INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 15 Control in Mining and Mineral Processing

CHAPTER 3

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

Pages 43–71

Edited by John D. Hedengren

How to cite this chapter. "Semi-Autogenous Mill Simulator for Impact-Zone Control Against a Moving Constraint," 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. 3, pp. 43–71.

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

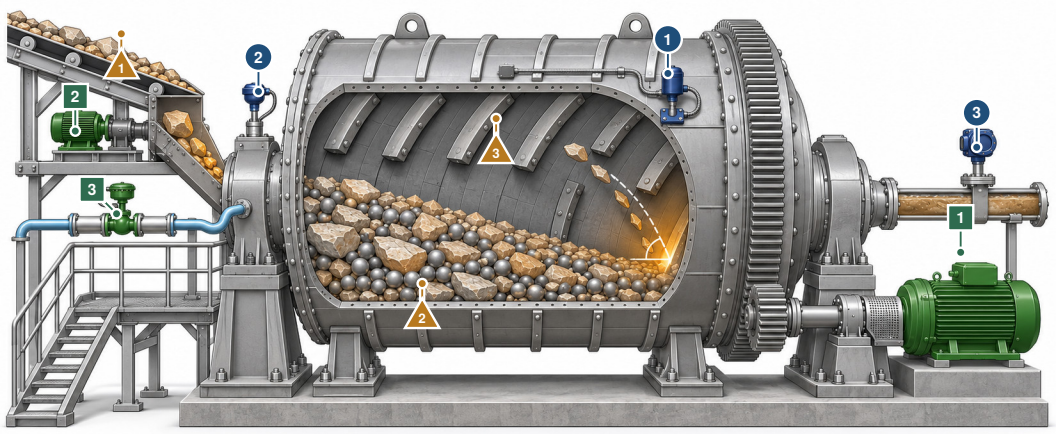
The app for this chapter
https://apmonitor.com/apps/sag-mill/

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

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

CHAPTER 3

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



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

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

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

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

IN PLAIN TERMS Aiming Falling Steel at the Rock

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

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

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

THREE WAYS TO READ THIS CHAPTER

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

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

Learning objectives

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

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

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

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

3.1 Introduction

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

DID YOU KNOW? One Percent of the Energy

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

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

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

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

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

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

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

3.2 Background

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

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

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

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

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

3.3 Process Model

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

3.3.1 The Design Point

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

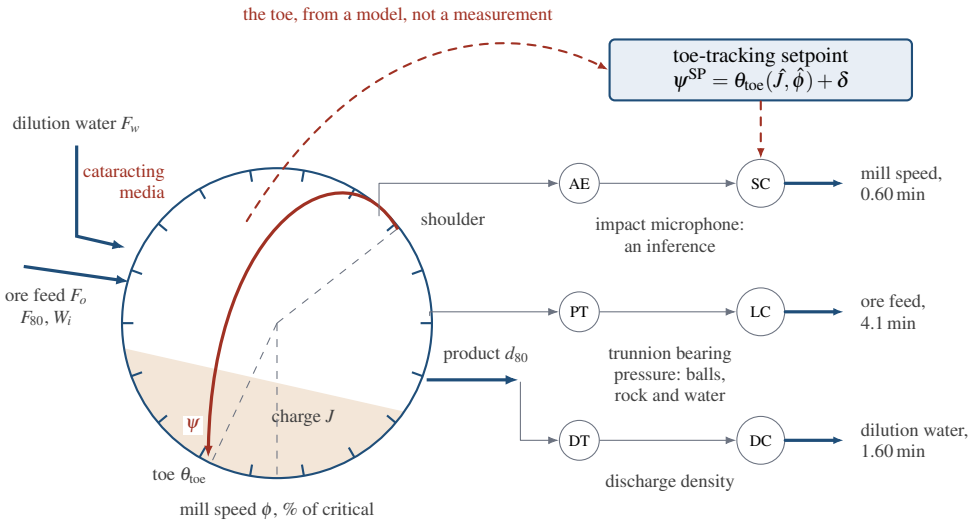


Figure 3.1 The simulated mill and its three loops. The impact-angle controller (SC) moves mill speed against an angle inferred from a shell-mounted microphone. The load controller (LC) moves ore feed against trunnion bearing pressure, which reads the mass of balls, rock and water together. The density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

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

3.3.2 Charge Geometry and the Position of the Toe

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

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

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

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

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

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

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

3.3.3 Ballistics of the Cataracting Media

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

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

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

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

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

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

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

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

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

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

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

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

3.3.4 Mill Power

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

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

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

3.3.5 Breakage, Product Size and Liner Wear

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

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

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

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

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

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

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

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

3.3.6 Inventory Balances

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

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

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

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

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

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

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

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

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

3.3.7 Measurements

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

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

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

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

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

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

3.3.8 Parameters

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

3.3.9 Initialization and Numerics

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

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

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

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

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

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

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

3.3.10 Assumptions

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

3.4 Nonlinear Characteristics and Dynamic Response

3.4.1 Process Gain Against Mill Speed

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

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

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

3.4.2 Power Against Filling

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

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

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

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

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

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

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

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

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

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

3.4.3 Open-Loop Response Times

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

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

3.5 Controller Design

3.5.1 Velocity Form PID Algorithm

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

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

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

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

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

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

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

3.5.2 The Three Loops and Their Tunings

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

3.5.3 The Supervisory Setpoint

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

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

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

3.5.4 Model Predictive Control

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

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

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

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

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

The quadratic program minimizes Equation (3.26).

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

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

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

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

3.6 Results

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

3.6.1 The Optimum and the Overload by Hand

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

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

3.6.2 A Setpoint Step Under Two Tunings

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

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

3.6.3 A Measurement That Is an Inference

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

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

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

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

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

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

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

3.6.4 A Worn Liner Moves the Plant and Not the Measurement

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

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

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

3.6.5 A Constraint That Moves While the Setpoint Does Not

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

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

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

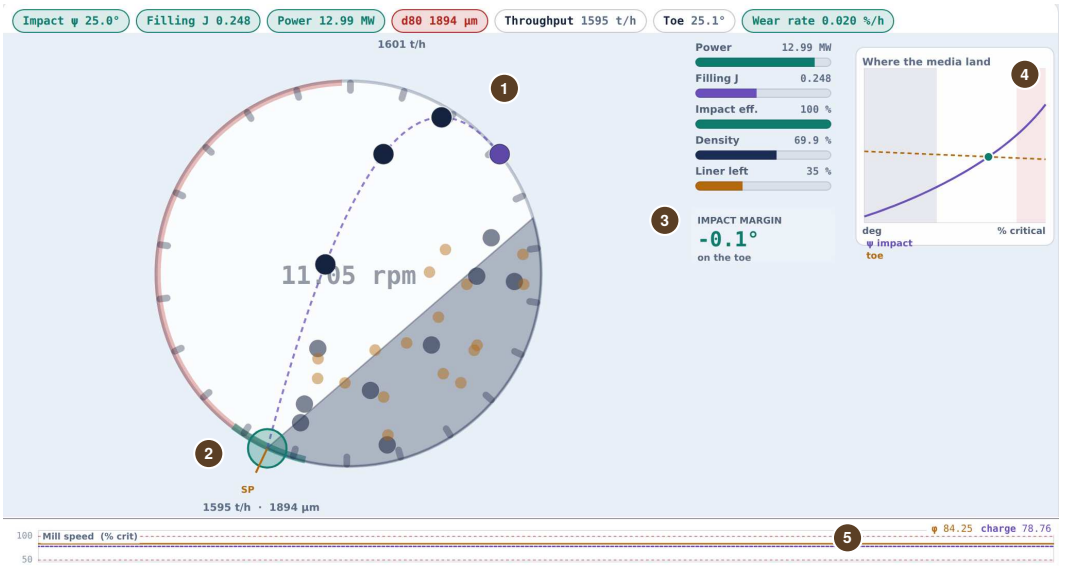


Figure 3.2 Scenario 8 at $t = 800$ min (13.33 h on the application clock) in automatic, with the liner 65 % worn at the start, cropped to the plant view and the mill speed strip. The mill is drawn in end view with the charge, the cataracting trajectory (1) from (3.7) and the impact point (2) marked against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

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

🕮 **Pause.** Try Exercise 3.3 before reading on.

3.6.6 Tracking the Toe on Three Disturbances

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

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

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

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

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

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

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

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

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

3.6.7 The Graded Shift and a Tuning That Makes It Worse

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

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

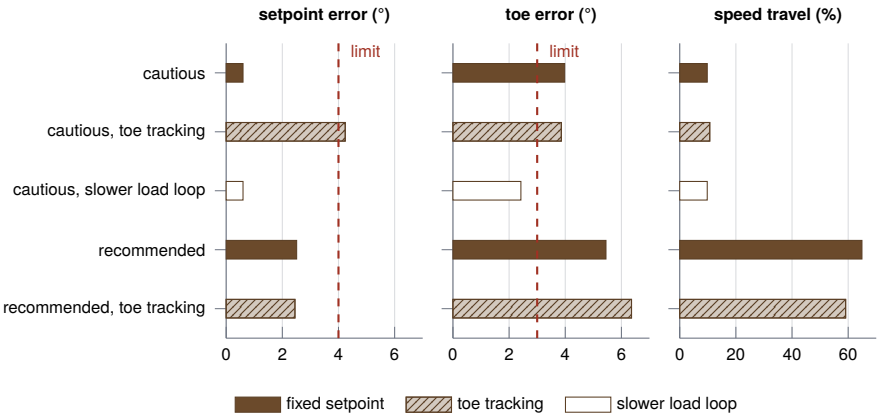


Figure 3.3 The graded shift of Table 3.9 for the five configurations on the loops' scale: the worst error from the setpoint, the worst error from the moving toe and the speed travel over 70 min. The dashed lines are the challenge's clauses of 4° and 3°. The sharp tuning, far off every scale, stays in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

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

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

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

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

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

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

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

3.6.8 Model Predictive Control Against the Three Loops

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

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

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

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

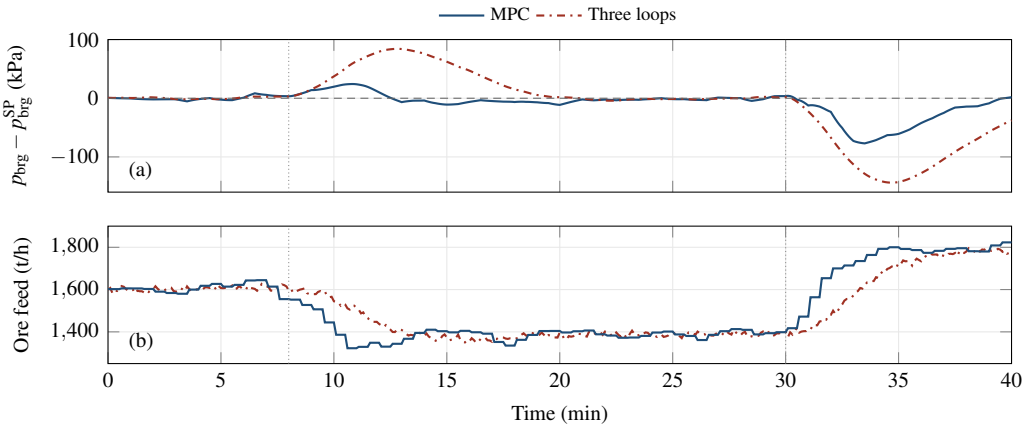


Figure 3.4 Scenario 11 under the MPC and under the three loops, both at the 5485 kPa load setpoint. (a) True bearing pressure over the setpoint. (b) Ore feed command. The MPC is the solid line (blue) and the three loops the dash-dotted line (red). The ore hardens at 8 min and softens at 30 min (vertical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

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

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

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

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

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

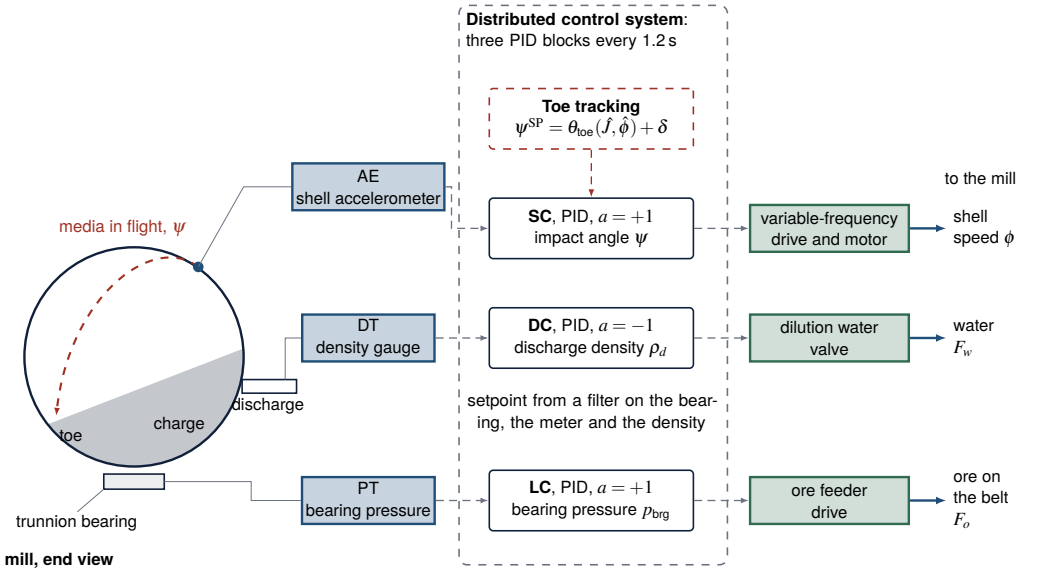


Figure 3.5 The three loops of the simulated mill drawn with their hardware. The shell accelerometer AE, the trunnion bearing pressure PT and the discharge density gauge DT feed three PID blocks, SC and LC reverse acting with $a = +1$ and DC direct acting with $a = -1$, which move the variable-frequency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

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

3.7 Technology in Practice

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

What Measures It

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

TECHNOLOGY Sensor: An Acoustic Pickup on the Mill Shell

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

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

IN PRACTICE Crash Stop and Grind-Out

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

What Moves It

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

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

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

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

Where the Controller Runs

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

What Is Modernizing It

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

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

3.8 Limitations

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

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

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

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

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

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

3.9 Discussion

Three results generalize past this mill.

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

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

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

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

3.10 Conclusion

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

KEY POINTS

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

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

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

3.11 Exercises

3.1 Concept check · 1 min

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

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

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

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

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

3.2 Concept check · 2 min

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

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

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

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

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

3.3 Concept check · 2 min

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

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

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

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

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

3.4 Theory and calculation · 6 min

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

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

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

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

3.5 Theory and calculation · 5 min

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

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

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

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

3.6 Theory and calculation · 6 min

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

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

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

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

3.7 App challenge · 8 min

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

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

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

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

3.8 App challenge · 8 min

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

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

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

3.9 App challenge · 9 min

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

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

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

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

3.10 Application area · 4 min

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

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

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

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

3.11 Application area · 4 min

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

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

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

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

3.12 Review the agent · 5 min

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