes 59.13 %, which plates 2304 kg from 100 to 168 h against 2569 kg in scenario 5, a loss of 10.3 % that grows no nodules and is therefore never flagged. The widest instantaneous gap of the run, 11.09 points at the physics step (8.73 in the 0.05 h log), opens at the reversal, when the estimate jumps while the plant has not yet moved. The two phases quantify the asymmetry that sizes the operating margin: past the stability limit the damage compounds and ends in a short that nothing undoes, and short of it the penalty is linear and recoverable. The reversal was moved from 60 to 100 h for this reason, because a pessimistic phase that begins before the short levels the nodules back down and hides the asymmetry.

Pause. Try Exercise 8.2 before reading on.

8.5.6 Dirty Electrolyte

Scenario 11 runs at 58 % and bypasses the return filter from 24 to 72 h, raising the suspended solids in the returning electrolyte from 2 to 40 mg/L. The solids in the cell rise from 4.3 to 27.4 mg/L, so the nucleation term of (8.3) adds nodule height at 8 % of the plating rate instead of about 1 %. Nothing the control system computes moves: the true and believed utilization agree, both loops hold their setpoints, and the growth exponent is the same as in a clean run. The nodules are 0.35 mm high at 72 h, but at three points above the stability limit every seed then doubles every 33 h, and the crop

Table 8.4 The graded 120 h crop. Rows after the first use constraint control and the tuned gains unless marked, written as u_{SP} (%), T_{SP} (°C) and G_{SP} (g/m^3) and set before the run. A capital letter marks a clause met (temperature, nodules, voltage, copper).

Configuration	Clauses	Cu (kg)	a (mm)	Peak V (V)	$\lvert\Delta T\rvert$ (°C)	Value ($)
As loaded	`t r V C`	4042	24.5	0.383	2.85	22396
60, 63, 4 (suggested)	`T R V C`	3857	0.33	0.368	1.17	32945
57, 63, 4	`T R V C`	3648	0.08	0.349	1.17	31153
62, 63, 4	`T R V C`	3998	1.04	0.381	1.17	34152
63, 63, 4	`T R V C`	4069	1.99	0.387	1.17	34758
64, 63, 4	`T r V C`	4140	3.93	0.393	1.17	22907
62, 63, 2	`T r V C`	3998	2.65	0.381	1.17	22131
62, 65, 4	`t R V C`	4176	1.57	0.386	3.00	**35661**
60, 63, 4, cautious	`t r V C`	3971	2.38	0.363	3.16	21977
Fixed current, tuned	`T r V C`	4042	25.8	0.386	1.18	22391

is downgraded at 149.6 h, three days after the filter came back, and harvested at 3.0 mm against 1.07 mm for the same case without the bypass.

> ⧖ HISTORY **Particles, Jets and Nodules (Before 1911)**
>
> By 1911 refiners knew that a higher current density improves the deposit only up to a point, beyond which it turns dark, loose and spongy. Jets of electrolyte aimed at the cathode renewed the solution at its surface, and in Graham's experiments a jet allowed a good deposit at 280 amperes per square foot, ruined at once when the jet was checked. Swan showed that solid particles in the electrolyte are one of the most fruitful causes of the nodular growths on deposited copper, and jets were liable to carry anode mud onto the cathode. Refineries stayed at 7.5 to 18 amperes per square foot [18].

The two remedies are not equivalent. Lowering the setpoint to 56 % only while the filter is out barely saves the crop, at 1.83 mm and $51 860, because the seeds are amplified after the filter returns rather than while it is out. Raising the leveler to 4 g/m^3 from 25 h lifts the stability limit to 56.5 %, stretches the doubling time to 65 h, and harvests the crop Grade A at 0.82 mm and $52 566, a smoother plate than the clean crop at the design dose, for 1.24 kg of additive and no copper. The constraint controller defends the constraint it computes, and a disturbance that spoils the product by another route has to be met by the operator who knows the mechanism.

8.5.7 Graded Cycle

The challenge runs one 120 h crop through a header that cools to 52 °C at 18 h, a purification circuit that drops to 40 kg/m^3 at 45 h, anodes that passivate at 72 h and recoveries at 96 and 104 h. It asks for the temperature within ±1.5 °C of setpoint throughout, nodules below the 2 mm line at harvest, a logged voltage never above 0.40 V and more than 3400 kg of copper on the case as given, and it loads on fixed current at 29.44 kA with the cautious gains.

Table 8.4 is a four-constraint problem in which the profit maximum sits outside the feasible set, which is the ordinary situation on a real plant and a rare one in a textbook.

As loaded the crop fails twice. Fixed current carries the cell across its stability limit when the header cools and the copper falls, so the cathode is downgraded at 87.3 h, and the cautious temperature loop lets the cell drift 2.41 °C low at 20.5 h and 2.85 °C high at 53.65 h with the heater inside its range. Tuning alone does not rescue it, because fixed current with the tuned gains holds the temperature within 1.18 °C and is still downgraded, at 86.0 h, and the suggested setpoints with the cautious gains

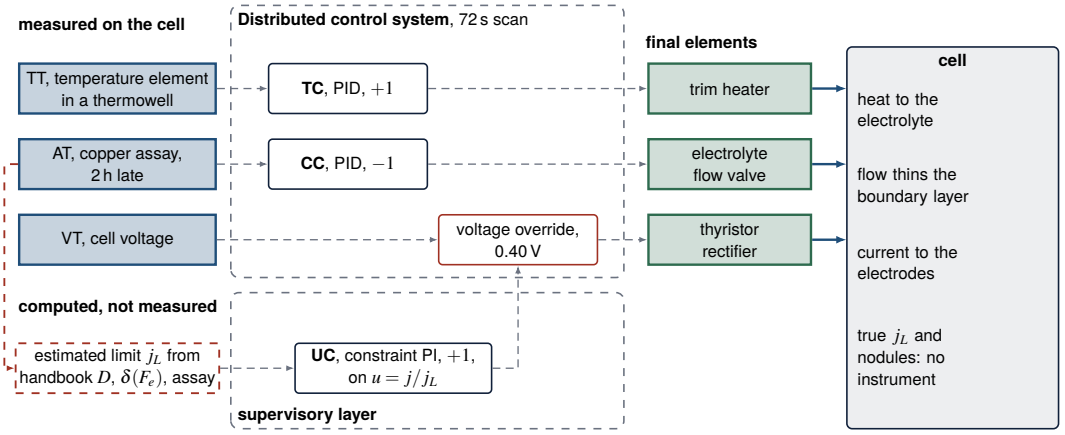


Figure 8.5 The loops of the simulated cell drawn with their hardware. The temperature element TT and the copper assay AT feed two PID blocks that move the trim heater and the electrolyte flow valve. In the supervisory layer the constraint controller UC acts on a limiting current density computed from a handbook diffusivity, a boundary-layer correlation and the assay (red dashes), and its command to the rectifier passes the voltage override, the one protection that reads a measurement.

fail both clauses. The suggested configuration, constraint control at 60 % and 63 °C with 4 g/m^3 of additive and the tuned gains, meets all four clauses. It leaves 1.67 mm of nodule margin, 32 mV of voltage margin and 0.33 °C of temperature margin, and revealing the suggested settings before Run or carrying them through Reset gives the same run.

The nodule clause sets the utilization. A setpoint of 63 % still passes at 1.99 mm, and 64 % is downgraded at 91.3 h. The additive moves that edge, because 62 % finishes at 1.04 mm with 4 g/m^3 and is downgraded at 104.2 h with the shipped 2 g/m^3. The 120 h ceiling is several points higher than the 168 h one, because the nodules have less time to compound. The temperature clause is scored from the first sample, so it limits how far the setpoint may move from the 62 °C at which the cell starts, and any setpoint at or above 64 °C fails at $t = 0$. The voltage clause never binds: the peak is 0.393 V at 64 %, the interlock never acts in any configuration tested, and switching it off leaves the suggested run unchanged.

The most valuable configuration tested, 62 % at 65 °C, is worth $35 661 and misses only the temperature clause. The best feasible one, 63 % at 63 °C, is worth 2.5 % less with 0.01 mm of nodule margin, and the suggested configuration is worth 7.6 % less with five sixths of the nodule allowance unspent, which is the margin the earlier scenarios argue for against an estimate that may be wrong.

8.6 Technology in Practice

Figure 8.5 draws the loops of Section 8.4 with their hardware.

What Measures It

The two regulatory loops act on ordinary measurements. The model reads the cell temperature through a thermowell lag of 0.1 h and the bulk copper from an assay that reaches the controllers 2 h after the sample is taken. The variable the constraint controller holds has no instrument at all. Its limit is computed, and the nodules it guards against are seen in a real tankhouse only once they short a cell.

TECHNOLOGY Sensor: Finding Shorts by Their Field and Heat

No instrument reads the limiting current density, but the damage done past it is found cell by cell. A shorted cathode carries abnormal current, and the traditional check is a hand-held gauss meter that senses the magnetic field around the cathode, often carried over the cells on a pole. Because shorts generate heat, infrared cameras with image processing are also used, though, like the gauss meter, they find short circuits and not open ones. A short is cleared by probing between the electrodes with a stainless steel rod and breaking off the nodule, often after lifting the cathode out, and many shorts are found only after significant damage has been done [23].

The model carries no short detection, so its scenarios run a shorted cell to the end of the cycle, where an operator would have found and cleared the short long before (Section 8.7).

What Moves It

The trim heater and the flow valve are first-order actuators, the flow held above its 0.4 m^3/h stop, and the rectifier is a lag of 0.02 h on the commanded current.

TECHNOLOGY Final Element: A Thyristor Rectifier

The rectifier turns the plant's alternating supply into the direct current that drives the cells in series. Large rectifiers deliver controlled currents in the kiloampere range, and thyristor phase-controlled units are the common type [24]. A controlled rectifier uses silicon-controlled rectifiers in place of diodes and sets its output by the moment in each cycle of the supply at which each one is triggered: the earlier the trigger, the longer it conducts [25]. Such units are judged on quality of control, harmonics, power factor, losses and efficiency [24]. A rectifier at its rating holds the current there, as the 42 kA rectifier of Table 8.1 does at 70 %.

The override lowers the command and never the measurement, so in scenario 7 the rectifier rides the 0.40 V limit at 29.5 kA until the anodes are cleaned.

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 [26]. Programmable logic controllers run PID loops as well, and the voltages of hundreds of electrolytic pots in the aluminum industry have been gathered to a central control computer [25]. The app runs its three PID blocks every 72 s, against a rectifier lag of 0.02 h and a thermal time scale of about 3 h. On such a platform TC and CC are PID function blocks, the limiting current density is a calculation block fed by the assay, and UC with the override is a supervisory application that writes the current setpoint of the rectifier.

IN PRACTICE What the Tankhouse Operator Watches

In practice the cell voltage of a copper refinery is normally held between 0.25 and 0.5 V, and a cell outside that range points to passivation of some of its anodes, which raises the energy used and lowers the cathode quality [27]. Short and open circuits are common in service. They come from poor anodes and cathodes, poor contact with the current supply, misaligned electrodes, and local differences in temperature, additive level or chemistry, so the cells are checked one by one and shorts are cleared by hand [23].

What Is Modernizing It

The first development is continuous monitoring of each electrode. Sensors that measure the magnetic field around the conductor carrying current to a cathode let a processor find open circuits as well as shorts, without anyone walking over the cells [23]. The second is the rectifier itself. Chopper rectifiers built from high-current insulated gate bipolar transistors are being considered as an alternative to the thyristor unit, and the pulse-width-modulated current-source rectifier of medium-voltage motor

drives is a future option [24]. The third is estimation. Inferential measurements are common wherever a controlled variable has no online sensor [5], and a moving horizon estimator that reconciled the voltage, the assay and the temperature could recover the diffusivity error of scenario 9 rather than absorb it, as Section 8.8 argues [28].

8.7 Limitations

There is one cell. A real tankhouse runs hundreds in electrical series, the current is common to all of them, and the crop is set by the worst cell rather than the average one. That single fact changes the control problem qualitatively, because the supervisory decision is about a distribution rather than a value, and nothing here represents it.

Impurity deportment is absent. Arsenic, antimony, bismuth and nickel decide whether the cathode passes a Grade A specification at least as often as physical roughness does, they interact with each other and with the additives, and their control is the principal reason a tankhouse has a purification circuit at all [17]. The app's quality variable is a nodule height and should be read as a physical proxy for grade.

Short-circuit detection and stripping are absent, and anode slime appears only as suspended solids with one nucleation coefficient for slime, fines and gypsum alike. Real short-circuit management is an infrared survey and a manual intervention, and the cell shorted in Figure 8.4 would have been found and cleared by an operator long before the cycle ended.

The boundary layer is a single power law in flow. Real mass transfer at a vertical cathode is driven by natural convection as well as forced flow, and the enhancement available from jetting or from air sparging is large enough to change the achievable current density substantially [16]. Equation (8.1) reproduces the right dependence and not the right number for any particular cell house.

The deposit law (8.3) is a linearized stability result for one representative nodule [3]. It holds for bumps smaller than the boundary layer and is stretched here all the way to the gap, so a real nodule much taller than the layer grows, and levels, more slowly than the model says, and a real cathode carries a population of nodules rather than one. The plate is graded on the nodules it carries at harvest, where a real overgrown nodule still holds the electrolyte it trapped, so the late rescues of scenario 6 are more generous than practice. The stability limit in a real cell depends on the additive package, the anode surface, the electrolyte impurities and the flow distribution, and locating it is an experimental campaign rather than a calculation [2].

> **OPEN CHALLENGE** **How a Nodule Crosses the Gap**
>
> Most studies of nodules describe their earliest stage, under 1 to 2 mm high. Experiments with seeded bumps found the height growing exponentially with time and suggested that only a nucleus of about 3 mm grows into a short within the operating time. Flakes of mold release agent from anode casting were found in nodules cut from shorted cells, but nodules without such a nucleus also form, for reasons not yet known [20]. Progress would be a measured account of what seeds the nodules that reach the anode, and a way to see them growing before they short.

8.8 Discussion

Three observations generalize beyond this cell.

The first is that an inferred controlled variable inherits every error in the inference, and the loop's own performance statistics say nothing about it. Table 8.3 has a controller with an excellent integrated error running a cell fifteen points above the utilization it believes. The remedy is not a better controller. It is either a measurement, or a periodic validation of the estimate against something that can be measured, or a margin large enough to absorb the estimate's uncertainty [5]. Validation against a late laboratory value is the bias update of Volume 5, Chapter 10 and Chapter 7. On this cell the only online candidate is the cell voltage, whose concentration overpotential in (8.6) makes it weakly informative about the true limit, and a moving horizon estimator that reconciled the voltage, the assay

and the temperature (Section 1.6.3) could recover the diffusivity error rather than absorb it [28]. The third remedy is priced by Table 8.1: the difference between 58.5 and 55 % utilization is $4054 per cycle, and that difference is the cost of the margin that protects against the estimate being wrong.

The second is that an instability that changes sign at a threshold turns a modest model error into a categorical outcome. Because the growth exponent of (8.3) is negative below the stability limit, a plant can operate for years with a badly calibrated estimator and discover it only when a change in flow, temperature or feed moves the true utilization across the line, after which the damage compounds out of sight until the last days of a crop. Constraints that are smooth degrade visibly, and constraints with a threshold do not. Scenario 11 adds the corollary that the threshold is not the only exposure: particles change how many seeds the instability has to amplify while every computed quantity stays where it was.

The third is that the profit maximum is often outside the feasible set. Table 8.4 shows the most valuable configuration violating a measured constraint, the best compliant one 2.5 % below it with almost no nodule margin, and the suggested one 7.6 % below it with five sixths of the nodule allowance unspent, which is the argument for putting the constraints into the optimization rather than leaving them to an override to discover [6, 19].

8.9 Conclusion

An electrorefining cell is controlled against a limiting current density that no instrument measures, and past a stability limit the nodules on the cathode grow in proportion to their own height, so the penalty for crossing it compounds. The app prices that shape over a 168 h cycle: 58.5 % of the limiting current density is worth $53 167 and finishes Grade A, 55 % is worth $49 113, and 66 % plates seventeen percent more copper and is worth twenty-four percent less. The central scenarios then make the estimate wrong rather than the controller. A handbook diffusivity twenty percent in error leaves a constraint controller holding a believed 58.99 % while the cell runs at 73.74 %, with two regulatory loops reporting integrated errors of 0.243 °C h and 3.737 kg h/m^3, unchanged from a run with the handbook right, and a cell shorted by nodules 58.2 h after the error appears. A bypassed filter downgrades a crop with no error in the estimate at all, and more leveler, not less current, saves it. The graded cycle places the profit maximum outside the feasible set, and the suggested configuration gives up 7.6 % of the most valuable one tested to meet all four constraints with margin. The app lets a reader make that trade, constraint control against an estimated limit, with realistic economics.

KEY POINTS

- More current makes more copper, but past a stability limit the nodules on the cathode grow in proportion to their own height.
- No instrument measures the limiting current density, so the constraint controller acts on an estimate and inherits every error in it.
- Over a 168 h cycle, 58.5 % of the limit finishes Grade A, while 66 % plates seventeen percent more copper and is worth twenty-four percent less.
- A handbook diffusivity twenty percent in error runs the cell at 73.74 % while the controller believes 58.99 %.
- A bypassed filter downgrades a crop with no error in the estimate, and more leveler saves it.

Where else this applies. Volume 5, Chapter 10 and Volume 5, Chapter 11 also control a variable that no instrument reaches, and Volume 4, Chapter 2 holds a rate against a limit with an override behind it.

References

[1] M. E. Schlesinger, M. J. King, K. C. Sole, and W. G. Davenport, *Extractive Metallurgy of Copper*, 5th ed. Oxford: Elsevier, 2011.

[2] M. J. Jaskuła, “Some remarks of the problem of current density optimization in copper electrorefining process,” *Electrochimica Acta*, vol. 28, no. 10, pp. 1395–1406, 1983.

[3] D. P. Barkey, R. H. Muller, and C. W. Tobias, "Roughness development in metal electrodeposition: II. Stability theory," *Journal of the Electrochemical Society*, vol. 136, no. 8, pp. 2207–2214, 1989.

[4] B. Veilleux, A.-M. Lafront, and E. Ghali, "Effect of thiourea on nodulation during copper electrorefining using scaled industrial cells," *Canadian Metallurgical Quarterly*, vol. 40, no. 3, pp. 343–354, 2001.

[5] L. Fortuna, S. Graziani, A. Rizzo, and M. G. Xibilia, *Soft Sensors for Monitoring and Control of Industrial Processes*, ser. Advances in Industrial Control. London: Springer, 2007.

[6] D. Hodouin, "Methods for automatic control, observation, and optimization in mineral processing plants," *Journal of Process Control*, vol. 21, no. 2, pp. 211–225, 2011.

[7] P. B. d. Moura Oliveira, J. D. Hedengren, and J. A. Rossiter, "Introducing digital controllers to undergraduate students using the TCLab Arduino kit," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 17 524–17 529, 2020.

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

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

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

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

[12] T. B. Braun, "Energy balances in the electrorefining and electrowinning of copper," *JOM*, vol. 33, no. 2, pp. 59–67, 1981.

[13] J. Jeswiet and A. Szekeres, "Energy consumption in mining comminution," *Procedia CIRP*, vol. 48, pp. 140–145, 2016.

[14] X. Ling, Z. H. Gu, and T. Z. Fahidy, "Effect of operating conditions on anode passivation in the electrorefining of copper," *Journal of Applied Electrochemistry*, vol. 24, no. 11, pp. 1109–1116, 1994.

[15] M. S. Moats and J. B. Hiskey, "The role of electrolyte additives on passivation behaviour during copper electrorefining," *Canadian Metallurgical Quarterly*, vol. 39, no. 3, pp. 297–306, 2000.

[16] H. Wang, Q. Wang, W. Xia, and B. Ren, "Effect of jet flow between electrodes on the cathode quality in copper electrorefining with high current density," *Metals*, vol. 8, no. 10, p. 833, 2018.

[17] B. A. Wills and J. A. Finch, *Wills' Mineral Processing Technology: An Introduction to the Practical Aspects of Ore Treatment and Mineral Recovery*, 8th ed. Oxford: Butterworth-Heinemann, 2015.

[18] "Copper," in *Encyclopædia Britannica*, 11th ed. Cambridge: Cambridge University Press, 1910, vol. 7, pp. 102–110, read in the Wikisource transcription (en.wikisource.org, 1911 Encyclopædia Britannica/Copper), accessed November 19, 2026.

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

[20] K. Adachi, Y. Nakai, S. Mitsuno, M. Miyamoto, A. Kitada, K. Fukami, and K. Murase, "The mechanism of nodular growth in copper electrorefining," *Journal of MMIJ*, vol. 136, no. 2, pp. 8–13, 2020, in Japanese, with an English abstract.

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

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

[23] E. Y. You, A. Berges-Dreyfous, D. B. George, D. Kim, and K. Zhou, "Monitoring electrolytic cell currents," U.S. Patent 7,445,696 B2, application filed March 17, 2005, 2008.

[24] J. Rodríguez, J. O. Pontt, C. Silva, E. P. Wiechmann, P. W. Hammond, F. W. Santucci, R. Alvarez, R. Musalem, S. Kouro, and P. Lezana, "Large current rectifiers: State of the art and future trends," *IEEE Transactions on Industrial Electronics*, vol. 52, no. 3, pp. 738–746, 2005.

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

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

[27] L. Junjie, D. Zhixuan, and Z. Yu, "Study on the optimal control of high lead anode copper electrolysis based on parallel flow technology," *Archives of Metallurgy and Materials*, vol. 70, no. 4, pp. 1965–1970, 2025.

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

8.10 Exercises

8.1 Concept check · 2 min

Which of these quantities does no instrument on the cell of Figure 8.1 measure, so that the loop holding it acts on a computed value?

(a) The cell temperature T, read through a thermowell
(b) The bulk copper concentration C, from the laboratory assay
(c) The cell voltage V, from the voltage transmitter
(d) The current density utilization $u = j/j_L$

8.2 Concept check · 1 min

Scenario 10 runs the same cell with its limiting-current estimate 12 % optimistic and then 12 % pessimistic. Why does the chapter conclude that a constraint controller should hold a margin below its estimated ceiling?

(a) Past the limit the nodules compound toward a short, while short of it only copper is lost
(b) The pessimistic phase plates less copper and also grows nodules, so both errors cost the same
(c) The optimistic phase plates more copper, so it is the cheaper of the two errors
(d) The voltage override cuts the current whenever the true utilization passes its setpoint

8.3 Concept check · 2 min

In scenario 8 the returning electrolyte drops from 44 to 37 kg/m^3 of copper. The copper loop runs with action -1. What does it do, and what does that do to the limiting current density?

(a) It opens the flow to bring more copper in, which thins the boundary layer and raises j_L
(b) It cuts the flow to keep the anode's surplus in the cell, which thickens the boundary layer and lowers j_L
(c) It cuts the flow, and the richer copper it retains raises j_L above its design value
(d) It does nothing until the assay arrives, because the loop is reverse acting

8.4 Theory and calculation · 6 min

Use Equation (8.1) with $n = 2$, $F = 96\,485\,\mathrm{C/mol}$, $D = 6.0 \times 10^{-10}\,\mathrm{m^2/s}$ and a molar mass of copper of 0.063 546 kg/mol.

(a) At design the bulk copper is 45.0 kg/m^3 and the boundary layer 140 µm. Compute the copper concentration in mol/m^3 to 1, j_L to 1 A/m^2, and the utilization at 320 A/m^2 to 0.1 %. *Checkpoint:* 708 mol/m^3; 586 A/m^2; 54.6 %

(b) The boundary layer follows $\delta_0 (F_{\mathrm{ref}}/F_e)^{n_\delta}$ with $F_{\mathrm{ref}} = 1.2\,\mathrm{m^3/h}$ and $n_\delta = 0.30$. By what factor, to 0.01, does it thicken when the flow closes to 0.4 m^3/h? *Checkpoint:* 1.39

(c) In scenario 8, j_L was 590.5 A/m^2 with 45.0 kg/m^3 and 139 µm. The upset leaves 39.0 kg/m^3 and 195 µm at the same temperature. Compute the new j_L within 2 A/m^2, the total cut in percent to 1, and the share of the cut due to the copper alone. *Checkpoint:* 365 A/m^2; 38 %; 13 points

8.5 Theory and calculation · 6 min

Use Equation (8.4) at the design temperature with $u_0 = 0.414$, $\Delta u_L = 0.17$, $G_{50} = 0.5\,\mathrm{g/m^3}$, $G_{\mathrm{org}} = 4.5\,\mathrm{g/m^3}$ and $\Delta u_{\mathrm{org}} = 1.0$ (the app's Show Model lists the last two).

(a) Compute the coverage θ to 0.001 and the stability limit u_s to 0.1 % at $G = 2$, 4 and 5 g/m^3. *Checkpoint:* 0.800, 55.0 %; 0.889, 56.5 %; 0.909, 45.7 %

(b) The cell runs at $u = 58$ % with the same current and boundary layer at 2 and at 4 g/m^3. Using Equation (8.3), compute the ratio of the growth exponents to 0.01. Table 8.2 gives $\sigma v = 0.022\,/\mathrm{h}$ at 2 g/m^3. What are the two doubling times, to 1 h? *Checkpoint:* 0.50; 32 h and 63 h

8.6 Theory and calculation · 6 min

The utilization loop sees an almost static plant: the rectifier lag is 0.02 h, and with j_L fixed, $u = 100\,I/(A_c j_L)$ in percent. Take $A_c = 92\,\mathrm{m}^2$ and $j_L = 586\,\mathrm{A/m}^2$.

(a) Compute the plant gain K_p in percent per ampere, to three significant figures. *Checkpoint:* 0.001 85 %/A

(b) For a PI controller with gain K_c and integral time τ_I on a static gain K_p, show that a setpoint step settles with the time constant $\tau_{cl} = \tau_I(1+K)/K$, where $K = K_c K_p$. Evaluate it, to 0.1 h, for the tuned set, 120 A per percent and 1 h. *Checkpoint:* 5.5 h

(c) Evaluate it for the cautious set of the graded cycle, 25 A per percent and 4 h, to 1 h, and compare it with the 120 h crop. *Checkpoint:* 90 h, most of the crop

8.7 App challenge · 8 min

Open scenario 5, whose name begins “Push it to 59%”, at the default seed. Run each part at 20 h per second until the run stops at the harvest, 168 h (the clock reads 7d 00:00), and read the Nodules and Copper this cycle tiles and the value in the harvest note.

(a) Run it as shipped, at 59 %. Also read the hour of the CATHODE DOWNGRADED mark, or the first row of the exported CSV whose `nodule_mm` reaches 2 (column Time (h)), within 0.1 h. *Checkpoint:* 3.18 mm; 6281 kg; $34 846; 152.6 h

(b) Press Reset, set Current density utilization setpoint to 58.5 % before Run, and run to the harvest. *Checkpoint:* 1.82 mm, Grade A; 6213 kg; $53 167

(c) By what percentage, to 0.1, does the 59 % crop exceed the 58.5 % crop in copper, and by what percentage, to 1, does it fall short in value? Explain the difference in one sentence. *Checkpoint:* 1.1 % more copper; 34 % less value

8.8 App challenge · 8 min

Open scenario 9, “The model is wrong, and nothing says so”, at the default seed. The true diffusivity falls 20 % below the handbook value at 8 h.

(a) Run at 20 h per second and pause at 150 h (the clock reads 6d 06:00), within 0.3 h. Read the Model error and Nodules tiles and the IAE tiles of the three loops: temperature and copper within 0.002, utilization within 5. *Checkpoint:* +14.7 points; 35.00 mm, shorted; 0.243, 3.737, 2093

(b) Open scenario 5, which loads the same settings without the diffusivity change, and read the same tiles at 150 h. *Checkpoint:* +0.0 points; about 1.85 mm; 0.243, 3.737, 30

(c) Which of these numbers could a real tankhouse compute, and what do they say about the two runs? *Checkpoint:* only the first two IAEs, which are identical

8.9 App challenge · 9 min

Open the challenge scenario, “one cathode cycle, one crop”, at its default seed. Run each part to the harvest at 120 h at 20 h per second. Read the challenge result in the progress panel (copper within 1 kg, nodules within 0.01 mm, peak voltage within 0.001 V, worst temperature error within 0.01 °C) and the value in the harvest note within $5.

(a) Run it as loaded. *Checkpoint:* fails: 4042 kg, 24.49 mm, 0.383 V, 2.85 °C; $22 396

(b) Reload the scenario with its reload button, press Reveal suggested tuning and run. *Checkpoint:* passes: 3857 kg, 0.33 mm, 0.368 V, 1.17 °C; $32 945

(c) Reload, press Reveal suggested tuning, then set Current density utilization setpoint to 63 % before Run, and run. By what percentage, to 0.1, is this crop worth more than the suggested

one, and how much nodule margin does it leave? Which would you run on a real tankhouse? *Checkpoint:* passes: 4069 kg, 1.99 mm, 0.387 V, 1.17 °C; $34 758; 5.5 %; 0.01 mm

8.10 Application area · 4 min

A tankhouse runs 400 cells in electrical series on one rectifier, so every cell carries the same current. The estimator describes the average cell, but a survey shows the worst cell's limiting current density 6 % below the average. Take the ceiling over a 168 h cycle as 58.8 % utilization, as in Section 8.5.1.

(a) What utilization setpoint, computed on the average cell, keeps the worst cell at its ceiling? Give it to 0.1 % and the value on the slider's 0.5 % grid at or below it. *Checkpoint:* 55.3 %; 55.0 %

(b) What share of the copper does that give up against running the average cell at its ceiling, to 1 %? *Checkpoint:* 6 %

8.11 Application area · 3 min

The cell runs at 0.285 V at design and at 0.309 V at 58.5 % utilization (Table 8.1). Take a cathode current efficiency of 0.965, electricity at $0.085 per kW h and copper at $8.60 per kilogram, as in the app's Show Model.

(a) Using Faraday's law, compute the specific energy at each voltage, to 1 kW h/t. *Checkpoint:* 249 and 270 kW h/t

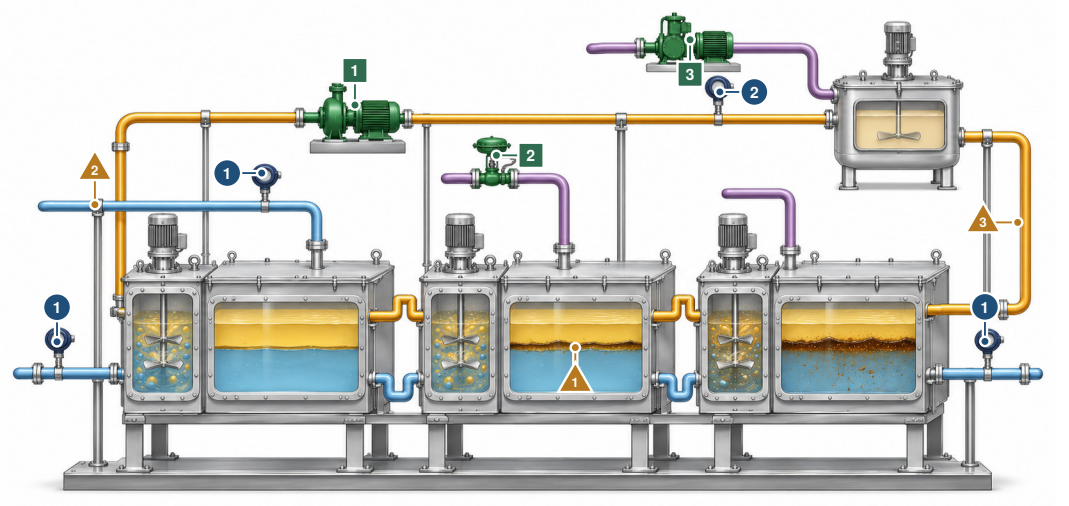
(b) What is the electricity cost per kilogram at 58.5 %, to $0.001, and its share of the copper value to 0.01 %? *Checkpoint:* $0.023 per kilogram; 0.27 %

8.12 Review the agent · 5 min

For scenario 11, an AI assistant writes: "The suspended solids only harm the cathode while the filter is bypassed, from 24 to 72 h. Lower the utilization setpoint from 58 to 56 % for exactly that window, which slows plating and therefore seeding. Raising the leveler is pointless, because the particles do not change the stability limit." Find the flaw and say how you would check it in the app.

CHAPTER 9

Two Time Scales and a Late Laboratory in Rare Earth Solvent Extraction



The controller moves 1 organic flow, held in ratio to the feed 2 scrub acid, held in ratio to the feed 3 caustic soda dose to the saponifier

The controller reads 1 feed, raffinate and product samples, analyzed every four hours 2 saponifier pH

What disturbs it 1 crud in one scrub stage 2 richer or leaner feed 3 extractant capacity loss

The organic picks up neodymium a little more readily than praseodymium in every unit, so the purity is made by the number of stages, and the days of metal they hold make the product slow to answer a change of feed. Illustration, not to scale.

IN PLAIN TERMS **Separating Two Nearly Identical Metals**

The main idea. Two rare metals used in strong magnets, neodymium and praseodymium, behave almost alike. A plant parts them in a row of seventy small tanks. In each tank an oily liquid takes up a little more neodymium than the watery liquid it meets.

The challenge. The tanks hold days of metal, so a change in what is fed reaches the product only after a day. It then keeps drifting for weeks. The only check on purity is a lab test every four hours that comes back late.

Why it matters. A plant that waits for that late test has lost days, so the flows must follow what is fed.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the fact box of the Introduction, the history box of Section 9.2, the try-it box of Section 9.4, the try-it and fact boxes of Section 9.6, 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 9.6, Section 9.7, Section 9.8 and the open challenge after it

The control question A seventy-stage neodymium and praseodymium cascade answers a change of scrub acid in about a day and a change of feed grade in weeks, and its only composition measurement is a laboratory analysis every four hours, reported an hour late. Which ratios, pairing, feedforward and loop speeds hold the product's purity and recovery, and where does the cascade's own capacity stop every controller?

Learning objectives

- Compute the relative gain of purity and recovery on the scrub and organic ratios, and show how it changes when the scrub ratio is keyed to the organic instead of the feed.
- Estimate the time scale of a feed-grade change from the inventory over the imbalance, and contrast it with the reflux response to a scrub change.
- Derive the material-balance feedforward from the feed analysis and say what the purity loop must still correct.
- Tune a product loop that acts once per laboratory report to the plant's day, and name its action in the convention of Chapter 1.
- Explain why 99.9 % purity at 97 % recovery is a capacity limit and not a stage limit.

Prerequisites Sections 1.4.4 (Kremser, Fenske, fast and slow modes), 1.4.5 (sampled assays), 1.4.6 (ratio and inferential control), 1.6.2 (velocity form and action) and 1.6.4 (scenarios). Section 1.6.3 is needed only for the predictive supervisor. Chapter 7 meets a late assay first.

The app https://apmonitor.com/apps/rare-earth-sx/

See also Volume 6, Chapter 3 for two quality loops on a sampled gauge with a negative relative gain on one pairing, and Volume 6, Chapter 2 for a pairing decided before any loop is tuned. Volume 5, Chapter 3 for composition control through a slow analyzer. Volume 3, Chapter 6 for a feedforward from the raw feed with the late product measurement as a slow trim.

9.1 Introduction

A 99.67 % neodymium product leaves the strip section of the simulated cascade at 74.8 kg/h. It is made from 2 m^3/h of a rare earth chloride solution that is 75 % neodymium and 25 % praseodymium, by 50 m^3/h of an organic phase that meets the aqueous in seventy mixer-settlers. In any one stage the organic carries neodymium in preference to praseodymium by a factor of 1.5, so a single contact of the feed would give an organic only about 82 % neodymium. The purity comes from the stage count, and the stage count makes the plant slow. The settlers and mixers hold 76.0 kmol of rare earth, 108.6 hours of feed.

The hard magnets that electric motors rely on include rare earth intermetallics, and their raw materials are expected to become a bottleneck that recycling is asked to relieve [1]. The separation of adjacent lanthanides is done industrially by solvent extraction, with up to hundreds of mixers and settlers assembled in series [2]. A dynamic equilibrium-stage model of a praseodymium and neodymium flowsheet predicted the steady profiles of a laboratory cascade well and its transients less well [3]. Control of industrial rare earth cascades has been posed as the control of the distribution profile of a component content along the cascade and studied with identified delays [4, 5], and dynamic models of mixer-settler units have been written for copper as well [6]. The control problem underneath has the shape of two-point composition control of a distillation column, where the choice of the configuration, which flows are manipulated and in what ratio, decides the interaction the controllers meet [7] and the relative gain array is the standard screen for it [8]. Volume 5, Chapter 3 closes a column's composition loop on a slow online analyzer, the same difficulty with minutes in place of hours.

DID YOU KNOW? Rare Earths Mined in 2024

In 2024 the world mined an estimated 390 000 tons of rare earths, counted as oxide, 270 000 of them in China. Neodymium oxide is priced at a purity of at least 99.5 %, the specification of this chapter, and averaged $56 per kilogram that year [9].

The simulator takes aim at two beliefs. The first is that a plant has one speed. Here a step in the scrub acid moves the purity 63 % of its way in 25.3 h at fixed flows, but a feed-grade step changes the material balance itself, and the purity of case 4 is still falling five days later toward a steady state of 90.9 %. The second is that the relative gain belongs to the plant. Keyed to the feed, the scrub ratio gives purity a relative gain of 3.92 on it and the crossed pairing −2.92, and keyed to the organic the same cascade gives 0.60 with both pairings positive. The same two lessons, a fast reflux mode beside a slow inventory mode and a pairing that depends on how each ratio is defined, apply to two-product distillation [7] and to other counter-current separations whose product is a small difference of two large flows.

These pieces seldom appear side by side. The industrial control studies work on plant data and identified models [4, 5], the laboratory model was checked on laboratory mixer-settlers [3], and the distillation results come from a different unit operation. This chapter tests what ratio stations, sampled proportional-integral (PI) loops, a pairing choice and a feedforward from the feed analysis can hold through feed changes, an extractant loss and crud, and where the cascade's own capacity stops every controller, using a 40/20/10 praseodymium and neodymium cascade simulated in a browser with its stage profile, laboratory samples and true products logged side by side. The control of rare earth cascades on monitoring stages and the dependence of the relative gain on the control configuration are both documented [4, 7], and the measurements of this chapter are properties of one teaching plant. The simulator is a two-phase stage stack with 2243 integrated quantities whose metal balances close to 10^{-11}, a steady-state solver that the dynamic stack reproduces to five decimals, and controllers confined to the pH, the flows and the delayed analyses. Section 9.2 follows the domain from ion exchange to counter-current cascades. Section 9.3 sets out the stage chemistry, the mixers, settlers and saponifier, the products and the measurements. Section 9.4 measures what each stage adds, the two time scales, the interaction of the two ratios and the capacity limit. Section 9.5 describes the ratio stations, the two sampled product loops, the feedforward and the optional predictive supervisor. Section 9.6 measures the ten cases, nine of which carry rules that grade them, with the challenge graded on its own verdict. Section 9.7 names the hardware of a real cascade, Section 9.8 states the limitations and Section 9.9 concludes.

9.2 From Ion Exchange to Counter-Current Cascades

The way from ion exchange to solvent extraction is told in the history box below, and the rest of this section follows the cascade theory, the extractant and the models.

HISTORY From Ion Exchange to Solvent Extraction (1947 to 1957)

The lanthanides were first isolated and purified by fractional crystallization, and ion exchange later replaced it as the way to make usable quantities of pure rare earths [10]. A paper of 1947 opened a series on the separation of rare earths by ion exchange with cerium and yttrium [11]. In the same years the solvent extraction of the trivalent rare earths was studied, and their order and degree of fractionation were reported in 1953 [12]. In 1957 the lanthanides were fractionally extracted as their dialkyl orthophosphates [13], the family of acidic organophosphorus cation exchangers to which the extractant of this chapter belongs. As demand grew, the low throughput and rising cost of ion exchange moved industrial separation to solvent extraction, which makes pure rare earths in large quantities, continuously and economically [10].

The design of a counter-current cascade for a pair with a small separation factor became a theory of its own, and a theory of counter-current extraction and its applications in the rare earth industry

Table 9.1 Variables, Their Roles and Their Values at the Design Point

Symbol	Role	Design	Range
x	CV, Nd purity, strip liquor	99.671 %	SP 99.0–99.95
R	CV, Nd recovery	98.75 %	SP 90–99.9
s	CV, saponified share	16.7 %	SP 10–18
W/F	MV, scrub per feed	1.173	0.8–1.6
O/A	MV, organic per feed	25	15–28
n_{OH}	MV, NaOH per organic (M)	0.167	0–0.18
S/O	set, strip per organic	0.012	0.010–0.03
z_{Nd}	DV, feed Nd share of the metals	75 %	60–90
F, c_{RE}	DV, feed flow, concentration	2 m^3/h, 0.35 M	1.5–2.5, 0.25–0.45
ℓ	DV, extractant capacity loss	0	0–30 %
κ	DV, crud in stage 55	0	0–90 %
$N_{k,n}$	states, mixer holdups (kmol)	profile	70 stages
$a_{k,n,j}$, $o_{k,n,j}$	states, settler cells (M)	profile	4 cells
$o_{k,sap}$	states, saponifier organic (M)	—	—
pH_n	output, measured pH of stage n	1.26 to −0.60	—
P	output, Nd product (kg/h)	74.8	—

was published in English in 1985 [14]. The simulator uses only the material balance any such cascade obeys, that the product draw is the extraction amount less the reflux returned by the scrub section. The phosphonic-acid extractant of the model, 2-ethylhexyl phosphonic acid mono-2-ethylhexyl ester, is the public name of a type sold as PC88A or P507, and it is described as the extractant of industrial rare earth separation facilities [3]. Reviews of rare earth solvent extraction list it among the commercially used extractants with D2EHPA [2], and separation factors of praseodymium and neodymium on it have been measured from process loadings down to trace concentrations [10, 15].

Process models followed the plants. A dynamic model of a praseodymium and neodymium flowsheet coupled equilibrium data with an acid balance and was checked against laboratory mixer-settlers [3]. Control of industrial cascades has been formulated as control of the profile of a component content along the cascade [4], and the identification of the process's delays for a neodymium content model has become a problem in its own right [5]. None of the sources consulted here describes the purity and recovery loops of a particular commercial plant, so the simulator's loop structure is a teaching design built on the column analogy rather than an account of how any plant is run.

9.3 The Cascade as Modeled

The cascade is a generic teaching design, drawn in Figure 9.1(a) in the arrangement of the application's Settlers view. Stages are numbered 1 at the raffinate end to 70 at the strip end, with 40 extraction, 20 scrub and 10 strip stages. The organic enters stage 1 from the saponifier and flows to stage 70, returning to the saponifier. The feed enters the mixer of stage 40 together with the scrub liquor leaving stage 41, and the aqueous flows down the extraction section to leave stage 1 as the raffinate. Scrub acid enters stage 60 and flows down the scrub section, and strip acid enters stage 70 and leaves stage 61 as the strip liquor, the neodymium product. The engine integrates in hours, flows are in m^3/h and concentrations in mol/L (kmol/m^3).

Table 9.1 gives every variable its role. The state is a holdup of praseodymium, neodymium, sodium and chloride in each mixer, four settler cells per phase per stage (four species in the aqueous, three in the organic) and the saponifier's organic, 2243 quantities in all, plus ten running totals of the flows and costs that cross the boundary.

9.3.1 Cation Exchange on a Dimeric Extractant

The extractant is held as dimer sites E, half its monomer concentration C_{HA}, and a trivalent metal takes three dimers and releases three hydrogen ions, while sodium takes one and releases one, as

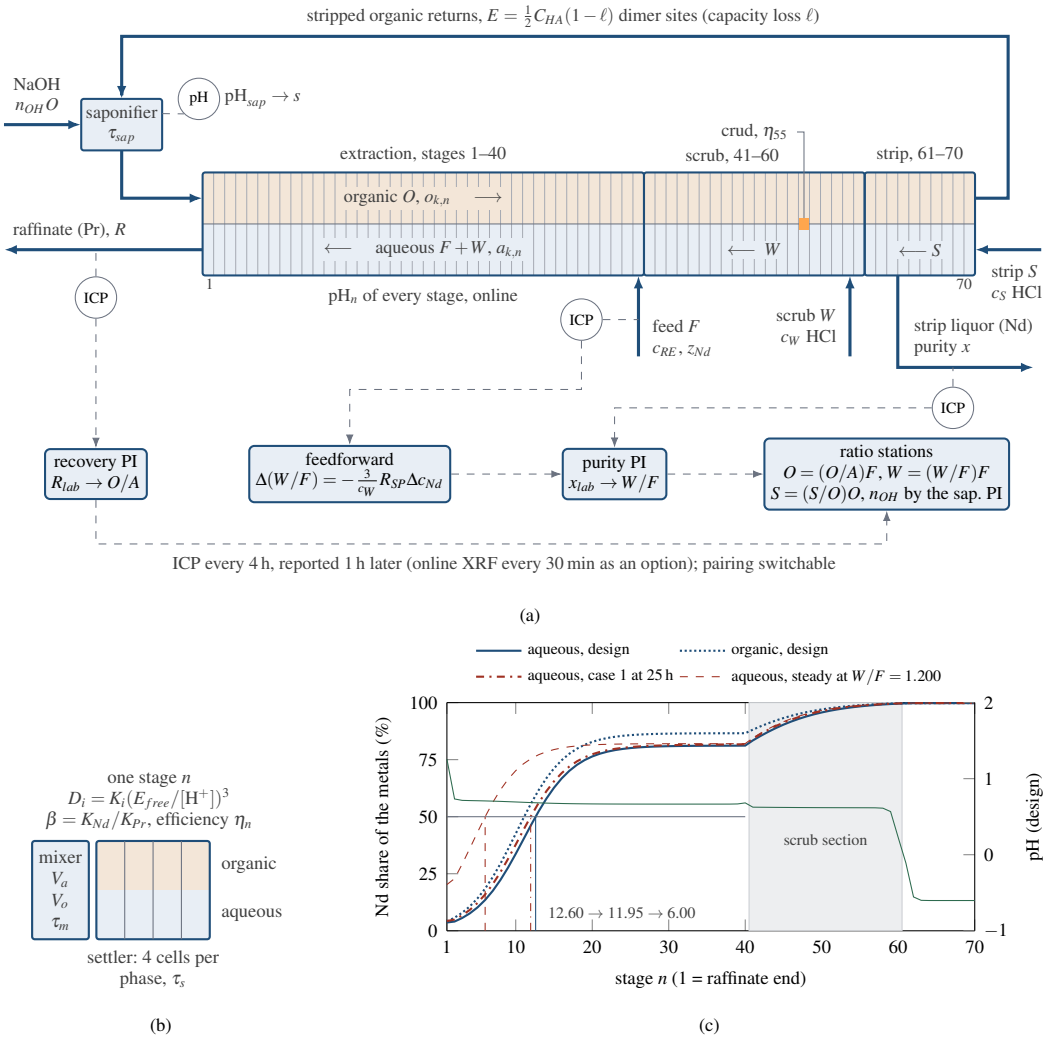


Figure 9.1 The 40/20/10 cascade as modeled. (a) The mixer-settler bank after the application's Settlers view, with the organic O from stage 1 to 70 and back through the saponifier, where NaOH at $n_{OH}O$ converts part of the extractant to its sodium form, the aqueous counter-current, the feed F at stage 40, scrub acid W at stage 60 and strip acid S at stage 70, the inductively coupled plasma (ICP) sample points on the feed, raffinate and strip liquor, the crud at stage 55, and the two sampled product loops with the feedforward (signals dashed). (b) One stage enlarged. (c) The neodymium share of each phase along the cascade at the design steady state (aqueous solid, organic dotted), the aqueous profile 24 hours after the scrub ratio of case 1 is raised from 1.173 to 1.200 (dash-dotted), and the steady state at 1.200 (dashed), with the stage where the aqueous is half neodymium (12.60, 11.95, 6.00, on the thin horizontal line at half) and the design pH (thin line, right axis). The scrub section is shaded. Process model only. Not to scale.

shown in Equation (9.1).

$$
\begin{aligned}
D_M &= \frac{o_M}{a_M} = K_M r^3, \qquad D_{Na} = K_{Na} r, \qquad r = \frac{E_{free}}{[H^+]} \\
E_{free} &= E - 3\sum_M o_M - o_{Na}, \qquad E = \tfrac{1}{2}C_{HA}(1-\ell) \\
[H^+] &= [Cl^-] - [Na^+] - 3\sum_M a_M
\end{aligned}
\tag{9.1}
$$

where M is Pr or Nd, a and o are aqueous and organic concentrations, chloride stays in the aqueous and the hydrogen ion follows from the aqueous charge balance. The separation factor $\beta = D_{Nd}/D_{Pr} = K_{Nd}/K_{Pr}$ is therefore the same at every stage, every acidity and every loading, and the organic holds at most $E/3 = 0.167\,\mathrm{M}$ of rare earth. Given a mixer's totals in its two phase volumes, the function $g(r) = E_{free}(r) - r[\mathrm{H}^+](r)$ is strictly decreasing, so the equilibrium is one root, found by a bracketed Newton iteration to a relative tolerance of 10^{-13} from the stage's last root. At the feed stage of the design point the root gives $[\mathrm{H}^+] = 0.208\,\mathrm{M}$ (pH 0.68), $E_{free} = 0.339\,\mathrm{M}$ and $D_{Nd} = 0.0906$ against $D_{Pr} = 0.0604$, and Equation (9.1) evaluated by hand from those concentrations agrees with the solver to six digits.

The neodymium constant comes from a laboratory measurement. A single-metal contact of 12.5 g/L neodymium with a 1 M extractant extracted 38 % at an equilibrium acidity of 0.13 M, a distribution ratio of 0.619 [10], and with the free dimer $(1 - 6 \times 0.0329)/2 = 0.40\,\mathrm{M}$ the expression $K_{Nd} = D[\mathrm{H}^+]^3/E_{free}^3$ gives 0.0212, used as 0.021. The same table's rows at 0.12 to 0.18 M acid give 0.019 to 0.044, and the constant falls as the acid rises, a dependence the model leaves out. The value $\beta = 1.5$ sits within the measured range. The thesis cites 1.5 to 1.6 from the literature for the pair, reports 1.29 to 1.59 in batch contacts across acidities and loadings (1.86 at 0.3 M acid) and 1.46 to 1.79 stage by stage in an 18-stage mixer-settler test [10], and contacts at trace concentrations in hydrochloric acid reached 1.75 to 2.05 [15].

9.3.2 Mixers, Settlers and the Saponifier

Each mixer holds the two phases at equilibrium in fixed volumes V_a and V_o, and a stage efficiency η acts as a bypass, as shown in Equation (9.2).

$$\begin{aligned} \frac{\mathrm{d}N_{k,n}}{\mathrm{d}t} &= \eta_n \left[A_n (a_{k,in} - a_k^*) + O(o_{k,in} - o_k^*) \right] \\ a_{k,out} &= \eta_n a_k^* + (1 - \eta_n)\, a_{k,in} \\ o_{k,out} &= \eta_n o_k^* + (1 - \eta_n)\, o_{k,in} \\ \frac{V_s}{4} \frac{\mathrm{d}c_j}{\mathrm{d}t} &= Q\,(c_{j-1} - c_j), \quad j = 1, \ldots, 4, \quad c_0 = c_{out} \end{aligned} \tag{9.2}$$

where $N_{k,n} = V_a a_k^* + V_o o_k^*$ is the mixer's holdup of species k, a^* and o^* its equilibrium, A_n the aqueous flow of the stage's section ($F + W$ in extraction, W in scrub, S in strip), and the settler of each phase is four well-mixed cells in series at constant volume V_s with $Q = A_n$ or O. Stage n takes its organic from the last organic cell of stage $n-1$ and its aqueous from the last aqueous cell of stage $n+1$, except that stage 70 takes the strip acid, stage 60 the scrub acid and stage 40 the mixture $(F a_F + W a_{41})/(F + W)$, with the feed composition $a_F = \left(c_{RE}(1 - z_{Nd}),\, c_{RE} z_{Nd},\, 0,\, 3c_{RE} + h_F\right)$ for Pr, Nd, Na and Cl. Volumes are sized at the design flows, $V_a = A_n \tau_m$, $V_o = O \tau_m$ and $V_s = Q \tau_s$. The crud event lowers one stage's efficiency, $\eta_{55} = \eta(1 - \kappa)$. The saponifier mixes the returning organic with the sodium that the NaOH loads, as shown in Equation (9.3).

$$\begin{aligned} V_{sap} \frac{\mathrm{d}o_{k,sap}}{\mathrm{d}t} &= O\,(o_{k,70} - o_{k,sap}) + \delta_{k,Na}\, \dot{n}_{NaOH} \\ \dot{n}_{NaOH} &= \min\left(n_{OH} O,\; 20 V_{sap} E_{free,sap}\right) \end{aligned} \tag{9.3}$$

where the cap keeps the dose within the free sites of about three minutes of organic, and the saponified share is $s = o_{Na,sap}/(2E)$, the fraction of the extractant monomer in the sodium form. The sodium form exchanges for rare earth in the first extraction stages without releasing acid, so the NaOH dose fixes how much rare earth the organic can take up.

9.3.3 Products, the Reagent Bill and the Steady State

The products are read at the last aqueous cell of stage 61 and of stage 1, as shown in Equation (9.4).

$$x = \frac{a_{Nd,61}}{a_{Nd,61} + a_{Pr,61}}, \qquad R = 1 - \frac{(F+W)\,a_{Nd,1}}{F\,a_{F,Nd}}$$
$$\frac{\$}{\mathrm{kg}} = \frac{p_N \dot{m}_{NaOH} + p_H \dot{m}_{HCl} + p_o e_o (F+W+S)}{S\,a_{Nd,61}\,M_{Nd}} \tag{9.4}$$

where $\dot{m}_{HCl} = M_{HCl}(W c_W + S c_S)$, e_o is the organic entrained per cubic meter of aqueous leaving and the prices p are teaching values. Recovery is measured from the raffinate's loss rather than from the product's neodymium, because the product-based yield counts the inventory the cascade gives up. In case 4 it reads up to 110.0 % at hour 26 while the raffinate-based recovery stays between 98.60 and 99.37 %. At the design point the bill is \$150/h of NaOH (8.35 kmol/h), \$69/h of hydrochloric acid and \$0.30/h of organic, \$2.93/kg of neodymium.

The same stage equations with the holdups removed give the steady state the cases start from and the Map view draws. Each stage's outlets are a function of its two inlets, and the residual $x - \Phi(x)$ over all stages is solved by pseudo-transient continuation with a block-tridiagonal Newton step whose Jacobian blocks are finite differences, with the organic recycle lagged one iteration. The design point converges to a residual of 1.4×10^{-13} in 49 iterations, and the steady balances close to 2.5×10^{-11} of the feed for praseodymium, 8.1×10^{-12} for neodymium and 2.1×10^{-13} kmol/h for the acid. The dynamic stack integrated for 6000 h after a scrub step and after a grade step reaches the solver's answers to five decimals, 99.69096 and 97.4675 % for the first and 90.89741 and 99.9998 % for the second.

9.3.4 Measurements

The controllers see a packet of measurements and nothing else, and a grep-level test in the application rejects controller code that reads a truth field. The pH of every stage's aqueous outlet is read every 5 min with a noise of 0.02, and the flows are metered exactly. The saponifier's pH carries 0.005 of noise and is converted to the saponified share through a declared titration curve, as shown in Equation (9.5).

$$\mathrm{pH}_{sap} = 6.5 + \log_{10}\frac{f}{1-f}, \quad f = \frac{o_{Na,sap}}{E}, \quad s_m = \frac{f}{2}$$
$$y_{lab}(t) = y(t_i) + \varepsilon_i, \qquad t_i + \theta \le t < t_{i+1} + \theta \tag{9.5}$$

where $t_i = i T_{lab}$ are the sampling instants, $T_{lab} = 4$ h and $\theta = 1$ h. The ICP analysis of the feed, raffinate and strip liquor reports the strip liquor's purity with 0.01 points of noise, the recovery computed from the sampled raffinate and feed with the metered flows (0.03 points), the feed's neodymium share and its neodymium concentration (0.5 %), each held until the next report. An online X-ray fluorescence option samples every 30 min, reports 6 min later and carries five times the ICP's purity noise. Section 1.4.5 gives the age of such a report. Every sample is kept with its true value at the drawing instant, and the challenge is graded on those true values.

9.3.5 Parameters and Numerics

Table 9.2 lists the parameters with their sources. The engine advances the stack by explicit Euler at $\Delta t = 20$ s, with concentrations clipped at zero and each stage's equilibrium root warm-started, and the ratio stations and the saponification loop act every 5 min. Halving the step moves the true purity and recovery of case 4 at hour 96 by 0.0012 and 0.0050 points. With one or four settler cells the steady purity agrees to eight digits, and the purity's 63 % time after a scrub step stays between 25.26 and 25.27 h at two, four and eight cells and at a 10 s step. Over 24 h with the capacity loss the sodium closes to 4.5×10^{-15} of the dose. With $\beta = 1$ both products carry the feed's 75 % to six decimals. Headless, the 120-hour challenge takes about 1 s (47 to 53 µs per step on two runs), and in the browser it plays in about 40 s at the top speed of 3 h per second.

Table 9.2 Parameters. S: From a Source. T: Teaching Value. A: Adjustable in the Parameters Panel or a Control

Symbol	Meaning	Value	
β	separation factor Nd/Pr	1.5	S, A
K_{Nd}	Nd constant of Eq. (9.1)	0.021	S, A
K_{Na}	sodium exchange constant	0.001	T
C_{HA}	extractant, as monomer	1.0 M	S, A
η	stage efficiency	0.95	T, A
τ_m, τ_s	mixer, settler residence	3, 12 min	S, A
τ_{sap}	saponifier residence	15 min	T, A
s	design saponification	16.7 %	S
c_W, c_S	scrub, strip acid (HCl)	3, 4 M	T, A
c_{RE}, h_F	feed metals, free acid	0.35, 0.01 M	S, T
z_{Nd}	feed Nd share	75 %	S
F	feed flow	2 m^3/h	T
—	organic capacity of the settlers	56 m^3/h (O/A 28)	T
—	saponification cap	18 %	T
—	titration center	pH 6.5	T
p_N, p_H, p_o	NaOH, HCl, organic prices	0.45, 0.20 $/kg, 3 $/L	T, A
e_o	organic entrained	20 mL/m^3	T, A
T_{lab}, θ	laboratory period, delay	4, 1 h	T, A
—	specification (purity, recovery)	99.5, 97 %	T

Sources: β, K_{Nd} (Table 4 of the thesis through its Eq. 4), the 1 M extractant, the 3-minute mixer needed for equilibrium, the 12-minute settler (a 480 mL settler at 40 mL/min of total flow), the 16.7 % saponification and the 50 g/L feed are taken from the thesis [10], and the 75/25 feed from the companion paper [3]. The thesis also states the manufacturers' guideline of at most 30 % of the stoichiometric loading [10]. On the 3:1 capacity of the model the design point loads 32.2 %, 15 % saponification would load 29.5 and the 18 % cap 34.2, so the cap is a teaching value near the guideline rather than the guideline itself.

9.3.6 Assumptions

Two metals stand for a plant that separates a dozen lanthanides pair by pair, and cerium, lanthanum and the heavy rare earths are absent. The stages are at equilibrium with an efficiency, so mass-transfer kinetics, a third phase, emulsions and aqueous carried in the organic are not modeled, and crud is only a loss of one stage's efficiency. Activities are concentrations and the separation factor and K_{Nd} are constants, while the measured constant falls with acidity. The scrub is hydrochloric acid, while the thesis's own flowsheet scrubbed with a neodymium solution in a 12-stage scrub circuit [10]. The settlers are four tanks in series at fixed volumes, an approximation of a plug delay that keeps the inventory closed when a flow changes. The titration curve, the prices and the entrainment are teaching values, and nothing in the model is plant data.

9.4 One Stage, Seventy Stages and Two Speeds

9.4.1 What Each Stage Adds

At total reflux the Fenske relation [16] of Section 1.4.4 would add $\ln\beta = 0.405$ to the log-odds $\ln(x/(1-x))$ of the purity at each stage. At the design ratios the steady-state solver gives 95.49, 97.69, 98.79, 99.37, 99.67, 99.83 and 99.91 % for 4, 8, 12, 16, 20, 24 and 28 scrub stages, a nearly straight line of 0.167 per stage. The scrub section runs far from total reflux. Its organic carries 2.681 kmol/h of metal into stage 41 and its aqueous returns 2.161 kmol/h as reflux, so the product is 0.520 kmol/h, a fifth of the extraction amount, and the reflux ratio is 4.15. The Kremser factor of a section at those flows (Equation (1.7) uses the same factor), $\beta W_m/O_m = 1.209$, gives $\ln 1.209 = 0.190$ per stage, and 0.181 with the stage efficiency taken as a bypass, against an organic log-odds gain of 0.185 to 0.195 per stage inside the section. The staircase's 0.167 is lower, likely because adding scrub stages at fixed flows also shifts the profile at the section's two ends, where the feed and strip entries break the

constant-factor pattern. At $\beta = 1.3$ the cascade lines out at 95.98 % and each stage adds 0.045, close to $\ln(1.3 \times 0.806) = 0.047$, and at $\beta = 1$ both products are the feed.

9.4.2 A Fast Mode and a Slow Mode

The cascade holds 76.0 kmol of rare earth. A change of the scrub ratio redistributes the reflux without changing the neodymium balance at the moment of the step. Raising W/F from 1.173 to 1.178 leaves the neodymium fed and drawn at 0.5250 kmol/h each, and the neodymium inventory needs only 2.57 kmol more to reach the new steady state. At fixed flows the purity moves 63 % of its way in 25.3 h and 95 % in 76 h, overshooting its final value by 1.3 % of the step near 200 h, while the recovery takes 390 h to move 63 %, as shown in Figure 9.2(a). On the engine, with the saponification loop and the instruments, the purity's 63 % time is 27.1 h.

A feed-grade step changes the balance itself. At 67.5 % neodymium the feed brings 0.4725 kmol/h while the unchanged draw takes 0.5250 kmol/h, and the new steady state holds 37.7 kmol less neodymium. Inventory over imbalance is 718 h. Measured, the purity of case 4 is flat for a day, crosses 99.5 % 81.7 h after the change, reaches 63 % of its fall at 877 h and is 91.80 % at 1000 h on its way to 90.90 %, as shown in Figure 9.2(b). The recovery moves first and the wrong way, because the feed changes at once and the raffinate does not. A settling time of about 30 h for a 10 % feed step would follow from the 18-stage laboratory rig of the thesis, which was run that long to reach steady state [10], but it does not carry to seventy stages with 108.6 hours of feed in them, so the cases teach the day-long reflux response and the weeks-long inventory response instead.

> **TRY IT IN THE APP** **4 · The Long Transient: A Leaner Feed Takes Days to Arrive**
>
> Load the scenario and press Run. Watch the purity trend. The feed turns leaner at hour 4, yet the purity stays flat for a day and only then begins to slide. The true purity crosses 99.5 % at hour 85.7, and the first laboratory report below it comes from the sample of hour 88.

The profile shows the slow mode too. In case 1 the scrub ratio is raised from 1.173 to 1.200 at hour 1. A day later the purity has risen from 99.671 to 99.733 %, most of its way, while the stage where the aqueous is half neodymium has moved only from 12.60 to 11.95 toward the 6.00 of the new steady state, as shown in Figure 9.1(c). The parting stage lies in the extraction section and follows the inventory, and it reads 11.24 at hour 49.

Pause. Try Exercise 9.2 before reading on.

9.4.3 Interaction, and What the Ratio Is Keyed To

The steady-state gains at the design point, from upward steps of 0.0005 in W/F and 0.02 in O/A, are $\partial x/\partial(W/F) = 0.0416$, $\partial x/\partial(O/A) = -0.00173$, $\partial R/\partial(W/F) = -2.55$ and $\partial R/\partial(O/A) = 0.142$ in fractions per unit ratio, and the relative gain of purity on the scrub ratio follows as shown in Equation (9.6).

$$\lambda_{11} = \left(1 - \frac{K_{12}K_{21}}{K_{11}K_{22}}\right)^{-1} = 3.92 \tag{9.6}$$

where K_{ij} is the gain of output i (purity, recovery) on input j (W/F, O/A). The knee of the purity and recovery trade is curved, so downward steps give 3.60 and central differences 3.76, and four other points on the knee give 3.72 to 3.82. The pairing of purity with the scrub ratio and recovery with O/A is the one to use, and the crossed pairing has $1 - \lambda_{11} = -2.92$. The sign test agrees. The gain of purity on O/A is -0.00172 open and $+0.00062$ with the recovery held by the scrub ratio, so it changes sign when the other loop closes [8]. Volume 6, Chapter 3 meets the same structure on a paper machine, two quality loops that act only when a sampled gauge reports and a relative gain that is negative on one pairing.

The number belongs to the handles as well as to the plant. With the scrub keyed to the feed, more organic raises the extraction amount while the reflux stays, so the product draw, their small difference,

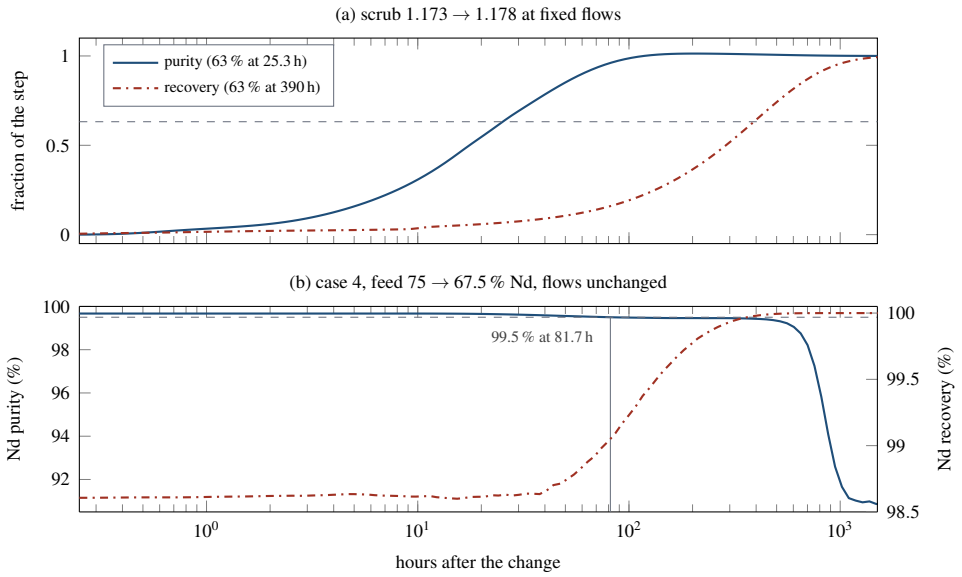


Figure 9.2 The two time scales. (a) The fraction of its step that the purity and the recovery have moved after the scrub ratio is raised from 1.173 to 1.178 at fixed flows and NaOH, on the dynamic stack alone, purity solid and recovery dash-dotted, with the 63 % line dashed. (b) Case 4 on the engine, the feed cut from 75 to 67.5 % neodymium with every flow unchanged, the true purity (solid, left axis) and recovery (dash-dotted, right axis) to 1500 h, with the 99.5 % line dashed.

grows and both products respond. Keyed to the organic, as scrub per cubic meter of organic, more organic brings more reflux with it, and the same cascade gives $\lambda_{11} = 0.62$ with upward steps and 0.60 by central differences, both pairings positive and neither rejected. Keying the ratio is the configuration choice of two-point distillation control, where the steady-state relative gain of the same column changes with the choice of flows and their ratios [7]. The application keys the scrub to the feed, the NaOH and the strip acid to the organic, and its pairing lesson is honest for those handles.

9.4.4 A Capacity Limit, Not a Stage Limit

The scrub ratio buys purity with recovery along each O/A curve, as shown in Figure 9.3. At 97 % recovery the purity is 99.04, 99.63, 99.70 and 99.82 % at O/A 20, 24, 25 and 28, and 99.87 % at O/A 28 with the 18 % saponification cap. A sweep of 49,771 steady states over O/A from 15 to 36, saponification from 10 to 24 % and four strip ratios finds the best purity at 97 % recovery within $O/A \leq 28$ and 18 % to be 99.866 %, so 99.9 % at 97 % is out of reach within the equipment limits. Past them it is reachable, at O/A 33 and the design saponification with $W/F = 1.625$ (99.902 %) or at 21 % saponification and O/A 26 (99.903 %). The limit is the extractant's and the settlers' capacity, not the stage count. The NaOH station stops at 0.18 kmol/m^3, the cap on the nameplate extractant, so the second point cannot be reached by hand. A station without that stop would reach it. The cheapest steady state that just meets 99.5 % and 97 % at the design saponification sits at O/A 23 and \$2.69/kg, against \$2.93/kg at the design point with its margins.

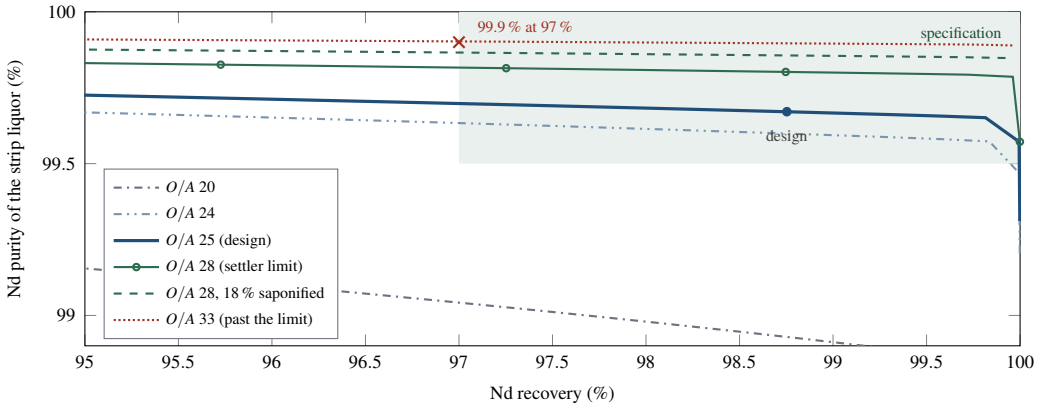


Figure 9.3 Purity against recovery from the steady-state solver at the design feed, each curve one O/A swept across the scrub ratio at 16.7 % saponification unless marked. The shaded corner is the teaching specification. The cross at 99.9 % purity and 97 % recovery lies above every curve inside the settlers' organic capacity and the saponification cap and below the curve at O/A 33.

9.5 Control as Implemented

9.5.1 Ratio Stations and Saponification

Every 5 min the first layer sets the flows from the metered feed, as shown in Equation (9.7).

$$\begin{aligned} O &= \min\big(56,\ (O/A)\,F_m\big), \qquad W = (W/F)\,F_m \\ S &= (S/O)\,O, \qquad n_{OH} = \mathrm{PI}_s(s_{SP} - s_m) \end{aligned} \tag{9.7}$$

with the NaOH commanded at $n_{OH}O$, so the NaOH and the strip acid follow the organic at once and the saponification PI ($K_c = 0.1$ kmol/m^3 per unit, $\tau_I = 0.3$ h) trims the NaOH per cubic meter between 0 and 0.18 on the share read from the saponifier's pH. The stations clamp O/A to 15–28 and W/F to 0.8–1.6, and the organic stops at 56 m^3/h, the settlers' capacity at the design feed.

9.5.2 Two Sampled Product Loops

Each product loop is the velocity-form PID of Section 1.6.2 with the derivative off, executed once per fresh analysis and holding between analyses, as shown in Equation (9.8).

$$\begin{aligned} e_k &= \sigma\,(y_{SP} - y_{lab,k}), \qquad \Delta t_k = t_k - t_{k-1} \\ S_k &= S_{k-1} + K_c(e_k - e_{k-1}), \qquad \hat{u}_k = S_k + I_{k-1} \\ u_k &= \min\big(u_{max} - u_{ff},\ \max(u_{min} - u_{ff},\ \hat{u}_k)\big) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I}e_k + \frac{u_k - \hat{u}_k}{T_t}\right]\Delta t_k \\ T_t &= \max(\tau_I,\ \Delta t_k) \end{aligned} \tag{9.8}$$

where t_k is the report time, σ is the action a of Equation (1.10), the error is in fractions with K_c entered per point, u is the ratio before the feedforward u_{ff}, and the applied ratio is $u_k + u_{ff}$. The integral used at report k was accumulated through report $k-1$, so a report's error reaches the ratio through the proportional term at once and through the integral four hours later. In the pairing of the relative gain, purity drives W/F and recovery O/A, both with $\sigma = +1$, reverse acting in the convention of Section 1.6.2, because each product rises with its ratio. The crossed pairing drives O/A from purity and W/F from recovery, both with $\sigma = -1$. Both are direct acting, because purity falls as O/A rises and recovery falls as W/F rises. The application's source labels these two loops direct

acting too, and its control panel names the pairing Purity→O/A. The shipped tuning is $K_c = 0.01$ per point with $\tau_I = 40\,\text{h}$ for purity and 0.05 with 400 h for recovery, and the cases set their own.

Pause. Try Exercise 9.3 before reading on.

Three details guard paths that no shipped case takes. The tracking time is $\max(\tau_I, \Delta t_k)$. With τ_I alone, a loop with $\tau_I = 1\,\text{h}$ at its scrub limit would be kicked from 1.6 to 1.23 at the next report while that report was still below setpoint. A pairing switch during a run re-initializes each PI on its new loop's error at the applied ratio. Without it, the first report after a switch would apply the jump between the two loops' proportional terms, $\Delta(W/F) = -0.1125$, against the -0.0130 applied. The NaOH station stops at its cap, as described above.

9.5.3 Feedforward From the Feed Analysis

The neodymium the product must carry is the neodymium fed times the recovery. Each mole of neodymium that the scrub acid returns down the scrub section takes three moles of hydrochloric acid, so at each feed analysis the scrub ratio moves as shown in Equation (9.9).

$$u_{ff} = -\frac{3}{c_W} R_{SP} \left(c_{Nd,lab} - c_{Nd,ref} \right) \tag{9.9}$$

where $c_{Nd,ref}$ is the feed analysis when the feedforward is switched on, which keeps the draw in step with the neodymium fed while the organic and the extraction amount stay. For the leaner feed of cases 4 and 6 the law asks $+0.0259$, and the part measured in a run is 0.0246 because the feed analysis carries its noise. The switch is bumpless, because the PI's integral takes over the feedforward's part when it is turned off. Volume 3, Chapter 6 makes the same move, a dose computed from the raw feed with the late product measurement demoted to a slow trim.

9.5.4 Predictive Ratio Supervision

The optional model predictive controller follows Section 1.6.3 and uses four local lag paths, rather than identifying a model inside a five-day case. Central steps of $\Delta W/F = \pm 0.0005$ and $\Delta O/A = \pm 0.02$ were observed for 1000 h at nominal chemistry with product loops manual and saponification automatic. Instrument precursors provide the identification outputs only. Online control uses the selected reported assays and metered flows. In input order $(W/F, O/A)$ and output order (purity, recovery), the fitted percentage-point gains and hour time constants are

$$\begin{aligned} K &= \begin{bmatrix} 4.1994 & -0.17229 \\ -249.108 & 13.6935 \end{bmatrix}, \\ \tau &= \begin{bmatrix} 26.9067 & 25.5122 \\ 507.155 & 484.893 \end{bmatrix}, \qquad G_{ij}(s) = \frac{K_{ij}}{1 + \tau_{ij} s}. \end{aligned} \tag{9.10}$$

These are finite-window dynamic fits, not established infinite-time gains. A fitted 0.2365 h delay on purity from scrub is omitted. RMS errors of the implemented paths are 0.00000765 and 0.0000169 purity points, and 0.00116 and 0.00292 recovery points. Saponifier noise remains in the experiments and cancels only in the central difference. The four states advance by exact exponential lags on measured ratios every 5 min.

Each fresh selected assay corrects an output bias by $b_i^{+} = b_i + 0.3[y_{i,lab} - \hat{y}_i(t_{sample}) - b_i]$, limited to ± 2 purity points and ± 5 recovery points. The nominal output at the sample time is interpolated from saved five-minute predictions. Comparing at report time would mix process motion with model error, the point of Equation (1.8). This is a bounded residual correction rather than a convergence guarantee for a changed gain [17]. No current stage profile, chemistry, capacity loss or unreported feed analysis enters the predictor.

Every 4 h, the quadratic program covers 480 h in 120 steps, with twelve graduated blocks. In percentage points and unscaled ratio units, the objective has output weights $(100, 1)$ and block-change

Table 9.3 The Ten Cases

#	Settings and events	Teaching point
1	Manual at design, learner raises W/F to 1.200	the profile and its pinch
2	Manual, β set to 1.30 and Reset	stages and separation factor
3	Manual, learner raises W/F to 1.185	the trade and its limit
4	Manual, feed 75 to 67.5 % Nd at 4 h	the long transient
5	Purity loop 1.5 per point, 4 h, feed to 82.5 % at 6 h	a loop on stale reports
6	Purity loop 0.05, 24 h, feed to 67.5 % at 6 h	the feedforward
7	Both loops, crossed pairing (1, 24 h and 0.05, 24 h)	the relative gain
8	Feedforward, purity loop 0.05, capacity −10 % at 6, crud at 30 h	the bank before the laboratory
9	Higher purity and lower recovery requested at 8 h	coordinated ratios
10	Challenge, 120 h, feed 82.5 % at 12 and 67.5 % at 48, crud at 72 h	the whole problem

Table 9.4 Case Outcomes Measured on the Engine

Case and condition	Result
1, W/F 1.200 at 1 h	parting 12.60 / 11.95 / 11.24 at 1, 25, 49 h, x 99.733 %, R 98.22 % at 25 h
2, β 1.5 / 1.3 / 1.0	x 99.67 / 95.98 / 75.00 %
3, W/F 1.185 at 1 h	x 99.700 / 99.710 %, R 98.54 / 98.43 % at 25 / 49 h, \$2.93 to \$3.04/kg
4, as shipped	first report after the step R 98.63 against 98.78 %, x below 99.5 % at 85.7 h, 99.476 % at 120 h
5, shipped / 0.05, 24 h	W/F 1.015–1.201 / 1.1665 to 1.1449, x 99.52–99.67 / 99.74 %
6, feedforward off / on	x 99.510 / 99.591 % at 96 h, \$2.94 / \$3.08 per kg over 96 h
7, crossed / paired	W/F 1.2013 / 1.155, O/A 25.472 / 24.646, x 99.702 / 99.661 % at 96 h
8, capacity loss	s_m 18.4 % at once, ICP R 98.20 % at 9 h, true R 87.17 % at 48 h

weights $(10^4, 10)$. Predicted purity and recovery floors are 99.5 and 97 %. One slack per output, shared across the horizon, costs $10^5\xi_x^2 + 10^4\xi_R^2 + 10^3(\xi_x + \xi_R)$, with ξ_x and ξ_R the slacks. The ridge is 10^{-8}. Ratio limits are $0.8 \le W/F \le 1.6$ and $15 \le O/A \le 28$, with the organic hydraulic cap retained, and block changes are at most 0.002 and 0.05. A manual product loop fixes its assigned handle. There is no reagent objective or feed-composition feedforward in this MPC, unlike the classical balance correction.

The ratio PIs track the applied commands. A rejected solver result holds both ratios on the first failed scan, then returns to sampled PI. Validation checks the whole planned sequence, diagnostics and bounds. A valid solve resumes prediction, and leaving the option clears its model history. These choices implement receding-horizon supervision [18] while the ratio stations and saponification loop remain active.

9.6 Ten Cases on the Running Cascade

The cases are listed in Table 9.3 and their outcomes in Table 9.4. Every value was measured with the application's engine of Section 1.6.1 stepped headless at its 20 s step and the default seed, with each case loaded as the browser loads it and every learner move made at the hour its text names. A measuring script reproduces every number the application's texts quote.

9.6.1 The Profile, the Staircase and the Trade

Case 1 opens at the design point with the metals parting at stage 12.60, the organic loaded to 32.2 % of its capacity at stage 40, and the measured pH at 1.26 in stage 1, about 0.7 through the extraction section, 0.22 at the top of the scrub section and below zero in the strip. Raising the scrub ratio to 1.200 moves the parting toward the raffinate end, to 11.95 a day later and 11.24 at hour 49, while the pH profile hardly changes, because the separation is an exchange of one metal for another at a nearly constant loading. Case 2 reads the staircase of the previous section on the Map and then lowers the separation factor to 1.30, where the shipped cascade lines out at 95.98 %. Case 3 raises the scrub ratio by one percent, to 1.185. The purity rises by four hundredths of a point, to 99.710 % at hour

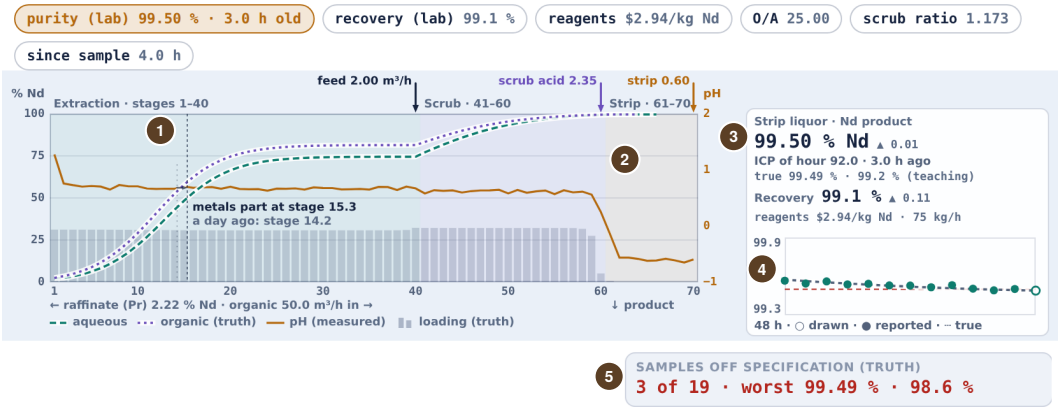


Figure 9.4 Case 4 at its hour-96 pause on the Cascade view, in Manual with every flow at the design ratios (O/A 25.00, scrub ratio 1.173), light theme, 1366 × 768, cropped to the plant view and the samples tile of the Learn panel. The feed fell from 75 to 67.5 % neodymium at hour 4. The profile shows the metals parting at stage 15.3 against 14.2 a day earlier (1), the organic loading near a third of its capacity and the measured pH, which falls below zero in the strip section (2). The products panel gives the ICP purity of 99.50 %, up 0.01 on the previous report, from the sample of hour 92 reported 3.0 h ago, the true 99.49 % and 99.2 %, and the recovery of 99.1 %, up 0.11 (3), with the samples of the last two days sliding toward the 99.5 % line (4). The badges give $2.94 of reagents per kilogram, and the scorecard three of 19 samples off specification with the worst at 99.49 % (5).

49, while the recovery falls to 98.43 % and is still falling at hour 121 (97.98 %), because one percent more reflux cuts a draw that is a fifth of the extraction amount by several percent.

9.6.2 A Leaner Feed by Hand

Case 4 cuts the feed to 67.5 % neodymium at hour 4 and leaves every flow alone. The first report after the change, the sample of hour 8 reported at hour 9, reads a recovery of 98.63 % against 98.78 % before, the wrong way, because the feed carries less neodymium before the raffinate has changed. The purity reads 99.666 % at hour 25, the true purity crosses 99.5 % at hour 85.7, and the first report below it is the sample of hour 88 (99.492 %). At the pause at hour 96 the true purity is 99.490 % with the recovery at 99.16 %, the parting has walked from 14.25 at hour 72 to 15.35, and three of the 19 samples drawn after the first day have fallen short of the specification, as shown in Figure 9.4. The unchanged draw takes more neodymium than arrives, the cascade gives up its inventory, and praseodymium creeps up the scrub section into the product.

🕮 **Pause.** Try Exercise 9.1 before reading on.

9.6.3 Loops on a Four-Hour Laboratory

Case 5 runs the purity loop at 1.5 per point with a 4 h integral time and raises the feed grade to 82.5 % at hour 6. Each report sees the result of moves made a day before and responds with another move. The moves do not grow from report to report, but the scrub ratio keeps swinging. From hour 6 to 96 the scrub ratio swings between 1.015 and 1.201, about 17 % of its middle, the pH of stage 1 between 1.18 and 2.08, and the true purity between 99.52 and 99.67 %. At 0.05 per point and 24 h the ratio moves smoothly from 1.1665 at hour 24 to 1.1449 at hour 96 and the purity settles near 99.74 %. The swinging loop stays inside the specification only because the plant is slow. On the online analyzer the same gain acts eight times as often and the scrub ratio swings between 0.899 and 1.304.

Case 6 repeats case 4's leaner feed at hour 6 with the detuned purity loop in Auto. Without the feedforward the loop sees nothing for a day and then chases the drift, the purity falling from 99.668 to

99.510 % by hour 96 (worst sample 99.509 %) while the scrub ratio creeps from 1.169 to 1.180. With it the scrub ratio steps to 1.197 when the leaner feed is reported at hour 9, the purity rises to 99.693 % at hour 24 and eases to 99.591 % at hour 96. The reagents over the 96 hours rise from $2.94 to $3.08 per kg, because without the feedforward the cascade was selling neodymium out of its own inventory.

> **TRY IT IN THE APP** **6 · A Feed Change With the Feedforward**
>
> Run to the pause at hour 96 with Feed feedforward Off, then switch it On, press Reset and run again. Off, the purity slides to 99.510 % by hour 96. On, the scrub ratio steps to 1.197 when the leaner feed is reported at hour 9, the purity ends at 99.591 %, and the reagents rise from $2.94 to $3.08 per kilogram.

9.6.4 The Pairing

Case 7 closes both loops with purity on O/A and recovery on the scrub ratio. Each loop undoes the other's move and asks for more. By hour 96 the scrub ratio has climbed from 1.173 to 1.2013 and O/A from 25.00 to 25.472 (1.2164 and 25.769 by hour 168), the purity sits at 99.702 % against its 99.6 % setpoint and the reagents cost $3.00/kg. Paired as the relative gain says, with gains of 0.05 and 1, the scrub ratio eases to 1.155 and O/A to 24.646, the purity comes down to 99.661 % with the recovery at 98.53 %, and the reagents fall to $2.90/kg. Along the crossed run the Map's relative gain, recomputed at the plant's present ratios, stays between 3.90 and 4.11. The drift is slow because the recovery answers in weeks, and slow is not stable.

9.6.5 Capacity Loss and Crud

Case 8 removes a tenth of the extractant's sites at hour 6. The saponifier reads 18.4 % within minutes against its 16.7 % setpoint, because the same sodium is a larger share of fewer sites, and the loop cuts the NaOH from 8.36 to 7.51 kmol/h by hour 24, which also cuts the rare earth the organic can take up. The pH of stage 1 passes below 1.20 at hour 7.1 and reads 1.09 at hour 9, when the first ICP report with the loss gives a recovery of 98.20 %. The true recovery is 97.20, 94.79 and 87.17 % at hours 12, 24 and 48, while the purity rises to 99.83 %. The crud at hour 30 costs about five hundredths of a point of purity (99.624 against 99.670 % at hour 48 on case 1's plant). Its stage matters. A 70 % efficiency loss in any of stages 5 to 40 costs 0.0001 points of steady purity, because the feed stage pinch leaves the extraction section stages to spare, while in stages 45 to 55 it costs 0.047 points of purity and 0.045 of recovery. Stage 30 is no praseodymium-only zone (its aqueous is 80.9 % neodymium), and it lies in the feed-stage pinch, so the crud event sits at stage 55.

9.6.6 Five Days, Two Feed Changes and Crud

The challenge runs 120 h with the feed at 82.5 % neodymium from hour 12, at 67.5 % from hour 48 and crud in stage 55 at hour 72. It is met when every ICP sample drawn after hour 24 has a true purity of at least 99.5 % and a true recovery of at least 97 %, and the reagents over the run cost at most $2.95/kg of neodymium. Changing a model parameter or switching the loops to the online analyzer voids the attempt. The configurations are compared in Table 9.5. As shipped, the purity loop alone at 0.5 per point and 6 h puts two of 25 samples short (worst 99.47 % at hour 76) at $2.85/kg. The suggested settings, the feedforward with the purity loop at 0.05 and 24 h, keep every sample at or above 99.55 % at $2.93/kg. The feedforward alone holds the purity but fails the cost at $3.01/kg, because nothing brings the scrub back down once the richer feed makes purity cheap, and the detuned loop alone is too slow for the leaner feed, with seven samples short. Adding the recovery loop to the suggested settings puts four samples short, because in five days it can only trade purity for NaOH. With the feedforward on, every purity gain from 0.05 to 0.3 passed at integral times of 12 and 24 h and every gain from 0.05 up failed at 6 h, so the integral time is what has to match the plant's day. At seeds 7 and 11 and at half the step the verdicts hold, the suggested settings passing at 99.547 to 99.557 % and $2.925 to $2.935 per kg, a margin of $0.015 to $0.025 per kg under the cap. On the

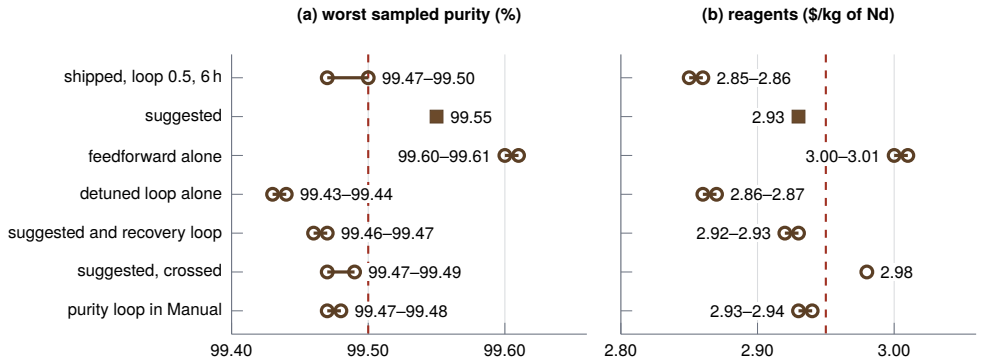


Figure 9.5 The challenge of Table 9.5 as a chart, the range over five seeds drawn from its low to its high value. Only the suggested settings (filled square) sit right of the purity floor and left of the cost cap (dashed lines). The feedforward alone holds the purity with the widest margin of any policy and fails on cost, while the shipped loop and the detuned loop alone, both without the feedforward, are the cheapest rows and fall short of the purity floor.

Table 9.5 The Challenge Over 120 h, Five Seeds Each

Policy (ICP only)	Pass	Short	Worst x (%)	$/kg
shipped, purity loop 0.5, 6 h	0/5	1–2	99.47–99.50	2.85–2.86
suggested, feedforward and 0.05, 24 h	5/5	0	99.55	2.93
feedforward alone (gain 0)	0/5	0	99.60–99.61	3.00–3.01
detuned loop alone	0/5	7	99.43–99.44	2.86–2.87
suggested and recovery loop (1, 24 h)	0/5	4	99.46–99.47	2.92–2.93
suggested, crossed pairing	0/5	2–3	99.47–99.49	2.98
purity loop in Manual	0/5	3	99.47–99.48	2.93–2.94

online analyzer even the shipped loop passes (worst 99.53 %, $2.88/kg), which is why the attempt is graded on the ICP. Figure 9.5 draws the table as a chart.

> **DID YOU KNOW? On Specification and Still Failing**
>
> From the simulator: in Table 9.5 the feedforward alone keeps the purity of the challenge highest, its worst sample at 99.60 to 99.61 %, yet fails all five seeds at $3.00 to $3.01 per kilogram. The suggested settings pass every seed with a worst sample of 99.55 % at $2.93.

9.6.7 Coordinated Ratios Against Sampled PI

Case 9 raises requested purity from 99.6 to 99.8 % and lowers requested recovery from 98.5 to 97.5 % at hour 8. Both loops are automatic, with identical seed, instruments and initial conditions for the two methods. The 120 h comparisons are in Table 9.6, whose rows run cases 9, 5, 6, 7 and 8 as loaded and the challenge with its suggested settings, with only the supervision changed. IAE integrates each reported product through the same metric used by the scorecard. Off-specification hours use the true purity, and cost is cumulative reagent expense per kilogram of Nd produced.

On the new requests, purity IAE falls from 13.11 to 9.24 point h and recovery IAE from 124.31 to 80.93 point h, but reagent cost rises from $2.977 to $3.097/kg. Neither run crosses the purity specification. The different response rates appear in Figure 9.6. Halving the integration step retains the ranking, with purity IAE 13.13 versus 9.26 point h. A six-setting hindsight search of the classical gains is retained in the evidence without claiming a globally optimal PI comparator.

Table 9.6 Matched 120-Hour Supervision Comparisons

Case, law	Purity IAE (point h)	Recovery IAE (point h)	Off (h)	$/kg
Setpoints, PI	13.11	124.31	0.00	2.977
Setpoints, MPC	9.24	80.93	0.00	3.097
Fast PI case, PI	3.16	120.42	0.00	2.520
Fast PI case, MPC	18.11	34.03	0.00	2.910
Feed change, PI	7.63	47.08	0.00	2.954
Feed change, MPC	6.47	33.25	0.00	2.991
Crossed PI, PI	11.92	15.20	0.00	3.012
Crossed PI, MPC	6.33	22.64	0.00	2.865
Capacity/crud, PI	24.74	1659.68	0.00	3.711
Capacity/crud, MPC	26.34	1899.15	0.00	3.864
Shift, suggested, PI	6.57	42.35	0.00	2.927
Shift, suggested, MPC	11.86	27.83	17.78	2.925

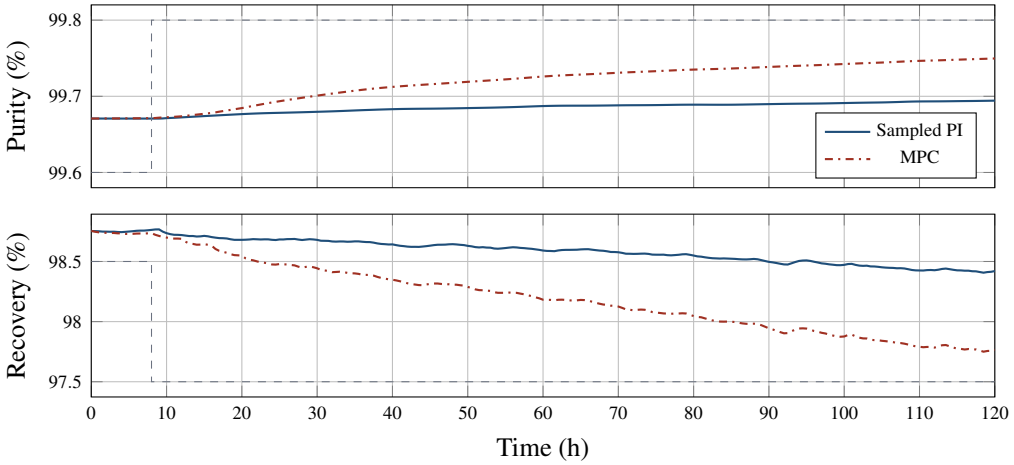


Figure 9.6 Purity and recovery following the paired setpoint changes of case 9, sampled PI solid and MPC dash-dotted. Thin dashed lines are the requests.

The benefit does not carry to the suggested shift. Compensated PI passes with no purity-specification crossing, while MPC spends 17.78 h below it and has worst sampled purity 99.443 %. The model has no feed-composition input. On the single feed change, adding the classical feedforward gives purity IAE 4.82 rather than MPC's 6.47 point h. With capacity loss and crud, worst sampled recovery is 69.93 versus 73.48 %. Bias correction does not repair the wrong dynamic gain. The soft product floors therefore cannot be interpreted as actual specification guarantees.

All ordinary comparison solves are accepted, 30 per 120 h active run. Chrome averages 3.98–4.67 ms per solve in the six layout captures, at most 0.117 % of one core at unit playback and 0.350 % at the top 3 h/s rung. Independent tests reject ten injected malformed or invalid solver results and reproduce two measurement rows exactly.

9.7 Technology in Practice

Figure 9.7 draws the loops of Section 9.5 with their hardware.

What Measures It

The cascade's composition is known only from analyses of samples. The model reports the purity, the recovery and the feed's neodymium from samples drawn every 4 h and reported 1 h later, reads

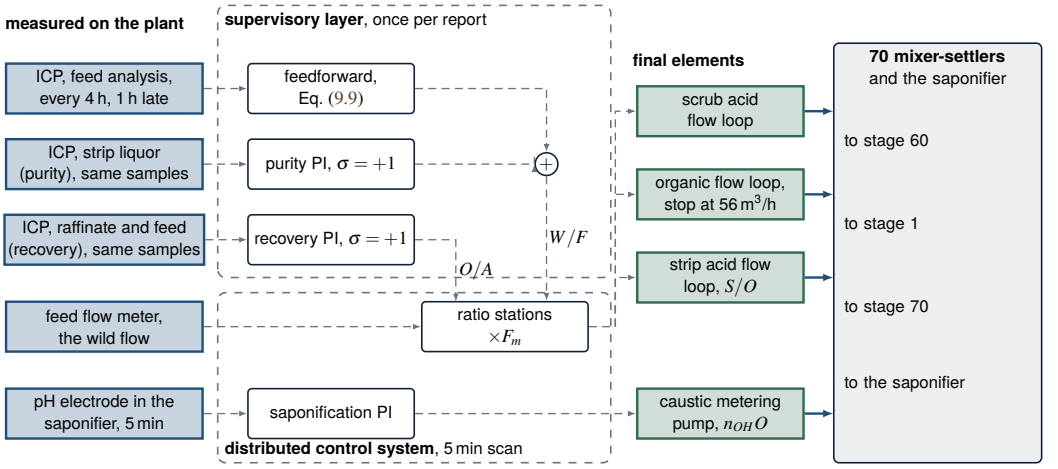


Figure 9.7 The loops of the simulated cascade drawn with their hardware. Three laboratory analyses of the same samples feed the supervisory layer, which acts once per report: the feedforward and the purity PI add up to the scrub ratio, and the recovery PI sets the organic ratio. The ratio stations of the regulatory layer multiply both by the metered feed, the wild flow, every 5 min, and the saponification PI on the saponifier's pH trims the caustic dose.

the pH of every stage every 5 min, and offers an online X-ray fluorescence option that samples every 30 min with five times the purity noise.

> **TECHNOLOGY Sensor: Plasma Emission Analysis in the Laboratory**
>
> The purity and recovery of this chapter come from inductively coupled plasma optical emission spectroscopy. A sample is sprayed into a plasma sustained by a radio-frequency coil, the excited atoms and ions emit light at discrete wavelengths, and optics separate the lines for a detector. Neodymium can be read at 430.358 nm and praseodymium at 417.939 nm. An organic sample cannot be analyzed directly, so its metal is first stripped into acid. In one laboratory study the samples were diluted 2 to 500 times into a calibration of 10 to 100 000 µg/L, a control standard every ten samples guarded against drift, spiked samples checked for interference, and mass balances on the stripped organic closed within 5 % [10].

The one-hour delay of Equation (9.5) stands for all of that work, the dilution, the stripping of an organic sample and the queue at the instrument, and the 0.01 points of purity noise stand for its error.

What Moves It

The model sets the organic, scrub, strip and caustic flows from the metered feed every 5 min, meters them exactly and stops the organic at the settlers' 56 m^3/h.

> **TECHNOLOGY Final Element: Ratio Stations and Metering Pumps**
>
> A ratio station takes the signal of the uncontrolled, or wild, flow, multiplies it by an adjustable constant and sends the product as the setpoint of the controller on the captive flow, so the proportion holds whatever the total [19]. In this cascade the feed is the wild flow, and the organic, the scrub acid and the strip acid are captive. Small flows of acid or caustic are commonly delivered by metering pumps, a pump, motor and drive in one package whose output runs from no flow at 4 mA to full flow at 20 mA. Many use a plunger driven by a solenoid and deliver discrete pulses rather than a continuous flow [19].

In the model each captive flow follows its setpoint at once, so the lag of a real flow loop and the pulses of a plunger pump are left out.

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 app the ratio stations and the saponification PI act every 5 min, while the two product loops act once per laboratory report, every 4 h, on a plant whose purity needs about a day to answer the scrub. On such a platform the ratio stations and the saponification loop are function blocks of the regulatory layer, and the product loops with the feedforward form a supervisory application that writes the two ratios.

> **IN PRACTICE** **Monitoring Stages and Experience**
>
> In industrial rare earth extraction, monitoring stages are set in each extraction and scrub section, and the flows of extractant and scrub liquor are adjusted from the deviation between the measured and the target content there. In one plant that content came from a daily laboratory analysis of field samples [21]. Dosing has been described as offline analysis, manual adjustment and experience control: the operator sets the flow rates from experience and inner PID loops hold them [22]. Saponification is among the largest reagent costs of a separation plant, and extractant makers recommend loading the organic to no more than 30 % of its stoichiometric capacity [10].

What Is Modernizing It

The first development is the soft sensor. Data-driven models estimate a hard-to-measure quality from variables that are easy to measure, and a recent one estimates the content at a monitoring stage from the extractant, scrub and feed flows and the feed composition, recorded every 30 min, against a laboratory that reports daily [21]. The second is machine vision: because the solution changes color with its praseodymium and neodymium content, color features of its image have been used to model the neodymium content of a production cascade [5]. The third is predictive control. An outer model predictive controller that sets the extractant and scrub flows over inner PID loops has been proposed in place of manual dosing, with an event trigger that updates the controller only when the process calls for it [22]. The predictive supervisor of Section 9.5.4 is a version of that idea, and its failure on the feed shift shows why such a controller still needs the feed analysis.

9.8 Limitations

Every flow, price and limit is a teaching value, and only the chemistry constants, the residence times, the saponification and the feed come from a source, a laboratory thesis rather than a plant. The constant K_{Nd} is held at one acidity while the measured value spans 0.019 to 0.044 over the process range. At the shipped ratios half of it gives 99.72 and 96.1 % and twice it 97.9 and 100 %, so the design ratios belong to the constant, and a re-tuned cascade would teach the same lessons at other ratios. The separation factor is the same at every stage and loading, while the sources measure 1.29 to 2.05 as those change [10, 15]. Two metals, equilibrium stages, settlers as four tanks and a hydrochloric acid scrub stand for a plant of many pairs, finite kinetics, real settler hydraulics and product reflux. Real feed tanks blend over hours, so a grade that steps at once is unrealistic, but the step isolates the inventory's response from everything else.

The organic loading has no consequence of its own. The 30 % guideline is met only approximately at design (32.2 %), and with the NaOH station at its cap on an extractant that has lost 30 % of its sites the loading reaches 50 % with no third phase or solubility penalty. The infeasibility of 99.9 % at 97 % holds only within the declared limits of $O/A \leq 28$ and 18 % saponification. The recovery loop barely acts within a case because the recovery responds in weeks, so the challenge is won without it. The relative gain is the model's at one feed and moves from 3.60 to 3.92 with the step direction. The laboratory has no sample-line lag and no outliers except a scripted one, and the classical controllers retain only their material balance. The added predictive supervisor uses the clean local model of Equation (9.10), which cannot infer changing gains or the current extraction inventory.

OPEN CHALLENGE **Measuring Composition Along the Cascade**

Rare earth component content is hard to measure quickly in a plant, where slow dynamics and a harsh environment stand in the way, and the offline analyses in wide use delay control and cost product consistency [21]. The solution changes color as its praseodymium and neodymium content changes, and color features of its image have served as inputs to a neural network model of neodymium content in a 30-stage plant, where each group of tanks adds a delay that most models ignore or hold constant [5]. Progress would be an online estimate of the stage profile, checked against the laboratory, that a controller could act on within hours rather than a day.

9.9 Conclusion

A praseodymium and neodymium cascade of 40 extraction, 20 scrub and 10 strip mixer-settlers is modeled here as a two-phase stage stack with cation exchange on a dimeric extractant, a charge-balance pH and a saponifier, whose metal balances close to 10^{-11} and whose steady-state solver the dynamic stack reproduces to five decimals. A separation factor of 1.5 adds 0.167 to the purity's log-odds per scrub stage at the design ratios, against 0.405 at total reflux. The scrub ratio moves the purity 63 % in about a day, while a feed-grade step drains an inventory of 108.6 hours of feed whose ratio to the new imbalance is 718 h, and the purity of case 4 crosses 99.5 % only at hour 86. Keyed to the feed the scrub gives a relative gain of 3.92 that rejects the crossed pairing, which in case 7 drives both ratios up for days, and keyed to the organic the same plant gives 0.60. The 99.9 % and 97 % pair is a capacity limit, reached at O/A 33 or at 21 % saponification. A feedforward from the feed analysis with a purity loop detuned to the plant's day passes five days at \$2.93/kg, where either alone fails. The chemistry constants come from laboratory data and the rest are teaching values. The predictive option reduces paired-setpoint IAE at greater reagent cost, but fails a shift the compensated PI passes. A fuller model would add an acidity-dependent extraction constant, product reflux in the scrub, a third rare earth and a soft sensor on the stage pH. The simulator shows that a separation with a fast and a slow end has to be controlled at the speed of each end, on what is measured, when it is measured.

KEY POINTS

- The purity comes from the stage count, and the stage count makes the plant slow.
- The scrub ratio moves the purity 63 % in about a day, while a feed-grade step drains an inventory whose ratio to the new imbalance is 718 h.
- Keyed to the feed, the scrub gives a relative gain of 3.92 that rejects the crossed pairing, and keyed to the organic the same plant gives 0.60.
- The 99.9 % and 97 % pair is a capacity limit, not a stage limit.
- A feedforward from the feed analysis with a purity loop detuned to the plant's day passes five days at \$2.93/kg, where either alone fails.

Where else this applies. Volume 5, Chapter 3 closes a composition loop on a slow analyzer, and Volume 3, Chapter 6 doses from the raw feed with the late product measurement as a slow trim.

References

[1] O. Gutfleisch, M. A. Willard, E. Brück, C. H. Chen, S. G. Sankar, and J. P. Liu, "Magnetic materials and devices for the 21st century: Stronger, lighter, and more energy efficient," *Advanced Materials*, vol. 23, no. 7, pp. 821–842, 2011.

[2] F. Xie, T. A. Zhang, D. Dreisinger, and F. Doyle, "A critical review on solvent extraction of rare earths from aqueous solutions," *Minerals Engineering*, vol. 56, pp. 10–28, 2014.

[3] K. L. Lyon, V. P. Utgikar, and M. R. Greenhalgh, "Dynamic modeling for the separation of rare earth elements using solvent extraction: Predicting separation performance using laboratory equilibrium data," *Industrial & Engineering Chemistry Research*, vol. 56, no. 4, pp. 1048–1056, 2017.

[4] H. Yang, F. Xu, R. Lu, and Y. Ding, "Component content distribution profile control in rare earth countercurrent extraction process," *Chinese Journal of Chemical Engineering*, vol. 23, no. 1, pp. 192–198, 2015.

[5] R. Lu, H. Liu, H. Yang, J. Zhu, and W. Dai, "Multi-delay identification of rare earth extraction process based on improved time-correlation analysis," *Sensors*, vol. 23, no. 3, p. 1102, 2023.

[6] C. M. Moreno, J. R. Pérez-Correa, and A. Otero, "Dynamic modelling of copper solvent extraction mixer–settler units," *Minerals Engineering*, vol. 22, no. 15, pp. 1350–1358, 2009.

[7] S. Skogestad, P. Lundström, and E. W. Jacobsen, "Selecting the best distillation control configuration," *AIChE Journal*, vol. 36, no. 5, pp. 753–764, 1990.

[8] 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.

[9] U.S. Geological Survey, "Mineral commodity summaries 2025: Rare earths," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[10] K. L. Lyon, "Separation of adjacent rare earth elements using solvent extraction," M.S. thesis, University of Idaho, Moscow, ID, 2016. [Online]. Available: https://www.lib.uidaho.edu/digital/etd/items/lyon_idaho_0089n_10815.html

[11] F. H. Spedding, A. F. Voigt, E. M. Gladrow, and N. R. Sleight, "The separation of rare earths by ion exchange. I. Cerium and yttrium," *Journal of the American Chemical Society*, vol. 69, no. 11, pp. 2777–2781, 1947.

[12] D. F. Peppard, J. P. Faris, P. R. Gray, and G. W. Mason, "Studies of the solvent extraction behavior of transition elements. I. Order and degree of fractionation of the trivalent rare earths," *The Journal of Physical Chemistry*, vol. 57, no. 3, pp. 294–301, 1953.

[13] D. F. Peppard, G. W. Mason, J. L. Maier, and W. J. Driscoll, "Fractional extraction of the lanthanides as their di-alkyl orthophosphates," *Journal of Inorganic and Nuclear Chemistry*, vol. 4, no. 5–6, pp. 334–343, 1957.

[14] Xu Guangxian, Li Biaoguo, and Yian Chunhua, "Theory of countercurrent extraction and its applications in rare earth extraction industry," in *New Frontiers in Rare Earth Science and Applications*. Elsevier, 1985, pp. 429–437.

[15] M. S. Safarzadeh, V. Agarwal, and L. Hayes, "New insights into the separation of Nd from Pr in hydrochloric and sulfuric acid solutions," *Polyhedron*, vol. 153, pp. 82–87, 2018.

[16] M. R. Fenske, "Fractionation of straight-run Pennsylvania gasoline," *Industrial & Engineering Chemistry*, vol. 24, no. 5, pp. 482–485, 1932.

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

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

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

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

[21] W. Dai, R. Lu, P. Chen, and H. Yang, "Novel dual residual-enhanced deep bidirectional LSTM network for soft sensing of rare earth component content," *Sensors*, vol. 26, no. 10, p. 3152, 2026.

[22] Z. Zhang, J. Zhu, C. Xiong, H. Yang, and F. Nie, "Constrained model predictive control based on event triggering for the rare earth extraction process," *ACS Omega*, vol. 8, no. 44, pp. 41 943–41 952, 2023.

9.10 Exercises

9.1 ○ Concept check · 2 min

In case 4 the feed grade falls from 75 to 67.5 % neodymium at hour 4 and no flow changes. The first laboratory report after the change, the sample of hour 8 reported at hour 9, reads a recovery of 98.63 % against 98.78 % before, although the recovery rises over the following days. Why does the first report move the wrong way?

(a) The leaner feed loads less neodymium on the organic, so more neodymium leaves in the raffinate at once

(b) The saponification loop cuts the NaOH when the feed changes, so the organic takes up less rare earth

(c) The recovery divides the raffinate's neodymium by the feed's, and the feed changes at once while the raffinate has not

(d) The recovery analysis carries noise, and the true recovery rises from the first hour while the report lags behind it

9.2 Concept check · 1 min

Why does the purity answer a step of the scrub ratio within about a day, but a step of the feed grade only over weeks?

(a) The scrub acid reaches the strip end at once, while the feed must travel through forty extraction stages

(b) The laboratory reports the scrub change at once, while the feed analysis is held for four hours

(c) The scrub step changes the separation factor, while the feed step leaves the separation factor unchanged

(d) The scrub step leaves the neodymium balance unchanged, while the feed step must drain tens of kilomoles at a small imbalance

9.3 Concept check · 2 min

In the crossed pairing of case 7 the purity analysis drives the O/A ratio. With the other loop open, the steady-state gain of purity on O/A is -0.00173 per unit ratio. In the convention of Section 1.6.2, what is this loop's action?

(a) Direct acting, $\sigma = -1$, because the purity falls as the organic ratio rises

(b) Reverse acting, $\sigma = +1$, because the purity loop on the scrub ratio is reverse acting

(c) Reverse acting, $\sigma = -1$, because the app names this pairing reversed

(d) Direct acting, $\sigma = +1$, because a negative relative gain flips the sign of the loop

9.4 Theory and calculation · 10 min

At the design point ($W/F = 1.173$, $O/A = 25$) the steady-state gains in fractions per unit ratio are $\partial x/\partial(W/F) = 0.0416$, $\partial x/\partial(O/A) = -0.00173$, $\partial R/\partial(W/F) = -2.55$ and $\partial R/\partial(O/A) = 0.142$. Give each relative gain to within 0.05.

(a) Compute λ_{11} of purity on the scrub ratio from Equation (9.6), and the relative gain of the crossed pairing. *Checkpoint:* $\lambda_{11} \approx 3.95$, crossed -2.95

(b) Key the scrub to the organic instead: $v = W/O$, so $W/F = v\,(O/A)$. Use the chain rule to find the gains of x and R on v at fixed O/A and on O/A at fixed v at the design point, and compute λ_{11} of purity on v. *Checkpoint:* gains 1.04, 0.000222, -63.75, 0.0224; $\lambda_{11} \approx 0.62$

(c) Which pairing does each keying reject, and why do the two numbers differ for one plant? *Checkpoint:* feed keyed rejects the crossed pairing, organic keyed rejects neither

9.5 Theory and calculation · 6 min

The design feed is 2 m^3/h at 0.35 M rare earth, 75 % of it neodymium, and the cascade holds 76.0 kmol of rare earth. Give each answer to within 1 %.

(a) How many hours of feed does the cascade hold? *Checkpoint:* 108.6 h

(b) The feed grade falls to 67.5 % with every flow unchanged, so the draw still takes 0.5250 kmol/h of neodymium. What is the imbalance, and what time does inventory over imbalance give for the 37.7 kmol the new steady state gives up? *Checkpoint:* 0.0525 kmol/h, 718 h

(c) A scrub step from 1.173 to 1.178 needs only 2.57 kmol more neodymium. Why does the same estimate not apply to it, and what does the chapter measure instead? *Checkpoint:* no imbalance at the step; 63 % in 25.3 h

9.6 Theory and calculation · 6 min

The feedforward of Equation (9.9) uses scrub acid at $c_W = 3\,\mathrm{M}$ and a recovery setpoint $R_{SP} = 98.5\,\%$. The feed is 0.35 M rare earth, and its grade falls from 75 to 67.5 % neodymium.

(a) Evaluate u_{ff} to four decimals. *Checkpoint:* $+0.0259$

(b) Show where the factor $3/c_W$ comes from, starting from the cut in the neodymium draw that keeps it equal to R_{SP} times the neodymium fed. *Checkpoint:* $\Delta W = 3R_{SP}F\,|\Delta c|/c_W$

(c) A run measures a feedforward part of 0.0246. The feed analysis reports its neodymium concentration with 0.5 % noise. How large is the noise of one reading of the design feed, 0.2625 M of neodymium, in M, and in the scrub ratio? Give two significant figures. *Checkpoint:* 0.0013 M, 0.0013 in the ratio

9.7 App challenge · 8 min

Open scenario 5, “The laboratory is four hours late: tuning a loop on stale reports”, at the default seed. It runs in Auto with Purity gain 1.5 per point and Purity integral time 4 h, and the feed grade rises to 82.5 % neodymium at hour 6. Run at 3 h per second to the pause at hour 96 (about 32 s) and export the CSV file. Read ratios within 0.002 and the purity within 0.01 points.

(a) Read the lowest and the highest `WF_cmd` from Time (h) 6 to 96, and give the swing as a percentage of their mean, to 1. *Checkpoint:* 1.015 and 1.201, 17 %

(b) Set Purity gain to 0.05 and Purity integral time to 24 h, press Reset, run to the pause at hour 96 and export again. Read `WF_cmd` at Time (h) 24 and 96, and the true purity at hour 96 on the tile “Nd product: purity · recovery · reagents”. *Checkpoint:* 1.1665, 1.1449, 99.74 %

(c) Both loops act once per report. Why does the second settle? *Checkpoint:* its integral time matches the plant’s day, not the report period

9.8 App challenge · 8 min

Open scenario 7, “The pairing: why the relative gain says purity goes on the scrub”, at the default seed. Purity drives O/A at gain 1 per point and recovery drives the scrub ratio at 0.05, both with 24 h integral times and both in Auto. Run at 3 h per second to the pause at hour 96 and export the CSV file. Read the scrub ratio within 0.0005, O/A within 0.01, and the true purity and recovery on the tile “Nd product: purity · recovery · reagents”.

(a) Read `WF_cmd` and `OA_cmd` at Time (h) 96 and the true purity. Which way did each ratio move from 1.173 and 25.00? *Checkpoint:* 1.2013 and 25.47, both up; 99.70 %

(b) Set Pairing to Purity→scrub, Purity gain to 0.05 and Recovery gain to 1, press Reset and run to the pause at hour 96. Read the same values and the true recovery. *Checkpoint:* 1.1550, 24.65, 99.66 %, 98.5 %

(c) Use the relative gain to explain why the crossed loops walked both ratios up while the purity moved away from its 99.6 % setpoint. *Checkpoint:* purity’s gain on O/A changes sign when the recovery loop closes

9.9 App challenge · 10 min

Open scenario 10, “Challenge · five days, two feed changes and crud”, at the default seed. It is met when every ICP sample drawn after hour 24 has a true purity of at least 99.5 % and a true recovery of at least 97 %, with the reagents at most $2.95 per kg of neodymium. Each run takes 40 s at 3 h per

second and pauses at hour 120 with its verdict on the Challenge target tile. Read the worst sample within 0.01 points and the cost within $0.01 per kg.

(a) Run it as loaded (purity loop at 0.5 per point and 6 h, feedforward off). Record the verdict, the samples short, the worst sample purity and the cost. *Checkpoint:* not met, 2 of 25, 99.47 %, $2.85/kg

(b) Set Feed feedforward to On, Purity gain to 0.05 and Purity integral time to 24 h, press Reset and run. *Checkpoint:* met, 0 of 25, 99.55 %, $2.93/kg

(c) Keep the feedforward and the gain of 0.05. Run once with the integral time at 6 h and once at 12 h, pressing Reset before each. Which passes, and which setting decides it? *Checkpoint:* 6 h: not met, 5 of 25, 99.49 %; 12 h: met, 99.52 %, $2.89/kg

9.10 Application area · 3 min

In case 8 a tenth of the extractant's sites are lost at hour 6. Put in order the three signals that show it, with the hour each first does: the saponified share read from the saponifier's pH, the pH of stage 1, and the laboratory's recovery. Then explain why the saponification loop, acting correctly on its own measurement, makes the recovery worse, and name one corrective move an engineer could make before the laboratory confirms the loss.

9.11 Application area · 3 min

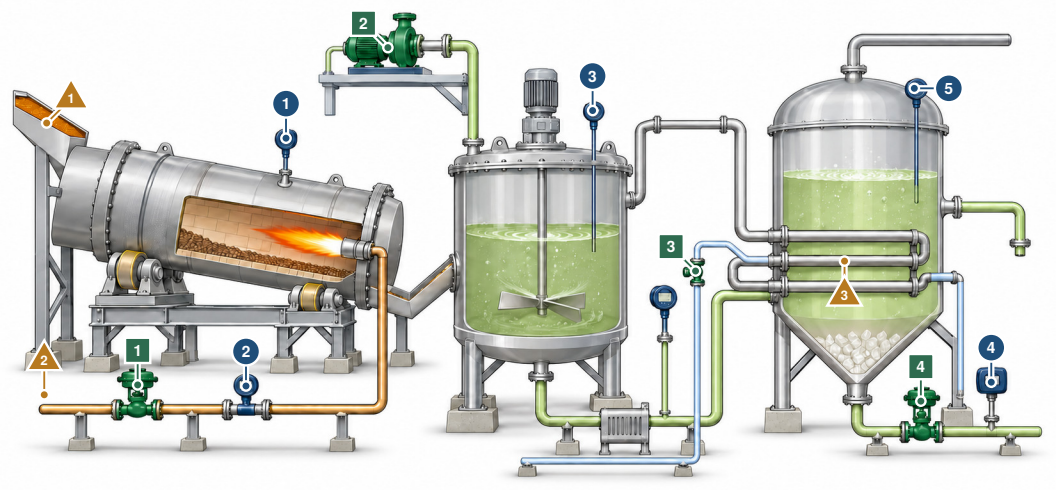
A plant manager asks for 99.9 % neodymium at 97 % recovery from the shipped cascade. Using Figure 9.3 and the sweep of Section 9.4, say why no setting of the ratios meets it within the equipment limits, name the two ways past those limits, and say what the model leaves out that an engineer must weigh before either.

9.12 Review the agent · 5 min

An AI assistant reviewing the challenge writes: "The suggested settings leave the recovery loop in Man. Close it too, at Recovery gain 1 per point and a 25 h Recovery integral time, so that both products are under feedback. Two loops always hold the specification better than one, and the reagent cost will barely change." Find the flaw, and say how you would check it in the app.

CHAPTER 10

A Lithium Hydroxide Refinery Simulator for Interacting Loops on Three Units in Series



The controller moves 1 the fuel valve of the kiln burner 2 the lime slurry pump 3 the steam valve that heats the evaporator 4 the product draw, which holds the level

The controller reads 1 the kiln bed temperature 2 the fuel flow, for the cascade 3 the pH electrode in the leach 4 the product strength analyzer, four minutes late 5 the evaporator level

What disturbs it 1 the ore grade of the feed 2 the gas supply pressure and heating value 3 scale on the heating tubes

Notice that one change in ore grade at the kiln feed reaches the kiln, the leach and the product at three different times, while each loop holds its own reading. Illustration, not to scale.

IN PLAIN TERMS **Three Steps in a Row, One Ore**

The main idea. A refinery turns crushed lithium ore into a strong lithium solution in three steps. A turning oven heats the ore, a stirred tank dissolves the lithium with lime, and a boiler drives off water. Each step has its own controller.

The challenge. Nothing is stored between the steps, so a change at the front reaches each controller at a different time. When the ore gets poorer, every controller holds its reading, yet the plant makes 27 percent less lithium. A probe that misreads the tank can even make the plant look better while it wastes lime.

Why it matters. A plant that watches only its controllers can lose output or money with no alarm.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the fact box of the Introduction, the history box of Section 10.2, the try-it boxes and the fact box of Section 10.6, 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 10.6, Section 10.7, Section 10.8 and the open challenge in it

The control question A calciner, a leach and an evaporative crystallizer run in series with no storage between them, so one disturbance reaches three loops at three different times. What does each loop see, what is a fuel flow cascade worth, and what does the plant lose while every loop holds its setpoint?

Learning objectives

- Identify each unit as a first-order plus dead-time plant from one manual step and turn the model into SIMC tunings.
- Predict which fuel-side disturbances a fuel flow cascade rejects, from what its inner loop measures.
- Trace an ore grade change through the three loops, with the sign and arrival time at each, and explain why the output falls with every loop on setpoint.
- Explain the interior maximum of the pH loop gain and where it makes a loop safe to tune.
- Detect a biased electrode or a scaling evaporator from evidence outside the loop that hides it.

Prerequisites Sections 1.6.2 (velocity form), 1.6.4 (scenarios and exports), 1.4.4 (leaching), 1.4.6 (cascade) and 1.4.7 (control layers of a circuit). Chapter 4 sorts a cascade's disturbances the same way.

The app https://apmonitor.com/apps/lithium-refinery/

See also Volume 6, Chapter 7 and Volume 6, Chapter 8 for plantwide structure, Volume 6, Chapter 4 for an evaporator held through a late analyzer. Volume 3, Chapter 2 and Volume 3, Chapter 4 for cascade by disturbance entry and fuel control on a fired heater. Volume 5, Chapter 3 and Volume 5, Chapter 11 for a slow analyzer and an unmeasured kiln product.

10.1 Introduction

A hard-rock lithium refinery converts spodumene concentrate into a lithium hydroxide solution in three steps. A rotary calciner heats the concentrate above 1000 °C, where the mineral converts from its dense alpha form to a porous beta form that reagents can reach [1, 2]. An alkaline leach dissolves the lithium out of the converted solid. An evaporative crystallizer concentrates the resulting liquor to the strength the downstream crystallization needs [3].

> ✦ DID YOU KNOW? **Lithium Mined in 2024**
>
> In 2024 the world produced an estimated 240 000 tons of lithium, counted as lithium content, 18 % more than in 2023. Australia, where the lithium comes from mineral operations, produced 88 000 tons of it. Batteries took an estimated 87 % of the lithium used [4].

Each unit has one controlled variable and one actuator, and each of the three loops is straightforward in isolation. The three are in series with no intermediate storage, so a disturbance entering at the front reaches each loop in turn at a time set by the units upstream of it. One fuel valve step reaches the spodumene conversion in 14.0 min, the kiln temperature the valve is nominally controlling in 29.3 min, and the product concentration two units downstream in 112.8 min, a factor of 8.1 between the first and the last. A controller tuned against the kiln loop alone therefore addresses one of three response times.

The model rests on standard process fundamentals. Conservation of energy on the kiln charge and conservation of species in the leach and the evaporator give one state equation per inventory, and the number of inventories in series sets the order of the response and the phase lag a controller has to survive [5, 6]. Each rate is written against a driving force, so the conversion follows an Arrhenius rate in bed temperature, the leach rate is first order in both unleached lithium and free hydroxide, and the evaporation rate is proportional to duty, which makes every loop gain a function of the operating point rather than a constant. Time-scale separation orders the resulting lags, with a 40 min solids residence in the kiln and a 44 min liquor residence in the evaporator against a 0.3 min valve and a 0.25 min electrode, which is what permits each unit to be identified as a first-order plus dead-time model [7].

The model and its scenarios teach open-loop identification from a reaction curve, process order and dead time, cascade structure and the disturbance classes an inner loop does and does not reject, nonlinear loop gain and the case for detuning, sampled and delayed measurements, integrating level control, anti-windup and bumpless transfer, the effect of instrument bias on a closed loop, and the distinction between a loop settling time and a plant settling time [5, 7]. Industry runs fired equipment on the flow cascade studied here [8], operates hydrometallurgical circuits against feed grade as the dominant unmeasured disturbance [9, 10], and pays for reagent consumption that no controlled variable reports.

This chapter describes a browser-hosted lithium hydroxide refinery simulator in enough detail to be rebuilt, measures what a fuel flow cascade is worth against three different fuel-side disturbances on the same kiln, and follows a single upstream ore grade change through three loops that all return to setpoint. Three measured outcomes carry the argument. The cascade reduces the integral error from a 20 % header pressure loss by a factor of 52 to 93, depending on where in the controller scan the loss arrives, and the same error from a 7 % heating value loss not at all. A grade change from 5.5 to 4.0 % Li_2O moves no controlled variable by more than a sixth of a pH unit and removes 27 % of the plant output over six hours. A pH electrode reading 0.70 low raises the extraction, the yield and the production while raising the reagent consumption by 24 %. The unit models, the cascade structure and the tuning rules are standard. What the app adds is an eleven-state model of three units in series, the identified model of each unit, and eleven scenarios with a completion requirement each, the last a graded challenge whose clauses are separable.

Section 10.2 reviews lithium refining, kiln control and plantwide control. Section 10.3 gives the equations of the calciner, the leach and the evaporative crystallizer. Section 10.4 identifies each unit from a step test, measures how fast one fuel valve reaches three plant variables and how the pH loop gain moves with the operating point. Section 10.5 sets out the loops and their tunings. Section 10.6 reports the cascade against three fuel-side disturbances, the grade change walking through the train, a throughput increase, evaporator scaling, a biased pH electrode and the graded run. Section 10.7 names the real instruments, final elements and platform of the loops. Sections 10.8 and 10.9 state what the model leaves out and which results carry over to other plants, and the Exercises close the chapter.

10.2 Background

Spodumene, $LiAlSi_2O_6$, carries 8.03 % Li_2O when pure, and a commercial concentrate assaying 5.5 % Li_2O is therefore roughly two thirds spodumene and one third gangue. The alpha form is essentially inert to acid and alkaline attack, and the industrial route begins by decrepitating it to the beta form above about 1000 °C [2]. The transformation proceeds through an intermediate gamma phase, and the reported activation energies vary with particle size, heating rate and impurities [1]. Downstream of the calciner the sulfate route and the alkaline route compete, and the flowsheets, reagent choices and impurity removal steps are surveyed for the lithium market as a whole [3]. Alkaline leaching of spodumene with hydroxide has been characterized against thermodynamic and surface data [11], and first-order dissolution kinetics of the form used here are standard for hydrometallurgical circuits [12].

Kiln control has a long history in the process industries. The first industrial application of fuzzy logic control was a cement kiln, chosen because the plant was too nonlinear, too slow and too poorly instrumented for the classical methods available at the time [13]. Model predictive control later displaced rule-based schemes on the same equipment, handling the kiln with the calciner and the cooler in one calculation [14], and Chapter 11 treats a rotary cement kiln whose air follows its fuel. On the modeling side, one-dimensional axial models of the bed and the counter-current gas have been developed and validated against plant data [15] and extended to the integrated preheater, calciner, kiln and cooler so that the recycles between them can be represented [16]. Those models are the reference against which the reduced kiln used here should be read.

> ⧖ HISTORY **Heat First, Then Acid (1948 to 1950)**
>
> A patent filed in 1948 and granted in 1950 describes the route that hard-rock lithium still follows. Its inventors found no acid strength, time or temperature that could take lithium out of spodumene as mined. They first heated the crushed ore to between about 1000 and 1350 °C, in a muffle, rotary or direct-fired furnace, to turn alpha-spodumene into beta, and found 1050 to 1100 °C ample. Any furnace would serve so long as uniform and adequate temperature control was possible. The beta form was then roasted with strong sulfuric acid at about 250 °C, never much above 300 °C, and the lithium sulfate leached out with water [17]. In 2019 this acid route was still the only one practiced on an industrial scale, and alkaline leaching was one of the alternatives proposed [18].

Plantwide control is the branch of the field that treats the interaction between units rather than the tuning of any one loop. The control structure design problem asks which variables to control, which to manipulate and how to pair them so that the plant, rather than the individual loop, meets its economic objective [5, 19]. Volume 6, Chapter 7 poses the same question for a reactor and a column joined by a recycle. Mineral processing plants make the question concrete, because feed grade and mineralogy vary on a time scale of hours, the measurements that would resolve them are sparse or offline, and reagent consumption is one of the two largest operating costs [9, 10]. The cascade structure is the standard local remedy for an actuator-side disturbance, and the fired-heater arrangement in which a temperature controller sets a fuel flow setpoint is the textbook case [5, 8].

However, the textbook justification for a cascade is usually presented as a single benefit, without separating the disturbances that enter inside the inner loop from those that do not. A fuel flow cascade rejects a header pressure change almost completely, because the inner loop measures the flow that the pressure changes. It rejects nothing at all of a heating value change, because the mass flow of gas is unaltered and no instrument in either loop observes the energy content. Volume 3, Chapter 2 sorts disturbances the same way on a level loop, and Chapter 4 on the feed pressure of a cyclone cluster. The application is built so that both disturbances can be applied to the same kiln under the same tuning, and so that the resulting numbers can be quoted separately.

10.3 Process Model

Figure 10.1 shows the train and its four loops.

10.3.1 Design Point and Overall Balance

The clock is in minutes. The plant takes 500 kg/min of concentrate at 5.5 % Li_2O and 3 % moisture. At design the calciner runs at 1090 °C on 26.5 kg/min of fuel and converts 93.5 % of the spodumene, the leach runs at pH 12.500 on 279.1 L/min of lime slurry, and the crystallizer holds 12.00 kg Li/m^3 at 77.1 % of steam duty on a 3.33 m level.

The lithium balance closes at every stage. The concentrate carries 12.39 kg/min of lithium, the calciner converts 11.59 kg/min of it into leachable form, the leach dissolves 11.04 kg/min, and the filter recovers 98.5 % of that as 10.87 kg/min of lithium in product liquor, an overall yield of 87.7 %. The application reports a leach extraction of 95.7 % on a concentration basis, which is 95.2 % on a mass basis because the reagent flow dilutes the slurry between the inlet and the outlet of the leach vessel.

10.3.2 Calciner

The kiln is one well-mixed zone carrying an energy balance on the charge and a first-order conversion. The energy balance is Equation (10.1).

$$M_k c_{p,k} \frac{\mathrm{d}T_k}{\mathrm{d}t} = \eta_f \Delta H_f \dot{m}_f - F_s c_{p,s}(T_k - T_{\mathrm{in}}) - UA_k(T_k - T_{\mathrm{amb}}) - \Delta H_c F_s \phi X_k - F_s w h_{\mathrm{vap}} \mu(T_k) \quad (10.1)$$

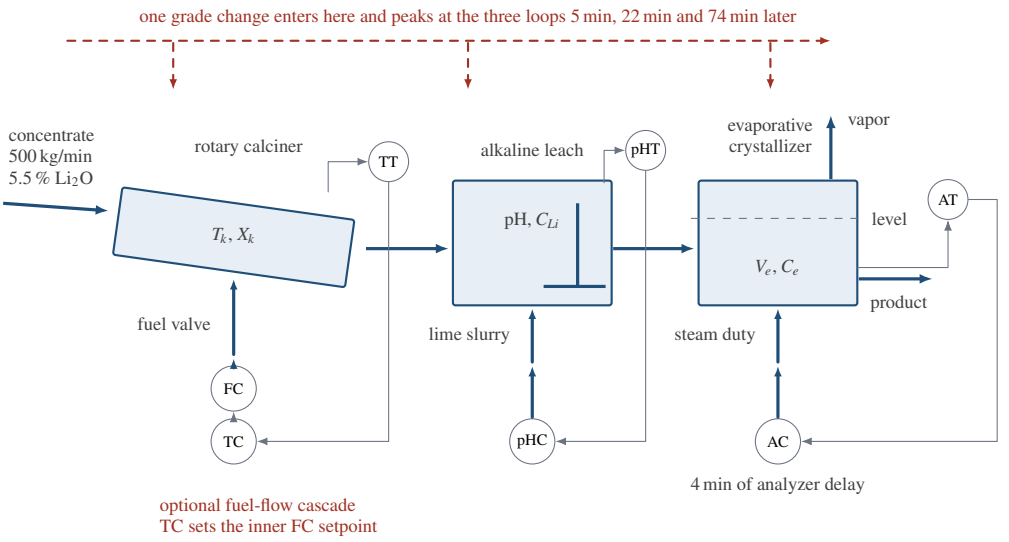


Figure 10.1 The simulated refinery. Three units in series, each with one loop. Kiln temperature goes to the fuel valve, optionally through a fuel flow controller in cascade. Leach pH goes to the lime slurry. Product concentration goes to the evaporator steam duty, through an analyzer with four minutes of transport delay. A fourth loop holds the crystallizer level. The dashed red path is the ore grade, which enters at the front and reaches the three controlled variables at three different times.

where $M_k c_{p,k} = 22\,000\,\mathrm{kJ/K}$ is the lumped thermal mass of the shell and charge, $\eta_f = 0.65$ is the burner efficiency to the bed, $\Delta H_f = 50\,000\,\mathrm{kJ/kg}$ is the fuel heating value, $\dot{m}_f$ is the fuel mass flow, F_s is the concentrate rate, $c_{p,s} = 1.05\,\mathrm{kJ/(kg\,K)}$, $UA_k = 220\,\mathrm{kJ/(min\,K)}$ is the shell loss coefficient, $\Delta H_c = 100\,\mathrm{kJ/kg}$ is the decrepitation enthalpy, $\phi = 0.685$ is the spodumene fraction of the concentrate at 5.5 % Li_2O, w is the moisture fraction, $h_{\mathrm{vap}} = 2450\,\mathrm{kJ/kg}$ and $\mu(T_k)$ is the fraction of the moisture driven off, equal to 1 at 100 °C and above and falling linearly to 0 at T_{in}, so that a kiln without fuel settles at the feed temperature. Every term is in kJ/min. At design the burner delivers 862 MJ/min against 559 of sensible heat, 234 of shell loss, 37 of moisture and 32 of reaction. The linearized thermal time constant is $M_k c_{p,k}/(F_s c_{p,s} + UA_k) = 29.5\,\mathrm{min}$.

Only the spodumene fraction of the concentrate absorbs the decrepitation enthalpy, so a lower ore grade is a smaller heat sink and drives the bed temperature up rather than down. That sign is the first of the three different disturbances one grade change delivers to this plant.

The conversion is a first-order reaction against a solids residence time, as shown in Equation (10.2).

$$\frac{\mathrm{d}X_k}{\mathrm{d}t} = k(T_k)\,(1 - X_k) - \frac{X_k}{\tau_s}, \qquad \tau_s = \frac{M_s}{F_s} \tag{10.2}$$

$$k(T_k) = k_{\mathrm{ref}}\,\mathrm{e}^{-\frac{E}{R}\left(\frac{1}{T_k} - \frac{1}{T_{\mathrm{ref}}}\right)} \tag{10.3}$$

where X_k is the fraction of the spodumene converted to the beta form, $M_s = 20\,000\,\mathrm{kg}$ is the kiln solids hold-up, $k_{\mathrm{ref}} = 0.475/\mathrm{min}$ at $T_{\mathrm{ref}} = 1100\,^{\circ}\mathrm{C}$, $E/R = 50\,500\,\mathrm{K}$ and both temperatures in (10.3) are absolute. The washout term X_k/τ_s is what makes the conversion less than unity, because material leaves the kiln after $\tau_s = 40\,\mathrm{min}$ whether or not it has reacted. At steady state $X_k = k\tau_s/(1 + k\tau_s)$, which gives 0.9355 at 1090 °C on $k = 0.363/\mathrm{min}$.

The conversion saturates against temperature over the operating range, at 0.780 at 1040 °C, 0.935 at 1090 °C and 0.982 at 1140 °C, and the overall plant yield follows it from 73.3 % to 92.0 % across the same span. The refractory limit is 1160 °C. Below about 1050 °C the kiln converts too little for

the leach to recover, and the loss is permanent because no downstream unit acts on unconverted alpha spodumene.

The fuel valve is a first-order lag of 0.3 min and the bed thermocouple carries 1.6 min of transport delay.

10.3.3 Leach

The leach is one stirred tank carrying three species balances and an energy balance. Extraction is first order in the unleached lithium and first order in the free hydroxide, as shown in Equations (10.4) to (10.6).

$$r = V_L\, k_L(T_L)\, C_s\, C_R \tag{10.4}$$

$$V_L \frac{\mathrm{d}C_s}{\mathrm{d}t} = F_{\text{in}} C_{s,\text{in}} - F_{\text{out}} C_s - r$$

$$V_L \frac{\mathrm{d}C_{Li}}{\mathrm{d}t} = r - F_{\text{out}} C_{Li} \tag{10.5}$$

$$V_L \frac{\mathrm{d}C_R}{\mathrm{d}t} = F_R C_{R,f} - F_{\text{out}} C_R - \nu\, r \tag{10.6}$$

where $V_L = 150\,\mathrm{m}^3$ is the reactor volume, C_s is the unleached lithium in the solids in kg Li/m^3, C_{Li} is the dissolved lithium, C_R is the free hydroxide in mol/m^3, $k_L = 0.0111\,\mathrm{m}^3/(\mathrm{mol\,min})$ at 95 °C with $E/R = 7000\,\mathrm{K}$, F_R is the lime slurry flow, $C_{R,f} = 6000\,\mathrm{mol/m}^3$ is its strength, and $\nu = 144.1\,\mathrm{mol/kg}$ is the hydroxide consumed per kilogram of lithium taken into solution, which is one mole of hydroxide per mole of lithium. The rate r is in kg/min, and $F_{\text{out}} = F_{\text{in}} + F_R$ carries the reagent dilution that separates the concentration-basis and mass-basis extractions.

The measurement is the logarithm of the free hydroxide, as shown in Equation (10.7).

$$\mathrm{pH} = 14 + \log_{10}\left(\frac{C_R + C_{\text{buf}}}{1000}\right) \tag{10.7}$$

where $C_{\text{buf}} = 0.5\,\mathrm{mol/m}^3$ is the residue alkalinity floor and the divisor converts mol/m^3 to mol/L. Extraction against pH saturates sharply, at 84.7 % at pH 11.9, 95.7 % at 12.5 and 98.9 % at 13.1, on lime rates of 235, 279 and 332 L/min. Above about pH 12.8 the extraction is within two points of its ceiling and further reagent buys almost nothing, which makes the pH setpoint an operating cost decision rather than a quality one.

The leach temperature is held near 95.8 °C by a jacket steam setting that is an operator input rather than a controlled variable. The electrode is a 0.25 min first-order lag with optional bias and noise.

10.3.4 Evaporative Crystallizer

The evaporator is a volume and a solute inventory, as shown in Equation (10.8).

$$\begin{aligned} \frac{\mathrm{d}V_e}{\mathrm{d}t} &= F_e - F_p - F_v, \qquad F_v = k_v Q_a \\ \frac{\mathrm{d}(V_e C_e)}{\mathrm{d}t} &= F_e C_{e,\text{in}} - F_p C_e \end{aligned} \tag{10.8}$$

where F_e is the filtrate feed, F_p is the product draw, F_v is the vapor rate, Q_a is the applied duty in MW behind a 0.8 min lag, and $k_v = 0.0955\,\mathrm{m}^3/(\mathrm{min\,MW})$. The vapor carries no lithium, which is why the liquor concentrates. Expanding the product rule gives the form integrated by the application, $\dot{C}_e = (F_e C_{e,\text{in}} - F_p C_e - C_e \dot{V}_e)/V_e$.

The level is an integrating state held by a dedicated draw controller that runs in both manual and automatic modes. It completes 63 % of a 0.3 m setpoint change in 2.5 min and 95 % in 4.3 min, then overshoots by 13.5 % and stays within 5 % of the new setpoint from 17.2 min. The concentration is

Table 10.1 First-order plus dead-time models identified from one step per unit (2 % fuel, 5 L/min lime, 1 % duty), measured at the transmitter. The two-point fit of the leach returns a negative dead time, so its entry bounds the delay before the transmitter first moves.

Loop	Gain	τ_p (min)	θ_p (min)
Kiln temperature	18.7 °C/%	28.4	1.86
Leach pH	0.0159 pH per L/min	55.0	< 0.3
Product strength	0.258 kg Li/m^3 per %	45.1	5.96

far slower. Its open-loop time constant is the liquor residence $V_e/F_p = 44$ min, and a step in duty moves it with a measured time constant of 45.1 min through an analyzer that adds 4.0 min of transport delay and a 1.2 min lag. Product below 11.5 kg Li/m^3 is unsaleable, and above the 15.5 kg Li/m^3 saturation line lithium hydroxide monohydrate deposits on the heating surfaces rather than in the crystallizer, so the concentration loop approaches a hard constraint through five minutes of dead time.

Scaling on those surfaces reduces the evaporation obtained per unit of duty over weeks of operation, which the application applies as a step in k_v within a single run.

10.3.5 Initialization, Numerics and Assumptions

The eleven states are the kiln temperature and conversion, the fuel valve position, the three leach concentrations, the leach temperature, the evaporator volume and concentration, the applied evaporator duty and the applied jacket steam. Initialization solves the three units in flow order at the given actuator settings, iterating the kiln because the reaction is a heat sink whose size depends on the temperature that drives it, and bisecting the leach hydroxide balance because the residual is monotone in C_R while a fixed-point substitution is not. The result is an exact equilibrium, so a run that is left alone does not drift. A manual-mode helper inverts the same steady-state solution to return the fuel, lime and duty that place all three controlled variables on their setpoints at the present feed, as far as the actuator ranges allow.

The engine of Section 1.6.1 integrates the model by fixed-step fourth-order Runge-Kutta at 0.01 min (0.6 s) and runs the controllers on a 0.1 min (6 s) scan, so every measurement is a sampled one and each loop carries half a scan of additional lag.

The calciner is one well-mixed zone rather than an axial profile, which is adequate for a control argument and not for a kiln design [15]. Impurity removal, the sulfate route and the final crystallization to lithium hydroxide monohydrate are outside the boundary.

10.4 Dynamic Analysis

10.4.1 Identified Models of the Three Units

Each unit is identified the way a commissioning engineer identifies one, by stepping its actuator in manual and reading the response off the transmitter. The two-point fits are given in Table 10.1.

None of the three identified plants is difficult to control in isolation. The kiln is close to a textbook first-order plus dead-time process with a dead-time fraction of 0.07. The leach carries no transport delay at all, only the electrode lag. The product loop has a dead-time fraction of 0.13 and approaches a hard saturation constraint rather than a graded one, and it is the loop that has to be detuned.

10.4.2 Propagation Speed Through the Train

The same actuator reaches three plant variables at three different times, as listed in Table 10.2.

Table 10.2 carries the measurement the application is organized around. The fuel valve reaches the conversion in fourteen minutes, the temperature it is nominally controlling in twenty-nine, and the product concentration two units downstream in nearly two hours, with ninety-five percent of the last taking four hours.

Table 10.2 Measured open-loop response of three plant variables to the same 5 % fuel valve step. The actuator is one valve. The step carries the bed past the 1160 °C refractory limit 38.8 min after it is made, which the application flags, so scenario 1 uses a 3 % step.

Variable	Unit	τ_{63} (min)	τ_{95} (min)
Spodumene conversion X_k	calciner	14.0	51.8
Kiln temperature T_k	calciner	29.3	87.8
Product concentration C_e	crystallizer	112.8	239.7

The conversion is faster than the temperature that drives it, which is a consequence of the saturation in Equation (10.2) rather than of any extra dynamics. The conversion has a time constant of $1/(k+1/\tau_s) = 2.6$ min about the design point, so it tracks the quasi-steady value $k\tau_s/(1+k\tau_s)$ with little lag. That value is a concave function of temperature, and when the bed has completed 63 % of a 95 °C rise the quasi-steady conversion has already completed 86 % of its own change. Approaching a saturating characteristic makes a downstream variable look faster than its driver, and the effect is not a modeling artifact.

10.4.3 Operating-Point Dependence of the pH Loop Gain

The pH loop is the one whose gain is not a constant, and the way it varies is not the textbook titration curve. Measured open-loop gains for a 5 L/min step are 0.0104 pH per L/min at pH 11.9, 0.0159 at 12.5, 0.0116 at 12.8, 0.0067 at 13.1 and 0.0034 at 13.4. The gain rises to a maximum at the design point and then falls by a factor of 2.4 by pH 13.1 and 4.7 by pH 13.4.

Two mechanisms compete. Below the design point almost all of the added hydroxide is consumed by (10.4), because the extraction is still climbing, so reagent buys extraction rather than free alkalinity and the pH moves little. Above it the extraction is saturated, the added hydroxide accumulates as C_R, and the logarithm of (10.7) compresses a tenfold concentration change into one pH unit. A controller tuned at the design point therefore carries about one fifth of its design loop gain by pH 13.4, which is where an instrument bias or a low ore grade will put it. The interior maximum makes the design point the safe place to tune, because the loop gain falls in both directions and the loop grows sluggish rather than oscillatory. Tuning the same loop at pH 13.1 and returning the plant to the design point reverses the argument, because the loop gain is then 2.4 times the value it was tuned for.

Scenario 2 shows the falling gain from two 30 L/min lime steps, from 279 to 309 and then to 339 L/min, each held 150 min. The hold has to be that long, because the pH covers two thirds of the first step only after 47 min. Read after 45 min, the two steps give 0.25 and 0.26 pH and the comparison reverses, while the 150 min holds give 0.38 and 0.25 against settled changes of 0.40 and 0.24.

10.5 Controller Design

10.5.1 Velocity Form PID Algorithm

All four loops, and the inner fuel flow loop when the kiln is in cascade, run the velocity-form PID of Section 1.6.2, Equations (1.10) to (1.13), with the derivative on the measurement [20] and a control period $\Delta t = 0.1$ min. The kiln, pH and product strength loops, and the inner fuel flow loop, are reverse acting with $a = +1$, because more fuel, lime or duty raises the measurement. The level loop is direct acting with $a = -1$, because more product draw lowers the level. The derivative of the kiln loop acts on a measurement filtered with $\tau_f = 0.5$ min, and τ_D is zero in every shipped and recommended tuning. The output is clipped to the actuator span and back-calculation with a tracking time $T_t = \tau_I$ corrects the integral alone. No rate limit and no further clamp act between the controller output and the actuator.

Among the runs reported here only the scaling case of scenario 9 puts an actuator on a limit, the evaporator duty at 227.3 min, and no reported number depends on the tracking time. Transfer

Table 10.3 Recommended and cautious tunings. The kiln loop follows the SIMC rules at $\tau_c = \theta_p$ and the other two are detuned from it.

	Recommended		Cautious	
Loop	K_c	τ_I (min)	K_c	τ_I (min)
Kiln temperature	0.42	15	0.08	40
Leach pH	240	25	40	60
Product strength	13	35	3	80
Inner fuel flow	3	0.4	3	0.4
Level	6	8	6	8

between Manual and Auto is bumpless in both directions, as Section 1.6.2 describes. With the kiln in single-loop mode the inner flow controller tracks the valve command in the same way. Output limits are 0 % to 100 % on the fuel valve and the evaporator duty, 0 L/min to 900 L/min on the lime slurry and 0 m³/min to 4 m³/min on the product draw.

10.5.2 Loop Pairing and Tunings

Kiln temperature goes to the fuel valve, leach pH to the lime slurry, product concentration to the steam duty and crystallizer level to the product draw. The level loop runs in both modes, because an inventory that nobody holds does not wait for an operator to select automatic. The recommended tunings are listed in Table 10.3 beside the cautious set that the application loads at start-up, in the three manual scenarios and in the challenge.

The SIMC rules give $K_c = \tau_p/(K_p(\tau_c + \theta_p))$ and $\tau_I = \min(\tau_p, 4(\tau_c + \theta_p))$ [21]. Applied to the kiln row of Table 10.1 with the standard aggressive choice $\tau_c = \theta_p$, they return $K_c = 0.41$ and $\tau_I = 14.9$ min against the 0.42 and 15 min recommended. The product loop is detuned to $\tau_c = 7.5$ min against a dead time of 5.96 min because it approaches the saturation constraint through that dead time, and the pH loop to $\tau_c = 14.1$ min because its identified dead time is under 0.3 min, where the aggressive choice $\tau_c = \theta_p$ would ask for a controller gain near 5800 L/min per pH unit against an actuator that spans 900 L/min. The integral times shipped for the two detuned loops, 25 and 35 min, are shorter than the 55 and 45 min that $\min(\tau_p, 4(\tau_c + \theta_p))$ returns for them.

The kiln loop is offered two ways. Single-loop, the temperature controller positions the fuel valve directly. In cascade, it sets a fuel flow setpoint and a fast flow controller positions the valve, which is the classical arrangement for a fired heater [8] and the one rotary kilns use [13, 14]. The inner loop settles in well under a minute against an outer time constant of 28.4 min, so the separation of time scales that cascade control requires is present by a factor of more than thirty [5].

10.6 Results

10.6.1 Cascade Against Three Fuel-Side Disturbances

Each disturbance is applied alone to the plant at its design steady state and followed for 300 min under both architectures at the same outer tuning. The runs use scenario 5 with the architecture selected before the run and the disturbance slider or setpoint moved at 20 min, on a controller scan. Results are given in Table 10.4.

Table 10.4, drawn as a chart in Figure 10.2, separates the two halves of the classical cascade argument, and the two are of very different size. Against a header pressure loss the cascade reduces the peak by a factor of 34 and the integral error by a factor of 93, because the inner loop measures the flow that the pressure changed and restores it to within 0.5 % in 0.6 min, before the bed temperature has moved more than 0.85 °C. The inner loop acts only at its 0.1 min scan, so a loss arriving just after a scan goes uncorrected for nearly a whole scan, and the factors fall to 18.9 and 52. Against a heating value loss the cascade reduces the peak by a factor of 1.05 and the integral error not at all, because

Table 10.4 Kiln temperature loop, single loop and cascade, one disturbance at a time from the design steady state. Peak is the largest absolute deviation from setpoint in °C and IAE is in °C min over 300 min. The setpoint row gives the overshoot past the new setpoint, because the peak error of a setpoint step is the step itself. A change that arrives one physics step after a scan raises the cascade header entries to 1.52 and 10.2 (1.56 and 28.8 at the cautious tuning) and each setpoint IAE by 2.7, and leaves the other entries unchanged.

	Single loop		Cascade	
Disturbance	Peak	IAE	Peak	IAE
Recommended tuning				
Setpoint +30 °C	5.15	117.7	3.75	**107.2**
Header pressure −20 %	28.79	526.3	**0.85**	**5.7**
Heating value −7 %	9.38	158.4	8.96	158.4
Cautious tuning				
Header pressure −20 %	79.25	7311	**0.88**	**16.0**

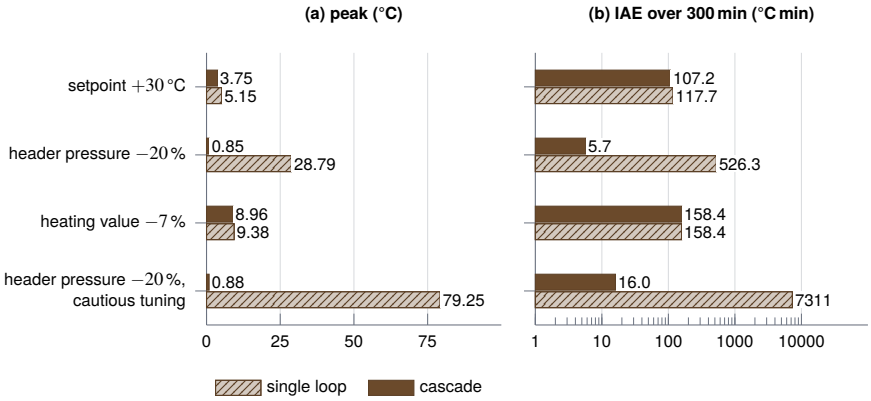


Figure 10.2 The kiln cascade of Table 10.4 as a chart, single loop hatched and cascade solid; the setpoint row's peak is the overshoot. Against a header pressure loss the cascade bar nearly vanishes at either tuning, while against a heating value loss the two bars are almost equal and the integral errors identical: the inner loop rejects only what its flow meter can see.

the mass flow of gas is unchanged, the inner loop sees nothing, and the whole energy error has to be found by the outer loop through the 1.86 min identified dead time. On a setpoint change it improves the integral error by a factor of 1.10, because removing the valve lag from the outer loop path is all that is left to remove and the 28.4 min thermal lag is untouched.

🕮 **Pause.** Try Exercise 10.1 before reading on.

The inner loop's rejection is nearly independent of the outer tuning. At the cautious tuning the same header dip costs 79.25 °C single loop and 0.88 °C in cascade, so detuning the outer loop by a factor of 5.3 in gain multiplies the single-loop peak by 2.8 and leaves the cascade peak within 0.03 °C of where it was. The two effects do not compose, because they act on different paths. Scenario 4 delivers the same comparison as one scripted sequence, a 25 % header loss at 20 min, its recovery at 90 min and a 10 % heating value loss at 130 min. The kiln falls 37.2 °C and then overshoots 43.1 °C single loop against 1.2 °C and 1.0 °C in cascade, while the heating value loss costs 11.7 °C single loop and 13.0 °C in cascade. The single-loop figure is the smaller only because that kiln is still 2.9 °C above setpoint from the overshoot when the loss arrives.

Table 10.5 A single upstream grade change walking through the train. Every controlled variable returns to setpoint.

t (min)	T_k (°C)	pH	Lime (L/min)	C_e (kg Li/m^3)	Output (kg/min)
19	1090.00	12.500	279.1	12.000	10.870
25	1091.26	12.562	264.1	12.000	10.774
40	1090.42	12.662	221.1	11.967	10.187
90	1090.01	12.531	203.0	11.787	8.773
220	1090.00	12.500	207.1	11.961	7.988
400	1090.00	12.500	207.1	11.998	7.935

> **TRY IT IN THE APP** **4 · Cascade on the Kiln**
>
> Run the four hours with Kiln architecture on Single loop. Then press Reset, switch it to Cascade and run again. The 25 % header pressure loss at minute 20 drops the kiln 37.2 °C on the single loop and 1.2 °C in cascade. The 10 % heating value loss at minute 130 costs 11.7 °C and 13.0 °C: the inner loop cannot see it.

10.6.2 Propagation of an Ore Grade Change

Scenario 6 drops the ore grade from 5.5 to 4.0 % Li_2O at twenty minutes with every setpoint and tuning unchanged. The trajectory is given in Table 10.5.

Table 10.5 records the central plantwide experiment, and every excursion in it is small. The kiln temperature peaks 1.26 °C high five minutes after the change and comes back. The pH peaks 0.163 high after twenty-two minutes and comes back. The product concentration reaches 0.214 kg Li/m^3 low after seventy-four minutes and comes back. Each loop does what it was designed to do, the largest excursion anywhere on the plant is a sixth of a pH unit, and the plant output falls from 10.87 to 7.93 kg/min, a reduction of 27 %.

> **TRY IT IN THE APP** **6 · The Ore Grade Falls**
>
> Load the scenario and press Run for four hours. Watch the lithium-through-the-plant trend. The grade drops at minute 20. The kiln temperature rises 1.26 °C and comes back, the pH rises 0.163 and comes back, and the product strength dips 0.214 kg Li/m^3 and comes back, yet the lithium recovered as product falls from 10.87 toward 7.93 kg/min and stays down.

The order of arrival is the instructive part. The kiln moves first, and it moves up, because the reaction is a heat sink and less spodumene is a smaller sink. The lime flow follows within about a minute, because less lithium in the feed is less hydroxide demand and the free alkalinity climbs, so the pH controller correctly cuts reagent by 26 % at the new steady state. The evaporator is last and is the unit under pressure, because weaker liquor has to be boiled harder to reach the same product strength, and its duty climbs from 77.1 to 86.6 %. The plant is still moving at 220 min and settles at about 400 min, six hours after an instantaneous change at the plant inlet.

Pause. Try Exercise 10.2 before reading on.

Nothing in this scenario is a control failure and nothing in it is remedied by a controller. What the scenario demonstrates is that a plant has a settling time that is not any loop's settling time, and that the correct response to a grade change is a production plan rather than a tuning change.

10.6.3 Grade Increase and Throughput Increase

Scenario 7 runs the same experiment upward, to 6.6 % Li_2O, and dilutes the lime slurry to 70 % strength at 150 min. Production settles at 12.93 kg/min and the product concentration peaks at 12.15 kg Li/m^3, well below the 15.5 kg Li/m^3 saturation line. The pH excursion of 0.251 is the largest

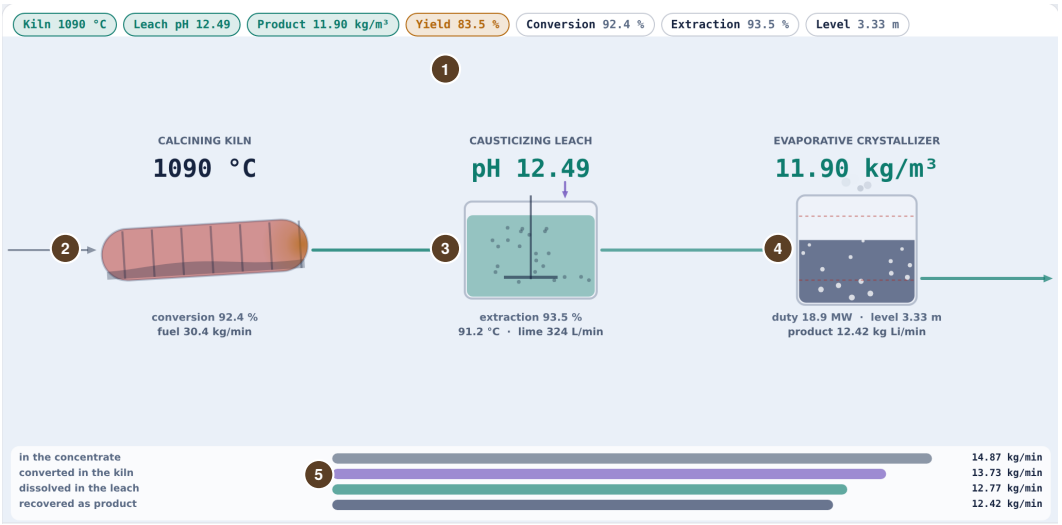


Figure 10.3 Scenario 8 at 2.08 h (124.8 min), 100 min after a twenty percent throughput increase, cropped to the plant view. All three loops are back at or near setpoint, with the kiln at 1090.0 °C, the leach at pH 12.492 and the product at 11.90 kg Li/m^3. (1) The badge row carries the three controlled variables and the yield. (2) The calcining kiln shows its bed temperature, conversion and fuel, (3) the leach its pH, extraction and lime flow, and (4) the evaporative crystallizer its product strength, duty and level. (5) The lithium balance beneath the flowsheet accounts for the metal at each stage: 14.87 kg/min in the concentrate, 13.73 converted in the kiln, 12.77 dissolved in the leach and 12.42 recovered as product, for a yield of 83.5 % against 87.7 before the change.

of the three feed-side scenarios, because the reagent loop is being asked for more lime at the moment each liter carries less hydroxide, and the lime flow ends at 472 L/min against 279 L/min at design. The flow controller has no defense against a reagent concentration change, because the disturbance is in the strength of the stream rather than in its rate.

Scenario 8 raises the concentrate feed by 20 % instead. The kiln temperature dips 17.71 °C while the fuel catches up and rises 17.95 °C when the feed is restored, the pH holds to 0.084 and the product concentration moves 0.921 kg Li/m^3, which is the largest single deviation the application produces from a normal operating change and nearly the whole of the graded run's ±1.0 kg Li/m^3 band. The dip that follows the push holds the product below the 11.5 kg Li/m^3 sales limit for 39.3 min. The kiln temperature loop holds 1090 °C without difficulty, and the conversion falls from 0.9355 to 0.9236 anyway, because the solids residence time falls from 40 min to 33.3 min and no controller measures conversion. That unmeasured loss carries through the train, and the yield falls from 87.7 % to 84.4 % with every controlled variable on setpoint, as shown in Figure 10.3. The crystallizer is the unit that reaches its capacity first, which is information required before a throughput increase rather than after one.

10.6.4 Evaporator Scaling

Scenario 9 steps the evaporator surface to 94, 86 and 72 % of clean duty at 30, 120 and 210 min. The concentration loop absorbs the first two steps, and for three hours the concentration stays within 0.41 kg Li/m^3 of setpoint while the duty command climbs. The duty command is the variable that carries the information. After the third step the product leaves the saleable range, below 11.5 kg Li/m^3, at 217.2 min, ten minutes before the duty reaches its limit at 227.3 min. The saturation flag therefore

Table 10.6 A pH electrode reading low. Every reported number improves. Values are instantaneous readings with the shipped noise seed, and the noise moves the lime flow by 6.9 L/min (one standard deviation).

t (min)	True pH	Measured	Lime (L/min)	Yield (%)
35	12.508	12.515	276.3	87.95
80	12.863	12.526	310.8	89.24
200	13.170	12.471	369.8	90.98
500	13.203	12.470	356.7	**91.26**

arrives too late to serve as the warning, and a plant that trends the manipulated variable against its clean value sees the loss of capacity hours earlier.

10.6.5 Bias on the pH Electrode

Scenario 10 puts a bias on the pH electrode, 0.35 low at forty minutes and 0.70 low at one hundred and sixty, on top of 0.04 of electrode noise from a fixed seed. The consequences are given in Table 10.6.

Table 10.6 records the effect of an electrode bias on the reported performance of the plant. The controller holds the reading at 12.50, so it adds lime until the true pH reaches 13.20, which is five times the free hydroxide the leach needs. Averaged over the first 40 min and over 300 min to 500 min, extraction rises from 95.8 to 99.1 % because the leach is in fact operating at a higher free hydroxide concentration. The yield rises from 87.8 to 91.2 % and production from 10.88 to 11.30 kg/min.

> ✦ DID YOU KNOW? **A Faulty Electrode That Looks Like an Improvement**
>
> From the simulator: at minute 500 of Table 10.6 the electrode shows pH 12.470, close to its setpoint, while the true pH is 13.203 and the lime runs at 356.7 L/min against 276.3 before the fault. The yield column climbs from 87.95 to 91.26 %, so the fault reads as a gain.

Every number on the daily report improves. The reagent consumption rises 24 %, from 280 to 348 L/min on the same averages (24.1 % to 25.1 % over eight other noise seeds), and it is the only symptom. The pH trend itself is flat, and a plant judged on yield alone records an improvement. Detecting the fault requires evidence from outside the loop, from a reagent consumption per tonne figure, a laboratory titration or a second electrode.

🕮 **Pause.** Try Exercise 10.3 before reading on.

10.6.6 Graded Run

The challenge starts from the cautious tunings with a single-loop kiln and applies a grade change to 4.6 % Li_2O at 20 min, a 7 % gas heating value drop at 70 min, an evaporator scaling step to 86 % of clean duty at 120 min and a grade recovery to 5.9 % at 170 min. The thermocouple, electrode and analyzer carry noise of 1.2 °C, 0.03 pH and 0.06 kg Li/m^3 from a fixed seed. The run is judged at 300 min on the largest deviation of each true controlled variable over the whole run, with the kiln within ±12 °C, the leach within ±0.15 pH, the product within ±1.0 kg Li/m^3, the kiln below the 1160 °C refractory limit and the product below the 15.5 kg Li/m^3 saturation line. Results with one setting changed at a time are given in Table 10.7.

Table 10.7 isolates each clause to one decision. The kiln loop's tuning owns the kiln clause, at 9.29 °C against 24.53 °C. The pH loop's tuning owns the pH clause, at 0.135 against 0.315, and the concentration loop's tuning owns the concentration clause, at 0.789 kg Li/m^3 against 1.942 kg Li/m^3. Influence runs only downstream, which is the one-way material coupling of Figure 10.1. The concentration loop leaves the pH clause unchanged to the third decimal, while the cautious pH loop moves the product deviation by 0.023 kg Li/m^3 and the cautious kiln loop moves the pH and product deviations by 0.006 and 0.037 kg Li/m^3.

Table 10.7 The graded run. Largest deviation from setpoint over the 300 min judged run, against bands of ±12 °C, ±0.15 pH and ±1.0 kg Li/m^3. Rows below Recommended change one setting of it.

Configuration	T_k (°C)	pH	C_e (kg Li/m^3)	Result
As loaded	24.66	0.358	1.980	fail
Recommended	**9.29**	**0.135**	**0.789**	pass
single-loop kiln	9.59	0.135	0.789	pass
cautious kiln loop	24.53	0.141	0.826	fail
cautious pH loop	9.29	0.315	0.766	fail
cautious concentration loop	9.29	0.135	1.942	fail

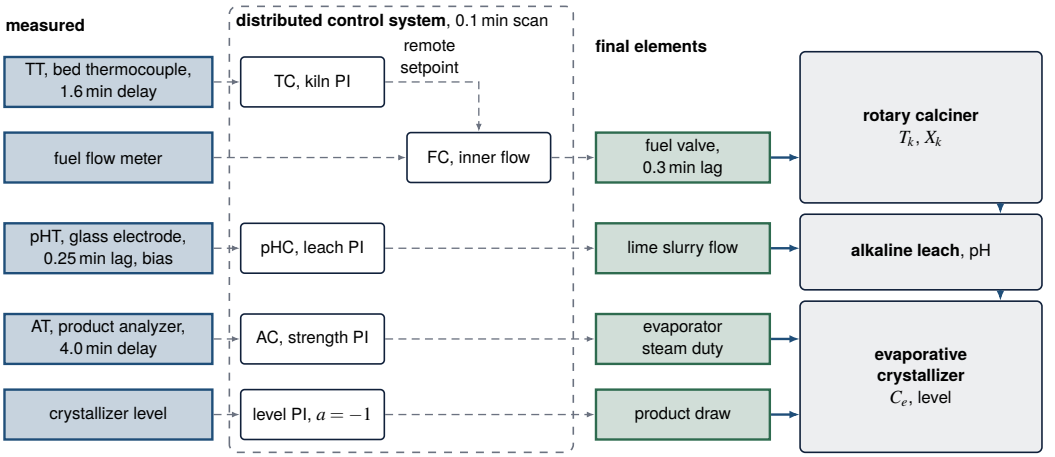


Figure 10.4 The loops of the simulated refinery drawn with their hardware. In cascade the kiln controller TC writes the setpoint of the fuel flow controller FC, which moves the valve; in single loop TC moves the valve directly. The three units downstream of the valve are joined in series by the material, with no storage between them, and every block runs on the same 0.1 min scan.

The cascade row is the one that qualifies Table 10.4. On this disturbance sequence the cascade changes the peak kiln excursion from 9.59 °C to 9.29 °C and nothing else, because the sequence contains a heating value change and no fuel supply disturbance, and because the outer loop at its recommended tuning absorbs what remains. The difference is smaller than the spread of the kiln deviation over eight other noise seeds, 8.36 °C to 10.29 °C in cascade. The cascade is worth a factor of 52 or more on integral error in Table 10.4 and three tenths of a degree here. Both numbers are correct and what differs is the disturbance.

The recommended configuration passes, and the pH clause is the one with the least margin. Its largest pH deviation, 0.135 at 190.5 min, comes twenty minutes after the grade recovery, and over eight other noise seeds it lies between 0.135 and 0.142, so the configuration passes on every seed tried with as little as 0.008 to spare. The product deviation peaks at 141 min, 21 min after the scaling step, because the evaporator's own response takes 45 min while the scaling step removes evaporation capacity in one move. As loaded, the run fails all three bands, the kiln by 12.7 °C after the heating value drop, the pH by 0.21 and the product by 0.98 kg Li/m^3.

10.7 Technology in Practice

Figure 10.4 draws the four loops of Section 10.5 and the inner fuel flow loop with their hardware.

What Measures It

Each unit has one measurement that its loop acts on. The kiln bed is read by a thermocouple. Its most common failure is to fail open, or burn out, and a furnace transmitter is set to read high when that happens, so the controller cuts the fuel rather than adding it [22]. The 1.6 min of transport delay on the bed reading in Section 10.3.2 stands in for the thermowell and the bed between the hot charge and the junction. The leach is read by a pH electrode, drawn in Figure 10.5.

> **TECHNOLOGY** **Sensor: The Glass pH Electrode**
>
> A thin bulb of glass that hydrogen ions can penetrate develops a voltage proportional to the difference between the pH of the process and that of a buffer sealed inside it, about 59 mV per pH unit at 25 °C and more when hot, so the transmitter needs the solution temperature. A reference electrode with a porous junction closes the circuit. The instrument is calibrated with two buffers, and a clean probe settles in each within about 30 s. The glass wears as it hydrates, so the slope falls with age, while stray potentials at the junction shift the zero, an offset of the kind scenario 10 applies. Sodium, potassium and lithium ions interfere at high pH [22].

The 0.25 min lag of the model's electrode stands for the response of the glass, and its optional bias for the shifted zero of an aged or fouled junction. The product strength comes from an analyzer that the model places 4.0 min downstream with a 1.2 min lag of its own.

What Moves It

The kiln's final element is the fuel valve, with a 0.3 min lag in the model, and in cascade a flow controller sits between it and the temperature controller.

> **TECHNOLOGY** **Final Element: The Fuel Valve Under a Flow Loop**
>
> The fuel enters through a control valve whose throttling changes with the pressure drop across it, more responsive near closed than near open. Placing a flow controller under the temperature controller confines that distortion to the flow loop, and a change in supply pressure is countered by the flow controller before the master sees it. The flow meter measures quantity, not energy: a heating value is computed from the composition of the gas. An air-to-open valve fails closed on loss of air, and the risks of the process, not habit, should decide the failure mode [22].

The lime slurry and the evaporator steam duty are the other two manipulated flows, and a draw on the product holds the crystallizer level. The model moves each of them through its own lag and clips it to its range, with no dead band or stiction.

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 [23]. In the app every controller, the inner flow loop included, acts on a 0.1 min scan against a kiln time constant of 28.4 min. On such a platform the four loops and the inner flow loop are PID function blocks of the regulatory layer, the cascade a remote setpoint from one block to the next.

> **IN PRACTICE** **Commissioning a Cascade**
>
> The secondary controller is tuned first, whenever a cascade is commissioned or optimized, because the master depends on it and it does not depend on the master. With the master in manual the secondary simply holds a constant setpoint. The secondary then has three modes: manual, automatic on a local setpoint, and cascade on the master's remote setpoint. A secondary left in automatic defeats the cascade, so it must stay in cascade mode for the strategy to work as designed [22].

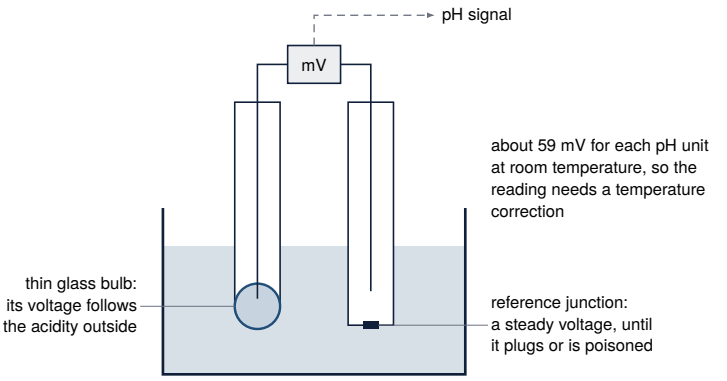


Figure 10.5 The glass pH electrode and its reference. Notice that the reading is a small voltage across the glass bulb, referred to a junction that can drift. A shift at the junction moves every reading by the same amount, which is the low bias of scenario 10 that the pH loop then hides.

What Is Modernizing It

The first development measures what the cascade cannot. A process gas chromatograph reads the composition of natural gas and a computer multiplies each part by its own heating value to give the heating value of the stream, which pipelines already use to price the gas [22]. Such a reading would show the heating value loss of Table 10.4 before the bed temperature does. The second is a new calciner. Fluidized heating of finely ground concentrate and microwave heating have been suggested to cut the excess heat of the rotary kiln [18]. The third is predictive control and estimation over the whole train, which Section 10.9 sets out: one calculation that holds the refractory limit and the saturation line as constraints, and an estimator that reconstructs the ore grade and the electrode bias [24].

10.8 Limitations

The calciner is one well-mixed zone with a single lumped conversion. A real rotary kiln has an axial temperature and conversion profile, a freeboard gas in counter-current, and a solids bed whose depth and residence time depend on rotation and slope, and integrated kiln simulators exist to capture those [15, 16]. Nothing here describes the condition of a kiln's hot end.

> **OPEN CHALLENGE** **The Energy of the Calciner**
>
> Calcination, acid digestion and purification are the stages that decide the energy use of a spodumene refinery, and the high energy consumption of the calcination step remains unsolved. A typical kiln is a fairly inefficient furnace fired on gas or coal, and its heat load has been put at 82 % above the thermodynamic need of the phase change, with the carbon dioxide that goes with it. Fluidized and microwave heating have been suggested [18]. Progress would be a calciner that reaches the same conversion of the ore on far less fuel, with the temperature control the chapter's kiln loop provides.

The phase transformation is one Arrhenius rate. The real conversion proceeds through an intermediate gamma phase whose extent depends on particle size, heating rate and impurities, and the reported activation energies span a wide range [1]. The application reproduces the saturating shape of Equation (10.2) and no particular kiln's numbers.

The leach is one tank with one lumped rate constant and no solids handling. A real circuit is several tanks with a filtration step, and a badly leached solid is the one that is difficult to filter. Extraction is first order in the manner used for other hydrometallurgical circuits [12] rather than shrinking-core. The reported extraction is on a concentration basis and therefore overstates the mass-basis extraction

by 0.5 percentage points at the design reagent rate, while the reported yield is computed from the mass balance and carries no such offset.

Impurity removal is entirely absent. On a real refinery the removal of aluminum, iron, calcium and magnesium between the leach and the crystallization is a substantial plant in its own right, and the product specification is dominated by those impurities rather than by concentration [3]. The product quality here is one number and should be read as a proxy for a duty rather than as a specification.

The crystallizer produces a concentrated liquor rather than crystals. There is no supersaturation, no nucleation and no crystal size distribution, so the saturation line is a limit rather than a mechanism, and the scaling disturbance is a step in evaporation per unit duty rather than a fouling model.

10.9 Discussion

Three results generalize beyond this refinery.

The first is that the value of a cascade depends on which disturbance it is asked to reject, and the two halves of the classical justification should be quoted separately. Table 10.4 shows the same inner loop worth a factor of 52 to 93 on integral error for a header pressure loss, 1.10 for a setpoint change and 1.00 for a heating value loss, and Section 10.6.6 shows it worth three tenths of a degree on a disturbance sequence containing neither. An inner loop rejects what enters inside its own measurement and nothing else. The corollary is that specifying a cascade requires naming the disturbances a plant actually faces, because the same structure is worth a factor of fifty or more against one of them and nothing against another.

The second is that a plant's settling time is not any loop's settling time. Table 10.5 has three loops with excursions of 1.26 °C, 0.163 pH and 0.21 kg Li/m^3, all returning to setpoint, on a plant that takes six hours to reach a new steady state 27 % below the old one. A regulatory scorecard records that day as compliant. Plantwide control exists because the scorecard is insufficient, and because the remedy is a change to the production plan that no local loop is permitted to make [5, 9, 19].

The third concerns instrument validation. Table 10.6 has a drifting electrode improving every figure the plant is judged on, because the controller compensating for it happens to push the process in a direction that helps. The only evidence is the reagent consumption, which is reported outside the control system and on a longer time scale. Instrument validation is therefore a control activity rather than a maintenance one [7], and reagent consumption is one of the largest operating costs a hydrometallurgical plant carries [10].

The documented model also supports control methods beyond the PID and cascade loops the application implements. The eleven differential states, the actuator limits and the two hard constraints match the problem form solved by nonlinear model predictive control, which would coordinate the fuel valve, the lime flow and the evaporator duty in one calculation, hold the 1160 °C refractory limit and the 15.5 kg Li/m^3 saturation line as constraints rather than as alarms, and run an economic objective on the gas and reagent bills against the lithium recovered, which is precisely the trade-off the pH setpoint sweep of Section 10.3.3 exposes and no local loop can resolve [24]. Moving horizon estimation over the same equations would reconstruct the ore grade and the evaporator surface condition from the kiln temperature, pH and product concentration histories, converting the two unmeasured disturbances of Table 10.5 and the scaling event into estimated quantities available to the controller. The same estimator would supply a continuous soft sensor for the spodumene conversion, which is measured nowhere and which falls by 0.012 during the throughput increase with every controlled variable on setpoint, and estimating the electrode bias of Table 10.6 as a parameter would surface that fault from inside the control system rather than from the reagent report. The Gekko Optimization Suite implements these methods in Python from model statements close to (10.1) through (10.8) [25].

10.10 Conclusion

A lithium hydroxide refinery puts three units in series with no storage between them, so one fuel valve reaches the spodumene conversion in 14.0 min, the kiln temperature it is controlling in 29.3 min and the product concentration two units downstream in 112.8 min. Each unit is identified as a first-order plus dead-time plant from one step test, and the three identified models set the tunings through the SIMC rules. A cascade on the kiln reduces the integral error of a header pressure loss by a factor of 52 to 93, a setpoint change by 1.10, a heating value loss by 1.00 and a graded disturbance sequence by three tenths of a degree of peak, so the disturbance rather than the structure decides what the cascade is worth. A single grade change at the front takes six hours to walk through the train, moves no controlled variable by more than a sixth of a pH unit, and removes 27 % of the plant output. An electrode reading 0.70 low raises the extraction, the yield and the production while raising the reagent consumption by 24 %, which is a failure that a plant measuring only its controlled variables is structurally unable to see. The app lets a reader run each of these results on a train whose constraint moves between units.

KEY POINTS

- With no storage between the units, one fuel valve reaches the conversion in 14.0 min, the kiln temperature in 29.3 min and the product in 112.8 min.
- A fuel flow cascade cuts the integral error of a header pressure loss by a factor of 52 to 93 and of a heating value loss not at all.
- A grade change to 4.0 % Li_2O moves no controlled variable by more than a sixth of a pH unit and removes 27 % of the output.
- An electrode reading 0.70 low raises the yield while the reagent rises 24 %.

Where else this applies. Chapter 4 and Volume 3, Chapter 2 sort a cascade's disturbances the same way, and Volume 6, Chapter 7 poses the plantwide question for a reactor and a column.

References

[1] A. A. Abdullah, H. C. Oskierski, M. Altarawneh, G. Senanayake, G. Lumpkin, and B. Z. Dlugogorski, "Phase transformation mechanism of spodumene during its calcination," *Minerals Engineering*, vol. 140, p. 105883, 2019.

[2] N. K. Salakjani, P. Singh, and A. N. Nikoloski, "Production of lithium: A literature review part 1: Pretreatment of spodumene," *Mineral Processing and Extractive Metallurgy Review*, vol. 41, no. 5, pp. 335–348, 2019.

[3] C. Dessemond, F. Lajoie-Leroux, G. Soucy, N. Laroche, and J.-F. Magnan, "Spodumene: The lithium market, resources and processes," *Minerals*, vol. 9, no. 6, p. 334, 2019.

[4] U.S. Geological Survey, "Mineral commodity summaries 2025: Lithium," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

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

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

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

[8] F. G. Shinskey, *Process Control Systems: Application, Design, and Tuning*, 4th ed. New York: McGraw-Hill, 1996.

[9] D. Hodouin, "Methods for automatic control, observation, and optimization in mineral processing plants," *Journal of Process Control*, vol. 21, no. 2, pp. 211–225, 2011.

[10] B. A. Wills and J. A. Finch, *Wills' Mineral Processing Technology: An Introduction to the Practical Aspects of Ore Treatment and Mineral Recovery*, 8th ed. Oxford: Butterworth-Heinemann, 2015.

[11] M. Samouhos, K. Spyrou, I. Pagonis, P. Oustadakis, and M. Taxiarchou, "Leaching of α-spodumene with sodium hydroxide: Optimization of lithium extraction using thermodynamic and XPS data," *Journal of Sustainable Metallurgy*, vol. 11, no. 3, pp. 2883–2897, 2025.

[12] J. McLaughlin and G. E. Agar, “Development and application of a first order rate equation for modelling the dissolution of gold in cyanide solution,” *Minerals Engineering*, vol. 4, no. 12, pp. 1305–1314, 1991.

[13] L. P. Holmblad and J.-J. Østergaard, “Control of a cement kiln by fuzzy logic techniques,” *IFAC Proceedings Volumes*, vol. 14, no. 2, pp. 809–814, 1981.

[14] K. S. Stadler, J. Poland, and E. Gallestey, “Model predictive control of a rotary cement kiln,” *Control Engineering Practice*, vol. 19, no. 1, pp. 1–9, 2011.

[15] K. S. Mujumdar and V. V. Ranade, “Simulation of rotary cement kilns using a one-dimensional model,” *Chemical Engineering Research and Design*, vol. 84, no. 3, pp. 165–177, 2006.

[16] K. S. Mujumdar, K. V. Ganesh, S. B. Kulkarni, and V. V. Ranade, “Rotary cement kiln simulator (RoCKS): Integrated modeling of pre-heater, calciner, kiln and clinker cooler,” *Chemical Engineering Science*, vol. 62, no. 9, pp. 2590–2607, 2007.

[17] R. B. Ellestad and K. M. Leute, “Method of extracting lithium values from spodumene ores,” U.S. Patent 2,516,109, application filed September 16, 1948, 1950, patented July 25, 1950.

[18] A. A. Abdullah, “Thermal treatment of spodumene ($LiAlSi_2O_6$) for lithium extraction,” Ph.D. dissertation, Murdoch University, Perth, Western Australia, 2019.

[19] S. Skogestad, “Control structure design for complete chemical plants,” *Computers & Chemical Engineering*, vol. 28, no. 1-2, pp. 219–234, 2004.

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

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

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

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

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

10.11 Exercises

10.1 ⌕ Concept check · 2 min

The kiln runs with the fuel flow cascade at the recommended tuning. Which fuel-side disturbance does the cascade remove almost entirely, and why?

(a) A loss of gas heating value, because the inner loop acts ten times a minute

(b) A loss of header pressure, because the inner loop measures the flow it changes

(c) Both losses equally, because the cascade takes the valve lag out of the outer loop

(d) Neither loss, because the outer loop still sets the fuel flow setpoint

10.2 ⌕ Concept check · 1 min

The concentrate grade falls from 5.5 to 4.0 % Li_2O with every loop in Auto. How do the kiln temperature and the lime flow first respond?

(a) The kiln cools and the lime flow rises

(b) The kiln cools and the lime flow falls

(c) The kiln heats and the lime flow falls

(d) The kiln heats and the lime flow rises

10.3 Concept check · 2 min

The pH electrode starts to read 0.70 low while the pH loop is in Auto at a setpoint of 12.50. Which describes the plant some hours later?

(a) The pH trend dips by 0.70 and the lime flow falls to restore it

(b) The true pH falls, so the extraction and the yield fall with it

(c) The true pH rises and the overdosed leach extracts less lithium

(d) The true pH rises, and extraction, yield and production rise with the lime

10.4 Theory and calculation · 6 min

The calciner conversion follows Equations (10.2) and (10.3) with $k_{\text{ref}} = 0.475/\text{min}$ at 1100 °C, $E/R = 50\,500\,\text{K}$ and a solids hold-up $M_s = 20\,000\,\text{kg}$.

(a) Compute k at 1090 °C to three significant figures and the steady conversion at a feed of 500 kg/min to four decimals. *Checkpoint:* 0.363/min; 0.9355

(b) The feed rises to 600 kg/min and the temperature loop holds 1090 °C. Compute the new residence time and steady conversion. *Checkpoint:* 33.3 min; 0.9236

(c) Linearize Equation (10.2) at fixed temperature about the design point and give the time constant of the conversion, to one decimal. Use it to explain why the conversion reaches 63 % of its change before the bed temperature does. *Checkpoint:* 2.6 min

10.5 Theory and calculation · 8 min

Use the identified models of Table 10.1 and the SIMC rules $K_c = \tau_p/(K_p(\tau_c + \theta_p))$ and $\tau_I = \min(\tau_p, 4(\tau_c + \theta_p))$.

(a) Tune the kiln loop with $\tau_c = \theta_p$. Give K_c to two decimals in % per °C and τ_I to one decimal. *Checkpoint:* 0.41; 14.9 min

(b) Tune the product strength loop with $\tau_c = 7.5\,\text{min}$. Give K_c to one decimal and τ_I to one decimal, and compare them with the recommended 13 and 35 min of Table 10.3. *Checkpoint:* 13.0; 45.1 min

(c) Tune the pH loop with $\tau_c = \theta_p = 0.3\,\text{min}$ (the bound of the table) and give K_c to the nearest 10 L/min per pH. What lime move would a pH error of 0.1 demand at once, and why is this τ_c unusable? *Checkpoint:* about 5770; about 577 L/min

10.6 Theory and calculation · 6 min

At design the calciner burns 26.53 kg/min of fuel with $\eta_f = 0.65$ and $\Delta H_f = 50\,000\,\text{kJ/kg}$. The feed is 500 kg/min at 25 °C with $c_{p,s} = 1.05\,\text{kJ/(kg K)}$ and 3 % moisture, $UA_k = 220\,\text{kJ/(min K)}$ with ambient at 25 °C, $T_k = 1090\,°\text{C}$, $X_k = 0.9355$ and $\Delta H_c = 100\,\text{kJ/kg}$ on a spodumene fraction $\phi = 5.5/8.03$.

(a) Compute the burner heat and the four sinks of Equation (10.1) in MJ/min, to one decimal, and check that they balance. *Checkpoint:* 862.2 in; 559.1, 234.3, 36.8, 32.0 out

(b) The grade falls to 4.0 % Li_2O ($\phi = 4.0/8.03$). At the same conversion, by how much does the reaction sink fall, and how far would the bed rise at fixed fuel if the sensible heat and the shell loss absorbed the difference? Give both to one decimal. *Checkpoint:* 8.7 MJ/min; 11.7 °C

(c) A run in Manual settles about 11 °C hot. Why is the true rise smaller than the estimate of (b)? *Checkpoint:* the hotter bed converts more

10.7 App challenge · 8 min

Open scenario 6, "The ore grade falls (5.5 → 4.0% Li_2O)", at the default seed, with every setting as loaded. Run it at 50 min per second to 400 min (the clock reads 6.67 h), then export the CSV. Its time column is Time (h), one row every 0.25 min.

(a) From the columns `kiln_C`, `ph` and `conc_kgm3`, give the largest deviation of each controlled variable from its setpoint after the 20 min grade change, with its sign and the time after the change at which it occurs (deviations within 0.01 in their own units, times within 0.5 min). *Checkpoint:* $+1.26\,^{\circ}\mathrm{C}$ at 5.3 min; $+0.163$ pH at 22.3 min; $-0.21\,\mathrm{kg\,Li/m^3}$ at 74.3 min

(b) Read `production_kgmin` and `li_feed_kgmin` at 19 min and at 400 min (within 0.005) and give the fall of each as a percentage of its 19 min value, to 0.1. *Checkpoint:* 10.870 to 7.935 kg/min, 27.0 %; 12.391 to 9.012 kg/min, 27.3 %

(c) Explain the order in which the three loops are disturbed, and why no retuning would recover the lost production. *Checkpoint:* kiln, leach, evaporator; the feed carries less lithium

10.8 App challenge · 8 min

Open scenario 4, whose name begins "Cascade on the kiln", at the default seed, and run at 20 min per second. The header pressure falls 25 % at 20 min and recovers at 90 min, and the heating value falls 10 % at 130 min.

(a) With Kiln architecture on Single loop as loaded, run to 240 min and export the CSV. From `kiln_C`, give the kiln deviation from 1090 °C at its lowest between 20 and 90 min, its highest between 90 and 130 min, its value at 130 min and its lowest after 130 min, each within 0.1 °C. *Checkpoint:* -37.2, $+43.1$, $+2.9$, -11.7

(b) Press Reset, select Cascade before pressing Run, run to 240 min and read the same four values. *Checkpoint:* -1.2, $+1.0$, 0.0, -13.0

(c) The heating value drop looks smaller under the single loop. Is the single loop better against it? Use the values at 130 min to explain. *Checkpoint:* no; it starts 2.9 °C hot

10.9 App challenge · 9 min

Open the challenge, "Challenge · take the plant through a grade change", at its default seed, and run each part at 20 min per second until it is judged at 300 min. Read the verdict line in the progress panel (kiln, pH and product deviations as it prints them).

(a) Run it as loaded. *Checkpoint:* fail: 24.7 °C, 0.36, 1.98 kg Li/m^3

(b) Reload it with its reload button, press Reveal suggested tuning and run. *Checkpoint:* pass: 9.3 °C, 0.13, 0.79 kg Li/m^3

(c) Reload, press Reveal suggested tuning, then set pH loop gain to 40 L/min per pH and pH loop integral time to 60 min before Run. Run, and explain which clauses moved and why the others did not. *Checkpoint:* fail: 9.3 °C, 0.31, 0.77 kg Li/m^3

10.10 Application area · 4 min

Production asks for 20 % more concentrate. Scenario 8 shows every loop back on setpoint while the yield falls from 87.7 to 84.4 %. As the process engineer, explain where the lithium is lost, which unit limits a further increase, and what you would measure or estimate before agreeing to the push.

10.11 Application area · 3 min

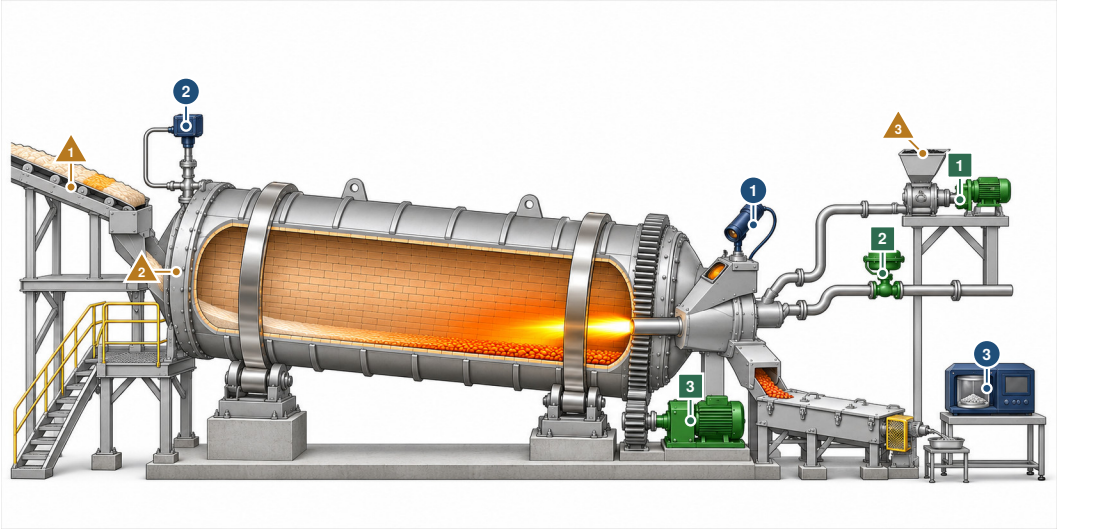
In scenario 9 the evaporator scales in three steps. An operator watches the product strength trend and the duty saturation flag. Why is that not enough, and what should the plant watch instead?

10.12 Review the agent · 5 min

An AI assistant reviews the pH loop and writes: "The process gain is 0.0159 pH per L/min at the design point but only 0.0034 at pH 13.4. A biased electrode can move the leach there, where the loop tuned at design will be 4.7 times too aggressive and oscillate. Retune at pH 13.4 so that the loop is safe everywhere." Find the flaw and say how you would check it in the app.

CHAPTER 11

Rotary Cement Kiln Simulator for Ratio Feedforward and Dead-Time-Limited Quality Control



The controller moves 1 the coal feeder, set by the temperature loop 2 the combustion air, set in ratio to the coal 3 the kiln drive, which sets the speed

The controller reads 1 the burning zone pyrometer 2 the oxygen analyzer at the feed end 3 the free lime test in the laboratory, about twenty minutes late

What disturbs it 1 the raw meal and how hard it burns 2 false air leaking in at the inlet seal 3 the heating value of the coal

Notice that the coal and the air enter together at the burner while the oxygen is read at the far end and the free lime only in the laboratory, which is why the air must follow the coal command rather than the oxygen reading. Illustration, not to scale.

IN PLAIN TERMS A Fast Flame and a Late Lab Test

The main idea. A cement kiln is a long, slowly turning steel tube. It bakes ground limestone and clay into hard lumps called clinker. Its flame answers a change in coal within seconds, but the lab test of the clinker comes back about twenty minutes late.

The challenge. If the coal goes up before the air does, the flame runs short of air and spoils the clinker. Setting the air from the coal order fixes that, but no setting makes a controller that waits for the lab test fast.

Why it matters. The kiln is where cement making burns its fuel, and clinker burned in a flame short of air cannot be repaired later.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	The first page, the fact box of Section 11.1, the history box of Section 11.2, the fact box and the two try-it boxes of Section 11.5, and the key points
🕮	**Study**	55 min	The whole chapter with its pause points and the twelve exercises
⛰	**Practice**	20 min	The first page, Section 11.5, Section 11.6, Section 11.7 and the open challenge in it

The control question A cement kiln's flame answers the coal in seconds while its free lime assay reports about twenty minutes late. Which control structure keeps the flame oxidizing through a firing move, and how fast can a loop closed on the late assay usefully be?

Learning objectives

- Explain why oxygen feedback alone cannot prevent the carbon monoxide excursion of a firing move, and why ratioing the air to the coal command removes it.
- Compute the excess air, the oxygen reading and the analyzer's error under false air from the combustion laws.
- Assemble the free lime loop's dead time budget and tune the loop by SIMC for its upper end.
- Judge a kiln configuration by its constraint tiles as well as its specific heat, using the graded challenge.
- Separate what a predictive controller gains from its optimization and what it gains from the information and limits it carries.

Prerequisites Sections 1.6.2 (velocity form and loop action), 1.6.4 (scenarios, tiles and exports), 1.4.5 (sampled assays) and 1.4.6 (ratio control). Chapter 7 meets the late assay first, and Chapter 10 the rotary kiln.

The app https://apmonitor.com/apps/cement-kiln/

See also Volume 3, Chapter 4 for cross-limiting and oxygen trim on a fired heater. Volume 5, Chapter 3 and Volume 5, Chapter 11 for a loop through a slow analyzer and an unmeasured kiln product. Volume 9, Chapter 4 for an offset-free predictive controller with a fuel bias estimate.

11.1 Introduction

A rotary kiln converts ground limestone and clay into cement clinker by carrying the material slowly down a rotating steel tube against a coal flame. Two time constants dominate its control. Raw meal takes 22.2 min to travel the kiln at 3.5 rpm, while the flame responds within seconds.

> ✦ DID YOU KNOW? **Four Billion Tonnes a Year**
>
> The world made an estimated 4.0 billion metric tons of cement in 2024, and the United States made its share at 99 plants [1]. Breaking down the limestone alone released 1.48 billion tonnes of carbon dioxide in 2017, about 4 % of the emissions from fossil fuels, before any kiln fuel is counted [2].

The solids transit and the laboratory delay are distinct: the meal takes about 22 min to travel the kiln, and the assay reports about 20 min after clinker leaves. The two time scales must be handled simultaneously. At the fast end, a change in coal rate reaches the burner tip after 0.4 min of transport and a 0.5 min feeder lag, and the resulting oxygen change reaches the back-end analyzer after a further 0.8 min of gas travel and a 0.3 min analyzer lag. Raising the burning zone setpoint by 15 °C gives the temperature controller a proportional step, so the coal command moves within a single 3 s scan while the combustion air is held by an oxygen controller that has not yet observed the change. The excess air ratio falls to 0.964, the carbon monoxide peaks at 611 ppm and spends 0.78 min above the 500 ppm alarm, and 899 ppm min of carbon monoxide leaves the back end. Clinker produced during the excursion is burned under a reducing atmosphere, and the quality loss is not recovered by subsequent operation.

No retuning removes the excursion, because the limit is the transport of gas from the burner to the analyzer rather than the controller settings. Quadrupling the oxygen loop gain and halving its integral time lowers the peak only to 528 ppm, while the oxygen loop's integrated error rises from 11.9 to 17.0 % min over the first hour. The disturbance, however, is the control system's own output, known exactly in the scan it is computed. Ratioing the air to the fuel command and assigning the oxygen controller to trim the ratio leaves the carbon monoxide at its 20 ppm floor and the minimum excess air at 1.085 through the identical move.

At the slow end the proportions reverse. Free lime, the residual uncombined calcium oxide on which clinker quality is judged, is a property of the clinker leaving the burning zone, sampled on a 10 min cycle and reported 15 min after sampling. The effective dead time of a loop closed on the assay is between 15.9 min and 25.9 min depending on where in the sampling cycle the disturbance lands, against a burning zone thermal time constant of 13.0 min. The 22.2 min solids transit is not part of that loop, because it delays only changes that enter at the feed end. A loop closed on the assay acts as a slow trim rather than a regulator, and the application quantifies the cost of tuning it aggressively.

Between the two extremes sits a combustion constraint. The carbon monoxide response to excess air is strongly asymmetric, with 20 ppm at $\lambda = 1.02$, 147 ppm at 1.00, 853 ppm at 0.95 and 1443 ppm at 0.92, while the oxygen falls smoothly and almost linearly across the same range. A penalty of that shape moves the operating point away from the thermodynamic optimum, and it is the reason a kiln runs with more excess air than combustion requires. The asymmetry also makes the oxygen measurement conditional on where the air enters. False air leaking through the kiln inlet seal carries oxygen that bypasses the flame, and at 4.5 kg/s of leakage the analyzer reads 4.06 % instead of 1.91 % while the true excess air is unchanged at 1.100.

The model rests on standard process fundamentals. Each zone carries an energy balance on the charge with counter-current gas exchange, the chemistry enters as two lumped conversions with Arrhenius rates, and the separation between a gas residence of seconds and a solids transit of tens of minutes justifies a quasi-steady gas phase [3, 4]. The model and its scenarios teach ratio and feedforward control, dead-time dominant loops, sampled and delayed measurements, anti-windup with bumpless transfer, and the identification of unmeasured disturbances [4, 5]. Industry runs fired equipment on the ratio-with-trim structure studied here, and kiln quality control operates against assay delays of the magnitude modeled [6, 7].

This chapter works with a rotary kiln simulator described in enough detail to be rebuilt. It compares ratio feedforward with pure oxygen feedback on the same firing move and the same false air leak, and it measures the effect of roughly twenty minutes of assay dead time on a quality loop closed around a burning zone whose own thermal time constant is thirteen minutes. Two measured outcomes carry the argument. The ratio structure removes the entire carbon monoxide excursion from a 15 °C firing move rather than reducing it, and a graded five hour run goes from 88.7 min of reducing flame to none when the structure is assembled correctly. The heat balance, the combustion arithmetic and the ratio structure are standard. What the application adds is an eighteen-state model, twelve graded scenarios and a dead time budget assembled from four separate contributions before any tuning is attempted. A model predictive controller of the coal and the air is offered beside the loops and measured against them on the same cases.

Section 11.2 reviews the background of cement kiln control. Section 11.3 sets out the five-zone model, its combustion and analyzer laws and its kinetics. Section 11.4 describes the burning zone loop, the ratio with oxygen trim, the free lime loop with its dead time budget and the predictive controller. Section 11.5 measures the combustion characteristic, the firing move under both air structures, false air, kiln speed and burnability, the graded challenge and the predictive controller against the loops. Section 11.6 sets the loops in their real hardware, and Sections 11.7 to 11.9 give the limitations, the discussion and the conclusion.

11.2 Background

Portland cement is made by heating a finely ground mixture of limestone and clay to roughly 1450 °C, at which temperature the lime combines with silica, alumina and iron oxide into the calcium silicate phases that give cement its strength. The chemistry proceeds as a sequence of steps. Water leaves, then the carbonate decomposes, then belite forms, then belite and the remaining free lime combine into alite in the presence of a melt whose quantity is fixed by the alumina and iron content of the raw mix [8]. Uncombined lime left in the clinker makes the cement unsound, so the residual free

lime is the number the product is graded on, and it is measured by chemical assay rather than by an instrument.

> ⧖ HISTORY **The Rotating Kiln Takes Over (About 1895 to 1910)**
>
> By 1910 the fixed kilns that had burned Portland cement were giving way to the rotary process. A chamber kiln needed about a week for one run and gave about 30 tons of clinker. A rotating steel cylinder 60 to 150 ft long, lined with refractory brick and carried on rollers at a slight angle, turned out 200 to 400 tons a week. Producer gas had fired the early cylinders with little success, petroleum served in the United States from about 1895, and by 1910 finely powdered coal blown in by air was almost universal. The hot clinker was cooled by a current of cold air, which went on, warmed, to the burner. Underburnt clinker held free lime that could slake slowly in the finished cement and expand enough to destroy the work [9].

The kiln is where the energy and most of the carbon dioxide are spent. Cement manufacture accounts for a large and growing share of global industrial carbon dioxide emissions, with the process emissions from calcination and the combustion emissions from firing contributing separately, and the historical accounting shows both rising with production through the twentieth century [2]. Specific thermal energy consumption is therefore the number the plant is judged on over a year, and the technical routes to improving it are well cataloged and largely exhausted [10]. The remaining margin is operational, and recovering it is a control problem.

Kiln control has an unusual place in the history of the field. The first industrial application of fuzzy logic control was a cement kiln, chosen precisely because the plant was too nonlinear, too slow and too poorly instrumented for the classical methods available at the time [11]. The choice reflected the process. A kiln has strong interaction between fuel, air, feed and speed, a quality measurement that arrives tens of minutes late, and an operator whose expertise consists largely of pattern recognition on a burning zone image. Rule-based and expert-system controllers dominated for two decades for that reason.

Model predictive control eventually displaced them, and reported industrial implementations handle the kiln together with the calciner and the cooler, using the free lime assay as an infrequent correction to an inferred quality rather than as a controlled variable [7]. On the modeling side, one-dimensional axial models of the bed and gas with the reaction sequence resolved have been developed and validated against plant data [12], and extended to the integrated preheater, calciner, kiln and cooler so that the recycles between them can be represented [13]. Those models are the reference against which a reduced model such as the one of this chapter should be read.

However, textbook presentations of feedforward begin from a measured external disturbance and a disturbance transfer function to invert [4]. That presentation does not cover the case where the disturbance is the control system's own output, where the feedback measurement is downstream of the actuator by a fixed transport time, and where the penalty for a late correction is asymmetric and permanent. The application is built to demonstrate that a well-tuned feedback loop is not the remedy for this class of fast disturbance. A refinery fired heater meets the same problem with a slow forced-draft fan and answers it with cross-limiting and oxygen trim in Volume 3, Chapter 4.

11.3 Process Model

The kiln is a 60 m tube of 4.2 m diameter inclined at 3.5 % and rotating at 3.5 rpm. Raw meal enters the high end at 860 °C and 90 % calcined, having already passed through a preheater tower and a precalciner that are boundary conditions here. Coal is fired at the low end into secondary air returning from the clinker cooler at 950 °C. The combustion gas travels counter-current to the meal and leaves the feed end at 1049 °C. The arrangement is drawn in Figure 11.1.

11.3.1 Units and Design Point

The clock runs in minutes. Solids and air are in t/h, energy in kJ, temperature in °C and oxygen in percent by volume dry. At the design condition the kiln takes 258 t/h of raw meal at 860 °C and 90 %

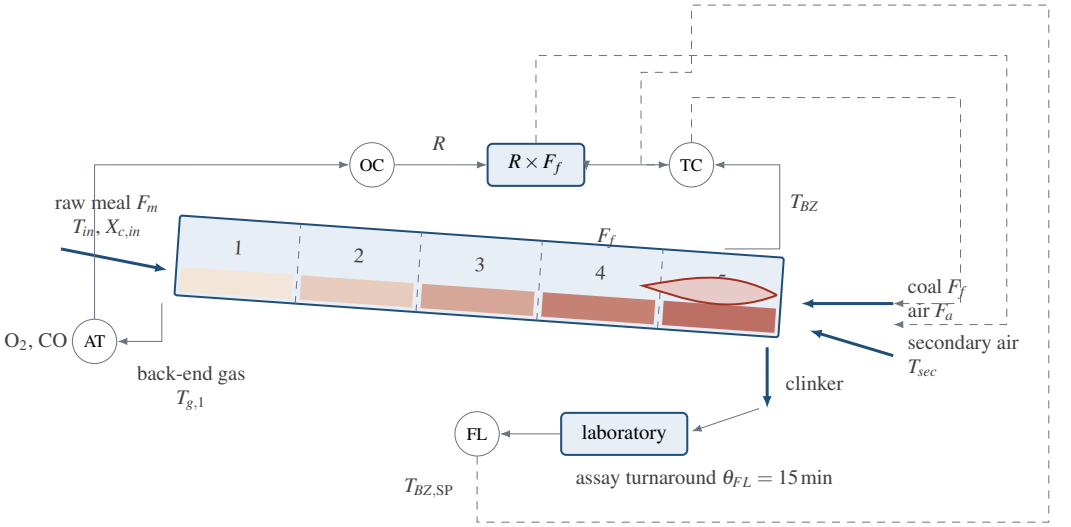


Figure 11.1 Rotary kiln as implemented, drawn on its incline with the five zones numbered from the feed end. Charge shading follows temperature. The temperature controller TC writes the coal rate F_f, which is multiplied by the ratio R to give the combustion air F_a in the same scan, while the oxygen controller OC moves R slowly from the back-end analyzer AT. The free lime assay FL arrives θ_{FL} after clinker leaves zone 5 and writes the burning zone setpoint $T_{BZ,\mathrm{SP}}$.

calcined, fires 11.033 t/h of coal with 127.434 t/h of air for an excess air ratio of 1.100, and turns at 3.5 rpm for a transit of 22.16 min. The charge profile is 826.3 °C, 838.3 °C, 891.5 °C, 1007.8 °C and 1429.8 °C and the gas profile is 1049 °C, 1135 °C, 1260 °C, 1440 °C and 1711 °C. Calcination is complete by zone 3 at $X_c = 0.995$ and clinkerization reaches $X_k = 0.982$ only in zone 5, giving 1.200 % free lime. The back end reads 1.910 % oxygen and 20 ppm carbon monoxide.

11.3.2 Transport

Residence time follows the standard rotary kiln correlation, and it is the only quantity in the model that kiln speed touches directly:

$$\tau = \frac{k_r L}{N D S}, \qquad \tau_z = \frac{\tau}{5}, \qquad M_z = \frac{F_m \tau}{300}, \tag{11.1}$$

where M_z is the charge hold-up per zone in tonnes. At design, $\tau_z = 4.43$ min and $M_z = 19.06$ t.

11.3.3 Charge Energy Balance

Each zone is a well-mixed lump whose effective thermal mass includes the refractory lining:

$$\left(M_z 10^3 C_{p,s}(1+f_{ref})\right)\frac{\mathrm{d}T_{s,i}}{\mathrm{d}t} = 60\Big[\dot{W}_s\left(T_{s,i-1}-T_{s,i}\right) + UA_i\left(T_{g,i}-T_{s,i}\right) + Q_{rad,i} - \Delta H_{cal}\, r_{cal,i} - \Delta H_{clk}\, r_{clk,i}\Big], \tag{11.2}$$

with $\dot{W}_s = F_m 10^3 C_{p,s}/3600$ the heat capacity flow of the solids in kW/K, and $r_{cal,i}$ and $r_{clk,i}$ the mass rates of calcination and clinkerization in zone i in kg/s, obtained from the kinetics of (11.10) and (11.11) applied to the zone hold-up. The factor $f_{ref} = 3.2$ is the ratio of refractory thermal mass to charge thermal mass, and it sets the burning zone time constant. Evaluating $M_z 10^3 C_{p,s}(1+f_{ref})/[60(\dot{W}_s + UA_5)]$ at the design point gives 13.0 min. The charge alone would respond in 3.1 min.

Table 11.1 Controlled, Manipulated and Disturbance Variables

Symbol	Meaning	Unit	Range
Controlled variables			
T_{BZ}	burning zone charge temperature	°C	1330–1520
O_2	back-end oxygen	%	0–8
FL	free lime in the clinker	%	0–5
Manipulated variables			
F_f	coal feed	t/h	4–18
F_a	combustion air	t/h	55–230
R	air to fuel ratio	t/t	8–16
$T_{BZ,\mathrm{SP}}$	burning zone setpoint (outer)	°C	1330–1520
N	kiln speed	rpm	1.5–5.5
Disturbance variables			
F_m	raw meal feed	t/h	180–340
β	raw mix burnability	—	0.55–1.40
c_f	coal calorific value	—	0.80–1.15
$X_{c,in}$	calcination degree from the precalciner	—	0.70–0.97
$\dot{m}_{fa}$	false air at the inlet seal	kg/s	0–8
T_{sec}	secondary air temperature	°C	700–1100
States			
$T_{s,i}$	charge temperature, zones 1–5	°C	—
$X_{c,i}$	calcination conversion, zones 1–5	—	—
$X_{k,i}$	clinkerization conversion, zones 1–5	—	—
F_f, F_a, N	actuator positions	—	—

The conductance profile rises steeply toward the discharge end, because radiation rather than convection is doing the work there:

$$UA_i = UA_0\, w_i, \qquad w = \{0.55,\ 0.70,\ 0.90,\ 1.25,\ 1.90\}. \tag{11.3}$$

11.3.4 Quasi-Steady Gas Profile

Gas residence in the kiln is seconds against a charge transit of 22 min, so the gas is treated as quasi-steady. Marching counter-current from the burning zone toward the feed end,

$$T_{g,i} = \frac{\dot{W}_g T_{g,i+1} + UA_i T_{s,i} + UA_{sh} T_a}{\dot{W}_g + UA_i + UA_{sh}}, \tag{11.4}$$

where $\dot{W}_g = (F_a + F_f)10^3 C_{p,g}/3600$ is the heat capacity flow of the combustion gas in kW/K and $T_a = 25\,^{\circ}\mathrm{C}$ is the ambient temperature behind the shell loss. The march starts from the flame temperature

$$T_{flame} = T_{sec} + \frac{f_g\, \eta(\lambda)\, Q_f}{\dot{W}_g}, \qquad Q_f = \frac{c_f\, \Delta H_f F_f\, 10^3}{3600}, \tag{11.5}$$

where Q_f is the firing heat release in kW and c_f is the calorific value multiplier of Table 11.1. The fraction $f_g = 0.6$ of the combustion heat is carried by the gas and the remainder is radiated directly to the charge in zone 5,

$$Q_{rad,5} = (1 - f_g)\, \eta(\lambda)\, Q_f. \tag{11.6}$$

Because the combustion air scales with the fuel, (11.5) makes the flame temperature nearly independent of firing rate. The coal rate changes the quantity of heat delivered rather than the flame temperature.

11.3.5 Combustion

Excess air, combustion efficiency and carbon monoxide are

$$\lambda = \frac{F_a}{\mathrm{AFR}\, F_f}, \qquad \eta(\lambda) = \min(1,\ \lambda) \tag{11.7}$$

$$\mathrm{CO} = 20 + k_{co}\,(\lambda_{ok} - \lambda)_+^{3/2} \tag{11.8}$$

with $\lambda_{ok} = 1.02$ and $k_{co} = 45\,000$ ppm. Below stoichiometric the combustion efficiency η equals λ, and the unburned fraction leaves as carbon monoxide. The three-halves power in (11.8) concentrates the carbon monoxide response in a narrow band below λ_{ok}, and (11.7) through (11.6) together make a reducing flame a cooler flame, because unburned fuel carries away part of the heat release.

The analyzer does not measure λ directly. False air entering at the inlet seal carries oxygen that bypasses the flame, and the instrument cannot distinguish where the oxygen entered:

$$\hat{\lambda} = \frac{F_a + 3.6\,\dot{m}_{fa}}{\mathrm{AFR}\, F_f}, \qquad \mathrm{O_2} = 21\left(1 - \frac{1}{\hat{\lambda}}\right)\%. \tag{11.9}$$

The leak joins the gas after the last zone, so $\dot{m}_{fa}$ enters (11.9) but neither (11.4) nor (11.5).

11.3.6 Reaction Kinetics

Calcination and clinkerization are first order in the unconverted fraction with Arrhenius rates:

$$\frac{\mathrm{d}X_{c,i}}{\mathrm{d}t} = \frac{X_{c,i-1} - X_{c,i}}{\tau_z} + k_{cal}(T_{s,i})(1 - X_{c,i}) \tag{11.10}$$

$$\frac{\mathrm{d}X_{k,i}}{\mathrm{d}t} = \frac{X_{k,i-1} - X_{k,i}}{\tau_z} + k_{clk}(T_{s,i})\,\beta\,(1 - X_{k,i}) \tag{11.11}$$

$$\mathrm{FL} = \mathrm{CaO}_{free}\left(1 - X_{k,5}\right). \tag{11.12}$$

The clinkerization activation temperature is 45 000 K, which gives (11.11) a threshold character. At the design zone residence the free lime of (11.12) is 9.3 % with zone 5 at 1300 °C, 4.2 % at 1350 °C and 1.4 % at 1420 °C, and it falls below the 2.0 % soundness limit only above about 1397 °C. The zone 5 conversion follows its temperature with a time constant $1/(1/\tau_z + k_{clk}\beta)$ of 0.083 min at design, so the free lime moves with the burning zone temperature within seconds. The burnability multiplier β stands for the raw mix chemistry, fineness and mineralogy together, and no continuous instrument on the kiln measures it.

11.3.7 Initialization

At reset the application solves (11.2) through (11.11) for their simultaneous steady state by damped substitution, sweeping the five zones and re-solving the gas profile on every pass until no charge temperature moves by more than 1×10^{-11} °C. The result is the profile quoted in Section 11.3.1 and the relative residual on every charge temperature and conversion derivative at $t = 0$ is below 1×10^{-13}. Actuator states start at the design point values of Section 11.3.1, the coal transport delay starts filled with the coal value, so a command at $t = 0$ reaches the burner 0.4 min later, and a setting chosen before a run equals the same setting carried through a reset. The automatic scenarios hold 1430 °C, 0.20 °C above the initialized burning zone, so the first scan adds 0.02 t/h of coal.

Table 11.2 Model Parameters, All Adjustable in the Application

Symbol	Meaning	Value	Unit
L	kiln length	60	m
D	kiln diameter	4.2	m
S	kiln slope	0.035	m/m
k_r	residence time coefficient	0.19	—
$C_{p,s}$	solids heat capacity	1.1	kJ/(kg K)
$C_{p,g}$	gas heat capacity	1.15	kJ/(kg K)
f_{ref}	refractory thermal mass	3.2	× charge
UA_0	conductance scale	18	kW/K
UA_{sh}	shell loss per zone	1.6	kW/K
T_a	ambient temperature	25	°C
f_g	flame heat carried by the gas	0.6	—
ΔH_f	coal calorific value	25000	kJ/kg
AFR	stoichiometric air-fuel ratio	10.5	kg/kg
λ_{ok}	excess air clearing CO	1.02	—
k_{co}	carbon monoxide scale	45000	ppm
CO_{lim}	carbon monoxide alarm	500	ppm
y_{CaCO_3}	carbonate fraction of the meal	0.78	—
ΔH_{cal}	calcination enthalpy	1780	kJ/kg
ΔH_{clk}	clinkerization enthalpy	100	kJ/kg
$k_{cal,0}$	calcination rate at 900 °C	0.9	1/min
E_{cal}/R	calcination activation temperature	25000	K
$k_{clk,0}$	clinkerization rate at 1450 °C	16.199	1/min
E_{clk}/R	clinkerization activation temperature	45000	K
CaO_{free}	lime available to combine	65	%
FL_{max}	free lime that makes cement unsound	2.0	%
T_{ring}	burning zone temperature starting a ring	1520	°C
τ_f	coal feeder lag	0.50	min
θ_f	coal transport dead time	0.40	min
τ_a	air damper lag	0.25	min
τ_N	kiln drive ramp	0.80	min
θ_{O_2}	oxygen analyzer transport delay	0.80	min
τ_{O_2}	oxygen analyzer response	0.30	min
τ_{BZ}	pyrometer filter	0.60	min
θ_{FL}	free lime analyzer turnaround	15	min
Δt_{FL}	free lime sample interval	10	min

11.3.8 Assumptions

1. **Five well-mixed zones.** The transport delay is produced by the staged geometry rather than by an explicit dead-time element, at the cost of a coarse temperature profile. Axial models resolve the bed and gas continuously [12], and a real kiln also has a bed depth and a refractory coating whose thickness changes over a campaign.

2. **The gas is quasi-steady.** No flame model, no gas composition, no radiation view factors, and no separation of convective from radiative transfer. The split between gas-borne and directly radiated heat is the single fitted fraction f_g and UA_i is a fitted profile.

3. **Two lumped reactions.** Real clinker chemistry is the belite and alite sequence in a melt whose quantity depends on the alumina and iron moduli [8]. Burnability is one multiplier standing in for all of it.

4. **The rest of the line is a boundary condition.** Preheater tower, precalciner and clinker cooler enter only as the inlet temperature, the inlet calcination degree and the secondary air temperature. The alkali and sulfur recycles that cause build-ups and blockages cannot appear, and integrated models exist precisely to capture them [13].

5. **No reducing-atmosphere chemistry.** Carbon monoxide is algebraic and the damage a reducing flame does to clinker quality is reported in the notes rather than modeled. Nitrogen oxides, which push the operating point the other way, are absent, so the application understates how narrow the acceptable combustion window is.

6. **Ring formation is an alarm.** Nothing builds up, restricts the kiln or changes the residence time. On a real kiln that feedback is among the most disruptive events in operation.

7. **Fixed-step integration.** Classical fourth-order Runge–Kutta at 0.01 min with a 3 s controller scan, so repeated runs reproduce exactly (Section 1.6.1). A physical kiln shows run-to-run variability.

11.4 Controller Design

Four feedback controllers are implemented, and Section 11.4.5 describes a model predictive controller that can take the coal and the air from them. The burning zone temperature drives the coal feeder, the back-end oxygen drives either the air damper or the air-fuel ratio, the free lime writes the burning zone setpoint, and the kiln speed is an operator setpoint with a first-order drive ramp. All use the same incremental, or velocity, form on a 3 s scan, computing an increment to the output each scan rather than the output itself.

11.4.1 Discrete PID Algorithm

Every loop runs the velocity-form PID of Section 1.6.2, Equations (1.10) to (1.13), with the error $e_k = a(\mathrm{SP}_k - \mathrm{PV}_k)$ and the action $a = \pm 1$. The derivative acts on the measurement passed through a 0.5 min first-order filter, and $\tau_D = 0$ in every scenario. Anti-windup is back-calculation with tracking time $T_t = \tau_I$, so the integral alone is corrected and the running sum of proportional and derivative increments is never truncated [14]. Each 3 s scan removes $\Delta t/\tau_I$ of the clipped excess, $1/180$ for the 9 min temperature loop. No run reported here reaches a coal, air, ratio or setpoint limit, so the anti-windup term is inactive in every result.

In manual mode, for the air loop that is not selected and for an advisory free lime loop, the algorithm runs the tracking step of Section 1.6.2, which sets the integral so that the output equals the value actually applied. In automatic mode the manual coal, air and burning zone stations are set to the controller outputs after every controller execution, so the transfer is bumpless in both directions. The temperature, oxygen and ratio trim loops are reverse acting with $a = +1$, because more coal heats the burning zone and more air, or a higher ratio, raises the oxygen. The free lime loop is direct acting with $a = -1$, because a hotter burning zone gives less free lime. The application labels it the same way.

11.4.2 Burning Zone Temperature Loop

Perturbing (11.2) through (11.6) about the design point gives 52.1 °C of burning zone temperature per t/h of coal. The path from the coal command passes a 0.5 min feeder lag, a 0.4 min coal transport delay, the 13.0 min thermal time constant of the zone itself and a 0.6 min pyrometer filter, so the zone alone would present roughly 13 min of lag with about a minute of apparent dead time. The zone does not answer alone, because the hotter gas also warms the four zones upstream, which pass warmer charge into it. With the air ratioed the true burning zone covers 63 % of a coal step after 30 min and 90 % after 77 min, while a copy of the model with zones 1 to 4 held at their design state covers 63 % after 15.8 min at 32.3 instead of 52.0 °C per t/h. Scenarios 3 to 10 use $K_c = 0.11\,\mathrm{t/h/^\circ C}$ and $\tau_I = 9\,\mathrm{min}$, which is the SIMC result [15] for the zone alone with a closed-loop time constant of about 1.3 min. With that tuning the burning zone covers 63 % of a 25 °C setpoint step in 3.7 min and overshoots by 3.9 °C.

11.4.3 Ratio Control with Oxygen Trim

The combustion air is governed by one of two selectable structures:

$$\text{feedback:}\quad F_a = \mathrm{PI}_{O_2}\left(\mathrm{O}_2^{\mathrm{SP}} - \hat{\mathrm{O}}_2\right) \tag{11.13}$$

$$\text{ratio:}\quad R = \mathrm{PI}_R\left(\mathrm{O}_2^{\mathrm{SP}} - \hat{\mathrm{O}}_2\right), \quad F_a = R\,F_{f,k} \tag{11.14}$$

In (11.14) the air is computed from the fuel command of the *same* scan, so it moves in the same 3 s the coal does, and the oxygen controller moves only the one parameter of the feedforward model. The loop that is not selected runs the tracking step continuously, so switching between them is bumpless. The product $R\,F_{f,k}$ is clipped to the air range of Table 11.1 after the trim, whose own limits apply to R, so a clipped air command would not reach the trim's integral. Scenarios 3 to 10 use $K_c = 1.5\,\mathrm{t/h/\%}$ with $\tau_I = 1.2\,\mathrm{min}$ for feedback air and $K_c = 0.15$ per percent with $\tau_I = 6\,\mathrm{min}$ for the ratio trim. The trim is deliberately five times slower than the loop it replaces, because it corrects a slowly drifting model parameter rather than a disturbance.

The structure of (11.14) is the standard industrial arrangement for fired equipment [6], and the reason it is standard is the timing argument rather than any property of the tuning. Ratio stations on other plants of this volume are compared in Section 1.4.6.

11.4.4 Free Lime Loop and Dead-Time Budget

The outer loop writes the burning zone setpoint from the laboratory result. The free lime of (11.12) belongs to the clinker leaving zone 5 and follows the burning zone temperature within seconds, so the 22.2 min transit of (11.1) lies outside this loop and delays only disturbances that enter at the feed end. The loop's effective dead time is the sum of four contributions, the sampled assay age of Section 1.4.5 plus the coal path:

$$\theta_{eff} = \underbrace{\theta_{FL}}_{15} + \underbrace{[0, \Delta t_{FL}]}_{0-10} + \underbrace{\theta_f}_{0.4} + \underbrace{\tau_f}_{0.5} = 15.9 - 25.9\,\mathrm{min}. \tag{11.15}$$

Against the 13.0 min zone time constant the ratio of dead time to lag is between 1.2 and 2.0, which places the loop in the dead-time dominant regime [4]. The free lime controller drives the closed temperature loop rather than the open zone, and that loop reaches 63 % of a setpoint step in 3.7 min, so the ratio for the loop as built is between four and seven. The delay panel of the application adds half a sample interval to the turnaround and reports 20 min at design, the age of the reported value against the clinker leaving the kiln, and shows the transit separately. Dead-time compensation such as the Smith predictor removes the delay from the characteristic equation only to the accuracy of the delay model [16], and the sampling contribution in (11.15) varies by 10 min from one disturbance to the next. Volume 5, Chapter 3 tunes a composition loop through a slow analyzer with the same budget. Scenarios 7 and 10 close the loop with $K_c = 15\,^\circ\mathrm{C/\%}$ and $\tau_I = 13\,\mathrm{min}$. SIMC on the 13.0 min zone,

with the upper end of (11.15) as the dead time and τ_c equal to it, gives the same integral time and a gain of 14 °C/% to 17 °C/% for the steady-state process gain of 0.015 %/°C to 0.018 %/°C between 1430 °C and 1455 °C and at the design point.

11.4.5 Model Predictive Control of Coal and Air

A fifth structure hands the coal and the combustion air to one model predictive controller, the receding-horizon calculation behind most industrial advanced control [17] and the option that Section 1.6.3 describes, and only scenario 11 loads with it selected. Every $\Delta t_m = 0.5$ min, or ten scans, a quadratic program plans both commands over $N_p = 120$ steps (60 min) in twelve move blocks with 24 decisions, and the kiln receives the first move alone. The plan holds the operator's burning zone setpoint, or with the free lime loop closed the temperature its own free lime model needs, and the air structure only names the loops that take over a rejected plan. The raw meal feed stays a disturbance and the kiln speed an operator setpoint, because the declared steady state reaches 1.20 % free lime with 10.018 t/h of coal at 1.5 rpm against 11.033 t/h at 3.5 rpm, and a plan free to slow the kiln would take that saving at a cost in bed depth and coating that the model does not contain.

The controller's model is not the five-zone kiln. Its burning zone is two first-order modes ($i = 1, 2$) driven by the coal at the burner and the applied air, as shown in Equation (11.16),

$$\begin{aligned} \tau_i \frac{\mathrm{d}x_i}{\mathrm{d}t} &= g_{F,i}\left(\hat{F}_f - F_f^0 + d\right) + g_{A,i}\left(\hat{F}_a - F_a^0\right) - x_i \\ \hat{T}_{BZ} &= T_{BZ}^0 + x_1 + x_2 \end{aligned} \tag{11.16}$$

where $\hat{F}_f$ is the coal command passed through the declared delay θ_f and lag τ_f, $\hat{F}_a$ is the air command passed through the damper lag τ_a, F_f^0 and F_a^0 are the design stations of Section 11.3.1, T_{BZ}^0 is the offset fixed when the estimate starts, and d is an input disturbance in t/h of coal. Manual step tests at the design point with the noise off, ± 0.25 t/h of coal and ± 4 t/h of air, give $\tau_1 = 17.0$ min and $\tau_2 = 42.3$ min with coal gains $g_{F,i}$ of 37.86 and 26.15 and air gains $g_{A,i}$ of −1.364 and 0.336 °C per t/h, leaving 0.27 % RMS of each response. An independent fit gives 16.3 and 41.5 min with steady gains of 64.02 and −1.030, against the plant's 64.02 and −1.037. The two modes are a reduction of a kiln whose linearization has a slowest real mode of 33.4 min and three complex pairs.

The oxygen and the free lime come from the declared laws. The analyzer's reading is Equation (11.9) without false air, linearized at every plan about the coal and air the model holds, and the declared delay θ_{O_2} and lag τ_{O_2} enter where the reading is compared with the model. A fourth output is the flame's own oxygen, $\mathrm{O}_{2,f} = 21(1 - \mathrm{AFR}\,\hat{F}_f/\hat{F}_a)$, computed from the plan's commands and the declared AFR. The free lime is Equation (11.11) for zone 5 alone at design burnability, taken in logarithms so that a change of burnability becomes an additive bias b_L, as shown in Equation (11.17),

$$\ln \widehat{\mathrm{FL}} = \ln \frac{\mathrm{CaO}_{free}}{1 + k_{clk}(\hat{T}_{BZ})\,\tau_z(N)} + b_L \tag{11.17}$$

where $\tau_z(N)$ is the zone residence of Equation (11.1) at the speed command. Against the kiln's steady state between 1367 and 1494 °C the law is within 1.72 % of the free lime. Every constant is fixed at load time, so a learner's change to the kiln's parameters acts on the plan as plant-model mismatch.

Three estimators keep the plan offset-free, each giving the model a disturbance of the kind the kiln receives [18], as in the offset-free controller of Volume 9, Chapter 4. A steady-state Kalman filter on the modes, the pyrometer filter and d, with variances of 2×10^{-5} (t/h)^2 per scan for d and 1 °C^2 for the pyrometer, reads any feed, calcination, coal or cooler change as a heat input and estimates 63 % of a 0.1 t/h step in d within 4.05 min. The oxygen bias b_{O_2} is the analyzer's reading less the model's prediction of it through that delay and lag, filtered over 2 min. At each assay b_L moves by $w_L = 0.6$ of the residual, taken against the model's free lime when the sample was cut, θ_{FL} before the report, which is the alignment of a slow, delayed quality measurement that multi-rate inferential control requires [19]. It is the bias update of Section 1.4.5, compared at the sampling instant, which Chapter 7

uses on a tailings assay. Compared at the report time instead, the free lime IAE of scenario 11 rises from 20.93 to 29.46 % min.

Each plan minimizes Equation (11.18),

$$J = \sum_{k=1}^{N_p} \left[\left(\hat{T}_{BZ,k} - T^* \right)^2 + q_O \left(\hat{\mathrm{O}}_{2,k} - \mathrm{O}_2^* \right)^2 \right] + \mu \sum_{j=1}^{12} \left(s_F \, \Delta F_{f,j}^2 + s_A \, \Delta F_{a,j}^2 \right) + \sum_{\ell=1}^{4} \left(\rho_\ell \, \varepsilon_\ell^2 + \sigma_\ell \, \varepsilon_\ell \right) \tag{11.18}$$

where $\hat{T}_{BZ,k}$ and $\hat{\mathrm{O}}_{2,k}$ are the predictions at step k, Δ is the change of a command from one block to the next, the first measured from the commands applied now, $q_O = 400$ per %^{2} against 1 per °C^2, $s_F = 4000$ and $s_A = 2$ per (t/h)^2, $\mu = 1$ is the move price, and $\varepsilon_\ell \geq 0$ are the slacks of four soft limits with prices ρ_ℓ and σ_ℓ. The target T^* is the operator's setpoint with the free lime loop advisory and, with it closed, the solution of Equation (11.17) for the free lime target, clipped to 1330 °C and to 5 °C under the ring limit. The oxygen target is $\mathrm{O}_2^* = \max\left[\mathrm{O}_2^{SP},\ \mathrm{O}_{2,min} + \max(0, b_{O_2})\right]$ with the floor $\mathrm{O}_{2,min} = 1\,\%$, which is $\lambda = 1.05$. The plan prices no coal, because two handles and two tracked outputs leave no steady degree of freedom, and a price of 640 per t/h per step saves 0.56 t of coal over scenario 11 while the free lime IAE rises to 40.86 % min.

The coal stays within 4 and 18 t/h and the air within 55 and 230 t/h on every block, the clamps of the loops, as hard bounds. Four limits are soft, and the shared solver allots each a single slack for the entire horizon rather than one per step: the analyzer's and the flame's oxygen above $\mathrm{O}_{2,min}$ at 1×10^6 per %^{2} and 1×10^5 per %, the free lime under the 2.0 % soundness limit with the loop closed at 1×10^5 and 1×10^4 in logarithmic units, and the burning zone between 1330 °C and the lower of 1520 °C and the ring limit at 1×10^4 and 1×10^3.

The raised target is what holds the flame under false air, as Section 11.5.6 measures. When the analyzer reads more than the plan's coal and air explain, b_{O_2} grows and O_2^* rises with it, so the analyzer is aimed at the floor plus the leak's share of the reading. At 150 min of scenario 11 the analyzer reads 3.47 % while the flame has 0.99 % ($\lambda = 1.050$), with $b_{O_2} = 2.31\,\%$ and a target of 3.29 %. The flame row cannot see a wrong stoichiometric ratio, because its excess air is computed from the commands, and it equals the true λ at every scan only because the simulated actuators do exactly what they are told. With the plant's AFR at 11.0 it reads 0.053 high and the flame reaches $\lambda = 0.995$ with 193 ppm of carbon monoxide.

A plan is kept only when the solver's own status is optimal, checked even when the solve reports success, every returned array has its shape and finite entries, every planned command lies inside its bounds and the first equals the applied move. Otherwise the loops take over with the learner's air structure and free lime loop until a later plan passes, starting from the applied coal, air and burning zone setpoint that they follow at every scan without preload. With the noise off, rejections forced at 60, 140 and 185.5 min of scenario 11 move the coal 0.002 to 0.009 t/h on the hand-over scan, and 0.134 t/h at 80 min, where the free lime loop answers an assay arriving on that scan. No run reported here rejected a plan.

The kiln model is integrated with the run's parameters, each instrument applies its own delay, lag and noise, and the plan's equations hold the declared values, which keeps plant, instruments and controller model separate. A replay test feeds the measurements of 150 min of scenario 11, logged with a different hidden kiln, to a fresh controller state and obtains all 3000 commands bit for bit, and at 23 captured decisions the Karush-Kuhn-Tucker conditions hold. A plan takes 0.85 to 0.96 ms on average in Chrome in the runs behind this chapter, 10.2 to 11.5 % of one core at the top speed of 60 min of kiln time per second.

Table 11.3 Combustion Air Swept at 11.033 t/h of Coal, Steady State

λ	F_a t/h	O_2 %	CO ppm	T_{BZ} °C	FL %
0.92	106.6	0.00	1443	1392.3	2.155
0.95	110.1	0.00	853	1411.1	1.606
1.00	115.8	0.00	147	1442.0	0.997
1.02	118.2	0.41	20	1439.6	1.035
1.05	121.6	1.00	20	1435.9	1.095
1.10	127.4	1.91	20	1429.8	1.200
1.20	139.0	3.50	20	1418.0	1.437
1.35	156.4	5.44	20	1401.4	1.861

11.5 Results

The application offers twelve scenarios in its selector, each isolating one mechanism: the kiln as a conveyor, the flame and the carbon monoxide cliff, the burning zone loop, oxygen feedback alone and the ratio against the same firing move, false air, the free lime loop, kiln speed, coal calorific value, raw meal chemistry, the predictive controller against the loops on one shift, and the graded challenge. Section 1.6.4 describes how a case is briefed, credited and exported. A case is credited when the learner performs its 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 all twelve cases accumulates toward certificates of progress and completion. Reset restarts a run from $t = 0$ and keeps every setting the learner changed, which is how scenarios 6 and 7 compare two configurations on the same disturbance, while reloading a scenario restores its own settings.

11.5.1 Steady-State Combustion Characteristic

Table 11.3 sweeps the combustion air at fixed coal. Oxygen falls smoothly and carbon monoxide does not. Between $\lambda = 1.02$ and $\lambda = 0.92$ the oxygen reading moves by 0.41 % and the carbon monoxide by 1423 ppm.

The burning zone column shows a second effect. The maximum lies near $\lambda = 1.00$ and the temperature falls on both sides. Above it, excess air is heated and carried up the kiln to the stack. Below it, part of the fuel leaves as carbon monoxide instead of releasing its heat, so a reducing flame is also a cooler flame. The kiln is nevertheless run at $\lambda = 1.10$, giving up 12.2 °C of burning zone temperature and 0.20 % of free lime, because of the carbon monoxide penalty on the fuel-rich side in the third column.

> **+ DID YOU KNOW? An Oxygen Reading That Stops at Zero**
>
> From the simulator: in Table 11.3 the oxygen reads 0.00 % at every excess air ratio from 1.00 down to 0.92, while the carbon monoxide climbs from 147 to 1443 ppm. Once the flame is short of air, the oxygen reading says so but cannot say by how much.

11.5.2 Ratio Feedforward Against a Firing Move

Scenarios 4 and 5 raise the burning zone setpoint by 15 °C at $t = 10$ min, which corresponds to the correction issued after a high free lime assay. Table 11.4 reports the identical move under the two air structures, feedback air in scenario 4 and the ratio in scenario 5. Peaks and times come from the 0.6 s integration step and the integrated errors from the application's scorecard.

The ratio structure removes the carbon monoxide excursion rather than reducing it, and the mechanism is timing rather than tuning. Under feedback the air cannot begin to move until the coal

Table 11.4 Burning Zone Setpoint Raised 15 °C at $t = 10$ min, Measured Over 60 min

Quantity	feedback air	ratio with trim
Peak carbon monoxide (ppm)	611	**20**
Minimum excess air λ	0.964	**1.085**
Time above 500 ppm (min)	0.78	**0.00**
Integrated CO above floor (ppm min)	899	**0**
Oxygen loop IAE (% min)	11.86	**3.25**
Burning zone IAE (°C min)	68.7	81.7

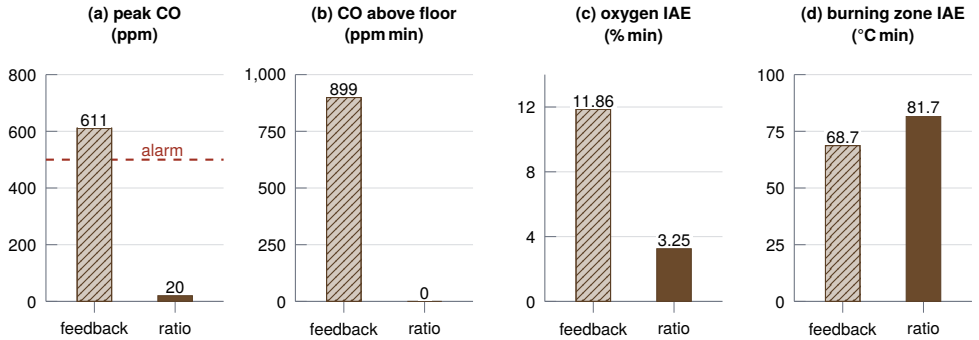


Figure 11.2 The firing move of Table 11.4 as a chart, feedback air hatched and the ratio with trim solid; the dashed line in (a) is the 500 ppm alarm. Notice that the ratio does not shrink the carbon monoxide bars, it removes them, and cuts the oxygen error too. The one bar that grows is the burning zone error in (d), the small price the text explains.

has finished its 0.4 min of transport and the gas has carried the change 0.8 min to the analyzer, so the air command first moves at 11.26 min, 1.26 min after the step, and the 0.3 min analyzer lag slows it further. The 0.5 min feeder lag keeps the extra coal arriving meanwhile, and the flame runs short of air, with λ below λ_{ok}, from 10.77 min to 13.66 min, 2.9 min in all. Under the ratio the air moves in the same 3 s scan.

The one column that degrades is the burning zone integrated error, by 19 %. Under feedback the lagging air leaves the extra coal burning with less excess air, a hotter flame that speeds the burning zone's rise. Under the ratio the air follows the coal command in the same scan, while the extra coal is still in transport and in the feeder lag, so the flame briefly runs with more excess air than at design, up to $\lambda = 1.235$, and the temperature loop carries slightly more of the correction. That cost is small against the 899 ppm min carbon monoxide excursion the structure eliminates. Figure 11.2 draws four rows of the table as bars.

Table 11.5 Kiln Speed Swept at Fixed Firing, Steady State

N (rpm)	τ (min)	T_{BZ} (°C)	FL (%)
2.5	31.0	1430.5	0.850
3.0	25.9	1430.1	1.025
3.5	22.2	1429.8	1.200
4.2	18.5	1429.5	1.445
5.0	15.5	1429.2	1.723

TRY IT IN THE APP 4 · Oxygen Feedback Alone, Against a Firing Move

Load the scenario and press Run. At minute 10 the burning zone setpoint rises 15 °C. Watch the Carbon monoxide and excess air trend: the coal moves at once, the air waits for the analyzer, and the carbon monoxide peaks near 611 ppm. Then load scenario 5, Ratio the air to the fuel, and watch the same trend stay at its 20 ppm floor.

Pause. Try Exercise 11.1 before reading on.

11.5.3 False Air and Analyzer Bias

Scenario 6 opens a false air leak of 4.5 kg/s at the kiln inlet seal at $t = 12$ min and repairs it at $t = 80$ min. At steady state the leak raises the analyzer from 1.910 % to 4.063 % oxygen with the true excess air unchanged at 1.100, because the oxygen entered downstream of the flame.

TRY IT IN THE APP 6 · False Air: The Analyzer Lies

Run the scenario with Air control on O_2 feedback. The leak opens at minute 12, the controller closes the air, and the carbon monoxide peaks at 1146 ppm. Press Reset, switch Air control to Ratio + trim and run past the repair at minute 80: the flame drifts only to an excess air of 0.971, with a peak of 504 ppm just after the repair.

Under feedback the controller acts on the biased reading, closes the damper, and drives the flame to a minimum excess air of 0.934 with a carbon monoxide peak of 1146 ppm at 35 min, above 500 ppm from 14 min after the leak for 39.9 min in all. Under the ratio the same bias enters through the trim, which moves five times more slowly, and the flame drifts to 0.971 with a peak of 504 ppm 0.8 min after the repair, when the analyzer registers it, crossing the alarm only for that last 0.8 min. The ratio structure does not protect the kiln against a biased analyzer. Left unrepaired, the leak takes both structures to the same state after 300 min, $\lambda = 0.963$ and about 635 ppm, because both obey the same analyzer. What the ratio changes is the speed of the failure. It converts an immediate loss of control into a slow drift of one model parameter, which is a more manageable failure, and it is the second reason the structure is standard.

Pause. Try Exercise 11.3 before reading on.

11.5.4 Kiln Speed, Burnability and the Free Lime Loop

Table 11.5 compares the two ways of moving free lime. Kiln speed changes the residence time of (11.1) and therefore the conversion of (11.11), and it does so without moving the burning zone at all.

Over the range, the free lime moves by 0.872 % while the burning zone moves by 1.3 °C. Speed is therefore an almost pure quality handle, and a fast one. In scenario 8 the free lime covers 63 % of its change 0.9 min after the speed step, against 3.6 min for the setpoint step of scenario 3. Kiln speed also sets throughput, so in practice it is committed to production rate rather than to quality.

Table 11.6 Challenge Scenario Over 300 min, Shipped Configuration Against the Suggested One

Quantity	shipped	suggested
Burning zone IAE (°C min)	6096.5	**674.5**
Oxygen IAE (% min)	338.5	**69.1**
Free lime IAE (% min)	242.7	**22.7**
Final free lime (%)	2.73	**1.20**
Final carbon monoxide (ppm)	3242	**20**
Final excess air λ	0.848	**1.035**
Time with a reducing flame (min)	88.7	**0.0**
Final specific heat (kJ/kg)	1600	1721

Burnability demonstrates why the assay loop is still required. Reducing β from 1.00 to 0.72 raises the free lime from 1.200 % to 1.646 % at unchanged firing, and the coal required to hold the burning zone at 1429.8 °C with the air fixed changes from 11.033 t/h to 11.024 t/h, a difference of 0.08 %. A burning zone controller is blind to it, and only the assay can see it.

Pause. Try Exercise 11.2 before reading on.

Scenario 7 closes the free lime loop with the budget of (11.15) in the path. The burnability step to 0.72 at 20 min reaches the controller with the result reported at 40 min, a dead time of 20 min in the middle of the budget. Over 6 h the loop accumulates 18.6 % min of free lime error with a burning zone error of 62.1 °C min, returns within 0.1 % of target 39 min after the step and stays within 0.02 % from 118 min after it. Raising K_c to 45 °C/% with $\tau_I = 12$ min produces an oscillation with a period of 57 min, close to three times that dead time, and in the sixth hour the free lime is still swinging between 0.92 % and 1.51 % while the burning zone setpoint has ranged over 1426 °C to 1485 °C. The loop therefore functions as a trim. The burning zone temperature holds the kiln, because it is fast and continuously measured, although it is not itself a product specification.

11.5.5 Challenge Scenario

The challenge, scenario 12, runs five hours with the raw meal stepped to 295 t/h at 20 min, burnability to 0.80 at 70 min, false air to 3 kg/s at 130 min and coal calorific value to 0.91 at 200 min. It opens with the air on feedback, the free lime loop open and the cautious tunings of scenario 1, which include $K_c = 0.025\,\mathrm{t/h/°C}$ and $\tau_I = 40$ min on the temperature loop. The pyrometer, the oxygen analyzer and the assay carry noise with standard deviations of 3 °C, 0.12 % and 0.05 %. Table 11.6 compares the shipped configuration with the suggested one, the ratio with trim, a closed free lime loop and the tunings of Section 11.4. The reducing flame time is the application's Reducing time tile, the minutes with carbon monoxide above the 500 ppm alarm.

The challenge passes when the burning zone stays within ±15 °C of setpoint, the carbon monoxide stays below 500 ppm throughout, the free lime returns below 2.0 % after each of the four upsets and the burning zone stays below the 1520 °C ring limit. The shipped configuration fails on three of the four limits. Its worst burning zone error is 44.1 °C, the carbon monoxide peaks at 3309 ppm and the free lime is still 2.7 % at 300 min after the calorific value drop at 200 min. The suggested configuration passes. Its worst error is 12.6 °C, the carbon monoxide peaks at 144 ppm, the free lime returns below 2.0 % after all four upsets with a peak of 1.48 %, and the burning zone peaks at 1449.6 °C, which leaves 2.4 °C and 356 ppm to spare.

The last two rows are related. The shipped configuration ends with the *better* specific heat consumption, 1600 kJ/kg against 1721 kJ/kg, and it achieves that by supplying insufficient combustion air for 88.7 min and producing clinker at 2.73 % free lime. Its cumulative coal is also lower, 68.7 kg per tonne of clinker against 71.4 kg. A scorecard reporting fuel per tonne alone would rank the

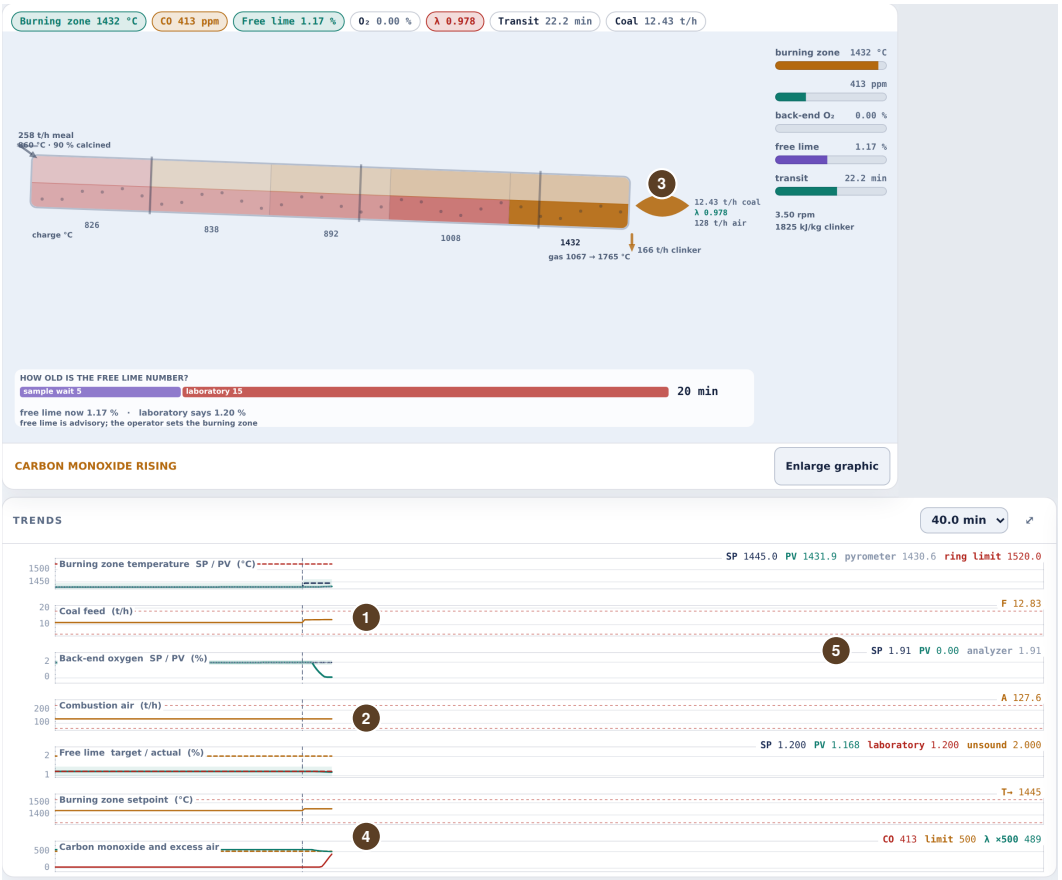


Figure 11.3 Scenario 4 at $t = 11.2$ min, cropped to the plant view and the trends, 1.2 min after the burning zone setpoint stepped from 1430 °C to 1445 °C, with the air on oxygen feedback alone. The coal command (1) rose from 11.05 t/h to 12.70 t/h within one scan and now stands at 12.83 t/h, with 12.43 t/h reaching the burner, while the combustion air (2) is still 127.6 t/h. The flame (3) is at $\lambda = 0.978$ with 413 ppm of carbon monoxide (4) and the back-end oxygen is 0.00 %, but the analyzer (5), 0.8 min of gas travel downstream, still reads 1.91 %.

shipped configuration higher. The application therefore reports carbon monoxide time and free lime beside the energy number, because a single figure of merit is blind to the constraint violations that reduce it.

Of the two structural choices, the air structure decides the outcome. The suggested tunings with feedback air still fail, with a carbon monoxide peak of 1329 ppm and 11.0 min of reducing flame, while the ratio with an advisory free lime loop passes, at a free lime error of 72.4 % min against 22.7 % min with the loop closed.

Figure 11.3 shows the kiln during scenario 4 at 11.2 min, after the coal has moved and before the air has.

11.5.6 Predictive Controller Against the Loops

The MPC is compared with the loops in Table 11.7. Every run lasts 300 min, or 360 min for scenario 10, at the case's own seed with its instruments, noise and events. The loops are the ratio with trim and the closed free lime loop at the case's tunings, or the challenge's suggested settings, while

Table 11.7 The MPC Against the Ratio With Trim and a Closed Free Lime Loop, With the Information and Limits It Carries Handed to the Loops

Case	Structure	FL IAE (% min)	λ_{min}	Under floor min (% min)	CO ppm (min)	Travel (t/h)
11	loops	22.48	0.948	191.5 (178.73)	891 (26.9)	100.5
	MPC	20.93	1.043	101.3 (1.48)	20 (0)	29.1
	MPC, no raised target	21.05	1.026	186.1 (10.82)	20 (0)	29.9
	MPC, assays at report	29.46	1.043	93.6 (1.58)	20 (0)	29.5
	loops, T*	19.70	0.941	190.8 (178.79)	1014 (26.4)	103.8
	loops, floor + cross	22.94	1.049	9.3 (0.05)	20 (0)	100.5
	loops, T*, floor + cross	19.60	1.049	9.2 (0.05)	20 (0)	103.5
	loops, hindsight	20.33	0.940	157.0 (106.75)	1040 (5.0)	188.5
6	loops	2.04	0.971	81.0 (58.91)	502 (0.3)	1.7
	MPC	0.39	1.050	8.5 (0.00)	20 (0)	0.6
	loops, floor + cross	0.87	1.050	0.3 (0.00)	20 (0)	0.6
7	loops	18.54	1.081	0	20 (0)	3.0
	MPC	15.12	1.096	0	20 (0)	1.7
	loops, T*	12.54	1.087	0	20 (0)	4.8
	loops, hindsight	13.24	0.998	0.9 (0.50)	169 (0)	10.9
8	loops	29.21	1.060	0	20 (0)	4.4
	MPC	4.70	1.094	0	20 (0)	3.7
	loops, T*	2.76	0.945	1.1 (0.83)	939 (0.3)	8.6
	loops, T*, floor + cross	2.68	1.082	0	20 (0)	8.3
9	loops	11.98	1.071	0	20 (0)	9.1
	MPC	8.40	1.093	0	20 (0)	7.3
	loops, T*	7.38	1.076	0	20 (0)	6.4
	loops, hindsight	3.13	1.057	0	20 (0)	10.5
10	loops	15.59	1.088	0	20 (0)	3.8
	MPC	13.75	1.096	0	20 (0)	2.7
	loops, T*	11.19	1.091	0	20 (0)	5.1
	loops, hindsight	10.82	1.021	0.7 (0.22)	20 (0)	13.5
12[a]	loops, met	22.67	1.000	142.3 (53.04)	144 (0)	100.6
	MPC, not met	23.58	1.044	82.1 (1.30)	20 (0)	29.7
	MPC, $\mu = 0.1$, met	20.94	1.046	85.6 (1.25)	20 (0)	56.9
	loops, T*, met	21.02	0.998	141.2 (53.34)	169 (0)	102.5
	loops, hindsight, met	19.08	0.983	104.0 (41.75)	339 (0)	185.7

FL IAE integrates the true free lime error, under floor is the time and depth of the flame's oxygen, $21(1 - 1/\lambda)$, below 1 %, CO is the peak with its minutes over the 500 ppm alarm, and travel is the summed movement of the coal command. Cases 6 to 10 carry no noise. Every MPC run made 600 plans (720 on case 10) and rejected none, at 0.79 to 0.84 ms per plan in Node.
[a]The challenge with its suggested settings, judged by its rule. The MPC with $\mu = 0.1$ is not shipped.

the MPC carries the settings shipped in scenario 11 into every case unchanged. Neither structure measures the feed, the meal or the coal. Three rows hand the loops what the plan carries. T* gives them the plan's burning zone target. Floor and cross-limit hold the trim's ratio above 1.05 AFR and ratio the air to the larger of the coal command and the modeled coal at the burner, so the air leads the coal up and lags it down, the cross-limited arrangement of fired equipment [6]. Hindsight is the best of 38 tunings of the three loops searched on that case after the fact. The measures come from the logged columns at 0.1 min, the sampling of the scorecard.

Scenario 11 hardens the raw meal to a burnability of 0.75 at 20 min, opens 4.5 kg/s of false air at 100 min and cuts the coal calorific value by 12 % at 180 min, with the challenge's instrument noise. The free lime peaks alike, at 1.629 and 1.631 %, and the IAE from 20 to 100 min is 11.77 under the MPC against 12.32 under the loops, as shown in Figure 11.4. The structures part at the leak. The trim believes the analyzer and takes the flame to $\lambda = 0.948$ at 165.5 min with 891 ppm of carbon monoxide, and the run ends at $\lambda = 0.990$ with 255 ppm, while the plan holds 1.043 or above and ends at 1.050. The plan burns 60.48 against 60.67 t of coal and moves it 29.1 against 100.5 t/h.

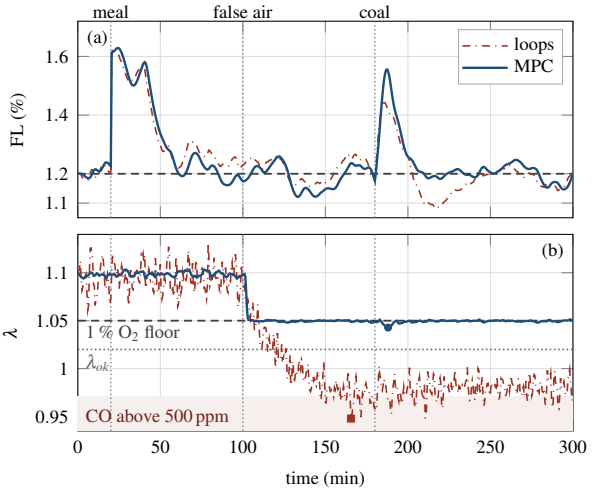


Figure 11.4 Scenario 11, the MPC against the ratio with trim and a closed free lime loop at the case's own seed, the MPC solid and the loops dash-dotted. (a) True free lime against its 1.20 % target, dashed. (b) The flame's lowest excess air per half minute against the MPC's floor at $\lambda = 1.05$, dashed, and λ_{ok}, dotted, with carbon monoxide above the alarm in the shaded band. The meal hardens at 20 min, false air enters at 100 min and the coal weakens at 180 min, at the dotted verticals. The lowest points, $\lambda = 1.043$ at 188.0 min (circle) and 0.948 at 165.5 min (square), are marked.

Removing one mechanism at a time places both gains outside the optimization. Without the raised oxygen target the plan's flame falls to $\lambda = 1.026$ with 186.1 min under the floor, while without the flame row it stays at 1.043 with 105.3 min, and the analyzer row never binds. A classical floor and cross-limit hold $\lambda = 1.049$ with 9.3 min under the floor against the plan's 101.3 min, since the plan rides a target placed exactly at the floor and is more than 0.05 % under it for only 4.1 min. The free lime gain is information. The loops handed T* reach 19.70 % min, and with the floor and cross-limit as well 19.60 % min at $\lambda = 1.049$, which beats the plan on free lime and on the flame together. What stays with the plan is calm, 29.1 against 103.5 t/h of coal travel.

The other cases repeat the pattern. Under the false air of scenario 6 the plan holds $\lambda = 1.050$ where the loops reach 0.971 and 502 ppm, and the floor and cross-limit hold 1.050 as well. On scenarios 7 to 10 the plan improves on the loops' free lime, the loops handed T* improve on the plan, and scenario 8 gives the plan its largest lead because the speed command enters its free lime law. There T* alone brings 939 ppm after the speed returns at 110 min, which the floor and cross-limit remove.

The losses are reported as measured. On free lime the plan loses to the loops handed T* on every case from 7 to 11 and in the challenge, and to the loops retuned in hindsight on scenarios 7, 9, 10, 11 and the challenge. The best of 360 settings on scenario 11 reaches 19.31 % min, with $\lambda = 0.920$ and 48.8 min over the alarm. On scenario 9 its free lime peaks at 1.562 against 1.468 %, because the move price that keeps the coal calm makes the plan slower than the PI on a heat load it cannot measure, and at a move price of 0.1 the peak is 1.495 %. At three further seeds of scenario 11 the plan reaches 21.13, 23.48 and 24.64 against the loops' 26.13, 27.54 and 22.32 % min, so its free lime lead holds on three of four seeds, while its flame stays at 1.043 or above on all four, and a halved integration step changes no ranking.

The same calm fails the challenge. With the suggested settings the plan's burning zone falls to 1414.3 °C against a target of 1431.3 °C at 26.1 min on the raw meal step, an error of 17.0 °C against the rule's 15 °C, while the loops' PI meets the rule at 12.6 °C. The feed step is a heat load that the plan sees only through its disturbance estimate and answers at the move price chosen on scenario 11. A move price of 0.1 would meet the rule at 14.8 °C with 56.9 against 29.7 t/h of coal travel, and

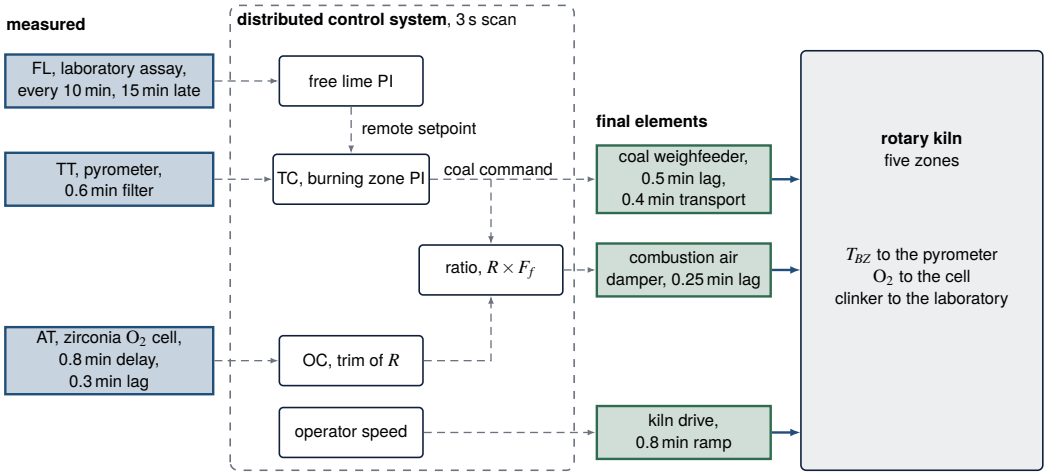


Figure 11.5 The loops of the simulated kiln drawn with their hardware. The temperature controller TC sets the coal weighfeeder, and the same coal command, multiplied by the ratio R, sets the combustion air in the same scan, while the oxygen controller OC only trims R. The free lime result writes the burning zone setpoint each time an assay arrives, and the kiln speed stays with the operator.

at the three further seeds the shipped plan misses at 15.9 and 15.7 °C and meets it at 14.0 °C. The limits the challenge does not score fall the other way, with the plan's flame 82.1 min and 1.30 % min under its floor and never below λ_{ok}, and the loops' flame 142.3 min and 53.04 % min under it, down to $\lambda = 1.000$ and below λ_{ok} for 19.9 min.

11.6 Technology in Practice

Figure 11.5 draws the loops of Section 11.4 with their hardware.

What Measures It

The burning zone is read by a pyrometer that looks at the hot charge from the burner end. The model's 0.6 min pyrometer filter stands for the smoothing such a reading gets. The oxygen at the back end can be read by a high-temperature zirconium oxide cell with platinum electrodes, one side in the gas and the other in heated reference air, drawn in Figure 11.6. The oxygen in the exhaust of an efficiently controlled burner sits near 2 %, and it matters for fuel efficiency and for nitrogen oxides alike [20]. The cell reads whatever oxygen reaches it, which is why the false air of Equation (11.9) biases it. The model's 0.8 min delay and 0.3 min lag stand for the gas travel and the cell's response. Free lime has no instrument at all: it is a laboratory assay.

> **TECHNOLOGY Sensor: The Burning Zone Pyrometer**
>
> A pyrometer reads temperature from the thermal radiation a hot body emits, which suits a burning zone no probe would survive. The instrument sees only a surface, and its main weakness is inaccuracy: the emittance of the target has to be found once against a contact reading and set in the instrument. Radiation reflected from other bodies taints the reading, and an instrument aimed through a gas may report the solid behind it [20]. A wrong emittance shifts every reading, and the loop then holds the wrong temperature.

What Moves It

The coal reaches the burner through a feeder that the temperature controller sets.

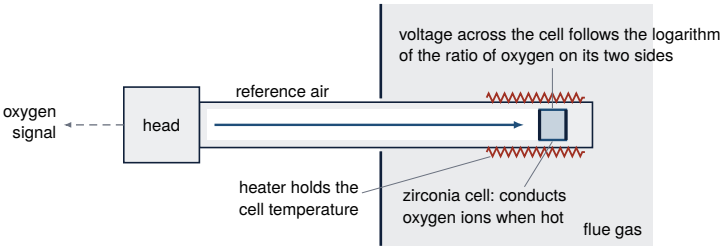


Figure 11.6 The zirconia oxygen cell. Notice that the voltage compares the oxygen in the gas at the probe with the reference air. Nothing in the cell knows whether that oxygen passed through the flame, so air leaking in at the inlet seal reads as excess air, the bias of scenario 6.

TECHNOLOGY Final Element: The Coal Weighfeeder

A weighfeeder carries the solid on a short conveyor belt whose middle section rests on rollers coupled to load cells. The load cells give the weight per length of belt, a speed sensor gives the belt speed, and their product is the mass flow, which the feeder's controller holds at a setpoint in tons per hour. The belt loading fluctuates in normal operation and the flow calculation lags the weighing, so the flow trend is noisy and slightly late [20]. The model's 0.5 min feeder lag and 0.4 min of transport to the burner stand for this path.

The combustion air is moved by a damper or fan with a 0.25 min lag in the model, and in the ratio structure its setpoint is the coal command times R. The kiln drive follows the operator's speed with a 0.8 min ramp. The model has no dead band or stiction.

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 [21]. In the app every loop acts on a 3 s scan against a burning zone time constant of 13.0 min, and the predictive controller plans every 0.5 min. On such a platform the temperature, oxygen trim and free lime loops are PID function blocks of the regulatory layer. The ratio is a multiplying block between the coal command and the air setpoint, and the free lime loop writes the temperature loop's remote setpoint each time an assay arrives.

IN PRACTICE Cross-Limiting the Air and the Fuel

Burner control in practice adds two selectors to the ratio. The air setpoint is the larger of the firing command and the measured fuel flow, and the fuel setpoint is the smaller of the firing command and the measured air flow. The air therefore leads when firing rises and the fuel leads when it falls, so the mixture runs lean, which is safe, rather than rich, which is dangerous. A rate limit on the firing command keeps a lean transient from blowing the flame out, and a multiplying function sets the air to fuel ratio [20]. The floor and cross-limit rows of Table 11.7 hand the kiln's loops a version of this.

What Is Modernizing It

The first development is supervisory control above these loops. Rule-based, fuzzy logic and expert systems try to act as the best human operator. One such system, first installed at a cement works in 1985, cut fuel use by nearly 8 %. Model predictive control with dynamic models of the kiln followed, and a kiln fitted in 1999 needed 4 % less energy. Typical savings of such systems are put at 2.5 to 5 % [22]. The second is measurement of the feed. Online analyzers give the chemical composition of the raw materials at once, so the blend can be corrected before it reaches the kiln [22], which meets the burnability disturbance of scenario 7 at its source. The third is the dynamic kiln model. One-dimensional models that track the individual clinker and fuel compounds are offered as

a stepping stone for developing process control on existing plants [23], the role the eighteen-state model of Section 11.3 plays here at a coarser scale.

11.7 Limitations

The reaction scheme is the largest simplification. Two lumped reactions with Arrhenius rates reproduce the shape of the conversion profile and the switch-like behavior of clinkerization, and they cannot represent the phase sequence, the melt or the effect of the alumina and iron moduli that decide burnability in practice [8]. The burnability multiplier β is a stand-in for all of it, which is adequate for a control argument about an unmeasured disturbance and inadequate for any statement about a particular raw mix.

The rest of the pyroprocessing line is a boundary condition. A real line recirculates alkalis and sulfur between the kiln and the preheater, and those cycles produce the build-ups, blockages and coating changes that dominate real kiln operation. Integrated simulators exist to capture exactly those recycles [13], and nothing in this application can substitute for one. Ring formation is an alarm here rather than a mechanism, so the kiln never restricts and the residence time never changes as a result.

Carbon monoxide is algebraic, and the clinker damage a reducing flame causes is asserted rather than computed. Nitrogen oxides are absent, which matters because they constrain the operating point from the opposite side and make the acceptable combustion window narrower than this application shows. The gas phase carries no composition at all, so the model can describe an oxygen reading and cannot describe a gas analysis.

The predictive controller's model leaves out what its estimators have to absorb. Its two modes have gains fixed at the design feed and firing, so a change of feed rate or operating point is model error. The kiln speed enters only the free lime law, although a speed step moves the burning zone by −1.80 °C per rpm at about 57 min before it settles at −0.74. Burnability reaches the plan only through the assays, false air only through the oxygen bias, and the true air-fuel ratio not at all, because the excess air the plan protects is computed from its commands with the declared stoichiometry. The raised target sits exactly at the floor rather than a margin above it, and the plan prices no fuel.

Finally, the free lime measurement is a clean delayed sample with optional noise. Real assays carry sampling error from a clinker stream that is not homogeneous, and disagreement between the laboratory result and the kiln condition is a routine operational problem that the application does not represent.

> ⚠ **OPEN CHALLENGE** **Clinker Quality That Holds Still**
>
> Ordinary Portland cement is about 90 % clinker, and its strength depends on the clinker's composition. Large swings in that composition limit how much supplementary material can replace clinker, because more clinker is needed to cover the spread in strength. Reducing the variation of clinker quality has been called a must for a more sustainable cement, through better quality predictions and through control [23]. The free lime loop shows why it is hard: the quality number arrives 15.9 to 25.9 min late. Progress would be a quality estimate at every scan that the burning zone loop could trust.

11.8 Discussion

The two ends of the kiln lead to opposite control conclusions, and the application presents both in one plant.

At the fast end the remedy is structural and complete. The disturbance is the control system's own output, it is known exactly and immediately, and using it removes the excursion rather than reducing it. The general form is that feedforward is built from whatever signal arrives with the disturbance, and feedback then corrects the model rather than the plant. The ratio is the model and 11.55 tonnes of air per tonne of coal is its one parameter. The same structure also degrades gracefully against a corrupted measurement, as Section 11.5.3 shows, which is a second argument that is usually absent from the textbook presentation.

At the slow end no structural remedy exists. Up to 26 min of assay dead time on a 13 min zone, driven through a temperature loop that answers in a few minutes, gives a loop that cannot be made fast, and (11.15) establishes that limit before any tuning is attempted. What follows is a division of roles. The burning zone temperature becomes the controlled variable because it is fast and measured, the free lime becomes a slow trim on its setpoint, and kiln speed becomes available as a faster quality handle. Industrial model predictive implementations arrive at the same division, using the assay to correct an inferred quality rather than to close a loop [7].

The model of Section 11.3 also supports control methods beyond the PID, ratio and trim loops and the linear predictive controller of Section 11.4.5 that the application implements. The eighteen differential states, algebraic gas profile and actuator limits match the problem form solved by nonlinear model predictive control, which would coordinate coal, air and kiln speed in one constrained calculation and could run an economic objective on the specific heat consumption of Table 11.6 while holding the carbon monoxide limit as a hard constraint rather than a scorecard entry [24]. Moving horizon estimation over the same equations would reconstruct burnability and false air from the temperature and oxygen histories, converting the two unmeasured disturbances of scenarios 6 and 10 into soft-sensed quantities, and Kalman filtering of the sampled assay would supply a continuous free lime estimate between laboratory results. The Gekko Optimization Suite implements these methods in Python from model statements close to (11.2) through (11.12) [25].

The third conclusion concerns performance metrics. Table 11.6 shows a configuration that wins on specific heat consumption and loses on every other measure, by supplying insufficient combustion air. Any single figure of merit on this plant can be improved by actuator action that the figure does not observe, which is the general reason that operating limits belong on the scorecard beside the objective rather than only in an alarm list [5].

11.9 Conclusion

A rotary kiln combines a combustion excursion lasting about three minutes and a quality loop carrying up to 26 min of dead time in one vessel. The fast problem has a complete structural remedy, because the air can be computed from a fuel command that is known in the scan it is issued, and the application measures that remedy as the difference between a 611 ppm carbon monoxide peak and none at all through the same firing move. The slow problem has no such remedy, and the dead-time budget of (11.15) establishes that limit before tuning is attempted rather than after a shift of oscillation. Between the two sits a combustion penalty so asymmetric that the kiln gives up 12.2 °C of burning zone temperature to stay clear of it, and an oxygen measurement that false air biases without any change in the flame's excess air. A graded five hour run assembled correctly removes 88.7 min of reducing flame and pays 121.6 kJ per kilogram of clinker for it, which is the trade-off the plant faces. The added predictive controller holds the flame at an excess air of 1.043 or more through false air where the ratio with trim falls to 0.948, mostly through its raised oxygen target, while the loops handed its burning zone target, a ratio floor and a cross-limit reach 19.60 against its 20.93 % min of free lime error.

KEY POINTS

- The air should follow the coal command, not the oxygen reading, which arrives only after the gas has traveled: a 15 °C firing move then leaves no carbon monoxide peak, against 611 ppm.
- False air biases the oxygen analyzer, and the ratio turns that failure into a slow drift.
- With 15.9 to 25.9 min of dead time on a 13.0 min zone, the free lime loop is a slow trim on the burning zone setpoint.
- A configuration can win on specific heat by running a reducing flame, so limits belong on the scorecard.

Where else this applies. Volume 3, Chapter 4 adds cross-limiting to the same ratio, and Volume 5, Chapter 3 tunes a loop through a slow analyzer.

References

[1] U.S. Geological Survey, "Mineral commodity summaries 2025: Cement," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[2] R. M. Andrew, "Global CO_2 emissions from cement production, 1928–2017," *Earth System Science Data*, vol. 10, no. 4, pp. 2213–2239, 2018.

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

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

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

[6] F. G. Shinskey, *Process Control Systems: Application, Design, and Tuning*, 4th ed. New York: McGraw-Hill, 1996.

[7] K. S. Stadler, J. Poland, and E. Gallestey, "Model predictive control of a rotary cement kiln," *Control Engineering Practice*, vol. 19, no. 1, pp. 1–9, 2011.

[8] W. Kurdowski, *Cement and Concrete Chemistry*. Dordrecht: Springer, 2014.

[9] "Cement," in *Encyclopædia Britannica*, 11th ed. Cambridge: Cambridge University Press, 1910, vol. 5, pp. 653–659, read in the Wikisource transcription (en.wikisource.org, 1911 Encyclopædia Britannica/Cement), accessed November 19, 2026.

[10] E. Worrell, L. Price, and C. Galitsky, "Energy efficiency improvement opportunities for the cement industry," Lawrence Berkeley National Laboratory, Tech. Rep., 2008.

[11] L. P. Holmblad and J.-J. Østergaard, "Control of a cement kiln by fuzzy logic techniques," *IFAC Proceedings Volumes*, vol. 14, no. 2, pp. 809–814, 1981.

[12] K. S. Mujumdar and V. V. Ranade, "Simulation of rotary cement kilns using a one-dimensional model," *Chemical Engineering Research and Design*, vol. 84, no. 3, pp. 165–177, 2006.

[13] K. S. Mujumdar, K. V. Ganesh, S. B. Kulkarni, and V. V. Ranade, "Rotary cement kiln simulator (RoCKS): Integrated modeling of pre-heater, calciner, kiln and clinker cooler," *Chemical Engineering Science*, vol. 62, no. 9, pp. 2590–2607, 2007.

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

[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] O. J. M. Smith, "Closer control of loops with dead time," *Chemical Engineering Progress*, vol. 53, no. 5, pp. 217–219, 1957.

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

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

[19] J. H. Lee and M. Morari, "Robust inferential control of multi-rate sampled-data systems," *Chemical Engineering Science*, vol. 47, no. 4, pp. 865–885, 1992.

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

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

[22] E. Worrell, C. Galitsky, and L. Price, "Energy efficiency improvement opportunities for the cement industry," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-72E, 2008.

[23] J. L. Svensen, W. R. L. da Silva, J. P. Merino, D. Sampath, and J. B. Jørgensen, "A dynamical simulation model of a cement clinker rotary kiln," arXiv:2405.03200, 2024.

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

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

11.10 Exercises

11.1 Concept check · 2 min

The burning zone setpoint steps up 15 °C and the temperature controller raises the coal command within one scan. Why does ratioing the combustion air to that command keep the carbon monoxide at its floor when no tuning of an oxygen feedback loop can?

(a) The coal command is known in the scan it is computed, so the air moves before any gas reaches the analyzer

(b) The ratio trim is tuned five times faster than the oxygen feedback loop it replaces

(c) The ratio structure raises the oxygen setpoint for the duration of every firing move

(d) The ratio removes the coal transport delay between the feeder and the burner

11.2 Concept check · 1 min

The raw meal becomes harder to burn, its burnability falling from 1.00 to 0.72, while the burning zone loop holds 1430 °C. Which reading is the first to show the change?

(a) The pyrometer, as the burning zone cools by about 10 °C

(b) The laboratory free lime, once a report covers clinker made after the change

(c) The oxygen analyzer, because a harder meal needs more combustion air

(d) The coal feed, which the temperature loop raises by about a tenth

11.3 Concept check · 2 min

A worn inlet seal admits 4.5 kg/s of false air while the air is on oxygen feedback. What does the controller do?

(a) It opens the damper further, because the leak dilutes the flame and cools it

(b) It does nothing, because the analyzer's 0.3 min lag filters the leak out

(c) It closes the damper and drives the flame's own excess air below stoichiometric

(d) It leaves the air alone and the burning zone loop cuts the coal instead

11.4 Theory and calculation · 6 min

Use Equation (11.1) and the parameters of Table 11.2 at the design feed of 258 t/h.

(a) Compute the transit τ, the zone residence τ_z and the zone hold-up M_z at 3.5 rpm, to two decimals. *Checkpoint:* 22.16 min; 4.43 min; 19.06 t

(b) Evaluate the burning zone time constant $M_z 10^3 C_{p,s}(1+f_{ref})/[60(\dot{W}_s + UA_5)]$ with $UA_5 = UA_0 w_5$ from Equation (11.3), and the same expression without the refractory, to one decimal. *Checkpoint:* 13.0 min; 3.1 min

(c) Scenario 8 raises the speed to 4.6 rpm. Compute the new transit to one decimal, and say why it does not enter the free lime loop's dead time. *Checkpoint:* 16.9 min

11.5 Theory and calculation · 7 min

Use Equations (11.7) to (11.9) with AFR = 10.5 at the design coal rate of 11.033 t/h.

(a) Compute λ and the oxygen reading at 127.434 t/h of air, and the carbon monoxide at $\lambda = 0.95$. *Checkpoint:* 1.100; 1.91 %; 853 ppm

(b) False air of 4.5 kg/s enters at the inlet seal. Compute $\hat{\lambda}$ and the analyzer reading with the air unchanged. *Checkpoint:* 1.240; 4.06 %

(c) An oxygen feedback loop drives the analyzer back to 1.91 % with the coal held at 11.033 t/h. Compute the air, the flame's λ and the carbon monoxide at steady state, and compare with the run in Section 11.5.3, which ends at $\lambda = 0.963$ and about 635 ppm. *Checkpoint:* 111.2 t/h; 0.960; 678 ppm

11.6 Theory and calculation · 7 min

The free lime loop's dead time is the budget of Equation (11.15). Use the SIMC rules $K_c = \tau_p/(K_p(\tau_c+\theta_p))$ and $\tau_I = \min(\tau_p, 4(\tau_c+\theta_p))$.

(a) Give the shortest and longest effective dead time and their ratios to the 13.0 min zone constant and to the 3.7 min the closed temperature loop takes to cover 63 % of a setpoint step, to one decimal. *Checkpoint:* 15.9 and 25.9 min; 1.2 and 2.0; 4.3 and 7.0

(b) Tune the loop on the 13.0 min zone with $\theta_p = \tau_c = 25.9$ min, for process gain magnitudes of 0.015 and 0.018 %/°C. Give K_c to one decimal and τ_I. *Checkpoint:* 16.7 and 13.9 °C/%; 13 min

(c) Which sign does the controller action take, and why? *Checkpoint:* $a = -1$

11.7 App challenge · 8 min

Run scenario 4, "Oxygen feedback alone, against a firing move", as shipped at the default seed, to 60 min at 5 min per second; pause within a minute of 60 min. Repeat with scenario 5, "Ratio the air to the fuel".

(a) For each run read the Reducing time tile and the burning zone IAE in the Learn panel, to ±0.5 °C min. *Checkpoint:* 0.8 and 0.0 min; 68.7 and 81.7 °C min

(b) Export each run's CSV and give the minimum of the `lambda` column to three decimals. *Checkpoint:* 0.964; 1.085

(c) By how many percent of scenario 4's value is the burning zone IAE higher under the ratio, to the nearest percent? Explain the increase from the excess air of the two runs during the rise. *Checkpoint:* 19 %

11.8 App challenge · 9 min

Load scenario 7, "Close the loop on free lime", at the default seed. Before Run, set Free lime gain to 45 and Free lime integral time to 12 min. Run to 360 min at 60 min per second and export the CSV, whose time column is in hours.

(a) From the `lime_pct` column after 120 min, estimate the oscillation period from successive highest peaks, to ±5 min. *Checkpoint:* about 57 min

(b) Give the lowest and highest `lime_pct` in the sixth hour, 300 to 360 min, to ±0.02, and the range of the burning zone setpoint `bz_sp_C` over the whole run, to ±1 °C. *Checkpoint:* 0.92 to 1.51 %; 1426 to 1485 °C

(c) The burnability step at 20 min first moves the setpoint at 40 min. Compare the period with that dead time and say what the shipped tuning of 15 and 13 min achieves instead. *Checkpoint:* about three dead times

11.9 App challenge · 9 min

Load the challenge at its default seed and run it as shipped to 300 min at 60 min per second. Then reload it, press Reveal suggested tuning before Run and run it again.

(a) For each run read the Reducing time, Specific heat, Coal per tonne and Free lime tiles at 300 min. *Checkpoint:* 88.7, 1600, 68.7, 2.73; 0.0, 1721, 71.4, 1.20

(b) A plant manager ranks configurations by specific heat alone. Which run wins, and what does the ranking hide? *Checkpoint:* the shipped run

(c) Keep the suggested tuning but set Air control back to O_2 feedback before Run. Read the Reducing time tile and say which of the two structural choices decides the challenge. *Checkpoint:* 11.0 min; the air structure

11.10 Application area · 3 min

Table 11.3 shows the burning zone hottest near $\lambda = 1.00$, yet the kiln runs at $\lambda = 1.10$.

(a) What does the operating point give up in burning zone temperature and free lime against $\lambda = 1.00$? *Checkpoint:* 12.2 °C; 0.20 %

(b) Give the reason a plant accepts that cost, and name one constraint the model leaves out that would narrow the window further. *Checkpoint:* carbon monoxide; nitrogen oxides

11.11 Application area · 4 min

Table 11.5 sweeps the kiln speed at fixed firing.

(a) From 3.5 to 4.2 rpm, how much does the free lime rise? Using the chapter's steady gain of 0.015 to 0.018 %/°C and 52.1 °C per t/h, how much hotter would the burning zone have to run, and how much more coal would that take? *Checkpoint:* 0.245 %; 14 to 16 °C; 0.26 to 0.31 t/h

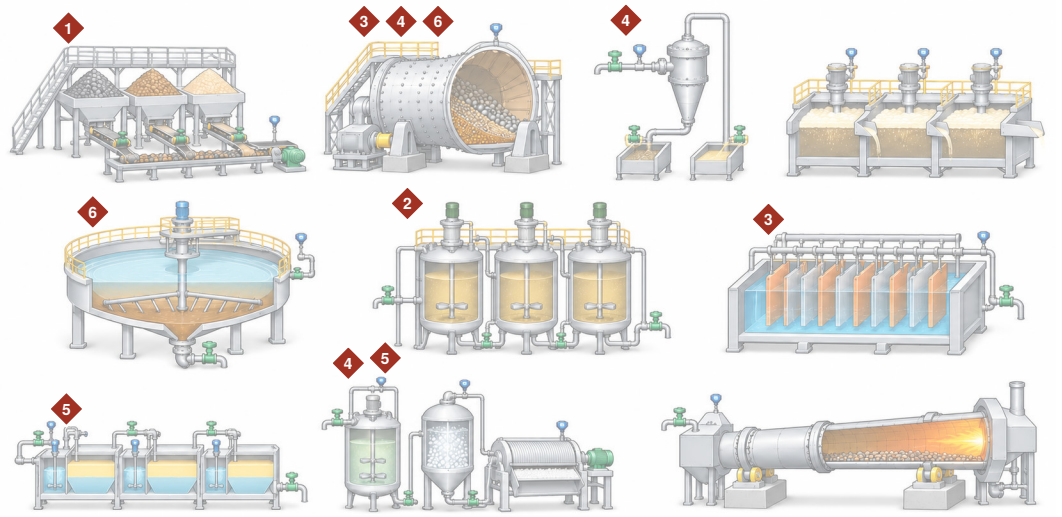
(b) Speed moves the free lime faster than the burning zone setpoint does. Why is it rarely used as a quality handle in practice? *Checkpoint:* throughput, bed depth, coating

11.12 Review the agent · 5 min

An AI assistant tunes the free lime loop: "The dead time is the 22.2 min kiln transit plus the 15 min turnaround, half a 10 min sample interval and 0.9 min of coal path, about 43 min. SIMC with $\tau_c = \theta_p$ on the 13 min zone and a gain of 0.016 %/°C gives $K_c \approx 9.4$ and $\tau_I = 13$ min. Expect a hot tuning to ring with a period of about two hours." Find the flaw and say how you would check it in the app.

CHAPTER 12

Summary and Outlook



Open gaps 1 arriving ore 2 blind estimates 3 inferred limits 4 hidden drift 5 inventories 6 protection

Notice that the mill carries three gaps, and that no gap marks the flotation cells or the kiln. Illustration, not to scale.

IN PLAIN TERMS **What Ten Mineral Plants Taught**

The main idea. Each chapter of this book ran one part of a mineral plant in a simulator, from the ore belt to the kiln. In every one the rock changed in ways that no instrument could see, and the test that judged the product came back late or never.

The challenge. A loop can hold its own reading steady while the plant drifts somewhere else. A late test also limits how hard a loop may correct.

Why it matters. The plants that did best measured what arrived and acted on it early, chose their targets with care, and let protection win when a limit was near.

Ten chapters followed ore from a belt scale to a cathode, a hydroxide product and a clinker. Each met the two problems of the volume's title. The ore changed without being measured for what mattered, and the quality that judged the plant arrived late or not at all. Table 1.1 of Chapter 1 listed, for each chapter, how the ore or feed varies, what is measured late or not at all, and the structure that handles it. This chapter adds what each structure 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, how the ore varies and what is late or never measured. The other columns name a structure that worked, one that failed, and the measured result with the table or section of the chapter that reports it.

The columns make five points, and Figure 12.1 draws one pair of results from each row.

Table 12.1 How Variable Ore and Late Assays Were Handled, Chapters 2 to 6

Chapter	Ore varies, and what is late	What worked	What failed	Measured result
2, ore feed	Calibrations, belt speed, an empty silo. The scale weighs the total and never the blend.	Speeds set through the calibrations. A bump test on a drifted feeder. Holding back ore that would overload the belt.	A master PI that chases transport dips. Every structure on a wrong calibration.	On a speed change 65.4 % on spec for the master PI against 78.5 for ratio feedforward and the MPC. A drifted feeder 37.6 to 40.0 % (Table 2.4). Shift over the rating 172 against 0 s (Table 2.5).
3, SAG mill	Hardness, ball charge, liner wear. The toe is inferred, the size 2 min late.	The cautious impact loop with a slower load loop that replaces lost steel without overshoot.	The recommended impact tuning on a noisy shift. Toe tracking on an inferred toe.	Two points of ball filling took the margin from +0.22 to −3.27° and the product from 2367 to 4136 µm with every loop on setpoint (Table 3.7). Shift toe error 2.42, 3.98 and 5.45° (Table 3.9).
4, hydro-cyclones	Coarseness, grindability, pump wear. Size 2 min late, circulating load never.	A pressure cascade under the size loop against pump wear. Commissioning from the level loop outward.	The same cascade against worn liners or a ninth cyclone. A tighter level loop.	Fineness IAE 154.9 to 3.2 for impeller wear, 25.1 to 73.1 for liners (Table 4.4). A level loop 7.8 times tighter raised the fineness IAE 7.6 % (Table 4.6).
5, flotation cell	Feed grade, throughput, blower. Grades 2.0 and 3.5 min late.	The grade loop alone at an economic setpoint, 0.55 m and 8.4 % copper, over a fast level loop.	An aggressive grade loop. A second quality loop on the collector. Tight regulation at the design depth.	Aggressive loop $3305 /h against 7270 at the same 76.3 % recovery (Table 5.3). Graded run 9890 suggested, 8580 at the as-found setpoints, 7072 with both quality loops (Table 5.4).
6, thickener	Feed solids, settleability. Settleability never, density, torque and bed in minutes.	A high select of density, torque and bed demands with external reset, and a feasible density setpoint.	A low select, which blocks the protection. The loaded setpoint and slow tuning.	Low select torque peak 109.9 % (Table 6.6). Graded shift torque 104.7 against 65.9 %, bed 3.280 against 2.567 m (Table 6.8).

Every loop on setpoint, and the plant somewhere else. The most repeated result of the volume is a plant that moved while its loops held. In Chapter 3 two points of ball filling left the impact angle, the load and the density on setpoint while the media landed inside the toe and the product coarsened by 75 %. In Chapter 10 an electrode reading 0.70 low raised the reagent 24 % while every number on the daily report improved. In Chapter 8 a diffusivity 20 % low left the temperature and copper loops with the integrated errors of a correct run while the cell shorted on nodules. In Chapter 4 the level and pressure loops held while the circulating load rose from 288 to 386 % (Section 4.5). A loop measures its own variable. In each case the variable that decided the result was one that no instrument read.

What the soft sensor cannot read arrives with the assay. A soft sensor removes the assay's dead time only for the disturbances among its inputs. In Chapter 7 the soft-sensor loop moved its cyanide setpoint 0.26 h after a measured feed grade step, and the assay loop 10.26 h after it. A lower leachability, which no input carries, first moved the setpoint at 20.1 h under either loop (Section 7.6.2). The bias update that joins model and assay is then a filter on stale, noisy data. A gain of one raised the actuator travel and the peak, and a gain of zero let the tailings sit above the limit for 172 of 200 hours. The same holds without a soft sensor. With its feedforward on, Chapter 9 passed its shift with

Table 12.2 How Variable Ore and Late Assays Were Handled, Chapters 7 to 11

Chapter	Ore varies, and what is late	What worked	What failed	Measured result
7, gold CIL	Grade, leachability, copper and sulfide parcels. Tailings sampled every 4 h, reported 6 h later.	A fixed cyanide setpoint on the graded week. A bias update gain near a third.	The soft-sensor cascade at its recommended target. No bias update. A bias gain of one.	Hours above 0.40 g/t were 0.0 fixed, 21.6 for the cascade and 172.0 with no update (Table 7.3). A gain of one raised the cyanide travel 63 % over 0.15 (Table 7.2).
8, copper refining cell	Purification upsets, passivation, dirty electrolyte. The limiting current never, copper 2 h late.	Utilization held on the inferred limit with margin, a voltage override, more leveler for a dirty electrolyte.	Fixed current. A diffusivity 20 % low in the estimate. Running for the most copper.	Believed 58.99, true 73.74 %, and the cell shorts (Table 8.3). At 66 % it plates 17 % more copper than at 58.5 and earns 24 % less (Table 8.1).
9, rare earth cascade	Feed grade and concentration. Purity and recovery from a laboratory 1 to 5 h old.	Feedforward from the feed analysis with a purity loop detuned to the plant's day, 0.05 per point and 24 h.	The purity loop at 0.5 per point and 6 h. The crossed pairing. Either part of the suggested structure alone.	Shift passed on 5 of 5 seeds at $2.93 per kilogram, the shipped loop on none, feedforward alone at $3.00 to 3.01 (Table 9.5). MPC 17.78 h below specification (Table 9.6).
10, lithium refinery	Ore grade, fuel header, heating value. Product strength read two units downstream.	A fuel flow cascade under the kiln temperature against header pressure. Each loop tuned for its own clause.	The cascade against a heating value loss. A pH loop that trusts a biased electrode.	Header loss IAE 526.3 against 5.7, heating value 158.4 under both (Table 10.4). An electrode 0.70 low raised the yield from 87.8 to 91.2 % and the reagent 24 % (Section 10.6).
11, cement kiln	Raw meal, coal heating value, false air. Free lime 15.9 to 25.9 min late.	Air ratioed to the coal command with oxygen trim. The free lime loop as a slow trim. A ratio floor and cross-limit.	Oxygen feedback alone. A hot free lime tuning against 20 min of dead time.	Carbon monoxide peak 611 against 20 ppm on one firing move (Table 11.4). Graded run with a reducing flame for 88.7 against 0.0 min, at 1600 against 1721 kJ/kg (Table 11.6).

every purity gain from 0.05 to 0.3 at integral times of 12 and 24 h, and failed with every gain at 6 h. A loop on a late measurement is tuned to the plant's response and the report interval, not to the size of the error.

Measure what arrives, feed it forward, trim from the assay. The structures that worked moved the actuator from something measured before it reached the product. The ore feed set each feeder through its calibration, the kiln ratioed its air to the coal command, and the rare earth cascade keyed its scrub to the feed analysis. A cascade is the same idea one step inside the plant, and its value depended on the path. The lithium kiln's flow loop cut the integral error of a header pressure loss by a factor of 52 to 93, depending on when the loss landed in the scan, and did nothing for a heating value loss. The hydrocyclone's pressure loop cut the fineness error of impeller wear 49 times and made worn liners about three times worse, because restoring the pressure drew dilution water and thinned the feed. The late assay then trims what the fast path cannot see.

The setpoint decided more than the tuning. Several chapters gained more from where a loop was set than from how it was tuned. The flotation cell regulated at its as-found setpoints with the best

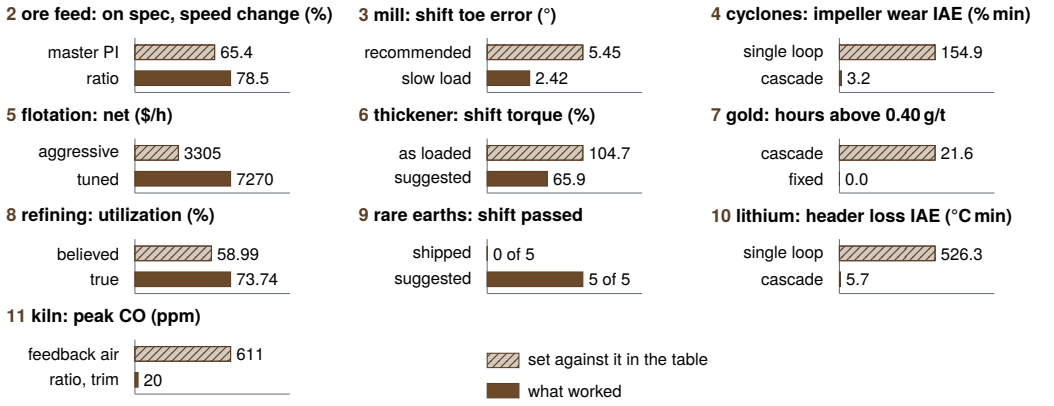


Figure 12.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 12.1 and 12.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 2, 5 and 9. In the panel of Chapter 8 the hatched bar is the utilization the cell believed and the solid bar the true one. Notice that the better result came from a change of structure, measurement or setpoint more often than from tuning alone.

depth tracking of its table and earned $1309 /h less than at 0.55 m. The refining cell plated the most copper at 66 % of its limiting current and earned the least. The thickener met its graded shift only with a density setpoint the rake could carry. The gold circuit held its tailings limit through the graded week with no outer loop at all, at a fixed 120 mg/L of cyanide, where the recommended cascade spent 21.6 h above it. Each optimum was flat or bounded by a limit that no instrument measured, so choosing it was an engineering decision and not a tuning result.

A faster loop bought nothing the disturbances could use. Tuning for a step response misled on plants whose disturbances are slow and whose measurements are noisy. The mill's recommended impact tuning, 15 times better on a setpoint step, had 3.3 times the integral error of the cautious tuning over the graded shift. The hydrocyclone's level loop, tightened, made the product worse. The kiln's free lime loop, pushed to a gain of 45 °C per percent, cycled with a period of 57 min on 20 min of dead time (Section 11.5.4). The flotation grade loop at $K_c = 12$ cycled on its own analyzer, sent the tailings valve through 2257 % of its range and cost 55 % of the cell's income.

12.2 Choosing a Method

The chapters suggest an order of questions for a circuit that is new to the reader. Each question names the chapter where its answer decided the result, and Figure 12.2 draws the same order as a flowchart.

1. *What in the ore varies, and what reads it, and when?* Fill in a row of Table 1.1 before choosing a structure. A scale that reads the total cannot see the blend (Chapter 2).

2. *Does the disturbance enter before a fast measurement?* Then a ratio, a feedforward or a cascade removes it (Chapter 11, Chapter 10). Check what the correction moves on its way, as the liners of Chapter 4 show.

3. *How old is the quality measurement when it is used?* Count the sampling and the report (Section 1.4.5). Set the integral time to the plant's response and the report interval (Chapter 9), and design the bias update as a filter (Chapter 7).

4. *Is the controlled variable inferred?* Name what the inference cannot see, and back the setpoint off by the error the estimate may carry (Equation (1.9), Chapter 8, Chapter 3).

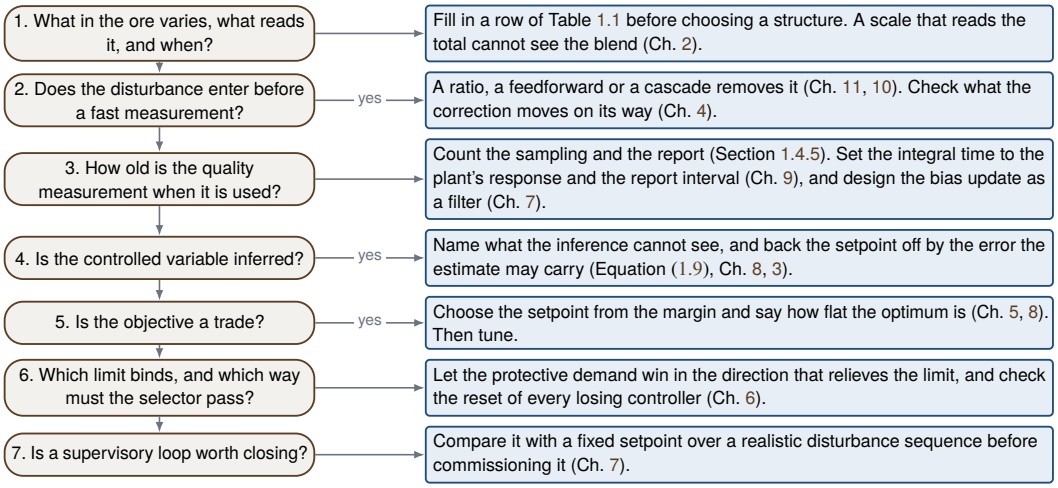


Figure 12.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the first three questions ask what varies, what is measured and when, and are answered before any loop is tuned.

5. *Is the objective a trade?* Choose the setpoint from the margin and say how flat the optimum is (Chapter 5, Chapter 8). Then tune.
6. *Which limit binds, and which way must the selector pass?* Let the protective demand win in the direction that relieves the limit, and check the reset of every losing controller (Chapter 6).
7. *Is a supervisory loop worth closing?* Compare it with a fixed setpoint over a realistic disturbance sequence before commissioning it (Chapter 7).

A predictive controller paid where it handled a constraint on material already in the plant. On the ore belt it was the only structure that kept the loading inside its rating through the infeasible shift under a temporary recipe, against 172 and 148 s over for the master PI and ratio feedforward (Table 2.5). On the kiln it held the flame at an excess air of 1.043 or more through false air where the ratio with trim fell to 0.948 (Table 11.7). Elsewhere most of its gain belonged to an operating point or to information the loops could be given. The flotation MPC's own optimum, entered as fixed stations, came within \$11 /h of it (Table 5.5). The mill's loops given the MPC's breakage estimate as a feedforward came within 0.4 % of its ore (Table 3.11). The kiln's loops given its burning zone target, a ratio floor and a cross-limit reached 19.60 against its 20.93 % min of free lime error. On the gold circuit a setpoint optimizer recovered all but \$1.9 per hour of the plan's gain, and the plan ran over the tank-five discharge limit for 28.3 h (Table 7.4). On the rare earth shift the MPC spent 17.78 h below specification where the compensated PI spent none. Before reaching for prediction, try the feedforward, the setpoint and the selector. Volume 9 treats that choice in full.

12.3 Where Practice Is Heading

Knowing the ore before the plant does. Geometallurgy has become a tool for reducing production risk and improving economic performance, and most of its definitions describe it as a bridge between geology and mineral processing [1]. Sensor-based ore sorting separates coarse ore before fine comminution. It can reject waste, divert material to different process lines or make concentrates at several points of a flowsheet, and its review opens by noting that deposit qualities are constantly decreasing [2]. Both move the measurement of the ore ahead of the circuit. In the editor's view this is

the most direct answer to the first problem of this volume. The scale of Chapter 2 weighs a total and never a blend, and every chapter after it learns of a harder or leaner ore from its effect.

> ⧖ HISTORY **2014 to 2022: The Instruments Arrive, the Assay Stays**
>
> The second history box of Chapter 1 followed instruments for the hidden plant from 1973 to 2004. In 2014 a thesis still called a measurement taken with the mill stopped the common way to find a mill's filling, and named acoustic, inductive proximity and conductive sensors as the new technologies [3]. A 2019 study watched the flow inside a hydrocyclone, hidden by its walls and the opaque slurry, through electrical impedance tomography at several hundred frames a second [4]. In 2022 a study found online elemental analyzers widely available yet under-used in mineral processing, where laboratory samples still guided control [5]. In the editor's reading the instruments have arrived, and in most loops of this volume the assay still decides.

Ore properties inferred from operating data. Where the ore is not measured ahead, it is increasingly inferred from the plant. A long short-term memory network trained on a year of operating data from each of two semi-autogenous mills forecast an operational relative hardness with accuracies above 80 % a short time ahead [6]. Chapter 3 gives the caution that goes with such estimates. Its toe, inferred from the bearing pressure with a declared ball charge, fell while the true toe rose, and a setpoint that tracked it passed the graded shift on none of eight noise seeds. In the editor's view an inferred ore property is a measurement with a model inside it, and it needs the statement of what it cannot see that Section 1.4.5 asks of a soft sensor.

More from the assays a plant already has. Froth image analysis has been used to identify operating regimes, to characterize froths and to build inferential sensors for chemical species in the froth, and it has been joined with multivariate statistical methods to monitor industrial platinum group metal flotation plants [7]. Data reconciliation that adds the stoichiometry of the minerals to the material balance estimates balanced mineral compositions, assays and flow rates, and the estimated mineral composition improves the evaluation of process performance [8]. In the editor's view such estimates will sit increasingly inside the quality loops, and Chapter 7 shows what they must still answer. Between two assays an estimate is only as good as its update rule, and there the update gain, not the model, decided whether the outer loop helped.

Economics at the top of the structure. A top-down plant-wide structure for a grinding circuit defines its economic cost with reference to the final product of the larger plant, and uses the mill's grind curves to set the setpoints of its economic controlled variables [9]. The flotation depth, the refining utilization and the thickener density of this volume were each worth more as a chosen operating point than as a tighter loop. In the editor's view the layer that chooses setpoints from metal prices and reagent costs will matter more as ore grades fall, and it needs the protective selectors below it to keep their place.

Learning methods, used with care. A review of machine learning in mineral processing classifies published work by the nature of its data (synthetic, laboratory or industrial), its level of success and its area. It notes that practitioners may be wary of techniques that overpromise, and proposes work on data collection, technique comparison, industrial participation and cost-benefit analysis [10]. Reinforcement learning has tuned the PID controllers of a multivariable grinding circuit on a nonlinear model verified with industrial data, with the agent's actions constrained to limit the risk of instability. In simulation the tuned controllers tracked setpoints better than manually tuned ones [11]. In the editor's view tuning will increasingly be proposed by software, and Chapter 3 shows the test such a proposal needs. Its recommended tuning won on a setpoint step and lost on the plant's own disturbances.

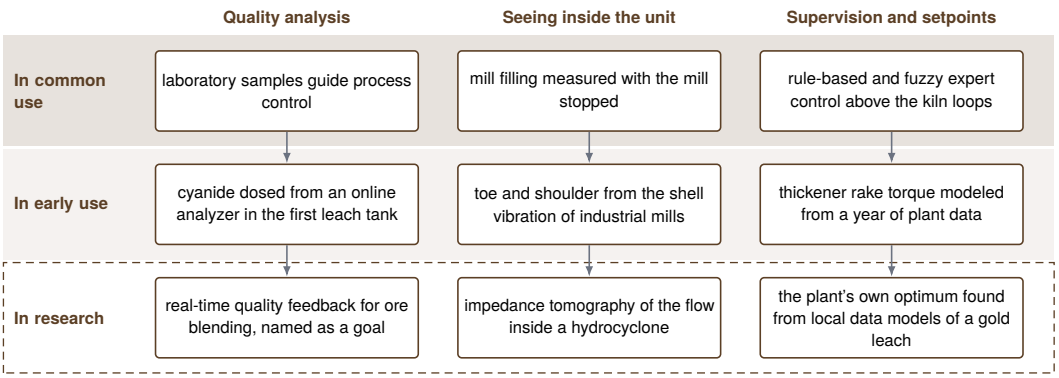


Figure 12.3 Developments in three threads. The top lane holds what the cited source calls standard or common practice, the middle lane what it reports in service at some sites or on plant data, and the bottom lane what it reports from simulation, laboratory work or as a need, in the editor's reading. Sources by thread: quality analysis [5, 12, 18], seeing inside the unit [3, 4] and supervision and setpoints [13, 16, 19]. Notice that two of the three threads move the measurement into the stream or the unit.

Analyzers move onto the stream. Cyanide in a gold leach is typically controlled from a measurement in the first leach tank, taken by hand every 2 to 12 hours or by an online instrument every 5 to 30 minutes, and control from the online measurement has been reported to cut cyanide consumption by at least five percent [12]. In cement plants, online analyzers let operators determine the chemical composition of the raw materials at once [13]. In the editor's view such analyzers are the first step toward the measurement that gap 1 of Section 12.5 asks for. Chapter 7 already reads its cyanide this way, and it is the tailings, sampled every four hours and reported six hours later, that still set the pace of its outer loop.

Cells and underflows watched continuously. Open and short circuits commonly occur in copper refining, and they have been found by walking over the cells with a hand-held gaussmeter. A patented system instead senses the magnetic field around the conductor that feeds each cathode and lets a processor find open circuits as well as shorts [14]. Rheometry has seen only limited use in process monitoring, and a 2019 study reported an automated capillary rheometer that takes a side stream of a thickener underflow, answering the need for a unit that measures a stream's flow behavior continuously and unattended [15]. In the editor's view both read a variable that this volume's loops leave to inference: the nodule that shorts the cell of Chapter 8 and the underflow that loads the rake of Chapter 6.

Protection written into the plan. Published work on paste thickeners has added a rake torque constraint to a predictive controller to prevent rake lifting and pumping problems, and has sampled the near future more finely than the distant one [16]. In rare earth extraction, where practice still rests on offline analysis, manual adjustment and experience, an outer predictive controller over inner PID loops has been proposed, with an event trigger that updates the controller only when the state of the process calls for it [17]. In the editor's view Chapter 6 shows what such a controller must still be compared against: a selector whose reset law, fixed twice, moved the comparison each time. Figure 12.3 sets the developments of this section 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.

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 every number to a run. Seven of the ten papers have a referee report. Five of these cover only the material added in the predictive-control rounds, and the rare earth report calls itself a self-review. The hydrocyclone, refining cell and lithium papers have none. Defects went into a log for each app, and the ten logs share no field. Eight give some entries a severity, 47 entries in all, and two give none, so this section reports cases rather than totals by kind. For this book one agent per chapter reran the numbers its exercises use and checked the app's helpers, completion steps and figures. The editor set the scope of each round and the rules on what each agent could change, and statements that runs contradicted were left for the editor to decide.

A toe that read the truth. The mill's toe tracking computed its setpoint from the run's true rock hold-up, ball charge and liner wear, which no instrument reports. Two mills with identical readings, one carrying two more points of ball filling and the other the same 40.3 t as rock, received setpoints 3.19° apart in one scan. With that knowledge toe tracking passed the graded shift. A technical review of the added predictive controller asked why toe tracking passed while the MPC, which had to estimate the same toe, did not. The fixed app aims at the toe its filter infers from the bearing pressure with the declared ball charge. No impact tuning then passed with toe tracking, the challenge's suggested settings became the slower load loop, and every toe-tracking number in the chapter moved. The volume's lesson that an inferred toe cannot see a lost ball charge exists because the leak was closed.

A soft sensor built on the plant's own parameters. The gold circuit's soft sensor, and the MPC that shared it, read the leach kinetics, the refractory gold and five more constants live from the run's parameters. Raising the leach rate constant by 20 % moved the estimate by −0.0181 g/t with every measurement fixed, and an edit at 30 h moved the cyanide setpoint at 30.11 h, before any assay. A truth audit found it. The constants now sit at the declared values, so an edit of the plant is model mismatch that the laboratory has to find, which is the problem the bias update exists to solve. No run moved. The review that followed found that the truth tests missed a path by which a new estimator could reach the plant's state, and five mutants were added to catch it.

A thirteen-minute constant that kept coming back. The kiln's texts gave the burning zone's answer to coal a 13 min time constant. That is the burning zone node's own constant. With the air ratioed, a coal step reaches 63 % of its change only after 30.1 min, because the zones upstream warm behind it. The technical review of the added MPC reported the claim, the paper stage timed the step, and the paper and the app's scenario texts were corrected. The figure had already reached Chapter 1, which the orchestrating agent corrected after the kiln chapter's report, and the kiln chapter's agent found one more line of the app's challenge text that still called the kiln a thirteen-minute plant.

A cause tested by removing it. The flotation paper blamed its limit cycle on the interaction between the grade loop and the froth depth, and called the grade tuning defensible for that loop alone. The book's agent held the froth depth exactly at setpoint at every physics step and ran the case again. The grade loop still cycled, from 0.74 to 8.80 % copper. The chapter now says the cycle belongs to the grade loop, and the series made it a rule to correct a cause that a run contradicts when no number has to change. Nine explanations in six chapters of this volume were rewritten under that rule or its predecessor for summaries. The hydrocyclone paper, for one, called a tuning unnecessarily sluggish that the measured plant gain makes 2.7 times more aggressive than intended.

A comparison that moved with the reset. The thickener's external reset was examined three times. A technical review found that the losing controllers settled at twice the offset the texts stated.

A later fix made the reset instant, and the book's agent then found that the instant reset lets two losing controllers take the pump in turn, which drove it from 48 to 100 % in six minutes. The app owner replaced the instant reset with finite back-calculation, and the selector now peaks at 85.45 % against 75.77 % for the MPC, with density IAE 1719 against 1777 (Table 6.7). Part of a predictive controller's measured advantage may belong to the structure it was compared with.

A factor of one hundred. The first mill model divided the liner slip by a hundred once instead of twice. At 65 % wear the effective speed was −277.9 % of critical and the power exactly zero, and the structural test suite passed, because nothing was undefined and no actuator left its bounds. The paper agent found it while writing the paper, when the worn-liner scenario diverged.

The book round. The ten chapter records list 158 rows of numeric verification, each a stated value or group of values set beside its reproduction. Seven records report no drift. The hydrocyclone record found a printed 2.7 that runs at 2.77 %, the refining record three roundings, and the thickener record numbers that the app owner's repair later regenerated. The others were kept as printed. Four papers named a loop's action opposite to the convention of Section 1.6.2. Four helpers set or promised what the actuators could not deliver. The cyanide helper wrote 633 kg/h to a pump that stops at 500, the kiln helper wrote 20 t/h of coal to a feeder that stops at 18, the lithium helper claimed three setpoints met when one was out of reach, and the mill's feed helper, pressed at the start of an overloaded case, brought the overload alarm at 4.3 instead of 22.7 min. Chapter 1 was written from the converted chapters before their agents had rerun them, and five of its statements were corrected from their reports. It had called the mill's impact angle inferred where a shell microphone measures it, given the lithium cascade a factor of 93 where the chapter measures 52 to 93, estimated about 25 scrub stages where the rare earth chapter uses 20, stated a gold peak tailings result more widely than the chapter shows, and carried the kiln's thirteen minutes.

Every error above was caught by a check that could fail: a hidden value moved, a step timed, a cause removed and the run repeated, a comparison rerun after a fix, a scenario measured for the paper. None was caught by rereading the text that contained it. Four working rules follow. Give a controller and a soft sensor only what the plant measures. Test a cause by removing it. Rerun a comparison after every fix to either side of it. Trace a number to its source before writing it a second time.

12.5 Gaps in Mineral 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. *Measuring the ore as it arrives.* The scale of Chapter 2 reads the total, and a feeder that lost 20 % of its gain left the blend wrong under every structure until a bump test found it. Hardness, grade and leachability reach the mill, the cells and the leach unmeasured. An instrument that reports them on the time scale of the circuit is not available for most ores. A 2022 study found online elemental analyzers widely available but under-used, and read copper in a suspension with a relative error of 6 % at a processing plant [5]. A 2026 review of ore blending names data from many sources that are hard to join, models too slow for real-time control and poor adaptability across ore types as the barriers to its industrial use [18]. It is hard because hardness and leachability are not elements that an analyzer reads. Progress would be a reading of the arriving ore fast enough for a blend or feedforward loop to close on.

2. *Soft sensors that state what they cannot see.* The gold soft sensor of Chapter 7 removed six hours of dead time for the feed grade and none for the leachability. Data-driven models are used

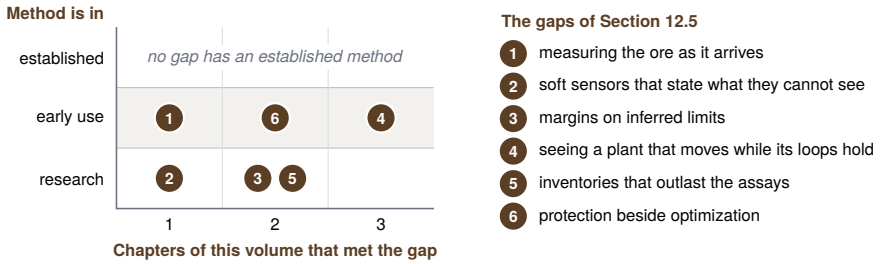


Figure 12.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: one for gap 1 (Chapter 2), one for gap 2 (7), two for gap 3 (3 and 8), three for gap 4 (3, 4 and 10), two for gap 5 (9 and 10) and two for gap 6 (3 and 6). Up is the maturity of the method available, 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 a method in early use.

extensively in industrial plants because an accurate mechanistic model of a gold leach is hard to build [19], and in rare earth extraction a soft sensor has estimated the content of a monitoring stage from flows and feed analyses recorded every 30 min, where the laboratory reports daily [22]. It is hard because a model cannot report a disturbance it was never given. An estimate needs, delivered with it, the list of disturbances it is blind to, and a rule that sets its bias update from the report interval and the assay noise.

3. *Margins on inferred limits.* The refining cell of Chapter 8 held a believed 59 % at a true 73.74, and the inferred toe of Chapter 3 moved the wrong way. Experiments on copper nodules suggested a threshold size above which a nodule leads to a short, and found nodules forming without the mold release agent seen in others, for reasons not yet known [23]. It is hard because the limit moves with what the ore and the anodes bring, and an estimate of it is checked only when it fails. A routine method that sets the back-off of Equation (1.9) from data, and revises it as the ore changes, is still a matter of judgment.

4. *Seeing a plant that moves while its loops hold.* The ball charge of Chapter 3, the electrode bias of Chapter 10 and the roping of Chapter 4 each left every loop on setpoint. The evidence was in an actuator's travel, a reagent per tonne or an unmeasured density. Monitoring that reads manipulated variables and consumptions against the ore, across a whole circuit, is not routine. The Open challenge box after this list takes the gap further.

5. *Inventories that outlast the assays.* The lithium train of Chapter 10 took six hours to settle after a grade change, and a leaner feed drained the rare earth cascade's inventory for days before the purity fell (Chapter 9). A rare earth extraction process has dozens or even hundreds of tanks, and models of it seldom consider the delays between them or take them as constant [24]. It is hard because a change shows in the product only after it has passed through the inventory. Design rules for loops whose plant settles over many report intervals, and for handing such changes to production planning, are thin.

6. *Protection beside optimization.* The comparison of Chapter 6 with an MPC moved with each of two fixes of its reset law, and the MPC of Chapter 3 raised a motor overload the loops did not. On one operating thickener the flocculant pump stops when the bed reaches a set level and density, as a safety feature [25], and a study of thickeners at several mines named control instrumentation that was not adequate or not working properly among the causes of their fluctuations [26]. It is hard because a protective layer is tested only when a limit is near. A tested implementation of selectors and external reset that an optimizer must sit above, with a standard check that protection keeps its authority, would let a comparison measure prediction and not the plumbing.

⚠ OPEN CHALLENGE Seeing a Plant That Moves While Its Loops Hold

Gap 4 is the gap most chapters met, and its hidden changes sit where instruments struggle. A monitoring device inside the shell of a running mill is virtually impossible, because the tumbling charge is so severe [27]. The flow inside a hydrocyclone is difficult or impossible to observe directly [4]. The voltage of a glass pH electrode decays with age, and the electrode is calibrated against two buffer solutions [28]. A bias of that kind, 0.70 low, raised the reagent of Chapter 10 by 24 % while every loop held. The change is hard to see because each loop absorbs it. Progress would be monitoring that compares actuator travel and reagent use with the ore across a circuit and warns before the assay does.

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 [29]. 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 [30]. 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 [31]. 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 [32]. 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 loop measures its own variable. The variable that decided the result was often one that no instrument read.
- A soft sensor removes the assay's dead time only for the disturbances among its inputs, and a loop on a late measurement is tuned to the plant's response and the report interval.
- The structures that worked moved the actuator from something measured before it reached the product, and the late assay trimmed what the fast path could not see.
- Where a loop was set often decided more than how it was tuned.
- Tuning for a step response misled on plants whose disturbances are slow and whose measurements are noisy.

References

[1] V. Lishchuk, P.-H. Koch, Y. Ghorbani, and A. R. Butcher, "Towards integrated geometallurgical approach: Critical review of current practices and future trends," *Minerals Engineering*, vol. 145, p. 106072, 2020.

[2] C. Robben and H. Wotruba, "Sensor-based ore sorting technology in mining—past, present and future," *Minerals*, vol. 9, no. 9, p. 523, 2019.

[3] S. B. Donkor, "On-line sensors for measuring the total ball and charge level in tumbling mills," MSc (Eng) thesis, University of Cape Town, Centre for Minerals Research, 2014, hdl.handle.net/11427/13225.

[4] C. J. Oates, J. Cockayne, R. G. Aykroyd, and M. Girolami, "Bayesian probabilistic numerical methods in time-dependent state estimation for industrial hydrocyclone equipment," *Journal of the American Statistical Association*, vol. 114, no. 528, pp. 1518–1531, 2019.

[5] L. Rybok, T. Kowoll, A. Thomas, A. Nad, and C. Bachmann, "Online elemental analysis for process control in the mineral processing industry," *TOS Forum*, vol. 2022, no. 11, p. 131, 2022.

[6] S. Avalos, W. Kracht, and J. Ortiz, "An LSTM approach for SAG mill operational relative-hardness prediction," *Minerals*, vol. 10, no. 9, p. 734, 2020.

[7] C. Aldrich and X. Liu, "Monitoring of flotation systems by use of multivariate froth image analysis," *Minerals*, vol. 11, no. 7, p. 683, 2021.

[8] M. Sadeghi, D. Hodouin, and C. Bazin, "Mineral processing plant data reconciliation including mineral mass balance constraints," *Minerals Engineering*, vol. 123, pp. 117–127, 2018.

[9] J. D. le Roux, S. Skogestad, and I. K. Craig, "Plant-wide control of grinding mill circuits: Top-down analysis," *IFAC-PapersOnLine*, vol. 49, no. 20, pp. 72–77, 2016.

[10] J. T. McCoy and L. Auret, "Machine learning applications in minerals processing: A review," *Minerals Engineering*, vol. 132, pp. 95–109, 2019.

[11] J. A. van Niekerk, J. D. le Roux, and I. K. Craig, "Reinforcement learning based automatic tuning of PID controllers in multivariable grinding mill circuits," *Control Engineering Practice*, vol. 165, p. 106522, 2025.

[12] W. A. M. van der Merwe, "The development of an online amperometric technique to measure free and WAD cyanide in gold plant leach liquors and effluent streams," M.Eng. dissertation, University of Cape Town, 2012.

[13] E. Worrell, C. Galitsky, and L. Price, "Energy efficiency improvement opportunities for the cement industry," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-72E, 2008.

[14] E. Y. You, A. Berges-Dreyfous, D. B. George, D. Kim, and K. Zhou, "Monitoring electrolytic cell currents," U.S. Patent 7,445,696 B2, application filed March 17, 2005, 2008.

[15] A. G. Chryss, A. Mönch, and K. Constanti-Carey, "Online rheology monitoring of a thickener underflow," in *Paste 2019: Proceedings of the 22nd International Conference on Paste, Thickened and Filtered Tailings*. Perth: Australian Centre for Geomechanics, 2019, pp. 495–504.

[16] J. Ruuska, E. Ruhanen, J. Kauppi, S. Kauvosaari, and M. Kosonen, "Multivariate linear regression model of paste thickener," in *Proceedings of the 61st SIMS Conference on Simulation and Modelling*, ser. Linköping Electronic Conference Proceedings, vol. 176, 2021, pp. 160–164.

[17] Z. Zhang, J. Zhu, C. Xiong, H. Yang, and F. Nie, "Constrained model predictive control based on event triggering for the rare earth extraction process," *ACS Omega*, vol. 8, no. 44, pp. 41 943–41 952, 2023.

[18] Z. Cai, H. Luo, and W. Yang, "The research progress and application status of ore blending technology: A review," *Physicochemical Problems of Mineral Processing*, vol. 62, no. 4, p. 226165, 2026.

[19] Q. Liu, D. He, Z. Wang, and J. Shi, "A data-driven iterative optimization compensation method based on PJIT-PLS for gold cyanidation leaching process," *IEEE Access*, vol. 7, pp. 138 683–138 693, 2019.

[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. Dai, R. Lu, P. Chen, and H. Yang, "Novel dual residual-enhanced deep bidirectional LSTM network for soft sensing of rare earth component content," *Sensors*, vol. 26, no. 10, p. 3152, 2026.

[23] K. Adachi, Y. Nakai, S. Mitsuno, M. Miyamoto, A. Kitada, K. Fukami, and K. Murase, "The mechanism of nodular growth in copper electrorefining," *Journal of MMIJ*, vol. 136, no. 2, pp. 8–13, 2020, in Japanese, with an English abstract.

[24] R. Lu, H. Liu, H. Yang, J. Zhu, and W. Dai, "Multi-delay identification of rare earth extraction process based on improved time-correlation analysis," *Sensors*, vol. 23, no. 3, p. 1102, 2023.

[25] M. Cooks, D. Mostert, L. Coetzee, and D. J. van Schalkwyk, "Thickener experiences at Langer Heinrich mine," in *Paste 2012: Proceedings of the 15th International Seminar on Paste and Thickened Tailings*. Perth: Australian Centre for Geomechanics, 2012, pp. 141–153.

[26] B. Pirouz, S. Javadi, and K. Seddon, "Thickener performance variability: underflow solids concentration and flowrate," in *Paste 2017: Proceedings of the 20th International Seminar on Paste and Thickened Tailings*. Beijing: University of Science and Technology Beijing, 2017, pp. 29–40.

[27] R. Rajamani, J. Delgadillo, and V. Duriseti, "Online SAG mill pulse measurement and optimization," University of Utah, Metallurgical Engineering Department, Semi-annual technical progress report DE-FC26-04NT42088, 2006, u.S. Department of Energy, Office of Scientific and Technical Information.

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

[29] 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.

[30] 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.

[31] 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.

[32] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

12.6 Exercises

12.1 Discussion · 8 min

An agent proposes to set the bias update gain of the gold circuit in Chapter 7 to one, "so that the soft sensor uses every assay in full, because the assay is the truth". Using Chapter 7 and Section 1.4.5, explain what the proposal overlooks, and say what evidence you would ask for before accepting any gain.

12.2 Discussion · 8 min

An agent adds a supervisory setpoint to a semi-autogenous mill. The impact angle is aimed at a toe that a model computes, and the agent reports that the graded shift of Chapter 3 now passes. Using Chapter 3 and Section 12.4.1, list what you would check before accepting the result, and say which check found the defect in the first app.

12.3 Discussion · 7 min

An agent recommends adding an inner loop under every quality loop of a circuit, because “a cascade rejects disturbances before they reach the product”. Using Chapter 4 and Chapter 10, say for which disturbances the claim held, for which it failed and why, and what test you would ask for.

12.4 Discussion · 8 min

A steady-state sweep by an agent finds that the graded crop of Chapter 8 is worth most at 62 % utilization and 65 °C, \$35 661, and it recommends that setting. Using Chapter 8 and Equation (1.9), discuss the recommendation.

12.5 Discussion · 8 min

An agent reports that the economic MPC of Chapter 5 earns \$1050 /h more than the shipped loops on scenario 11 and recommends an MPC for every cell. Using Section 12.2 and the chapters it cites, list the comparisons you would ask for, and say what each did in this volume.

12.6 Discussion · 7 min

A colleague argues that Chapter 1 could quote the papers without rerunning anything, because every paper was refereed. Using Section 12.4.1, answer the argument, and propose a routine that would have caught the kiln’s thirteen-minute constant at its first use.

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) 83.3 kg/m, 42 s. (b) 480 t/h at once, and 52.5 s.

1.3 (a) $g = 0.595$, 555.6 t/h. (b) 58.3 t/h, 19.4 % of the product. (c) 9.9 min.

1.5 (a) 86.98 % recovery, 2.37 % mass pull. (b) 83.71 %.

1.7 18.18 % in one tank and 6.40 % in three tanks of the same total residence.

1.9 (a) 2 to 10 h, 6 h on average. (b) 11.5 h, and 57.7 % of the noise spread.

Chapter 2. Ore Blending Through Speed-Dependent Delays

2.1 (d). The scale reads loading times speed, and the ore already on the belt keeps its 125 kg/m: $3.6 \times 1.5 \times 125 = 675\,\mathrm{t/h}$.

2.3 (b). The loading bound just past C's chute would be exceeded by C's extra ore stacked on B's ore still in transit.

2.5 (a) 45, 56.25 and 45 %. (b) 901 t/h at 12.1:60.6:27.3. (c) 846.7 t/h in both cases, set by the belt.

2.7 (a) 61, 30 and 11 s (± 1 s). (b) 940, 1020 and 1080 t/h; 4, 8 and 6 t/h per percent. (c) The drive lag, the span filling and the scale filter add about two seconds to the travel.

2.9 (a) 8.07 t, 516 s, 193 t, not met. (b) 847 t/h, 0 s, 216 t, met. (c) 18 s, met. Under the temporary recipe C's increase lands on B's last ore.

2.11 It is a product decision made in advance: temporary recipe when the mill tolerates a B-free blend for a few minutes, preserve when off-recipe ore does more harm than lost throughput.

Chapter 3. Semi-Autogenous Mill Impact-Zone Control

3.1 (c). Equation (3.19) makes the bearing pressure affine in $M_b + M_c + M_f + M_w$.

3.3 (a). The bearing pressure cannot tell steel from rock, and the toe of Equation (3.4) rises with the filling.

3.5 (a) $J^{*} = 0.4315$, 249.7 t of rock. (b) 0.534, because the charge mass grows with J from a fixed ball and water mass. (c) 14.8 t of rock and 183 kPa.

3.7 (a) 25.16, 35.99, 62.20 and 77.53°. (b) 3.61 and 5.11 °/%, against local gains of 3.456 at the design point and 4.393 to 5.276 between 86 and 90 %. (c) 31.69 against 34.64°: the meter averages over 0.40 min.

3.9 (a) 0.6° and 4.0°, a failure of the ±3° toe clause, and 147.4 t at 0.783 h. (b) 0.6° and 2.4°, a pass, and 135.1 t at 0.918 h, for 1.9 % less ore.

3.11 (a) 40.3 t and 498 kPa. (b) $J = 0.218$, and the toe falls by about 3.4°. (c) Raise the load setpoint by about 498 kPa at the same time, to 5985 kPa on the control's 5 kPa grid.

Chapter 4. Hydrocyclone Classification Cascade Control

4.1 (c). The pressure transmitter PT reads it. The other three are computed by the model or set as a disturbance.

4.3 (a). A cascade removes only what enters its inner measurement, and the coarse feed reaches the product through the recycle and the cut, not through the pressure.

4.5 (a) 6.0 m^3 and 0.26 min. (b) 0.139. (c) 1.37 min.

4.7 (a) 3.2 % min and 0.1 min. (b) 154.9 % min and 33.3 min. (c) About 49.

4.9 (a) Fails: pressure 13.3 kPa, level 13.6 %, 0 of 4; 964.9 % min. (b) Passes: 3.1 kPa, 3.4 %, 4 of 4, underflow at most 69.5 %; 149.4 % min; 14.9 min. (c) Fails with 3 of 4; 284.0 % min; the feed step at 8 min.

4.11 (a) 1.26 and 2.02. (b) 35.7 % more power for a cut 11.3 % finer.

Chapter 5. Froth Flotation and an Economic Setpoint

5.1 (b). The concentrate analyzer AT reads it, 2.0 min late. The other three are model quantities.

5.3 (d). The average value must exceed \$9500 /h, with the depth within 0.05 m of setpoint and the grade above 4.5 % copper.

5.5 (a) 20 m^2, 0.167 min/m, 48.1/s. (b) 16.7 % and 0.149 m, a thinning of 0.151 m. (c) The damper lag and the liquid that leaves while the holdup rises.

5.7 (a) −0.111 m, 3.67 %, IAE 0.294. (b) −0.127 m, 3.21 %, IAE 0.727. (c) About 1.14: the holdup moves the interface faster than any valve can answer.

5.9 (a) Fails: 0.118 m, \$8282 /h, 7.0 min. (b) Passes: 0.039 m, \$9890 /h, 0.0 min. (c) Fails: 0.037 m, \$8580 /h; about \$1310 /h.

5.11 (a) \$50 /h per g/t, and the extra collector costs only \$50 /h of the \$3001 /h lost. (b) The tailings assay improves at every step; the dose belongs on the margin, near 30 g/t here.

Chapter 6. Thickener Override Selector Control

6.1 (a). The selector must pass the demand that moves the pump in the direction that relieves the limit, and for torque and bed that is more pump.

6.3 (c). The Langmuir curve has saturated, and above 2.2 times the reference dose the overdose penalty forces a deeper bed for the same density.

6.5 (a) 85.4 %. (b) 2.9 and 16.1 points of torque per point of solids, a ratio of 5.5. (c) 66.41 %, below the model's value, which is the safe side.

6.7 (a) 62.81 %, 4.97 h, 1.691 m, 59.6 %. (b) 64.87 %, 6.08 h, 1.960 m, 74.2 %, constant. (c) 0.28 against 0.33.

6.9 (a) 3.02 %, 104.7 %, 351 min, 3.28 m: fails on density, torque and clarity. (b) 1.30 %, 65.9 %, 43 min, 2.57 m: passes. (c) 1.67 %, 71.5 %, 94 min, 2.72 m: fails on density and clarity.

6.11 (a) 4.1 and 10.4 points. (b) $75 + 10.4 + 3.6 = 89.0$, so 3.0 points remain. At 88 % the same overshoot and noise reach the trip.

Chapter 7. Gold Carbon-in-Leach Inferential Control

7.1 (d). The estimate reads only measured inputs, all held on setpoint, so it waits for the bias update from the first assay that sees the slower ore.

7.3 (b). A loop holds the point it measures. The last of the gold dissolves in tank five, which runs short of cyanide.

7.5 (a) \$31.2 per hour. (b) \$5.24 per hour, so it must lower the tailings by about 0.17×10^{-3} g/t. (c) The step to 200 mg/L pays (0.20×10^{-3} g/t per mg/L), the step to 260 mg/L does not (0.08×10^{-3} g/t per mg/L).

7.7 (a) 20.3 h and 10.3 h, about 10.3 and 0.3 h after the step. (b) 428 kg/h and 296 kg/h. (c) About 150 mg/L at 40 h, settling near 140 mg/L.

7.9 (a) Not met on the tailings clause, 0.433 g/t. (b) Met, 0.525 kg/t, 21.6 h above 0.40 g/t, peak 0.452 g/t. (c) Met, 0.535 kg/t, never above 0.40 g/t.

7.11 (a) 39 %. (b) Err high: the penalty below the optimum is about forty times the prize at it, and in scenario 8 tank five falls below the 9.8 floor at 15.0 h and reaches 9.50 while the head tank stays within 0.04 of 10.40.

Chapter 8. Copper Electrorefining at an Inferred Limit

8.1 (d). The limiting current density is computed from a handbook diffusivity, a flow correlation and the assay, so u is an estimate.

8.3 (b). The loop is direct acting: a low assay cuts the flow, and the closed valve costs more of j_L than the leaner copper does.

8.5 (a) 0.800 and 55.0 %, 0.889 and 56.5 %, 0.909 and 45.7 %. (b) A ratio of 0.50, and about 32 and 63 h.

8.7 (a) 3.18 mm, 6281 kg, \$34 846, downgraded at 152.6 h. (b) 1.82 mm and Grade A, 6213 kg, \$53 167. (c) 1.1 % more copper and 34 % less value: the crop is graded whole at harvest.

8.9 (a) Fails with 4042 kg, 24.49 mm, 0.383 V, 2.85 °C, \$22 396. (b) Passes with 3857 kg, 0.33 mm, 0.368 V, 1.17 °C, \$32 945. (c) Passes with 4069 kg, 1.99 mm, \$34 758: 5.5 % more value for 0.01 mm of margin.

8.11 (a) 249 and 270 kW h/t. (b) \$0.023 per kilogram, 0.27 % of the copper value.

Chapter 9. Rare Earth Solvent Extraction Control

9.1 (c). In Equation (9.4) the feed's neodymium falls at the step while the raffinate's has not yet changed, so the same loss is a larger share of a smaller feed.

9.3 (a). Action $+1$ is reverse acting for a measurement that rises with its manipulated variable, and purity falls as O/A rises, so the loop is direct acting with $\sigma = -1$.

9.5 (a) 0.70 kmol/h of feed, 108.6 h. (b) 0.0525 kmol/h and 718 h. (c) The scrub step leaves the neodymium fed and drawn equal, and the purity moves 63 % in 25.3 h.

9.7 (a) 1.015 to 1.201, about 17 %. (b) 1.1665, 1.1449 and 99.74 %. (c) The purity answers the scrub ratio with a 63 % time of about a day, and an integral time of 4 h adds to the ratio for six reports before the first move shows.

9.9 (a) Not met, 2 of 25 short, 99.47 %, \$2.85/kg. (b) Met, 0 of 25, 99.55 %, \$2.93/kg. (c) 6 h fails with 5 of 25 short (99.49 %), 12 h passes (99.52 %, \$2.89/kg): the integral time decides.

9.11 The best within $O/A \leq 28$ and 18 % saponification is 99.866 %. It is reached at O/A 33 or at 21 % saponification, which need more settler capacity or load the extractant past the guideline, where the model has no third phase.

Chapter 10. Lithium Refinery: Three Units in Series

10.1 (b). The inner loop measures fuel flow, which a header pressure loss changes and a heating value loss does not.

10.3 (d). The loop holds the reading, so the true pH settles near 13.20, and every reported figure improves while lime use rises by about a quarter.

10.5 (a) $K_c = 0.41$, $\tau_I = 14.9$ min. (b) $K_c = 13.0$, $\tau_I = 45.1$ min; the shipped integral is shorter. (c) $K_c \approx 5770$, a 577 L/min move for 0.1 pH, against a 900 L/min span.

10.7 (a) Kiln $+1.26$ °C at about 5 min, pH $+0.163$ at about 22 min, product -0.21 kg Li/m^3 at about 74 min. (b) Production falls 27.0 %, feed lithium 27.3 %. (c) The disturbance travels with the material, and the loss is in the feed, not in any loop.

10.9 (a) Fails all three bands, kiln ± 24.7, pH ± 0.36, product ± 1.98. (b) Passes, ± 9.3, ± 0.13, ± 0.79. (c) Fails on pH, ± 0.31; the kiln is unchanged and the product moves slightly, because influence runs only downstream.

10.11 The product sits on setpoint while the duty absorbs the scaling, and it goes off specification before the duty saturates. Trend the duty against its clean value.

Chapter 11. Rotary Cement Kiln Ratio Feedforward

11.1 (a). The disturbance is the control system's own output, so the air can move with it instead of after the analyzer sees it.

11.3 (c). The analyzer reads oxygen that never passed the flame, so the controller removes real combustion air to bring the reading down.

11.5 (a) 1.100, 1.91 %, 853 ppm. (b) 1.240, 4.06 %. (c) 111.2 t/h, $\lambda = 0.960$, 678 ppm; the run's coal rises, which lifts λ slightly.

11.7 (a) Reducing time 0.8 against 0.0 min; IAE 68.7 against 81.7 °C min. (b) 0.964 against 1.085. (c) 19 %; the lagging air of scenario 4 lets the flame run hotter during the rise.

11.9 (a) Shipped 88.7 min, 1600 kJ/kg, 68.7 kg/t, 2.73 %; suggested 0.0 min, 1721 kJ/kg, 71.4 kg/t, 1.20 %. (b) The shipped run, by starving the flame and making unsound clinker. (c) 11.0 min, a failure; the air structure decides.

11.11 (a) 0.245 %; 14 to 16 °C; 0.26 to 0.31 t/h. (b) Speed is committed to production rate, and it also moves bed depth, heat transfer and coating, which the model does not contain.

Chapter 12. Summary and Outlook

12.1 A good answer notes that each assay is six hours old and carries the laboratory's noise, so a gain of one steps the estimate by that noise and by what the plant did six hours earlier. On the feed grade step it raised the cyanide travel to 5472 against 3356 kg/h at 0.15, and the peak tailings to 0.3158 against 0.3089 g/t. Without noise it still doubled the travel. A gain of zero let the tailings sit above the limit for 172 of 200 hours. It asks for travel and peak over several noise seeds and treats the gain as a filter set by the report interval and the noise.

12.3 A good answer states that a cascade removes only what enters before its inner measurement. Impeller wear fell from 154.9 to 3.2 % min, and a header pressure loss by a factor of 52 to 93. Coarse feed, feed rate, grindability and heating value were not helped. Worn liners and a ninth cyclone were made about three times worse, because restoring the pressure drew dilution water and thinned the feed. It asks for one disturbance at a time under both structures at the same outer tuning, with the secondary variables trended.

12.5 A good answer asks for comparators that receive what the MPC receives. The flotation optimum entered as fixed stations came within $11 /h, and the grade loop alone as shipped already recovered 933 of the 1050. The mill's loops given the breakage estimate came within 0.4 %. The kiln's loops given the target, a floor and a cross-limit beat the plan on free lime. It asks for the same seeds and events, the constraints a plant cares about, actuator travel, and a check of the structure compared with, as the thickener's reset showed.

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.

activated carbon Porous carbon granules that take dissolved gold out of a leach slurry and are later stripped and reused. *Chapter 7*.

alkaline leach Dissolving a metal with a basic solution rather than an acid; the lithium refinery leaches with lime. *Chapter 10*.

analyzer, on-stream An instrument that samples plant streams in turn and reports a size or a composition every few minutes. *Chapter 5*.

anode passivation A film that forms on a refining anode and blocks it, so the cell needs a higher voltage to pass the same current. *Chapter 8*.

apex The small opening at the bottom of a hydrocyclone through which the coarse, heavy product leaves. *Chapter 4*.

assay A chemical analysis of a sample that reports how much of an element it holds, often hours after the sample was taken. *Chapter 1*.

back-calculation A way to stop windup: when a controller output is limited, its integral is pulled toward the output actually applied. *Chapter 3*.

back-off A margin by which a setpoint is kept below a limit, so that measurement or model error cannot carry the plant past it. *Chapter 8*.

belt scale A weigher under a conveyor that multiplies the load on a short section of belt by the belt speed to give a flow. *Chapter 1*.

bias update A correction added to a soft sensor and adjusted each time a laboratory result arrives, so the estimate follows the plant. *Chapter 1*.

bump test A small deliberate step on one input, used to see how much and how fast the plant answers. *Chapter 2*.

burnability How easily a cement raw mix turns into clinker; a hard-burning mix needs a hotter kiln. *Chapter 11*.

carbon-in-leach (CIL) Gold extraction in which activated carbon sits in the leach tanks themselves and picks up the gold as it dissolves. *Chapter 7*.

cascade control Two loops in series, in which the outer loop sets the setpoint of a faster inner loop. *Chapter 1*.

cathode nodule A lump that grows on a refining cathode and can reach across to the anode and short the cell. *Chapter 8*.

circulating load The coarse material a classifier sends back to the mill, given as a share of the fresh feed. *Chapter 1*.

classification Splitting particles by size, or by size and density, into a coarse and a fine stream. *Chapter 4*.

clinker The hard nodules that come out of a cement kiln and are ground with gypsum to make cement. *Chapter 11*.

closed grinding circuit A mill whose coarse product is separated and sent back to it until it is fine enough to leave. *Chapter 1*.

collector A reagent that makes the valuable mineral water-repellent so that it sticks to air bubbles in flotation. *Chapter 5*.

comminution Breaking rock into smaller pieces by crushing and grinding, the largest user of energy in a concentrator. *Chapter 3*.

concentrate The product of a separation that carries most of the valuable mineral in a small share of the mass. *Chapter 1*.

constraint control Control that holds a plant close to a limit, such as a current or a torque, instead of at a fixed target. *Chapter 8*.

critical speed The mill speed at which the charge would stick to the shell and stop tumbling; mills run at a share of it. *Chapter 3*.

crud A dirty layer of solids and emulsion that collects between the two liquids of a solvent extraction settler. *Chapter 9*.

cut size The particle size that a classifier sends to its two products in equal shares. *Chapter 4*.

cyanide, free The cyanide in a leach solution that is not bound to metals and is available to dissolve gold. *Chapter 7*.

cyanide, weak acid dissociable (WAD) Cyanide bound loosely to metals such as copper; it is toxic and is released when the solution is mildly acidic. *Chapter 7*.

dead time The time between a change at a plant input and the first sign of it at the measurement. *Chapter 5*.

density gauge An instrument that reads the solids content of a slurry in a pipe, often by how much radiation the slurry absorbs. *Chapter 6*.

dilution water Water added to a mill or a sump to set how thick the slurry is. *Chapter 4*.

discrete element method (DEM) A computer method that follows every ball and rock in a mill as a separate body to predict how the charge moves. *Chapter 3*.

economic optimum The operating point that earns the most, found by trading the value of one product against the cost of another. *Chapter 1*.

electrorefining Purifying a metal by dissolving impure anodes in an electrolyte and plating pure metal on cathodes with an electric current. *Chapter 8*.

entrainment Waste particles carried into a flotation froth with the water rather than lifted by bubbles. *Chapter 5*.

external reset Feeding a selector's chosen output back to each losing controller so its integral does not wind up while it waits. *Chapter 1*.

extractant The chemical in the organic liquid of solvent extraction that picks up metal ions from the water. *Chapter 9*.

false air Outside air that leaks into a kiln at its seals and dilutes the flue gas without passing through the flame. *Chapter 11*.

feedforward Moving a handle as soon as a disturbance is measured, before the controlled variable has changed. *Chapter 9*.

final control element The device a controller moves to act on the plant, such as a valve, a pump, a feeder or a drive. *Chapter 1*.

flocculant A polymer added to a thickener feed that binds fine particles into larger flocs so they settle faster. *Chapter 6*.

flotation Separating minerals by making the valuable one stick to air bubbles that rise into a froth. *Chapter 1*.

flowsheet A drawing of a plant's units and the streams that join them, from feed to products. *Chapter 1*.

free lime Lime left unreacted in cement clinker; too much shows an underburnt kiln and spoils the cement. *Chapter 11*.

froth depth The thickness of the froth layer on top of a flotation cell, set by the pulp level below it. *Chapter 5*.

gain scheduling Changing a controller's settings with the operating point, because the plant's response changes there. *Chapter 3*.

grade The share of the valuable element in a stream, such as the percent of copper in a concentrate. *Chapter 1*.

grade-recovery curve The trade in flotation: pulling more mineral into the froth raises recovery but lowers the concentrate grade. *Chapter 5*.

hydrocyclone A cone into which slurry is pumped tangentially, so that it whirls and sends coarse particles down and fine ones up. *Chapter 4*.

impact angle The angle on the mill shell at which the falling charge lands, inferred from the sound the impacts make. *Chapter 3*.

impact microphone A microphone beside a mill that listens to the charge hitting the shell, from which the impact angle is computed. *Chapter 3*.

inductively coupled plasma (ICP) A laboratory instrument that heats a sample in a plasma and reads its elements from the light they give off. *Chapter 9*.

inferential control Controlling a variable that no instrument reads, by acting on an estimate computed from what is measured. *Chapter 1*.

instrument bias A steady error in a reading, such as an aged pH electrode that reads every value too low. *Chapter 10*.

integrating process A process, such as a level in a sump, that keeps drifting while inflow and outflow differ. *Chapter 4*.

Kalman filter An estimator that blends a model's prediction with noisy measurements, each weighted by how much it is trusted. *Chapter 3*.

leaching Dissolving a metal out of ore or concentrate with a chemical solution. *Chapter 1*.

limiting current density The highest current per area at which metal ions can reach a cathode; near it, deposits grow rough. *Chapter 8*.

liner wear The slow loss of the steel lifters inside a mill, which changes how the charge is lifted and where it lands. *Chapter 3*.

mass balance Accounting that what enters a unit equals what leaves it plus what accumulates inside. *Chapter 1*.

mill filling The share of a mill's volume taken by balls, rock and slurry; mill power peaks at a middle filling. *Chapter 3*.

mixer-settler A solvent extraction stage in which two liquids are stirred together and then left to separate by density. *Chapter 9*.

model predictive control (MPC) A controller that uses a model to plan several future moves against limits and applies the first, then plans again. *Chapter 2*.

moving horizon estimation Estimating a plant's state by fitting a model to a window of recent measurements, repeated as the window moves. *Chapter 7*.

ore blending Drawing ore from several piles in set

proportions so the mill receives a steady mixture. *Chapter 2*.

overflow The fine or clear stream that leaves the top of a hydrocyclone or a thickener. *Chapter 1*.

override control A protective loop that takes over a shared handle through a selector when a limit is approached. *Chapter 1*.

oxygen trim A slow loop that adjusts the air to fuel ratio from an oxygen reading in the flue gas. *Chapter 11*.

pH A measure of how acidic or basic a solution is; low numbers are acidic, high numbers basic. *Chapter 7*.

PID control The common feedback law that acts on the error, its sum over time and its rate of change. *Chapter 2*.

rake torque The twisting load on a thickener's rake drive, which rises as the settled bed grows deeper and denser. *Chapter 6*.

rare earth elements A group of chemically similar metals, such as neodymium and praseodymium, used in strong permanent magnets. *Chapter 9*.

ratio control Moving one flow in a fixed proportion to another, measured flow. *Chapter 1*.

recovery The share of the valuable element in the feed that ends up in the product. *Chapter 1*.

rectifier The power supply that turns alternating current into the direct current a refining tankhouse needs. *Chapter 8*.

roping A hydrocyclone fault in which the underflow leaves as a thick rope and coarse particles escape to the overflow. *Chapter 4*.

rotary kiln A long, slightly sloped steel tube lined with brick that turns slowly while a flame heats the solids moving through it. *Chapter 10*.

saponification Partly neutralizing the extractant with caustic soda before it meets the feed, which sets how strongly it extracts. *Chapter 9*.

selector A block that passes the higher or the lower of two or more controller outputs to a shared actuator. *Chapter 1*.

semi-autogenous mill (SAG) A large rotating drum that grinds ore mostly by the ore itself, helped by a modest load of steel balls. *Chapter 3*.

separation factor How strongly an extractant prefers one metal over another; near one, many stages are needed. *Chapter 1*.

settling velocity The speed at which particles sink through a liquid; it depends on their size, density and flocculation. *Chapter 6*.

slurry A mixture of fine solid particles and water that can be pumped like a liquid. *Chapter 1*.

soft sensor A calculation that estimates an unmeasured variable from measured ones, continuously. *Chapter 1*.

solvent extraction Moving metal ions from a water solution into an immiscible organic liquid, stage by stage, to separate or purify them. *Chapter 1*.

spodumene The hard-rock mineral from which much of the world's lithium comes; it must be heated before acid can leach it. *Chapter 10*.

tailings The waste stream left after the valuable mineral has been taken out. *Chapter 1*.

tankhouse The hall of electrolytic cells in which metal is refined or won by current. *Chapter 8*.

thickener A wide, shallow tank in which solids settle out of a slurry, giving a dense underflow and clear water. *Chapter 1*.

toe of the charge The lowest point at which the tumbling charge meets the mill shell, where falling rock should strike. *Chapter 3*.

transport delay The time material takes to travel from one point to another, such as along a conveyor belt. *Chapter 1*.

trunnion bearing pressure The oil pressure in a mill's bearings, which rises with the weight of everything inside the mill. *Chapter 3*.

two-product formula A formula that computes recovery from three assays of feed, concentrate and tailings. *Chapter 1*.

underflow The coarse or dense stream that leaves the bottom of a hydrocyclone or a thickener. *Chapter 1*.

variable-speed drive An electronic drive that sets the speed of a motor, and so of a feeder, pump or mill. *Chapter 1*.

weighfeeder A short conveyor on load cells that both weighs and delivers a set flow of solids, such as coal to a burner. *Chapter 11*.

X-ray fluorescence (XRF) analyzer An analyzer that reads the elements in a sample from the X-rays they give back when struck by X-rays. *Chapter 5*.

Index of Control Topics

AI agents
  causal claims, 267
  engineering practice, 266
  open problems, 270
  reproduction, 268
  truth audit, 267
  verification record, 267
  while studying, 17
arrival cost, 147
assay
  sampled, 9, 142, 193

back-calculation, 55, 81, 101, 144
back-off, 181
belt scale, 11
belt weigher
  verification, 37
bias update, 9, 102, 142
bump test, 30
bumpless transfer, 15, 55, 101, 240

cascade control, 10, 72, 140
  commissioning, 89, 225
  disturbance entry, 214
  disturbance path, 261
  fuel flow, 225
  inner-loop speed, 81
circulating load, 6, 78
combustion
  excess air, 238
commissioning order, 90
constraint back-off, 62
constraint control, 169
  inferred limit, 167
control layers, 13

data export, 16
data reconciliation, 264
dead time, 100, 140
  dominant, 232
dead-time compensation, 24
  history, 24
discrete element method
  history, 46
disturbance
  entry point, 73

economic optimum, 7, 95
economic setpoint, 261

feasibility supervisor, 29
feedforward, 198
  from a measured feed, 261
  from the controller's own output, 244
feedforward control, 90
  model-based, 55
  ore properties, 161
Fenske relation, 194
final control element, 11
first-order plus dead time
  identification, 217
flotation
  grade and recovery, 7, 97
flow scheme, 4
froth imaging, 264

gain scheduling, 52
geometallurgy, 263

Hammerstein model, 102

inferential control, 138, 140
  constraint, 11
  unmeasured variable, 260
inferential measurement, 43, 45, 167
instrument bias, 223
integrating process, 81
inventory control, 195

Kalman filter, 16, 148
  disturbance states, 126
  input disturbance, 56

laboratory analysis, 9, 142, 187
leaching, 8
limit cycle, 104
loop interaction, 72, 95

machine learning
  mineral processing, 264
manual helpers, 82
mass balance
  from assays, 6
material tracking, 30
measurement

history, 3
mineral processing
control design procedure, 262
open problems, 268
model error, 176
model predictive control
against simpler control, 263
beside a selector, 125
constraints, 265
economic, 101, 145
flotation, 111
in the apps, 15
move blocking, 31, 55, 101, 146
offset-free, 56
soft constraints, 126, 198
successive linearization, 55
transport delay, 31
moving constraint, 43
moving horizon estimation, 16, 147

nodule growth
open problem, 181

offset-free tracking, 102, 145
on-stream analyzer
cyanide, 265
maintenance, 110
online analysis
history, 264
operating point
gain and time constant, 52
ore sorting, 263
override control, 10, 117
voltage interlock, 173
oxygen trim, 241

pairing, 195
pH control
nonlinear gain, 218
pH measurement
glass electrode, 225
PID control
anti-windup, 15, 30, 55, 123, 144
controller action, 218, 240
direct acting, 124, 172, 197
velocity form, 15, 54, 80, 101, 144, 240
PID tuning
SIMC rule, 81, 101
slow disturbances, 262
Ziegler–Nichols rules, 55
plantwide control, 211
plant settling time, 221
positive feedback, 53
process monitoring
electrorefining, 265
open problem, 270

ratio control, 11, 30, 197
air to fuel, 232
cross-limiting, 252
ratio station, 204
real-time optimization
gold leaching, 161
receding horizon, 101
recycle
positive feedback, 78
relative gain array, 100, 195
negative, 195

sampled measurement
laboratory assay, 241
sampled-data control, 197
scenarios, 16
completion rules, 16
scorecard, 16
secant gain, 52
sedimentation
open problems, 133
selector, 10
high select, 117
separation factor, 9, 194
SIMC tuning, 173, 219
simulation engine, 14
simulator framework, 13
soft constraint, 102, 146
soft sensor, 9, 138, 142
bias update, 260
rare earth extraction, 206
solvent extraction, 8
step test, 100
on a nonlinear plant, 54
supervisory control, 55, 98, 161

temperature measurement
pyrometer, 251
thickener, 8
thickener design
history, 119
time-scale separation, 54, 100, 144, 187, 195
toe tracking, 55
transport delay, 5
belt coordinate, 26

speed-dependent, 22
two-point composition control, 187
two-product formula, 7

ultimate gain, 55

variable ore, 11
variable-speed drive, 11

weighfeeder, 252

Index of Applications

activated carbon, 139
  history, 142
alkaline leach, 216
anode passivation, 175
apex, 86

bed pressure, 132
bed residence time, 118
belt conveyor, 22
belt scale, 5, 23
  principle, 36
belt weigher, 24
  history, 24
Bond's third theory, 49
bubble surface area flux, 99
burnability, 238

carbon monoxide, 234
carbon-in-leach, 138
cataracting, 47
cathode nodule, 170
  history, 178
cell voltage
  practice, 180
cement clinker, 232
cement raw mix, 23
charge motion, 46
classification, 72
clinker quality
  variability, 253
closed grinding circuit, 6, 72
collector
  coverage, 99
  overdose, 104
comminution, 44
compression zone, 118
control valve
  slurry service, 110
copper electrorefining, 167
  history, 169
critical speed, 47
crud, 201
cut size, 77
cyanicide, 142
cyanide analyzer, 158
cyanide control
  practice, 160
cyanide destruction, 160
cyclone feed sump, 76

density gauge
  gamma, 131
dilution water, 76

empty-silo policy, 30
entrainment, 99
evaporative crystallizer, 216
evaporator scaling, 222
extractant
  capacity loss, 201
  phosphonic acid, 190

false air, 234
feeder, 22
feedwell, 122
fire assay, 139
flocculant, 121
flooding, 99
flowsheet, 4
free lime, 234
Freundlich isotherm, 140
froth depth, 97
  measurement, 110
froth flotation, 7, 95
  history, 97
froth recovery, 99

gas holdup, 99
gold cyanidation, 138
  adsorption on carbon, 140
  history, 140
grade-recovery curve, 97
grinding mill, 43

hydrocyclone, 72
  cut-size correlation, 77
  density effect, 89
  history, 75
hydrogen cyanide, 142

impact angle, 44
impact margin, 45
impact microphone, 50

history, 46
principle, 65
inductively coupled plasma, 204
instrument, 11

leach train, 8, 142
leveling agent, 171
limiting current density, 169
liner wear
charge slip, 48
lithium hydroxide refinery, 211

metering pump, 159
mill filling, 47
crash stop, 65
mill overload, 53, 57
mill power, 48
mineral processing, 1
history, 3
milestones, 2
mixer-settler, 187
monitoring stage, 205
morphological stability, 170

neodymium, 188

on-stream analyzer, 100
ore blending, 22
online analysis, 38
overflow clarity, 121
oxygen analyzer, 238

particle size analyzer, 89
partition curve, 78
praseodymium, 188
pressure transmitter
remote seal, 87
pulp level
air effect, 112

rake torque, 119
rare earth elements, 187
history, 189
raw meal, 233
rectifier, 180
roping, 78, 86
rotary kiln
calciner, 214
cement, 232
energy use, 226
history, 235
residence time, 236
speed, 246
rougher flotation cell, 95

saponification, 197
Sauter mean bubble diameter, 99
semi-autogenous mill, 43
monitoring, 66
settling velocity
hindered, 121
short-circuit detection, 180
slurry pump
variable speed, 132
solvent extraction, 8, 187
spodumene, 211
history, 214
sulfidic ore, 151
suspended solids, 172

tailings thickener, 117
tankhouse, 167
thickener, 8, 117
history, 118
toe of the charge, 45
trunnion bearing pressure, 45

variable-frequency drive, 37, 65, 88

water recovery, 118

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.