INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 15 Control in Mining and Mineral Processing

CHAPTER 2

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

Pages 22–42

Edited by John D. Hedengren

How to cite this chapter. "Three Feeders and One Belt Scale: Ore Blending Through Speed-Dependent Delays," 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. 2, pp. 22–42.

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

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

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

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

CHAPTER 2

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



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

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

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

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

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

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

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

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

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 2.2, the fact boxes of the Introduction and of Section 2.6, the try-it boxes of Section 2.4, and the key points
Study	45 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 2.6, Section 2.7, Section 2.8 with its open challenge

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

Learning objectives

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

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

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

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

2.1 Introduction

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

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

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

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

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

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

2.2 Belt Scales, Raw-Mix Proportioning and Delayed Feed

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

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

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

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

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

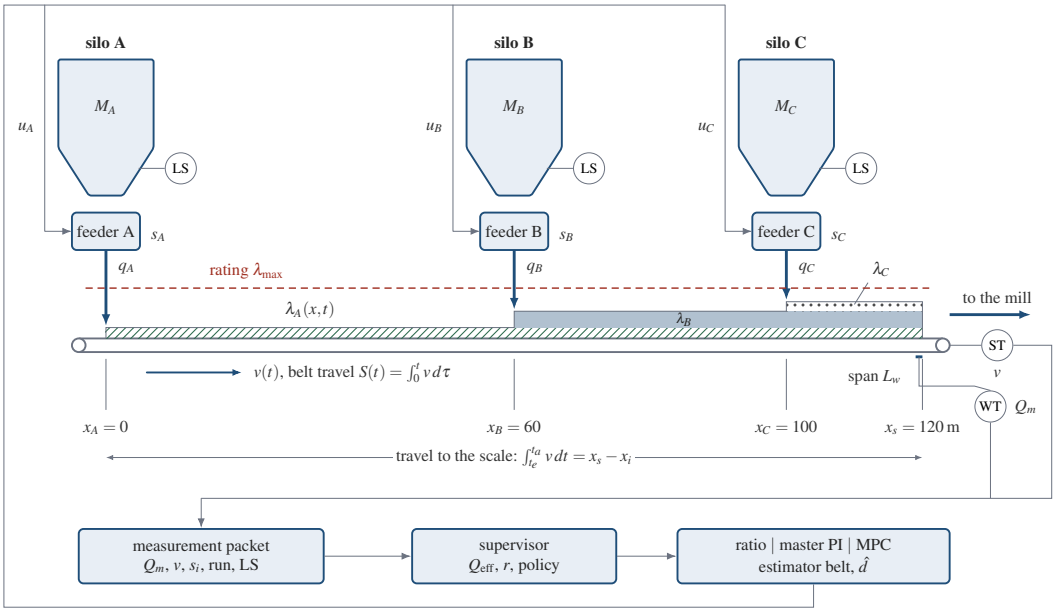


Figure 2.1 Three-feeder ore conveyor as modeled, positions to scale along the belt. Silo inventories M_i feed drives at actual speed s_i, which deliver q_i onto the belt at x_i. Source-resolved loadings $\lambda_i(x,t)$ ride the belt at speed $v(t)$ to the weighbridge of span L_w at x_s, where the weight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

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

2.3 The Conveyor as Modeled

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

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

2.3.1 Drives, Silos and Delivery

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

Table 2.1 Variables, Their Roles and Design-Point Values

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

Equation (2.1).

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

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

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

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

2.3.2 The Belt in Its Own Coordinate

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

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

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

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

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

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

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

Pause. Try Exercise 2.1 before reading on.

2.3.3 The Belt Scale

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

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

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

2.3.4 Parameters, Initialization and Numerics

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

2.3.5 Assumptions

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

2.4 Dynamic Character

2.4.1 One Measurement, Three Delays

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

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

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

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

2.4.2 A Speed Change Is a Flow Disturbance

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

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

2.4.3 What One Total Cannot Say

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

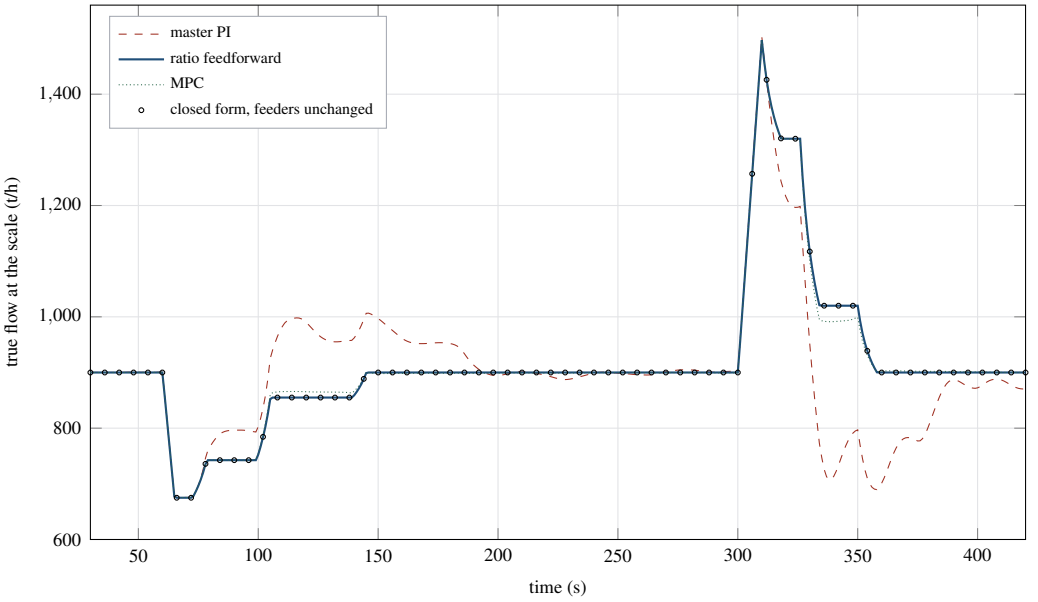


Figure 2.2 True flow at the scale in case 6, belt slowed from 2.0 to 1.5 m/s at 60 s and speeded to 2.5 m/s at 300 s. Master PI dashed, ratio feedforward thick solid, MPC dotted. The circles are Equation (2.4) with the feeders held at their design flows. Ratio feedforward and MPC lie on or near the closed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

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

2.5 Implemented Control System

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

2.5.1 Supervisor

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

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

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

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

2.5.2 Ratio Feedforward and the Master PI

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

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

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

2.5.3 Estimator, Blend Confidence and Bump Test

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

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

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

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

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

2.5.4 Model Predictive Control

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

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

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

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

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

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

Table 2.3 The Nine Cases

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

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

2.6 Nine Cases, Measured

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

2.6.1 Code Verification

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

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

2.6.2 Allocation and Tuning

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

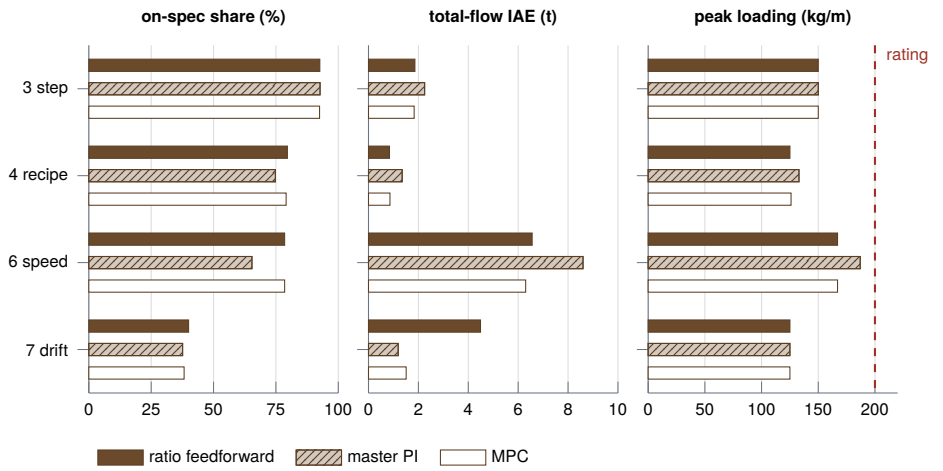


Figure 2.3 Three columns of Table 2.4, means over three paired seeds: ratio feedforward solid, master PI hatched, MPC open. The dashed line is the belt's 200 kg/m loading rating. Notice that the master PI loses on-spec share in cases 4 and 6, where it chases dips the transport makes, yet has the lowest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

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

2.6.3 Controller Comparison

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

Table 2.4 Structures Compared, Means Over Three Paired Seeds

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

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

+ DID YOU KNOW? Feedback That Makes the Blend Worse

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

2.6.4 Calibration, Empty Silos and Capacity

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

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

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

2.6.5 The Challenge

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

Table 2.5 Challenge Verdicts Over 1200 s

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

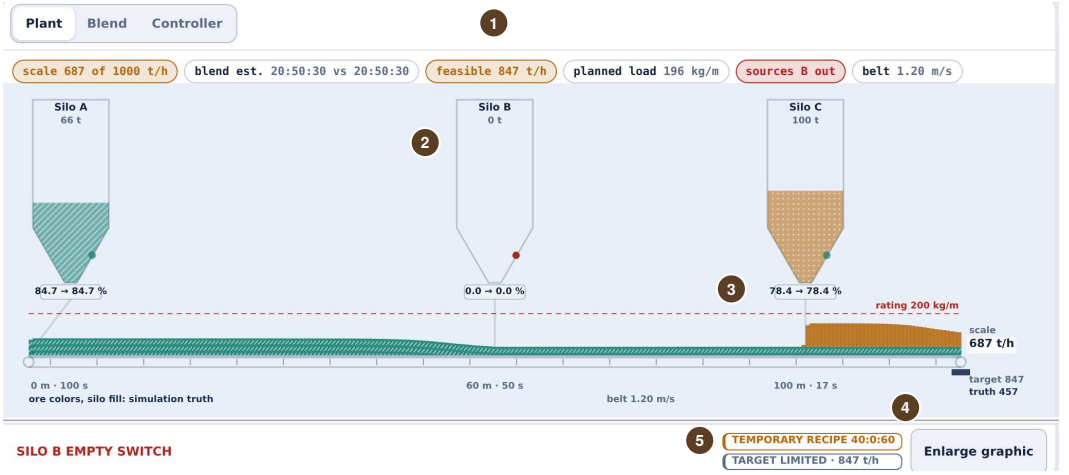


Figure 2.4 The challenge at $t = 640$ s on the Plant view with the suggested settings (MPC, supervisor On, temporary recipe 40:0:60 labeled), cropped to the plant view and its status line, with the belt at 1.20 m/s. (1) The effective target is 847 t/h against 1000 t/h asked, and the MPC's planned loading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

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

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

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

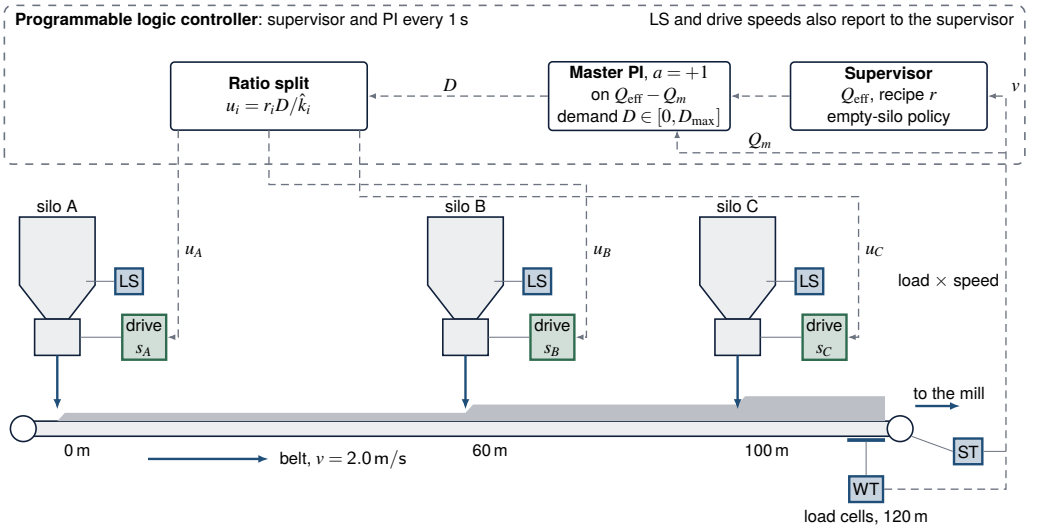


Figure 2.5 The master loop of the simulated conveyor drawn with its hardware. The belt scale's load cells WT and its speed sensor ST give the flow Q_m, the empty switches LS and the drives' own speed feedback complete the measurement packet, and the supervisor and the master PI, reverse acting with $a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

2.7 Technology in Practice

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

What Measures It

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

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

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

What Moves It

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

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

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

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

Where the Controller Runs

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

IN PRACTICE How a Belt Weigher Is Verified

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

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

What Is Modernizing It

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

2.8 Limitations

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

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

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

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

2.9 Conclusion

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

KEY POINTS

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

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

References

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

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

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

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

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

[6] J. E. Normey-Rico, T. L. M. Santos, R. C. C. Flesch, and B. C. Torrico, "Control of dead-time process: From the Smith predictor to general multi-input multi-output dead-time compensators," *Frontiers in Control Engineering*, vol. 3, p. 953768, 2022.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

2.10 Exercises

2.1 Concept check · 2 min

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

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

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

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

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

2.2 Concept check · 1 min

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

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

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

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

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

2.3 Concept check · 2 min

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

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

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

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

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

2.4 Theory and calculation · 6 min

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

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

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

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

2.5 Theory and calculation · 6 min

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

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

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

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

2.6 Theory and calculation · 8 min

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

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

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

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

2.7 App challenge · 8 min

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

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

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

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

2.8 App challenge · 8 min

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

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

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

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

2.9 App challenge · 8 min

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

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

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

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

2.10 Application area · 4 min

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

2.11 Application area · 3 min

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

2.12 Review the agent · 5 min

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