INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 15 Control in Mining and Mineral Processing

CHAPTER 7

Soft Sensing and Inferential Control in a Gold Carbon-in-Leach Simulator

Pages 138–166

Edited by John D. Hedengren

How to cite this chapter. "Soft Sensing and Inferential Control in a Gold Carbon-in-Leach Simulator," 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. 7, pp. 138–166.

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

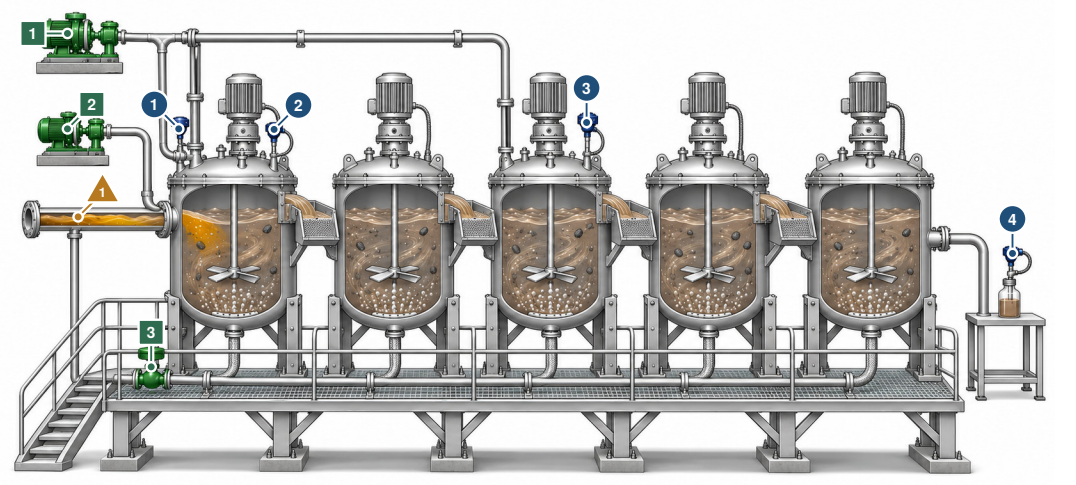
The app for this chapter
https://apmonitor.com/apps/gold-cil/

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

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

CHAPTER 7

Soft Sensing and Inferential Control in a Gold Carbon-in-Leach Simulator



The controller moves 1 cyanide dosing pump, to tanks 1 and 3 2 lime dosing pump 3 air valve
The controller reads 1 free cyanide analyzer 2 pH probe 3 dissolved oxygen probe
4 tailings sample for the laboratory

What disturbs it 1 changes in the ore: grade, leaching rate, copper, sulfides

The gold left in the slurry leaving the last tank decides what the circuit earns, and the laboratory reports it hours after the slurry has gone. Illustration, not to scale.

IN PLAIN TERMS **Steering by a Report That Comes Late**

The main idea. Gold is taken out of crushed rock by stirring it with a weak cyanide solution in a row of five tanks. Grains of charcoal in the tanks soak up the dissolved gold, and more cyanide frees more gold but costs more.

The challenge. The plant is paid on the gold left in the waste, which a laboratory reports hours late. A quick estimate from a computer model helps, but it misses a change in the ore until the laboratory corrects it.

Why it matters. Too little cyanide throws gold away, and too much wastes reagent that then has to be destroyed.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history and fact boxes of Section 7.2, the history box of Section 7.3, the try-it boxes of Sections 7.6.2 and 7.6.4, the fact box of Section 7.6.5, and the key points
Study	70 min	The whole chapter with its pause points and the twelve exercises
Practice	30 min	The first page, Section 7.6, Section 7.7, Section 7.8, and the open challenges after it and in Section 7.9

The control question A carbon-in-leach circuit is paid on tailings gold that a laboratory reports six to ten hours late, on a train that takes most of a day to answer. Should an outer loop move the cyanide setpoint, and on the assay or on a soft sensor the assay corrects?

Learning objectives

- Explain why a loop on a late, sampled assay acts in steps, and when a soft sensor is faster.
- Compute the time constant and the noise of a bias update, and choose its gain.
- Show why holding the head-tank setpoints does not hold the tailings.
- Judge from a measured week whether a supervisory loop beats a fixed setpoint.
- Separate what an economic predictive controller gains from its target and from its prediction.

Prerequisites Sections 1.4.4, 1.4.5, 1.4.6 and 1.6.2, and Section 1.6.3 for the predictive controller and the estimator. Chapter 5 has an analyzer delay and an economic setpoint on a faster unit.

The app https://apmonitor.com/apps/gold-cil/

See also Volume 5, Chapter 10, Volume 5, Chapter 4 and Volume 5, Chapter 3 for late quality measurements. Volume 9, Chapter 10 and Volume 9, Chapter 4 for supervisory setpoints and offset removal.

7.1 Introduction

A carbon-in-leach circuit dissolves gold from ground ore with cyanide and adsorbs it onto activated carbon in the same tanks. The reagent decision is a cyanide addition rate. The variable the circuit is paid on is the gold left in the tailings, and the interval between the two is longer than in any other application in this series.

The plant contributes about half of that interval. A step in the sodium cyanide addition reaches sixty-three percent of its effect on the free cyanide concentration of the first tank in 2.68 h and on the tailings grade in 11.52 h, with ninety-five percent of the tailings response taking 20.46 h, because five stirred tanks in series each hold about 3.14 h of slurry and each is a first-order lag on the one before it.

The measurement contributes the remainder. Tailings grade is a fire assay taken every four hours and reported six hours later, so the most recent available number describes the circuit as it was between six and ten hours earlier (Section 1.4.5 explains the age of a sampled assay). No continuous instrument for that variable exists.

A controller closed on that measurement is not poorly tuned. It is structurally slow, and the standard industrial remedy is an inferential estimate. This chapter works with a simulated carbon-in-leach circuit that supplies such an estimate, shows three distinct ways in which it fails, and measures whether closing an outer loop around it improves operation at all.

The circuit shows process order through five first-order lags in series, dead time and sampled measurement through the laboratory assay, cascade structure and time-scale separation between the secondary and outer loops, inferential control through the soft sensor, discrete bias correction as a filter with its own dynamics, and actuator travel as a cost that tuning must account for. Section 7.2 places the problem in mineral processing practice. Section 7.3 gives the five-tank model, the soft sensor and its bias update, and Section 7.4 ranks the three reagent setpoints at steady state. Section 7.5 describes the secondary loops, the outer loop, the predictive controller and the moving horizon estimator. Section 7.6 closes the outer loop on the assay and on the soft sensor (Section 7.6.1), meets unmeasured ore changes (Section 7.6.2), varies the bias update gain (Section 7.6.3), passes copper and sulfide parcels (Section 7.6.4), runs the graded week (Section 7.6.5), and puts the predictive controller (Section 7.6.6) and the estimator (Section 7.6.7) in place of the loop and the update. Section 7.6.5 holds the central result, that on a week of difficult ore the recommended outer loop saves little reagent and costs most of a day above the tailings limit. Section 7.7 names the hardware of a real circuit, Sections 7.8 to 7.10 state the model's limits, what generalizes and what was found, and Section 7.11 holds the exercises.

7.2 Background

Cyanidation followed by adsorption onto activated carbon is the dominant hydrometallurgical route for gold recovery, and carbon-in-leach places the carbon in the same tanks as the leach so that dissolution and adsorption proceed together [1, 2]. Gold dissolution requires cyanide and dissolved oxygen together, and the alkalinity that keeps the cyanide in its ionic form is maintained with lime. Three reagents therefore set the leaching rate, and all three are measurable in the head tank on continuous instruments.

> ⧖ HISTORY **A Dilute Cyanide Solution (1887)**
>
> In 1887 John Stewart MacArthur and Robert and William Forrest, of Glasgow, patented in Britain a process that dissolves gold from powdered ore with a solution of a cyanide. Kept extremely dilute, the solution took the gold and silver in preference to the base metals, and stirring the ore with it in an agitated vessel sped the process. Their United States patent was granted in May 1889. A second patent of the same year dealt with ores holding partly oxidized pyrites, which absorbed much of the cyanide. The ore was first neutralized with an alkali or an alkaline earth, then leached, and the gold was precipitated from the solution on zinc [3, 4]. The lime addition of this chapter does the same work.

The variable that carries the economics is not measurable on those instruments. Mineral processing plants routinely regulate the fast intermediate variables while the grade or recovery that determines revenue arrives from a laboratory hours later, and the resulting mismatch between the control layer and the economic layer is a long-standing problem in the field [5]. Where an instrument is missing, the choices are to add one, to accept the delay, or to infer the variable from what is already measured, and the economics of that choice have been posed as a sensor placement problem [6].

> ✦ DID YOU KNOW? **The Cyanide That Was Not Measured (1995)**
>
> In 1995 a change to copper-rich ore at a mine in Australia raised the cyanide bound to metals, the weak-acid-dissociable cyanide, to 350 parts per million in its tailings dam, and about 2700 birds died. The mine measured only free cyanide, so the rise went undetected [7].

Inferential control takes the third route. A model computes an estimate of the unmeasured controlled variable from measurements that are available now, the controller closes on the estimate, and the delayed laboratory result is used only to correct the estimate's offset [8, 9]. Section 1.4.6 describes the structure as practiced, and Volume 5, Chapter 10 applies it to a dryer whose product moisture is inferred from a soft sensor that a sampled and delayed laboratory corrects. The structure is a cascade in the usual sense. The secondary loops reject fast disturbances in variables that respond in hours, and the outer loop trims their setpoints against a variable that responds in half a day [8, 10]. Cascade control earns its keep when the secondary loop is several times faster than the primary one, and the ratio of about four to one between the cyanide response and the tailings response places this circuit inside that condition.

Dead time is what makes the alternative unattractive. A proportional-integral controller closed directly on a measurement whose dead time is comparable to the dominant plant time constant must be detuned severely to remain stable, and the detuning removes the disturbance rejection the loop was installed to provide [11, 12]. Tuning rules that account for dead time explicitly, such as the internal model control family and the SIMC rules, make the trade-off visible in the gain itself [13, 14].

The physical fundamentals underneath the simulator are conservation laws applied to each tank. Component mass balances on the unleached gold in the solids, the dissolved gold, the gold loaded on carbon, the free cyanide and the alkalinity give five first-order ordinary differential equations per tank, with the flow terms setting the residence time and the reaction terms setting the driving force. Adsorption is driven by departure from a Freundlich equilibrium loading, and dissolution is driven by the unleached gold remaining above a refractory floor. Dissolved oxygen equilibrates in minutes

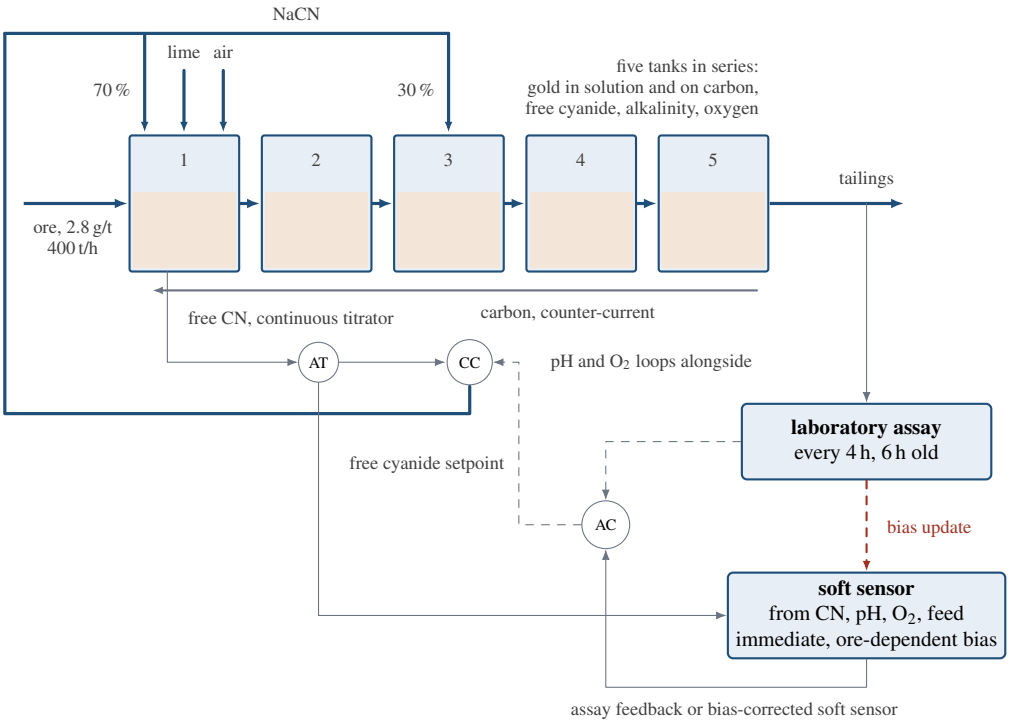


Figure 7.1 The simulated carbon-in-leach train. Three secondary loops hold free cyanide on the sodium cyanide addition, head-tank pH on the lime addition and dissolved oxygen on the air rate, all on continuous instruments. Seventy percent of the sodium cyanide enters tank one and thirty percent enters tank three. An outer controller (AC) moves the free cyanide setpoint against a tailings grade that is not measured, taking it either from the laboratory assay, which is sampled every four hours and reported six hours later, or from a soft sensor that is available immediately and carries a structural error whenever the ore changes. The dashed path from the laboratory assay to the soft sensor, drawn red, is the bias update that reconciles the two.

against tank residence times of hours, so time-scale separation permits it to be treated algebraically rather than as an additional state.

7.3 Process Model

Figure 7.1 shows the train, the three secondary loops and the two competing outer architectures.

7.3.1 Design Point

The clock is in hours. Five tanks of 2000 m^3 treat 400 t/h of ore at 2.8 g/t gold, giving a slurry residence time of 3.14 h per tank and 15.7 h over the train. At design the addition rates are 198.1 kg/h of sodium cyanide, 66.3 kg/h of lime and 100 % air, which hold 120 mg/L of free cyanide, pH 10.40 and 6.0 mg/L of dissolved oxygen in the head tank. The unleached gold in the solids falls from 1.099 g/t in tank one to 0.045 g/t in tank five and the dissolved gold from 0.031 mg/L to 0.0001 mg/L, so the tailings assay 0.295 g/t at 89.5 % recovery, and the loaded carbon leaves tank one at 4058 g/t. Sodium cyanide consumption at the design point is 0.495 kg/t. Seventy percent of the sodium cyanide is added to tank one and thirty percent to tank three, and all of the lime enters tank one, so the free cyanide falls from 120 mg/L in tank one to 89.7 mg/L in tank five and the pH from 10.40 to 10.05.

7.3.2 Leach and Adsorption Kinetics

Gold dissolution is first order in the unleached leachable gold with a rate constant that depends on all three reagents [15], and the cyanide and oxygen terms are saturating, as shown in Equation (7.1). The leach train of Section 1.4.4 has the same first-order dissolution at a fixed rate, and Equation (1.6) gives its residue.

$$r = k \frac{[CN^-]}{K_{CN} + [CN^-]} \cdot \frac{[O_2]}{K_O + [O_2]} \cdot f(\mathrm{pH}) \ (Au - Au_\infty) \tag{7.1}$$

where r is the specific dissolution rate in g/t per hour, k is the intrinsic rate constant of 0.947/h, K_{CN} and K_O are the half-rate concentrations of 60 mg/L and 3.0 mg/L, $f(\mathrm{pH})$ is a dimensionless factor with a maximum at pH 10.4, and Au_∞ is the refractory fraction. The refractory fraction is 0.25 g/t of the 2.8 g/t feed and is the floor that no reagent addition reaches. The pH factor falls away on both sides of its maximum, which is what produces the interior optimum of Table 7.1, and the underlying surface chemistry is more involved than either of the saturating factors [16].

Free cyanide is consumed by the dissolving gold, by cyanicides in the ore and by volatilization as hydrogen cyanide, the last of which depends on pH through the fraction of the free cyanide present in the protonated form, as shown in Equation (7.2).

$$f_{HCN} = \frac{1}{1 + 10^{\mathrm{pH} - 9.31}} \tag{7.2}$$

where 9.31 is the pK_a of hydrogen cyanide and f_{HCN} is dimensionless. At the design pH of 10.40 the protonated fraction is 7.5 %, and at the safety floor of pH 9.8 it is 24.5 %, so the pH loop protects the cyanide loop as well as the operators. Soluble copper raises the cyanicide term by more than a factor of two [17] and is the disturbance of Section 7.6.4.

Carbon adsorption is driven by the departure of the loading from a Freundlich equilibrium and proceeds at a finite film transfer rate, counter-current to the slurry, so tank five contacts the freshest carbon and holds the lowest solution grade.

> ⧖ HISTORY **Carbon Joins the Leach (1880 to 1973)**
>
> Wood charcoal was patented by Davis in 1880 to scavenge gold from chlorination solutions, and once cyanidation was established it was patented for cyanide solutions as well. At a mine in Western Australia in 1917 it showed only a very low capacity for gold, and the gold could be recovered only by burning the charcoal, so zinc precipitation prevailed. In 1950 Zadra eluted gold from activated carbon with a hot alkaline cyanide solution and electrowon it, which let the carbon be reused, and the first large carbon-in-pulp plant was installed in South Dakota in 1973 [18]. Carbon-in-leach puts the carbon in the leach tanks themselves and suits ores that leach fast or rob gold from the solution [19].

7.3.3 Tailings Grade Estimation

Three values of the tailings grade are computed and reported separately.

The first is the value computed by the model, which is available to the application and to no instrument on a real circuit. The second is the laboratory assay, which is that value sampled every four hours and reported six hours later. The third is the soft sensor, which evaluates (7.1) forward from the measured reagent concentrations and the feed grade and reports immediately. The soft sensor assumes that every tank sits at the head-tank cyanide, pH and oxygen, which the profile of Section 7.3.1 shows to be false, so it is structurally optimistic and its accuracy is set by its assumption about the ore. A bias term reconciles the estimate with the assay, as shown in Equation (7.3).

$$b \leftarrow b + \beta \left(Au_{\mathrm{lab}} - \hat{Au}_{\mathrm{soft}} - b \right) \tag{7.3}$$

where b is the bias in g/t, Au_{lab} is the assay just reported, $\hat{Au}_{\mathrm{soft}}$ is the soft sensor output recorded at the time the sample was cut rather than the current one, and β is a dimensionless operator setting.

Table 7.1 Steady state against each secondary setpoint, with the other two holding. 0.25 g/t of the feed is refractory and no reagent reaches it.

Setpoint	Value	Tailings (g/t)	Recovery (%)	NaCN (kg/t)	Margin ($/h)
Free cyanide	60 mg/L	0.3503	87.49	0.255	76007
	120 mg/L	0.2950	89.46	0.495	77423
	200 mg/L	0.2789	90.04	0.813	**77516**
	260 mg/L	0.2742	90.21	1.051	77357
pH	9.60	0.3967	85.83	0.511	74268
	10.40	0.2950	89.46	0.495	77423
	10.55	0.2922	89.57	0.494	**77505**
	11.20	0.3521	87.43	0.487	75604
Oxygen	3.0 mg/L	0.3446	87.69	0.493	75940
	6.0 mg/L	0.2950	89.46	0.495	77423
	7.2 mg/L	0.2883	89.70	0.496	**77502**

It is the update of Equation (1.8) with the weight λ written β, and the app labels it λ, the Soft sensor bias update gain. The update is applied once per assay, so (7.3) is an exponential filter on the assay innovation with a time constant of $-T_s/\ln(1-\beta)$ for a sampling interval T_s of 4 h, which approaches T_s/β for a small gain. The update runs only while the outer loop reads the soft sensor, under the soft sensor loop or the predictive controller of Section 7.5.2. The application supplies a default of $\beta = 0.40$ and the graded week uses 0.35. Sections 7.6.2 to 7.6.5 measure what happens at $\beta = 0$, 0.15, 0.35 and 1.0.

🕮 **Pause.** Try Exercise 7.2 before reading on.

7.3.4 Numerical Implementation and Assumptions

Initialization marches the five tanks forward from the feed at the given addition rates and sweeps the counter-current carbon backward until neither direction moves, so the returned state is an exact equilibrium. Helper routines return the lime addition for a chosen pH and the cyanide addition for a chosen free cyanide concentration, both by bisection.

Gold is one lumped leachable class plus a refractory floor rather than a particle size distribution, carbon is one loading per tank with no regeneration circuit, and cyanide destruction, elution and the gold room are outside the boundary. Integration is fixed-step fourth-order Runge-Kutta at 0.01 h with a 3 min controller scan, as the engine of Section 1.6.1 runs it. The free cyanide titrator reports the head-tank concentration 0.40 h late, the pH electrode is a first-order lag of 0.05 h, and the cyanide pump, lime feeder and air damper are first-order lags of 0.08 h, 0.12 h and 0.06 h.

7.4 Effect of the Reagent Setpoints on Leach Performance

Three reagent setpoints are available to an operator and their steady-state effects differ in shape.

Table 7.1 contains three different shapes. The cyanide response saturates, because raising the setpoint from 200 mg/L to 260 mg/L lowers the tailings grade by 0.0047 g/t and raises the reagent consumption by 0.238 kg/t, so the operating margin has a broad flat maximum near 200 mg/L. The oxygen response saturates more sharply and the blower reaches its 250 % ceiling at 7.2 mg/L. The pH response has an interior optimum at 10.55, because the leach rate falls above that value while the free cyanide is lost as hydrogen cyanide below it.

The pH optimum is shallow and sits close to the design value. It is worth $82 per hour against the design pH of 10.40 and $3237 per hour against pH 9.60, so the penalty for operating low is roughly

forty times the prize for finding the optimum. The lower bound on pH is a safety constraint at 9.8, below which hydrogen cyanide evolves in quantity. The upper bound is economic and no instrument on the circuit reports it.

None of these setpoints is the difficulty. The difficulty is that the tailings grade is the variable the mine is paid on, it responds in half a day, and it is measured four times a shift with a six-hour reporting delay.

7.5 Controller Design

Three secondary loops and one outer loop use the velocity-form PID law of Section 1.6.2, Equations (1.10) to (1.13) [10]. Each scan computes an increment to the reagent addition rather than the rate itself, and the derivative is taken on the measurement, so a setpoint move produces no derivative kick. The controller scan is $\Delta t = 0.05\,\mathrm{h}$, the measurement passes through a first-order filter of 0.2 h for the derivative alone, and only the cyanide loop offers a derivative time, which is zero as supplied. The output is clipped to the actuator span and, where a loop declares one, moved no more than $r\,\Delta t$ per scan by a rate limit of r per hour. Only the outer loop has a rate limit, $r = 8\,\mathrm{mg/L}$ per hour. Anti-windup is back-calculation on the integral alone with a tracking time $T_t = \tau_I$, the last term of Equation (1.13). While the output is clipped or rate limited it drives the integral toward the value that reproduces the emitted output, removing $\Delta t / \tau_I$ of the excess each scan, and the running sum of proportional and derivative increments is never truncated. With $T_t = \tau_I$ and no derivative, the integral relaxes toward the emitted output with time constant τ_I instead of integrating the error while the ramp delivers the proportional demand, so the integral does not wind up behind the rate limit. In manual, and in the outer loop whenever it is switched out, the integral tracks so that the output equals the operator's value exactly, which gives bumpless transfer, and in automatic the manual stations follow the controller outputs after every execution. The secondary loops are tuned by the SIMC rules [14]. Each has the action $a = +1$, reverse acting in the convention of Section 1.6.2, because more reagent raises its own measurement. Free cyanide is controlled by the sodium cyanide addition between 0 and 500 kg/h, head-tank pH by the lime addition between 0 and 600 kg/h and dissolved oxygen by the air rate between 20 % and 250 %, all three on continuous instruments.

7.5.1 Time-Scale Separation Between the Loops

The three secondary loops close on measurements that respond in hours and are available continuously. A step in sodium cyanide reaches sixty-three percent of its effect on the head-tank free cyanide in 2.68 h and ninety-five percent in 7.90 h, and a step in lime reaches sixty-three percent of its effect on head-tank pH in 3.18 h, both measured in the application from a ten percent step in manual. Both responses are close to first order, but the titrator adds 0.40 h of dead time to the cyanide measurement, so the cyanide loop sees its own step reach sixty-three percent only after 3.11 h.

The outer loop closes on a variable whose response to the same actuator takes 11.52 h to sixty-three percent and 20.46 h to ninety-five percent, and which is not measured. The ratio between the two responses is about 4.3 to one on the plant alone before the assay is considered, which places the arrangement inside the condition under which cascade control is normally recommended [8]. Satisfying that condition is not sufficient, and Section 7.6.5 reports why.

The outer controller moves the free cyanide setpoint against the tailings grade in the ordinary cascade arrangement, and four configurations are available: the outer loop switched out, the outer loop closed on the laboratory assay, the outer loop closed on the soft sensor with the bias update of (7.3), and the predictive controller of Section 7.5.2. The app's Outer loop selector names them Local, Assay trim, Soft sensor and MPC. The first three configurations are compared in Section 7.6, and the fourth in Section 7.6.6. The outer controller is direct acting ($a = -1$) in the convention of Section 1.6.2, as the app's hint also names it, because raising the cyanide setpoint raises the leach rate and therefore lowers the tailings grade, so a grade above target calls for more cyanide. Its output

is held between 40 mg/L and 400 mg/L and rate limited to 8 mg/L per hour so that a new assay does not step the secondary setpoint.

7.5.2 Model Predictive Control

The fourth outer-loop setting replaces the outer PI controller with an economic receding-horizon quadratic program [20, 21] that moves the free cyanide setpoint, which the cyanide loop follows as before. Section 1.6.3 describes the predictive option of the apps in general, and this section gives what is particular to this one. The pH and oxygen loops are unchanged. The three configurations above stay the application's default and the starting structure of every scenario but scenario 11, which opens on the predictive controller. While the plan acts, the outer PI controller tracks the setpoint it writes, so a hand-over in either direction carries no bump.

Controller model. The steady part of the model is the soft sensor itself. Its estimate $w(c)$, which is (7.1) evaluated with every tank at the head-tank conditions and with the soft sensor's structural error included, is computed at the measured pH, oxygen, feed grade and throughput from the soft sensor's own constants, which are the plant's declared values, and it is linearized about the measured free cyanide c_0 at every decision by central differences of ±5 mg/L, which give a slope w' and a curvature w''. Ahead of it sits the cyanide loop, identified from setpoint steps of ±10 mg/L at the tuning of scenario 11 with noise off as a first-order lag of 0.48 h behind 0.66 h of dead time (0.475 h and 0.645 h in an independent refit, with a residual of 1.05 % of the response), and carried as one decision period of delay followed by a lag $\tau_{CN} = 0.74$ h, whose mean delay of 1.14 h is close to the refit's 1.12 h. Behind it sit five equal lags, one per tank, at the residue time constant of the soft model's own tank balance, $\tau_r = \tau_{res}/(1+\tau_{res}k_s)$, where $\tau_{res} = 3.14$ h is the slurry residence per tank (the τ_t of Section 1.4.4) and k_s is the soft model's leach rate at the measured point, the rate constant k times the three factors of Equation (7.1), 0.421/h at design, so that $\tau_r = 1.353$ h. These lags reproduce a cyanide step from the soft model's output to the true tailings with a residual of 15.2 %, against 27.7 % for lags at the residence time, and refitting the gain alone, to 1.22, lowers the residual to 3.8 %, so what remains is a gain error. The model is given in Equation (7.4).

$$
\begin{aligned}
d_{k+1} &= u_k \\
\tau_{CN}\,\dot{c} &= d - c \\
\tau_r\,\dot{\ell}_1 &= w(c_0) + w'\,(c - c_0) - \ell_1 \\
\tau_r\,\dot{\ell}_i &= \ell_{i-1} - \ell_i, \qquad i = 2,\ldots,5 \\
\hat{A}u_k &= \ell_{5,k} + b
\end{aligned}
\tag{7.4}
$$

where u_k is the free cyanide setpoint decided at step k in mg/L, d is the setpoint decided one period earlier, c is the measured head-tank free cyanide, ℓ_i are the lags in g/t, b is the bias of (7.3), and $\hat{A}u_k$ is the predicted tailings grade. The continuous part is discretized exactly with the input held over each period and rebuilt at every decision. Between decisions the lags are driven at every controller scan by the soft sensor's output, whichever outer loop acts, so the last lag is the soft estimate delayed through the tanks. The bias is the soft sensor's own, updated by (7.3) with the same gain β and the same residual aligned to the time the sample was cut, which makes the prediction offset-free through an output disturbance model [22] that moves only when an assay is reported. Volume 9, Chapter 4 removes offset instead with a Kalman estimator that carries a fuel bias, a disturbance at the input. A result that arrives between two decisions moves b at once, and the next decision plans with it. The feed grade, throughput, pH and oxygen are held at their measured values over the horizon, because nothing in the circuit announces their future.

The soft sensor's structural error is therefore the controller model's. The steady slope of the true tailings against free cyanide is 1.25 times the model's at design, 1.84 times on the ore of scenario 11 at 200 mg/L and 2.0 times on the ore of scenario 6, because the model holds every tank at the head-tank

cyanide while the cyanide falls along the train. The bias corrects the level of the estimate and not its slope.

Objective and constraints. Every $T_d = 0.4\,\mathrm{h}$, the titrator cycle of eight controller scans, the plan covers $N = 60$ steps (24 h) in ten move blocks of $1, 1, 1, 2, 2, 3, 4, 6, 9$ and 31 steps, and the first move of the minimizer of Equation (7.5) is sent and held until the next decision.

$$\begin{aligned} J = \sum_{k=1}^{N} \left[SP_{Au}\hat{Au}_k + c_{CN}u_k + \tfrac{1}{2}SP_{Au}w''\,(u_k - c_0)^2 \right] T_d \\ + SP_{Au}T_{lag}\hat{Au}_N + s_m \sum_j \Delta u_j^2 + \sum_{k\in\mathscr{K}} \left(\rho_2 \varepsilon_k^2 + \rho_1 \varepsilon_k \right) \end{aligned} \tag{7.5}$$

Here S is the measured ore rate, 400 t/h at design, and P_{Au} is the declared gold price of \$78 per gram, so 0.001 g/t of tailings is worth \$31.2 per hour. The cyanide price c_{CN}, in dollars per hour per mg/L of setpoint, is the declared sodium cyanide price of \$2.40 per kilogram times the measured addition per mg/L in the head tank (the applied addition over the titrator reading, each filtered over 2 h, 1.651 kg/h per mg/L at design), plus the declared destruction price of \$3.50 per kilogram on the fraction 0.7471 of the head-tank cyanide that reaches tank five at the declared defaults, carried by 488.9 m³/h of solution, which totals \$5.24 per hour per mg/L at design. With the curvature term the unconstrained optimum is $u^* = c_0 - (SP_{Au}w' + c_{CN})/(SP_{Au}w'')$, the setpoint at which the next mg/L buys as much gold as it costs. The terminal term, with $T_{lag} = T_d + \tau_{CN} + 5\tau_r$ (7.9 h at design), credits the gold that a late move has still to deliver after the horizon, so that a move held for 24 h is credited 24.2 h of its gold rather than 16.6 h. The move price s_m is \$0.50 per (mg/L)² at each block boundary j.

The tailings target Au_{SP} becomes a soft ceiling, $\hat{Au}_k \le Au_{SP} + \varepsilon_k$, judged at every sixth step ($\mathscr{K} = \{1, 7, \ldots, 55\}$, one point every 2.4 h), each with a slack $\varepsilon_k \ge 0$ of its own priced at $\rho_2 = 10^{10}$ per (g/t)² and $\rho_1 = 10^7$ per g/t. An excess is therefore paid for at the step where it occurs, and an excess that the plant cannot avoid late in the day does not make the same excess free in the near hours. The linear price is 800 times the gold that 1 g/t of tailings over one step is worth at design. Judging every step gives the same net margin to within \$1.7 per hour on every case of Table 7.4 but costs a plan 1.5 ms in Chrome instead of about 0.2 ms. The setpoint is bounded hard below by the outer loop's 40 mg/L and above by the lesser of 400 mg/L and the reach of the pump, its 500 kg/h maximum divided by the measured addition per mg/L, about 303 mg/L at design, above which the head tank cannot follow the setpoint. The first move is bounded hard by the outer loop's own ramp, 3.2 mg/L per decision, and the later blocks are free. The ten inputs and the ten slacks form a quadratic program whose ridge is set to 10^{-6} explicitly, because the solver's default scales with the largest diagonal entry, 2×10^{10} from the ceiling's price, and the resulting ridge of 20 on the absolute setpoint pulls the end of the plan at hour 90 of scenario 11 from 161.8 mg/L to 128.6 mg/L against the economics. The horizon and the move price were set once on scenario 11, where horizons from 8 to 48 h fall within \$4.9 per hour of one another and move prices from 0.1 to 10 times the shipped value within \$1.5, and they are not retuned for any other case.

Acceptance, fallback and cost. The controller accepts a plan when the solver's status is optimal, the arrays have the expected dimensions, every planned setpoint lies within its bounds, the first planned setpoint equals the move sent and lies within the ramp, and the objective value, the slack and the predicted tailings are all finite. Otherwise the soft sensor loop, whose integral has tracked the applied setpoint, takes the setpoint, a fallback mark enters the historian, and the next accepted plan takes it back. The first plan is solved at $t = 0$. Failures forced at six instants of scenario 11 with noise off move the setpoint by at most 0.12 mg/L at the hand-over and by at most the 3.2 mg/L ramp at the take-back. In Chrome a plan takes 211 to 220 µs on average over the 600 plans of scenario 11, 0.055 % of one core at one plant hour per second.

Three models kept apart. The plant is the five-tank model of Section 7.3 with its true leachability, cyanicides, acid and oxygen demands and cyanide profile. The identified models are the cyanide loop's lag and dead time and the five residue lags. The controller's model is (7.4) on the soft sensor's linearization, fed only the measured free cyanide, pH, oxygen, feed grade and throughput, the applied setpoint and addition, the laboratory through the bias, the declared prices and the soft sensor's own constants, held at the plant's declared values, so that an edit of the plant's parameters reaches the controller only through its measurements. Both controllers hold their own copy of seven constants (the rate constant, the two half-rate concentrations, the pH optimum and width, the refractory gold and the ore's specific gravity) at the declared values, so a parameter changed in the simulated plant becomes a model mismatch that the laboratory has to reveal, as it would on an operating circuit. The application's tests confirm that it never reads the leachability, a parcel, the true tailings or the plant state, by source inspection, by seventeen deliberately altered copies of the code that the inspection rejects, and by raising the hidden tank-five residue by thirty percent at an assay instant, after which the bias, the plans and the setpoint stay bit-identical through the next two decisions.

7.5.3 Moving Horizon Estimation of the Bias

The bias b read by the soft sensor loop and by the predictive controller can instead come from a moving horizon estimator (MHE) [23], offered as a second laboratory update beside (7.3) and selected by the app's Lab update switch (Section 1.6.3 introduces the estimator options of the apps). It acts only under those two outer loops, because the other two use no bias. The update of (7.3) stays the application's default and the setting of every scenario but scenario 12, and the soft sensor, the predictive controller and the loops are unchanged, reading the estimate wherever they read the bias of the update. Correcting a leaching soft sensor from its residual against measured data, and choosing how that correction adapts, is an established problem in gold cyanidation [24], and the estimator offers a second answer to it beside the fixed-gain update.

The estimator advances one step per laboratory sample and places each sample at the instant it was cut rather than the instant its assay is reported. A result that arrives 6 h late therefore enters at the step it describes with no delay states, because an estimator that refits a batch of past measurements can place a delayed one where it belongs [25], and the bias in force is the estimate at the newest cut, since a random walk predicted across the six hours in transit stays where it is. Volume 5, Chapter 4 reconciles a residence-time predictor with an analyzer and a laboratory that report late in the same spirit. The model is given in Equation (7.6).

$$x_{k+1} = x_k + (r_{k+1} - r_k) + w_k, \qquad y_k = x_k + v_k \tag{7.6}$$

where x_k is the true tailings grade at cut k in g/t, y_k is the assay of that sample, v_k is the laboratory error, w_k is the change of the bias between consecutive samples, and r_k is the model's tailings grade at the cut, so that $b_k = x_k - r_k$. The reference r_k is not the soft sensor's value at the cut, which (7.3) uses, but that output passed through the five residue lags τ_r of (7.4), the lag ℓ_5 that the predictive controller's estimator integrates at every scan under every outer loop. It is what the model says the tanks had delivered when the sample was cut. Over the last $N = 12$ transitions, which hold thirteen results and span 48 h at the 4 h sampling interval, the estimate solves the quadratic program of Equation (7.7).

$$\begin{aligned} \min_{x_s, w_s, \ldots, w_{k-1}} \quad & \frac{(x_s - \bar{x}_s)^2}{\bar{P}} + \sum_{j=s}^{k-1} \frac{w_j^2}{\sigma_w^2} + \sum_{j=s}^{k} \frac{(y_j - x_j)^2}{\sigma_{lab}^2} \\ \text{subject to } (7.6), \quad & x_j \geq A u_\infty, \quad \left| w_j \right| \leq 3\sigma_w \end{aligned} \tag{7.7}$$

Here $s = k - N$, and the first term is the arrival cost in its filtering form, which carries every older result through $\bar{x}_s$, the estimator's own earlier estimate moved one step through the model, and $\bar{P}$, the a-priori variance of the Kalman recursion run beside it [23]. The laboratory repeatability $\sigma_{lab} =$

0.015 g/t is declared, five percent of the 0.3 g/t design tailings, and the assay noise that a scenario sets is never read. The drift σ_w follows from a rule fixed before any run, that the steady Kalman gain of the random walk in (7.6) equal the default bias gain $\beta = 0.4$. That requires $\sigma_w^2/\sigma_{lab}^2 = \beta^2/(1-\beta)$, so $\sigma_w = 0.00775$ g/t per sample, and $\bar{P}$ starts at its steady value $\sigma_w^2/\beta = 1.5 \times 10^{-4}$ (g/t)^2, where it stays. The bounds are declared as well. The floor $Au_\infty = 0.25$ g/t is the refractory gold of (7.1), below which no tailings can fall, and $3\sigma_w = 0.0232$ g/t per sample is the three-sigma limit of the drift model. That limit shapes the fitted path rather than capping the move that one result causes, so an assay 0.15 g/t high on scenario 12 moves the estimate by 0.0535 g/t with no bound active. No upper bound is set, since the feed grade moves the tailings and a constant ceiling would never bind.

While no bound is active, a moving horizon estimate with this arrival cost equals the Kalman filter's for any window length, and an independent Kalman filter reproduces every update of scenarios 5 and 9 to 12 to within 4.4×10^{-16} g/t. With these weights the gain is 0.4 from the first sample, and the estimator then reduces to (7.3) at $\beta = 0.4$ with r_k in place of the soft sensor's static value at the cut. Handed that static value as its reference, it reproduces the bias update at $\beta = 0.4$ to within 1.2×10^{-15} g/t over 240 h of scenario 12. What the estimator adds to the update is therefore the lagged reference, the two bounds and the window, and Section 7.6.7 removes each in turn.

An update is used only when the solver reports it optimal and the estimate and the window are finite and on or above the floor. Otherwise the update of (7.3), at the operator's β, moves the bias for that result, a fallback mark enters the historian, and the next accepted update takes the bias back. Taking the bias over, on a switch to the estimator, at the start of a run or after a fallback, is ramped at the drift limit per sampling interval, 2.9×10^{-4} g/t per controller scan. An iteration limit, a non-finite estimate and an estimate below the floor, each forced at the 128 h result of scenario 12 with noise off, are rejected with the update's bias reproduced exactly, and the estimator retakes the bias over 26 scans while the cyanide setpoint moves at most 0.4 mg/L per scan, the outer loop's own ramp. The estimator reads the assay, the time its sample was cut, the model's tailings from its own history and its frozen constants, and never the true tailings, the ore or a plant parameter. Source inspection in the application's tests rejects sixteen deliberately altered copies that break this rule, and an edit of each of 42 plant parameters and 11 disturbance settings at 30 h, under both outer loops, never moves the estimate before it moves a measurement.

Each update is a quadratic program in thirteen unknowns, the first state of the window and twelve drifts, solved by the same dense active-set solver as the plan, once per 4 h of plant time. On the browser's frame loop at one plant hour per second an update takes 227 to 300 µs on average over two measurements, under 0.01 % of one core, against 0.11 % to 0.12 % for the plans of the predictive controller. At the top speed of 2000 plant hours per second the estimator takes 0.8 % to 1.7 % of a core and the plans 24.5 %, a load that predates the estimator and keeps the page below the requested speed with either update. Every timing was taken while other work loaded the host.

7.6 Results

The app presents thirteen scenarios in order: the train walked by hand, the three reagent handles ranked, the three secondary loops, the outer loop on the assay and then on the soft sensor, three ore disturbances, the two failure modes of the soft sensor, the predictive controller, the moving horizon estimator and a graded challenge. Each has a completion rule judged on the run, as Section 1.6.4 describes, and the challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt. Reset restarts a case from $t = 0$ with the reader's settings kept, while the reload button restores the scenario's own. Every result below is a scenario as it loads, with its own seed, events and noise, unless the text names a change.

7.6.1 Outer Loop on the Assay and on the Soft Sensor

Scenarios 4 and 5 apply the same disturbance to the same plant, a feed grade step from 2.8 to 3.2 g/t at 10 h, with the outer loop closed first on the assay and then on the soft sensor, and run for 200 h.

Both configurations hold the tailings grade. The peak is 0.3018 g/t on the assay and 0.2991 g/t on the soft sensor, and both return to the 0.295 g/t target at a recovery of 90.78 %. The soft sensor loop moves the cyanide setpoint past 122 mg/L 0.26 h after the step, because the feed grade is one of its inputs, and the assay loop 10.26 h after it, once the first assay containing the richer ore has been reported. The difference also appears in the actuator. The assay loop moves the cyanide addition through 428 kg/h of cumulative travel over the run and the soft sensor loop through 296, a factor of 1.4, because the assay loop corrects a six-hour-old error in four-hour steps and overshoots its final setpoint of 140 mg/L to 150 mg/L, while the soft sensor loop approaches it without overshoot.

That comparison supports treating an inferential estimate as a measurement and judging it on the same terms as one.

7.6.2 Unmeasured Ore Disturbances

Scenario 6 reduces the ore leachability to 0.82 at 12 h without changing any quantity the soft sensor reads. The estimate is computed from the reagent concentrations and the feed grade, all of which are held on setpoint, so it stays at 0.295 g/t while the true tailings grade reaches 0.320 g/t by 20 h and peaks at 0.3249 g/t at 26 h. The first assay that contains the change is reported at 20 h, and the cyanide setpoint first moves at 20.1 h, as it does in the same case with the outer loop on the assay. For this class of disturbance the soft sensor is no faster than the assay. The loop recovers at a cost of 745 kg/h of cyanide travel by 110 h, when the scenario's next disturbance arrives, two and a half times the travel of the whole feed grade case of Scenario 5.

> **TRY IT IN THE APP** **6 · The Disturbance the Soft Sensor Cannot See**
>
> Load the scenario and press Run with the outer loop on the soft sensor, as it loads. Watch the three tailings values after the harder ore arrives at 12 h. The soft sensor stays at 0.295 g/t while the true grade reaches 0.320 g/t by 20 h, and the cyanide setpoint does not move until the first assay that holds the change is reported.

Scenario 9 introduces the same effect deliberately, by biasing the soft sensor fourteen percent optimistic at 12 h. The estimate and the assay diverge by 0.0403 g/t at that instant, the bias update brings the true tailings grade back within 0.002 g/t of target by 79 h, and the cyanide travel over the 240 h run reaches 2193. The loop returns to setpoint only because the assay remains available to correct it. With $\beta = 0$ the tailings settle at 0.354 g/t, a recovery of 87.4 %, while the estimate, which the loop holds near its target, reads 0.297 g/t at the 140 h recalibration.

A soft sensor removes dead time from the loop and introduces a model error that the loop cannot distinguish from a process change. Whether the exchange is favorable depends on the bias update, and the bias update is limited by the assay interval.

Pause. Try Exercise 7.1 before reading on.

7.6.3 Effect of the Bias Update Gain

Scenario 10 sets $\beta = 1$, so that each assay replaces the bias outright, and applies the feed grade step of Scenario 5 at 12 h with 0.02 g/t of noise on the assay.

Table 7.2 quantifies the failure mode. Each assay is six hours old and carries the laboratory's error, so replacing the bias with it steps the estimate by that error and by whatever the plant did six hours earlier, and the outer controller passes each step on to the cyanide setpoint. With the scenario's noise the setpoint at $\beta = 1$ wanders over 101 mg/L against 53 mg/L at $\beta = 0.15$, the actuator travel is fourteen percent higher than at 0.35 and sixty-three percent higher than at 0.15, and the peak tailings grade is 0.0055 g/t and 0.0069 g/t worse. The reagent consumption does not change. Without assay noise a gain of one still doubles the travel with no measurable change in the tailings, so the dead time alone costs actuator wear and the noise costs control. Over six other noise seeds the travel at $\beta = 1$

Table 7.2 The feed grade step of Scenario 10 under three bias gains over 240 h, with the scenario's assay noise and without it. The setpoint range is taken from 24 h onward.

β	Assay noise (g/t)	CN travel (kg/h)	NaCN (kg/t)	CN setpoint (mg/L)	Peak tailings (g/t)
1.0	0.02	5472	0.512	78–179	0.3158
0.35	0.02	4803	0.519	93–171	0.3103
0.15	0.02	3356	0.517	101–154	0.3089
1.0	0	425	0.568	126–140	0.2990
0.35	0	220	0.568	130–140	0.2986

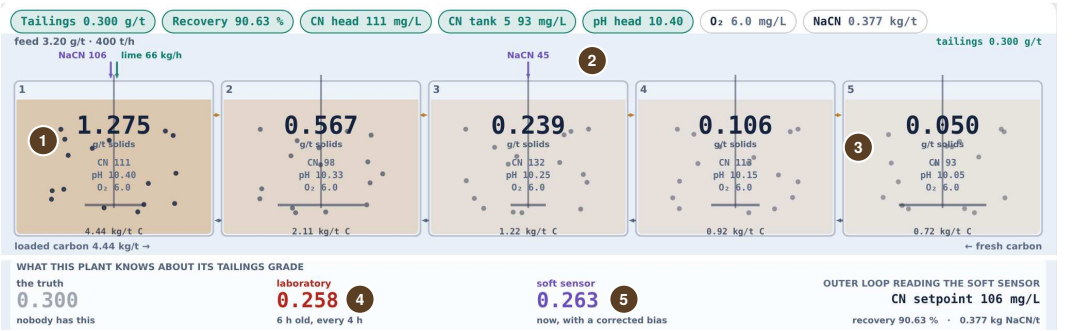


Figure 7.2 Scenario 10 at 170 h, cut to the plant view and the tailings panel, with the bias update replacing the bias outright at every noisy assay. (1) The head tank, where the sodium cyanide and the lime enter and the free cyanide is held. (2) The second cyanide addition, to tank three. (3) Tank five, the last stage of leaching. The three values of the tailings grade are displayed together: the value computed by the model at 0.300 g/t, (4) the laboratory assay at 0.258 g/t reported six hours after the sample was cut, and (5) the soft sensor at 0.263 g/t, 0.037 g/t below the truth. The free cyanide loop is holding its moving setpoint, now 106 mg/L.

exceeds that at 0.35 by 14 to 23 percent and that at 0.15 by 56 to 94 percent, and the peak is higher at $\beta = 1$ than at 0.15 on every seed.

The configuration the application supplies uses a bias gain near a third, which filters the assay across roughly three samples and gives the correction a time constant of about 9 h. Setting the update gain is therefore a filter design problem, and the sampling interval and the assay noise rather than the controller tuning set the answer.

7.6.4 Copper and Sulfide Parcels

Scenario 7 sends a copper parcel through the circuit from 12 h to 120 h, multiplying the cyanicide term by 2.2. The head-tank cyanide loop holds the free cyanide within 10.2 mg/L of its setpoint, but the free cyanide in tank five falls from 89.7 mg/L to 53.2 mg/L by 28 h, so the last stage of leaching slows while every setpoint is met and the tailings grade rises to 0.3084 g/t. The soft sensor loop recovers the tailings by raising the cyanide setpoint to 160 mg/L, which takes the addition to 311 kg/h and the reagent intensity from 0.495 kg/t to 0.777 kg/t. At steady state with the setpoint held at 120 mg/L the same parcel would cost $466 per hour of operating margin. Moving half of the addition to tank three before the run holds tank five at or above 76.3 mg/L through the parcel and the tailings at or below 0.299 g/t. Holding tank one on setpoint says nothing about tank five, which is where the last of the gold has to dissolve.

Table 7.3 Outer loop ablation over the graded week of 200 h. Every row after the first uses the suggested secondary tuning and setpoints (pH 10.5, 6.5 mg/L of oxygen). Above is the time the tailings spend above 0.40 g/t, CN error is the largest departure of the free cyanide from its setpoint, and the last column gives the application's verdict.

Outer loop	Above (h)	Peak (g/t)	NaCN (kg/t)	CN error (mg/L)	Met
As shipped, CN at 120 mg/L	80.9	0.4368	0.529	8.5	no
None, CN at 120 mg/L	**0.0**	0.3982	0.535	6.1	yes
None, CN at 125 mg/L	**0.0**	0.3929	0.557	6.2	yes
Soft sensor, target 0.36	21.6	0.4521	0.525	10.1	yes
Soft sensor, target 0.33	**0.0**	0.3911	0.716	18.2	no
Assay, target 0.36	52.7	0.5232	0.559	9.7	yes
Soft sensor, $\beta = 0$	172.0	0.5887	0.237	9.5	no

TRY IT IN THE APP **7 · A Copper Parcel Eats the Cyanide**

Run the scenario and watch the dashed tank-five trace on the cyanide loop. The head tank stays within 10.2 mg/L of its setpoint, yet tank five falls from 89.7 to 53.2 mg/L by 28 h. Then reset, open Show Model, move the share of the cyanide added to the head tank from 0.70 to 0.50 and run again: tank five now stays at or above 76.3 mg/L.

Scenario 8 sends a sulfidic parcel at 12 h, which multiplies the oxygen demand by 2.4 and, 0.1 h later, the acid demand by three. The air command reaches its 250 % ceiling at 13.7 h and the dissolved oxygen settles at 5.32 mg/L against its 6.0 mg/L setpoint. The lime addition rises from 66.3 kg/h to 82.3 kg/h and holds the head tank within 0.04 of pH 10.40, but the pH in tank five falls below 9.8 at 15.0 h and reaches 9.50, where the alkalinity is exhausted, by 17.3 h. The tailings grade reaches 0.334 g/t at 29 h, and the outer loop recovers it only by raising the cyanide setpoint to 367 mg/L, with the sodium cyanide pump at its 500 kg/h limit from 115 h and 1.109 kg/t consumed over the 240 h run. The soft sensor reads the head-tank pH, which is held, and so reads 0.034 g/t optimistic at 28 h. A loop defends the point at which it measures, and where the variable matters along the whole train a single addition point is not a substitute for a tighter loop.

Pause. Try Exercise 7.3 before reading on.

7.6.5 Outer Loop Ablation Over the Graded Week

The graded challenge runs 200 h against harder ore at 20 h (leachability 0.80), a copper parcel at 60 h, an acid parcel at 105 h and a richer feed of 3.3 g/t at 150 h, with noise on the titrator, the pH electrode and the assay, starting on local control with cautious tunings. It is scored on the free cyanide within ±12 mg/L of its setpoint and the head-tank pH at or above 9.8 at every integration step, the tailings below 0.40 g/t immediately before 60 h, 105 h and 150 h and at 200 h, and the cumulative consumption under 1.05 kg/t at the end.

Table 7.3 is the central measurement of this chapter, with the hours above 0.40 g/t and the peak taken at the integration step. With the settings as shipped the circuit fails, because the tailings are 0.409 g/t before the richer feed and 0.433 g/t at the end, while the cyanide (8.5 mg/L), pH (10.32) and reagent (0.529 kg/t) clauses hold. Figure 7.3 draws the hours above the limit and the reagent of every row.

With the suggested secondary tuning, no outer loop and the free cyanide held at its design 120 mg/L, the tailings grade never exceeds 0.40 g/t and the circuit uses 0.535 kg/t of sodium cyanide. The soft sensor cascade at a target of 0.36 g/t also passes as scored, at 0.525 kg/t, two percent less reagent, but it spends 21.6 h above the limit after the harder ore and peaks at 0.4521 g/t, because the

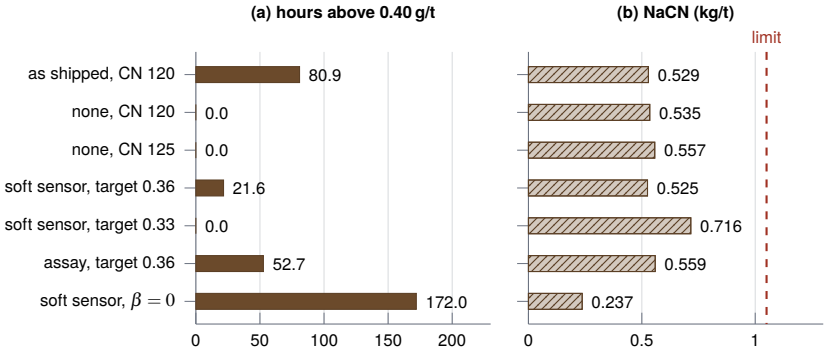


Figure 7.3 The graded week of Table 7.3 as a chart. A fixed cyanide setpoint with no outer loop spends no hours above the tailings limit, while the recommended soft sensor cascade saves a little reagent and spends 21.6 h above it. The loop with no bias update uses the least reagent of any row and the most hours above the limit.

challenge judges the tailings only at its checkpoints. As scored, its worst cyanide error is 10.14 mg/L, the worst checkpoint tailings 0.379 g/t, the lowest pH 10.40 and the consumption 0.525 kg/t, which passes with 1.9 mg/L and 0.021 g/t to spare, and the app gives the same run whether the suggested settings are revealed before Run or carried through Reset. Tightening the target to 0.33 g/t keeps the tailings under the limit throughout and takes the consumption to 0.716 kg/t, a third more than the fixed setpoint, and the richer feed then drives the sodium cyanide pump to its 500 kg/h limit for 7.6 h, so the free cyanide ends 18.2 mg/L below its setpoint and the cyanide clause fails. On the shipped noise seed no target from 0.300 to 0.380 g/t, in steps of 0.005, both keeps the tailings under the limit and passes. Targets up to 0.330 break the cyanide clause, and targets from 0.335 to 0.375 pass with 5.4 h to 43.1 h above the limit. Every row with the suggested tuning moves the cyanide addition through 10 031 kg/h to 12 670 kg/h, most of it in answer to titrator noise.

The comparison depends partly on the noise. Over six other noise seeds the fixed setpoint at 120 mg/L exceeds the limit once, for 5.1 h with a peak of 0.4004 g/t, while 125 mg/L stays at or below 0.395 g/t on every seed at 0.556 kg/t to 0.557 kg/t. The cascade at 0.36 spends 21 h to 38 h above the limit on every seed, and at 0.33 it both holds the limit and passes on four of the six. A fixed setpoint must also stay within 12 mg/L of the 120 mg/L it starts from, because the cyanide clause counts the error that a setpoint step itself creates, as the challenge text states.

The last row shows the behavior of the outer loop when its estimate is not corrected. At $\beta = 0$ the soft sensor remains optimistic through the ore change, the controller acts as though the tailings grade were acceptable, it reduces the cyanide consumption to 0.237 kg/t, and the circuit spends 172 h of the 200 h week above the tailings limit. That row is the argument for the bias update and against reading the reagent number in isolation.

> **+ DID YOU KNOW? The Cheapest Week Was the Worst**
>
> From the simulator: with no bias update, the soft sensor loop of Table 7.3 uses the least reagent of any row, 0.237 kg/t, and keeps the tailings above the 0.40 g/t limit for 172 h of the 200 h week.

The excursion of the recommended cascade has a specific cause. Its target lies 0.065 g/t above the tailings grade at the start, so at $K_c = 2500$ (mg/L)/(g/t) the proportional term alone demands a cyanide setpoint about 160 mg/L below the one the loop inherits, and the rate limit carries the setpoint down to 48 mg/L by 10 h and to its 40 mg/L floor by 18.6 h. The harder ore at 20 h then changes a property the soft sensor cannot see (Section 7.6.2), so the loop learns of it only from the assays, and the tailings pass 0.40 g/t at 25.8 h and return under it at 47.4 h. The setpoint stays low because

the error stays large, not because the integral winds up: at 18.6 h the integral part of the output is 76 mg/L and the proportional part cancels it on a soft sensor reading 0.03 g/t below the target. As counterfactual model runs, shortening the tracking time T_t of Equation (1.13) from 14 h to 2 h leaves 20.4 h above the limit, and only a one-scan tracking time removes the excursion on the shipped seed (peak 0.399 g/t, 0.491 kg/t), by discarding the proportional demand of every rate-limited move. That variant delays the Scenario 5 setpoint past 122 mg/L from 10.26 h to 12.91 h and still exceeds the limit for 22 h to 32 h on four of the six other seeds, so the application retains $T_t = \tau_I$ to preserve proportional response through rate-limited moves. A fixed setpoint at 120 mg/L instead runs the tailings near 0.29 g/t before the harder ore and stays under the limit through all four disturbances, with 0.002 g/t to spare after the richer feed. The value of an outer loop on an operating circuit comes from slow drift in the properties its estimator sees, and this week opens with a step in one that it does not see.

7.6.6 Model Predictive Control Against the Soft Sensor Loop

Table 7.4 runs scenarios 5 to 11 and the challenge, each from its own settings with its own seed, events, noise and duration, under four outer loops. The first is the structure the case ships with, which is the soft sensor loop, or local control on the challenge as shipped. The second is the predictive controller of Section 7.5.2 with the settings of scenario 11 and nothing retuned. The third gives the soft sensor loop the predictive controller's economics without its prediction, because every 0.4 h a setpoint optimizer moves the loop's tailings target to the steady optimum of the same soft model at the same prices, capped at the case's own target. The fourth is local control at the best fixed cyanide setpoint of a scan from 80 to 400 mg/L in steps of 10, chosen after the run and therefore a bound that no causal controller attains. No row receives a preview, and every row is scored on the application's net margin, gold at $78 per gram less sodium cyanide, lime, air and cyanide destruction.

The predictive controller earns more than the shipped structure on every case except scenario 6, from $10.3 per hour on scenario 8 to $188.6 on scenario 10 and $904.1 on the suggested challenge settings, and it keeps the true tailings under the target throughout scenarios 5, 9 and 10, where the loop spends 188.1, 74.3 and 200.1 h above it. The hours over the target favor the plan by construction, because a regulator that holds the target sits on it and exceeds it often by a small margin, at most 0.0041 g/t over 188.1 h on scenario 5, while the plan settles below it, so the depth is the fairer comparison. Almost all of the gain belongs to the economic target rather than to the prediction. The soft sensor loop with an optimized target comes within $12.9 per hour of the predictive controller on scenario 5, $13.9 on scenario 7, $1.6 on scenario 9, $1.9 on scenario 11 and $51.1 on the suggested challenge, and it earns $19.7 more on scenario 6, $2.9 more on scenario 8 and $1.7 more on scenario 10. The rule that a tailings target encodes is too lean in cyanide at the declared prices. After the richer feed of scenario 11 the soft sensor loop holds 0.295 g/t at 139 mg/L, while the plan settles at 162 mg/L with the tailings at 0.289 g/t, within 0.2 mg/L of the unconstrained optimum of 161.6 mg/L evaluated by hand from (7.5) at hour 90.

On scenario 8 the tailings ceiling cannot be met, and every judged step at which it is broken carries its own price. The plan therefore asks for all the cyanide the pump can deliver: the setpoint first reaches 300 mg/L at 44.9 h and stays between 309.8 and 310.5 mg/L from 60 h to the end, equal to the plan's own final setpoint, with the pump at its limit for 67.4 h against 124.9 h for the soft sensor loop, and the true tailings still end at 0.298 g/t. A single slack shared by the whole horizon, the form the controller used before, priced only the worst step there, held the applied setpoint between 267.6 and 275.3 mg/L from 60.1 h while its own final setpoint stood at least 35.5 mg/L higher, and earned $75.3 per hour more, because the cyanide that the ceiling demands buys less gold than it costs.

The fixed setpoint chosen in hindsight beats the predictive controller on every case, by $8.3 to $72.7 per hour, by $132.4 on scenario 8 and by $224.3 on the suggested challenge, where 230 mg/L fails the cyanide clause by its own initial step of 110 mg/L. The best fixed setpoints lie at 170 to 230 mg/L, above the plan's, because the plant's tailings respond to cyanide 1.25 to 2.0 times as strongly as the soft model says (Section 7.5.2), and the bias that corrects the model's level cannot

Table 7.4 Outer Loops Compared on Net Margin, Tailings and Cyanide. Over Is the Time the True Tailings Spend Above the Target in Force (0.36 g/t on the Challenge's Suggested Settings, 0.295 g/t Elsewhere) and Depth the Largest Excess, Discharge the Time Tank Five's Free Cyanide Spends Above 200 mg/L, and Travel the Summed Movement of the Cyanide Setpoint. Fixed Setpoints Are Chosen in Hindsight

Case	Outer loop	Net ($/h)	Over (h)	Depth (g/t)	NaCN (kg/t)	Discharge (h)	Travel (mg/L)	Plans, fallbacks
5, richer feed	Soft sensor (shipped)	89 170.7	188.1	0.0041	0.566	0.0	54	—
200 h	MPC	89 232.9	0.0	0	0.663	0.0	39	500, 0
	Soft sensor, optimized target	89 220.0	0.0	0	0.638	0.0	85	—
	Fixed 180 mg/L	89 241.2	0.0	0	0.739	0.0	0	—
6, harder ore	Soft sensor (shipped)	83 697.9	148.8	0.0299	0.852	29.5	389	—
240 h	MPC	83 679.5	94.3	0.0167	0.942	101.6	806	600, 0
	Soft sensor, optimized target	83 699.2	147.6	0.0221	0.891	29.3	355	—
	Fixed 200 mg/L	83 752.2	145.1	0.0137	0.816	0.0	0	—
7, copper parcel	Soft sensor (shipped)	77 324.5	84.7	0.0134	0.612	0.0	163	—
240 h	MPC	77 382.9	46.1	0.0034	0.688	0.0	415	600, 0
	Soft sensor, optimized target	77 369.0	103.4	0.0066	0.670	0.0	135	—
	Fixed 180 mg/L	77 405.1	0.0	0	0.796	0.0	0	—
8, sulfidic parcel	Soft sensor (shipped)	76 330.2	228.0	0.0390	1.109	141.0	337	—
240 h	MPC	76 340.5	224.1	0.0257	1.168	186.2	190	600, 0
	Soft sensor, optimized target	76 343.4	225.8	0.0333	1.112	139.4	336	—
	Fixed 220 mg/L	76 472.9	220.2	0.0127	0.895	0.0	0	—
9, structural error	Soft sensor (shipped)	77 405.9	74.3	0.0222	0.507	0.0	475	—
240 h	MPC	77 517.2	0.0	0	0.631	0.0	262	600, 0
	Soft sensor, optimized target	77 515.6	0.0	0	0.612	0.0	247	—
	Fixed 170 mg/L	77 537.9	0.0	0	0.695	0.0	0	—
10, $\beta = 1$	Soft sensor (shipped)	89 045.4	200.1	0.0208	0.512	0.0	1562	—
240 h	MPC	89 234.0	0.0	0	0.710	0.0	849	600, 0
	Soft sensor, optimized target	89 235.7	0.0	0	0.698	0.0	849	—
	Fixed 180 mg/L	89 243.0	0.0	0	0.739	0.0	0	—
11, richer feed,	Soft sensor	89 000.1	157.0	0.0245	0.695	0.0	303	—
then harder ore	MPC (shipped)	89 025.8	38.3	0.0168	0.783	28.3	452	600, 0
240 h	Soft sensor, optimized target	89 023.9	69.1	0.0179	0.749	0.0	278	—
	Fixed 190 mg/L	89 054.0	69.1	0.0101	0.778	0.0	0	—
Challenge,	Soft sensor (suggested)	79 144.1	106.0	0.0921	0.525	0.0	1299	—
suggested	MPC	80 048.2	16.0	0.0064	0.686	0.0	344	500, 0
settings,	Soft sensor, optimized target	79 997.1	26.3	0.0102	0.662	0.0	1128	—
200 h	Fixed 230 mg/L	80 272.5	0.0	0	1.012	0.0	0	—
Challenge,	Local (shipped)	78 994.9	180.0	0.1418	0.529	0.0	0	—
shipped	MPC	79 888.6	178.6	0.0590	1.042	0.0	370	500, 0
settings,	Soft sensor, optimized target	79 497.4	179.8	0.0923	0.668	0.0	245	—
200 h	Fixed 260 mg/L	79 907.9	173.9	0.0622	1.078	0.0	0	—

correct its slope. At hour 90 of scenario 11 the plant's own steady optimum is 178 mg/L against the plan's 161.8 mg/L. The soft sensor's structural error, inherited by the controller's model, is therefore what the hindsight setpoint exploits. In a counterfactual run that hands the plan the plant's true steady slope and curvature at every decision, information no operating circuit possesses, the predictive controller comes within $1.5 per hour of the hindsight setpoint on scenario 5, $7.4 on scenario 9, $10.2 on scenario 10 and $22.2 on scenario 11, and passes it on scenario 7 by $3.6 and on the suggested challenge by $14.2. Gaps of $65 and $126 per hour remain on scenarios 6 and 8, because the hindsight setpoint knows the unseen ore and the parcel from the first hour.

Scenario 11 applies the richer feed of scenario 5 at 12 h and a harder ore of leachability 0.85 from 100 h to 170 h, runs 240 h, and opens on the predictive controller, as shown in Figure 7.4. Neither structure sees the harder ore before the laboratory does. The loop first moves at 108.2 h and the

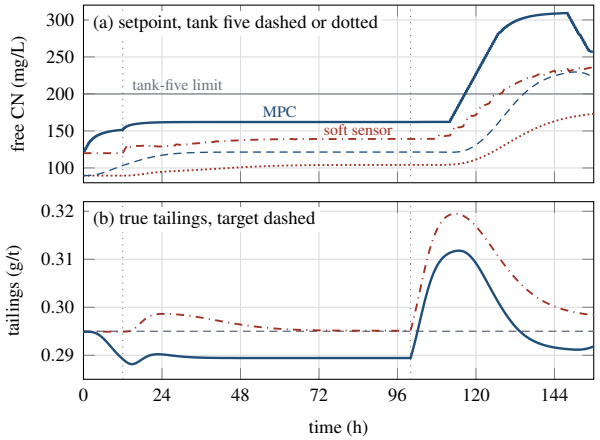


Figure 7.4 Scenario 11 from the start to 156 h under the predictive controller (solid and dashed lines, drawn blue) and the soft sensor loop (dash-dotted and dotted lines, drawn red). (a) The free cyanide setpoint (solid for the predictive controller, dash-dotted for the loop) and tank five's free cyanide (dashed and dotted) against its 200 mg/L discharge limit (thin gray line). (b) The true tailings (solid and dash-dotted) against the 0.295 g/t target (gray dashed). The vertical dotted lines mark the richer feed at 12 h and the harder ore at 100 h. The plan's setpoint peaks at 309 mg/L, the reach of the pump, at 147.7 h, and at 156 h it is back at 257 mg/L with the true tailings at 0.292 g/t, under the target since 133.5 h, while the loop's setpoint is 236 mg/L with the tailings at 0.298 g/t.

plan at 112.1 h, but once the bias grows the ceiling binds and the plan moves faster, past 200 mg/L at 116.5 h against 127.4 h for the loop, and further, to a peak of 309.2 mg/L at 147.7 h, at the reach of the pump, against the loop's 243.5 mg/L at 180.1 h. The true tailings are already back under the target from 133.5 h, while the soft estimate with its bias still reads 0.2953 g/t at 156 h, because the bias carries assays of samples cut six hours earlier and the model's slope on this ore is 1/1.84 of the plant's. Over the 240 h the plan earns \$89,025.8 per hour against \$89,000.1 on 0.783 kg/t of sodium cyanide against 0.695 kg/t, and within the harder ore, from 100 to 170 h, it earns \$89,158.3 against \$89,153.4 while the optimized target earns \$89,197.8.

The plan's price cannot be had by lowering the loop's target once. With the soft sensor loop's target lowered by hand to the 0.289 g/t at which the plan settles, scenario 11 earns \$89,003.5 per hour over the 240 h on 0.816 kg/t of sodium cyanide, against the plan's 0.783 kg/t, and so recovers \$3.4 of the plan's \$25.7. Local control at fixed setpoints of 160, 190 and 220 mg/L earns \$89,029.1, \$89,054.0 and \$89,023.1. The gain requires a target that moves with the conditions, as the optimized target of Table 7.4 does at \$89,023.9.

The predictive controller loses scenario 6 by \$18.4 per hour, \$83,679.5 against \$83,697.9, although its tailings spend 94.3 h above the target against 148.8 h at a depth of 0.0167 g/t against 0.0299 g/t. The harder ore at 12 h is invisible to the model until the assays report it, and once the bias grows the plan holds up to 309 mg/L from 44.0 to 74.2 h with the true tailings under the target, for the same two reasons as in scenario 11, the old samples in the bias and a slope half the plant's. After the scenario's richer feed at 110 h the plan runs at the reach of the pump, 309.8 mg/L at most, and over 110 to 160 h it earns \$88,982.8 per hour against the loop's \$89,044.1. The true slope does not remove the loss, since it raises scenario 6 by only \$8.1 per hour.

The plan also breaks a limit it does not model. The 200 mg/L discharge limit on tank five's free cyanide is an alarm that no shipped loop holds and the scorecard does not count, and the objective prices the cyanide that reaches tank five only through its destruction cost. Under the predictive controller tank five exceeds the limit for 28.3 h on scenario 11, 101.6 h on scenario 6 and 186.2 h on scenario 8, against 0, 29.5 and 141.0 h under the soft sensor loop, with peaks of 230, 233 and

211 mg/L, and on every other case neither structure reaches it. That discharge is a cost of the predictive controller as configured.

Under the suggested challenge settings the plan passes, with the free cyanide within 9.9 mg/L of its setpoint, the checkpoint tailings at or below 0.365 g/t, the pH at or above 10.40 and 0.69 kg/t of sodium cyanide, and it spends 16.0 h above its 0.36 g/t target against 106.0 h for the cascade of Table 7.3. With the shipped settings, whose cyanide loop keeps the cautious tuning, the plan moves the setpoint faster than that loop follows and the free cyanide strays 77.2 mg/L from it, so the challenge fails on the cyanide clause instead of the tailings clause on which local control fails.

7.6.7 Moving Horizon Estimation Against the Bias Update

Table 7.5 runs the soft sensor's two failure modes, the structural error of scenario 9 and the over-fast update of scenario 10, under the soft sensor loop and the predictive controller, each case from its own settings, seed, noise and 240 h duration. Row 1 is the bias update at the case's gain ($\beta = 0.4$ on scenario 9 and 1 on scenario 10). Row 2 is the estimator of Section 7.5.3 with the settings of scenario 12, retuned for no case. Row 3 is the bias update retuned once by the estimator's own rule, $\beta = 0.4$. The rows marked 4 remove one mechanism at a time: the bounds (4a), the window, cut to one transition, the smallest the solver accepts (4b), and the lagged reference, replaced by the soft sensor's value at the cut (4d). Row 4c is a bias update at $\beta = 0.4$, written without the estimator and fed the lagged reference, and row 4e gives the shipped update at the case's own gain that same reference.

No bound binds in any of these runs, so removing the bounds (4a), cutting the window to one transition (4b), or both, leaves each run unchanged to every printed digit, and the plain bias update at $\beta = 0.4$ on the lagged reference (4c) reproduces the estimator on every case. As shipped, the estimator is therefore a bias update at $\beta = 0.4$ against a lagged reference, and every difference from row 1 comes from that gain and that reference. Given the static reference instead (4d), it is row 3.

On scenario 10 under the soft sensor loop the estimator earns \$89,074.1 per hour against \$89,045.4 for the shipped update at $\beta = 1$, a gain of \$28.7, of which \$25.1 belongs to the gain, since the update at 0.4 earns \$89,070.5, and \$3.6 to the lagged reference. It also halves the soft sensor's error (0.0100 against 0.0209 g/t) and cuts the setpoint travel by a third (1062 against 1575 mg/L), while the shipped gain on the lagged reference (4e) stays close to row 1 at \$89,051.6. Under the predictive controller the estimator earns \$89,225.4 against \$89,234.0, with half the error (0.0106 against 0.0196 g/t) and less than half the travel (409 against 849 mg/L), and row 3 earns \$89,222.4, so a steadier bias is worth no margin to the plan on this case.

On scenario 9 the estimator loses. The soft sensor's fourteen percent error is the sensor's own and nothing in the plant moves when it appears, but the lagged reference treats it as though it had passed through the tanks and corrects it with their lag. Under the soft sensor loop the estimator earns \$77,386.3 per hour against \$77,405.9, the true tailings spend 127.7 h above the target against 74.4 h and peak at 0.3266 g/t against 0.3172 g/t, the highest of any outer loop on this case, and the error rises from 0.0119 to 0.0144 g/t. Under the predictive controller it earns \$77,507.7 against \$77,517.2, and tank five's free cyanide passes the 200 mg/L discharge limit for 2.1 h, with a peak of 200.5 mg/L, where no shipped structure reaches the limit on this case.

Judged as an estimator, the lagged reference helps where the soft sensor's own inputs move and hurts where its error is its own. Over the scenarios outside Table 7.5, the soft sensor's error under its own loop falls from 0.0008 to 0.0004 g/t on the feed grade step of scenario 5 and by at most 0.0002 g/t on scenarios 6, 7, 8 and 11, rises from 0.0162 to 0.0177 g/t on scenario 12, and falls from 0.0300 to 0.0210 g/t on the challenge's suggested settings. The estimator's own estimate of the tailings at each cut, which has the assays on both sides of it, is closer still, 0.0053 g/t on scenario 9 and 0.0109 g/t on scenario 10 under the soft sensor loop. None of these errors is ever fed back.

The window and the bounds act only with a noisier laboratory. With the assay noise raised to 0.05 g/t at $\beta = 0.4$, which no scenario ships, the floor is active in 28 of the 59 updates of scenario 9 and 17 of scenario 10 under the predictive controller. There it cuts the largest error of the soft sensor

Table 7.5 The Bias Update and the Moving Horizon Estimator on the Soft Sensor's Two Failure Modes, 240 h. Over Is the Time the True Tailings Spend Above the 0.295 g/t Target, Counted at Every Integration Step, Peak the Largest True Tailings, Tank Five the Largest Free Cyanide in Tank Five and Discharge the Time It Spends Above 200 mg/L, Error the RMS Error of the Soft Sensor With Its Bias Against the True Tailings After 12 h, and Travel the Summed Movement of the Cyanide Setpoint at Every Step. No Bound Is Active in Any of the 59 Updates of Any Row

Case	Loop	Row	Net ($/h)	Over (h)	Peak (g/t)	Tank five (mg/L)	Discharge (h)	Error (g/t)	Travel (mg/L)
9, structural error	Soft sensor	1 bias update, $\beta = 0.4$	77 405.9	74.4	0.3172	126.2	0.0	0.0119	480
		2 estimator	77 386.3	127.7	0.3266	147.1	0.0	0.0144	505
		3 bias update, $\beta = 0.4$	77 405.9	74.4	0.3172	126.2	0.0	0.0119	480
		4a no bounds	77 386.3	127.7	0.3266	147.1	0.0	0.0144	505
		4b one transition	77 386.3	127.7	0.3266	147.1	0.0	0.0144	505
		4c update on lagged reference	77 386.3	127.7	0.3266	147.1	0.0	0.0144	505
		4d static reference	77 405.9	74.4	0.3172	126.2	0.0	0.0119	480
	Predictive	1 bias update, $\beta = 0.4$	77 517.2	0.0	0.2950	173.0	0.0	0.0111	262
		2 estimator	77 507.7	0.0	0.2950	200.5	2.1	0.0143	327
		3 bias update, $\beta = 0.4$	77 517.2	0.0	0.2950	173.0	0.0	0.0111	262
		4a no bounds	77 507.7	0.0	0.2950	200.5	2.1	0.0143	327
		4b one transition	77 507.7	0.0	0.2950	200.5	2.1	0.0143	327
		4c update on lagged reference	77 507.7	0.0	0.2950	200.5	2.1	0.0143	327
		4d static reference	77 517.2	0.0	0.2950	173.0	0.0	0.0111	262
10, $\beta = 1$	Soft sensor	1 bias update, $\beta = 1$	89 045.4	200.1	0.3158	124.4	0.0	0.0209	1575
		2 estimator	89 074.1	195.1	0.3098	119.1	0.0	0.0100	1062
		3 bias update, $\beta = 0.4$	89 070.5	197.0	0.3102	120.5	0.0	0.0104	1070
		4a no bounds	89 074.1	195.1	0.3098	119.1	0.0	0.0100	1062
		4b one transition	89 074.1	195.1	0.3098	119.1	0.0	0.0100	1062
		4c update on lagged reference	89 074.1	195.1	0.3098	119.1	0.0	0.0100	1062
		4d static reference	89 070.5	197.0	0.3102	120.5	0.0	0.0104	1070
		4e $\beta = 1$ on lagged reference	89 051.6	203.5	0.3122	123.6	0.0	0.0198	1558
	Predictive	1 bias update, $\beta = 1$	89 234.0	0.0	0.2950	157.6	0.0	0.0196	849
		2 estimator	89 225.4	0.0	0.2950	159.4	0.0	0.0106	409
		3 bias update, $\beta = 0.4$	89 222.4	0.0	0.2950	176.2	0.0	0.0108	417
		4a no bounds	89 225.4	0.0	0.2950	159.4	0.0	0.0106	409
		4b one transition	89 225.4	0.0	0.2950	159.4	0.0	0.0106	409
		4c update on lagged reference	89 225.4	0.0	0.2950	159.4	0.0	0.0106	409
		4d static reference	89 222.4	0.0	0.2950	176.2	0.0	0.0108	417
		4e $\beta = 1$ on lagged reference	89 234.3	0.0	0.2950	157.6	0.0	0.0195	876

from 0.0656 to 0.0469 g/t on scenario 9 and from 0.0653 to 0.0469 g/t on scenario 10, against the same estimator without bounds, and the net margin does not change at the first decimal. Under the soft sensor loop the floor binds three times on scenario 9 and twice on scenario 10 and changes the rms error by less than 0.001 g/t.

Scenario 12 sets both effects in one run. The soft sensor loop runs at $\beta = 0.4$ with 0.02 g/t of assay noise, the feed grade rises from 2.8 to 3.2 g/t at 12 h, and the soft sensor reads fourteen percent optimistic from 120 h, as shown in Figure 7.5. Between 12 and 120 h the estimator's soft sensor is off by 0.0097 g/t rms against 0.0105 g/t and by more than 0.01 g/t for 38.2 h against 42.3 h. Between 120 and 240 h it is off by 0.0225 g/t against 0.0200 g/t, for 59.5 h against 54.6 h, and the true tailings peak at 0.3458 g/t against 0.3367 g/t, higher than under any shipped structure. Over the 240 h the loop earns $88,994.5 per hour on the estimator against $89,024.4 on the update, and no bound is active in 59 updates. The curves in Figure 7.5 are drawn from engine data exported at 0.2 h intervals, not from screen captures.

Across every existing scenario and the challenge on its suggested settings, each run under the four shipped outer loops and under the soft sensor loop and the predictive controller with the estimator, the head-tank pH minima are identical in every row. Beyond scenario 9, the estimator is less safe than every shipped structure on four cases. The challenge on its suggested settings under the soft

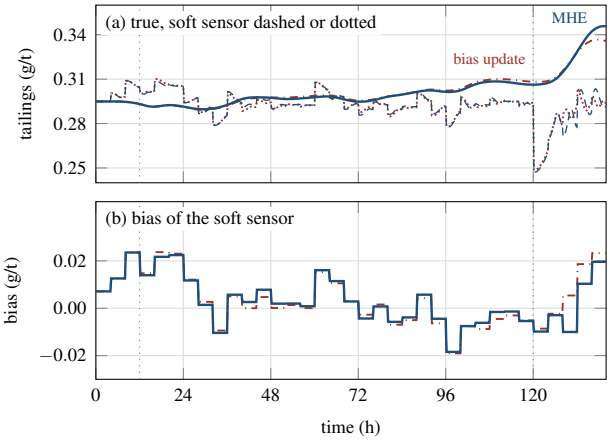


Figure 7.5 Scenario 12 from the start to 140 h with the moving horizon estimator (solid and dashed lines, drawn blue) and the bias update (dash-dotted and dotted lines, drawn red). (a) The true tailings (solid for the estimator, dash-dotted for the update) and the soft sensor with its bias (dashed and dotted). (b) The bias in force (solid and dash-dotted). The vertical dotted lines mark the richer feed at 12 h and the soft sensor's fourteen percent error from 120 h. At 140 h the true tailings are 0.346 g/t with the estimator and 0.336 g/t with the update, while the soft sensor reads 0.292 g/t and 0.293 g/t.

sensor loop spends 112.5 h above its target against 106.0 h, although its peak is lower (0.4294 against 0.4521 g/t) and so is its time above 0.40 g/t (17.7 against 21.6 h), and tank five's largest free cyanide under the predictive controller is higher by 0.2 mg/L to 1.8 mg/L on scenarios 6, 7 and 10 (232.9 against 232.7, 115.0 against 114.8 and 159.4 against 157.6 mg/L). Against the bias update under the same loop it also spends more hours above the target on scenario 7 (108.2 against 84.7 h under the soft sensor loop and 53.4 against 46.1 h under the predictive controller) and under the predictive controller on scenarios 6 and 11 (98.9 against 94.4 h, 41.4 against 38.3 h). It spends fewer hours above the discharge limit on scenarios 6, 8 and 11 (80.1 against 101.6 h, 130.7 against 141.0 h, 24.7 against 28.3 h).

7.7 Technology in Practice

Figure 7.6 draws the loops of Section 7.5 with their hardware.

What Measures It

The model reduces the free cyanide titrator to a reading of the head tank 0.40 h late, a cycle inside the range that online analyzers report.

> **TECHNOLOGY** **Sensor: An Online Free Cyanide Titrator**
>
> An online free cyanide analyzer filters a solids-free sample of solution from the leach pulp and titrates it with silver nitrate. Silver ions bind the free cyanide until an excess appears, and a silver electrode in the cell detects that end point as a change in potential. Such an analyzer reports every 5 to 30 minutes, against 2 to 12 hours for a sample taken and titrated by hand. Copper is its weakness: part of the cyanide bound to copper is titrated with the free cyanide and makes the end point hard to find and inaccurate [7]. On a copper-rich ore the reading can therefore stand above the true free cyanide, and the loop holds the wrong value.

The pH electrode is a first-order lag of 0.05 h and the oxygen probe reads the head tank continuously. The variable the circuit is paid on, the tailings grade, has no continuous instrument at all. The

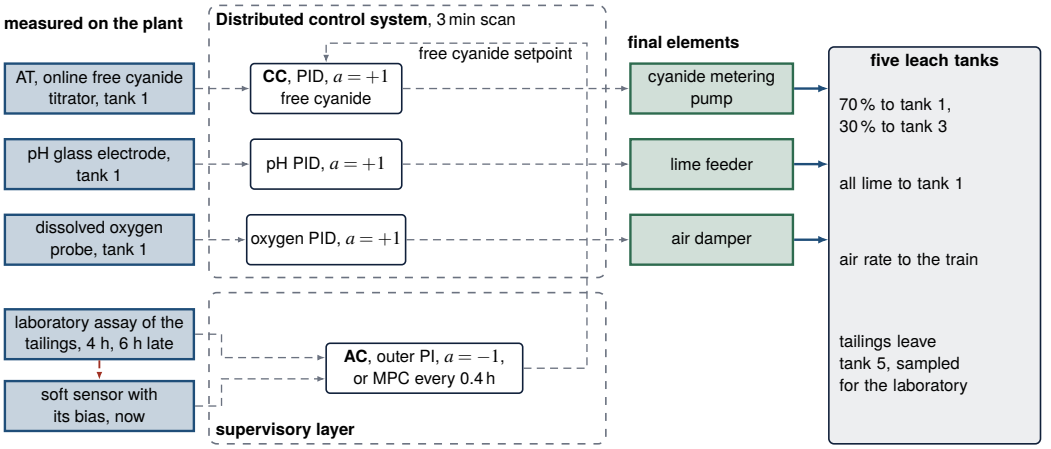


Figure 7.6 The loops of the simulated circuit drawn with their hardware. The online titrator AT, the pH electrode and the oxygen probe in tank one feed three PID blocks that move the cyanide metering pump, the lime feeder and the air damper. In the supervisory layer the outer controller AC, or the predictive controller in its place, reads the laboratory assay or the soft sensor whose bias the assay corrects (dashed), and writes the free cyanide setpoint of CC.

model takes it from a laboratory assay of a sample cut every four hours and reported six hours later, and from the soft sensor of Equation (7.3).

What Moves It

The model makes the cyanide pump, the lime feeder and the air damper first-order lags of 0.08 h, 0.12 h and 0.06 h, with the cyanide addition held between 0 and 500 kg/h and split seventy to thirty between tanks one and three.

> **TECHNOLOGY** **Final Element: A Cyanide Metering Pump**
>
> A metering pump is a pump mechanism, motor and drive electronics in one package, made to inject small, controlled flows of a chemical. A 4 to 20 mA signal sets its delivery from zero to full flow, and a knob on the pump sets what full flow is. Many metering pumps drive a plunger with a solenoid at a variable rate, so they deliver discrete pulses rather than a continuous stream [26]. A positive-displacement pump discharges a fixed volume with every turn and locks if its discharge is blocked [26]. A pump that has lost its delivery shows on a trend as a controller output that keeps rising while the free cyanide falls.

The sulfidic parcel and the graded week at the tighter target both drive the modeled pump to its 500 kg/h limit, the saturation that the full-flow setting of a real pump imposes in the same way.

Where the Controller Runs

Control loops in a plant are often nested or cascaded, with the setpoint of one loop based on the variable of another, and with cycle times from milliseconds to minutes. A distributed control system uses a supervisory loop to mediate a group of localized controllers [27], and programmable logic controllers run PID loops as well [26]. The app runs its PID blocks on a 3 min scan, short against the 2.68 h cyanide response and the 11.52 h tailings response. On such a platform the three secondary loops are PID function blocks, the soft sensor and its bias update are a calculation block that runs once per assay, and the outer loop or the predictive controller is a supervisory application that writes the free cyanide setpoint every 0.4 h.

> **IN PRACTICE Where the Cyanide Is Measured**
>
> In practice the free cyanide is measured in the first leach tank, by hand every 2 to 12 hours or with an online analyzer every 5 to 30 minutes. The consumption can change within two hours when the ore mineralogy changes, so the faster measurement holds the setpoint more tightly and sends less cyanide to the tailings [7]. The International Cyanide Management Code asks a mine to minimize its cyanide use and so limit the cyanide in its mill tailings (Standard of Practice 4.2), and to monitor the effects of cyanide on wildlife and on surface and ground water (Standard of Practice 4.9) [28].

What Is Modernizing It

The first development is the online analyzer. Control of the cyanide addition from an online measurement has been reported to cut cyanide consumption by at least five percent against manual control, and mines that sign the International Cyanide Management Code are required to measure the weak-acid-dissociable cyanide online as well [7]. The second is the data-driven soft sensor. Because an accurate mechanistic model of a gold leach is hard to build, data-driven modeling is used extensively in industrial plants, with methods from partial least squares to neural networks [29], where the soft sensor of this chapter evaluates the leach kinetics instead. The third is optimization of the setpoint itself, by real-time optimization of the leach under disturbance and model uncertainty [29] or by the economic predictive controller and moving horizon estimator of Sections 7.5.2 and 7.5.3.

7.8 Limitations

Gold is one leachable class and one refractory floor. Real ore has a distribution of liberation and grain size, so the leach curve is a spectrum of rate constants rather than a single one, and the refractory fraction itself depends on the grind [1]. The application represents that a slower ore raises the tailings grade and does not predict by how much for a particular deposit.

Carbon is one loading per tank with no regeneration circuit. Elution, acid washing, kiln regeneration and carbon attrition are all absent, and carbon management is a substantial part of the control task on an operating circuit [2]. Preg-robbing by carbonaceous ore, which is the failure mode that makes some deposits uneconomic, has no representation.

Cyanide destruction and the tailings storage facility are outside the boundary. The application reports a detoxification load and does not model it, and the environmental constraint on an operating circuit is at least as binding as the reagent cost.

The soft sensor is a steady-state evaluation of the same leach kinetics the plant model uses, with its constants fixed at the plant's declared values, differing in the ore parameters, in the assumption of a uniform profile along the train and in omitting the gold that leaves in solution. A soft sensor on an operating circuit is a fitted model of a different form, and its failure modes include extrapolation outside the data on which it was fitted, which cannot be staged here. What the application does represent is that an inference is no better than its most uncertain input, and that the most uncertain input is usually the ore [6].

The predictive controller inherits the soft sensor's model and its limits. Its slope of tailings against cyanide is 1/1.25 of the plant's at design and about half of it on a harder ore, no bias can correct a slope, and a fixed setpoint chosen in hindsight therefore beats it by \$8.3 to \$132.4 per hour. It sees an ore change only through the laboratory, which costs it scenario 6. Its ceiling carries a slack at every judged step, so a ceiling that cannot be met costs the plan at every step where it is broken, and on scenario 8 the plan runs the pump at its limit for 67.4 h, ends with the true tailings at 0.298 g/t and earns \$75.3 per hour less than the earlier form with one slack over the day, which postponed that cyanide. Honoring a ceiling the pump cannot reach is thus a design choice with a measured price. The plan holds the measured feed grade, throughput, pH and oxygen constant over its horizon, and it carries no constraint on tank five's free cyanide, which exceeds the 200 mg/L discharge limit for 28.3 h on scenario 11, 101.6 h on scenario 6 and 186.2 h on scenario 8.

The moving horizon estimator inherits the soft sensor's model through its reference and adds the tanks' lag to it. It treats every error as one that passed through the train, so an error at the sensor itself is corrected with that lag, which makes it less safe than every shipped structure on scenario 9. Its only unknown is an additive bias, so the soft sensor's understated slope remains. Its weights are declared rather than fitted, a laboratory repeatability of 0.015 g/t where scenarios 10 and 12 carry 0.02 g/t of assay noise, and because neither bound binds on a shipped case, the window and the bounds add nothing there that a bias update on the same reference lacks.

> **OPEN CHALLENGE** **Reaching the Plant's Own Optimum**
>
> A setpoint that is optimal on a model of a gold cyanidation leach is often not optimal on the plant. An accurate mechanistic model is hard to build for a long process with complex chemistry, and the plant drifts as its equipment ages and its ore changes. Published remedies fit local models from plant data near the current point and step the setpoint toward the plant's own optimum while the residue grade meets its limit, and they have been shown in simulation [29]. Progress would be such a method proved on an operating circuit.

7.9 Discussion

Three results generalize beyond this circuit.

The first is that a soft sensor is a measurement and should be evaluated as one. Sections 7.6.1 and 7.6.2 evaluate it that way, finding that it removes six hours of dead time and introduces a model error of up to 0.04 g/t, and that the second figure is tolerable only because an assay exists to correct it. A soft sensor without a reconciling measurement is a model running open loop rather than a faster measurement.

> **OPEN CHALLENGE** **Seeing the Ore Before It Leaches**
>
> Sensors on the ore feed can reveal trends in its grade and in how its gold occurs, and such readings could inform the cyanide addition. An operating plant still judges its recovery from samples, drawn in one study twice a day from every tank of a carbon-in-leach circuit for 209 days, and the variability of the feed and the uncertainty of sampling blur what they show [19]. Progress would be a measurement of how fast the incoming ore will leach, made before it reaches the first tank.

The second is that the reconciliation has dynamics of its own. Table 7.2 shows a bias update that is too fast passing the assay dead time and the laboratory's noise back into the loop, with the cyanide setpoint wandering over 101 mg/L against 53 mg/L at a gain of 0.15. The correct gain defines a filter length, and the sampling interval and the assay noise rather than the controller set that length. The moving horizon estimator of Table 7.5 leads to the same place, because with no bound active it is a filter of that gain and its result is decided by the gain and the reference it compares against, not by its window.

The third concerns when a supervisory layer is justified. An outer loop is worth closing when the plant drifts more slowly than the loop can act and through variables its estimate sees. On Table 7.3 the recommended cascade saves two percent of reagent and spends 21.6 h above the tailings limit, because its target spends the margin that a fixed setpoint keeps and the first disturbance is one its estimator cannot see. The secondary loops already reject what is available to be rejected. A supervisory layer that the disturbances do not justify is a common and costly choice [5, 12]. Volume 9, Chapter 10 asks the same question of an aeration plant, where a slow effluent ammonia loop sets a fast dissolved oxygen setpoint.

7.10 Conclusion

A carbon-in-leach circuit places 11.52 h of plant response and six hours of laboratory turnaround between a reagent decision and the evidence that judges it. The application integrates a five-tank leach and adsorption train, holds free cyanide, alkalinity and dissolved oxygen on continuous instruments,

and reports the tailings grade as a computed value, a delayed assay and a soft sensor together. The soft sensor removes the assay dead time and introduces an error of up to 0.04 g/t that only the laboratory detects, which makes the bias update the governing design decision: at a gain of one it passes the assay dead time and noise back into the loop and raises the actuator travel by sixty-three percent over a gain of 0.15, and at zero it starves the leach for 172 h of a 200 h week.

On the graded week the outer loop is hardly worth closing, because a fixed cyanide setpoint holds the tailings under 0.40 g/t throughout at 0.535 kg/t of reagent, while the recommended cascade saves two percent of the reagent and exceeds the limit for 21.6 h, and a target tight enough to stay under the limit costs a third more reagent and runs the cyanide pump into its limit. The excursion comes from the proportional response to the target itself rather than from integral windup. The result does not show that inferential control is inappropriate. It shows that whether to close a supervisory loop is a question that precedes the question of what to close it on.

An economic predictive controller on the soft sensor’s model earns $25.7 per hour more than the soft sensor loop on scenario 11, but a setpoint optimizer feeding that loop recovers all but $1.9 of it, a fixed setpoint chosen in hindsight earns $28.2 more than the plan, and the plan exceeds the tank-five discharge limit for 28.3 h where the loop does not. A moving horizon estimator of the bias, whose bounds never bind on a shipped case, is a bias update at a gain of 0.4 against a lagged model of the tanks, so its $28.7 per hour on the over-fast update belongs mostly to its gain, and it loses the structural error of scenario 9. Every result here except the setpoint optimizer’s is a shipped scenario at its own seed, so a reader can reproduce it in the app and change one setting at a time.

KEY POINTS

- The gold left in the tailings decides what the circuit earns, and a laboratory reports it six hours late.
- A soft sensor removes the assay dead time and adds an error of up to 0.04 g/t that only the laboratory detects.
- The bias update gain is a filter design, set by the sampling interval and the assay noise.
- On the graded week a fixed cyanide setpoint holds the tailings under 0.40 g/t, while the recommended cascade spends 21.6 h above the limit.
- Whether to close a supervisory loop precedes the question of what to close it on.

Where else this applies. A late, sampled quality measurement corrected into an inferential estimate also governs Volume 5, Chapter 10 and Volume 5, Chapter 4.

References

[1] J. O. Marsden and C. I. House, *The Chemistry of Gold Extraction*, 2nd ed. Littleton, CO: Society for Mining, Metallurgy, and Exploration, 2006.

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

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

[4] ——, “Process of separating gold and silver from ore,” U.S. Patent 418,137, application filed April 4, 1889, 1889, patented December 24, 1889.

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

[6] E. Mazzour, D. Hodouin, and S. Makni, “Optimal sensor implementation in metallurgical plants: an application to a generic mineral separation plant,” *International Journal of Mineral Processing*, vol. 69, no. 1–4, pp. 185–203, 2003.

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

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

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

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

[11] ——, "The future of PID control," *Control Engineering Practice*, vol. 9, no. 11, pp. 1163–1175, 2001.

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

[13] D. E. Rivera, M. Morari, and S. Skogestad, "Internal model control: PID controller design," *Industrial & Engineering Chemistry Process Design and Development*, vol. 25, no. 1, pp. 252–265, 1986.

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

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

[16] G. Senanayake, "Kinetics and reaction mechanism of gold cyanidation: Surface reaction model via Au(I)–OH–CN complexes," *Hydrometallurgy*, vol. 80, no. 1–2, pp. 1–12, 2005.

[17] M. Adams, R. Lawrence, and M. Bratty, "Biogenic sulphide for cyanide recycle and copper recovery in gold-copper ore processing," *Minerals Engineering*, vol. 21, no. 6, pp. 509–517, 2008.

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

[19] H. J. Glass, R. Abdulwahed, and S. Iyakwari, "Monitoring gold recovery by sampling a carbon-in-leach plant," in *Proceedings of the 57th International October Conference on Mining and Metallurgy*, Bor, Serbia, 2026, pp. 177–181.

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

[21] M. Ellis, H. Durand, and P. D. Christofides, "A tutorial review of economic model predictive control methods," *Journal of Process Control*, vol. 24, no. 8, pp. 1156–1178, 2014.

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

[23] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: stability and moving horizon approximations," *IEEE Transactions on Automatic Control*, vol. 48, no. 2, pp. 246–258, 2003.

[24] Q. Liu, Z. Wang, Y. Liang, Z. Wang, F. Peng, and D. He, "Soft sensor modeling and adaptive model correction strategy for gold cyanide leaching process," in *2019 Chinese Control And Decision Conference (CCDC)*. IEEE, 2019, pp. 6065–6070.

[25] R. Dubois, S. Bertrand, and A. Eudes, "Performance evaluation of a moving horizon estimator for multi-rate sensor fusion with time-delayed measurements," in *2018 22nd International Conference on System Theory, Control and Computing (ICSTCC)*. IEEE, 2018, pp. 664–669.

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

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

[28] International Cyanide Management Institute, "The International Cyanide Management Code: principles and standards of practice for mining operations," cyanidecode.org/about-the-cyanide-code/the-cyanide-code, 2026, accessed November 19, 2026.

[29] Q. Liu, D. He, Z. Wang, and J. Shi, "A data-driven iterative optimization compensation method based on PJIT-PLS for gold cyanidation leaching process," *IEEE Access*, vol. 7, pp. 138 683–138 693, 2019.

7.11 Exercises

7.1 Concept check · 2 min

In scenario 6 the ore leachability falls to 0.82 at 12 h while the feed grade, the cyanide, the pH and the oxygen stay where they were. With the outer loop on the soft sensor, the cyanide setpoint first moves at 20.1 h, the same time as with the outer loop on the assay. Why is the soft sensor no faster here?

(a) The titrator's 0.40 h dead time delays the cyanide reading the soft sensor uses
(b) The outer loop's rate limit holds the setpoint until its error has grown past a threshold
(c) The soft sensor passes its estimate through the five tank lags before the loop reads it
(d) Leachability is not one of its inputs, so only an assay that contains the change can move it

7.2 Concept check · 1 min

Each assay updates the bias by Equation (7.3). Which value of the soft sensor is the assay compared with?

(a) Its value stored at the instant the sample was cut, six hours before the report
(b) Its value at the instant the assay is reported, the latest estimate available
(c) Its average over the four hours between the last two samples that were cut
(d) Its value at the report, passed through the five residue lags of the tanks

7.3 Concept check · 2 min

In scenario 7 a copper parcel multiplies the cyanicide term by 2.2. The head-tank cyanide loop stays within 10.2 mg/L of its setpoint, yet the tailings rise to 0.3084 g/t. Which explanation is right?

(a) The titrator's delay lets the head-tank cyanide wander far from its setpoint between readings
(b) The parcel consumes cyanide along the train, so tank five falls to 53.2 mg/L while tank one is held
(c) The copper acidifies the later tanks, and the head-tank pH electrode cannot see their pH fall
(d) The soft sensor stops reading the feed grade while the parcel passes through the train

7.4 Theory and calculation · 6 min

The tailings are sampled every $T_s = 4$ h and each assay is reported 6 h after its sample was cut. The bias update of Equation (7.3) runs at gain β.

(a) How old is the assay in hand at any instant, at least, at most and on average? *Checkpoint:* 6 h, 10 h, 8 h.

(b) Find the time constant of the bias update at $\beta = 0.15$, 0.35 and 0.40, to 0.1 h. *Checkpoint:* 24.6 h, 9.3 h, 7.8 h.

(c) The assay carries white noise of spread 0.02 g/t. Using Section 1.4.5, find the spread it leaves in the bias at $\beta = 0.15$ and 1. Then take a single assay 0.02 g/t high with nothing else changing: how far does the outer loop's proportional term, at $K_c = 2500\,\mathrm{mg/L}$ per g/t, move the cyanide setpoint at each gain, and how long does the 8 mg/L per hour ramp take to deliver the larger move? *Checkpoint:* 0.0057 g/t and 0.020 g/t; 7.5 mg/L and 50 mg/L; about 6.3 h.

7.5 Theory and calculation · 6 min

At design the circuit treats 400 t/h of ore, gold is worth \$78 per gram and sodium cyanide costs \$2.40 per kilogram. Holding one more mg/L of free cyanide in the head tank takes 1.651 kg/h of sodium cyanide, and a fraction 0.7471 of the head-tank cyanide reaches tank five and leaves in 488.9 $\mathrm{m^3/h}$ of solution, where its destruction costs \$3.50 per kilogram.

(a) What is 0.001 g/t of tailings worth per hour, to the dime? *Checkpoint:* \$31.2 per hour.

(b) What does one more mg/L of setpoint cost per hour, to the cent, and how much must it lower the tailings to pay for itself, to two significant figures? *Checkpoint:* \$5.24 per hour; about 0.17×10^{-3} g/t per mg/L.

(c) Use the cyanide rows of Table 7.1 to say whether the steps from 120 to 200 mg/L and from 200 to 260 mg/L pay. *Checkpoint:* The first pays, the second does not.

7.6 Theory and calculation · 5 min

The soft sensor assumes every tank sits at the head-tank conditions. At design the head tank holds 120 mg/L of free cyanide, 6.0 mg/L of oxygen and pH 10.40, where $f(\mathrm{pH}) = 1$. Take $k = 0.947/\mathrm{h}$, $K_{\mathrm{CN}} = 60\,\mathrm{mg/L}$, $K_{\mathrm{O}} = 3.0\,\mathrm{mg/L}$, a residence of 3.14 h per tank, five tanks, a feed of 2.8 g/t and 0.25 g/t of refractory gold.

(a) Find the soft model's leach rate k_s at these conditions, to three decimals. *Checkpoint:* 0.421/h.

(b) Find the tailings grade the soft sensor computes, to three decimals, with Equation (1.6) applied to the leachable gold, and compare it with the true 0.295 g/t. *Checkpoint:* 0.288 g/t, about 0.007 g/t optimistic.

(c) Find the residue time constant τ_r of one tank and the sum of the five, to 0.01 h and 0.1 h. *Checkpoint:* 1.35 h and 6.8 h.

7.7 App challenge · 8 min

Open scenario 4, "Close the loop on the laboratory assay", and then scenario 5, "Close it on a soft sensor instead", each as loaded and at the default seed. Both step the feed grade from 2.8 to 3.2 g/t at 10 h. Run each for 200 h at 500 h per second and export the CSV, whose time column is Time (h).

(a) Find the first row in which `cn_setpoint_mgL` exceeds 122 mg/L in each run, within 0.1 h, and how long after the feed step each one comes. *Checkpoint:* 20.3 h and 10.3 h; about 10.3 and 0.3 h after the step.

(b) Read the MV travel tile of the free cyanide loop on the scorecard at 200 h in each run, within 2 kg/h. *Checkpoint:* 428 kg/h and 296 kg/h.

(c) Find the largest cyanide setpoint in scenario 4 and when it occurs, within 1 mg/L and 0.5 h, and the setpoint at 200 h. Explain the overshoot. *Checkpoint:* About 150 mg/L at 40 h; about 140 mg/L.

7.8 App challenge · 9 min

Open scenario 10, "A bias update fast enough to bring the dead time back", at the default seed. It runs the soft sensor loop with Soft sensor bias update gain 1.0, a Tailings assay noise of 0.02 g/t and a feed grade step at 12 h. Run each part for 240 h at 500 h per second. Read the free cyanide loop's MV travel tile, within 2 kg/h, and the NaCN tile, and from the CSV the lowest and highest `cn_setpoint_mgL` at or after 24 h, within 1 mg/L.

(a) Run it as loaded. *Checkpoint:* 5472 kg/h; 0.512 kg/t; 78 to 179 mg/L.

(b) Press Reset, set Soft sensor bias update gain to 0.15 before Run, and run. *Checkpoint:* 3356 kg/h; 0.517 kg/t; 101 to 154 mg/L.

(c) Reload the scenario with its reload button, set Tailings assay noise (Advanced) to 0 before Run, keep the gain at 1.0, and run. Read the travel. What does the dead time alone cost, and what does the noise cost? *Checkpoint:* 425 kg/h.

7.9 App challenge · 9 min

Open the graded challenge, "a week of difficult ore", at its default seed. Run each part for the full 200 h at 500 h per second, and read the challenge result in the progress panel and the NaCN tile. Use the CSV to count the hours with `tails_gt` above 0.40 g/t, at 0.1 h a row, within 0.2 h.

(a) Run it as loaded. Which clause fails? *Checkpoint:* Not met: the tailings, 0.433 g/t; cyanide 8.5 mg/L, pH 10.32.

(b) Reload, press Reveal suggested tuning and run. Count the hours above 0.40 g/t and read the peak of `tails_gt`. *Checkpoint:* Met: 10.1 mg/L, 0.379 g/t, pH 10.40; 0.525 kg/t; 21.6 h; 0.452 g/t.

(c) Press Reset, set Outer loop to Local before Run, leaving the free cyanide setpoint at 120 mg/L, and run. Compare with (b). *Checkpoint:* Met: 6.1 mg/L, 0.396 g/t; 0.535 kg/t; no hour above 0.40 g/t.

7.10 Application area · 4 min

A plant meters all of its cyanide into the first tank and holds the head-tank free cyanide on a tight loop.

(a) Using scenario 7 of Section 7.6.4, say what happens in tank five when a copper-bearing ore arrives, and what changing the split of the addition between tanks one and three achieves. *Checkpoint:* Tank five falls to 53.2 mg/L; with half in tank three it stays at or above 76.3 mg/L.

(b) The predictive controller of Section 7.6.6 pushes tank five's free cyanide above its 200 mg/L discharge limit for 28.3 h on scenario 11. What would you add before running such a plan on a plant? *Checkpoint:* A constraint or override on tank five's cyanide.

7.11 Application area · 4 min

The head-tank pH is held by lime at a setpoint chosen by the operator.

(a) Use Equation (7.2) to find the share of the free cyanide present as hydrogen cyanide at pH 9.5. *Checkpoint:* 39 %.

(b) The margin optimum lies at pH 10.55, \$82 per hour above the design pH of 10.40, and pH 9.60 costs \$3237 per hour. On which side of the optimum would you err, and what does the sulfidic parcel of scenario 8 add to the argument? *Checkpoint:* The high side; tank five can fall below the floor while the head tank is held.

7.12 Review the agent · 5 min

An AI assistant summarizes Table 7.5: "On scenario 10 the moving horizon estimator earns \$28.7 per hour more than the shipped bias update, because its 48-hour window and its bounds reject noisy assays that the update believes. Replace the bias update with the estimator on every circuit." Find the flaw and say how you would check it in the app.