INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 15 Control in Mining and Mineral Processing

CHAPTER 1

Introduction: Variable Ore and Late Assays Along the Circuit

Pages 1–21

Edited by John D. Hedengren

How to cite this chapter. "Introduction: Variable Ore and Late Assays Along the Circuit," in *Control in Mining and Mineral Processing*, John D. Hedengren, Ed., vol. 15 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 1, pp. 1–21.

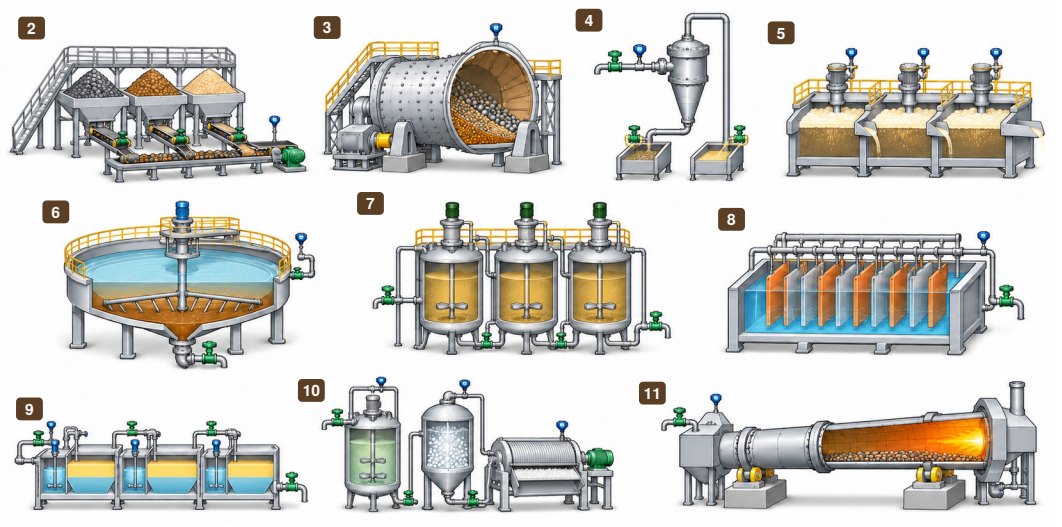
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v15-mining-minerals/

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

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

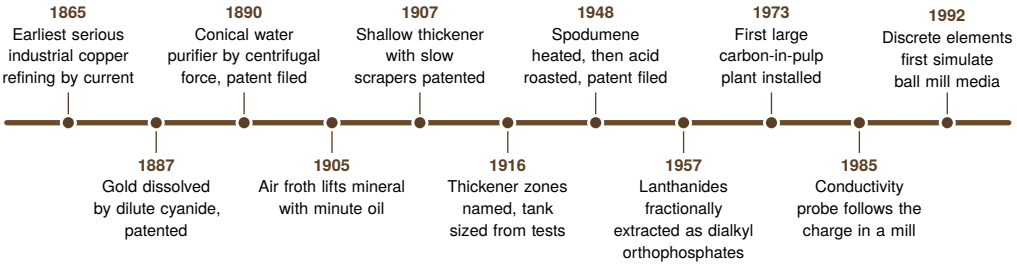


Figure 1.1 Milestones of mineral processing and its measurement. The dates are from [3] (1865), [4] (1887), [5] (1890), [6] (1905), [7] (1907), [8] (1916), [9] (1948), [10] (1957), [11] (1973), [12] (1985) and [13] (1992). A United States patent is dated by its first filing. Notice that every unit of 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

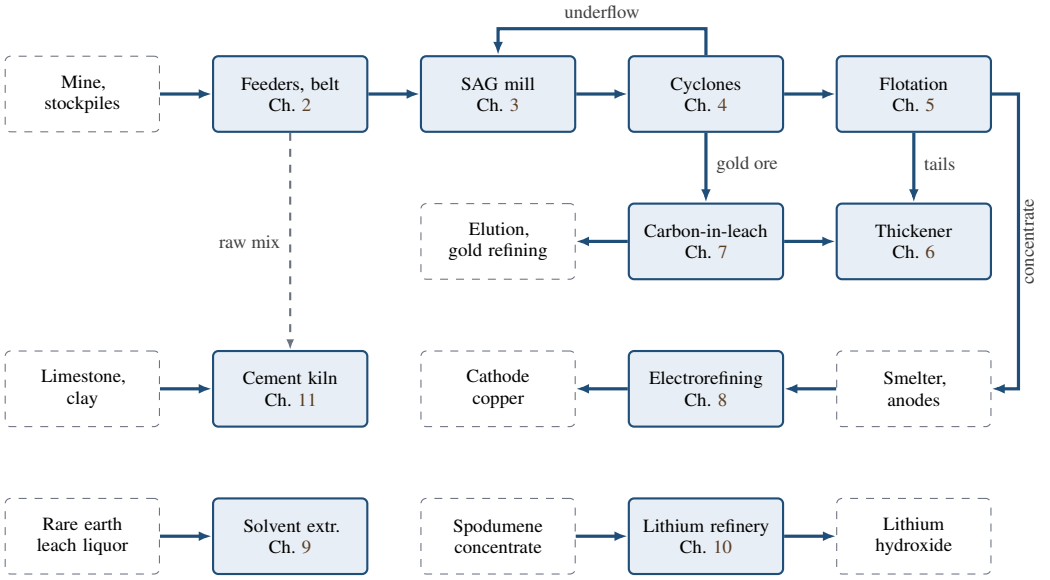


Figure 1.2 The Plants of the Volume on the Path From Ore to Product. Filled Blocks Are the Plants of Chapters 2 to 11, Dashed Blocks Lie Outside the Volume. The Top Two Rows Are a Concentrator, With Thickener Water Returning to the Mill (Not Drawn). The Dashed Arrow Marks That Feeders and a Belt Scale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

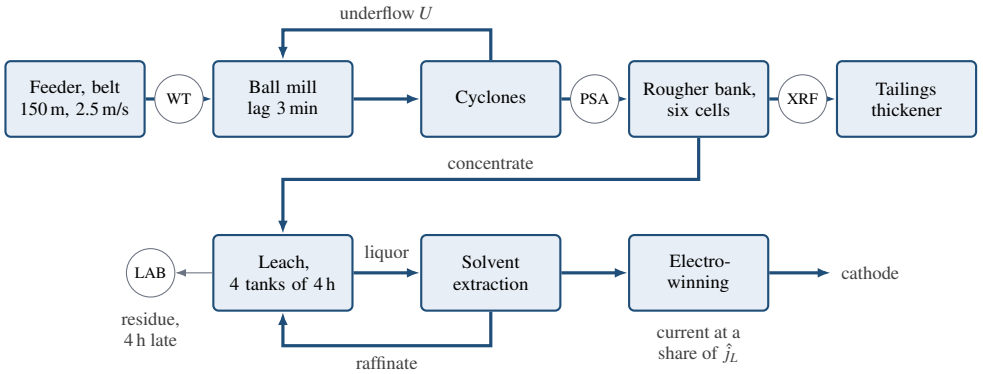


Figure 1.3 The Running Example C-1. Circles Mark the Measurements on the Product Path: the Belt Scale (WT), a Particle Size Analyzer (PSA) on the Cyclone Overflow, an On-Stream X-Ray Fluorescence Analyzer (XRF) on the Flotation Streams, and a Laboratory Assay (LAB) of the Leach Residue, Sampled Every 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.

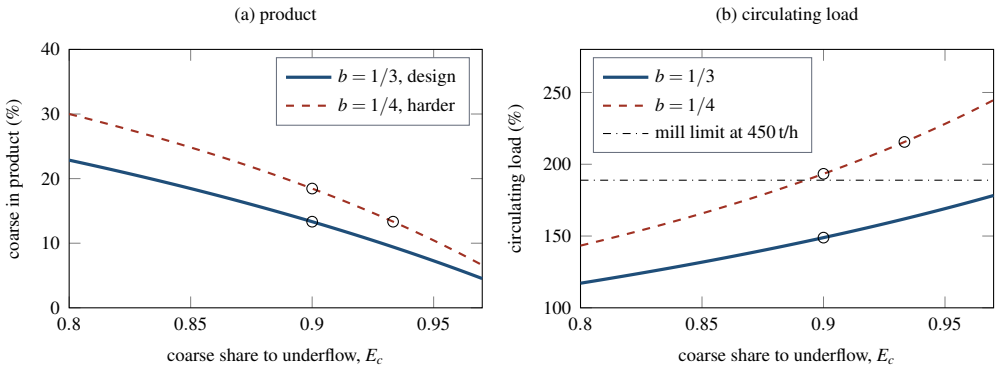


Figure 1.4 C-1's Closed Grinding Circuit at 450 t/h of Fresh Ore. (a) The Coarse Share of the Product and (b) the Circulating Load Against the Share of Coarse the Cyclones Return, for the Design Ore (Thick Solid) and a Harder Ore (Dashed). Circles Mark the Design, the Harder Ore at the Design Cut, and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

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

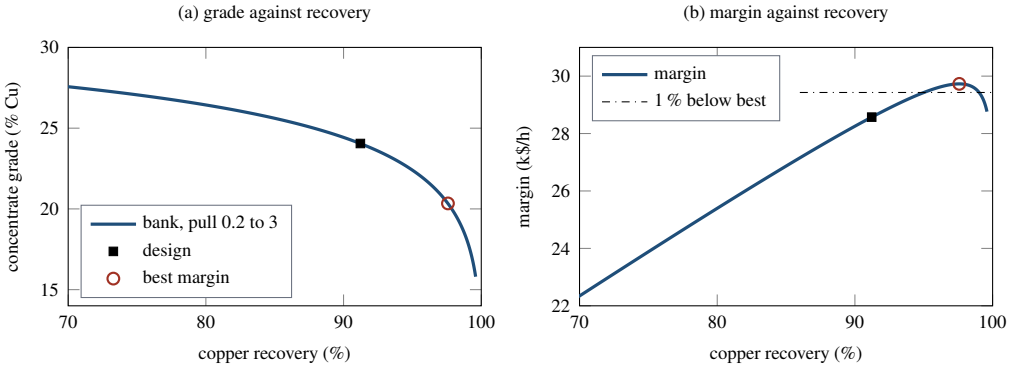


Figure 1.5 C-1's Rougher Bank as the Pull Factor Moves From 0.2 to 3. (a) Concentrate Grade Against Copper Recovery. (b) The Margin, Copper at $8000/t Less $250 per Tonne of Concentrate. The Square Marks the Design and the Circle the Best Margin. Every Point Above the Dash-Dotted Line Is Within 1 % of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

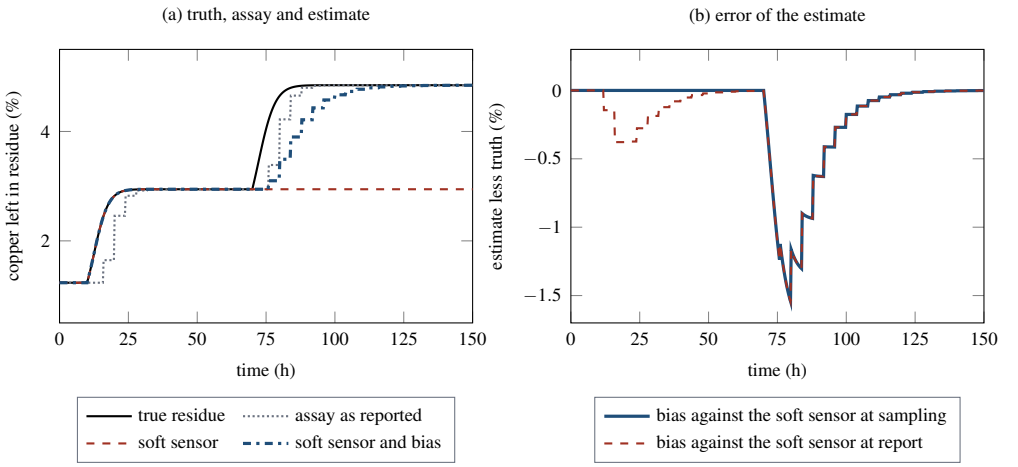


Figure 1.6 C-1's Leach Residue With the Train 5 K Cooler From 10 h (Measured) and an Ore That Leaches 20 % More Slowly From 70 h (Not Measured). (a) The True Residue (Solid), the Assay Held From Each Report (Dotted), the Soft Sensor (Dashed) and the Soft Sensor With the Bias Update at $\lambda = 0.35$ (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

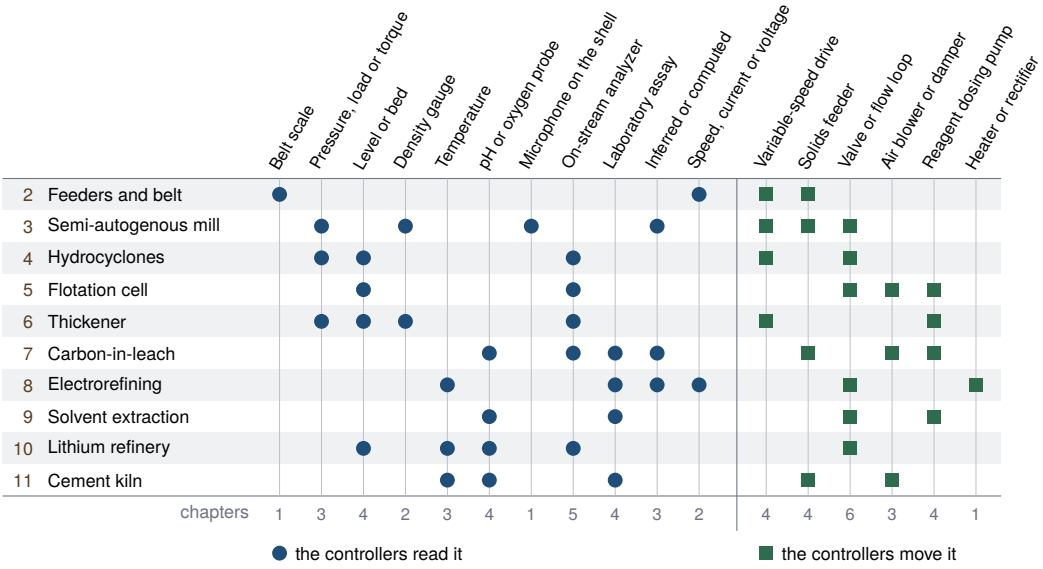


Figure 1.7 Sensors and final elements by chapter. A round mark means that the chapter's controllers read the device, a square mark that they move it. A mark is set only where the chapter's variables table, control structure or technology section names the device, and a flow station counts as a valve 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.

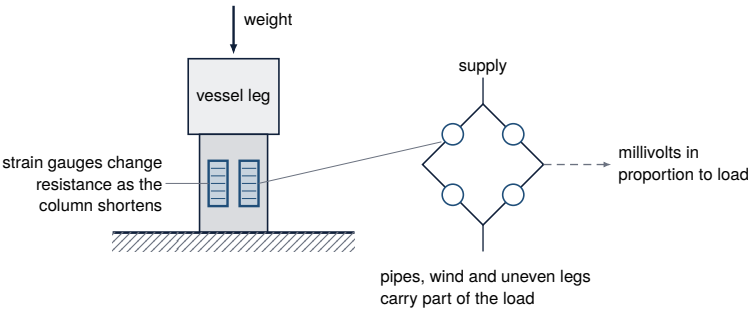


Figure 1.8 Principle of a strain-gauge load cell, drawn under the leg of a vessel. Notice that the cell reads only the load that passes through it. On a belt scale the cells carry the rollers of the weighing section, so they weigh whatever lies on that section, from whichever feeder it came [37].

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.

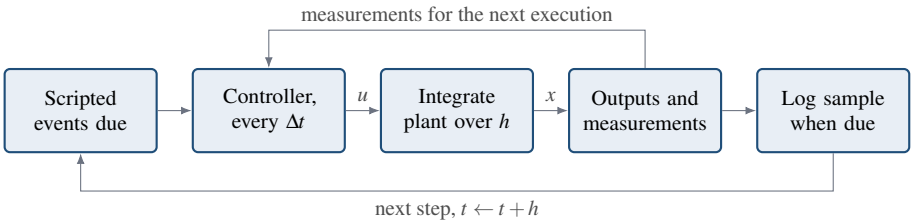


Figure 1.10 Work of the Engine in One Integration Step

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