INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 15 Control in Mining and Mineral Processing

CHAPTER 5

A Browser-Based Froth Flotation Simulator for Economic Setpoint Selection and Loop Interaction

Pages 95–116

Edited by John D. Hedengren

How to cite this chapter. "A Browser-Based Froth Flotation Simulator for Economic Setpoint Selection and Loop Interaction," 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. 5, pp. 95–116.

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

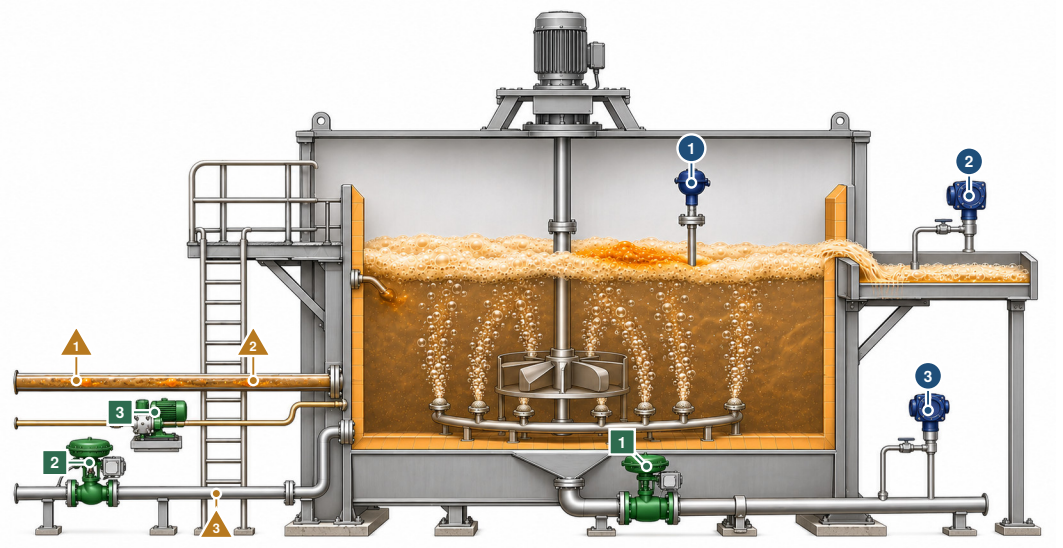
The app for this chapter
https://apmonitor.com/apps/froth-flotation/

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

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

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.

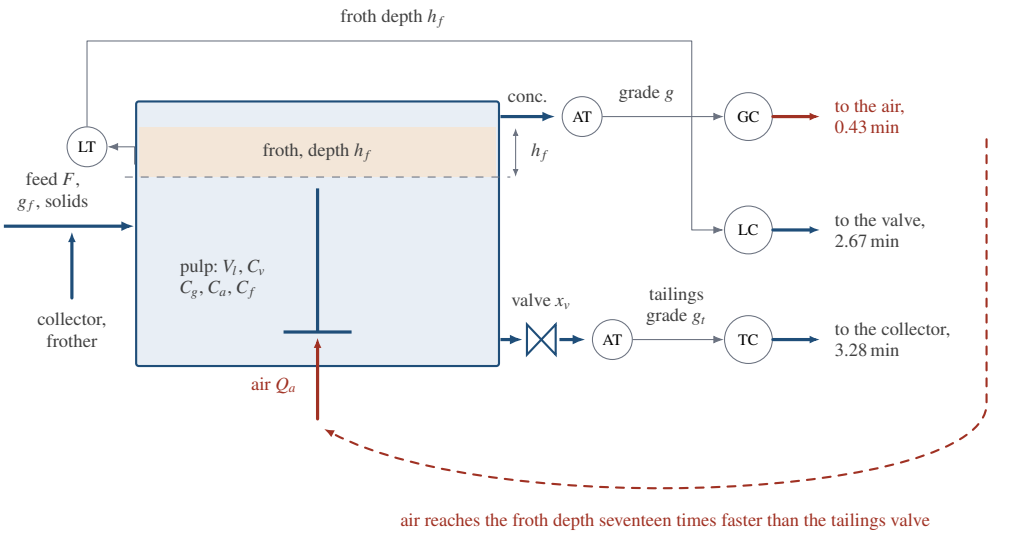


Figure 5.1 The simulated cell. The level controller (LC) holds froth depth on the tailings valve, the grade controller (GC) holds concentrate grade on the air rate, and the tailings-grade controller (TC) holds the tailings assay on the collector dose. Each output arrow carries the measured open-loop 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

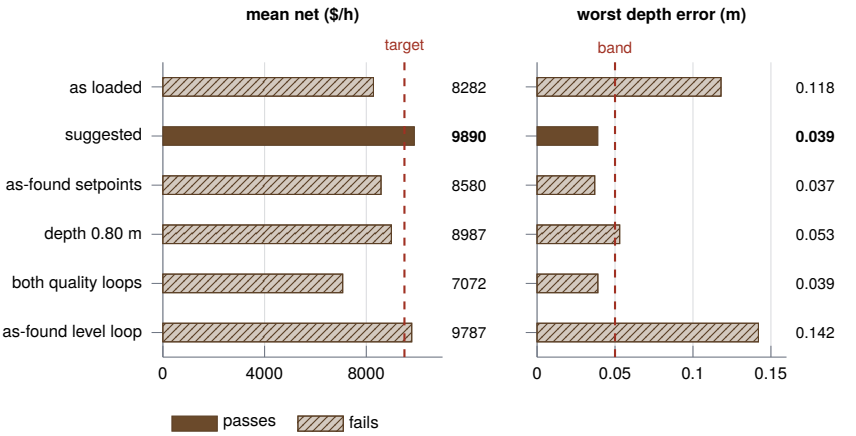


Figure 5.3 The graded scenario of Table 5.4: mean net value against the $9500 /h target, and worst froth depth error against the 0.05 m band. Only the suggested configuration (solid) clears both. Notice that the as-found setpoints hold the depth best of all and still fall short of the target, while the 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

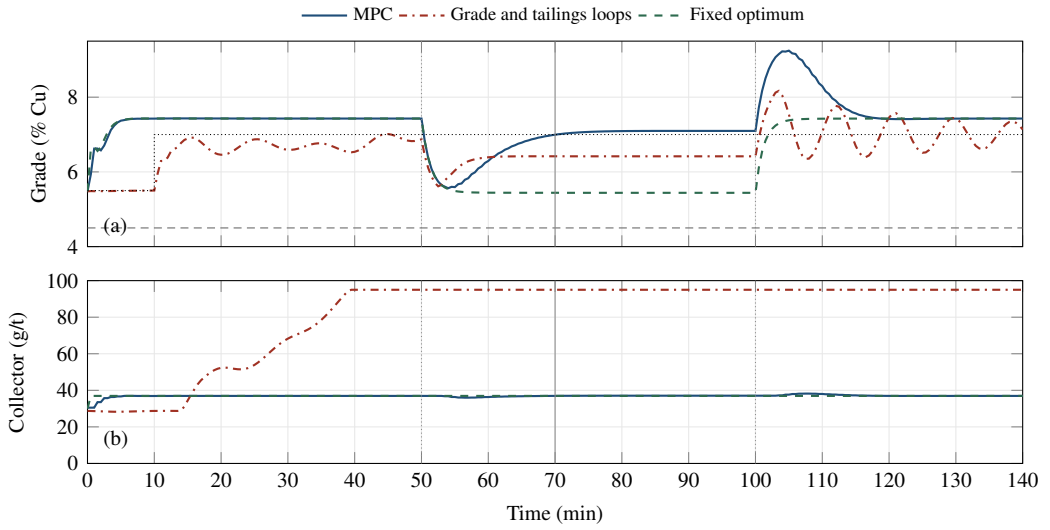


Figure 5.4 Scenario 11 under the MPC (solid), the shipped grade and tailings loops (dash-dotted), and the MPC model's optimum as fixed stations (thick dashes). (a) True concentrate grade with the request (dotted) and the 4.5 % minimum (thin gray dashes). (b) Collector dose. The lean feed lasts from 50 to 100 min (dotted), and the solid line marks 70 min.

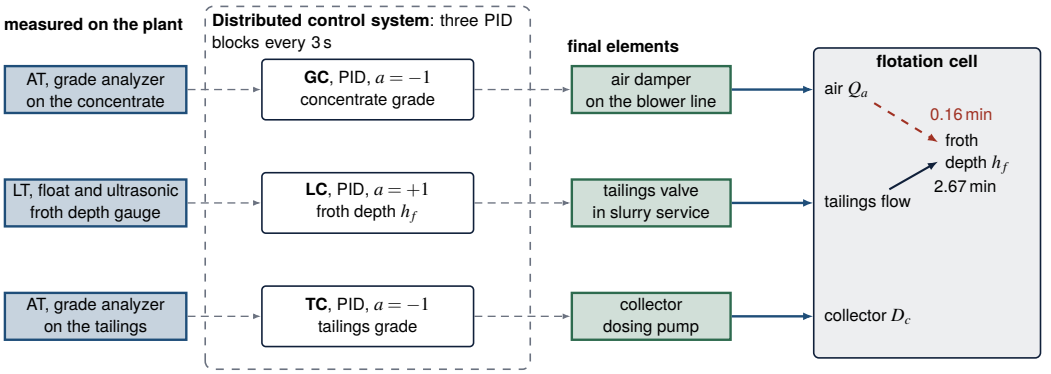


Figure 5.5 The three loops of the simulated cell drawn with their hardware. The froth depth gauge LT and the two grade analyzers AT feed three PID blocks. LC is reverse acting and moves the tailings valve, GC and TC are direct acting and move the air damper and the collector dosing pump. The dashed path 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.